

Multiobjective Optimization

- The problem of empirically optimizing a code can be posed as the mathematical optimization problem
- Increasingly, multiple metrics are of interest simultaneously
- When the relative weights or constraints on these objectives are not known at search time, autotuning becomes a *multiobjective optimization problem*:

$$\min_{x \in \mathcal{X}} F(x) = [F_1(x), \dots, F_p(x)]$$

THE OBJECTIVES $F_i(x)$

- $F_1(x), \dots, F_p(x)$ are p **possibly conflicting objectives** that need to be optimized simultaneously
- Can capture average, median, quantile (e.g., worst-case) empirical performance
- Often stochastic/noisy (from measurement and/or run)
- Depends on machine and input size (or distribution over inputs)
- Assumes no *a priori* weights available for the objectives
- Examples: run time, power, energy, failure rate

THE DECISIONS x


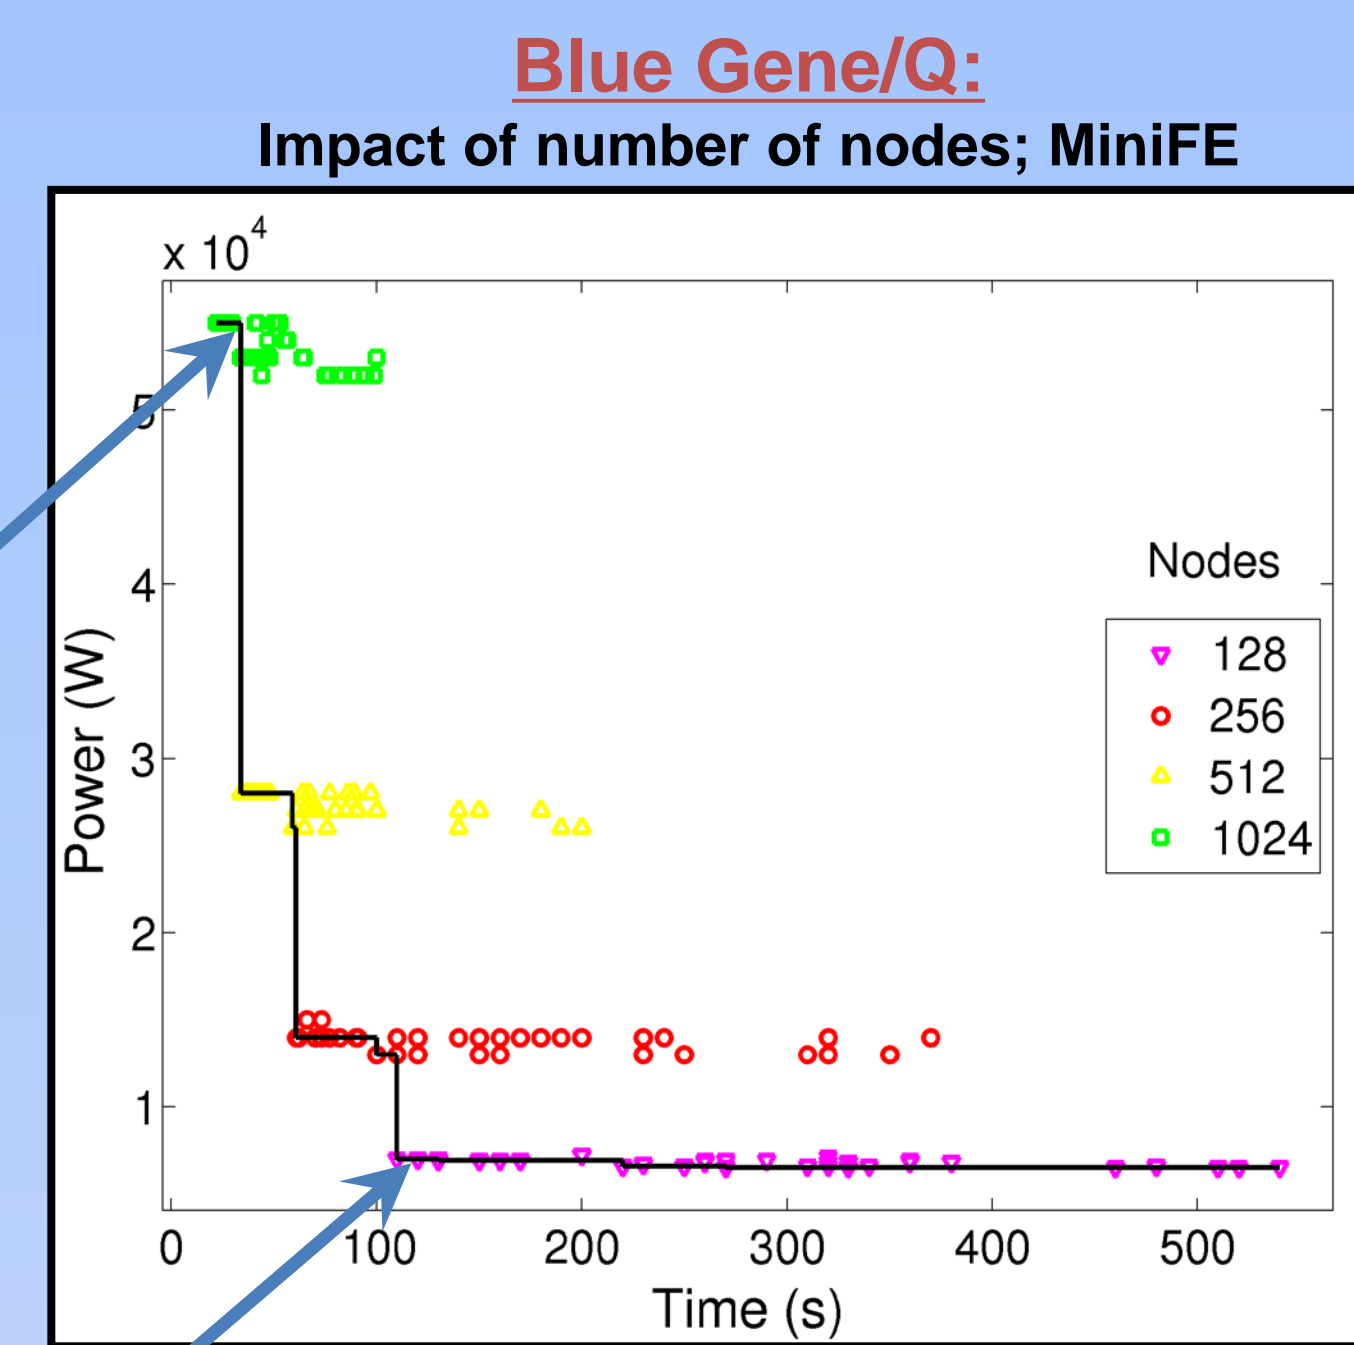
- Integer (unroll factor, register tiling, +examples)
- "Continuous" (algorithmic parameters, internal tolerances)
- Each x generates a code variant (e.g., through source-to-source or compiler-based transformation)

