
User Guide of OmpSs-2

Release OmpSs-2 2024.05

BSC Programming Models

Jul 16, 2024

CONTENTS

1	Installation of OmpSs-2	3
1.1	Installation of ovni (optional)	3
1.2	Installation of Nanos6	4
1.3	Installation of nOS-V	6
1.4	Installation of NODES	7
1.5	Installation of LLVM-based compiler	8
1.6	Installation of Mercurium legacy compiler	9
2	LLVM-based compiler	11
2.1	Compiling OmpSs-2 programs using the LLVM-based compiler	11
2.2	Problems with the LLVM-based compiler	13
2.3	LLVM-based compiler OmpSs-2 support	14
3	Nanos6 Runtime Options	15
3.1	Executing and controlling number of CPUs	15
3.2	Runtime configuration options	15
3.3	Runtime variants	16
3.4	Task data dependencies	17
3.5	Task scheduler	18
3.6	Task worksharings	18
3.7	Stack size	18
3.8	Dynamic Load Balancing (DLB)	18
3.9	CPU manager policies	19
3.10	Benchmarking, instrumenting and debugging	19
3.11	Throttle	28
3.12	Hardware counters	29
3.13	OmpSs-2@Cluster	29
4	NODES Runtime Options	31
4.1	Executing and controlling number of CPUs	31
4.2	Debugging	31
4.3	Generating ovni traces	31
5	nOS-V Library	33
5.1	Linking with nOS-V	33
5.2	Runtime configuration	33
5.3	Debugging	36
5.4	Generating ovni traces	37
6	Mercurium Legacy Compiler Options	39
6.1	Mercurium drivers	39

6.2	Common compilation flags	39
6.3	Compile OmpSs-2 programs	40
6.4	Problems during compilation	41
	Index	43

This document is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).



The information included in this document is provided “as is”, with no warranties whatsoever, including any warranty of merchantability, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. The document is not guaranteed to be complete and/or error-free at this stage and it is subject to changes without further notice.

This document shall never be construed as a commitment by the Barcelona Supercomputing Center. The center will not assume any responsibility for errors or omissions in it. You can send comments, corrections and/or suggestions to pm-tools at bsc.es. This document is provided for informational purposes only. Due to the dynamic nature of the contents is strongly recommended to check if there is an updated version in the following link:

Note: Find a PDF version of this document at <https://pm.bsc.es/ftp/ompss-2/doc/user-guide/OmpSs-2-User-Guide.pdf>

INSTALLATION OF OMPSS-2

The first step is choosing a directory where you will install OmpSs-2. In this document this directory will be referred to as the `INSTALLATION_PREFIX` directory. We recommend you to set an environment variable `INSTALLATION_PREFIX` with the desired installation directory. For instance:

```
$ export INSTALLATION_PREFIX=$HOME/installation/ompss-2
```

1.1 Installation of ovni (optional)

The `ovni` project implements a fast instrumentation library that records small events during the execution of programs to later investigate how the execution happened using `Paraver`.

The Nanos6 runtime system can be configured with `ovni` to extract execution traces. Building Nanos6 with `ovni` is optional but highly recommended. This page is just a quick summary of the installation of `ovni`. For a more detailed information check [ovni GitHub](#) and [ovni online documentation](#).

1. Get an `ovni` release tarball from <https://github.com/bsc-pm/ovni>.
2. Unpack the tarball and enter the just created directory:

```
$ tar -xvf ovni-xxx.tar.gz  
$ cd ovni-xxx
```

3. Export a target installation directory for `ovni`:

```
$ export OVNI_PREFIX=$HOME/installation/ovni
```

4. Create and move to a build directory:

```
$ mkdir build  
$ cd build
```

5. Configure it with CMake:

```
$ cmake -DCMAKE_INSTALL_PREFIX=$OVNI_PREFIX
```

6. Build and install:

```
$ make  
$ make install
```

Note: ovni may need other packages as dependencies (e.g., an MPI installation) and may use several options with the CMake command. Please check the [ovni online documentation](#) for more details.

1.2 Installation of Nanos6

Nanos6 is a runtime that implements the OmpSs-2 parallel programming model, developed by the [Programming Models group](#) at the [Barcelona Supercomputing Center](#).

Nanos6 can be obtained from the [GitHub public repository](#) or by contacting us at pm-tools@bsc.es.

1.2.1 Build requirements

To install Nanos6 the following tools and libraries must be installed:

1. automake, autoconf, libtool, pkg-config, make and a C and C++ compiler
2. boost \geq 1.59
3. hwloc
4. numactl

1.2.2 Optional libraries and tools

In addition to the build requirements, the following libraries and tools enable additional features:

1. [Extrac](#) to generate execution traces for offline performance analysis with [Paraver](#)
2. [CUDA](#) to enable CUDA tasks
3. [PGI](#) or [NVIDIA HPC-SDK](#) to enable OpenACC tasks
4. [PAPI](#) to generate real-time statistics of hardware counters
5. [PQOS](#) to generate real-time statistics of hardware counters
6. [DLB](#) to enable dynamic management and sharing of computing resources
7. [jemalloc](#) to use jemalloc as the default memory allocator, providing better performance than the default glibc implementation. Jemalloc software must be compiled with `--enable-stats` and `--with-jemalloc-prefix=nanos6_je_` to link with the runtime
8. [ovni](#) (\geq 1.5.0) to generate execution traces for performance analysis with [Paraver](#)

1.2.3 Build procedure

Nanos6 uses the standard GNU automake and libtool toolchain. When cloning from a repository, the building environment must be prepared through the command below. When the code is distributed through a tarball, it usually does not need that command.

```
$ ./autogen.sh
```

Note: The `autogen.sh` script was recently introduced. Please use this script instead of the `autoreconf` command; otherwise the autotools configuration will fail.

Then execute the following commands:

```
$ ./configure --prefix=$INSTALLATION_PREFIX ...other options...
$ make
$ make install
```

where `$INSTALLATION_PREFIX` is the directory into which to install Nanos6.

The configure script accepts the following options:

1. `--with-nanos6-clang=prefix` to specify the prefix of the LLVM/Clang installation that supports OmpSs-2
2. `--with-nanos6-mercurium=prefix` to specify the prefix of the Mercurium installation
3. `--with-boost=prefix` to specify the prefix of the Boost installation
4. `--with-papi=prefix` to specify the prefix of the PAPI installation
5. `--with-libnuma=prefix` to specify the prefix of the numactl installation
6. `--with-extrae=prefix` to specify the prefix of the extrae installation
7. `--with-dlb=prefix` to specify the prefix of the DLB installation
8. `--with-papi=prefix` to specify the prefix of the PAPI installation
9. `--with-pqos=prefix` to specify the prefix of the PQoS installation
10. `--with-cuda[=prefix]` to enable support for CUDA tasks
11. `--enable-openacc` to enable support for OpenACC tasks; requires PGI compilers
12. `--with-pgi=prefix` to specify the prefix of the PGI compilers installation, in case they are not in `$PATH`
13. `--enable-monitoring` to enable monitoring and predictions of task/CPU/thread statistics
14. `--enable-chrono-arch` to enable an architecture-based timer for the monitoring infrastructure
15. `--with-ovni=prefix` to specify the prefix of the ovni installation and enable the ovni instrumentation

The `hwloc` dependency is mandatory, and, by default, an internal `hwloc` is embedded to the Nanos6 library. This behavior can be modified through the `--with-hwloc` option, which can take the following values:

- `--with-hwloc=embedded`: The `hwloc` is built and embedded into the Nanos6 library as an internal module. This is useful when user programs may have third-party software (e.g., MPI libraries) that depend on a different `hwloc` version and may conflict with the one used by Nanos6. In this way, the `hwloc` library is internal and is only used by Nanos6. This is the **default** behavior if the option is not present, or no value is provided. See `nanos6-build-embedding` for more information
- `--with-hwloc=pkgconfig`: The `hwloc` is an external installation and Nanos6 should discover it through the `pkg-config` tool. Make sure to set the `PKG_CONFIG_PATH` if the `hwloc` is not installed in non-standard directories
- `--with-hwloc=<prefix>`: A prefix of an external `hwloc` installation

The `jemalloc` dependency is optional but highly **recommended**. This allocator significantly improves the performance of the Nanos6 runtime by optimizing the memory allocations. By default, an internal `jemalloc` is embedded to the Nanos6 library. This behavior can be modified through the `--with-jemalloc` option, which can take the following values:

- `--with-jemalloc=embedded`: The jemalloc is built and embedded into Nanos6 as an internal library. The building process installs the jemalloc headers and libraries in `$INSTALLATION_PREFIX/deps/nanos6/jemalloc` and dynamically links our runtime against the jemalloc library. This is the **default** behavior if the option is not provided. See `nanos6-build-embedding` for more information
- `--with-jemalloc=<prefix>`: A prefix of an external jemalloc installation configured with the `--enable-stats` and `--with-jemalloc-prefix=nanos6_je_` options
- `--with-jemalloc=no` or `--without-jemalloc`: Disable the jemalloc allocator (not recommended)

The location of an external hwloc can be retrieved through `pkg-config` when specifying `--with-hwloc=pkgconfig`. If it is installed in a non-standard location, `pkg-config` can be told where to find it through the `PKG_CONFIG_PATH` environment variable. For instance:

```
$ export PKG_CONFIG_PATH=/apps/HWLOC/2.0.0/INTEL/lib/pkgconfig:$PKG_CONFIG_PATH
```

The `--with-cuda` flag is needed to enable CUDA tasks. The location of CUDA can be retrieved automatically, if it is in standard system locations (`/usr/lib`, `/usr/include`, etc), or through `pkg-config`. Alternatively, for non-standard installation paths, it can be specified using the optional `=prefix` of the parameter.

