OSELAS.Support
OSELAS.Training
OSELAS.Development
OSELAS.Services

Quickstart Manual
OSELAS.BSP( )
Phytec phyCORE-MPC5200B-IO

PHYTEC

Pengutronix

Contents

Part I

OSELAS Quickstart for Phytec phyCORE-MPC5200B-IO

1 Getting a working Environment

1.1 Download Software Components

In order to follow this manual, some software archives are needed. There are several possibilities how to get these: either as part of an evaluation board package or by downloading them from the Pengutronix web site.

The central place for OSELAS related documentation is http://www.oselas.com. This website provides all required packages and documentation (at least for software components which are available to the public).

To build OSELAS.BSP-Phytec-phyCORE-12, the following archives have to be available on the development host:

- ptxdist-1.99.12.tgz
- ptxdist-1.99.12-patches.tgz
- OSELAS.BSP-Phytec-phyCORE-12.tar.gz
- OSELAS.Toolchain-1.99.3.2.tar.bz2

If they are not available on the development system yet, it is necessary to get them.

1.2 PTXdist Installation

The PTXdist build system can be used to create a root filesystem for embedded Linux devices. In order to start development with PTXdist it is necessary to install the software on the development system.

This chapter provides information about how to install and configure PTXdist on the development host.

1.2.1 Main Parts of PTXdist

The most important software component which is necessary to build an OSELAS.BSP() board support package is the `ptxdist` tool. So before starting any work we'll have to install PTXdist on the development host.

PTXdist consists of the following parts:

The `ptxdist` Program: `ptxdist` is installed on the development host during the installation process. `ptxdist` is called to trigger any action, like building a software packet, cleaning up the tree etc. Usually the `ptxdist` program is used in a *workspace* directory, which contains all project relevant files.

A Configuration System: The config system is used to customize a *configuration*, which contains information about which packages have to be built and which options are selected.

Patches: Due to the fact that some upstream packages are not bug free – especially with regard to cross compilation – it is often necessary to patch the original software. PTXdist contains a mechanism to automatically apply patches to packages. The patches are bundled into a separate archive. Nevertheless, they are necessary to build a working system.

Package Descriptions: For each software component there is a "recipe" file, specifying which actions have to be done to prepare and compile the software. Additionally, packages contain their configuration sniplet for the config system.

Toolchains: PTXdist does not come with a pre-built binary toolchain. Nevertheless, PTXdist itself is able to build toolchains, which are provided by the OSELAS.Toolchain() project. More in-deep information about the OSELAS.Toolchain() project can be found here: http://www.pengutronix.de/oselas/toolchain/index_en.html

Board Support Package This is an optional component, mostly shipped aside with a piece of hardware. There are various BSP available, some are generic, some are intended for a specific hardware.

1.2.2 Extracting the Sources

To install PTXdist, at least two archives have to be extracted:

ptxdist-1.99.12.tgz The PTXdist software itself.

ptxdist-1.99.12-patches.tgz All patches against upstream software packets (known as the 'patch repository').

ptxdist-1.99.12-projects.tgz Generic projects (optional), can be used as a starting point for self-built projects.

The PTXdist and patches packets have to be extracted into some temporary directory in order to be built before the installation, for example the `local/` directory in the user's home. If this directory does not exist, we have to create it and change into it:

```
$ cd
$ mkdir local
$ cd local
```

Next steps are to extract the archives:

```
$ tar -zxf ptxdist-1.99.12.tgz
$ tar -zxf ptxdist-1.99.12-patches.tgz
```

and if required the generic projects:

```
$ tar -zxf ptxdist-1.99.12-projects.tgz
```

If everything goes well, we now have a PTXdist-1.99.12 directory, so we can change into it:

```
$ cd ptxdist-1.99.12
$ ls -lF
total 487
drwxr-xr-x  13 jb users   1024 Mar 23 13:25 ./
drwxr-xr-x  22 jb users   3072 Mar 23 13:25 ../
-rw-r--r--   1 jb users    377 Feb 23 22:23 .gitignore
-rw-r--r--   1 jb users  18361 Apr 24  2003 COPYING
-rw-r--r--   1 jb users   3731 Mar 11 18:09 CREDITS
-rw-r--r--   1 jb users 115540 Mar  7 15:25 ChangeLog
-rw-r--r--   1 jb users     58 Apr 24  2003 INSTALL
-rw-r--r--   1 jb users   2246 Feb  9 14:29 Makefile.in
-rw-r--r--   1 jb users   4196 Jan 20 22:33 README
-rw-r--r--   1 jb users    691 Apr 26  2007 REVISION_POLICY
-rw-r--r--   1 jb users  54219 Mar 23 10:51 TODO
drwxr-xr-x   2 jb users   1024 Mar 23 11:27 autoconf/
```

```
-rwxr-xr-x   1 jb users      28 Jun 20  2006 autogen.sh*
drwxr-xr-x   2 jb users    1024 Mar 23 11:27 bin/
drwxr-xr-x   6 jb users    1024 Mar 23 11:27 config/
-rwxr-xr-x   1 jb users  226185 Mar 23 11:27 configure*
-rw-r--r--   1 jb users   12390 Mar 23 11:16 configure.ac
drwxr-xr-x   2 jb users    1024 Mar 23 11:27 debian/
drwxr-xr-x   8 jb users    1024 Mar 23 11:27 generic/
drwxr-xr-x 164 jb users    4096 Mar 23 11:27 patches/
drwxr-xr-x   2 jb users    1024 Mar 23 11:27 platforms/
drwxr-xr-x   4 jb users    1024 Mar 23 11:27 plugins/
drwxr-xr-x   6 jb users   30720 Mar 23 11:27 rules/
drwxr-xr-x   7 jb users    1024 Mar 23 11:27 scripts/
drwxr-xr-x   2 jb users    1024 Mar 23 11:27 tests/
```

1.2.3 Prerequisites

Before PTXdist can be installed it has to be checked if all necessary programs are installed on the development host. The configure script will stop if it discovers that something is missing.

The PTXdist installation is based on GNU autotools, so the first thing to be done now is to configure the packet:

```
$ ./configure
```

This will check your system for required components PTXdist relies on. If all required components are found the output ends with:

```
[...]
checking whether /usr/bin/patch will work... yes

configure: creating ./config.status
config.status: creating Makefile
config.status: creating scripts/ptxdist_version.sh
config.status: creating rules/ptxdist-version.in

ptxdist version 1.99.12 configured.
Using '/usr/local' for installation prefix.

Report bugs to ptxdist@pengutronix.de
```

Without further arguments PTXdist is configured to be installed into /usr/local, which is the standard location for user installed programs. To change the installation path to anything non-standard, we use the --prefix argument to the configure script. The --help option offers more information about what else can be changed for the installation process.

The installation paths are configured in a way that several PTXdist versions can be installed in parallel. So if an old version of PTXdist is already installed there is no need to remove it.

One of the most important tasks for the configure script is to find out if all the programs PTXdist depends on are already present on the development host. The script will stop with an error message in case something is missing. If this happens, the missing tools have to be installed from the distribution befor re-running the configure script.

When the configure script is finished successfully, we can now run

```
$ make
```

Pengutronix

All program parts are being compiled, and if there are no errors we can now install PTXdist into it's final location. In order to write to `/usr/local`, this step has to be performed as user *root*:

```
$ sudo make install
[enter password]
[...]
```

If we don't have root access to the machine it is also possible to install PTXdist into some other directory with the `--prefix` option. We need to take care that the `bin/` directory below the new installation dir is added to our `$PATH` environment variable (for example by exporting it in `~/.bashrc`).

