Design of an Interconnect Architecture and Signaling Technology for Parallelism in Communication

Jongsun Kim, *Member, IEEE*, Ingrid Verbauwhede, *Senior Member, IEEE*, and
Mau-Chung Frank Chang, *Fellow, IEEE*

Abstract—The need for efficient interconnect architectures beyond the conventional time-division multiplexing (TDM) protocol-based interconnects has been brought on by the continued increase of required communication bandwidth and concurrency of small-scale digital systems. To improve the overall system performance without increasing communication resources and complexity, this paper presents a cost-effective interconnect architecture, communication protocol, and signaling technology that exploits parallelism in board-level communication, resulting in shorter latency and higher concurrency on a shared bus or link: the proposed source synchronous CDMA interconnect (SS-CDMA-I) enables dual concurrent transactions on a single wire line as well as flexible input/output (I/O) reconfiguration. The SS-CDMA-I utilizes 2-bit orthogonal CDMA coding and a variation of source synchronous clocking for multilevel superposition; a single 3-level SSCDMA-I line operates as if it consists of dual virtual time-multiplexed interconnects, which exploits communication parallelism with a reduced number of pins, wires, and complexity. The unique multiple access capability of the SSCDMA-I improves real-time communication between multiple semiconductor intellectual property (IP) blocks on a shared link or bus by reducing the bus contention interference from simultaneous traffic requests and by taking advantage of shorter request latency. The prototype transceiver chip is implemented in 0.18- μm CMOS and the 10-cm test PC board system achieves an aggregate data rate of 2.5 Gb/s/pin between four off-chip (2Tx-to-2Rx) I/Os.

Index Terms—Bandwidth, bus, input/output (I/O) interface, inter-chip communication, interconnect, latency.

I. INTRODUCTION

GENERALLY, the digital system performance consists of two parts: computation performance and communication performance. Thanks to the development in the CMOS technology scaling, the computation performance of a chip has been increased dramatically. As computation performance goes up, the required communication bandwidth (or throughput) needs also increase with the same rate. However, the number of input/output (I/O) pins and the total I/O bandwidth of a single chip

Manuscript received June 30, 2006; revised January 27, 2007. This work was supported in part by the National Science Foundation (NSF), by the Semiconductor Research Corporation (SRC), and by the Defense Advanced Research Projects Agency (DARPA).

J. Kim was with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA. He is now with Samsung, Gyunggi-Do 463-779, South Korea (e-mail: jskim@ee.ucla.edu).

I. Verbauwhede is with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA and also with the Electrical Engineering Department, K. U. Leuven, B-3001 Leuven, Belgium (e-mail: ingrid@ee.ucla.edu; mfchang@ee.ucla.edu).

M.-C. F. Chang is with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA (e-mail: mfchang@ee.ucla.edu).

Digital Object Identifier 10.1109/TVLSI.2007.900739

Fig. 1. (a) Number of pins of a chip package. (b) Off-chip I/O speed (multidrop memory bus versus differential point-to-point link). (c) Sustained (or effective) bandwidth versus burst transfer size (= packet size) of a packet-based serial link at ideal peak bandwidth of 1 Gb/s/pin.

have been scaling much more slowly. Therefore, the communication between multiple on- or off-chip semiconductor intellectual property (IP) blocks is becoming a dominant cost, performance, and power factor in modern digital systems. The conventional interconnection architectures are under ever-increasing pressure to perform faster and more efficiently.

To make a simple, cost-effective and compact size digital system, techniques for sharing communication resources (pads, pins, and metal wires), such as using shared buses, have been an indispensable choice over the last 30 years. However, in modern pin-limited chips and systems, traditional communication resource sharing brings severe circuit-level overhead (limited bandwidth) as well as system- and architecture-level overhead (the typical shared bus architectures usually have scheduling and arbitration phase overhead). As shown in Fig. 1, the 2005 ITRS roadmap [1] predicts that the number of application-specific integrated circuit (ASIC) pins and microprocessor pins will increase 6% and 10% per year, respectively. It also predicts that the high-performance off-chip I/O speed will increase 25% per year up to over 18 GHz in 2013 by using differential point-to-point interconnects. On the other hand, the multipoint memory bus speed will remain at 1 GHz in 2013. This seems to

indicate that the available off-chip interconnect solution beyond multi-gigabits per second per pin will be high-speed serial links using point-to-point connection with advanced equalization schemes. However, the off-chip point-to-point I/O throughput of Fig. 1(b) is not the effective bandwidth but the theoretical peak bandwidth when the burst read/write data transfer size is ideally long (e.g., data streaming). In reality, as shown in Fig. 1(c), the sustained (or effective) off-chip data bandwidth of a packet-based serial link is severely decreased when the burst transfer size (or packet size) is short. For example, when the transferred data size is 4, 16, and 32 words (1 word = 8 bytes), the sustained data bandwidth is only 28.6%, 61.5%, and 76.2% of the ideal peak bandwidth, respectively. These results are obtained by the following simple equations assuming packet overhead = 10 bytes and extra latency = 0

$$\begin{aligned} &\text{Sustained (or effective) data bandwidth} \\ &= \frac{\text{data size}}{\text{total latency}} \end{aligned} \quad (1)$$

$$\begin{aligned} &\text{Total latency} \\ &= \left\lceil \frac{(\text{packet overhead} + \text{data size})}{\text{peak bandwidth}} \right\rceil + \text{extra latency} \end{aligned} \quad (2)$$

where packet size = packet overhead + data size. Effective bandwidth reduction is mainly due to the transmission of extra information (i.e., packet overhead) such as packet types, source/destination addresses, and error correcting (e.g., the typical one-way read packet overhead of RapidIO [2] is about 12 bytes).

In addition to this packet overhead, the sustained bandwidth will be further reduced due to the extra latency factors including arbitration (or scheduling) overhead, bus contention, interrupts, and initial setup delay. In particular, the memory subsystem latency (such as DRAM core access latency) is another key latency factor that decreases the sustained bandwidth in shared-memory systems. Also, the complicated equalization scheme of a serial link increases the latency of the transceiver circuits (Tx/Rx). It may introduce timing overhead when the Tx-to-Rx channel flight time is equivalent to or smaller than the Tx/Rx circuit latency. Therefore, even the high-speed serial links (mostly based on packet protocols) may have limited capability to increase system performance in latency-sensitive short-distance applications (such as memory access) due to the reduced effective data bandwidth and the increased network latency. Also, the back-to-back request latency problem still remains. These will be discussed more in Sections II and III.

To achieve higher computational performance at the system level, instead of just increasing processor clock speed, parallelism in computation is widely used for chip multiprocessing (CMP) and symmetric multiprocessor (SMP) system designs. These systems achieve computational concurrencies by using scheduling techniques such as multitasking or multithreading [3]. However, the on- and off-chip interconnect communication architectures are still mostly based on conventional time-division multiplexing (TDM)-based communication protocols [typically known as time-division multiple access (TDMA) or time-interleaving] that do not allow real communication parallelism. So, the key motivation of this paper is to explore how to exploit parallelism and concurrency in communication without

increasing resources and complexity. This paper suggests how the standard I/O interfaces should be modified for higher system performance with less cost. Using our results, designers of systems with multiple heterogeneous semiconductor-based processing devices will be able to take advantage of parallelism at both the computation and communication levels.

In order to realize this goal, this paper presents a new interconnect architecture and signaling technology called source synchronous CDMA interconnect (SSCDMA-I) that can enable latency-aware system design [4]–[6]. The single 3-level pulse-amplitude modulation (3-PAM) SSCDMA-I (bus or link) can operate as if it consists of two virtual TDM-based interconnects (TDM-I). So, dual on- or off-chip transceivers can occupy a shared physical wire interconnect simultaneously but they are separated from each other by using a set of 2-bit orthogonal codes. This means that two semiconductor IP cores can request and utilize the shared bus or link simultaneously without any arbitration phase or contention. By adopting this technology to system interconnect design, a new type of latency-aware system bus, processor, or memory bus architecture will be available.

