PGAS Model for the Implementation of Scalable Cluster Systems*

Juan A. Villar, Francisco J. Alfaro, José L. Sánchez
DSI - Univ. of Castilla-La Mancha
02071 - Albacete, Spain
{juanan, falfaro, jsanchez}@dsi.uclm.es

José Duato
DISCA - Tech. Univ. of Valencia
46022 - Valencia, Spain
jduato@disca.upv.es

Abstract

This paper introduces an extended version of the traditional Partitioned Global Address Space (PGAS) model, for the implementation of scalable cluster systems, that the HyperTransport Consortium Advanced Technology Group (ATG) is working on. Using the Simics and GEMS simulators, we developed a software module that approximates the behaviour of a PGAS cluster. This approach mainly provides the simplest mechanism to evaluate how much the PGAS infrastructure will affect overall the application performance. The aim of this work is to study the feasibility of the ATG's PGAS model, for running applications with high memory requirements. It is not focused on measure the performance of applications, although some performance measures are given. Such a model, will let manufacturers build clusters that enable the execution of these applications, in such a way that it will be impossible to run them in a single processor without adding more memory, or in a multi-processor without modifying the source code of the application.

1. Introduction

Traditionally, shared memory systems have been used to run applications requiring a high memory space. However, such as systems do not scale more than tens of processors. As memory requirements of applications and the number of applications that run concurrently on computers have been increasing, designers have made proposals to partially solve the lack of memory on computer systems. In this way, modern operating systems provide advanced virtual memory managers that solve the lack of memory. These managers utilize secondary devices for freeing contents of memory whenever it is necessary. This approach provides a simple way of running applications that have a running image size bigger than available physical memory size. When an application is running and the system is low of memory, the virtual memory manager can evict it of a special device called a "swap device", or swap, to free memory. This technique is called swapping [17], and there are several approaches in the literature about it. Just to mention a few, Unix-based systems use a separate swap partition type that is hosted in the user file system. In contrast, Windows uses a user-space file that is hosted inside the file system, while the MacOS X operating system can use partitions and files.

Nevertheless, swapping has several drawbacks. The first one comes from the access time of the swapping device, which is usually a hard disk, therefore, one or more orders of magnitude higher than the memory access time. The second drawback is thrashing, which occurs when the memory manager evicts parts of the running image of a process and, after a time, it reallocates those parts in memory again. Thus, solving the lack of memory always involves a high run time.

The HyperTransport Consortium Advanced Technology Group (ATG) is working on an extended version of the traditional PGAS programming model [5]. Such a model will let manufacturers build clusters with PGAS native support. In this paper we carry out an assessment of "rough" DPGAS model through Simics [11] and GEMS [12] simulators. However, this work has not attempted to conduct a study using a hardware implementation, because the ATG has not completed the specification of its DPGAS model yet, and therefore it is impossible to achieve that kind of evaluation. The behaviour of the final system will be similar (bridging the gap) to the behaviour in a cluster with DPGAS native support.

The AMD's Opteron processor can be used as a commodity to build clusters. This processor includes the memory controller on-die, in such a way that all memory is accessible from one memory controller. The Opteron processors use the AMD's HyperTransport protocol [9] for communicating with each other. Moreover, the HyperTransport

protocol enables CPUs to directly connect the Opteron HyperTransport link to add-in card subsystems via the HTX connector [7], which is placed in the motherboard.

In addition, the increasing use of interconnection networks to intercommunicate the processor and memory, such as AMD HyperTransport might allow the construction of scalable systems, from hundreds to thousands of nodes composed of multicore processors. Therefore, the Opteron processors will provide the basis to build the clusters with DPGAS native support.

The remainder of this paper is structured as follows: In the next section we present a summary of the related work. The details of the proposed model are explained in section 3, while the section 4 details the simulation scenarios and the results obtained. In the last section, we conclude and provide a brief overview of the future work.

