



X-KAAPI Fortran programming interface

Thierry Gautier, Fabien Lementec, Vincent Faucher

► To cite this version:

Thierry Gautier, Fabien Lementec, Vincent Faucher. X-KAAPI Fortran programming interface. [Technical Report] RT-0429, INRIA. 2011, pp.20. hal-00727795

HAL Id: hal-00727795

<https://inria.hal.science/hal-00727795>

Submitted on 4 Sep 2012

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.



X-KAAPI Fortran programming interface

Thierry Gautier, Fabien Lementec, Vincent Faucher

**TECHNICAL
REPORT**

N° 429

November 2011

Project-Teams MOAIS



X-KAAPI Fortran programming interface

Thierry Gautier*, Fabien Lementec†, Vincent Faucher‡

Project-Teams MOAIS

Technical Report n° 429 — November 2011 — 17 pages

Abstract: This report defines the X-KAAPI Fortran programming interface.

Key-words: parallel computing, X-KAAPI, Fortran

* INRIA

† INRIA

‡ CEA, DEN, DANS, DM2S, SEMT, DYN, F-91191 Gif-sur-Yvette, France

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

X-KAAPI Fortran programming interface

Résumé : Ce rapport décrit l'interface Fortran du logical X-KAAPI.

Mots-clés : environnement de programmation parallèle, X-KAAPI, Fortran

Contents

1	Software installation	4
2	Initialization and termination	5
3	Library versioning	6
4	Concurrency	7
5	Adaptive grains	8
6	Performance	10
7	Independent loops	11
8	Dataflow programming	13
9	Parallel regions	15
10	Synchronization	16
11	Error codes	17

1 Software installation

X-KAAPI is both a programming model and a runtime for high performance parallelism targeting multicore and distributed architectures. It relies on the work stealing paradigm. X-KAAPI was developed in the MOAIS INRIA project by Thierry Gautier, Fabien Le Mentec, Vincent Danjean and Christophe Laferrière in the early stage of the library.

In this report, only the programming model based on the C API is presented. The runtime library comes also with a full set of complementary programming interfaces: C, C++ and STL-like interfaces. The C++ and STL interfaces, at a higher level than the C interface, may be directly used for developing parallel programs or libraries.

Supported Platforms

X-KAAPI targets essentially SMP and NUMA platforms. The runtime should run on any system providing:

- a GNU toolchain (4.3),
- the pthread library,
- Unix based environment.

It has been extensively tested on the following operating systems:

- GNU-Linux with x86_64 architectures,
- MacOSX/Intel processor.

There is no version for Windows yet.

X-KAAPI Contacts

If you wish to contact the XKaapi team, please visit the web site at:

<http://kaapi.gforge.inria.fr>

2 Initialization and termination

2.1 Synopsis

```
INTEGER*4 FUNCTION KAAPIF_INIT(INTEGER*4 FLAGS)
INTEGER*4 FUNCTION KAAPIF_FINALIZE()
```

2.2 Description

KAAPIF_INIT initializes the runtime. It must be called once per program before using any of the other routines. If successful, there must be a corresponding *KAAPIF_FINALIZE* at the end of the program.

2.3 Parameters

- *FLAGS*: if not zero, start only the main thread to avoid disturbing the execution until tasks are actually scheduled. The other threads are suspended waiting for a parallel region to be entered (refer to *KAAPIF_BEGIN_PARALLEL*).

2.4 Return value

Refer to the *Error codes* section.

2.5 Example

Refer to examples/kaapif/foreach

```
PROGRAM MAIN
  INTEGER*4 E

  E = KAAPIF_INIT(1)
  ...
  E = KAAPIF_FINALIZE()

END PROGRAM MAIN
```

3 Library versioning

3.1 Synopsis

```
SUBROUTINE KAAPIF_GET_VERSION(CHARACTER(40) VERSION)
```

3.2 Description

KAAPIF_GET_VERSION retrieves a string describing the runtime version. Currently, the string is the git last commit hash in base64.

3.3 Parameters

- *VERSION*: An array large enough to hold the version.

3.4 Return value

None.

3.5 Example

Refer to `examples/kaapif/get_version`

```
PROGRAM MAIN

    CHARACTER(40) VERSION

    CALL KAAPIF_INIT(1)
    CALL KAAPIF_GET_VERSION(VERSION)
    CALL KAAPIF_FINALIZE()

    WRITE(*, *) 'version: ', VERSION(1:40)

END PROGRAM MAIN
```

4 Concurrency

4.1 Synopsis

```
INTEGER*4 FUNCTION KAAPIF_GET_CONCURRENCY()  
INTEGER*4 FUNCTION KAAPIF_GET_THREAD_NUM()
```

4.2 Description

Concurrency related routines.

