CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. 2006; **00**:1–12 *Prepared using* cpeauth.cls *[Version: 2002/09/19 v2.02]*

A Framework for Performance Analysis of Co-Array Fortran

Bernd Mohr[1,*], Luiz DeRose[2] and Jeffrey Vetter[3]

[1] *Forschungszentrum Jülich, ZAM, Jülich, Germany*
[2] *Cray Inc., Mendota Heights, MN, USA*
[3] *Oak Ridge National Laboratory, Oak Ridge, TN, USA*

SUMMARY

Co-Array Fortran is a parallel programming language for scientific applications that provides a very intuitive mechanism for communication, and especially, one-sided communication. Despite the benefits of this integration of communication primitives with the language, analyzing the performance of CAF applications is not straightforward, which is due, in part, to a lack of tools for analysis of the communication behavior of Co-Array Fortran applications. In this paper, we present an extension to the KOJAK toolkit based on a source-to-source translator that supports performance instrumentation, data collection, trace generation, and performance visualization of Co-Array Fortran applications. We illustrate this approach with a performance visualization of a Co-Array Fortran version of the Halo kernel benchmark using the VAMPIR event trace visualization tool.

KEY WORDS: performance analysis, Co-Array Fortran

Introduction

Co-Array Fortran (CAF) [14] extends Fortran 2003 providing a simple, explicit notation for data decomposition, communication, and synchronization, expressed in a natural Fortran-like syntax. These extensions provide a straightforward and powerful paradigm for parallel programming of scientific applications based on one-sided communication. One of the major problems facing CAF users is the lack of Performance analysis tools for investigating the communication and synchronization behavior of their application. This deficiency is due in part to the fact that communication operations in CAF programs are not expressed through function calls, as in MPI, or via directives that are executed by a run-time library, as in OpenMP. In contrast, CAF communication operations are integrated into the language, and, on certain platforms like the Cray X1, they are implemented via remote memory access instructions provided by the hardware.

*Correspondence to: Bernd Mohr, Forschungszentrum Jülich, ZAM, 52425 Jülich, Germany. email: b.mohr@fz-juelich.de

Copyright © 2006 John Wiley & Sons, Ltd.

For MPI applications, performance data collection is, in general, facilitated by the existence of the MPI profiling interface (PMPI), which is used by most MPI tools [2, 8, 16]. Similarly, performance measurement of OpenMP applications can be done by instrumenting the calls to the runtime library [1, 4, 5]. However, with the challenge of CAF communication primitives being integrated into the language, and potentially implemented with special hardware instructions, the instrumentation of these communication primitives requires a similarly integrated approach.

In order to address this problem, we first defined PCAF, an interface specification of a set of routines intended to monitor all important aspects of CAF applications. Then, we extended the OPARI source-to-source instrumentation tool [12] to search for CAF constructs and to generate instrumented source code with the appropriate PCAF calls. Finally, we implemented the PCAF interface for the KOJAK measurement system [15] enabling it to trace CAF communication and synchronization instructions. With this extension, the KOJAK measurement system is able to support performance instrumentation and performance data collection of CAF applications, generating trace files that can be analyzed with the VAMPIR event trace visualization tool [13].

The remainder of this paper is organized as follows. First, we present an overview of Co-Array Fortran. Then, we briefly describe the KOJAK performance measurement and analysis environment. The main part of the article describes our approach for performance instrumentation and measurement of Co-Array Fortran applications. Finally, we discuss performance visualization with an example using the Halo kernel benchmark code.

An Overview of Co-Array Fortran

Co-array Fortran [14] is a parallel programming language extension to Fortran 2003 currently under investigation for inclusion in the next Fortran language standard. At the highest level, CAF uses a Single Program Multiple Data (SPMD) model to allow multiple copies (*images*) of a program to execute asynchronously. Each image contains its own private set of data objects. When data objects are distributed across multiple images, the array syntax of CAF uses an additional trailing subscript in square brackets to allow explicit access to remote data (as shown in Figure 2), and it is referred to as the *co-dimension*. Data references that do not use these square brackets are strictly local accesses.The CAF compiler translates these remote data accesses into underlying communication mechanisms for each target system. CAF also includes intrinsic routines to synchronize images, to return the number of images, and to return the index of the current image. Besides functions for delimiting a critical region, CAF provides four different forms of a barrier synchronization:

SYNC_ALL(): a global barrier where every image waits for every other image.