THE CONSTRAINTS \mathcal{X}

- Ensuring feasibility of transformation
- Correctness of output, maximum temperature, etc.

Illustrative example on three objectives on IBM BG/Q

- Minimizing run time is conflicting with power consumption
- Minimizing energy consumption conflicts with run time



Multi objective optimization concerns the study of optimizing two or more objectives simultaneously.

Even if there is a unique optimal (software/hardware) decision when any of the objectives is considered in isolation, there may be an entire set of solutions when the objectives are considered collectively.

GOAL:

Develop multi objective optimization framework that allows exploration of the tradeoffs


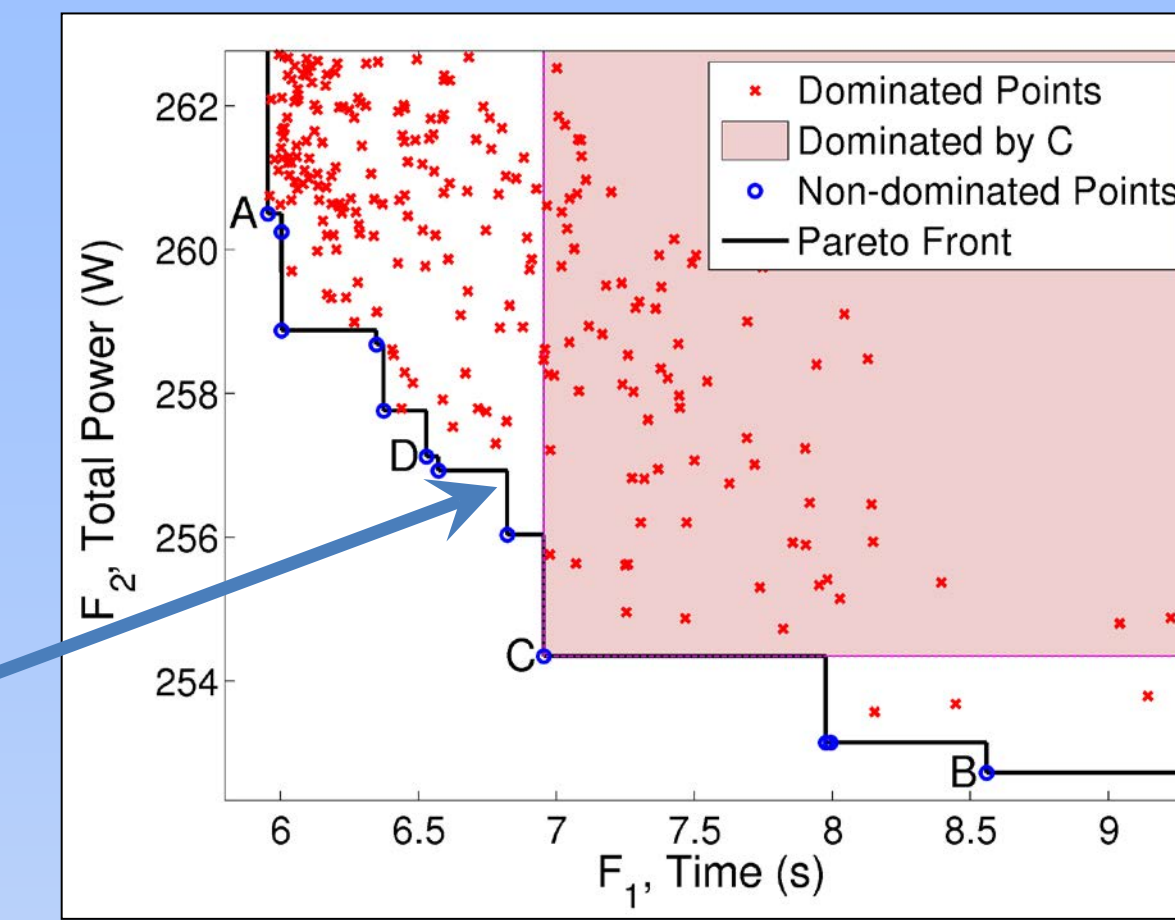
Existence of these tradeoffs can motivate hardware designers to expose a richer and more appropriate set of knobs to future administrators and software designers

A framework that is sufficiently general and can be easily extended to incorporate new hardware-and software-based power/energy knobs as they become available


Pareto Optimality

Arises when several objectives need to be optimized simultaneously

- Code variants now live both in a decision space and in an objective space



- Pareto front contains significantly richer information than one obtains from single-objective formulations
- Code variants for which no other variant is better in all objectives are said to be **non-dominated or Pareto optimal**
- For search algorithms, only certain regions of the objective space are of interest
- The ideal and nadir point define the range of objective that include all possible optimal tradeoffs
- When multiple objectives are **not competing** the Pareto front corresponds to a single point, which simultaneously minimizes both objectives



- For many time-power-energy multi-objective problems, there can be measurement error in each objective
- Consequently, we have a **relaxed Pareto front** that potentially consists of a cloud of points

