

NewMadeleine: a Fast Communication Scheduling Engine for High Performance Networks

Olivier Aumage, Elisabeth Brunet, Nathalie Furmento, Raymond Namyst

► **To cite this version:**

Olivier Aumage, Elisabeth Brunet, Nathalie Furmento, Raymond Namyst. NewMadeleine: a Fast Communication Scheduling Engine for High Performance Networks. [Research Report] RR-6085, INRIA. 2007. <inria-00122723v2>

HAL Id: inria-00122723

<https://hal.inria.fr/inria-00122723v2>

Submitted on 8 Jan 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

**NEWMADELEINE: *a Fast Communication
Scheduling Engine for High Performance Networks***

Olivier AUMAGE and Élisabeth BRUNET and Nathalie FURMENTO and Raymond NAMYST

N° 6085

December 2006

Thème NUM



*Rapport
de recherche*

NEWMADELEINE: a Fast Communication Scheduling Engine for High Performance Networks*

Olivier AUMAGE and Élisabeth BRUNET and Nathalie FURMENTO and
Raymond NAMYST

Thème NUM — Systèmes numériques
Projet Runtime

Rapport de recherche n° 6085 — December 2006 — 13 pages

Abstract: Communication libraries have dramatically made progress over the fifteen years, pushed by the success of cluster architectures as the preferred platform for high performance distributed computing. However, many potential optimizations are left unexplored in the process of mapping application communication requests onto low level network commands. The fundamental cause of this situation is that the design of communication subsystems is mostly focused on reducing the latency by shortening the critical path. In this paper, we present a new communication scheduling engine which dynamically optimizes application requests in accordance with the NICs capabilities and activity. The optimizing code is generic and portable. The database of optimizing strategies may be dynamically extended.

Key-words: Communication Scheduling, High Performance Networks, Dynamic Optimizations

* This work was funded through the ANR-05-CIGC-11 grant.

NEWMADELEINE: Un Moteur d'Ordonnement pour Communications Rapides sur Réseaux à Haute Performance

Résumé : Les bibliothèques de communication ont fait d'énorme progrès ces quinze dernières années, en partie du fait du succès des architectures de grappes pour le calcul distribué haute performance. Néanmoins, plusieurs techniques d'optimisation n'ont pas été intégrées n'étant pas considérées comme essentielles à la réduction des temps de latence. Dans cet article, nous présentons un nouveau moteur d'ordonnement de communications qui optimise dynamiquement les requêtes des applications en fonction des caractéristiques et de la disponibilité des cartes réseaux. Le code d'optimisation est générique et portable. De nouvelles stratégies d'optimisation peuvent être ajoutées dynamiquement.

Mots-clés : Ordonnement des Communications, Réseaux à Haute Performance, Optimisations Dynamiques

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 5 |
| 2 | New optimization potentials for communications | 5 |
| 3 | The NEWMADELEINE communication engine | 6 |
| 3.1 | Optimization window | 6 |
| 3.2 | Strategies, tactics and optimization selection | 6 |
| 3.3 | Architecture | 7 |
| 3.4 | Application Programming Interfaces | 8 |
| 4 | Highlights of the implementation | 8 |
| 5 | Evaluation | 8 |
| 5.1 | Overhead of NEWMADELEINE | 9 |
| 5.2 | Aggregation of small messages | 9 |
| 5.3 | Optimization of derived datatypes | 10 |
| 6 | Related works | 11 |
| 7 | Conclusion and Future Work | 11 |

1 Introduction

The success of cluster architecture as the most widespread platform for high performance computing mainly comes from their aggressive performance/cost ratio. However, it is also the result of the substantial progress made in the field of fast cluster interconnects. Hardware improvements towards lower latency and higher bandwidth have been followed by huge progress at the software level. Using a large panel of mechanisms such as user-mode communications, zero-copy transactions and communication operation offload, the critical path in sending and receiving a packet has been drastically reduced. The cost is now only a few hundreds of microprocessor cycles per transaction with libraries such as Elan/Quadrics [8] or MX/Myrinet [6]. Even more impressive is the fact that some recent implementations of the MPI standard such as MPICH2 [2] or OpenMPI [4], which have been carefully designed to directly map *basic* point-to-point requests onto the underlying low-level interfaces, almost exhibit themselves the same level of performance.