2. Related Work

HyperTransport (HT) is an interconnection technology which enable connecting the processors among each other and with the I/O devices. It provides an extremely low latency, high bandwidth and excellent scalability. Moreover, the definition of the HTX connector allows co-processing and acceleration based on ASIC or FPGA technologies. In particular, it is receiving a highlighted interest of the community because it makes easy to reduce power consumption by the use of accelerators.

Partitioned Global Address Space languages combine a Single Program Multiple Data (SPMD) programming model with a global address space, which is logically partitioned to give each thread a portion of shared memory to which it has affinity [19]. In the SPMD model, a fixed number of threads are created at program startup, and every thread runs the same program. Each thread has both a space for private local memory and some partition of the shared space to which it has affinity. A private object may only be accessed by its corresponding thread, whereas all threads can read or write any object in the shared address space. The partitioning of the shared space into regions with logical affinity to threads allows programmers to explicitly control data layout, which is then used by the runtime system to map threads and their associated data to processors: on a distributed memory machine, the local memory of a processor holds both the thread's private data and the shared data with affinity to that thread.

The HTC Advanced Technology Group (ATG) [1] is working on developing proposals for HyperTransport to define an address space globally and dynamic partitioning (DPGAS) for using in scalable clusters. The idea is not new, but it comes from the existing PGAS models [5]. In the bibliography several PGAS applications can be found, for example, Unified Parallel C [6] to define models of programming languages. In addition, developers of these languages have tools like GASnet [4] which is a communication interface for programming languages such as Unified Parallel C. GASnet is a language independent of the network that allows the definition of libraries providing global addressing. GASnet is inspiring the work of HTC Advanced Technology Group for the implementation of PGAS [18] in a native way.

The ATG is proposing the mechanisms and abstractions that will allow the construction of clusters using Opteron processors in common. The motherboard containing Opteron processors will support the HTX connector. By plugging extension cards on the HTX connector will allow the formation of a cluster of motherboards. The memory controller of each Opteron will divide the whole range of physical addresses in regions and distribute them among the memory of other Opterons. Subsequently, the memory controllers will be able to access remote regions of memory transparently. Following this approach, the Opteron processors can avoid the use of a device that swapping up for the lack of memory, since access to a memory on a memory controller will be lower than access to a local secondary storage device, and hence system performance will be enhanced.

In such systems, the physically addressable memory in all nodes is part of a global address space with non-uniform access time from any specific node. From the perspective of a node, the global address space is composed of local partitions and remote partitions where the former can be accessed with the lowest latency and the latter can be accessed with larger and possibly non-uniform latencies (across distinct partitions). In this model, a local partition refers to DRAM accessed through a tightly integrated memory controller.

The Computer Architecture Group at the University of Heidelberg in Germany has the expertise to design complex hardware/software systems. The HTX-Board [8], a contribution of this group, provides a convenient and efficient way to evaluate user specific devices connected to the Hypertransport connector standardised under the name of HTX-Connector. In [16] they published the architecture and mechanisms of the HTX-board. Subsequently, in [10] they have introduced a novel communication engine in combination with the HyperTransport interface. They provide an excellent prototype to get real-world measurements connecting two Opteron through two HTX-boards. Moreover, they also show the initial latencies of sending a HT packet on a HT link and propose some optimisations that can minimise that latency (e.g. doubling the HT clock frequency or migrating FPGA to ASIC technology).
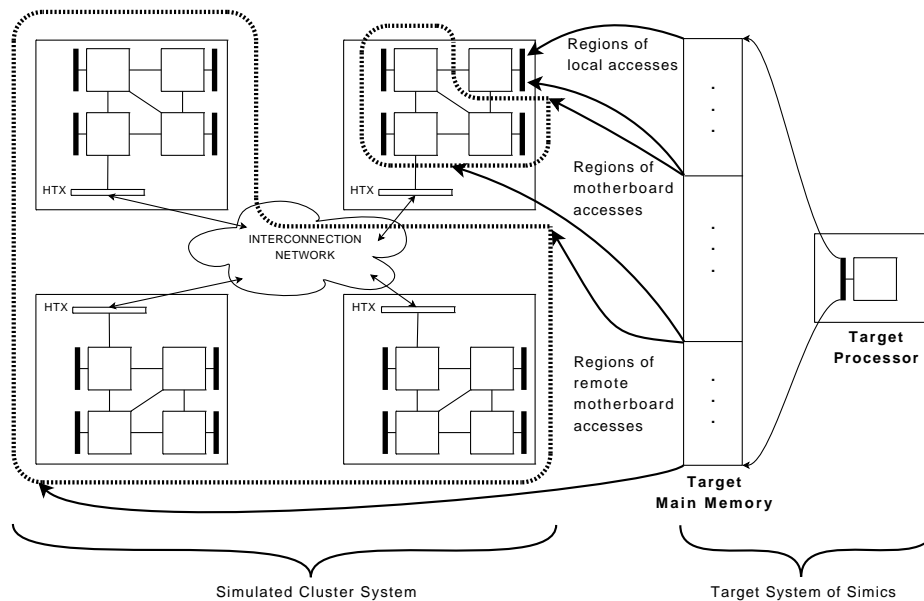
Figure 1. Diagram of the simulated cluster.

