Hybrid Wire-Surface Wave Architecture for One-to-Many Communication in Networks-on-Chip

Ammar Karkar[1,2], Nizar Dahir[1,2], Ra'ed Al-Dujaily[2], Kenneth Tong[3], Terrence Mak[4], and Alex Yakovlev[1]

[1]School of Electrical and Electronic Engineering, Newcastle University, UK, Email: {a.j.m.karkar, nizar.dahir, alex.yakovlev}@newcastle.ac.uk
[2]IT Research Centre, University of Kufa, Iraq, Email: {ammar.karkar,nizar.dahir, raed.aldujaily}@uokufa.edu.iq
[3]Department of Electrical and Electronic Engineering, UCL, London, UK, Email: K.tong@ucl.ac.uk
[4]Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, Email: stmak@cse.cuhk.edu.hk

***Abstract*—Network-on-chip (NoC) is a communication paradigm that has emerged to tackle different on-chip challenges and has satisfied different demands in terms of high performance and economical interconnect implementation. However, merely metal based NoC pursuit offers limited scalability with the relentless technology scaling, especially in one-to-many (1-to-M) communication. To meet the scalability demand, this paper proposes a new hybrid architecture empowered by both metal interconnects and Zenneck surface wave interconnects (SWI). This architecture, in conjunction with newly proposed routing and global arbitration schemes, avoids overloading the NoC and alleviates traffic hotspots compared to the trend of handling 1-to-M traffic as unicast. This work addresses the system level challenges for intra chip multicasting. Evaluation results, based on a cycle-accurate simulation and hardware description, demonstrate the effectiveness of the proposed architecture in terms of power reduction ratio of 4 to 12X and average delay reduction of 25X or more, compared to a regular NoC. These results are achieved with negligible hardware overheads.**

I. Introduction

Integrated circuit technology processes, forced by growing market demands, are scaling rapidly and are causing an intensification of current and future System-on-Chip (SoC) in terms of transistor density and functional complexity. As a result, the number of integrated intellectual property (IPs) cores inside a single SoC has increased dramatically and led the research community [1], [2] and industry [3] to adopt NoC as the underlying communication structure. This is especially true for chip multiprocessors (CMPs) that were introduced to provide near linear performance over complexity improvements (PollackâĂŹs rule), while maintaining lower power and frequency budget [4]. CMP performance and power consumption depend both on NoC and cache coherence protocols. These protocols rely highly on an underlying communication fabric to provide one-to-many (1-to-M) communication.

Related work in NoC struggled to achieve 1-to-M latency and energy close to wire-latency and wire-energy [5], [6]. This will not be sufficient in the near future with the projected issues in regular metal based NoCs since it struggles to match the needed scalability, especially for global communication in terms of latency and energy (J/b) [7]. These challenges inspired many studies to look for alternative communication fabrics such as radio frequency interconnects (RF-I) [8] and optical interconnects [9]. These interconnects seem to be far from ideal due to complexity, power consumption and/or area overhead [7]. Zenneck surface wave (SW) is an emerging interconnect that could be the optimum solution for mitigating

978-3-9815370-2-4/DATE14/ © 2014 EDAA

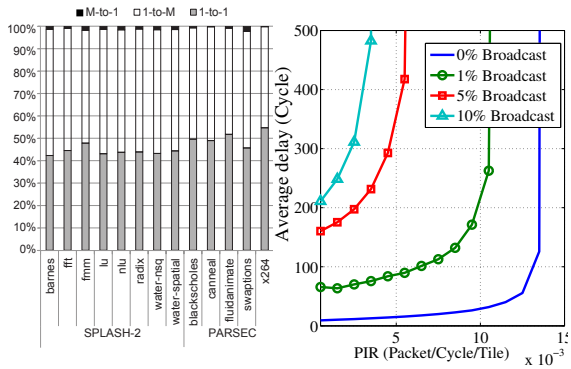

(a) (b)

Fig. 1: (a) 1-to-M, M-to-1 and 1-to-1 traffic percentage in different CMP real benchmark when broadcast-based cache coherence protocol is used [5], (b) NoCs simulation shows an increase in average delay and fast network saturation when broadcast treated as unicast.

global 1-to-M communication issues. This technology is an electromagnetic wave that propagates and is guided through an interface between different media surfaces. The major contributions of this paper are:

- Introduces a hybrid wire-SW interconnects architecture that exploit SW remarkable features for 1-to-M.
- Proposes routing and arbitration mechanisms for 1-to-M traffic that maximize the hybrid architecture utilization.
- Evaluates rigorously the proposed architecture which found to surpass the previous work.

