

Evaluation of Java Message Passing in High Performance Data Analytics

Saliya Ekanayake, Geoffrey Fox
School of Informatics and Computing
Indiana University
Bloomington, Indiana, USA
{sekanaya, gcf}@indiana.edu

Abstract—**TODO**.

Index Terms—**TODO**

I. INTRODUCTION

MESSAGE passing has been the *de-facto* model in realizing distributed memory parallelism [1] where Message Passing Interface (MPI) with its implementations such as MPICH2 [2] and OpenMPI [3] have been successful in producing high performance applications [4]. We use message passing with threads [5] in our data analytics applications on Windows High Performance Computing (HPC) environments [6] using MPI.NET [7] and Microsoft's Task Parallel Library (TPL) [8]. However, the number of available Windows HPC systems are limited and our attempts to run these on traditional Linux based HPC clusters using Mono – a cross platform .NET development framework – have been unsuccessful due to poor performance. Therefore, we decided to migrate our applications to Java for reasons 1) productivity offered by Java and its ecosystem; and 2) emerging success of Java in HPC [9]. In this paper, we present our experience in evaluating the performance of our parallel deterministic annealing clustering program (DAVS) [10] and micro benchmark results for two Java message passing frameworks – OpenMPI with Java binding and FastMPJ [11] – compared against native OpenMPI and MPI.NET. The results show clear improvement of application performance over C# with MPI.NET and near native performance in micro benchmarks.

II. RELATED WORK

Message passing support for Java can be classified as pure Java implementations or Java bindings for existing native MPI libraries (i.e. wrapper implementations). Pure Java implementations advocate portability, but may not be as efficient as Java bindings that call native MPI (see III.D and [11]). There are two proposed Application Programming Interfaces (API) for Java message passing – mpiJava 1.2 [12] and Java Grande Forum (JGF) Message Passing interface for Java (MPJ) [13]. However, there are implementations that follow custom API as well [9].

Performance of Java MPI support has been studied with different implementations and recently in [11, 14]. The focus of these studies is to evaluate MPI kernel operations and little or no information given on applying Java MPI to scientific applications. However, there is an increasing interest [15] on using MPI with large scale data analytics frameworks such as Apache Hadoop [16]. MR+ [17] is a framework with similar intent, which allows Hadoop MapReduce [18] programs to run on any cluster under any resource manager while providing capabilities of MPI as well.

III. TECHNICAL EVALUATION

The DAVS code is about 15k lines of C# code and to evaluate performance on Java we used a combination of commercially available code converter [19] and carefully inspected manual rewrites to port the C# code to Java. Furthermore, we performed a series of serial and parallel tests to confirm correctness is preserved during the migration, prior to evaluating performance.

Our interest in this experiment was to compare application performance when run on Linux based HPC clusters against results on Windows HPC environments. We noticed from initial runs that two of the MPI operations – allreduce, and send and receive – contribute to the most of inter-process communication. Therefore, we extended the evaluation by performing micro benchmarks for these, which further supported our choice to use Java.

A. Computer Systems

We used two Indiana University clusters, Madrid and Tempest, and one FutureGrid [20] cluster – India, as described below.

Tempest: 32 nodes, each has 4 Intel Xeon E7450 CPUs at 2.40GHz with 6 cores, totaling 24 cores per node; 48 GB node memory and 20Gbps Infiniband (IB) network connection. It runs Windows Server 2008 R2 HPC Edition – version 6.1 (Build 7601: Service Pack 1).

Madrid: 8 nodes, each has 4 AMD Opteron 8356 at 2.30GHz with 4 cores, totaling 16 cores per node; 16GB node memory and 1Gbps Ethernet network connection. It runs Red Hat Enterprise Linux Server release 6.5

FutureGrid (India): 128 nodes, each has 2 Intel Xeon X5550 CPUs at 2.66GHz with 4 cores, totaling 8 cores per

```

Input: maxMsgSize // maximum message size in bytes

msgSizeL = 8192 // messages to be considered large
numLtr = 1000 // iterations for small messages
numLtrL = 100 // iterations for large messages
skip = 200 // skip this many for small messages
skipL = 10 // skip this many for large messages

comm = MPI_COMM_WORLD
me = MPI_Comm_rank(comm)
size = MPI_Comm_size(comm)

sbuff[maxMsgSize/4] // float array –initialized to 1.0
rbuff[maxMsgSize/4] // float array –initialized to 0.0

For i = 1 to i ≤ maxMsgSize/4
  If i > msgSizeL
    skip = skipL
    numLtr = numLtrL
  MPI_Barrier(comm)
  duration = 0.0
  For j = 0 to j < numLtr + skip
    t = MPI_Wtime()
    MPI_Allreduce(sbuff,rbuff,i
                  MPI_SUM,comm)

    If j ≥ skip
      duration += MPI_Wtime() - t
    MPI_Barrier(comm)
    j = j + 1
  latency = duration/numLtr
  avgTime = MPI_Reduce(latency, MPI_SUM, 0)
  If me == 0
    Print(i, avgTime)
  i = i * 2
MPI_Barrier(comm)

```

Fig. 1 Pseudo code for allreduce benchmark

node; 24GB node memory and 20Gbps IB network connection. It runs Red Hat Enterprise Linux Server release 5.10.

