

Reducing Power and Area by Interconnecting Memory Controllers to Memory Ranks with RF Coplanar Waveguides on the Same Package

Mario D. Marino and Kevin Skadron
Dept. of Computer Science, University of Virginia
{mdm9uw,skadron}@cs.virginia.edu

Abstract

The physical channel is the element that consumes the largest amount of power in a traditional memory controller (MC). Wired-RF can potentially decrease the amount of power dissipated by replacing the physical memory channel by an RF-channel, just as optical memory systems do by replacing the physical memory channel by an optical-channel.

Considering that RF transmission can potentially consume less power than a traditional bus for on-chip distances, we propose to replace the traditional digital MC physical channel by coupling RF transmitters (TX), receivers (RX), an RF quilt-packaging coplanar waveguide (CPW), and a quilt-to- to interconnect MCs and memory ranks on the same package in a multicore. We evaluate the proposed solution in terms of power and area employing ITRS [1] and RF predictions[17]. Preliminary estimation shows that the proposed RF interface is able to save up to 57.3% in terms of area and up to 78.2% in terms of power consumption for next processor generations. Furthermore, considering a fixed area budget of one MC as a reference, the proposed interface can improve bandwidth up to 2.2x for an 8-core multiprocessor with 3 MCs and, assuming a fixed power budget of one MC, the proposed interface can improve bandwidth of up to 2.4x.

I. INTRODUCTION

The use of excessive power in the physical memory channel is one of the factors that limits bandwidth scaling. For example, Corona [18] estimates an excessive power of 160W assuming the estimation of 2 mW/Gb/s by Palmer [15] for a bandwidth of 10 TB/s. Thus, saving interconnection power in a physical memory channel interconnection needs to be investigated.

Wired-RF is an alternative just as optics to improve bandwidth over the traditional electrical solution. Both technologies share similar core concepts of modulation and a low communication media; furthermore, both solutions can reduce the number of I/O pins [18][9][8], and, finally, use less power than electrical interconnects [18][5]. At the RC domain, wires can carry multiple signals through different carriers, and signals can travel at the speed of light [7]. With the same goal of advanced optical memory systems which estimates to achieve around 0.078 pJ/bit [18], recent studies show the potential of wired-RF in reducing memory interconnection power, by reducing the memory bus power to around 2.5 pJ/bit for a 8.4Gbps memory [4]. Since the integration with wired-RF is a natural extension of CMOS, and given the power gains potentially achieved, we believe that this solution can be better coupled to CMOS circuits than optics as a technique to improve bandwidth while reducing power.

RF can be potentially used to save power when compared to optical and electrical for distances from 1mm up to 30 cm [17]. For an on-package scenario under this distance constraint, we observe that the power reduction can be potentially obtained in two different sub-scenarios, i.e., interconnecting processor and memory in a 3D stack or coplanar configurations. Fabricating vertical RF-lines such as CPW or transmission lines is still challenging since the way the materials are deposited to form the lines cause dispersion, distortion, and reflection effects [13]. Moreover, in a 3D scenario, even considering a large number of DRAM rank layers, the distance that the signals travel remains small, thus the use of RF might not be beneficial when it comes to power reduction; instead, traditional electrical transmission would be more appropriate. A current scenario where RF can possibly benefit power reduction is a coplanar layout disposition, with a processor die connected to a set of rank memories on the same package using a RF waveguide such as quilt-packaging [11]. Moreover, the Intel study by Polka[16] noticed that average length between processor and memory-stacks in a multi-chip-package (MCP) is around 7mm, which confirms that RF can be applied within this distance range.

We assume that the length of the interconnection from the MCs plus the length of the signals through the RXs, TXs, Quilt-packaging [11]¹, until the center of the DRAM ranks is larger than 1mm, so that RF power benefits are likely to be achieved. We also presume that the DRAM ranks are placed around the processor die but on the same package such as multi-chip-package (MCP). Figure 1 (a) illustrates this scenario.

We propose to replace the traditional physical memory bus channel by a set of TXs, RXs, the quilt packaging, and the CPW interconnection. In this case, each set of TXs and RXs connects each MC to each DRAM rank. Furthermore, we define the term

¹Quilt-packaging interconnection is [11] a low cost CPW which was designed to be used at RF frequencies. In this technique, a CPW between two dies facing each other is formed by extending the on-die interconnection of each one. Most importantly, quilt-packaging minimizes capacitance, inductance and reflection effects, which are extremely important in the RF scenario.

