
CUBE3 — User Manual

Generic Display for Application Performance Data

Version 3.2 / August 18, 2009

Fengguang Song, Felix Wolf, Farzona Pulatova, Markus Geimer, Daniel Becker, Brian Wylie

Copyright © 2008 University of Tennessee
Copyright © 2008-2009 Forschungszentrum Jülich GmbH

Contents

1	Introduction	3
2	Using the Display	4
2.1	Basic Principles	4
2.2	GUI Components	5
2.2.1	Tree Browsers	6
2.2.2	Menu Bar	7
2.2.3	Color Legend	8
2.2.4	Status Bar	8
2.2.5	Context Menus	8
2.3	Topology Display	9
2.3.1	Topology Menu Bar	10
3	Performance Algebra	10
3.1	Difference	10
3.2	Merge	11
3.3	Mean	11
4	Creating CUBE Files	11
4.1	CUBE API	12
4.1.1	Metric Dimension	12
4.1.2	Program Dimension	13
4.1.3	System Dimension	14
4.1.4	Virtual Topologies	15
4.1.5	Severity Mapping	15
4.1.6	Miscellaneous	16
4.1.7	Writer Library in C	17
4.2	Typical Usage	19

Abstract

CUBE is a generic presentation component suitable for displaying a wide variety of performance metrics for parallel programs including MPI and OpenMP applications. Program performance is represented in a multi-dimensional space including various program and system resources. The tool allows the interactive exploration of this space in a scalable fashion and browsing the different kinds of performance behavior with ease. CUBE also includes a library to read and write performance data as well as operators to compare, integrate, and summarize data from different experiments. This user manual provides instructions of how to use the CUBE display, how to use the operators, and how to write CUBE files.

The CUBE3 implementation has incompatible API and file format to preceding versions.

1 Introduction

CUBE (CUBE Uniform Behavioral Encoding) is a generic presentation component suitable for displaying a wide variety of performance metrics for parallel programs including MPI [2] and OpenMP [3] applications. CUBE allows interactive exploration of a multidimensional metric space in a scalable fashion. Scalability is achieved in two ways: hierarchical decomposition of individual dimensions and aggregation across different dimensions. All metrics are uniformly accommodated in the same display and thus provide the ability to easily compare the effects of different kinds of program behavior.

CUBE has been designed around a high-level data model of program behavior called the CUBE *performance space*. The CUBE performance space consists of three dimensions: a metric dimension, a program dimension, and a system dimension. The metric dimension contains a set of metrics, such as communication time or cache misses. The program dimension contains the program's call tree, which includes all the call paths onto which metric values can be mapped. The system dimension contains all the control flows of the program, which can be processes or threads depending on the parallel programming model. Each point (m, c, l) of the space can be mapped onto a number representing the actual measurement for metric m while the control flow of process/thread l was executing call path c . This mapping is called the *severity* of the performance space.

Each dimension of the performance space is organized in a hierarchy. First, the metric dimension is organized in an inclusion hierarchy where a metric at a lower level is a subset of its parent, for example, communication time is below execution time. Second, the program dimension is organized in a call-tree hierarchy. Flat profiles can be represented as multiple trivial call trees consisting only of a single node. Finally, the system dimension is organized in a multi-level hierarchy consisting of the levels: machine, SMP node, process, and thread.

CUBE also includes a library to read and write instances of the previously described data model in the form of an XML file. The file representation is divided into a *metadata* part and a *data* part. The metadata part describes the structure of the three dimensions plus the definitions of various program and system resources. The data part contains the actual severity numbers to be mapped onto the different elements of the performance space.

The display component can load such a file and display the different dimensions of the performance space using three coupled tree browsers (Figure 1). The browsers are connected so that the user can view one dimension with respect to another dimension. For example, the user can click on a particular metric and see its distribution across the call tree. If the CUBE file contains topological information, the distribution of the performance metric across the topology can be examined using the CUBE topology view. Furthermore, the display is augmented with a source-code display

that can show the exact position of a call site in the source code.

As performance tuning of parallel applications usually involves multiple experiments to compare the effects of certain optimization strategies, CUBE includes a new feature designed to simplify cross-experiment analysis. The CUBE algebra [4] is an extension of the framework for multi-execution performance tuning by Karavanic and Miller [1] and offers a set of operators that can be used to compare, integrate, and summarize multiple CUBE data sets. The algebra allows the combination of multiple CUBE data sets into a single one that can be displayed like the original ones.

The following sections explain how to use the CUBE display, how to create CUBE files, and how to use the algebra and other tools.

2 Using the Display

This section explains how to use the CUBE display component. After a brief description of the basic principles, different components of the GUI will be described in detail.

2.1 Basic Principles

The CUBE display consists of three tree browsers, each of them representing a dimension of the performance space (Figure 1). The left tree displays the metric dimension, the middle tree displays the program dimension, and the right tree displays the system dimension. The nodes in the metric tree represent metrics. The nodes in the program dimension can have different semantics depending on the particular view that has been selected. In Figure 1, they represent call paths forming a call tree. The nodes in the system dimension represent machines, nodes, processes, or threads from top to bottom.

