

CUBE 4.3.4 – Cube Tools Developer Guide

Commented example of a CUBE tool

April 3, 2016

The Scalasca Development Team
scalasca@fz-juelich.de

Copyright

Copyright © 1998–2015 Forschungszentrum Jülich GmbH, Germany

Copyright © 2009–2015 German Research School for Simulation Sciences GmbH, Jülich/Aachen, Germany

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the names of Forschungszentrum Jülich GmbH or German Research School for Simulation Sciences GmbH, Jülich/Aachen, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Contents

Copyright	iii
1 Makefile for provided examples	1
1.1 Quick info about makefile.	1
1.2 Commented source	1
2 Examples of using C++ library	3
2.1 Commented source	3
3 Examples of using CUBE c-writer library	11
3.1 Commented source	11

1 Makefile for provided examples

1.1 Quick info about makefile.

Here we provide a small example of a makefile, which is used to compile and build examples delivered with CUBE. Similar makefiles can be used by developers to compile and build own CUBE tools.

1.2 Commented source

First we specify the installation path of CUBE and its "cube-config" script. This script delivers correct flags for compiling and linking, paths to the CUBE tools and GUI. (besides of another useful technical information)

```
CUBE\_DIR = /path/CubeInstall  
CUBE\_CONFIG = $(CUBE\_DIR)/bin/cube-config
```

Additionally we specify CPPFLAGS and LDFLAGS to compile and link examples.

```
CPPFLAGS = $(shell $(CUBE\_CONFIG) --cube-cxxflags)  
CFLAGS = $(shell $(CUBE\_CONFIG) --cubew-cxxflags)  
CLDFLAGS = $(shell $(CUBE\_CONFIG) --cubew-ldflags)  
CPPLDFLAGS = $(shell $(CUBE\_CONFIG) --cube-ldflags)
```

Here a compiler is selected to compile and build the example.

```
\textcolor{preprocessor}{# GNU COMPILER }  
CXX = g++  
CC = gcc -std=c99  
MPICXX= mpicxx
```

We define explicit suffixes for an executable file, created from C source, from c++ source and an MPI executable. If one develops a tool, which is using MPI, it is useful (sometimes) to define a special suffix for automatic compilation.

```
.SUFFIXES: .c .o .cpp .c.exe .cpp.exe .c.o .cpp.o .mpi.o .mpi.cpp  
.PHONY: all clean
```

Object files of examples and their targets

Chapter 1. Makefile for provided examples

```
\textcolor{preprocessor}{# Object files}
OBJS = cube\_example.cpp.o \(\backslashbackslash\)
      cubew\_example.c.o

TARGET = cube\_example.cpp.exe \(\backslashbackslash\)
        cubew\_example.c.exe
```

Automatic rule for the compilation of every single C++ source into .o file and for building targets.

```
%.cpp.o : %.cpp
    $(CXX) -c $< -o $@ $(CPPFLAGS)

%.cpp.exe : %.cpp.o
    $(CXX) $< -o $@ $(CPPLDFLAGS)
```

Automatic rule for the compilation of every single C++ with MPI source into .o file and for building targets.

```
%.mpi.o : %.mpi.cpp
    $(MPICXX) -c $< -o $@ $(CPPFLAGS) $(CFLAGS)

%.mpi.exe : %.mpi.o
    $(MPICXX) $< -o $@ $(CLDFLAGS)
```

Automatic rule for the compilation of every single C source into .o file and for building targets.

```
%.c.o : %.c
    $(CC) -c $< -o $@ $(CFLAGS)

%.c.exe : %.c.o
    $(CC) $< -o $@ $(CLDFLAGS)
```

```
\textcolor{preprocessor}{#-----}
\textcolor{preprocessor}{# Rules}
\textcolor{preprocessor}{#-----}
```

```
all: $(TARGET)
```


2 Examples of using C++ library

Present example shows in several short steps the main idea of creating a cube file using C++ library and obtaining different values from this cube file.

