Energy model accuracy

Linux Plumbers Conference 2022

Power Management and Thermal Control

Morten Rasmussen

14/09/2022

© 2022 Arm

Motivation

- CPU power and performance characteristics can be represented in the Linux kernel by an Energy Model.

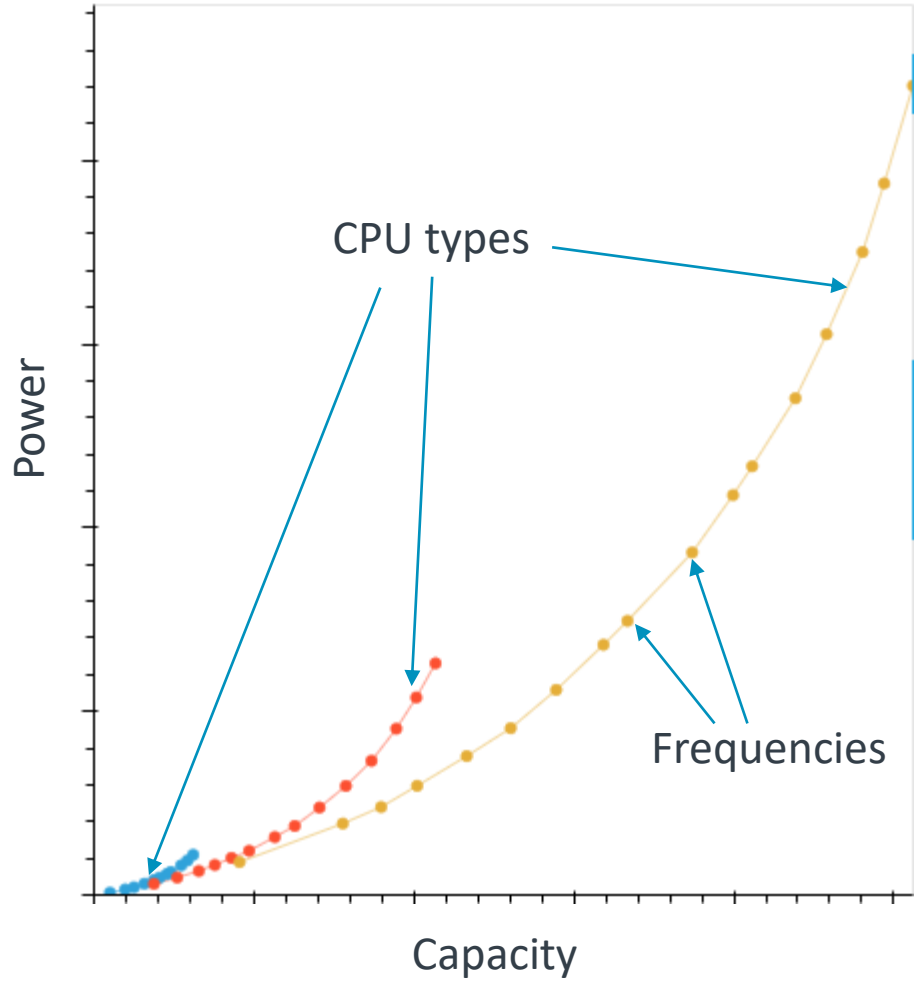

© 2022 Arm

Motivation

- Device power and performance characteristics can be represented in the Linux kernel by an Energy Model.
- History:
  - The Energy Model framework was originally introduced with Energy Aware Scheduling
  - Enables the scheduler to make better scheduling decisions by reasoning about relative energy efficiency.
  - Static data assuming insignificant variation with factors such as workload (instruction mix) and temperature.
  - Introduced to be better than no information at all, so error margins could be tolerated.
- Today:
  - Demand for accurate energy predictions driven by more choices (big.LITTLE -> 3-gear) and temperature impact.
    - The 20% error margins have been reduced significantly in the meantime.

© 2022 Arm

Measured energy model

- Actual power and performance measurement on single 3-gear device for four benchmarks.
  - Each coloured curve represents the DVFS curve of each CPU type.
  - Significant power and performance variation with workload.

