Can User-Level Protocols Take Advantage of Multi-CPU NICs?

Piyush Shivam
Dept. of Comp. & Info. Sci.
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210
shivam@cis.ohio-state.edu

Pete Wyckoff
Ohio Supercomputer Center
1224 Kinnear Road
Columbus, OH 43212
pw@osc.edu

Dhabaleswar Panda
Dept. of Comp. & Info. Sci.
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210
panda@cis.ohio-state.edu

Abstract

User-level protocols and their implementations on smart and programmable network interface cards (NICs) have been alleviating the communication bottleneck for high speed interconnets . Most of the user-level protocols developed so far have been based on single-CPU NICs. One of the more popular current generation Gigabit Ethernet NICs includes two CPUs, though. This raises an open challenge whether performance of user-level protocols can be improved by taking advantage of a multi-CPU NIC. In this paper, we analyze the intrinsic issues associated with such a challenge and explore different parallelization and pipelining schemes to enhance the performance of our earlier developed EMP protocol for single-CPU Alteon NICs. Performance evaluation results indicate that parallelizing the receive path of the protocol can deliver 964 Mbps of bandwidth, close to the maximum achievable on Gigabit Ethernet. This scheme also delivers up to 8% improvement in latency for a range of message sizes. Parallelizing the send path leads to 17% improvement in bidirectional bandwidth.

1. Introduction

High-performance computing on a cluster of workstations requires that the communication latency be as small as possible . This has led to the development of a range of user-level network protocols: FM [7] for Myrinet, UNet [14] for ATM and Fast Ethernet, GM [2] for Myrinet, our recent work on EMP [10] for Gigabit Ethernet, etc. [1] During the last few years, the designs and developments related to user-level protocols have been brought into an industry standard in terms of the *Virtual Interface Architecture* (VIA) [13]. An extension to the VIA interface is already included in the latest *InfiniBand Architecture* (IBA) [3] as the Verbs layer.

As processor technology is moving towards gigahertz speeds and network technology is moving towards 10–30 Gbits/sec [3] it is becoming increasingly important to exploit the capabilities of the NIC to achieve the best possible communication performance. In current generation systems, the PCI bus serves as a fundamental limitation to achieving better communication performance. This aspect is being alleviated in the IBA standard. Thus, the design of the NICs and their interfaces are getting increased attention.

Most of the older and current generation NICs support only one processor. Thus, to the best of our knowledge, all user-level communication protocols including VIA implementations have been centered around single-CPU NICs. One popular current generation NIC design is the two-CPU core from Alteon [6] for Gigabit Ethernet. This leads to the following interesting challenges:

1. Can user-level protocols be better implemented by taking advantage of a multi-CPU NIC?

2. What are alternative strategies to parallelization and pipelining of user-level protocols with a two-CPU NIC, and what are the intrinsic issues?

3. How much performance benefit can be achieved with such parallelization and pipelining?

In this paper, we analyze, design, implement, and evaluate a parallel version of the user-level protocol layer with two-CPU Alteon NICs for Gigabit Ethernet. We enhance the *Ethernet Message Passing* (EMP) [10] protocol which we have recently developed for Gigabit Ethernet using Alteon NICs. EMP was developed by taking into account only one of the two available CPUs in the NIC. First we analyze the send and receive paths of the EMP messaging layer to determine the costs associated with the basic steps. Next, we analyze the challenges involved in parallelizing and/or pipelining user-level protocols for the two-CPU Alteon NIC. This leads to four alternative enhancements: splitting up the send path only (SO), splitting up the receive path only (RO), splitting both the send and receive

paths (SR), and assigning dedicated CPUs for send and receive (DSR).

We implement these strategies on our cluster testbed with 933 MHz Intel PIII systems and evaluate their performance benefits. The best results are obtained with the RO scheme for unidirectional traffic, giving a small message (10 bytes) latency of 22.62 us and bandwidth of 964 Mbps. This is compared to the base case latency of 24.31 us (a gain of 7.0%) and bandwidth of 840 Mbps. For large messages the latency improvement was around 8.3%. For bidirectional traffic the best results were achieved with the SO scheme where the total bandwidth peaked at 1100 Mbps as compared to 940 Mbps in the base case, a gain of 17%.

2. Architectural overview

Alteon Web Systems, now owned by Nortel Networks, produced a Gigabit Ethernet network interface chipset based around a general purpose embedded microprocessor design which they called the Tigon2.

