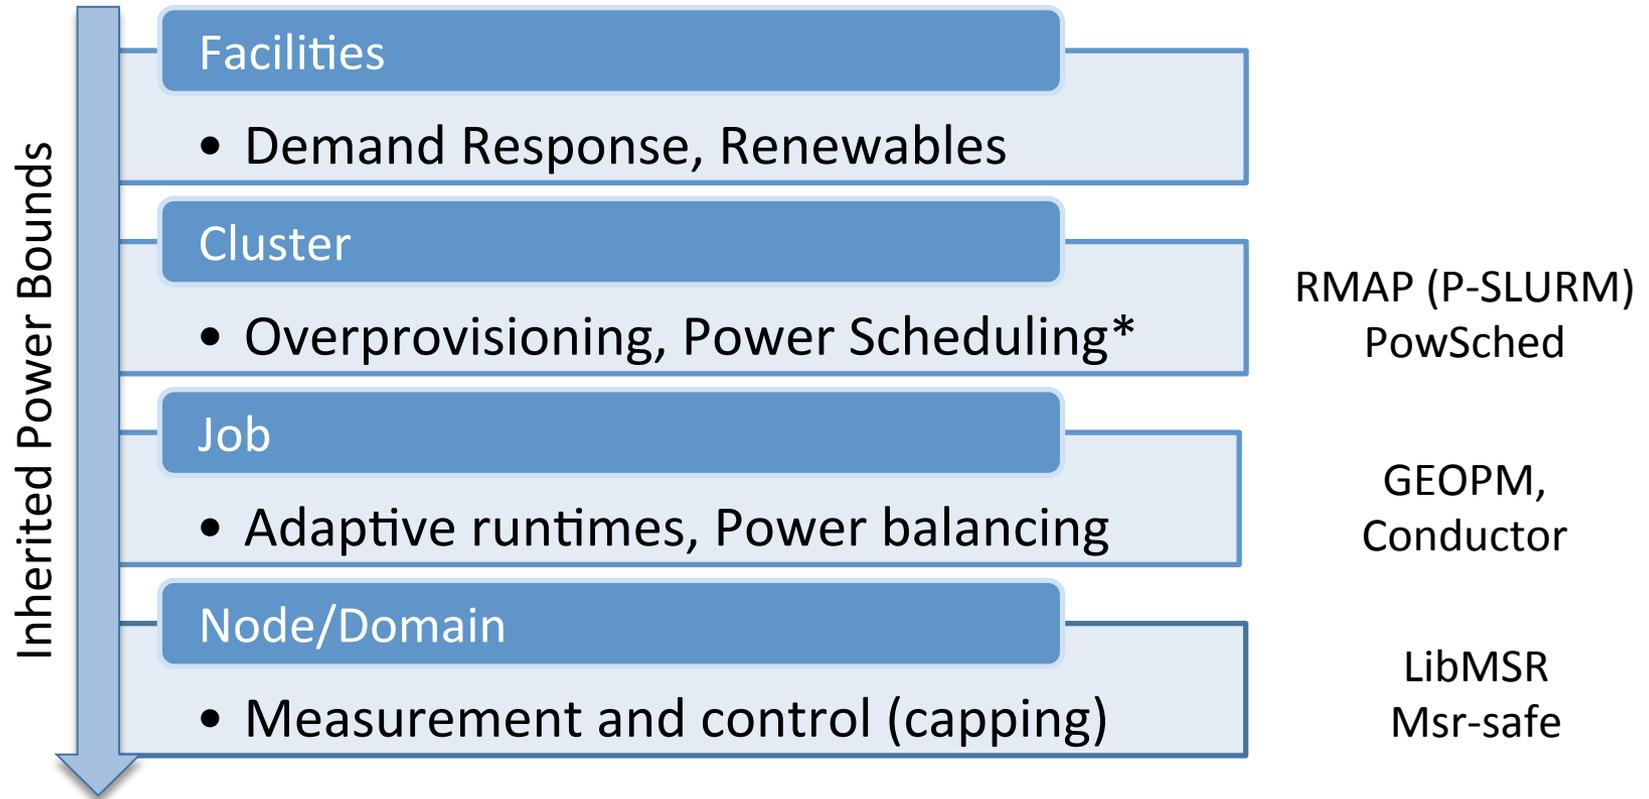


Interfacing GEOPM with PowerStack

Aniruddha Marathe
Tapasya Patki
Barry Rountree
Martin Schulz



Power Stack: Enabling Efficient Power Management Through Hierarchical Design

