
CSAIL

Computer Science and Artificial Intelligence Laboratory


Massachusetts Institute of Technology

Tests for Monsoon Instruction Subset 1

J. Young

1990, September

Computation Structures Group
Memo 307



The Stata Center, 32 Vassar Street, Cambridge, Massachusetts 02139

**LABORATORY FOR
COMPUTER SCIENCE**



**MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY**

Tests for Monsoon Instruction Subset 1

Computation Structures Group Memo 307 (Rev. 1)
September 11, 1990

Jonathan Young

This report describes research done at the Laboratory of Computer Science of the Massachusetts Institute of Technology. Funding for the Laboratory is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-84-K-0099.

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Tests for Monsoon Instruction Subset 1

Jonathan Young

September 11, 1990

1 Introduction

This document describes a series of tests which are designed to test increasingly more complicated portions of the Monsoon Macroarchitecture Specification [2]. Our goal is to run these tests on the hardware simulator, on MINT and other software emulators, and eventually on the actual hardware when it arrives.

This is an evolving document. The tests currently described herein test only those instructions in instruction subset 1 (IS1) [1].

2 The tests

All tests begin execution with a single token in the pipeline and terminate with the pipeline idle. Our goal has been to create small tests (no more than 30 instruction and 100 pipeline beats). The RTS test takes 106 pipeline beats, and the LOCK test takes over 200 beats, but the rest of the tests are fairly small.

To each test corresponds a "base number" n . The code for each test generally starts at $\#xn0$, the frame pointer at $\#xn0$, and I-structure memory (if used) at $\#xn00$, the initial token for the ISTR test should be $\langle L, \#x50, \#x50 \rangle$ (pointer to $\#x500$).

Test	Base	Instructions tested
FORK-JOIN	1	one fork, one join, identity
FLOATS	2	fadd, fsub, fmul, fdiv, fltp, feqp
BOOLS	4	and, or, xor, ash and .nl1
ISTR	5	faif, istr, %%istr
SEND	6	ct, aoct
	7	(reserved for SEND)
RTS	8	stake, id.st1, id.rnl1, ap
MISC	$\#xA$	remov, swt
MDEF	$\#xB$	multiply-deferring i-structures
LOCK	$\#xC$	multiply deferring locks
PLTP	$\#x10$	PLT, PLP, %%PLMEM, %%PLMEM1
SVC	$\#x11$	synchronous traps

Note that the graphs presented at the end of this document depend crucially on the ability to predict which of two tokens will be executed first. In particular, the fanout ($\mathbf{fo1.m2}$) instruction produces tokens at both dest1 and dest2. Most of these graphs assume that the token at dest1 is produced first. *Not all of these graphs are completely consistent with the macroarch yet.*

The graphs also assume that the token queues are configured as a *stack*. On the gate-level simulator this assumption was not the case.

3 FORK-JOIN

This tests the basic execution of the machine: monadic and dyadic matching, the identity alu operation, and the generation of 0, 1, or 2 output tokens. The first instruction split the input token into two, which

are joined by the second instruction and terminated by **stop**, the third instruction. Note that since **stop** is the last instruction in the current microcode, this also enforces the declaration of unused opcodes and second-level-decode entries.

4 FLOATS

This test pushes a known value (5.5) through 4 floating-point ALU operations (**fadd**, **fsub**, **fmul**, and **fdiv**) in parallel. We also exercise different matching and output modes: first, we match with an absolutely-addressed constant and issue one token (**.n1**); second, we issue 2 tokens (**.n2**); third, we do a frame-relative dyadic match which issues 2 tokens (**.n2**). Finally, we do a frame-relative dyadic match which issues only 1 token; these 4 tokens are fed into two floating compares which in turn feed a logical **and**. The result should be boolean false.

5 BOOLS

This test exercises the floating-point comparison and (integer) boolean operators.

6 ISTR

This test executes several two-phase memory transactions - three I-fetches and two I-stores, to be precise. We test I-fetch from both present and empty locations, as well as I-store to both empty and deferred. (We do not test the full functionality of the **%istr** instruction, only that need for subset 0.)

This test is superseded by the MDEF test. As written, this test will not run under IS1.

7 SEND

This test uses the **ct** and **aoct** instructions to test a rudimentary function-call sequence.

8 RTS

This tests the instructions needed to implement a prototypical run-time system on the Monsoon machine which allocates storage (i.e. **get-context** and **get-aggregate**) but cannot deallocate it. The instructions particularly needed are **stake** (take on read-only, spin on empty) and **id.st1c** ("put").

This test executes two **get-context** and two **get-aggregate** operations in parallel; one of each spins for several cycles while the critical section of the other one is executed.

Note that because of the limitations of IS0, we are still taking advantage of the pointer/continuation duality.

9 MISC

This final test exercises "everything else" in IS0. We test dyadic matching when the left token arrives first, and we test the switch (**swt**), fix (**fix**), and remove (**remov**) opcodes. We also test loading a global constant.

10 MDEF

This tests multiply deferring I-structures. Three fetches occur before any store.

We do not yet test delays.

11 LOCK

This tests locks - resources to which only one process may have access at a time. Access is gained by a TAKE and released by a PUT instruction; additional instructions in user code (TAKE-AUX and TAKE-AUX1) are required to allow queueing of multiple takes.

There are a total of 5 takes in this test, labeled A, B, C, D, and E. The approximate order of events is as follows: TAKE(A), TAKE(B), TAKE(C), PUT (returns to C), C sends retake to B. Meanwhile, TAKE(D) and TAKE(E) happen, so when B does the retake - TAKE(B) - it also defers, getting a pointer to E. *This is where the two deferred lists are merged*: B sends the E continuation to A, who stores it. Eventually, more puts happen, and everyone gets their chance to play with the resource.

12 PLTP

This test runs the PLT (processor-local take) and PLP (processor-local put) instructions. These instructions do a two-phase-transaction-style synchronized read (write) of a location *without giving up the thread on the processor*, for use in exceptions. The values are constrained to have only one reader and one writer, like locks, but additionally, neither operation can defer.

13 SVC

This tests the SVC (synchronous trap) instruction and accompanying support instructions, %EXC1 and %EXC2. When used at the entry point to an exception, they store the continuation and XA in appropriate locations in the frame, passing XB to the next instruction in the exception handler.

This exception handler adds A and B and returns the value to IP+1.L (this is the conventional place for exceptions to return values).

14 Testing Limitations

The simulation of only 600 instructions on our machine must inherently leave something out. In particular, we did not test the FALU operations in any comprehensive manner, but only enough to establish that the correct operation was being performed on the correct (e.g. not flipped) arguments. We do not test every second-level-decode of every macroinstruction; we intentionally take advantage of the microcode structure to test matching modes and ALU operations separately.

15 Appendix: Test Graphs

The dataflow graphs for the tests, annotated to include each token expected during execution, appear on the following pages.