Conditions for Energy Time Tradeoff


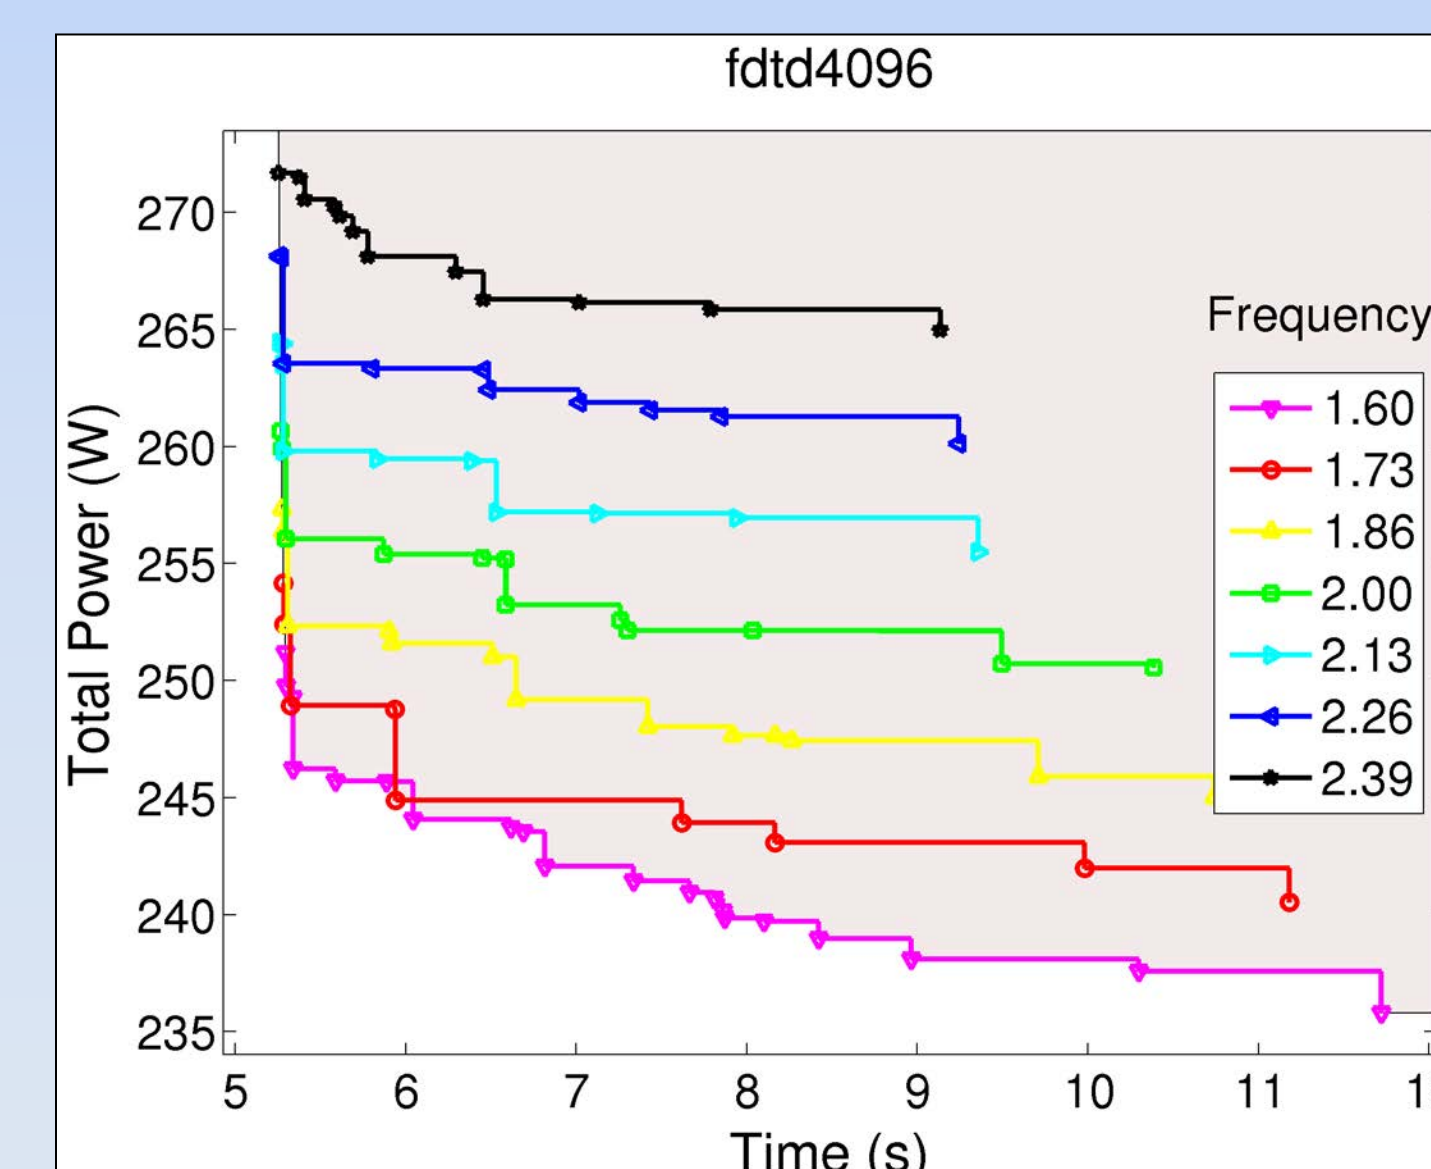
- Because of the relationship between power and energy, all points on the energy-time Pareto front have a corresponding point on the power-time Pareto front
- Number of non-dominated points for energy-time is bounded by the number of non-dominated points for power-time
- A necessary condition for x to be a non-dominated point on the energy-time Pareto front is