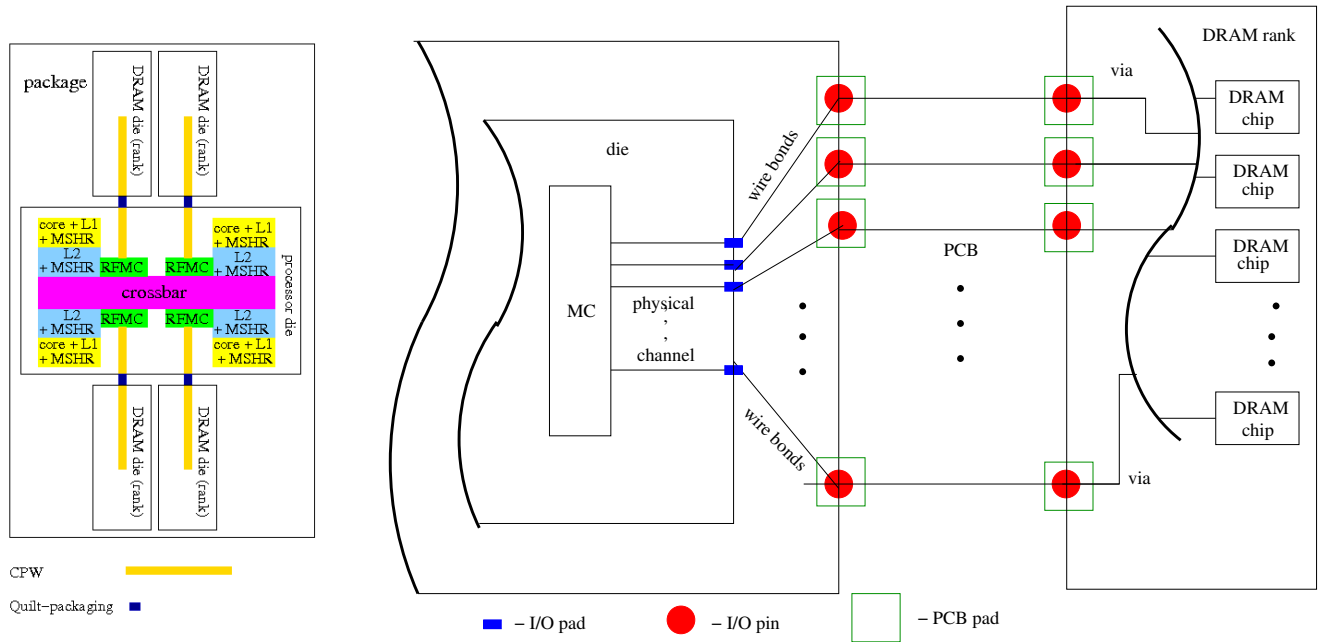


Figure 1: (a) processor and DRAM on package (not on scale); (b) signal path from MC to DRAM rank

RFMC-quilt to specify that in the proposed solution the MC is coupled to TXs, RXs, and the quilt-packaging interconnection. The main contributions of this paper are as follows:

- propose an on-die processor-memory organization suited to RF
- evaluate the power and area gains when replacing the typical MC and its physical memory bus channel by the RFMC-quilt, coupled with TXs, RXs, and the quilt-packaging interconnection for different technology generations.
- given the potential power and area gains proportioned by the proposed solution, we investigate the architectural benefits in terms of performance for an 8-000 multicore.

Section 1 describes the design challenges when coupling the MC with the RF elements. Section 2 discusses the issues that traditional electrical interconnections have if they are set to work at RF frequencies. Section 3 explains how quilt-packaging is used as interconnection. Section 4 presents the proposed organization. Section 5 evaluates the power and area savings by replacing the traditional electrical physical memory channel by quilt-packaging. Finally, section 6 concludes the paper.