In our previous work [5], a point-to-point serial link interconnect using self-synchronized clock-data recovery (CDR) circuits was introduced. It featured on-chip I/O reconfigurability between two serial link chips. The feasibility of reconfigurable memory systems was studied in [6]. Reconfigurable interconnect for next generation systems (RINGS) was introduced in [7]. This paper specifically focuses on CMOS technology-based synchronous, low-latency, bi-directional, point-to-point or multipoint I/O interface design for advanced parallelism in short range off-chip communications. It also discusses the detailed transceiver circuit design. The measurements show the capability of transmitting aggregate data rates of 2.5 Gb/s/pin with real-time I/O reconfigurability for dynamic bus channel allocation capability. The proposed 3-PAM SSCDMA-I improves system performance, especially in cases where real-time communication is required [4], [6].

This paper is organized as follows. Section II presents the overview of various on- and off-chip interconnect architectures, communication protocols, and signaling technologies. Section III discusses the issues for next generation high-performance interconnection system design. Section IV presents the architecture and protocol of the proposed 3-PAM SSCDMA-I, compares the bus transaction efficiency, and discusses signaling related issues. Section V describes detailed I/O transceiver design and gives experimental results. Section VI discusses conclusions.

II. OVERVIEW OF INTERCONNECT ARCHITECTURES, PROTOCOLS, AND SIGNALING TECHNOLOGIES

In this section, we briefly discuss various physical on- and (short-range) off-chip interconnect architectures which specifies the physical organization of the interconnection network; communication protocols which control network resource arbitration and utilization; and signaling technologies used in circuit level design. The proper choice of efficient interconnection network types, communication arbitration protocol and signaling technology is essential in meeting strict performance, power, and cost requirements of most digital systems today.

A. On-Chip Communication

On-chip communication has been traditionally based on half-duplex or full-duplex metal wire interconnects using either point-to-point links or shared buses. These conventional nonhierarchical structures are not efficient in system-on-chip devices with numerous semiconductor IP cores. In many cases today, hierarchical bus structures using multiple bridges are applied to build complex on-chip networks for SOC's with billions of transistors. This hierarchical approach increases bus utilization at the cost of the increased network latency as well as the complicated arbitration protocol overhead.

In this Section, we first discuss the communication protocols of on-chip interconnection architectures. There are basically three classes of on-chip communication or arbitration protocols to utilize the available channel bandwidth more efficiently: 1) static order scheduling; 2) (fixed or dynamic) time-driven scheduling; and 3) (static or dynamic) priority-driven scheduling [12]. All of these protocols do not allow multiple concurrent transactions on a shared bus or link with high-speed and low bit-error-rates. These arbitration schemes determine access rights to shared physical channels in serialized manners by the progression of time. This is because conventional signaling technologies cannot support real-time multiple access on a shared transmission medium simultaneously. Instead, they allow fixed or variable time slices to each transceiver device on a shared medium and supports fixed- or variable-block transfers and pipelined split transactions (typically known as interleaved bus). Therefore, all of the previous three traditional arbitration protocols inherently belong to the category of TDM that allows only a single transaction at a time. We classify all the previously mentioned arbitration protocols as a "TDM-based protocol" or "TDM protocol" and use this term as a reference in this paper. Generally, TDM or TDMA is the most simple and predictable protocol although it has low efficiency and long latency as well.

Here are two practical examples for widely used on-chip bus standards based on TDM protocols: advanced microcontroller bus architecture (AMBA) [18] and IBM CoreConnect [8]. AMBA, including AMBA high-speed-bus (AHB) and advanced system bus (ASB), is an established open standard for ARM processors and system-on-chip (SoC) design. Both AMBA and CoreConnect support split transactions (or pipelined buses) and burst transfers as well as multilayer architectures by using switches or bridges. Unlike the traditional interconnect architectures based on circuit switching, network-on-chip (NOC) is a new emerging approach based on packet switching by sending packets over the network to achieve higher scalability than ordinary buses and larger flexibility than ordinary point-to-point links [21], [22]. An NOC consists of multiple point-to-point links (based on TDM protocols) interconnected by packet-based switches or routers.

In the area of signaling technologies, wide parallel buses using rail-to-rail full swing CMOS signaling have been widely used as on-chip interconnects. However, as technology scaling continues today, especially in global cross-chip buses with repeaters, the excessive signaling power, energy, delay, and signal integrity problems are becoming a major bottleneck for CPU, memory, and SoC design. Performance degradation

coming from the limited metal wire interconnect bandwidth, wire delay, skew, cross-talk, and synchronization problems between numerous IP blocks on a single chip is critical. Although using copper (Cu) wires with low resistivity and low- κ (< 3.0) dielectric materials could be preferred physical solutions for today and the near term, the increased global wire delay and skew problems coming from continuous process scaling are impacting performance seriously in sub-100-nm technology generations. One of the reasons is that the resistivity of the smallest pitch copper global wire is expected to increase about 40% by 2010 [1]. Therefore, in order to reduce signaling power and energy dissipation and support the growing bandwidth demand, serial link-type differential low-swing (0.1–0.5 V) signaling is an unavoidable trend in high speed on-chip global communications [10], [19]. There are other research using on-chip transmission lines as interconnects to reduce power and delay as well as to increase bandwidth up to 8 Gb/s [20].

B. Short Range Off-Chip Communication

Most communications in short range (in-the-box) chip-to-chip interconnects (such as CPU-to-memory buses and backplanes) are based on the nonpreemptive TDM-based protocol used in traditional on-chip communication. Regardless of their different signaling schemes, these interconnect topologies usually adopt switched link or bus architectures. There are numerous types of electrical signaling technology today. To achieve higher network bandwidth, recent high-performance interconnection networks are performed by point-to-point link-based packet switch networks rather than traditional multipoint bus-based switch networks.

The recent high speed signaling standards are usually based on current- or voltage-mode 2-PAM binary (or NRZ) signaling [23]: high-speed transceiver logic (HSTL), stub-series terminated logic (SSTL), Gunning transceiver logic (GTL), Rambus signaling level (RSL), low-voltage differential signaling (LVDS), and current mode logic (CML). Usually, single-ended multipoint bus interface technologies tend not to support over 1 GHz due to their small noise margin issues from channel discontinuities, attenuation, and distortion. For example, memory interfaces-based on multipoint bus architecture have adopted LVTTTL (3.3 V) for SDRAM and SSTL₂ (2.5 V) for DDR, and SSTL₁₈ (1.8 V) for DDR2. DDR2 can operate up to a data rate of 800 Mb/s/pin. The graphic memory I/Os based on point-to-point link architecture adopt SSTL₂ for GDDR, SSTL₁₈ for GDDR2, and pseudo-open-drain logic (PODL) for GDDR3. GDDR3 now offers bandwidth of up to 64 GB/s with a data rate of 2 Gb/s/pin.¹ Recently, LVDS is widely used for many high-speed serial link applications. It is based on point-to-point current-mode differential signaling with a simple 100- Ω termination resistor at the receiver. It can be used for driving cables of up to 15 m with a data rate of typically less than 1 Gb/s/pair. LVDS families include flat panel display links (FPD-Links), LVDS display interfaces (LDIs), digital crosspoint switches, and SCI processor interconnect. CML is also based on point-to-point differential signaling with parallel on-chip input and output terminations (50 Ω) at the

¹[Online]. Available: <http://www.samsung.com/Products/Semiconductor/GraphicsMemory/index.html>

