Mixed Wire and Surface-wave Communication Fabrics for Decentralized On-Chip Multicasting

Ammar Karkar^{1,2}, Kin-Fai Tong³, Terrence Mak⁴, and Alex Yakovlev¹

¹School of Electrical and Electronic Engineering, Newcastle University, UK, Email: {a.j.m.karkar, alex.yakovlev}@newcastle.ac.uk

²IT Research Centre, University of Kufa, Iraq, Email: ammar.karkar@uokufa.edu.iq

³Department of Electrical and Electronic Engineering, UCL, London, UK, Email: K.tong@ucl.ac.uk

⁴Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, Email: stmak@cse.cuhk.edu.hk

Abstract-Network-on-chip (NoC) has emerged to tackle different on-chip challenges and has satisfied different demands in terms of high performance, economical and reliable interconnect implementation. However, a merely metal-based interconnect reaches performance bound with the relentless technology scaling. Especially, it displayed a bottleneck to meet the communication bandwidth demand for multicasting. This paper proposes a novel hybrid architecture, which improves the on-chip communication bandwidth significantly using mixed wires and surface wave interconnects (SWI) fabrics. In particular, the bandwidth of multicasting can be drastically improved. We introduce a decentralized arbitration method to fully utilize the slack-time scheduling with deadlock-free flow control. Evaluation results, based on a cycleaccurate and hardware-based simulation, demonstrate the effectiveness of the proposed architecture and methods. Compared to a wire-based NoC, the mixed fabric approach can achieve an improvement in power reduction and communication speed up to 63% and 12X, respectively. These results are achieved with almost negligible hardware overheads. This new paradigm efficiently addresses the emerged challenges for on-chip communications.

I. INTRODUCTION

Due to growing market demands, integrated circuit technology processes are scaling rapidly, 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, leading the research community [1], [2] and industry [3] to adopt the NoC as the underlying communication structure. This is especially true for chip multiprocessors (CMPs) that were introduced to provide nearlinear performance improvements when complexity increased (Pollack's rule), while maintaining lower power and frequency budget [4]. CMP performance and power consumption depend both on NoCs and cache coherence protocols.

Cache coherence protocols depend on a range of one-tomany (1-to-M) communication patterns such as multicast invalidation requests (directory-based protocols) and broadcasting ordering tokens (broadcast-based protocols) [5], [6]. Cache coherence protocols (especially broadcast-based) produce relatively high multicast ratio over the total packet injection rate (PIR), up to 52.4% [6], [5]. For instance, Fig. 1a shows a high 1-to-M ratio for the token coherence protocol for a



Fig. 1: (a) 1-to-M, M-to-1 and 1-to-1 traffic percentage in different CMP real benchmark when a broadcast-based cache coherence protocol is used [6], (b) our 6×4 NoC simulation that shows fast NoC saturation is proportional to 1-to-M ratio in the overall traffics in case of software multicast.

range of applications [6]. This could be catastrophic for global coherence unless the interconnect fabric supports the 1-to-M communication. Therefore, the trend in cache coherence protocols design is to mitigate (because they are unable to eliminate) 1-to-M communication. As a result any promising solutions that require high ratio of 1-to-M and/or large multicast destination groups, are avoided [5], [6]. Therefore, there is a need to eliminate these constraints and improve performance by proposing a 1-to-M interconnect architecture.

NoCs designers conventionally treat 1-to-M traffic patterns as repeated unicast traffic (software multicast) [2]. This basic handling will have a dramatic effect on the NoC, for the following reasons: (1) 1-to-M increases congestion on the source node of this traffic and thus creates bottlenecks, (2) causes a poor quality of service (QoS) due to the queueing of repeated unicast packets on the same communication fabric, (3) power consumption is increased due to retransmitting the same data but to different destinations. As a result, even small percentage of 1-to-M traffic will have severe affects on NoC performance and cost, as shown in Fig. 1b.

