

Lecture 24

Thread-Level Speculation

Reference: "Compiler Optimization of Scalar Value Communication Between Speculative Threads", by Antonia Zhai, Christopher B. Colohan, J. Gregory Steffan and Todd C. Mowry. ASPLOS, 2002.

Automatic Parallelization

Proving independence of threads is hard:


- complex control flow
- complex data structures
- pointers, pointers, pointers
- run-time inputs

How can we make the compiler's job feasible?

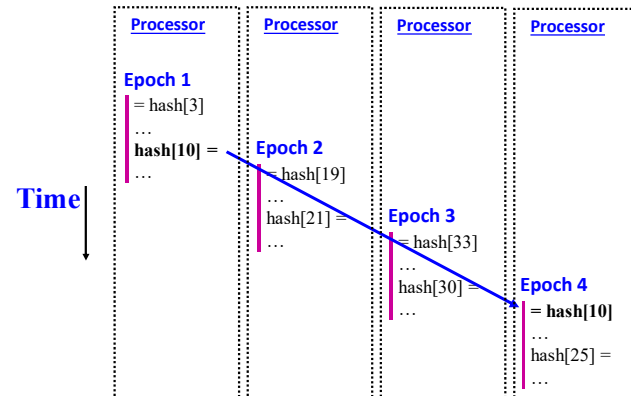
Thread-Level Speculation (TLS)

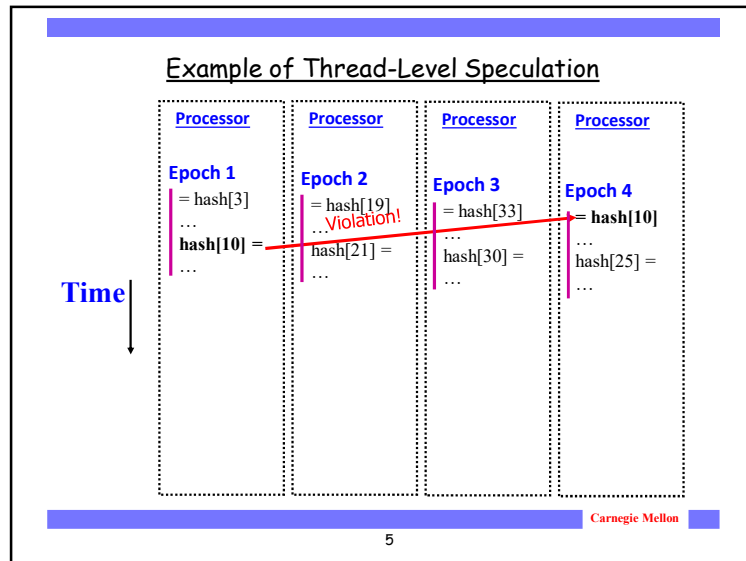
Example

```
while (...){
  x = hash[index1];
  ...
  hash[index2] = y;
  ...
}
```

