Open Your AIoT gateway: A Journey into Ubuntu 24.04 with lightweight desktop and NPU power

UbuCon Korea 2025, Wig Cheng (8/10/2025)

About Me

Wig Cheng

IEI Integration Corp.
Android OS Engineering Manager

Skills

ARM Android BSP, Yocto BSP, Ubuntu/Debian OS
U-boot/Kernel Development
RISC-V / ARM MCU Development

Open Source Contribution

U-Boot Upstream
Linux Kernel Upstream
Google AOSP Upstream
Niryo Robotic arm – ROS Stack

Open Source Communities

Kakip AI-SBC Community - Maintainer
OpenEPD – Maintainer
Google Developer Groups Taoyuan Taiwan – AOSP Instructor

Github: wigcheng

Agenda

01 What is the Gateway

02 Make an Ubuntu rootfs with lightweight desktop

03 AIoT with NPU accelerations

04 MCU develop environment

05 Conclusion & Future work

What is the Gateway

What is the Gateway

Industrial Gateway

Connects OT (Operational Tech) with IT (Information Tech) systems.

Connect PLCs, SCADAs, CNCs, legacy industrial equipment.

Protocol: Modbus, OPC-UA
Application: Factory

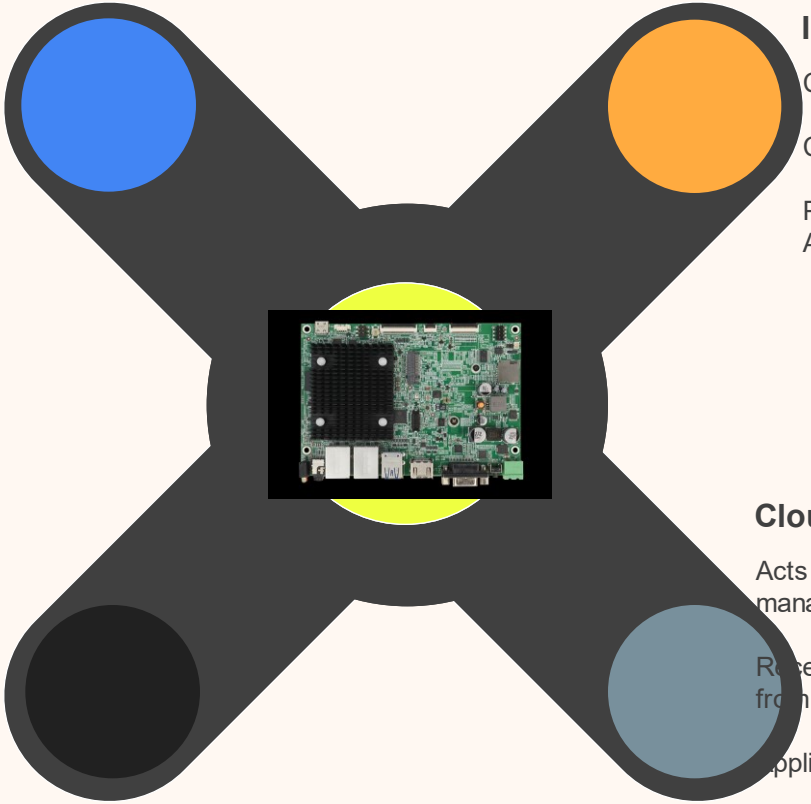

IOT Gateway

Connects end devices/sensors to the Internet

Connects Various sensors, wearables.

Protocol: BLE, MQTT, Zigbee
Application: Basic edge computing

AIOT Gateway

Performs AI inference at the edge;

Same as IoT, but often **cameras**, **microphones** for AI inference.

Protocol: V4L2, ALSA
Application: "Smart" edge computing

Cloud Gateway

Acts as a single entry point for cloud services; manages & authenticates connections.

Receives data from IoT/AIoT gateways or directly from devices.

Application: International data transferring

Make an Ubuntu rootfs with lightweight desktop

Platforms

- WAFER-IMX8MP
  - NXP IMX8MP
  - CPU: Cortex-A53 + Cortex-M7
  - GPU: Vivante GC7000UL
  - NPU: VeriSilicon VIPNano-SI+.8002 (**2.3TOPS**)

- FRDM-IMX93
  - NXP IMX93
  - CPU: Cortex-A55 + Cortex-M33
  - GPU: N/A
  - NPU: ARM **Ethos-U65** microNPU (**0.5TOPS**)

- RZ/T2H EVK
  - Renesas RZ/T2H
  - CPU: Cortex-A55 + Cortex-R8
  - GPU: N/A
  - NPU: N/A

Debootstrap

- **Definition**: A command-line tool used to install a **minimal but fully functional** Debian, Ubuntu, or derivative base system into a specified directory.

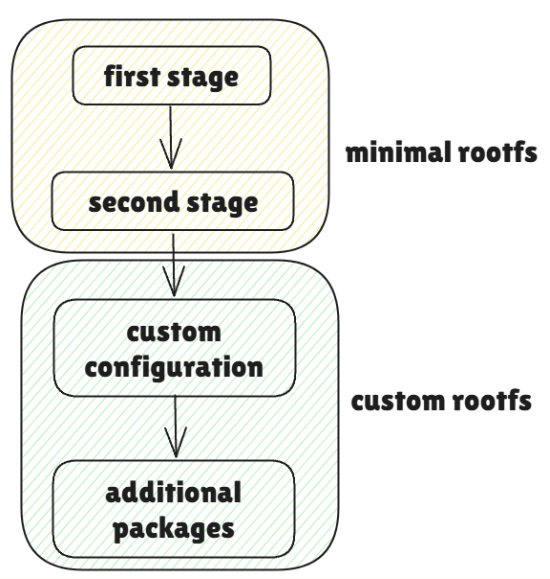

- Ex: ARM64 Ubuntu 24.04 **minimal rootfs**

```
$ sudo apt update
$ sudo apt install debootstrap
$ sudo apt install qemu-system-arm qemu-user-static
$ sudo debootstrap --arch=arm64 --
keyring=/usr/share/keyrings/ubuntu-archive-
keyring.gpg --verbose --foreign noble rootfs

$ sudo cp /usr/bin/qemu-aarch64-
static rootfs/usr/bin

$ sudo LANG=C chroot rootfs
/debootstrap/debootstrap --second-stage
```

