
OmpSs@FPGA User Guide

Release 2.1.1

BSC Programming Models

Apr 25, 2022

CONTENTS

1	Compile OmpSs@FPGA programs	3
1.1	Binaries	3
1.2	Bitstream	3
1.2.1	HW Instrumentation	4
1.2.2	Shared memory port	4
1.3	Boot Files	4
2	Running OmpSs@FPGA Programs	5
2.1	Nanos++ FPGA Architecture options	5
2.1.1	Nanos++ FPGA Architecture options	5
3	Create boot files for ultrascale	7
3.1	Prerequisites	7
3.1.1	Petalinux installation	7
3.2	Petalinux project setup	7
3.2.1	Unpack the bsp	8
3.2.2	[Optional] Fix known problems in AXIOM-ZU9EG-2016.3 project	8
3.2.3	[Optional] Modify the FSBL to have the Fallback system	8
3.2.4	Configure petalinux	9
3.2.5	Configure linux kernel	9
3.3	Petalinux (2016.3) build for a custom hdf	10
3.3.1	Add missing nodes to device tree	10
3.3.2	Build the Linux system	10
3.3.3	[Optional] Build PMU Firmware	11
3.3.4	Create BOOT.BIN file	11
3.4	Petalinux (2018.3) build for a custom hdf	11
3.4.1	Add missing nodes to device tree	11
3.4.2	Build the Linux system	12
3.4.3	Create BOOT.BIN file	12
4	FAQ: Frequently Asked Questions	13
4.1	What is OmpSs?	13
	Index	15

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. Barcelona Supercomputing Center will not assume any responsibility for errors or omissions in this document. Please send comments, corrections and/or suggestions to ompss-fpga-support at bsc.es. This document is provided for informational purposes only.

Note: There is a PDF version of this document at <http://pm.bsc.es/ftp/ompss-at-fpga/doc/user-guide-2.1.1/OmpSsFPGAUserGuide.pdf>

COMPILE OMPSS@FPGA PROGRAMS

To compile an OmpSs@FPGA program you should follow the general OmpSs compilation procedure using the Mercurium compiler. More information is provided in the OmpSs User Guide (<https://pm.bsc.es/ftp/ompss/doc/user-guide/compile-programs.html>). The following sections detail the specific options of Mercurium to generate the binaries, bitstream and boot files.

1.1 Binaries

There are two specific Mercurium front-ends for the FPGA devices:

- `fpgacxx` for C++ applications.
- `fpgacc` for C applications.

1.2 Bitstream

Note: Mercurium expects the Accelerator Integration Tool (AIT, formerly autoVivado) to be available on the PATH, if not the linker will fail. Moreover, AIT expects VivadoHLS and Vivado to be available in the PATH.

Warning: Sourcing the Vivado `settings.sh` file may break the cross-compilation toolchain. Instead, just add the directory of vivado binaries in the PATH.

To generate the bitstream, you should enable the bitstream generation in the Mercurium compiler (using the `-bitstream-generation` flag) and provide it the FPGA linker (aka AIT) flags with `--Wf` option. If the FPGA linker flags does not contain the `-b` (or `--board`) and `-n` (or `--name`) options, Mercurium will not launch AIT.

For example, to compile the `dotproduct` application, in debug mode, for the Zedboard, with a target frequency of 100Mhz, you can use the following command:

```
arm-linux-gnueabi-fpgacc --debug --ompss --bitstream-generation \  
  src/dotproduct.c -o dotproduct-d \  
  --Wf, "--board=zedboard,--clock=100,--name=dotproduct,--task_manager"
```

1.2.1 HW Instrumentation

You can use the `--instrument` (or `--instrumentation`) option of Mercurium to enable the HW instrumentation generation. The instrumentation can be generated and not used when running the application, but if you generate the bitstream without instrumentation support you will not be able to instrument the executions in the FPGA accelerators. Note that the application binary also has to be compiled with the `--instrument`

(or `--instrumentation`) option.