3. Model

Because of the fact that the ATG has not completed the specification of its model DPGAS yet, it is impossible to carry out an evaluation using a hardware system. However, it is feasible to carry out an approximate evaluation of the DPGAS model by simulation.

In order to simulate the cluster, Simics 2.1 and GEMS 2.2.19 simulators were used. In the Simics context, two fundamentals terms are always used:

- The computer on which we are running Simics is referred to *host* system.

- The computer simulated by Simics is referred to *target* system. Specifically, it simulates the cluster with DPGAS support.

Our approach consists in simulating the execution of a sequential application as if it would be running on a cluster with DPGAS support just using the execution-driven Simics simulator. In that hypothetical cluster, any processor might issue requests of the global address space. The global address space will consist of the joint of every memory in the cluster. From the processors point of view, most of the physical memory in the cluster can be accessed, except some private areas that will not be allowed to access. The Figure 1 explains how a cluster can be simulated starting from a simulated processor in Simics.

Without lossing generality, we run a sequential application in the target system and we process every message going inside the system. The messages are deliverated by the simulator as they are in a usual execution, but the delay of every message is customised depending on the type, source and destination of the message.

In the AMD's whitepaper [15] a suite of benchmarks is examined to illustrate their performance and scalability in single, multi-processor, and cluster configurations. Its results clearly show the exceptionally responsiveness of an Opteron-based NUMA support system. Specifically, it shows the cost for one processor for accessing to the shared memory in a four AMD Opteron motherboard. The latencies are given depending on the distance between the transmitter and receiver processors.

Regarding the application, it must be a benchmark that makes an intense use of memory. Also, a sequential applicaction is preferred in order to avoid any dependency produced by a parallel execution. Stream [13] is a well-known benchmark that measures bandwidth sustainable by ordinary user programs, and not the theoretical peak bandwidth that vendors advertise [14]. Moreover, it is used by AMD to measure the performance of the memory of their processors [2].

The benchmark performs functions with matrices that are stored in memory. The functions are executed several times. When the benchmark concludes, it returns the rate of

traffic data of memory in MB/seconds and execution time (average, minimum and maximum) in seconds, for each function. Both performance indexes are the most relevant in this kind of benchmarking. The execution time gives the global performance measure and the traffic rate offers the real load of the memory system.

4. Evaluation

In this section, we start describing the simulation model we have used to carry out our experiment. Then, we present the results we have obtained and some comments about them.

4.1. Simulation Model

In all the simulations, our customized GEMS module is loaded. It is responsible for managing all the messages sent by the memory system. We assume that the target processors utilised in this work are used to build systems with a coherent shared memory, similar to the Opteron processors. Therefore, the GEMS module has to manage the messages caused by the memory coherence protocols. The study of memory coherence protocols is out of the scope of this paper. Thus the simulations have been carried out with a single processor in order to reduce the influence of these protocols.

