

The Arm laptop project

Johan Hovold

Hovold Consulting AB

Kernel Recipes, Paris

25 September 2023

Introduction

- Windows on Arm (WoA)
- Laptops built on (Qualcomm) 64-bit Arm SoCs
- Quiet and power efficient
- Non-standard boot chain
- Customised OS

Introduction

- Windows on Arm (WoA)
- Laptops built on (Qualcomm) 64-bit Arm SoCs
- Quiet and power efficient
- Non-standard boot chain
- Customised OS
- What would it take to run Linux on them?

Project overview

- Linaro project
- Funded by Arm and Qualcomm
- Hardware from Qualcomm and Lenovo
- Proof-of-concept: Mainline Linux on Arm laptop (built for Windows)
- How far could we get?

People

- Johan Hovold
- Björn Andersson
- Linaro's Qualcomm landing team
- Building on work done by Qualcomm, Linaro and many others

Talk outline

- Background
- Status
- Work in progress
- Open issues
- Setup instructions
- Distro support

Previous work

- Linaro and Arm collaboration
- Try to get Linux to boot on WoA laptops
 - ASUS NovaGo TP370QL
 - HP Envy x2
 - Lenovo Mixx 630
 - Lenovo Yoga C630
- Enable basic features
- Investigate options for distro support (e.g. ACPI vs DT)
- <https://github.com/aarch64-laptops/>
- #aarch64-laptops on OFTC

Lenovo ThinkPad X13s

- Snapdragon 8xc Gen3 Compute (sc8280xp)
- 13.3" FHD (1920x1200) display
- 32 GB LPDDR4 (up to)
- 1 TB NVMe SSD (up to)
- Adreno 690 GPU
- 2 speakers, 2 microphones, headphone jack
- 5 MP camera
- Wi-Fi 6E (802.11ax)
- Bluetooth 5.1
- 5G modem (optional)
- 2 x USB-C 3.2 Gen 2
- 49.5 Wh battery



Boot firmware

- Qualcomm and Windows on Arm
- UEFI
 - Runtime services not available after boot (e.g. EFI variables)
- ACPI vs DT
 - Supporting Qualcomm's non-standard ACPI not feasible
- DtbLoader.efi
 - Ship devicetree blob with UEFI firmware
 - Allows for generic distro installers
- Hypervisor
 - UEFI and Linux starts in EL1
 - No virtualisation

Work outline

- Feature implementation and mainlining
 - Reverse engineering
 - Refactoring (technical debt)
 - Firmware requests
- Bug fixing
- Patch review
- Patch tracking
- Work-in-progress (WIP) branches
- Support (#aarch64-laptops)

WIP branches

- Important fixes and features under development
- Minimal `johan_defconfig`
- Rebased on RC kernels
- Several regressions found and fixed early
 - Example: IRQ software resend in 6.5-rc1
- 41 branches so far, latest:
 - <https://github.com/jhovold/linux/tree/wip/sc8280xp-v6.6-rc2>
- Announced on `#aarch64-laptops`
 - Testing by early adopters
 - Repackaged with distro config (e.g. by steev)
 - Base for distro images (e.g. Ubuntu)

Device firmware

- Use Windows firmware files for bringup
- Work with Lenovo and Qualcomm to get firmware released
- Wi-Fi board file not compatible with ath11k firmware
 - Took one year to get calibration data released
- Everything but video acceleration now in `linux-firmware.git`

Mainline feature support

- 6.0
 - Backlight
 - CPUfreq
 - Keyboard
 - Remoteproc
 - Touchpad
 - Touchscreen
 - USB
 - Watchdog
- 6.1
 - System suspend
- 6.2
 - Modem
 - NVMe SSD
 - PCIe (x4)
 - Thermal sensors
- 6.3
 - Battery
 - External display
 - Internal display (*)
- 6.4
 - Bluetooth
 - IOMMU (**)
 - RTC
 - Touchpad (alternate)
 - Wi-Fi
- 6.5
 - Audio
 - GPU
 - USB-C orientation switching

Work in progress

- Camera
 - Raw sensor data
 - Software processing (e.g. debayering)
- EFI variables¹ (merged for 6.7)
 - Maximilian Luz
- Fingerprint reader¹
- Video acceleration¹
- USB-C muxing (4-lane DisplayPort)
- Performance optimisation (e.g. better memory bus scaling)
- Power consumption

¹Supported in WIP branches

Power consumption

- Idle: 3.2 W (15 h)²
- Suspend: 1.7 W (29 h)
 - Suspend to idle
 - Not yet hitting deepest low-power state

²Backlight at 66%, PCIe ASPM, Xorg, Wi-Fi

Audio

- Speakers, microphones, headphone jack and DisplayPort working in 6.5
- Late fix for speaker and headphone distortion
- Kernel, topology and UCM files need to be updated in lock step
- Known issues
 - Pops and clicks (partial fix in WIP branch)
 - Microphone distortion