Relevant studies in NoC have struggled to achieve 1-to-M latency and energy close to wire-latency and wire-energy [6], [5]. This will not be sufficient in the near future given the projected issues with 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 have inspired many researchers to look for alternative communication fabrics, such as radio frequency (RF) interconnects (RF-I) [8], wireless NoC (WiNoC) [9] and optical interconnects [10]. However, these interconnects

We would like to acknowledge the EPSRC who supported this research through grants UNCOVER (EP/K001698/1) and GAELS (EP/I038551/1). Also, we would like to acknowledge the Croucher Foundation Start-up Grant and Natural Science Foundation of China (No. 61376024) for supporting the research. Lastly, we would like to acknowledge the HCED financial support from Iraq government.

seem far from ideal due to their complexity, power consumption and/or area overheads [7]. Zenneck surface wave (SW) is an emerging interconnect that could be the optimum solution for mitigating global 1-to-M communication issues [11], [12], [13]. This technology is an electromagnetic wave that propagates and is guided through an interface between different media surfaces. A previous work suggested hybrid W-SWI architecture which was focusing on broadcasting (1to-all) using centralized arbitration [13]. However, this paper proposes a more efficient routing scheme that would improve the SWI layer utilization under multicast (1-to-M). The major contributions of this paper are:

- A novel hybrid wire-SW interconnects architecture with stretched multicast is proposed. It exploits the combination of wire and SW in the NoC architecture and maximizes the performance for on-chip multicast.
- A novel decentralized arbitration scheme is proposed which can maximize the slack-time scheduling in a multiple-source and multiple-destination scenario.
- It has been shown that the proposed arbitration and routing method is deadlock-free using newly proposed ID-tagging scheme.
- The proposed architecture is rigorously evaluated. Comparing to the previous approaches, it can achieve an improvement for power reduction and communication speed up to 63% and 12X, respectively.

II. MIXED WIRE-SW INTERCONNECT ARCHITECTURE

Mixed interconnect architectures (that combine shared medium bus and the indirect network such as mesh) showed previously limited scalability in performance critical SoC due to the limitation of the metal-based bus layer [2]. This 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].

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 [14], [15] and the excellent scalability and performance features of SWI for on-chip global communication, along with on-chip implementation considerations have been discussed in [11], [12]. Thus, this section offers a brief description of SWI and then focuses on its merits in 1-to-M communication. The SW 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 [14], [15]. The surface can be engineered by altering its dimensions, and the materials of the chosen conductor and/or dielectric to achieve the required characteristic impedance (Z_0) . Moreover, a dielectric-coated metal flat surface has been developed successfully using a dielectric material overlaid with conductor sheet. The designed surface can be implemented off-chip to reduce die area overhead. An already implemented transceiver with FDMA designed for a wave-guided RF signals is chosen [16]. Also, an integration of a transducer is needed to launch the waved signal into the surface [14], which can be fabricated separately and then using flip-chip bonding and the through-silicon-via (TSV) technique to connect it to the integrated transceiver, see Fig. 2. More details about SW propagation can be found in [14], [15].



Fig. 2: Integrated transceiver and integrated transducer (inverted quarter-wavelength monopole) stacked over the designed surface for generating on-chip surface wave [14].



Fig. 3: Illustrative example showing that inserting two SWI channels in the proposed hybrid wire-SWI multilayer-network can increases the overall NoC bisection. (a) Conventional mesh NoC, and (b) two layers of interconnections, including wire and SWI.

SWI interconnect offers natural fanout features. For instance, the E-field decay rate in SW from the source horizontally along the boundary should be around $(1/\sqrt{d})$, where d is the distance from the source [15]. On the other hand, vertically, the decay is exponential away from the boundary. This allows less power dissipation for far larger coverage areas (up to 10cm) than the regular WiNoC where it propagates the signal up to 23mm [9], [17] because wireless signal dissipation via antennas and free space. However, both WiNoC and the SWI signals transmits in all directions (over the surface for the SWI) at a speed close to the speed of light if we assume the WiNoC antenna radiation pattern is 360°. Thus, SWI can fanout the signal across the chip in one clock cycle with competitive levels of power consumption and circuit complexity compared to other emerging interconnects. For example, in addition to complexity, compatibility and power issues [10], optical interconnects need extra hardware to fanout the signal to different directions and receivers such as splitters and combiners that decay the optical signal. On the other hand, aside from the cross-talk, RF transmission lines (TL) forking requires stubs (impedance discontinuity) that would cause lower propagation velocity and signal reflections unless careful matching circuit designed at each drop point [18]. Therefore, to avoid using a tree of TLs many designs have proposed a worm or cycle layout of this thick wires to pass through all the nodes which would add nontrivial area overheads, signal decay and latency. Even for the WiNoC, the integrated antennas are considered an area hungry components [9], [17]. Therefore, SWI may be favoured supplementary interconnect for future CMPs that depend highly on 1-to-M.