The installation is now done, so the temporary folder may now be removed:

```
$ cd ../../
$ rm -fr local
```

1.2.4 Configuring PTXdist

When using PTXdist for the first time, some setup properties have to be configured. Two settings are the most important ones: Where to store the source packages and if a proxy must be used to gain access to the world wide web.

Run PTXdist's setup:

```
$ ptxdist setup
```

Due to PTXdist is working with sources only, it needs various source archives from the world wide web. If these archives are not present on our host, PTXdist starts the `wget` command to download them on demand.

Proxy Setup

To do so, an internet access is required. If this access is managed by a proxy `wget` command must be adviced to use it. PTXdist can be configured to advice the `wget` command automatically: Navigate to entry *Proxies* and enter the required addresses and ports to access the proxy in the form:

$$\text{<protocol>://<address>:<port>}$$

Source Archive Location

Whenever PTXdist downloads source archives it stores these archives in a project local manner. If we are working with more than one project, every project would download its own required archives. To share all source archives between all projects PTXdist can be configured to use only one archive directory for all projects it handles: Navigate to menu entry *Source Directory* and enter the path to the directory where PTXdist should store archives to share between projects.

Generic Project Location

If we already installed the generic projects we should also configure PTXdist to know this location. If we already did so, we can use the command `ptxdist projects` to get a list of available projects and `ptxdist clone` to get a local working copy of a shared generic project.

Pengutronix.

Navigate to menu entry *Project Searchpath* and enter the path to projects that can be used in such a way. Here we can configure more than one path, each part can be delemited by a colon. For example for PTXdist's generic projects and our own previous projects like this:

```
/usr/local/lib/ptxdist-1.99.12/projects:/office/my_projects/ptxdist
```

Leave the menu and store the configuration. PTXdist is now ready for use.

1.3 Toolchains

Before we can start building our first userland we need a cross toolchain. On Linux, toolchains are no monolithic beasts. Most parts of what we need to cross compile code for the embedded target comes from the *GNU Compiler Collection*, gcc. The gcc packet includes the compiler frontend, gcc, plus several backend tools (cc1, g++, ld etc.) which actually perform the different stages of the compile process. gcc does not contain the assembler, so we also need the *GNU Binutils package* which provides lowlevel stuff.

Cross compilers and tools are usually named like the corresponding host tool, but with a prefix – the *GNU target*. For example, the cross compilers for ARM and powerpc may look like

- `arm-softfloat-linux-gnu-gcc`
- `powerpc-unknown-linux-gnu-gcc`

With these compiler frontends we can convert e.g. a C program into binary code for specific machines. So for example if a C program is to be compiled natively, it works like this:

```
$ gcc test.c -o test
```

To build the same binary for the ARM architecture we have to use the cross compiler instead of the native one:

```
$ arm-softfloat-linux-gnu-gcc test.c -o test
```

Also part of what we consider to be the "toolchain" is the runtime library (libc, dynamic linker). All programs running on the embedded system are linked against the libc, which also offers the interface from user space functions to the kernel.

The compiler and libc are very tightly coupled components: the second stage compiler, which is used to build normal user space code, is being built against the libc itself. For example, if the target does not contain a hardware floating point unit, but the toolchain generates floating point code, it will fail. This is also the case when the toolchain builds code for i686 CPUs, whereas the target is i586.

So in order to make things working consistently it is necessary that the runtime libc is identical with the libc the compiler was built against.

PTXdist doesn't contain a pre-built binary toolchain. Remember that it's not a distribution but a development tool. But it can be used to build a toolchain for our target. Building the toolchain usually has only to be done once. It may be a good idea to do that over night, because it may take several hours, depending on the target architecture and development host power.

1.3.1 Using Existing Toolchains

If a toolchain is already installed which is known to be working, the toolchain building step with PTXdist may be omitted.

> The OSELAS.BoardSupport() Packages shipped for PTXdist have been tested with the OSE-LAS.Toolchains() built with the same PTXdist version. So if an external toolchain is being used which isn't known to be stable, a target may fail. Note that not all compiler versions and combinations work properly in a cross environment.

Every OSELAS.BoardSupport() Package checks for its OSELAS.Toolchain it's tested against, so using a different toolchain vendor requires an additional step:

Open the OSELAS.BoardSupport() Package menu with:

```
$ ptxdist platformconfig
```

and navigate to `architecture ---> toolchain` and `check for specific toolchain vendor`. Clear this entry to disable the toolchain vendor check.

1.3.2 Building a Toolchain

PTXdist handles toolchain building as a simple project, like all other projects, too. So we can download the OSELAS.Toolchain bundle and build the required toolchain for the OSELAS.BoardSupport() Package.

A PTXdist project generally allows to build into some project defined directory; all OSELAS.Toolchain projects that come with PTXdist are configured to use the standard installation paths mentioned below.

All OSELAS.Toolchain projects install their result into `/opt/OSELAS.Toolchain-1.99.3/`.

> Usually the `/opt` directory is not world writeable. So in order to build our OSELAS.Toolchain into that directory we need to use a root account to change the permissions. PTXdist detects this case and asks if we want to run `sudo` to do the job for us. Alternatively we can enter:
> `mkdir /opt/OSELAS.Toolchain-1.99.3`
> `chown <username> /opt/OSELAS.Toolchain-1.99.3`
> `chmod a+rwx /opt/OSELAS.Toolchain-1.99.3`.

We recommend to keep this installation path as PTXdist expects the toolchains at `/opt`. Whenever we go to select a platform in a project, PTXdist tries to find the right toolchain from data read from the platform configuration settings and a toolchain at `/opt` that matches to these settings. But that's for our convenience only. If we decide to install the toolchains at a different location, we still can use the *toolchain* parameter to define the toolchain to be used on a per project base.

1.3.3 Building the OSELAS.Toolchain for OSELAS.BSP-Phytec-phyCORE-12

To compile and install an OSELAS.Toolchain we have to extract the OSELAS.Toolchain archive, change into the new folder, configure the compiler in question and start the build.

The required compiler to build the OSELAS.BSP-Phytec-phyCORE-12 board support package is

```
powerpc-603e-linux-gnu_gcc-4.3.2_glibc-2.8_binutils-2.18_kernel-2.6.27-sanitized
```

So the steps to build this toolchain are:

```
$ tar xf OSELAS.Toolchain-1.99.3.2.tar.bz2
$ cd OSELAS.Toolchain-1.99.3.2
$ ptxdist select ptxconfigs/↵
        powerpc-603e-linux-gnu_gcc-4.3.2_glibc-2.8_binutils-2.18_kernel-2.6.27-sanitized.ptxconfig
$ ptxdist go
```

At this stage we have to go to our boss and tell him that it's probably time to go home for the day. Even on reasonably fast machines the time to build an OSELAS.Toolchain is something like around 30 minutes up to a few hours.

Measured times on different machines:

- Single Pentium 2.5 GHz, 2 GiB RAM: about 2 hours

- Turion ML-34, 2 GiB RAM: about 1 hour 30 minutes

- Dual Athlon 2.1 GHz, 2 GiB RAM: about 1 hour 20 minutes

- Dual Quad-Core-Pentium 1.8 GHz, 8 GiB RAM: about 25 minutes

Another possibility is to read the next chapters of this manual, to find out how to start a new project.

When the OSELAS.Toolchain project build is finished, PTXdist is ready for prime time and we can continue with our first project.

1.3.4 Protecting the Toolchain

All toolchain components are built with regular user permissions. In order to avoid accidential changes in the toolchain, the files should be set to read-only permissions after the installation has finished successfully. It is also possible to set the file ownership to root. This is an important step for reliability, so it is highly recommended.