GPU

- Adreno 690 support in Linux 6.5 and Mesa 23.1
 - Rob Clark
 - Important fix in Mesa 23.1.4 (Firefox crash)
- Extracted a690 GMU firmware not (yet) in linux-firmware
 - Driver will use a660 GMU firmware in 6.6 (and 6.5.3)
 - Use symlink before 6.5.3

```
# ln -s a660_gmu.bin a690_gmu.bin
```

Bluetooth

- Unique device address (BD_ADDR) stored externally
- Need details from Qualcomm (or reverse engineer)
- No support in bluetoothd for setting custom or random address
 - <https://github.com/bluez/bluez/issues/107>
- Must be set manually (or using systemd service) for now
- Similar problem with Wi-Fi MAC address
 - Falls back to random address

```
# btmgmt --index 0 public-addr 00:11:22:33:44:55
```

Remote processors

- Qualcomm protection-domain mapper daemon (pd-mapper)
 - <https://github.com/andersson/pd-mapper>
- Reads configuration shipped with remoteproc firmware
- Tells remote processors which services to start
- Tells kernel where to find services
 - Audio
 - Battery
 - USB-C Alternate Mode
- Battery may not charge if pd-mapper not running
- Depends on Qualcomm IPC Router name service (qrtr-ns)
 - <https://github.com/andersson/qrtr>
- Functionality should be moved to kernel and devicetree

Second-source devices

- X13s comes with one of two touchpad controllers
- Boot firmware should determine which one is populated
- Update the devicetree accordingly

Touchpad devicetree

```
&i2c21 {
    touchpad@15 {
        compatible = "hid-over-i2c";
        reg = <0x15>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "disabled";
    };
    touchpad@2c {
        compatible = "hid-over-i2c";
        reg = <0x2c>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "disabled";
    };
};
```

Touchpad devicetree

```
&i2c21 {
    touchpad@15 {
        compatible = "hid-over-i2c";
        reg = <0x15>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "disabled";
    };
    touchpad@2c {
        compatible = "hid-over-i2c";
        reg = <0x2c>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "okay";
    };
};
```

Second-source devices (cont.)

- X13s firmware does not support this
- Linux needs to probe both touchpads
- HID driver already checks if a device is present
- Used for optionally populated devices (e.g. the X13s touchscreen)

HID driver probe

```
int i2c_hid_core_probe(struct i2c_client *client, ...)
{
    ...
    i2c_hid_core_power_up();

    /* Make sure there is something at this address */
    ret = i2c_smbus_read_byte(client);
    if (ret < 0) {
        i2c_hid_core_power_down();
        return -ENODEV;
    }
    hid = hid_allocate_device();
    ...
    hid_add_device(hid);
    return 0;
}
```


Touchpad devicetree

```
&i2c21 {
    touchpad@15 {
        compatible = "hid-over-i2c";
        reg = <0x15>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "okay";
    };
    touchpad@2c {
        compatible = "hid-over-i2c";
        reg = <0x2c>;
        interrupts = <182 IRQ_TYPE_LEVEL_LOW>;
        vdd-supply = <&vreg_3p3>;
        status = "okay";
    };
};
```

Asynchronous probe woes

- Devices can be probed in parallel

Asynchronous probe woes

- Devices can be probed in parallel
- Interrupt mapping race in IRQ core
 - Affects all shared interrupts
 - Two mappings for the same interrupt line
 - Touchpad interrupts dismissed as spurious
 - Irqdomain locking fixes merged in 6.4

Asynchronous probe woes

- Devices can be probed in parallel
- Interrupt mapping race in IRQ core
 - Affects all shared interrupts
 - Two mappings for the same interrupt line
 - Touchpad interrupts dismissed as spurious
 - Irqdomain locking fixes merged in 6.4
- "Shared" resources claimed temporarily
 - Check existence before acquiring resources
 - Interrupt regression in 6.6-rc1
 - Problem: A reset GPIO can not be shared

Asynchronous probe woes

- Devices can be probed in parallel
- Interrupt mapping race in IRQ core
 - Affects all shared interrupts
 - Two mappings for the same interrupt line
 - Touchpad interrupts dismissed as spurious
 - Irqdomain locking fixes merged in 6.4
- "Shared" resources claimed temporarily
 - Check existence before acquiring resources
 - Interrupt regression in 6.6-rc1
 - Problem: A reset GPIO can not be shared
- Some resources claimed before HID driver probes
 - Move pin configuration to parent bus node (hack)
 - Or allow sharing?

