Homework 3, due: 02/10

MATH 9830, Spring 2015 Timo Heister, `heister@clemson.edu`

0. Before you start:
   - Read the description of tutorial step-3, optionally watch the linked videos.
   - Submit your code via email. Print/write down all output mentioned below.
1. The residual in step m (measured in the A-norm) of the conjugate gradient algorithm satifies
$$\|x_m - x\|_A \leq 2\left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^m \|x_0 - x\|_A,$$
where $Ax = b$ and κ is the condition number of the (preconditioned) matrix. Argue why we can expect to do $O(\sqrt{\kappa})$ iterations to reach a typical residual assuming $1 \ll \kappa$. Truncating the Taylor expansion of $\log(1+x)$ might be helpful.
2. Create a modified version of step-3 to test the CG algorithm in deal.II:
   (a) Create a loop in run() to test 3,4,5,6,7 global refinemements of the `hyper_cube`.
   (b) Solve the linear system with a relative residual of 10^{-8} (use `system_rhs.l2_norm()*1e-8` as tolerance) with CG and the following preconditioners: a) no preconditioner, b) PreconditionJacobi, c) PreconditionSSOR, d) SparseMIC (this is incomplete Cholesky). Remember to start from a zero solution each time.
   (c) Create a table with the number of unknowns and the number of iterations for each method (see `solver_control.last_step()`).
   (d) For each method estimate the complexity as $O(n^c)$ with some constant c.
   (e) Assuming the condition number of the matrix is $\kappa = O(h^{-2})$, are your results consistent with question 1?
   (f) State the total cost for the best preconditioner as a power of n using your knowledge about the cost of one CG iteration.