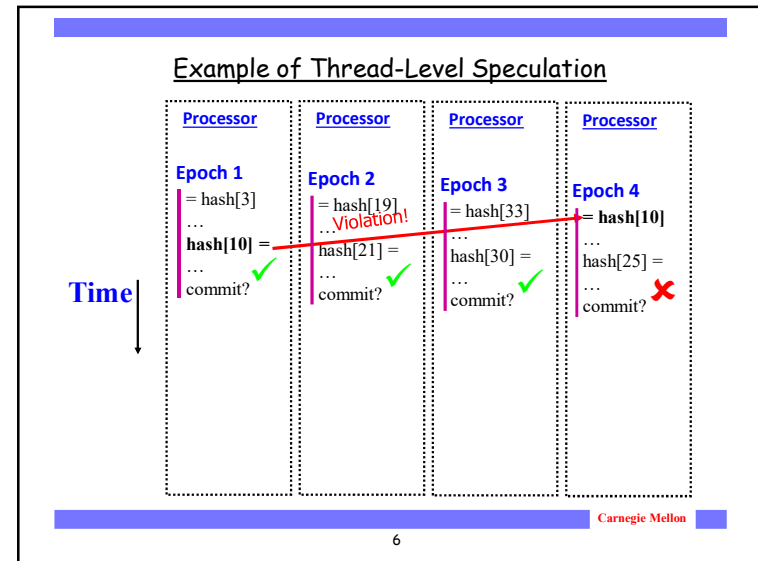


Example of Thread-Level Speculation

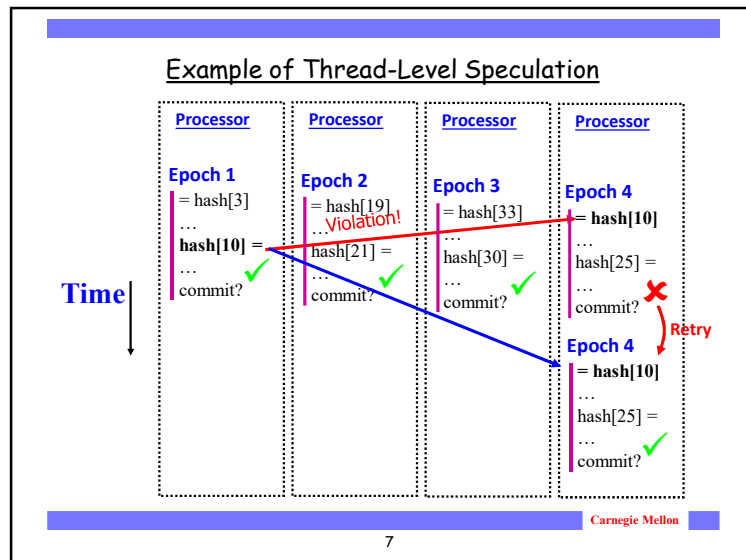




5



6



7

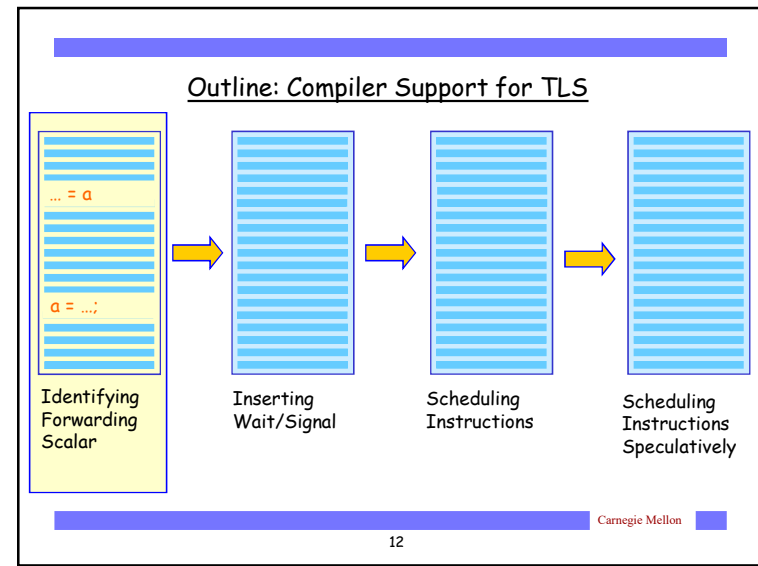
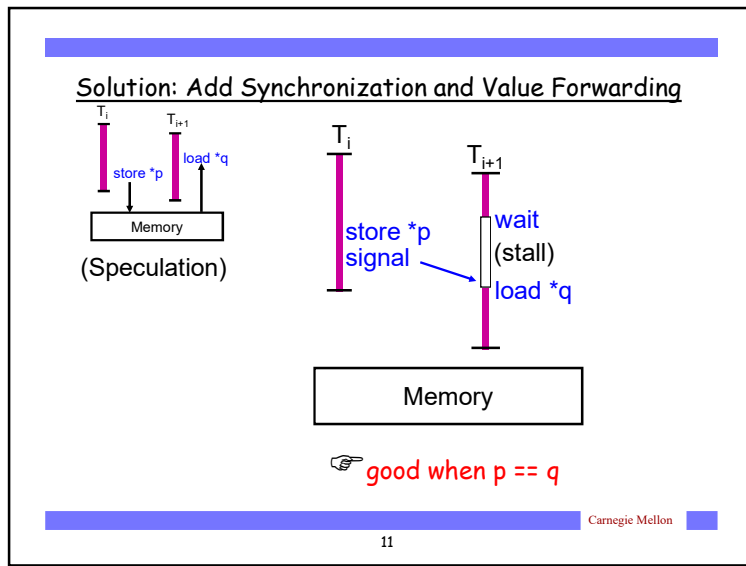
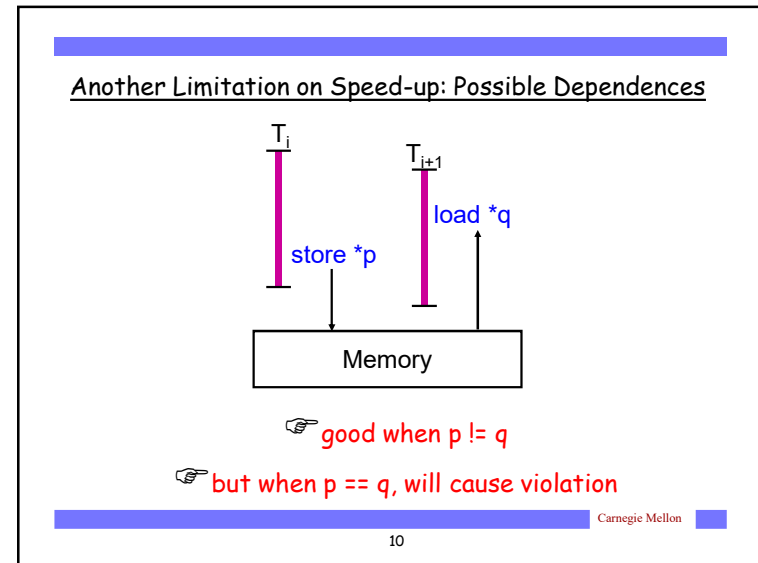
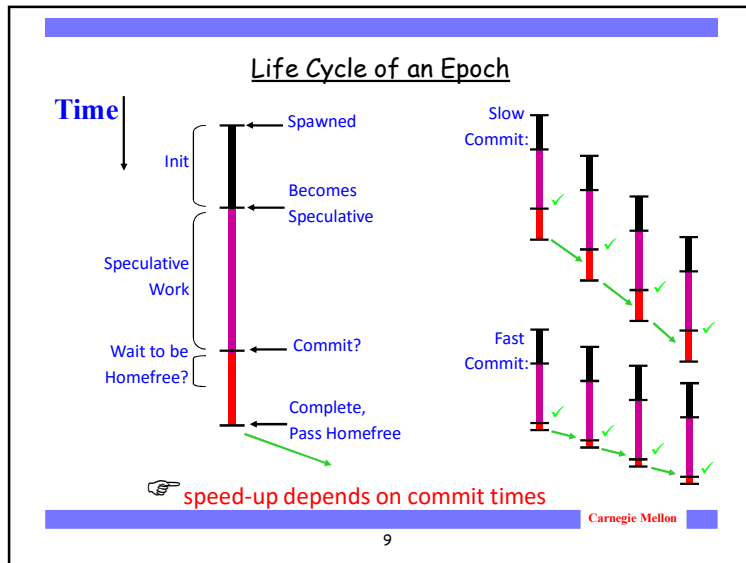
System Requirements for Thread-Level Speculation

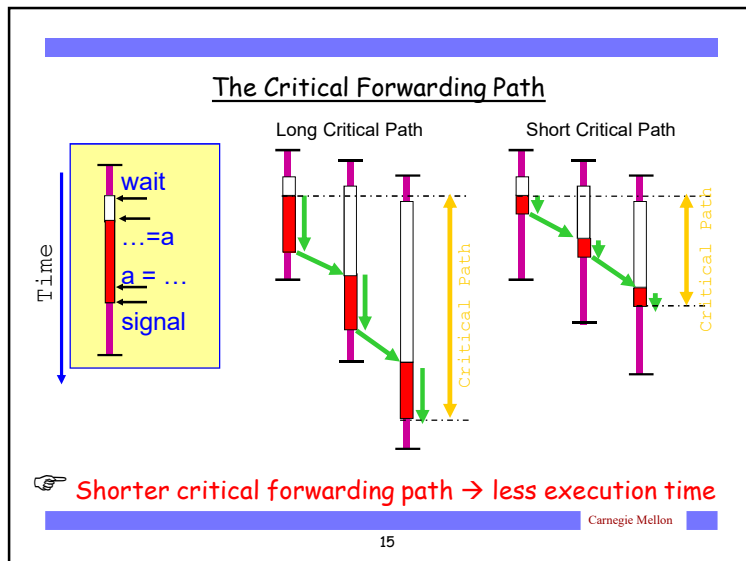
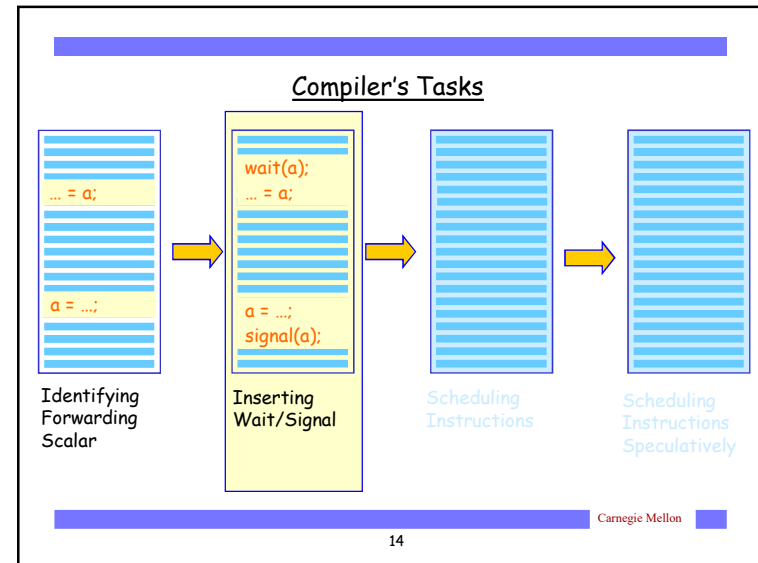
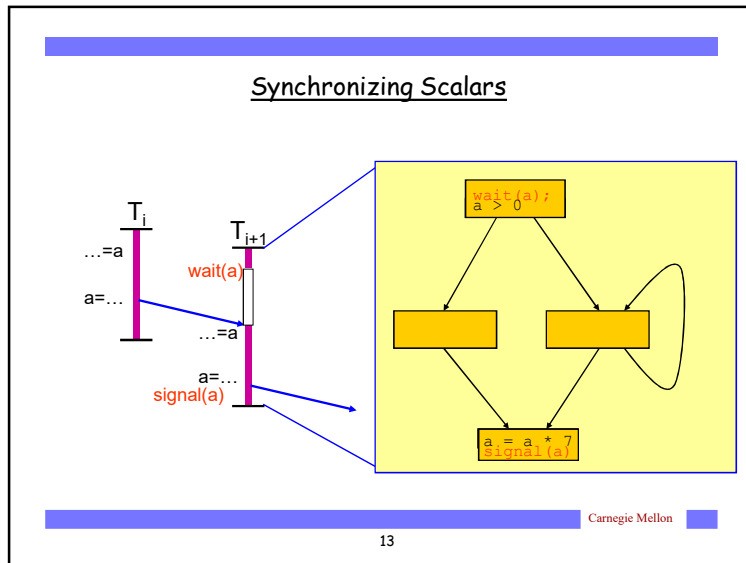
System requirements:

- 1) Detect data dependence violations
 - extend invalidation-based cache coherence
- 2) Buffer speculative modifications
 - use the caches as speculative buffers

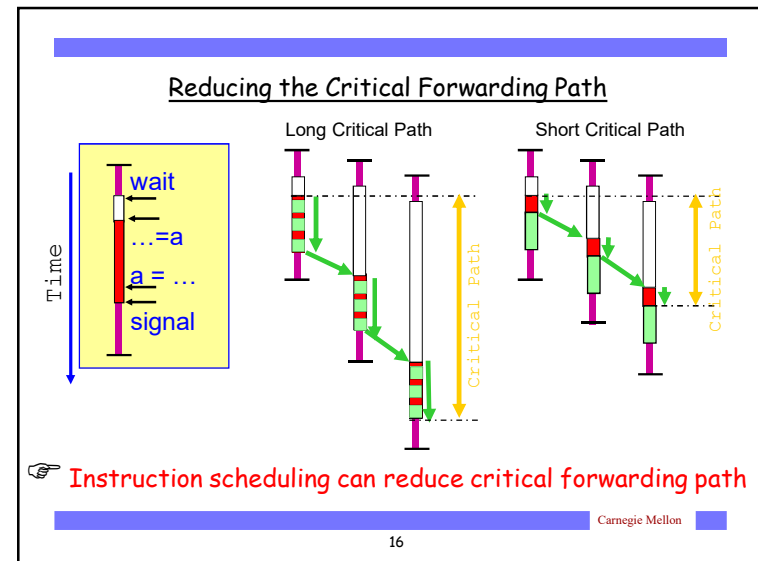
Carnegie Mellon

8

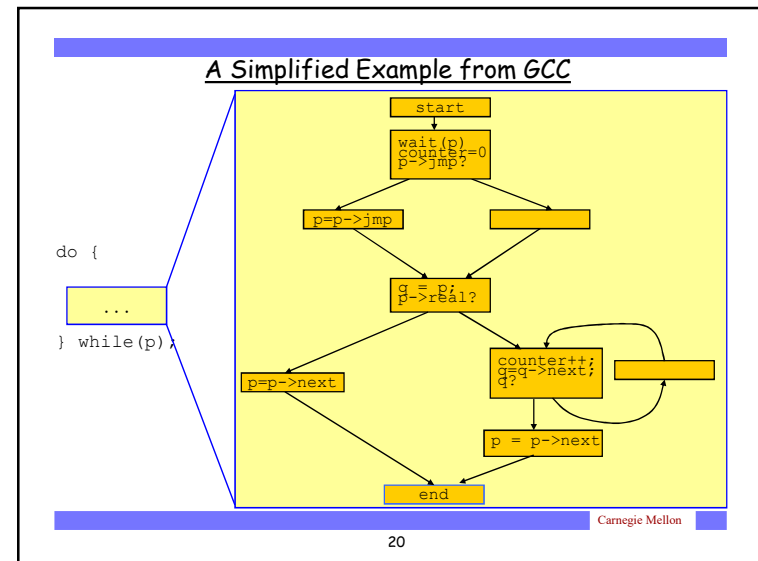
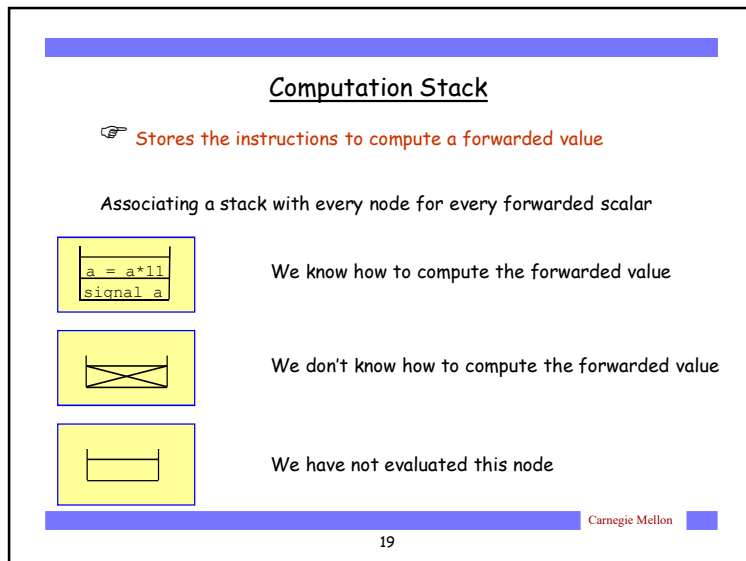
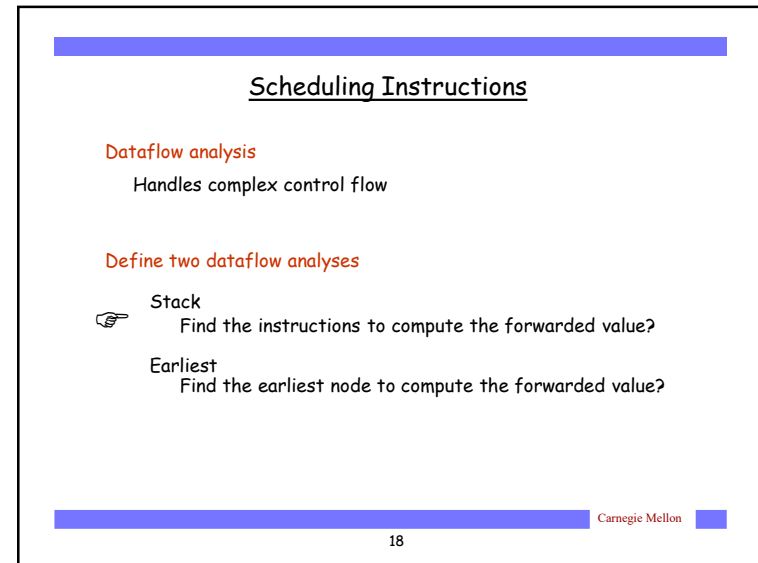
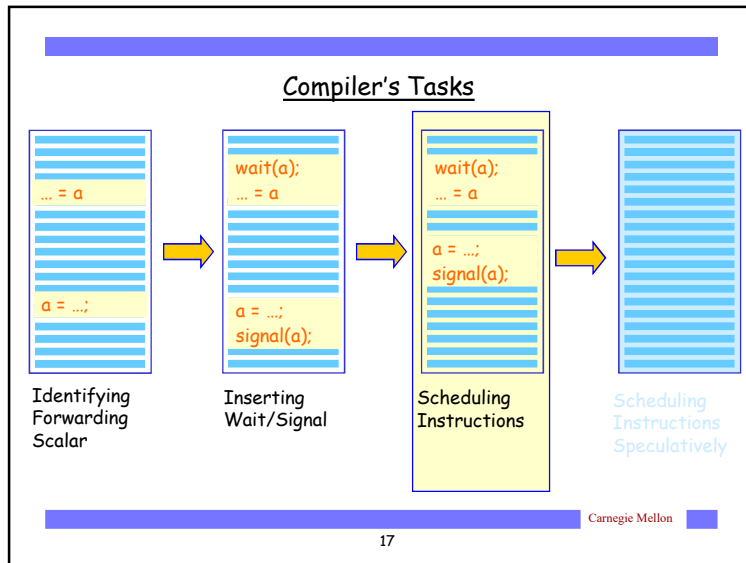