1.3.5 Building Additional Toolchains

The OSELAS.Toolchain-1.99.3.2 bundle comes with various predefined toolchains. Refer the `ptxconfigs/` folder for other definitions. To build additional toolchains we only have to clean our current toolchain project, removing the current `selected_ptxconfig` link and creating a new one.

```
$ ptxdist clean
$ rm selected_ptxconfig
$ ptxdist select ptxconfigs/any_other_toolchain_def.ptxconfig
$ ptxdist go
```

All toolchains will be installed side by side architecture dependent into directory

`/opt/OSELAS.Toolchain-1.99.3/architecture_part`.

Different toolchains for the same architecture will be installed side by side version dependent into directory

`/opt/OSELAS.Toolchain-1.99.3/architecture_part/version_part`.

2 Building phyCORE-MPC5200B-IO's root filesystem

2.1 Extracting the Board Support Package

In order to work with a PTXdist based project we have to extract the archive first.

```
$ tar -zxf OSELAS.BSP-Phytec-phyCORE-12.tar.gz
$ cd OSELAS.BSP-Phytec-phyCORE-12
```

PTXdist is project centric, so now after changing into the new directory we have access to all valid components.

```
total 60
-rw-r--r--  1 jbe ptx   598 May 28 20:54 AUTHORS
-rw-r--r--  1 jbe ptx 10601 May 28 21:46 ChangeLog
-rw-r--r--  1 jbe ptx  1313 May 28 20:54 Kconfig
-rw-r--r--  1 jbe ptx   285 May 28 20:54 TODO
drwxr-xr-x 11 jbe ptx  4096 May 28 21:45 configs
drwxr-xr-x  3 jbe ptx  4096 May 28 21:19 documentation
drwxr-xr-x 12 jbe ptx  4096 May 28 20:53 local_src
drwxr-xr-x 14 jbe ptx  4096 May 28 20:54 patches
drwxr-xr-x  3 jbe ptx  4096 May 28 20:54 platforms
drwxr-xr-x  8 jbe ptx  4096 May 28 20:54 projectroot
drwxr-xr-x  3 jbe ptx  4096 May 28 20:54 protocols
drwxr-xr-x  4 jbe ptx  4096 May 28 20:54 rules
drwxr-xr-x  3 jbe ptx  4096 May 28 20:54 tests
```

Notes about some of the files and directories listed above:

ChangeLog Here you can read what has changed in this release. Note: This file does not always exist.

documentation If this BSP is one of our OSELAS BSPs, this directory contains the Quickstart you are currenly reading in.

configs A multiplatform BSP contains configurations for more than one target. This directory contains the platform configuration files.

projectroot Contains files and configuration for the target's runtime. A running GNU/Linux system uses many text files for runtime configuration. Most of the time the generic files from the PTXdist installation will fit the needs. But if not, customized files are located in this directory.

rules If something special is required to build the BSP for the target it is intended for, then this directory contains these additional rules.

patches If some special patches are required to build the BSP for this target, then this directory contains these patches on a per package basis.

tests Contains test scripts for automated target setup.

Pengutronix.

2.2 Selecting a Software Platform

First of all we have to select a software platform for the userland configuration. This step defines what kind of applications will be built for the hardware platform. The OSELAS.BSP-Phytec-phyCORE-12 comes with a predefined configuration we select in the following step:

```
$ ptxdist select configs/ptxconfig
info: selected ptxconfig:
      'configs/ptxconfig'
```

2.3 Selecting a Hardware Platform

Before we can build this BSP, we need to select one of the possible platforms to build for. In this case we want to build for the phyCORE-MPC5200B-IO:

```
$ ptxdist platform configs/phyCORE-MPC5200B-IO-1.99.12-1/platformconfig
info: selected platformconfig:
       'configs/phyCORE-MPC5200B-IO-1.99.12-1/platformconfig'
```

Note: If you have installed the OSELAS.Toolchain() at its default location, PTXdist should already have detected the proper toolchain while selecting the platform. In this case it will output:

```
found and using toolchain:
'/opt/OSELAS.Toolchain-1.99.3/powerpc-603e-linux-gnu/↵
    gcc-4.3.2-glibc-2.8-binutils-2.18-kernel-2.6.27-sanitized/bin'
```

If it fails you can continue to select the toolchain manually as mentioned in the next section. If this autodetection was successful, we can omit the steps of the section and continue to build the BSP.

In the unified OSELAS.BSP-Phytec-phyCORE-12, one included platform can use more userland features than another. For example platforms with graphic features will also build graphic support, but platforms sans display do not need it. To speed up compilation for specific platforms PTXdist provides collections, to reduce the amount of programs to be compiled for specific cases.

To reduce the package count for the phyCORE-MPC5200B-IO-1.99.12-1 run:

```
$ ptxdist collection configs/configs/collectionconfig-headless
info: selected collectionconfig:
       'configs/configs/collectionconfig-headless'
```

2.4 Selecting a Toolchain

If not automatically detected, the last step in selecting various configurations is to select the toolchain to be used to build everything for the target.

```
$ ptxdist toolchain /opt/OSELAS.Toolchain-1.99.3/powerpc-603e-linux-gnu/↵
      gcc-4.3.2-glibc-2.8-binutils-2.18-kernel-2.6.27-sanitized/bin
```

2.5 Building the Root Filesystem

Now everything is prepared for PTXdist to compile the BSP. Starting the engines is simply done with:

```
$ ptxdist go
```

PTXdist does now automatically find out from the `selected_ptxconfig` and `selected_platformconfig` files which packages belong to the project and starts compiling their *targetinstall* stages (that one that actually puts the compiled binaries into the root filesystem). While doing this, PTXdist finds out about all the dependencies between the packets and brings them into the correct order.

While the command `ptxdist go` is running we can watch it building all the different stages of a packet. In the end the final root filesystem for the target board can be found in the `platform-phyCORE-MPC5200B-IO/root/` directory and a bunch of *.ipk* packets in the `platform-phyCORE-MPC5200B-IO/packages/` directory, containing the single applications the root filesystem consists of.

2.6 Building an Image

After we have built a root filesystem, we can make an image, which can be flashed to the target device. To do this call

```
$ ptxdist images
```

PTXdist will then extract the content of priorly created *.ipk* packages to a temporary directory and generate an image out of it. PTXdist supports following image types:

- **hd.img:** contains grub bootloader, kernel and root files in a ext2 partition. Mostly used for X86 target systems.
- **root.jffs2:** root files inside a jffs2 filesystem.
- **uRamdisk:** a u-boot loadable Ramdisk
- **initrd.gz:** a traditional initrd RAM disk to be used as initrdramfs by the kernel
- **root.ext2:** root files inside a ext2 filesystem.
- **root.squashfs:** root files inside a squashfs filesystem.
- **root.tgz:** root files inside a plain gzip compressed tar ball.

The to be generated Image types and addtional options can be defined with

```
$ ptxdist platformconfig
```

Then select the submenu "image creation options". The generated image will be placed into `platform-phyCORE-MPC5200B-IO/images/`.

Only the content of the *.ipk* packages will be used to generate the image. This means that files which are put manually into the `platform-phyCORE-MPC5200B-IO/root/` will not be enclosed in the image. If custom files are needed for the target, install it with ptxdist.

Now that there is a root filesystem in our workspace we'll have to make it visible to the phyCORE-MPC5200B-IO. There are two possibilities to do this:

Figure 2.1: Booting the root filesystem, built with PTXdist, from the host via network and from flash.