2.1 Commented source

Include standard c++ header

```
....  
\textcolor{preprocessor}{\#include <iostream>}  
\textcolor{preprocessor}{\#include <string>}  
\textcolor{preprocessor}{\#include <vector>}  
\textcolor{preprocessor}{\#include <fstream>}
```

Include CUBE headers. Notice, that all C++ headers of Cube framework have a prefix CubeXXX.h.

```
\textcolor{preprocessor}{\#include "Cube.h"}  
\textcolor{preprocessor}{\#include "CubeMetric.h"}  
\textcolor{preprocessor}{\#include "CubeCnode.h"}  
\textcolor{preprocessor}{\#include "CubeProcess.h"}  
\textcolor{preprocessor}{\#include "CubeNode.h"}  
\textcolor{preprocessor}{\#include "CubeThread.h"}  
\textcolor{preprocessor}{\#include "CubeCartesian.h"}
```

Define namespaces, where we do operate, standard templates library and cube namespace.

```
\textcolor{keyword}{using namespace }std;  
\textcolor{keyword}{using namespace }cube;
```

Start main and declare needed variables. We want to create three metrics and some regions.

```
\textcolor{keywordtype}{int} main(\textcolor{keywordtype}{int} argc, \textcolor{keywordtype}{char}* a  
    Metric *met0, *met1, *met2;  
    Region *regn0, *regn1, *regn2;  
    Cnode *cnode0, *cnode1, *cnode2;  
    Machine* mach;  
    Node* node;  
    Process* proc0, *proc1;  
    Thread *thrd0, *thrd1;
```

Enclose everything in a *try-catch* block, because the cube library throws exceptions, if something goes wrong.

```
\textcolor{keywordflow}{try}
\{
```

Create cube object. This means we create a new empty Cube.

```
Cube cube;
```

Specify general properties of the cube object.

```
\textcolor{comment}{// Specify mirrors (optional)}
cube.def\_mirror(\textcolor{stringliteral}{"http://icl.cs.utk.edu/software/kojak/"});
cube.def\_mirror(\textcolor{stringliteral}{"http://www.fz-juelich.de/jsc/kojak/"});
```

```
\textcolor{comment}{// Specify information related to the file (optional)}
cube.def\_attr(\textcolor{stringliteral}{"experiment time"}, \textcolor{stringliteral}{"November 1st 2000"});
cube.def\_attr(\textcolor{stringliteral}{"description"}, \textcolor{stringliteral}{"a simple example"});
```

```
cube.set\_statistic\_name(\textcolor{stringliteral}{"statisticstat"});
cube.set\_metrics\_title(\textcolor{stringliteral}{"Metrics"});
cube.set\_calltree\_title(\textcolor{stringliteral}{"Calltree"});
cube.set\_systemtree\_title(\textcolor{stringliteral}{"Systemtree"});
```

Now we start to define dimensions of the cube.

First we define the *metric* dimension. Notice, that metrics build a tree and parents have to be declared before their children.

Every metric can be of one of two kinds: *inclusive* or *exclusive*.

If one omits `CUBE_INCLUSIVE` or `CUBE_EXCLUSIVE`, cube automatically creates an *exclusive* metric. One cannot change this type later.

Every metric needs a display name, an unique name, type of values (`INTEGER`, `DOUBLE`, `MAXDOUBLE`, `MINDOUBLE`, others), units of measurement, value (usually empty string), URL, where one can find the documentation about this metric, description and its parent in the metric tree.

The cube returns a pointer on struct *cube_metric*, which has to be used for saving or reading values from the cube.

```
\textcolor{comment}{// Build metric tree }
met0 = cube.def\_met(\textcolor{stringliteral}{"Time"}, \textcolor{stringliteral}{"Uniq\_name1"}, \
    \textcolor{stringliteral}{"@mirror@patterns-2.1.html#execution"},
    \textcolor{stringliteral}{"root node"}, NULL, CUBE\_INCLUSIVE); \textcolor{comment}{// }
met1 = cube.def\_met(\textcolor{stringliteral}{"User time"}, \textcolor{stringliteral}{"Uniq\_name2"},
    \textcolor{stringliteral}{"http://www.cs.utk.edu/usr.html"},
```

```
        \textcolor{stringliteral}{2nd level"}, met0); \textcolor{comment}{// without usi
met2 = cube.def\_met(\textcolor{stringliteral}{System time"}, \textcolor{stringliteral}{Uniq\_name3
        \textcolor{stringliteral}{http://www.cs.utk.edu/sys.html"},
        \textcolor{stringliteral}{2nd level"}, met0); \textcolor{comment}{// without usi
```

Then we define another dimension, the "call tree" dimension. This dimension gets defined in a two-step way:

1. One defines a list of regions in the instrumented source code;
2. One builds a call tree with the regions defined above;

Then we define the `call tree` dimension. This dimension gets defined in a two-step way:

1. One defines a list of regions in the instrumented source code;
2. One builds a call tree with regions defined in the previous step.