B. The Proposed Architecture

The SWI has significant advantages over many interconnects fabrics, as mentioned earlier. However, as all RFbased interconnects, it suffers from limitations in terms of shared media and limited ranges of frequencies. These make it infeasible to completely replace metal wire interconnects in the near future [9]. Moreover, unlike wire-based global communication, wire-based local communication seems to scale well with technology scaling [7]. In addition, this type of interconnect has the cheapest implementation cost compared to other fabrics. Therefore, the best solution would be to combine both interconnects, metal and SW, in one hybrid wire-SW interconnects in a multi-layered network architecture; in short W-SWI, as shown in Fig. 3. The first layer is a regular K-ary 2-mesh topology, and the second layer is the surface wave bus topology where a number of nodes can transmit to all other nodes in the NoC. As a result, this architecture offers a natural fanout feature which is otherwise lost when interconnects move from the bus to the NoC. Moreover, it substantially increases the network bisection [1]. Therefore, the reliability of the overall system is improved by increasing the network bisection, which makes it more robust to any type of 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 the fanout feature, all the routers in the NoC can receive information through SWIs. 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 to as master nodes, while the rest are referred to as slave nodes. Master nodes are distributed so that average hop count from all slaves to the nearest master is at a minimum. In this paper the single-chip cloud computer (SCC) [3] is adopted as the regular mesh NoC baseline architecture (Mesh of short). However, in W-SWI architecture a sixth port needs to be added to the router with all related control circuits. Also the crossbar switch size needs to be adjusted according to the new added port. This port is linked to a transceiver with either Tx/Rx for master capability or just Rx for slave.

III. HYBRID-ARCH 1-TO-M ROUTING AND ARBITRATION A. Stretched Multicast

To improve SWI utilization and avoid centralized arbitration, this paper proposes a technique that allows a master to transmit a multicast flow to any available slaves. In other words, a master does not have to wait until allocating all the requested slaves (multicast group) to avoid deadlock as suggested previously [13]. To mathematically express the problem in this approach, assume a master (M_i) is sending a request vector (RV_i) to the Global multi-resource arbiter (GMA) to allocate a set of multicast-group ($MG_i = \{S_x, S_y, ..., S_z\}$ and $MG_i \subset \{S_0, ..., S_N\}$, where S is a slave, N is the NoC size). The probability that a slave $(S_j \in MG_i)$ is free in time slot t is denoted by $P_t(S_i)$. Therefore, the probability that the M_i request $(Rg_{i,j})$ will be granted is the intersection probability: $P_t(S_x, S_y, ..., S_z)$, where $P_t(S_x, S_y, ..., S_z) \leq P_t(S_j)$ for all $S_j \in MG_i$. This will keep free requested slaves ideal until all the members of MG_i are free.

To overcome this problem, this paper proposes a stretched multicast where a master can transmit to any number of free slaves and retransmit it to the rest later. This achieved using decentralized arbitration realized by any simple independent fair arbiter in each slave. Round-Robin (RR) has been chosen since it provides stronger fairness than rotating or random arbiters and less circuit complexity than matrix and queuing arbiters [1]. There are many possible scenarios where this scheme will show better contention handling, fairness and higher SWI utilization. For instance, Fig. 4 shows a comparison example of contention handling between centralized (W-SWI-C) with two



Fig. 4: Example shows that traffic flows from masters, M_1 , M_2 , M_3 and M_4 , are multiplexed in SWI. (a) W-SWI-C architecture (centralized arbitration), and (b) W-SWI-D architecture (decentralized arbitration and stretched multicast) where M is a master, S is a slave and T represents the time series.

virtual channels in the port linked to the SW (VCSW) and decentralized (W-SWI-D). Clearly, even though we assumed that the W-SWI-C manage to allocate all MG_i at T_1 , it offers less fairness and the multiplexing between flows is limited to the packet level. In contrast, although the same flit might be retransmitted (up to N_M , where N_M is the number of master nodes), the stretched multicast offers higher fairness and utilization of the SWI. This is due to reallocating slaves on flit base prevents flows from blocking each other. The blocking period (T_b) can be predicted to be $T_b \geq Packet_{size} \times N_v$, where N_v is number of VC. This is due to any congestion in the master (The channel owner) or the slave would that cause an idle time slot(s). In addition, stretched multicast can improve quality of service (QoS) in terms of the approximation of the flows average delays.