The host system is a SUN W2100Z workstation that has two Opteron processors at 2.4 GHz and a DDR-400 memory of 4 GB. SuSe 10.2 is used as operating system. The target system is the *sarek* preinstalled system of Simics which is a UltraSPARC processor at 75 MHz. Solaris is used as operating system and an amount of 512 MB of memory is configured.

The assumption that the whole memory of the cluster is 512 MB seems initially nonsense. However, the aim of this work has never been to propose a detailed DPGAS simulation model because the ATG has not finished its DPGAS model yet. Considering an increase of the target memory size, that is the memory size of the cluster, requires to increase the Stream benchmark size and then it causes an exponential simulation time growth. Even though the results remain representative because the benchmark spreads the accesses out the memory address space.

Additionally, the parameters of the Stream benchmark have been set to achieve the following behaviour:

1. Stream runs two series of functions (copy, scale, add and triad). Previous tests had proved that increasing the number of series does not alter the final outcome in the absence of processes that interfere with Stream.

2. The size of Stream is 460 Mbytes. This size was chosen because it represents 90% of the available simulated memory (512 MB).

3. The GEMS simulator requires to set the latencies in the target processor cycles, so a 2.4 GHz processor was considered for the translation from real nanoseconds to simulated cycles.

It must be noticed that the latencies in our simulations are considered as an approximation. However, we have selected values that reflect real systems:

- When an access is destined for memory allocated in the same motherboard, a latency of 115 ns (nanoseconds) for each access has been considered [15].

- When an access is destined for memory allocated in a remote motherboard, the latency has been deduced from the proposed delays in [10, 16]. We assume all the improvements suggested by [10, 16] like ASIC technology instead of using FPGA technology, doubling HyperTransport link frequency, and a 16-bit HyperTransport link width. In this case, the calculation of the latencies is:

  - Due to the all technology improvements, [10] claims that a total latency of 130 ns for the transmission and expects a fixed latency (for local CPU and remote memory controller) of about 300 ns. The functionality of [10] is exactly the behavior of a remote write (transfer of data from one node to another). Then we assume a latency of 430 ns for a Write operation of 64 byte payload.

  - The read is much more difficult. The Opteron used in the [16] paper only issues 32 bit read operations, due to the access granularity of 64 bit (a limitation of the K10 architecture). Hence, we assume a series of 8 consecutive operations have to be issued to fetch the total amount of 64 byte. This is the reason to assume that the Read operation takes 4880 ns.

- When an access is destined for memory mapped in swap device, we assume a latency of 45600 ns as it is suggested in [3] for enterprise class harddisks.

Regarding to simulation scenarios, we have considered the following scenarios:

- Local scenario (1P): It represents a desktop system with just one processor and one motherboard.

- Shared scenario (4P): It represents a server system, commonly known as a shared memory multiprocessor. There are four processors ensembled in one motherboard.

(a) 1P-SWAP

(b) 4P-SWAP-U

(c) 4P-SWAP-D
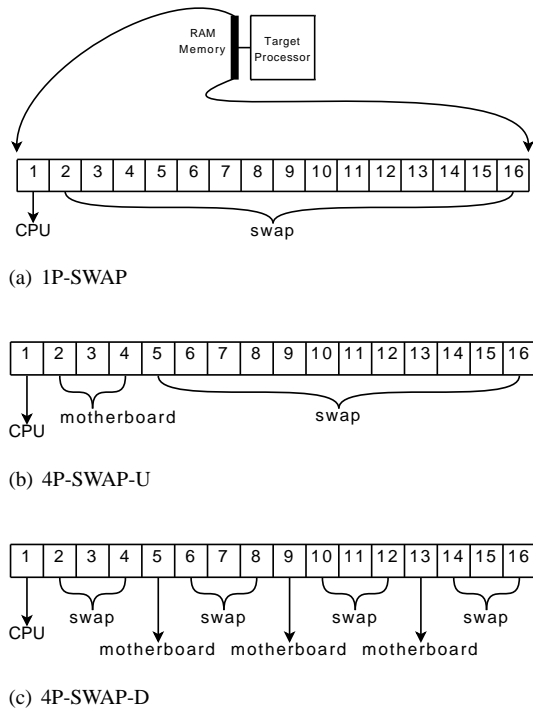
Figure 2. Distribution of target memory into regions and their access type.