16 Acknowledgements

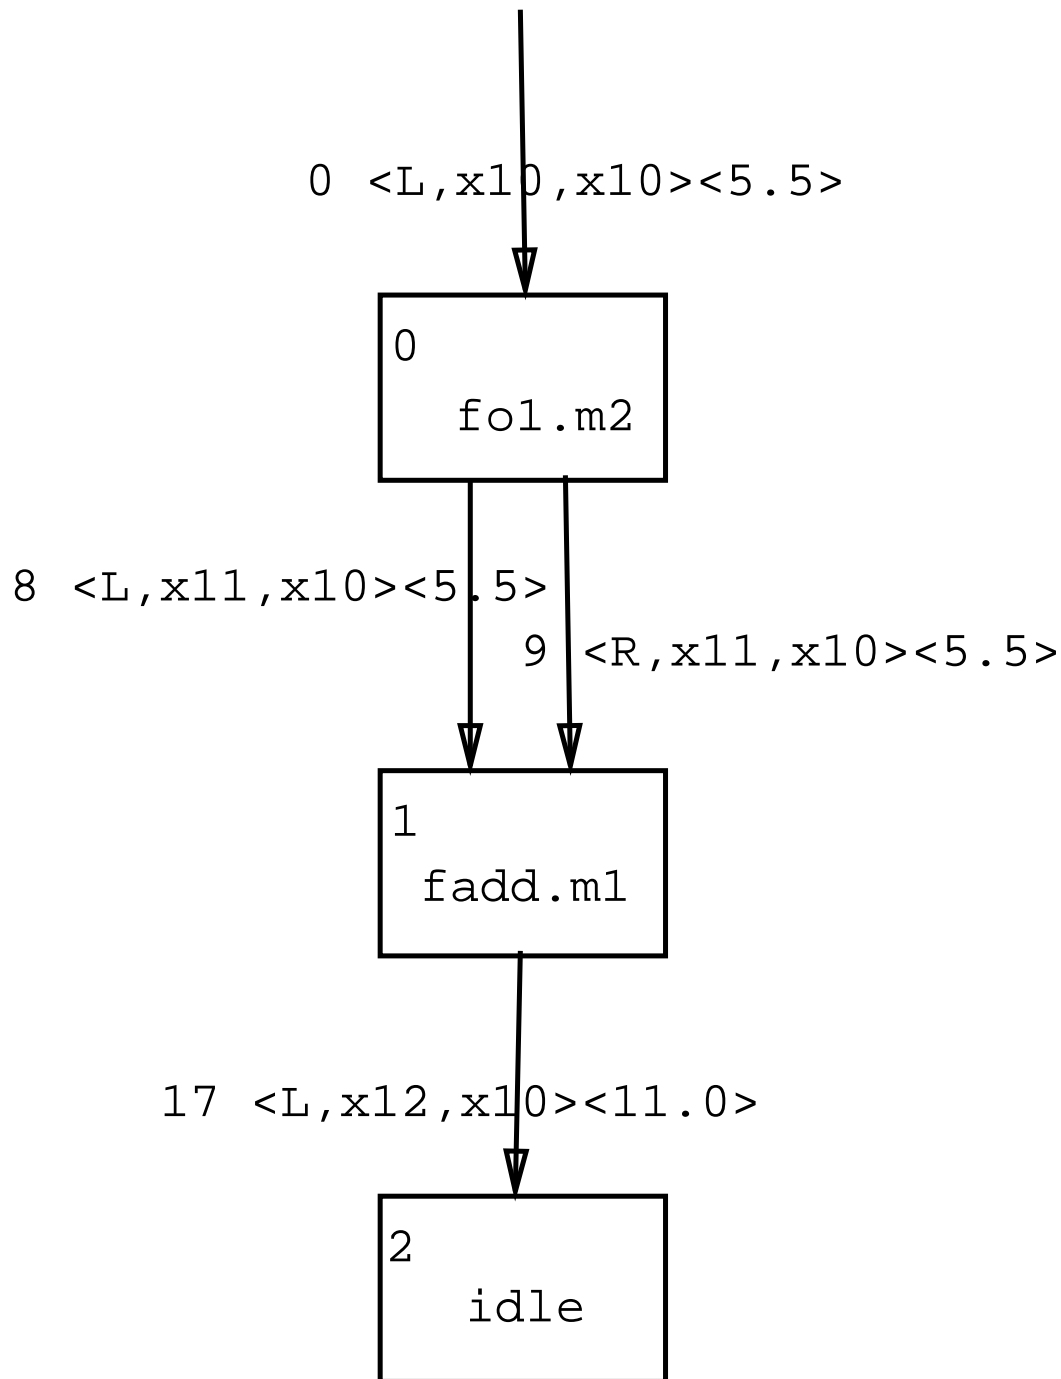
This document is the combined effort of several people. I am particularly indebted to Madhu Sharma, who drew the pictures and debugged my code, and to Ralph Tiberio, who first executed these graphs on the hardware simulator.

References

- [1] M. J. Beckerle and J. H. Young. Monsoon instruction subsets. Computation Structures Group Memo 306a, Massachusetts Institute of Technology Laboratory for Computer Science, Cambridge MA, 1988.

- [2] G. M. Papadopoulos and K. R. Traub. Monsoon macroarchitecture reference manual. Computation Structures Group Memo 306, Massachusetts Institute of Technology Laboratory for Computer Science, Cambridge MA, 1990.