The `--enable-openacc` flag is needed to enable OpenACC tasks. The location of PGI compilers can be retrieved from the `$PATH` variable, if it is not specified through the `--with-pgi` parameter.

Optionally, if you passed any valid LLVM/Clang or Mercurium installation, you can execute the Nanos6 tests by running:

```
$ make check
```

Embedding software dependencies

As mentioned above, there are some software dependencies that may be embedded into Nanos6. This is the case for `hwloc` and `jemalloc`, which will be embedded by default. The sources of these embedded dependencies are taken from the `deps` sub-directory in this repository. Inside the sub-directory, there is a default `hwloc` and `jemalloc` source tarballs. These tarballs are automatically extracted into `deps/hwloc` and `deps/jemalloc` by our `autogen.sh` script.

These are the source packages that are then built when choosing `--with-hwloc=embedded` or `--with-jemalloc=embedded`. You may change the embedded software version by placing the desired tarball inside the `deps` folder and re-running `autogen.sh` with the option `--embed-<SOFTWARE> <VERSION>`. Currently, `<SOFTWARE>` can be `hwloc` or `jemalloc` and `<VERSION>` should be the desired version number of that software. For instance, a valid option could be:

```
./autogen.sh --embed-jemalloc 5.3.0
```

For the moment, the tarballs must follow the format `deps/<SOFTWARE>-<VERSION>.tar.gz`.

1.3 Installation of nOS-V

The nOS-V runtime is a lightweight tasking library with co-execution capabilities, which is designed to be used as a core for other more complex parallel runtimes and implements the nOS-V tasking API.

nOS-V can be obtained from the [GitHub public repository](#) or by contacting us at pm-tools@bsc.es.

1.3.1 Build requirements

The following software is required to build and install nOS-V:

1. automake, autoconf, libtool, make and a C11 compiler
2. libnuma

1.3.2 Optional libraries and tools

The following software is required to enable additional optional features:

1. [ovni](#) to instrument and generate execution traces for offline performance analysis with [Paraver](#)
2. [PAPI](#) to generate real-time statistics of hardware counters

1.3.3 Build procedure

nOS-V uses the standard GNU automake and libtool toolchain. When cloning from a repository, the building environment must be prepared through the command below. When the code is distributed through a tarball, it usually does not need that command.

```
$ autoreconf -fiv
```

Then execute the following commands:

```
$ ./configure --prefix=$INSTALLATION_PREFIX ...other options...
$ make
$ make install
```

where `$INSTALLATION_PREFIX` is the directory into which to install nOS-V.

The configure script accepts the following options:

1. `--with-ovni=prefix` to specify the prefix of the ovni installation and enable the ovni instrumentation
2. `--with-libnuma=prefix` to specify the prefix of the libnuma installation
3. `--with-papi=prefix` to specify the prefix of the PAPI installation and enable monitoring features
4. `CACHELINE_WIDTH=width` to specify the length in bytes of a cache line in the target architecture

After building, nOS-V tests can be executed using:

```
$ make check
```

1.4 Installation of NODES

The NODES (nOS-V based OmpSs-2 DEpendency System) runtime is a library designed to work on top of the nOS-V runtime. It includes most of the functionalities from its predecessor, Nanos6, whilst leaving the interaction with the system to nOS-V. NODES implements the OmpSs-2 parallel programming model, developed by the [Programming Models group](#) at the [Barcelona Supercomputing Center](#).

NODES can be obtained from the [GitHub public repository](#) or by contacting us at pm-tools@bsc.es.

1.4.1 Build requirements

The following software is required to build and install NODES:

1. automake, autoconf, libtool, pkg-config, make and a C and C++11 compiler
2. boost >= 1.71
3. nOS-V

1.4.2 Optional libraries and tools

The following software is required to enable additional optional features:

1. ovni (>= 1.5.0) to instrument and generate execution traces for offline performance analysis with [Paraver](#)

1.4.3 Build procedure

NODES uses the standard GNU automake and libtool toolchain. When cloning from a repository, the building environment must be prepared through the command below. When the code is distributed through a tarball, it usually does not need that command.

```
$ autoreconf -fiv
```

Then execute the following commands:

```
$ ./configure --prefix=$INSTALLATION_PREFIX ...other options...
$ make
$ make install
```

where \$INSTALLATION_PREFIX is the directory into which to install NODES.

The configure script accepts the following options:

1. --with-nosv=prefix to specify the prefix of the nOS-V installation
2. --with-boost=prefix to specify the prefix of the Boost installation
3. --with-ovni=prefix to specify the prefix of the ovni installation and enable the ovni instrumentation
4. --with-nodes-clang=prefix to specify the prefix of the LLVM/Clang installation that supports OmpSs-2

Optionally, passing a valid LLVM/Clang installation when configuring enables executing the NODES tests by running:

```
$ make check
```

1.5 Installation of LLVM-based compiler

The LLVM website describes a list of [build requirements of LLVM](#).

You should be able to compile and install the LLVM-based compiler with the following commands:

```
$ cmake -S llvm -B build \  
        -DCMAKE_BUILD_TYPE=Release \  
        -DCMAKE_INSTALL_PREFIX=$INSTALLATION_PREFIX \  
        \
```

(continues on next page)

(continued from previous page)

```

-DLLVM_ENABLE_PROJECTS=clang \
-DLLVM_INSTALL_TOOLCHAIN_ONLY=ON \
-DCLANG_DEFAULT_OMPSS2_RUNTIME=libnanos6 \
-DCLANG_DEFAULT_NANOS6_HOME=$INSTALLATION_PREFIX \
-DCLANG_DEFAULT_NODES_HOME=$INSTALLATION_PREFIX
$ cd build
$ make
$ make install

```

The LLVM/Clang can compile OmpSs-2 code and generate the binaries for both Nanos6 and NODES runtime systems. The support for OmpSs-2 and their runtime systems is always enabled. Deciding whether an OmpSs-2 binary is generated for a specific runtime is done when compiling the user application with clang.

During the building stage, LLVM/Clang only needs to know where to find the runtime libraries. The following CMake options are available to specify the path to the runtime libraries and which is the default runtime for OmpSs-2 binaries:

- `CLANG_DEFAULT_NANOS6_HOME` to specify the default path to the Nanos6 installation. If not provided, the **default** path is the installation prefix of LLVM/Clang. This Nanos6 path can be overridden when compiling OmpSs-2 applications by defining the `NANOS6_HOME` environment variable
- `CLANG_DEFAULT_NODES_HOME` to specify the default path to the NODES installation. If not provided, the **default** path is the installation prefix of LLVM/Clang. This NODES path can be overridden when compiling OmpSs-2 applications by defining the `NODES_HOME` environment variable
- `CLANG_DEFAULT_OMPSS2_RUNTIME` to specify which runtime library is targeted by default if the user did not specify which runtime when compiling an application. The accepted values are `libnanos6` (**default**) and `libnodes`

More details about customizing the LLVM build can be found in the [LLVM website](#)

1.6 Installation of Mercurium legacy compiler

Important: Mercurium is the OmpSs-2 legacy compiler and is unsupported now. **You do not have to install it to use OmpSs-2.** We recommend to use the LLVM-based compiler. Check `InstallationLLVM` and *LLVM-based compiler*.

You can find the build requirements, the configuration flags and the instructions to build Mercurium in the following link: <https://github.com/bsc-pm/mcxx>

You should be able to compile and install Mercurium with the following commands:

```

$ autoreconf -fiv
$ ./configure --prefix=$INSTALLATION_PREFIX --enable-ompss-2 --with-nanos6=$INSTALLATION_
↪PREFIX
$ make
$ make install

```


LLVM-BASED COMPILER

This section describes how to use the LLVM-based compiler to compile OmpSs-2 programs. This compiler is based on the open-source [LLVM infrastructure](#) and is licensed under the [LLVM license](#).

2.1 Compiling OmpSs-2 programs using the LLVM-based compiler

2.1.1 C/C++ programs

Following is a very simple OmpSs-2 program in C:

```
/* test.c */
#include <stdio.h>

int main(int argc, char *argv[])
{
    int x = argc;
    #pragma oss task inout(x)
    {
        x++;
    }
    #pragma oss task in(x)
    {
        printf("argc + 1 == %d\n", x);
    }
    #pragma oss taskwait
    return 0;
}
```

Compile it using clang and the `-fompss-2` flag:

```
$ clang -o test -fompss-2 test.c
```

Important: Do not forget the flag `-fompss-2` in both **compilation** and **linking** stages of your application. Otherwise, your application will not be compiled with parallel support or not linked to the tasking runtime library.

And run it:

```
$ ./test
argc + 1 == 2
```

Use the driver clang++ to compile C++ applications:

```
/* test.cpp */
#include <iostream>

template <typename T>
void test(T x)
{
    T t = x;
    #pragma oss task inout(x)
    {
        x *= 2;
    }
    #pragma oss task in(x)
    {
        std::cout << t << " * 2 == " << x << "\n";
    }
    #pragma oss taskwait
}

int main(int argc, char *argv[])
{
    test(42);
    test(84.1);
    return 0;
}
```

Compile it using clang++ and the -fompss-2 flag:

```
$ clang++ -o test -fompss-2 test.cpp
```

And run it:

```
$ ./test
42 * 2 = 84
84.1 * 2 = 168.2
```

2.1.2 Choosing between Nanos6 and NODES runtimes

As stated in InstallationLLVM, our LLVM/Clang can compile OmpSs-2 programs targeting the Nanos6 or the NODES runtime systems. LLVM/Clang follows the order below to decide which runtime system to target when compiling an OmpSs-2 program. The list is ordered from most to least priority:

1. The LLVM/Clang's -fompss-2 flag accepts an optional value to explicitly choose the runtime library when compiling an application. The accepted values are libnanos6 and libnodes. For instance, the flag -fompss-2=libnanos6 will make the application to target the Nanos6 runtime
2. Otherwise, LLVM/Clang reads the OMPSS2_RUNTIME environment variable when compiling or linking OmpSs-2 applications. Similarly, the accepted values for the envar are libnanos6 and libnodes
3. Lastly, LLVM/Clang uses the default runtime specified through the CLANG_DEFAULT_OMPSS2_RUNTIME CMake option at the building stage. As shown in InstallationLLVM, the libnanos6 is the default if not specified

Notice that is user's responsibility to provide the same runtime value when compiling and linking all source code files of the same application.