For example, the previous compilation command with the instrumentation available will be:

```
arm-linux-gnueabi-hf-fpgacc --instrument --ompss --bitstream-generation \  
src/dotproduct.c -o dotproduct-d \  
--Wf, "--board=zedboard,--clock=100,--name=dotproduct,--task_manager"
```

1.2.2 Shared memory port

By default, Mercurium generates an independent port to access the main memory for each task argument. Moreover, the bit-width of those ports equals to the argument data type width. This can result in a huge interconnection network when there are several task accelerators or they have several non-scalar arguments.

This behavior can be modified to generate unique shared port to access the main memory between all task arguments. This is achieved with the `fpga_memory_port_width` option of Mercurium which defines the desired bit-width of the shared port. The value must be a common multiple of the bit-widths for all task arguments.

The usage of the Mercurium variable to generate a 128 bit port in the previous dotproduct command will be like:

```
arm-linux-gnueabi-hf-fpgacc --ompss --bitstream-generation \  
src/dotproduct.c -o dotproduct-d \  
--variable=fpga_memory_port_width:128 \  
--Wf, "--board=zedboard,--clock=100,--name=dotproduct,--task_manager"
```

1.3 Boot Files

Some boards do not support loading the bitstream into the FPGA after the boot, therefore the boot files should be updated and the board rebooted. This step is not needed for the z7000 family of devices as the bitstream can be loaded after boot. AIT supports the generation of boot files for some boards but the step is disabled by default and should be enabled by hand.

First, you need to set the following environment variables:

- `PETALINUX_INSTALL`. Petalinux installation directory.
- `PETALINUX_BUILD`. Petalinux project directory. See *Create boot files for ultrascale* to have more information about how to setup a petalinux project build.

Then you can invoke AIT with the same options provided in `--Wf` and the following new options: `--from_step=boot` `--to_step=boot`. Also, you may directly add the `--to_step=boot` option in `--Wf` during the Mercurium launch.

RUNNING OMPSS@FPGA PROGRAMS

To run an OmpSs@FPGA program you should follow the general OmpSs run procedure. More information is provided in the OmpSs User Guide (<https://pm.bsc.es/ftp/ompss/doc/user-guide/run-programs.html>).

2.1 Nanos++ FPGA Architecture options

The Nanos++ behavior can be tuned with different environment options. They are summarized and briefly described in the Nanos++ help (`nanox --help`). The FPGA architecture options are also available at:

2.1.1 Nanos++ FPGA Architecture options

The Nanos++ behavior can be tuned with different environment options. They are summarized and briefly described in the Nanos++ help, the FPGA architecture section is shown below:

```
FPGA specific options
NX_ARGS options
  --fpga-alloc-align [=]<integer + suffix>
    FPGA allocation alignment (def: 16)
  --fpga-alloc-pool-size [=]<integer + suffix>
    FPGA device memory pool size (def: 512MB)
  --fpga-create-callback --no-fpga-create-callback
    Register the task creation callback during the plugin initialization (def: false,
↪ automatically enabled when needed)
  --fpga-create-callback-disable --no-fpga-create-callback-disable
    Disable the registration of the task creation callback to handle task creation,
↪ from the FPGA (def: false)
  --fpga-disable --no-fpga-disable
    Disable the support for FPGA accelerators and allocator
  --fpga-enable --no-fpga-enable
    Enable the support for FPGA accelerators and allocator
  --fpga-finish-task-burst [=]<integer>
    Max number of tasks to be finalized in a burst when limit is reached (def: 8)
  --fpga-helper-threads [=]<integer>
    Defines de number of helper threads managing fpga accelerators (def: 1)
  --fpga-hybrid-worker --no-fpga-hybrid-worker
    Allow FPGA helper thread to run smp tasks (def: enabled)
  --fpga-idle-callback --no-fpga-idle-callback
    Perform fpga operations using the IDLE event callback of Event Dispatcher (def:
↪ enabled)
  --fpga-max-pending-tasks [=]<integer>
    Number of tasks allowed to be pending finalization for an fpga accelerator (def:
↪ 4)
```

(continues on next page)