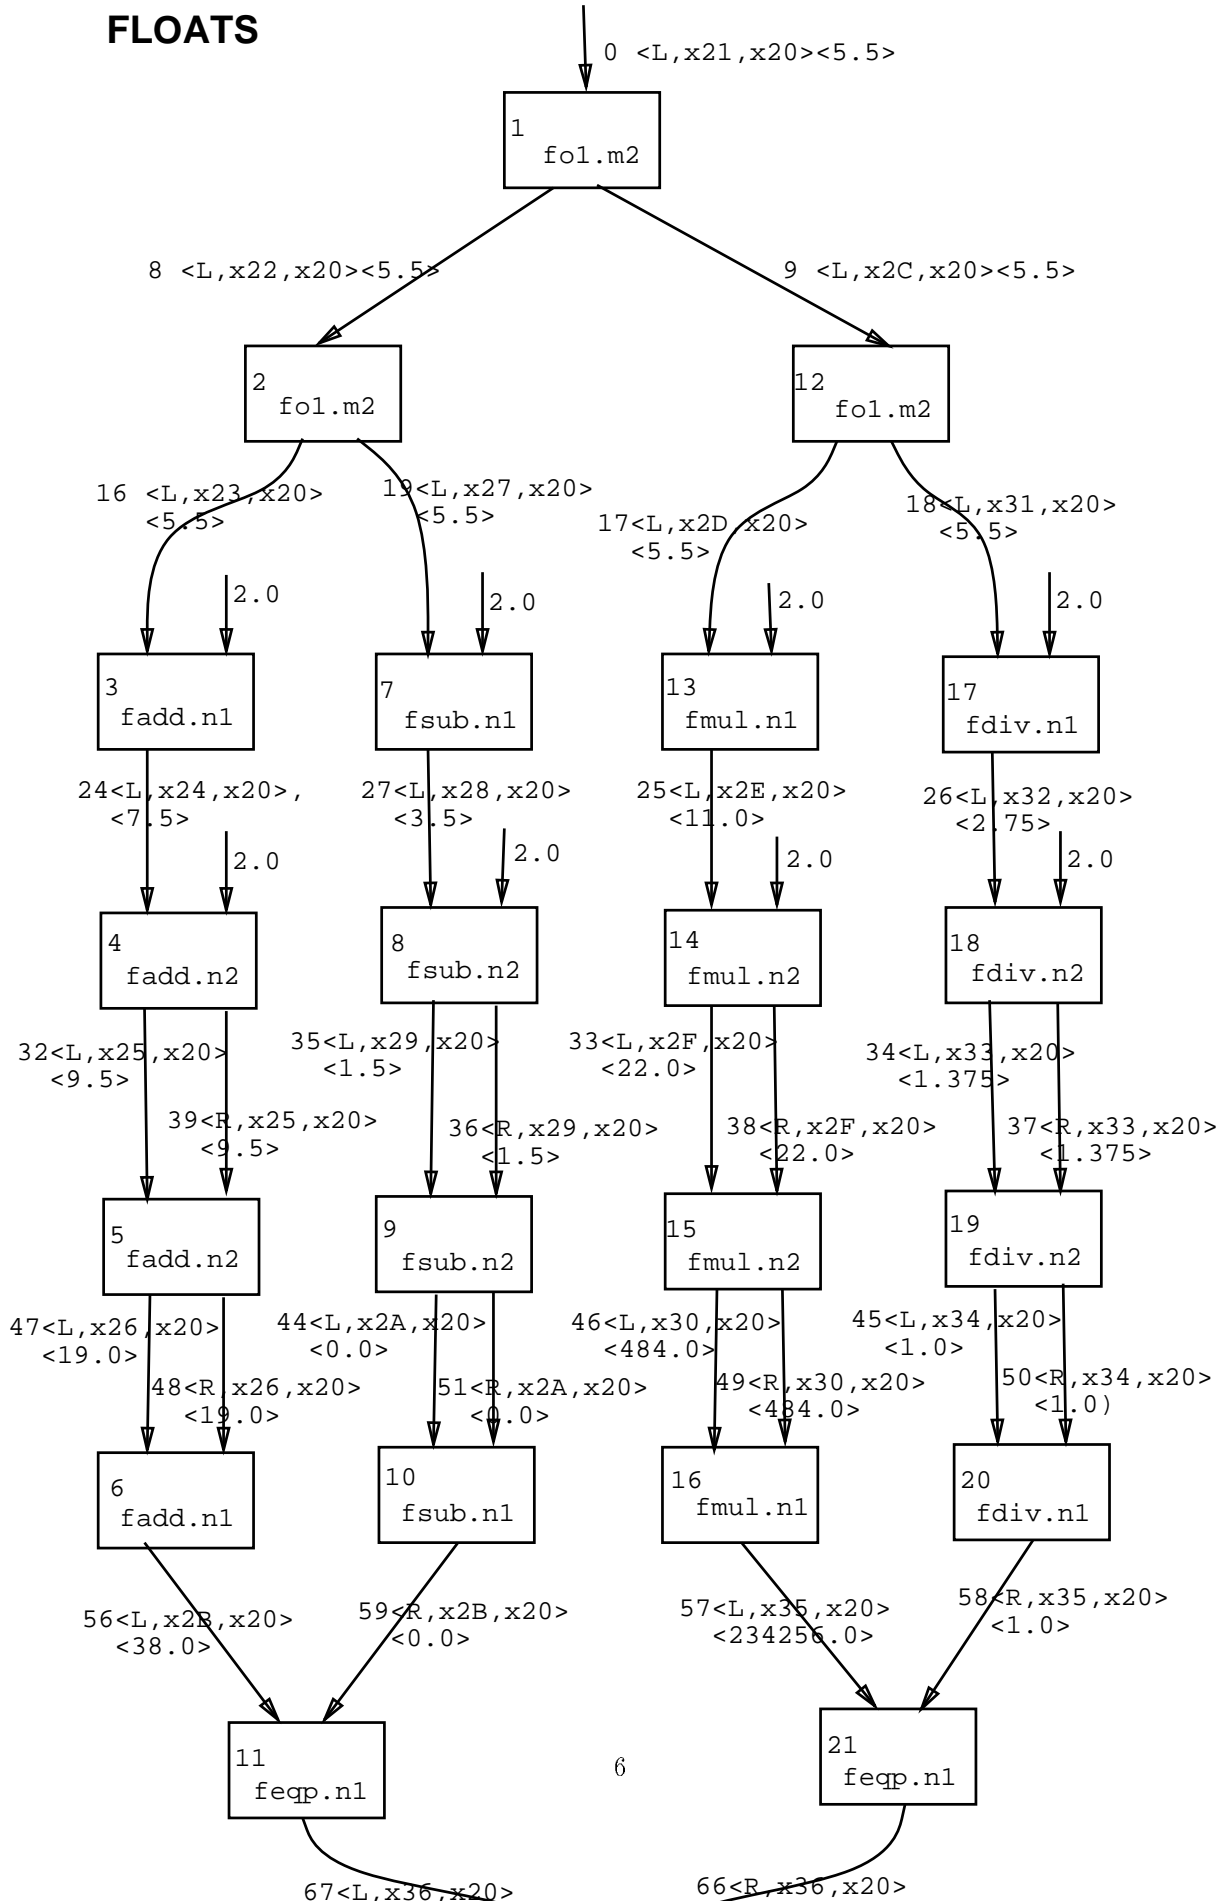
FORK-JOIN