Once a runtime has been chosen, clang has to decide which is the path where it can find that runtime library. The compiler uses the runtime library directory following these rules from most to least priority:

1. LLVM/Clang reads the NANOS6_HOME or NODES_HOME environment variables to know the installation path of Nanos6 and NODES, respectively. If defined, the default paths to the the runtimes provided at the building stage are ignored
2. Otherwise, the paths provided through the CLANG_DEFAULT_NANOS6_HOME and CLANG_DEFAULT_NODES_HOME CMake options at the building stage are used. See InstallationLLVM

The LLVM/Clang compiler defines the `_OMPSS_2` preprocessor macro when compiling an OmpSs-2 program, as required by the [OmpSs-2 Specification](#). Additionally, the compiler defines the `_OMPSS_2_NANOS6` or `_OMPSS_2_NODES` preprocessor macros depending on the chosen runtime. This way, the programmer can determine the target runtime from the application's code, e.g., for including runtime-specific headers.

2.2 Problems with the LLVM-based compiler

While we put big efforts to make a reasonably robust compiler, you may encounter a bug or problem when using it. There are several errors of different nature that you may run into:

- The compiler ends abnormally with an internal error telling you to report a bug.
- The compiler does not crash but gives an error on your input code and compilation stops, as if your code were not valid.
- The compiler forgets something in the generated code and linking fails.
- Compilation succeeds but the program crashes at runtime.

2.2.1 How can you help us to solve the problem quicker in the LLVM compiler?

If clang crashes, it will generate a reproducer automatically with a message like this:

```
*****
PLEASE ATTACH THE FOLLOWING FILES TO THE BUG REPORT:
Preprocessed source(s) and associated run script(s) are located at:
clang: note: diagnostic msg: /tmp/t-56517d.c
clang: note: diagnostic msg: /tmp/t-56517d.sh
clang: note: diagnostic msg:
*****
```

Send us an email to `pm-tools` at `bsc.es` with those two files attached. Reproducer files are usually very large. Please compress them with `gzip` (or `bzip2` or any similar tool).

If clang does not crash, then follow the next steps to obtain a reproducer:

1. Figure out the compilation command of the file that fails to compile. Make sure you can replicate the problem using that compilation command alone.
2. Add the option `-gen-reproducer` to the compilation command. Clang will generate a reproducer. Send us the reproducer files.

If for some reason the option `-gen-reproducer` does not work for you, use the following steps:

1. Figure out the compilation command of the file that fails to compile. Make sure you can replicate the problem using that compilation command alone.

2. If your compilation command includes `-c`, replace it by `-E`. If it does not include `-c` simply add `-E`.
3. If your compilation command includes `-o file` (or `-o file.o`) replace it by `-o file.ii`. If it does not include `-o`, simply add `-o file.ii`.
4. Now run the compiler with this modified compilation command. It should have generated a `file.ii`.
5. Send us that file. Please compress it with `gzip` (or `bzip2` or any similar tool)

2.3 LLVM-based compiler OmpSs-2 support

The LLVM-based compiler for OmpSs-2 does not support all the features of the legacy Mercurium compiler.

2.3.1 Unsupported features

Feature	Notes and workarounds	Planned as future work?
OpenMP compatibility mode	LLVM already supports OpenMP and it would conflict in the compiler itself	No
Fortran support	Planned once the LLVM's flang project reaches sufficient maturity	Yes

NANOS6 RUNTIME OPTIONS

This section describes how to run OmpSs-2 applications and which runtime options are available.

3.1 Executing and controlling number of CPUs

Nanos6 applications can be compiled and executed in this way:

```
# Compile OmpSs-2 program with LLVM/Clang
$ clang -fompss-2 app.c -o app

# Execute on all available cores of the current session
$ ./app
```

The number of cores that are used is controlled by running the application through the `taskset` command. For instance:

```
# Execute on cores 0, 1, 2 and 4
$ taskset -c 0-2,4 ./app
```

3.2 Runtime configuration options

The behaviour of the Nanos6 runtime can be tuned after compilation by means of a configuration file. All former Nanos6 environment variables are obsolete and will be ignored by the runtime system. Currently, the supported configuration file format is [TOML v1.0.0-rc1](#). The default configuration file is named `nanos6.toml` and can be found in the `share` directory of the Nanos6 installation:

```
$INSTALLATION_PREFIX/share
```

Note: The `nanos6.toml` file was recently moved from `$INSTALLATION_PREFIX/share/doc/nanos6/scripts` to the `$INSTALLATION_PREFIX/share` directory

To override the default configuration, it is recommended to copy the default file and change the relevant options. The first configuration file found will be interpreted, according to the following order:

1. The file pointed by the `NANOS6_CONFIG` environment variable.
2. The file `nanos6.toml` found in the current working directory.
3. The file `nanos6.toml` found in the installation path (default file).

The configuration file is organized into different sections and subsections. Option names have the format `<section>[.<subsection>].<option_name>`. For instance, the dependency system that the runtime should load is specified by the `version.dependencies` option. Inside the TOML configuration file, that option is placed inside the `version` section as follows:

```
[version]
  dependencies = "discrete"
```

Alternatively, if the configuration has to be changed programatically and creating new files is not practical, the configuration options can be overridden using the `NANOS6_CONFIG_OVERRIDE` environment variable. The content of this variable has to be in the format `option1=value1,option2=value2,option3=value3,...`, providing a comma-separated list of assignments.

For example, you can run the following command to change the dependency implementation and use CTF instrumentation:

```
NANOS6_CONFIG_OVERRIDE="version.dependencies=discrete,version.instrument=ctf" ./omps-
↪program`
```

By default, the runtime system emits a warning during initialization when it detects the definition of irrelevant environment variables that start with the `NANOS6` prefix. The exceptions are the previous two variables `NANOS6_CONFIG` and `NANOS6_CONFIG_OVERRIDE`, but also the `NANOS6_HOME` variable. This latter is not relevant for the Nanos6 runtime system but it is often defined by users and it is relevant for the OmpSs-2 LLVM-based compiler. The warning can be disabled by setting the `loader.warn_envvars` configuration option to `false`.

3.3 Runtime variants

There are several Nanos6 runtime variants, each one focusing on different aspects of parallel executions: performance, debugging, instrumentation, etc. OmpSs-2 applications do not require recompiling their code to run with instrumentation, e.g., to extract *Extræ* traces or to generate additional information. This is instead controlled through configuration options, at run-time. Users can select a specific Nanos6 variant when running an application by setting the `version.instrument`, `version.debug` and `version.dependencies` configuration variables. This section explains the different values for these variables.

The instrumentation is specified by the `version.instrument` configuration variable and can take the following values:

`version.instrument = none`

This is the **default** value and does not enable any kind of instrumentation. This is the variant that should be used when executing performance experiments since it is the one that adds no instrumentation overhead. Continue reading this section for more information about performance runs with Nanos6.

`version.instrument = ovni`

Instrumented with `ovni` to generate *Paraver* traces. See *Generating ovni traces*.

`version.instrument = extræ`

Instrumented to produce *Paraver* traces. See: *Generating Extræ traces*.

`version.instrument = ctf`

Instrumented to produce CTF traces and convert them to the *Paraver* format. See: *Generating CTF traces*.

`version.instrument = verbose`

Instrumented to emit a log of the execution. See: *Verbose instrumentation*.

`version.instrument = lint`

Instrumented to support the `OmpSs-2@Linter` tool.

Note: The `graph` and `stats` instrumentations have been recently removed and are no longer available

By default, Nanos6 loads runtime variants that are compiled using high optimization flags and most of the internal assertions turned off. This is the configuration that should be used along with the default `none` instrumentation for benchmarking experiments. However, these are not the only options that provide the best performance. See [Benchmarking](#) for more information.

Sometimes it is useful to run an OmpSs-2 program with debug information for debugging purposes (e.g., when a program crashes). The runtime system provides the option `version.debug` to load a runtime variant that has been compiled without optimizations and with all internal assertions turned on. The default value for this option is `false` (no debug), but can be changed to `true` to enable the debug information. Please note that the runtime system will significantly decrease its performance when enabling this option. Additionally, all instrumentation variants have their optimized and debug variants.

Finally, the last configuration variable used to specify a runtime variant is the `version.dependencies`, which is explained in the next section.

3.4 Task data dependencies

The Nanos6 runtime has support for different dependency implementations. The `discrete` dependencies are the default dependency implementation. This is the most optimized implementation but it does not fully support the OmpSs-2 dependency model since it does not support region dependencies. In the case the user program requires region dependencies (e.g., to detect dependencies among partial overlapping dependency regions), Nanos6 provides the `regions` implementation, which is completely spec-compliant. This latter is also the only implementation that supports `OmpSs-2@Cluster`.

The dependency implementation can be selected at run-time through the `version.dependencies` configuration variable. The available implementations are:

`version.dependencies = "discrete"`

Default and optimized implementation not supporting region dependencies. Region syntax is supported but will behave as a discrete dependency to the first address. Scales better than the default implementation thanks to its simpler logic and is functionally similar to traditional OpenMP model.