B. Software Environments

We used .NET 4.0 runtime and MPI.NET 1.0.0 for C# based tests. DAVS Java version uses a novel parallel tasks library called Habanero Java library from Rice University [21, 22], which requires Java 8. Therefore, we used an early access (EA) release – build 1.8.0-ea-b118. Early access releases may not be well optimized, but doing a quick test with threading switched off we could confirm DAVS performs equally well on stable Java 7 and Java 8 EA.

There have been several message passing frameworks for Java [23], but due to the lack of support for IB network and to other drawbacks discussed in [11], we decided to evaluate OpenMPI with its Java binding and FastMPJ, which is a pure Java implementation of mpiJava 1.2 [12] specification.

OpenMPI’s Java binding [24] is an adaptation from the original mpiJava library [25]. However, OpenMPI community has recently introduced major changes to its API, and internals, especially removing MPI.OBJECT type and adding support for direct buffers in Java. These changes happened while we were evaluating DAVS, thus we tested OpenMPI Java binding in one of its original (nightly snapshot version 1.9a1r28881) and updated forms (source tree revision 30301). We will refer to these as OMPI-nightly and OMPI-trunk for simplicity.

C. MPI Micro Benchmarks

We based our experiments on Ohio MicroBenchmark (OMB) suite [26], which is intended to test native MPI implementations’ performance. Therefore, we implemented the selected allreduce, and send and receive tests in all three Java MPI flavors and MPI.NET, in order to test Java and C# MPI implementations.

Fig. 1 presents the pseudo code for OMB test. Note the syntax of MPI operations do not adhere to a particular language and the number of actual parameters are cut short for clarity. Also, depending on the language and MPI framework used, the implementation details such as data structures used and buffer allocation were different from one another.

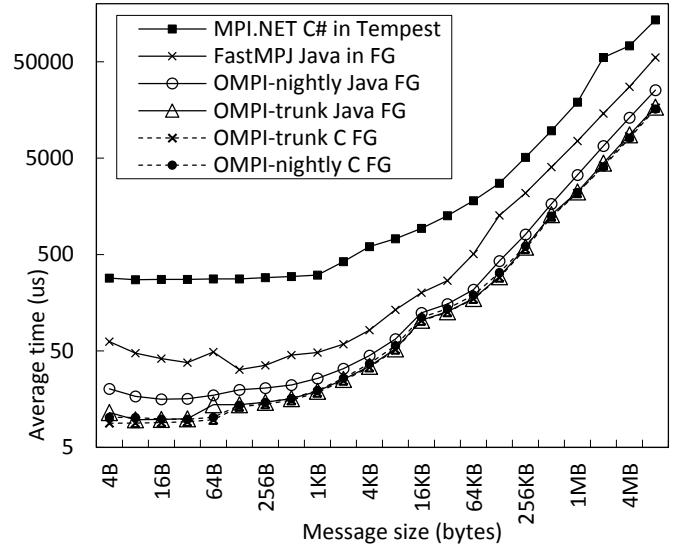


Fig. 2. Performance of MPI allreduce operation

Fig. 2 shows the results of allreduce benchmark for different MPI implementations. These are averaged values over patterns 1x1x8, 1x2x8, and 1x4x8 where pattern format is number of threads per process x number of processes per node x number of nodes (i.e. TxPxN). The best performance came with C versions of OpenMPI, but interestingly OMPI-trunk Java performance overlaps on these indicating its near zero overhead. The older, OMPI-nightly Java performance is near as well, but shows more overhead than its successor. FastMPJ performance is better than MPI.NET, but slower than OpenMPI versions. The slowest performance came with MPI.NET, which may be improved with further tuning, but as our focus was to evaluate Java versions we did not proceed in this direction.

We experienced a similar pattern with MPI send and receive (Fig. 3) where OMPI-trunk Java performance overlaps with results from native MPI as shown in Fig. 4.