tarball the rootfs.tgz from ***rootfs folder***

Live-Build

- **Bootstrap**： This is the initial phase of populating the chroot directory with packages to make a barebones Debian system
- **Chroot**： completes the construction of chroot directory, populating it with all of the packages listed in the configuration, along with any other materials. Most customization of content occurs in this stage
- **Binary**： **builds a bootable image**, using the contents of the chroot directory to construct the root filesystem for the Live system, and including the installer and any other additional material on the target medium outside of the Live system's filesystem

- Ex: ARM64 Ubuntu 24.04 minimal rootfs

  $ sudo apt update

  $ sudo apt install live-build

  $ sudo apt install qemu-utils qemu-system qemu-user

```
sudo lb config \
--architecture arm64 \
--bootstrap-qemu-arch arm64 \
--bootstrap-qemu-static /usr/bin/qemu-aarch64-static \
--archive-areas "main restricted universe multiverse" \
--parent-archive-areas "main restricted universe multiverse" \
--mirror-bootstrap "http://ports.ubuntu.com/ubuntu-ports/" \
--parent-mirror-bootstrap "http://ports.ubuntu.com/ubuntu-ports/" \
--mirror-chroot "http://ports.ubuntu.com/ubuntu-ports/" \
--parent-mirror-chroot "http://ports.ubuntu.com/ubuntu-ports/" \
--mirror-chroot-security "http://ports.ubuntu.com/ubuntu-ports/" \
--parent-mirror-chroot-security "http://ports.ubuntu.com/ubuntu-ports/" \
--mirror-binary "http://ports.ubuntu.com/ubuntu-ports/" \
--parent-mirror-binary "http://ports.ubuntu.com/ubuntu-ports/" \
--mirror-binary-security "http://ports.ubuntu.com/ubuntu-ports/" \
--parent-mirror-binary-security "http://ports.ubuntu.com/ubuntu-ports/" \
--keyring-packages ubuntu-keyring \
--system normal \
--mode ubuntu \
--chroot-filesystem ext4 \
--distribution noble
```

ARM32 V.S. ARM64

- On WAFER-IMX8MP platforms, the performance of the ARM64 rootfs is about **30%** better than that of the ARM32 rootfs.
  - ARM64 inference FPS using TF-Lite: 24FPS
  - ARM32 inference FPS using TF-Lite: 16FPS
- Raspberry PI face this problem before

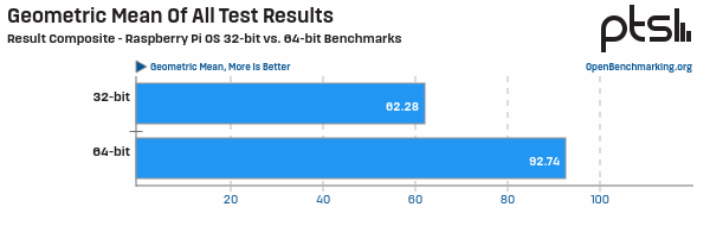

[# from Phoronix]

- Why we need ARM32 rootfs?
  For some special case such as.....

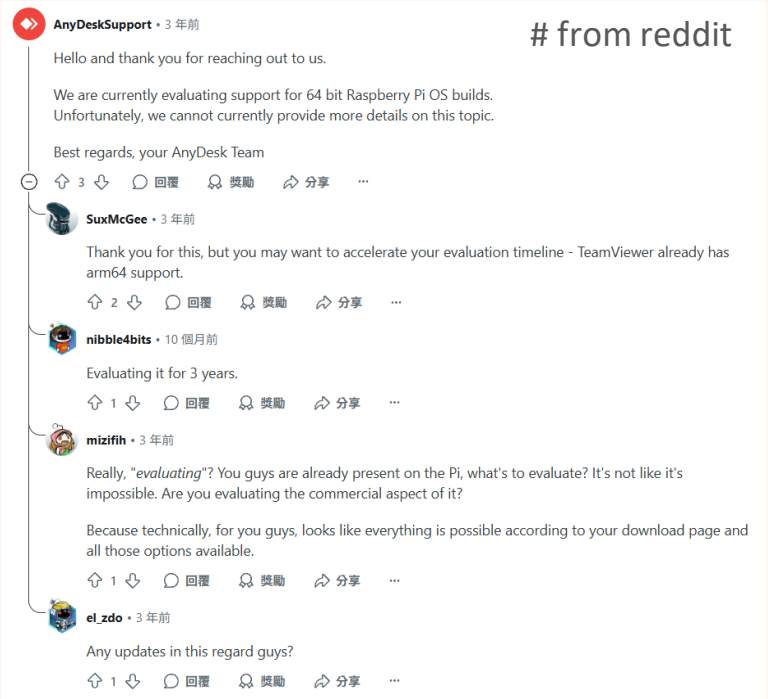

LXQt Desktop

- Necessary packages in qemu environment
  - $ apt -y install lxqt (1.4.0)
  - $ apt -y install slim
- Auto login configuration (in qemu as well)
  - $ sed -i 's/#auto_login\s\+no/auto_login yes/' /etc/slim.conf
  - $ sed -I 's/#default_user\s\+simone/default_user ubuntu/' /etc/slim.conf