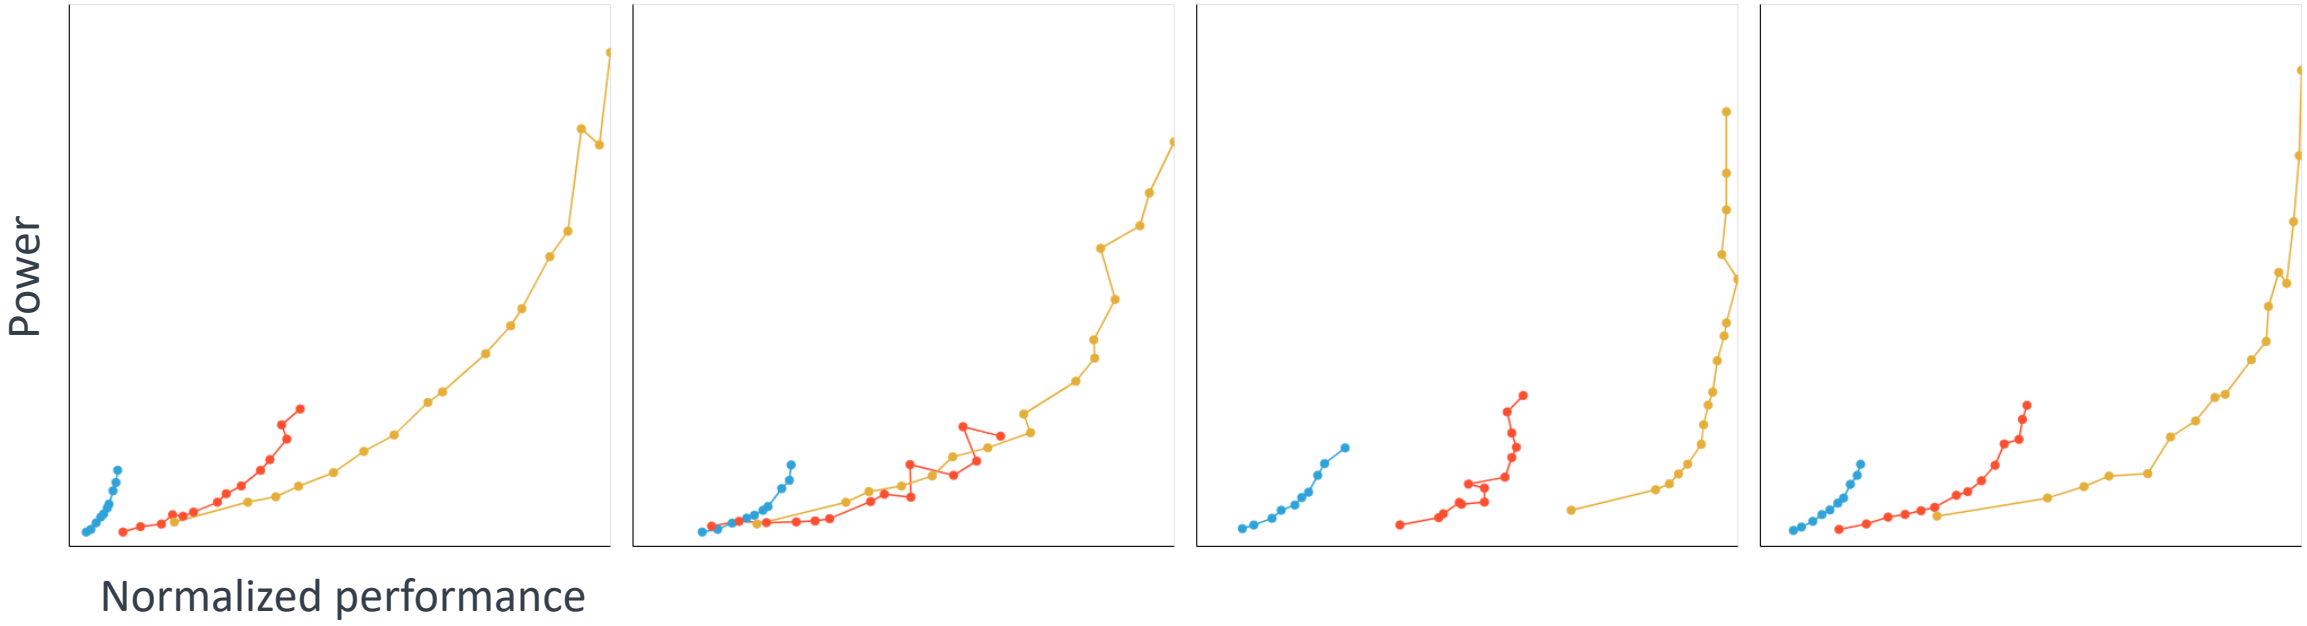

© 2022 Arm

Measured energy model

- Actual power and performance measurement on single 3-gear device for four benchmarks.
  - Each coloured curve represents the DVFS curve of each CPU type.
  - Significant power and performance variation with workload.