B. Deadlock-free Flow Control

The SWI can be represented logically as a bus topology with multiple master nodes (Tx/Rx) and multiple slave nodes (Rx). Each master has its own dedicated physical channel (frequency channel). However, each slave can listen to one master at a time, which creates a competition that escalates as the network size and master number is increased. In addition, in case of multicast these masters became the forking point which might cause a multicast dependency. Fig. 5 shows an example of multicast dependency where two masters M_1 and M_2 are requesting slaves $\{D_1, D_2, D_3\}$ and $\{D_1, D_3, D_4\}$, respectively. In this example, M_1 manages to allocate $\{D_2, D_3\}$ and now waiting for D_1 , while M_2 allocated D_1, D_4 and now waiting for D_3 . Consequently, this cycle dependency will cause a deadlock since each master will not release the slaves that has been allocated unless it delivers the flit to the remaining multicast group. Therefore, previous work proposed a centralized global multi-resources arbiter (GMA) which avoids this problem by allocating all requested slaves as a group and release them as group once the whole wormhole flow (packet) has been delivered [13].

In order to break the multicast dependencies without the need to allocate all the requested slaves (MG_i) at a given time slot, each master virtually should have its own non-blocking channel for every slave. This will allow wormhole switching on packet level yet with cut-through switching on the flit level. In previous work, the allocation of slave's VC was done dynamically which gives the advantage of having $N_v \leq N_M$



Fig. 5: An example of the deadlock created by the multicast dependency. M_1 allocates a subset of the multicast group $\{D_2, D_3\}$ that is requested by M_2 while waiting for M_2 to release the rest of the multicast group slave (D_1) .

[13]. However, the drawbacks of these blocking channels are the need to allocate all MG_i to avoid multicast dependency and the limitation to a wormhole switching. The other design option is to have each master have its own statically allocated virtual link. Therefore, since each master already has its own physical channel frequency, the virtually non-blocking channel can be achieved at the slave router by either using statically allocated VC where each master transmits via one VC ($N_v = N_M$). The other solution is ID-tagging based flow control(Tag) that developed in this paper. This is simply achieved by tagging each flit (F) with transmitter master ID (M_x) so that $Tag = M_x : Tag, M_x \in \{1, ..., N_M\}$. Then at the reservation table (RT) the allocation entry is distinguished based on the input buffer (i), Tag and the input VC (V_i , if there is VCSW). For simplicity the tag-ID remain unchanged in the output port while drained for the local processing element (PE). For example, assume a router micro-architecture that designed to provide two virtual non-blocking channels by using either ID-tagging or statically allocated VCs. Obviously, ID-tagging allows the designers to have virtual-channelless design which requires less area and power overheads. For this example if alterations are limited to SWI port, reduction in area $\sim 2\%$, and power $\sim 5\%$. Moreover, Tagging gives the freedom to use VC for other purposes such as multiplexing flows from the same master $(V\bar{C}S\bar{W})$ in case of a congestion.

To prove that this scheme breaks the multicast dependencies on the SWI layer, assume we have requests $(Rq_{i,j} \in RV_i : Rq_{i,j} \in \{0,1\})$ for a M_i to any slave $(S_j \in \{1,...,N\})$. These requests should be granted within finite time (T_f) . This is to prove that $\forall Rq_{i,j} : Rq_{i,j} = Gr_{i,j} = 1$ within T_f , where $Gr_{i,j}$ is grant signal. Given the probability $(P_t(S_j))$ of a slave local arbiter (RR) grants a master request at next time slot T_n is $P_{T_n}(S_j) = \frac{1}{N_M}$. In worst case scenario M_i has just been granted in T_{n-1} and all masters are requesting S_j at T_n . Thus, $P_{T_{n+N_M}}(j) = 1$ so that maximum waiting is $= N_M$. Therefore, the serving time ST_i needed for a flit to be delivered to all slaves and ejected from the output buffer in M_i is:

$$ST_i = max \Big(\bigcup_{T_n = 1 \to Max(T_n)} \{T_n \times (Rq_{i,0}, ..., Rq_{i,N})\}\Big)$$
(1)

where $max(T_n) = N_M$ for all $Rq_{i,j}$, $j \in \{1, ..., N\}$, therefore $1 \leq ST_i \leq N_M$. For instance, if M_1 send requests $\{Rq_{1,1}, Rq_{1,2}, Rq_{1,3}\}$ at T_0 , which has been granted in: $\{T_3, T_3, T_1\}$. Then according to Equation. 1, $ST_1 = T_3$.

