

SME-C v0.5

C99 with pragma and API for parallel execution, streaming,
processor mapping and communication generation

SMECY Artemis European Project

Rémi BARRÈRE (Remi.Barrere@thalesgroup.com)
Marcel BEEMSTER (marcel@ace.nl)
Ronan KERYELL (Ronan.Keryell@silkan.com)

December 10, 2012

Contents

1	Introduction	2
1.1	SMECY programming model	2
1.2	SME-C Streaming Model	4
1.2.1	Background	4
1.2.2	Goals	4
1.2.3	Status as of June 2011	4
1.2.4	Future Extensions (not Implementation Related)	5
1.3	Reference documents	5
2	Intermediate representation use cases	5
2.1	Direct programming	6
2.2	System high-level synthesis	6
2.3	Hardware high-level synthesis	6
3	Exemples	6
3.1	Program with contiguous memory transfers	6
3.2	Program with non-contiguous memory transfers	9
3.3	Pipelined example	10
3.4	Remapping example	10
4	Description of high level process structure	13
5	SMECY embedded C language	13
6	Description of the SMECY directives	14
6.1	OpenMP support	14
6.2	Mapping on hardware	14
6.3	Producer/consumer information	15
6.3.1	Function arguments	15
6.3.2	Global variables	16

6.4	Stream programming	16
6.5	Labelling statement	16
6.6	Remapping specification	16
6.7	Hardware specific pragma	16
7	SMECY high-level APIs	17
7.1	OpenMP	17
7.2	MultiCore Association APIs	17
7.3	Multicore Communication API (MCAPI)	17
7.3.1	MTAPI	17
7.3.2	MRAPI	17
7.4	NPL API	18
7.5	EdkDSP/ASVP API	18
7.6	OpenCL	18
8	Some design patterns for STHORM SME-C	18
8.1	Same computation on all the PEs	18
8.2	Same computation on all the PEs of a cluster	19
8.3	Stream computation through the PEs of a cluster	20
8.4	Systolic computation with the PEs of a cluster	21
8.5	Round robin computation on the clusters	23
9	Conclusion	24

1 Introduction

In the SMECY project we want to use C source code as a portable intermediate representation (IR) between tools from high-level tools down to lower-level tools because of its good trade-off between expressiveness and readability, without compromising portability.

The targets envisioned in SMECY are heterogeneous multicore systems with shared memory or not, with various hardware accelerators, such as ASIC, ASIP, FPGA, GPU, specialized vector processors, partially reconfigurable accelerators...

Unfortunately, since it is undecidable to get high-level properties from such programs, we use decorations to help tools to understand program behaviour and generate codes for some hardware targets. We try to keep clear decorations, easy to understand, so that SMECY C (summarized as SME-C in the following but also referred as IR-1 in the project) can also be used as a programming language.

A SMECY program contains various functions that may be executed on various processors, accelerators, GPU... that may consume and produce data from different physical memory spaces.

Since we want to express also performance on given platforms, we keep the opportunity to have platform-specific pragma and API or specialized intrinsic types and functions, for example to express use of special hardware accelerator functions or operations.

The hardware specific pragma and intrinsics are to be defined between software and hardware partners involved in various use cases. But it may not be possible to address all the programming models and platforms envisioned in the project.

The description of a reference compiler, `smecc`, is given in a companion report, “Some implementation details from the `smecc` compiler”. Several tools producing or consuming SME-C have been developed in the SMECY project.

1.1 SMECY programming model

The programming model is based on C processes, with a virtual shared memory and threads *à la* OpenMP. Since we may have quite asynchronous processes in a real application description

or at the back-end level in the execution model, we can cope indeed with different C processes communicating with an API.

We add mapping information stating on which hardware part a function is to be placed and run.

The programming model exposed in the following is based on an OpenMP SMP model because of its (rather) simple readability, elegance, old background and wide acceptance. We make the hypothesis that a SMECY program is a correct OpenMP program that can be executed in sequential with a C OpenMP-free compiler (just by ignoring OpenMP `#pragma`) and in parallel on a SMP target (such an ARM or *x86* machine) by using an OpenMP compiler *with the same semantics*. Since we use C (by opposition to a DSL) as an internal representation, we choose this behaviour to stick to standard behaviour as much as possible. This is known as the sequential equivalence. Since we can cope different results for performance reasons from executing in parallel non associative floating point operations, we deal with only *weak* sequential equivalence instead of *strong* sequential equivalence.

Of course, this model is incompatible with a real hardware target envisioned in SMECY, so we need to add hints in the code explaining memory dependencies at the function call levels. Since it is quite difficult to describe general dependencies, we approximate memory dependencies with rectangles (and more generally hyperparallelepiped in any dimension) that can be read, written or both. We think these abstractions are good trade-offs between expressiveness (and what a programmer can endure...) and hardware capabilities. Even if there is some similitude with HPF (High Performance Fortran) or XMP pragmas, we do not deal with strides¹.

With this information, the tools can guess the communication to generate between the different functions and memory spaces to emulate the global OpenMP memory semantics.

The neat side effect is that we have the same global program executed on all the platforms (sequential, real OpenMP and SMECY) with the same semantics and we can see the sequential version as a functional simulator of the SMECY application and the real OpenMP version as a parallelized quicker version of this simulator.

It is also easy to debug the application, but also all the SMECY tools used or developed in the project.

To be able to address real hardware from the C level with special needs:

- to specify hardware register names;
- define input/output routines specifically but in portable way;
- define fixed point arithmetic computations with a given precision;
- specific data size;
- accumulator register (DSP...);
- different memory spaces that can be chosen specifically to optimize storage and speed (DSP, hardware accelerators with scratch pad memory...);
- saturated arithmetic.

For all these we rely on the TR 18037 Embedded C standard, supported for example by ACE tools.

To describe different processes communicating together in an asynchronous way, we do not have anymore a sequential equivalence and then do not use pragma to express this. So we use a simple communication, synchronization and threading API. Since we target embedded systems with a light efficient implementation, we can rely on such standard API as the ones of the MultiCore Association: MCAPI (communication), MRAPI (synchronization) and MTAPI (threading).