However, for that matter, only the very basic point-to-point messaging requests are considered first-class citizens. More complex requests such as non-contiguous messages are left mostly unattended, and even more so are the irregular and multi-flow communication schemes. The purpose of the NEWMADELEINE communication engine introduced in this paper is to address this situation thoroughly. With a carefully designed architecture, explained in Section 3, the NEWMADELEINE optimization layer delivers much better performance on complex communication schemes with a negligible overhead on basic single packet point-to-point requests as it will be shown in Section 5. Through MAD-MPI, our simple, straightforward proof-of-concept implementation of a subset of the MPI API, we show that MPI applications can also benefit from the NEWMADELEINE communication engine.

2 New optimization potentials for communications

The art of mapping simple contiguous point-to-point requests from the MPI API onto fast, low latency networks arguably reached its climax with MPI implementations such as MPICH2 [2] or OpenMPI [4], in terms of pure efficiency. However, needs of real-life applications in communication requests not always meet such a narrow focus. The *message layout* may be more complex. Messages may be made of multiple fragments disseminated irregularly in memory. The preferred *optimization strategy* may differ from favoring the latency, and instead favoring the bandwidth may be a better bet for applications using a remote storage system. A preference for communication overlap may be more suitable for computing intensive applications. Applications may even have need for different optimization strategies at different stages or for different requests.

Even the MPI API itself is not always the most suited API. We showed in the past ([1]) that applications and programming environments that make use of high-level communication protocols (DSM, RPC, etc.) may draw a substantial benefit from using a communication interface more powerful than the regular MPI. Whilst this interface behaves optimally for the classical *ping-pong* tests and the simple regular communication patterns, the overall lack of expressiveness of MPI does not capture subtle informations such as the dependencies between the multiple fragments (service request, arguments, targeted object) of a remote method invocation.

As a result, many communication schemes and needs are left mostly unexplored. These schemes —and combination of these— represent a rich field of optimization potentials. Our new NEWMADELEINE communication engine has been designed with the purpose of harvesting

these potentials without incurring a prohibitive penalty for basic communication requests. Our engine dynamically optimizes the communication flow or multiple flows (increasingly found in nowadays composite applications). It may decide to accumulate packets in order to make use of some gather/scatter capabilities or to aggregate several short requests into a single larger one if constraints allow it, or to reorder packets, or even to favor an earlier delivery of high priority fragments (such as a RPC service id, needed for preparing the data areas to receive the service arguments). For such classes of applications characterized by complex, irregular communication schemes, it is indeed not a matter of simply translating communication operations into calls to the network driver, but rather to dynamically *interpret*, *reorganize* and *optimize* the communication operation flow.

3 The NEWMADELEINE communication engine

We present here the main features of NEWMADELEINE ¹, our new communication library.

3.1 Optimization window

To be efficient, dynamic optimizations for multi-packets and multi-flow communication schemes require the use of an optimization window that accumulates packets. Traditionally, communication libraries, being synchronous, tightly link the communication requests to the application workflow, and therefore transmit incoming packets immediately to the lower network layer without any accumulation. A communication library that would choose to accumulate packets with the aim of possible optimizations would lead to increased latency and possibly cause deadlocks. To avoid such drawbacks, it is necessary to untie the processing of the communication requests from the application workflow, and instead to link this processing to the activity of the NICs (Network Interface Card).

While the NICs are busy, NEWMADELEINE keeps accumulating packets in its optimization window. As soon as a NIC becomes idle, the optimization window is analyzed so as to create a new ready-to-send packet to be transferred through the card. This approach does thus not have the cons of the synchronous solutions. It also brings two advantages: NICs are exploited at their maximum (they are not overloaded when there is a high demand of transfers and under exploited when there is not) and the communication optimizations are made *just-in-time* so they closely fit the ongoing communication scheme at any given time.