(continued from previous page)

```

--fpga-max-threads-callback [=]<integer>
    Max. number of threads concurrently working in the FPGA IDLE callback (def: 1)
--fpga-num [=]<integer>
    Defines de number of FPGA acceleratos to use (def: #accels from libxtasks)
Environment variables
NX_FPGA_ALLOC_ALIGN = <integer + suffix>
    FPGA allocation alignment (def: 16)
NX_FPGA_ALLOC_POOL_SIZE = <integer + suffix>
    FPGA device memory pool size (def: 512MB)
NX_FPGA_DISABLE = yes/no
    Disable the support for FPGA accelerators and allocator
NX_FPGA_ENABLE = yes/no
    Enable the support for FPGA accelerators and allocator
NX_FPGA_FINISH_TASK_BURST = <integer>
    Max number of tasks to be finalized in a burst when limit is reached (def: 8)
NX_FPGA_HELPER_THREADS = <integer>
    Defines de number of helper threads managing fpga accelerators (def: 1)
NX_FPGA_HYBRID_WORKER = yes/no
    Allow FPGA helper thread to run smp tasks (def: enabled)
NX_FPGA_MAX_PENDING_TASKS = <integer>
    Number of tasks allowed to be pending finalization for an fpga accelerator (def: ↵
↪4)
NX_FPGA_MAX_THREADS_CALLBACK = <integer>
    Max. number of threads concurrently working in the FPGA IDLE callback (def: 1)
NX_FPGA_NUM = <integer>
    Defines de number of FPGA acceleratos to use (def: #accels from libxtasks)

```

CREATE BOOT FILES FOR ULTRASCALE

The newer versions of the Accelerator Integration Tool (AIT, formerly autoVivado) support the automatic generation of boot files for some boards. This includes the steps in *Petalinux (2016.3) build for a custom hdf* or *Petalinux (2018.3) build for a custom hdf*, which are the ones repeated for every BOOT.BIN generation. The steps in *Petalinux project setup* are needed to setup the petalinux project build environment. Assuming that you have a valid petalinux build, you can use the ait functionality with the following points:

- Add the option `--to_step=boot` when calling ait.
- Provide the Petalinux installation and project directories using the following environment variables:
 - `PETALINUX_INSTALL` Petalinux installation directory.
 - `PETALINUX_BUILD` Petalinux project directory.

The following sections explain how to build a petalinux project and how to generate a BOOT.BIN using this project.

3.1 Prerequisites

- Petalinux installer (<https://www.xilinx.com/products/design-tools/embedded-software/petalinux-sdk.html>).
- Vivado handoff file (hdf) from a synthesized Vivado project.

3.1.1 Petalinux installation

Petalinux is installed running its auto-installer package:

```
./petalinux-v2016.3-final-installer.run
```

After installation, you should source the petalinux environment file. Usually, this needs to be done every time you want to run from a new terminal. Note that the petalinux settings may change the ARM cross compilers breaking the `OmpSs@FPGA tool-chain`.

```
source <petalinux install dir>/settings.sh
```

3.2 Petalinux project setup

The following steps should be executed once. After them, you will be able to build different boot files just using the AIT option or executing the steps in any of the following sections: *Petalinux (2016.3) build for a custom hdf* (for Petalinux 2016.3) or *Petalinux (2018.3) build for a custom hdf* (for Petalinux 2018.3).

3.2.1 Unpack the bsp

Unpack the bsp to create the petalinux project.

```
petalinux-create -t project -s <path to petalinux bsp>
```

3.2.2 [Optional] Fix known problems in AXIOM-ZU9EG-2016.3 project

Here are some patches for known problems:

Problem downloading Root FS

There is a problem during the BSP build when the scripts try to download the rootfs image from the Axiom webservers. The problem is that the download script checks that the server is alive with a `ping` and the AXIOM server is not responding to such type of traffic.

Patch file (axiom_bsp_patch_00.diff)

3.2.3 [Optional] Modify the FSBL to have the Fallback system

We developed a modification of Xilinx First Stage Boot Loader (FSBL) to support a fallback boot to a valid known BOOT.BIN file. More information in:

FSBL Fallback Mechanism

We developed a modification of Xilinx First Stage Boot Loader (FSBL) to support a fallback boot to a valid known BOOT.BIN file. The idea is to have a mechanism to allow the boards boot using a known BOOT.BIN file after a failed boot due to the usage of a wrong/corrupted BOOT.BIN. Fallback mechanism operational flow is the following:

- Does the `FLBK.TXT` file exist in the root of the `BOOT` partition?
- Yes. Enter in fallback mode and boot using the `FLBK.BIN` file
- No. Create the `FLBK.TXT` file and follow with a regular boot (usually using the `BOOT.BIN` file).
- The OS should mount the boot partition and remove the `FLBK.TXT` file after each boot (aka a successful boot).