¹Indeed, by using some higher dimensional arrays than the arrays used in the application, you can express them... So may be we should allow to express them in the syntax?

As modern programs need a clear documentation, such as with Doxygen marking style providing meta information on different elements of the program, we can use this (hopefully correct) information to help the compiling process itself.

Since tools are to be oriented for a specific target, taking into account various hardware and compilation parameters that are to be kept orthogonally to application sources, those parameters are kept aside in some description files that flow between tools. These files may be represented with an XML syntax (using a SMECY naming space) or by even simpler format (JSON, YAML...).

1.2 SME-C Streaming Model

This section describes the state of the newly developed Streaming Annotation to C for use in SMECY. This is a work in progress, so many imperfections still need to be fixed and many extensions are both desirable and possible.

1.2.1 Background

In several of the SMECY applications, it appears to be a suitable parallel computation model to exploit application parallelism. This is the case in particular for two of the Use-Case-A applications, the M5 protocol analyser and the OFDM application.

Streaming has the advantage that it fits well to parallelizing applications that process data in a pipelined fashion, while there may still be strong data dependencies that require data to be processed sequentially at specific points in the pipeline. This is opposed to a data-parallel model, where such dependencies do not exist.

Streaming can exploit both a coarse-grained level of parallelism and parallelism at a fine-grained level. At the fine-grained level, the overhead of passing data between processing nodes in the pipeline must of course be minimized. For fine-grained parallelism, this may require hardware support. To achieve good load balancing, it is important that the processing nodes have a comparable grain size. If one node required much more processing than the others, it becomes the bottleneck and no parallelism can be exploited.

Streaming has the advantage of data locality. Data is passed around in the distributed point-to-point network. Such networks can be implemented far more efficiently than, for example, a shared memory connection between the processing nodes.

1.2.2 Goals

1. The SME-C streaming model is an annotation on top of a valid sequential program. Thus, the program should also work when the streaming annotations are ignored. This eases program development because it allows the validation of the program in an sequential environment.
2. In principle, a streaming program can also be written using OpenMP, but it requires the explicit coding of the communication between nodes. In the SME-C streaming model, the communication is derived from the program source and generated by the compiler. This is very important for the parallel performance tuning of the application because it allows to experiment with different load partitioning without having to re-program process communication.
3. The SME-C streaming model is mapped onto a few basic parallel machine primitives for process creation and communication. This eases the task of retargeting to different target architectures, with shared or distributed memory between the nodes.

1.2.3 Status as of June 2011

ACE has currently implemented a source-to-source compiler that accepts C programs with the streaming annotation and produces a partitioned program with separate processes (nodes) for

each of the nodes in the stream. In addition, it implements the communication links between the nodes.

The generated code includes library calls to implement the low-level tasks of process creation and synchronization. This small library of about 5 calls is currently implemented on top of (shared memory) POSIX pthreads. It is not hard to retarget this library to different underlying run-time systems.

To be used, the streaming model requires a part of the target application to be rewritten into a particular form, using a while loop and the SME-C stream annotations. Only this part needs to be passed through the source-to-source compiler. Hence, large parts of the application remain unmodified and do not need to be processed by the stream compiler.

Stream termination is currently not handled well. Stream termination has to be mapped from a sequential to a distributed decision process and the design and implementation of that is still to be done.

1.2.4 Future Extensions (not Implementation Related)

The highest priority for extension is to provide a mechanism for stream termination. The challenge here is to make the mechanism such that it still allows a natural programming style under sequential program interpretation.

To facilitate additional parallel performance tuning, a mechanism must be design to allow node replication. It would allow for a single node (that turns out to be a bottleneck) to be replicated into multiple nodes that run on multiple processors. Obviously this complicates the generated communication primitives.

Instead of generating communication primitives, the compiler can also limit itself to only partitioning and generate a communication graph for subsequent processing by tools such as SPEAR and BIPS in the SMECY project.

An extension is needed to pass (fixed) parameters and initial values into the streaming nodes that do not turn into communication.

For TVN's H264 application, certain compute intensive loops rely on array processing. Although there is parallelism in these loops, they cannot be easily mapped to a fully data-parallel implementation because of data dependencies. Given the nature of these data dependencies, it seems to be possible to transform them into a stream-processing model.

1.3 Reference documents

Besides the SMECY documentation, the reader should be knowledgeable of some work of the ISO/IEC JTC1/SC22/WG14 committee on the C language standard:

- ISO/IEC 9899 - Programming languages - C (Technical Corrigendum 3 Cor. 3:2007(E)) <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1256.pdf>
- TR 18037: Embedded C <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1169.pdf>
- Future C1X standard <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1548.pdf>

and other standards such as

- MCA (MultiCore Association) API: MCAPI, MRAPI & MTAPI, www.multicore-association.org
- HPF (High Performance Fortran) <http://hpff.rice.edu/versions/hpf2>

2 Intermediate representation use cases

To update, with OpenMP 4, XMP, C11, C++11..., OpenACC, OpenHMP...

Insert here the SMECY use-cases to develop generic use-cases

2.1 Direct programming

A programmer can program her application directly in SME-C with OpenMP and SMECY-specific pragma and API to target a SMECY platform.

It can be at rather high-level, by using only high-level pragma, or rather at a lower level, by using different communicating processes with the API and even specialized API and pragma for specific hardware.

A process written in C code with pragma and API can express a global host controlling process of an application or a local program in a specialized processor. And we may have many of such processes to express different producer and consumer Kahn processes interacting through a NoC in an asynchronous way.

2.2 System high-level synthesis

A compiler can take a sequential plain C, MATLAB, Scilab, NumPy or another language code, analyze and parallelize the code by adding automatically parallel and mapping pragma. This can be seen as a high-level synthesis at system level.

A tool such as Par4All can do these kinds of transformation.