TABLE I
COMPARISON OF EMERGING HIGH-SPEED SYSTEM INTERCONNECT STANDARDS WITH THE PROPOSED SSCDMA-I

	PClexpress	Hyper Transport	RapidIO	InfiniBand	SSCDMA-I
Signaling (structure)	2-PAM differential LVDS (point-to-point link)	2-PAM differential LVDS (point-to-point link)	2-PAM differential LVDS (point-to-point link)	2-PAM differential LVDS (point-to-point link)	3-PAM single ended (multi-point bus or point-to-point link)
Application	Chip-to-chip, Board-to-board, out-of-box	Universal chip-to-chip	Chip-to-chip, board-to-board via backplane	Clusters, Data-center fabric	Universal Chip-to-chip
Channel length (media)	Inches to meters (PCB + cable)	Inches (PCB)	Inches (PCB, fiber, cable)	Meters to km (PCB, fiber, cable)	Inches (PCB)
Data rate (Gbps)	2.5	1.6	2.0 (source syn.) 3.125 (self syn.)	2.5	2.5
Clocking	Self-synchronized	Source synchronous	Serial: self-syn. Parallel: source syn.	Self-synchronized	Serial: self-syn. Parallel: source syn.
Encoding	8B10B	8B10B	8B10B	8B10B	2-bit orthogonal

near and far end of the link, which minimizes the number of external device for channel termination and biasing. Due to its simplicity and high-speed capability, CML is popular today for networking physical layer (PHY), and high-speed SerDes applications at a data rate of about 3 Gb/s or up to 10 Gb/s (e.g., XAUI: 10-Gb attachment unit interface¹) at the cost of complicated feed-forward or decision-feedback equalization circuit overhead.

The system interconnects are also in a transition status. The peripheral component interconnect (PCI) bus has been the most successful local interconnect in multipoint and parallel bus implementation for attaching peripheral devices to a PC motherboard. (PCI: 32- or 64-bit width 33 MHz, peak 133, or 266 MB/s, PCI-X: 133 MHz with peak 1066 MB/s). Also, PC cluster (a group of loosely coupled computers) systems, which are connected through system area networks such as Myricom Myrinet and Dolphin PCI-SCI [25], are commonly based on PCI or PCI-X. In order to achieve higher communication bandwidth and design flexibility, other new intra-system communication standards such as PCI Express [26], HyperTransport [27], and RapidIO [28] are all based on packet switched serial/parallel point-to-point link architectures. PCI Express (supported by Intel) is now becoming a new standard in backplane transceivers and local I/Os for personal computers [26]. HT (originated by AMD) is usually used to PC processor to chipset connections as well as game consoles [27]. RapidIO is an intra-system interconnect targeting for high-performance embedded equipment market [23], [28]. On the other hand, InfiniBand [29] is an inter-system interconnect targeting for switched fabric server I/Os for the data center, clustering and internet computing environments. These new techniques require the use of switching devices capable of interchanging time frames for split transaction. Thus these have additional latency overhead and complexity in control and clock synchronization. All of these four new standards adopt basically the same LVDS technology as their signaling scheme as well as the 8B10B encoding technique to reduce inter-symbol interference (ISI) and to achieve DC balanced signals for easy self-clock recovery. Table I shows a brief comparison of emerging high-speed system interconnect standards [17] and the proposed 3-PAM SSCDMA-I.

III. DESIGN OF ELECTRICAL INTERCONNECTS FOR HIGH PERFORMANCE

Depending on application and system requirements, there are basically three types of solutions for physical metal wire interconnection and communication among multiple IPs on a chip or a board. The first simple solution is to use dedicated “lines” to build a simple point-to-point link topology. This can be extended to build a mesh-, star-, or ring-type interconnection network by adopting circuit switching. This is, however, inefficient and wasteful when applied to large networks, since most interconnects of the links would be in idle mode. This is because, in circuit switching, a connection is preestablished prior to the start of transmission. Circuit switching allocates a dedicated circuit and line connection between every two nodes across a selected path. The second solution is to use a multipoint bus architecture based on time-multiplexed circuit switches. This is usually the easiest, cheapest, and most widely used interconnection network with limited length and bandwidth.

The third solution is to apply advanced switching technology to the previous point-to-point links, which is usually based on packet switching [9] rather than circuit switching to utilize bandwidth more efficiently. In packet switched networks, no connection is setup prior to the start of transmission of packets. Since each packet contains a header which identifies the packet’s source and target address information, a packet can be switched and routed in the network. The packet-based high-speed point-to-point serial link is a new trend to enable a high bandwidth interconnection network between multiple IPs. Multiple parallel serial links, sometimes referred to as a serial bus, are used to increase the bandwidth of a network. However, designers should notice that the packet-switched serial link or bus inherently has longer network latency, since each switch or router on the network begins to forward the packet only after the whole packet has been received.

There are some additional variants of switched interconnection networks. For instance, the crossbar switch is a well-known example of space-division circuit switching. Time-division circuit switching (based on typical TDM protocol) has been widely used in short distance communications (such as buses). Packet switching can be applied to multidrop buses. Also,

packet switching is now applied to global inter-cluster communication for NOCs [10] and for many emerging high-speed serial I/O standards such as PCIe, HT, and RapidIO. All of the previous conventional switching-based bus or link architectures adopt the TDM protocol.

Traditional TDM-based on-chip and chip-to-chip (in-the-box) interconnects are a cost and performance bottleneck in many digital applications today. This is because the channel bandwidth and latency of typical links and buses are major limiting factors for system performance. This is particularly true where multiple processors communicate with a main memory through a shared memory bus. In a typical multipoint bus with multiple heterogeneous network client devices (e.g., processors, memories, DSPs, and other IP cores), two transmitters cannot transmit during the same time period since the access to a bus is controlled by the progression of time (as the name TDM or TDMA signifies). Therefore, some method of arbitration or scheduling for interleaved bus operations is required. In these conventional TDM- or TDMA-based interconnect (TDM-I or TDMA-I) architectures, the requirements of client-devices' communication resource (i.e., bus or link) sharing are fixed and nonconcurrent: they cannot handle multiple real-time data flows on a shared interconnect simultaneously. Also, the communication resource allocation usually takes additional request latency due to scheduling, bus contention, and queuing delays on a shared bus or a link. Therefore, if a shared memory multiprocessor system utilizes the same conventional TDM-based memory bus that was used for a single core system, this system may show degradation in performance as the number of processing cores increases.

In addition, as already shown in Fig. 1(c), the ideal theoretical peak bandwidth of a bus or link can be achieved only when the burst transfer size of a message or a packet is ideally long enough [11]. The typical bus or link is a two-way communications medium and an ideally long read/write packet without interrupt is usually unacceptable. Thus, the sustained bandwidth is usually much less than the ideal peak bandwidth. This means that although the moving trend in interconnect architectures, from parallel multipoint I/Os to serial high-speed point-to-point link structures, is unavoidable, the benefit of high theoretical peak bandwidth of a serial link [see Fig. 1(b)] does not occur in practice. This means the communication latency is as important as the communication bandwidth for small-scale digital system performance. Depending on the application, the packet or message size of the interconnection network and the network system's bandwidth and latency sensitivity should be taken into account. However, current research in interconnect system design has only focused on bandwidth rather than latency.

Therefore, as system design methodology moves from the traditional computation-centric design to communication-aware design [12], instead of just increasing the communication bandwidth there is great demand for improving the communication efficiency without increasing resource, cost, and complexity. Supporting flexibility and concurrency in handling numerous data transactions simultaneously to reduce communication latency is a key concern. Our proposed SSCDMA-I reduces system overhead due to contention latency and provides concurrency not found in traditional interconnect networks.