FSBL Patch

FSBL Patch file (fsbl_patch_v010.diff)

In addition to the patch, the read-only filesystem protection must be disabled before the `BOOT.BIN` generation to allow the fallback mechanism work. This can be done by editing the `xparameters.h` file (`components/bootloader/zynqmp_fsbl/zynqmp_fsbl_bsp/psu_cortexa53_0/include/xparameters.h`) and removing any definition of `FILE_SYSTEM_READ_ONLY` pre-processor variable. Note that this header is re-generated/updated by petalinux tools in some steps. We need to ensure that it is properly edited when the bootloader is compiled, usually during the `petalinux-build` step.

System cleanup service

```
Systemctl service (remove-fsbl-flbk.service)
```

This service mounts the boot partition and removes the FSBL.TXT file created by the BSC FSBL. If the file is not removed, the next boot will use the FLBK.BIN file instead of BOOT.BIN. To install the service, copy the service file in the /etc/systemd/system/ folder and enable it with the following commands (they may require root privileges):

```
systemctl daemon-reload
systemctl enable remove-fsbl-flbk.service
systemctl start remove-fsbl-flbk.service
```

3.2.4 Configure petalinux

Run petalinux configuration. No changes need to be made to petalinux configuration, but this step has to be run.

```
export GIT_SSL_NO_VERIFY=1 #Ignore broken certificates
petalinux-config
```

After configuration this step, petalinux will download any needed files from external repositories.

3.2.5 Configure linux kernel

To enter the kernel configuration utility, run:

```
petalinux-config -c kernel
```

[Optional] Enable Xilinx DMA driver

Note: This step is only needed when the the use of DMA engines is desired.

Xilinx driver support has to be enabled in order to support Xilinx DMA engine devices. Usually, this is not needed as OmpSs@FPGA does not make use them to send tasks, neither information, between the host and the FPGA device. It can be enabled in: Device drivers → DMA Engines Support → Xilinx axi DMAS

Fix old kernels

In petalinux <2017, there is a known problem in the Xilinx DMA implementation. To fix it, download `xilinx_dma.c` and replace it in `<project_dir>/build/linux/kernel/download/linux-4.6.0-AXIOM-v2016/drivers/dma/xilinx/xilinx_dma.c`, when using a remote kernel, otherwise in `<petalinux install dir>/components/linux-kernel/xlnx-4.6/drivers/dma/xilinx/xilinx_dma.c`.

[Optional] Increase the CMA (Contiguous Memory Area)

You may want to increase the CMA size. It is used by Nanos++ as memory for the FPGA device copies. Its size can be set in: Device drivers → Generic Driver Options → DMA Contiguous Memory Allocator

3.3 Petalinux (2016.3) build for a custom hdf

Once petalinux project is setup, you can update it to contain a custom bitstream with your hardware. This steps can be repeated several times without executing again the steps in the *Petalinux project setup* section. Moreover, AIT supports the automatic execution of the following steps as explained in the beginning of this page.

First, you need to import the hardware description file (hdf) in the petalinux project. This is done executing the following command in the root directory of the petalinux project build.

```
petalinux-config --get-hw-description <path to application hdf file>
```

3.3.1 Add missing nodes to device tree

Some nodes should be added to the device tree before compiling it.

misc_clk_0

Edit the file `./subsystems/linux/configs/device-tree/pl.dtsi` to add or edit the node `misc_clk_0`. It should have the following contents (ensure that clock-frequency is correctly set):

```
misc_clk_0: misc_clk_0 {
    compatible = "fixed-clock";
    #clock-cells = <0>;
    clock-frequency = <200>;
};
```

pl_bsc.dtsi

AIT will generate a `pl_bsc.dtsi` file in the main Vivado project folder. This file contains the missing nodes in the `amba_pl` based on your application build. This file must be copied in `./subsystems/linux/configs/device-tree/` folder and included in `./subsystems/linux/configs/device-tree/system-conf.dtsi` file.