First one defines regions.

Every region has a name, start and end line, URL with the documentation of the region, description and source file (module). Regions build a list, therefore no "parent-child" relation is given.

The cube returns a pointer on Object *Region*, which can be used later for the calculations, visualization or access to the data.

```
\textcolor{comment}{// Build call tree}
\textcolor{keywordtype}{string} mod    = \textcolor{stringliteral}{"/ICL/CUBE/example.c"};
regn0  = cube.def\_region(\textcolor{stringliteral}{main"}, 21, 100, \textcolor{stringliteral}{""},
regn1  = cube.def\_region(\textcolor{stringliteral}{foo"}, 1, 10, \textcolor{stringliteral}{""}, \te
regn2  = cube.def\_region(\textcolor{stringliteral}{bar"}, 11, 20, \textcolor{stringliteral}{""}, \t
```

Then one defines an actual dimension, the `call tree` dimension.

The call tree consists of so called CNODEs. Cnode stands for "call path".

Every cnode gets as a parameter a region, source file (module), its id and parent cnode (caller).

Parent cnodes have to be defined before their children. Region might be entered from different places in the program, therefore different cnodes might have same region as a parameter.

```
cnode0 = cube.def\_cnode(regn0, mod, 21, NULL);
cnode1 = cube.def\_cnode(regn1, mod, 60, cnode0);
cnode2 = cube.def\_cnode(regn2, mod, 80, cnode0);
```

The last dimension is the `system tree` dimension. Currently CUBE defines the system dimension with the fixed hierarchy: MACHINE → NODES → PROCESSES → THREADS

It leads to the fixed sequence of calls in the system dimension definition:

1. First one creates a root for the system dimension : *Machine*. *Machine* has a name and description.
2. *Machine* consists of *Nodes*. Every *Node* has a name and a *Machine* as a parent.
3. On every *Node* run several *cube_processes* (as many cores are available). *Process* has a name, MPI rank and *Node* as a parent.
4. Every *Process* spawns several (one or more) *Threads* (OMP, Pthreads, Java Threads). *Thread* has a name, its rank and *Process* as a parent.

The cube returns a pointer on *CubeMachine*, *Node*, *Process* or *Thread*, which has to be used later to define further level in the system tree or to access the data in the cube.

```
\textcolor{comment}{// Build location tree}
mach = cube.def\_mach(\textcolor{stringliteral}{"MSC"}, \textcolor{stringliteral}{""});
node = cube.def\_node(\textcolor{stringliteral}{"Athena"}, mach);
proc0 = cube.def\_proc(\textcolor{stringliteral}{"Process 0"}, 0, node);
proc1 = cube.def\_proc(\textcolor{stringliteral}{"Process 1"}, 1, node);
thrd0 = cube.def\_thrd(\textcolor{stringliteral}{"Thread 0"}, 0, proc0);
thrd1 = cube.def\_thrd(\textcolor{stringliteral}{"Thread 1"}, 1, proc1);
```

After the dimensions are defined, one fills the cube object with the data. Every data value is specified by three "coordinates": (*Metric*, *Cnode*, *Thread*)

Note, that *Machine*, *Node* and *Process* are not a "coordinate". They are used only to build up the physical construction of the machine.