SYNC_ALL(<wait list>): a global barrier where every image waits only for the listed images.

SYNC_TEAM(<team>): a barrier where a team of images wait for every other team member.

SYNC_TEAM(<team>, <wait list>): a barrier where a team of images wait for a subgroup of the team members.

CAF was originally developed on the Cray-T3D, and, as such, it is very efficient on platforms that support one-sided messaging and fast barrier operations. On systems with globally addressable

Copyright © 2006 John Wiley & Sons, Ltd.

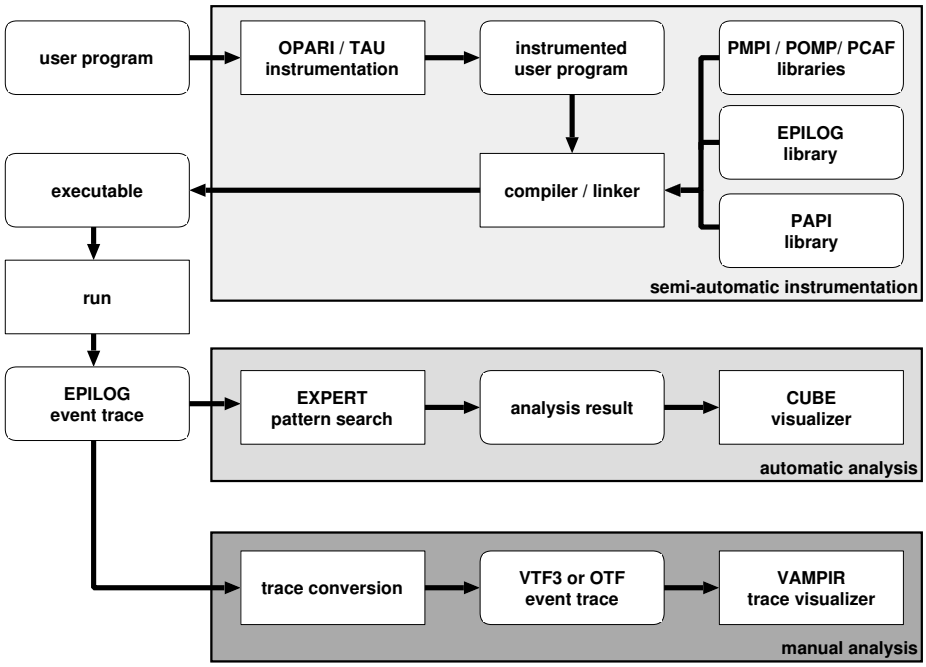

Figure 1. KOJAK OVERALL ARCHITECTURE.

memory, such as the Cray X1 or the SGI Altix 3700, these mechanisms may be as simple as load and store memory references. By contrast, on distributed memory systems that do not support efficient Remote Direct Memory Access (RDMA), these mechanisms can be implemented in MPI, SHMEM, or some lower level library like GasNET or ARMCI.

The KOJAK Measurement System

The KOJAK performance-analysis tool environment provides a complete tracing-based solution for automatic performance analysis of MPI, OpenMP, SHMEM, or hybrid applications using combinations of these three paradigms running on parallel computers. KOJAK describes performance problems using a high-level of abstraction in terms of execution patterns that result from an inefficient use of the underlying programming model(s). KOJAK's overall architecture is depicted in Figure 1. Tasks and components are represented as rectangles and their inputs and outputs are represented as boxes with rounded corners. The arrows illustrate the whole performance-analysis process from instrumentation to result presentation.

Copyright © 2006 John Wiley & Sons, Ltd.

The KOJAK analysis process is composed of two parts: a semi-automatic, multi-level instrumentation of the user application followed by an automatic analysis of the generated performance data. The first part is considered semi-automatic because it requires the user to slightly modify the makefile.

To begin the process, the user supplies the application's source code, written in either C, C++, or Fortran, to OPARI, which is a source-to-source translation tool. OPARI performs automatic instrumentation of OpenMP constructs and redirection of OpenMP-library calls to instrumented wrapper functions on the source-code level, based on the POMP OpenMP monitoring API [11]. Later, in Section "Performance Instrumentation," we describe how we extended OPARI for instrumentation of CAF programs with the appropriate PCAF calls.