`version.dependencies = "regions"`

Supporting all dependency features. Default implementation in `OmpSs-2@Cluster` installations.

In cases where an OmpSs-2 program requires region dependency support, we recommended to add the declarative directive `assert` in any of the program source files, as shown below. Then, before the program is started, the runtime system will check whether the loaded dependency implementation is `regions` and will abort the execution if it is not `true`.

```
#pragma oss assert("version.dependencies==regions")

int main() {
    // ...
}
```

Notice that the `assert` directive could also check whether the runtime is using `discrete` dependencies.

3.5 Task scheduler

The scheduling infrastructure provides the following configuration options to modify the behavior of the task scheduler:

scheduler.policy (default: fifo)

Specifies whether ready tasks are added to the ready queue using a FIFO (`fifo`) or a LIFO (`lifo`) policy.

scheduler.immediate_successor (default: 0.75)

Probability of enabling the immediate successor feature to improve cache data reutilization between successor tasks. If enabled, when a CPU finishes a task it starts executing the successor task (computed through their data dependencies).

scheduler.priority (default: true)

Boolean indicating whether the scheduler should consider the task priorities defined by the user in the task's priority clause.

3.6 Task worksharings

Important: Task worksharings, which were implemented by the `task for` clause, are no longer part of the OmpSs-2 specification.

3.7 Stack size

By default, Nanos6 allocates stacks of 8 MB for its worker threads. In some codes this may not be enough. For instance, when converting Fortran codes, some global variables may need to be converted into local variables. This may increase substantially the amount of stack required to run the code and may surpass the space that is available.

To solve that problem, the stack size can be set through the `misc.stack_size` configuration variable. Its value is expressed in bytes but it also accepts the K, M, G, T and E suffixes, that are interpreted as power of 2 multipliers. For instance:

```
[misc]
stack_size = "16M"
```

3.8 Dynamic Load Balancing (DLB)

DLB is a library devoted to speed up hybrid parallel applications and maximize the utilization of computational resources. More information about this library can be found at the following link: <https://pm.bsc.es/dlb>.

To enable DLB support for Nanos6, a working DLB installation must be present in your environment. Configuring Nanos6 with DLB support is done through the `--with-dlb` flag, specifying the root directory of the DLB installation.

After configuring DLB support for Nanos6, its enabling can be controlled at run-time through the `dlb.enabled` configuration variable. To run Nanos6 with DLB support then, this variable must be set to true (`dlb.enabled=true`), since by default DLB is disabled.

Once DLB is enabled for Nanos6, OmpSs-2 applications will benefit from dynamic resource sharing automatically. Assuming that DLB has been explicitly enabled at the configuration file, the following example showcases the executions of two applications that share the available CPUs between them:

```
# Run the first application using 10 CPUs (0, 1, ..., 9)
$ taskset -c 0-9 ./merge-sort.test &

# Run the second application using 10 CPUs (10, 11, ..., 19)
$ taskset -c 10-19 ./cholesky-fact.test &

# Now those applications should be running while sharing resources
# ...
```

3.9 CPU manager policies

Nanos6 offers different policies when handling CPUs through the `cpumanager.policy` configuration variable:

cpumanager.policy = idle

Activates the idle policy, in which idle threads halt on a blocking condition, while not consuming CPU cycles.

cpumanager.policy = hybrid

Activates the hybrid policy, in which idle threads continue spinning for a specific number of iterations while consuming CPU cycles and then halt on a blocking condition.

cpumanager.policy = busy

Activates the busy policy, in which idle threads continue spinning and never halt, consuming CPU cycles.

cpumanager.policy = lewi

If DLB is enabled, activates the LEnd When Idle policy, which is similar to the idle policy but in DLB mode. In this policy, idle threads lend their CPU to other runtimes or processes.

cpumanager.policy = greedy

If DLB is enabled, activates the greedy policy, which disables lending CPUs from the process' mask, but allows acquiring and lending external CPUs.

cpumanager.policy = default

Fallback to the **default** implementation. If DLB is disabled, this policy falls back to the **hybrid** policy, while if DLB is enabled it falls back to the **lewi** policy.

Furthermore, the number of busy-wait iterations performed in the hybrid policy can be controlled by the following configuration variable:

cpumanager.busy_iters = X

This variable indicates the collective number of busy iterations performed by all CPUs. Thus, when setting a value of X , idle threads will spin for $X/numCPUs$ iterations before halting. By default this number is set to 240K iterations collectively.

3.10 Benchmarking, instrumenting and debugging

As previously stated, there are several Nanos6 runtime variants, each one focusing on different aspects of parallel executions: performance, debugging, instrumentation, etc. OmpSs-2 applications do not require recompiling their code to run with instrumentation, e.g., to extract Paraver traces or to generate additional information. This is instead controlled through configuration options, at run-time. Users can select a specific Nanos6 variant when running an application by setting the `version.instrument`, `version.debug` and `version.dependencies` configuration variables. The next subsections explain the details of the different variants and options of Nanos6.

Important: The ovni instrumentation is the recommended variant to extract Paraver traces. Our support for Extrae and CTF instrumentations is deprecated, so they will not include new further features.

3.10.1 Benchmarking

Nanos6 loads an optimized runtime variant if the `version.debug` option is disabled (by default is disabled). The optimized runtime variants are compiled with high optimization flags, they do not perform validity checks and they do not provide debug information. The benchmarking of OmpSs-2 applications should be performed with an optimized version (`version.debug=false`) and no instrumentation (`version.instrument=none`) to get optimal performance.

Additionally, Nanos6 offers an extra performance option called Turbo, which enables even more aggressive optimizations. This option is disabled by default but it can be enabled by setting `turbo.enabled=true`. Currently, this option enables two Intel® floating-point (FP) unit optimizations in all tasks: flush-to-zero (FZ) and denormals are zero (DAZ). However, please note these FP optimizations could alter the precision of floating-point computations.

We recommend to combine a optimized version without instrumentation (`version.debug=false` and `version.instrument=none`) with the Turbo feature (`turbo.enabled=true`), the discrete dependency implementation (`version.dependencies=discrete`) and the Jemalloc memory allocator to obtain the best performance.

Jemalloc must be enabled at configuration time by passing `--with-jemalloc` to the Nanos6 configure script. See the Nanos6 building requirements to check which are the requirements for supporting jemalloc.

3.10.2 Generating ovni traces

Nanos6 can generate execution traces with the ovni library, which generates lightweight binary traces, and it is possible to mix ovni-instrumented libraries together with an OmpSs-2 program and obtain a single coherent trace.

To enable the generation of ovni traces, Nanos6 must be built with the `--with-ovni` option, and without `--disable-ovni-instrumentation`. The application must run with the `version.instrument=ovni` configuration option. The trace will be left in a `ovni/` directory, which can be transformed into a Paraver trace with the `ovniemu` utility. The Paraver configuration files (views) can be found in the `ovni/cfg` directory.

See the [ovni documentation](#) for more details.

The level of detail can be controlled with the `instrument.ovni.level` configuration option, a higher number includes more events but also incurs in a larger performance penalty.

3.10.3 Generating Extrae traces

To generate a Paraver trace using Extrae, the `version.instrument` configuration variable must be set to `extrae` before running the application.

By default, the runtime will generate a trace as if `EXTRAE_ON` was set to 1. In addition, the `EXTRAE_CONFIG_FILE` environment variable can be set to an Extrae configuration file for fine tuning, for instance, to enable recording hardware counters. See: the [Extrae documentation](#).

The amount of information generated in the Extrae traces can be controlled through the `instrument.extrae.detail_level` configuration option. Its value determines a level of detail that goes from 0, which is the least detailed, up to 8. The default level is 1. Lower levels incur in less overhead and produce smaller traces. Higher levels have more overhead, produce bigger traces, but are more precise and contain more information. The information generated at each level is incremental and is the following:

Level 0

Basic level

Includes basic information about the runtime state and the execution of tasks.

Counters about the number of tasks in the system are approximate and are sampled at periodic intervals.

Level 1

Default level

Counters about the number of tasks in the system are precise in terms of time and value.

Adds communication records that link the point of instantiation of a task to the point where they start their execution.

Adds communication records that link the point where a task get blocked to the point where a task unblocks it (makes it ready), and from that point to the point where the task actually resumes its execution.

Adds communication records that show task dependency relations. The set of predecessors a task that is shown is limited to the set of tasks that have not finished their execution once said task has been instantiated. The links go from the end of the execution of a predecessor to the start of the execution of the successor.

Levels 2 to 7

Unused

Currently they do not add further information.

Level 8

Very detailed level

Adds communication records that link the end of the execution of a task to the point where its parent returns from a taskwait. The trace only contains links for the tasks that have not finished once their parent enters the taskwait. This information is only available at level 8, since it may make the trace significantly bigger.

The runtime installation contains a set of already made Paraver configuration files at the following subdirectory:

```
$INSTALLATION_PREFIX/share/doc/nanos6/paraver-cfg/nanos6
```

Support for hardware counters is enabled through file specified in the `EXTRAE_CONFIG_FILE` environment variable. The procedure is explained in the [Extræe documentation](#). However, the `instrument.extrae.as_threads` configuration variable must also be set to `true`. This is a temporary measure that is needed to produce correct hardware counter information. The resulting trace will expose the actual runtime threads, as opposed to the CPU view that is generated by default.

3.10.4 Generating CTF traces

The CTF instrumentation backend is a performance evaluation and debugging tool aimed for Nanos6 users and developers alike. Its purpose is to record the most relevant information of a Nanos6 execution efficiently and to provide means for visually inspecting it offline. The Nanos6 CTF backend is simple (no installation prerequisites) and mostly lock-free (see the section “Implementation details” for more information).