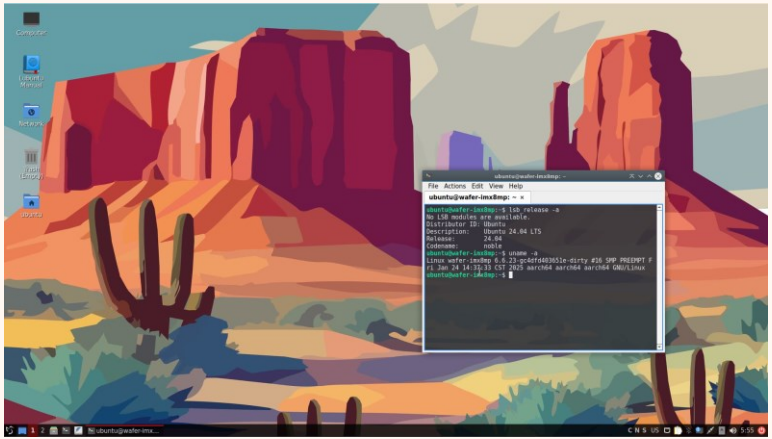

LXQt

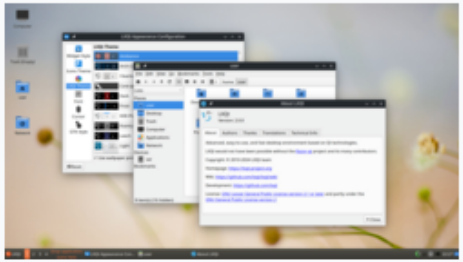

LXQt 2.0.0 Screenshot

Original author(s)	Hong Jen Yee (“PCMan”)
Developer(s)	The LXQt team[1]
Initial release	July 21, 2013; 11 years ago
Stable release	2.2.0[2] / 17 April 2025
Repository	github.com/lxqt
Written in	C++, C, Qt
Operating system	Unix-like
Available in	Multilingual
Type	Desktop environment
License	GPL, LGPL
Website	lxqt-project.org

Home | UbuCon Korea 2025 - Chromium

UbuCon Korea 2025 - Eve | Home | UbuCon Korea 20

2025.ubuntu-kr.org/ko/

UbuCon Korea 2025

UbuCon Korea 2025

2025년 8월 10일

한국마이크로소프트 12-13층

우분투한국커뮤니티의 연례 행사인 우부콘 코리아(UbuCon Korea)는 매년 국내 우분투 사용자와 기여자를 연결하는 가장 큰 행사 입니다. 오는 여름, 개발자, 엔지니어, 크리에이터, 연구원, 기업가 그리고 우분투를 좋아하는 여러분 모두를 우분투한국커뮤니티의 20주년을 축하하는 자리에 초대합니다!

참가등록 | 후원사로 참여하기 | 얼리버드 할인코드 받기

우분투와 오픈소스 공통점 하나로 모두가 모이는 행사

우부콘 코리아(UbuCon Korea)는 우분투와 오픈소스라는 공통점 하나로 매년 다양한 배경을 가진 사람들이 모여 우분투와 오픈소스로 무엇을 할 수 있고, 또 어떻게 우분투 프로젝트와 커뮤니티에 기여할 수 있는지 함께 모색하는 우분투한국커뮤니티의 연례 행사 입니다.

이 행사에 와도 좋을지 확신이 서지 않나요? 작년 현장 스케치를 영상으로 만나보세요.

주요 연사

Hanku Lee

Grégory Schiano

ubuntu@frdm-imx93: ~

File Actions Edit View Help

ubuntu@frdm-imx93: ~

```
ubuntu@frdm-imx93:~$ neofetch
                    .-/+oossssoo+/-.               ubuntu@frdm-imx93
                `:+ssssssssssssssssss+:`           ----------------
              -+ssssssssssssssssssyyssss+-         OS: Ubuntu 24.04 LTS aarch64
            .ossssssssssssssssssdMMMNysssso.       Host: NXP i.MX93 11X11 FRDM board
           /ssssssssssshdmmNNmmyNMMMMhssssss/      Kernel: 6.12.3-gbd7a9d3d93d5
          +ssssssssshmydMMMMMMMNddddyssssssss+     Uptime: 42 mins
         /sssssssshNMMMyhyyyyhmNMMMNhssssssss/     Packages: 1595 (dpkg), 8 (snap)
        .ssssssssdMMMNhsssssssssshNMMMdssssssss.   Shell: bash 5.2.21
        +sssshhhyNMMNyssssssssssssyNMMMysssssss+   Resolution: 1920x1080
        ossyNMMMNyMMhsssssssssssssshmmmhssssssso   DE: LXQt 1.4.0
        ossyNMMMNyMMhsssssssssssshmmmhssssssso     WM: Xfwm4
        +sssshhhyNMMNyssssssssssssyNMMMysssssss+   WM Theme: Breeze
        .ssssssssdMMMNhsssssssssshNMMMdssssssss.   Theme: Yaru [GTK3]
         /sssssssshNMMMyhyyyyhdNMMMNhssssssss/     Icons: Yaru [GTK3]
          +ssssssssshmydMMMMMMMNddddyssssssss+     Terminal: qterminal
           /ssssssssssshdmNNNNmyNMMMMhssssss/      CPU: (2)
            .ossssssssssssssssssdMMMNysssso.       Memory: 839MiB / 1940MiB
              -+ssssssssssssssssssyyyssss+-
                `:+ssssssssssssssssss+:`
                    .-/+oossssoo+/-.

ubuntu@frdm-imx93:~$
```