The Tigon chip is a 388-pin ASIC consisting of two MIPS-like microprocessors running at 88 MHz, an internal memory bus with interface to external SRAM, a 64-bit, 66 MHz PCI interface, and an interface to an external MAC. The chip also includes an instruction and data cache, and a small amount of fast per-CPU "scratchpad" memory. Hardware registers can be used by the processor cores to control the operation of other systems on the Tigon, including the PCI interface, a timer, two host DMA engines, transmit and receive MAC FIFOs, and a DMA assist engine. Our particular cards have 512 kB of external SRAM, although implementations with more memory are available.

The hardware provides a single semaphore which can be used to synchronize the two CPUs. Each CPU has its own register which it writes with any value to request ownership of the semaphore, then must loop until a read from the semaphore register is non-zero, indicating successful ownership. This is the only general locking mechanism available at the NIC. For more details see [9].

3. Overview of the EMP protocol

In this section we provide an overview of the implementation of the EMP protocol [10]. We first provide an overview of the basic steps. Next, we discuss these steps in detail. Finally, we present a timing analysis of these steps on a single-CPU NIC. The description of the steps and the timing analysis will help us understand the challenges involved in parallelizing the EMP protocol.

3.1. Basic protocol steps

Here we outline the basic steps happening at the sending side and the receiving side of the EMP protocol. The

sending side performs the following steps.

Send bookkeeping. Send bookkeeping refers to the operations which take place for preparing the frame for being sent. The bookkeeping operations can be outlined as:
Handle posted transmit descriptor: This step is initiated by the host which operates asynchronously with the NIC. The introduction of each new transmit request leads to the rest of the operations. This operation takes place per message.
Message fragmentation: The host desires to send a message, which is a user-space entity corresponding to some size of the application's data structures. The NIC must fragment this into frames, which is a quantity defined by the underlying ethernet hardware as the largest quantum of data which can be supported in the network, 1500 bytes in our system.
Initialize transmission record: Each message which enters the transmit queue on the NIC is given a record in a NIC-resident table which keeps track of the state of that message including how many frames, a pointer to the host data, which frames have been sent, which have been acknowledged, the message recipient, and so on.

Transmission. The steps involved in transmission are:
DMA from HOST to the NIC: Along with the two DMA channels the NIC also contains the "DMA Assist" state machine to help off-load some of the tasks from the NIC processor. Once the bookkeeping steps for the frame are over, the DMA assist engine will queue a request for data from the HOST. When the transfer has completed, it will automatically tell the MAC to send the frame. The transfer is made in the send buffer which is updated after each transfer. This set of operations takes place for every frame and hence will take more time for large message sizes.
MAC to wire: The NIC uses MAC transmit descriptor to keep track of frames being sent to the serial Ethernet interface. The MAC is responsible for sending frames to the external network interface by reading the associated MAC transmit descriptor and the frame from the local memory buffer. This operation happens per frame and each frame uses one MAC descriptor, hence the overhead incurred will increase with increasing message size.

Receive acknowledgment. Once the sender knows that the receiver has successfully received the frames it can release the resources related to the sent data. Receive acknowledgment introduces only minimal per-frame overhead, again, because acknowledgment is a process which applies only to groups of frames [10].

Similarly the receiving side steps are:

Receive bookkeeping. The receive bookkeeping refers to the operations which need to be performed before the frame can be sent to the host. These operations are:
Handle preposted receive descriptor: This step is initiated by the host for all the messages expected to arrive in future.

Here the state information which is necessary for matching an incoming frame is stored at the NIC. If a frame arrives and does not find a matching preposted descriptor, it is dropped to avoid buffering at the NIC [10].

Classify frame: This step looks at the header of each incoming frame and identifies if it is a header frame, data frame, acknowledgment frame or negative acknowledgment. It also identifies the preposted receive to which the incoming frame belongs by going through all the preposted records. In the process it also identifies if the frame has already arrived and, if so, drops it. Classify frame is performed for every frame and hence the overhead per message increases with increasing message size.

Process frame: Once the frame has been correctly identified in the previous step, the frame header information (message sequence number, frame sequence number, etc.) is stored in the receive data structures for reliability and other bookkeeping purposes. After recording this information, this step also initiates the DMA of the incoming frame data. This step is done per frame and the overhead increases as the message size increases.

Receiving. The step comprising the actual receiving process involves the following operations:

Wire to MAC: Similar to transmission, the NIC uses MAC receive descriptors to keep track of frames being received from the serial Ethernet interface. Error conditions are monitored during frame reception and reported to the firmware through the status word located in the descriptors. Before the data is given to the NIC the 32-bit CRC is verified and noted in the status word.