1. Making the root filesystem persistent in the onboard media.
2. Booting from the development host, via network.

Figure 2.1 shows both methods. The main method used in the OSELAS.BSP-Phytec-phyCORE-12 BSP is to provide all needed components to run on the target itself. The Linux kernel and the root filesystem is persistent in the media the target features. This means the only connection needed is the nullmodem cable to see what is happening on our target. We call this method *standalone*.

The other method is to provide all needed components via network. In this case the development host is connected to the phyCORE-MPC5200B-IO with a serial nullmodem cable **and** via ethernet; the embedded board boots into the bootloader, then issues a TFTP request on the network and boots the kernel from the TFTP server on the host. Then, after decompressing the kernel into the RAM and starting it, the kernel mounts its root filesystem via NFS (Network File System) from the original location of the `platform-phyCORE-MPC5200B-IO/root/` directory in our PTXdist workspace.

The OSELAS.BSP-Phytec-phyCORE-12 provides both methods. The latter one is especially for development purposes, as it provides a very quick turnaround while testing the kernel and the root filesystem.

This chapter describes how to set up our target with features supported by PTXdist to simplify this challange.

2.7 Target Side Preparation

The phyCORE-MPC5200B-IO uses U-Boot as its bootloader. U-Boot can be customized with environment variables and scripts to support any boot constellation. OSELAS.BSP-Phytec-phyCORE-12 comes with a predefined environment setup to easily bring up the phyCORE-MPC5200B-IO.

Usually the environment doesn't have to be set manually on our target. PTXdist comes with an automated setup procedure to achieve a correct environment on the target.

Due to the fact that some of the values of these U-Boot environment variables must meet our local network environment and development host settings we have to define them prior to running the automated setup procedure.

Note: At this point of time it makes sense to check if the serial connection is already working, because it is essential for any further step we will do.

We can try to connect to the target with our favorite terminal application (`minicom` or `kermit` for example). With a powered target we identify the correct physical serial port and ensure that the communication is working. Make sure to leave this terminal application to unlock the serial port prior to the next steps.

To set up development host and target specific value settings, we run the command

```
$ cd OSELAS.BSP-Phytec-phyCORE-12
$ ptxdist boardsetup
```

We navigate to "Network Configuration" and replace the default settings with our local network settings. In the next step we also should check if the "Host's Serial Configuration" entries meet our local development host settings. Especially the "serial port" must correspond to our real physical connection.

When everything is set up, we can "Exit" the dialog and and save our new settings.

Now the command

```
$ ptxdist test setenv
```

will automatically set up a correct default environment on our phyCORE-MPC5200B-IO. We have to powercycle our target to make this step happen.
It should output lines like these when it was successful:

```
==============================
Please power on your board now!
==============================


Logging into U-Boot.....................OK
Setting new environment.................OK
Test finished successfully.
```

Note: If it fails, reading `platform-phyCORE-MPC5200B-IO/test.log` will give further information about why it has failed. Also extending the command line shown above by a `--debug` can help to see whats going wrong.

⚠️ Users reported this step could fail if the Linux system running PTXdistis a virtual machine as guest in an operating system from Redmont. In this case it seems at least one of the two OSes is eating up characters sent to the serial line. Pengutronix recommends running PTXdiston a real Linux system.

2.8 Stand-Alone Booting Linux

To use the the target standalone, the rootfs has to be made persistent in one of the onboard suported media of the phyCORE-MPC5200B-IO. The following sections describe the steps necessary to bring the rootfs into the onboard NOR type flash.

Only for preparation we need a network connection to the embedded board and a network aware bootloader which can fetch any data from a TFTP server.

After preparation is done, the phyCORE-MPC5200B-IO can work independently from the development host. We can "cut" the network (and serial cable) and the phyCORE-MPC5200B-IO will continue to work.

Pengutronix.

2.8.1 Development Host Preparations

On the development host a TFTP server has to be installed and configured. The exact method to do so is distribution specific; as the TFTP server is usually started by one of the inetd servers, the manual sections describing `inetd` or `xinetd` should be consulted.

Usually TFTP servers are using the `/tftpboot` directory to fetch files from, so if we want to push kernel images into this directory we have to make sure we are able to write there. As the access permissions are normally configured in a way to let only user **root** write to `/tftpboot` we have to change it. The boardsetup scripts coming with this BSP expect write permission in TFTP directory!

We can run a simple:

```
$ touch /tftpboot/my_file
```

to test if we have permissions to create files in this directory. If it fails we have to ask the administrator to grant these permissions.

Note: We must `/tftpboot` part of the command above with our local settings.

2.8.2 Preparations on the Embedded Board

To boot phyCORE-MPC5200B-IO stand-alone, anything needed to run a Linux system must be locally accessible. So at this point of time we must replace any current content in phyCORE-MPC5200B-IO's flash memory.

But first we must create the new root filesystem image prepared for its usage on the phyCORE-MPC5200B-IO:

```
$ cd OSELAS.BSP-Phytec-phyCORE-12
$ ptxdist images
```

To simplify this step, OSELAS.BSP-Phytec-phyCORE-12 comes with an automated setup procedure for this step. To use this procedure we run the command:

```
$ ptxdist test flash
```

Note: This command requires a serial and a network connection. The network connection can be cut after this step.

This command will automatically write a root filesystem to the correct flash partition on the phyCORE-MPC5200B-IO. It only works if we previously have set up the environment variables successfully (described at page 15). The command should output lines like this when it was successful:

```
==============================
Please power on your board now!
==============================

Logging into U-Boot....................OK
Flashing kernel........................OK
Flashing rootfs........................OK
Flashing oftree........................OK
Test finished successfully.
```

Note: If it fails, reading `platform-phyCORE-MPC5200B-IO/test.log` will give further information about why it has failed.

2.8.3 Booting the Embedded Board

To check that everything went successfully up to here, we can run the *boot* test.

```
$ ptxdist test boot
=============================
Please power on your board now!
=============================


Checking for U-Boot.....................OK
Checking for Kernel.....................OK
Checking for init.......................OK
Checking for login......................OK
Test finished successfully.
```

This will check if the environment settings and flash partitioning are working as expected, so the target comes up in stand-alone mode up to the login prompt.

Note: If it fails, reading `platform-phyCORE-MPC5200B-IO/test.log` will give further information about why it has failed.

After the next reset or powercycle of the board, it should boot the kernel from the flash, start it and mount the root filesystem also from flash.

Note: The default login account is `root` with an empty password.

2.9 Remote-Booting Linux

The next method we want to try after building a root filesystem is the network-remote boot variant. This method is especially intended for development as everything related to the root filesystem happens on the host only. It's the fastest way in a phase of a project, where things are changing frequently. Any change made in the local `platform-phyCORE-MPC5200B-IO/root/` directory simply "appears" on the embedded device immediately.

All we need is a network interface on the embedded board and a network aware bootloader which can fetch the kernel from a TFTP server.

2.9.1 Development Host Preparations

If we already have booted the phyCORE-MPC5200B-IO locally (as described in the previous section), all of the development host preparations are done.

If not, then a TFTP server has to be installed and configured on the development host. The exact method of doing this is distribution specific; as the TFTP server is usually started by one of the `inetd` servers, the manual sections describing `inetd` or `xinetd` should be consulted.

Usually TFTP servers are using the `/tftpboot` directory to fetch files from, so if we want to push data files to this directory, we have to make sure we are able to write there. As the access permissions are normally configured in a way to let only user **root** write to `/tftpboot` we have to change it. If we don't want to change the permission or if its disallowed to change anything, the `sudo` command may help.

```
$ sudo cp platform-phyCORE-MPC5200B-IO/images/linuximage /tftpboot/uImage-pcm032
```

The NFS server is not restricted to a certain filesystem location, so all we have to do on most distributions is to modify the file /etc/exports and export our root filesystem to the embedded network. In this example file the whole work directory is exported, and the "lab network" between the development host is 192.168.23.0, so the IP addresses have to be adapted to the local needs:

```
/home/<user>/work 192.168.23.0/255.255.255.0(rw,no_root_squash,sync)
```

Note: Replace <user> with your home directory name.