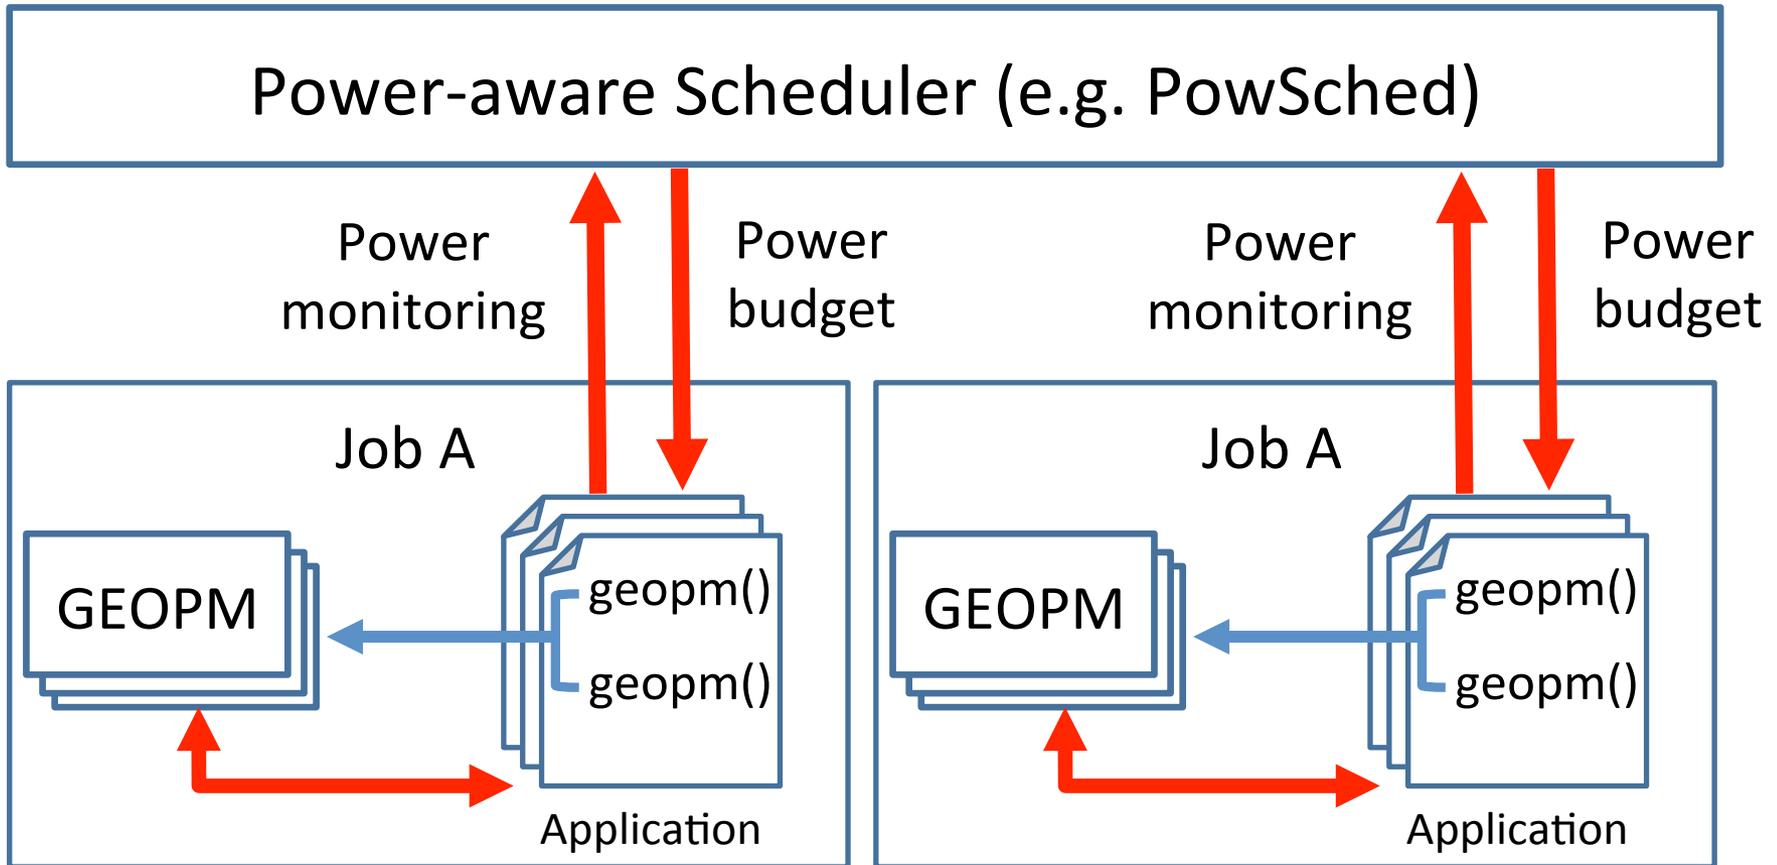


*For the scope of this talk, power scheduling does not represent job schedulers such as SLURM or FLUX, but an independent entity that manages cluster power (e.g., PowSched).