NIC to HOST: Here the "DMA Assist" engine comes into play exactly like in the transmit case but in the reverse direction.

Send acknowledgment. This step involves a combination of bookkeeping and transmission. The acknowledgment is sent as a single frame with some control information but no data. Hence the overhead involved in this step is not as large as that for any data frame. Moreover, this does not involve per-frame overhead because an acknowledgment is sent only for complete groups of frames.

3.2. Timing analysis

We did a complete profiling of our protocol to find out how much time is spent in each of the steps. As we discussed, each of the steps consists of one or more operations. But for the sake of clarity we are showing only the timings for the major steps. Table 1 shows the analysis. The experimental platform is described in Section 6.

Receive bookkeeping is more expensive than send bookkeeping because while sending, the frames are sent in order but they can arrive out of order on the receive side (due to switch dropping and reordering of frames). So extra effort is needed per frame to accept these out of order frames and put them in the correct order. Also, the tag matching for MPI happens at the receiver.

Table 1. Timing analysis for the major functional operations.

Operation	Time (us)
Send bookkeeping	5.25
Transmission	5.50
Receive acknowledgment	5.75
Recv bookkeeping	10.50
Receiving	2.75
Send acknowledgment	2.50

4. Challenges of using multi-CPU NIC

In order to take advantage of a multi-CPU NIC, the basic protocol steps need to be distributed across the two processors. However, these steps need to share some common state information at some point in the execution. Typically, NICs have limited hardware resources to assist in this operation without introducing additional overhead. Here, we take a critical look at the limitations of the Alteon NIC and the potential alternatives for achieving our objective.

4.1. NIC constraints

The Alteon NIC does not provide hardware support for concurrency. There is only one lock, hence fine-grained parallelism is expensive. Coarse-grained parallelism is inappropriate for the kind of operations performed at the NIC, due to its limited resources. Shared resources (MAC, DMA) do not have hardware support for concurrency, and use the only available lock, thus overloading that single semaphore.

4.2. Achieving concurrency

There are many scenarios where even after distributing the send and receive functions on different processors, the sending and receiving state information needs to be shared.

To minimize sharing of such state one may keep separate data structures for send bookkeeping and receive bookkeeping so that both the operations can happen in parallel without needing to access the other's data structure. However, this cannot be guaranteed for every case.

One way to solve this problem would be to share the data structures across the CPUs. However this would mean that each access to the data structure requires synchronization, which would be very expensive as the data structures are accessed frequently.

To reduce the synchronization overhead, the bookkeeping data structures can be fine-grained so that locking one data structure does not lead to halting of other operations which can proceed using other unrelated data structures.

One may also accomplish synchronization by allocating a special region in the NIC SRAM where one CPU would write the data needed by the other CPU, which would then

read the common data from there. This might be a better option because in this case the overhead is generated only when there is a need for sharing data between the CPUs.

4.3. Exploiting pipelining and parallelization

Amdahl's law states that the speed-up achievable on a parallel computer can be significantly limited by the existence of a small fraction of inherently sequential code which cannot be parallelized. In any reliable network protocol there will be a lot of steps which have to be executed sequentially. In fact, serially constrained operations become the norm. This puts a limit on the amount of work which can be scheduled in parallel. This limitation forces us to think about the underlying implementation and make appropriate changes so that we can perform the maximum number of operations in parallel. In addition to parallelization, pipelining can also be exploited, where the operations happen one after another but not in parallel. In this paper, we explore both pipelining and parallelization to enhance the performance of user-level protocols with multi-CPU NICs.

5. Schemes for parallelization and pipelining

In this section, we propose and analyze alternative schemes to enhance the performance of the EMP protocol with the support of a two-CPU NIC. The basic approach was to distribute the major steps of send and receive paths to achieve a balance of work on the two processors. This break-up was done with the goal of achieving pipelining or parallelism—whichever would be possible. We tried to achieve the latter as much as possible but were limited by the inherent sequentiality of the protocol in many cases.

We analyzed the send path and the receive path for parallelization based on our timing analysis and recognized the following four alternatives.

SO: The send path only is split-up across the NIC CPUs.
RO: The receive path only is split-up across the NIC CPUs.
DSR: Dedicated processors for send and receive paths.
SR: Both send and receive path are split-up.