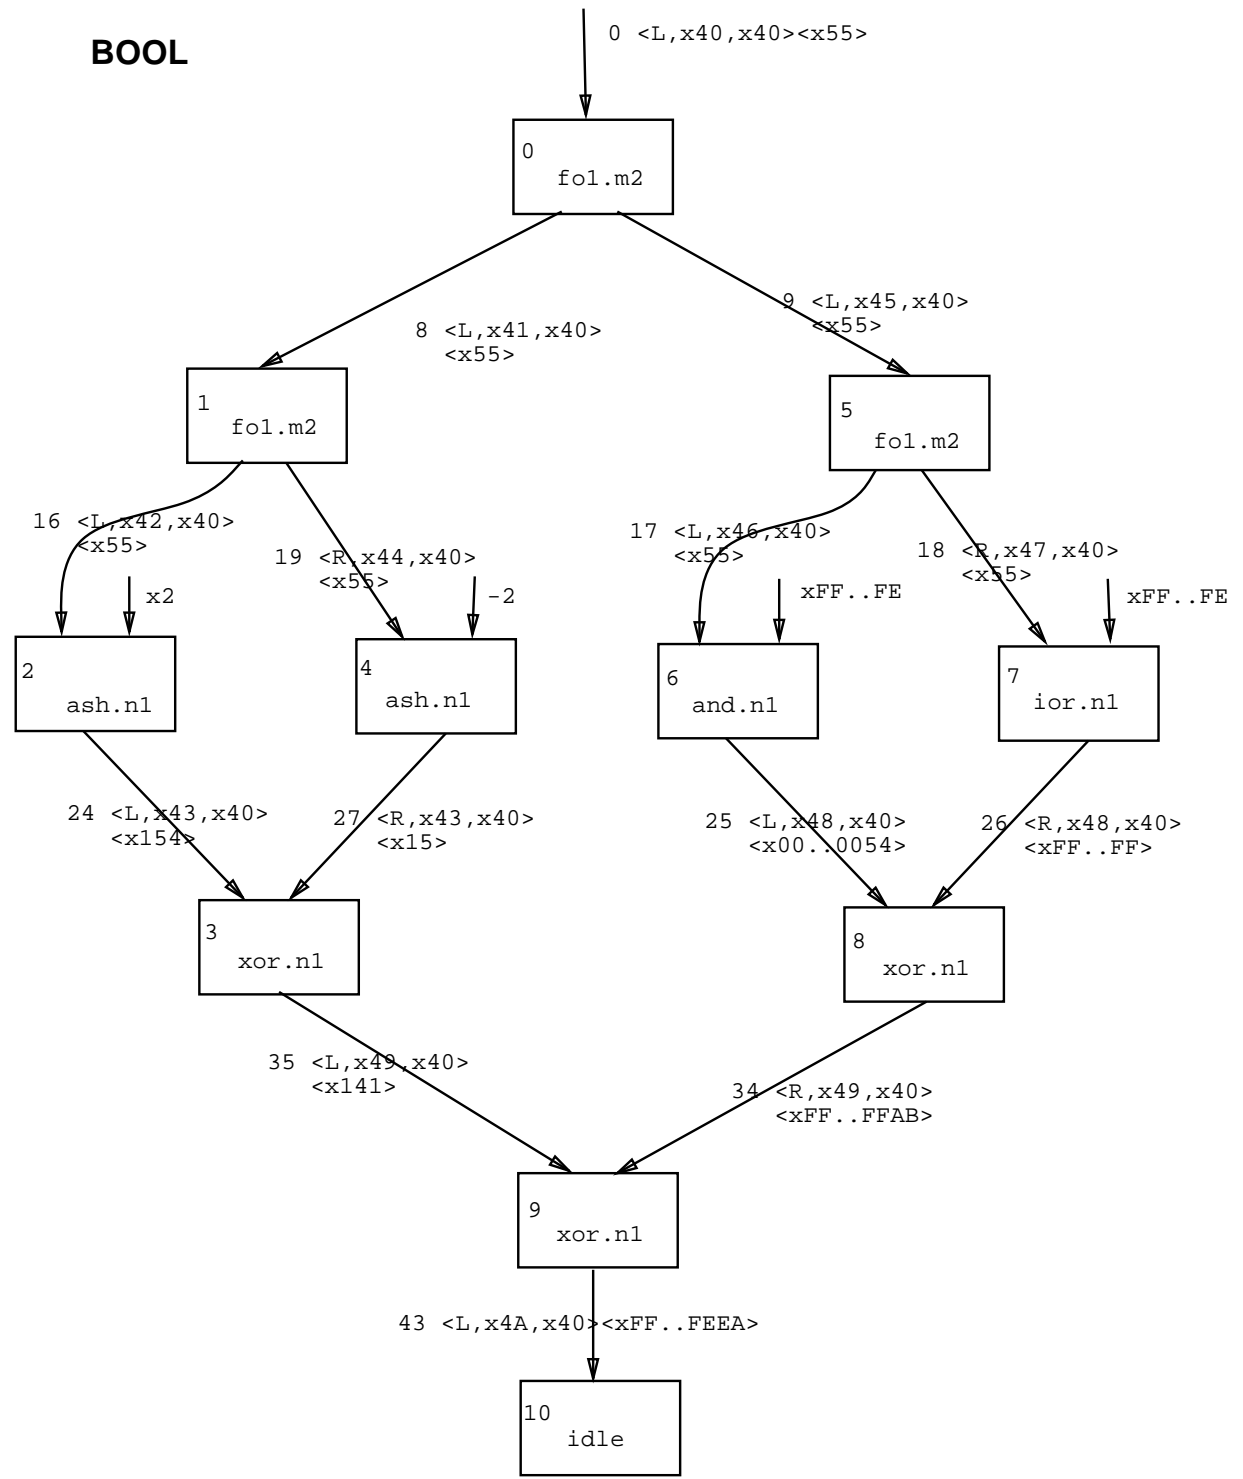


5.5in

Figure 1: Fork-Join

