

ViChaR: A Dynamic Virtual Channel Regulator for Network-on-Chip Routers*

Chrysostomos A. Nicopoulos, Dongkook Park, Jongman Kim,

N. Vijaykrishnan, Mazin S. Yousif[†], Chita R. Das

Dept. of CSE, The Pennsylvania State University
University Park, PA 16802, USA
{nicopoul,dpark,jmkim,vijay,das}@cse.psu.edu

[†]Corporate Technology Group, Intel Corp.
Hillsboro, OR 97124, USA
mazin.s.yousif@intel.com

Abstract

The advent of deep sub-micron technology has recently highlighted the criticality of the on-chip interconnects. As diminishing feature sizes have led to increases in global wiring delays, Network-on-Chip (NoC) architectures are viewed as a possible solution to the wiring challenge and have recently crystallized into a significant research thrust. Both NoC performance and energy budget depend heavily on the routers' buffer resources. This paper introduces a novel unified buffer structure, called the dynamic Virtual Channel Regulator (ViChaR), which dynamically allocates Virtual Channels (VC) and buffer resources according to network traffic conditions. ViChaR maximizes throughput by dispensing a variable number of VCs on demand. Simulation results using a cycle-accurate simulator show a performance increase of 25% on average over an equal-size generic router buffer, or similar performance using a 50% smaller buffer. ViChaR's ability to provide similar performance with half the buffer size of a generic router is of paramount importance, since this can yield total area and power savings of 30% and 34%, respectively, based on synthesized designs in 90 nm technology.

1. Introduction

Rapidly diminishing feature sizes into the nanoscale regime have resulted in dramatic increases in transistor densities. While gate delays are scaling down accordingly, wiring delays are, in fact, increasing; as wire cross-sections decrease, resistance increases. This undesirable behavior has transformed the interconnect into a major hindrance. A signal would require multiple clock cycles to traverse the length of a large System-on-Chip (SoC). To combat the delay issues emanating from slow global wiring, researchers have proposed the use of packet-based communication networks, known as Networks-on-Chip (NoC) [1-4]. NoCs, much like macro networks, can scale efficiently as the number of nodes (i.e. processing elements) increases. Besides performance, current designs indicate an additional alarming trend pertaining to the on-chip interconnect: the

chip area and power budgets are increasingly being dominated by the interconnection network [5-7]. As the architectural focus shifts from monolithic, computation-centric designs to multi-core, communication-centric systems, communication power has become comparable to logic and memory power, and is expected to eventually surpass them [6]. This ominous trend has been observed by several researchers [1, 5, 8] and the realization of its ramifications has fueled momentum in investigating NoC architectures. Researchers have proposed sophisticated router architectures with performance enhancements [9], area-constrained methodologies [7], power-efficient and thermal-aware designs [5, 10], and fault-tolerant mechanisms [11].

It is known that router buffers are instrumental in the overall operation of the on-chip network. However, of the different components comprising the interconnection fabric of SoCs, buffers are the largest leakage power consumers in an NoC router, consuming about 64% of the total router leakage power [12]. Similarly, buffers consume significant dynamic power [8, 13] and this consumption increases rapidly as packet flow throughput increases [13]. In fact, it has been observed that storing a packet in a buffer consumes far more energy than transmitting the packet [13]. Furthermore, the area occupied by an on-chip router is dominated by the buffers [2, 14, 15]. Consequently, buffer design plays a crucial role in architecting high performance and energy efficient on-chip interconnects, and is the focus of this paper.

1.1. Importance of Buffer Size and Organization

Decreasing the buffer size arbitrarily to reclaim silicon area and minimize power consumption is not a viable solution, because of the intricate relationship between network performance and buffer resources. Buffer size and management are directly linked to the flow control policy employed by the network; flow control, in turn, affects network performance and resource utilization. Whereas an efficient flow control policy enables a network to reach 80% of its theoretical capacity, a poorly implemented policy would result in a meager 30% [16]. Wormhole flow control [17] was introduced to improve performance through finer-granularity buffer and channel control at the flit level instead of the packet level (a flit is the smallest unit of flow control; one packet is composed of a number of

* This research was supported in part by NSF grants CCR-0208734, EIA-0202007, CCF-0429631, CNS-0509251, CRI-0454123, CAREER 0093085, SRC grant 00541, and a grant from DARPA/MARCO GSRC.




Figure 1. Limitations of a Statically Assigned Buffer Organization (H=Head flit, D=Data/Middle flit, T=Tail flit)

flits). This technique relaxes the constraints on buffer size at each router, allowing for a more efficient use of storage space than store-and-forward and virtual cut-through [18] switching. However, the channel capacity is still poorly utilized; while the buffers are allocated at the flit level, physical paths are still allocated at the packet level. Hence, a blocked packet can impede the progress of other packets waiting in line and may also cause multi-node link blocking (a direct consequence of the fact that the flits of a single packet are distributed across several nodes in wormhole routers). To remedy this predicament, Virtual Channel (VC) flow control [19] assigns multiple virtual paths (each with its own associated buffer queue) to the same physical channel. It has been shown that VC routers can increase throughput by up to 40% over wormhole routers without VCs [16]. As a side bonus, virtual channels can also help with deadlock avoidance [20]. The work in this paper assumes, without loss of generality, the use of VC-based wormhole flow control, which suits the low buffer requirements of NoC routers.

The way virtual channels – and hence buffers – are organized within a router is also instrumental in optimizing performance. The number of VCs per physical channel and the VC depth are two parameters that form an elaborate interplay between buffer utilization, throughput and latency. Researchers in the macro-network field have identified the decisive role of virtual channel organization in overall system performance [21, 22]. Detailed studies of the relation between virtual channels and network latency indicate that for low traffic intensity, a small number of VCs can suffice. In high traffic rates, however, increasing the number of VCs is a more effective way of improving performance than simply increasing the buffer depth [22]. Under light network traffic, the number of packets traveling through a router is small enough to be accommodated by a limited number of VCs; increasing the number of VCs yields no tangible benefits. Under high traffic, many packets are contending for router resources; increasing VC depth will not alleviate this contention because of Head-of-Line (HoL) blocking. Increasing the number of VCs, though, will allow more packets to share the physical channels. This dichotomy in VC organization implies that routers with fixed buffer structures will either be underutilized or will underperform under certain traffic conditions, as illustrated in the examples of Figure 1. This figure highlights the weaknesses of statically-partitioned buffers.

Since buffer resources come at a premium in resource-constrained NoC environments (they consume valuable power and silicon real-estate), it is imperative to limit the buffer size to a minimum without severely affecting performance. This objective function can only be achieved through the use of efficient management techniques which optimize buffer utilization. Since size and organization are design-time decisions, they cannot be dynamically changed during operation based on observed traffic patterns. However, the use of a carefully designed buffer controller can significantly affect the efficiency of storing and forwarding of the flits. Therefore, the throughput of a switch can be maximized through dynamic and real-time throttling of buffer resources.

1.2. A Dynamic Virtual Channel Regulator

Given the aforementioned significance of the NoC buffers in the area, power and performance triptych, we thereby introduce ViChar^{*}: a dynamic Virtual Channel Regulator, which dispenses VCs according to network traffic. The ViChar module is a very compact unit operating at the granularity of one router input/output port; therefore, a conventional 5-port NoC router would employ five such units to oversee buffer management.

ViChar's operation revolves around two intertwined concepts which constitute the *two fundamental contributions* of this work:

(1) ViChar uses a Unified Buffer Structure (UBS), instead of individual and statically partitioned First-In-First-Out (FIFO) buffers. While the unified buffer concept is not new, in this work we are revisiting the concept within the confinements of the strict resource limitations of on-chip networks. This is the first attempt to incorporate a unified buffer management in NoC routers. The new flavor in our endeavor stems from a fundamentally different implementation approach: we introduce a novel, table-based design which provides single-clock operation without incurring prohibitive overhead. Most importantly though, it enables the use of a flexible and dynamically varying virtual channel management scheme, thereby replacing the conventional, static resource allocation.