Input: `maxMsgSize` // maximum message size in bytes

```

msgSizeL = 8192 // messages to be considered large
numltr = 1000 // iterations for small messages
numltrL = 100 // iterations for large messages
skip = 200 // skip this many for small messages
skipL = 10 // skip this many for large messages
comm = MPI_COMM_WORLD
me = MPI_Comm_rank(comm)
size = MPI_Comm_size(comm)

```

```

sbuff[maxMsgSize] // byte array -initialized to 1.0
rbuff[maxMsgSize] // byte array -initialized to 0.0

```

```

For i = 0 to i ≤ maxMsgSize
  If i > msgSizeL
    skip = skipL
    numltr = numltrL
  MPI_Barrier(comm)
  duration = 0.0
  If me == 0
    For j = 0 to j < numltr + skip
      If j == skip
        t = MPI_Wtime()
        MPI_Send(sbuff, i, 1, 1)
        MPI_Recv(rbuff, i, 1, 1)
        duration += MPI_Wtime() - t
      Else If me == 1
        For j = 0 to j < numltr + skip
          MPI_Recv(rbuff, i, 0, 1)
          MPI_Send(sbuff, i, 0, 1)
        If me == 0
          latency = duration / 2 * numltr
          Print(i, latency)
          i = i == 0 ? 1 : i * 2
    MPI_Barrier(comm)

```

Fig. 3 Pseudo code for send and receive benchmark

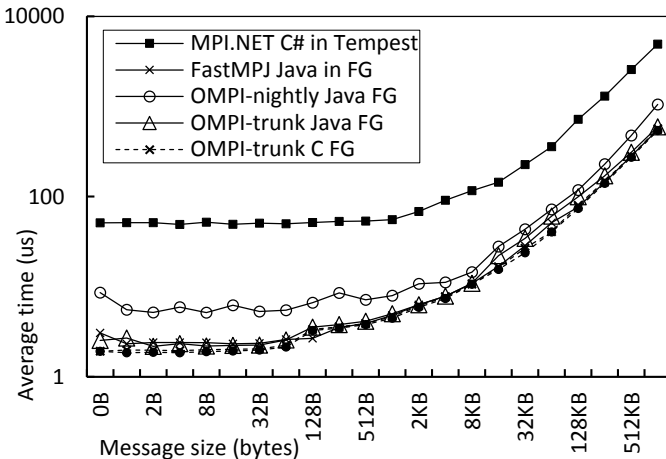


Fig. 4 Performance of MPI send and receive operations

D. Application Performance

We decided to apply DAVS to the same “peak-matching” problem [10] that its C# variant used to solve, so we could verify accuracy and compare performance. We performed clustering of the LC-MS data [10] in two modes – Charge2 and Charge5 – where the former processed 241605 points and found on average 24.5k clusters. Charge5 mode handled 16747 points producing an average of 28k clusters. These modes exercises different execution flows in DAVS where Charge2 is more intense in both computation and communication than Charge5. DAVS supports threading too, but we are still working on this in Java versions. Therefore, we do not include a performance comparison with threads in this paper.




Fig. 5 DAVS Charge5 performance

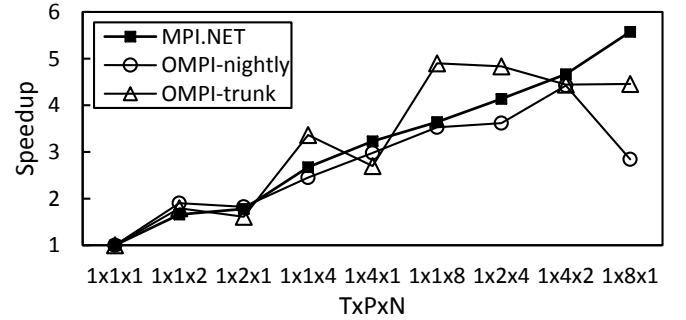


Fig. 6 DAVS Charge5 speedup

Fig. 5 and Fig. 6 show Charge5 performance and speedup under different MPI libraries. OMPI-trunk happens to give the best performance and it achieves nearly double the performance in all cases compared to MPI.NET. Charge2 performance and speedup show similar results as given in Fig. 7 and Fig. 8. Note we could not test pattern 1x8x1 due to insufficient memory.




Fig. 7 DAVS Charge2 performance




Fig. 8 DAVS Charge2 speedup

SUMMARY

Scientific applications written in C, C++, and Fortran have embraced MPI since its inception and various attempts have been made over the years to establish this relationship for applications written in Java. However, only few implementations such as OpenMPI and FastMPJ are in active development with support for fast interconnect systems. OpenMPI in particular has recently introduced improvements to its Java binding to close the gap between Java and native performance. The kernel benchmarks we performed agree with this and depicts latest OpenMPI Java binding as the best among selected Java MPI implementations.

Our aim of this effort has been to migrate existing C# based code to Java in hope of running on traditional HPC clusters while utilizing the rich programming environment of Java. The initial runs of DAVS show promising performance and we expect to complement this work by adding support for threads in near future.