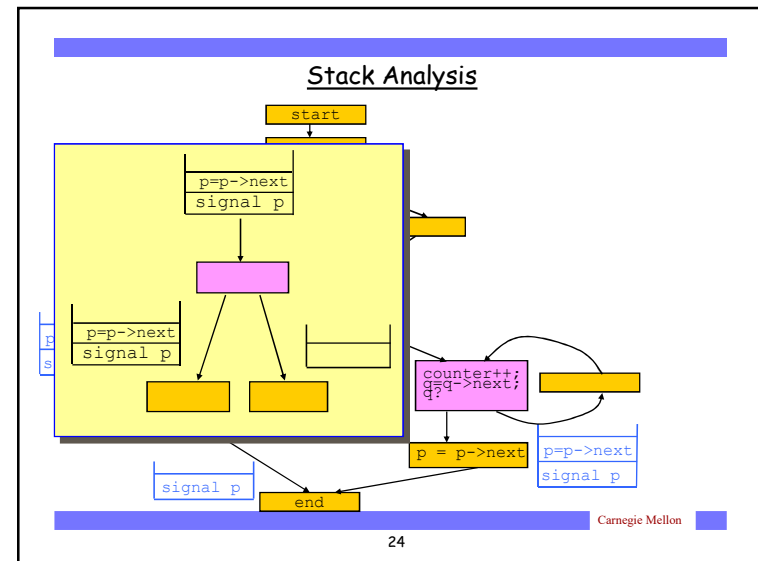
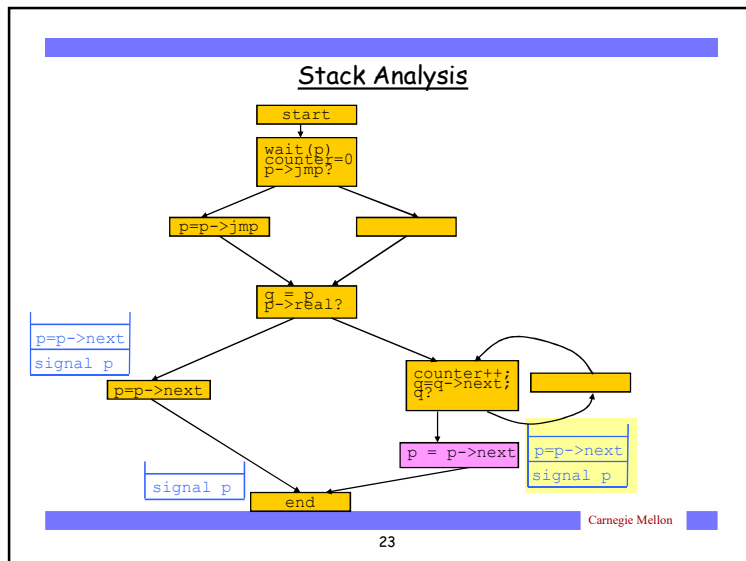
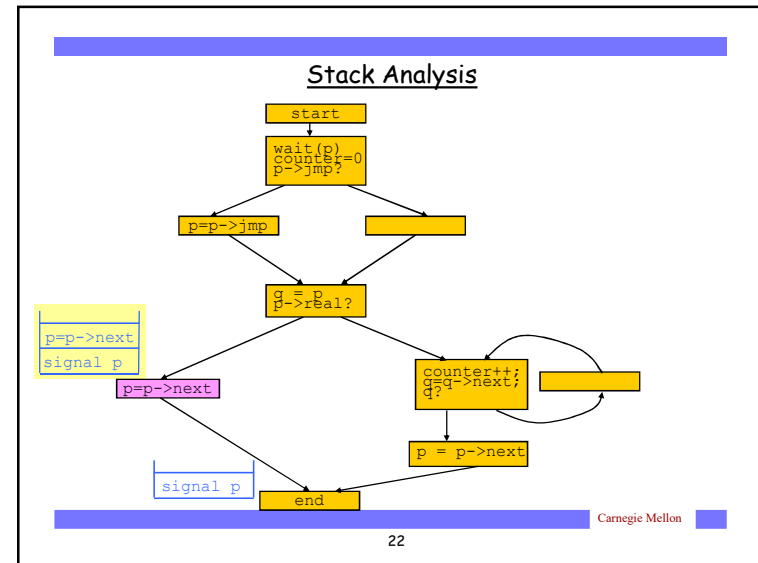
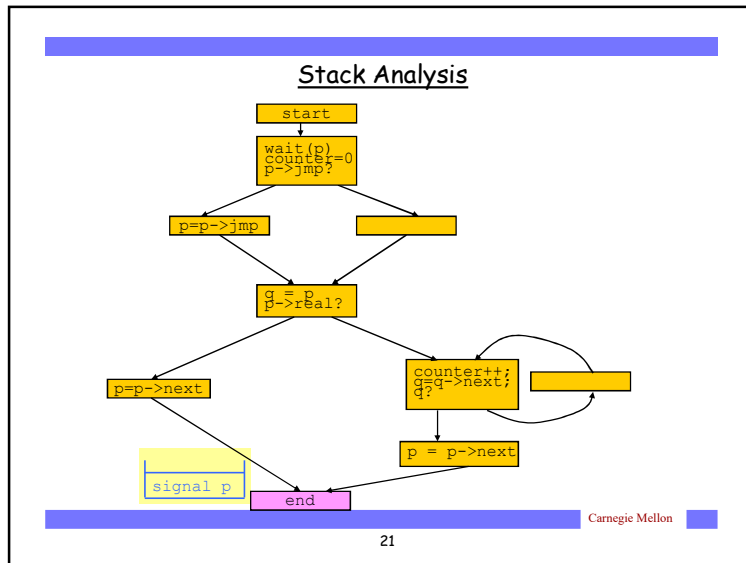


