Energy efficient low level dvfs for HPC applications

Georges Da Costa, Jean-Marc Pierson
IRIT, Toulouse University

Green Days @ Paris

Plan

Context

To optimize a computing center:

- Gather insight on running applications
- Choose how to act
  - depends on application
  - More precise : phase of application
- Act (change frequency, switch on/off parts of nodes,...)

Optimize : reduce energy consumption at the same performance

Ignorance is bliss, really?

- *-AAS (PAAS, IAAS,...) leads to ignorance
- Ignorance leads to errors
- Errors lead to inefficiency

Focus :

How to optimize a computing center while knowing nothing?

Know your enemy

- What we know
  - HPC applications
  - Goal: Save energy
  - No impact on performance (SLA,...)
- Name your weapon (constraints)
  - Minimum impact of monitoring
  - Closed application, no source
  - Even full OS freedom (Grid'5000, VMs)

Is it so important?

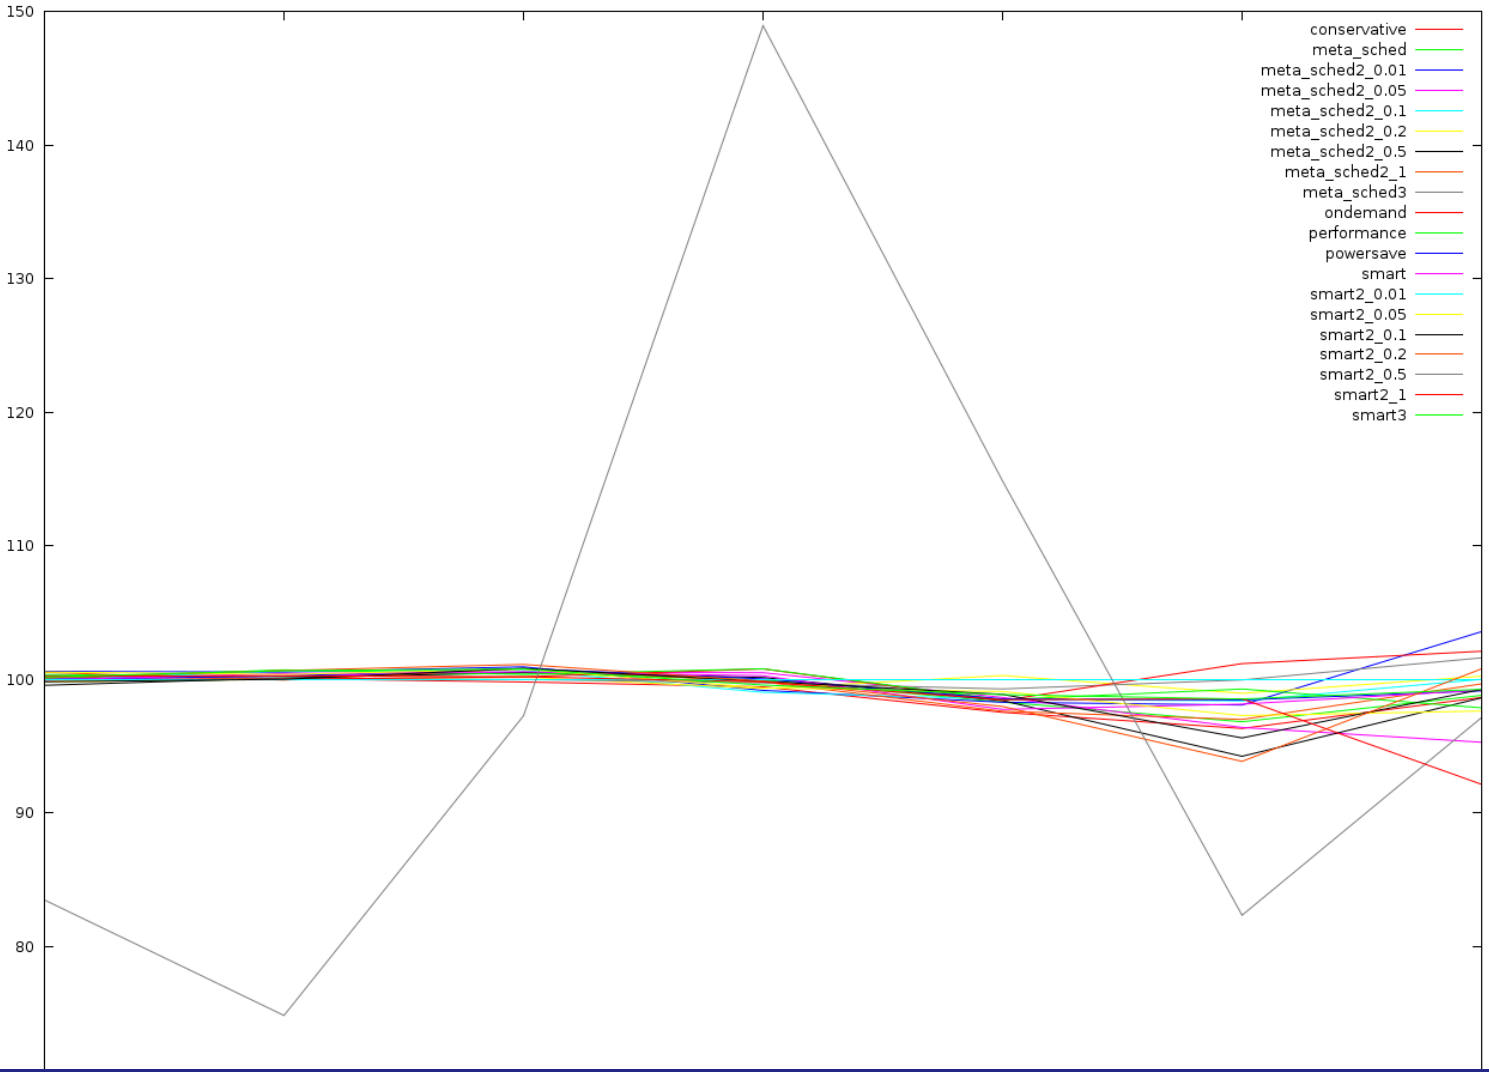

Plan

Who's Who

- Which application is running?
  - Ask the developer, but
    - Depends on library
    - Can cheat (if accounting is related to it)
    - Computers work for us, not the opposite
  - Application is not important, its behavior is!
    - Two different apps can have the same impact
    - System changes can have the same impact

We need Run-Time Behavioral Detection

BigBrother is watching

Run-time detection

- Behavioral pattern
- Extract information
  - Fine grained : performance counters, system values
  - Coarse grained : network, disk, energy consumption
- Classical remark: impact of the monitoring infrastructure