3.2 Strategies, tactics and optimization selection

The optimization window being built, one has now to deal with the actual optimization work. Let us consider the state of the physical multiplexing units at a given time. If at least one of the multiplexing units is idle, one has to assign it some work. The assignment is done by calling an *optimization function* to elect the next request to be submitted to each idle unit. In doing so, it may select a packet to be sent from the optimization window, or for instance, synthesize a request out of several packets from that window. A wide panel of arguments may be used as an input to the optimizing function: the number of packets in the window, the specific characteristics of each packet (destination, flow tag, length, sequence number, dependency attributes), the nominal and functional characteristics of the underlying network, possibly some hints given by the application itself with respect with the packet scheduling policy, and the number and state of multiplexing units at this point in time.

¹This library is a new evolution of the communication library known as MADELEINE.

Given such a large amount of potential parameters, several optimization tactics may be available, but the optimal combination of these tactics is a difficult problem. We thus propose a (dynamically in the future) selectable optimization function instead of a fixed optimizing heuristic. The optimization function is to be selected among an extensible and programmable set of *strategies*. Each tactic applies some elementary optimizing operations selected from the panel of usual operations toward some particular optimizing goal.

While any multiplexing unit is available, the communication requests are just accumulated. Another possibility would be to prepare a single *ready-to-send* packet to anticipate for any upcoming completion of the current activity of one of the multiplexing unit and immediately re-feed it once it becomes idle. A third possibility would be to run the optimization function unconditionally once the packet backlog has reached a predefined threshold length.

3.3 Architecture

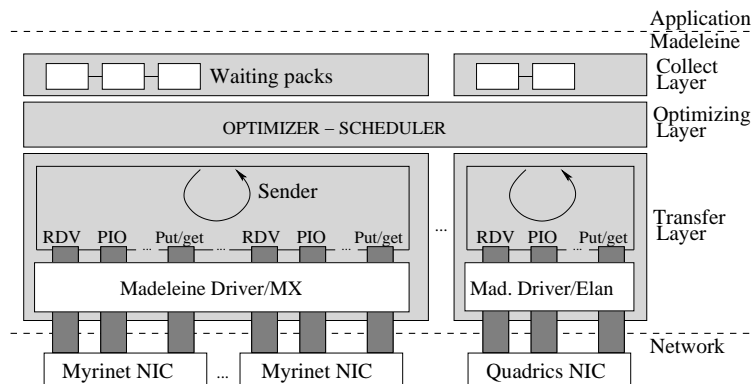


Figure 1: NEWMADELEINE architecture.

The NEWMADELEINE architecture is organized in three layers as shown on Figure 1, an application data collect layer, an optimizing and scheduling layer and a transfer layer that controls the network cards. Only the transfer layer is NIC-specific.

The transfer layer mimics a process scheduler, which when called by a processor, will select the new ready process to be run. Indeed, the transfer layer controls the activities of the NICs, and requests from the upper layer a new optimized packet to be sent, as soon as a card becomes idle.

The optimizing and scheduling layer is queried by the transfer layer when a NIC gets idle and needs to be re-filled. If any packet has been previously prepared (see Section 3.2), it is then responsible for analyzing the backlog of accumulated packets in order to build a request both compatible with the application constraints and efficient performance-wise. This request is then immediately submitted to the idle NIC.

The collect layer is in charge of registering the pieces of data submitted by the various communication flows of the application as well as the meta-data necessary in their identification by the receiving side (tag number, sender id, sequence number). Once encapsulated, and in order to load balance the packets among the available NICs (possibly from heterogeneous technologies), the collected pieces of data are inserted onto a dedicated list for a specific network technology selected by the application or (by default) on the common list for automatized load-balancing among all the NICs (possibly from heterogeneous technologies).

