Hochschule
München
University of
Applied Sciences

Department of Computer
Science and Mathematics

How accurate are current tools and models for estimating software energy consumption?

Prof. Dr. Andreas Brunnert

Software Energy Consumption

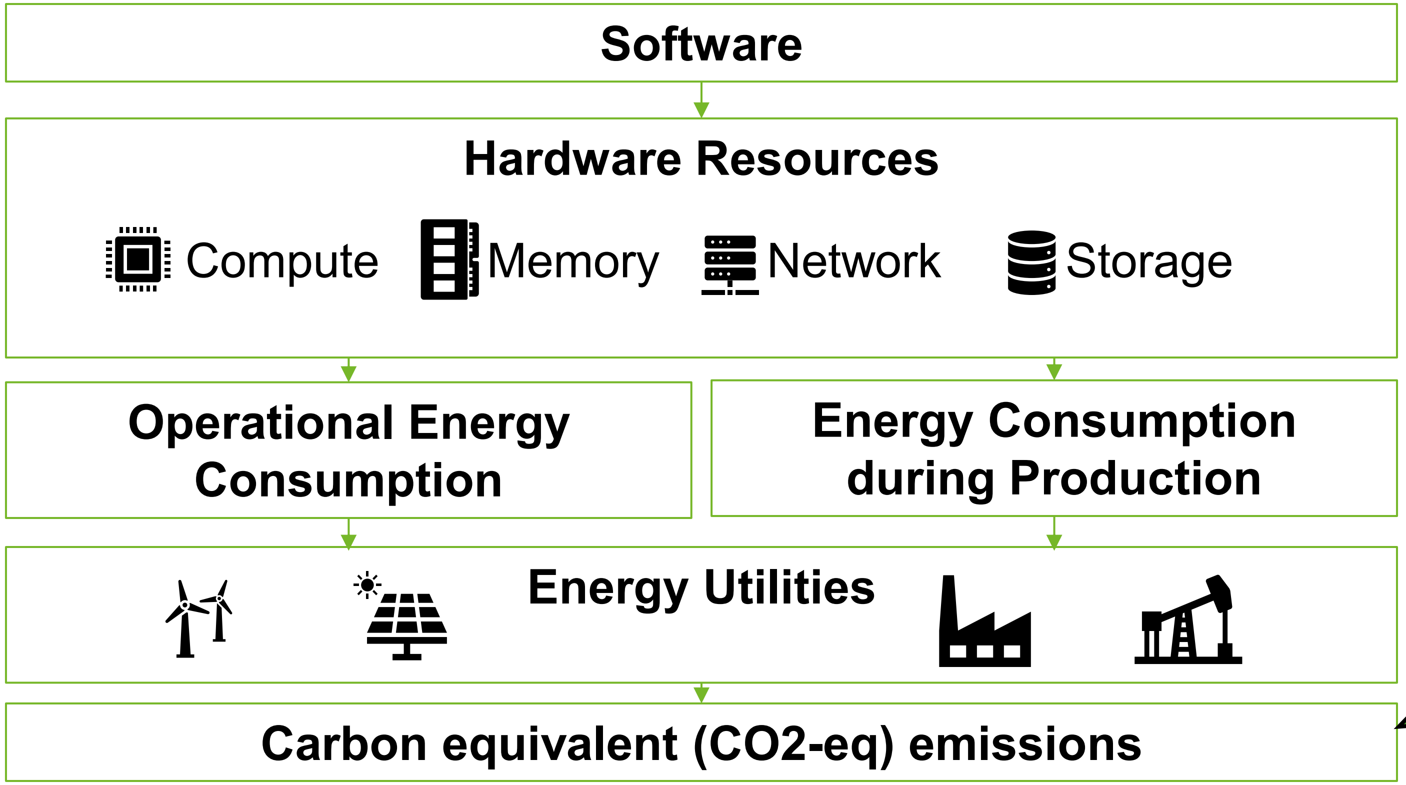

Focus of this talk

If you get energy wrong, the indirect effects derived from it, such as carbon emissions, will also be incorrect.