(2) ViChar provides each individual router port with a variable number of VCs, each of which is dispensed dynamically according to network traffic conditions. This translates into fewer but deeper VCs

* The name ViChar was intentionally chosen to echo the word Vicar, who is someone acting as a substitute or agent for a superior.

under light traffic, and more but shallower VCs under heavy traffic. This attribute successfully marries two contradicting buffer organizations, which are impossible to combine in conventional, statically-allocated buffers. Furthermore, ViChaR's dynamic allocation scheme ensures a smooth continuum between these two extremes (few/deeper VCs versus more/shallower VCs) as the network intensity fluctuates.

The proposed ViChaR architecture and a generic buffer architecture were synthesized in 90 nm technology. Area and power extracts indicate that a modest overhead in area and power due to more elaborate control logic in ViChaR is amortized by greater area and power gains through the use of fewer virtual channel arbiters. Thus, overall ViChaR allows for 4% area reduction and incurs a minimal 2% power increase compared to an equal-size generic buffer implementation. Further, simulations with a cycle-accurate NoC simulator under various traffic patterns show that ViChaR reduces network latency by 25% on average compared to a generic buffer of equal size. Alternatively, ViChaR can achieve the same performance as a conventionally buffered router by using 50% less buffer space. This result is of *profound significance* because it allows the designer to *reduce the buffer size by half without affecting performance*. This yields net area and power benefits of around 30% and 34%, respectively, over the entire NoC router.

The rest of the paper is organized as follows: a summary of related work is presented in Section 2, the proposed ViChaR architecture is analyzed in Section 3, simulation results are discussed in Section 4, and the concluding remarks are given in Section 5.

2. Related Work in Buffer Design

Interest in packet-based on-chip networks has rapidly gained momentum over the last few years, and analysis and optimization of on-chip interconnect architectures have garnered great attention. In this section, we focus solely on buffer related aspects. Within the realm of on-chip buffer design, both size and organization have been shown to be directly related to network performance [14]. Buffer sizing in particular has been investigated in [14, 15]. However, these papers adopt a static approach, where optimal buffer sizes are pre-determined at design-time based on a detailed analysis of application-specific traffic patterns. The sizing is optimal for only one particular application and one hardware mapping. However, a technique to dynamically alter the buffer organization at run-time is more desirable for a general purpose and reconfigurable SoC executing different workloads. A dynamic scheme would maximize utilization regardless of the traffic type in the NoC.

Centralized buffer organizations have been studied extensively in the macro-network realm, but the solutions proposed are not amenable to resource constrained on-chip implementations. In particular, a unified and dynamically-allocated buffer structure was originally presented in [23] in the form of the Dynamically Allocated Multi-Queue (DAMQ) buffer. However, whereas the DAMQ architecture was part of a single-chip communication coprocessor for multi-computer systems, the proposed implementation in this

paper is aimed at area- and power-constrained, ultra-low latency on-chip communication. This profoundly affected our design considerations as follows:

(1) The DAMQ used a fixed number of queues (i.e. virtual channels) per input port. Specifically, four queues were used, one for each of three output ports and a local processor interface. Consequently, all packets in the same queue had to obey the FIFO order, i.e. all packets in the same queue could still get stuck behind a blocked packet at the head of the queue.

(2) The control logic of the DAMQ buffer was very complex, relying on a system of linked lists to organize the data path. These linked lists were stored in pointer registers which had to be updated constantly. This caused a three-cycle delay for every flit arrival/departure, mainly because data had to be moved between pointer registers, and a so-called "free list" had to be updated (a linked list keeping track of available buffer slots) [24]. This three-cycle delay – while acceptable for inter-chip communication – would prove intolerable in an on-chip router.

The DAMQ project spawned a few other designs, which aimed to simplify the hardware implementation and lower overall complexity. Two notable examples of these designs were the DAMQ with self-compacting buffers [25] and the Fully Connected Circular Buffer (FC-CB) [26]. Both designs have less overhead than the linked-list approach of [23] by employing registers, which selectively shift some flits inside the buffer to enable all flits of one VC to occupy a contiguous buffer space. The FC-CB design [26] improves on [25] by using a circular structure, which shifts in only one direction and ensures that any flit will shift by at most one position each cycle. However, the FC-CB has two main disadvantages when applied to an on-chip network. First, being fully connected, it requires a $P^2 \times P$ crossbar instead of the regular $P \times P$ crossbar for a P -input switch. Such large and power-hungry crossbars are unattractive for on-chip routers. Second, the circular shifter allows an incoming flit to be placed anywhere in the buffer and requires selective shifting of some parts of the buffer while leaving the rest of the buffer undisturbed. This functionality inflicts considerable increases in latency, area and power over a simple, non-shifting buffer implementation, like the proposed ViChaR design. The overhead is due to the large MUXes which are required between each buffer slot to enable both shifting and direct input.

The circular-shift buffer of the FC-CB was implemented in Verilog HDL and synthesized in 90 nm commercial TSMC libraries to assess its viability in on-chip designs. The circular buffer implementation of the FC-CB increases the datapath delay by 26% compared to ViChaR's stationary (i.e. non-shifting) buffer. Increases in datapath delay may affect the pipeline period in deeply pipelined router designs; a longer period will adversely affect throughput. Moreover, the FC-CB's large MUXes incur an increase of approximately 18% in buffer area. More importantly, though, the continuous shifting of the FC-CB buffer every clock cycle (assuming continuous incoming traffic) increases the dynamic power budget by 66%. Obviously, this overhead renders the FC-CB implementation

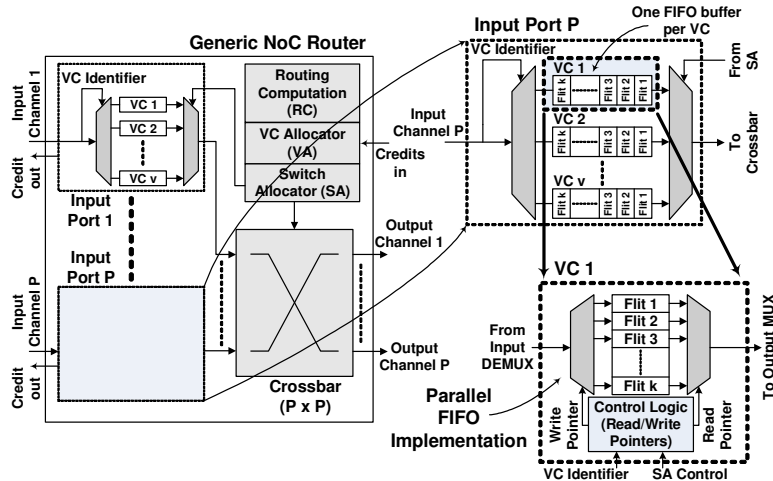


Figure 2. A Generic NoC Router Architecture

unattractive for on-chip applications. Finally, the FC-CB still works with a fixed number of VCs, just like the DAMQ design. In this paper, we will show that a dynamically variable number of VCs optimizes performance.

The notion of dynamically allocating VC resources based on traffic conditions was presented in [27], through the VCDAMQ and DAMQ-with-recruit-registers (DAMQWR) implementations. However, both designs were coupled to DAMQ underpinnings; hence, they employed the linked-list approach of the original DAMQ, which is too costly for an on-chip network. Nevertheless, the work of [27] highlighted the significance of dynamic allocation of buffer resources, which forms the premise of the design proposed in this paper.

Finally, the Chaos router [28] and BLAM routing algorithm [29] provide an alternative technique to saving buffer space. They employ packet misrouting, instead of storage, under heavy load. However, randomized (non-minimal) routing may make it harder to meet strict latency guarantees required in many NoCs (e.g., multimedia SoCs). Moreover, these schemes do not support dynamic VC allocation to handle fluctuating traffic.

