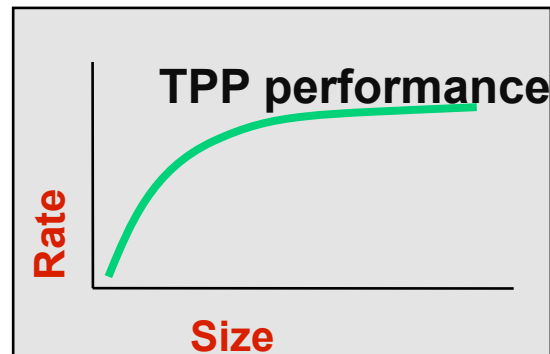


# LINPACK BENCHMARK AND SOME ISSUES

Jack Dongarra, UTK/ORNL

# HPL

- Linpack benchmark uses  $O(n^2)$  data and  $O(n^3)$  operations.
- If we look at the performance as a function of size we see something like this.



- So you want to run a large a problem as you can on your machine to get the most performance.

# Running HPL

3

- The LANL RoadRunner HPL run took about 2 hours.
- They ran a size of  $n=2.3 \times 10^6$
- At ORNL they have more memory, 300 TB, and they wanted to run a problem which used most of it. A matrix of size  $n = 4.7 \times 10^6$
- This run took about 18 hours!!

# In 4 Years ...

4

- Have a 5 Pflop/s system
- If memory goes up by a factor of 5 we will be able to do a problem of size  $n = 33.5 \times 10^6$
- Running at 5 Pflop/s the benchmark run will take 2.5 days to complete
- Clearly we have a issue here

# We Have to Do Something

5


- One of the positive aspects of Linpack is that it stresses the system.
- Run only a portion of the run.
- But for how long?
  - ▣ 4 hours? 6 hours? 12 hours?
- Have to check the results for numerical accuracy.

$$\frac{\|Ax - b\|}{(\|A\|\|x\| + \|b\|)n \varepsilon} = O(1)$$

# Over the Course of the Run

6

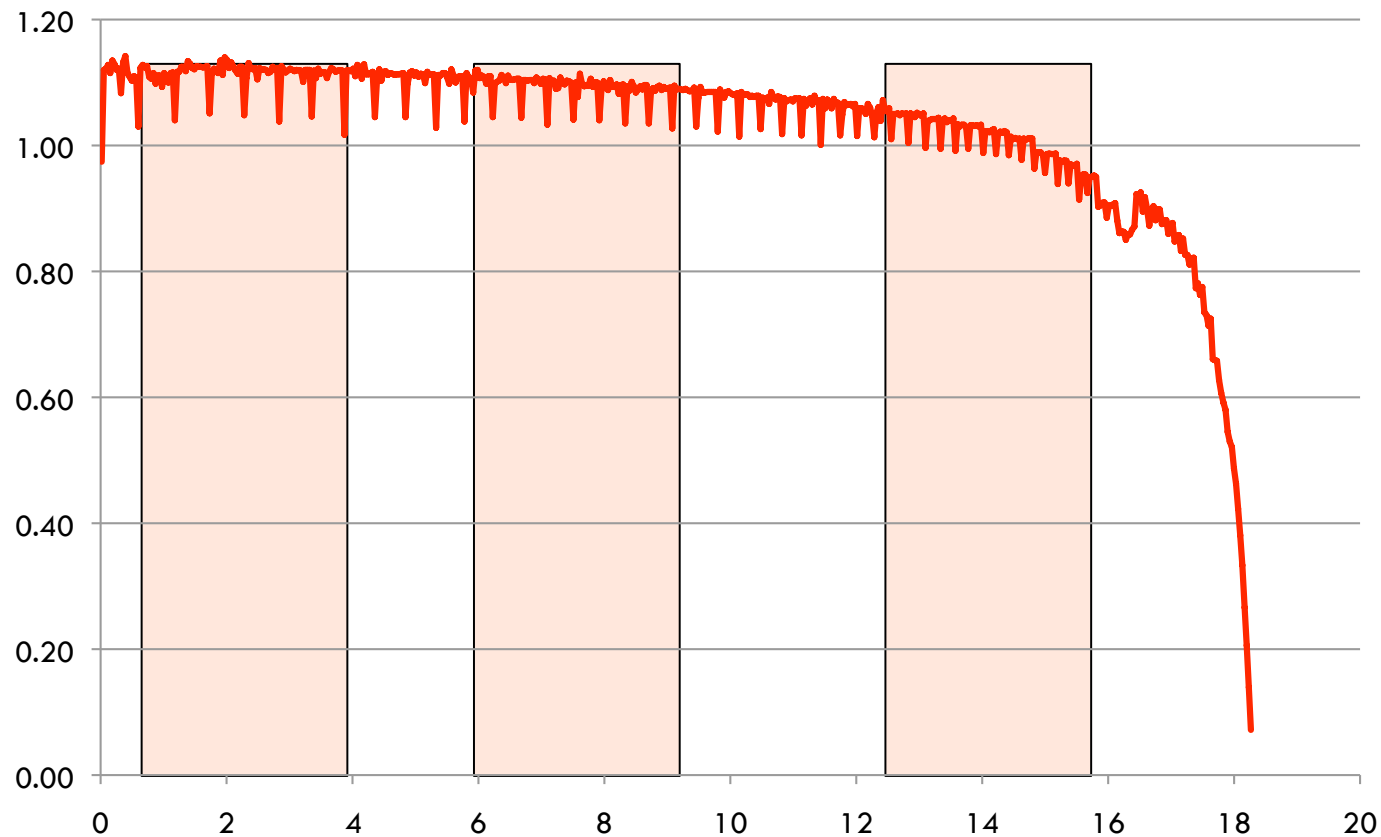
- Can't just start the run and stop it.
- The performance will vary over the course of the run.



