CS4414. In class assignment: estimate what is feasible.

Suppose you want to solve a linear system $Ax = b$, where A is a non-sparse NxN matrix. What is the largest N your laptop can handle in 24 hrs, using Mathematica's LinearSolve[] with all default settings?

1 What to submit

Individual work: each student submits his her own work at the end of class.

2 Key Points

1. laptops vary by about factor of 2 in clock speed, and by about a factor of 4 in the number of cores, so should not expect more than an order of magnitude range. Class solutions: almost 5 orders.
2. Begin with looking up which methods Mathematica uses. LAPACK. Good.
3. Show Mathematic script, keep adding features.
4. How to generate sparse matrices. Random. Pros and Cons.
5. start 2x2. Check what it looks like
6. check residual to make sure LinerSolve[] is working.
7. Check Condition number. $C = 10^k$, loose k digits of precision. Is this OK for your problem?
8. use unix time. Pros and cons.
9. First 2 points. N=1000 t=0.384 sec. N=2000 t=0.455. Extrapolates to 0.071 sec per 1000. N= 24 x 3600 /0.071 x1000 per day.
10. Wrong. need more points
11. Do not use blind Fit. $O(N^3)$ is generic.
12. Keep checking residuals. Why is this not a guarantee?
13. Still not right. Memory limit.

14. Use Mathematica built in function and unix top.

15. Xmgrace to fit. Mention log-log plots to analyze $y = ax^k$.

16. Once you hit memory limit (8G on my laptop), performance deteriorates. This occurs at N=32,000 or so, takes 45 mins.

17. In the end, can not do more than N=40,000 or so.