3. The Proposed Dynamic Virtual Channel Regulator (ViChaR)

3.1. A Baseline NoC Router

A generic NoC router architecture [9] is illustrated in Figure 2. The router has P input and P output channels/ports. In most implementations, $P=5$; four inputs from the four cardinal directions (North, East, South and West) and one from the local Processing Element (PE). The Routing Computation unit, RC, is responsible for directing the header flit of an incoming packet to the appropriate output Physical Channel/port (PC) and dictating valid Virtual Channels (VC) within the selected PC. The routing is done based on destination information present in each header flit, and can be deterministic or adaptive. The Virtual channel Allocation unit (VA) arbitrates amongst all packets

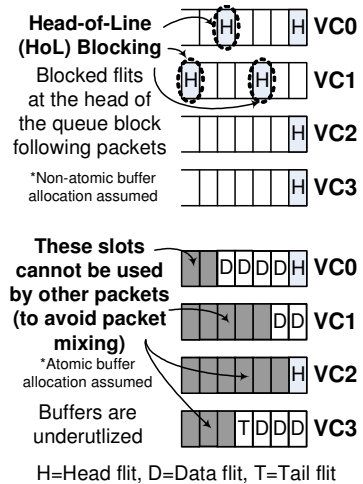


Figure 3. Limitations of Existing FIFO Buffers

requesting access to the same VCs and decides on winners. The Switch Allocation unit (SA) arbitrates amongst all VCs requesting access to the crossbar and grants permission to the winning flits. The winners are then able to traverse the crossbar and are placed on the respective output links. Simple router implementations require a clock cycle for each component within the router. Lower-latency router architectures parallelize the RC, VA and SA using speculative allocation [30], which predicts the winner of the VA stage and performs SA based on that. Further, look-ahead routing can also be employed to perform routing of node $i+1$ at node i . These two modifications have led to two-stage, and even single-stage [9], routers, which parallelize the various stages of operation.

So far, as a result of scarce area and power resources and ultra-low latency requirements, on-chip routers have relied on very simple buffer structures. In the case of virtual channel-based NoC routers, these structures consist of a specified number of FIFO buffers per input port, with each FIFO corresponding to a virtual channel. This is illustrated in Figure 2. Such organization amounts to a static partitioning of buffer resources. Hence, each input port of an NoC router has v virtual channels, each of which has a dedicated k -flit FIFO buffer. Current on-chip routers have small buffers to minimize their overhead; v and k are usually much smaller than in macro networks [9]. The necessity for very low latency dictates the use of a parallel FIFO implementation, as shown in the bottom right of Figure 2. As opposed to a serial FIFO implementation [31], the parallel flavor eliminates the need for a flit to traverse all slots in a pipelined manner before exiting the buffer [31]. This fine-grained control requires more complex logic, which relies on read and write pointers to maintain the FIFO order. Given the small sizes of on-chip buffers, though, the inclusion of a parallel FIFO implementation is by no means prohibitive. The buffers within an NoC router can be implemented as either registers or SRAM/DRAM memory [32, 33]. However, given the relatively small buffer sizes employed, it is more reasonable to use small registers as buffers to avoid the address decoding/encoding latencies of big memories

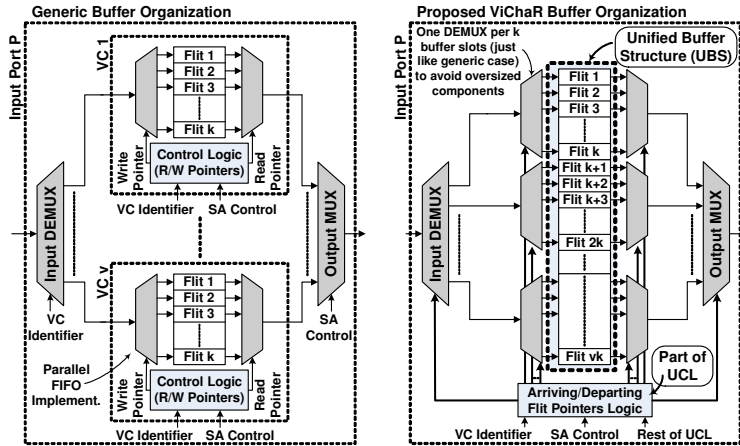


Figure 4. The Proposed ViChar Architecture

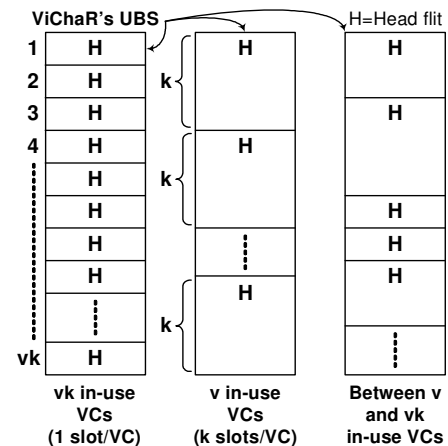


Figure 5. Possible VC Configurations in ViChar

and the access latencies associated with global bitlines/wordlines [32]. To this extent, the NoC buffers in this paper were implemented as registers.

FIFO buffers in statically assigned buffer structures have two inherent disadvantages. First, a packet at the head of a VC whose designated output port is busy will block all subsequent packets in that VC from being transmitted (assuming non-atomic buffer allocation) even if their designated output ports are free. This Head-of-Line (HoL) blocking can severely affect network performance in congested conditions, similar to the previously discussed DAMQ. This scenario is illustrated at the top of Figure 3. Second, if only part of a packet occupies a VC buffer at a given time, then any vacant slots in that buffer cannot be reassigned to a new packet for as long as that VC is reserved by the partial packet to avoid packet/message mixing. Thus, a VC buffer may only be occupied by a single header flit because the remaining flits happen to be blocked in preceding routers due to congestion. In such a scenario, the remaining free slots in the buffer cannot be assigned to other packets until the tail flit of the current packet releases the VC. This attribute of FIFO buffers can lead to substantial under-utilization of the buffers, as shown at the bottom of Figure 3, and cripple network performance.

3.2. The ViChar Architecture

Figure 4 illustrates the buffer organization of a conventional NoC router (left) and our proposed alterations (right). The crux of ViChar is composed of two main components: (1) the *Unified Buffer Structure (UBS)*, shown in Figure 4, and (2) the associated control logic, called *Unified Control Logic (UCL)*.

Figure 4 shows only one of the five sub-modules of UCL, the *Arriving/Departing Flit Pointer Logic*. This sub-module constitutes the interface between the UBS and the UCL; the UCL controls the unified buffer (UBS) through the *Arriving/Departing Flit Pointer Logic* module. A top-level block diagram of the entire ViChar architecture is shown in Figure 6. This figure illustrates all five of the UCL sub-modules: (1) the *Arriving/Departing Flit Pointers Logic*, (2) the *Slot Availability Tracker*, (3) the *VC Availability Tracker*, (4) the *VC Control Table*, and (5) the *Token (VC) Dispenser*. The operation of each component and the

interaction between the UBS and its controlling entity (the UCL) are described in detail in section 3.2.2. All five modules function independently and in parallel, which is of critical importance to the ultra-low latency requirements of the router. The UCL components work in tandem with the unified buffer (UBS), providing dynamic allocation of both virtual channels and their associated buffer depth. As illustrated in Figure 6, the two main ViChar components (UBS and UCL) are logically separated into two groups: the unified buffer (UBS) and two of the five UCL modules (the *Arriving/Departing Flit Pointers Logic* and the *Slot Availability Tracker*) are situated at the input side of the router (i.e. to accept all incoming flits), while the remaining modules of the control logic (UCL) are responsible for the VC arbitration of all flits destined to a particular output port. Based on incoming traffic and information from the Slot and VC Availability Trackers, the *Token (VC) Dispenser* grants VC IDs to new packets accordingly. The *VC Control Table* is the central hub of ViChar's operation, keeping track of all in-use VCs and a detailed status of the unified buffer (UBS). When flits arrive and/or depart, the *Arriving/Departing Flit Pointers Logic* controls the UBS's MUXes and DEMUXes in accordance with the *VC Control Table*.