2.3 Hardware high-level synthesis

A compiler can take a program with SMECY pragma and compile any call with a mapping of a given kind into some hardware configuration or program to be executed instead of the function and an API call to use this hardware part from the host program.

Since the pragma are designed to be concretely compilable, such a tool should be easy to do with a simple compilation framework, such as ROSE Compiler.

The SMECY API and intrinsics are chosen to be mapped quite straightforwardly to real hardware functions by the back-end.

3 Exemples

3.1 Program with contiguous memory transfers

During C memory transfer, if we work on arrays with the last dimension taken as a whole, the memory is contiguous and the programs often work even there are some aliasing such as using a 2D array zone as a linearized 1D vector.

The following program exposes this kind of code where some work sharing is done by contiguous memory blocks.

```
1 /* To compile this program on Linux, try:
```

```
    make CFLAGS='-std=c99 -Wall' pragma_example
```

```
    To run:
```

```
6 ./pragma_example; echo $?  
    It should print 0 if OK.
```

```
    You can even compile it to run on multicore SMP for free with
```

```
11 make CFLAGS='-std=c99 -fopenmp -Wall' pragma_example
```

```
    To verify there are really some clone() system calls that create the threads:  
    strace -f ./pragma_example ; echo $?
```

```
16 You can notice that the #pragma smecy are ignored (the project is  
    on-going :-)) but that the program produces already correct results in
```

```

    sequential execution and parallel OpenMP execution.

    Enjoy!
21     Remi.Barrere@thalesgroup.com
        Ronan.Keryell@hpc-project.com
        for ARTEMIS SMECY European project.
    */
26     #include <stdbool.h>

    /* function Gen
31     Example of old C89 array use-case where the size is unknown. Note that
        this implies some nasty access linearization with array with more than
        1 dimension.
    */
    void Gen(int *out, int size) {
36     // Can be executed in parallel
    #pragma omp parallel for
        for (int i = 0; i < size; i++)
            out [i] = 0;
    }
41

    /* function Add

        Nice C99 array with dynamic size definition. Note this implies having
46     array size given first
    */
    void Add(int size, int in[size], int out[size]) {
        // Can be executed in parallel
    #pragma omp parallel for
51     for (int i = 0; i < size; i++)
        out [i] = in [i] + 1;
    }

56     /* function Test */
    bool Test(int size, int in[size]) {
        bool ok = true;
        /* Can be executed in parallel, ok is initialized from global value and
            at loop exit ok is the  $\&\&$  operation between all the local ok
61     instances: */
    #pragma omp parallel for reduction(&&:ok)
        for (int i = 0; i < size; i++)
            /* We cannot have this simple code here:
                if (in [i] != 2)
66     exit(-1) ;
                because a loop or a fonction with exit() cannot be executed in parallel.

                Proof: there is a parallel execution interleaving that may execute
                some computations in some threads with a greater i that the one
                executing the exit() done on another thread. So the causality is
71     not respected.

                Anyway, in an heterogenous execution, just think about how to
                implement the exit() operating system call from an

```

```

76     accelerator... No hope. :-)

        So use a reduction instead and return the status for later
        inspection:
        */
81     ok &= (in[i] == 2);

        // Return false if at least one in[i] is not equal to 2:
        return ok;
    }
86

    /* main */
    int main(int argc, char* argv[]) {
        int tab[6][200];
91     // Gen is mapped on GPP 0, it produced (out) an array written to arg 1:
    #pragma smecy map(GPP, 0) arg(1, [6][200], out)
        /* Note there is an array linearization here, since we give a 2D array
        to Gen() that uses it . This is bad programming style, but it is just
        to show it can be handled in the model :-) */
96     Gen((int *) tab, 200*6);

        // Launch different things in parallel:
    #pragma omp parallel sections
    {
101     // Do one thing in parallel...
    #pragma omp section
    {
        /* Map this "Add" call to PE 0, arg 2 is communicated as input as an
        array of "int [3][200]", and after execution arg 3 is
106     communicated out as an array of "int [3][200]"

        Note the aliasing of the 2 last arguments. Just to show we can
        handle it. :*/
    #pragma smecy map(PE, 0) arg(2, [3][200], in) arg(3, [3][200], out)
111     Add(200*3, (int *) tab, (int *) tab);
    }
        // ...with another thing
    #pragma omp section
    {
116     /* Map this "Add" call to PE 1, arg 2 is communicated as input as an
        array of "int [3][200]" from address tab[3][0], that is the
        second half of tab, and after execution arg 3 is communicated out
        as an array of "int [3][200]", that is the second half of tab

121     Note the aliasing of the 2 last arguments. Just to show we can
        handle it. :*/
    #pragma smecy map(PE, 1) arg(2, [3][200], in) \
        arg(3, [3][200], out)
        Add(200*3, &tab[3][0], &tab[3][0]);
126     }
    }

        // Launch different things in parallel:
    #pragma omp parallel sections
131     {
    #pragma omp section
    {

```



```

136     #pragma smecy map(PE, 2) arg(2, [2][200], in) arg(3, [2][200], out)
        Add(200*2, (int *) tab, (int *) tab);
    }
    #pragma omp section
    {
141     #pragma smecy map(PE, 3) arg(2, [2][200], in) arg(3, [2][200], out)
        Add(200*2, &tab[2][0], &tab[2][0]);
    }
    #pragma omp section
    {
146     #pragma smecy map(PE, 4) arg(2, [2][200], in) arg(3, [2][200], out)
        Add(200*2, &tab[4][0], &tab[4][0]);
    }
    }
    // An example where arg 2 is just used as a whole implicitly:
    #pragma smecy map(GPP, 0) arg(2, in)
    bool result = Test(200*6, (int *) tab);
151    // Return non 0 if the computation went wrong:
    return !result;
}