Nanos6 stores traces in the Common Trace Format (CTF). CTF traces can be directly visualized with tools such as `babeltrace1` or `babeltrace2` (raw command line inspector). But it is recommended first to convert them to Paraver traces using the provided `ctf2prv` command. Although Nanos6 requires no special packages to write CTF traces, the `ctf2prv` converter needs `python3` and the `babeltrace2` python bindings.

The Nanos6 CTF backend will transparently collect hardware counters information if the Nanos6 Hardware counters infrastructure has been enabled. See the User Guide for more information on how to enable hardware counters.

This backend objective is not to replace Extræe, but to provide a minimum set of features for easy Nanos6 introspection. More sophisticated analysis (such as task dependency graphs) will have to be performed with Extræe as usual.

Implementation details

Nanos6 keeps a per-core buffer to store CTF events before flushing them to the storage device. Therefore, events generated by Nanos6 Worker threads do not require to take a lock. However, events triggered by a thread not directly controlled by Nanos6 share a buffer, and it is necessary to take a lock.

When a buffer is full, it is needed to flush it to the storage device. Flushing is done synchronously (i.e. a worker thread that generates a tracepoint will have to flush the buffer if there no free space left) and it might affect the execution workflow. Flushing operations are also recorded and can be inspected with the “CTF flush buffers to disk” view.

Traces are written by default under the `TMPDIR` environment variable or under `/tmp` if not set. Traces written to `/tmp` are kept in RAM memory (see `tmpfs` for more information) and flushing translates to a memory copy operation. When an application execution finishes, Nanos6 copies the trace to the current directory. The configuration file provides an option `instrument.ctf.tmpdir` to override the default location.

Usage

To generate a CTF trace, run the application with the `version.instrument` configuration variable set to `ctf`.

This will create a `trace-<app_name>-<app_pid>` folder in the current directory, hereinafter referred to as `$TRACE` for convenience. The subdirectory `$TRACE/ctf` contains the ctf trace as recorded by Nanos6.

By default, Nanos6 will convert the trace automatically at the end of the execution unless the user explicitly sets the configuration variable `instrument.ctf.converter.enabled = false`. The converted Paraver trace will be stored under the `$TRACE/prv` subdirectory. The environment variable `CTF2PRV_TIMEOUT=<minutes>` can be set to stop the conversion after the specified elapsed time in minutes. Please note that the conversion tool requires python3 and the `babeltrace2` package.

Additionally, there is command to manually convert a trace:

```
$ ctf2prv $TRACE
```

which will generate the directory `$TRACE/prv` with the Paraver trace.

Linux Kernel tracing

The requirements for tracing kernel events are: - Linux Kernel \geq 4.1.0 - Set `/proc/sys/kernel/perf_event_paranoid` to `-1` - Grant the current user read permissions for `/sys/kernel/debug/tracing` or `/sys/kernel/tracing` (tracefs standard mountpoints)

The Linux Kernel provides a set of events (also named tracepoints). To enable kernel events tracing in Nanos6, the user needs to specify which events wants to collect through the `nanos6` configuration file. The option `instrument.ctf.events.kernel.presets` allows to specify a list of presets that ease selecting events for supported Paraver views. For instance, try to set `instrument.ctf.events.kernel.presets=[preemption]` to enable preemption-related events.

Please check the Paraver views description below for the list of presets that each view requires.

If the user does not want to enable some of the events defined in a preset, it can blacklist them in the `instrument.ctf.events.kernel.exclude` list.

Additionally, the user can provide a file path in `instrument.ctf.events.kernel.file` with a list of raw kernel events to enable. An example file named `kernel_events.conf` follows.

```
$ cat kernel_events.conf
sched_switch
sys_enter_open
```

(continues on next page)

(continued from previous page)

```
sys_exit_open
# mm_page_alloc # commented lines are not enabled
```

A list of events that your system supports can be obtained with the command:

```
$ perf list tracepoint | grep Tracepoint | awk '{sub(/.*:\/, "", $1); print $1}'
```

Or by inspecting the contents of tracefs, usually mounted by default at `/sys/kernel/tracing/events` or `/sys/kernel/debug/tracing/events`. Also, Nanos6 creates the file `$TRACE/nanos6_kerneldefs.json` which contains a human-readable format definition of all supported system events. Unfortunately, there is no official Linux description of each event at the moment; please check the Linux Kernel source code for a definition of the events you are interested into.

If too many events are enabled, tracing might fail due to reaching the limit of open files in the system (one file descriptor is opened per event and core). In that case, please, increase the open file limit, reduce the number of events and/or the number of cores.

The CTF kernel trace is stored under `$TRACE/ctf/kernel` whilst the Nanos6 internal events trace are stored in `$TRACE/ctf/ust`. Nanos6 will convert the generated CTF kernel trace (both user and kernel) into Paraver at the end of the execution, as usual.

Paraver views

A number of Paraver views are provided with each Nanos6 distribution under the directory:

```
$INSTALLATION_PREFIX/share/doc/nanos6/paraver-cfg/ctf2prv
```

Some views refer to “Nanos6 core code”. This encompasses all code that it is not strictly user-written code (task code). This includes, for instance, the worker thread main loop (where tasks are requested constantly) the Nanos6 initialization code or the shutdown code.

The Paraver minimum resolution is 1 pixel. Hence, if the current zoom level is not close enough, multiple events might fall into the same pixel. In this situation, Paraver must choose which of these events must be colored. Several visualization options to tackle this problem can be found under the “Drawmode” option in the context menu (right mouse click on a Paraver view). By default, most views have the “Random” mode, which means that drawn events are selected randomly. This is useful to avoid giving priority to an event or another, but the Paraver user must be aware that when zooming in (or out) events might seem to “move” due to another set of events being selected randomly after zooming. In the particular case of graphs (such as the “Number of” views), the Drawmode “Maximum” option must be selected if searching for maximum peaks, otherwise the event with the maximum value might not be selected to be drawn and the peak might remain hidden. Similarly, “Minimum” must be selected if searching for minimum peaks.

Please, note that some views might complain about missing events when opened in Paraver. It is likely that the missing event is “Runtime: Busy Waiting” as it is only generated when running Nanos6 under the busy policy. You can safely ignore the message. If unsure, you can try disabling the “Runtime: Busy Waiting” in your view, save the cfg, and reload the cfg again.

Do not use the Extrae cfg’s views as some event identifiers are not compatible.

Tasks

Shows the name (label) of tasks executed in each core.

Tasks and runtime

Shows the name (label) of tasks executed in each core. It also displays the time spent on Nanos6 Runtime (shown as “Runtime”) and threads busy waiting (shown as “busy wait”) while waiting for ready tasks to execute.

Task source code

Shows the source code location of tasks executed in each core. It also displays the time spent on Nanos6 Runtime (shown as “Runtime”) and threads busy waiting (shown as “busy wait”) while waiting for ready tasks to execute.

Hardware counters

Shows the collected hardware counters (HWC) information. Please note that you must enable HWC collection in Nanos6 first. See the User Guide for more details.

The HWC information is collected per each task burst (a task executing without interruption such as blocking) and per also inside Nanos6 core code. By default, it displays PAPI_TOT_INS. Each HWC is displayed as a different Extrae Event. If you want to inspect another counter, please, modify the “Value to display” subview of the “Hardware counters” view to display the appropriate Extrae Event (HWC event) under Paraver’s “Filter -> Events -> Event Type -> Types” menu.

Task id

Shows the unique id of the tasks executed in each core.

The default colouring of this view is “Not Null Gradient Color”, which eases identifying the execution progress of tasks as they were created. The id’s 0, 1, 2 and 3 are reserved for Idle, Runtime, Busy Wait and Task, respectively. Hence, it might be interesting to manually set the Paraver’s “Semantic Minimum” value to 4, which will not include these values for the gradient colouring. This is the value by default, but if you perform a semantic adjustment, you will need to change the “Semantic Minimum” manually again.

It might also be interesting to draw this view as “Code Color” which will make it easier to spot different consecutive tasks. Under this colouring scheme, it is easier to see the time spent running Nanos6 core code shown as “Runtime” and when threads perform a “busy wait” while waiting for more tasks to execute.

Thread id

Shows the thread Id (TID) of the thread that was running in each core. Be aware that this view only shows the thread placement according to Nanos6 perspective, not the OS perspective. This means that even if this view shows a thread running uninterruptedly in a core for a long time, the system might have preempted the Nanos6 thread by another system thread a number of times without this being displayed in the view.

Runtime status simple

Shows the a simplified view of the runtime states. It displays the time spent running Nanos6 core code shown as “Runtime”, when threads perform a “busy wait” while waiting for more tasks to execute and when the runtime is running task’s code shown as “Task”.

Runtime subsystems

Shows the activity of the most interesting Nanos6 subsystems. In essence, a coarse-grained classification of the time spent while running Nanos6 core code. The displayed subsystems include: Task creation/initialization, dependency registration/unregistration, scheduling add/get tasks and polling services.

Number of blocked tasks

Shows a graph of tasks in the blocked state. Blocked tasks are tasks that started running at some point and then stopped running before completing. This might block, for instance, due to a taskwait directive.

You might want to change paraver’s “Drawmode” option to “Maximum” or “Minimum” if searching for maximum or minimum peaks. By default, it is set to “Maximum”, which means that minimum peaks might be hidden if the view is too much zoomed out.

Number of running tasks

Shows a graph of tasks being executed by some worker thread.

Number of created tasks

Shows a graph with the count of total created tasks. A task is created when its Nanos6 data structures are allocated and registered within the Nanos6 core.

Number of created workers

Shows a graph of created worker threads by Nanos6. Once a worker thread is created, it is not destroyed until

the end of the application's execution.