For each of these alternatives, we illustrate how different components (steps) are distributed over two processors at both the sending and receiving sides. We compare our schemes with the base case scheme where all the steps happen at the same processor.

SO. The split-up of the send path in SO happens as shown in Table 2. Here, we aim to achieve pipelining by running the bookkeeping phase of a later message with the transmission phase of an earlier message for a unidirectional flow. There is some parallelism also happening at the receiver between send ack (2.50 us) and recv dma (1.25 us). The receive path for SO remains the same as in the base case. One needs to distinguish the difference between the

Table 2. Component distribution - sending side(top), receiving side (bottom) for unidirectional traffic.

	A	(us)	B	us)
SO	send bkkeep	5.25	txmission	5.50
	recv ack	5.75		
RO	send bkkeep	5.25		
	txmission	5.50		
	recv ack	5.75		
DSR	send bkkeep	5.25	recv ack	3.25
	txmission	5.50		
	recv ack	2.50		
SR	send bkkeep	5.25	txmission	5.50
	recv ack	5.75		
	A	(us)	B	(us)
SO	recv bkkeep	6.25	send ack	2.50
	recv frame	4.25		
	recv dma	2.75		
RO	recv bkkeep	6.25	recv frame	4.25
	send ack	2.50	recv dma	2.75
DSR	send ack	2.50	recv bkkeep	6.25
			recv frame	4.25
			recv dma	2.75
SR	recv bkkeep	6.25	recv frame	4.25
			send ack	2.50
			recv dma	2.75

receive path and receive side. The receive path is made up of receive bookkeeping and actual receiving (DMA). The send acknowledgment on the receiver is a part of the send path since the send ack involves steps used in sending and not receiving.

RO. The split-up of functions in RO happens as shown in Table 2. The send ack (2.50 us) happens in parallel with the recv dma (2.75 us) and a part of recv bkkeep (4.25 us). The split-up of receive bookkeeping helps in achieving pipelining also. We are able to achieve a very good balance of functions on the receiving side. The send path remains the same as in the base case. Again, similar to the receive path scenario one needs to distinguish between the send path and sending side. The sending side has a receive step happening which is a part of the receive path.

DSR. In this case, we are dedicating one CPU each for the send path and the receive path on the sending as well as receiving side. This way we achieve an almost complete split of the send and receive paths. The recv ack step is split on the sending side because a part of it needs to update the send data structures and hence it is scheduled at the send processor. The functions are distributed as shown in Table 2.

SR. Here we combine the optimized send path and receive path together to see if we can benefit from the overall op-

timization of the protocol. It is a combination of SO and RO as depicted in Table 2. We hope to gain from the benefits of pipelining on the send side and parallelization on the receive side.

6. Implementation and performance

For the Gigabit Ethernet tests, we used two dual 933 MHz Intel PIII systems, built around the ServerWorks LE chipset which has a 64-bit 66 MHz PCI bus, and unmodified Linux 2.4.2. Our NICs are Netgear 620, which have 512 kB of memory. The machines were connected back-to-back with a strand of fiber.

6.1. Exploiting the NIC hardware capability

To solve the problem of synchronization we allocated a special common area in the NIC SRAM though which the CPUs can communicate common data. The benefits of such an approach were discussed in section **??**.

We developed a pair of calls, `spin_lock` and `spin_unlock`, which are used to gain exclusive access for protected code regions. We would have preferred to have multiple points of synchronization to implement object-specific locking, but the hardware provides exactly one point for inter-CPU synchronization through a semaphore. Thus accesses to protected regions become potentially very expensive due to high contention for this single lock.

The other communication mechanism we used was to set bits in the event register of each processor. These calls use spin locks to guarantee exclusive access to the event register, whereby one CPU sets a bit in the event register of the other. The second CPU will notice this event in its main priority-based dispatch loop, clear the bit, and process the event.

Running two processors simultaneously puts more load on the memory system in the NIC. We attempt to alleviate this pressure somewhat by moving frequently used variables to the processor-private "scratchpad" memory area in each CPU. This small region (16 kB on cpu A, 8 kB on cpu B) also has faster access times, so we put frequently-called functions there too.

6.2. Results and discussion

In this section we analyze the results derived from the alternatives discussed so far. We tested each of our alternatives for unidirectional as well as bidirectional flows. For unidirectional flow, we evaluated latency as well as bandwidth. For bidirectional flow, we evaluated bandwidth. We get better performance than the base case (single CPU per NIC) by using at least one alternative in each of the cases.