```

3.2 Program with non-contiguous memory transfers

In the following example, we apply different computations on square pieces of the image, that do not have contiguous representation in memory. That is why we need to express restrictions on the use of the whole array.

```

/* To compile this program on Linux, try:
2
    make CFLAGS='-std=c99 -Wall' example_2D

    To run:
    ./example_2D; echo $?
7
    It should print 0 if OK.

    You can even compile it to run on multicore SMP for free with

12
    make CFLAGS='-std=c99 -fopenmp -Wall' example_2D

    To verify there are really some clone() system calls that create the threads:
    strace -f ./example_2D ; echo $?

    You can notice that the #pragma smecy are ignored (the project is
17
    on-going :-)) but that the program produces already correct results in
    sequential execution and parallel OpenMP execution.

    Enjoy!

22
    Ronan.Keryell@hpc-project.com
    for ARTEMIS SMECY European project.
*/

#include <stdlib.h>
27 #include "example_helper.h"

// Problem size
enum { WIDTH = 500, HEIGHT = 200 };

```

32

```

/* The main host program controlling and representing the whole
   application */
int main(int argc, char* argv[]) {
37  int image[HEIGHT][WIDTH];
    unsigned char output[HEIGHT][WIDTH];

    // Initialize with some values
    init_image(WIDTH, HEIGHT, image);
42
    #pragma omp parallel sections
    {
        // On one processor
        // We rewrite a small part of image:
47  #pragma smecy map(PE, 0) arg(3, inout, [HEIGHT][WIDTH] \
        / [HEIGHT/3:HEIGHT/3 + HEIGHT/2 - 1] \
        [WIDTH/8:WIDTH/8 + HEIGHT/2 - 1])
        square_symmetry(WIDTH, HEIGHT, image, HEIGHT/2, WIDTH/8, HEIGHT/3);

52    // On another processor
    #pragma omp section
        // Here let the compiler to guess the array size
    #pragma smecy map(PE, 1) arg(3, inout, / [HEIGHT/4:HEIGHT/4 + HEIGHT/2 - 1] \
        [3*WIDTH/8:3*WIDTH/8 + HEIGHT/2 - 1])
57    square_symmetry(WIDTH, HEIGHT, image, HEIGHT/2, 3*WIDTH/4, HEIGHT/4);

        // On another processor
    #pragma omp section
        // Here let the compiler to guess the array size
62  #pragma smecy map(PE, 1) arg(3, inout, / [2*HEIGHT/5:2*HEIGHT/5 + HEIGHT/2 - 1] \
        [WIDTH/2:WIDTH/2 + HEIGHT/2 - 1])
        square_symmetry(WIDTH, HEIGHT, image, HEIGHT/2, WIDTH/2, 2*HEIGHT/5);
    }
    // Here there is a synchronization because of the parallel part end
67
    // Since there
    normalize_to_char(WIDTH, HEIGHT, image, output);

    write_pgm_image("2D_example-output.pgm", WIDTH, HEIGHT, output);
72
    return EXIT_SUCCESS;
}

```

3.3 Pipelined example

The example shows a producer/consumer program. At every communication step, a full array of 128 integers is passed from the 'Producer' to the 'Consumer'. The streaming compiler finds out the sizes of the communication buffers from the program.

3.4 Remapping example

Some information can be in a given layout but needed in another layout to be used by a specific hardware accelerator.

```

#include <stdlib.h>
2 #include "example_helper.h"

```

```

1  #include <stdlib.h>
   #include <stdio.h>

   #define buffer_length 128
   typedef int data_buff[buffer_length];
6
   /* Produce a random buffer

      @param[out] data_buffer is an array initialized with random numbers
   */
11 void Produce(data_buff data_buffer) {
    for(int i = 0; i < buffer_length; i++)
        /* Note that rand() is not thread-safe but it is OK for this
           example */
        data_buffer[i] = rand();
16 }

   /* Compute the average value of an array and display it

      @param[in] data_buffer is the array to analyze
21 */
   void Consume(data_buff data_buffer) {
    double average = 0;
    for(int i = 0; i < buffer_length; i++)
        /* Note that rand() is not thread-safe but it is OK for this
           example */
26         average += data_buffer[i];
        // Normalize:
        average /= RANDMAX;
        average /= buffer_length;
31 printf("Average = %f\n", average);
    }

   int main() {
36     data_buff data_buffer;
        /* This while-loop is indeed to be executed in a pipelined way according
           to the following pragma: */
        #pragma smecy stream_loop
        while(1) {
41             // This pragma is optional indeed:
            #pragma smecy stage
                Produce(data_buffer);
            #pragma smecy stage
                Consume(data_buffer);
46         }
        return 0;
    }
}

```

Figure 1: Example of pipelined streamed loop.

```

// Problem size
enum { WIDTH = 500, HEIGHT = 200, LINE_SIZE = 100 };

7 /* Apply some pixel value inversion in a 1D array
   */
void
invert_vector(int line_size,
              int input_line[line_size],
12             int output_line[line_size]) {
    for(int i = 0; i < line_size; i++)
        output_line[i] = 500 - input_line[i];
}

17
/* The main host program controlling and representing the whole
   application */
int main(int argc, char* argv[]) {
    int image[HEIGHT][WIDTH];
22    unsigned char output[HEIGHT][WIDTH];

    // Initialize with some values
    init_image(WIDTH, HEIGHT, image);

27    // Draw 70 horizontal lines and map operation on 8 PEs:
    #pragma omp parallel for num_threads(8)
        for(int proc = 0; proc < 70; proc++)
            // Each iteration is on a different PE in parallel:
32    #pragma smecy map(PE, proc & 7) \
        arg(2, in, [1][LINE_SIZE]) \
        arg(3, out, [1][LINE_SIZE])
        // Invert an horizontal line:
        invert_vector(LINE_SIZE,
37                     &image[HEIGHT - 20 - proc][WIDTH/2 + 2*proc],
                     &image[HEIGHT - 20 - proc][WIDTH/2 + 2*proc]);

    /* Here we guess we have 5 hardware accelerators and we launch
       operations on them: */
    #pragma omp parallel for num_threads(5)
42    for(int proc = 0; proc < 5; proc++) {
        /* This is need to express the fact that our accelerator only accept
           continuous data but we want apply them on non contiguous data in
           the array */
        int input_line[LINE_SIZE];
47        int output_line[LINE_SIZE];
        /* We need to remap data in the good shape. The compiler should use
           the remapping information to generate DMA transfer for example and
           remove input_line array */
        SMECY_remap_int2D_to_int1D(HEIGHT, WIDTH, HEIGHT/3, 30 + 20*proc,
52                                LINE_SIZE, 1, image,
                                LINE_SIZE, input_line);
        // Each iteration is on a different PE in parallel:
        #pragma smecy map(PE, proc) arg(2, in, [LINE_SIZE]) arg(3, out, [LINE_SIZE])
        invert_vector(LINE_SIZE, input_line, output_line);
57        SMECY_remap_int1D_to_int2D(LINE_SIZE, output_line,
                                    HEIGHT, WIDTH, HEIGHT/3, 30 + 20*proc,
                                    LINE_SIZE, 1, image);
    }
}