C++ cube library allows random access to the data, therefore no specific restrictions about sequence of calls are made.

```
cube.set\_sev( met0, cnode0, thrd0, 12. );
cube.set\_sev( met0, cnode0, thrd1, 11. );
cube.set\_sev( met0, cnode1, thrd0, 5. );
cube.set\_sev( met0, cnode1, thrd1, 6. );
cube.set\_sev( met0, cnode2, thrd0, 4.2 );
cube.set\_sev( met0, cnode2, thrd1, 3.5 );

cube.set\_sev( met1, cnode0, thrd0, 4. );
cube.set\_sev( met1, cnode0, thrd1, 4. );
cube.set\_sev( met1, cnode1, thrd0, 3. );
cube.set\_sev( met1, cnode1, thrd1, 3.2 );
cube.set\_sev( met1, cnode2, thrd0, 0.2 );
cube.set\_sev( met1, cnode2, thrd1, 0.5 );

cube.set\_sev( met2, cnode0, thrd0, 2. );
cube.set\_sev( met2, cnode0, thrd1, 2.4 );
cube.set\_sev( met2, cnode1, thrd0, 1.2 );
cube.set\_sev( met2, cnode1, thrd1, 1.2 );
cube.set\_sev( met2, cnode2, thrd0, 0.01 );
cube.set\_sev( met2, cnode2, thrd1, 0.03 );

cube.set\_sev( met3, cnode0, thrd0, 10 );
```

```

cube.set\_sev( met3, cnode0, thrd1, 20 );
cube.set\_sev( met3, cnode1, thrd0, 800 );
cube.set\_sev( met3, cnode1, thrd1, 700 );
cube.set\_sev( met3, cnode2, thrd0, 9300 );
cube.set\_sev( met3, cnode2, thrd1, 9500 );

```

There are different calls to get a value from cube:

1. `\textcolor{keywordtype}{double} value1 = cube.get_sev(met0, cnode0, thrd0);`

Get a double representation of a single value for a given metric, call path and thread. This call is backward compatible to cube3 and delivers an exclusive value along the call tree.

2. `\textcolor{keywordtype}{double} value2 =
cube.get_sev(met0, cube::CUBE_CALCULATE_INCLUSIVE,
cnode0, cube::CUBE_CALCULATE_INCLUSIVE,
thrd0 , cube::CUBE_CALCULATE_INCLUSIVE);`

Get a double representation of a single value for a given metric, call path and thread. Here one specifies kind of value along every three. Notice, that this time it delivers inclusive value along the call tree (Metric values are naturally inclusive and threads are leaves in the current model of the system tree)

3. `\textcolor{keywordtype}{double} value3 =
cube.get_sev(met0, cube::CUBE_CALCULATE_INCLUSIVE,
cnode0, cube::CUBE_CALCULATE_INCLUSIVE,
node , cube::CUBE_CALCULATE_INCLUSIVE);`

Get a double representation of a single value for a given metric, call path and whole node. It means, aggregation over all threads of all processes of this node is made.

4. `\textcolor{keywordtype}{double} value4 =
cube.get_sev(met0, cube::CUBE_CALCULATE_INCLUSIVE,
cnode0, cube::CUBE_CALCULATE_INCLUSIVE
);`

Get a double representation of a single value for a given metric, call path, aggregated over whole system tree using algebra of the metric value.

5. `\textcolor{keywordtype}{double} value5 =
cube.get_sev(met0, cube::CUBE_CALCULATE_EXCLUSIVE
);`

Get a double representation of a single value for a given metric (exclusive along metric tree), aggregated over the whole call tree and system tree using algebra of the metric value.

There are several another additional calls of the same type. Please see documentation (Cube library source).

CUBE can carry a set of so called "topologies": mappings $\text{THREAD} \rightarrow (x, y, z, \dots)$

Then GUI visualize every value (*Metric, Cnode, Thread*) for the selected metric and cnode as a 1D, 2D or 3D set of points with the different colors.