Instrumentation of user functions is done either during compilation by a compiler-supplied instrumentation interface or on the source-code level using TAU [2]. TAU is able to automatically instrument the source code of C, C++, and Fortran programs using a preprocessor based on the PDT toolkit [10].

Instrumentation for MPI events is accomplished with a wrapper library based on the PMPI profiling interface. All MPI, OpenMP, CAF and user-function instrumentation calls the EPILOG run-time library, which provides mechanisms for buffering and trace-file creation. The application can also be linked to the PAPI library [3] for collection of hardware counter metrics as part of the trace file. At the end of the instrumentation process, the user has a fully instrumented executable.

Running this executable generates a trace file in the EPILOG format. After program termination, the trace file is fed into the EXPERT analyzer. (See [15] for details of the automatic analysis, which is outside of the scope of this paper.) In addition, the automatic analysis can be combined with a manual analysis using VAMPIR [13], which allows the user to investigate the patterns identified by EXPERT in a time-line display via utilities that convert the EPILOG trace file into the VAMPIR VTF3 and OTF formats.

Performance Instrumentation and Measurement Approach

In this section, we describe the event model that we use to describe the behavior of CAF applications, and the approach we take to instrument CAF programs and to collect the necessary measurement data.

An Event Model of CAF

KOJAK uses an event-based approach to analyze parallel programs. A stream or trace of events allow it to describe the dynamic behavior of an application over time. If necessary, execution statistics can be calculated from that trace. The events represent all the important points in the execution of the program. Our CAF event model is based on KOJAK's basic model for one-sided communication [6]. We extended KOJAK's existing set of events, which cover describing the begin and end of user functions and MPI and OpenMP related activities, with events for representing the execution of CAF programs, namely the begin and end of CAF synchronization primitives and the begin and end of remote load and store operations.

Table I summarizes KOJAK's event types used for analyzing CAF applications. For each of these events, we collect a time stamp and location. For CAF synchronization functions, we also record which function was entered or exited. For the barrier routines, we also collect the group of images which

Copyright © 2006 John Wiley & Sons, Ltd.

Table I. KOJAK's event types for CAF programs

Abstraction	Event type	Type specific Attributes
Entering / leaving a region (e.g., a function)	ENTER	region id
	EXIT	region id
Leaving a collective CAF synchronization regions	CAFCEXIT	region id, team id, wait group id
Start / end of RMA one-sided transfers	PUT_1TS	rma id, length, destination loc
	PUT_1TE	rma id, length, source loc
	GET_1TS	rma id, length, destination loc
	GET_1TE	rma id, length, source loc

participate in the barrier and the group of images waited for, if applicable. Finally, for loads and stores, we collect the amount of data transferred (i.e., the number of array elements) and the source or destination of the transfer, as well as a unique RMA operation identifier that provides an easier mapping of #_1TE to the corresponding #_1TS events in the later analysis stage.

The event model is also the basis for the instrumentation and measurement. The events and their attributes specify the elements of CAF programs that need to be instrumented and the data that has to be collected.

Performance Instrumentation

Instrumentation of CAF programs can be done on either of two levels, depending on how CAF is implemented on a specific computing platform. On systems where CAF constructs and API calls are translated into calls to a run-time library, these calls could easily be instrumented by traditional techniques (e.g., linking a pre-instrumented run-time library or instrumenting the calls with a binary instrumentation tool). However, for systems like the Cray X1, where the CAF communication is executed via hardware instructions, this approach is not always possible. Therefore, we extended OPARI, KOJAK's source-to-source translation tool, to also locate and instrument all CAF constructs of a program.

As Fortran statements can occupy more than one source line, OPARI reads a program line by line, taking into consideration continuation lines and multiple statements on the same source line. Each line is scanned for occurrences of CAF constructs and synchronization calls (but ignoring comments and contents of strings). CAF constructs can be located by looking for pairs of brackets ([...]) which are not used otherwise in Fortran, except for array constructors in the 2003 Fortran standard. This exception can be easily handled.