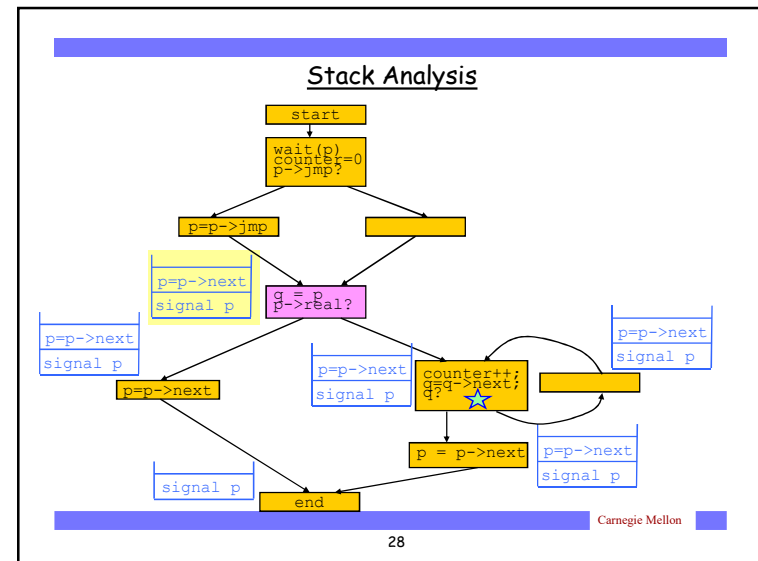
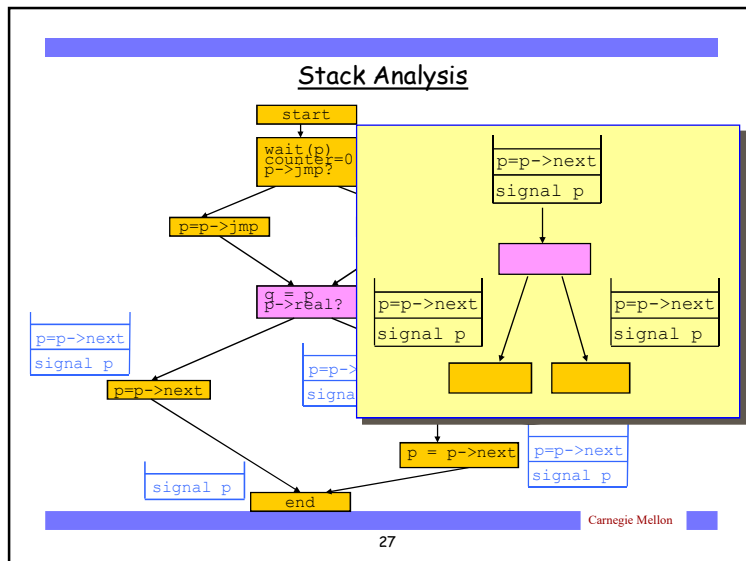
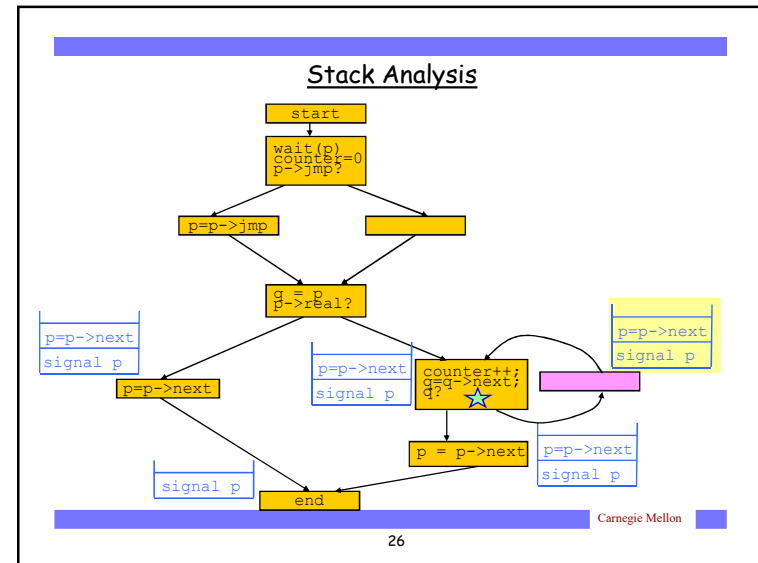
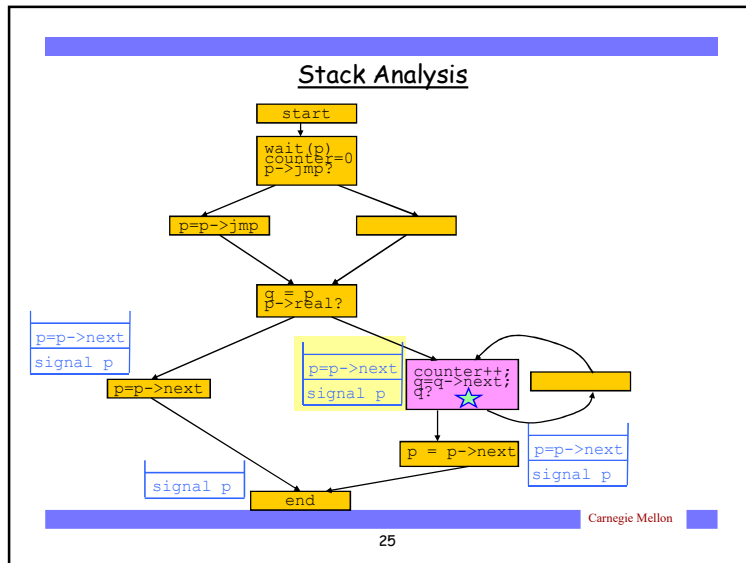
Shorter critical forwarding path → less execution time