FLOATS





7.0in

Figure 3: Bool

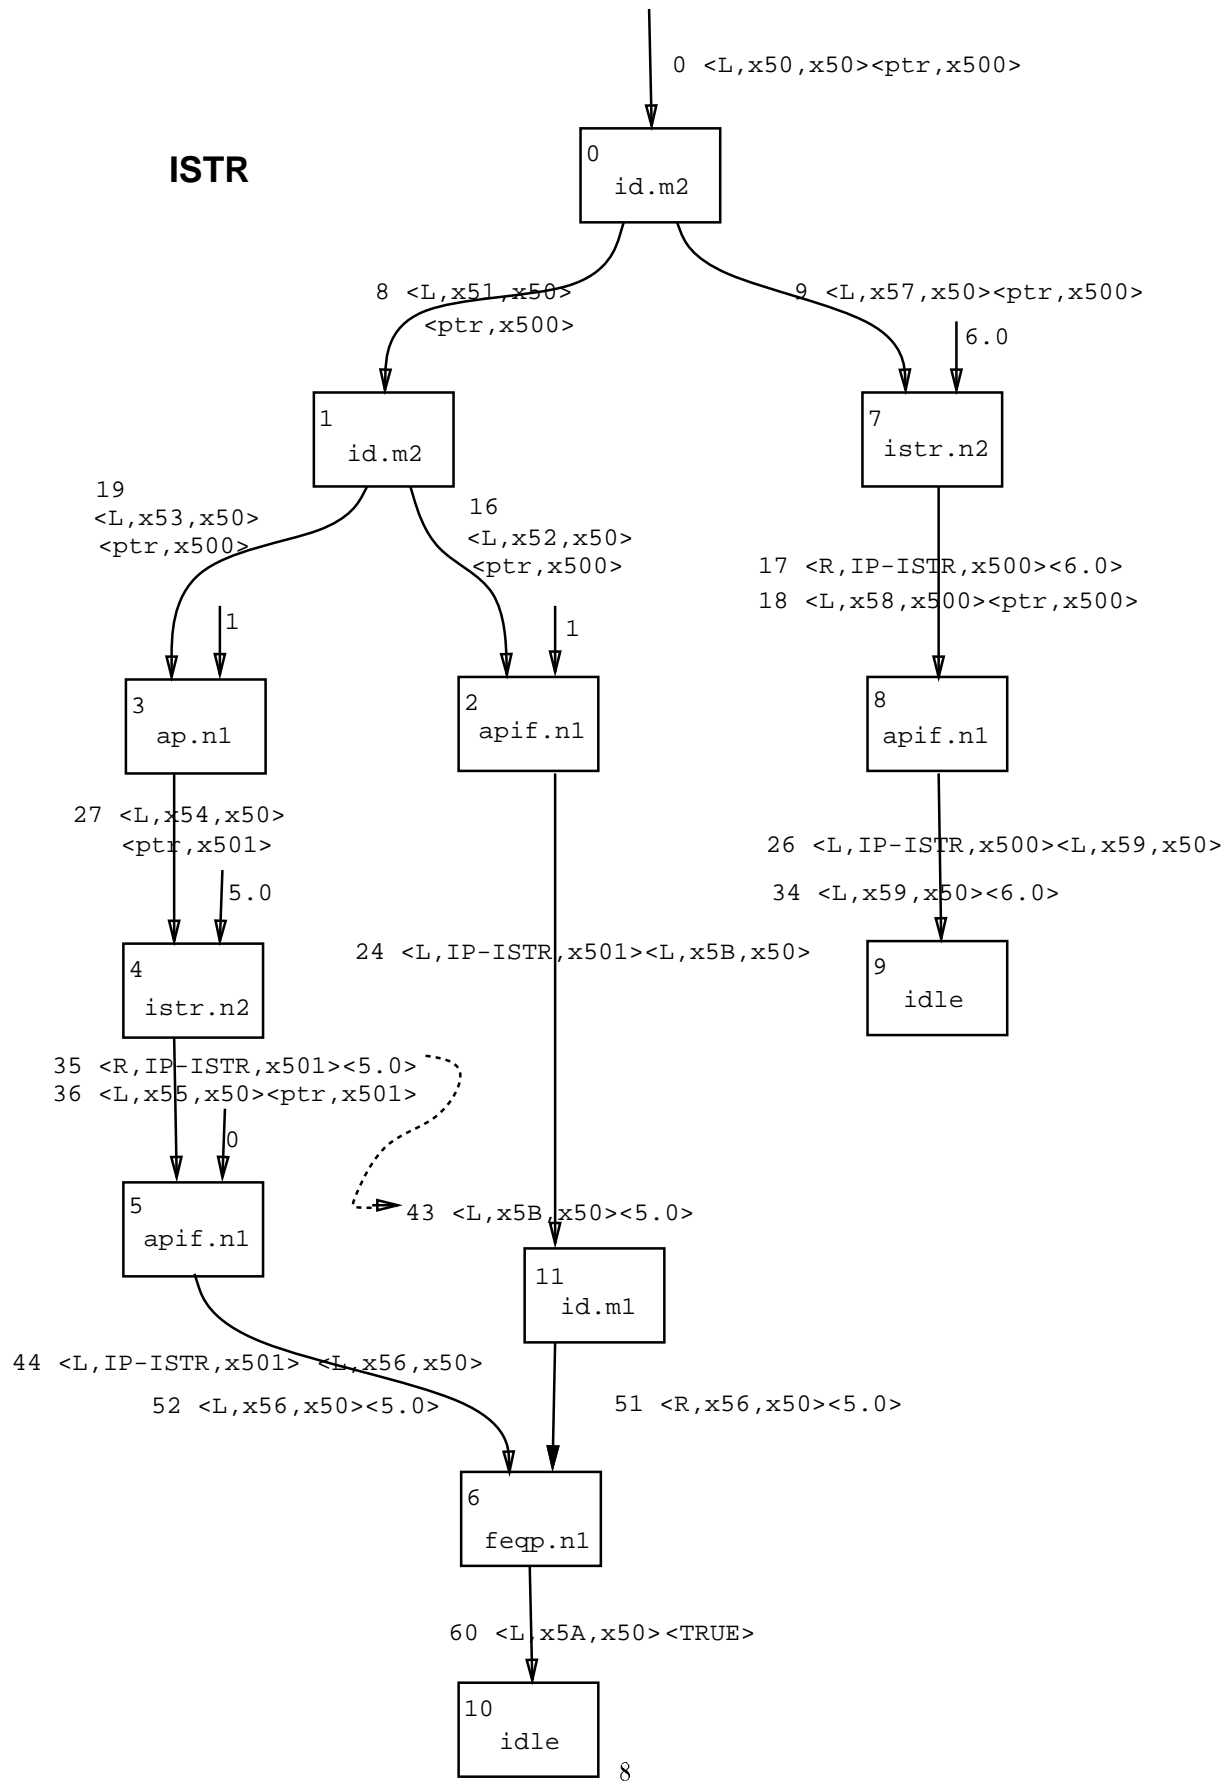