3.4 Application Programming Interfaces

The communication engine `NEWMADELEINE` provides several interfaces. The first interface is similar to the interface of the former `MADELEINE` library, it allows to incrementally build messages. With this interface, a `NEWMADELEINE` message is made of several pieces of data, located anywhere in user-space. The message is initiated and finalized with a synchronization barrier call.

Thereafter, in order to exhibit the performance of `NEWMADELEINE` with MPI applications, we have implemented a subset of the MPI standard on top of `NEWMADELEINE`. This implementation called `MAD-MPI` is based on the point-to-point nonblocking posting (`isend`, `irecv`) and completion (`wait`, `test`) operations of MPI, these four operations being directly mapped to the equivalent operations of `NEWMADELEINE`.

`MAD-MPI` also implements some optimizations mechanisms for derived datatypes [3]. MPI derived datatypes deal with noncontiguous memory locations. The advanced optimizations of `NEWMADELEINE` allowing to reorder packets should lead to a significant gain when sending and receiving data based on derived datatypes.

4 Highlights of the implementation

A `NEWMADELEINE` prototype has been implemented over `GM/MYRINET`, `MX/MYRINET`, `ELAN/QUADRICS`, `SISCI/SCI` and `TCP/ETHERNET`. The implementation of each corresponding *transfer layer* consists in a minimal network API (initialisation, closing, sending, receiving and polling methods) for which each function is, at the best, a direct call to the adequate function of the network driver. In addition, some information are collected such as the threshold for the *rendez-vous* protocol or the availability of the *gather/scatter* or as well the remote direct access (RDMA) functionality.

Scheduling strategies are as well based on a minimal interface and are independent from the network technology. Developing a new strategy only requires to write a few methods such as an initialisation method, and a request method which returns the next communication request to be sent or received. Information about the underlying network can be obtained in a generic manner through a specific API. Thus, any strategy can be directly combined with any network protocol supported by `NEWMADELEINE`. Currently, the `NEWMADELEINE` scheduler offers two strategies: an *aggregation* one which accumulates communication requests as long as the cumulated length does not require to switch to the *rendez-vous* protocol and a *multi-rails* one which balances the communication flow over the set of available NICS, possibly by splitting messages in a heterogeneous manner if necessary.

5 Evaluation

In this section, we present the results obtained by comparing our lightweight implementation of MPI (`MAD-MPI`) against dedicated implementations of MPI. The first test evaluates the software overhead introduced by `MAD-MPI` and the following ones highlight the gain obtained by reordering communication requests.

All our experiments have been carried out on the same platform, a set of two dual-core 1.8 GHz `OPTERON` boxes with 1MB of L2 cache and 1GB of main memory. The OS kernel is Linux version 2.6.17. Nodes are interconnected through `MYRI-10G` NICs with the `MX1.2.0` driver and through `QUADRICS QM500` NICs with the `ELAN` driver.

5.1 Overhead of NEWMADLEINE

To evaluate the overhead of the NEWMADLEINE scheduling engine under situations where no optimization is possible, we use a MPI ping-pong program exchanging single-segment messages (i.e. contiguous arrays of bytes). Figures 2(a) and 2(b) shows the transfer times obtained with MPICH-MX, OPENMPI-MX 1.1 and with MAD-MPI over MYRI-10G and Figures 2(c) and 2(d) ones obtained with MPICH-QUADRICS and MAD-MPI over QUADRICS.