Power Stack Open Questions

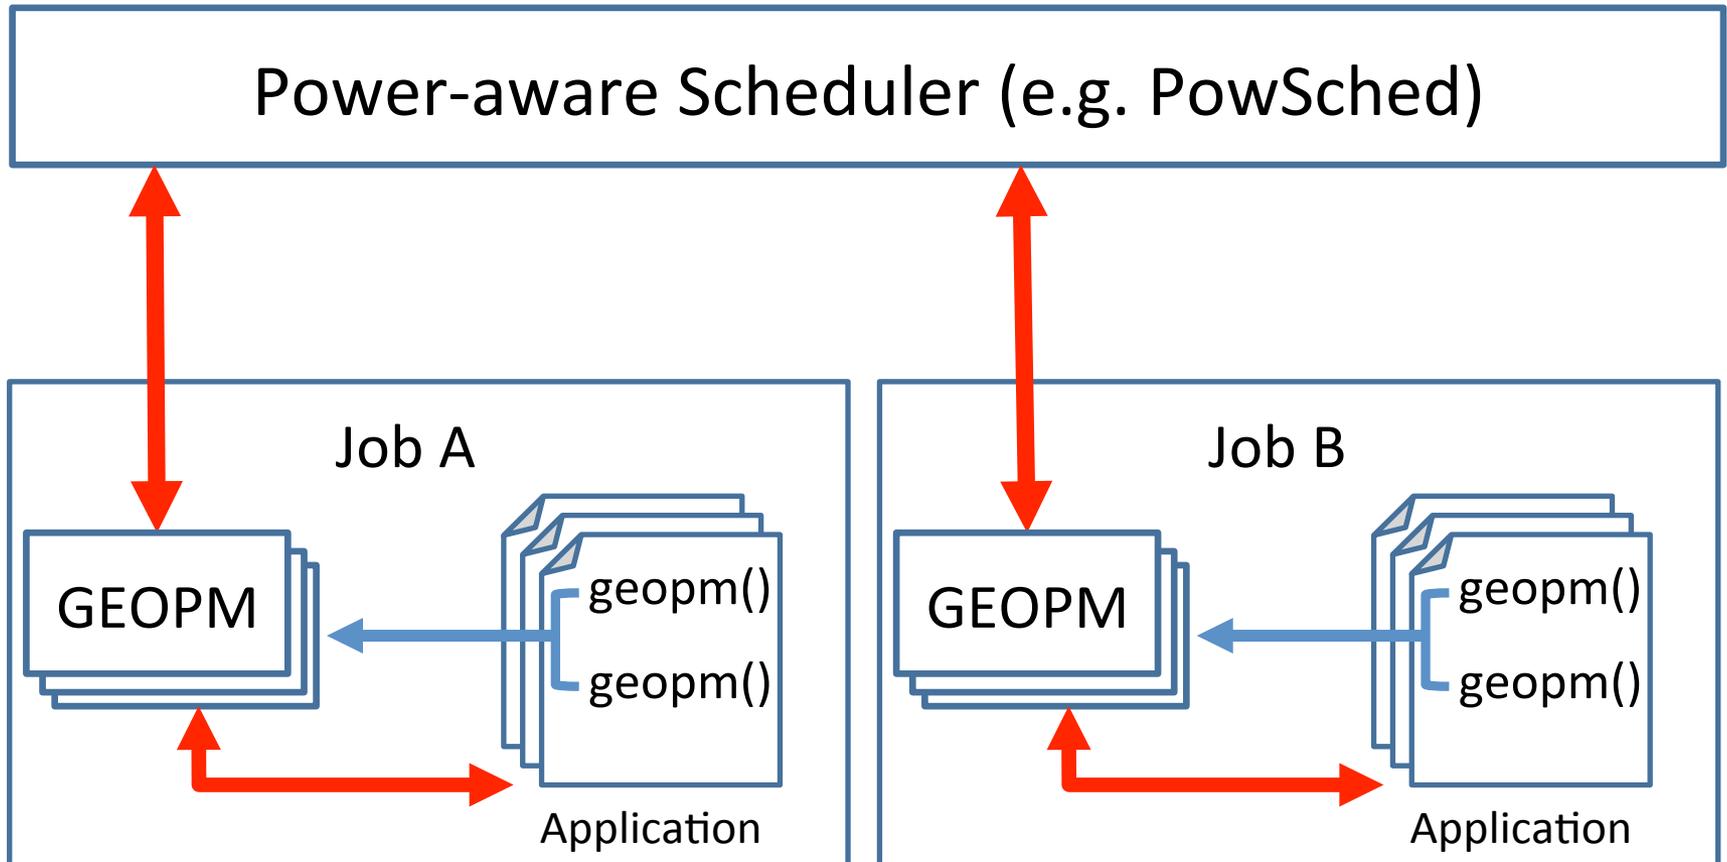
1. **Who should apply** the power limit at the node level: power scheduler or Individual GEOPM instances?

