

Mixed-Criticality Systems with Partial Lockdown and Cache Reclamation Upon Mode Change

K. Bletsas^{*†}, M. A. Awan^{*†}, P. F. Souto^{*‡}, B. Åkesson^{*†} and E. Tovar^{*†}


* CISTER/INESC-TEC Research Centre, Porto, Portugal

† ISEP, Polytechnic Institute of Porto, Portugal

‡ University of Porto, Faculty of Engineering, Portugal

The classic Vestal model with mode changes

- In each mode, only tasks of a respective criticality or higher execute.
- Different WCET estimates for the same task in different modes.
 - By techniques with corresponding confidence levels.
- When a task overruns its WCET estimate for that mode, a mode change occurs (e.g., $L \rightarrow H$, with two modes).
- Essentially, the processor resources originally for L-tasks are repurposed for the H-tasks upon mode change.



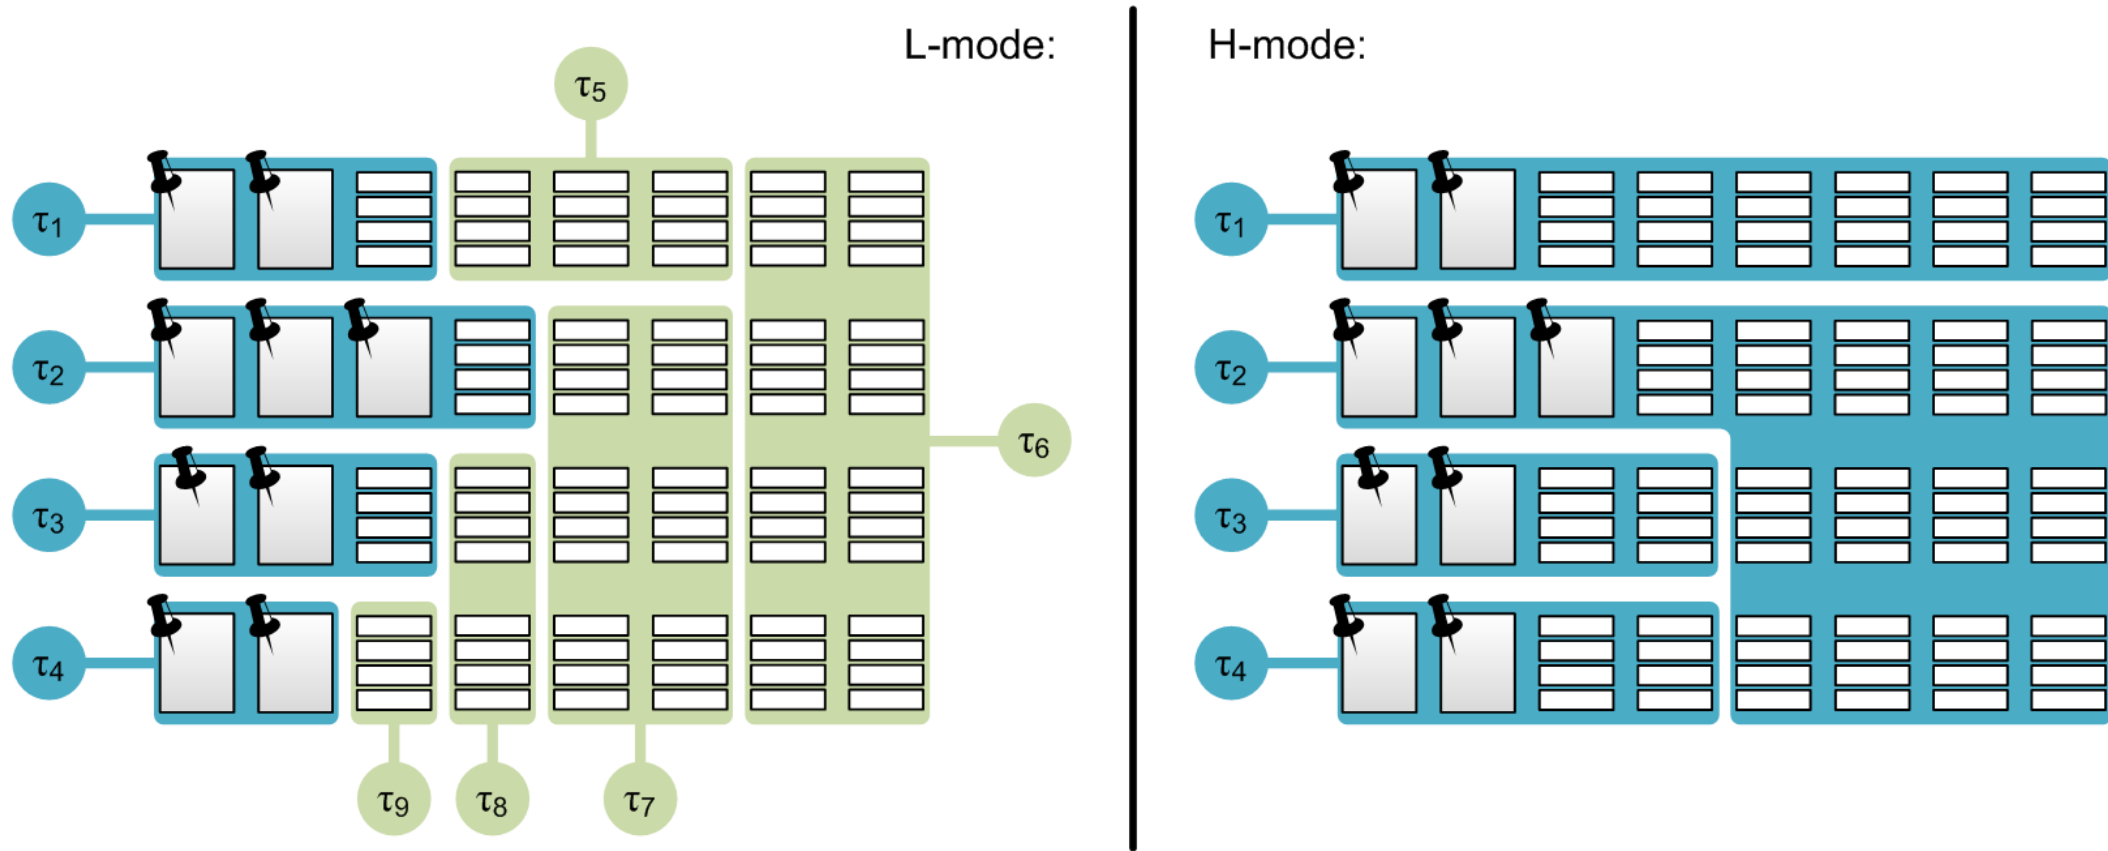
Idea: Reclaim more kinds of resources at mode change!

- Explored in our paper at the main track
- Resource targeted: the shared last-level cache of a multicore.
- The cache is partitioned to the tasks
- At mode change the cache partitions of the L-tasks are reclaimed and given to the H-tasks as additional resources.
 - The cache partitions can either be used to lock the hottest pages of the respective task in place (e.g., using Colored Lockdown);
 - Or they can be used dynamically by the task (e.g., under LRU policy).
- This WiP paper considers a more refined arrangement.

Partial Lockdown

- H-tasks:
 - Part of their cache partition is used for locked hot pages.
 - The rest of an H-task's cache partition is used dynamically
- L-tasks:
 - Their partitions are only used dynamically.
- At mode change:
 - The cache reclaimed from the L-tasks is used to enlarge the H-tasks' dynamically used partitions.

Illustration of the arrangement



Tradeoffs and challenges

- WCET estimates depend on the mode and cache resources assigned.
 - Tractable and accurate parametric WCET estimation needed.
- How much cache per task to use for locking hot pages and how much to use dynamically?
 - Interesting and complex tradeoffs in terms of derivable WCET estimates, that extend to schedulability analysis.
- Reclaiming only dynamically used cache limits mode change overheads.
 - No unlocking, bringing in new page and locking anew needed.