II. WHY THE CURRENT DIGITAL PATH DOES NOT WORK AT RF FREQUENCIES?

In a traditional processor-to-memory path, signals have to traverse a long way composed by a sequence of electrical elements. Assuming that we have an on-chip memory controller, the physical transmission part of MC is connected to vias, which are extended and routed through the microprocessor through repeaters until they are connected to the inner bond leads. Then, to get to the external world, the inner bond leads are connected to the outer bond leads. After that, electrical wires connect the outer bond leads to the corresponding I/O pads at the substrate [11]. Then, small soldered balls on the I/O pads are used to physically attach the bottom part of the processor to a soldered ball of the flip-chip interconnection. After going through the soldered ball, the signal achieves the carrier. Then, the bottom side of the package carrier gets connected to the motherboard PCB or buses. Finally, after traveling a certain distance and having repeaters restoring the signal, the electrical signal gets the memory slots, which are attached to the memory ranks themselves. Figure 1 (b) has the complete path of the signals.

Coupling a traditional Flip-chip package with an RF interconnection is challenging because all of the elements in the signal path have different impedances. If this traditional interface were coupled to RF at RF frequencies, the different impedances of each signal-path element would cause reflection, distortion, and interference effects, thus causing signal loss and bandwidth degradation, consequently requiring signal regeneration and increasing power consumption. These sources of bandwidth and power inefficiency motivate the use of an appropriate RF interconnection such as quilt-packaging.

III. QUILT-PACKAGING USED AS AN RF PROCESSOR-TO-MEMORY INTERCONNECTION

CPW is a type of waveguide designed to conduct RF waves. It is composed of a central metallic strip line placed on top of a dielectric plane. The metallic line is separated by two different slits (empty regions) from a ground plane. Figure 2 (a) illustrates a CPW with a W-wide and separation from the metallic line by S-wide slits.

Quilt-packaging [11][20] was designed to address general losses at high frequencies. Quilt-packaging is a CPW transmission line designed to eliminate the traditional signal path by connecting two dies directly using such a waveguide. The quilt-package

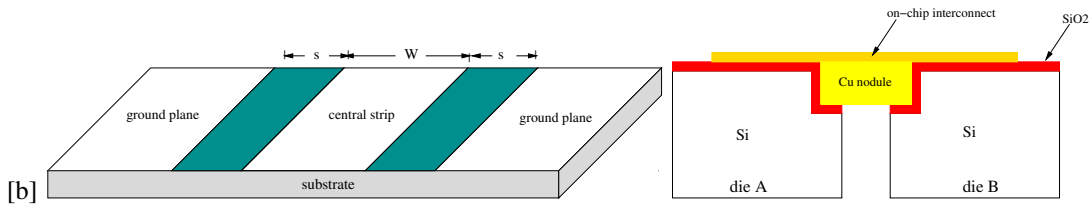


Figure 2: (a) and (b): CPW and quilt implementation extracted from [11]

interconnection general idea can be seen in figure 2 (b). By extending the interconnection of two (or more) dies facing each other, a CPW interconnection is formed. To improve the discontinuity between the on-chip interconnect and the copper nodule, a tapered waveguide nodule pattern was also included. As a result, coupled with transmitters and receivers, the amount of elements that the signal has to traverse when going from one die to another is drastically reduced. As a consequence, quilt-packaging reports a remarkable loss of 0.1dB [11], hence extremely suitable to be used at high frequencies. According to [11], the prototype has the following parameters: resistivity of 10 ohm.cm, thickness of the substrate 100um, SiO₂ thickness of 1um, on-chip Cu interconnect is 1um thick, spacing between the signal and CPW ground plane is 8um, and the on-chip interconnect has a length of 200um, and copper nodule length is 100um – 80um embedded in the silicon substrate and 20um extended outside the chip edge, so that the distance between chips is 40um.

There are several reasons for employing quilt-packaging in order to connect the MCs to each DRAM rank. First, it was designed and simulated for RF frequencies up to 200GHz. Furthermore, a prototype was built and tested for RF frequencies up to 60 GHz. Although the fabrication process of quilt-packaging is different from typical microprocessor fabrication processes, ITRS RF scaling predictions are still valid if applied to transceivers and receivers coupled to quilt-packaging since the typical distances that the signal has to traverse are typically on-chip ones [17]. Once quilt-packaging was developed for connecting dies on a package, its dimensions are much smaller than FR4 external board ones, which are typically used to connect different chips [13][4] thus consuming smaller power. Finally, due to its design and fabrication technique, quilt presents low losses (around 0.1dB), thus potentially providing much larger bandwidth and also contributing to have a low power usage.