Software Energy Consumption

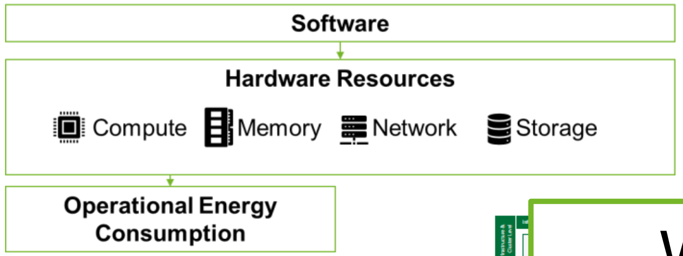

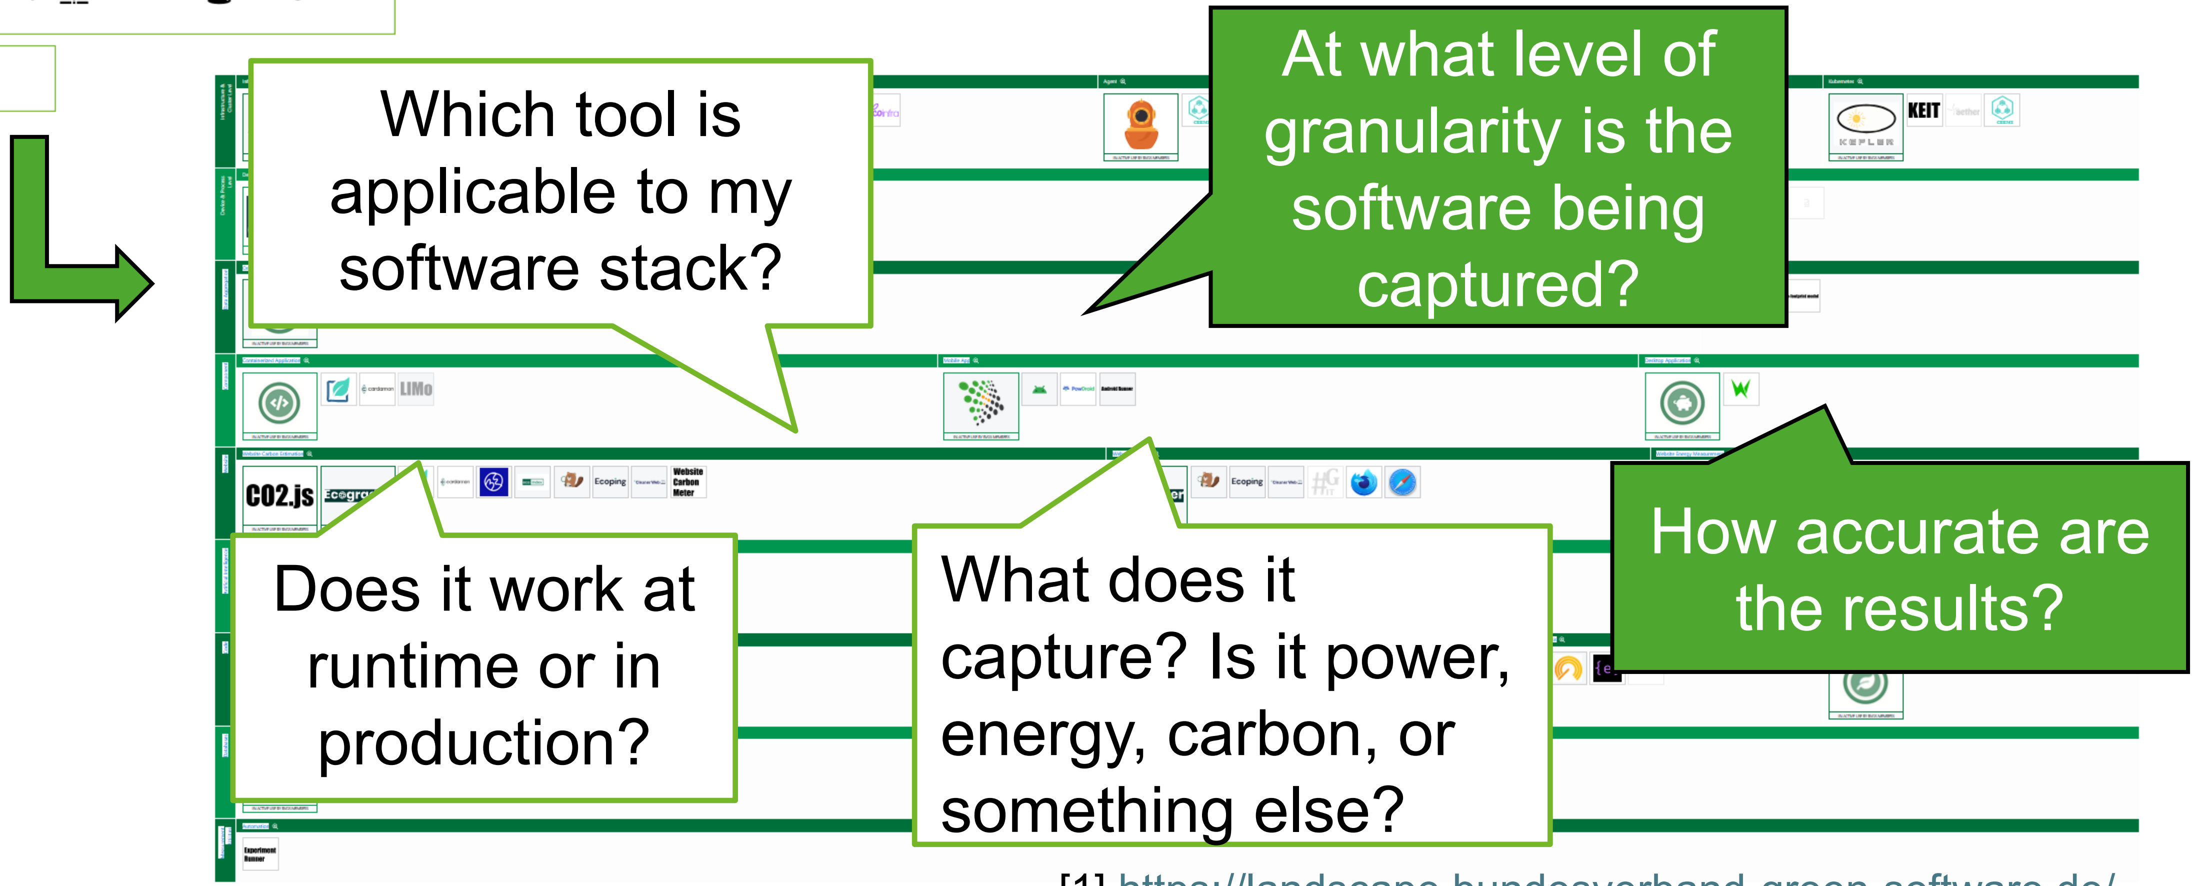

[1] https://landscape.bundesverband-green-software.de/

Software Energy Consumption

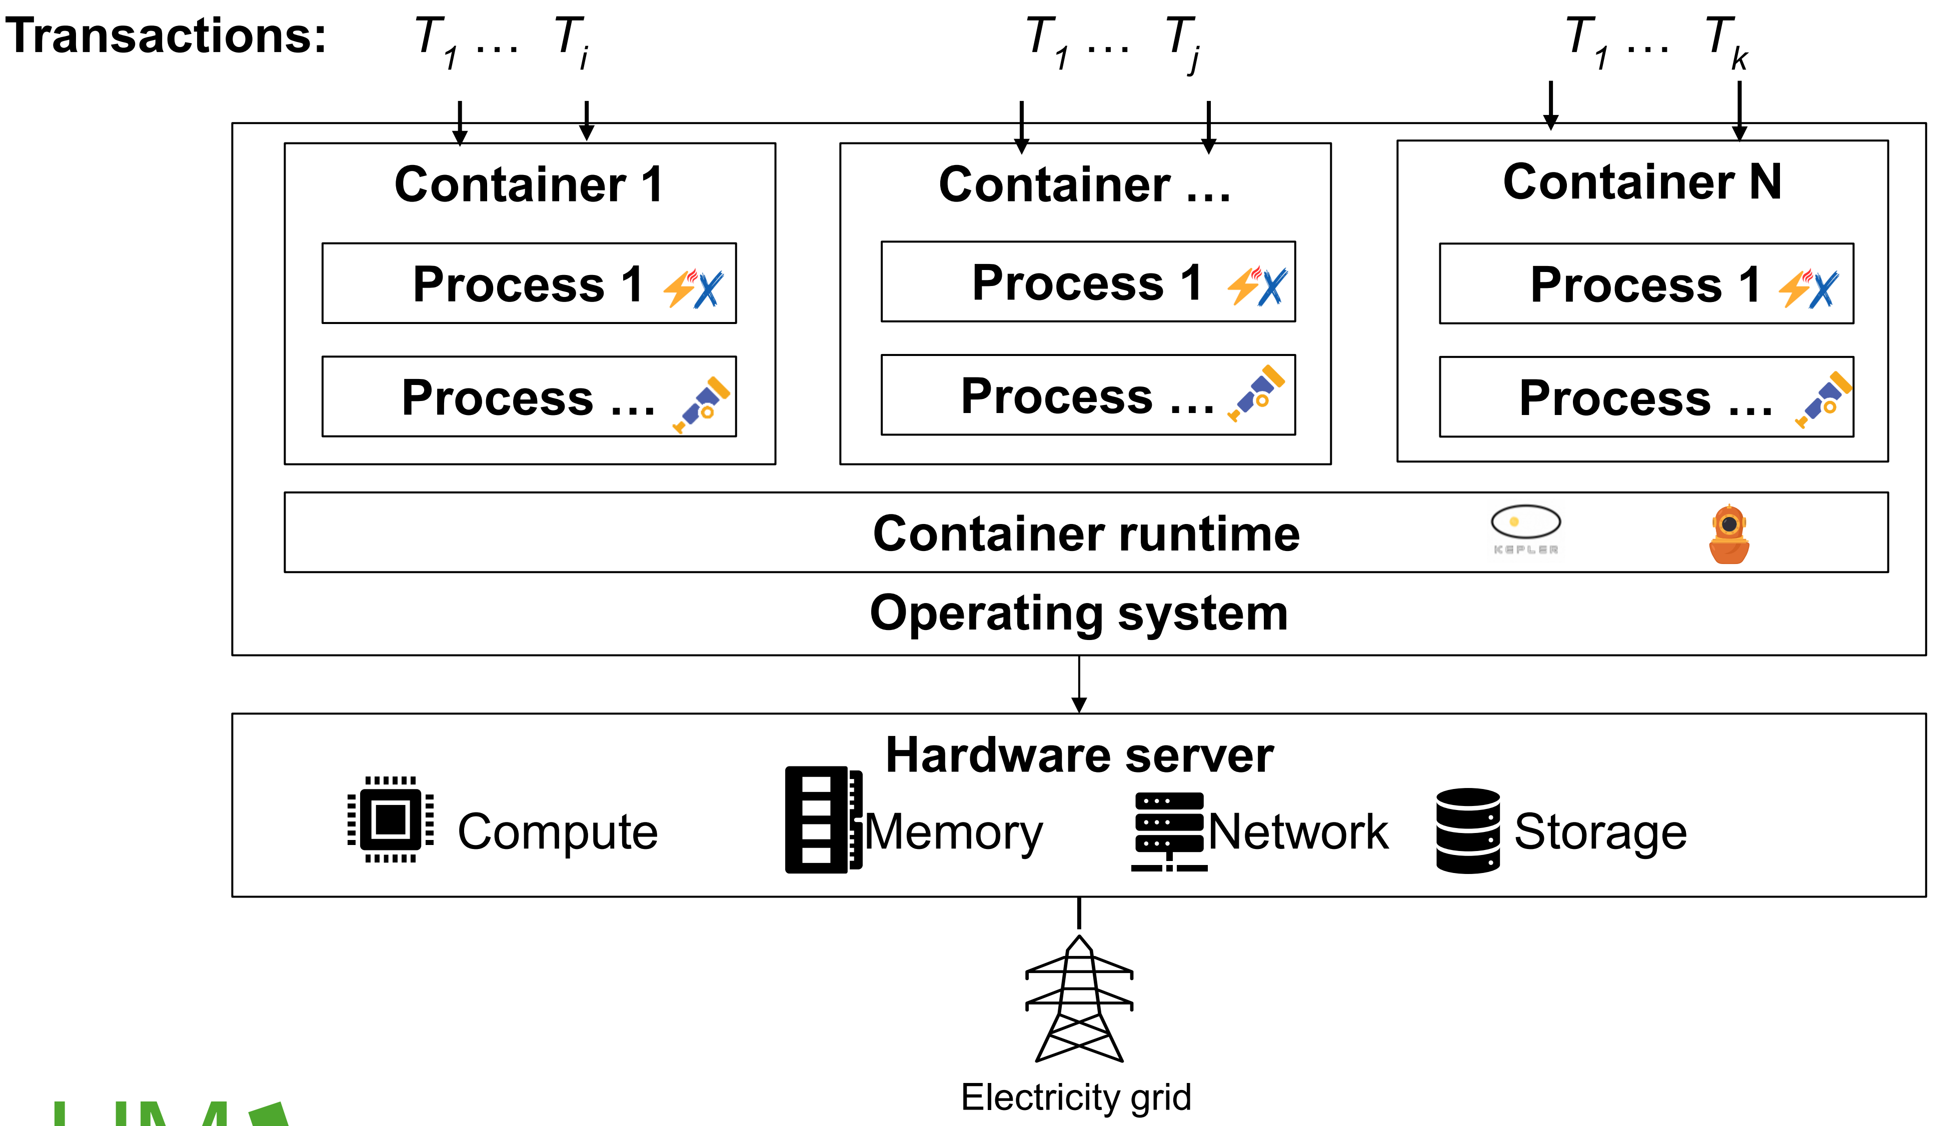

Transaction-level attributions (e.g., using JoularJX [5] or the Open-Telemetry Java Agent Extension (OTJAE) [6])

Process-level attributions (e.g., using Kepler [3], OTJAE [6], JoularJX [5] or Scaphandre [5])

Container-level attributions (e.g., using Kepler [3] or Scaphandre [4])

Hardware measurements (e.g., using hardware interfaces such as IPMI, Intel RAPL or Nvidia NVML)

External measurements (e.g., using energy meters between the power plug and the server)

Software Energy Consumption

JoularJX attributes RAPL power values to individual transactions.

OpenTelemetry Java Agent Extension (OTJAE): Derives transaction power consumption solely based on formulas.

Transaction-level attributions (e.g., using JoularJX [5] or the Open-Telemetry Java Agent Extension (OTJAE) [6])

Kepler, Scaphandre, and JoularJX attribute RAPL power values to individual processes

OpenTelemetry Java Agent Extension (OTJAE): Derives process power consumption solely based on formulas.

Process-level attributions (e.g., using Kepler [3], OTJAE, JoularJX [5] or Scaphandre [4])

Kepler and Scaphandre: attribute RAPL power values to individual containers

Container-level attributions (e.g., using Kepler [3] or Scaphandre [4])

Intel RAPL (accessible via Powercap [2]) provides continuous energy readings for the processor sockets.