2.9.2 Preparations on the Embedded Board

We already provided the phyCORE-MPC5200B-IO with the default environment at page 15. So there is no additional preparation required here.

2.9.3 Booting the Embedded Board

The default environment settings coming with the OSELAS.BSP-Phytec-phyCORE-12 has the possibility to boot from the internal flash or from the network. The definition what should happen after power on is made with the environment variable bootcmd. The default setting is run bcmd_flash and will boot from flash.

To change this behavior we have to change the value of the bootcmd environment variable to run bcmd_net.

```
uboot> setenv bootcmd 'run bcmd_net'
uboot> saveenv
```

The next time the target will start it will use the network based booting mechanism.

3 Accessing Peripherals

The following sections provide an overview of the supported hardware components and their corresponding operating system drivers. Further changes can be ported on demand of the customer.

Phytec's phyCORE-MPC5200B-IO starter kit consists of the following individual boards:

1. The phyCORE-MPC5200B-IO module itself, containing the MPC5200B processor, RAM, flash and several other peripherals.
2. The starter kit baseboard (PCM973).

To achieve maximum software re-use, the Linux kernel offers a sophisticated infrastructure, layering software components into board specific parts. The OSELAS.BSP( ) tries to modularize the kit features as far as possible; that means that when a customized baseboards or even customer specific module is developed, most of the software support can be re-used without error prone copy-and-paste. So the kernel code corresponding to the boards above can be found in the kernel source tree at:

1. `arch/powerpc/platforms/mpc5200_simple.c` for the processor module
2. `arch/powerpc/boot/dts/pcm032.dts` for the platform description

In fact, software re-use is one of the most important features of the Linux kernel and especially of the PowerPC port, which always had to fight with an insane number of possibilities of the System-on-Chip CPUs.

⚠️ Note that the huge variety of possibilities offered by the phyCORE modules makes it difficult to have a completely generic implementation on the operating system side. Nevertheless, the OSELAS.BSP( ) can easily be adapted to customer specific variants. In case of interest, contact the Pengutronix support (support@pengutronix.de) and ask for a dedicated offer.

The following sections provide an overview of the supported hardware components and their operating system drivers.

3.1 NOR Flash

Linux offers the Memory Technology Devices Interface (MTD) to access low level flash chips, directly connected to a SoC CPU.

Modern kernels offer a method to define flash partitions on the kernel command line, using the `mtdparts` command line argument:

```
mtdparts=physmap-flash.0:256k(uboot)ro,128k(ubootenv),2M(kernel),-(root)
```

This line, for example, specifies several partitions with their size and name which can be used as `/dev/mtd0`, `/dev/mtd1` etc. from Linux. Additionally, this argument is also understood by reasonably new U-Boot bootloaders, so if there is any need to change the partitioning layout, the U-Boot environment is the only place where the layout has to be changed.

From userspace the NOR flash partitions can be accessed as

- /dev/mtdblock0 (e.g. U-Boot partition)
- /dev/mtdblock1 (e.g. U-Boot environment partition)
- /dev/mtdblock2 (e.g. Kernel partition)
- /dev/mtdblock3 (e.g. Linux rootfs partition)

Note: This is an example only. The partitioning on our phyCORE-MPC5200B-IO target can differ from this layout.

Only the `/dev/mtdblock3` on the phyCORE-MPC5200B-IO has a filesystem, so the other partitions cannot be mounted into the rootfs. The only way to access them is by pushing a prepared flash image into the corresponding `/dev/mtd` device node.

3.2 Serial Units

The MPC5200B SoC supports up to 6 so called PSC units with configureable serial protocolls. On the phyCORE-MPC5200B-IO two PSC units are configured to work as serial interfaces: PSC3 and PSC6. They are accessible as regular serial TTYs at `/dev/ttyPSC0` and `/dev/ttyPSC1`.

Note: /dev/ttyPSC0 (PSC3) is used as the default serial console.

3.3 SRAM Memory

The phyCORE-MPC5200B-IO is shipped with a 2 MiB SRAM. We can use it as plain memory or we can put a filesystem on top of it and use it as a regular part of the root filesystem.

- /dev/mtdblock7 is the full SRAM partition.

To create a filesystem we type:

```
~# mkfs.minix -n 30 /dev/mtdblock7
704 inodes
2048 blocks
Firstdatazone=26 (26)
Zonesize=1024
Maxsize=268966912
```

Note: The output of the `mkfs.minix` my differ due to different SRAM sizes.

And to mount it:

```
~# mount -t minix /dev/mtdblock7 /mnt
```

3.4 CAN Bus

The phyCORE-MPC5200B-IO provides a CAN feature, which is supported by drivers using the (currently work-in-progress) proposed Linux standard CAN framework "Socket-CAN". Using this framework, CAN interfaces can be programmed with the BSD socket API.

Pengutronix

Configuration happens within the script `/etc/network/can-pre-up`. This script will be called when `/etc/init.d/networking` is running at system start up. To change default used bitrates on the target change the variables **CAN_0_BITRATE** and/or **CAN_1_BITRATE** in `/etc/network/can-pre-up`.

For a persistent change of the default bitrates change the local `projectroot/etc/network/can-pre-up` instead and rebuild the BSP.

> ⚠️ The Socket-CAN API is still work in progress and was submitted to the upstream kernel maintainers in part only.

3.4.1 About Socket-CAN

The CAN (Controller Area Network[1]) bus offers a low-bandwidth, prioritised message fieldbus for communication between microcontrollers. Unfortunately, CAN was not designed with the ISO/OSI layer model in mind, so most CAN APIs available throughout the industry don't support a clean separation between the different logical protocol layers, like for example known from ethernet.

The *Socket-CAN* framework for Linux extends the BSD socket API concept towards CAN bus. It consists of

- a core part (candev.ko)
- chip drivers (e. g. mscan, sja1000 etc.)

So in order to start working with CAN interfaces we'll have to make sure all necessary drivers are loaded.

Starting and Configuring Interfaces from the Command Line

If all drivers are present in the kernel, "ifconfig -a" shows which network interfaces are available; as Socket-CAN chip interfaces are normal Linux network devices (with some additional features special to CAN), not only the ethernet devices can be observed but also CAN ports.

For this example, we are only interested in the first CAN port, so the information for can0 looks like

```
# ifconfig can0
can0 Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
     UP RUNNING NOARP  MTU:16  Metric:1
     RX packets:35948 errors:0 dropped:0 overruns:0 frame:0
     TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
     collisions:0 txqueuelen:10000
     RX bytes:243744 (238.0 KiB)  TX bytes:2 (2.0 B)
     Interrupt:145 Base address:0x900
```

The output contains the usual parameters also shown for ethernet interfaces, so not all of these are necessarily relevant for CAN (for example the MAC address). These parameters contain useful information:

---

[1]ISO 11898/11519

Field	Description
cano	Interface Name
NOARP	CAN cannot use ARP protocol
MTU	Maximum Transfer Unit, always 8
RX packets	Number of Received Packets
TX packets	Number of Transmitted Packets
RX bytes	Number of Received Bytes
TX bytes	Number of Transmitted Bytes
errors...	Bus Error Statistics

Table 3.1: CAN interface information

Inferfaces shown by the "ifconfig -a" command can be configured with `canconfig`. This command adds CAN specific configuration possibilities for network interfaces, similar to for example `iwconfig` for wireless ethernet cards.