# How to Capture Performance?

7


- Take a window of performance and use it.
- But what window?



# How to Capture Performance?

8

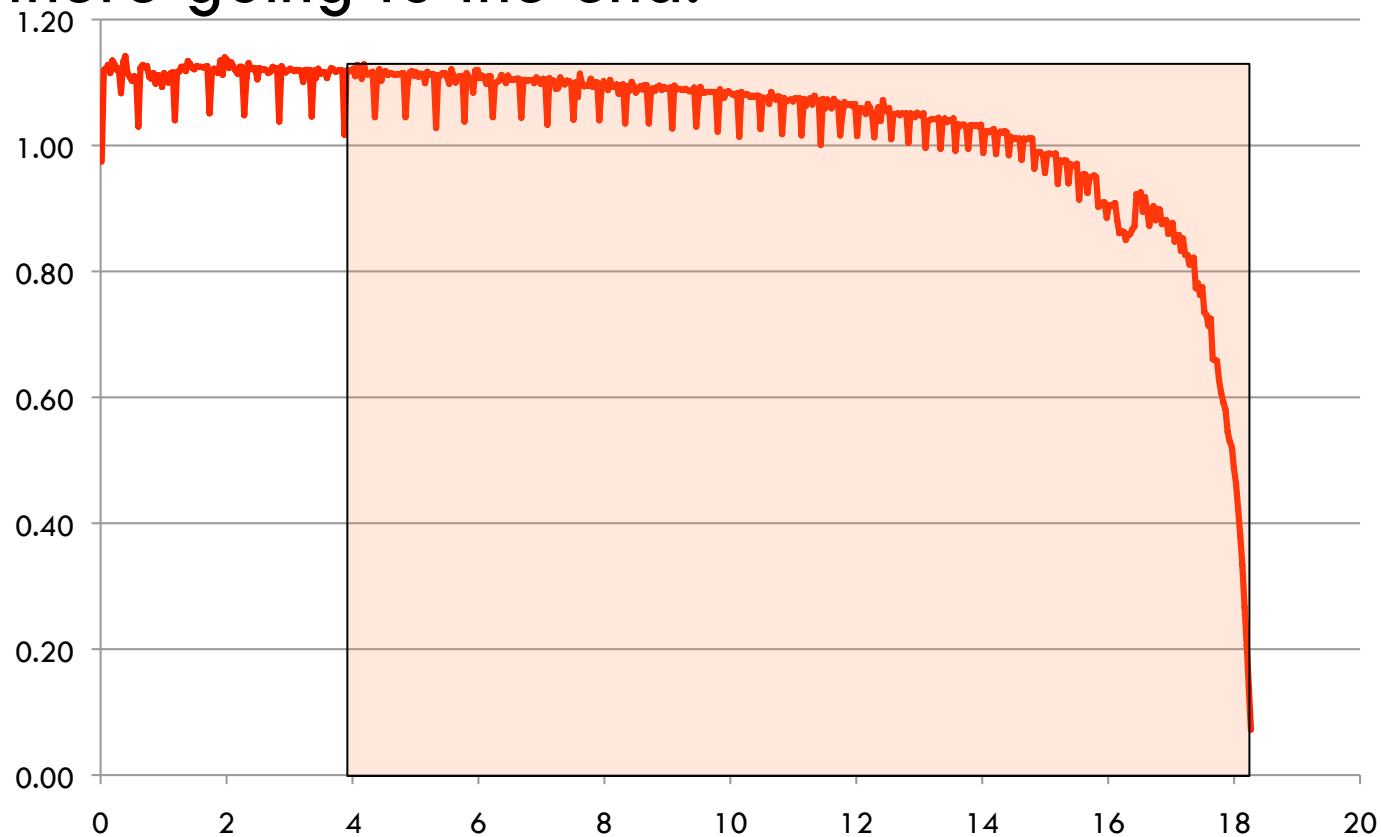
- Should we do sampling, like Monte Carlo?



# How to Capture Performance?

9


- Figure out the point to start (say what would have been 4 hours into the run) and begin the timing there going to the end.



# How to Capture Performance?

10

- Determine the section where the computation and communications for the execution reflect a completed run.



# HPL

- We are planning on making changes and will probably be ready after ISC in Hamburg.