Hardware measurements (e.g., using hardware interfaces such as IPMI, Intel RAPL or Nvidia NVML)

Continuous power readings for each power plug
→ Ground truth

External measurements (e.g., using energy meters between the power plug and the server)

Experiment Setup

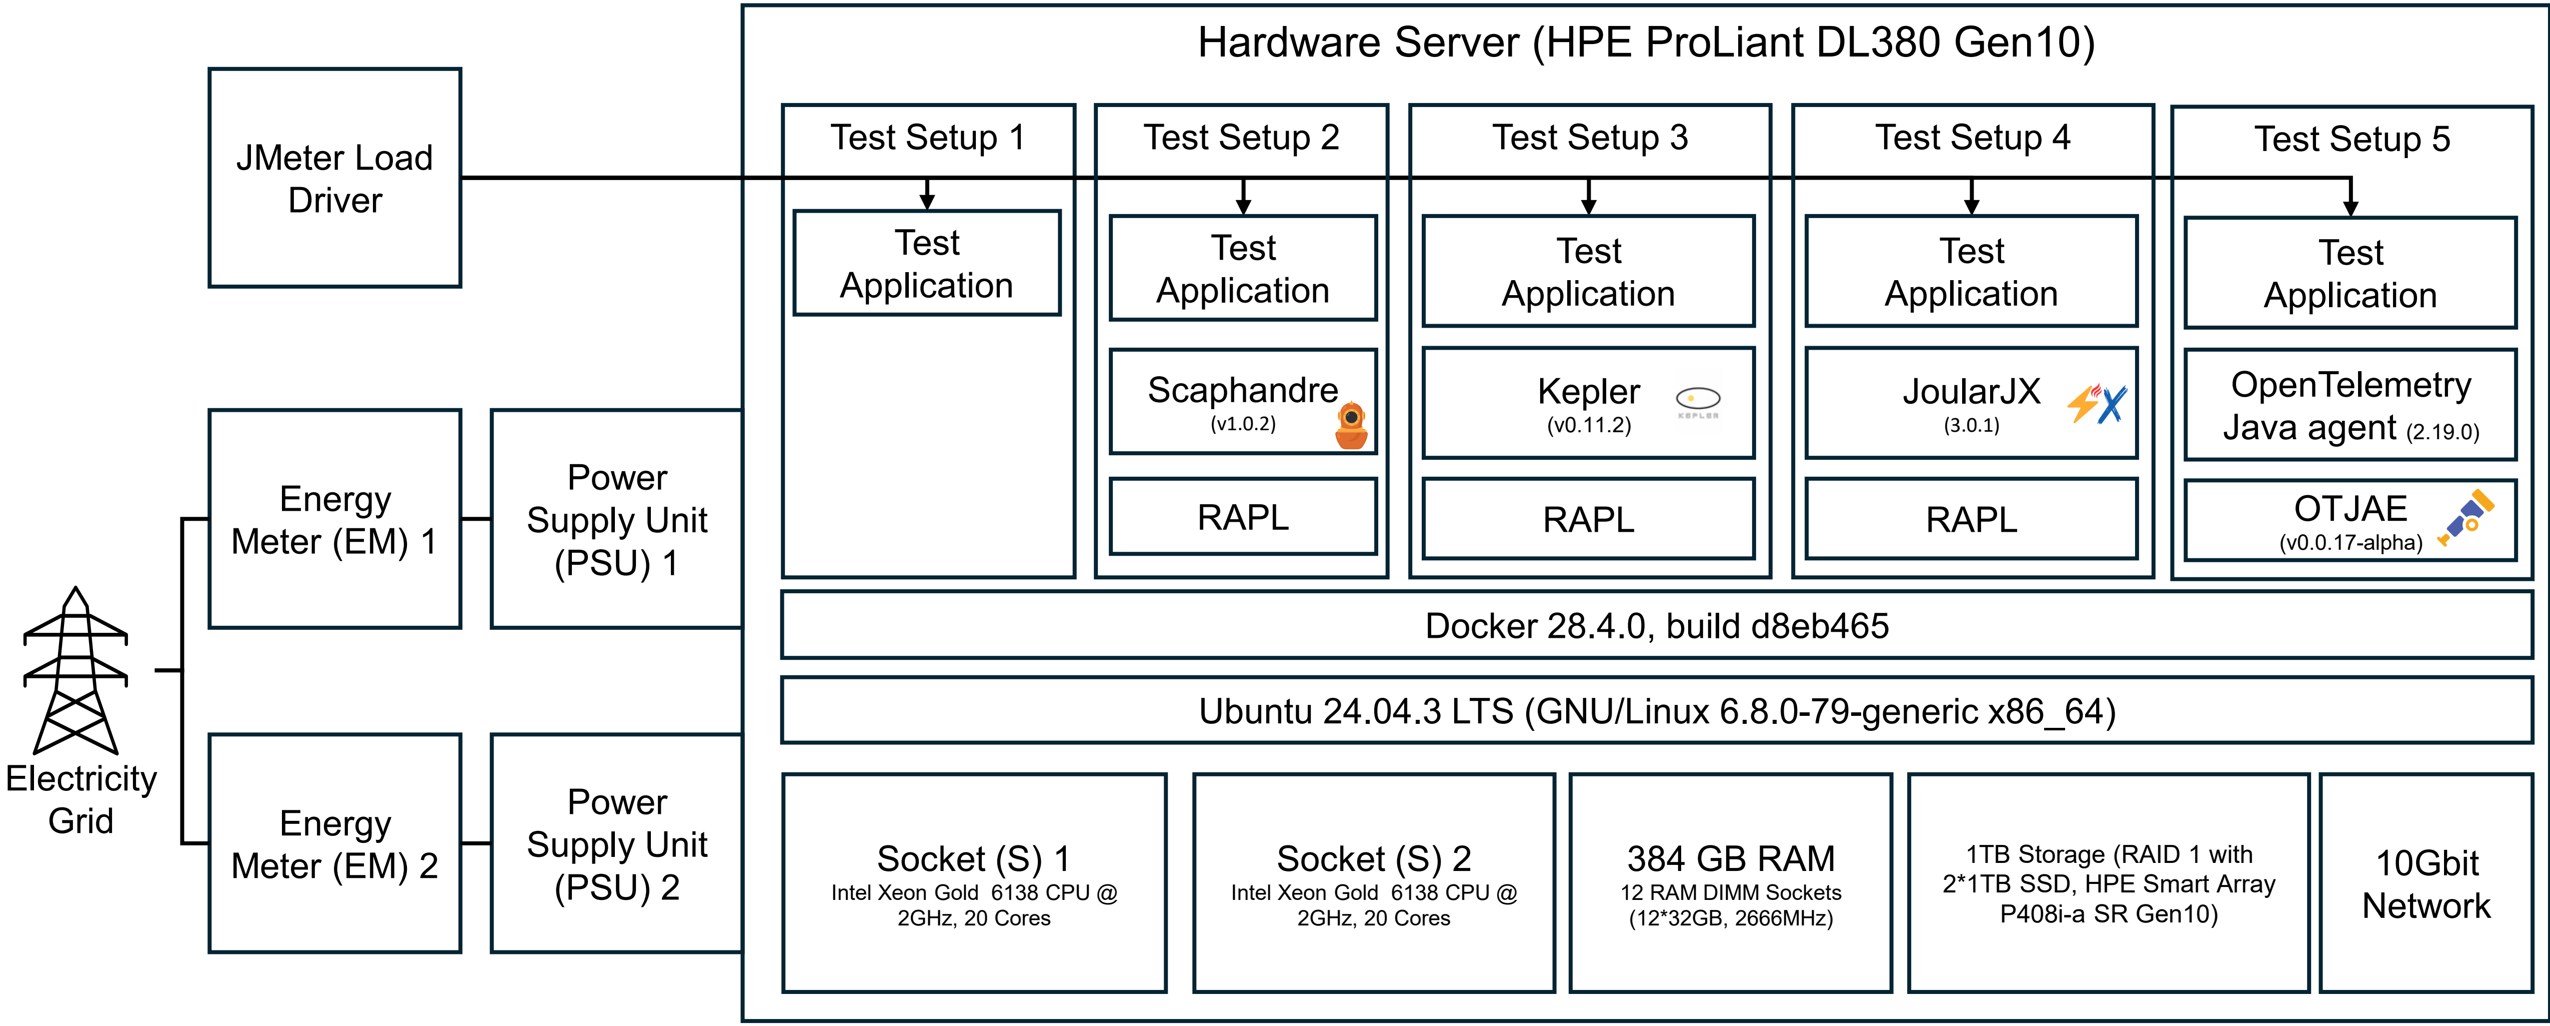

Experimental Results – External vs. Internal Power

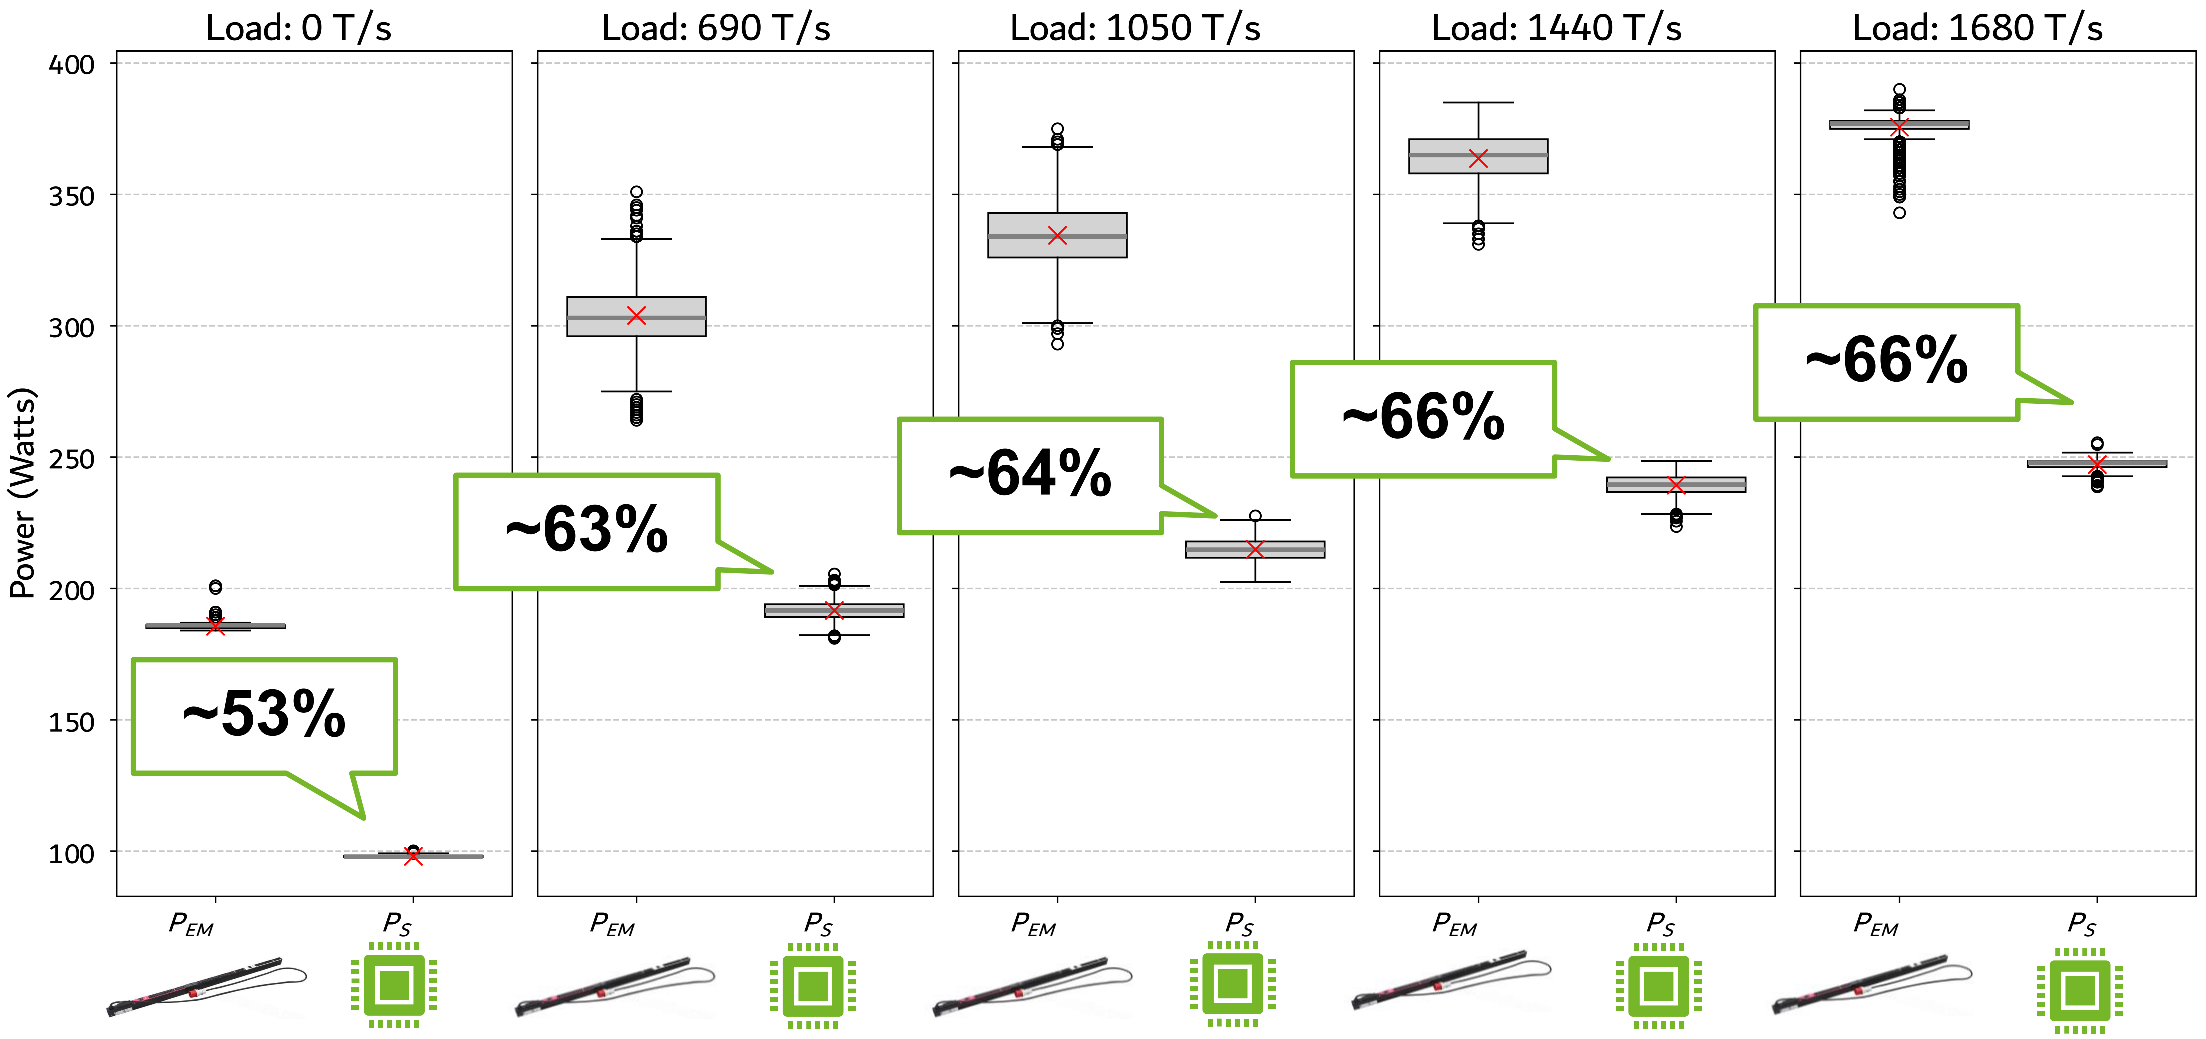

Experimental Results – External vs. Internal Power

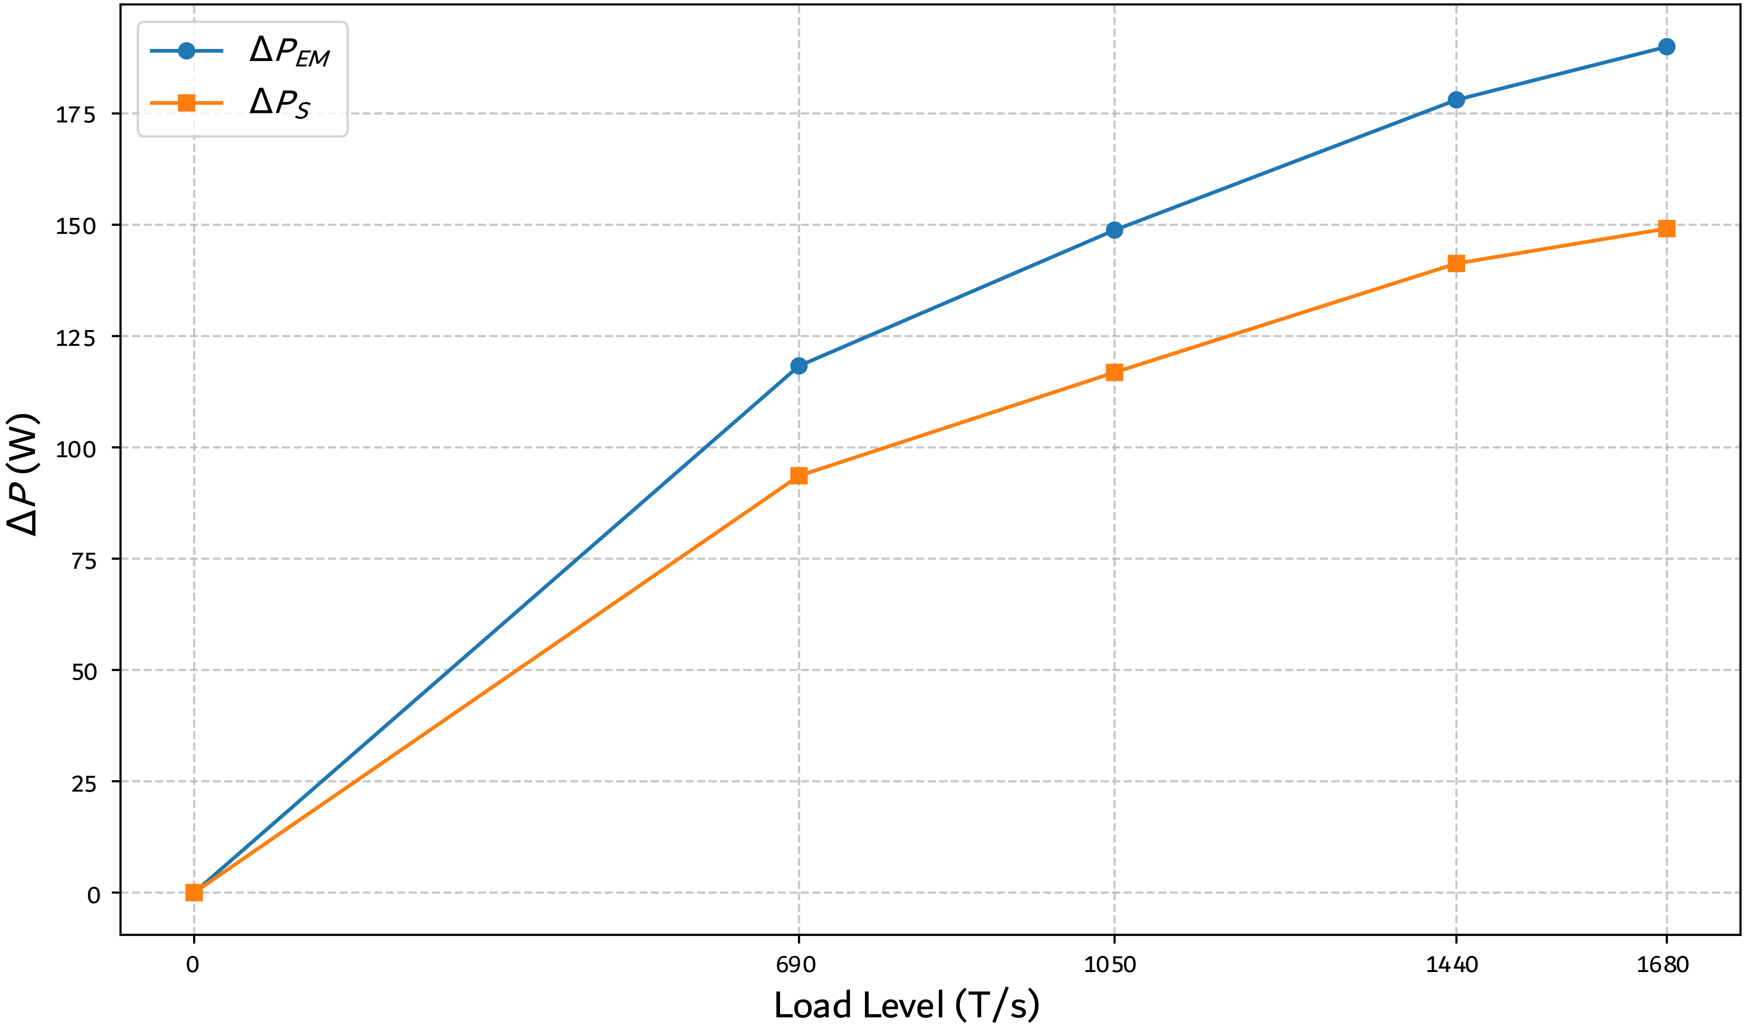

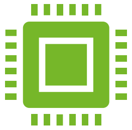

Experimental Results – Accuracy Overview

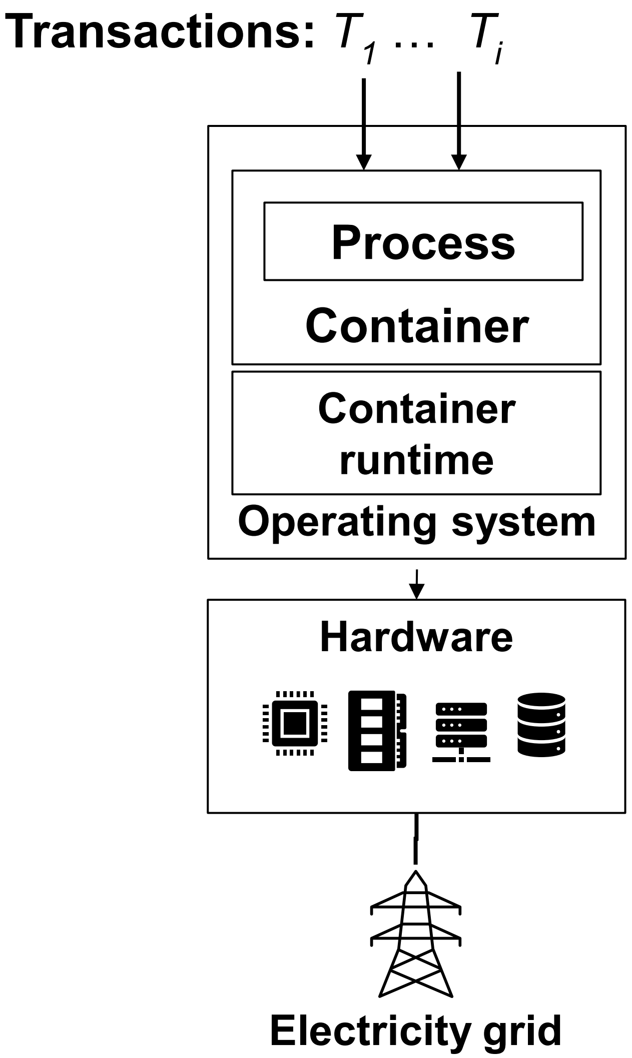

Transaction-level attributions: ~45 - 64%