The baudrate for cano can now be changed:

```
# canconfig can0 bitrate 250000
```

and the interface is started with

```
# ifconfig can0 up
```

Using the CAN Interfaces from the Command Line

After successfully configuring the local CAN interface and attaching some kind of CAN devices to this physical bus, we can test this connection with command line tools.

The tools `cansend` and `candump` are dedicated to this purpose.

To send a simple CAN message with ID 0x20 and one data byte of value 0xAA just enter:

```
# cansend can0 --identifier=0x20 0xAA
```

To receive CAN messages run the `candump` command:

```
# candump can0
interface = can0, family = 29, type = 3, proto = 0
<0x020> [1] aa
```

The output of `candump` shown in this example was the result of running the `cansend` example above on a different machine.

See cansend's and candump's manual pages for further information about using and options.

3.5 Network

The phyCORE-MPC5200B-IO module features ethernet OnChip, which is being used to provide the `eth0` network interface. The interface offers a standard Linux network port which can be programmed using the BSD socket interface.

3.6 USB Host Controller Unit

The phyCORE-MPC5200B-IO supports a standard OHCI Rev. 1.0a compliant host controller onboard for low and full speed connections. Up to two ports are supported by this CPU.

Note: The physical layer is not part of the phyCORE-MPC5200B-IO, so it depends on the base board, how many ports can be used.

The PCM973 baseboard only supports one working USB connector (the lower one). The second connector (the upper one) acts as power supply only (+5V up to 500mA). Refer to the hardware manual for further details.

3.7 I²C Master

The MPC5200B processor based phyCORE-MPC5200B-IO supports a dedicated I²C controller onchip. The kernel supports this controller as a master controller.

Additional I²C device drivers can use the standard I²C device API to gain access to their devices through this master controller. For further information about the I²C framework see `Documentation/i2c` in the kernel source tree.

3.7.1 I²C Realtime Clock RTC8564

Due to the Real Time Clock framework of the kernel the RTC8564 clock chip can be accessed using the same tools as for any other real time clock.

Date and time can be manipulated with the `hwclock` tool, using the `-w` (systohc) and `-s` (hctosys) options. For more information about this tool refer to the manpage of hwclock.

OSELAS.BSP-Phytec-phyCORE-12 tries to set up the date at system startup. If there was a powerfail `hwclock` will state:

```
pcf8564 1-0051: low voltage detected, date/time is not reliable.
pcf8564 1-0051: retrieved date/time is not valid.
```

In this case set the date manually (see `man date`) and run `hwclock -w -u` to store the new date into the RTC8564.

3.7.2 I²C Device 24W32

This device is a 4 kiB non-volatile memory for general purpose usage.

This type of memory is accessible through the sysfs filesystem. To read the EEPROM content simply `open()` the entry `/sys/bus/i2c/devices/1-0052/eeprom` and use `fseek()` and `read()` to get the values.

3.8 GPIO

Like most modern System-on-Chip CPUs, the MPC5200B has numerous GPIO pins. Some of them are unaccessible for the userspace as Linux drivers use them internally. Others are also used by drivers but are exposed to userspace via sysfs. Finally, the remaining GPIOs can be requested for custom use by userspace, also via sysfs.

Refer to the in-kernel documentation `Documentation/gpio.txt` for complete details how to use the sysfs-interface for manually exporting GPIOs.

3.9 Watchdog

The internal watchdog will be activated when an application opens the device /dev/watchdog. Default timeout is 60 seconds. An application must periodically write to this device. It does not matter what is written. Just the interval between these writes should not exceed the timeout value, otherwise the system will be reset.

For testing the hardware, there is also a shell command which can do the triggering:

```
~# watchdog -t <trigger-time-in-seconds> /dev/watchdog
```

This command is part of the busybox shell environment. Keep in mind, that it should only be used for testing. If the watchdog gets fed by it, a crash of the real application will go unnoticed.

3.10 ATA IDE/CompactFlash Card

The phyCORE-MPC5200B-IO can handle IDE/ATA devices up to PIO mode 4. MPC5200B's ATA driver is a so called *pata* driver. It enables access to these devices through the SCSI layer. All connected ATA/IDE devices will occur as /dev/sd* device nodes.

3.11 FPGA Support

The phyCORE-MPC5200B-IO is shipped with an Altera FPGA of type **Cyclone II EP2C8F256C8N**. This FPGA is a general purpose device with no special dedication when shipped. It's up to you to blow life into it with your own firmware, for example to do some high speed signal processing. The OSELAS.BSP-Phytec-phyCORE-12 provides a mechanism to load the FPGA firmware while the whole system is already running.

3.11.1 General

This BSP is shipped with a generic FPGA firmware download driver. At runtime its not loaded by default. Instead its up to the user to start the download manually.

Putting the mechanism to work is as easy as loading the module with parameter firmware=<filename.rbf>. The firmware must be in RBF (Altera *Raw Binary Format*) and stored in the place where the firmware agent of our embedded system will expect it. (It depends on the configuration of our udev daemon. Usually it should be /lib/firmware/)

In practice, the module is loaded like in this example: ~# modprobe fpga firmware=my_own_firmware.rbf

The firmware download driver assumes specific hardware connections between processor and FPGA which are shown in the following table. The user does not need to take care about these when using the standard hardware.

⚠️ A few data lines are shared with the PSC unit 6 running as a UART. From the system's view it occurs as /dev/ttyPSC1. While downloading the firmware into the FPGA this UART cannot transfer data. After downloading it can act as an UART again.

Additionaly GPIO6 is used to switch the PSC6 lines between UART and FPGA mode.

If we want to use this download feauture we must ensure the follwing bridges are set:

See phyCORE-MPC5200B-IO's datasheet for FPGA's pin assignment.

FPGA Pin	MPC5200B Pin	Pin Loc.	Pin Name	Pin Conf.
nCONFIG	UART6_CTS_TTL	PSC6_1	GPIO_WKUP_5	out
CONFIG_DONE	UART6_RXD_TTL	PSC6_0	GPIO_WKUP_4	in
nSTATUS	GPIO7		GPIO_WKUP_7	in
DCLK	UART6_RTS_TTL	PSC6_3	GPIO_IRDA_1	out
DATA0	UART6_TXD_TTL	PSC6_2	GPIO_IRDA_0	out

Table 3.2: FPGA pin connections

Jumper	shorten	function
J8	1-2	FPGA passive mode
J11	1-2	FPGA control line routing

Table 3.3: FPGA download enabling

3.11.2 Demo

The BSP is shipped with a simple demo. To download it into the FPGA you should enter:

```
~# modeprobe fpga firmware=MPC_to_WBSlave.rbf
```

It supports five 32 bit registers at the baseaddess CS3 (chip select) is using.

- Register 0 at CS3 baseaddress + 0x00

- Register 1 at CS3 baseaddress + 0x04

- Register 2 at CS3 baseaddress + 0x08

- Register 3 at CS3 baseaddress + 0x0C

- Register 4 at CS3 baseaddress + 0x10 (mirrored up to the end of CS3 address space)

In the first step to access the FPGA registers we do not need a device driver. Due to PowerPC architecture maps devices in memory we can use a small tool called `memedit` for testing.

