

Imperial College London  
Department of Computing

**Computer Systems (M2)**  
Pentium programming – lab exercise

You'll learn more about assembly language programming as the course progresses. In the meantime you're encouraged to run the following program. To do so, you'll need to login to one of the CSG machines running **Linux** (not Windows!). You can do this exercise individually or with others. It shouldn't take more than 40 minutes to complete. Have fun!!

\* Type the program in the box below into a file called `hello.s` **Do not make any mistakes!!!!**

<pre>segment .data msg     db 'Hello world!',0xA len     equ \$-msg  segment .text global _start  _start:     mov  eax, 5 outer:     mov  ebx, 1000000000  inner:     dec  ebx     jg   inner     dec  eax     jg   outer      mov  eax, 4     mov  ebx, 1     mov  ecx, msg     mov  edx, len     int  0x80      mov  eax, 1     mov  ebx, 0     int  0x80</pre>	<pre>switch to data segment declare and initialise variable msg set constant len = number of bytes in msg  switch to text (i.e. code) segment make _start visible outside of this file  program starts here number of times to repeat outer loop  repeat inner loop 1 billion times. type 1 followed by 9 zeros - do not type any more zeros!  execute this &amp; next instruction 1 billion times jump if ebx greater than zero to label 'inner' decrement eax outer loop counter jump if eax greater than zero to label 'outer'  linux system call 4, i.e. write () file descriptor 1, i.e. standard output address of variable 'msg' number of bytes in message to write interrupt Linux, i.e. Linux will write the message  linux system call 1 i.e. exit () error code 0, i.e. no errors interrupt Linux, i.e. Linux will exit the program</pre>
* Assemble into an object file version with: <code>nasm -f elf hello.s</code>	nasm is the Netwide assembler. The command will produce an object file named <code>hello.o</code> if there are no errors in file <code>hello.s</code>
* Then link into an executable program with: <code>ld -s -o hello hello.o</code>	ld is the Linux object file 'linker' which can (amongst other things) link several object files into one executable program.
Run the program with: <code>hello</code> or <code>./hello</code>	The program executes over 10 billion Pentium instructions! 5 billion <code>dec</code> instructions and 5 billion <code>kg</code> instructions.
* Find the size of the executable file with: <code>wc -c hello</code>	This is probably the smallest complete program you'll ever create during your degree!
* Find the size of the code and data with: <code>size hello</code> Why is there a difference between the sum of these sizes and the size of the executable program file?	The code size is given in the 'text' column! Try the command <code>file hello</code> for fun also.
* Run and time the program with: <code>/usr/bin/time -p hello</code> If you run the program several times, the times may differ. Why?	For a more verbose timing output use the command: <code>/usr/bin/time -v hello</code> View the cpu type using <code>more /proc/cpuinfo</code>
* Login to other CSG Linux machines with different processors and compare the speed.	For names of machines, go to the csg home page and click on Computer List on the panel. Try older/slower machines