In addition to the previous architecture-level issues, there are other signaling technology related issues (including ISI, crosstalk, simultaneous switching Ldi/Dt noise, excessive I/O signaling energy, and EMI) in high-speed interconnect design that limit communication performance. Frequency-dependent channel losses such as skin effect and dielectrical loss as well as dispersion, reflection, and attenuation on the channel also should be considered. Unequalized signaling systems, which are usually used for multipoint buses, can only tolerate about a 3-dB loss. Equalized signaling systems using preemphasis transmitters or feedback equalization receivers can tolerate up to a 10- to 20-dB loss. Timing uncertainty issues such as jitter, skew, and synchronization problems are also critical. All of these signaling related issues may incur communication errors and increase bit error rate (BER) due to the limited signal-to-noise ratio (SNR). For chip-to-chip communication applications, the reasonable value of BER is about 10^{-15} (at least less than 10^{-12}). A higher BER may necessitate heavy protocol overhead to deal with retransmission and error corrections. This extra protocol overhead is hard to implement in latency-sensitive short-distance applications.

IV. PROPOSED 3-PAM SSCDMA INTERCONNECT (SSCDMA-I)

In traditional PCs, the performance of memory bus subsystems can be increased by widening the bus channels or increasing the bus speed. These bandwidth-aware approaches, however, increase cost and complexity. Furthermore, communication latency and concurrency problems still exist and thus limit the overall system performance. Using multiple channels and a small burst size usually improves concurrency at the expense of throughput. In interconnect systems using packet-based protocols, the packet overhead and guard time between packets in adjacent time slots become more significant as time slots become smaller. If these configurations utilize a high-speed narrow bus or link which requires larger time slots or packet size, they usually suffer from longer channel request latency [13]. Here, the channel request latency is the amount of delay time that is due to stalling for memory or system bus resource allocations. Although the typical split transaction bus or pipelined bus can offer higher bandwidth by using packets, it usually has higher latency than nonsplit transaction buses [11]. On the other hand, the nonsplit transaction buses using fixed-assignment TDM protocol is not efficient for bursty or sporadic traffic, since the bus is not released for the next request until the last word is transferred [14].

Therefore, the traditional interconnect systems (mostly based on typical TDM protocol-based buses or links) are bandwidth-aware rather than latency-aware, which inherently have long request latency for concurrent or back-to-back requests due to the bus contention and queuing delays: the high sustained or effective data bandwidth (or throughput) is rarely achieved. However, in the proposed SSCDMA-I bus or link, two traffic requests can be processed and delivered simultaneously through a shared interconnect. This is a latency-aware, as well as bandwidth-aware, design approach that can enable parallelism in communication and, therefore, reduce the concurrent request latency overhead effectively.

Fig. 2. Bus architectures for simultaneous transferring two separate data packets $[D0(t), D1(t)]$. (a) Conventional TDM-I bus with two channels. (b) Conventional split-transaction TDM-I bus using (de)multiplexing with one channel. (c) Proposed 3-PAM SSCDMA-I bus with one channel.

A. Comparison of Bus Transaction Efficiency and Communication Protocol

In this section, we describe the SSCDMA-I protocol and compare the transaction efficiency of the high-speed serial/parallel bus architectures through examples shown in Figs. 2 and 3. Between the two important performance components of bandwidth (or throughput) and latency, we focus on communication latency and assume the bandwidth per wire line is the same for all cases. We also assume that packet protocols are used here, since packet-level data transfer is more efficient than bit- or byte-level transfer for achieving a high sustained channel bandwidth. Generally, communication latency (typically round-trip latency) is defined as the amount of time it takes to complete a request over an interconnection network: the time from the source device sending a request packet to the destination device to the source device receiving a first part of the response packet [11]. If this first part of the response packet is the critical word, then this latency is called round-trip critical word latency ($t_{r_critical}$). The unidirectional critical word latency $t_{critical}$, which is like a half-part of $t_{r_critical}$, is defined for a reply of a read request. End-to-end latency is defined as the time between sending the source packet to the receiving of the whole response packet. Both the transmission time t_{trans} (which is equal to the size of the transmitted packet or message divided by the communication bandwidth) and the flight time t_f (caused by the propagation delay through an interconnect) are part of the total latency component. A larger packet will take longer t_{trans} than

Fig. 3. Concurrent transaction for transmitting two data packets $[D0(t), D1(t)]$. (a) conventional TDM-I bus using two channels. (b) TDM-I split-transaction bus using one channel. (c) Proposed 3-PAM SSCDMA-I bus using one channel.

a short packet. In this paper, we define the “bus contention latency (t_{buscon})” as the additional delay time caused by bus con-

Fig. 4. (a) Comparison of critical word latency (back-to-back requests in 8-bit wide TDM-I and SSCDMA-I memory buses with $N = 4$). (b) Burst length (or minimum packet size) versus critical word latency (neglecting t_{RAC} for simplicity).

tention or bus saturation. This is the time that is caused by stalls in the interconnection network while waiting for the release of a shared bus or link. This is crucial in system design where processor cores share the same main memory and I/Os through a common interconnect. Therefore, the unidirectional communication latency $t_{u_latency}$ and the unidirectional critical word latency $t_{critical}$ of a typical memory (DRAM) access using a conventional TDM-based memory bus can be defined as follows:

Single read request:

$$t_{u_latency} = t_{initial} + t_f + t_{trans} + t_{DRAM} \quad (3)$$

Single read request:

$$t_{critical} = t_{initial} + t_f + t_{min} + t_{DRAM} \quad (4)$$

Back-to-back read request:

$$t_{critical} = t_{initial} + t_f + t_{min} + t_{trans_prev} + t_{DRAM} \quad (5)$$

Back-to-back read request:

$$t_{buscon} = t_{min} + t_{trans_prev} \quad (6)$$

where $t_{initial}$ is the initial delay time for bus arbitration (reply for a request) and initial setup time for transmission, t_{min} is the minimum size of a packet (usually one word), t_{trans_prev} is the previous packet's transmission time, and t_{DRAM} is the DRAM core access time.

Though the bus contention latency t_{buscon} must be taken into account for realistic latency and bandwidth calculation and performance estimation, it is often ignored. However, depending on the application, the bus contention latency may be much higher than other latency factors. In conventional TDM-I buses, the bus contention latency, which increases the unidirectional critical

word latency of a back-to-back read request, may decrease the performance of latency-sensitive digital systems. For instance, there are two types of requests that affect system performance when processors are accessing a shared memory. The first type is the latency-sensitive request, which requires that the critical word arrives as soon as possible. This means that it does not need an excessive amount of data at each instance and throughput is a secondary concern. Message passing between processors or shared memory access in a multicore processor system is an example of this kind of request. The second type is the bandwidth-sensitive request that requires minimum end-to-end latency to receive the whole requested packets. The impact of this bus contention latency will be discussed more in Figs. 3 and 4.

Suppose the IP block 0 (IP0), consisting of two processing units (cores A and B), is connected to the IP1 (with cores C and D) via a shared interconnect such as a PCB bus or a serial link. As shown in Figs. 2(a) and 3(a), two interconnect channels are required for two core devices (core A and B) to transmit two separate packets ($D0(t)$ and $D1(t)$) of IP0 simultaneously in a conventional TDM-I architecture using typical binary signaling (i.e., 2-PAM or 2-level NRZ using square pulses). It is also possible to transfer those two packets through a single channel by using split transaction or time multiplexing packet switching. This technique allows the bus to transfer $D0(t)$ and $D1(t)$ packet in back-to-back pipelined bus cycles by using a multiplexer (mux) and a demultiplexer (demux) as shown in Figs. 2(b) and 3(b). However, this split transaction bus technique introduces bus contention latency t_{buscon} and, therefore, increases the channel request latency $Ta1$ ($= 5$ time frames) of Fig. 3(b). $Ta1$ is defined as the delay experience when fetching

the critical word B1 of data $D1(t)$ when block transfers (e.g., minimum burst length = 4 time frames) are required. Thus, $Ta1$ is same as the definition of bus contention latency t_{buscon} of (6). So, when there are back-to-back requests, the amount of the added bus contention latency is proportional to the previous packet size or t_{trans_prev} . Trying to reduce $Ta1$ by reducing the packet size is inefficient due to additional overhead for headers and trailers in the packets to distinguish information and tolerate timing uncertainty. Although the latency $Ta1$ may be reduced somewhat by increasing the bus speed, it brings extra cost and circuit overhead and the bus contention latency still remains. This illustrates the inherent limits of conventional TDM-I architectures using traditional protocol and signaling schemes in dealing with multiple concurrent data transfers. It cannot meet the need to increase the number of IPs that can communicate with each other simultaneously in real-time communication systems. Overcoming this problem usually introduces the overhead of inserting a hierarchy of bus structures that introduce cross-hierarchical bus latencies or adding extra dedicated lines.