About LXQt

LXQt
Version: 1.4.0

About | Authors | Thanks | Translations | Technical Info

LXQt Desktop Toolbox - Technical Info

LXQt About Version	1.4.0
LXQt Version	1.4.0
Qt	5.15.13
Build type	Release
System Configuration	/etc/xdg
Share Directory	/usr/share/lxqt
Translations	/usr/share/lxqt/lxqt/translations /usr/local/share/lxqt/translations /usr/share/lxqt/translations

Copy to clipboard

× Close

1 2 Home | UbuCon Kor... ubuntu@frdm-imx9... About LXQt

CNS US 8:17

AIoT with NPU accelerations

Backport upstream IMX8MP NPU driver

- Special thanks...

Tomeu Vizoso
NPU driver Engineer
Maintainer for RK3588, **IMX8MP** NPU upstream driver

IMX8MP etnaviv driver
Upstream Kernel 6.10 backport to
Vendor Kernel 6.6.23

drivers/gpu/drm/etnaviv/etnaviv_gpu.c

```
&gpu_3d {
»   status = "okay";
+»  status = "disabled";
};

&gpu_2d {
»   status = "okay";
+»  status = "disabled";
};

&ml_vipsi {
+»  compatible = "vivante,gc";
»   status = "okay";
};

&mix_gpu_ml {
»   status = "okay";
+»  status = "disabled";
};
```

drm/etnaviv: fix tx clock gating on some GC7000 variants
ManMower authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Restore some id values
Christian Gmeiner authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: add sensitive state for PE_RT_ADDR_4_PIPE(3, 0|1) address
Christian Gmeiner authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Expose a few more chipspecs to userspace
tomeuv authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: disable MLCG and pulse eater on GPU reset
lynxeye-dev authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Drop the 'len' parameter of etnaviv_iommu_map() function
Sui Jingfeng authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Clean up etnaviv_gem_get_pages
Sui Jingfeng authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Add a helper to get the first available GPU device node
Sui Jingfeng authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Convert to platform remove callback returning void
Uwe Kleine-König authored and wigcheng committed on Dec 31, 2024

drm/sched: Convert the GPU scheduler to variable number of run-queues
Luben Tuikov authored and wigcheng committed on Dec 31, 2024

drm/etnaviv: Add helper functions to create and destroy platform device

Library porting - IMX8MP NPU

- MESA 3D – TensorFlow-Lite Delegate
  - Meson build (1.3.2)

```
# Install build dependencies
~ # apt-get -y build-dep mesa
~ # apt-get -y install git cmake

# Download sources
~ $ git clone https://gitlab.freedesktop.org/mesa/mesa.git

# Build Mesa
~ $ cd mesa
mesa $ meson setup build -Dgallium-drivers=etnaviv -Dvulkan-drivers= -Dteflon=true
mesa $ meson compile -C build
```

  - TensorFlow Lite 2.13
  - python 3.10

```
--- a/src/gallium/drivers/etnaviv/etnaviv_ml_nn_v8.c
+++ b/src/gallium/drivers/etnaviv/etnaviv_ml_nn_v8.c
@@ -211,7 +211,8 @@ static uint32_t calculate_bias_correction(struct etna_ml_subgraph *subgraph, con
 static void
 append_bits(uint32_t value, size_t size, struct bitstream *bitstream)
 {
-   assert(value < 1 << size);
+   //assert(value < 1 << (long)size);
+   assert((uint64_t)value < (1ULL << size));
    if (!size)
       return;
    bitstream->buffer |= (uint64_t)value << bitstream->bits_in_buffer;
```

ARM32 rootfs need tweak the source code

```
ubuntu@wafer-imx8mp:~/mesa$ python3.10 src/gallium/frontends/teflon/tests/classification.py        -i ~/grace_hopper.bmp        -m src/gallium/targets/teflon/tests/mobilenet_v1_1.0_224_quant.tflite        -l src/gallium/frontends/
teflon/tests/labels_mobilenet_quant_v1_224.txt
0.874510: military uniform
0.031373: Windsor tie
0.015686: mortarboard
0.011765: bulletproof vest
0.007843: bow tie
time: 223.657ms CPU
ubuntu@wafer-imx8mp:~/mesa$ python3.10 src/gallium/frontends/teflon/tests/classification.py        -i ~/grace_hopper.bmp        -m src/gallium/targets/teflon/tests/mobilenet_v1_1.0_224_quant.tflite        -l src/gallium/frontends/
teflon/tests/labels_mobilenet_quant_v1_224.txt        -e build/src/gallium/targets/teflon/libteflon.so
Loading external delegate from build/src/gallium/targets/teflon/libteflon.so with args: {}
0.870588: military uniform
0.031373: Windsor tie
0.011765: mortarboard
0.007843: bow tie
0.007843: bulletproof vest
time: 7.963ms NPU
```

TensorFlow-Lite demo

- https://github.com/tomeuv/TensorFlow-Lite-Object-Detection-on-Android-and-Raspberry-Pi/tree/teflon-demo
- Command:
  $ export CAMERA_NUM=3
  $ python3 TFLite_detection_webcam.py --modeldir=./ --edgetpu
  (adopts coco_ssd_mobilenet_v1_1.0_quant model)