Users can perform two types of actions: selecting a node or expanding/collapsing a node. The expansion/collapsion behavior for the system tree is different from the other trees because either all entities of a given level are expanded or none.

Each node is associated with a metric value, which is called the *severity* and is displayed simultaneously using a numerical value as well as a colored square. Colors enable the easy identification of nodes of interest even in a large tree, whereas the numerical values enable the precise comparison of individual values. The sign of a value is visually distinguished by the relief of the colored square. A raised relief indicates a positive sign, a sunken relief indicates a negative sign.

A value shown in the metric tree represents the sum of a particular metric for the entire program, that is, across all call paths and the entire system. A value shown in the call tree represents the sum of the selected metric across all processes or threads for a particular call path. A value shown in the system tree represents the selected metric for the selected call path and a particular system resource. Briefly, a tree is always an aggregation of all of its neighbor trees to the right.

Note that all the hierarchies in CUBE are inclusion hierarchies, meaning that a child node represents a part of the parent node. For example, the metric hierarchy might display cache misses as a child node of cache accesses because the former event is a subset of the latter event. Similarly, in Figure 2 the call path *main* contains the call paths *main-foo* and *main-bar* as child nodes because their execution times are included in their parent's execution time.

The severity displayed in CUBE follows the principle of *single representation*, that is, within a




Figure 1: CUBE display window.

tree each fraction of the severity is displayed only once. The purpose of this display strategy is to have a particular performance problem to appear only once in the tree and, thus, help identify it more quickly. Therefore, the severity displayed at a node depends on the node's state, whether it is expanded or collapsed. The severity of a collapsed node represents the whole subtree associated with that node, whereas the severity of an expanded node represents only the fraction that is not covered by its descendants because the severity of its descendants is now displayed separately. We call the former one *inclusive* severity, whereas we call the latter one *exclusive* severity.



Figure 2: Node of the call tree in collapsed or expanded state.

For instance, a call tree may have a node *main* with two children *main-foo* and *main-bar* (Figure 2). In the collapsed state, this node is labeled with the time spent in the whole program. In the expanded state it displays only the fraction that is spent neither in *foo* nor in *bar*. Note that the label of a node does not change when it is expanded or collapsed, even if the severity of the node changes from exclusive to inclusive or vice versa.

2.2 GUI Components

The GUI consists of a menu bar, three tree browsers, a color legend, and a status bar. In addition, some tree browsers provides a context menu associated with each node that can be used to access node-specific information.

2.2.1 Tree Browsers

The tree browsers are controlled by the left and right mouse buttons. The left mouse button is used to select or expand/collapse a node. The right mouse button is used to pop up a context menu with node-specific information, such as online documentation.

A label in the metric tree shows a metric name. A label in the call tree shows the last callee of a particular call path. If you want to know the complete call path, you must read all labels from the root down to the particular node you are interested in. After switching to the region-profile view (see below), labels in the middle tree denote regions depending on their level in the tree. A label in the system tree shows the name of the system resource it represents, such as a node name or a machine name. Processes and threads are usually identified by a number, but it is possible to give them specific names when creating a CUBE file. The thread level of single-threaded applications is hidden. Note that all trees can have multiple root nodes.

Each tree view has its own drop-down menu, where it is possible to change the way the severity values are displayed. The options include: *absolute value* (default), a *root percentage*, a *selection percentage*, an *external percentage*, a *peer percentage*, or a *peer distribution*. The last two options are only available for the system tree. The absolute value is the real value measured. When displaying a value as a root percentage, the percentage refers to the value shown at the root of the metric tree when it is in collapsed state. However, both absolute mode and root percentage mode have the disadvantage that values can become very small the more you go to the right, since aggregation occurs from right to left. To avoid this problem, the user can switch to selection percentage. Then, a percentage in the right or middle tree always refers to the selection in the neighbor to the left, that is, a percentage in the system dimension refers to the selection in the program dimension and a percentage in the program dimension refers to the selected metric dimension. In this mode the percentages in the middle and right tree always sum up to one hundred percent. Furthermore, to facilitate the comparison of different experiments, users can choose the external percentage mode to display percentages relative to another data set. The external percentage mode is basically like the normal percentage mode except that the value equal to 100% is determined by another data set. The peer percentage mode shows the percentage relative to the maximum amount of peer values (all entities of the current leaf level), depending on the current expansion depth. The severity values for the non-peer nodes are shown as N.A. The peer distribution mode shows the percentage relative to the maximum and non-zero minimum amount of peer values, depending on the current expansion depth. The non-peer node severity values and all peers with exact zero values are shown as N.A. Note that in the absolute mode, all values are displayed in scientific notation. To prevent cluttering the display, only the mantissa is shown at the nodes with the exponent displayed at the color legend.

Each tree view also has a status bar, where the left section shows the selected absolute value and the percentage relative to 100% as defined in the selected percentage mode and the right section shows the value or range according to which colors are assigned depending on the selected mode.