Process-level attributions: ~16 - 65%

Container-level attributions: ~16 - 65%

Hardware measurements: ~53 - 66%

External measurements: 100%

Experimental Results – Accuracy Overview per Tool

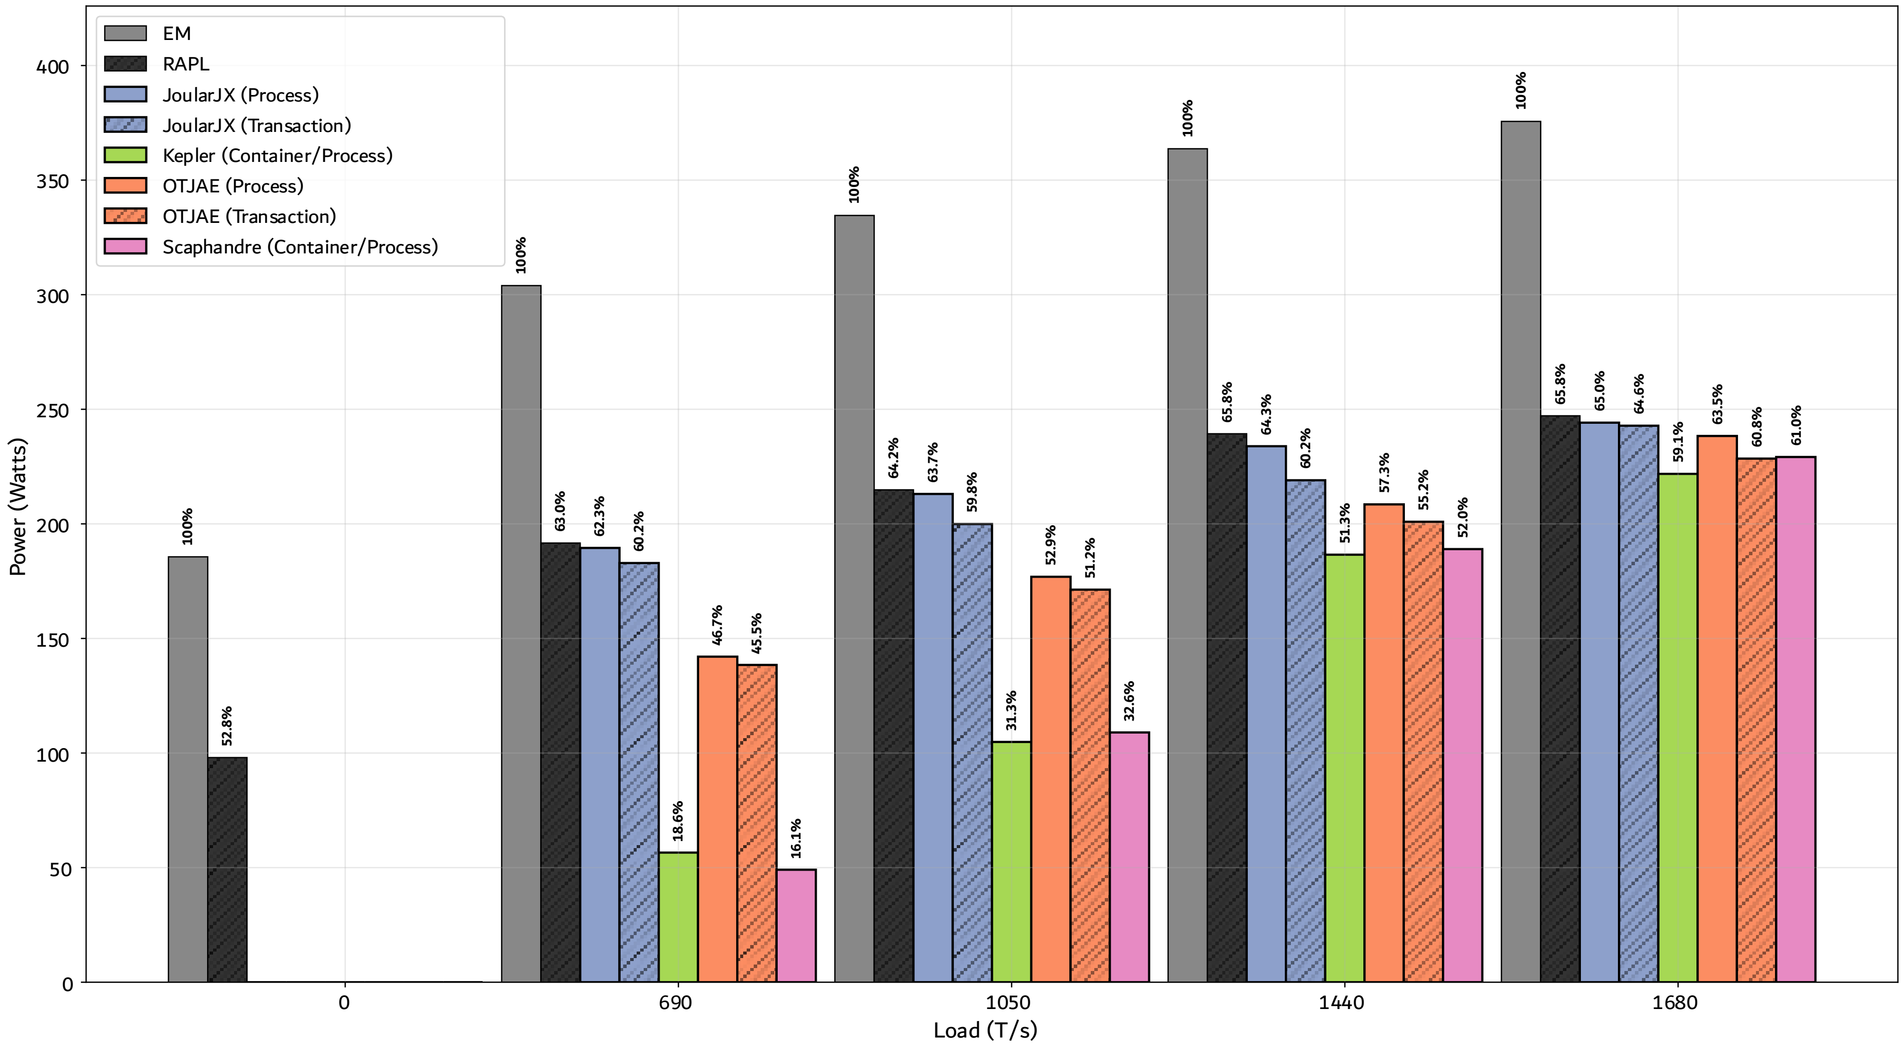

Summary and Conclusions

- **Accuracy of tools:**
  - Container/process level: 16–65%
  - Transaction level: 45–64%
  - Kepler and Scaphandre show a significant underestimation at low CPU utilization
    - → as this is very common scenario in data centers you must account for this using idle power distribution
- **Limitations of RAPL-based tools:**
  - Capture only CPU & DRAM power