The first word of the statement determines whether it is a declaration line, an allocation statement involving Co-Arrays, or a statement containing a remote load or store operation. For CAF declarations and allocations, OPARI collects attributes like array dimensions, and lower and upper bounds for later use. The allocate and deallocate statements that involve Co-Arrays, as well as automatic deallocation at

Copyright © 2006 John Wiley & Sons, Ltd.

```
integer :: me, num, left
integer :: A(1024,1024)[*]
. . .
me = this_image()
num = num_images()
left = me - 1
if ( left < 1 ) left = num



A(me,::2)[left] = me

call sync_all()
```

(a)

```
integer :: me, num, left
integer :: A(1024,1024)[*]
. . .
me = this_image()
num = num_images()
left = me - 1
if ( left < 1 ) left = num
call PCAF_rma_store_begin(-1+left, &
    1 * max((ubound(A,2)-         &
             lbound(A,2)+2)/2,0))
A(me,::2)[left] = me
call PCAF_rma_store_end(-1+left)
call PCAF_sync_all()
```

(b)

Figure 2. (a) Example of a CAF source code and (b) OPARI instrumented version

the end of subprograms, also contain internal synchronization and therefore would be good candidates for some additional form of instrumentation, probably insertion of a call before and after the statement, or at the end of a subprogram when there is automatic deallocation. OPARI currently does not support this functionality, but it could be easily extended to do so.

The handling of statements containing remote memory operations is more complex. First, all operations are located in the line. If it is an assignment statement and the operation appears before the assignment operator, it is a store operation. In all other cases it is a load. OPARI determines which CAF array is referenced by the operation, the number of elements transferred (by parsing the index specification), and the source or destination of the transfer (determined by the expression inside the brackets). Simple assignment statements containing a single remote memory operation are instrumented by inserting calls to the corresponding PCAF monitoring functions before and after the statement, which get passed in the attributes determined by OPARI. In case of more complex statements where a remote memory operation cannot be easily separated out and wrapped by the measurement calls, or when it is necessary to keep instrumentation overhead low, OPARI uses the single call version of the PCAF remote memory access monitoring functions (instead of separate begin and end calls) and inserts them either before (for loads) or after (for stores) the statement for each identified remote memory access operation.

Finally, OPARI scans the line for calls to CAF synchronization routines, and replaces them by calls to PCAF wrapper functions that will execute the original call in addition to collecting all important attributes.

Figure 2(b) shows the instrumented source code generated for the example in Figure 2(a). In this example, there is a two-dimensional array `A`, which is distributed on all processors. In the CAF statement `A(me,::2)[left] = me`, each processor updates the odd entries of the row corresponding to its image in the left neighbor array with its index, and then waits on a barrier.

Copyright © 2006 John Wiley & Sons, Ltd.

Prepared using cpeauth.cls

Remote Memory Access Monitoring Routines

```
  SUBROUTINE PCAF_rma_store_begin(dest, nelem)
  SUBROUTINE PCAF_rma_store_end(dest)
  SUBROUTINE PCAF_rma_store(dest, nelem)
  SUBROUTINE PCAF_rma_load_begin(src, nelem)
  SUBROUTINE PCAF_rma_load_end(src)
  SUBROUTINE PCAF_rma_load(src, nelem)
where  INTEGER, INTENT(IN) :: dest, src, nelem
```

CAF **Synchronization Wrapper Routines**

```
  SUBROUTINE PCAF_sync_all()
  SUBROUTINE PCAF_sync_all(wait)
  SUBROUTINE PCAF_sync_team(team)
  SUBROUTINE PCAF_sync_team(team, wait)
  SUBROUTINE PCAF_sync_file(unit)
  SUBROUTINE PCAF_sync_memory()
  SUBROUTINE PCAF_start_critical()
  SUBROUTINE PCAF_end_critical()
where  INTEGER, INTENT(IN) :: unit
       INTEGER, INTENT(IN) :: wait(:), team(:)
```

Figure 3. PCAF Measurement Function Interface Specification

OPARI identifies the CAF statement, and adds a begin and end instrumentation event. The call to indicate the beginning of the event contains the destination of the store (normalized to the range 0 to `num_images()-1`) and the number of array elements being transferred; the end call only gets passed the destination. The barrier call (`sync_all`) is translated into a call of the corresponding *wrapper* function.