Finding patterns, NPB example, CG

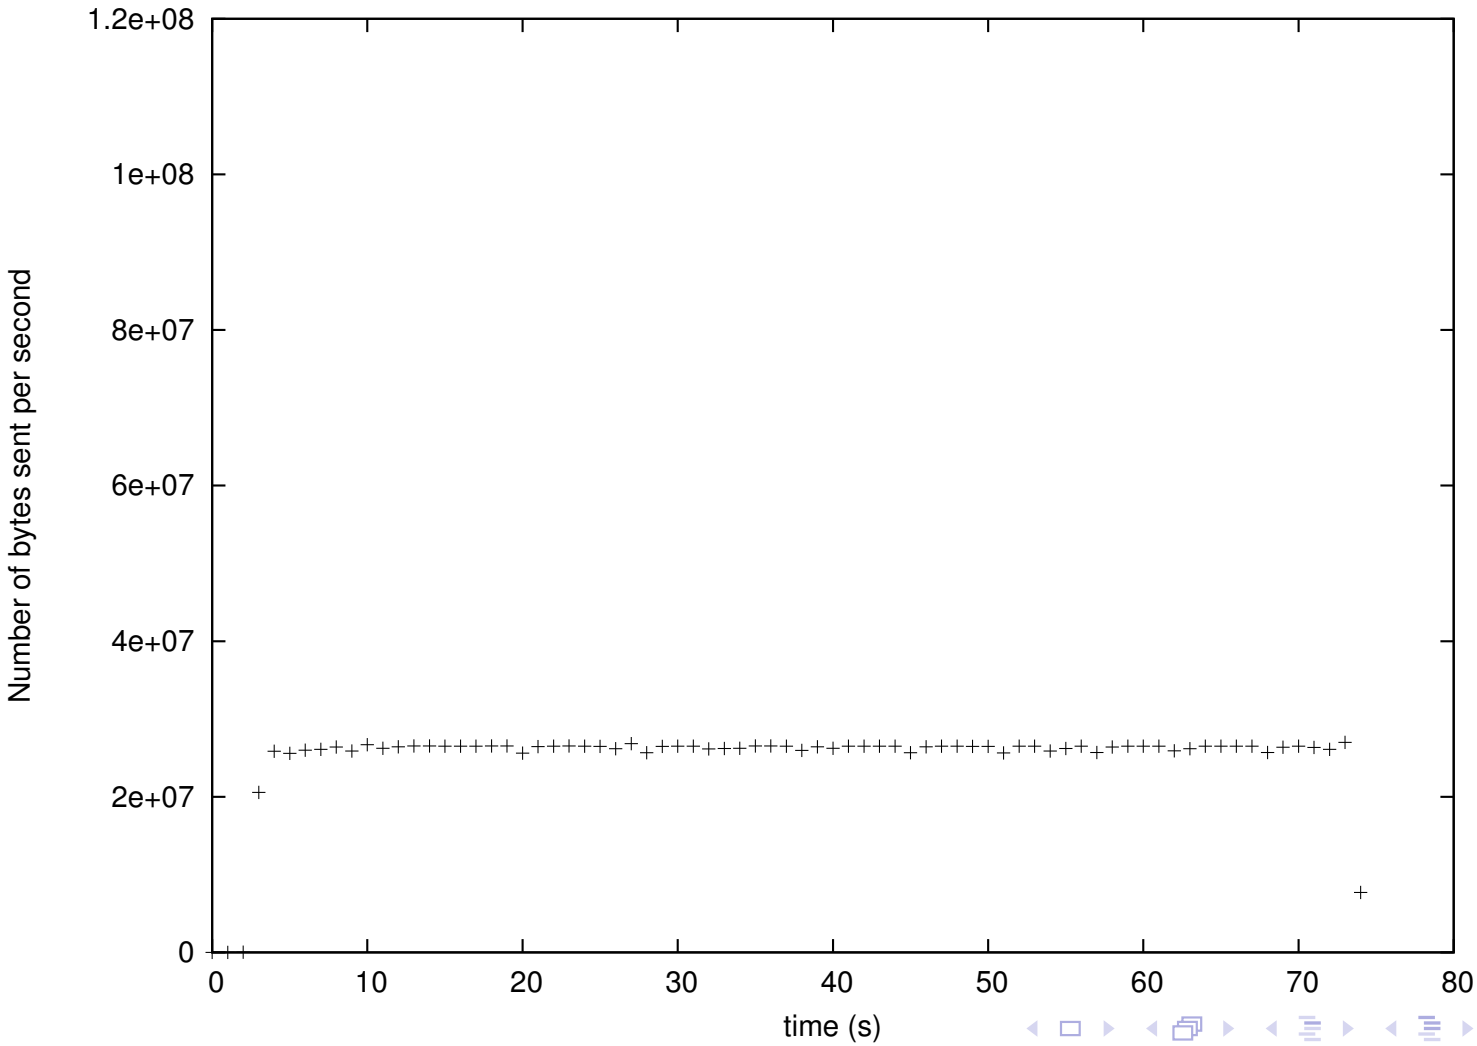

Finding patterns, NPB example, FT

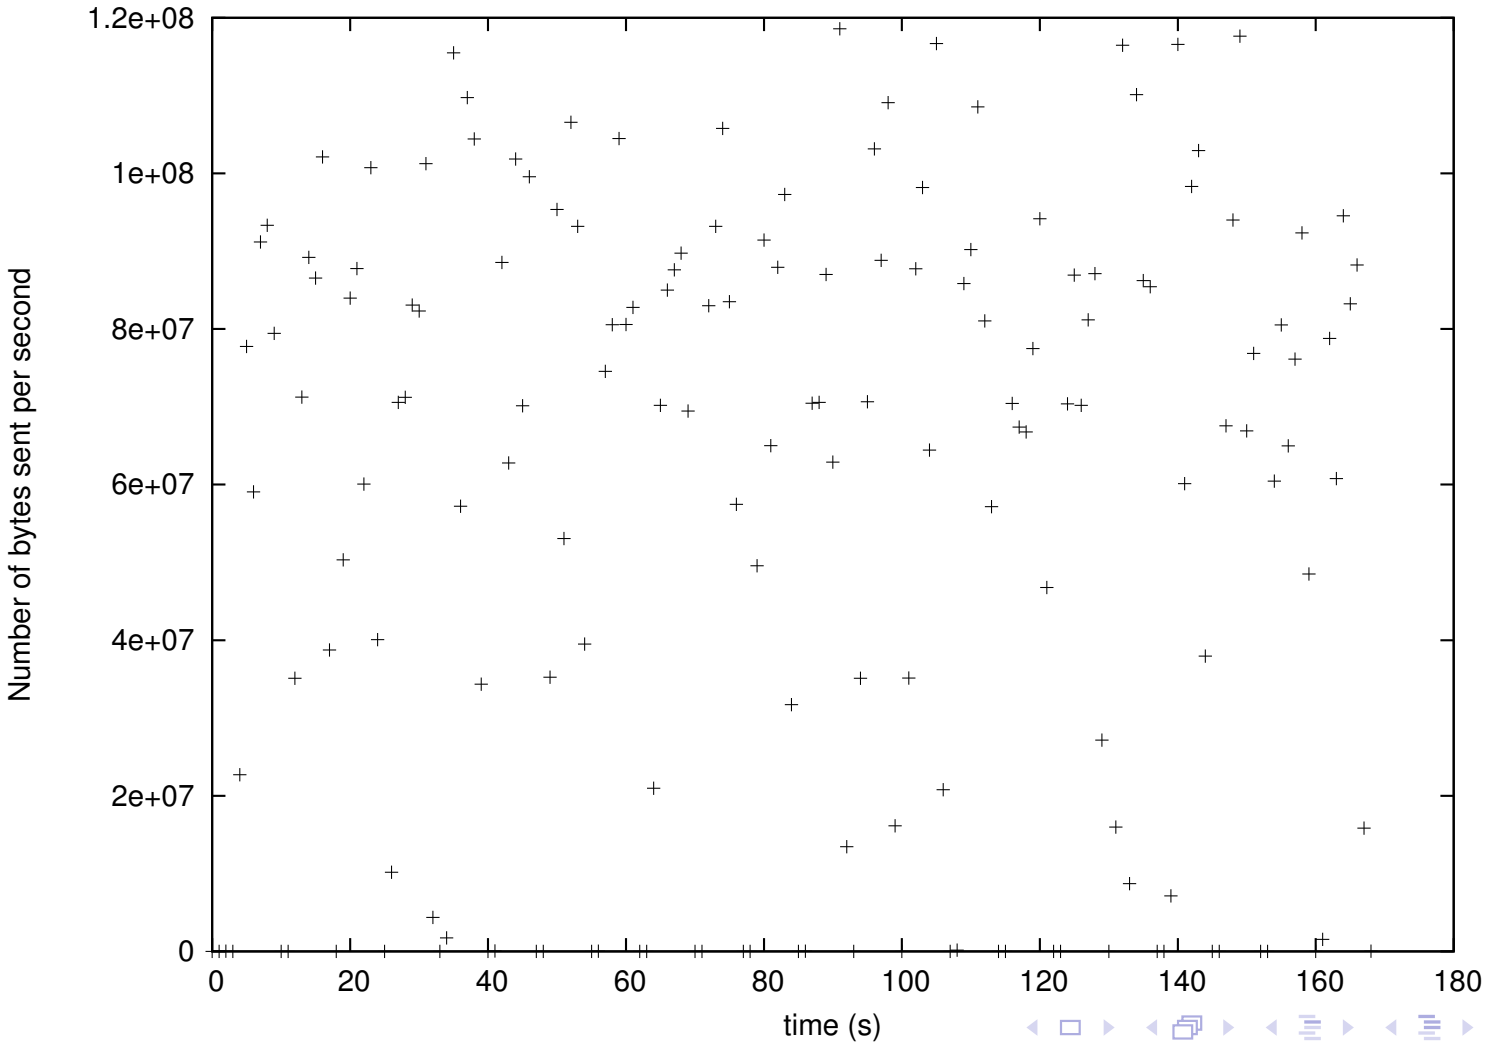