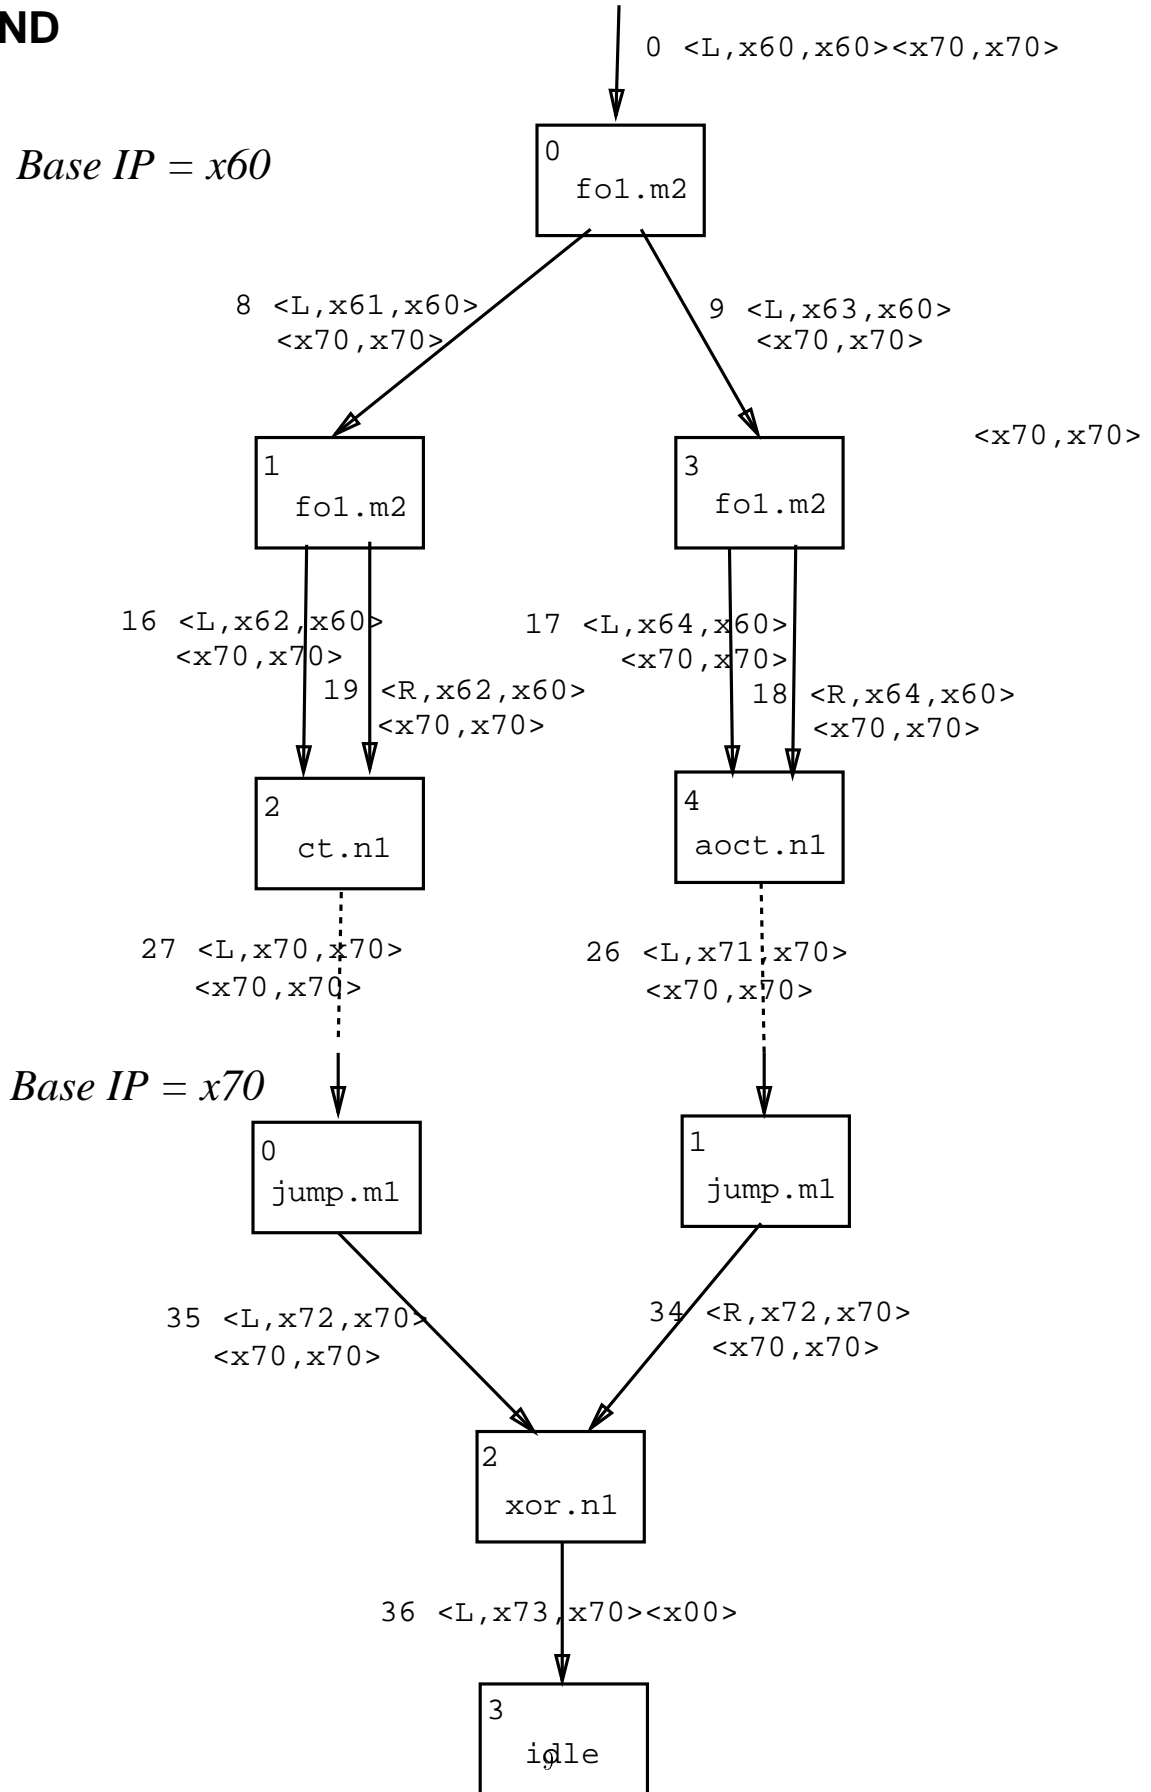


