Über die Bedeutung des Turbomodus aktueller Mehrkernprozessoren für Betriebssysteme

Jan Schönherr
Jan Richling

Kommunikations- und Betriebssysteme
Technische Universität Berlin

Gero Mühl

Architektur von Anwendungssystemen
Universität Rostock

Matthias Werner

Professur für Betriebssysteme
Technische Universität Chemnitz

Agenda

- Turbo mode
  - Definition
  - Implementations
- Handling by operating systems
- Currently unused opportunities
  - Performance
  - Energy efficiency
  - Multi-socket systems
- Summary

Turbo mode

- Ability to distribute power budget dynamically
  - (realized in hardware)

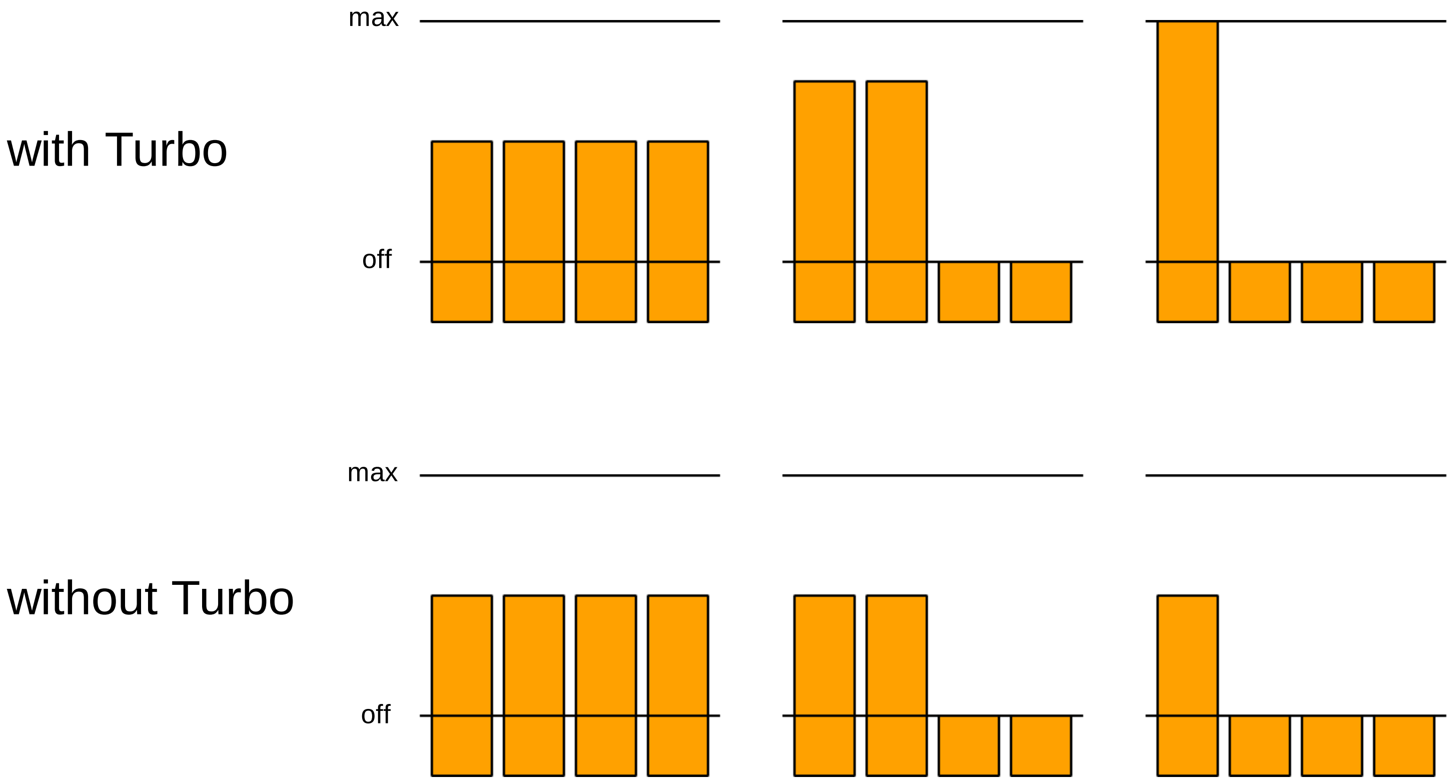

Intel’s Turbo Boost Technology

- Processor monitors itself
  - Power, current, temperature

- Highest frequency within physical margins is selected
  - Additionally: artificial limit on the number of active cores
  - Nominal freq. <= actual freq. <= Turbo freq.