First one specifies a number of dimensions (any number is supported, currently shown are only the first three), a vector with the sizes in every dimension and its periodicity and creates an object of class *Cartesian*

```
\textcolor{comment}{// building a topology}
\textcolor{comment}{// create 1st cartesian.}
\textcolor{keywordtype}{int} ndims = 2;
vector<long> dimv;
vector<bool> periodv;
\textcolor{keywordflow}{for} ( \textcolor{keywordtype}{int} i = 0; i < ndims; i++) \{
    dimv.push\_back(5);
    \textcolor{keywordflow}{if} ( i % 2 == 0)
        periodv.push\_back( \textcolor{keyword}{true} );
    \textcolor{keywordflow}{else}
        periodv.push\_back( \textcolor{keyword}{false} );
\}

Cartesian* cart = cube.def\_cart(ndims, dimv, periodv);
```

The coordinates one defines like a vector and creates a mapping.

```
\textcolor{keywordflow}{if} (cart != NULL)
\{
    vector<long> p[2];
    p[0].push\_back(0);
    p[0].push\_back(0);
    p[1].push\_back(2);
    p[1].push\_back(2);
    cube.def\_coords(cart, thrd1, p[0]);
\}
```

One can define names for every dimension.

```
vector<string> dim\_names;
dim\_names.push\_back( \textcolor{stringliteral}{ "X" } );
dim\_names.push\_back( \textcolor{stringliteral}{ "Y" } );
cart->set\_namedims( dim\_names );
```

In the same way one can create any number of topologies. They are shown in the GUI.

```
ndims = 2;
vector<long> dimv2;
vector<bool> periodv2;
\textcolor{keywordflow}{for} ( \textcolor{keywordtype}{int} i = 0; i < ndims; i++ )
\{
    dimv2.push\_back( 3 );
    \textcolor{keywordflow}{if} ( i % 2 == 0 )
    \{
        periodv2.push\_back( \textcolor{keyword}{true} );
    \}
```

```
\textcolor{keywordflow}{else}
\{
  periodv2.push\_back( \textcolor{keyword}{false} );
\}
\}

Cartesian* cart2 = cube.def\_cart( ndims, dimv2, periodv2 );
cart2->set\_name( \textcolor{stringliteral}{"Application topology 2"} );
\textcolor{keywordflow}{if} ( cart2 != NULL )
\{
  vector<long> p2[ 2 ];
  p2[ 0 ].push\_back( 0 );
  p2[ 0 ].push\_back( 1 );
  p2[ 1 ].push\_back( 1 );
  p2[ 1 ].push\_back( 0 );
  cube.def\_coords( cart2, thrd0, p2[ 0 ] );
  cube.def\_coords( cart2, thrd1, p2[ 1 ] );
\}
cart2->set\_dim\_name( 0, \textcolor{stringliteral}{"Dimension 1"} );
cart2->set\_dim\_name( 1, \textcolor{stringliteral}{"Dimension 2"} );
```

To save cube file on disk, one calls the following method. Extension ".cubex" will be added automatically. There is a corresponding method `openCubeReport (...)`;

```
cube.writeCubeReport( \textcolor{stringliteral}{"cube-example"} );
```

There are some routines for the conversion from cube3 to cube4 called.

```
\textcolor{comment}{// Output to a cube file}
ofstream out;
out.open(\textcolor{stringliteral}{"example.cube"});
out << cube;

\textcolor{comment}{// Read it (example.cube) in and write it to another file (example2.cube)}
ifstream in(\textcolor{stringliteral}{"example.cube"});
Cube cube2;

in >> cube2;
ofstream out2;
out2.open(\textcolor{stringliteral}{"example2.cube"});

out2 << cube2;
```

Here we catch all exceptions, thrown by the cube. We print the exception message and finish the program.

```
\} \textcolor{keywordflow}{catch} (RuntimeError e)
```

```
\{  
    cout << \textcolor{stringliteral}{"Error: "} << e.get\_msg() << endl;  
}  
  
\}
```


3 Examples of using CUBE c-writer library

Present example shows in several short steps the main idea of creating a cube file using C writer library.

In this example we do not show the optimization, which is needed to prevent unnecessary seeks while writing.

3.1 Commented source

Include standard c header

```
\textcolor{preprocessor}{\#include <stdio.h>}
\textcolor{preprocessor}{\#include <stdint.h>}
\textcolor{preprocessor}{\#include <stdlib.h>}
\textcolor{preprocessor}{\#include <unistd.h>}
\textcolor{preprocessor}{\#include <errno.h>}
\textcolor{preprocessor}{\#include <string.h>}
```

Include CUBE headers.