6.2.1. Unidirectional traffic

The latency is determined by halving the time to complete a single ping-pong test. The "ping" side posts two descriptors: one for receive, then one for transmit, then a busy-wait loop is entered until both actions are finished by the NIC. Meanwhile the "pong" side posts a receive descriptor, waits for the message to arrive, then posts and waits for transmission of the return message. This entire process is iterated 10 000 times from which an average round-trip time is produced, then divided by two to estimate one-way latency.

The unidirectional throughput is calculated from one-way sends with a trailing return acknowledgment. The user-level receive code posts as many receive descriptors as possible (about 400), and continually waits for messages to come in and posts new receives as slots become available. The transmit side posts two transmit descriptors so that the NIC will always have something ready to send, and loops waiting for one of the sends to complete then immediately posts another to take its place. Each transmit is known to have completed because the receiving NIC generates an acknowledgment message which signals the sending NIC to inform the host that the message has arrived. This is iterated 10 000 times to generate a good average.

SO. The unidirectional bandwidth (Figure 1) is the same as in the base case. On the sending side the benefit is obtained due to pipelining (between send bkkeep and txmission) and parallelism (recv ack and txmission).

However the gains are offset by two factors. First, the overhead in inter-CPU communication. Next, while comparing the receive and the send path, we can see that the receive path has more overhead than the send side. The sending side cannot send at any rate since it will swamp the receiver. It waits for the acknowledgment from the receiving side for a certain number of messages (two in our case) before it sends out more messages. Thus whatever speedup which can be gained due to pipelining or parallelism is limited by the reception of acknowledgment from the receiver which is again dependent on receive processing. Since the receive processing is happening in the same way as the base case except for a very small amount of parallelism (which is offset by the inter-CPU communication overhead), the pipelining/parallelization does not demonstrate much benefit for the SO case.

RO. The unidirectional bandwidth (Figure 1) is much better than the base case. In fact it reaches up to 99.78% of the theoretical throughput limit on Gigabit Ethernet (taking into account the required preamble and inter-frame gap and our protocol headers). This happens because on the receiving side, we are able to parallelize send ack on CPU A and recv frame and recv DMA on CPU B (Table 2). In addition, the

Figure 1. Latency & Bandwidth comparisons for unidirectional and bidirectional traffic. The x-axis indicates message size in kilobytes. The y-axis shows latency(rightmost) and bandwidth.

distribution of jobs on the receiving side is well balanced, resulting in both the CPUs being occupied most of the time.

Since the receive path has been parallelized effectively we are able to offset the inter-CPU communication overhead and achieve almost the maximum possible bandwidth. This implies that the receive side is the bottleneck, a fact confirmed by the SO case where we left the receive side almost unaltered and did not achieve any benefits even though we had pipelined and parallelized the send path.

By looking at Figure 1 we observe that even the latency improves for RO parallelism. For a 10-byte messages we obtained a a gain of about 7%. For a message size of 14 kB we were able to achieve a latency improvement of about 8.3%, indicating that the rate of latency improvement increases with increasing message size. The large message case results are in [9].

DSR. The unidirectional bandwidth (Figure 1) is marginally better than the base case. Here the scenario is very similar to the SO case, with the receive side being the bottleneck. However, since on the receive side we do schedule send ack to happen on a different CPU, we are able to see the marginal improvement in bandwidth numbers. The improvement is marginal because send ack is a very small portion of the receive side processing.

By looking at Figure 1 we can conclude that latency also benefits with this approach. For a 10-byte message, we get a latency improvement of 6%.

SR. This alternative gives the best unidirectional bandwidth. One is able to achieve almost complete utilization of Gigabit Ethernet's bandwidth. The results are very similar to the RO case but one can see the benefits of pipelining/parallelizing the send path also in the SR case (Fig-

ure 1).

6.2.2. Bidirectional traffic

Bidirectional throughput is calculated in a manner similar to the unidirectional throughput, except both sides are busy sending to each other. After the startup pre-posting of many receive descriptors, the timer is started on one side. Then two messages are initiated at each side, and a main loop is iterated 10 000 times which consists of four operations: wait for the oldest transmit to complete, wait for the oldest receive to complete, post another transmit, post another receive. Using one application rather than two on each host ensures that we do not suffer from operating system scheduler decisions.

Bidirectional traffic is more complex than unidirectional traffic. A lot more steps are happening at the sender and the receiver as compared to the unidirectional case. These steps are send bookkeeping, transmission, acknowledgment receive, sending acknowledgment, receive bookkeeping and dmaing of the receive as well as send data.

