

2011 Paper 9 Question 8

Optimising Compilers

- (a) Explain what is meant by “instruction scheduling” in a compiler. Indicate what properties of architectures make it beneficial. Are some architectures too simple, or too complex, for it to be useful? [3 marks]
- (b) Sketch an algorithm that schedules instructions within a basic block for an architecture of your choice on which scheduling is useful. Indicate its time complexity $O(f(n))$ stating to what ‘ n ’ refers. What additional complications arise in choosing which instruction to emit first from a basic block? [5 marks]
- (c) Indicate how your algorithm could be modified to deal with a *static multiple-issue* VLIW-style architecture. The input basic block contains simple instructions, but the output basic block contains wide instructions consisting of two simple instructions (either of which may be NOP). These two instructions execute in parallel.

Now explain how your algorithm can schedule the following sequence of instructions to execute in parallel, or detail any adjustments necessary to permit this.

```
st  r1,4(r8)
ld  r2,0(r8)
```

Explain how *alias analysis* or *points-to analysis* might provide information to the instruction scheduler to enable a wider range of load/store pairs to be re-scheduled. [6 marks]

- (d) Briefly summarise a source-level approach to alias or points-to analysis, for example Andersen’s algorithm. Your summary should emphasise key choices, for example any data-structures used; detailed algorithms are not required. [6 marks]