Number of running workers

Shows a graph of Nanos6 worker threads that are allowed to run on a core. Please note that even if Nanos6 allows a worker to run, another system thread might have temporarily preempted that worker. System preemptions are not displayed in this view.

Number of blocked workers

Shows a graph of worker threads that are blocked (no longer running), either because they are idle or because they are waiting for some event. This view only counts workers that blocked due to Nanos6 will. If a worker that is running a task blocks because of a blocking operation performed by the task code (hence, outside the scope of Nanos6) it will not be shown in this view.

CTF flush buffers to disk

Eventually, the Nanos6 buffers that hold the events are filled and need to be flushed to disk. When this happens, the application's execution might be altered because of the introduced overhead. This flushing can occur in a number of places within the Nanos6 core, either because the buffers were completely full or because Nanos6 decided to flush them before reaching that limit.

This view shows exactly when the flushing happened. If attempting to write the event A into the Nanos6 event buffer triggers a flush, the produced Nanos6 trace will first show the event that triggered the flush followed by the flushing events.

Paraver views for kernel events

All the views listed below require Linux Kernel events. Please, check the section "Linux Kernel tracing" for more details.

The Paraver views that depend on kernel events can be found under:

```
$INSTALLATION_PREFIX/share/doc/nanos6/paraver-cfg/nanos6/ctf2prv/kernel
```

Kernel preemptions

Uses Linux Kernel events to show preemptions affecting the cores where the traced application runs. Preemptions include interruptions and both user and kernel threads. Interrupts will be shown as "IRQ" and "Soft IRQ". IRQs are run in interrupt context. Soft IRQs are deferred interrupt work that run out of interrupt context.

This view requires the Linux Kernel `preemption` preset (interrupts and threads) or the `context_switch` preset (only threads). If using the `context_switch` preset, Paraver might warn about not finding interrupt events, but it is safe to inspect the view.

Kernel thread ids

Uses Linux Kernel events to show the Thread id (TID) of the running thread in each core used by Nanos6. This view requires the Linux Kernel `context_switch` preset.

Kernel system calls

Uses Linux Kernel events to show the system calls performed by each thread on cores used by Nanos6. In this view, all threads but the traced application will be displayed as "Other threads" to enhance the readability of system calls.

This view requires the Linux Kernel `syscall` preset to trace all system calls. If the user only wants to trace a set of syscalls, it can instead specify the `context_switch` preset and provide a list of system call events manually using the `instrument.ctf.events.kernel.file` option in the Nanos6 configuration file.

3.10.5 Verbose instrumentation

To enable verbose logging, run the application with the `version.instrument` configuration variable set to `verbose`.

By default it generates a lot of information. This is controlled by the `instrument.verbose.areas` configuration variable, which can contain a list of areas. The areas are the following:

AddTask

Task creation.

Blocking

Blocking and unblocking within a task through calls to the blocking API.

ComputePlaceManagement

Starting and stopping compute places (CPUs, GPUs, etc).

DependenciesByAccess

Dependencies by their accesses.

DependenciesByAccessLinks

Dependencies by the links between the accesses to the same data.

DependenciesByGroup

Dependencies by groups of tasks that determine common predecessors and common successors.

LeaderThread

Execution of the leader thread.

LoggingMessages

Additional logging messages.

TaskExecution

Task execution.

TaskStatus

Task status transitions.

TaskWait

Entering and exiting taskwaits.

ThreadManagement

Thread creation, activation and suspension.

UserMutex

User-side mutexes (critical).

The case is ignored, and the `all` keyword enables all of them. Additionally, an area can have the `!` prepended to it to disable it. For instance, `areas = ["AddTask", "TaskExecution" , "TaskWait"]` is a good starting point.

The default value is:

```
[instrument]
[instrument.verbose]
  areas = ["all", "!ComputePlaceManagement", "!DependenciesByAccess", "!
↪DependenciesByAccessLinks",
          "!DependenciesByGroup", "!LeaderThread", "!TaskStatus", "!ThreadManagement"]
```

By default, events are recorded with their timestamp and ordered accordingly. This can be disabled by setting the `instrument.verbose.timestamps` configuration variable to `false`.

The output is emitted by default to standard error, but it can be sent to a file by specifying it through the `instrument.verbose.output_file` option. Also the `instrument.verbose.dump_only_on_exit` can be set to `true` to delay the output to the end of the program to avoid getting it mixed with the output of the program.

3.10.6 Debugging

By default, the runtime is optimized for speed and will assume that the application code is correct. Hence, it will not perform most validity checks. To enable validity checks, run the application with the `version.debug` configuration variable set to `true`. This will enable many internal validity checks that may be violated when the application code is incorrect.

To debug an application with a regular debugger, please compile its code with the regular debugging flags and also the `-keep` flag (Mercurium). This flag forces Mercurium to dump the transformed code in the local file system, so that it is available for the debugger.

The runtime system has a component named “loader” that is the responsible for selecting a runtime variant that meets the requirements set by the user through the configuration variables. The loader is also responsible for calling some functions that initialize the runtime variant and convert the `main` function into a regular task. To debug problems due to the installation, run the application with the `loader.verbose` configuration variable set to `true`.

3.10.7 Runtime loader verbosity

The `loader.verbose` configuration option controls the verbosity of the Nanos6 loader. By default (`false` value), it is quiet. If the configuration variable has value `true` it will emit to standard error the actions that it takes and their outcome.

By default, the loader will attempt to load the actual runtime library from the path determined by the operating system (taking into account the `rpath` and the `LD_LIBRARY_PATH` environment variable). If it fails to load the library, then it will attempt to locate the library at the same location as the Nanos6 loader.

The default search path can be overridden through the `loader.library_path` configuration variable. If it exists the first attempt at loading the runtime will be performed at the directory specified in that variable. The loader does not accept multiple directories in that variable.

The Nanos6 loader resolves the addresses of the API functions to the actual runtime implementation. In addition, it also checks for the implementation of some features, and if they are not found, it will either complain or emit a warning and fall back to a compatible but less powerful implementation. More specifically, the loader accepts running applications that make use of weak dependencies and will fall back to strong dependencies if the runtime does not have support for them.

3.10.8 Reporting runtime information

Information about the runtime may be obtained by running the application with the `loader.report_prefix` configuration variable, or by invoking the following command:

```
$ nanos6-info --runtime-details
Runtime path /opt/nanos6/lib/libnanos6-optimized.so.0.0.0
Runtime Version 2017-11-07 09:26:03 +0100 5cb1900
Runtime Branch master
Runtime Compiler Version g++ (Debian 7.2.0-12) 7.2.1 20171025
Runtime Compiler Flags -DNDEBUG -Wall -Wextra -Wdisabled-optimization -Wshadow -
↪ fvisibility=hidden -O3 -flto
Initial CPU List 0-3
NUMA Node 0 CPU List 0-3
Scheduler priority
Dependency Implementation regions (linear-regions-fragmented)
Threading Model pthreads
```

The `loader.report_prefix` configuration variable may be defined as an empty string or it may contain a string that will be prepended to each line. For instance, it can contain a sequence that starts a comment in the output of the program. Example:

```
$ NANOS6_CONFIG_OVERRIDE="loader.report_prefix=#" ./app
Some application output ...
# string version 2017-11-07 09:26:03 +0100 5cb1900 Runtime Version
# string branch master Runtime Branch
# string compiler_version g++ (Debian 7.2.0-12) 7.2.1 20171025 Runtime_
↳ Compiler Version
# string compiler_flags -DNDEBUG -Wall -Wextra -Wdisabled-optimization -Wshadow -
↳ fvisibility=hidden -O3 -flto Runtime Compiler Flags
# string initial_cpu_list 0-3 Initial CPU List
# string numa_node_0_cpu_list 0-3 NUMA Node 0 CPU List
# string scheduler priority Scheduler
# string dependency_implementation regions (linear-regions-fragmented)
↳ Dependency Implementation
# string threading_model pthreads Threading Model
```

3.11 Throttle

There are some cases where user programs are designed to run for a very long time, instantiating in the order of tens of millions of tasks or more. These programs can demand a huge amount of memory in small intervals when they rely only on data dependencies to achieve task synchronization. In these cases, the runtime system could run out of memory when allocating internal structures for task-related information if the number of instantiated tasks is not kept under control.

To prevent this issue, the runtime system offers a *throttle* mechanism that monitors memory usage and stops task creators while there is high memory pressure. This mechanism does not incur too much overhead because the stopped threads execute other ready tasks (already instantiated) until the memory pressure decreases. The main idea of this mechanism is to prevent the runtime system from exceeding the memory budget during execution. Furthermore, the execution time when enabling this feature should be similar to the time in a system with infinite memory.

The throttle mechanism requires a valid installation of Jemalloc, which is a scalable multi-threading memory allocator. Hence, the runtime system must be configured with the `--with-jemalloc` option. Although the throttle feature is disabled by default, it can be enabled and tuned at runtime through the following configuration variables:

throttle.enabled (default: false)

Boolean variable that enables the throttle mechanism.

throttle.tasks (default: 5.000.000)

Maximum absolute number of alive childs that any task can have. It is divided by 10 at each nesting level.

throttle.pressure (default: 70)

Percentage of memory budget used at which point the number of tasks allowed to exist will be decreased linearly until reaching 1 at 100% memory pressure.

throttle.max_memory (default: available physical memory / 2)

Maximum used memory or memory budget. Note that this variable can be set in terms of bytes or in memory units. For example: `throttle.max_memory=50GB`.

3.12 Hardware counters

Nanos6 offers an infrastructure to obtain hardware counter statistics of tasks with various backends. The usage of this API is controlled through the Nanos6 configuration file. Currently, Nanos6 supports the PAPI, RAPL and PQoS backends. All the available hardware counter backends are listed in the default configuration file, found in the scripts folder. To enable any of these, modify the `false` fields and change them to `true`. Specific counters can be enabled or disabled by adding or removing their name from the list of `counters` inside each backend subsection.