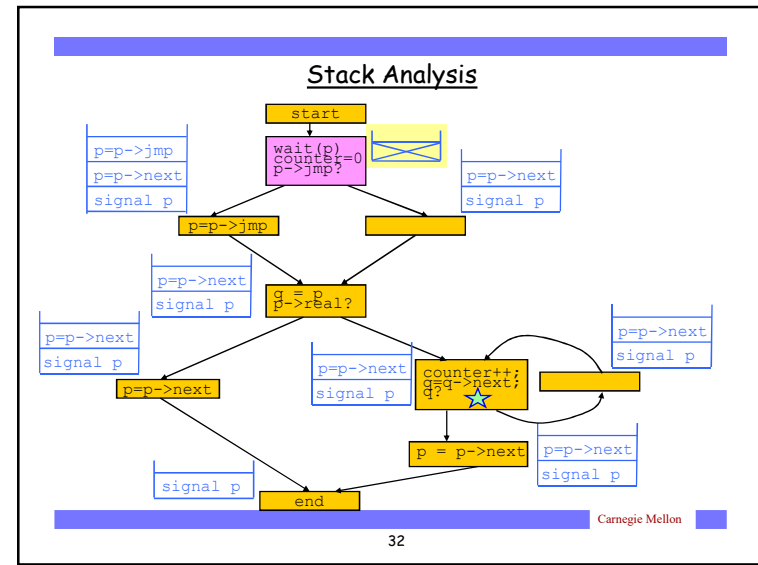
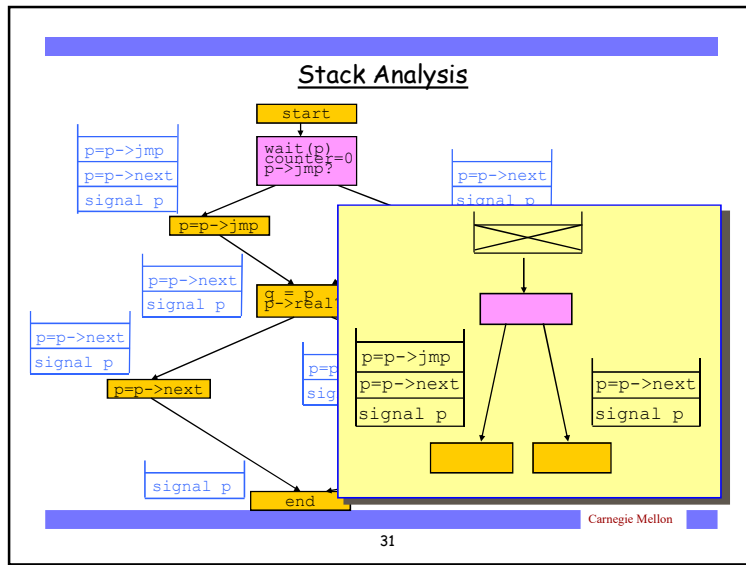
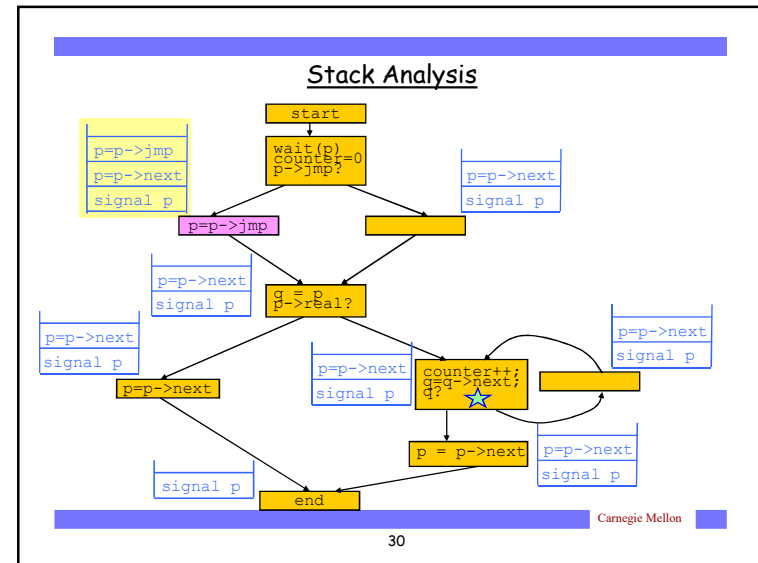
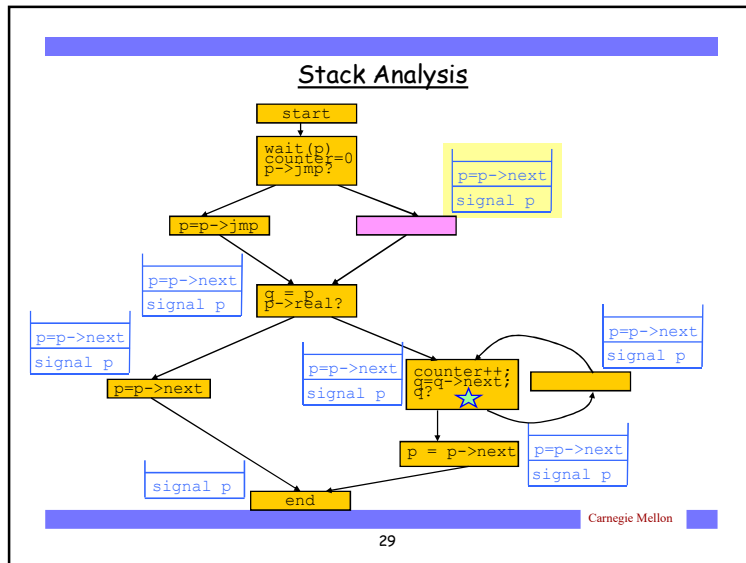


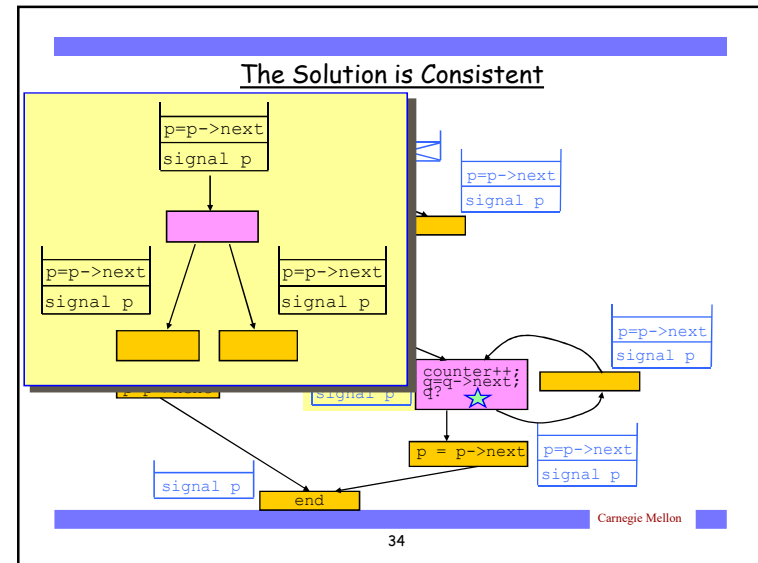
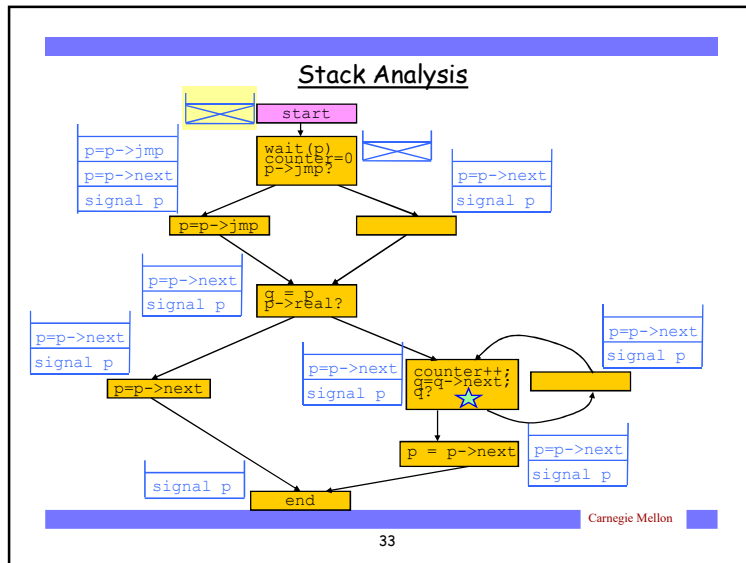
Instruction scheduling can reduce critical forwarding path