- Remote scenario (16P): It represents a enterprise system or a cluster. A total of sixteen processors are distributed between four motherboards of four processors per motherboard that are interconnected using the HTX connectors.

A group of extra three scenarios have been selected to evaluate the loss of performance due to the utilisation of swapping. The Figure 2 shows the distribution of the target memory into regions and the access type that is associated with each region. Each region corresponds to a contiguous physical memory partition controlled by a single node. In our test, we assumed all regions are of the same size. These extra scenarios are based on the previous scenarios and they consist in distributing the regions of memory between swapping devices, as if they were accesses to secondary devices and therefore such accesses were applied an extra delay. All the extra scenarios assume a partitioning of the memory in 16 regions. The main characteristics are:

- 1P-SW.- It is basically the 1P scenario, but the accesses from the second to the last regions are accesses to a swap device (see Figure 2(a)).

- 4P-SW-U.- It is similar to the 4P scenario. It assigns uniformly the swap regions (see Figure 2(b)).

- 4P-SW-D.- Similar to the 4P-SW-U scenario, but the swap regions are interleaved on the memory (see Figure 2(c)).

4.2. Simulation Results

The aim of this work has never been focused on the performance of the application, but the behaviour of a cluster with native DPGAS support (scenario 16P) to run applications. Note that it will be impossible to run these applications in a single processor (scenario 1P) or multi-processor (scenario 4P) due to the memory requirements of these applications. Note also that the only alternative in these cases is the use of swapping devices, which is considered in the 1P-SW and 4P-SW scenarios.

The results have shown that scenario 16P are always a better choice than extra scenarios that implement swapping. Of course, the performance of the unfeasible scenarios 1P and 4P is much better than the performance of the 16P scenario, but this one is the best option for those applications with high memory requirements.

Mainly, it is interesting to know how the execution time evolves. Figure 3 depicts that the execution time increases for 4P and 16P on average 0.54 seconds (2.68%) and 40.76 seconds (202.68%) with regard to the 1P scenario. Because of it is a memory benchmark, it is also interesting to know how the traffic memory evolves, so the Figure 3 depicts that the performance of the memory for 4P and 16P decreases 0.53 MB/s (2.63%) and 13.46 MB/s (66.93%), regarding the performance achieved by scenario 1P. When extra scenarios are studied, the results are even worse, as it is expected. Both, the execution time and the memory traffic, fall dramatically for all scenarios.

5. Conclusions and Future Work

This paper presents the results of the preliminary assessment of the work in progress that is made by the ATG. Because of the ATG has not completed the specification of its DPGAS model yet, it is not possible to carry out a hardware evaluation. However, we have performance a simulation-driven study.

Firstly, we developed a module of the GEMS simulator for tracking the memory requests and customizing their latency. By this module, we could simulate approximately the behaviour of an application running in a cluster with DPGAS support. This cluster will run any application with high memory requirements if they do not exceed the whole memory of the cluster, because the application will be spread out the DRAM memory of the processors in the cluster.

As it was explained, swapping can solve the lack of memory, but the application performace falls dramatically

(a) Increasing of execution time with regard to scenario 1P (numerical and logarithmic scale).

(b) Decreasing of mean traffic with regard to scenario 1P (numerical and logarithmic scale).

Figure 3. Performance results of the Stream application in each scenario.

as the lack of memory increases. This paper has introduced the PGAS model as one alternative to swapping. The results have demostrated that the DPGAS model will never be a better option than having enough memory in the processor that runs the applications, because a lot of time will spend in accesses to remote memories. However, it will be always better than using swapping, because the latencies of the inter-memory communications are always lower than accesses to swapping.

