

# End-to-end performance optimization

## Measurement and Analysis

- I/O, memory, GPU
- GPU: CUDA, OpenCL
- Hybrid probe-based and sampling-based
- Instrumentation for source, binary, APIs, and runtime
- Parameter-based and context-based
- Provenance capture

## Modeling

- Performance bounds
  - Roofline
  - Static analysis
- Single-node performance
- Communication

## Autotuning

- Mixed programming models
- Nonvolatile memories
- Whole program
- Multi-objective Optimization
  - Performance
  - Energy
  - Resilience