Finding patterns, NPB example, SP

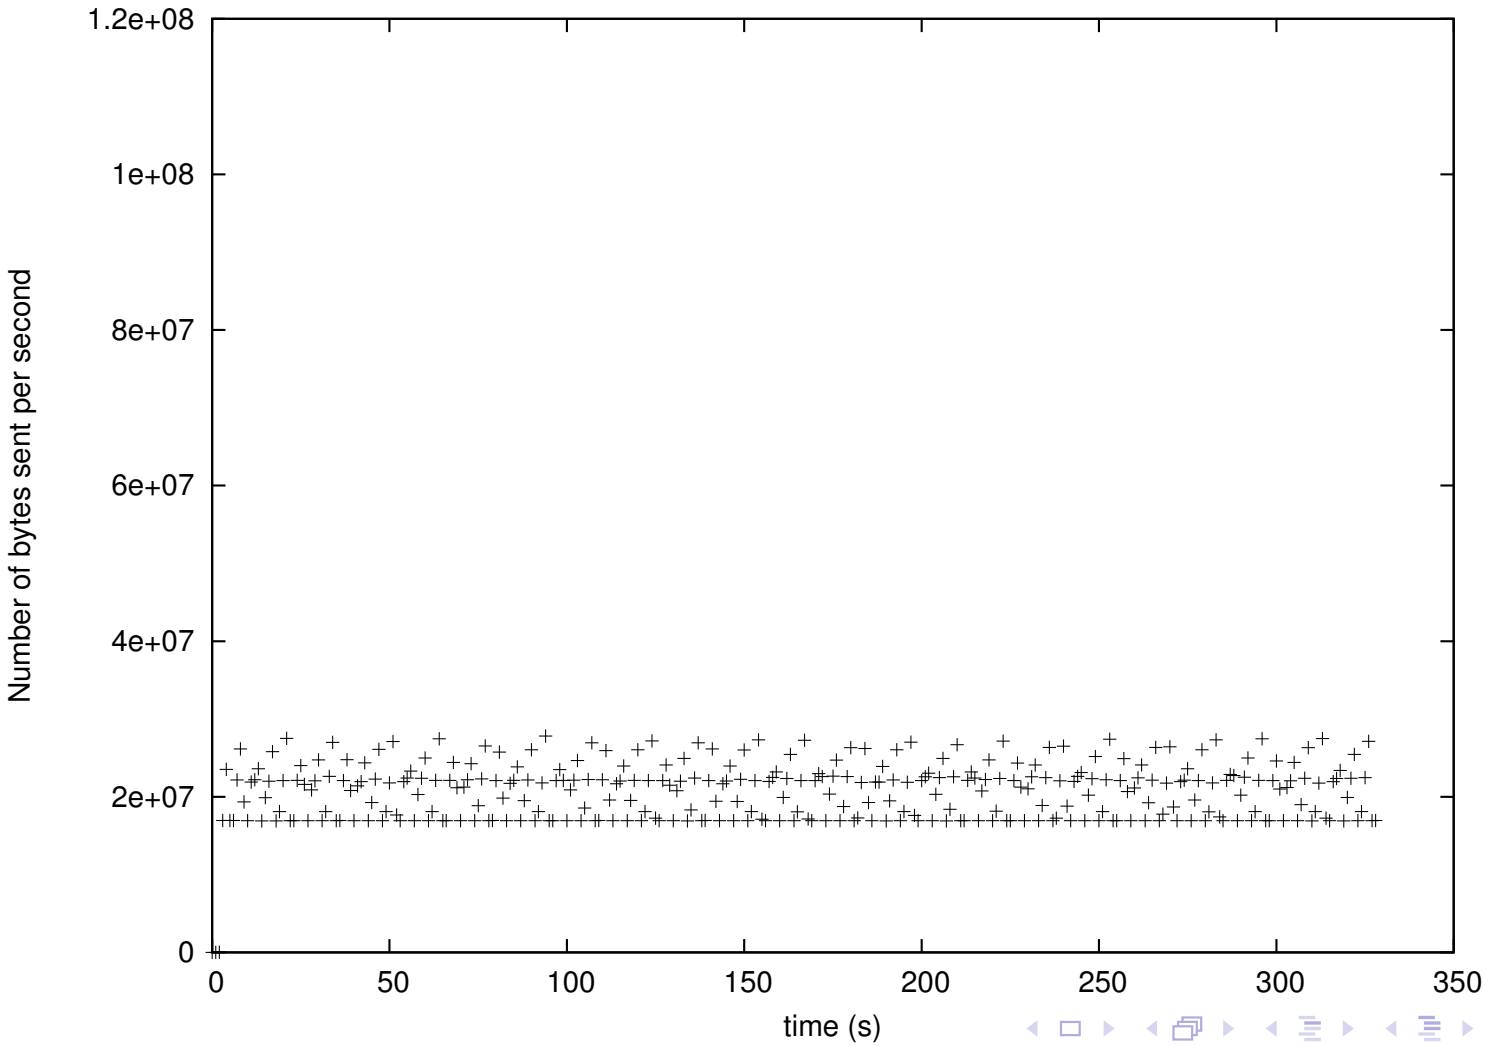

Approach

Once we have a model

- Step 1: Measure some characteristics
- Step 2: Compare to reference
- Step 3: Categorize application (or phase)

Usecase : Nas Parallel Benchmark

- Different type of workload
- Representative of HPC applications
- Seven benchmarks

Network and energy (left measured, right computed by R)