It is important to realize that the UBS is physically identical to the generic buffer structure: the v independent k -flit FIFO buffers of a traditional implementation are simply *logically* grouped in a single vk -flit entity (the UBS in Figure 4). Hence, other than the control logic, there is no additional hardware complexity, since the vk -flit UBS is NOT a large, monolithic structure; it groups the existing buffers together, and it is only through the use of its control mechanism (the UCL) that the buffers appear as a *logically unified structure*. As shown in Figure 4, UBS retains the same number of MUXes/DEMUXes as the generic implementation, i.e. one MUX/DEMUX per k flits, to avoid large (and hence slower) components.

3.2.1. Variable Number of Virtual Channels. Whereas a conventional NoC router can support only a fixed, statically assigned number of VCs per input port (namely v , as shown in Figure 4), the ViChar architecture can have a variable number of assigned VCs, based on

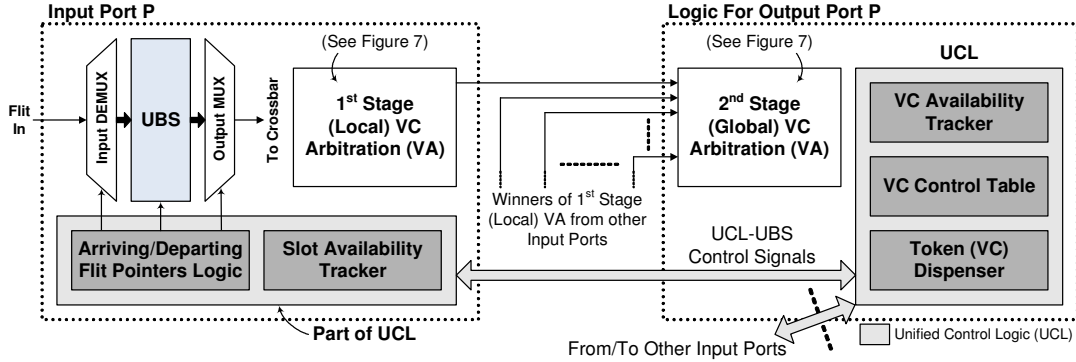


Figure 6. ViChaR Block Diagram (Only One of P Ports Shown Here)

network conditions. *ViChaR* assigns at most one packet to each VC so as to enable fine flow control granularity; on the contrary, the sharing of a single VC by multiple packets can lead to situations where a blocked packet impedes the progress of another packet which happens to use the same VC (known as HoL blocking, as described in Section 3.1). A vk -flit *ViChaR* structure can support anywhere between v VCs (when each VC occupies the maximum of k flits) and vk VCs (when each VC occupies the minimum of 1 flit) at any given time under full load. This variability in the number of in-use VCs is illustrated in Figure 5. To aid understanding, each VC in Figure 5 is shown to occupy a contiguous space in the buffer (UBS); in reality, however, this may not be the case because the UBS allows the VCs to include non-consecutive buffer slots (this fact will be explained in more detail in Section 3.2.2). Hence, the system can support a variable number of in-flight packets per port, dynamically allocating new VCs when network conditions dictate it. Dynamic variability in in-flight messages can increase throughput under heavy traffic.

As a result of its unified buffer and dynamic behavior, the *ViChaR* structure alters the Virtual channel Allocation (VA) logic of the router. Since the router function may return multiple output VCs restricted to a single physical channel [30], two arbitration stages are required in both the generic and *ViChaR* cases, as shown in Figure 7. In the generic case, the first stage reduces the number of requests from each input VC to one (this ensures the request of a single VC at a particular output port by each input VC). Subsequently, the winning request from each input VC proceeds to the second arbitration stage. Details of the VA operation are omitted for brevity, but can be found in [30].

In the proposed *ViChaR* architecture, VA takes a different approach due to the dynamic VC allocation scheme: the first arbitration stage reduces the number of requests for a particular output port to one request per input port. The generic router (Figure 7(a)) requires $v:1$ arbiters, since the number of VCs supported is fixed to v . *ViChaR*, on the other hand, supports anywhere between v and vk VCs per port at any given time. To accommodate the worst case scenario (i.e. vk in-flight VCs), *ViChaR* needs larger $vk:1$ arbiters in stage 1 of the allocation (Figure 7(b)). The second arbitration stage in *ViChaR* produces a winner for each output port among all the competing input ports. Therefore, while the proposed *ViChaR* architecture uses larger Stage 1

arbiters ($vk:1$ vs. $v:1$), it uses much smaller and fewer Stage 2 arbiters. The reason for the simplified second stage is that *ViChaR* dynamically allocates VCs as needed, instead of accepting requests for specific VCs (which would necessitate one arbiter per output VC, just like the generic case). It is this attribute that helps the *ViChaR* implementation incur only a slight increase in power consumption (and even achieve a small area decrease), compared to a generic architecture, as will be shown shortly.

The variable number of VCs supported by *ViChaR* also necessitates bigger arbiters in the first stage of Switch Allocation (SA), as shown in Figure 8. Similar to VA, switch allocation is performed in two stages. The first stage accounts for the sharing of a single port by a number of VCs. Again, *ViChaR* needs larger $vk:1$ arbiters. The second stage arbitrates between the winning requests from each input port (i.e. P ports) for each output port; thus, it is the same for both architectures. The *ViChaR* overhead due to the bigger stage-1 SA arbiters (illustrated in Table 1's detailed breakdown) is almost fully amortized by the bigger savings resulting from the smaller VA stage discussed previously.

To analyze the area and power overhead, NoC routers with (a) a generic buffer and (b) the proposed *ViChaR* buffer were implemented in structural Register-Transfer Level (RTL) Verilog and then synthesized in Synopsys Design Compiler using a TSMC 90 nm standard cell library. The resulting designs operate at a supply voltage of 1 V and a clock frequency of 500 MHz. The routers have 5 input ports (i.e. $P=5$), 4 VCs per input port (i.e. $v=4$), each VC is four-flit deep (i.e. $k=4$), and each flit is 128 bits long. Both area and power estimates were extracted from the synthesized router implementations. A comparison of the area and power overhead of the two schemes is shown in Table 1. Note that both routers have equal buffer space ($vk=16$ buffer slots per input port) for fairness. It is evident that while *ViChaR* incurs an overhead in terms of control logic and switch allocation (SA), this overhead is over-compensated (in terms of area) by a larger reduction in the VA logic. Thus, the *ViChaR* model provides area savings of around 4%. In terms of power, *ViChaR* consumes slightly more power (1.75%). This power increase, however, is negligible compared to the performance benefits of *ViChaR*, as will be demonstrated in Section 4.