  - Miss network, storage → measure only 53–66% of total system power.
- **Model-based approaches:**
  - OTJAE shows accuracy comparable to RAPL-based tools
  - Promising for cloud environments without hardware access.
    - → this enables the possibility to more easily represent additional resources
- **Open Challenges:**
  - We lack good models for non-compute-bound workloads. These include workloads relating to memory, storage and network energy consumption.
  - The way in which resources are split is often not user-friendly (e.g. how is idle power draw split among components?).
  - Further analysis is required to improve the accuracy of attributing concurrent workloads to the same system.

References

- **Tools:**
  - [1] https://landscape.bundesverband-green-software.de/
  - [2] Powercap: https://www.kernel.org/doc/html/next/power/powercap/powercap.html
  - [3] Kepler: https://github.com/sustainable-computing-io/kepler
  - [4] Scaphandre: https://github.com/hubblo-org/scaphandre
  - [5] JoularJX: https://github.com/joular/joularjx
  - [6] OTJAE: https://github.com/RETIT/opentelemetry-javaagent-extension
- **Replication Package and Data:**
  - https://github.com/hm-green-it-lab/jss2025
- **Further reading:**
  - Software Energy Consumption: https://luiscruz.github.io/2023/05/13/energy-units.html
  - Kepler paper (note: architecture change since v0.10.0): https://ieeexplore.ieee.org/document/10254956
  - Kepler since v0.10.0: https://github.com/sustainable-computing-io/kepler/blob/v0.11.2/docs/developer/power-attribution-guide.md
  - JoularJX paper: https://ieeexplore.ieee.org/document/9826760
  - OTJAE paper: https://dl.gi.de/items/3cbc03f9-64b5-41a8-be00-d45cea2412cb
  - OTJAE / JoularJX evaluation paper: https://dl.acm.org/doi/10.1145/3696630.3728709

Hochschule
München
University of
Applied Sciences

Department of Computer
Science and Mathematics

Thank you very much for your attention!
Any questions?

brunnert@hm.edu