IV. USING RF TO INTERCONNECT A MC AND A DRAM RANK

Figure 4 illustrates how the proposed interface is placed between the MC and the DDR2 memory. For this project, on the memory side, we place the transmitters and receivers at the memory ranks (center). As another design choice, they could be placed at the memory banks; in this case, it would be necessary to replicate them in each of the banks in order to receive the memory requests, which potentially increases area and power. We leave a more detailed design exploration of the placement of the transceivers and receivers at the memory side as a future work.

After leaving the MC, a memory controller request has to go through a sequence of elements composed by bus, I/O pads, soldered balls, I/O pins, socket, PCB, and finally the memory rank, where, after traversing the memory rank pads, the signals go through busses and finally achieve each of the memory banks. To have signals traversing in both directions, the RF interface has transmitters and receivers on the MC and on each rank. In the proposed interface, after a MC request, the signals go through the transmitters – where they are converted to analog waves, then they go through the quilt-packaging interconnection, and finally the receivers at the ranks – where the analog waves are converted back to digital signals; after the receivers, the signal is sent through busses and achieve a memory rank. The signal does traverse the same path in the opposite direction, when a memory rank responds to a MC request: the rank response (memory burst) is sent to the transmitters placed at the memory rank. After that, the transmitters convert the response to waves, which go through the quilt-packaging interconnection,

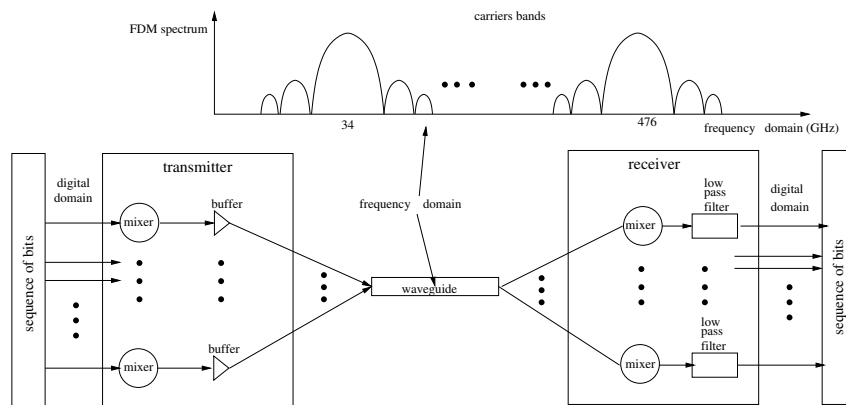


Figure 3: extracted and modified from [12]

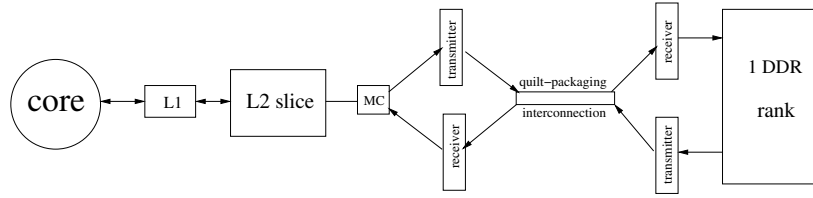


Figure 4: RF processor-memory organization

and finally get in the receivers on the MC side. These receivers convert down the waves into electrical signals, which finally achieve the MC. Figure 3 shows the described path.

For the rest of the paper, when we refer to the term RFMC we refer to the MC interfaced with the DRAM ranks as they are in the proposed RF interface, i.e., with RF transmitters, receivers and the quilt-package interconnection. The process of up-conversion and down-conversion of data signals to analog waves and vice versa involves signal delay going through transmitters and receivers. For example, for an RF transmission line, the typical range of these delays is around 200 picoseconds (ps) [12]. ITRS predictions [1] indicate the number of carriers and data rate scales with the advance of technology as showed in table I. First, we notice that modulation allows a scalable data rate per wire. Considering the total data rate per wire seen in this table, RF allows us to go from 30 to 140 Gbits/s, which are large enough to support data transfers of typical DDR3 data rates from 10GB/s to 17GB/s [2] using only one wire; that is generically how one quilt-packaging line can supply enough memory bandwidth to the DRAM ranks; we leave this exploration as a future investigation.