Figure 4: I-structures

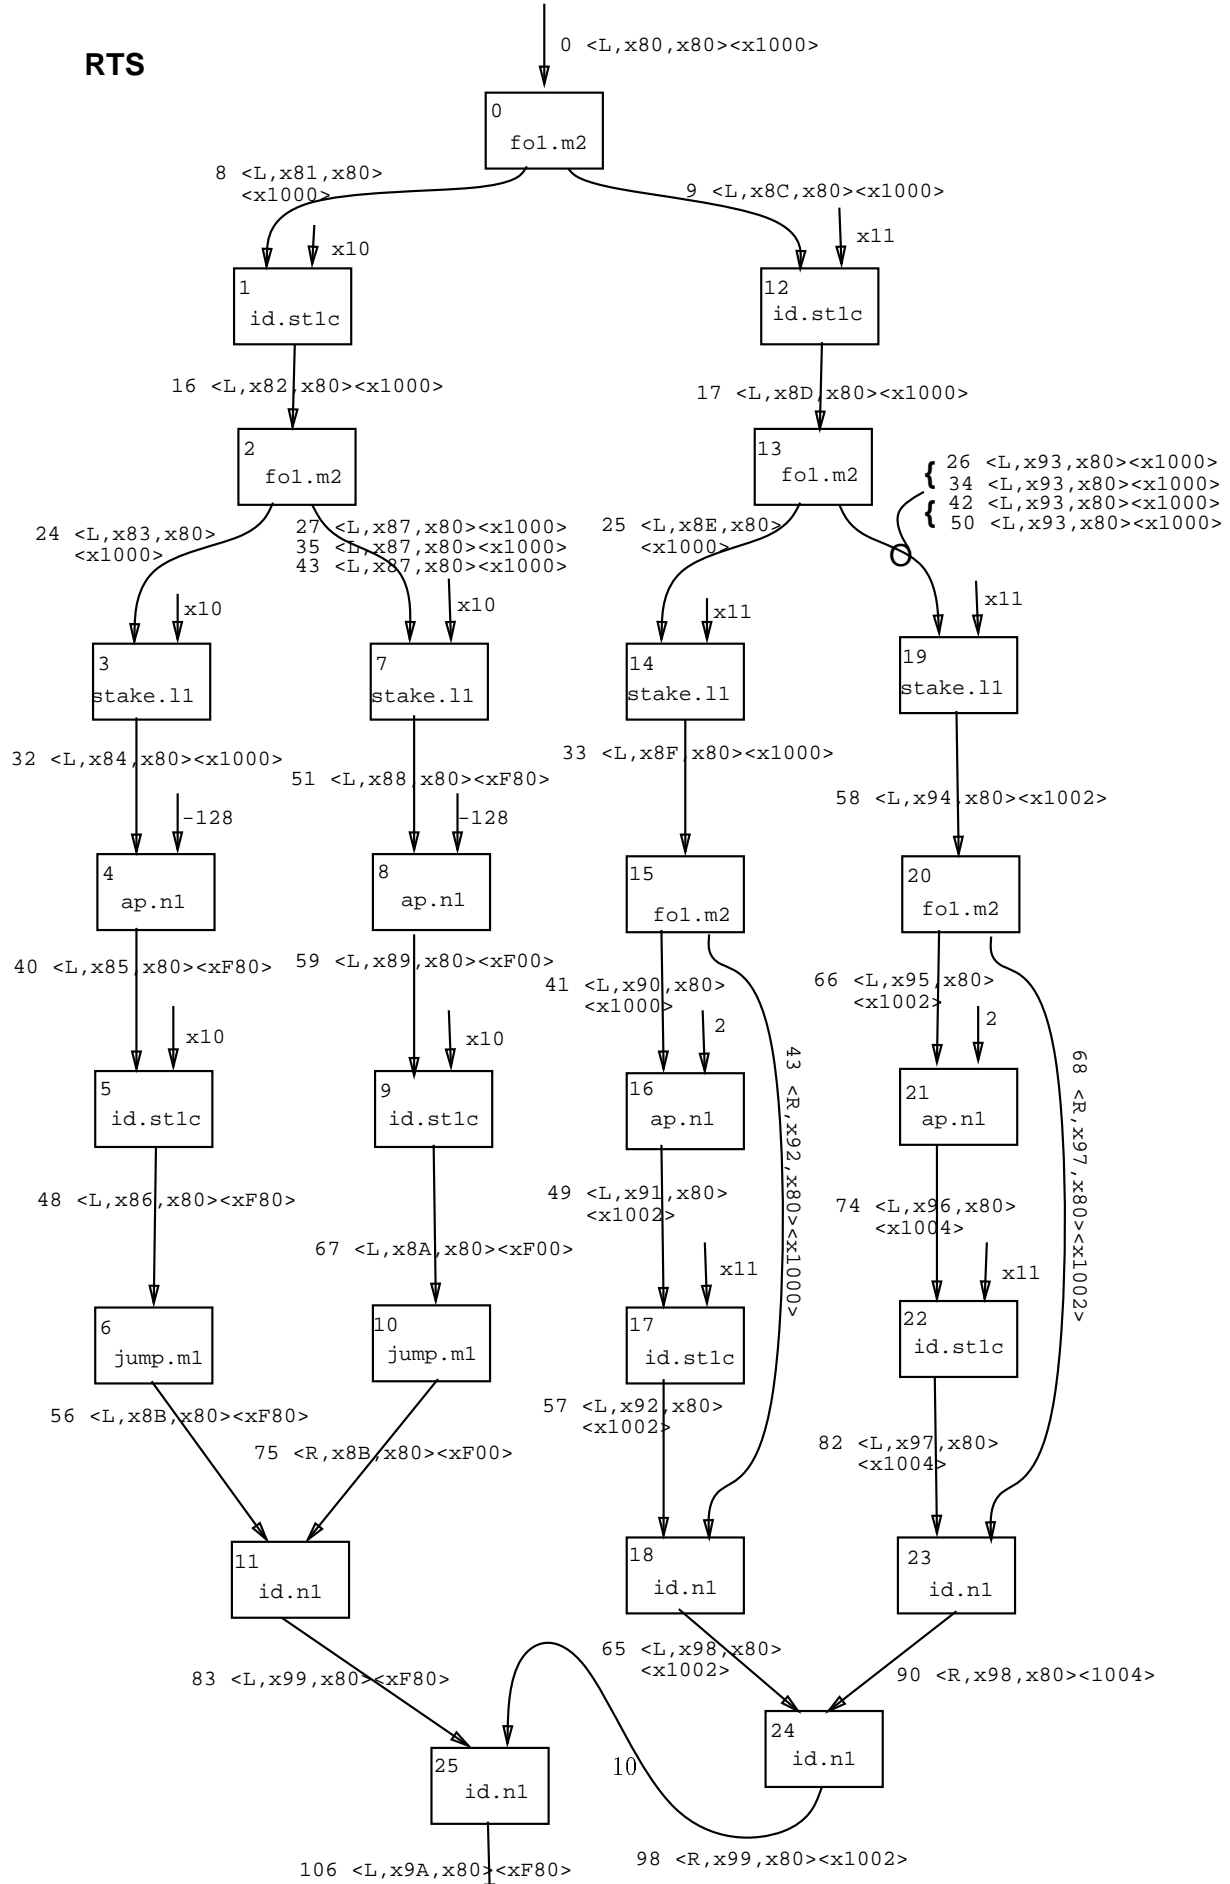
SEND

Base IP = x60

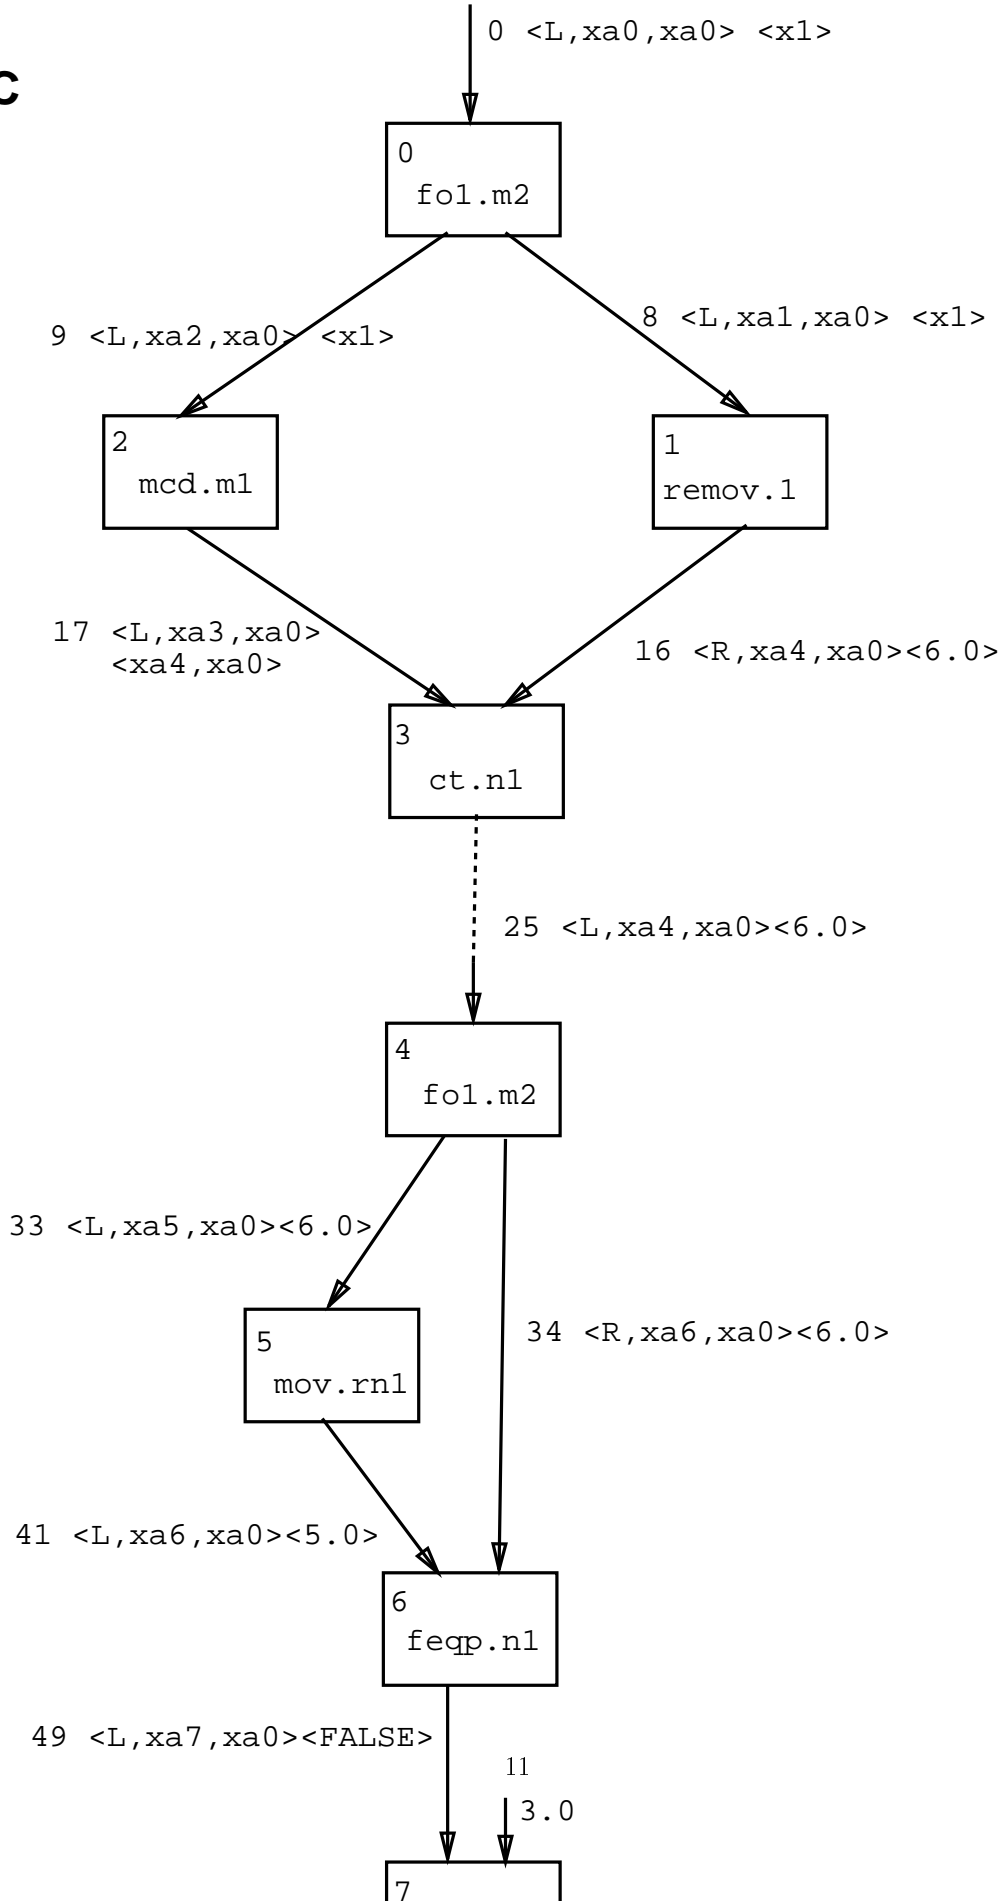