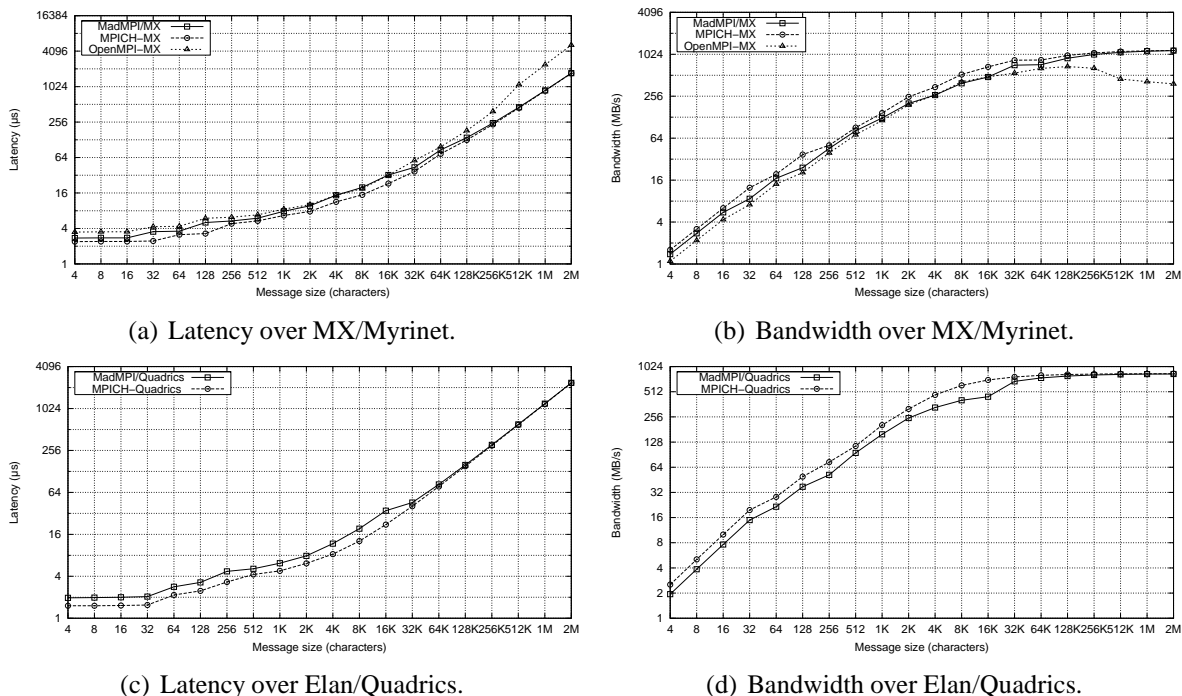



Figure 2: Raw point-to-point ping-pong.

On both networks, MAD-MPI introduces a constant overhead of less than $0,5 \mu s$ and reaches 1155 Mbytes/s in bandwidth over MYRI-10G and 835 Mbytes/s over QUADRICS. In the both case, this small overhead actually results from different factors. On one hand, an extra header is systematically added to the data by NEWMADLEINE for allowing the reordering and the multiplexing of the packets. Thus, the exchanged packets are slightly larger with NEWMADLEINE than with MPICH-MX. On the other hand, the NEWMADLEINE scheduler introduces some extra operations on the critical path to inspect the “ready list” of packets and check if some optimization action would be appropriate. The performance target of NEWMADLEINE not being regular communication schemes, we consider that this latency overhead of MAD-MPI remains reasonable.


5.2 Aggregation of small messages

To evaluate the benefits of using a communication scheduler able to perform optimizations over the whole communication flow between a pair of processes, we compare the performance of a multi-segments ping-pong program with several implementations of MPI. Each “ping” is a series of independent MPI_Isend operations that use separate MPI communicators (to demonstrate that the scope of MAD-MPI optimizations is really global). The MAD-MPI implementation was

configured to use the *aggregation* optimization strategy. Results obtained over MX-10G and QUADRICS are reported on Figure 3.



(a) 8-segments ping-pong - Latency over MX/Myrinet. (b) 16-segments ping-pong - Latency over MX/Myrinet.



(c) 8-segments ping-pong - Latency over Elan/Quadracs. (d) 16-seg. ping-pong-Latency over Elan/Quadracs.

Figure 3: Performance of a ping-pong program featuring multi-segments messages.