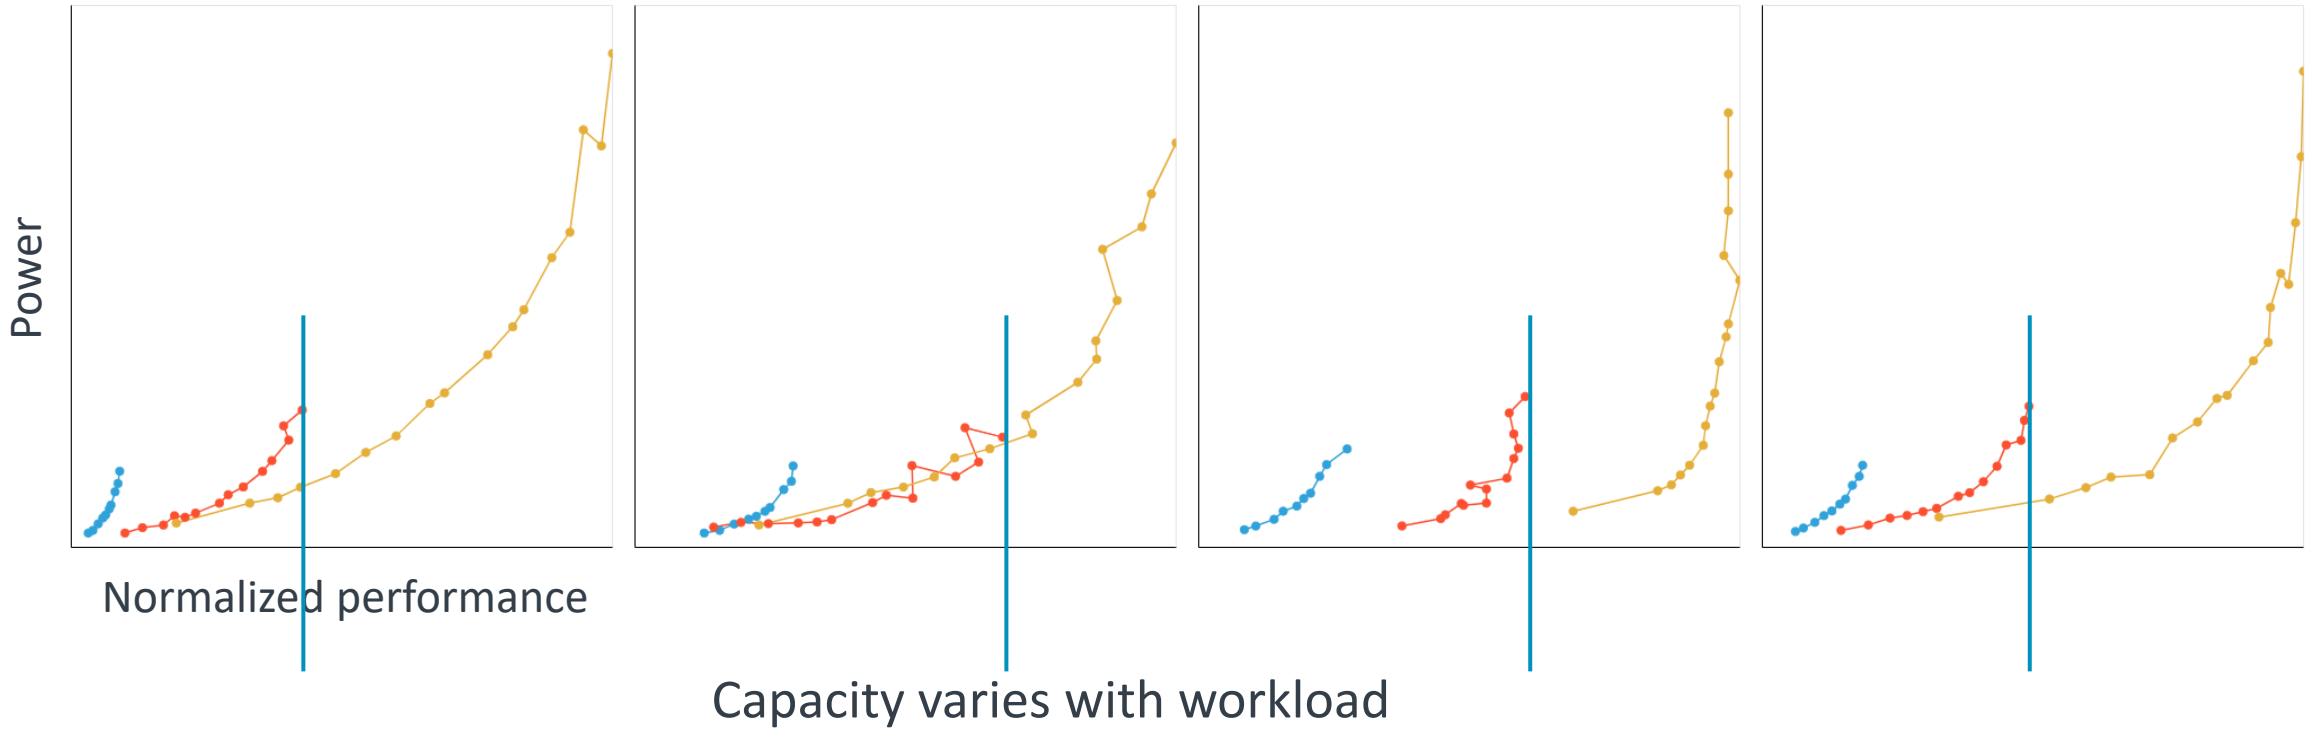

© 2022 Arm

Measured energy model

- Actual power and performance measurement on single 3-gear device for four benchmarks.
  - Each coloured curve represents the DVFS curve of each CPU type.
  - Significant power and performance variation with workload.