Notice, that CUBE4 c-writer headers got an prefix `cubew_XXX.h`

```
\textcolor{preprocessor}{\#include "cubew\_cube.h"}
```

Start main and define a name of the cube file. Extension ".cubex" will be append automatically.

```
\textcolor{keywordtype}{int} main(\textcolor{keywordtype}{int} argc, \textcolor{keywordtype}{char}* a
\{
```

```
    \textcolor{keywordtype}{char} cubefile[12] = \textcolor{stringliteral}{"simple-cube"};
```

Create the structure of the cube. CUBE_MASTER defines, that in parallel MPI environment this cube (usually rank 0) writes all parts of the cube (anchor, indexes and data). The last argument is ignored in current version of CUBE c-writer.

```
cube_t* cube=cube_create(\textcolor{stringliteral}{"example"}, CUBE_MASTER, CUBE_FALSE);
\textcolor{keywordflow}{if} (!cube) \{
    fprintf(stderr,\textcolor{stringliteral}{"cube_create failed!\(\\backslash)n"});
    exit(1);
\}
```

Specify general properties of cube object.

```
cube_def_mirror(cube, \textcolor{stringliteral}{"http://icl.cs.utk.edu/software/kojak/"});
cube_def_mirror(cube, \textcolor{stringliteral}{"http://www.fz-juelich.de/jsc/kojak/"});
cube_def_attr(cube, \textcolor{stringliteral}{"description"}, \textcolor{stringliteral}{"a simple
cube_def_attr(cube, \textcolor{stringliteral}{"experiment time"}, \textcolor{stringliteral}{"Nove

cube_set_statistic_name(cube, \textcolor{stringliteral}{"mystat"});
```

Now we start to define the dimensions of the cube.

First we define `metric` dimension. Notice, that metrics build a tree and parents have to be declared before their children.

Every metric can be of two kinds: *inclusive* or *exclusive*.

Every metric needs a display name, an unique name, type of values (INTEGER, DOUBLE, MAXDOUBLE, MINDOUBLE, others), units of measurement, value (usually empty string), URL, where one can find the documentation about this metric, description and its parent in the metric tree.

The cube returns a pointer on structure `cube_metric`, which has to be used for saving or reading values from the cube.

```
cube_metric *met0, *met1, *met2;
met0 = cube_def_met(cube, \textcolor{stringliteral}{"Time"}, \textcolor{stringliteral}{"Uniq_name"},
    \textcolor{stringliteral}{"@mirror@patterns-2.1.html#execution"},
    \textcolor{stringliteral}{"root node"}, NULL, CUBE_METRIC_EXCLUSIVE);
met1 = cube_def_met(cube, \textcolor{stringliteral}{"User time"}, \textcolor{stringliteral}{"Uniq_name"},
    \textcolor{stringliteral}{"http://www.cs.utk.edu/usr.html"},
    \textcolor{stringliteral}{"2nd level"}, met0, CUBE_METRIC_INCLUSIVE);
met2 = cube_def_met(cube, \textcolor{stringliteral}{"System time"}, \textcolor{stringliteral}{"Uniq_name"},
    \textcolor{stringliteral}{"http://www.cs.utk.edu/sys.html"},
    \textcolor{stringliteral}{"2nd level"}, met0, CUBE_METRIC_EXCLUSIVE);
```

Then we define the `calltree` dimension. This dimension gets defined in a two-step way:

1. One defines a list of regions in the instrumented source code;
2. One builds a call tree with the regions defined in the previous step.

First one defines the regions.

Every region has a name, start and end line, URL with the documentation of the region, description and source file (module). Regions build a list, therefore no "parent-child" relation is given.

The cube returns a pointer on structure *cube_region*, which can be used later for the calculations, visualization or access to the data.

```
\textcolor{keywordtype}{char}* mod = \textcolor{stringliteral}{"/ICL/CUBE/example.c"};
cube\_region *regn0, *regn1, *regn2;
regn0 = cube\_def\_region(cube, \textcolor{stringliteral}{"main"}, 21, 100, \textcolor{stringliteral}
regn1 = cube\_def\_region(cube, \textcolor{stringliteral}{"<<init>>foo"}, 1, 10, \textcolor{stringlit
regn2 = cube\_def\_region(cube, \textcolor{stringliteral}{"<<loop>>bar"}, 11, 20, \textcolor{stringlit
```

Then one defines an actual dimension, the call tree dimension.