Therefore, the design of cost-effective on- and off-chip system interconnects that can allow multiple concurrent data transactions (without increasing resources) and change the nature of traditional bus request protocol is essential [4], [6]. As shown in Figs. 2(c) and 3(c), the proposed 3-PAM SSCDMA-I bus allows the four devices (Tx0/Tx1 of core A/B to Rx0/Rx1 of core C/D) to access the shared interconnect simultaneously: two packets are transmitted on a single wire trace concurrently. This means that a single SSCDMA-I line consists of dual virtual TDM-I channels capable of operating simultaneously but isolated by each other. The communication protocol of each virtual SSCDMA channel is like a TDM protocol. As shown in Fig. 3(c) of SSCDMA-I bus, although each transmitter's effective data transmission time doubles due to 2-bit orthogonal coding, the aggregate data rate per wire line is same as that of conventional TDM-I bus shown in Fig. 3(b). On the other hand, the request latency of $Ta2$ (= 2 time frames) is 60% smaller than the $Ta1$ (= 5 time frames) shown in Fig. 3(b). This is because the SSCDMA-I has no bus contention latency for every two requests. The request latency of packet B is reduced dramatically at the cost of increasing the transmission window of packet A. The unidirectional critical word latency $t_{critical}$ and the bus contention latency t_{buscon} of a SSCDMA-based memory bus can be defined as follows.

SSCDMA-I back-to-back read request

$$t_{critical} = t_{initial} + t_f + 2^*t_{min} + t_{DRAM}. \quad (7)$$

SSCDMA-I back-to-back read request

$$t_{buscon} = 2^*t_{min} \quad (8)$$

where the $t_{buscon} = Ta2$ and t_{min} is same as that of a TDM-I memory bus.

To compare the critical word latencies, Fig. 4(a) illustrates a typical shared memory access scenario in a single- or multicore system. The bidirectional memory bus implementations based on a 3-PAM SSCDMA-I and a conventional 2-PAM TDM-I are compared. Both memory busses have the same number of data and address lines with the same aggregate bandwidth per wire

line. In a conventional TDM-I bus, if two read requests (address 1 and 2) arrive at memory at the same time or back-to-back, the second request must stall until the first request finishes using the shard interconnect. If we assume that the bus arbitration process takes zero processor cycles, the transmission of an 8 byte cache line takes 4 cycles through the TDM bus and 8 cycles through the SSCDMA bus, respectively. The memory will take $t_{DRAM} = t_{RAC} = 8$ cycles to read data from the memory banks. [For Fig. 4(b), we assume $t_{RAC} = 0$ and $t_{initial} = 0$ for simplicity.]

So, in the case of an 8-bit wide TDM-I memory bus shown in Fig. 4(a), when the DRAM tries to transmit a data packet of one word (= 8 bytes) size, it takes four clock cycles to transfer each 8-byte packet by using dual edges of the clock. However, this split transaction technique increases the channel request latency of the second request further if block transfers (e.g., minimum burst length > 1 word) are required. For example, if the packet size is not one word but N words, the latency experienced by request 2 is correspondingly increased to 4^*N clock cycles.

As shown in Fig. 4, if we assume the first word of each packet is the critical word and the packet size (or minimum burst length) $N = 4$, the TDM-I bus takes a latency of $Ta1 = 20$ clock cycles (= $4^*(N + 1) = 20$, where the number of data channel $C = 8$) to fetch the critical word in addition to the common DRAM core access latency $t_{DRAM} = t_{RAC}$ of eight cycles. On the other hand, in the case of SSCDMA-I memory bus, it has much smaller fixed request latency of $Ta2 = 8$ clock cycles. It is 60% faster than the conventional TDM-I bus to fetch the same critical word. Fig. 4(b) shows the burst length versus critical word latency of the second request when data bus channel width C is 8- and 16-bit wide. As the channel becomes narrower (16-bit $Ta1'$ to 8-bit $Ta1$), the critical word latency $Ta1$ of TDM-I increases more rapidly. As the burst length N of the data packet increases, the gap between the latency $Ta1$ of TDM-I and $Ta2$ of SSCDMA-I is increased. This example illustrates that a typical TDM-I memory bus has a much longer critical word latency than the SSCDMA-I memory bus. As a result, the unique dual concurrent transaction characteristic of the SSCDMA-I bus makes it possible to achieve communication concurrency as well as reduced back-to-back request critical word latency.

In addition, with shared memory multicore SMPs as an example, the SSCDMA-I memory bus enables advanced critical-word first fetch inherently and allows each core to continue execution without waiting for the full cache line block to be loaded. This is essential since the CPU core usually requires only one word (e.g., 8 bytes) of the full cache line block at a time and memory access is mostly read dominant [11].

B. Dual Concurrent Transactions for Advanced Parallelism and Concurrency in Communication

Fig. 5 shows the detailed simultaneous multiple access operation for dual concurrent bus transactions by using a single 3-PAM SSCDMA-I between two off-chip slaves (Tx0 at point A and Tx1 at point B) and a master (Rx0 and Rx1 at point C) as an example. To apply parallelism in board-level communications, 2-bit orthogonal CDMA coding is used for encoding and decoding of the base-band data. Source synchronous clocking

Fig. 5. Simultaneous multiple access for two concurrent bus transactions on a 3-PAM SSCDMA-I bus. (a) Interconnect system block diagram. (b) Detailed operation between two slaves (Tx0, Tx1) and a master (Rx0, Rx1).

is used to enable 3-level superposition with a proper timing by delay-locked loops (DLLs). Matched resistive terminations are assumed to be used on both ends of the transmission line and a separate clock line that travels in two directions (clock to/from master) on the board is also used. Each transmitter and receiver is assigned a 2-bit orthogonal code, which is a sequence of numbers called chips. By using the 2-bit Walsh codes $C0$ (for Tx0 and Rx0) or $C1$ (for Tx1 and Rx1), the encoded data output $cd0(t) (= D0(t) \oplus C0)$ is driven by the output driver of Tx0 and transferred to a 2-level signal (primed $cd0(t)'$) with a small swing of $V_s/2$ and then transmitted through the channel at point A.

In Fig. 5(b), if the Tx0 attempts to send binary data by using $C0 = [1 \ 1]$, it sends output signal with a swing of 0 or $V_s/2$. After the channel flight time delay of $tf1$, the transmitted signal $cd0(t)'$ arrives at point B and Tx1 sends out $cd1(t)'$ signal (similarly encoded by $C1 = [1 \ 0]$) with swing of 0 or $V_s/2$ at that time. Superposition of the two 2-level traveling-wave signals ($cd0(t)'$ + $cd1(t)'$) results in maximum swing of about V_s ($= V_s/2 + V_s/2$) on the transmission line. The master receives the signal by source synchronous clocking which removes the board-level skew between the clock and data and enables multi-level superposition without distortion. This ensures that Walsh codes are perfectly orthogonal to each other during transmission, since all the transmitted off-chip signals are in perfect synchronization. Thus, if $C0$ and $C1$ are exactly orthogonal, then the decoder output of Rx0 is described as follows:

$$[(D0(t) \oplus C0) + (D1(t) \oplus C1)] \bullet C0 = D0(t) \oplus C0. \quad (9)$$

Fig. 6. (a) Transmitter circuit Tx. (b) Transmitter timing diagram.