```

```

62 // Convert int image to char image:
   normalize_to_char(WIDTH, HEIGHT, image, output);

   write_pgm_image("remapping_example-output.pgm", WIDTH, HEIGHT, output);
67 return EXIT_SUCCESS;
   }

```

4 Description of high level process structure

Here should be described the metadata on the process organization.

5 SMECY embedded C language

We take as input C99 (ISO/IEC 9899:2007) language with extensions for embedded systems (TR 18037).

Refer to these documents for more information.

```

__thread
Thread-Local Storage
_Thread_local
Doc. No.: WG14/N1364
Date: 2008-11-11
Reply to: Clark Nelson

```

<http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1364.htm>

```

C++
ISO/IEC JTC1 SC22 WG21 N2659 = 08-0169 - 2008-06-11 proposal
http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2008/n2659.htm

```

<http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1351.pdf>
 WG14 N1351

C Language support for multiprocessor application environments.

Walter Banks
 Byte Craft Limited
 Canada
 February 2009

```

ISO/IEC JTC1 SC22 WG14 N1275
Date: 2007-10-20
Reference number of document:
ISO/IEC TR 18037
Committee identification: ISO/IEC JTC1 SC22 WG14
SC22 Secretariat: ANSI
http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1275.pdf

```

It should be able to describe process instance, interconnection...

comment the following

WG14 N1386
Additions to ISO/IEC TR 18037 to support named execution
space.
Walter Banks
Byte Craft Limited
Canada
April 2009
Named execution addition to IEC/ISO 18037
<http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1386.pdf>

6 Description of the SMECY directives

The generic format of SMECY code decorations are language dependent, because if here we describe a SMECY IR implementation based on C, it is indeed more general.

- In C/C++:

```
#pragma smecy clause[[,clause]... newline
```

We can use \ at the end of line for continuation information.

- In Fortran:

```
!$smecy clause[[,clause]... newline
```

Use & at the end of line for continuation information.

- In other languages: use #pragma equivalent, if not available, use comments *à la* Fortran. For example in Python:

```
#$smecy clause[[,clause]... newline
```

Use also & at the end of line for continuation information

In implementations that support a preprocessor, the `_SMECY` macro name is defined to have the decimal value `yyyymm` where `yyyy` and `mm` are the year and month designations of the version of the SMECY API that the implementation supports. If this macro is the subject of a `#define` or a `#undef` preprocessing directive, the behavior is unspecified.

6.1 OpenMP support

SMECY is based on OpenMP 3.1. Since a SMECY platform is made at least from a (SMP) control processor, any OpenMP compliant program can run on it anyway. The SMECY tools can use more or less information from available OpenMP decorations available in the code.

The on-coming version 4 dealing more with SIMD instruction and heterogeneous computing seems promising as a starting point for a next version of SME-C.

6.2 Mapping on hardware

The mapping of a function call on a specific piece of hardware can be specified with

```
#pragma smecy map(hardware[, unit][, unit]...)  
some_function_call(...);
```

- *hardware* is a symbol representing a hardware component of a given target such as CPU, GPP, GPU, PE... They are target specific.

- *unit* are optional multidimensional instance number for a specific hardware part. This is typically an integer starting a 0 or for ST THORM 2 integers, the first one is the cluster number and the second one is the processor element number inside the cluster. This hardware number can be an expression of the environment to be able to have a loop managing different accelerators.

We can add an `if(scalar-expression)` to predicate hardware launching according some runtime expression to choose between hardware or local software execution, as in OpenMP with the same syntax. The idea is to be able to do a software execution if the data to process is too small compared to the latency of an hardware accelerator.

Recursion is not supported on hardware-mapped functions. If there are functions called from hardware-mapped functions, they will be automatically inlined (so no recursion allowed in them either). If a function is mapped to a more programmable hardware (GPP), recursions in these called functions should be allowed.

By default, call to hardware accelerators are synchronous, so you may launch the call into an OpenMP thread or into another MCA API process. But simpler way is to use the `async` keyword to launch in an asynchronous way, such as:

```
#pragma smecy map(...) async
```

But the it may be useful to way later for the production of an accelerator to be ready. This is done with the `wait` pragma such as in:

```
#pragma smecy wait(PE,2)
```

6.3 Producer/consumer information

To generate hardware communications where there is only a function call, the compiler need to figure out what is the memory zone used to execute the function and then what memory zone in written by the function *and* that will be used later in the program². From these information, copy-in and copy-out operations can be generated.

6.3.1 Function arguments

```
#pragma smecy arg(arg_id, arg_clause[, arg_clause]...)
    some_function_call(...);
```

Direction directive defines how the data flows between the caller and the function:

- `in`: the argument is used by the function;
- `out`: the argument is used by the function;
- `inout`: the argument is used and produced by the function;
- `unused`: the argument is not used and produced by the function.

Argument layout specifies how the argument is used in the function with:

- an optional *array_size_descriptor* such as `[n] [m]` expressing that the data is used from the callee point of view as such an array starting at the address given in parameter. If not specified, all the caller argument is used;
- an optional *array_range_descriptor* restriction such as `/[n:m] [2] [3:7]` expressing that the data is used from the callee point of view as an array with only this element ranges used. If not specified, all the array is used according to its size specified or not. If only some ranges are lacking, all the matching dimension is used. For example `/[4] []` matches the column 4 of an array.