This paper is organized as follows: the next section highlights the motivation of this work. Sec.III presents a hybrid wire-surface wave interconnect architecture. Sec.IV proposes a broadcasting routing scheme and global arbitration for the proposed architecture. Sec.V shows area overheads and demonstrates and discusses the performance and power results obtained from the developed cycle accurate simulation. Sec.VI presents the conclusion and draws directions for future work.

II. Motivation

Cache coherence protocols depend on a range of 1-to-M communication patterns with variant percentages such as multicasting invalidation requests (Directory-based protocols) and broadcasting ordering tokens (Broadcast-based protocols) [6], [5]. Cache coherence broadcast-based protocols (such as token coherence) offer less hardware overhead (directory storage that scale with number of cores) and low latency unlike other cache coherence protocols [5]. However, in these protocols the broadcasting ratio to the total packet injection

rate (PIR) considered to be very high (average 14.3% [5]). For instance, Fig.1a shows 1-to-M high ratio for the Token Coherence protocol for a range of applications [5]. This could be a nightmare for global coherence unless interconnects fabric supports the 1-to-M communication. Nonetheless, NoCs conventionally treat 1-to-M traffic patterns as repeated unicast traffic. This basic handling will have a dramatic effect on the NoC due to the following reasons: (1) 1-to-M increases congestion on the source node of this traffic (router, network interface and links) and thus creates a bottleneck, (2) causes a poor QoS due to the queueing of repeated unicast packets on the same communication fabric, (3) increases power consumption due to retransmitting the same data, but to different destinations. Thus, even small percentage of this traffic will reflect badly on NoC performance and cost, as shown in Fig.1b.

Consequently, the trend in cache coherence protocols design is to mitigate, but unable to eliminate, 1-to-M communication. As a result, Broadcast-based protocols or any other promising solutions that require high ratio of 1-to-M and/or large multicasting destination groups are avoided. In this study we try to eliminate these constrains through proposing an interconnect architecture empowered by emerging surface wave technology.

III. Hybrid Wire-SW interconnect Architecture

Hybrid interconnect architecture (that combine shared medium bus and the indirect network) shown previously limited scalability in performance critical SoC due to the limitation of the metal-wire based bus layer [2]. Hybrid network architecture could retain the broadcasting capability of the buses and reduce inter-node average hop count while maintaining high interconnect scalability if high performance interconnect such as SWI is adopted as bus system [2].

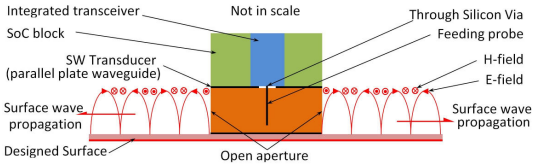

Fig. 2: Shows the need for integrated transceiver and integrated transducer (inverted quarter-wavelength monopole) stacked over the designed surface.

A. SWI and Fanout Feature

SWI is an emerging interconnect that is introduced for low power, plane and 1-to-M communication pattern. This technology, in conjunction with experiment results, has been presented in [10], [11] and the excellent scalability and performance features of SWI for on-chip global communication along with on-chip implementation consideration have been discussed in [12], [13]. Thus, this section offers a brief description of SWI and then focuses on merits of SWI in 1-to-M communication. The Zenneck surface wave is an inhomogeneous plane electromagnetic wave supported by a surface. The designed surface medium consists of a dielectric layer placed over a corrugated conductor layer [10], [11]. The surface can be engineered by altering its dimensions and materials of conductor and/or dielectric used to achieve the needed characteristic impedance (Z_0). Moreover, a dielectric-coated metal flat surface has been developed successfully using a dielectric martial overlaid with conductor sheet. The designed surface can be implemented off-chip to reduce die area overhead. A transceiver with FDMA designed for wave guided signal is chosen [8] to transfer electrical signal into RF signal. Also, an integration of a transducer linked to the transceiver is needed to launch the waved signal into the surface [10] as shown in Fig.2. More details about surface wave propagation can be found in [10], [11].