4.3 Return value

KAAPIF_GET_CONCURRENCY returns the number of parallel thread available to the X-KAAPI runtime.

KAAPIF_GET_THREAD_NUM returns the current thread identifier. Note it should only be called in the context of a X-KAAPI thread.

For both functions, a negative value means an error occurred. Refer to the *Error codes* section.

4.4 Example

Refer to examples/kaapif/foreach

5 Adaptive grains

5.1 Synopsis

```

INTEGER*4 FUNCTION KAAPIF_SET_GRAINS
(
  INTEGER*4 PAR_GRAIN,
  INTEGER*4 SEQ_GRAIN
)

```

5.2 Description

KAAPIF_SET_GRAINS sets the adaptive loop grains. Grains are used to amortize the act of extracting work for both the parallel (ie. during a steal) and sequential (ie. during a pop) executions. Guessing those grains is problem specific. The general intuition is that they should be inversely proportionnal to a single task processing time (given a task processing time is constant across execution). The parallel grain should be greater than the sequential grain, since a steal operation requires a bigger task to be amortized.

Note that this routine sets global variables used by a subsequent calls to *KAAPIF_FOREACH* family functions.

To avoid reentrancy issues, *KAAPIF_SET_GRAINS* should be called just before the corresponding *KAAPIF_FOREACH* function.

5.3 Parameters

- *PAR_GRAIN*: below this size, a task cannot be split for subsequent parallel execution. default to 32.
- *SEQ_GRAIN*: the size used by the sequential execution to extract work from its local workqueue. default to 16.

5.4 Return value

Refer to the *Error codes* section.

5.5 Example

```

PROGRAM MAIN
  INTEGER*4 E

  ...
  E = KAAPIF_INIT(1)

```

```
E = KAAPIF_SET_GRAINS(32, 32)
E = KAAPIF_FOREACH(...)
...
```

```
END PROGRAM MAIN
```

6 Performance

6.1 Synopsis

```
REAL*8 FUNCTION KAAPIF_GET_TIME()
```

6.2 Description

Capture the current time. Used to measure the time spent in a code region.

6.3 Parameters

None.

6.4 Return value

The current time, in microseconds.

6.5 Example

Refer to examples/kaapif/foreach

```
PROGRAM MAIN
  REAL*8 START
  REAL*8 STOP
  INTEGER*4 E

  E = KAAPIF_INIT(1)
  START = KAAPIF_GET_TIME()
  ...
  STOP = KAAPIF_GET_TIME()
  E = KAAPIF_FINALIZE()

  WRITE(*, *) STOP - START

END PROGRAM MAIN
```

7 Independent loops

7.1 Synopsis

```

INTEGER*4 FUNCTION KAAPIF_FOREACH
(
  INTEGER*4 FIRST, INTEGER*4 LAST,
  INTEGER*4 NARGS,
  BODY,
  ...
)

INTEGER*4 FUNCTION KAAPIF_FOREACH_WITH_FORMAT
(
  INTEGER*4 FIRST, INTEGER*4 LAST,
  INTEGER*4 NARGS,
  BODY,
  ...
)

```

7.2 Description

Those routines run a parallel loop over the range $[FIRST, LAST]$ (note this is an **inclusive** interval). The loop body is defined by *BODY* whose arguments are given in parameters. It must have the following prototype:

```
SUBROUTINE BODY(I, J, TID, ...)
```

- $[I, J]$ the subrange to process (note that interval is inclusive)
- *TID* the thread identifier

7.3 Parameters

- *FIRST*, *LAST*: the iteration range indices, inclusive.
- *NARGS*: the argument count
- *BODY*: the function body to be called at each iteration
- ...: the arguments passed to *BODY*. For *KAAPIF_FOREACH_WITH_FORMAT*, refer to the *KAAPIF_SPAWN* documentation.

7.4 Return value

Refer to the *Error codes* section.

