

INTERPROCESSOR COMMUNICATION SEEN AS LOAD-STORE INSTRUCTION GENERALIZATION

Manolis G.H. Katevenis[†]

Institute of Computer Science, Foundation for Research and Technology - Hellas (FORTH-ICS),
Vassilika Vouton, Heraklion, Crete, GR-70013 Greece - Member of HiPEAC
<http://archvlsi.ics.forth.gr/~kateveni/>

ABSTRACT

This paper presents the current (2007) author's views and opinions on interprocessor communication (IPC) and how it should evolve in future multiprocessors, with an attempt to define an IPC architecture that is uniformly extensible from small-scale chip multiprocessors (CMP) to large-scale multi-chip parallel systems. We adopt the viewpoint that all IPC is based on data transfers through an (on-chip or multi-chip) interconnection network; all processors, acceleration engines, and independent memories interface to this network through *lightweight* and *tightly-coupled* network interfaces (NI). Just like traditional (RISC) processor-to-memory communication is via load and store instructions and atomic operation primitives, future IPC should be based on the corresponding data transfer primitives: remote read DMA, remote write DMA, and remote queues. We assume a global address space. By examining the traditional functions of address translation and network routing (and perhaps even caching), we observe that these are closely related, and thus will probably have to be provided through a unified mechanism, whose purpose should be to support data *migration* and *protection*; flow control and congestion management may have to be closely linked to this same mechanism. Two interesting possibilities are to architect NI's so that they detect some events and cause corresponding actions, and to implement cache coherence, when and where provided, on top of the above hardware primitives provided by the NI's.

1. INTERPROCESSOR COMMUNICATION IN PAST AND FUTURE MP'S

Digital processors increasingly work in cooperation, rather than in isolation. Chip multiprocessors (CMP) [1] [2] [3] contribute centrally to this trend. *Interprocessor Communication (IPC)* is the means by which processor cooperation is implemented. The purpose of this paper is to reflect upon future IPC architectures and to provide this author's opinion about them, even if parts of this opinion are based on intuition.

Past systems provided either *fast but non-scalable* IPC through coherent caches, or *scalable but long-latency* IPC through general-purpose interconnection networking. There was an underlying –sometimes in-admitted– assumption that large-scale *and* high-performance IPC was either not needed or not feasible or prohibitively expensive. This paper opposes that view, and tries to outline a *unifying* architecture for high performance IPC both at the small and at the large scale.

1.1. Communication to Computation Ratio

The ratio of (remote) communication to (local) computation varies widely from application to application, just like, e.g., some applications perform no floating-point calculations at all while other applications

[†]also with the Dept. of Computer Science, University of Crete, Greece




Fig. 1. Network Interfaces of the past and their overheads

critically depend on fast floating-point hardware for their performance. Interprocessor communication was expensive, relative to local computation, in the computer systems of the past half century, especially in large-scale multiprocessors (MP's larger than typical cache-coherence diameters). Thus developers architected their applications so as to avoid frequent or extensive communication, and so as to batch communicated data into coarse-grain units.

Some designers observe this property of typical current applications, and conclude that there is no need to support high ratios of communication-to-computation, or fine-grain communication patterns –or at least not so for communication by means other than coherent caches. Such an attitude, however, clearly creates a chicken-and-egg effect: applications will not use extensive, fine-grain communication as long as architectures do not support it, and new architectures will not support it because applications do not use it! In this author's opinion, new architectures should not omit or delay the exploitation of opportunities created by the new technology in the area of high-throughput, low-latency, and fine-grain communication, at all scales –both on-chip and system-wide. Such architectural support for advanced IPC will enable a whole new breed of parallel applications.

1.2. IPC Overheads in Cluster Computing using Network Interfaces of the Past

Large-scale multiprocessing is based, since a decade ago, on cluster computing. In clusters, interprocessor communication (IPC) has to occur through the I/O subsystem, which was designed for the peripheral devices of the past that were much slower than processors or modern interconnection networks. Figure 1 illustrates this placement of the network interface (NI) and its associated overheads; these overheads are listed below together with the methods to eliminate them, based on the research of the last 15 years. Advanced architectures have adopted some such solutions, and widespread adoption must occur in all forthcoming architectures.