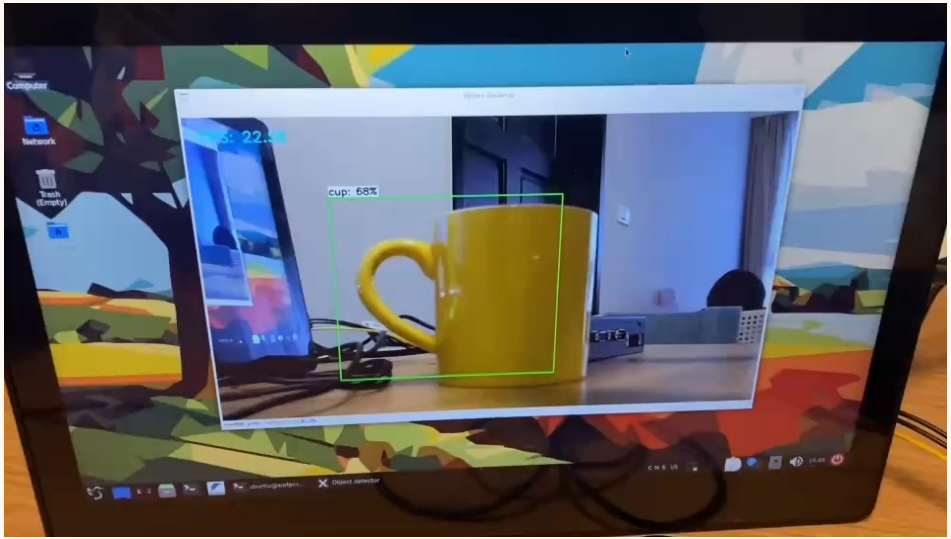

Library porting - IMX93 NPU

- Not an upstream driver, it's propietary driver
- Necessary repos (Kernel 6.6.36)
  - Vela 3.12.0
  - TensorFlow Lite 2.16
  - Python 3.12
  - ethosu_delegate 3.16
  - ethos-u-firmware

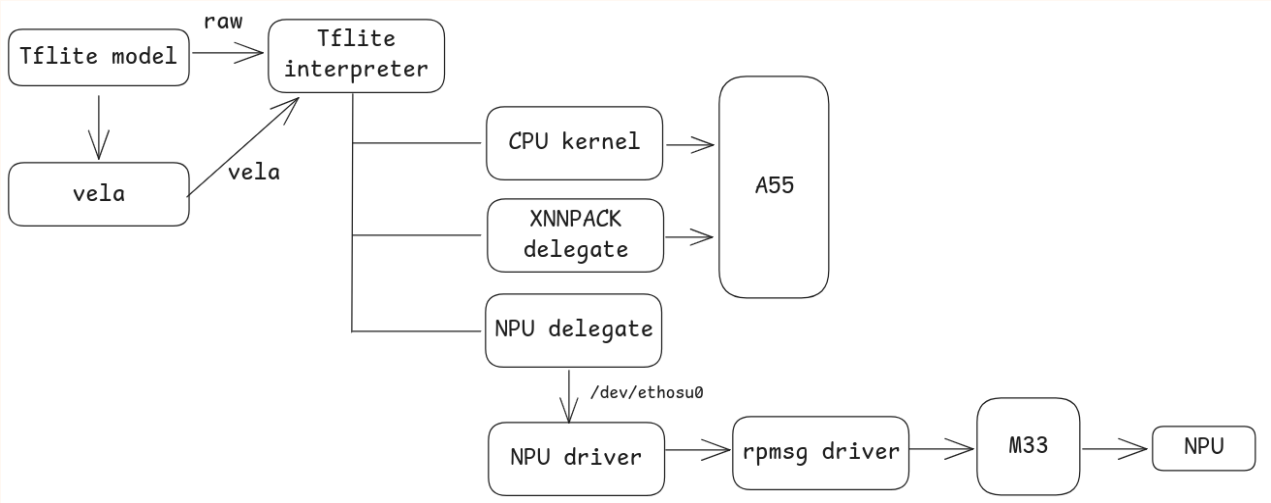

TensorFlow-Lite demo

```
ubuntu@frdm-imx93:~/ai$ python3.12 ./classification.py -i grace_hopper.bmp -l labels_mobilenet_quant_v1_224.txt -m mobilenet_v1_1.0_224_quant.tflite
INFO: Created TensorFlow Lite XNNPACK delegate for CPU.
0.870588: military uniform
0.031373: Windsor tie
0.011765: mortarboard
0.007843: bow tie
0.007843: bulletproof vest
time: 51.716ms
```

CPU

```
ubuntu@frdm-imx93:~/ai$ python3.12 ./classification.py -i grace_hopper.bmp -l labels_mobilenet_quant_v1_224.txt -m mobilenet_v1_1.0_224_quant.tflite -e /usr/lib/aarch64-linux-gnu/libethosu_delegate.so
Loading external delegate from /usr/lib/aarch64-linux-gnu/libethosu_delegate.so with args: {}
INFO: Ethosu delegate: device_name set to /dev/ethosu0.
INFO: Ethosu delegate: cache_file_path set to .
INFO: Ethosu delegate: timeout set to 60000000000.
INFO: Ethosu delegate: enable_cycle_counter set to 0.
INFO: Ethosu delegate: enable_profiling set to 0.
INFO: Ethosu delegate: profiling_buffer_size set to 2048.
INFO: Ethosu delegate: pmu_event0 set to 0.
INFO: Ethosu delegate: pmu_event1 set to 0.
INFO: Ethosu delegate: pmu_event2 set to 0.
INFO: Ethosu delegate: pmu_event3 set to 0.
INFO: EthosuDelegate: 31 nodes delegated out of 31 nodes with 1 partitions.
0.874510: military uniform
0.031373: Windsor tie
0.015686: mortarboard
0.011765: bulletproof vest
0.007843: bow tie
time: 3.838ms
```