Asynchronous probe woes

- Devices can be probed in parallel
- Interrupt mapping race in IRQ core
 - Affects all shared interrupts
 - Two mappings for the same interrupt line
 - Touchpad interrupts dismissed as spurious
 - Irqdomain locking fixes merged in 6.4
- "Shared" resources claimed temporarily
 - Check existence before acquiring resources
 - Interrupt regression in 6.6-rc1
 - Problem: A reset GPIO can not be shared
- Some resources claimed before HID driver probes
 - Move pin configuration to parent bus node (hack)
 - Or allow sharing?
- Extend devicetree specification?
 - Mark devices as mutually exclusive
 - Force sequential probe

Bootloader handover

- Resources may have been (left) enabled by boot firmware
 - Clocks
 - Regulators
 - Power domains
 - Interconnects
- Handover should be seamless (e.g. EFI framebuffer)
- Disable any resources not needed to save power

Handover problems

- Unused clocks and power domains disabled at late init
 - Breaks display handover when display driver built as module
 - Workaround: `clk_ignore_unused` and `pd_ignore_unused`
 - Ignoring wastes power
 - Fails to disable unused clocks for modular clock drivers

Handover problems

- Unused clocks and power domains disabled at late init
 - Breaks display handover when display driver built as module
 - Workaround: `clk_ignore_unused` and `pd_ignore_unused`
 - Ignoring wastes power
 - Fails to disable unused clocks for modular clock drivers
- Unused regulators disabled 30 seconds after late init
 - Breaks early console unless display driver in `initramfs`
 - Example: Rescue shell or full-disk encryption

Handover problems

- Unused clocks and power domains disabled at late init
 - Breaks display handover when display driver built as module
 - Workaround: `clk_ignore_unused` and `pd_ignore_unused`
 - Ignoring wastes power
 - Fails to disable unused clocks for modular clock drivers
- Unused regulators disabled 30 seconds after late init
 - Breaks early console unless display driver in `initramfs`
 - Example: Rescue shell or full-disk encryption
- Interconnect state synced when all consumers' drivers are bound
 - Power state left at maximum if driver fails to probe
 - Example: Firmware missing (video acceleration)
 - Example: Driver disabled

Disabling unused resources

- Need common mechanism for disabling unused resources
- Could be built on device links and sync-state mechanism, which triggers when all consumers' drivers are bound
- Also need timer for consumers that fail to probe?
- Timer extended when a driver is bound (cf. deferred probe)
- Trigger state sync on expiry
- Problem: Need to pause timer in rescue shell
- Alternative: User space triggers sync when booted
 - Use existing `state_synced` sysfs interface?

Future work

- Camera ISP
- Bluetooth BD_ADDR and Wi-Fi MAC address
- Hibernation
- Keyboard special keys (e.g. mic mute)
- Thermal throttling, non-CPU (e.g. GPU, DSP, charger, radio)
- Trusted Platform Module (TPM)
- USB Power Delivery (USB-PD)
- Virtualisation

Future work

- ~~Camera-ISP~~
- Bluetooth BD_ADDR and Wi-Fi MAC address
- Hibernation
- Keyboard special keys (e.g. mic mute)
- Thermal throttling, non-CPU (e.g. GPU, DSP, charger, radio)
- Trusted Platform Module (TPM)
- USB Power Delivery (USB-PD)
- ~~Virtualisation~~

Getting started with X13s

- Linux 6.5 (mainline or WIP)
 - <https://github.com/jhovold/linux/tree/wip/sc8280xp-v6.5>
 - `make johan_defconfig`
- `linux-firmware-20230919`
 - `ln -s a660_gmu.bin a690_gmu.bin`
 - `qcvss8280.mbn` (with WIP branch)
- `alsa-ucm-conf 1.2.10` + volume fixes
 - <https://github.com/alsa-project/alsa-ucm-conf/pull/335>
- Mesa 23.1.4
- Qualcomm `pd-mapper` and `qrtr-ns` daemons
 - <https://github.com/andersson/pd-mapper>
 - <https://github.com/andersson/qrtr>

Kernel command line

- `clk_ignore_unused` (generic issue)
- `pd_ignore_unused` (generic issue)
- `arm64.nopauth` (firmware bug)
- `efi=noruntime` (firmware bug)
- `iommu.passthrough=0 iommu.strict=0` (USB throughput)
- `pcie_aspm.policy=powersupersave` (PCIe ASPM)

Initramfs

- Modules needed to boot
 - `nvme`, `phy-qcom-qmp-pcie`, `pcie-qcom`
- Modules needed for early console (optional)
 - `i2c_hid_of`, `i2c-qcom-geni`
 - `leds-qcom-lpg`, `pwm-bl`
 - `qrtr`, `pmic-glink-altmode`, `gpio-sbu-mux`, `phy-qcom-qmp-combo`
 - `panel-edp`, `msm`, `phy-qcom-edp`
- Depends on config
- Documented in commit message for `johan_defconfig`
- Missing display dependency can break EFI framebuffer

Distro support

- <https://launchpad.net/~ubuntu-concept/+archive/ubuntu/x13s>
- https://fedoraproject.org/wiki/Thinkpad_X13s
- <https://wiki.debian.org/InstallingDebianOn/Thinkpad/X13s>

Summary

- Lenovo ThinkPad X13s well supported by mainline Linux 6.5
- More features under way (e.g. camera)
- Working on improving suspend power consumption

- For updated status, known issues and other resources, see:
 - <https://github.com/jhovold/linux/wiki/X13s>

Thanks!

johan@hovoldconsulting.com

johan@kernel.org