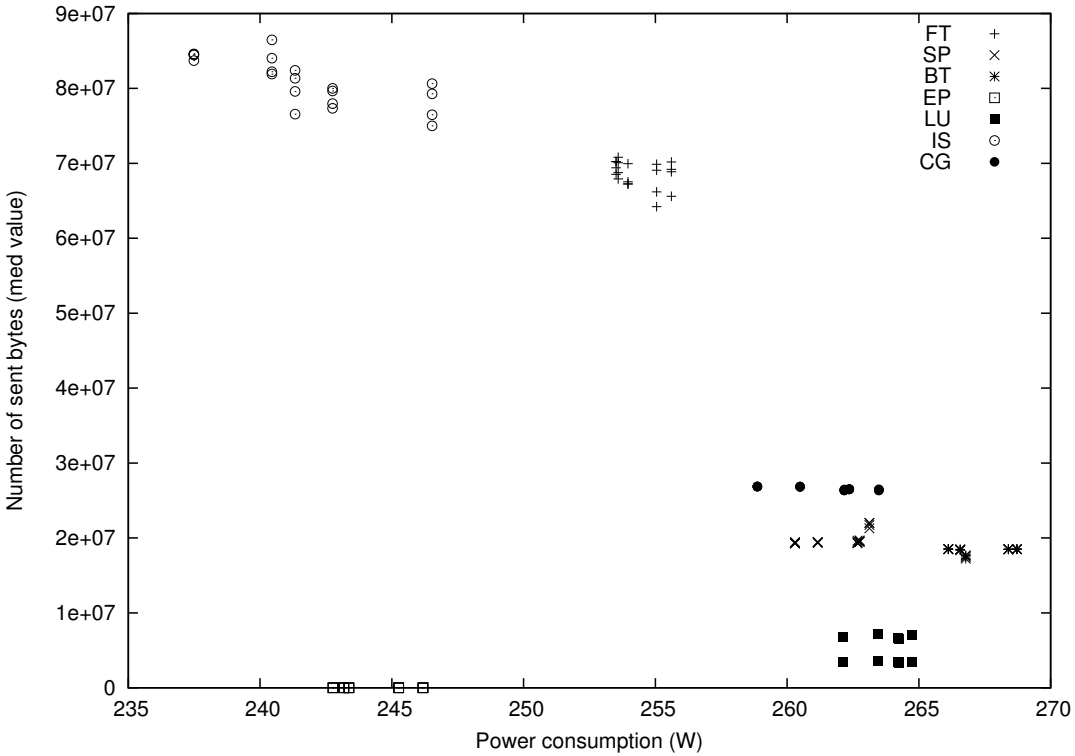

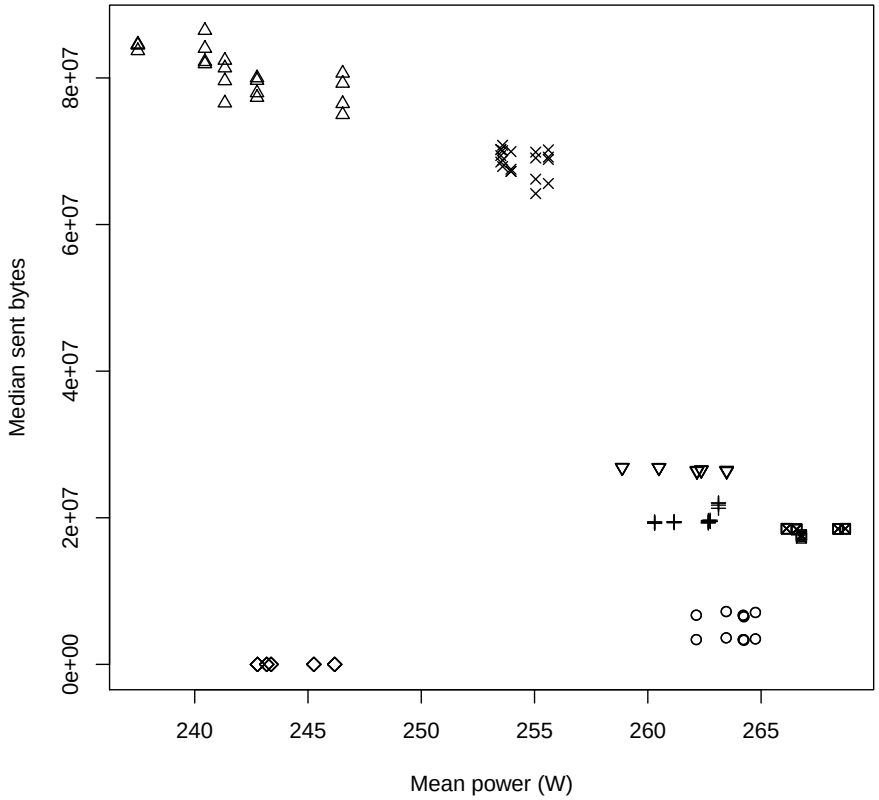

Energy efficient dvfs

- Categorize application (phase)
- Apply a rule-based algorithm
- Change processor speed to min or max

Can take into account several objectives depending on rules

- Energy only
- Energy with taking into account performance
- ...

Plan

With more control comes more efficiency

Governors

- Needs control over OS
- Take low level information
- Take decision on frequency