- Transparent for OS
  - Cannot explicitly select frequencies above nominal frequency
  - Cannot (easily) determine actual frequency

  - On/Off via P-state mechanism

AMD’s Turbo Core

- Simple logic depending on the number of active cores
  - If enough cores are sleeping, the remaining ones are boosted
- No intermediate frequencies between nominal and turbo frequency
- On/off via MSR
- Only cores in highest OS visible P-state are boosted
- Easy to determine actual frequency
- Deterministic (compared to Intel)

Current OS support for Turbo

- None :-(

- Not entirely true
  - Both, Turbo Boost and Turbo Core, are enabled by existing mechanisms
  - Power saving logic increases effectiveness of Turbo
    - OSs try to keep individual cores idle (to some extend)

- Current situation
  - Turbo mode always utilized in lightly loaded scenarios
  - Turbo mode never (completely) utilized in fully loaded scenarios
  - CPUs reach their TDP limit more often
  - Normally reduced energy efficiency

Turbo for selected processes

- Easy for AMD processors
  - Enable Turbo
  - Put enough cores to sleep
  - Set some cores to the highest P-state
  - Set remaining cores to another P-state
- Advantages
  - Easy to realize in OSs
    - Can reuse existing affinity management
- Disadvantages
  - Only one voltage domain in AMD processors
    - Not most energy efficient solution for non-boosted cores

Turbo for selected processes

- In case of only one voltage and/or frequency domain
  - Group processes into sets
  - At some time
    - Enable turbo
    - Put enough cores to sleep
    - Execute set of selected processes
  - At another time
    - Disable turbo
    - Select an energy efficient frequency for the set
    - Execute the set

- Advantages
  - Most energy efficient solution under given performance constraints
- Disadvantages
  - Requires a new (or heavily modified) scheduler