- Each I/O operation required the intervention of the operating system in order to provide virtualization of the I/O devices, i.e. transparent sharing of these devices among multiple user processes. Traditional system call overhead is in the hundreds to thousands of clock cycles. To eliminate this, multiple methods for (protected) user-level access to the I/O devices have been demonstrated [4].
- I/O transfers occurred over the I/O bus, with latencies in the tens to hundreds of processor cycles –quite slower than the memory bus. To reduce this overhead, researchers have proposed to tightly couple the network interface to the cache hierarchy [5].
- I/O data had to be copied multiple times –typically at least once between network interface and memory, and once between kernel and user address space. To eliminate this overhead, zero-copy communication protocols have been developed [6].

allows the option to implement some layers of the cache coherence protocols in software, using the underlying data transport primitives provided by the NI hardware.

- When network interfaces are provided *both* next to each processor (level-1 NI) and at the chip boundary (level-2 NI), the IPC protocols of large-scale systems can be *seamlessly extended* to the case where the communicating processors happen to reside on the same chip. These protocols can now feature low-overhead and low-latency, owing to the tight coupling of L1 NI's to the processors. At the same time, the high cost of off-chip communication (increased round-trip times, hence increased buffering) needs only be borne by the L2 NI's.

2. REMOTE DMA: THE GENERALIZATION OF LOAD/STORE INSTRUCTIONS

Communication occurs by transferring (copying) data from a source to a destination. Each transfer operation may concern various amounts of data: a single bit (rarely), or a single byte (inefficient, but used for text of media types), or a few bytes (one word, as with load and store instructions), or tens of bytes (e.g. a cache line), or hundreds of bytes (a small data structure?), or thousands of bytes (e.g. a small page), or millions or more of bytes (e.g. a large page or file). Hardware architects and application developers choose this granularity based on a well-known tradeoff: fine granularity for improved flexibility and space economy, versus coarse granularity for better amortization of operation overhead over a larger amount of work.

2.1. Granularity of Data Transfers as a function of Distance

Within the processor, the overhead of each transfer is the use of one word-wide path for one clock cycle; hence, transfers are performed at the word granularity. Between the processor and its local memory (e.g. level-1 cache, when caches are used), the transfer overhead is again the use of one word-wide path for one clock cycle, plus the consumption of one (word-wide) load or store instruction, plus the transmission, translation (TLB), and checking (against cache tags) of the (word-wide) address. In this case, single-word granularity is again used, for simplicity and flexibility reasons, although the overhead (on the order of two words –one instruction plus one address) could justify a granularity of a few data words per load/store instruction (and some processors do in fact provide such multi-word instructions). Between levels of the cache and memory hierarchy, each transfer incurs the cost of a few clock cycles in the cache controller, a bus transaction (a network-on-chip (NoC) packet, in future systems), and possibly a DRAM row activation (a few tens of nanoseconds). The transfer granularity, here, is the cache line –usually 32 to 128 bytes– which uses a typically 8-byte-wide datapath for 4 to 16 clock cycles; again, granularity is commensurate to overhead.

Besides the above, the only other type of data transfers supported by traditional computer systems were input/output (I/O) transfers, which occurred through the slow path discussed in section 1.2. Under these circumstances, the only other data transfers that traditional systems supported had to be on the granularity of tens of KBytes for them to achieve substantial efficiency.

In the future computing systems, where massive multiprocessing will be the default, all interprocessor communication, including to the storage and I/O devices, will be through the interconnection network (both on-chip and off-chip). The basic overhead of a data transfer through the interconnection network, in terms of buffer space and transmission capacity utilization, is the per-packet header, when compared to the packet's payload; these headers (plus CRC and other line overhead) are, today, on the order of 16 to 32 bytes, per packet. Other overheads, in terms of processing cost and complexity, are the per-packet routing, scheduling, and flow control decisions; there are known techniques, today, to limit all these to a few clock cycles per packet.




Fig. 4. Remote DMA requires single sender per receive buffer, or previous synchronization.

provides RDMA, we consider register-based transfers to only offer a very small added value: given all the other network overheads, the performance gain from saving a few register-memory transfers is quite limited. However, we do consider that RDMA initiation should be as fast as a few accesses to *local* memory (e.g. L1 cache). Remote DMA initiation requires two address parameters, A_{src} and A_{dst} . For the protection and migration support mechanism to properly operate, these arguments should be (global) *virtual* addresses; this is further discussed in section 4.

