

Performance-aware cluster scheduling

Diana Andreea Popescu, Andrew W. Moore (University of Cambridge)

- Scheduling based on resource requirements vs. based on desired application performance
- Flip the viewpoint -> knowing the desired performance, what resources does the application need?
- Work focused on networking requirements, in particular latency
- Preliminary results show between 13% and 42% improvement in average application performance

Contact: diana.popescu@cl.cam.ac.uk