Figure 1 shows considerable improvement for all the alternatives. This is so because of more number of steps and hence more opportunities to parallelize/pipeline in a bidirectional traffic as compared to the unidirectional traffic. In the bidirectional traffic we encounter the same interplay of factors which was present in the unidirectional traffic. However, for large message sizes the inter-CPU communication overtakes the gain. This does not happen in the unidirectional case because the amount of inter-CPU communication is less as compared to the bidirectional case. The detailed analysis for each case can be looked in [9].

7. Related work

All the previous efforts to parallelize the network protocols have been focussed on WAN protocols such as TCP/IP and Symmetric Multi Processor (SMP) systems. These approaches involve processing multiple packets in parallel [5], executing different layers in parallel for TCP/IP like protocols [15] and breaking the protocol functions across multiple processors [8].

It is to be noted that all these research directions focus on host-based layered protocols and try to exploit SMP systems. The focus of our work is how to take advantage of multi-CPU NIC to enhance user-level protocols. The key idea is to free the host CPUs as much as possible to devote its cycles to the parallel application itself.

8. Conclusions and future work

In this paper, we have presented how to take advantage of a multi-CPU NIC, as available in Alteon NIC core implementations, to improve point-to-point communication performance on Gigabit Ethernet. We have considered our earlier developed EMP protocol (valid for single-CPU NIC) and analyzed different alternatives to parallelize and pipeline different steps of the communication operation. The study shows that parallelizing the receive path can deliver benefits for unidirectional as well as bidirectional traffic. In fact, this scheme allows us to reach the theoretical throughput of the medium.

As a result of our investigations into work distribution strategies on the multi-CPU Alteon NIC, we have determined multiple promising paths for future study. Currently the split-up of tasks in done at compile time. We would like to produce a truly dynamic event scheduling system, where the next available event is handled by either processor when it becomes free.

We are also exploring the benefits of multi-CPU NICs to support collective communication operations efficiently. We plan to perform application-level performance evaluation for our scheme. Another interesting direction is to explore the required architectural support at the NIC to parallelize both point-to-point and collective communication operations.

References

[1] M. Banikazemi, V. Moorthy, L. Herger, D. Panda, and B. Abali. Efficient Virtual Interface Architecture support for the IBM SP switch-connected NT clusters. In *IPDPS*, May 2000.

[2] N. Boden, D. Cohen, and R. Felderman. Myrinet: a gigabit per second local-area network. *IEEE Micro*, 15(1):29, February 1995.

[3] Infiniband. http://www.infinibandta.org.

[4] MVIA. http://www.nersc.gov/research/FTG/via, 1998.

[5] E. Nahum, D. Yates, J. Kurose, and D. Towsley. Performance issues in parallelized network protocols. In *Proceedings of the First USENIX Symposium on OSDI*, pages 125–137, November 1994.

[6] Netgear. http://www.netgear.com/adapters_main.asp.

[7] S. Pakin, M. Lauria, and A. Chien. High performance messaging on workstations: Illinois Fast Messages (FM) for Myrinet, 1995.

[8] T. Porta and M. Schwartz. A high-speed protocol parallel implementation: design and analysis. In *Fourth IFIP TC6.1/WG6.4 International Conference on High Performance Networking*, pages 135–150, December 1992.

[9] P. Shivam, P. Wyckoff, and D. Panda. Can user level protocols take advantage of multi-CPU NICs? Technical Report OSU-CISRC-08/01-TR09, The Ohio State University, 2001.

[10] P. Shivam, P. Wyckoff, and D. Panda. EMP: Zero-copy OS-bypass NIC-driven gigabit ethernet message passing. In *Proceedings of SC01*, November 2001.

[11] E. Speight, H. Abdel-Shafi, and J. Bennett. Realizing the performance potential of a virtual interface architecture. In *Proceedings of the International Conference on Supercomputing*, June 1999.

[12] TechFest. http://www.techfest.com/networking/lan/ethernet2.htm.

[13] VI. http://www.viarch.org, 1998.

[14] T. von Eicken, A. Basu, V. Buch, and W. Vogels. U-net: A user-level network interface for parallel and distributed computing. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles*, December 1995.

[15] M. Zitterbart. High-speed protocol implementations based on a multiprocessor architecture. In *Protocols for High Speed Networks*, pages 151–163, 1999.