The FPGA is connected to CS3 und CS4 of the processor's Local Plus Bus. To get the physical mapping we can run:

```
~# cat /proc/iomem
80000000-9fffffff : /pci@f0000d00
a0000000-afffffff : /pci@f0000d00
f0000600-f000060f : mpc5200_wdt
f0000900-f000097f : can
  f0000900-f000097f : mpc52xx-mscan
f0000980-f00009ff : can
  f0000980-f00009ff : mpc52xx-mscan
f0001000-f00010fe : ohci_hcd
f0001200-f0001283 : bestcomm-core
f0002400-f0002457 : mpc52xx_psc_uart
f0002c00-f0002c57 : mpc52xx_psc_uart
f0003000-f00033ff : mpc52xx-fec
f0003a00-f0003a7f : mpc52xx_ata
```

```
f0008000-f000bfff : bestcomm-core
fbe00000-fbffffff : sram
fe000000-ffffffff : fe000000.flash
```

This areas of interest are here the area#1 and area#2. They are reserved for the FPGA access. area#1 is connected to CS3, area#2 is connected to CS4.

To gain access to this areas we are using the memedit tool on the target:

```
~# memedit /dev/mem
```

This tool is an interactive one, like the command line. So the next step is to enter a command to map the physical memory space into the memory of the memedit process.

After we know FPGA's physical mapping, we now can map this area to gain access to the FPGA itself. When our mapping of CS3/4 differs we have to enter other address values here. The default should be:

```
-> map 0xf9e00000 0x200000
```

We now can read and write into this area and - if the FPGA is ready to work - we can access its registers.

```
-> md 0x0 0x20
00000000:  00000000 00000000 00000000 00000000 ................
00000010:  00000000 00000000 00000000 00000000 ................
```

After system reset everything should be zero. We do some writings to check the registers:

```
-> mm 0x0 0x10
new values:
0x00000000: 00000010
-> mm 0x4 0x20
new values:
0x00000004: 00000020
-> mm 0x8 0x30
new values:
0x00000008: 00000030
-> mm 0xC 0x40
new values:
0x0000000c: 00000040
-> mm 0x10 0x50
new values:
0x00000010: 00000050
```

And try to read back the written values:

```
-> md 0x0 0x20
00000000:  00000010 00000020 00000030 00000040 ....... ...0...@
00000010:  00000050 00000050 00000050 00000050 ...P...P...P...P
```

We see the described behaviour of FPGA's default firmware. Four single registers at offset 0x0,0x4,0x8 and 0xC, and a mirrored register until the end of FPGA's area starting at offset 0x10.

4 Special Notes

4.1 Analysing the CAN Bus Data Transfer

The OSELAS.BSP-Phytec-phyCORE-12 BSP comes with the standard *pcap* library and `tcpdump` tool. Both are capable of analyzing CAN data transfer wich includes time stamping.

We set up the CAN interface(s) as usual and use it in our application. With `tcpdump` we can sniff at any point of time the data transferred on the CAN line.

To do so, we simply start tcpdump:

```
# tcpdump -i can0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on can0, link-type LINUX_CAN (Linux CAN), capture size 68 bytes
```

Whenever there is any traffic on the line, `tcpdump` will log it to stdout. We will generate some traffic by using the `cansend` command:

```
# cansend can0 -i 0x12 0x0f 0xf0 0x10 0x01
```

For this data, `tcpdump` will output:

```
00:15:52.482066 CAN Out ID:00000012 PL_LEN:4 PAYLOAD: 0x0f 0xf0 0x10 0x01
```

The log `tcpdump` generates consist of six fields:

1. `00:15:52.482066` is the timestamp this data was on the line. Its format is HH:MM:SS:TTTTTT, with TTTTTT as second's fraction

2. `CAN` interface type

3. `Out` message'S data direction on this interface

4. `ID:00000012` CAN message ID

5. `PL_LEN:4` byte count of message data

6. `PAYLOAD: 0x0f 0xf0 0x10 0x01` the payload data

Some notes:

- The data direction field could be `Out` or `In`
- The CAN message ID encodes some additional info into higher bit values:

    - Bit 31 encodes an extended frame. If this bit is set, an extended message frame was on the line
    - Bit 30 encodes an RTR frame. If this bit is set, a remote transmission message frame was on the line

    The message ID resides in the lower bits of this field

- The PAYLOAD field could be empty, when there were no data elements in the message

Pengutronix

4.2 Realtimetests

The aim of this chapter is to determine the reactivity of the phyCORE MPC5200B Module using the realtime patches of the linux kernel.

4.2.1 Methods and used tools

The leading feature of a realtime operating system is the predictability of its time response. The latency between a task being scheduled and its real dispatch on the target system gives the main indication of the efficiency of the system in a realtime environment. While the latency is measured, the system has to handle different load scenarios with a high system load.

For the determination of its efficiency, the tool cyclictest was used. The tool determines the latency between the scheduled starting time and the real dispatch of one or more threads.

The measurement was repeated with three load scenarios. These load scenarios are chosen with following requirements:

- the usual load in the default usecase

- the usual bottleneck capabilities of a realtime os

With this assumption the following loadscenarios are chosen

- load of the schedulers: The tool hackbench is used. Hackbench is generating a defined number of grouped processes which, dependent on the group, are communicating over a defined socket.

- load of the processor cache and the memory interface: The tool lat_mem_rd of the benchmark toolkit lmbench is used. With its reads and writes of the memory contents, the processor cache and memomory interfaces are constantly claimed.

- network I/O load: A host computer in the same network as the tested target sends icmp packets in very short intervals to the tested target. (flood ping)

- memory I/O load: The flash memory of the target is written and read in a constant loop with random data.

4.2.2 Test platforms and environment

The used kernel and root filesystem for the platform are generated with the version 12 of the OSELAS.BSP-phyCORE BSP. The linux kernel ist supplemented with the preempt-rt patch to achieve the full realtime benefit. When the measurements were made the up-to-date kernel version of the rt-preempt patches was 2.6.29.5-rt22. Therefore the kernel version in the BSP was adapted to the 2.6.29 kernel.

The test platform starts its kernel out of the flash memory. The root filesystem is also flashed onto the onboard flash memory.

Besides the test applications, the following userspace daemons are running on the target:

- udevd

- syslogd

- sshd

- telnetd

- pure-ftpd (SERVER)

4.2.3 Test preparation and execution

The testcycles lasted – depending on the load scenario – 30 minutes. After the determination of the minimum-, average- and maximum latency, the histogram of the statistical distribution of each measurement results between the minimum and maximum measurement results was established.

Before the tests could be executed, the target system needed to be optimised for the real time operation. Therefore the rt process scheduling time limit had to be disabled. This mechanism limits the scheduling time of the realtime processes and ensures the addressability of the system also with a CPU load of 100%. For the testcases this mechanism is more constrictive then constructive and it generates big deviations on high load. The mechanism can be controlled with the proc Interface. A value of -1 turns this mechanism of.

```
# echo -1 > /proc/sys/kernel/sched_rt_runtime_us
```

The measurement on the target is started with one thread of the priority of 80 in the scheduler sched_fifo. The sched_fifo associated threads and processes are treated as realtime operations. The priority of 80 ensures that the event handling of the task is treated with higher priority than softinterrupts in the system, whose standard priority is 50. Besides this, the measurement is started with the parameter -häd enables the histogram functionality for each started measurement cycle. On the target the measurement is started with the following parameters:

```
# cyclictest -p 80 -n -t 1 -D 30m -h 150
```

For the explanation of the funtionality and the interpretation of the measurement results please refer to the chapter 4.2.1.

4.2.4 Testresults

In the following chapters the measurement results ordered by the load scenarios will be explained.

4.2.4.1 Load scenario: Load of the scheduler

The load of the scheduler is simulated with the tool hackbench:

```
# while [ true ]; do hackbench 10; done
```

Following you see a represenation of the results put in a graph and in table form.

	[µs]
Max.	107
Avg.	52
Min.	28

Table 4.1: Results with scheduler load

The figure 4.1 shows the exposition of the measurement deviation.