Ongoing fellowship

Preliminary results

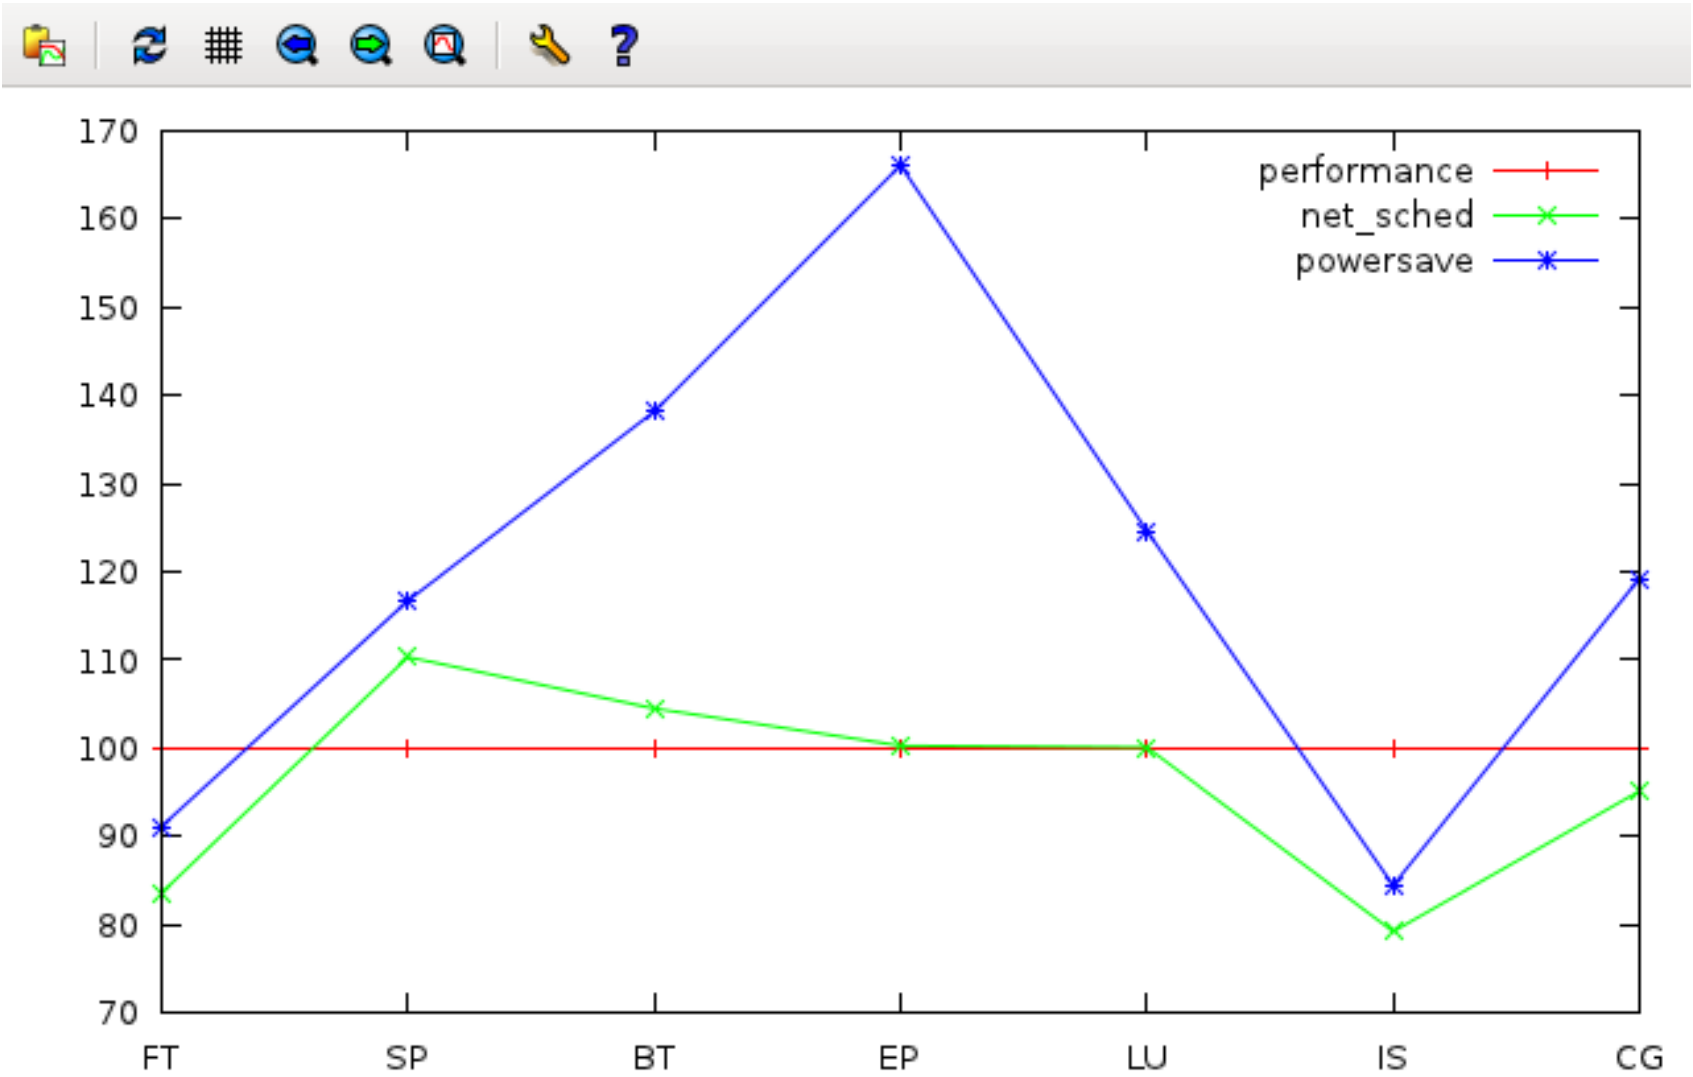

Conclusion

- Application characterization is possible with no impact!
- Using it we can optimize resources usages and reduce energy (J)
- Also possible to use at the kernel level.