²If a function produces something not used later, it is useless to get it back.

The more precise this description is and the less data transfers occur.

If the argument usage is independent from the call site they can be specified only at the function definition level instead of at each call site..

6.3.2 Global variables

Right now we do not deal with sharing information through global variables, because it is more difficult to track. Only function parameters are used to exchange information.

But we can imagine to map global variables with this clause:

```
#pragma smecy global_var(var, arg_clause[, arg_clause]...)
    some_function_call(...);
```

6.4 Stream programming

It is possible to stream a `while` loop in several pipeline stages that execute in parallel and pass information between stages/

The two pragmas used here are:

`#pragma smecy stream_loop:` this indicated the following `while` loop must be turned into a stream of processes;

`#pragma smecy stage:` this acts as a separator between groups of statements and define the boundary of pipeline stages. Only data passing over these separators is turned into communication.³ Note that the first stage pragma can be eluded since the begin of the loop body define a stage up to the next stage pragma.

You can refer to example shown on figure 1.

6.5 Labelling statement

Since ACE mentioned an interest to name part of a program from other external tools, such as to do some fine mapping, a labeling pragma has been added to name⁴ statements:

```
#pragma smecy label(name)
    any_statement;
```

6.6 Remapping specification

Since we always want a sequential equivalence, that means that the sequential code representing the computation on an accelerator really consume The easy HPF

To finish

6.7 Hardware specific pragma

³It used to be `#pragma smecy stream_node(n)` as an instruction to define both a stage node and a way to name it for another later use by ACE. But after København meeting, it appears it is better to use another way to label things only when needed. See § 6.5.

⁴Should it be an atom, a string, a number?

To be defined in collaboration with the various hardware suppliers of the project (ST P2012/STHOR. EdkD-SP/ASVP...).

7 SMECY high-level APIs

7.1 OpenMP

Since we support OpenMP pragma, we also support OpenMP API that allows for example:

- getting/setting the number of threads;
- getting the number of available processors on the current domain;
- manipulating locks.

7.2 MultiCore Association APIs

Since in this project we deal with communicating processes, asynchronous communications, synchronization, etc. we need an API to be able to express them from inside the processes. We rely on a standard API instead of reinventing on again. The MultiCore Association designed some APIs specially targeted at low resource embedded systems, so we rely on the MCAPI, MTAPI and MRAPI for low memory footprint light-weight communications, threading and synchronization.

There is a free implementation based on POSIX so we can have a working version on any decent operating system for testing the SME-C programs. Since this is a reference implementation based on Linux *pthread*s, it can also be used as an example to port to the various available hardware.

During the SMECY project, ST & CEA has phased out NPL, their own library to program the STHORM, because it was equivalent to the MCA standard API. So for a message passing interface API, the MCA APIs have to be used on this platform.

Right now there are 3 different specifications we can use in the project.

7.3 Multicore Communication API (MCAPI)

This is the main and first produced API by the MCA⁵ dealing with communications between processes.

This standard define basically 3 kinds of communication channels:

- messages: connection-less datagrams, similar to UDP datagrams in IP networking;
- packet channels: connection-oriented, unidirectional, FIFO packet streams;
- scalar channel: connection-oriented, single-word, unidirectional, FIFO scalar streams.

A communication entity in MCAPI is a node, that can be gathered in some domains to add hierarchy, and can represent a process on a processor or a hardware accelerator for example.

Please refer to the “Multicore Communications API (MCAPI) Specification V2.015” document for more information.

7.3.1 MTAPI

To complete

7.3.2 MRAPI

To complete

⁵Note that this MCAPI should not be confused with the MCA APIs, that are the APIs in general designed by the MultiCore Association, even if the WWW site and the documents of this association are not always very clear... This was confusing during the København meeting.

7.4 NPL API

NPL is the API defined to program ST P2012/STHORM in a native way. Since it is rather at the same level of MCAPI/MRAPI/MTAPI, it should be easy to implement one above the other. Since the MCA APIs are standard, we think that it is commercially interesting to provide a MCA APIs over NPL or other to widen P2012/STHORM usage.

Indeed, in 10/2011 this API seems to have been phased out by ST & CEA. So this part is kept for its history interest and MCAPI is used in the SMECY project as a lower intermediate representation, the IR-2.

7.5 EdkDSP/ASVP API

7.6 OpenCL

Since ST P2012/STHORM can be programmed in OpenCL which is also a programming API, a C process can use OpenCL orthogonally with other API. A kernel launching is done by defining the kernel source, the memory zone to use and to transfer and the different parameters of the kernel.

Refer to OpenCL documentation for more information.

Define here what is useful in the project.

8 Some design patterns for STHORM SME-C

In this section we provide some typical use-cases for the ST STHORM platform to illustrate how SME-C can be used to program some common application patterns.

8.1 Same computation on all the PEs

```
#define N 1000
2 #define NB_CLUSTERS 4
#define NB_PES 16
#define MIN(x, y) ((x) < (y) ? x : y)

#include <stdio.h>
7
/* Initialize an array between 2 given lines */
void init_array(int a[N][N], int begin, int end) {
    for (int i = begin; i < end; i++)
        for (int j = 0; j < N; j++)
12     a[i][j] = 2*i + 3*j;
}

int main() {
17     int a[N][N];

    int slice = N/(NB_CLUSTERS*NB_PES);
    /* Launch enough OpenMP thread to control all the fabric:

        Assume that the runtime allows enough threads with nested
22     parallelism */
    #pragma omp parallel for num_threads(NB_CLUSTERS)
        for (int cluster = 0; cluster < NB_CLUSTERS; cluster++)
    #pragma omp parallel for num_threads(NB_PES)
        for (int pe = 0; pe < NB_PES; pe++) {
27     /* So now the iterations should be distributed with 1
        iteration/thread, on NB_CLUSTERS*NB_PES threads.
```

```

        Distribute the initialization on all the fabric: */
        int global_pe = cluster*NBPES + pe;
32     int begin = slice*global_pe;
        int end = MIN(N, slice*(global_pe + 1));
#pragma smecy map(STHORM, cluster, pe)
#pragma smecy arg(a, out, /[[begin:end-1][[]])
        init_array(a, begin, end);
37     }

        printf("a[27][42] = %d\n", a[27][42]);

        return 0;
42 }