NPU

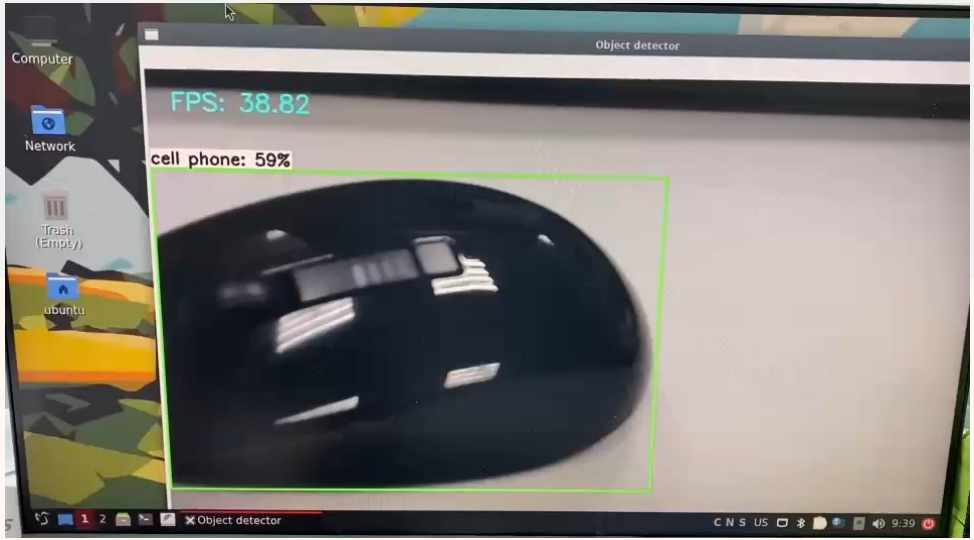

Single software stack for IMX8MP and IMX93

- Ubuntu 24.04
  - Default is python 3.12
  - Manual install python 3.10 in qemu environment
    # add-apt-repository -y ppa:deadsnakes/ppa
    # apt -y install python3.10 python3.10-dev python3-pip
  - Install specific packages using ***pip3 tool***
    # **python3.10** -m pip install --break-system-packages <target package==version>
    # **python3.12** -m pip install --break-system-packages <target package==version>