Figure 4.1: Deviation of the MPC5200B with schedular load

4.2.4.2 Load scenario: Load of the memoryinterface

The laod of the memory interface and the processor cache is simulated with the tool lat_mem_rd. It is started in a manner to generate the highest possible memory load on the target:

```
# while [ 1 ]; do lat_mem_rd 48; done
```

Following you see a representation of the results put in a graph and in table form.

	[µs]
Max.	144
Avg.	31
Min.	20

Table 4.2: Results with memoryinterface load

The figure 4.2 shows the exposition of the measurement deviation.

Pengutronix.

Figure 4.2: Deviation of the MPC5200B with memoryinterface load

4.2.4.3 Load scenario: Network I/O load

The network I/O load is generated with a separate host computer:

```
# sudo ping -f \${TARGET_IP}
```

Following you see a represenation of the results put in a graph and in table form.

	[µs]
Max.	102
Avg.	45
Min.	18

Table 4.3: Results with network I/O load

The figure 4.3 shows the exposition of the measurement deviation.

Pengutronix

Figure 4.3: Deviation of the MPC5200B with network I/O load

4.2.4.4 Loadscenario: Memory I/O load

To generate the I/O load on the memory, the program dd is used to read and write two mebibytes of random data on the flash memory. The MPC5200B uses NOR flash.

```
# while [ 1 ]; do rm -f tmp.pattern; dd if=/dev/urandom of=\tmp.pattern count=2048;
  dd if=tmp.pattern of=/dev/null; done
```

Following you see a represenation of the results put in a graph and in table form.

	[µs]
Max.	108
Avg.	35
Min.	18

Table 4.4: Results with memory I/O load

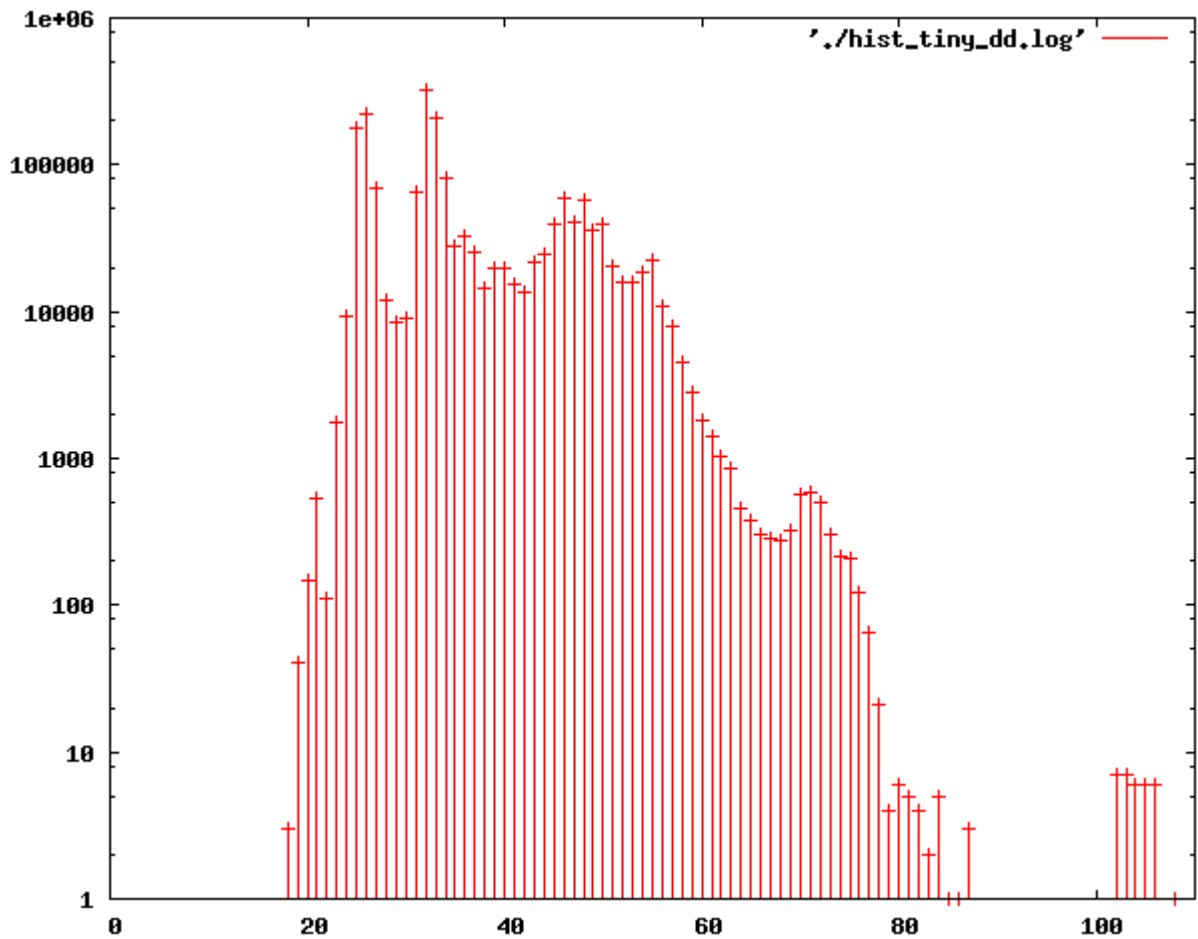The figure 4.4 shows the exposition of the measurement deviation.

Figure 4.4: Deviation of the MPC5200B with memory I/O load

5 Getting help

Below is a list of locations where you can get help in case of trouble. For questions how to do something special with PTXdist or general questions about Linux in the embedded world, try these.

5.1 Mailing Lists

5.1.1 About PTXdist in Particular

This is an English language public mailing list for questions about PTXdist. See

<div align="center">

http://www.pengutronix.de/mailinglists/index_en.html

</div>

how to subscribe to this list. If you want to search through the mailing list archive, visit

<div align="center">

http://www.mail-archive.com/

</div>

and search for the list *ptxdist*. Please note again that this mailing list is just related to the PTXdist as a software. For questions regarding your specific BSP, see the following items.

5.1.2 About Embedded Linux in General

This is a German language public mailing list for general questions about Linux in embedded environments. See

<div align="center">

http://www.pengutronix.de/mailinglists/index_de.html

</div>

how to subscribe to this list. Note: You can also send mails in English.

5.2 News Groups

5.2.1 About Linux in Embedded Environments

This is an English newsgroup for general questions about Linux in embedded environments.

<div align="center">

comp.os.linux.embedded

</div>

5.2.2 About General Unix/Linux Questions

This is a German newsgroup for general questions about Unix/Linux programming.

<div align="center">

de.comp.os.unix.programming

</div>

5.3 Chat/IRC

About PTXdist in particular

irc.freenode.net:6667

Create a connection to the **irc.freenode.net:6667** server and enter the chatroom **#ptxdist**. This is an English room to answer questions about PTXdist. Best time to meet somebody there is at European daytime.

5.4 phyCORE-MPC5200B-IO Support Mailing List

OSELAS.Phytec@pengutronix.de

This is an english language public maillist for all BSP related questions specific to Phytec's hardware. See web site

http://www.pengutronix.de/mailinglists/index_en.html

5.5 Commercial Support

You can order immediate support through customer specific mailing lists, by telephone or also on site. Ask our sales representative for a price quotation for your special requirements.

Contact us at:

Pengutronix
Peiner Str. 6-8
31137 Hildesheim
Germany
Phone: +49 - 51 21 / 20 69 17 - 0
Fax: +49 - 51 21 / 20 69 17 - 55 55

or by electronic mail:

sales@pengutronix.de