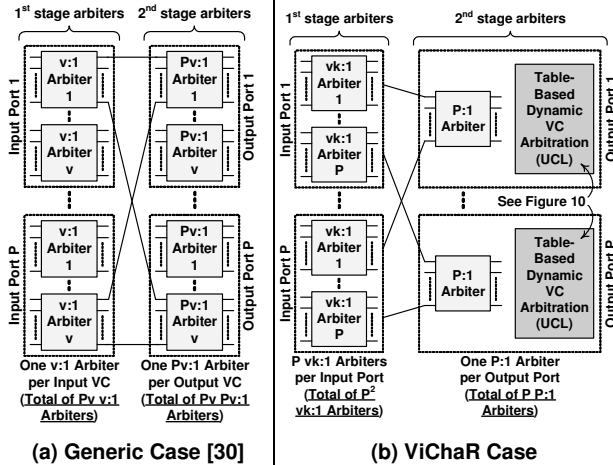


Figure 7. Virtual Channel Arbitration (VA)

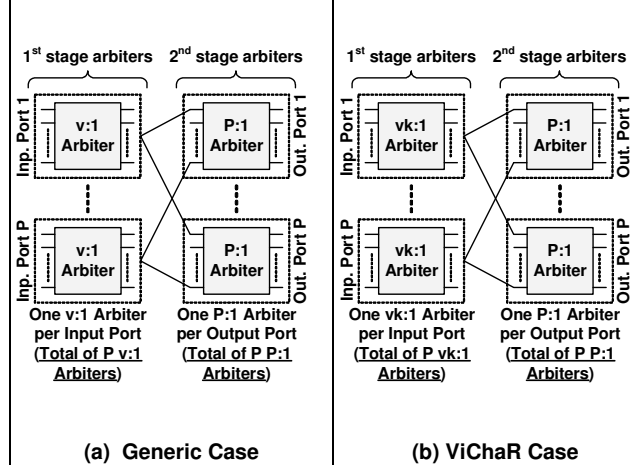


Figure 8. Switch Allocation (SA)

3.2.2. ViChaR Component Analysis. The key challenges in designing ViChaR were to *avoid* (a) deepening the router's pipeline, and (b) decreasing the operating frequency. To circumvent the multi-cycle delay induced by a linked-list approach [23] to ViChaR, we opted instead for a *table-based* approach, as illustrated in Figure 10. This logic is required for each output port in the router. Following is a break-down of the control logic (UCL) sub-modules of the proposed ViChaR architecture:

VC Control Table: The VC Control Table (see Figure 10) forms the core of the control logic of ViChaR. It is a compact table, holding the slot IDs of all flits currently in the buffers, which are requesting the particular output port (e.g. West). Note that since the number of buffer slots in on-chip routers is resource-constrained, the size of the table is minimal, as demonstrated by the low overhead in the control logic in Table 1. The VC Control Table is organized by VC ID, with each VC having room for at most a single packet. Without loss of generality, in this work we assumed a packet to consist of four flits: a Head flit, two Data (middle) flits, and a Tail flit. The packet size is assumed to be constant, but the table can trivially be changed to accommodate a variable-sized packet protocol. As seen in the VC Control Table box of Figure 10 (right-hand side), the VCs can include non-consecutive buffer slots (e.g. VC1 comprises of slots 2, 4, 6 and 7) of the South input port (i.e. flits arriving from the South). This attribute allows full-flexibility in buffer utilization and avoids the issues encountered in statically-allocated buffers. VC3 only occupies one slot (10) in Figure 10. In a static buffer, 3 additional slots would have to be reserved for the remaining flits of VC3; those slots would remain unused if the remaining flits happened to be blocked in previous routers. Instead, in ViChaR those slots can be used by other VCs, thus maximizing the buffer utilization. Furthermore, the use of a table-based controller makes the management of a variable number of VCs very easy: non-used VCs are simply NULLed out in the VC Control Table (e.g. VC4 in Figure 10).

Arriving/Departing Flit Pointers Logic: The Flit Pointers Logic directly controls the Input and Output MUXes/ DEMUXes of the unified buffer (UBS), as

illustrated in Figure 9, and is directly linked to the VC Control Table module. Once a flit departs, its location in the VC Control Table is invalidated by asserting a NULL bit. There is a set of such pointers for each VC in the table. However, the overhead is minimal due to the simplicity of the pointer logic; both Departing and Arriving Flit Pointers are implemented in combinational logic and simply have to observe the non-NULL locations in their VC. For example, the Departing Flit pointer points at the first non-NULL location (in its particular VC) starting from the left of the table, as shown on the right side of Figure 10 for VC2 (in the VC Control Table box). If all the entries in a single row of the VC Control Table are NULL, then the VC must be empty; thus, the pointer logic releases the VC by notifying the VC Availability Tracker (Release Token signal in Figure 9). When a new flit arrives, the pointer logic guides the flit to the appropriate slot in the unified buffer (UBS), based on the flit's VC ID and information from the Slot Availability Tracker. Finally, newly arrived header flits in the UBS can request an output VC by first undergoing local (1st stage) arbitration (top right of Figure 9), and then global (2nd stage) arbitration (bottom left of Figure 10).

VC and Slot Availability Trackers: The VC Availability Tracker simply keeps track of all the VCs in the VC Control Table that are not used. The Token (VC) Dispenser dynamically assigns VCs to new incoming packets based on information provided by the VC Availability Tracker. Similarly, the Slot Availability Tracker keeps track of all the UBS slots which are not in use. When a new flit arrives, it is stored into a slot indicated by the Slot Availability Tracker. The VC and Slot Availability Trackers are functionally identical. They consist of a small table, as shown at the bottom right of Figure 9 (Slot Availability Tracker) and the top left of Figure 10 (VC Availability Tracker). Each row of the table corresponds to one VC ID (in the VC Availability Tracker) or one buffer slot (in the Slot Availability Tracker). For each entry in the table, one bit indicates that the VC/Slot is available (logic 1) or occupied (logic 0). Both trackers have a pointer which points to the top-most available entry. If all VCs are occupied (i.e. all-zero table in the VC Availability

Table 1. Area and Power Overhead of the ViChar Architecture.

The results in this table assume equal-size buffers for both router designs. However, ViChar's efficient buffer management scheme allows for a 50% decrease in buffer size with no performance degradation (see Section 4). In such a case, area and power is reduced by 30% and 34%, respectively, over the whole router.

Component (one input port)	Area (in μm^2)	Power (in mW)
ViChar Table-Based Contr. Logic	12,961.16	5.36
ViChar Buffer Slots (16 slots)	54,809.44	15.36
ViChar VA Logic	27,613.54	8.82
ViChar SA Logic	6,514.90	2.06
TOTAL for ViChar Architecture	101,899.04	31.60
Generic Control Logic	10,379.92	5.12
Generic Buffer Slots (16 slots)	54,809.44	15.36
Generic VA Logic	38,958.80	9.94
Generic SA Logic	2,032.93	0.64
TOTAL for Gen. Architecture	106,181.09	31.06
ViChar Overhead / Savings	- 4,282.05	+ 0.54
	SAVINGS	OVERHEAD

Tracker), the Token (VC) Dispenser stops granting new VCs to requesting packets. Similarly, an all-zero table in the Slot Availability Tracker implies a full buffer (UBS); this is reflected in the credit information sent to adjacent routers. The functionality of the VC/Slot Availability Trackers is implemented in combinational logic, similar to the Flit Pointers Logic described above.

Token (VC) Dispenser: The Token (VC) Dispenser interfaces with the $P:1$ Arbiter and is responsible for dispensing free VCs to requesting packets. VCs here are like tokens; they are granted to new packets and then returned to the dispenser upon release. The flow diagram of the Dispenser's operation is illustrated on the right-hand side of Figure 10. Based on information provided by the VC Availability Tracker, the Token Dispenser decides whether to grant a VC or not. The VC dispenser keeps checking for possible deadlock situations among the in-use VCs. Deadlocks may occur in networks which employ adaptive routing schemes. If a pre-specified time threshold is exceeded, the Token Dispenser can channel an existing VC into one of the escape VCs to break the deadlock. As a proof of concept of ViChar's functionality, the experiments in this paper use deterministic (XY) routing, which is inherently deadlock-free. However, ViChar was designed to operate under adaptive routing schemes as well. Therefore, the Token (VC) Dispenser needs to account for possible deadlock situations. Toward that extent, a number of VCs can be designated as "escape", or "drain" channels to provide deadlock recovery in adaptive routing algorithms (escape channels employ a deterministic routing algorithm to break the deadlock) [34]. The Dispenser needs to switch deadlocked flits into these escape channels if there is a need. One experiment in Section 4.2 validates the effectiveness of this technique under adaptive routing.