After opening a data set the middle panel shows the call tree of the program. However, a user might wish to know which fraction of a metric can be attributed to a particular region regardless of from where it was called. In this case, the user can switch from the call-tree mode (default) to the region-profile mode (Figure 3). In the region-profile mode, the call-tree hierarchy is replaced with a source-code hierarchy consisting of two levels: region, and subregions. The subregions, if applicable, are displayed as a single child node labeled *subregions*. A *subregions* node represents all regions directly called from the region above. In this way, the user is able to see which fraction

of a metric is associated with a region exclusively, that is, without its regions called from there.

2.2.2 Menu Bar

The menu bar consists of three menus, a file menu, a view menu, and a help menu.

File

The file menu can be used to open and close a file and to exit CUBE. It also allows users to add additional mirrors to the existing ones.

View

The view menu can be used to set a reference data set for the external percentage mode.

If one or more virtual topologies have been defined in the CUBE file, and if the user clicks on the topology tab in the GUI, the *Topology* menu item will be enabled. Otherwise it is disabled. After selecting topology tab, the Cartesian-selection dialog pops up if the CUBE file has multiple topologies. Through this dialog, users can choose a specific topology view to display in a topology tab next to the system tree tab. Please refer to Section 2.3 for detailed information.

Help

Currently, the help menu provides only an About dialog with release information.




Figure 3: CUBE flat profile.

2.2.3 Color Legend

The color is taken from a spectrum ranging from blue to red representing the whole range of possible values. To avoid an unnecessary distraction, insignificant values close to zero are displayed in dark gray. Exact zero values just have the background color.

2.2.4 Status Bar

The numbers $m \times n$ indicate that there are m processes and for each process there are at most n threads in the execution.

2.2.5 Context Menus

All tree views provide a context menu that can be used to obtain specific information on each node. The context menu is accessible via the right mouse button. It displays all or a subset of the options described below.

The call tree has a context menu consisting of two levels. The first-level menu items are *Call site* and *Called region*. Choosing the *Call site* menu shows the information related to the call site, and choosing the *Called region* menu shows the information related to the region being called by the call site (i.e., the callee).

Location: Displays the source-code location of a program resource in textual form (i.e., at which line and in what module). In the module-profile and region-profile modes, it always refers to the location of its associated region. In the call-tree mode, a call-tree node is usually associated with two entities: a callsite and the region called by the callsite. By entering a specific level of the context menu: *Callsite* or *Called region*, users are able to check either the associated call site's or the called region's location. For the call site, it shows the call site's location where it has been called or its calling region's location if the line number of the call site is undefined. For the called region, it shows the location of the region being called by the call site.

Source code: Displays and highlights the source code of a program resource in the source code browser. In the module-profile and region-profile modes, it always shows and highlights the source code of its associated region. In the call-tree mode, since each call-tree node has a context menu of two levels, by choosing the *Call site* menu it displays and highlights the source code of the call site or the block of source code of the calling region. And by choosing the *Called region* menu it displays and highlights the block of code of the region being called by the call site. Note that not all data sets provide sufficient line-number information to show the correct section of the source code.

Online description: Both metrics and regions can be linked to an online description. For example, metrics might point to an online documentation explaining their semantics, or regions representing library functions might point to the corresponding library documentation.

Info: A brief description of the selected node supplied by the CUBE data set.

2.3 Topology Display

In many parallel applications, each process (or thread) communicates only with a limited number of processes. The parallel algorithm divides the application domain into smaller chunks known as sub domains. A process usually communicates with processes owning sub domains adjacent to its own. The mapping of data onto processes and the neighborhood relationship resulting from this mapping is called *virtual topology*. Many applications use one or more virtual topologies (Figure 4) specified as one-, two- or three-dimensional Cartesian grids. The CUBE topology display shows performance data mapped onto the Cartesian topology of the application. The corresponding grid is specified by two parameters: number of dimensions and size of each dimension.




Figure 4: Topology Display

The display consists of a drop-down menu and the actual Cartesian grid. The Cartesian grid is presented by planes stacked on top of each other in a three dimensional projection. The number of planes depends on the number of dimensions in the grid. Each plane is divided into squares. The number of squares depends on the dimension size. Each square represents a system resource (e.g a process) of the application and has a coordinate associate with it.

The grid displays the severity of the selected metric in the selected call path for each system resource participating in the application's topology. The severity is represented as a color. A system resource might not be a part of the application's virtual topology or may have a zero value for a metric. Therefore, it is sometimes possible to have some uncolored squares in the grid picture.

2.3.1 Topology Menu Bar

The menu related to Topology is located in the View Menu. It consists of three submenus: a view menu, a geometry menu, and a zoom menu.

View: The view menu can be used to choose one of the three possible orientations of the grid. The coordinate axes at the bottom of the picture indicate the direction of X, Y and Z dimensions in the three-dimensional space. In case of one- or two- dimensional grids, users are provided with only one orientation of the grid.