C. 1-to-M Routing Scheme

Routing techniques are required to direct and deliver 1to-M traffic to its final destination(s) with minimum cost and better utilization of the proposed architecture. Previous



Fig. 6: Demonstration of 1-to-M W-SWI routing schemes, which is an improved tree-based with branching only at SWI.

Procedure: Slave interfacej: local node ID
for $i = 1 \rightarrow N_M$ do
$R_i = Rx(RV_i[j]);$
for all the $i \in \{1, \dots, N\}$ if all $Cr_{i+1} = 0$ do
$\begin{vmatrix} i = RR(R_1,, R_1); \\ Tx(Gr_{i,j} = 1); \end{vmatrix}$
if $Rx(Flit) == 1$ then
$Tx(G_{i,j}=0);$
end
$ \begin{array}{c c} \textbf{Procedure: Master interface} & i: \ local \ node \ ID \\ \textbf{forall the} \ P \in \{N, E, S, W, L, SW\} \ \textbf{do} \\ & \textbf{if} \ Reserve(Input - port, Destination, MAB) == SW \ \textbf{then} \\ & \textbf{forall the} \ j \in \{0, \dots, N\} \ \textbf{do} \\ & RV_i[j] = RV_i[j] \lor MAB[j]; \end{array} $
end
end end
$\begin{vmatrix} i \\ end \\ end \\ for all the i \in \{0, \dots, N\} \ IF any \ Br(Gr[i]) = -1 \ do$
$\begin{vmatrix} & & end \\ & end \\ end \\ forall the \ j \in \{0,, N\} \ IF \ any \ Rx(Gr[j]) == 1 \ do$
$\begin{vmatrix} & \text{end} \\ & \text{end} \\ & \text{end} \\ & \text{forall the } j \in \{0,, N\} \text{ IF any } Rx(Gr[j]) == 1 \text{ do} \\ & Tx(Flit = FIFO()); \\ & BV([i] = M AB[j] = 0; \\ & \text{end} \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{forall the } SV([i]) = M AB[j] = 0; \\ & \text{for } SV([i]) = M AB[j] = 0;$
$ \begin{vmatrix} end \\ end \end{vmatrix} $ end forall the $j \in \{0,, N\}$ <i>IF any</i> $Rx(Gr[j]) == 1$ do Tx(Flit = FIFO()); $RV_i[j] = MAB[j] = 0;$ forall the $j \in \{0,, N\}$ <i>IF all</i> $MAB[j] = -0$ do
end end forall the $j \in \{0,, N\}$ IF any $Rx(Gr[j]) == 1$ do Tx(Flit = FIFO()); $RV_i[j] = MAB[j] = 0;$ forall the $j \in \{0,, N\}$ IF all $MAB[j] == 0$ do FIFO.eject; end
end end forall the $j \in \{0,, N\}$ <i>IF any</i> $Rx(Gr[j]) == 1$ do Tx(Flit = FIFO()); $RV_i[j] = MAB[j] = 0;$ forall the $j \in \{0,, N\}$ <i>IF all</i> $MAB[j] == 0$ do <i>FIFO.eject</i> ; end end

Algorithm 1: Master/slave communicate protocol using SW interconnects

related work can be classified into: path-based and tree-based schemes. In the former, a 1-to-M packet is submitted to the members of a multicast group sequentially [19]. This method shows very significant delivery latency. On the other hand, in the tree-based scheme, 1-to-M traffic is propagated in a tree route embedded in the NoC [6], [5]. This alleviates traffic load and reduces energy consumption for the interconnect fabric. Nonetheless, this scheme still magnifies the issues with the regular wire-based NoC since it increases the load and hotspots.