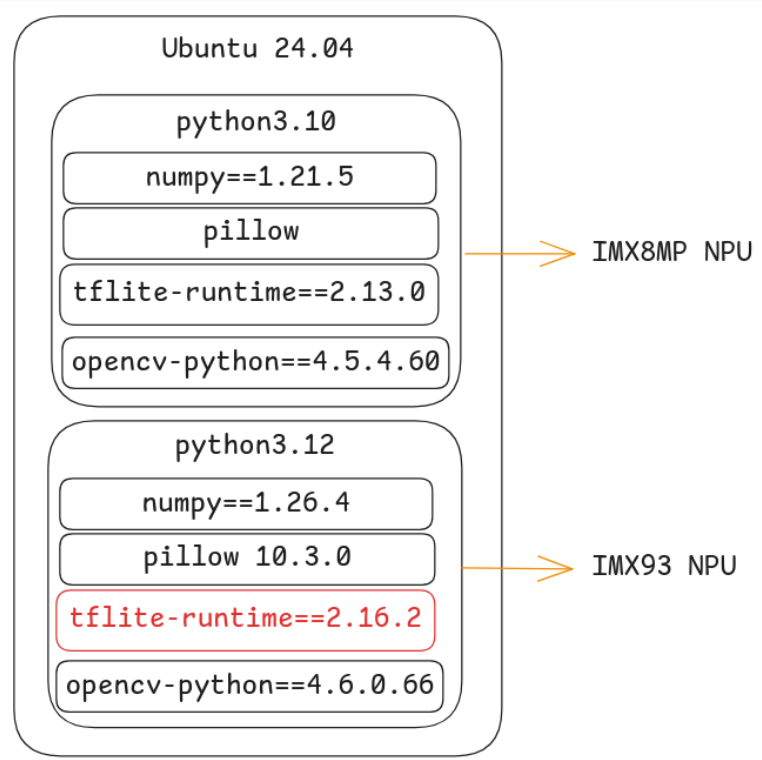

MCU develop environment

An e-paper example for IMX93

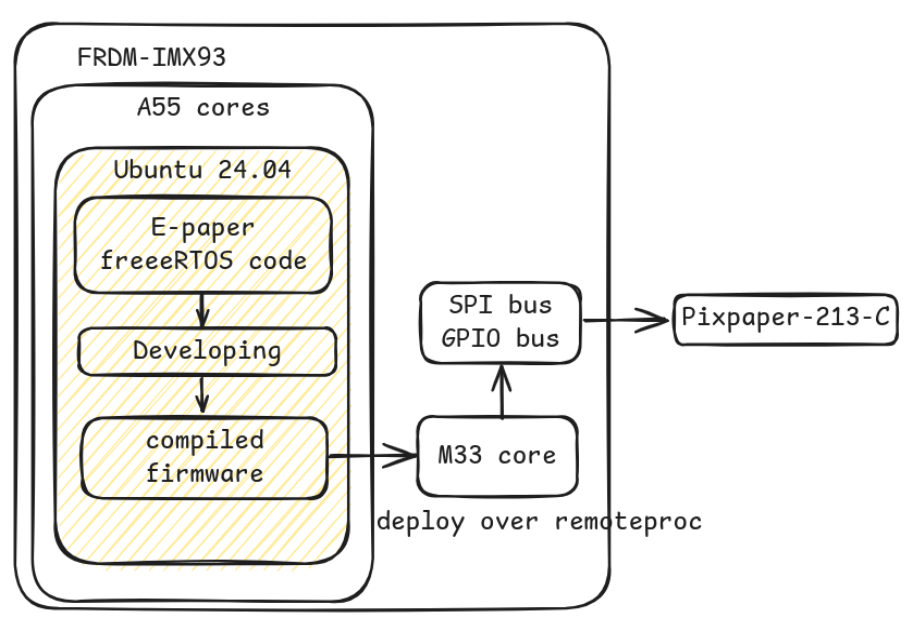

- Cortex M33 MCU belongs to ARM32
  - Toolchain
    - Development on host PC
      `arm-gnu-toolchain-12.2.mpacbti-rel1-x86_64-arm-none-eabi.tar.xz`
    - Development on FRDM-IMX93
      `arm-gnu-toolchain-12.2.mpacbti-rel1-aarch64-arm-none-eabi.tar.xz`
  - Prepare FreeRTOS source code
    - Write example code and Makefile

      $ export ARMGCC_DIR=/opt/m33/toolchain

      $ cd <source directory>

      $ ./build_release.sh

      output file: e-paper.elf

Demo

```
Assign elf load path and filename
$ echo -n /home/ubuntu/ > /sys/module/firmware_class/parameters/path
$ echo -n e-paper.elf > /sys/class/remoteproc/remoteproc0/firmware

Running / Stop the elf program
$ echo -n start > /sys/class/remoteproc/remoteproc0/state
$ echo -n stop > /sys/class/remoteproc/remoteproc0/state
```

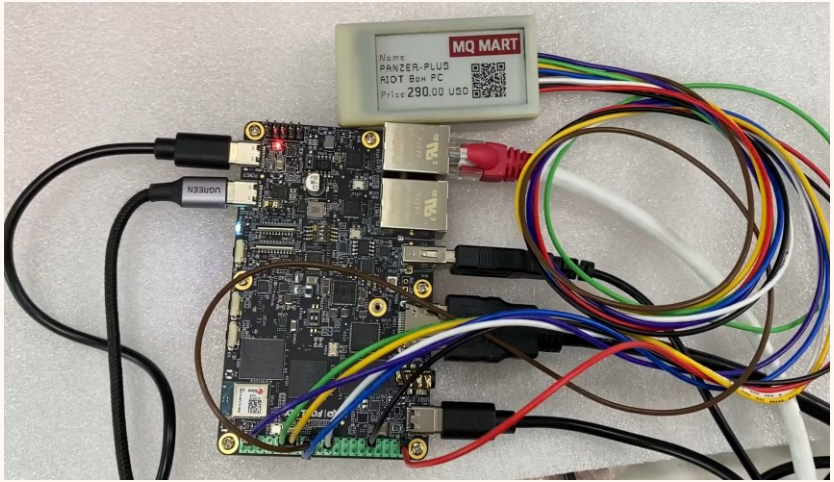

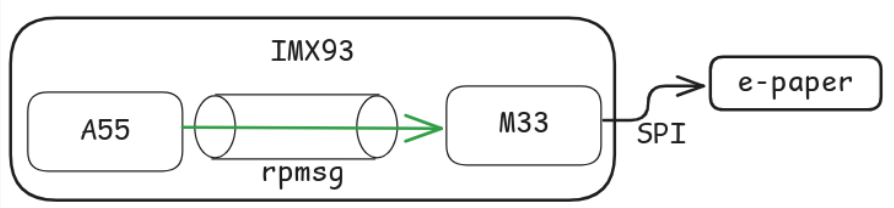

Conclusion and future work

Conclusion

- Generate Ubuntu/Debian rootfs methods
  - Debootstrap: Easy, but need write the build architecture
  - Live-Build: Difficult, but more flexible for large architecture
- ARM64 platform
  - The performance of ARM64 rootfs is better than ARM32 rootfs
  - LXDE, XFCE, and LXQt comparison:
    LXQt is the best compatibility and stability.
- NPU
  - TOPS is not the sole indicator of performance, the optimization of NPU drivers and libraries also plays a critical role.
    - IMX8MP (2.3 TOPS): peak 24 FPS => But the good news is.....
    - IMX93 (0.5 TOPS): peak 40 FPS

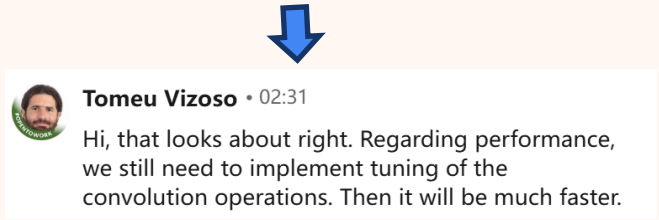

Our Ubuntu BSP

- Download link
  - https://github.com/QNAP-android-internal/ubuntu-classic-imx.git
- Compile command
  - **WAFER-IMX8MP:** make rootfs PLATFORM="wafer-imx8mp"
  - **FRDM-IMX93:** make rootfs PLATFORM="frdm-imx93"
  - **RZ/T2H EVK:** WIP