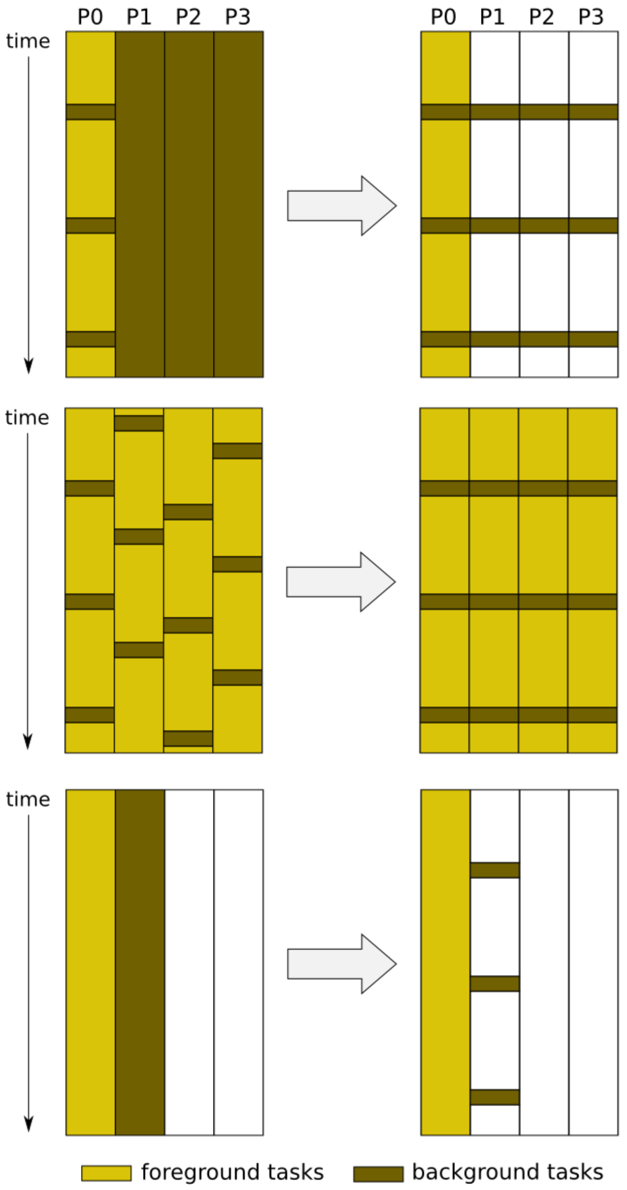

Preliminary evaluation

> Foreground tasks take full advantage of Turbo
> Enough background tasks compensate additional energy consumption