For some programs, complete instrumentation of all CAF constructs may be unnecessary or too intrusive; therefore, the amount of instrumentation can be controlled via command line options to OPARI, allowing to disable instrumentation by groups (e.g., synchronization functions or remote memory transfers).

Performance Measurement

Finally, the KOJAK measurement system was extended by implementing the necessary PCAF monitoring functions and wrapper routines and adding support for the handling of the new remote memory access event types. We chose to implement our approach within the KOJAK framework, as KOJAK is very portable and supports all major HPC computing platforms. Also, this way, we could re-use many of KOJAK's features like event trace buffer management, generation, and conversion. Finally, it allows us not only to analyze plain CAF applications but also hybrid programs using any

Copyright © 2006 John Wiley & Sons, Ltd.

combination of MPI, OpenMP, SHMEM, and CAF. A separate, new instrumentor just for CAF would probably be problematic in this respect, as the modifications done by two independent source-to-source preprocessors could conflict.

The PCAF interface is shown in Figure 3. The first group of functions are used for the instrumentation of CAF remote load and store operations. As explained in the last section, for each case, there is a ”begin/end” two call version as well as a single call one. The second group of functions are wrapper functions, one for each intrinsic synchronization and utility function defined for CAF. The list shown in Figure 3 is not complete in that regard; it only includes those functions which were supported by the Cray Fortran compiler we used for our experiments. Especially, the current draft standard for CAF [7] (as of February 2006) specifies a slightly different list of synchronization functions as well as new CAF collective functions like `CO_PRODUCT` or `CO_SUM`. Implementing support for all CAF intrinsic functions once the CAF standard stabilizes is very easy.

Since this monitoring API is open, and OPARI is a stand-alone tool, other performance analysis projects could use this infrastructure to also support CAF. For example, it would be very easy to implement a version of the PCAF monitoring library which (instead of tracing) just collects basic statistics (number of RMA transfers, amount of data transferred) for each participating image. Ideally, in the future, CAF compilers could support this interface directly.

Performance Visualization

For illustration of our performance analysis approach, we ran the Halo kernel benchmark on the Cray X1 system at the Oak Ridge National Laboratory, using 16 and 64 processors. The Halo benchmark simulates a halo border exchange with the four different synchronization methods CAF provides. The exchange procedure has the following outline:

```
S1          HINS(1:3*n)   = HOEW(1:3*n)
S2          CALL synchronization method
S3          HONS(1:n)     = HINS(1:n)[MYPEN]
S4          HONS(n+1:3*n) = HINS(n+1:3*n)[MYPES]
S5          CALL synchronization method
S6          HIEW(1:3*n)   = HONS(1:3*n)
S7          CALL synchronization method
S8          HOEW(1:n)     = HIEW(1:n)[MYPEW]
S9          HOEW(n+1:3*n) = HIEW(n+1:3*n)[MYPEE]
S10         CALL synchronization method
```

During each iteration, the following events from our event model occur: `S2` a synchronization call; `S3` a remote load of n elements from the north neighbor; `S4` a remote load of $2n$ elements from the south neighbor; `S5` a synchronization call; `S7` another synchronization call; `S8` a remote load of n elements from the west neighbor; `S9` a remote load of $2n$ elements from the east neighbor; and finally `S10` a synchronization call. For each synchronization method, this procedure is repeated 5 times per iteration, with 10 iterations being executed with n varying from 2 to 1024 in powers of 2.

Copyright © 2006 John Wiley & Sons, Ltd.