Geometry: Due to varying dimension sizes, planes in the grid might overlap with each other and the size of the squares might be too small to recognize their color. This may pose a problem for the user to view the topology information effectively. The geometry menu circumvents this problem by providing options to scale the picture in various ways. The *Angle* option helps the user to adjust the skew of the three-dimensional projection. The *Plane Distance* option helps to adjust the inter-plane distance. The *Plane Length* option helps users scale the area of each plane.

Zoom: The zoom menu can be used to zoom-in or zoom-out on the grid.

3 Performance Algebra

As performance tuning of parallel applications usually involves multiple experiments to compare the effects of certain optimization strategies, CUBE offers a mechanism called *performance algebra* that can be used to merge, subtract, and average the data from different experiments and view the results in the form of a single “derived” experiment. Using the same representation for derived experiments and original experiments provides access to the derived behavior based on familiar metaphors and tools in addition to an arbitrary and easy composition of operations. The algebra is an ideal tool to verify and locate performance improvements and degradations likewise. The algebra includes three operators *diff*, *merge*, and *mean* provided as command-line utilities which take two or more CUBE files as input and generate another CUBE file as output. The operations are closed in the sense that the operators can be applied to the results of previous operations. Note that although all operators are defined for any valid CUBE data sets, not all possible operations make actually sense. For example, whereas it can be very helpful to compare two versions of the same code, computing the difference between entirely different programs is unlikely to yield any useful results.

3.1 Difference

Changing a program can alter its performance behavior. Altering the performance behavior means that different results are achieved for different metrics. Some might increase while others might decrease. Some might rise in certain parts of the program only, while they drop off in other parts. Finding the reason for a gain or loss in overall performance often requires considering the performance change as a multidimensional structure. With CUBE’s difference operator, a user can view this structure by computing the difference between two experiments and rendering the derived result experiment like an original one. The difference operator takes two experiments and computes a derived experiment whose severity function reflects the difference between the minuend’s severity and the subtrahend’s severity.

Usage: cube3_diff [-o output] [-c] [-C] [-h] minuend subtrahend

- o Name of the output file (default: diff.cube)
- c Do not collapse system dimension, if experiments are incompatible
- C Collapse system dimension!
- h Help; Output a brief help message.

3.2 Merge

The merge operator's purpose is the integration of performance data from different sources. Often a certain combination of performance metrics cannot be measured during a single run. For example, certain combinations of hardware events cannot be counted simultaneously due to hardware resource limits. Or the combination of performance metrics requires using different monitoring tools that cannot be deployed during the same run. The merge operator takes an arbitrary number of CUBE experiments with a different or overlapping set of metrics and yields a derived CUBE experiment with a joint set of metrics.

Usage: cube3_merge [-o output] [-c] [-C] [-h] cube ...

- o Name of the output file (default: merge.cube)
- c Do not collapse system dimension, if experiments are incompatible
- C Collapse system dimension!
- h Help; Output a brief help message.

3.3 Mean

The mean operator is intended to smooth the effects of random errors introduced by unrelated system activity during an experiment or to summarize across a range of execution parameters. The user can conduct several experiments and create a single average experiment from the whole series. The mean operator takes an arbitrary number of arguments.

Usage: cube3_mean [-o output] [-c] [-C] [-h] cube ...

- o Name of the output file (default: mean.cube)
- c Do not collapse system dimension, if experiments are incompatible
- C Collapse system dimension!
- h Help; Output a brief help message.

4 Creating CUBE Files

The CUBE data format in an XML instance [5]. The CUBE library provides an interface to create CUBE files. It is a simple class interface and includes only a few methods. This section first describes the CUBE API and then presents a simple C++ program as an example of how to use it.

4.1 CUBE API

The class interface defines a class `Cube`. The class provides a default constructor and fourty methods. The methods are divided into four groups. The first three groups are used to define the three dimensions of the performance space and the last group is used to enter the actual data. In addition, an output operator `<<` to write the data to a file is provided.

4.1.1 Metric Dimension

This group refers to the metric dimension of the performance space. It consists of a single method used to build metric trees. Each node in the metric tree represents a performance metric. Metrics have different units of measurement. The unit can be either “`sec`” (i.e., seconds) for time based metrics, such as execution time, or “`occ`” (i.e., occurrences) for event-based metrics, such as floating-point operations. During the establishment of a metric tree, a child metric is usually more specific than its parent, and both of them have the same unit of measurement. Thus, a child performance metric has to be a subset of its parent metric (e.g., system time is a subset of execution time).

```
Metric* def_met (const std::string &disp_name, const std::string &uniq_name,
                const std::string &dtype, const std::string &uom,
                const std::string &val, const std::string &url,
                const std::string &descr, Metric* parent);
```

Returns a metric with display name `disp_name`, unique name `uniq_name` and description `descr`. `dtype` specifies the data type, which can either be “`INTEGER`” or “`FLOAT`”. `uom` is the unit of measurement, which is either “`sec`” for seconds or “`occ`” for number of occurrences. The `val` field specifies whether there is any data available for this particular metric. It can either be “`VOID`” (no data available, metric will not be shown in CUBE) or an empty string (metric will be shown and data is present). `parent` is a previously created metric which will be the new metric’s parent. To define a root node, use `NULL` instead. `url` is a link to an HTML page describing the new metric in detail. If you want to mirror the page at several locations, you can use the macro `@mirror@` as a prefix, which will be replaced by an available mirror defined using `def_mirror()` (see Section 4.1.6).

```
const std::vector<Metric*>& get_metv () const;
```