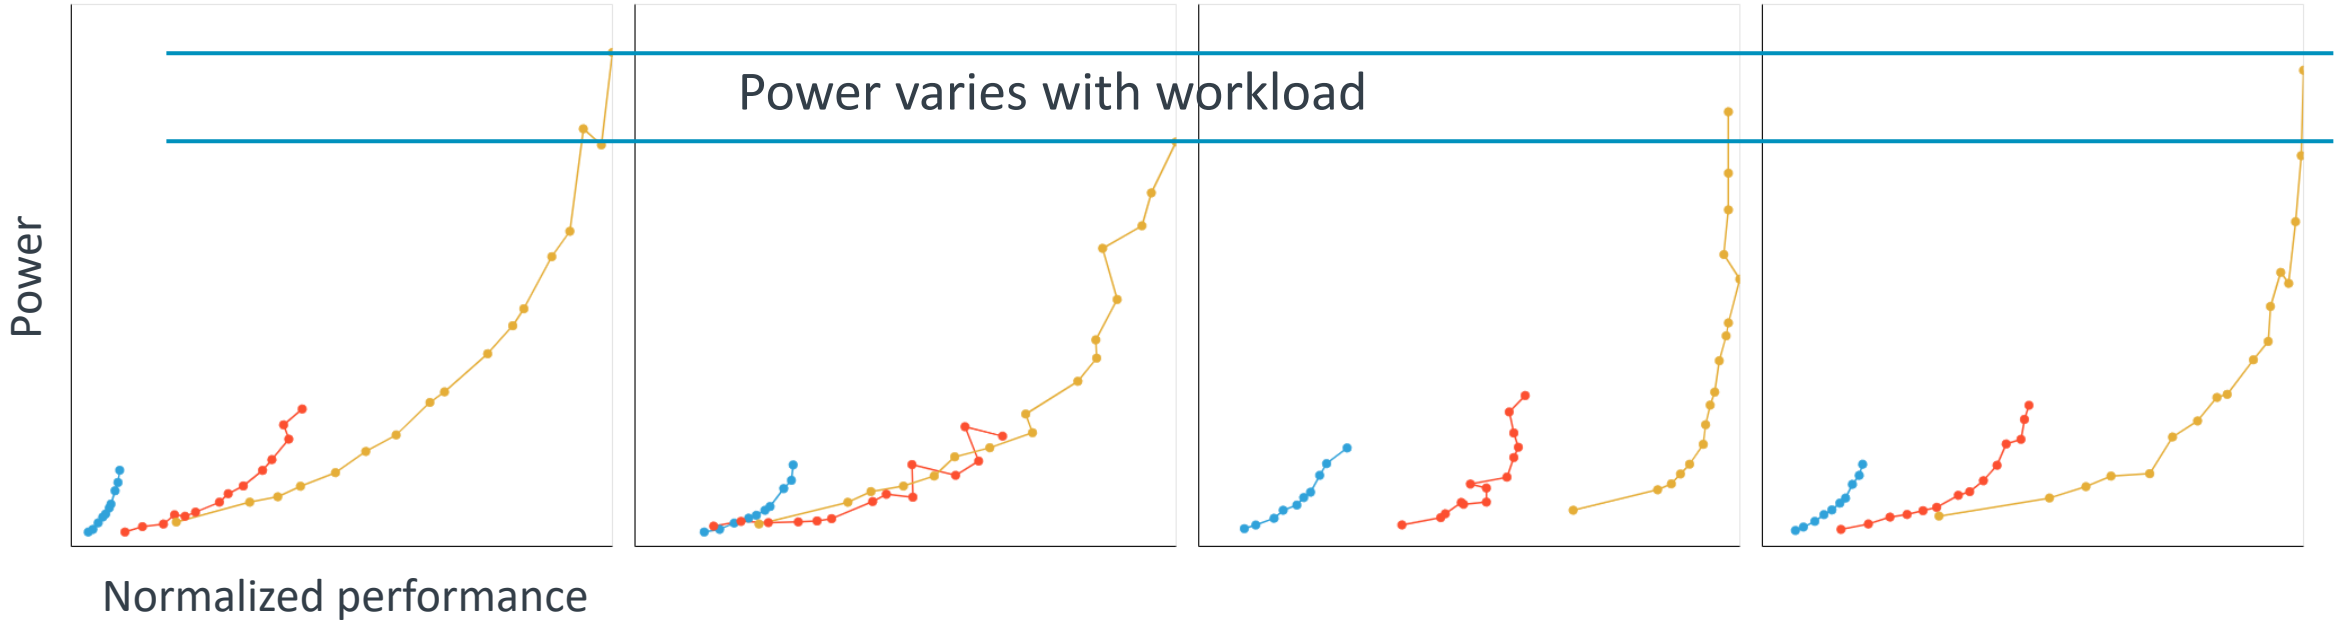

© 2022 Arm

Modifying the energy model dynamically

- The EM is essential for EAS to reason about performance domains (DVFS) and consequences of task placement.
- Can we dynamically update and/or overlay the EM to enhance EAS predictions?
  - Multiple EMs selected by user-space?
    - Only one active at the time, not helping mixed workloads but simpler to implement.
  - Temperature correction factor or dynamic update?
    - Discussed yesterday in Android MC.
  - Per-task workload correction factor?
    - Single default EM but data modified using correction factor at each use.
  - Other ideas?

© 2022 Arm

arm

Thank You
Danke
Gracias
Grazie
谢谢
ありがとう
Asante
Merci
감사합니다
धन्यवाद
Kiitos
شكرًا
ধন্যবাদ
תודה

© 2022 Arm

The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

www.arm.com/company/policies/trademarks

© 2022 Arm