Assuming that no deadlock situation exists, the Token Dispenser can proceed with its normal operation. The Dispenser grants new VCs on a First-Come-First-Served (FCFS) basis; if a new header flit wins the VC arbitration and the VC Availability Tracker indicates that a free VC is available, then the new packet will be granted a new VC. The Dispenser does not give priority

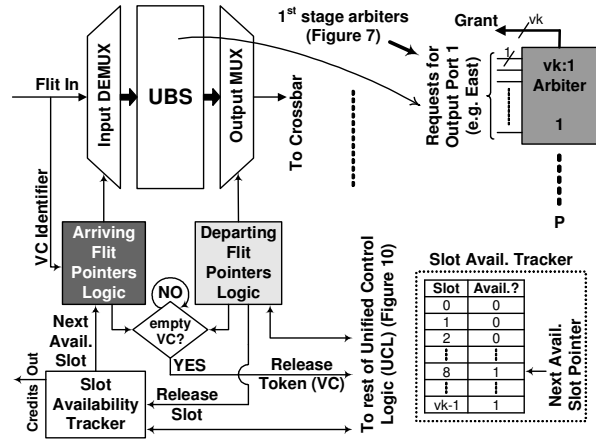


Figure 9. The ViChar UBS Architecture (One Input Port Shown)

to flits of existing VCs. In principle, a more elaborate mechanism could be used for dispensing new VCs, which would monitor on-going traffic and reach a decision based on some quantitative metric or prior traffic history. However, given the highly restrictive objective function of minimal area, power and latency budgets in the design of on-chip networks, such complex monitoring mechanisms were deemed infeasible. After all, ViChar was architected to operate within one clock cycle. The use of an FCFS scheme in the Token Dispenser turns out to be very efficient at maximizing performance. The Dispenser is able to self-throttle the dispensing of new VCs based on traffic conditions: if more packets request a channel (high traffic) more VCs are dispensed; if fewer packets are present (low traffic) fewer VCs are granted and more buffer depth is allotted to existing VCs.

ViChar's Effect on the Router Pipeline: The control logic (UCL) of ViChar was designed in such a way as to decouple the operation of the sub-modules from each other. Thus, sub-modules are kept compact and can all operate in parallel, hence completing the entire operation in a single clock cycle. This is a significant improvement over the three-clock cycle delay of [23]. Figure 11 shows the pipeline stages of both a generic and the ViChar router pipelines. As previously mentioned, the ViChar architecture modifies the VA and SA stages (Stages 2 and 3 in Figure 11). The dark-colored boxes indicate the components modified/added in the ViChar structure as compared to the generic case. As shown in the figure, the additional hardware operates in parallel without affecting the critical path of the router. This fact is also verified by our studies of the critical path delays of all major components of the router architectures (extracted from the synthesized designs). In both cases, the bottleneck that determines the minimum clock period is the arbitration logic (for the VA and SA stages, as shown in Figure 7 and Figure 8, respectively). All the components of the ViChar router remain within the slack provided by the slower arbiters. Hence, the ViChar architecture does not affect the pipeline depth or the clock frequency. Furthermore, since ViChar does not create any interdependencies between pipeline

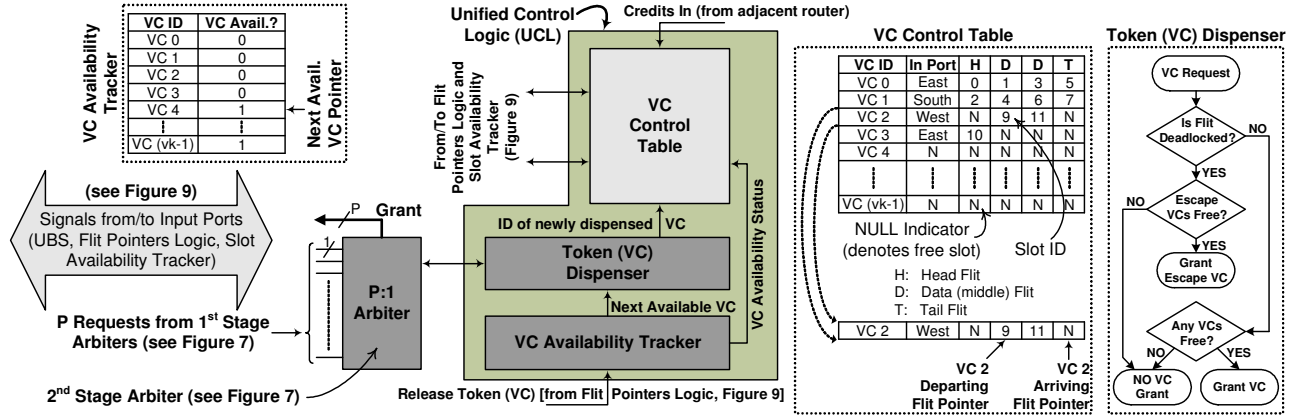


Figure 10. ViChar Table-Based UCL Architecture (Logic For One of P Ports Shown)

stages, it can also be used in speculative router architectures which minimize the pipeline length.

4. Simulation Results

4.1. Simulation Platform

A cycle-accurate on-chip network simulator was used to conduct detailed evaluation of the architectures under discussion. The simulator operates at the granularity of individual architectural components. The simulation test-bench models the pipelined routers and the interconnection links. All simulations were performed in a 64-node (8x8) MESH network with 4-stage pipelined routers. Each router has 5 physical channels (ports) including the PE-router channel. The generic router (shown as "GEN" in results graphs) has a set of 4 virtual channels per port. Each VC holds four 128-bit flits (i.e. a total of $5 \times 4 \times 4 = 80$ buffer slots). The ViChar router ("ViC" in results graphs) has a 16-flit unified buffer per port (i.e. a total of $5 \times 16 = 80$ buffer slots, just like the generic case). One packet consists of four flits. The simulator keeps injecting messages into the network until 300,000 messages (including 100,000 warm-up messages) are ejected. Two network traffic patterns were investigated: (1) Uniform Random (UR), where a node injects messages into the network at regular intervals specified by the injection rate, and (2) Self-Similar (SS), which emulates internet and Ethernet traffic. For destination node selection, two distributions were used: (1) Normal Random (NR), and (2) Tornado (TN) [35]. In all cases, except one, deterministic (XY) routing and wormhole switching were employed. One experiment used minimal adaptive routing to evaluate the systems in a deadlock-prone environment. Single link traversal was assumed to complete within one clock cycle at 500 MHz clock frequency. Both dynamic and leakage power estimates were extracted from the synthesized router designs and back-annotated into the network simulator.

4.2. Analysis of Results

Our simulation exploration starts with a latency comparison between a conventional, statically assigned buffer architecture and the proposed ViChar implementation. We first assume that both designs have equal-sized buffers; specifically, 16-flit buffers per input port (i.e. a total of 80 flits per NoC router). In the

generic design (GEN), the 16 buffer slots are arranged as 4 VCs, each with a 4-flit depth. ViChar (ViC), on the other hand, can dynamically assign its 16 buffer slots to a variable number of VCs, each with a variable buffer depth. Figure 12(a) and Figure 12(b) show the average network latency (in clock cycles) as a function of injection rate (in flits/node/cycle) for Uniform Random (UR) and Self-Similar (SS) traffic patterns, respectively. The graphs include results for both Normal Random (NR) and Tornado (TN) source-destination selection patterns. In all cases, ViChar substantially outperforms the generic architecture; by 28% (NR) and 24% (TN) on average for Uniform Random traffic, and 25% (NR) and 18% (TN) for Self-Similar traffic. More importantly, though, ViChar saturates at higher injection rates than the generic case.