2.3. Cache Protocols on top of RDMA?

In a multiprocessor with coherent caches, interprocessor communication is traditionally performed through the cache coherence protocol. Observe, however, that the handling of a cache read miss involves the equivalent of a remote read DMA for a block of size equal to one cache line. Also, a cache line write-back operation is the equivalent of a remote write DMA for a block of the same size. Thus, *RDMA-capable hardware may constitute an appropriate substrate for cache protocols to run on top of*. Section 3 (remote dequeue operation, and sections 3.4, 3.5, 3.6) discusses the possibility for network interfaces to trigger actions in response to events. Such a capability could be used to detect cache miss or cache coherence events and trigger corresponding cache line transfers via RDMA.

2.4. One-to-One Communication using Remote DMA

Figure 4 illustrates the RDMA operation, particularly in an environment where multiple parallel transfers exist, and the packets of each transfer may be routed through different paths (“adaptive” or “multipath” routing). For multiple senders, $P1$ and $P2$, to be sending to a same receiver, $P3$, the receiver must have set up *separate memory areas* where each transfer is taking place –otherwise the synchronization overhead between $P1$ and $P2$ would be excessive. Section 3.2 comments on what happens in situations where such separate buffer pre-allocation would be too expensive.

Multipath (adaptive) routing is desirable because it greatly improves network performance; however, multipath routing causes out-of-order delivery –a headache that many architects want to avoid. Remote DMA matches well with multipath routing: since each packet specifies its own destination address (and since destination regions are non-overlapping), it does *not matter* which packet arrives first and which one second –each packet goes and finds its own destination, and when all packets arrive the “puzzle” that they form will have been built. The only problem that remains is to detect when all packets belonging to a same DMA “session” have arrived. In systems guaranteeing in-order delivery, receiver software can detect that by polling a flag in the last byte of the DMA region. Systems that allow out-of-order delivery must provide other mechanisms to detect DMA completion, e.g. counting the number of bytes that have been received.

triggered by the (remote) readers and not by the queue itself. Figure 5 illustrates this by showing dequeue requests being sent by potential readers *at the time when* these readers wish to perform a dequeue. In the figure, requests from processors $P3$ and $P4$ happen to arrive at the same time; each of them, *atomically* dequeues a different (head) packet ($m1$ and $m2$ in the figure, in an unpredictable order).

Implementations of remote readers may vary. One alternative is to treat incoming request packets as *active messages* [10]: upon arrival, the local processor performs the dequeue operation and returns the dequeued message to the requestor. Another alternative is for the network interface *hardware* itself to interpret such request packets and perform the dequeue and dispatch operations [11]; this is analogous to having a finite-state machine (FSM) within the NI play the role of a (simple) processor running the active message code. We referred to such a NI capability as “triggering actions in response to events” in section 2.3, and we further use it in sections 3.4, 3.5, and 3.6. The next subsections describe or comment on several applications of (remote) queues in (asynchronous) I/O processing or in interprocessor communication and synchronization.

3.1. Queues as Generalization of Single-Item Communication Buffers

The most primitive form of producer-consumer communication is illustrated by the traditional teletype (tty) style I/O interface. Consider, for example, a keyboard-input interface consisting of one data register and one status register; the essential part of the status register is the *empty/full bit*, informing the receiver whether or not a new data byte has arrived in the data register since the previous one was read. Such a *single-item* interface imposes a tight interlocking of producer and consumer rates: if the consumer is too slow during even a short time period (a single byte-production period), a data-overflow error occurs. Conversely, if the consumer is too fast, it has to wait, using either polling (busy-wait) or an interrupt mechanism based on the empty/full bit.

Queues are the generalization of single-item interfaces; the larger they can grow, the more freedom they allow on producer-consumer rate-difference fluctuation. When there is a single producer and a single consumer process, it suffices to grow the *data* part of the interface to a multi-item form, e.g. a circular-buffer queue. The *control* part of the interface can still use a single-item form: a single head-tail pointer pair suffices. On the other hand, when there are multiple producers or multiple consumers, the control part of the interface must be generalized as well. The next subsection discusses the multiple-producer case; the multiple-consumer case is analogous, using the remote-dequeue operation, as discussed above.