Neither the MPICH nor the OPENMPI try to aggregate individual messages submitted in a short time interval. Nevertheless, the MPICH-MX and MPICH-QUADRICS implementations are able to pipeline the transfer of a series of messages in a very efficient manner. Despite this fact, the MAD-MPI implementation outperforms MPICH and OPENMPI thanks to its aggressive optimizer which is able to coalesce packets even if they belong to different logical communication flows (i.e. MPI communicators). We can observe that MAD-MPI is up to 70 % faster than other implementations of MPI over MX-10G, and up to 50 % faster than MPICH over QUADRICS.

5.3 Optimization of derived datatypes

We now present the performance of our optimization mechanisms when using MPI derived datatypes. We use a ping-pong program which exchanges arrays of a given indexed datatype. The datatype describes a sequence of two data blocks, one small block (64 bytes) followed by a large data block (256 KBytes).

In order to process a derived datatype communication request, MPICH copies all the data fragments into a new contiguous buffer and sends the obtained buffer in a unique transaction, using the *rendez-vous* protocol if necessary. Data are received in a temporary memory area before being dispatched to their final destination [5]. This behaviour is certainly optimized when dealing with a small overall data size as the *memcpy* operations for each the data blocks will cost less than the multiple communication operations. However, the cost of a memory copy operation being proportional to the size of the data, this behaviour is no longer optimized when dealing

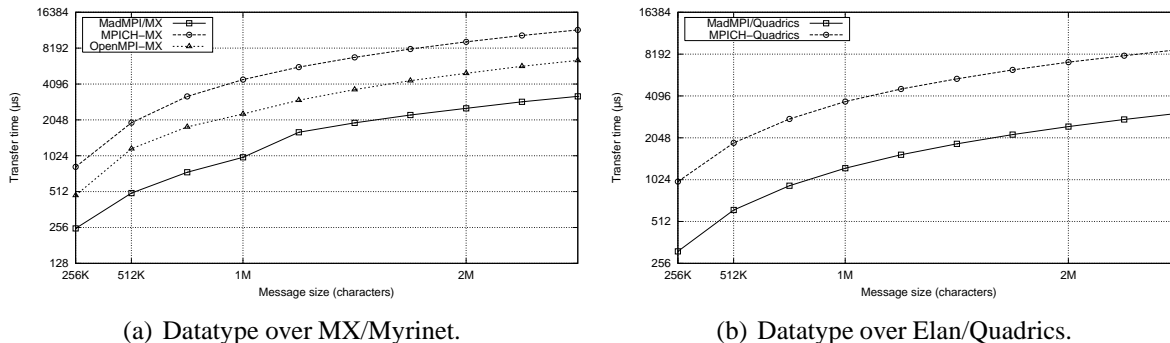


Figure 4: Indexed datatype.

with bigger blocks. In the absence of related documentation, we guess that OPENMPI has the same behaviour. On the opposite, MAD-MPI uses an algorithm which generates an individual communication request for each block, allowing the underlying communication layer to perform any appropriate optimization. We used a scheduling strategy which aggregates all the small blocks (using messages reordering) with the *rendez-vous* requests of the large blocks, hence the large blocks are directly received at their final destination, and the whole transfer is made with a zero-copy technique.

This strategy decreases significantly the time transfer of indexed datatypes build upon small and large fragments. On our example (see Figure 4), we can observe a gain of about 70 % in comparison with MPICH and about 50 % with OPENMPI over MX and until about 70 % versus MPICH over QUADRICS.

6 Related works

This section briefly presents the recent research efforts in designing communication supports for high performance networks.

MPICH2-NEMESIS [2] is to our knowledge the current best implementation of the MPI standard both in term of bandwidth and latency. Nevertheless, its implementation currently focusses on the performance of individual transfers. Thus, no message reordering or multiplexing is used. PMV2 [10] is used as the multi-network communication subsystem for the YAMPPI [7] MPI implementation by the same team. As such, it also focuses on achieving low latency, as does MPICH2-NEMESIS. VMI 2 [9] is a low-level communication library dedicated to address fault tolerance issues. It provides multi-rails capabilities on top of homogeneous and heterogeneous networks. It also implements some NIC scheduling but does not provide multiplexing of communication flows. It is therefore mostly oriented towards improving bandwidth. MADELEINE 3 [1] is the preceding major version of NEWMADLEINE. It already was performing some very basic request optimizations such as aggregating packets without inter-dependencies. However, MADELEINE 3 was only sending header-less packets by design, making it impossible to implement non-deterministic sender actions such as opportunistic multi-flow aggregation or packet reordering.