```

8.2 Same computation on all the PEs of a cluster

```

#define N 1000
#define NB_CLUSTERS 4
3 #define NBPES 16
#define MIN(x, y) ((x) < (y) ? x : y)

#include <stdio.h>

8 /* Initialize an array between 2 given lines */
void init_array(int a[N][N]) {
#pragma omp parallel for
    for (int i = 0; i < N; i++)
        for (int j = 0; j < N; j++)
13     a[i][j] = 2*i + 3*j;
}

void mult(int a[N][N], int begin, int end, int fact) {
    for (int i = begin; i < end; i++)
18     for (int j = 0; j < N; j++)
        a[i][j] *= fact;
}

// Compute a distribution by stripe on the PEs:
23 #define ITER_BEGIN N/(NB_CLUSTERS*NBPES)*(cluster*NBPES + pe)
#define ITER_END MIN(N, N/(NB_CLUSTERS*NBPES)*(cluster*NBPES + pe + 1))

int main() {
    int a[N][N];
28

    // Initialize in parallel on the Cortex A9:
    init_array(a);

    /* Then do some computations with the fabric
33

        Launch enough OpenMP thread on the Cortex A9 to control all the
        clusters */
#pragma omp parallel for num_threads(NB_CLUSTERS)
    for (int cluster = 0; cluster < NB_CLUSTERS; cluster++)
38     switch(cluster) {
        case 0:
            // The cluster 0:
#pragma omp parallel for num_threads(NBPES)
                for (int pe = 0; pe < NBPES; pe++)

```

```

43 #pragma smecy map(SIHORM, cluster , pe)
    #pragma smecy arg(a, inout , / [begin:end - 1] [])
        mult(a, ITER_BEGIN, ITER_END, -1);
        break;

48     case 1:
        // The cluster 1:
    #pragma omp parallel for num_threads(NB_PES)
        for (int pe = 0; pe < NB_PES; pe++)
    #pragma smecy map(SIHORM, cluster , pe)
53 #pragma smecy arg(a, inout , / [begin:end - 1] [])
        mult(a, ITER_BEGIN, ITER_END, 3);
        break;

        case 2:
58     // The cluster 2:
    #pragma omp parallel for num_threads(NB_PES)
        for (int pe = 0; pe < NB_PES; pe++)
    #pragma smecy map(SIHORM, cluster , pe)
    #pragma smecy arg(a, inout , / [begin:end - 1] [])
63 mult(a, ITER_BEGIN, ITER_END, -7);
        break;

        case 3:
        // The cluster 3:
68 #pragma omp parallel for num_threads(NB_PES)
        for (int pe = 0; pe < NB_PES; pe++)
    #pragma smecy map(SIHORM, cluster , pe)
    #pragma smecy arg(a, inout , / [begin:end - 1] [])
        mult(a, ITER_BEGIN, ITER_END, 9);
73     break;
    }

    printf("a[27][42] = %d\n", a[27][42]);

78     return 0;
}

```

8.3 Stream computation through the PEs of a cluster

```

1 #define NB_CLUSTERS 4
  #define NB_PES 16

  #include <stdio.h>

6 /* Compute a polynomial with HÅrner method in a pipelined way as a stream

    ((ax + b)x + c)x + d)x + e
    */

11 /* A stage of a systolic HÅrner polynomial computation */
void stage(int *poly, int *x, int *input,
           int *pass_through, int *output) {
    // Compute one HÅrner factor y = ax + b
16 *output = *input * *x + *poly;
    // Propagate x down the stream:
    *pass_through = *x;
}

```

```

}

21 void output(int *v) {
    // In real code there is no such IO available on STHORM...
    printf("Polynomial = %d\n", *v);
}

26

int main() {
    /* Values where to compute the polynomial */
    int x[] = { 0, 1, -1, 2, 4, 5, 10 };
31 /* The polynomial is  $-4x^3 + x^2 + 2x + 3$  : */
    int poly[] = {0, -4, 1, 2, 3};
    int v[] = {0, 0, 0, 0, 0, 0};
    /* To propagate the x unchanged through the pipeline: */
    int path_through[] = {0, 0, 0, 0, 0};

36 #pragma smecy stream_loop
    for (int i = 0; i < sizeof(x)/sizeof(x[0]); i++) {
        /* The following could be done with macros */
41 #pragma smecy map(STHORM, 0, 0) arg(1, in) arg(2, in) arg(3, in) \
        arg(4, out) arg(5, out)
        stage(&poly[0], &x[i], &v[0], &path_through[0], &v[1]);
    #pragma smecy stage map(STHORM, 0, 1) arg(1, in) arg(2, in) arg(3, in) \
        arg(4, out) arg(5, out)
        stage(&poly[1], &path_through[0], &v[1], &path_through[1], &v[2]);
46 #pragma smecy stage map(STHORM, 0, 2) arg(1, in) arg(2, in) arg(3, in) \
        arg(4, out) arg(5, out)
        stage(&poly[2], &path_through[1], &v[2], &path_through[2], &v[3]);
    #pragma smecy stage map(STHORM, 0, 3) arg(1, in) arg(2, in) arg(3, in) \
        arg(4, out) arg(5, out)
51 stage(&poly[3], &path_through[2], &v[3], &path_through[3], &v[4]);
    #pragma smecy stage map(STHORM, 0, 4) arg(1, in) arg(2, in) arg(3, in) \
        arg(4, out) arg(5, out)
        stage(&poly[4], &path_through[3], &v[4], &path_through[4], &v[5]);
    #pragma smecy stage map(STHORM, 0, 5) arg(1, in) arg(2, in) arg(3, in) \
56 arg(4, out) arg(5, out)
        output(&v[5]);
    }
    return 0;
}
}

```

8.4 Systolic computation with the PEs of a cluster

```

#define N 1000
#define NB_CLUSTERS 4
#define MIN(x, y) ((x) < (y) ? x : y)

5 #include <stdio.h>

/* Compute a polynomial with Horner method in a pipelined way with a
   systolic line of PE

10 (((ax + b)x + c)x + d)x + e

This is a spatial developping of the streaming example in
stream_cluster_PEs.c. So read stream_cluster_PEs.c first for a gentle
introduction.

