The duration of this examination is 50 minutes

Please answer *all three* questions

This examination is partly based on the Intel Nehalem processor architecture, as described in the article "Nehalem: Intel's Future Processor and System", by David Kanter (Real World Technologies, 04-02-2008), which you should have available to you in the examination. Where the article is incomplete, you are invited to speculate using your understanding of the underlying architectural principles.

1 a Which parts of Nehalem's branch predictor need to be replicated for each thread?

b Nehalem's branch prediction may lead to branch execution time ranging from zero cycles to many. What is the branch execution time when a level-1 BTB hit occurs?

c What is the worst case?

d What intermediate cases might occur?

The four parts carry, respectively, 20%, 20%, 20%, and 40% of the marks.

2 a What happens when Nehalem's ROB is full?

b What *two* things happen in the Nehalem microarchitecture when an instruction is committed?

c Nehalem's level-1 data cache has an access latency of four cycles. If this were increased, for example to six, its capacity could be much larger. What would be the disadvantages of doing this?

The three parts carry, respectively, 20%, 30%, and 50% of the marks.

3 Consider the following code fragment:

```
      float A[N],B[N],C[N];
S1:   for (i=0; i<N; i++) {
S2      if (A[i] > 0) {
S3:       C[i] = A[i] + B[i];
        }
      }
```

a Explain (in outline terms) how this loop could be executed using SSE instructions.

b It has been proposed that the SSE-like instructions in some future designs might operate on much longer registers.

(i) What performance advantage might arise, in suitable applications?

(ii) What problem arises with the example above? How might it be minimised, in suitable applications?

The two parts carry equal marks.