V. EVALUATION AND ANALYSIS

MC are basically composed of three different parts [10], (i) the front engine (FE), that receives requests from the memory, (ii) the transaction engine (TE), that translates the memory requests into control and physical memory requests, and finally, the physical transmission (PHY), composed by control and data physical channels.

We employed McPAT [10] to derivate the area and power of the FE and TE for both MC and RFMC assuming a 16-buffer entry and both MCs clocked at 2.0 GHz. In order to estimate the area of the RFMCs, we aggressively assume that a RFMC is basically composed by removing the PHY part of a MC plus the RX and TX elements.

We evaluate the energy consumed by memory channel in terms of energy per bit (mW/Gbits/s). To estimate it, we assumed that one rank has a fixed data rate budget of 5 GBytes/s and it is 64-bit wide, with memory bus clocked at 333 MHz; the baseline is composed by one on-package DRAM rank with this specification and one MC. Given this memory data rate budget, to calculate the power employed in the RFMC elements, we calculated the amount of power used by the TXs, RXs, CPW rank-to-quilt interconnection, and quilt. We presume that all the interconnection path between the RFMC to the center of the DRAM rank was designed to match quilt-interconnection in order to avoid reflection and insertion losses. Furthermore, to the best of our knowledge, up to this point there are no TX and RX circuits designed for quilt and coupled or not to the CPW extension. We then premise that these elements behave such as the RF-I interconnection proposed by F. Chang [6], which was proposed for CPWs and transmission lines. In addition, the power estimation considered RFMC-to-DRAM-rank on-package ranges of distances from 1 to 10mm. Due to the low-loss insertion provided by quilt along a short part of the signal path and given the absence of TX/RX circuits designed for quilt, we conservatively adopt a 10% reduction on the power supplied by the amplifier (TX side). Finally, we estimate the energy used along the MC and the respective DRAM rank by considering: (i) the same distance range adopted for the RFMC case; (ii) power reduction of 10% on each technology generation [17]; and finally, that (iii) the distance versus energy projections by Tam [17] are also valid to other than 16nm technology generations.

Figure 5 (a) shows the estimated area for the different technology generations. The left bars illustrate the area distribution of each of the MC components, while the right ones represent the area distribution regarding the RFMC components. FE and TE occupy the same amount of area in both bars because in both MC and RFMC implementations FE and TE are equivalent, i.e., they have similar dimensions.

Moreover, as expected by shrinking the transistor size in each technology generation, we notice that each of the individual components of either the MC or the RFMC has a reduction in size of up to circa 42% between two consecutive technology

technology (nm)	data rate per band (Gbits/s)	#carriers per wire to match DRAM (mW)	Power TX + RX	Energy per bit (pJ)	Area (Tx + Rx) (mm ²)
45	7	6	28.14	0.67	0.00690
32	8	5	24	0.60	0.00495
22	9	5	23.4	0.53	0.00439

Table I: RF parameters - extracted from [12][17] - modified with 10% power amplifier reduction and to match DRAM bandwidth