SWI interconnect offers natural fanout features. For instance, according to [11], the E-field decay rate in Zenneck surface wave from the source horizontally along the boundary should be around $(1/\sqrt{d})$, where d is the distance from the source. On the other hand, vertically the decay is exponential away from the boundary. This allows less power dissipation for far larger coverage areas than the regular wireless RF [14]. Moreover, SWI electromagnetic signal propagates (in all direction of the surface) in speed close to the speed of light. This makes SWI able to fanout the signal cross the chip in one clock cycle with competitive power consumption and circuit complexity over other emerging interconnects. For example, Optical interconnects (in addition to complexity, compatibility and power issues [9]) need extra hardware to offer fanout feature such as splitters and combiners that decay the optical signal. On the other hand, RF transmission lines forking requires stubs (impedance discontinuity) that would cause lower propagation velocity and signal reflections unless careful matching circuit designed at each drop point [15]. Therefore, to avoid using tree of transmission lines many designs proposed a worm or cycle layout of this thick wires to pass through all the nodes which would add nontrivial area overhead, signal decay and latency. This make SWI favourable supplementary interconnect for applications depend highly on 1-to-M, such as future CMP with cache coherence protocols.

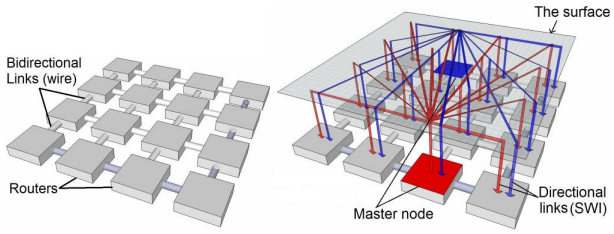

(a) Basline-Arch (b) Hybrid-Arch

Fig. 3: Illustration example showing that inserting two SWI channels in the proposed hybrid wire-SWI multilayer-network increases the overall NoC bisection.(a) conventional mesh NoC, (b) connections of both layers metal wire and SWI.

B. The Proposed Architecture

SWI has significant advantages over many interconnects fabrics as mentioned earlier. However, as all RF-based interconnects, it suffers from limitations in terms of shared media and limited ranges of frequencies which make it infeasible to replace metal wire interconnect completely in the near future [14]. Moreover, wire-based local communication seems to scale well with technology scaling unlike wire-based global communication [7]. In addition, this interconnects has the cheapest implementation cost compared to other fabrics. Therefore, the best solution is to combine both interconnects, metal and surface wave, in one hybrid wire-SW interconnects

(multi-layered network) architecture (in short Hybrid-Arch), as shown in Fig.3. The first layer is regular K-ary 2-mesh topology and the second layer is the surface wave bus topology where number of master is based on the number of different frequency channels available. Therefore, this topology offers natural fanout feature that has been lost when interconnects move from bus to NoC. Moreover, it substantially increases the network bisection [1]. Therefore, reliability of the overall system is improved by increasing the network bisection, which makes it more robust to failure that may isolate part of the NoC. As a result, this architecture will increase the network capacity to higher PIR and offer higher fault tolerance.

In order to preserve fanout feature, all the routers in the NoC can receive information through SWI. However, fewer nodes have the transmission capability to achieve lower area and circuit overhead and also remain constrained by the limited frequency bandwidth. These nodes will be referred as master nodes, while the rest are referred as slave. Master nodes are distributed so that average hop count from all slaves to the nearest master node is minimum. In this paper Single-Chip Cloud Computer (SCC) [3] is adopted as the baseline architecture (Baseline-Arch). However, in Hybrid-Arch a sixth port needs to be added to the router with all related control circuits. Also crossbar switch size needs to be adjusted according to the new added port. This port is linked to a transceiver with either Tx/Rx capability or just Rx.