Returns a vector with all metrics in the CUBE object.

```
const std::vector<Metric*>& get_root_metv () const;
```

Returns a vector with all roots of the metric dimension in the CUBE object.

```
Metric* get_met (const std::string& uniq_name) const;
```

Returns a metric with the given `uniq_name`. Returns `NULL` if the CUBE object doesn’t contain a metric with this name.

```
Metric* get_root_met (Metric * met);
```

Returns the root metric for the given metric `met`.

4.1.2 Program Dimension

This group refers to the program dimension of the performance space. The entities presented in this dimension are *region*, *call site*, and *call-tree node* (i.e., call paths). A region can be a function, a loop, or a basic block. Each region can have multiple call sites from which the control flow of the program enters a new region. Although we use the term call site here, any place that causes the program to enter a new region can be represented as a call site, including loop entries. Correspondingly, the region entered from a call site is called *callee*, which might as well be a loop. Every call-tree node points to a call site. The actual call path represented by a call-tree node can be derived by following all the call sites starting at the root node and ending at the particular node of interest. The user can choose among three ways of defining the program dimension:

1. Call tree with line numbers
2. Call tree without line numbers
3. Flat profile

A call tree with line numbers is defined as a tree whose nodes point to call sites. A call tree without line numbers is defined as a tree whose nodes point to regions (i.e., the callees). A flat profile is simply defined as a set of regions, that is, no tree has to be defined.

```
Region* def_region (const std::string &name, long begln, long endln,
                  const std::string &url, const std::string &descr,
                  const std::string &mod);
```

Returns a new region with region name *name* and description *descr*. The region is located in the module *mod* and exists from line *begln* to line *endln*. *url* is a link to an HTML page describing the new region in detail. For example, if the region is a library function, the *url* can point its documentation. If you want to mirror the page at several locations, you can use the macro `@mirror@` as a prefix, which will be replaced by an available mirror defined using `def_mirror()` (see Section 4.1.6).

```
Cnode* def_cnode (Region* callee,
                 const std::string &mod, int line,
                 Cnode* parent);
```

Returns a new call-tree node representing a call from call site located at the line *line* of the module *mod*. The call tree node calls the callee *callee* (i.e., a previously defined region). *parent* is a previously created call-tree node which will be the new one's parent. To define a root node, use `NULL` instead. This method is used to create a call tree with line numbers.

```
Cnode* def_cnode (Region* region,
                 Cnode* parent);
```

Defines a new call-tree node representing a call to the region *region*. *parent* is a previously created call-tree node which will be the new one's parent. To define a root node, use `NULL` instead. Note that different from the previous `def_cnode()`, this method is used to create a call-tree without line numbers where each call-tree node points to a region.

To define a call tree with line numbers use `def_cnode(Region*, string, int...)`. To define a call tree without line numbers use `def_cnode(Region*, Cnode*)` instead. To create a flat profile use neither one — just defining a set of regions will be sufficient.

```
const std::vector<Region*>& get_regv () const;
```

Returns a vector with all regions in the CUBE object.

```
const std::vector<Cnode*>& get_cnodev () const;
```

Returns a vector with all call-tree nodes in the CUBE object.

```
Cnode* get_cnode (Cnode & cn) const;
```

Search a call-tree node *cn*. Returns NULL if the CUBE object does not contain the given call-tree node.

4.1.3 System Dimension

This group refers to the system dimension of the performance space. It reflects the system resources which the program is using at runtime. The entities present in this dimension are *machine*, *node*, *process*, and *thread*, which populate four levels of the system hierarchy in the given order. That is, the first level consists of machines, the second level of nodes, and so on. Finally, the last (i.e., leaf) level is populated only by threads. The system tree is built in a top-down way starting with a machine. Note that even if every process has only one thread, users still need to define the thread level.

```
Machine* def_mach (const std::string &name, const std::string &desc);
```

Returns a new machine with the name *name* and description *desc*.

```
Node* def_node (const std::string &name, Machine* mach);
```

Returns a new (SMP) node which has the name *name* and which belongs to the machine *mach*.

```
Process* def_proc (const std::string &name, int rank,  
                  Node* node);
```

Returns a new process which has the name *name* and the rank *rank*. The rank is a number from $0 - (n - 1)$, where *n* is the total number of processes. MPI applications may use the rank in `MPI_COMM_WORLD`. The process runs on the node *node*.

```
Thread* def_thrd (const std::string &name, int rank,  
                 Process* proc);
```

Defines a new thread which has the name *name* and the rank *rank*. The rank is a number from $0 - (n - 1)$, where *n* is the total number of threads spawned by a process. OpenMP applications may use the OpenMP thread number. The thread belongs to the process *proc*.

```
const std::vector<Sysres*>& get_sysv () const;
```