Figure 12(c) shows the buffer occupancy at injection rates between 0.25 and 0.35 (i.e. before the onset of saturation). Higher buffer occupancy indicates network blocking. ViChar is clearly much more efficient at moving flits through the router; the buffer occupancy of a 16-flit/port ViChar design is considerably lower than an equal-size static configuration. *Buffer occupancy alone, however, is not an indicative metric, since it does not relay any information about network latency.* To validate ViChar's highly efficient buffer management scheme, its latency at these smaller buffer sizes should also be investigated. To that extent, Figure 12(d) and Figure 12(e) depict how the latency of ViChar at various buffer sizes compares to the latency of the generic architecture with a fixed 16-flit/port buffer size. It is evident from the graphs that *the UBS can achieve similar performance with less than half the buffer size of*

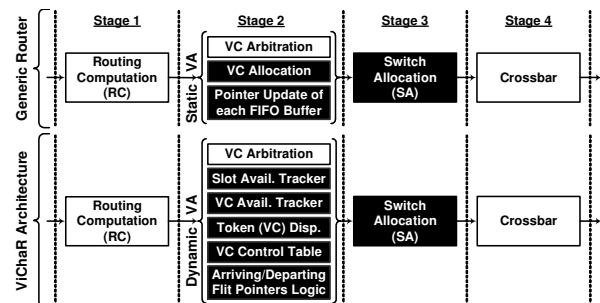


Figure 11. Generic & ViChar NoC Router Pipelines




Figure 12. Average Latency, % Buffer Occupancy, and Average Power Consumption Simulation Results

the generic architecture. This is of profound importance, since buffers dominate the area and power budgets of NoC routers; reducing the buffer size by 50% will yield significant savings. An alternative way to visualize this phenomenon is illustrated in Figure 12(f). This graph shows how the latency of ViChaR at various buffer sizes compares to the latency of the generic architecture with a 16-flit/port buffer size (horizontal dashed line) at an injection rate of 0.25. ViChaR has higher latency only when its buffer size drops below 8 flits per port. On the other hand, Figure 12(g) shows that decreasing the buffer size in a generic, statically assigned buffer structure always degrades performance.

Following on the very encouraging result that ViChaR can achieve similar performance as a conventional buffer by using only half the buffers, Figure 12(h) shows the total average power consumption of the 8x8 MESH network for different buffer configurations. For equally sized configurations, ViChaR consumes slightly more power than a conventional buffer structure. At injection rates up to 0.3, ViChaR consumes about 2% more power, corroborating the results of Table 1. At higher injection rates (when the network saturates), excessive switching activity causes this difference to grow a bit more, even

though it never exceeds 5%. However, since ViChaR's efficiency allows us to halve the buffer resources with no discernible effect on performance, the overall power drops by about 34% (ViC-8 in Figure 12(h)) for equivalent performance. Similarly, the area occupied by the router decreases by around 30%, based on synthesis results. These decreases can lead to more power- and area-efficient SoCs.

Figure 12(i) compares average network latency under minimal adaptive routing to validate ViChaR's effectiveness in handling deadlocks. Escape (drain) channels, which employ deterministic (i.e. deadlock-free) routing, were used in both the generic and ViChaR architectures to break deadlocks. Evidently, ViChaR was able to handle all deadlock situations while significantly outperforming the conventional design.

Figure 13(a) and Figure 13(b) present another metric of network performance, namely throughput (in flits per cycle). These graphs follow the same trend as the latency experiments, with ViChaR clearly outperforming a conventional buffer structure. Figure 13(c) includes the throughput of two different (but of equal size) generic configurations: 4 VCs each with a 3-flit depth, and 3 VCs with a 4-flit depth. The graph indicates that while varying the statically-assigned VC configuration of a

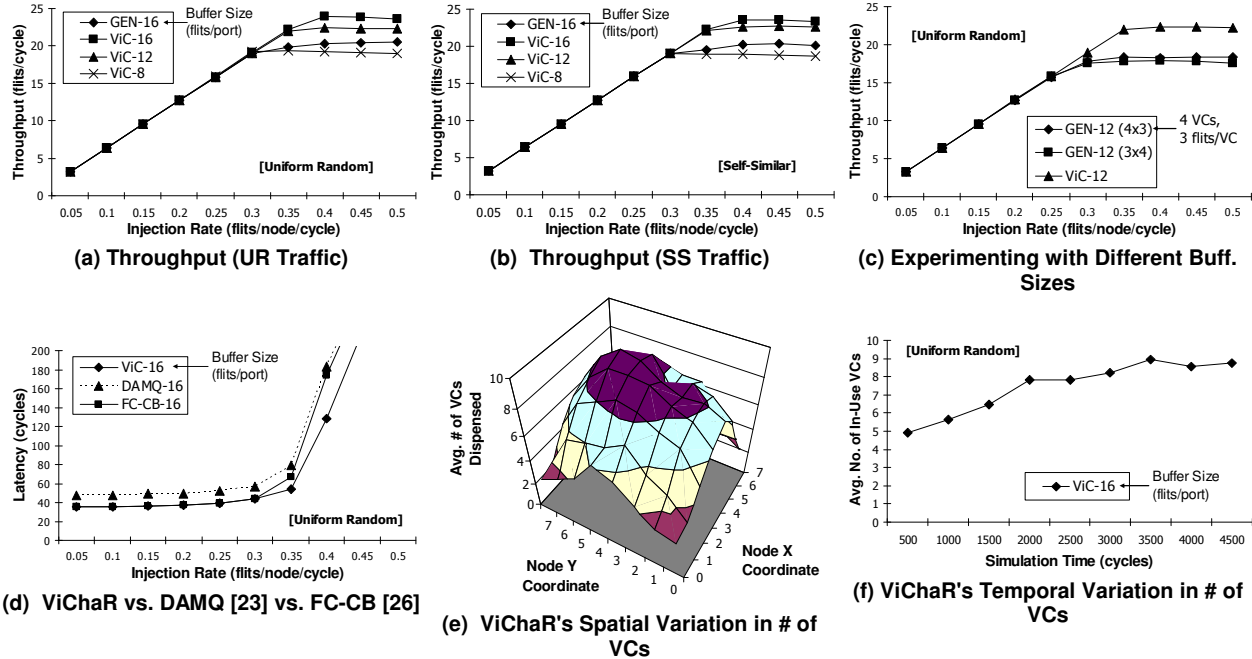


Figure 13. Simulation Results Demonstrating ViChaR's Efficient Virtual Channel Management Scheme

generic buffer does affect throughput, its performance still trails that of the dynamically variable design of ViChaR.

In the related work section (Section 2), we analyzed in detail why the unified buffers of the DAMQ [23] and FC-CB [26] would underperform compared to ViChaR's dynamic design. Both the DAMQ and FC-CB structures were implemented and incorporated into our cycle-accurate simulator. Figure 13(d) shows how all designs fare against each other. DAMQ loses out because of its 3-cycle buffer delay, and its fixed number of VCs, as previously explained. For a fair comparison, we assumed that the FC-CB design completes its buffer management procedure in one clock cycle (just like ViChaR). As seen in Figure 13(d), at low injection rates, the FC-CB's performance is almost identical to ViChaR's. However, as network traffic increases, FC-CB's performance starts to degrade compared to ViChaR. This is attributed to FC-CB's fixed number of VCs (i.e. just like DAMQ). Under heavier traffic loads, ViChaR's ability to dynamically dispense more VCs helps improve performance quite drastically. Note also that both FC-CB and DAMQ would incur much higher area and power penalties (as explained in Section 2). In terms of throughput (not shown here), ViChaR's improvement over DAMQ and FC-CB is a more modest 5% (on average). However, ViChaR would perform substantially better in latency-critical applications.

Finally, Figure 13(e) depicts the spatial variation in the number of VCs used in the 8x8 MESH, while Figure 13(f) shows the temporal variation over simulation time. The average number of VCs used varies continuously according to network traffic. Figure 13(e) shows the average number of VCs dispensed at each node of the 8x8 MESH network over the whole simulation time at an injection rate of 0.25. As expected, the nodes situated at

the middle of the network exhibit higher congestion; ViChaR successfully self-throttled its resources by granting more VCs in these nodes in order to optimize performance. In Figure 13(f), as the network fills up with packets, the average number of VCs used over all nodes increases accordingly to handle the traffic. These results validate the effectiveness of our FCFS scheme employed in the Token (VC) Dispenser (Section 3.2.2).

