

Agenda for today

- Discussion of chapter on performance analysis and measurements (04b-perf)

Feedback: some questions (1)

- What are “macro“- and “microBenchmarks?
- When is one or the other appropriate?
- Which method should be used to compute the mean of relative figures?
- How does a profiling tool like `gprof` work? Can it be used for OS profiling?
- How about a Performance Monitoring Unit (PMU)?
- How can a pseudo random generator be coerced to produce identical sequences of “random“ numbers?
- How can this be useful (i.e. it isn't really random...)?
- In a (micro-)benchmark, how to compensate loop overhead?