Call tree consists of so called CNODEs. Cnode stands for "call path".

Every cnode gets as a parameter a region, source file (module), its id and parent cnode (caller).

Parent cnodes have to be defined before their children. Region might be entered from different places in the program, therefore different cnodes might have same region as a parameter.

```
cube\_cnode *cnode0, *cnode1, *cnode2;
cnode0 = cube\_def\_cnode\_cs(cube, regn0, mod, 21, NULL);
cnode1 = cube\_def\_cnode\_cs(cube, regn1, mod, 60, cnode0);
cnode2 = cube\_def\_cnode\_cs(cube, regn2, mod, 80, cnode0);
```

CUBE4 supports two kind of parameters of a cnode: numeric and string parameter. Every cnode can carry any number of both of them.

```
cube\_cnode\_add\_numeric\_parameter(cnode0, \textcolor{stringliteral}{"Phase"}, 1);
cube\_cnode\_add\_numeric\_parameter(cnode0, \textcolor{stringliteral}{"Phase"}, 2);
cube\_cnode\_add\_string\_parameter(cnode0, \textcolor{stringliteral}{"Iteration"}, \textcolor{stringlit
cube\_cnode\_add\_string\_parameter(cnode2, \textcolor{stringliteral}{"Etappe"}, \textcolor{stringlit
```

The last dimension is the system tree dimension. Currently CUBE defines the system dimension with the fixed hierarchy: MACHINE → NODES → PROCESSES → THREADS

It leads to the fixed sequence of calls in the system dimension definition:

1. First one creates a root for the system dimension : *cube_machine*. Machine has a name and description.
2. Machine consists of *cube_nodes*. Every *cube_node* has a name and a *cube_↔ machine* as a parent.
3. On every *cube_node* run several *cube_processes* (as many cores are available). *cube_process* has a name, MPI rank and *cube_node* as a parent.
4. Every *cube_process* spawns several (one or more) *cube_threads* (OMP, Pthreads, Java Threads). *cube_thread* has a name, its rank and *cube_process* as a parent.

The cube returns a pointer on *cube_machine*, *cube_node*, *cube_process* or *cube_thread*, which has to be used later to define further level in the system tree or to access the data in the cube.

```
cube\_machine* mach = cube\_def\_mach(cube, \textcolor{stringliteral}{\"MSC<<juelich>>\"}, \textcolor{stringliteral}{\"Athena<<juropa>>\"}, mach);
cube\_node* node = cube\_def\_node(cube, \textcolor{stringliteral}{\"Athena<<juropa>>\"}, mach);
cube\_process* proc0 = cube\_def\_proc(cube, \textcolor{stringliteral}{\"Process 0<<master>>\"}, 0, n);
cube\_process* proc1 = cube\_def\_proc(cube, \textcolor{stringliteral}{\"Process 1<<worker>>\"}, 1, n);
cube\_thread* thrd0 = cube\_def\_thrd(cube, \textcolor{stringliteral}{\"Thread 0<<iterator>>\"}, 0, n);
cube\_thread* thrd1 = cube\_def\_thrd(cube, \textcolor{stringliteral}{\"Thread 1<<solver>>\"}, 1, n);
```

CUBE can carry a set of so called "topologies": mappings $\text{THREAD} \rightarrow (x, y, z, \dots)$

Then the GUI is used to visualize every value (*cube_metric*, *cube_cnode*, *cube_thread*) for selected metric and cnode as a 1D, 2D or 3D set of points with the different colors.