7 Conclusion and Future Work

In this paper, we present a new low-level communication engine for high speed networks. NEWMADLEINE features a programmable packet scheduler that can use a wide range of “*just-in-*

time” optimization strategies. Data segments can be aggregated into the same physical packet even if they belong to different logical channels (e.g. different MPI communicators). If needed, they can be reordered (to maximize the number of aggregation operations), or even sent out-of-order. Large data segments can also be split on the sending side (and later reassembled on the receiving side) into several chunks that may be sent through different networks.

To demonstrate the relevance of using such a communication engine for building communication libraries such as MPI implementations, we developed a simple optimization strategy that aggressively tries to aggregate small packets (data and control) whenever possible. The experiments conducted over MYRI-10G and QUADRICS networks show that the gain obtained when transferring complex (i.e. non contiguous) messages within MPI applications can be substantial. We also showed that the overhead of the NEWMADELEINE optimizer is small when running applications with no opportunity for communication optimizations.

We are currently designing more powerful optimization strategies to efficiently exploit multiple, heterogeneous physical networks within the same application. The NEWMADELEINE architecture is particularly well suited to the implementation of greedy load-balancing strategies over multiples network interface cards. In the short term, we also plan to port a full featured MPI implementation such as MPICH2 [2] or OpenMPI [4] on top of NEWMADELEINE, to evaluate the impact of aggressive communication optimizations in real applications.

References

- [1] O. Aumage, L. Bougé, A. Denis, L. Eyraud, J.-F. Méhaut, G. Mercier, R. Namyst, and L. Prylli. A Portable and Efficient Communication Library for High-Performance Cluster Computing. *Cluster Computing*, 5(1):43–54, 2002.
- [2] D. Buntinas, G. Mercier, and W. Gropp. Design and Evaluation of Nemesis, a Scalable, Low-Latency, Message-Passing Communication Subsystem. In *The Sixth IEEE International Symposium on Cluster Computing and the Grid (CCGRID '06)*, pages 521–530, Washington, DC, USA, 2006. IEEE Computer Society.
- [3] N. Furmento and G. Mercier. Optimisation Mechanisms for MPICH-Madeleine. Technical Report 0306, INRIA, July 2005. Also available as LaBRI Report 1362-05.
- [4] R. L. Graham, T. S. Woodall, and J. M. Squyres. Open MPI: A Flexible High Performance MPI. In *The 6th Annual International Conference on Parallel Processing and Applied Mathematics*, 2005.
- [5] W. Gropp, E. Lusk, and D. Swider. Toward faster packing and unpacking of mpi datatypes. <http://www.mcs.anl.gov/mpi/mpich1/papers/index.html>.
- [6] Myricom Inc. Myrinet EXpress (MX): A High Performance, Low-level, Message-Passing Interface for Myrinet, 2003. <http://www.myri.com/scs/>.
- [7] Y. Ishikawa. YAMPPI Official Home Page. <http://www.il.is.s.u-tokyo.ac.jp/yampii>.
- [8] Quadrics Ltd. Elan Programming Manual, 2003. <http://www.quadrics.com/>.
- [9] S. Pakin and A. Pant. VMI 2.0: A Dynamically Reconfigurable Messaging Layer for Availability, Usability, and Management. In *The 8th International Symposium on High Performance Computer Architecture (HPCA-8)*, 2002.
- [10] T. Takahashi, S. Sumimoto, A. Hori, H. Harada, and Y. Ishikawa. PM2: High performance communication middleware for heterogeneous network environments. In *SC'00*, pages 52–53, 2000.



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399