$$P(x) \leq \frac{P(x^{(1)})T(x^{(1)})}{T(x)}$$

- To observe tradeoff, the power savings must outpace the product of idle power and relative slow-down:

$$P(x^{(1)}) - P(x) \geq \frac{T(x) - T(x^{(1)})}{T(x^{(1)})} P_I$$

EX: Intel Xeon E5530: Impact of frequency scaling; SPAPT ftdtd kernel

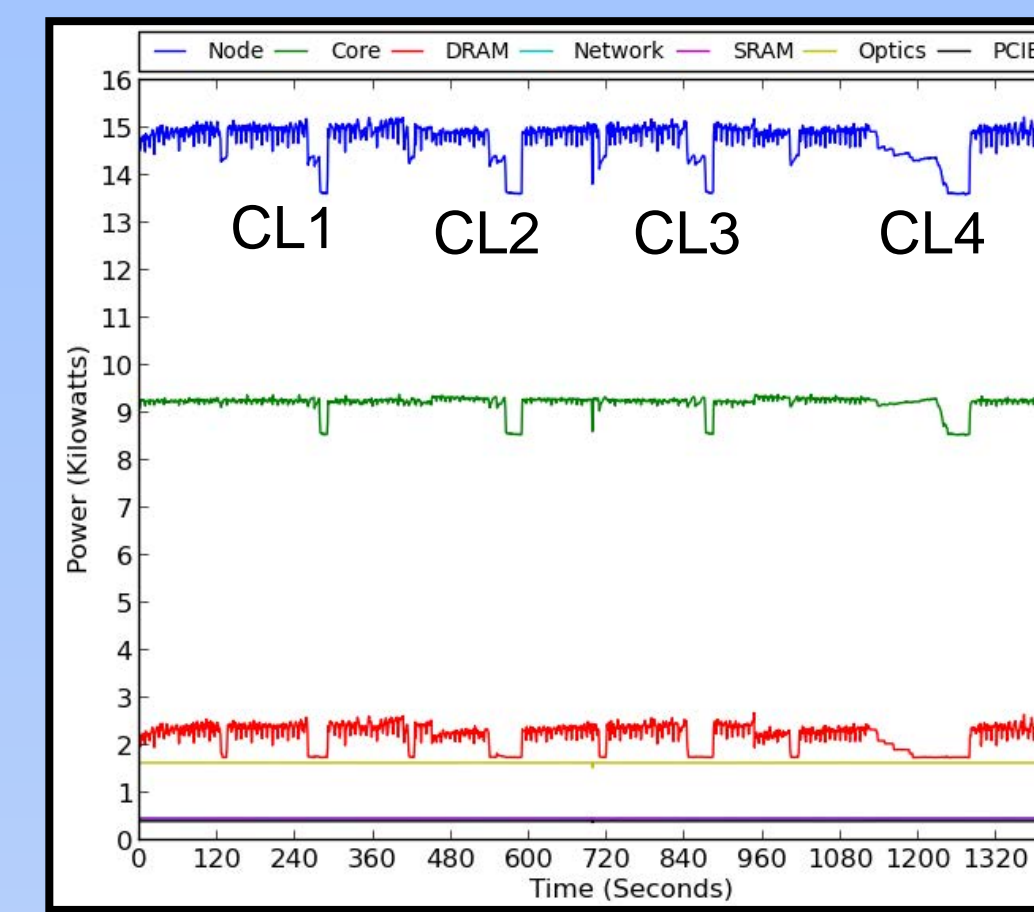


Tradeoff Studies


EX: Blue Gene/Q: Energy-time tradeoff with multilevel FTI checkpoint library; CORAL benchmarks