The routing for the proposed architecture is an improved tree-based scheme where the embedded tree path forks at one point (nearest master), as shown in Fig. 6. 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 for the following reasons: (1) each node simply needs to be aware of the nearest master then direct the traffic to it using any partially adaptive routing algorithm. Hence, there is no need for extra circuitry or complicated algorithms; (2) packets will be replicated only at the destination routers. This will reduce power consumption by eliminating the need for duplicated traffic to travel through costly intermediate resources such as wires and routers. In order to direct the packet to the multicast group members (MG_i) , each multicast packet header must have multicast-address-bits (MAB). This header field is a bit vector where each bit is representing a salve (S_i) and it is set if $S_i \in MG_i$.

To understand the communication protocol taking place at

the SWI layer, Algorithm 1 shows the procedure of masters and slaves interfaces. A master (M_i) interface sends a RV_i via SWI that consists of a request $(Rq_{i,j})$ to each slave $(S_j|\forall j:$ MAB(j) = 1). At the slave end, the local RR arbiter will determine to which master it will listen to in the next time slot and send the $Gr_{i,j}$ signal. Both $Rq_{i,j}$ and $Gr_{i,j}$ are transmitted via SWI. When the RR grants the request, the data handshake phase starts by sending the data flit to all slaves responded to the request and waits for them to acknowledge the reception before resetting the requests intended to them. The adopted handshaking protocol is a Non-return-to-zero protocol [20]. Then, if all MG_i members received the multicast flit, it will be ejected.

IV. SYSTEM LEVEL EVALUATION AND DISCUSSION *A. Performance Evaluation*

This section demonstrates results obtained from our cycle accurate NoC simulator which was built by modifying the existing Noxim simulator [21] for the W-SWI. In addition, a state-of-the-art tree-based on-chip multicast scheme was replicated in order to compare it with the proposed architecture, which is the virtual circuit tree (VCT) [5]. This technique is chosen because it is efficient and simple where a minimum of modifications to the baseline router architecture are required except for the VCT table and the control circuits for forking (one flit/cycle). The packet size of 12 flit is chosen to demonstrate the behaviour of the proposed architecture under a wormhole flow control. However, due to the fact that the VCT is not designed to handle wormhole traffic (deadlock occurrence), four flits packet size were chosen when evaluating against VCT. Also, a VCT with 512 entries had been considered in our evaluations of VCT since it offers good performance/cost level in most cases [5]. The number of master nodes in W-SWI is based on the available frequency range for 45nm, which is estimated to be 4 channels (adding 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. The evaluation is done under uniform traffic where the multicast source and group members were selected randomly.

The results show significant improvements in performance in terms of average delay. For instance, Fig. 7a shows much less average delay for W-SWI-D over the VCT even for zeroload-latency (ZLL), \sim 12X less. This is simply due to the fact that W-SWI-D is utilizing the SWI short-cuts to reach all the multicast destinations in one hop, which reduce the load on the NoC. On the other hand, Fig. 7b shows comparison between W-SWI-C and the proposed decentralized arbitration and routing scheme (W-SWI-D) for different VC sizes. For router architecture simplicity, all ports has the same VC size including the ports linked to SWI. Obviously, the proposed W-SWI-D shows better performance than W-SWI-C with one VC(improvements $\sim 10\%$). This is due to the fact that with one VC, ID-tagging allows each master to have a virtually nonblocking channel. Moreover, performance results show that even for higher N_v , the W-SWI-D is better before the NoC saturation (improvements: $\sim 6\%$ for two VC and $\sim 14\%$ for four VC at $PIR = 1.5 \times ZLL$). As a result of higher SWI utilization (up to 2.5% increase in the flit number transmitted via SWI). However, at the saturation point the multiplexing delay between masters starts to overcome the improvements



Fig. 7: System level evaluation of average delay results for NoC (6×6) with 10% multicast PIR of the total PIR: (a) VCT compared to W-SWI-D, (b) W-SWI-C versus W-SWI-D with different VC numbers.

TABLE I: Results improvements over Baseline architecture (Mesh) comparison between W-SWI-C and the proposed W-SWI-D with 10% multicast ratio.