Returns a vector with all system resources (e.g. node, thread, process) available in the CUBE object.

```
const std::vector<Machine*>& get_machv () const;
```

Returns a vector with all machines in the CUBE object.

```
const std::vector<Node*>& get_nodev () const;
```

Returns a vector with all nodes of all machines in the CUBE object.

```
const std::vector<Process*>& get_procv () const;
```

Returns a vector with all processes in the CUBE object.

```
const std::vector<Thread*>& get_thrdrv () const;
```

Returns a vector with all threads in the CUBE object.

```
Machine * get_mach (Machine & mach) const;
```

Search for the machine `mach` in the CUBE object. Returns NULL if the CUBE object does not contain the given machine.

```
Node *get_node (Node & node) const ;
```

Search for the node `node` in the CUBE object. Returns NULL if the CUBE object does not contain the given node.

4.1.4 Virtual Topologies

Virtual topologies are used to describe adjacency relationships among machines, SMP nodes, processes or threads. A topology usually consists of a single class of entities such as threads or processes. The CUBE API provides a set of functions to create Cartesian topologies and to define the machine/SMP node/process/thread mappings onto coordinates. Note that the definition of virtual topologies is optional.

```
Cartesian* def_cart (long ndims, const std::vector<long>& dimv,  
                   const std::vector<bool>& periodv);
```

Defines a new Cartesian topology. `ndims` and `dimv` specify the number of dimensions and the size of each dimension. `periodv` specifies the periodicity for each dimension. Currently, the maximum value for `ndims` is three.

```
void def_coords (Cartesian* cart, Sysres* sys,  
               const std::vector<long>& coordv);
```

Maps a specific system resource onto a Cartesian coordinate. The system resource `sys` may be a machine, SMP node, process or a thread. It is not recommended to map a mixed set of entities onto one topology (e.g., machines and threads are located in the same topology). The parameter of `cart` has been defined by the above `def_cart()` method.

```
const std::vector<Cartesian *>& get_cartv () const ;
```

Returns a vector of all cartesian topologies available in the CUBE object.

```
const Cartesian * get_cart ( int i) const ;
```

Returns in `i`-th topology in the CUBE object.

4.1.5 Severity Mapping

After the establishment of the performance space, users can assign severity values to points of the space. Each point is identified by a tuple (`met`, `cnode`, `thrd`). The value should be inclusive with respect to the metric, but exclusive with respect to the call-tree node, that is it should not cover its children. The default severity value for the data points left undefined is zero. Thus, users only need to define non-zero data points.

```
void set_sev (Metric* met, Cnode* cnode,  
             Thread* thrd, double value);
```

Assigns the value `value` to the point `(met, cnode, thrd)`.

```
void add_sev (Metric* met, Cnode* cnode,  
             Thread* thrd, double value);
```

Adds the value `value` to the present value at point `(met, cnode, thrd)`.

The previous two methods `set_sev()` and `add_sev()` are intended to be used when the program dimension contains a call tree and not a flat profile. As the flat profile does not require the definition of call-tree nodes, the following two functions should be used instead:

```
void set_sev (Metric* met, Region* region,  
             Thread* thrd, double value);
```

Assigns the value `value` to the point `(met, region, thrd)`.

```
void add_sev (Metric* met, Region* region,  
             Thread* thrd, double value);
```

Adds the value `value` to the present value at point `(met, region, thrd)`.

```
double get_sev ( Metric * met, Cnode * cnode, Thread * thrd) const;
```

Returns the value for the point `(met, cnode, thrd)`.

4.1.6 Miscellaneous

Often users may want to define some information related to the CUBE file itself, such as the creation date, experiment platform, and so on. For this purpose, CUBE allows the definition of arbitrary attributes in every CUBE data set. An attribute is simply a key-value pair and can be defined using the following method:

```
void def_attr (const std::string &key, const std::string &value);
```

Assigns the value `value` to the attribute `key`.

CUBE allows using multiple mirrors for the online documentation associated with metrics and regions. The `url` expression supplied as an argument for `def_metric()` and `def_region()` can contain a prefix `@mirror@`. When the online documentation is accessed, CUBE can substitute all mirrors defined for the prefix until a valid one has been found. If no valid online mirror can be found, CUBE will substitute the `./doc` directory of the installation path for `@mirror@`.

```
void def_mirror (const std::string &mirror);
```

Defines the mirror `mirror` as potential substitution for the URL prefix `@mirror@`.

```
std::string get_attr(const std::string &key) const;
```

Returns the attribute in the CUBE object stored for the given `key`.

```
const std::map<std::string, std::string> get_attrs() const;
```

Returns all attributes associated to the CUBE object as a map.

```
const std::vector<std::string>& get_mirrors() const;
```

Returns all mirrors defined in the CUBE object.

```
int get_num_thrd() const;
```

Returns the maximal number of threads per process in the CUBE object.