Next we showcase a simplified version of the hardware counter section of the configuration file, where the PAPI backend is enabled with counters that monitor the total number of instructions and cycles, and the PAPI backend is enabled as well:

```
[hardware_counters]
  [hardware_counters.papi]
    enabled = true
    counters = ["PAPI_TOT_INS", "PAPI_TOT_CYC"]
  [hardware_counters.rapl]
    enabled = true
```

3.13 OmpSs-2@Cluster

In order to enable OmpSs-2@Cluster support, you need a working MPI installation in your environment that supports multithreading, i.e. `MPI_THREAD_MULTIPLE`. Nanos6 needs to be configured with the `--enable-cluster` flag.

For more information on how to write and run cluster applications see [Cluster.md](#).

NODES RUNTIME OPTIONS

This section describes how to run OmpSs-2 applications using NODES, and which runtime options are available.

4.1 Executing and controlling number of CPUs

NODES applications can be compiled and executed in this way:

```
# Compile OmpSs-2 program with LLVM/Clang
$ clang -fompss-2=libnodes app.c -o app

# Execute on all available cores of the current session
$ ./app
```

The number of cores that are used is controlled by running the application through the `taskset` command. For instance:

```
# Execute on cores 0, 1, 2 and 4
$ taskset -c 0-2,4 ./app
```

4.2 Debugging

By default, NODES is optimized to execute applications as efficiently as possible in terms of performance, and will assume that the application code is correct. Thus, it will not perform most runtime validity checks. To enable validity checks, users must compile NODES enabling debug options, by passing the `--enable-debug` flag at configure time. This will enable many internal validity checks that may be violated when the application code is incorrect.

To debug an application with a regular debugger, please compile its code with the regular debugging flags.

4.3 Generating ovni traces

NODES can generate execution traces with the ovni library, which generates lightweight binary traces, and it is possible to mix ovni-instrumented libraries together with an OmpSs-2 program and obtain a single coherent trace.

To enable the generation of ovni traces, NODES must be configured with the `--with-ovni` option. Once NODES has been built with ovni support, it is up to the user to enable it, as it is disabled by default. To enable ovni instrumentation, the `NODES_OVNI` environment variable must be set as follows: `export NODES_OVNI=1`.

The trace will be left in the `ovni/` directory, which can be transformed into a Paraver trace with the `ovniemu` utility. The Paraver configuration files (views) can be found in the `ovni/cfg` directory.

See the [ovni documentation](#) for more details.

NOS-V LIBRARY

This section describes how to use and configure the nOS-V library.

5.1 Linking with nOS-V

nOS-V applications can be linked by leveraging the `-lnosv` linker flag. The API is available in the `nosv.h` header.

5.2 Runtime configuration

nOS-V comes with an extensive list of tuning, configuration and customization features. The main way to change this configuration is through the use of the `nosv.toml` configuration file. When a nOS-V program calls `nosv_init()`, the library will search in the following order for a suitable configuration file:

1. The file specified through the `NOSV_CONFIG` environment variable
2. A file named `nosv.toml` in the current directory
3. The default configuration file, located in the installation directory

After the configuration file has loaded, the override configuration is parsed, which is a set of comma-separated TOML-like variables specified in the `NOSV_CONFIG_OVERRIDE` environment variable. This is an example using the override mechanism to change some configurations:

```
$ export NOSV_CONFIG_OVERRIDE="shared_memory.name=example,turbo.enabled=true"
$ ./nos-v-application
```

Configurations are described in the default `nosv.toml` file, which is reproduced here:

```
# Default nOS-V configuration file.
# Note that every process in the same nOS-V instance should load an identical
↪configuration file,
# otherwise, the behaviour is unspecified

# Shared memory configuration
[shared_memory]
    # Name of the shared memory section. This variable acts as an "id" for each nOS-V
    ↪instance
    # Two applications with the same shared memory name will share a nOS-V instance,
    ↪while different names
    # will result in independent nOS-V instances. Note that both the name of the shared
```

(continues on next page)

(continued from previous page)

```

↪memory as well as the
    # isolation level have to match for nOS-V instances to connect
    name = "nosv"
    # Isolation level of nOS-V. Possible values are:
    # "process" - Each process has a separate nOS-V instance
    # "user" - nOS-V instances are shared for the same user
    # "group" - nOS-V instances are shared for the same group
    # "public" - nOS-V instances are shared for the whole system
    isolation_level = "user"
    # Size of the shared memory section. Choose powers of two
    size = "2G"
    # Start of the shared memory mapping
    # This option should be memory address codified in hexadecimal (start with 0x)
    start = 0x0000200000000000

# Scheduler configuration
# These parameters allow to fine tune the scheduler's behaviour
[scheduler]
    # Number of logical CPUs mapped to a single scheduler SPSC for ready tasks
    # Minimum is 1, there is no maximum. Try to choose numbers that divide evenly the
↪number
    # of CPUs in the system.
    # A lower number will yield more scheduling throughput and is adequate when using
    # multiple task creator processes.
    # A higher number will yield lower scheduling latency and is adequate when using
↪one or few
    # task creator processes
    cpus_per_queue = 1
    # Number of tasks that are grabbed off of a queue when processing ready tasks.
    # Lowering the number may be better if there are very few ready tasks during
↪execution.
    # Increasing the batch may help if your CPU has bad single-thread performance.
    # In general, the default value of 64 should result in a good trade-off
    queue_batch = 64
    # Scheduler quantum in ns. This is a guideline for how long should we execute one
↪process' tasks
    # until we have to switch to the next one. However, as nOS-V tasks are not
↪preemptable, it isn't enforced.
    # This parameter is specially relevant for the nosv_schedpoint function, which will
↪try to schedule
    # each quantum ns.
    # A lower value will cause more inter-process context switches but may provide more
↪uniform progress,
    # while a higher value will minimize context switches but may stall applications
↪for longer
    quantum_ns = 20000000 # nanoseconds
    # Size of the queues that are used to place tasks into the scheduler
    # A good value should be a multiple of queue_batch
    in_queue_size = 256
    # Enable the immediate successor policy, which controls wether nOS-V honors the
↪NOSV_SUBMIT_IMMEDIATE
    # to execute a task immediately after the current one

```

(continues on next page)

(continued from previous page)

```

immediate_successor = true

# CPU Governor configuration
# Controls the policy that nOS-V follows to block idle CPUs to save energy and resources
[governor]
  # There is a choice between three different governor policies:
  # - hybrid: CPUs will spin for governor.spins before going to sleep when no work is
↳available
  # - idle: CPUs will sleep immediately when no work is available
  # - busy: CPUs will never sleep
  # In general, use idle when targeting minimum power usage, and busy when targeting
↳maximum performance
  # The default is hybrid as it provides a good balance between power and performance.
  policy = "hybrid"
  # Number of times a CPU will spin without sleeping in the "hybrid" policy.
  # When "idle" or "busy" are selected, this setting is ignored
  spins = 10000

# CPU Manager
[cpumanager]
  # CPU binding for the execution. The valid values are "inherit" which inherits the
↳cpu mask of the
  # first process that initializes the runtime, or a CPU mask in hexadecimal with the
↳format "0xFFFF"
  binding = "inherit"

# Affinity
[affinity]
  # Default affinity for all created tasks. This will be the default value for nosv
↳task affinity in the
  # current process. This setting is normally used in conjunction with "binding" and
↳makes it possible to
  # easily combine multiple processes in the same nOS-V instance.
  # Values can be: "all" or "cpu-X", "numa-X", "complex-X" where "X" is the physical
↳index of the resource.
  # For example, "numa-1" will set affinity of all tasks to the NUMA Node 1.
  default = "all"
  # Default affinity policy, either "strict" or "preferred". When the default
↳affinity is "all", this value
  # is not used.
  default_policy = "strict"
  # It is also possible to define the composition of the NUMA Nodes, which will
↳override the autodiscovered
  # composition that nOS-V can do through libnuma.
  # It should be defined as a 2-dimensional array, which has one row per NUMA node,
↳and each column is one
  # core that belongs to that NUMA Node. Use the system ids for the cores. Consider
↳the following example
  # in a system with 2 NUMA Nodes, each one with 4 cores:
  # numa_nodes = [[0, 1, 2, 3], [4, 5, 6, 7]]

```

(continues on next page)

```

# Debug options
[debug]
    # Dump all the configuration options nOS-V is running with, its final parsed values.
    ↪and the
    # path of the config file being used
    dump_config = false

# Hardware Counters configuration
[hwcounters]
    # Whether to print verbose information if a backend is enabled
    verbose = false
    # The enabled HWCOUNTER backends. Possible options: "papi", "none"
    backend = "none"
    # The list of PAPI counters to read. By default only "PAPI_TOT_INS" and "PAPI_TOT_
    ↪CYC"
    papi_events = [
        "PAPI_TOT_INS",
        "PAPI_TOT_CYC"
    ]

# Enabling turbo will cause nOS-V to set architecture-specific optimization flags on
    ↪every created
# or attached threads. In x86, this will cause the FTZ and DAZ flags of the SSE FPU to
    ↪be enabled,
# causing a significant performance increase in floating-point applications, but
    ↪disabling IEEE-754
# compatibility.
[turbo]
    enabled = false

# Monitoring capabilities and configuration.
[monitoring]
    enabled = false
    # Whether to print verbose information if monitoring is enabled
    verbose = false

# Instrumentation. Only applies if nOS-V was configured with --with-ovni
[instrumentation]
    # Select the instrumentation to be used. Current options are "none" or "ovni".
    version = "none"

```

5.3 Debugging

Although nOS-V APIs have error checking, it may be useful to build the library with debug symbols enabled in order to pinpoint bugs and issues. To do so, you can use the `--enable-debug` flag at configure time.