Scheduling Instructions

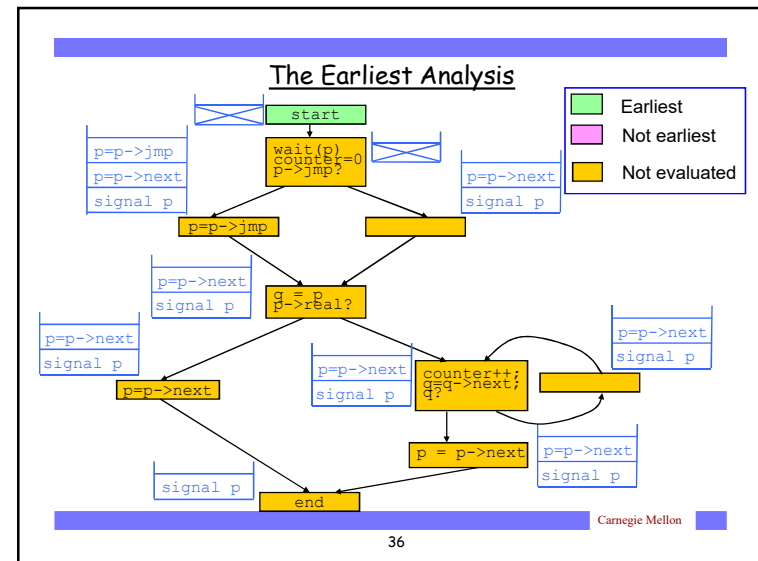
Dataflow analysis
Handles complex control flow

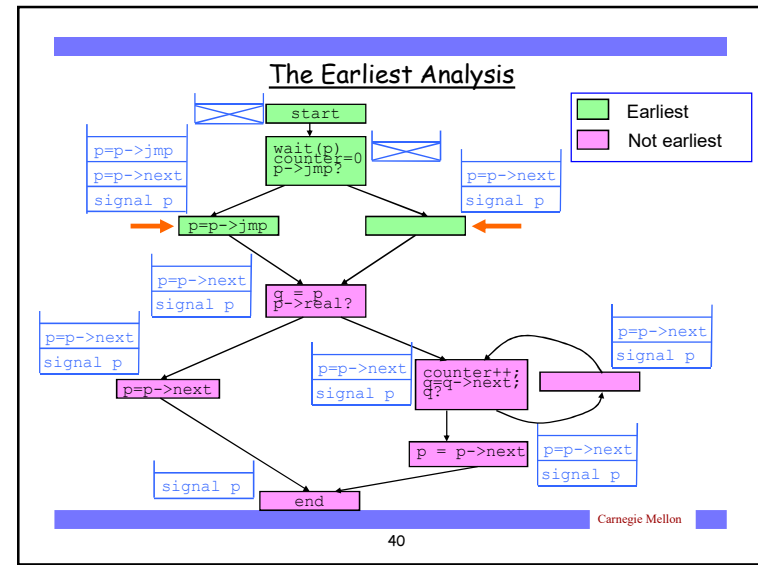
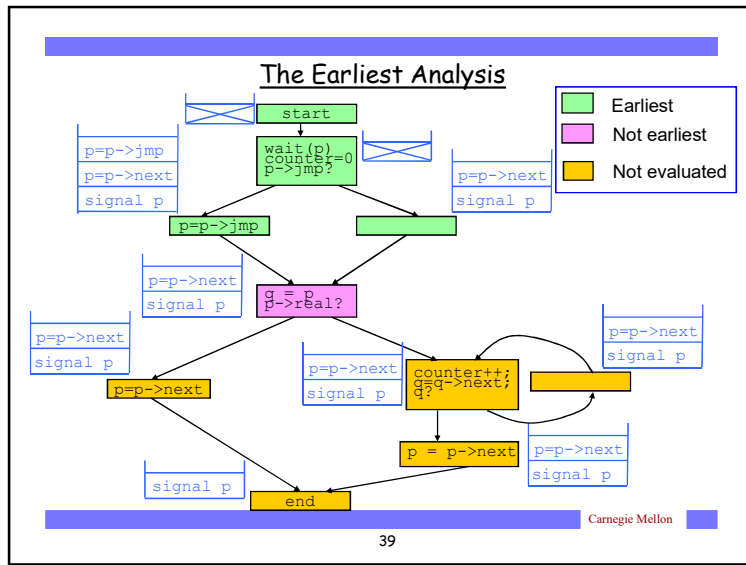
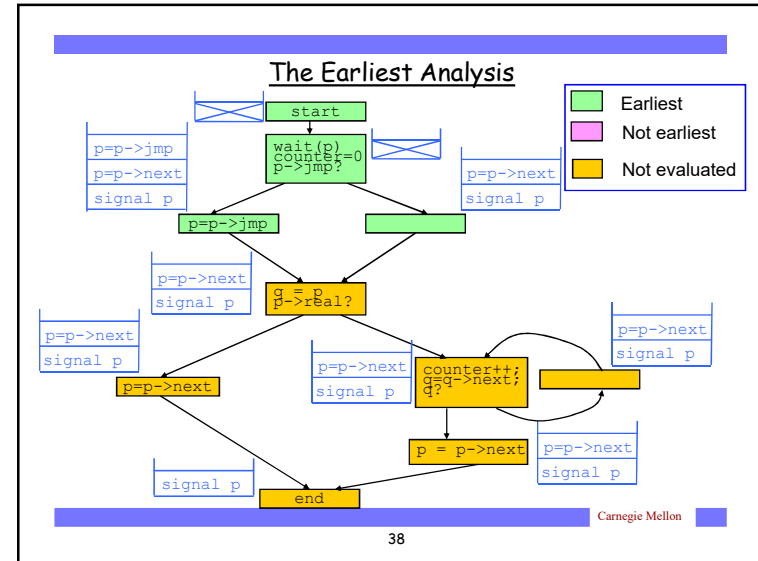
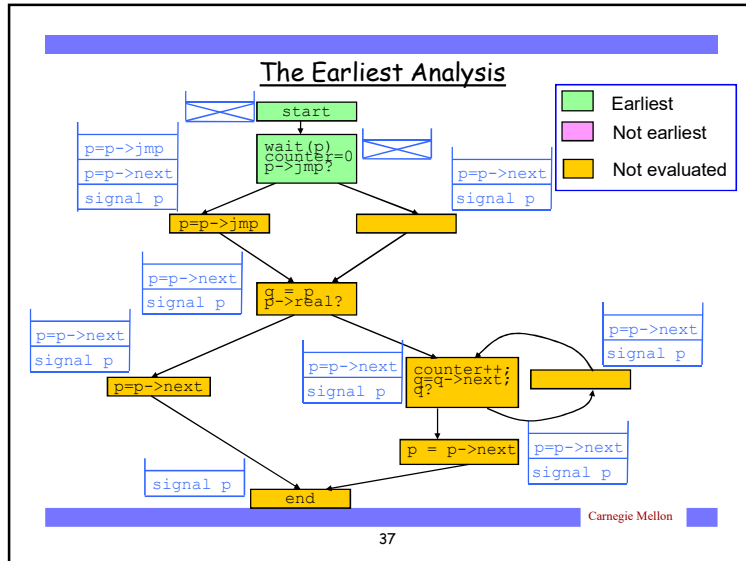
Define two dataflow analyses