5. Conclusions

The continuing technology shrinkage into the deep sub-micron era has magnified the delay mismatch between gates and global wires. Wiring will significantly affect design decisions in the forthcoming billion-transistor chips, whether these are complex heterogeneous SoCs, or Chip Multi-Processors (CMP). Networks-on-Chip (NoC) have surfaced as a possible solution to escalating wiring delays in future multi-core chips. NoC performance is directly related to the routers' buffer size and utilization. In this paper, we introduce a centralized buffer architecture, called the Virtual Channel Regulator (ViChaR), which dynamically allocates virtual channels and buffer slots in real-time, depending on traffic conditions. Unlike current implementations, the ViChaR can dispense a variable number of VCs at any given time to maximize network throughput.

Simulation results using a cycle-accurate network simulator indicate *performance improvement of around 25%* under various traffic patterns, as compared to a conventional router with equal buffer size, with a modest 2% power increase. Most importantly, though, ViChaR is shown to achieve performance similar to that of a generic router, while using a 50% smaller buffer. This attribute is a testament to ViChaR's efficient dynamic buffer management scheme, and is a result of utmost

significance in the NoC arena. Synthesized designs in 90 nm technology indicate that *decreasing the ViChaR's buffer size by 50% leads to area and power savings of 30% and 34%, respectively, with no degradation in performance.*

For future work, we intend to evaluate the performance of ViChaR using workloads and traces from existing System-on-Chip architectures.

6. References

- [1] L. Benini and G. D. Micheli, "Networks on Chips: A New SoC Paradigm," *IEEE Computer*, vol. 35, pp. 70-78, 2002.
- [2] W. J. Dally and B. Towles, "Route Packets, Not Wires: On-Chip Interconnection Networks," in *Proceedings of the Design Automation Conference (DAC)*, 2001.
- [3] A. Hemani, A. Jantsch, S. Kumar, A. Postula, J. Oberg, M. Millberg, and D. Lindqvist, "Network on chip: An architecture for billion transistor era," in *Proceedings of the IEEE NorChip Conference*, 2000.
- [4] P. Guerrier and A. Greiner, "A generic architecture for on-chip packet-switched interconnections," in *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pp. 250-256, 2000.
- [5] S. Li, L. S. Peh, and N. K. Jha, "Dynamic voltage scaling with links for power optimization of interconnection networks," in *Proceedings of the 9th International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 91-102, 2003.
- [6] S. Heo and K. Asanovic, "Replacing global wires with an on-chip network: a power analysis," in *Proceedings of the 2005 International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 369-374, 2005.
- [7] R. Kumar, V. Zyuban, and D. M. Tullsen, "Interconnections in multi-core architectures: understanding mechanisms, overheads and scaling," in *Proceedings of the 32nd International Symposium on Computer Architecture (ISCA)*, pp. 408-419, 2005.
- [8] W. Hangsheng, L. S. Peh, and S. Malik, "Power-driven design of router microarchitectures in on-chip networks," in *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 105-116, 2003.
- [9] R. Mullins, A. West, and S. Moore, "Low-latency virtual-channel routers for on-chip networks," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pp. 188-197, 2004.
- [10] L. Shang, L. S. Peh, A. Kumar, and N. K. Jha, "Thermal Modeling, Characterization and Management of On-Chip Networks," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, pp. 67-78, 2004.
- [11] R. Marculescu, "Networks-on-chip: the quest for on-chip fault-tolerant communication," in *Proceedings of the IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, pp. 8-12, 2003.
- [12] C. Xuning and L. S. Peh, "Leakage power modeling and optimization in interconnection networks," in *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 90-95, 2003.
- [13] T. T. Ye, L. Benini, and G. De Micheli, "Analysis of power consumption on switch fabrics in network routers," in *Proceedings of the 39th Design Automation Conference (DAC)*, pp. 524-529, 2002.
- [14] G. Varatkar and R. Marculescu, "Traffic analysis for on-chip networks design of multimedia applications," in *Proceedings of the 39th Design Automation Conference (DAC)*, pp. 795-800, 2002.
- [15] H. Jingcao and R. Marculescu, "Application-specific buffer space allocation for networks-on-chip router design," in *Proceedings of the IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pp. 354-361, 2004.
- [16] L. S. Peh and W. J. Dally, "A delay model for router microarchitectures," *IEEE Micro*, vol. 21, pp. 26-34, 2001.
- [17] W. J. Dally and C. L. Seitz, "The torus routing chip," *Journal of Distributed Computing*, vol. 1(3), pp. 187-196, 1986.
- [18] P. Kermani and L. Kleinrock, "Virtual cut-through: a new computer communication switching technique," *Computer Networks*, vol. 3(4), pp. 267-286, 1979.
- [19] W. J. Dally, "Virtual-channel flow control," in *Proceedings of the 17th Annual International Symposium on Computer Architecture (ISCA)*, pp. 60-68, 1990.
- [20] W. J. Dally and C. L. Seitz, "Deadlock-free message routing in multiprocessor interconnection networks," *IEEE Transactions on Computers*, vol. C-36(5), pp. 547-553, 1987.
- [21] Y. M. Boura and C. R. Das, "Performance analysis of buffering schemes in wormhole routers," *IEEE Transactions on Computers*, vol. 46, pp. 687-694, 1997.
- [22] M. Rezazad and H. Sarbazi-azad, "The effect of virtual channel organization on the performance of interconnection networks," in *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium*, 2005.
- [23] Y. Tamir and G. L. Frazier, "High-performance multiqueue buffers for VLSI communication switches," in *Proceedings of the 15th Annual International Symposium on Computer Architecture (ISCA)*, pp. 343-354, 1988.
- [24] G. L. Frazier and Y. Tamir, "The design and implementation of a multiqueue buffer for VLSI communication switches," in *Proceedings of the IEEE International Conference on Computer Design (ICCD)*, pp. 466-471, 1989.
- [25] J. Park, B. W. O'Krafka, S. Vassiliadis, and J. Delgado-Frias, "Design and evaluation of a DAMQ multiprocessor network with self-compacting buffers," in *Proceedings of Supercomputing*, pp. 713-722, 1994.
- [26] N. Ni, M. Pirvu, and L. Bhuyan, "Circular buffered switch design with wormhole routing and virtual channels," in *Proceedings of the International Conference on Computer Design (ICCD)*, pp. 466-473, 1998.
- [27] Y. Choi and T. M. Pinkston, "Evaluation of queue designs for true fully adaptive routers," in *Journal of Parallel and Distributed Computing*, vol. 64(5), pp. 606-616, 2004.
- [28] S. Konstantinidou and L. Snyder, "The Chaos router," *IEEE Transactions on Computers*, vol. 43, pp. 1386-1397, 1994.
- [29] M. Thottethodi, A. R. Lebeck, and S. S. Mukherjee, "BLAM: a high-performance routing algorithm for virtual cut-through networks," in *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 10 pp., 2003.
- [30] L. S. Peh and W. J. Dally, "A delay model and speculative architecture for pipelined routers," in *Proceedings of the 7th International Symposium on High Performance Computer Architecture (HPCA)*, pp. 255-266, 2001.
- [31] A. V. Yakovlev, A. M. Koelmans, and L. Lavagno, "High-level modeling and design of asynchronous interface logic," *IEEE Design & Test of Computers*, vol. 12, pp. 32-40, 1995.
- [32] H. Jingcao and R. Marculescu, "Energy- and performance-aware mapping for regular NoC architectures," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, pp. 551-562, 2005.
- [33] W. Hangsheng, L. S. Peh, and S. Malik, "A technology-aware and energy-oriented topology exploration for on-chip networks," in *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pp. 1238-1243 Vol. 2, 2005.
- [34] J. Duato, "A new theory of deadlock-free adaptive routing in wormhole networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 4, pp. 1320-1331, 1993.
- [35] S. Arjun, W. J. Dally, A. K. Gupta, and B. Towles, "GOAL: a load-balanced adaptive routing algorithm for torus networks," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pp. 194-205, 2003.