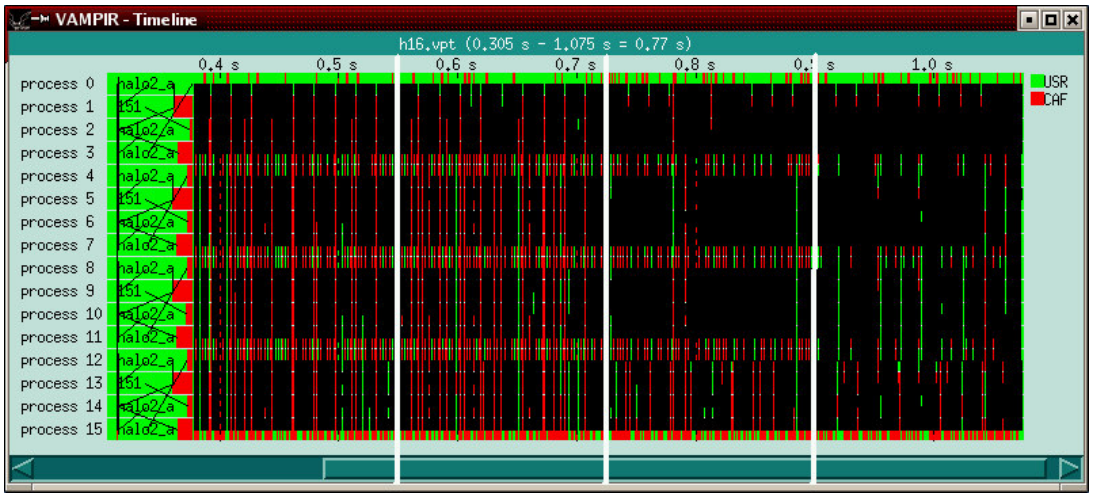

sync_all sync_all(wait) sync_team sync_team(wait)

Figure 4. Timeline view of the complete Halo benchmark using 16 processors

The original Fortran 77 code for the benchmark actually used separate loops for each of the CAF assignment statements (`S1, S3, S4, S6, S8, S9`). For example, statement `S3` was coded as

```
...
DO i=1,n
    HONS(i) = HINS(i)[MYPEN]
END DO
...
```

Changing the code to use Fortran 90 array syntax had several advantages: First, instrumenting and measuring the improved version generates considerably less overhead (2 PCAF calls instead of $2 \times n$). For the Halo benchmark, which is an extreme case as it only consists of CAF remote memory accesses and synchronization calls, it meant that we could not successfully perform a measurement of the original code, as we ran out of disk space for the resulting huge trace files. Secondly, the compact array syntax form very likely has more performance stability across implementations and architectures. Ideally, a smart compiler would generate the same code for both forms, however, using the compact form ensures that even a non-optimizing compiler generates code for one remote block-transfer instead of a loop of a lot of short remote memory operations. Especially on a system which does not have RDMA hardware support, this can make a huge difference. Finally, we think that the improved code is easier to understand and to maintain.

Figure 4 shows the timeline view of the Halo benchmark running with 16 processors. The four phases of the code (marked with white lines in the figure) can easily be identified due to the different communication behavior of each of the synchronization methods. The communication pattern between

Copyright © 2006 John Wiley & Sons, Ltd.

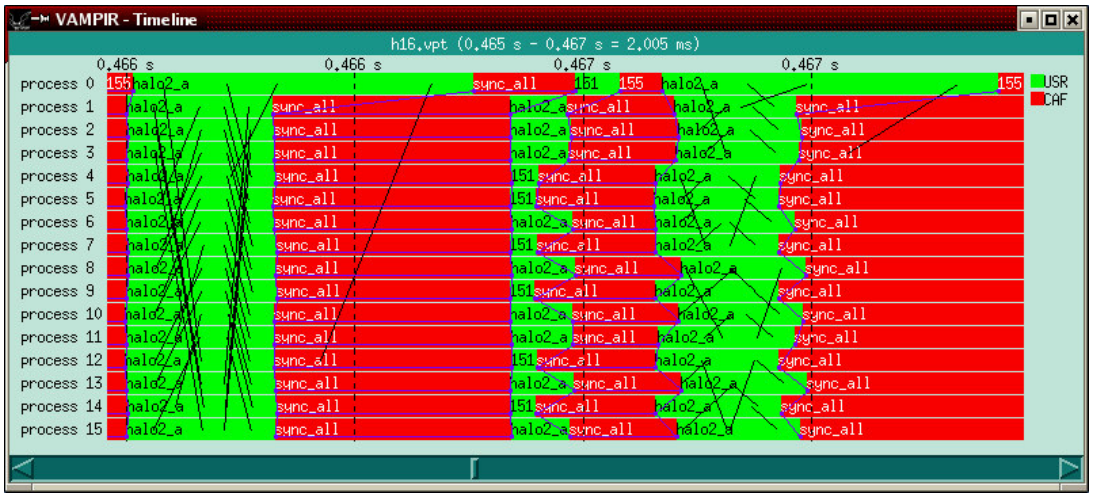

(a) One exchange using `sync_all`

(b) One exchange using `sync_team(wait)`

Figure 5. Partial timeline views of the Halo benchmark

processors, as well as the amount of data exchanged, can be observed with the pair-wise communication statistics view, shown in Figure 6 (left).