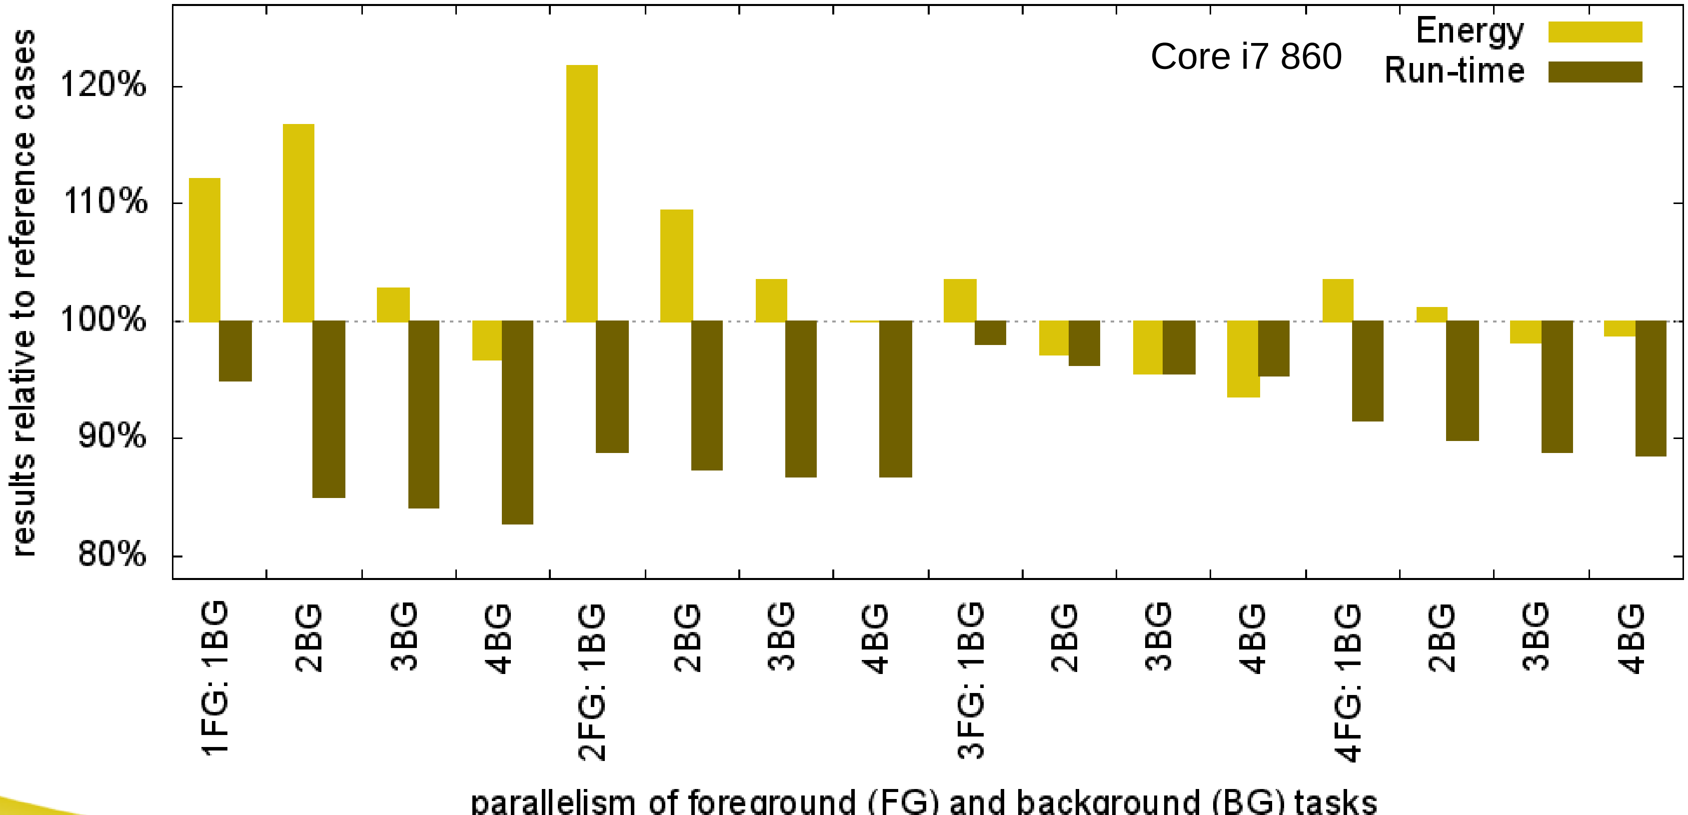

Energy efficient frequencies

- Energy efficient frequency depends on
  - Currently executed workload
  - Processor

- CPU-bound
  - Profits from higher frequency
  - Until voltage is raised beyond a certain level
- Memory-bound
  - Profits from lower frequencies
  - Until static power consumption prevails

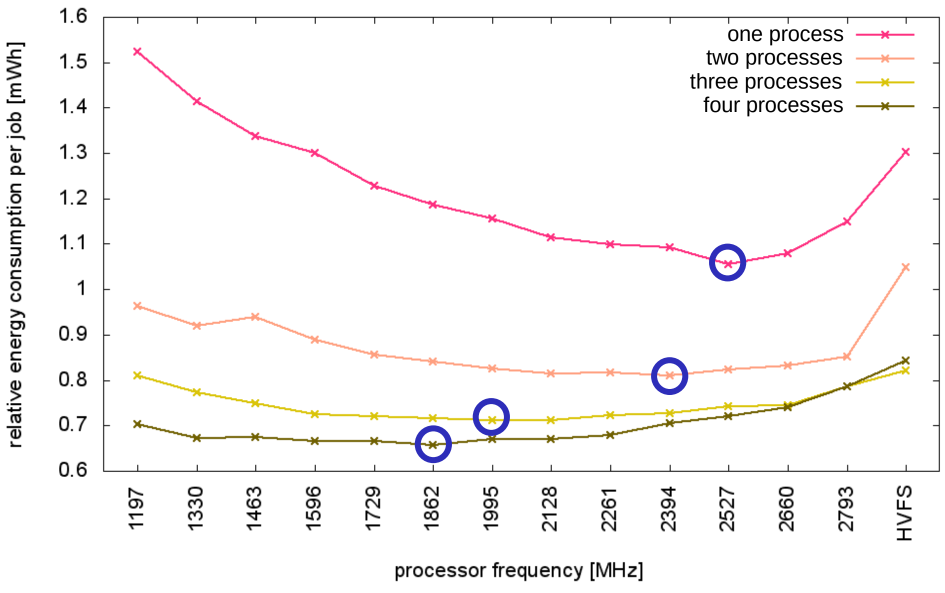

- Unverified theory: Processors of the same family exhibit similar behavior, i. e. the same energy efficient frequency

=> For some workloads and processors the energy efficient frequency is a turbo frequency

Multi-socket systems

- Additional scheduling considerations
  - Spread load to utilize Turbo?
  - Or concentrate load to conserve energy?

  - "Power" as a resource
    - Similar to cache or memory bandwidth

- (Intel) Processors are no longer symmetric
  - Some can hold a higher frequency in certain workloads than others

Summary

> Turbo mode achieves higher performance in many workloads out of the box – at the cost of a much lower energy efficiency

> Adequate handling can deliver even higher performance in some workloads and regain energy efficiency

> Current and future work
  > Adequate Turbo support for Linux
  > Application <=> OS interfaces