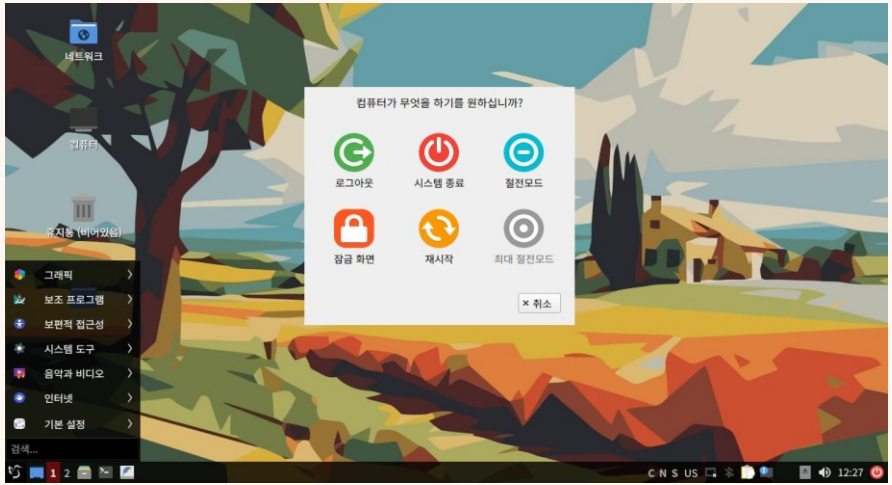

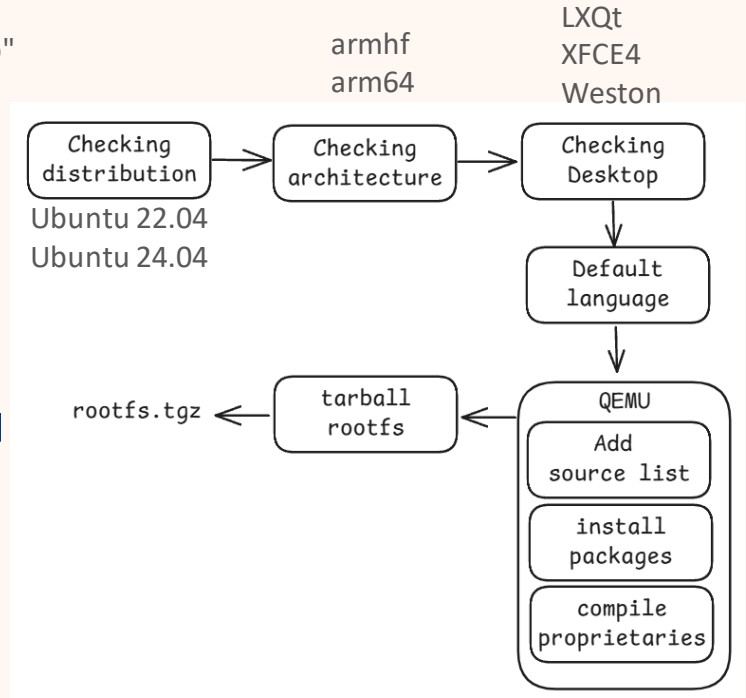

Future work

- Keep improving Ubuntu BSP
  - Build rootfs tarball → Build runtime image
  - Dynamically switch between mainline and vendor kernel
  - IMX8MP: Keep backport the mainline driver
  - IMX93: Observing the mainline driver => **Good news from 7/23**
- More platform support
  - RZ/T2H EVK
  - Other low cost platforms

```
From: Rob Herring (Arm) @ 2025-07-22 22:58 UTC (permalink / raw)
  To: Tomeu Vizoso, Krzysztof Kozlowski, Conor Dooley, Oded Gabbay,
      Maarten Lankhorst, Maxime Ripard, Thomas Zimmermann, David Airlie,
      Simona Vetter, Sumit Semwal, Christian König, Robin Murphy,
      Steven Price
  Cc: devicetree, linux-kernel, dri-devel, linux-media, linaro-mm-sig

The Arm Ethos-U65/85 NPUs are designed for edge AI inference
applications[0].

The driver works with Mesa Teflon. WIP support is available here[1]. The
UAPI should also be compatible with the downstream driver stack[2] and
Vela compiler though that has not been implemented.

Testing so far has been on i.MX93 boards with Ethos-U65. Support for U85
is still todo. Only minor changes on driver side will be needed for U85
support.

A git tree is here[3].

Rob

[0] https://www.arm.com/products/silicon-ip-cpu?families=ethos%20npus
[1] https://gitlab.freedesktop.org/tomeu/mesa.git ethos
[2] https://gitlab.arm.com/artificial-intelligence/ethos-u/
[3] git://git.kernel.org/pub/scm/linux/kernel/git/robh/linux.git ethos

Signed-off-by: Rob Herring (Arm) <robh@kernel.org>
---
Rob Herring (Arm) (2):
      dt-bindings: npu: Add Arm Ethos-U65/U85
      accel: Add Arm Ethos-U NPU driver

 .../devicetree/bindings/npu/arm,ethos.yaml         |  79 +++
 MAINTAINERS                                        |   9 +
 drivers/accel/Kconfig                              |   1 +
 drivers/accel/Makefile                             |   1 +
 drivers/accel/ethos/Kconfig                        |  10 +
 drivers/accel/ethos/Makefile                       |   4 +
 drivers/accel/ethos/ethos_device.h                 | 186 ++++++
 drivers/accel/ethos/ethos_drv.c                    | 412 ++++++++++++
 drivers/accel/ethos/ethos_drv.h                    |  15 +
 drivers/accel/ethos/ethos_gem.c                    | 707 ++++++++++++++++++++
 drivers/accel/ethos/ethos_gem.h                    |  46 ++
 drivers/accel/ethos/ethos_job.c                    | 527 ++++++++++++++++
 drivers/accel/ethos/ethos_job.h                    |  41 ++
 include/uapi/drm/ethos_accel.h                     | 262 ++++++++
 14 files changed, 2300 insertions(+)
---
base-commit: 19272b37aa4f83ca52bdf9c16d5d81bdd1354494
change-id: 20250715-ethos-3fdd39ef6f19

Best regards,
```

Thank you!!!