NoC	VC	PIR	improvements over Mesh					
size		$(\times 10^{-4})$	Average delay (x)		Power (%)			
			W-SWI-C	W-SWI-D	W-SWI-C	W-SWI-D		
	1	15	9.64	10.61	63.34	63.64		
6x6	2	35	10.15	11.15	55.2	56.03		
	4	45	10.09	11.77	53.57	55.24		
	1	17	8.65	9.6	56.84	57.24		
8x8	2	29	8.46	9.38	52.25	53.56		
	4	33	8.61	10.58	51.41	52.94		
	1	12	8.24	9.16	55.48	55.77		
10x10	2	22	7.14	8.21	50.14	51.11		
	4	24	7.55	9.69	49.56	50.99		

in the SWI utilization. Thus, This architecture is favourable for low or medium load applications. Moreover, W-SWI-D improvements ratio over the Mesh compared to W-SWI-C seems to be steady (up to 2X at normal load), even when the NoC size increased as shown by the Table I.

B. Power Reduction

This section presents the evaluation results of an important cost metric for future interconnects, and for 1-to-M in particular, which is power consumption. The router's static and dynamic power is calculated using the Orion 2.0 [22] 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 modelled based on the analytical model introduced previously [11]. The GMA [13] and the local RR arbiter was designed using Verilog and then synthesized using the 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 were used in Noxim to calculate the overall NoC power consumption.

Table I shows the power consumption saving ratios for both W-SWI-C and the proposed W-SWI-D over the Mesh for different NoC sizes and N_v . Evidently, both architectures demonstrate significant improvements in the NoC power consumption reduction ratio over Mesh. In addition, even though the W-SWI-D may retransmit the same flit many times (up to N_M) via SWI it achieves power saving (1 - 2%) due to reduction of the flit delivery time. These findings prove that the W-SWI-D is more effective in mitigating 1-to-M communication issues in terms of power consumption.

NoC component		Area per item (mm^2) for $45nm$ technology				
Component	No	Mesh	W-SWI-C	W-SWI-D	VCT [5]	
Router	24	1.0853	1.5124	1.5931	1.0853	
Transmitter	4	-	0.1558	0.1558	-	
Receiver	24	-	0.0083	0.0083	-	
Global arbiter	1	-	0.0552	-	-	
Local arbiter	24	-	-	0.0002	-	
VCT table	24	-	-	-	2.3608	
Wire Links	1	13.653	13.653	13.653	13.653	
Total extra area over M	Mesh=	= (all compo-	11.13	13.01	56.66	
nents × their No.)- M	lesh a	trea (mm^2)				
NoC/SCC-die area (%)	7	8.96	9.29	16.99	

TABLE II: Area overhead evaluation for W-SWI-C, proposed W-SWI-D and VCT-512 over Baseline architecture (Mesh).

C. Area Overheads Evaluation

It is essential to evaluate chip area overheads for the extra on-chip circuits required for the proposed architecture. Firstly, for transceivers it is assumed that the active area calculated in previous research [8] is the only part scaled down when moving 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, where the active area is proportional to the square of the scaling factor. Secondly, the area for extra router port (buffer, crossbar and related circuits) and the VC allocation unit for the W-SWI-D are calculated using the Orion 2.0 [22] model as $0.427mm^2$ and $0.0813mm^2$, respectively. In addition, the global arbiter (GMA) and the local arbiter (RR) was designed using Verilog and then synthesized using Synopsys design compiler and mapped onto the PDK 45nm technology library to calculate its area $0.0114mm^2$ and $0.0114mm^2$, added to it the TxRx estimated area $0.0438mm^2$. Likewise, the VCT with a 512 entry/source lookup table was designed using Verilog and then synthesised to be $2.3608mm^2$ per router.

Table II shows area considerations for the Mesh, the W-SWI-C, the proposed W-SWI-D, VCT and WiNoC. Obviously, the W-SWI-D area overhead is higher than the W-SWI-C ($\sim 2.3\%, \sim 2\%$, respectively). This is mostly due to increasing the VC allocation unit area in all routers to implement the ID-tagging scheme. However, compared to having regular VC, the area overhead would be much higher, see section III-B. These area overheads is considered negligible compared to 1-to-M wire-based NoC such as VCT ($\sim 9\%$).