```

```

15  */

    /* Values where to compute the polynomial */
    int x[] = { 0, 1, -1, 2, 4, 5, 10 };
20  /* The polynomial is  $-4x^3 + x^2 + 2x + 3$  : */
    int poly[] = {0, -4, 1, 2, 3};

    /* A stage of a systolic Horner polynomial computation */
    void stage(int *poly, int *x, int *input,
25         int *pass_through, int *output) {
        // Compute one Horner factor  $y = ax + b$ 
        *output = *input * *x + *poly;
        // Propagate x down the stream:
        *pass_through = *x;
30 }

    void output(int *v, int t, int pe) {
        // In real code there is no such IO available on STHORM...
35     printf("Polynomial at t=%d with pe=%d = %d\n", t, pe, *v);
    }

    int main() {
40     /* Since the 2 following arrays are read and written at the same time to
        move data through the pipeline, increase the dimension to store
        current t and next t+1 values. Put the time dimension first to limit
        cache line false sharing: */
        int v[2][sizeof(poly)/sizeof(poly[0]) + 1];
45     /* To propagate the x unchanged through the pipeline: */
        int path_through[2][sizeof(poly)/sizeof(poly[0])];

        // Launch the thread only once
        #pragma omp parallel num_threads(sizeof(poly)/sizeof(poly[0]) + 1)
50     /* The time loop.
        It takes len(x)+len(poly) cycles to go through the pipeline */
        for (int t = 0;
            t < sizeof(x)/sizeof(x[0]) + sizeof(poly)/sizeof(poly[0]);
            t++) {
55     /* Toggle between 0 and 1 in phase oposition: */
        int current = t & 1;
        int next = ~t & 1;

        /* Spatially spread the computation on the PEs, +1 is for for a PE doing the
60     output: */
        #pragma omp for schedule(static, 1)
            for (int pe = 0; pe < sizeof(poly)/sizeof(poly[0]) + 1; pe++) {
                if (pe == 0) {
                    // Special case for the first stage which taps in the input
65                    if (t < sizeof(x)/sizeof(x[0])) {
                        v[current][0] = 0;
                        // Only run when there are some data to read
                    #pragma smecy map(STHORM, 0, 0) arg(1, in) arg(2, in) arg(3, in) \
70                        arg(4, out) arg(5, out)
                        stage(&poly[0], &x[t], &v[current][0],
                            &path_through[next][0], &v[next][1]);
                }
            }
    }

```

```

    }
    else if (pe == sizeof(poly)/sizeof(poly[0])) {
75     /* The PE doing the output is the last stage: it runs only after
        the data got time to pass through the whole pipeline */
        if (t >= pe)
            /* The output is done on the host, so no #pragma map here
            output(&v[current][5], t, pe);
80     }
        else {
            /* The normal computation pipeline stage, with the general schedule
            if (t >= pe && t < pe + sizeof(x)/sizeof(x[0])) {
85     #pragma smecy stage map(STHORM, 0, pe) arg(1, in) arg(2, in) arg(3, in) \
                arg(4, out) arg(5, out)
                stage(&poly[pe], &path_through[current][pe - 1], &v[current][pe],
                    &path_through[next][pe], &v[next][pe + 1]);
            }
        }
90     /* There is an implicit barrier here
    }
}
return 0;
}

```

8.5 Round robin computation on the clusters

```

1  #define N 10
   #define NB.CLUSTERS 4
   #define NB.PES 16
   #define MIN(x, y) ((x) < (y) ? x : y)

6  #include <stdio.h>

   int get_ticket() {
       static int t = 0;
11  int ticket;
       /* In case this is called from several threads. Avoid a flush, anyway
   #pragma omp atomic capture
       ticket = t++;
       return ticket;
16 }

   int get_data(t) {
       //sleep(t&1);
21  /* In a real application, get a radar signal time slice for example */
       return t*2;
   }

26 int compute(int d, int cluster, int pe) {
       /* In a real application, do a computation on the data */
       return d*cluster + pe;
   }

31

   int main() {
       /* Launch all the threads to control the clusters only once */

```

```

36  #pragma omp parallel num_threads(NB.CLUSTERS)
    {
        for (int i = 0; i < N; i++) {
            /* Execute 1 iteration per thread and there will be some ordered
               statement.

41             It is useless to wait at the end of the iterations, but it looks
               like a nowait here break the ordered. Compiler bug? */
            #pragma omp for schedule(static, 1) ordered
                for (int cluster = 0; cluster < NB.CLUSTERS; cluster++) {
46                 /* Get an ID in the order of the sequential iteration. Remove the
                   ordered if it is not needed. */
                    int t;
                    #pragma omp ordered
                        {
                            t = get_ticket();
51                        }
                            int d = get_data(t);
                            int r;
                            #pragma omp parallel for num_threads(NB.PES) reduction(+:r)
                                for (int pe = 0; pe < NB.PES; pe++)
56 #pragma smecy map(SIHORM, cluster, pe)
                                    r += compute(d, cluster, pe);
                                    /* Produce the result in order. Remove the ordered if it is not
                                       needed. */
                                #pragma omp ordered
36                                    {
61                                        printf("Cluster %d produced %d for ticket %d\n", cluster, r, t);
                                        }
                                    }
                                }
                            }
66    }
    }

```

9 Conclusion

Keep it
simple