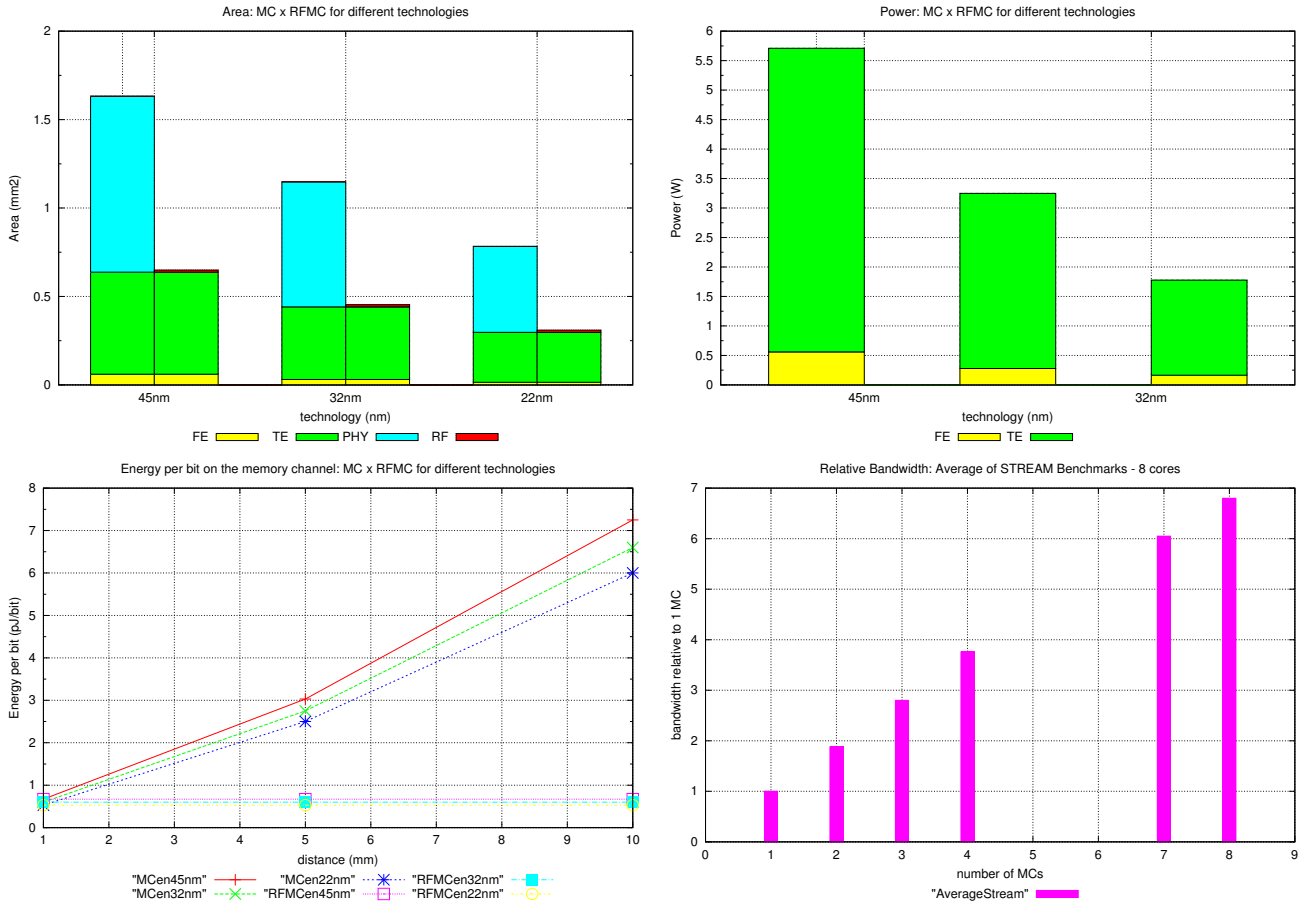


Figure 5: (a) area of the different technologies; the left bar represents a MC, while the right bar an RFMC ; (b) power of MC and RFMC, and (c) power of different memories; (d) average relative bandwidth of STREAM benchmarks; the baseline has 3MCs; RFMC within the area budget has equivalent 7 MCs and RF within the power budget has equivalent 8MCs

generations. From figure 5(a) we observe that we can save up to 57.3% considering the area reduction of RF; RXs, TXs, and quilt-interconnection smaller dimensions are responsible for this reduction. As a consequence, if we assume the area of a traditional MC as the area budget, we could potentially fit up to 2.4x more RFMCs, i.e., for an 8-core processor with area reserved for 3 traditional MCs, we can fit up to 7.2 RFMCs within the same area budget. Furthermore, we performed an experiment to investigate this benefit using M5 [3] and DRAMsim simulators [19] with STREAM benchmark [14]. The parameters of this experiment are placed in table II. In this experiment, we varied the number of MCs for 1,2,4, and 8 for an 8-OOO multicore and investigated the relative bandwidth benefit. We have not considered any RF latency benefits, which we leave for further investigation. Figure 5 (d) shows the bandwidth benefits; we extrapolated the results for 3 MCs and 7 MCs and included them in this figure. We observed that the RFMC (which corresponds to 7 MCs) has up to 2.2x more bandwidth compared to the baseline version with 3 MCs.