V. CONCLUSION AND FUTURE WORK

This paper proposes a mixed W-SWI architecture with effective stretched multicast scheme that is able to tackle onchip 1-to-M communication efficiently and gives an attractive area/performance trade-off. SW low power dissipation, high signal propagation speed and fanout capability all help to significantly mitigate the 1-to-M communication issues that the NoCs-based CMP in particular suffer from. In addition, newly introduced decentralized deadlock-free routing techniques such as stretched multicast and ID-tagging flow control for this architecture and 1-to-M traffic are discussed. The evaluation results show significant improvements in terms of average delay and power consumption with a relatively small die area penalty compared to other architectures. Future work should include the investigation of many-to-one traffic patterns.

REFERENCES

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

- [2] J. Duato, S. Yalmanchili, and L. Ni, Interconnection Networks : an Engineering Approach. Los Alamitos, Calif.: IEEE CS Press, 1997.
- [3] P. Salihundam, S. Jain, T. Jacob, 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] N. Jerger, L.-S. Peh, and M. Lipasti, "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.
- [6] T. Krishna, L.-S. Peh, B. M. Beckmann, and S. K. Reinhardt, "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, (NY, USA), pp. 71–82, ACM, 2011.
- [7] Semiconductor Industry Association, "ITRS: International Technology Roadmap for Semiconductors ." http://www.itrs.net/reports.html [online], 2009.
- [8] M. C. F. Chang, J. Cong, A. Kaplan, C. Liu, M. Naik, J. Premkumar, G. Reinman, E. Socher, and S.-W. Tam, "Power reduction of cmp communication networks via rf-interconnects," in *Microarchitecture*, 2008. MICRO-41. 2008 41st IEEE/ACM International Symposium on, pp. 376–387, Nov 2008.
- [9] S. Deb, A. Ganguly, P. Pande, B. Belzer, and D. Heo, "Wireless noc as interconnection backbone for multicore chips: Promises and challenges," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, vol. 2, pp. 228–239, June 2012.
- [10] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proceedings of the IEEE*, vol. 97, pp. 1166 –1185, july 2009.
- [11] A. Karkar, R. Al-Dujaily, A. Yakovlev, K. Tong, and T. Mak, "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.
- [12] A. Karkar, J. Turner, K. Tong, R. AI-Dujaily, T. Mak, A. Yakovlev, and F. Xia, "Hybrid wire-surface wave interconnects for next-generation networks-on-chip," *Computers Digital Techniques, IET*, vol. 7, pp. 294– 303, November 2013.
- [13] A. Karkar, N. Dahir, R. Al-Dujaily, K. Tong, T. Mak, and A. Yakovlev, "Hybrid wire-surface wave architecture for one-to-many communication in networks-on-chip," in *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*, pp. 1–4, March 2014.
- [14] J. Turner, M. Jessup, and K. 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.
- [15] J. Hendry, "Isolation of the zenneck surface wave," in Antennas and Propagation Conference (LAPC), 2010 Loughborough, pp. 613 –616, nov. 2010.
- [16] M.-C. Chang, V. Roychowdhury, L. Zhang, H. Shin, and Y. Qian, "Rf/wireless interconnect for inter- and intra-chip communications," *Proceedings of the IEEE*, vol. 89, pp. 456–466, Apr 2001.
- [17] A. Ganguly, K. Chang, S. Deb, P. Pande, and *et al.*, "Scalable hybrid wireless network-on-chip architectures for multicore systems," *Computers, IEEE Transactions on*, vol. 60, pp. 1485–1502, Oct 2011.
- [18] W. J. Dally and J. W. Poulton, *Digital systems engineering*. New York, NY, USA: Cambridge University Press, 1998.
- [19] X. Lin, P. McKinley, and L. Ni, "Deadlock-free multicast wormhole routing in 2-d mesh multicomputers," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 5, pp. 793–804, Aug 1994.
- [20] D. Kinniment, Synchronization, Arbitration and Choice. England: Wiley Publishing, 2007.
- [21] F. Fazzino, M. Palesi, and D. Patti, "Noxim: Network-on-chip simulator."
- [22] A. Kahng, B. Li, L.-S. Peh, and K. Samadi, "Orion 2.0: A power-area simulator for interconnection networks," *Very Large Scale Integration* (VLSI) Systems, IEEE Transactions on, vol. 20, pp. 191–196, Jan 2012.