REFERENCES

- [1] LUSK, E. and YELICK, K. LANGUAGES FOR HIGH-PRODUCTIVITY COMPUTING: THE DARPA HPCS LANGUAGE PROJECT. *Parallel Processing Letters*, 17, 01 (2007), 89-102.
- [2] Laboratory, A. N. *MPICH2*. City.
- [3] Gabriel, E., Fagg, G., Bosilca, G., Angskun, T., Dongarra, J., Squyres, J., Sahay, V., Kambadur, P., Barrett, B., Lumsdaine, A., Castain, R., Daniel, D., Graham, R. and Woodall, T. *Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation*. Springer Berlin Heidelberg, City, 2004.
- [4] Gropp, W. Learning from the Success of MPI. In *Proceedings of the 8th International Conference on High Performance Computing* (2001). Springer-Verlag, [insert City of Publication],[insert 2001 of Publication].
- [5] Ekanayake, S. *Survey on High Productivity Computing Systems (HPCS) Languages*. Pervasive Technology Institute, Indiana University, Bloomington, 2013.
- [6] Qiu, J., Beason, S., Bae, S.-H., Ekanayake, S. and Fox, G. Performance of Windows Multicore Systems on Threading and MPI. In *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing* (2010). IEEE Computer Society, [insert City of Publication],[insert 2010 of Publication].
- [7] Gregor, D. and Lumsdaine, A. Design and implementation of a high-performance MPI for C# and the common language infrastructure. In *Proceedings of the Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming* (Salt Lake City, UT, USA, 2008). ACM, [insert City of Publication],[insert 2008 of Publication].
- [8] Daan Leijen, J. H. *Optimize Managed Code For Multi-Core Machines*. City.
- [9] Taboada, G. L., Touri, J., #241, Ram, #243 and Doallo, n. Java for high performance computing: assessment of current research and practice. In *Proceedings of the Proceedings of the 7th International Conference on Principles and Practice of Programming in Java* (Calgary, Alberta, Canada, 2009). ACM, [insert City of Publication],[insert 2009 of Publication].
- [10] Fox, G., Mani, D. R. and Pyne, S. *Parallel deterministic annealing clustering and its application to LC-MS data analysis*. IEEE, City, 2013.

- [11] Expósito, R., Ramos, S., Taboada, G., Touriño, J. and Doallo, R. FastMPJ: a scalable and efficient Java message-passing library. *Cluster Computing*(2014/02/06 2014), 1-20.
- [12] Bryan Carpenter, G. F., Sung-Hoon Ko and Sang Lim. *mpiJava 1.2: API Specification*. 1999.
- [13] Carpenter, B., Getov, V., Judd, G., Skjellum, A. and Fox, G. MPJ: MPI-like message passing for Java. *Concurrency: Practice and Experience*, 12, 11 (2000), 1019-1038.
- [14] Taboada, G. L., Tourino, J. and Doallo, R. *Performance analysis of Java message-passing libraries on fast Ethernet, Myrinet and SCI clusters*. City, 2003.
- [15] Squyres, J. *Resurrecting MPI and Java*. City, 2012.
- [16] Tom White *Hadoop: The Definitive Guide*. Yahoo Press; Second Edition edition, 2010.
- [17] Ralph H. Castain, W. T. *MR+ A Technical Overview*. 2012.
- [18] Dean, J. and Ghemawat, S. MapReduce: Simplified Data Processing on Large Clusters. *Sixth Symposium on Operating Systems Design and Implementation*(2004), 137-150.
- [19] Inc., T. S. S. *C# to Java Converter*. City.
- [20] von Laszewski, G., Fox, G. C., Fugang, W., Younge, A. J., Kulshrestha, A., Pike, G. G., Smith, W., Vo, x, ckler, J., Figueiredo, R. J., Fortes, J. and Keahey, K. *Design of the FutureGrid experiment management framework*. City, 2010.
- [21] Cav, V., #233, Zhao, J., Shirako, J. and Sarkar, V. Habanero-Java: the new adventures of old X10. In *Proceedings of the Proceedings of the 9th International Conference on Principles and Practice of Programming in Java* (Kongens Lyngby, Denmark, 2011). ACM, [insert City of Publication],[insert 2011 of Publication].
- [22] Sarkar, V. a. I., Shams Mahmood *HJ Library*. City.
- [23] Taboada, G. L., Ramos, S., Exp, R. R., #243, sito, Touri, J., #241, Ram, #243 and Doallo, n. Java in the High Performance Computing arena: Research, practice and experience. *Sci. Comput. Program.*, 78, 5 (2013), 425-444.
- [24] Project, T. O. M. *FAQ: Where did the Java interface come from?*, City.
- [25] Baker, M., Carpenter, B., Fox, G., Hoon Ko, S. and Lim, S. *mpiJava: An object-oriented java interface to MPI*. Springer Berlin Heidelberg, City, 1999.
- [26] Laboratory, T. O. S. U. s. N.-B. C. and (NBCL) *OMB (OSU Micro-Benchmarks)*. City.