

Computer Science 206P – Scientific Computing (Wayne Hayes)

Assignment #1: Machine Representation of real numbers

Write a program to determine the following values on the openlab machines: the number of bits in the mantissa (aka “significant”), the machine epsilon, and the smallest and largest possible represented values. Do it for the following floating point types: single precision, double precision, Intel extended precision (Python float32, float64, and longdouble/float128, respectively), and if your language supports it, quadruple precision. You may assume all of these are a power of 2, but you may *not* use any library routines that automatically compute these values for you. Present the results in a table that has as many lines as there are types, and order the lines from least to most accurate representation. Your output table must have **exactly 7 columns, with no header line, separated by TABS (not spaces)**: (1) language (repeated for each row), (2) name of datatype, (3) size of the datatype in bytes, (4) number of mantissa bits determined by your program, (5) machine epsilon, and finally (6) smallest and (7) largest representable positive numbers. All the latter values should be printed using the normal output routine (eg, 10^{-7} is printed 1e-7 or 1E-07 or something very similar) using at least 6 figures of precision. If you are using C or C++ and using “printf” to output, then note that “long double” must be output as “%Llf”. There is a script on openlab called /home/cs206p/bin/a1syntax that will check that your output satisfies the I/O spec.