

# Abstracting platform resources using power & performance domains over SCMI

Nikunj Kela Engineer, Senior Staff Qualcomm Innovation Center, Inc.

#### Agenda

- Motivation
- Achieving abstraction using SCMI
- Our architecture overview
- Status and plan



#### Resources, a typical peripheral driver interface with



& Linaro Connect Madrid 2024

#### What are the problems?

- The Linux drivers become too complex
  - Need to interface with clock, regulator, interconnect, GPIO/Pin-control, PHY etc. frameworks
- Too many platform drivers needed for a single peripheral to work
  - Significant development efforts
  - Significant bring up efforts
  - Significant upstreaming efforts
  - Difficult to maintain so many drivers

🗞 Linaro Connect 🗾 🛛

## SCMI(System Control and Management Interface)

- A set of operating system-independent software interface that are used in system management
- SCMI is split into two layers
  - o Protocols
    - The services you want to provide
    - e.g., clock, voltage, power-domain, reset, performance etc.
  - o Transports
    - The mechanism by which the parameters/results are communicated between the caller of the interface(an Agent) and the implementer (the platform)
    - e.g. mailbox, smc/hvc call, virtio etc.



#### SCMI to achieve abstraction



🗞 Linaro Connect 🗾

#### Can we do better?



#### Our approach

- Define logical/modelled power and performance domains
- Peripheral drivers can now request Firmware to configure device Performance Level thus avoiding multiple requests for individual resources
- Use SCMI power and performance protocols to achieve
  - Better abstraction (Interconnect Bandwidth and PHY can also be abstracted)
  - Reduce # of requests to Firmware



🗞 Linaro Connect

## Continue...(Modelled performance domains)

- In Linux, the performance domains have been used for compute(CPU) devices only so far
- We extend its usage to peripheral devices
- Define performance domains to perform
  - Dynamic clocks voting i.e. set\_rate operations
  - Dynamic interconnect bandwidth voting
  - Dynamic changes to PHY modes
  - Dynamic changes to baud rates in case of UART
- Define performance levels for the peripherals
  - For baud-rates in UART, perf levels can be defined as
    - LEVEL\_9600, LEVEL\_115200, etc.
  - What these levels translate to is hidden in the Firmware



source: wikimedia commons license: <u>CC BY-SA 2.0 Deed | Attribution-ShareAlike 2.0 Generic | Creative Commons</u> no changes were made



## Continue...(Modelled power domains)





source: wikimedia commons license: <u>CC BY-SA 2.0 Deed | Attribution-ShareAlike 2.0 Generic | Creative Commons</u> no changes were made

🗞 Linaro Connect

#### Code Flow on Linux



## Changes pushed in Linux Kernel

- SCMI performance driver is available to work with peripheral devices by extending genpd framework [1]
- OPP framework is extended to work with SCMI performance driver [2]
- Helper APIs are now available so the consumer drivers can avoid the manual attach/detach in case of multiple power-domains [3]
  - [1]: https://lore.kernel.org/all/20230825112633.236607-1-ulf.hansson@linaro.org/
  - [2]: https://lore.kernel.org/all/20230925131715.138411-1-ulf.hansson@linaro.org/
  - [3]: https://lore.kernel.org/all/20240130123951.236243-1-ulf.hansson@linaro.org/

#### Our Architecture – Firmware Managed Resources

- Dedicated SCMI channels (shmem + doorbells) for each peripheral device
- FW VM has multiple vCPUs; the requests can be served in parallel



## How to figure fw-managed setup in drivers

- Since we are using power and perf protocol(and not clocks, voltage etc.), peripheral drivers need to know if running in a fw-managed setup or locally managed setup
- Check on the presence/absence of the clocks(and/or other resources) in DT 1.
  - **Pros**: don't need anything additional 0
  - **Cons**: may work for some driver; clock maybe optional on some platforms; may not scale; 0 insufficient condition to ascertain fw-managed setup
- Use a new compatible for each device(e.g. "qcom,ufshc-fw-managed") 2.
  - **Pros**: cleaner and explicit way to determine fw-managed setup 0
  - **Cons**: HW IP is not changed so may not be warranted; 0
- 3. Use vendor DT property for devices (e.g. "qcom,firmware-managed-resources")
  - **Pros**: cleaner and explicit way to determine fw-managed setup 0
  - **Cons**: repetition of the same property in device tree nodes; vendor specific so 3<sup>rd</sup> party 0 drivers may not accept it

Linaro Connect Madrid 2024

#### Continue...

- 4. Use a vendor agnostic DT property for devices (e.g. "firmware-managed-resources")
  - **Pros**: cleaner and explicit way to determine fw-managed setup
  - **Cons**: repetition of the same property in device tree nodes; different vendors may abstract differently
- 5. Use a board(platform)level compatible
  - **Pros**: common way to all devices
  - **Cons**: each driver will need to match for board(platform) level DT; might not scale
- Upstream discussion(not concluded) on this can be found [4]
- We will conform to upstream suggestions once the discussion concludes

[4]: <u>https://lore.kernel.org/all/be31801e-bb21-426b-f7aa-2b52727de646@quicinc.com/</u>

& Linaro Connect Madrid 2024

#### Status and plan

- Firmware VM side development is 90% complete
- Most of the Linux peripheral device driver changes are under internal review
- SCMI transport driver changes for Qualcomm hypervisor is available in upstream [5]
- We plan to send peripheral driver changes in upstream soon

[5]: https://lore.kernel.org/all/20231009191437.27926-1-quic\_nkela@quicinc.com/





## Thank you

MAD24-310