The inner product $C0 \bullet C1$ (which is represented as a dot operator by multiplying two code sequences, chip-by-chip, and adding the results) is zero. This means that the modulated spreading signals can coexist on a single line without causing interference, which enables a new way of dual-concurrent transaction signaling. At the master point C, the transmitted 3-level analog signals are digitized by the ADCs and then despread by the same orthogonal codes $C0$ or $C1$. The two decoded outputs, X and Y , are integrated by current integrators and generate differential small swing signals X' for even and Y' for odd data. After the sense-amp operation, the original data packets $D0(t)$ and $D1(t)$ are recovered. Detailed transceiver circuit design, orthogonal coding, and the reconfigurable I/O operation are explained in Section V. We can also apply this SSCDMA-I to simultaneous bidirectional signaling (SBD) on a point-to-point link, assuming that the length of the physical transmission channel is well matched with the transmission cycle time (or bandwidth) to avoid eye distortion which is common in all conventional SBD schemes [15].

C. Signaling Issues and Comparison

Current- or voltage-mode 2-PAM binary signaling has been used most widely in applications for high-speed on- and off-chip communications due to its simplicity, high noise immunity, and reliability. Four-PAM multilevel signaling has also been introduced, since its lower baud rate can double the eye width and also reduce frequency-dependent channel losses without increasing bus speed [16]. However, for a given power, 4-PAM is more susceptible to channel interference noise such as crosstalk, since the multiple signal levels are more closely spaced than 2-PAM, resulting in higher BER. Also, the effective eye width is further reduced due to the transitions between signal levels that are not adjacent. In any case, neither 2- nor 4-PAM are reconfigurable and simultaneous multiple access is not allowed.

Fig. 7. (a) Receiver circuit Rx. (b) Timing diagram. (c) Current integrator circuit.

In addition, concurrent bus transaction without adding separate dedicated lines is impossible.

From a signaling point of view, the 3-PAM SSCDMA-I can be regarded as a compromise between the conventional 2- and 4-PAM TDM-I schemes. Table II compares the output signal waveforms (or symbols) of typical 2- and 4-PAM TDM drivers and the 3-PAM SSCDMA-I driver, transmitting 2 bits of data through a single line at a time. We assume that all of the previous three schemes use current-mode incident wave signaling with matched channel terminations. 2-PAM TDM-I uses one reference voltage (V_{ref}) and two series data eyes with swing amplitude of V_s at frequency of f_{clk} . 4-PAM TDM-I uses three reference voltages (V_{refh} , V_{refm} , V_{refl}) and three parallel stacked eyes of $V_s/3$ at half frequency $f_{clk}/2$. 3-PAM SSCDMA-I utilizes two reference voltages (V_{refh} , V_{refl}) and four series and parallel eyes of $V_s/2$ at f_{clk} . Comparisons reveal a tradeoff between noise and timing margin. Compared with 2-PAM, 3-PAM SSCDMA has the same timing margin and weaker noise margin. Compared with 4-PAM, 3-PAM SSCDMA is stronger in noise margin due to the smaller number of levels but weaker in timing margin since it requires bandwidth expansion factor of two for orthogonal coding. By assuming the use of source synchronous clocking for each scheme, timing uncertainty issues such as clock jitter and skew become a common problem for all cases.

The signaling power of the previous three current mode drivers has no relation to channel frequency. Therefore, as shown in Table II, the two typical buses using 2- and 4-PAM require N -bit wide separate lines for concurrent N -bit bus transactions with a signaling power dissipation of $N \cdot V_s \cdot I$ regardless of peak bandwidth values (I is the average current

TABLE II
SIGNALING COMPARISON OF 2- AND
4-PAM TDM-I AND THE PROPOSED 3-PAM SSCDMA-I

	(a) 2-PAM TDM-I	(b) 4-PAM TDM-I	(c) 3-PAM SSCDMA-I
Signal Waveform (for 2-bit transmitting)			
Signal Levels	2	4	3
Number of Eyes	2	3	4
Channel Swing	V_s	$V_s/3$	$V_s/2$
# of wires for N-bit concurrent bus transactions	N	N	$N/2$
Channel Frequency	f_{clk}	$f_{clk}/2$	f_{clk}
Signaling power for N-bit concurrent bus transactions	$N \cdot V_s \cdot I$	$N \cdot V_s \cdot I$	$N/2 \cdot V_s \cdot I$
Multiple Access Capability	No	No	Yes
Reconfigurable I/O operation	No	No	Yes

of the output drivers). However, the 3-PAM SSCDMA bus requires only $N/2$ -bit wide lines for the same bus transactions and thus the total channel signaling power ($= N/2 \cdot V_s \cdot I$) can be reduced by up to 50% due to the fewer number of I/O lines. We assume that all the previous output drivers use the same maximum channel swing of V_s (although the noise margin is different) and the effect of the predrivers of the output drivers is negligible.

Fig. 8. Simulated 3-PAM SSCDMA-I bus with real-time I/O reconfigurability at an aggregate data rate of 2.0 G/s/pin.

V. I/O TRANSCIEVER ARCHITECTURE

A. Circuit Design

Unlike wireless CDMA transceivers, the 3-PAM SSCDMA-I transceiver does not require a complicated up- and down-conversion process, which enables less design complexity and lower power consumption. Fig. 6 shows the transmitter circuit (Tx) and its timing diagram. The encoder is a simple XOR gate and the output driver is a typical current mode open-drain structure. The encoder outputs ($cd_{odd}(t)$, $cd_{even}(t)$) occupy twice the bandwidth than the base-band signal $D0(t)$

$$cd_{odd}(t) = \{D0(t)_{odd} \oplus C0[1]\} \quad (10)$$

$$cd_{even}(t) = \{D0(t)_{even} \oplus C0[0]\}. \quad (11)$$

The encoder transforms a 1-bit base-band data “1” into a 2-bit coded signal “0 1” as shown in Fig. 6(b) by using orthogonal codewords of two digits ($C = [1 \ 0]$ in this example). The modulated data $cd0(t) (= cd_{odd}(t) + cd_{even}(t))$ drives the output driver with the dual edge of a differential clock ($clk/clkb$ from transmit DLL), creating the small swing 2-level output signal $cd0(t)'$. In this timing diagram, we assume the gate delay is zero for simplicity. In a case when another transmitter is activated with a different code C1, the two 2-level small swing signals ($cd0(t)'$ and $cd1(t)'$) form a 3-level signal waveform on

the channel as depicted in Fig. 5(b). Although these spreading signals seem to overlap, they are orthogonal and linearly independent in the time domain. $C0[1 : 0]$ and $C1[1 : 0]$ are the 2-bit orthogonal Walsh codes, $[1 \ 1]$ and $[1 \ 0]$, respectively, (0 is used instead of -1 for CMOS circuit implementation).