IV. Hybrid-Arch 1-to-M Routing and Arbitration

Routing and arbitration techniques are required to direct and deliver 1-to-M traffic to its final destinations with minimum cost and better utilization of the proposed architecture. Previous related work can be classified into two classes: Path-based and Tree-based. In Path-based, a 1-to-M packet is submitted to members of a multicast group sequentially [16]. This method shows variant and very significant delivering latency. On the other hand, in Tree-based, 1-to-M traffic is propagated in an embedded tree route in the network and forked only at the disjoint points in the path [5], [6]. This way it alleviates traffic load and reduces energy consumption for interconnect fabric. Nonetheless, this traffic pattern still magnifies issues of the regular wire metal NoC as it increases the load and hotspots.

The routing scheme for the proposed architecture is an improved Tree-based where the embedded tree path forks at one point (nearest master). From this, traffic flow is delivered to all the leaves (slaves) in one hop. Thus, it provides higher efficiency in handling 1-to-M traffic since: (1) each node simply needs to be aware of the nearest master and directs the 1-to-M traffic using any partially adaptive routing algorithm. Hence, there is no need for extra circuit or complicated algorithms to build the multicasting tree path; (2) packets will be replicated only at the destination routers. This will reduce power consumption by eliminating the need of duplicated traffic to travel through costly intermediate resources such as wires and routers. In order to direct the packet to the multicast group members, each multicast packet header must have Multicast-Address-Bits (MAB). This header field is a bit vector where each bit is representing a node and it is set if the node is a multicasting member. In addition, it will be used along with the source ID to identify a 1-to-M request. This request is sent to the global arbiter.

The proposed global arbitration scheme achieves the best legal match in two cycles (given that requested resources are free) and cosist of three stages. First stage (1) is a request masking to validate if master request is possible for any of the VSWI by comparing its MAB with already reserved resources in reservation table. The next stags represent the lonely output allocator [1] that achieve best legal match requests-to-VSWI by minimising conflict between requests over VSWIs. Then at final stage, the winner request (the oldest for each VSWI) will be selected and stored in a reservation table. The size of this table is proportional to NoC size, master nodes number (SWI channels) and the VSWI numbers. The results of the last stage will be transmitted as a grant signal through SWI specific frequency channel. More detailed about the design consideration and the structure of this arbiter will be presented in our future work.

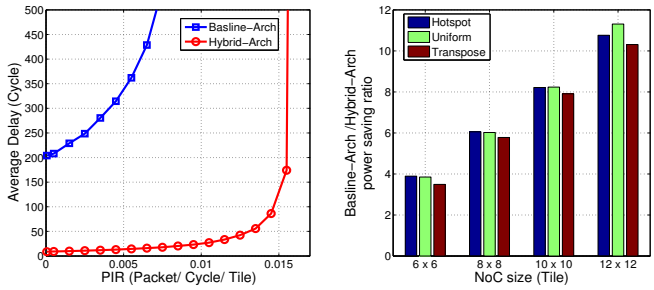

(a) (b)

Fig. 4: System level evaluation results: (a) Average delay Baseline-Arch, Hybrid-Arch1 of NoC (6×4), (b) Communication power saving ratio of the Hybrid-Arch over the Baseline-Arch for different network sizes and traffic scenarios under broadcast ratio 10%.