Figure 5 (a) and Figure 5 (b) show a section of the timeline corresponding to a full exchange (one call to the exchange subroutine) for `sync_all` and `sync_team(wait)` synchronization methods,

Copyright © 2006 John Wiley & Sons, Ltd.

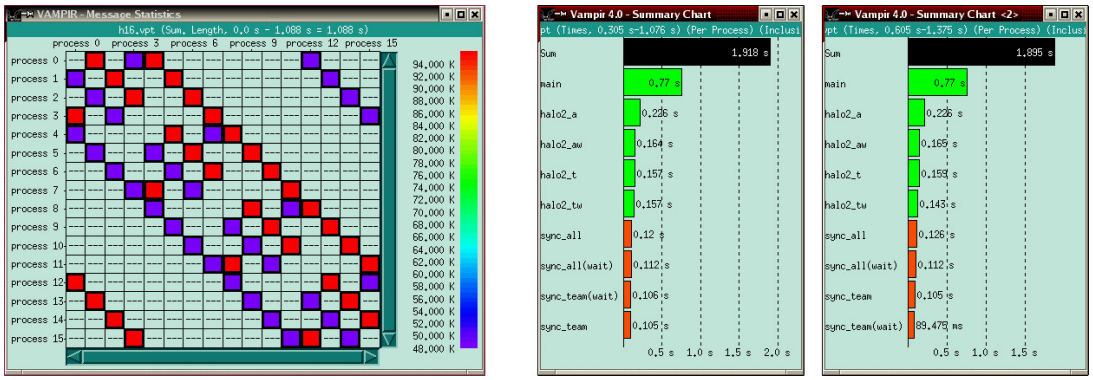

(16 processor run) (64 processor run)

Figure 6. Message statistics view of the Halo benchmark using 16 processors (left) and Summary Chart View of Function times running on 16 and 64 processors (right)

respectively. We observe that the region corresponding to the `sync_team(wait)` synchronization method is much more irregular (unsynchronized) than the one for the `sync_all`, where the waiting times are longer, due to the global synchronization.

Finally, on Figure 6 (right), we observe the time spent on each synchronization method for the 16 and 64 processors runs, respectively. We notice that with the increase of number of processors, the `sync_team(wait)` method performs significantly better in this benchmark than the `sync_all` method, going from about 10% faster with 16 processors to about 30% faster with 64 processors.

Conclusion

The CAF parallel programming language extends Fortran 2003 providing a simple technique for accessing and managing distributed data objects. This language-level abstraction hides much of the complexity of managing communication, but, unfortunately, this also makes diagnosing performance problems much more difficult. In this paper, we have proposed one approach to solve this problem. Our solution uses a portable source-to-source translator to allow performance instrumentation, data collection, trace generation, and performance visualization of Co-Array Fortran applications, implemented as an extension of the KOJAK performance analysis toolset. We illustrated this approach with performance visualization of a Co-Array Fortran version of the Halo kernel benchmark, running on a Cray X1 system and using the VAMPIR event trace visualization tool. Our initial results are promising; we can obtain statistical quantification and graphical presentation of CAF communication and synchronization characteristics.

In view of our experience with analyzing CAF applications with the KOJAK toolkit, we are considering several extensions to our work: First, we need to determine the benefits of this approach

Copyright © 2006 John Wiley & Sons, Ltd.

for real-world CAF and hybrid MPI, SHMEM, and CAF applications. Second, it would be interesting to extend KOJAK's automated analysis to also cover CAF constructs. Last, we need to synchronize our work with some recent work on a proposed performance tool interface for global address space languages [9].

ACKNOWLEDGEMENT

The authors would like to thank Bill Long, from Cray Inc., for his very helpful explanations and clarifications of the CAF specification.