- Analytical models for expected run time and energy consumption for multilevel checkpointing
- Characterize the Pareto-optimal solution set and investigate the tradeoffs time and energy
- Power consumption measurements of large-scale executions on an IBM Blue Gene/Q with several applications
- Analyzed several system-level parameters for multilevel checkpointing that can potentially impact the tradeoffs

Component level power profiling



Effect of different level 4 power



Effect of idle power


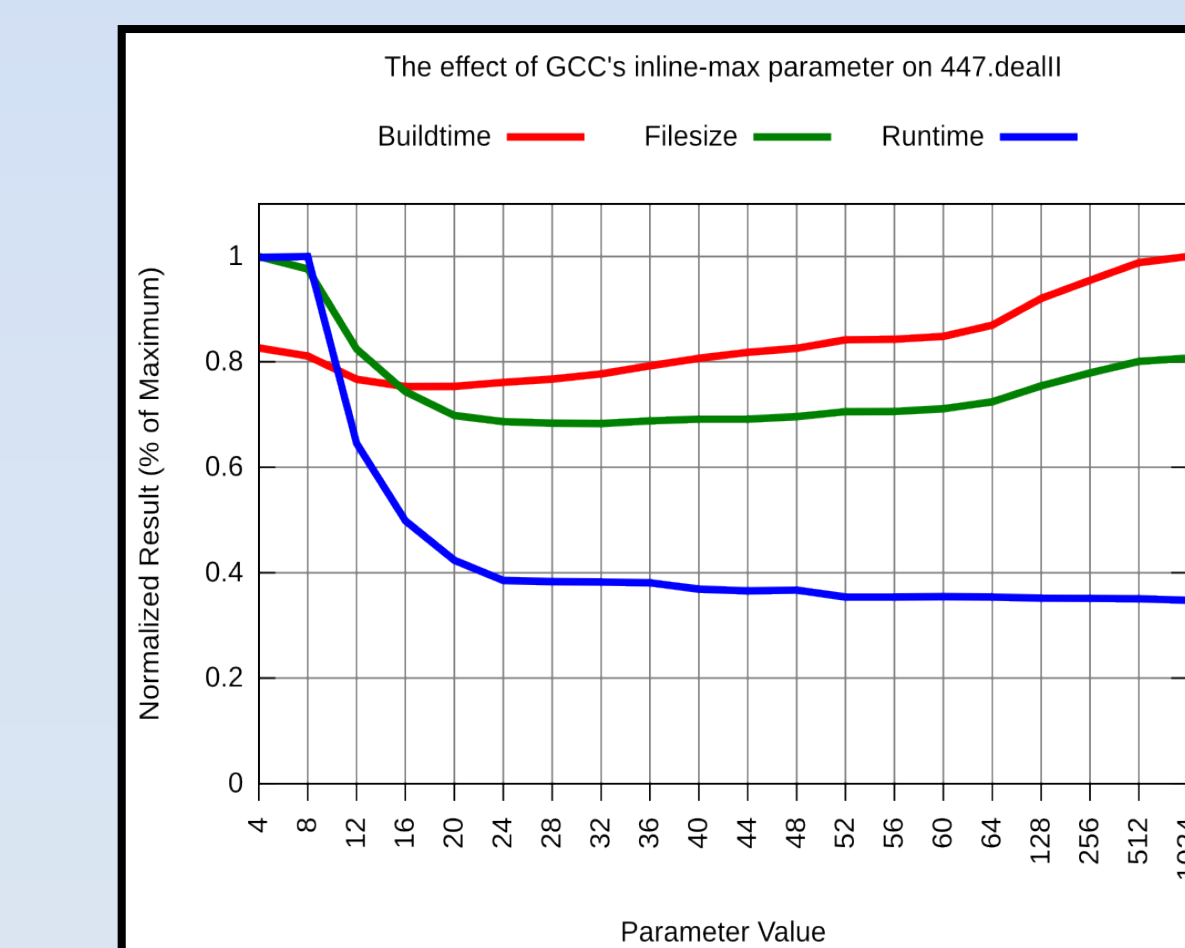
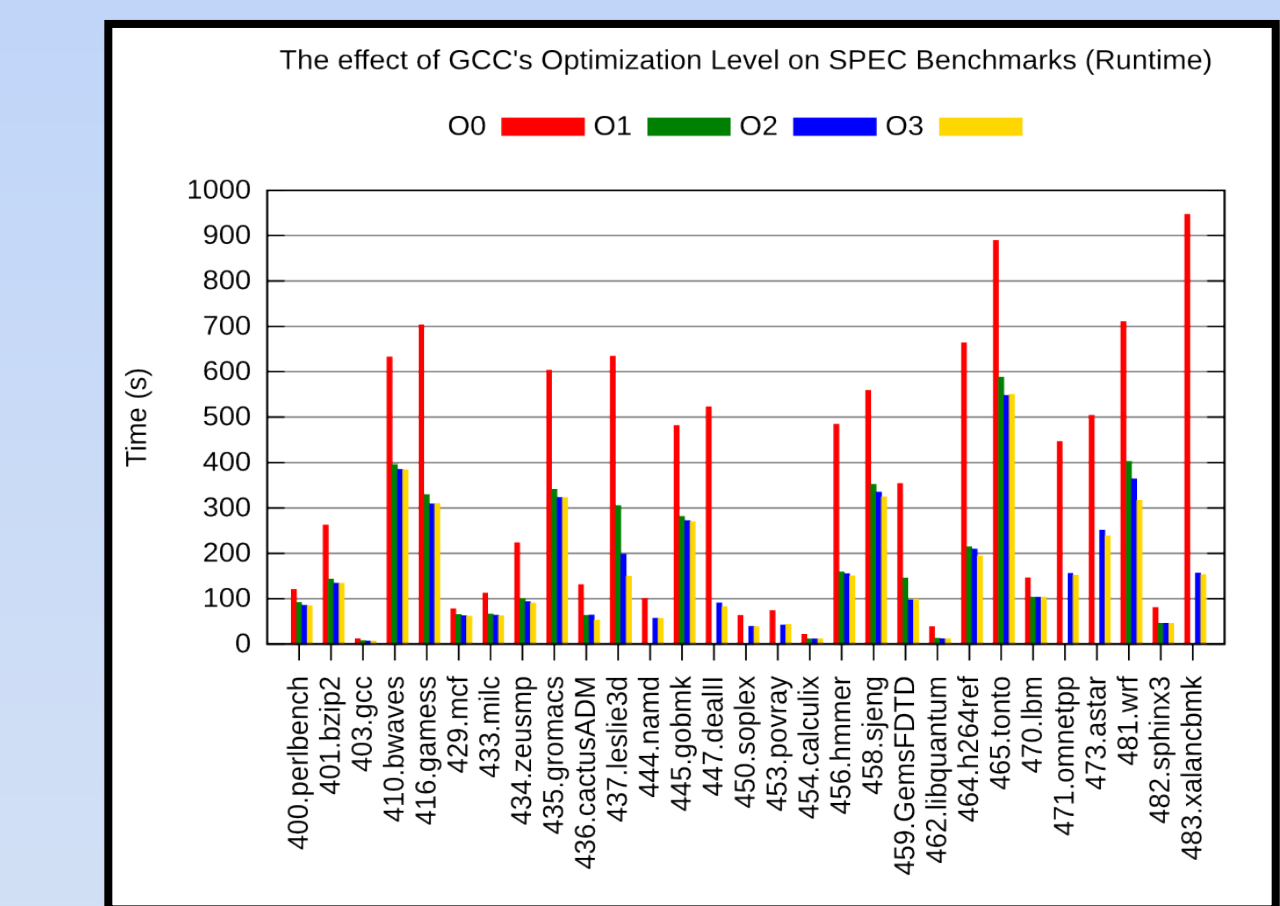
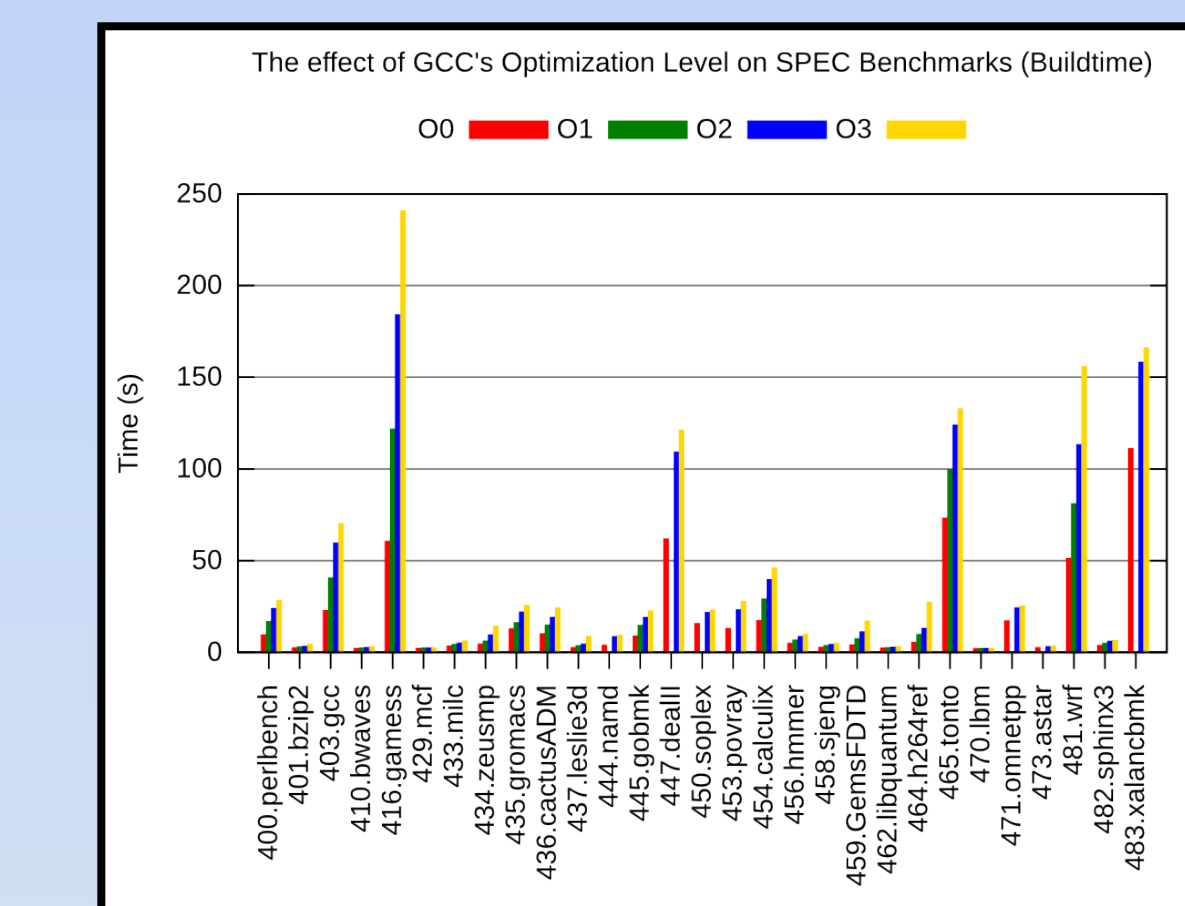