Fig. 7(a) shows the receiver circuit Rx, which consists of two 3-level (or 1.5-bit) interleaving ADCs with two dc references, a decoder, two current integrators, and two sense amplifier based flip-flops (SAs). In order to recover the two separate data signal ($D0(t)$, $D1(t)$) from the composite spectrum, the ADC converts the 3-level signal into the thermometer code and then the decoder de-spreads the signal by using the same 2-bit orthogonal codes as was illustrated in Fig. 5(b) with (9). Since the inner product $C0 \bullet C1$ is logically zero, the base-band data can be recovered easily by choosing the proper orthogonal codes, $C0$ for $D0(t)$ and $C1$ for $D1(t)$. The decoder consists of only four parallel XOR gates and two multiplexers, which minimizes the circuit overhead and internal delay. As shown in the timing diagram of Fig. 7(b), when the decoder output (X , Y) has more 1's than 0's in two integration periods (equals to one cycle), the integrator generates 1. Gate delays are assumed to be zero in this figure for simplicity. Instead of using a digital comparator, the analog current integrator, shown in Fig. 7(c), is used to reduce the recovering delay and remove the glitch errors. It determines the value of the decoded signal at every two integration periods

Fig. 9. Measured 3-PAM SSCDMA-I signaling with four eye diagram for an aggregate data rate of 2.5 Gb/s/pin.

with differential output. The ADCs and decoder are synchronized with the receive DLL clock (clk/clkb) and the integrators use the half frequency clock (hclk/hclkb). If we implement the two receivers (Rx0 and Rx1) in the same chip, the two 3-level interleaving ADCs can be shared.

B. Experimental Results

Fig. 8 shows the simulated simultaneous multiple access operation of the 2 Gb/s/pin 3-PAM SSCDMA-I bus for 2Tx-to-2Rx dual concurrent bus transactions through the 10-cm PCB trace shown in Fig. 5(b) with $t_{f1} = 0.4$ ns and $t_{f2} = 0.6$ ns. Tx0 is transmitting $D0(t)$ with a data rate of 1 Gb/s and Tx1 is transmitting $D1(t)$ with the same data rate. Point C shows the transmitted 3-level signal waveforms at the receiver side channel. Then the decoder outputs (X , Y) are integrated for each two time periods. By simply changing the orthogonal codes (e.g., $C0$ to $C1$ for Rx0 and $C1$ to $C0$ for Rx1) at the reconfiguring point, the outputs of Rx0 and Rx1 are changed from $D0(t)$ to $D1(t)$ and $D1(t)$ to $D0(t)$, simultaneously. This reconfigurable I/O operation is achieved without latency. This means the configuration of interconnection between I/Os on a shared bus can be changed by choosing different orthogonal codes in real time. This software-based I/O channel reconfiguration enables better bus utilization.

Fig. 9 shows the measured 3-PAM SSCDMA-I 4-eye diagram on the PCB test board for an aggregate data rate of 2.5 Gb/s/pin. The four-chip test board system consists of two transmitters and two receivers on a 10-cm long 50- Ω FR4 PC board. An eye opening of more than 100 mV was measured. The transceiver, fabricated in 0.18- μ m CMOS, active area is $340 \times 160 \mu\text{m}^2$. With a 1.8-V supply, the average power consumption of the two transmitters (each with a swing of 300 mV) and two receivers are 9.6 and 23.4 mW, respectively. The two 50- Ω terminators dissipate 7.2 mW for typical data patterns with balanced 1's and 0's. Table III summarizes the performance of the 3-PAM SSCDMA-I transceiver chipset. Fig. 10 shows the second transceiver chip microphotograph fabricated in 0.10- μ m CMOS Samsung DRAM process.

TABLE III
PERFORMANCE SUMMARY

Supply Voltage	1.8 V
Process	0.18- μ m CMOS
Aggregate Data Rate	2.5 Gb/s/pin with four (2Tx-to-2Rx) chip simultaneous multiple access
Channel Interface Frequency	1.25 GHz
Signaling Level	3-level ($V_s = 600\text{mV}$)
Channel Termination Power	7.2 mW
Transmitter (x 2ea) Power	9.6 mW
Receiver (x 2ea) Power	23.4 mW
Active Transceiver Area	$340 \times 160 \mu\text{m}^2$

Fig. 10. Transceiver chip microphotograph using 0.10- μ m CMOS Samsung DRAM process.

VI. CONCLUSION

To alleviate heavy traffic load in traditional TDM-based bus or link architectures, this paper presented a latency-aware interconnect architecture, protocol, and circuit techniques to exploit parallelism in communication. By utilizing 2-bit orthogonal CDMA coding techniques and source synchronous clocking for multilevel superposition, the proposed SSCDMA-I (bus or link) enables dual concurrent transactions on a shared interconnect without bus contention. The single 3-PAM SSCDMA-I operates as if it consists of dual virtual TDM-I channels operating individually without interference; This achieves lower back-to-back request latency, higher concurrency and bus utilization, and lower arbitration overhead than typical TDM-I buses. The interconnect architectures, communication protocols, bus transaction efficiency, and signaling issues have been compared with various conventional TDM-I schemes. The prototype transceiver chip simulated and fabricated in 0.18- and 0.10- μ m CMOS technologies and tested in a 10-cm 50- Ω PC board achieves an aggregate data rate of 2.5 Gb/s/pin and allows real-time I/O reconfigurability between four (2Tx-to-2Rx) off-chip I/Os.

ACKNOWLEDGMENT

The authors would like to thank Samsung, Gyeonggi-Do, South Korea, for fabricating the chip.

REFERENCES

- [1] Int. Technology Roadmap for Semiconductors, "ITRS roadmap," (2005). [Online]. Available: <http://www.itrs.net/Common/2005ITRS/Home2005.htm>