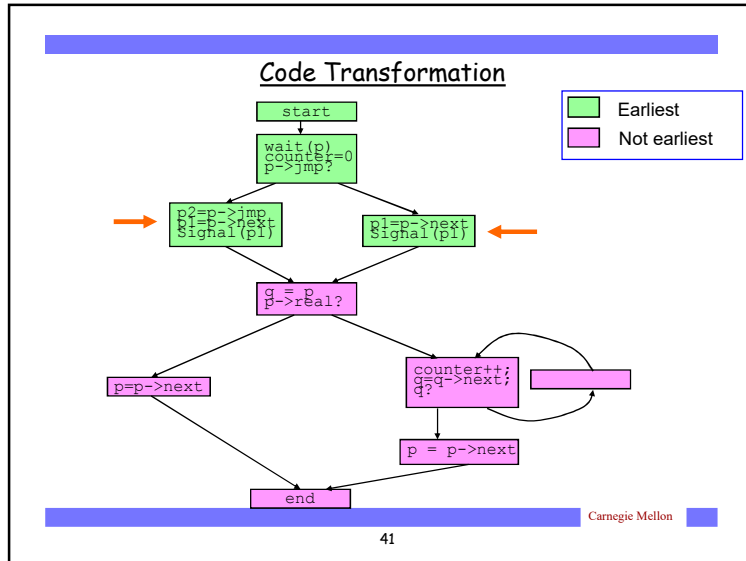
Stack
Find the instructions to compute the forwarded value?

Earliest
Find the earliest node to compute the forwarded value?

35 Carnegie Mellon







Experimental Framework

Benchmarks

- from SPECint95 and SPECint2000, -O3 optimization

Underlying architecture

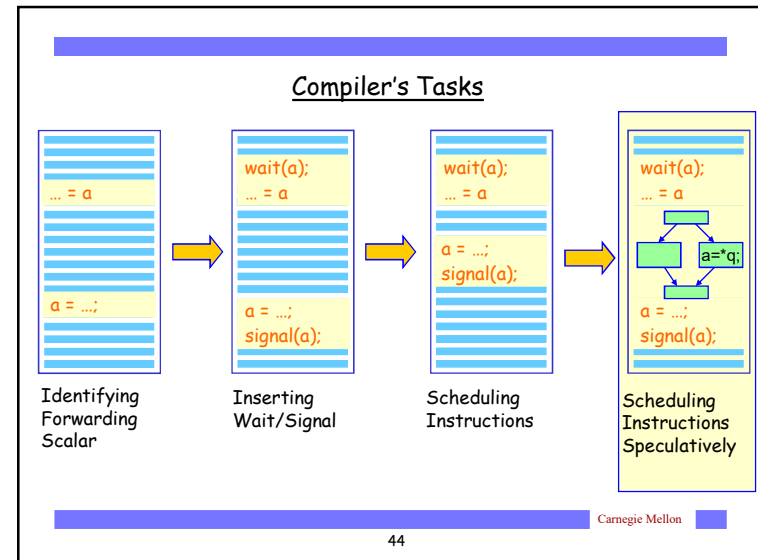

- 4-processor, single-chip multiprocessor
- speculation supported by coherence


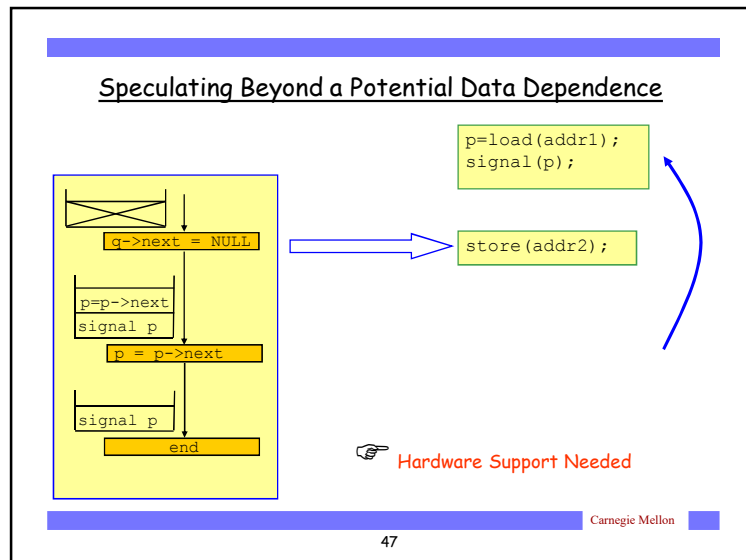

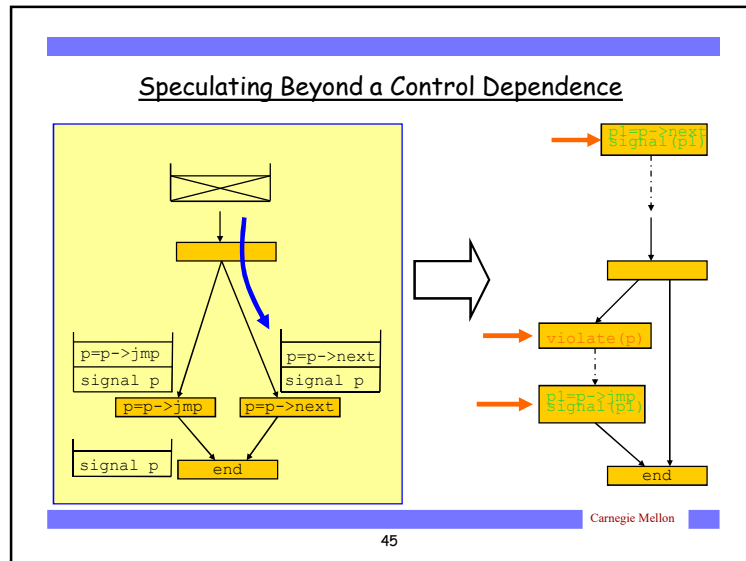
Simulator

- superscalar, similar to MIPS R10K
- models all bandwidth and contention

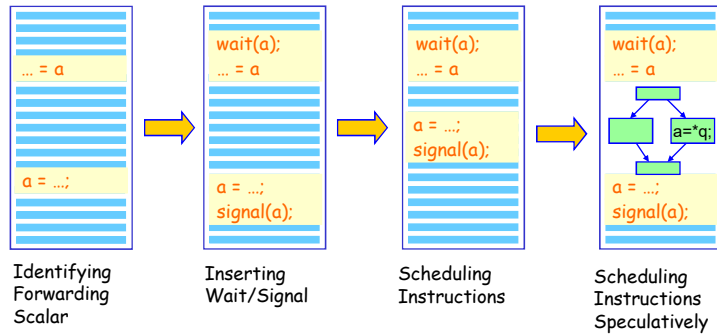
detailed simulation!

42 Carnegie Mellon





Summary: Compiler Support for Thread-Level Speculation



49

Carnegie Mellon

Conclusions

Instruction scheduling for reducing synchronization

- Is effective in reducing critical forwarding path
 - Performance improved by 18%
- Is beneficial to handle complex control flow, such as inner loops
 - Improved GCC by 3%
- Gives additional benefit with speculative instruction scheduling
 - One biggest benefactor is GCC, performance improved by 18%
 - Some hardware support needed



Critical forwarding path can be addressed by the compiler

50

Carnegie Mellon

What's Next

- Project Milestone Reports due midnight Monday
- Exam on Wednesday

51

Carnegie Mellon