4.1.7 Writer Library in C

In order to create data files, another possibility is to use the C version of the CUBE writer API. The interface defines a struct `cube_t` and provides the following functions:

```
cube_t* cube_create();
```

Returns a new CUBE structure.

```
void cube_free(cube_t* c);
```

Destroys the given CUBE structure.

```
cube_metric* cube_def_met (cube_t* c, const char* disp_name,  
                           const char* uniq_name, const char* dtype,  
                           const char* uom, const char* val,  
                           const char* url, const char* descr,  
                           cube_metric* parent);
```

Returns a new metric structure.

```
cube_region* cube_def_region (cube_t* c, const char* name, long begin,  
                              long endl, const char* url,  
                              const char* descr, const char* mod);
```

Returns a new region.

```
cube_cnode* cube_def_cnode_cs (cube_t* c, cube_region* callee,  
                               const char* mod, int line,  
                               cube_cnode* parent);
```

Returns a new call-tree node structure with line numbers.

```
cube_cnode* cube_def_cnode (cube_t* c, cube_region* callee,  
                             cube_cnode* parent);
```

Returns a new call-tree node structure without line numbers.

```
cube_machine* cube_def_mach (cube_t* c, const char* name  
                             const char* desc);
```

Returns a new machine.

```
cube_node* cube_def_node (cube_t* c, const char* name,  
                          cube_machine* mach);
```

Returns a new node.

```
cube_process* cube_def_proc (cube_t* c, const char* name,  
                              int rank, cube_node* node);
```

Returns a new process.

```
cube_thread* cube_def_thrd (cube_t* c, const char* name,  
                             int rank, cube_process* proc);
```

Returns a new thread.

```
cube_cartesian* cube_def_cart (cube_t* c, long ndims,  
                               long int* dimv, int* periodv);
```

Defines a new Cartesian topology.

```
void cube_def_coords (cube_t* c, cube_cartesian* cart,  
                    cube_thread* thrd, long int* coord);
```

Maps a thread onto a Cartesian coordinate.

```
void cube_set_sev (cube_t* c, cube_metric* met, cube_cnode* cnode,  
                 cube_thread* thrd, double value);
```

Assigns the severity value to the point (met, cnode, thrd). Can only be used after metric, cnode and thread definitions are complete. Note that you can only use either the region or the cnode form of these calls, but not both at the same time.

```
double cube_get_sev (cube_t* c, cube_metric* met, cube_cnode* cnode,  
                   cube_thread* thrd);
```

Returns the severity of the point (met, cnode, thrd).

```
void cube_set_sev_reg (cube_t* c, cube_metric* met, cube_region* reg,  
                     cube_thread* thrd, double value);
```

Assigns the severity value to the point (met, reg, thrd). Can only be used after metric, regino and thread definitions are complete. Note that you can only use either the region or the cnode form of these calls, but not both at the same time.

```
void cube_add_sev (cube_t* c, cube_metric* met, cube_cnode* cnode,  
                 cube_thread* thrd, double value);
```

Adds the severity value to the present value at point (met, cnode, thrd). Can only be used after metric, cnode and thread definitions are complete. Note that you can only use either the region or the cnode form of these calls, but not both at the same time.

```
void cube_add_sev_reg (cube_t* c, cube_metric* met, cube_region* reg,  
                     cube_thread* thrd, double value);
```

Adds the severity value to the present value at point (met, reg, thrd). Can only be used after metric, region and thread definitions are complete. Note that you can only use either the region or the cnode form of these calls, but not both at the same time.

```
void cube_write_all (cube_t* c, FILE* fp);
```

Writes the entire CUBE data to the given file. This basically corresponds to calling cube_write_def() and cube_write_sev_matrix().

```
void cube_write_def (cube_t* c, FILE* fp);
```

Writes the definitions part of the CUBE data to the given file. Should only be used after definitions are complete.

```
void cube_write_sev_matrix (cube_t* c, FILE* fp);
```

Writes the severity values part of the CUBE data to the given file. Should only be used after severity values are completely set. Unset values default to zero.

```
void cube_write_sev_row (cube_t* c, FILE* fp,  
                       cube_metric* met,  
                       cube_cnode* cnode,  
                       double* sevs);
```

```

1      void foo() {
        ...
10     }
11     void bar() {
        ...
20     }
21     int main(int argc, char* argv) {
        ...
60     foo();
        ...
80     bar();
        ...
100    }

```

Figure 5: Target-application source code `example.c`

Writes the given severity values of (`met`, `cnode`) for all threads to the given file. This can be used instead of `cube_write_sev_matrix()` to incrementally write parts of the severity matrix.

```
void cube_write_finish (cube_t* c, FILE* fp);
```

Writes the end tags to a file. Must be called at the very end before closing the file, but only when incrementally writing the severity matrix using `cube_write_sev_matrix()`. When using `cube_write_sev_matrix()` to write the severity matrix in one chunk, calling this function is not needed.