REFERENCES

1. E. Ayguadé, M. Brorsson, H. Brunst, H.-C. Hoppe, S. Karlsson, X. Martorell, W. E. Nagel, F. Schlimbach, G. Utrera, and M. Winkler. OpenMP Performance Analysis Approach in the INTONE Project. In *Proceedings of the Third European Workshop on OpenMP - EWOMP'01*, September 2001.
2. R. Bell, A. D. Malony, and S. Shende. A Portable, Extensible, and Scalable Tool for Parallel Performance Profile Analysis. In *Proceedings of Euro-Par 2003*, pages 17–26, 2003.
3. S. Browne, J. Dongarra, N. Garner, G. Ho, and P. Mucci. A Portable Programming Interface for Performance Evaluation on Modern Processors. *The International Journal of High Performance Computing Applications*, 14(3):189–204, Fall 2000.
4. J. Caubet, J. Gimenez, J. Labarta, L. DeRose, and J. Vetter. A Dynamic Tracing Mechanism for Performance Analysis of OpenMP Applications. In *Proceedings of the Workshop on OpenMP Applications and Tools - WOMPAT 2001*, pages 53 – 67, July 2001.
5. L. DeRose, B. Mohr, and S. Seelam. Profiling and Tracing OpenMP Applications with POMP Based Monitoring Libraries. In *Proceedings of Euro-Par 2004*, pages 39–46, September 2004.
6. Marc-André Hermanns, Bernd Mohr, and Felix Wolf. Event-based Measurement and Analysis of One-sided Communication. In *Proceedings of Euro-Par 2005*, September 2005.
7. International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). *ISO/IEC/JTC1/SC22/WG5-N1649 - WG5 Standing Document 5*, January 2006. http://www.nag.co.uk/sc22wg5/WG5Projects.
8. S. Kim, B. Kuhn, M. Voss, H.-C. Hoppe, and W. Nagel. VGV: Supporting Performance Analysis of Object-Oriented Mixed MPI/OpenMP Parallel Applications. In *Proceedings of the International Parallel and Distributed Processing Symposium*, April 2002.
9. Adam Leko, Dan Bonachea, Hung-Hsun Su, Bryan Golden, Hans Sherburne, and Alan D. George. GASP: A Performance Tool Interface for Global Address Space Languages. Technical report, Electrical Engineering Dept., University of Florida and Computer Science Dept., University of California at Berkeley, November 2005.
10. K. A. Lindlan, Janice Cuny, A. D. Malony, S. Shende, B. Mohr, R. Rivenburgh, and C. Rasmussen. A Tool Framework for Static and Dynamic Analysis of Object-Oriented Software with Templates. In *Proceedings of Supercomputing 2000*, November 2000.
11. B. Mohr, A. Mallony, H.-C. Hoppe, F. Schlimbach, G. Haab, and S. Shah. A Performance Monitoring Interface for OpenMP. In *Proceedings of the fourth European Workshop on OpenMP - EWOMP'02*, September 2002.
12. Bernd Mohr, Allen Malony, Sameer Shende, and Felix Wolf. Design and Prototype of a Performance Tool Interface for OpenMP. *The Journal of Supercomputing*, 23:105–128, 2002.
13. W. Nagel, A. Arnold, M. Weber, H.-C. Hoppe, and K. Solchenbach. Vampir: Visualization and Analysis of MPI Resources. *Supercomputer*, 12:69–80, January 1996.
14. R. W. Numrich and J. K. Reid. Co-Array Fortran for Parallel Programming. *ACM Fortran Forum*, 17(2), 1998.
15. Felix Wolf and Bernd Mohr. Automatic Performance Analysis of Hybrid MPI/OpenMP Applications. *Journal of Systems Architecture, Special Issue 'Evolutions in parallel distributed and network-based processing'*, 49(10–11):421–439, November 2003.
16. C. Wu, A. Bolmarcich, M. Snir, D. Wootton, F. Parpia, A. Chan, E. Lusk, and W. Gropp. From trace generation to visualization: A performance framework for distributed parallel systems. In *Proceedings of Supercomputing 2000*, November 2000.

Copyright © 2006 John Wiley & Sons, Ltd.