V. System Level Evaluation and Discussion

A. Performance Improvements

This section presents results obtained from our cycle accurate NoC simulator which is built by modifying the existing Noxim simulator [17] for the Hybrid-Arch and the Baseline-Arch. The packet size of 12 flit is chosen to demonstrate the behaviour of the proposed architecture under a wormhole flow control. The number of master nodes is based on the available frequency range for $45nm$, which is estimated to be 4 channels (added to it the frequencies specified for control signals). However, this range is scaling with technology [8]. As a result, the number of master nodes is increased when simulating larger NoCs assuming the technology has been scaled too. Results show significant improvements in performance in terms of average delay. For instance, Fig.4a shows much less average delay for Hybrid-Arch over the Baseline-Arch even for zero-load-latency (around 25X less). Moreover, proposed architecture offers around twice (2X) the saturated PIR. This is simply due to the fact that Baseline-Arch is replicating the broadcast traffic to all destinations, thus increasing the load on the NoC. Also, packets no longer need to travel through routers and costly wire links if they use the SWI highway to reach all the broadcast destinations in one hop.

B. Power Reduction

This section presents evaluation results of an important cost metric for future interconnects and 1-to-M in particular, which is the power consumption. The router's static and dynamic power is calculated using the Orion 2.0 [18] model, and power dissipation for wire links is calculated for the X ($3.6mm$) and

the Y ($5.2mm$) directions. The transceiver (TxRx) power consumption projection [8] is used ($24mW$ per sub-channel). SWI power dissipation is also modulated based on the analytical model introduced in [12]. In addition, the global arbiter was designed using Verilog and then synthesized using Synopsys Design Compiler and mapped onto the PDK 45nm technology library to calculate its dynamic power ($4.8mW$) and leakage power ($59.3\mu W$). Then all these values are used in Noxim to calculate the overall NoC power consumption. Fig.4b shows power consumption saving ratios for the proposed Hybrid-Arch over the Baseline-Arch for different NoC sizes and traffics. It obviously demonstrates significant improvements in NoC power consumption reduction ratio. For instance, power ratio of Baseline-Arch to Hybrid-Arch starts from more than double (4X) and increases up to 12X as we increase the NoC size. These findings prove that Hybrid-Arch has a remarkable scalability feature and effectiveness in mitigating 1-to-M communication issues.

TABLE I: Area overhead evaluation for Hybrid-Arch, Baseline-Arch and RF-I with transmission lines (TL). Notice that some components are needed only in the master nodes

NoC component		Area per item (mm^2) for $45nm$ technology		
Component	No.	Baseline-Arch	Proposed Hybrid-Arch	RF-I with TL
Router	24	1.0853	1.5124	1.5124
Transmitter	4	-	0.1558	0.1558
Receiver	24	-	0.0083	0.0083
Gloable Arbiter	1	-	0.0552	0.0552
TL Link	1	-	-	0.7608
Wire Link	1	13.653	13.653	13.653
Total extra area over Baseline-Arch (all components × their No.)- Basline Arch area		-	11.13	11.9
Total die area(mm^2)		567.1		
NoC area (%)		7	8.96	9.1

C. Area Overheads Evaluation

It is essential to evaluate chip area overheads for the extra on-chip circuits required for our architecture. Firstly, for transceivers we assume that the active area calculated in [8] is the only part scaled down when we move to $45nm$ technology, while the passive parts remain almost the same since they are proportional to the channels' operational frequency range. Therefore, the projected transmitter area is ($4870\mu m^2$) per sub-channel, while the projected receiver area is ($260\mu m^2$) per sub-channel (the active area is proportional to the square of the scaling factor). Secondly, the extra router port (buffer, crossbar and related circuits) area is calculated using the Orion 2.0 [18] model to be ($0.427mm^2$). In addition, the global arbiter designed using Verilog and then synthesized using Synopsys Design Compiler and mapped onto the PDK $45nm$ technology library to calculate its area ($0.0114mm^2$), added to it the TxRx estimated area ($0.0438mm^2$). On the other hand, the value of RF-I's transmission line with a pitch of $12\mu m$ considered to be routed through the chip (NoC size 6×4) as (U) shape passing through all nodes to calculate its area [8]. Table I shows area considerations for the Baseline-Arch, proposed Hybrid-Arch and transmission line RF-I. Obviously, most of the extra area overhead for the Hybrid-Arch around (1.9%) is due to the extra router port, which could be less if smaller buffers and smaller number of VC's for each port were used. Moreover, Hybrid-Arch actually saves the area comparing to RF-I due to the need for transmission lines in high metal layer.