First one specifies a number of dimensions (any number is supported), a vector with the sizes in every dimension and its periodicity and creates a structure of type *cube_cartesian*

```
\textcolor{keywordtype}{long} dimv0[NDIMS] = \{ 5, 5 \};
\textcolor{keywordtype}{int} periodv0[NDIMS] = \{ 1, 0 \};
cube\_cartesian* cart0 = cube\_def\_cart(cube, NDIMS, dimv0, periodv0);
cube\_cart\_set\_name(cart0, \textcolor{stringliteral}{\"Application Topology 1\"});
```

The coordinates are defined like a vector and create a mapping.

```
\textcolor{keywordtype}{long} coordv[NDIMS] = \{ 0, 0 \};
cube\_def\_coords(cube, cart0, thrd1, coordv);
```

```
\textcolor{keywordtype}{long} dimv1[NDIMS] = \{ 3, 3 \};
\textcolor{keywordtype}{int} periodv1[NDIMS] = \{ 1, 0 \};
cube\_cartesian* cart1 = cube\_def\_cart(cube, NDIMS, dimv1, periodv1);
cube\_cart\_set\_name(cart1, \textcolor{stringliteral}{\"MPI Topology 3\"});
```

```
\textcolor{keywordtype}{long} coordv0[NDIMS] = \{ 0, 1 \};
\textcolor{keywordtype}{long} coordv1[NDIMS] = \{ 1, 0 \};
cube\_def\_coords(cube, cart1, thrd0, coordv0);
cube\_def\_coords(cube, cart1, thrd1, coordv1);
```

The same way one can create any number of topologies. They are shown in the GUI.

```
\textcolor{keywordtype}{long} dimv2[4] = \{ 3, 3, 3, 3 \};
\textcolor{keywordtype}{int} periodv2[4] = \{ 1, 0, 0, 0 \};
cube\_cartesian* cart2 = cube\_def\_cart(cube, 4, dimv2, periodv2);
```

```
\textcolor{keywordtype}{long} coordv20[4] = \{ 0, 1, 0, 0 \};
\textcolor{keywordtype}{long} coordv21[4] = \{ 1, 0, 0, 0 \};
cube\_def\_coords(cube, cart2, thrd0, coordv20);
cube\_def\_coords(cube, cart2, thrd1, coordv21);
cube\_cart\_set\_name(cart2, \textcolor{stringliteral}{\"Second\"});
```

```
\textcolor{keywordtype}{long} dimv3[14] = \{ 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3 \};
\textcolor{keywordtype}{int} periodv3[14] = \{ 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0 \};
cube\_cartesian* cart3 = cube\_def\_cart(cube, 14, dimv3, periodv3);
```

```
\textcolor{keywordtype}{long} coordv32[14] = \{ 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 2 \};
\textcolor{keywordtype}{long} coordv33[14] = \{ 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0 \};
cube\_def\_coords(cube, cart3, thrd0, coordv32);
cube\_def\_coords(cube, cart3, thrd1, coordv33);
cube\_cart\_set\_name(cart3, \textcolor{stringliteral}{\"Third\"});
```

Once the dimensions are defined, one fills the cube object with the data. Definition of topologies can be done after filling the cube.

Every data value is specified by three "coordinates": (*cube_metric*, *cube_cnode*, *cube_thread*)

Note, that *cube_machine*, *cube_node* and *cube_process* are not a "coordinate". They are used only to build up the physical construction of the machine.

Actual writing is done metric-wise and row-wise. First all values of one metric written, then the next metric and so on. No mixing of metrics in this sequence is allowed.

Cube writes data row-wise. It means, for a given cnode, one has to provide an array of values, written in the order of threads in the system dimension.

```
\textcolor{keywordtype}{double} sev1[2];

sev1[0]=123.4;
sev1[1]=567.9;
cube\_write\_sev\_row(cube, met0, cnode2, sev1);

sev1[0]=1123.4;
sev1[1]=2567.9;
cube\_write\_sev\_row(cube, met0, cnode1, sev1);

sev1[0]=-1123.4;
sev1[1]=3567.9;
cube\_write\_sev\_row(cube, met0, cnode0, sev1);

sev1[0]=-123.4;
sev1[1]=-567.9;
cube\_write\_sev\_row(cube, met1, cnode0, sev1);

uint64\_t sev2[2];
sev2[0]=23;
sev2[1]=26;
cube\_write\_sev\_row(cube, met2, cnode2, sev2);

printf(\textcolor{stringliteral}{\"Test file %s complete.\\(\\backslash\\)n\"}, cubefile);

cube\_free(cube);
\textcolor{keywordflow}{return} 0;
\}
```