Figure 5 (b) shows the typical power range used by FE and TE elements, used in both MC and RFMC implementations. We observe that the power utilized in both parts decreases with the technology improvement due to the smaller power used by

Core	2.0 GHz, OOO-Core, 4-wide issue, branch predictor = bimodal
L1 cache	32kB dcache + 32 kB icache ; associativity = 2, MSHR = 8, latency = 0.5 ns
private L2 slice	1MB/per core ; associativity = 8, MSHR = 8; latency = 2.0 ns
RF crossbar	latency = 1 cycle, 2.0 GHz frequency
MC	1 MC/core, 2.0 GHz, on-chip, input buffer = 16; 8MCs or 8RFMCs trans. queue size = 16/MC
1 DRAM rank per MC	DDR2 667MHz, tcas=trcd=trp=15ns; tras=45ns based on Micron [2] MT8JTF12864AZ
STREAM [14] benchmark average	64 million elements

Table II: parameters of the modeled architecture and STREAM benchmark

smaller transistors across the generations. Figure 5 (c) illustrates the energy versus distance estimation for different technologies; considering an average distance of 5mm, we can save an average of 78.2% of PHY part energy, (the RF elements RXs, TXs, and quilt are more power-efficient than the traditional electrical TX, RX, and busses within this distance range). Another way to interpret these energy savings is as if, given the power budget of a MC as a baseline, we were able to fit up to 4.6x RFMCs. In order to estimate the benefits on the number of MCs in terms of performance, we considered a typical multicore on the market with 8 cores and 3 MCs. In this case, we could potentially fit up to 13.8 MCs with the same power budget. However, instead of 13.8 MCs, we conservatively assume 8 RFMCs to follow a 1:1 proportion between cores and MCs, and evaluate the performance gains of this benefit. Figure 5(d) shows that the RFMC version (which corresponds to 8 MCs) is up to 2.4x faster than the baseline with only 3MCs.

VI. CONCLUSIONS AND FUTURE WORK

We proposed to replace the physical channel that connects the MC and memory DRAM ranks by coupling the MC with RF elements such as RXs, TXs and a quilt-package coplanar waveguide. As a preliminary result, we potentially observed area and power reductions of up to 57.3% and 78.2% respectively when compared to a traditional MC. Assuming these gains in terms of power and area, the proposed technique showed a bandwidth improvement of up to 2.4x.