VI. Conclusion and Future Work

This paper proposes a hybrid wire-SWI architecture for on-chip communication that is able to tackle the 1-to-M traffic efficiently and gives an attractive area/performance trade-off. Zenneck surface wave low power dissipation, high signal propagation speed and fanout capability all help to significantly mitigate the 1-to-M communication issues that NoCs based CMP in particular suffer from. In addition, routing and arbitration techniques for this architecture and 1-to-M traffic have been proposed and discussed. Evaluation results show significant improvements in terms of average delay, saturated PIR and power consumption with a relatively small die area penalty compared to related work. Future work includes investigating many-to-one traffic patterns.

References

[1] W. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, 2004.

[2] J. Duato and *et al.*, *Interconnection networks [electronic resource]: an engineering approach*. Morgan Kaufmann, 2003.

[3] P. Salihundam and *et al.*, "A 2 tb/s 6 4 mesh network for a single-chip cloud computer with dvfs in 45 nm cmos," *Solid-State Circuits, IEEE Journal of*, vol. 46, pp. 757 –766, april 2011.

[4] S. Borkar, "Thousand core chips: a technology perspective," in *Proceedings of the 44th annual Design Automation Conference*, DAC '07, (New York, NY, USA), pp. 746–749, ACM, 2007.

[5] T. Krishna and *et al.*, "Towards the ideal on-chip fabric for 1-to-many and many-to-1 communication," in *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-44 '11, (New York, NY, USA), pp. 71–82, ACM, 2011.

[6] N. Jerger and *et al.*, "Virtual circuit tree multicasting: A case for on-chip hardware multicast support," in *Computer Architecture, 2008. ISCA '08. 35th International Symposium on*, pp. 229–240, 2008.

[7] Semiconductor Industry Association, "ITRS: International Technology Roadmap for Semiconductors ." http://www.itrs.net/reports.html [online], 2009.

[8] M.-C. Chang and *et al.*, "Power reduction of cmp communication networks via rf-interconnects," in *Microarchitecture. MICRO-41. 41st IEEE/ACM International Symposium on*, pp. 376 –387, nov. 2008.

[9] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proceedings of the IEEE*, vol. 97, pp. 1166 –1185, july 2009.

[10] J. Turner, M. Jessup, and K.-F. Tong, "A novel technique enabling the realisation of 60 GHz body area networks," in *Wearable and Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*, pp. 58 –62, may 2012.

[11] J. Hendry, "Isolation of the zenneck surface wave," in *Antennas and Propagation Conference (LAPC), 2010 Loughborough*, pp. 613 –616, nov. 2010.

[12] A. Karkar and *et al.*, "Surface wave communication system for on-chip and off-chip interconnects," in *Proceedings of the Fifth International Workshop on Network on Chip Architectures*, NoCArc '12, (New York, NY, USA), pp. 11–16, ACM, 2012.

[13] A. Karkar and *et al.*, "Hybrid wire-surface wave interconnects for next-generation networks-on-chip," *IET Computers and Digital Techniques*, vol. 7, pp. 294–303(9), November 2013.

[14] K. O and *et al.*, "The feasibility of on-chip interconnection using antennas," in *Computer-Aided Design, 2005. ICCAD-2005. IEEE/ACM International Conference on*, pp. 979 – 984, nov. 2005.

[15] W. J. Dally and J. W. Poulton, *Digital systems engineering*. New York, NY, USA: Cambridge University Press, 1998.

[16] X. Lin and *et al.*, "Deadlock-free multicast wormhole routing in 2-d mesh multicomputers," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 5, no. 8, pp. 793–804, 1994.

[17] F. Fazzino, M. Palesi, and D. Patti, "Noxim: Network-on-chip simulator."

[18] A. B. Kahng and *et al.*, "Orion 2.0: A power-area simulator for interconnection networks," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1 –5, 2011.