

PCI Express® driver for our DESY AMC® boards

Most of our recently developed AMC® boards like the DAMC-FMC2ZUP or the DAMC-FMC1Z7IO are based on modern Xilinx® SoC devices. To implement the PCI Express® endpoint and the corresponding logic (register access and DMA) we use Xilinx® XDMA.

To provide seamless integration between our boards and Linux®-based computers, we have packaged the XDMA driver as a Dynamic Kernel Module Support (DKMS) package under a permissive open-source license.

The source code for the package can be found on our GitHub page: <https://github.com/MicroTCA-Tech-Lab/xdma-metapackage>

Once installed, the package provides a kernel module (xdma.ko) and some command line utilities (xdma-reg-rw, xdma-dma-from-device, xdma-dma-to-device). Shown below is an example of a DMA transfer between the on-board DDR4 and the CPU memory, where a throughput of 4.3 GB/s is reached (with PCIe gen 3 x8 link):

```
$ xdma-dma-from-device \
  -d /dev/xdma/slot2/c2h0 \
  -a 0x400000000 \
  -s $((1024*1024*1024)) \
  -v -c 10
#0: CLOCK_MONOTONIC 0.368244053 sec. read 1073741824/1073741824 bytes
#1: CLOCK_MONOTONIC 0.227387605 sec. read 1073741824/1073741824 bytes
#2: CLOCK_MONOTONIC 0.231948013 sec. read 1073741824/1073741824 bytes
#3: CLOCK_MONOTONIC 0.230998660 sec. read 1073741824/1073741824 bytes
#4: CLOCK_MONOTONIC 0.231771950 sec. read 1073741824/1073741824 bytes
#5: CLOCK_MONOTONIC 0.232103944 sec. read 1073741824/1073741824 bytes
#6: CLOCK_MONOTONIC 0.230952715 sec. read 1073741824/1073741824 bytes
#7: CLOCK_MONOTONIC 0.231591264 sec. read 1073741824/1073741824 bytes
#8: CLOCK_MONOTONIC 0.231777930 sec. read 1073741824/1073741824 bytes
#9: CLOCK_MONOTONIC 0.231258505 sec. read 1073741824/1073741824 bytes
** Avg time device /dev/xdma/slot2/c2h0, total time 2448034639 nsec, avg_time =
244803456.000000, size = 1073741824, BW = 4386.138184
/dev/xdma/slot2/c2h0 ** Average BW = 1073741824, 4386.13818
```

Additionally, a udev script assures that device names match the slot number in a MicroTCA® crate, similar to the Consistent Network Device Naming convention for network adapters. Shown below is the output from the MCH (on the left) and the listing of the directories with driver devices (on the right):

> show_fru	\$ ls -l /dev/xdma
13 AMC9 M4 DAMC-DS812ZUP	slot10
14 AMC10 M4 DAMC-FMC1Z7IO	slot11
15 AMC11 M4 DAMC-FMC2ZUP	slot12
16 AMC12 M4 DAMC-FMC2ZUP	slot9

The driver can then be used from C++ and Python (we provide an example together with DAMC-FMC2ZUP and DAMC-FMC1Z7IO) and it is also supported by the latest ChimeraTK release. The .deb package for Ubuntu 20.04 LTS is also available in the [DOOCS Debian repository](#) and can be installed with one command:

```
sudo apt install xdma-dkms
```

PCI Express® and PCIe® are registered trademarks of PCI-SIG. Xilinx® and Zynq® are registered trademarks and UltraScale+™ is a trademark of Xilinx, Inc. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. MicroTCA® is a registered trademark of PCI Industrial Computer Manufacturers Group. AdvancedMC® is trademark of the PCI Industrial Computer Manufacturers Group.

DESY

Deutsches Elektronen-Synchrotron
Notkestr. 85 • 22607 Hamburg
mtca-techlab@desy.de
techlab.desy.de

microTCA
TECHNOLOGY LAB