VII. ACKNOWLEDGEMENT

We would like to thank Prof Gary Bernstein and Prof Patrick Fay from University of Notre Dame for the fruitful discussions about quilt-packaging interconnection. We would also like to thank Prof Mircea Stan in helping discussing RF issues. This work was supported in part by NSF grant nos. CNS-0916908, CCF-0551630, and CNS-0615277. We would also like to thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] ITRS HOME. Accessed date: 09/10/2010 ; <http://www.itrs.net/>.
- [2] Micron manufactures DRAM components and modules and NAND Flash. Accessed date: 02/10/2010 ; <http://www.micron.com/>.
- [3] Binkert, Nathan L. and Dreslinski, Ronald G. and Hsu, Lisa R. and Lim, Kevin T. and Saidi, Ali G. and Reinhardt, Steven K. The M5 Simulator: Modeling Networked Systems. *IEEE Micro*, 26(4):52–60, 2006.
- [4] Byun, G. and Kim, Y. and Kim, J. and Tam, S.-W. and Cong, J. and Reinman, G. and Chang, M.-C. F. . An 8.4Gb/s 2.5pJ/b Mobile Memory I/O Interface Using Bi-directional and Simultaneous Dual (Base+RF)-Band Signaling. *IEEE*, 2011.
- [5] Chang, M.-C. Frank and Cong, Jason and Kaplan, Adam and Liu, Chunyue and Naik, Mishali and Premkumar, Jagannath and Reinman, Glenn and Socher, Eran and Tam, Sai-Wang. Power reduction of CMP communication networks via RF-interconnects. In *MICRO '08: Proceedings of the 2008 41st IEEE/ACM International Symposium on Microarchitecture*, pages 376–387, Washington, DC, USA, 2008. IEEE Computer Society.
- [6] Chang, Mau-Chung Frank and Vwani P. Roychowdhury and Liyang Zhang and Hyunchol Shin and Yongxi Qian. RF/Wireless Interconnect for Inter- and Intra-Chip Communications. 89(4):456–466, Apr 2001.
- [7] Chang, Richard T. and Talwalkar, Niranjan and Yue, Patrick and Wong, S. Simon. Near Speed-of-Light Signaling Over On-Chip Electrical Interconnects. *Journal of Solid-State Circuits*, 38(5):834–838, 2003.
- [8] Kim, J. and Xu, Z. and Frank, M.C. Reconfigurable Memory Bus Systems using Multi-Gbps/pi CDMA I/O Transceivers. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, volume 2, pages 33–36, Bangkok, Thailand, 2003. IEEE Computer Society.
- [9] Kirman, Nevin and Kirman, Meyrem and Dokania, Rajeev K. and Martinez, Jose F. and Apsel, Alyssa B. and Watkins, Matthew A. and Albonese, David H. Leveraging Optical Technology in Future Bus-based Chip Multiprocessors. In *MICRO 39: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 492–503, Washington, DC, USA, 2006. IEEE Computer Society.
- [10] Li, Sheng and Ahn, Jung Ho and Strong, Richard D. and Brockman, Jay B. and Tullsen, Dean M. and Jouppi, Norman P. McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Micro-42: Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 469–480, New York, NY, USA, 2009. ACM.
- [11] Liu, Qing. *QUILT PACKAGING: A NOVEL HIGH SPEED CHIP-TO-CHIP COMMUNICATION PARADIGM FOR SYSTEM-IN-PACKAGE*. PhD thesis, Notre Dame, Indiana, USA, December 2007. Chair-Jacob, Bruce L.
- [12] M. Frank Chang and Jason Cong and Adam Kaplan and Mishali Naik and Glenn Reinman and Eran Socher and Sai-Wang Tam. CMP Network-on-Chip Overlaid With Multi-Band RF-interconnect. In *HPCA*, pages 191–202, 2008.
- [13] Mau-Chung Frank Chang and Ingrid Verbaauwhede and Charles Chien and Zhiwei Xu and Jongsun Kim and Jenwei Ko and Qun Gu and Bo-cheng Lai. Advanced RF/Baseband Interconnect Schemes for Inter- and Intra-ULSI Communications. 52:1271–1285, Jul 2005.
- [14] McCalpin, John D. Memory Bandwidth and Machine Balance in Current High Performance Computers. *IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter*, pages 19–25, December 1995.
- [15] Palmer, R. and Poulton, J. and Fuller, A. and Greer, T. and Eyles, J. and Dally, W. and Horowitz, M. A 14-mW 6.25-Gb/s Transceiver in 90-nm CMOS. *IEEE J. Solid-State Circuits*, 42(12):2745–2757, 2007.
- [16] Polka L. A. et al. Package Technology to Address the Memory Bandwidth Challenge for Tera-scale Computing. *Intel Technology Journal*, 11(3):197–206, 2007.
- [17] Tam, Sai-Wang and Chang, M.-C., and Cong, J. and Reinman, G. and Socher, E. RF-Interconnect for Future Network-on-Chip. pages 255–280, 2011.
- [18] Vantrease, Dana and Schreiber, Robert and Monchiero, Matteo and McLaren, Moray and Jouppi, Norman P. and Fiorentino, Marco and Davis, Al and Binkert, Nathan and Beausoleil, Raymond G. and Ahn, Jung Ho. Corona: System Implications of Emerging Nanophotonic Technology. In *ISCA '08: Proceedings of the 35th International Symposium on Computer Architecture*, pages 153–164, Washington, DC, USA, 2008. IEEE Computer Society.
- [19] Wang, David and Ganesh, Brinda and Tuaycharoen, Nuengwong and Baynes, Kathleen and Jaleel, Aamer and Jacob, Bruce. DRAMsim: a memory system simulator. *SIGARCH Comput. Archit. News*, 33(4):100–107, 2005.
- [20] Wayne L. Buckhanan, Michael Niemier, and Gary H. Bernstein. Bridging the HPC Processor-Memory Gap with Quilt Packaging. In *Micro/Nano Symposium (UGIM)*, 2010 18th Biennial University/Government/Industry.