7.5 Example

Refer to examples/kaapif/foreach

Refer to examples/kaapif/foreach_with_format

```
! computation task entry point
SUBROUTINE OP(I, J, TID, ARRAY)
  DO K = I, J
    ! process ARRAY(K)
    ...
  END DO
  RETURN
END

PROGRAM MAIN
  INTEGER*4 E

  ...
  ! apply the OP routine on ARRAY[1:SIZE]
  E = KAAPIF_FOREACH(1, SIZE, 1, OP, ARRAY)

  ...
  ! version with format
  ! as above, ARRAY is the only argument
  !
  E = KAAPIF_FOREACH_WITH_FORMAT
  (
    ! iterated range and argument count are the same as above
    1, SIZE, 1, OP,
    ! ARRAY an array of SIZE double
    ! ARRAY elements are read and written by the task
    ARRAY, KAAPIF_TYPE_DOUBLE, SIZE, KAAPIF_MODE_RW
  )
  ...

END PROGRAM MAIN
```

8 Dataflow programming

8.1 Synopsis

```

INTEGER*4 FUNCTION KAAPIF_SPAWN
(
  INTEGER*4 NARGS,
  BODY,
  ...
)

```

8.2 Description

Create a new computation task implemented by the function *BODY*.

BODY is called with the user specified arguments, there is no argument added by XKAAPI:

```

SUBROUTINE BODY(ARG0, ARG1, ...)

```

Each task parameter is described by 4 successive arguments including:

- the argument *VALUE*,
- the parameter *TYPE*,
- the element *COUNT*,
- the access *MODE*.

TYPE is one of the following:

- KAAPIF_TYPE_CHAR=0,
- KAAPIF_TYPE_INT=1,
- KAAPIF_TYPE_REAL=2,
- KAAPIF_TYPE_DOUBLE=3.

If a parameter is an array, *COUNT* must be set to the array size. For a scalar value, it must be set to 1.

MODE is one of the following:

- KAAPIF_MODE_R=0 for a read access,
- KAAPIF_MODE_W=1 for a write access,
- KAAPIF_MODE_RW=2 for a read write access,
- KAAPIF_MODE_V=3 for a parameter passed by value.

8.3 Parameters

- *NARGS*: the argument count.
- *BODY*: the task body.
- ...: the *VALUE*, *TYPE*, *COUNT*, *MODE* tuple list.

8.4 Return value

Refer to the *Error codes* section.

8.5 Example

Refer to examples/kaapif/dfg

```

! computation task entry point
SUBROUTINE OP(A, B)
  ! task user specific code
  ...
  RETURN
END

PROGRAM MAIN
  INTEGER*4 E
  ...
  ! spawn a task implemented by the OP routine
  E = KAAPIF_SPAWN(2, OP,
    ! argument[0]
    & 42
    & KAAPIF_TYPE_DOUBLE,
    & 1,
    & KAAPIF_MODE_V,
    ! argument[1]
    & 42,
    & KAAPIF_TYPE_DOUBLE,
    & 1,
    & KAAPIF_MODE_V)
  ...
END PROGRAM MAIN

```

9 Parallel regions

9.1 Synopsis

```
INTEGER*4 FUNCTION KAAPIF_BEGIN_PARALLEL()  
INTEGER*4 FUNCTION KAAPIF_END_PARALLEL(INTEGER*4 FLAGS)
```

9.2 Description

KAAPIF_BEGIN_PARALLEL and *KAAPIF_END_PARALLEL* mark the start and the end of a parallel region. Regions are used to wakeup and suspend the X-KAAPI system threads so they avoid disturbing the application when idle. This is important if another parallel library is being used. Whether threads are suspendable or not is controlled according by the *KAAPIF_INIT* parameter.

9.3 Parameters

- *FLAGS*: if zero, an implicit synchronization is inserted before leaving the region.

9.4 Return value

Refer to the *Error codes* section.

9.5 Example

Refer to examples/kaapif/dfg

```
PROGRAM MAIN  
  INTEGER*4 E  
  
  ...  
  E = KAAPIF_BEGIN_PARALLEL()  
  ...  
  E = KAAPIF_END_PARALLEL(1)  
  ...  
  
END PROGRAM MAIN
```

10 Synchronization

10.1 Synopsis

```
INTEGER*4 FUNCTION KAAPIF_SYNC()
```

10.2 Description

Synchronize the sequential with the parallel execution flow. When this routine returns, every computation task has been executed and memory is consistent for the processor executing the sequential flow.

10.3 Return value

Refer to the *Error codes* section.

10.4 Example

Refer to examples/kaapif/dfg

```
PROGRAM MAIN
  INTEGER*4 E

  ...
  E = KAAPIF_SYNC()
  ...

END PROGRAM MAIN
```

11 Error codes

When indicated, a routine may return one of the following error code:

- KAAPIF_SUCCESS=0: success
- KAAPIF_ERR_FAILURE=-1: generic error code
- KAAPIF_ERR_EINVAL=-2: invalid argument
- KAAPIF_ERR_UNIMPL=-3: feature not implemented



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-0803