- [2] S. Fuller, *RapidIO: The Embedded System Interconnect*. New York: Wiley, 2005.
- [3] W. Stallings, *Computer Organization and Architecture; Designing for Performance*. Englewood Cliffs, NJ: Prentice-Hall, 2003.
- [4] J. Kim, Z. Xu, and M. F. Chang, "A 2-Gbps/pin source synchronous CDMA bus interface with simultaneous multi-chip access and reconfigurable I/O capability," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, 2003, pp. 317–320.
- [5] Z. Xu, H. Shin, J. Kim, M. F. Chang, and C. Chien, "Giga bit/s CDMA interconnect transceiver chip-set with multilevel signal data recovery for re-configurable VLSI system," in *IEEE Int. Solid-State Circuits Conf. Tech. Dig.*, 2003, pp. 322–323.
- [6] J. Kim, Z. Xu, and M. F. Chang, "Re-configurable memory bus systems using multi-Gbps/pin CDMA I/O transceivers," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2003, pp. II-33–36.
- [7] I. Verbaauwhede and M.-C. F. Chang, "Reconfigurable interconnect for next generation systems," in *Proc. ACM Int. Workshop Syst. Level Interconnect Prediction (SLIP)*, 2002, pp. 71–74.
- [8] IBM, Hopewell Junction, NY, "IBM CoreConnect bus architecture," (1999). [Online]. Available: <http://www-03.ibm.com/chips/products/coreconnect/>
- [9] B. A. Forouzan, *Data Communications and Networking*. New York: McGraw-Hill, 2004.
- [10] K. Lee, S. Lee, and H. Yoo, "Low-power network-on-chip for high-performance SoC design," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 14, no. 2, pp. 148–160, Feb. 2006.
- [11] J. L. Hennessy and D. A. Patterson, *Computer Architecture; A Quantitative Approach*. San Mateo, CA: Morgan Kaufmann, 2003.
- [12] A. A. Jerraya and W. Wolf, *Multiprocessor Systems-on-Chips*. San Mateo, CA: Morgan Kaufmann, 2005.
- [13] V. Cuppu and B. Jacob, "Concurrency, latency, or system overhead: Which has the largest impact on uniprocessor DRAM-system performance," in *Proc. 28th Annu. Int. Symp. Comput. Arch.*, 2001, pp. 62–71.
- [14] B. Sklar, *Digital Communications; Fundamentals and Applications*. Englewood Cliffs, NJ: Prentice-Hall, 2001.
- [15] T. Takahashi, T. Muto, Y. Shirai, F. Shirotori, Y. Takada, A. Yamagiwa, A. Nishida, A. Hotta, and T. Kiyuna, "110-GB simultaneous bidirectional transceiver logic synchronized with a system clock," *IEEE J. Solid-State Circuits*, vol. 34, no. 11, pp. 1526–1532, Nov. 1999.
- [16] J. Zerbe, P. Chau, C. Werner, T. Thrush, H. Liaw, B. Garlepp, and K. Donnelly, "1.6 Gb/s/pin 4-PAM signaling and circuits for a multidrop bus," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 752–760, May 2001.
- [17] HyperTransport Consortium, "Testing HyperTransport backgrounder," (2001). [Online]. Available: http://www.hypertransport.org/docs/dg/HT_Backgrounder_final_v8.pdf
- [18] ARM, "AMBA Specification Rev.2," (1999). [Online]. Available: http://www.arm.com/products/solutions/AMBA_Spec.html
- [19] R. Ho, K. Mai, and M. Horowitz, "Efficient on-chip global interconnects," in *Proc. Symp. VLSI Circuits*, 2003, pp. 271–274.
- [20] A. P. Jose, G. Patounakis, and K. L. Shepard, "Near speed-of-light on-chip interconnects using pulsed current-mode signaling," in *Proc. Symp. VLSI Circuits*, 2005, pp. 108–111.
- [21] W. Dally *et al.*, "Route packets, not wires: On-chip interconnection networks," in *Proc. Design Autom. Conf.*, 2001, pp. 684–689.
- [22] D. Wiklund, S. Sathe, and D. Liu, "Benchmarking of on-chip interconnection networks," in *Proc. Int. Conf. Microelectron.*, 2004, pp. 621–624.
- [23] T. Granberg, *Digital Techniques for High-Speed Design*. Englewood Cliffs, NJ: Prentice-Hall, 2004.
- [24] Myricom Corp., Arcadia, CA, "Myricom homepage," (1994). [Online]. Available: <http://www.myricom/>
- [25] Dolphin Corp., Marlborough, MA, "PCI SCI-64 adapter card," (1996). [Online]. Available: http://www.dolphinics.no/dics_pci_sci64.htm
- [26] A. Bhatt, "Creating a third generation I/O interconnect," (2004). [Online]. Available: <http://www.intel.com/technology/pciexpress/devnet/docs/WhatIsPCIExpress.pdf>
- [27] HyperTransport Consortium, "HyperTransport consortium," (2006). [Online]. Available: <http://www.hypertransport.org/tech/index.cfm>
- [28] RapidIO, "RapidIO technology overview and applications," (2005). [Online]. Available: http://www.rapidio.org/education/documents/RapidIO_Overview-Apps_v07.pdf
- [29] InfiniBand, Beaverton, OR, "InfiniBand architecture overview," (2002). [Online]. Available: <http://www.infinibandta.org/itinfo/>

Jongsun Kim (S'01–M'07) received the B.S. degree in electronics engineering from Hanyang University, Seoul, South Korea, in 1992, the M.S. degree in electrical and computer engineering from Pohang University of Science and Technology (POSTECH), Pohang, North Gyeongsang, South Korea, in 1994, and the Ph.D. degree in electrical engineering from the University of California, Los Angeles (UCLA), in 2006.

Since March 2007, he has been with Samsung Electronics, Gyeonggi-Do, Korea. From 1994 to 2001, he was with Samsung Electronics, where he worked on the design and development of Synchronous DRAM, SGRAM, Rambus DRAM, and other specialty DRAMs. He was a Postdoctoral Fellow at UCLA from 2006 to 2007. His research interests include the area of mixed-signal integrated circuits and systems design. His current research area includes high-speed signaling and clocking circuits for next-generation DRAM, low-power memory, and advanced I/O interconnect circuits and systems for chip-to-chip communications.

Ingrid Verbaauwhede (M'93–SM'00) received the Ph.D. degree from the Katholieke Universiteit Leuven (K.U. Leuven), Leuven, Belgium, in 1991.

She is currently a Professor and co-director of the Computer Security and Industrial Cryptography (COSIC) Laboratory with the Electrical Engineering Department, K.U. Leuven and an Adjunct Associate Professor with the Department of Electrical Engineering, University of California (UCLA), Los Angeles.

She was a Lecturer and Visiting Research Engineer with University of California, Berkeley, from 1992 to 1994. From 1994 to 1998, she was a Principal Engineer first with TCSI, Berkeley, CA, and then with Atmel, Berkeley, CA. She was an Associate Professor with UCLA in 1998 and with K.U. Leuven in 2003. She has been active in many conferences, where she was the Program Chair and the General Chair for ISLPED in 2002 and 2003, respectively. She was a member of the executive committee of the 42nd and 43rd DAC as the design community chair. She is a member of the executive committee of DATE 2007. She is the Program Chair of the 2007 Cryptographic Hardware and Embedded Systems (CHES) Conference. For more information on her research, see www.emsec.ee.ucla.edu or www.esat.kuleuven.be/cosic.

Mau-Chung Frank Chang (F'96) received the B.S. degree in physics from the National Taiwan University, Taipei City, Taiwan, R.O.C., in 1972, the M.S. degree in material science from the National Tsing-Hua University, Hsinchu City, Taiwan, R.O.C., in 1974, and the Ph.D. degree in electrical engineering from the National Chiao-Tung University, Hsinchu City, Taiwan, R.O.C., in 1979.

He is currently a Full Professor with the Electrical Engineering Department and Director of the High Speed Electronics Laboratory, University of California (UCLA), Los Angeles. Before joining UCLA, (1983–1997) he was the Assistant Director and Department Manager of the High Speed Electronics Laboratory, Rockwell Science Center, Thousand Oaks, CA. In this tenure, he successfully developed and transferred the AlGaAs/GaAs HBT technology from the research laboratory to the production line (Conexant Systems). The HBT production has now grown into a multibillion dollar business worldwide. During his career, his research work has been mostly in the development of high-speed semiconductor devices and integrated circuits for mixed signal communication and sensor system applications. He was the principal investigator at Rockwell to lead the US-DARPA ADC and DAC development for direct conversion transceiver (DCT) and digital radar receivers (DRR) systems. He was the inventor of the multi-I/O, re-configurable RF/wireless interconnects based on FDMA and CDMA multiple access algorithms for inter- and intra-ULSI communications. He was also pioneered in demonstrating the first cantilever-type RF MEMS switch for microwave/millimeter-wave transmission applications. Recently, he has led the first demonstration of a dual mode (CDMA/AMPS) power amplifier based on the Si/SiGe HBT technology for wireless handset transmitter applications; the first K-band (24–27 GHz) RFIC in CMOS and the first CMOS ADC with 1 GHz instantaneous bandwidth. He has authored or coauthored over 200 technical papers, 10 book chapters,

edited 1 book, and holds 18 U.S. patents. Since 1997, Dr. Chang has founded a communication chip-design company (G-Plus, Westwood, CA, recently acquired by SST) and served as the Board Director in two GaAs manufacturing companies, including GCS (Torrance, CA) and GCT (Hsinchu, Taiwan). Both GCS and GCT are regarded as world-class pure-player foundry houses for microwave and millimeter-wave integrated circuits (MMIC) production.

Dr. Chang is a co-editor of the IEEE TRANSACTIONS ON ELECTRON DEVICES and was a guest editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, in 1991

and 1992, respectively, and the *Journal of High-Speed Electronics and Systems* in 1994. He was a recipient of the Rockwell's Leonardo Da Vinci (Engineer of the Year) Award in 1992, the National Chiao-Tung University's Distinguished Alumni Award in 1997, and the National Tsing-Hua University's Distinguished Alumni Award in 2002. He was named an IEEE Fellow in 1996 for his pioneering contributions in ultra-high-speed HBT integrated circuit development, and was honored with the IEEE David Sarnoff Award in 2006.