4.2 Typical Usage

A simple C++ program is given to demonstrate how to use the CUBE write interface. Figure 6 shows the corresponding CUBE display. The source code of the target application is provided in Figure 5.

```

// A C++ example using CUBE write interface
#include <cube3/Cube.h>
#include <string>
#include <fstream>

using namespace std;
using namespace cube;

int main(int argc, char* argv[]) {
    Cube cube;

    // Specify mirrors (optional)
    cube.def_mirror("http://icl.cs.utk.edu/software/kojak/");
    cube.def_mirror("http://www.fz-juelich.de/jsc/kojak/");

```

```

// Specify information related to the file (optional)
cube.def_attr("experiment time", "September 27th, 2006");
cube.def_attr("description", "a simple example");

// Build metric tree
Metric* met0 = cube.def_met("Time", "Time", "FLOAT", "sec", "",
                           "@mirror@patterns-2.1.html#execution",
                           "root node", NULL); // using mirror
Metric* met1 = cube.def_met("User time", "User Time", "FLOAT", "sec", "",
                           "http://www.cs.utk.edu/usr.html",
                           "2nd level", met0); // without using mirror
Metric* met2 = cube.def_met("System time", "System Time", "FLOAT", "sec", "",
                           "http://www.cs.utk.edu/sys.html",
                           "2nd level", met0); // without using mirror

// Build call tree
string mod = "/ICL/CUBE/example.c";
Region* regn0 = cube.def_region("main", 21, 100, "", "1st level", mod);
Region* regn1 = cube.def_region("foo", 1, 10, "", "2nd level", mod);
Region* regn2 = cube.def_region("bar", 11, 20, "", "2nd level", mod);

Cnode* cnode0 = cube.def_cnode(regn0, mod, 21, NULL);
Cnode* cnode1 = cube.def_cnode(regn1, mod, 60, cnode0);
Cnode* cnode2 = cube.def_cnode(regn2, mod, 80, cnode0);

// Build system resource tree
Machine* mach = cube.def_mach("MSC", "");
Node* node = cube.def_node("Athena", mach);
Process* proc0 = cube.def_proc("Process 0", 0, node);
Process* proc1 = cube.def_proc("Process 1", 1, node);
Thread* thrd0 = cube.def_thrd("Thread 0", 0, proc0);
Thread* thrd1 = cube.def_thrd("Thread 1", 1, proc1);

// Build 2D Cartesian a topology (a 5x5 grid)
int ndims = 2;
vector<long> dimv;
vector<bool> periodv;
for (int i = 0; i < ndims; i++) {
    dimv.push_back(5);
    if (i % 2 == 0)
        periodv.push_back(true);
    else
        periodv.push_back(false);
}
Cartesian* cart = cube.def_cart(ndims, dimv, periodv);
vector<long> coord0, coord1;
coord0.push_back(0);
coord0.push_back(0);
coord1.push_back(3);
coord1.push_back(3);
// map the two threads onto the above 2 coordinates
cube.def_coords(cart, thrd0, coord0);
cube.def_coords(cart, thrd1, coord1);

```

```

// Severity mapping
cube.set_sev(met0, cnode0, thrd0, 4);
cube.set_sev(met0, cnode0, thrd1, 4);
cube.set_sev(met0, cnode1, thrd0, 4);
cube.set_sev(met0, cnode1, thrd1, 4);
cube.set_sev(met0, cnode2, thrd0, 4);
cube.set_sev(met0, cnode2, thrd1, 4);
cube.set_sev(met1, cnode0, thrd0, 1);
cube.set_sev(met1, cnode0, thrd1, 1);
cube.set_sev(met1, cnode1, thrd0, 1);
cube.set_sev(met1, cnode1, thrd1, 1);
cube.set_sev(met1, cnode2, thrd0, 1);
cube.set_sev(met1, cnode2, thrd1, 1);
cube.set_sev(met2, cnode0, thrd0, 1);
cube.set_sev(met2, cnode0, thrd1, 1);
cube.set_sev(met2, cnode1, thrd0, 1);
cube.set_sev(met2, cnode1, thrd1, 1);
cube.set_sev(met2, cnode2, thrd0, 1);
cube.set_sev(met2, cnode2, thrd1, 1);

// Output to a cube file
ofstream out;
out.open("example.cube");
out << cube;
}

```




Figure 6: Display of example.cube

References

- [1] K. L. Karavanic and B. Miller. A Framework for Multi-Execution Performance Tuning. *Parallel and Distributed Computing Practices*, 4(3), September 2001. Special Issue on Monitoring Systems and Tool Interoperability.
- [2] Message Passing Interface Forum. *MPI: A Message Passing Interface Standard*, June 1995. <http://www.mpi-forum.org>.
- [3] OpenMP Architecture Review Board. *OpenMP Application Program Interface — Version 2.5*, May 2005. <http://www.openmp.org>.
- [4] F. Song, F. Wolf, N. Bhatia, J. Dongarra, and S. Moore. An Algebra for Cross-Experiment Performance Analysis. In *Proc. of ICPP 2004*, pages 63–72, Montreal, Canada, August 2004.
- [5] World Wide Web Consortium. *Extensible Markup Language (XML) 1.0 (Second Edition)*, October 2000. <http://www.w3.org/TR/REC-xml>.