For example, it will be located in `test_ait/Vivado/test/` folder if the project name is `test`.

3.3.2 Build the Linux system

When the project is correctly updated, you can build it with the following commands:

```
petalinux-build
```

Error in fsbl compilation

In some cases, fsbl compilation triggered during the petalinux build can fail. This is due to a bad cleanup from previous compilation. In this cases, a complete fsbl cleanup and a new build must be performed. Note, that this extra cleanup may collision with the steps described in *[Optional] Modify the FSBL to have the Fallback system*.

```
petalinux-build -c bootloader -x mrproper
petalinux-build
```

3.3.3 [Optional] Build PMU Firmware

Run hsi (included in petalinux and Xilinx SDK).

```
hsi
```

Inside hsi run

```
set hwdsgn [open_hw_design <hardware.hdf>]
generate_app -hw $hwdsgn -os standalone -proc psu_pmu_0 -app zynqmp_pmufw -compile -
↳sw pmufw -dir <dir_for_new_app>
```

Warning: As of vivado 2016.3 pmu firmware breaks Trezz's TEBF0808 boot

3.3.4 Create BOOT.BIN file

```
petalinux-package --force --boot --fsbl images/linux/zynqmp_fsbl.elf --fpga <path to_
↳application bit file> --u-boot images/linux/u-boot.elf
cp BOOT.BIN images/linux/image.ub <path to boot partition>
```

When using PMU firmware, pmu binary has to be included in boot.bin file. To do so, add the `--pmufw <pmufw.elf>` argument to the `petalinux-package` command.

3.4 Petalinux (2018.3) build for a custom hdf

Once petalinux 2018.3 project is setup, you can update it to contain a custom bitstream with your hardware. This steps can be repeated several times without executing again the steps in the [Petalinux project setup](#) section. Moreover, AIT supports the automatic execution of the following steps as explained in the beginning of this page.

First, you need to import the hardware description file (hdf) in the petalinux project. This is done executing the following command in the root directory of the petalinux project build.

```
petalinux-config --get-hw-description <path to application hdf file>
```

3.4.1 Add missing nodes to device tree

Some nodes should be added to the device tree before compiling it.

pl_bsc.dtsi

AIT will generate a `pl_bsc.dtsi` file in the main Vivado project folder. This file contains the missing nodes in the `amba_pl` based on your application build. For example, it will be located in `test_ait/Vivado/test/` folder if the project name is `test`. The contents of such file must be placed at the end of `./project-spec/meta-user/recipes-bsp/device-tree/files/system-user.dtsi` file. Note that any remaining contents from a previous build must be removed before. The following command will append the `pl_bsc.dtsi` content at the end of `system-user.dtsi` file:

```
cat <path to vivado project>/pl_bsc.dtsi >>project-spec/meta-user/recipes-bsp/device-
↳tree/files/system-user.dtsi
```

3.4.2 Build the Linux system

When the project is correctly updated, you can build it with the following commands:

```
petalinux-build
```

3.4.3 Create BOOT.BIN file

```
petalinux-package --force --boot --fsbl images/linux/zynqmp_fsbl.elf --fpga <path to_  
↪application bit file> --u-boot images/linux/u-boot.elf  
cp BOOT.BIN images/linux/image.ub <path to boot partition>
```


FAQ: FREQUENTLY ASKED QUESTIONS

4.1 What is OmpSs?

OmpSs is an effort to integrate features from the StarSs programming model developed at BSC into a single programming model. In particular, our objective is to extend OpenMP with new directives to support asynchronous parallelism and heterogeneity (devices like GPUs, FPGAs). Then, OmpSs@FPGA is the extension of OmpSs tools to fully support FPGA devices.

Note: For more information about OmpSs programming model refer to <https://pm.bsc.es/ompss>

- genindex

INDEX

B

boot
 ultrascale, 6

C

compile
 OmpSs@FPGA, 1

F

FAQ, 12
 About OmpSs, 13
FSBL Fallback Mechanism, 8

N

Nanos++ FPGA Architecture options, 5

O

OmpSs@FPGA
 compile, 1
 running, 4

P

Problem downloading Root FS, 8

R

running
 OmpSs@FPGA, 4

U

ultrascale
 boot, 6