- CL 1: Local checkpoint
- CL 2: Local checkpoint + Partner-copy
- CL 3: Local checkpoint + RS coding
- CL 4: Parallel file system

Increasing level 4 power relative to other levels has a significant impact on the observed tradeoff

- Ratio of checkpoint and idle power
- Significant impact on the tradeoff

EX: Three simultaneous objectives: Build time, file size, and execution time



Our studies show that in some settings objectives of interest can be strictly correlated and there is a single, "ideal" decision point; in others, significant tradeoffs exist.

Future Investigations

- Develop multi objective optimization algorithms for autotuning search
- Identify appropriate use cases
- Study other tradeoffs:
 - Resilience versus memory footprint;
 - Memory footprint versus execution time; Memory footprint versus energy

"Can search algorithms save large-scale automatic performance tuning?" Balaprakash, Wild, & Hovland, *IWAPT 2011*.
 "Online adaptive code generation and tuning." Tiwari & Hollingsworth, *IPDPS 2011*.
 "Energy-Performance Tradeoffs in Multilevel Checkpoint Strategies." Balaprakash, Wild, & Hovland et al, Preprint ANL/MCS-P5138-0514, 2014.
 "Multi-objective optimization of HPC kernels for performance, power, and energy" Balaprakash, Tiwari, & Wild, *PMBS*, 2013