Power Monitoring and Limiting



Simultaneous power budget enforcement degrades efficiency

Interfacing GEOPM: Proposed Solution



Interface to exchange job power usage and power budget between (modified) power-aware scheduler and GEOPM

Power Stack Open Questions

1. **Who should apply** the power limit at the node level: the power scheduler or Individual GEOPM instances?
 - Exploratory evaluation shows only GEOPM should apply the power limit
 - Development of the proposed interface between GEOPM and the power scheduler

2. **How much power** budget should the power scheduler assign to GEOPM?
 - Or how much power should GEOPM request?

3. **How often** should the resource scheduler and GEOPM
 - a. Sample power usage?
 - b. Update power schedule?

On-going and Future Exploratory Work

- Simulation of power-aware job scheduling with GEOPM for hardware overprovisioned systems
 - Explore potential areas to improve the efficiency of Power Stack
- We are working with Intel on the development and deployment of GEOPM on LLNL systems.



Relevant Efforts in ECP Projects

- STPM14: ECP-PowerSteering
 - Migrate configuration exploration phase from Conductor into GEOPM
 - Main contributor: Aniruddha Marathe
- STSS08: ECP-Argo-GRM
 - Static power-aware plugin for Flux, later to be used for Argo to interface with GEOPM
 - Main contributor: Tapasya Patki

Collaborations

- **LLNL**

- David Boehme (Caliper)
- Kathleen Shoga (LC)
- Stephanie Labasan

- **Research Labs/Industry**

- Intel, USA
- LRZ/Intel Munich, Germany

- **Academic**

- University of Arizona: David Lowenthal, Samuel Cotter
- University of Georgia: Shelby Funk, William Whiteside
- University of Oregon: Srinivasan Ramesh, Dan Ellsworth (past)
- TU Munich: Martin Schulz

- **External Projects**

- **Argo**

- Swann Perarnau
- Kamil Iskra
- Kazutomo Yoshii
- Rinku Gupta
- Pete Beckman

- **LLNL Flux Team**