3.2. Many-to-One Communication

Section 2.4 and figure 4 dealt with multiple producers sending data to a common consumer processor. In order for data transfers to proceed independent of each other (i.e. without needing prior synchronization), separate buffer areas must have been preallocated for each communicating pair. Each buffer area has to have its own control structure, e.g. a head-tail pointer pair, or empty/full bits per sub-block. There are two costs associated with such a setup. One overhead is for the receiver to poll the several buffer areas in order to discover where new data have arrived; we discuss this in section 3.4 below.

Another overhead occurs when the number of *potential* producers is much larger than the number of actual senders. In a system containing e.g. thousands of communicating processors (or threads), we may not know a priori which processor will want to send information to a given receiver P_r , and it would be too expensive to preallocate separate buffer areas at P_r for each and every other processor in the system. *Remote queues* (RQ) for *control* information is the proper solution in this case: a single RQ is set up at P_r for receiving communication requests by other processors. For every such request received, P_r allocates a dedicated buffer area, and responds with a pointer (in global address space) to this area. The requesting processor, then, uses remote DMA to put its data into this buffer.

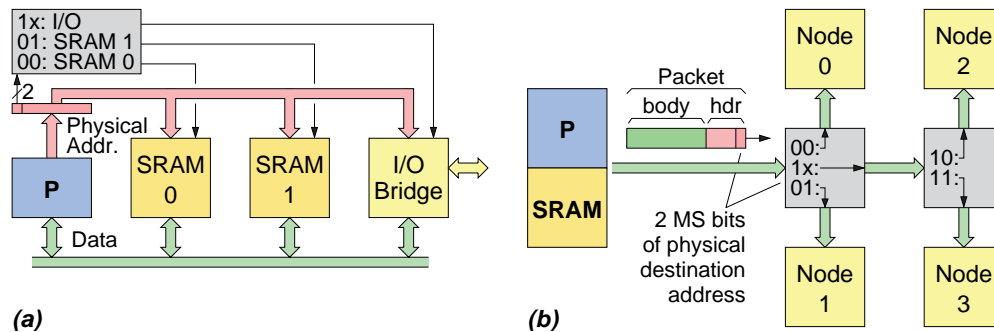


Fig. 6. (a) Physical address decoding in a uniprocessor; (b) geographical address routing in a MP.

the data structure; request messages are remotely enqueued into it. Either a local thread or the local NI interprets these request messages: for each of them, ship the data structure (e.g. via RDMA), then wait for a response (equivalent to releasing the lock), then service another request message, and so on.

3.6. Barrier Synchronization

A small-scale barrier synchronization can be implemented with one node (the “root”) collecting phase-completion notifications from all participating nodes, and then in turn notifying all of them when that happens. A large-scale barrier should use a tree of collecting nodes and a tree of notifying nodes, in order to avoid the root becoming a communication bottleneck. Collecting and counting notifications –either at a tree node or at the root– can be conveniently done through a queue: each participant remotely enqueues a phase-completion notification. A computation thread, running on the processor that holds the queue, dequeues and counts the messages; when the prescribed count is reached, the thread sends a notification message to its parent node in the tree, or, if it is the root, to all its children. Alternatively, the barrier could be implemented completely in hardware, if the NI’s can trigger actions in response to events, as discussed above: each queue may be set up so that it counts packet arrivals. In the present setup, an event should be signaled not for every packet arrival but when the number of arrivals reaches a prescribed count (just like a RDMA completion event occurs when the number of byte arrivals reaches a prescribed count). When this “event” occurs, the NI generates and sends one or more notification messages to other appropriate queues.

4. NETWORK ROUTING AS GENERALIZATION OF ADDRESS DECODING

Load and store instructions specify a memory address each, thus directing the data transfer to a specific memory word in a specific memory area. Figure 6(a) illustrates this “directioning” in a traditional uniprocessor, with a simple example of a memory system consisting of two SRAM chips (or blocks) and one (memory-mapped) I/O device. The most significant (MS) address bits are decoded and select a chip, then inside each chip a similar addressing tree directs the data transfer to a specific word in it. If we consider the data bus as a network connecting the processor to each word in memory, then the role of the memory address is to route the “access packet” to a specific destination reachable through that network. In this example, the address under consideration is the *physical* address of the load or store instruction, because it is the address that leads the transfer to a specific physical location in the memory or I/O circuitry.