Furthermore, it is also possible to build nOS-V with Address Sanitizer enabled. In this mode, the runtime will assist ASAN by poisoning areas of the shared memory when they are not in use, which may be useful to find memory bugs in user programs that end up corrupting nOS-V memory. To enable ASAN mode, you can use the `--enable-asan` flag at configure time.

5.4 Generating ovni traces

nOS-V can generate execution traces with the ovni library, which generates lightweight binary traces. Traces can then be transformed into a Paraver trace. It is worth noting that multiple nOS-V programs running in the same or different instances can be combined to create a coherent system-wide trace.

To enable the generation of ovni traces, nOS-V must be configured with the `--with-ovni` option. Once nOS-V has been built with ovni support, it is up to the user to enable it, as it is disabled by default. To enable ovni instrumentation a user must set the `instrumentation.version` configuration option to `ovni`.

The trace will be left in the `ovni/` directory, which can be transformed into a Paraver trace with the `ovniemu` utility. The Paraver configuration files (views) can be found in the `ovni/cfg` directory.

See the [ovni documentation](#) for more details.

MERCURIUM LEGACY COMPILER OPTIONS

Important: Mercurium is the OmpSs-2 legacy compiler and is unsupported now. We recommend to use the LLVM-based compiler instead. Check *LLVM-based compiler*.

This section describes how to use the legacy Mercurium compiler to compile OmpSs-2 programs.

6.1 Mercurium drivers

The list of available drivers can be found here: https://github.com/bsc-pm/mcxx/blob/master/doc/md_pages/profiles.md

6.2 Common compilation flags

Usual flags like `-O`, `-O1`, `-O2`, `-O3`, `-D`, `-c`, `-o`, ... are recognized by Mercurium.

Almost every Mercurium-specific flag is of the form `--xxx`.

Mercurium drivers are deliberately compatible with `gcc`. This means that flags of the form `-fXXX`, `-mXXX` and `-Wxxx` are accepted and passed onto the backend compiler without interpretation by Mercurium drivers.

Warning: In GCC a flag of the form `-fXXX` is equivalent to a flag of the form `--XXX`. This is **not** the case in Mercurium.

6.2.1 Getting command line help

You can get a summary of all the flags accepted by Mercurium using `--help` with any of the drivers:

```
$ mcc --help
Usage: mcc options file [file..]
Options:
  -h, --help           Shows this help and quits
  --version            Shows version and quits
  --v, --verbose       Runs verbosely, displaying the programs
                       invoked by the compiler
  ...
```

6.2.2 Passing vendor-specific flags

While almost every gcc of the form `-fXXX` or `-mXXX` can be passed directly to a Mercurium driver, some other vendor-specific flags may not be well known or be misunderstood by Mercurium. When this happens, Mercurium has a generic way to pass parameters to the backend compiler and linker.

--Wn,<comma-separated-list-of-flags>

Passes comma-separated flags to the native compiler. These flags are used when Mercurium invokes the backend compiler to generate the object file (.o)

--Wl,<comma-separated-list-of-flags>

Passes comma-separated-flags to the linker. These flags are used when Mercurium invokes the linker

--Wp,<comma-separated-list-of-flags>

Passes comma-separated flags to the C/Fortran preprocessor. These flags are used when Mercurium invokes the preprocessor on a C or Fortran file.

These flags can be combined. Flags `--Wp,a` `--Wp,b` are equivalent to `--Wp,a,b`. Flag `--Wnp,a` is equivalent to `--Wn,a` `--Wp,a`

Important: Do not confuse `--Wl` and `--Wp` with the gcc similar flags `-Wl` and `-Wp` (note that gcc ones have a single -). The latter can be used with the former, as in `--Wl,-Wl,muldefs`. That said, Mercurium supports `-Wl` and `-Wp` directly, so `-Wl,muldefs` should be enough.

6.3 Compile OmpSs-2 programs

For OmpSs-2 programs that run in SMP or NUMA systems, you do not have to do anything. Just pick one of the drivers above.

Following is a very simple OmpSs-2 program in C:

```
/* test.c */
#include <stdio.h>

int main(int argc, char *argv[])
{
    int x = argc;
    #pragma oss task inout(x)
    {
        x++;
    }
    #pragma oss task in(x)
    {
        printf("argc + 1 == %d\n", x);
    }
    #pragma oss taskwait
    return 0;
}
```

Compile it using `mcc`:

```
$ mcc -o test --ompss-2 test.c
```

Important: Do not forget the flag `--ompss-2` otherwise your program will be compiled without parallel support.

And run it:

```
$ ./test
argc + 1 == 2
```

6.3.1 Compile OpenMP programs with OmpSs-2 support

By passing the `--openmp-compatibility` flag, Mercurium will support some constructs of OpenMP (e.g. `taskloop`, `parallel for`, etc.).

Warning: Experimental flag. There are no guarantees the generated code behaves as the OpenMP program.

6.4 Problems during compilation

While we put big efforts to make a reasonably robust compiler, you may encounter a bug or problem with Mercurium.

There are several errors of different nature that you may run into:

- Mercurium ends abnormally with an internal error telling you to open a ticket.
- Mercurium does not crash but gives an error on your input code and compilation stops, as if your code were not valid.
- Mercurium does not crash, but gives an error involving an `internal-source`.
- Mercurium generates wrong code and native compilation fails on an intermediate file.
- Mercurium forgets something in the generated code and linking fails.

6.4.1 How can you help us to solve the problem quicker?

In order for us to fix your problem we need the *preprocessed* file. If your program is C/C++ we need you to do:

1. Figure out the compilation command of the file that fails to compile. Make sure you can replicate the problem using that compilation command alone.
2. If your compilation command includes `-c`, replace it by `-E`. If it does not include `-c` simply add `-E`.
3. If your compilation command includes `-o file` (or `-o file.o`) replace it by `-o file.ii`. If it does not include `-o`, simply add `-o file.ii`.
4. Now run the compiler with this modified compilation command. It should have generated a `file.ii`.
5. These files are usually very large. Please compress them with `gzip` (or `bzip2` or any similar tool).

Send us an email to `pm-tools@bsc.es` with the error message you are experiencing and the (compressed) preprocessed file attached. If your program is Fortran just the input file may be enough, but you may have to add all the `INCLUDED` files and modules.

B

benchmarking
 Nanos6, 20

build procedure
 Nanos6, 4
 NODES, 8
 nos-v, 7

build requirements
 Nanos6, 4
 NODES, 8
 nos-v, 7

C

clang
 flags, 11
 OmpSs-2, 11

clang;, 8

cluster
 Nanos6, 29

compilation
 problems, 41

compile
 OmpSs-2, 39

CPUManager
 Nanos6, 19

CTF
 instrumentation, Nanos6, 21

D

debug
 instrumentation, Nanos6, 27
 instrumentation, NODES, 31
 instrumentation, nos-v, 36

dependencies
 Nanos6, 17

DLB
 Nanos6, 18

drivers
 Mercurium, 39

E

Extrae

instrumentation, Nanos6, 20

H

HWCounters
 Nanos6, 29

I

imcc, 39

imcxx, 39

imfc, 39

installation
 LLVM, 8
 Nanos6, 4
 NODES, 7
 nos-v, 6
 OmpSs-2, 3
 ovni, 3

installation; legacy compiler
 Mercurium, 9

instrumentation
 Nanos6 CTF, 21
 Nanos6 debug, 27
 Nanos6 Extrae, 20
 Nanos6 ovni, 20
 Nanos6 runtime information, 27
 Nanos6 verbose, 26
 NODES debug, 31
 NODES ovni, 31
 nos-v debug, 36
 nos-v ovni, 37

L

LLVM
 installation, 8

LLVM compiler
 problems, 13

LLVM;, 8

M

mcc, 39

mcxx, 39

Mercurium

- common flags, 39
- drivers, 39
- help, 39
- installation; legacy compiler, 9
- vendor-specific flags, 39

mfc, 39

N

Nanos6

- benchmarking, 20
- build procedure, 4
- build requirements, 4
- cluster, 29
- CPUManager, 19
- CTF instrumentation, 21
- debug instrumentation, 27
- dependencies, 17
- DLB, 18
- Extrae instrumentation, 20
- HWCounters, 29
- installation, 4
- optional build requirements, 4
- ovni instrumentation, 20
- runtime information instrumentation, 27
- scheduling, 18
- taskfor, 18
- Throttle, 28
- verbose instrumentation, 26

NODES

- build procedure, 8
- build requirements, 8
- debug instrumentation, 31
- installation, 7
- optional build requirements, 8
- ovni instrumentation, 31

nos-v

- build procedure, 7
- build requirements, 7
- debug instrumentation, 36
- installation, 6
- optional build requirements, 7
- ovni instrumentation, 37
- runtime library, 33

O

OmpSs-2

- clang, 11
- compile, 39
- installation, 3
- runtime system, 15, 31

optional build requirements

- Nanos6, 4
- NODES, 8
- nos-v, 7

ovni

- installation, 3
- instrumentation, Nanos6, 20
- instrumentation, NODES, 31
- instrumentation, nos-v, 37

P

problems

- compilation, 41
- LLVM compiler, 13

R

runtime information

- instrumentation, Nanos6, 27

runtime library

- nos-v, 33

runtime system

- OmpSs-2, 15, 31

S

scheduling

- Nanos6, 18

T

taskfor

- Nanos6, 18

Throttle

- Nanos6, 28

V

verbose

- instrumentation, Nanos6, 26

X

xlmcc, 39

xlmcxx, 39

xlmfc, 39