As future work it is interesting to keep updated of all the work that is done by the ATG, for example, the real implementation of the DPGAS support, the future improvements of the HTX-board, and the specification of the HTX connector.

Acknowledgements

The authors want to give special thanks to Dr. Holger Fröning of the University of Heidelberg for his helpful comments and assistance for the development of this work.

References

[1] HyperTransport Consortium Advanced Technology Group. http://www.hypertransport.org.

[2] Second-Generation AMD Opteron Processor Industry Standard Server Benchmarks. http://www.amd.com/us-en/Processors/SellAMDProducts/0,,30_177_4458_3505%5E8796%5E8800%5E15227,00.html.

[3] Enterprise-class versus Desktop class Hard Drives, Revision 1.0, April 2008.

[4] D. Bonachea. Gasnet Specification, version 1.1, Report No. UCB/CSD-02-1207, October 2002.

[5] P. Charles, C. Grothoff, V. Saraswat, and et. al. X10: An object-oriented approach to nonuniform cluster computing (OOPSLA'05). In *Proceedings of the 20th annual ACM SIG-PLAN Conference on object oriented programming, systems, languages, and applications*, 2005.

[6] T. El-Ghazawi, W. Carlson, T. Sterling, and K. Yelick. *UPC: Distributed Shared Memory Programming*. 2005.

[7] D. Emberson and D. O'Flaherty. HTX Specification for HyperTransport 3.0 Daughtercards and ATX/EATX Moth-

erboards, Document HTC20080701-00030-0001. Technical report, HyperTransport Consortium, June 2008.

[8] H. Fröning, M. Nüssle, D. Slogsnat, H. Litz, and U. Brüning. The HTX-Board: A Rapid Prototyping Station. In *Proceeding of 3rd Annual FPGAWorld Conference, Stockholm, Sweden*, November 2006.

[9] HyperTransport Consortium. Hypertransport Specification, Revision 3.0, April 2007.

[10] H. Litz, H. Froening, M. Nuessle, and U. Bruening. VELO: A Novel Communication Engine for Ultra-Low Latency Message Transfers. In *Proceedings of 37th International Conference on Parallel Processing (ICPP-08), Portland, Oregon, USA*, September 2008.

[11] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hållberg, J. Högberg, F. Larsson, A. Moestedt, and B. Werner. Simics: A Full System Simulation Platform Computer. 35(2), February 2002.

[12] M. M. Martin, D. J. Sorin, B. M. Beckmann, M. R. Marty, M. Xu, A. R. Alameldeen, K. E. Moore, M. D. Hill, and D. A. Wood. Multifacet's General Execution-driven Multiprocessor Simulator (GEMS) Toolset. *Computer Architecture News (CAN)*, September 2005.

[13] J. D. McCalpin. STREAM: Sustainable Memory Bandwidth in High Performance Computers. Technical report, University of Virginia (USA), 2007. `http://www.cs.virginia.edu/stream`.

[14] S. A. McKee. Reflections on the Memory Wall. In *Conference Computing Frontiers*, 2004.

[15] D. O'flaherty and M. Goddard. AMD Opteron Processor Benchmarking for Clustered Systems. Technical report, Advanced Micro Devices, July 2003.

[16] D. Slogsnat, A. Giese, M. Nüssle, and U. Brüning. An open-source HyperTransport core. *ACM Trans. Reconfigurable Technol. Syst.*, 1(3):1–21, September 2008.

[17] W. Stallings. *Operating Systems: Internals and Design Principles (6th Edition)*. Prentice Hall, April 2008.

[18] S. Yalamanchili, J. Young, J. Duato, and F. Silla. A Dynamic, Partitioned Global Address Space Model for High Performance Clusters. Document GIT-CERCS-08-01, School of Electrical and Computer Engineering (Georgia Institute of Technology (USA); Universidad Politecnica de Valencia (Spain), 2008.

[19] K. Yelick and et. al. Productivity and Performance Using Partitioned Global Address Space Languages. In *Proceeding of International Workshop on Parallel Symbolic Computation (PASCO'07)*, July 2007.