Figure 6(b) illustrates the analogous situation for IPC –rather than load or store instructions– in a multiprocessor. A network packet, carrying part of a remote DMA or a remote enqueue message, specifies

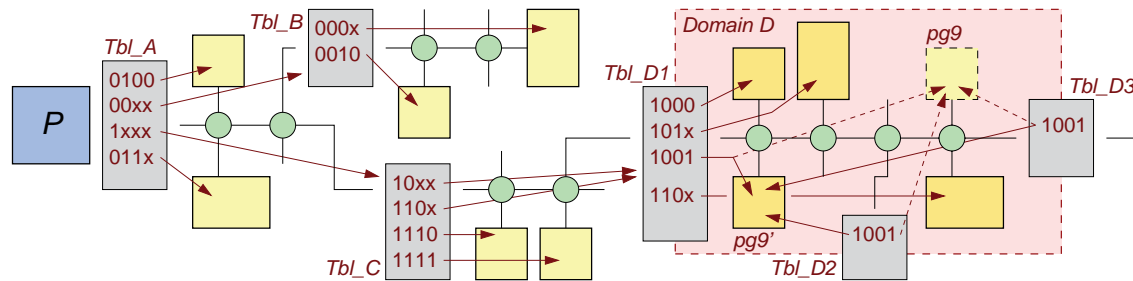


Fig. 8. Progressive translation/routing: packets carry virtual addresses, tables provide physical route (address) for the next few steps.

the translation table by restricting the number of locations that each logical item is allowed to migrate to (using set associativity); on the other hand, it introduces the cost of *searching* for the desired item by performing multiple tag comparisons –often in parallel, consuming energy. Cache lookup is *not* scalable to cases where data are allowed to migrate far away: it requires a broadcast of the search tag and multiple comparisons.

4.2. Progressive Translation for Localized Updates upon Migration

Old uniprocessors performed virtual-to-physical address translation at only one place –the processor’s TLB. Newer systems need multiple TLB’s –one per processor, plus a TLB in the network interface, in order to provide user-level access to the NI: the user must provide *virtual* addresses as arguments, and the NI has to perform the (protected) translation to physical addresses for the operation to be executed. When a data item (e.g. a page) migrates, *all* copies of the translation table or portions thereof (e.g. TLB’s) need to be updated, which is a problem as this number increases. For scalable multiprocessing, this system has to be revised: we cannot afford all virtual-to-physical address translation to occur at the source node of an IPC packet, because that would mean that all nodes (may) know the current (physical) location of all (virtual) data objects, hence all nodes may have to be notified when an object is migrated.

Figure 8 illustrates a multi-level translation scheme that can solve this problem; let us call it *progressive translation*. Each packet (imagine that it starts from processor P) may pass through several translation tables on its way to its destination. The packet *always* carries its *virtual* destination address. Each translation (routing) table directs the packet for a few more steps (through a few more switches –the round devices in the figure) towards its final destination. Each table contains accurate (definitive) location (routing) information for the data items (e.g. pages) that are close to it, but only approximate information for those that are further away: “go to that table and ask there”, it says...

Imagine, in figure 8, that (virtual) page 1001 migrates from (physical) page-frame $pg9$ to frame $pg9'$, and that, while doing so, it stays within the same *domain* (neighborhood) D . Then, only the translation/routing tables at the entries of and inside D need to be updated –assuming that all external tables merely point to an entry point into D , rather than specifying exact location within D .

Figure 8 also illustrates a couple of practical aspects of progressive translation. The function of the tables can be interpreted to be to provide *run-time configurable* network routing, rather than fixed “geographic” routing: each destination address is allowed to change its location in the net. However, it would be too expensive, in general, to provide this capability inside each network switch –especially so inside the switches of networks-on-chip (NoC). The figure shows individual switches as little circles, and shows translation/routing tables positioned every few switches –but not in every single switch. Each table provides *physical* route information through a sub-net –but not through the entire, system-wide network.

6. ACKNOWLEDGMENTS

The viewpoint and opinions expressed in this paper have been reached by the author following the discussions and the work performed within the “Scalable Computer Architecture (SARC)” integrated project #27648 of FP6, supported by the European Commission. These discussions have been led and steered by Stamatis Vassiliadis; the project itself owes its existence to Stamatis Vassiliadis. Stamati, *thank you!* –I miss you; our project and our scientific community misses you; Samos is missing you. It is too bad that you departed so soon –you could have helped us much more, we would have enjoyed your company much more. Καλό ταξίδι αγαπημένε συμπατριώτη, συνεργάτη, φίλε –farewell dearest compatriot, colleague, friend; we will remember you forever!

Other colleagues who helped formulate these views include Stamatis Kavadias, Georgi Gaydadjiev, Michael Papamichael, Angelos Bilas, Babak Falsafi, Christos Sotiriou, Spyros Lyberis, Vasilis Papaefstathiou, George Kalokerinos, Manolis Marazakis, Angelos Ioannou, Jose Duato, and Ian Johnson. Many thanks to all of them!

7. REFERENCES

- [1] D. Pham et al., “The design and implementation of a first-generation CELL processor,” in *Proc. IEEE Int. Solid-State Circuits Conference (ISSCC)*, Feb. 2005.
- [2] P. Kongetira, K. Aingaran, and K. Olukotun, “Niagara: A 32-way multithreaded SPARC processor,” *IEEE Micro*, vol. 25, no. 2, pp. 21–29, Mar. 2005.
- [3] Intel, “World’s first quad-core processors for desktop and mainstream servers,” in <http://www.intel.com/quad-core/>.
- [4] S. Mukherjee and M. Hill, “A survey of user-level network interfaces for system area networks,” in *Tech. Report 1340, Computer Sci. Dept., Univ. of Wisconsin, Madison USA*, 1997.
- [5] S. Mukherjee, B. Falsafi, M. Hill, and D. Wood, “Coherent network interfaces for fine-grain communication,” in *Proc. 23rd Int. Symposium on Computer Architecture (ISCA’96)*, Philadelphia, PA USA, May 1996, pp. 247–258.
- [6] D. Dunning, G. Regnier, G. McAlpine, D. Cameron, B. Shubert, F. Berry, A. M. Merritt, E. Gronke, and C. Dodd, “The virtual interface architecture,” *IEEE Micro*, vol. 18, no. 2, pp. 66–76, 1998.
- [7] G. Gaydadjiev, “Personnal communication,” 2006.
- [8] S. Keckler, W. Dally, D. Maskit, N. Carter, A. Chang, and W. Lee, “Exploiting fine-grain thread level parallelism on the MIT Multi-ALU processor,” in *Proc. 25th Int. Symposium on Computer Architecture (ISCA’98)*, June 1998, pp. 306–317.
- [9] E. Brewer, F. Chong, L. Liu, S. Sharma, and J. Kubiawicz, “Remote queues: Exposing message queues for optimization and atomicity,” in *Proc. 7th ACM Symposium on Parallel Algorithms and Architectures (SPAA’95)*, Santa Barbara, CA USA, June 1995, pp. 42–53.
- [10] T. von Eicken, D. Culler, S. Goldstein, and K. Schauer, “Active messages: A mechanism for integrated communication and computation,” in *Proc. 19th Int. Symposium on Computer Architecture (ISCA’92)*, Gold Coast, Australia, May 1992, pp. 256–266.
- [11] S. Kavadias, “Network interface support for synchronization primitives,” in *HiPEAC ACACES summer school Poster Session, L’Aquila, Italy*, July 2006.
- [12] M. Katevenis, I. Mavroidis, G. Sapountzis, E. Kalyvianaki, I. Mavroidis, and G. Glykopoulos, “Wormhole IP over (connectionless) ATM,” *IEEE/ACM Trans. Networking*, vol. 9, no. 5, pp. 650–661, Oct. 2001.
- [13] M. Ruiz-Sanchez, E. Biersack, and W. Dabbous, “Survey and taxonomy of IP address lookup algorithms,” *IEEE Network*, vol. 15, no. 2, pp. 8–23, Mar. 2001.