Base IP = x70

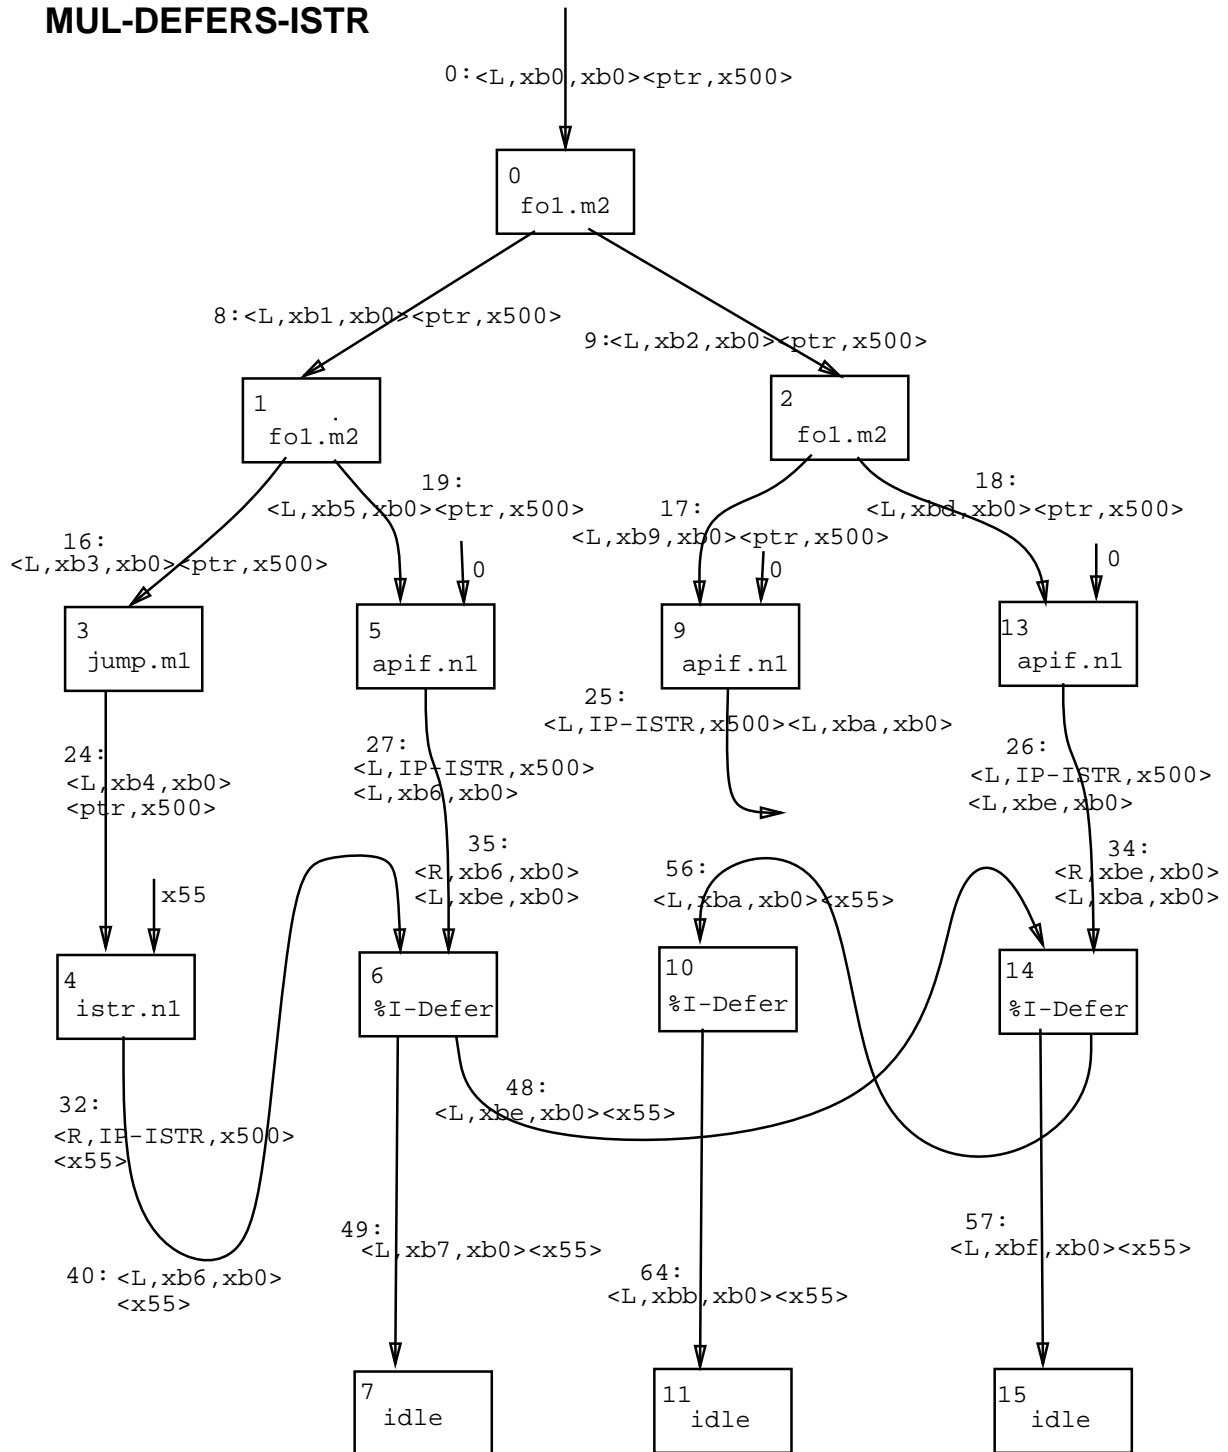
RTS



MISC



MUL-DEFERS-ISTR



9.5in

Figure 8: Multiple Deferring I-Structures

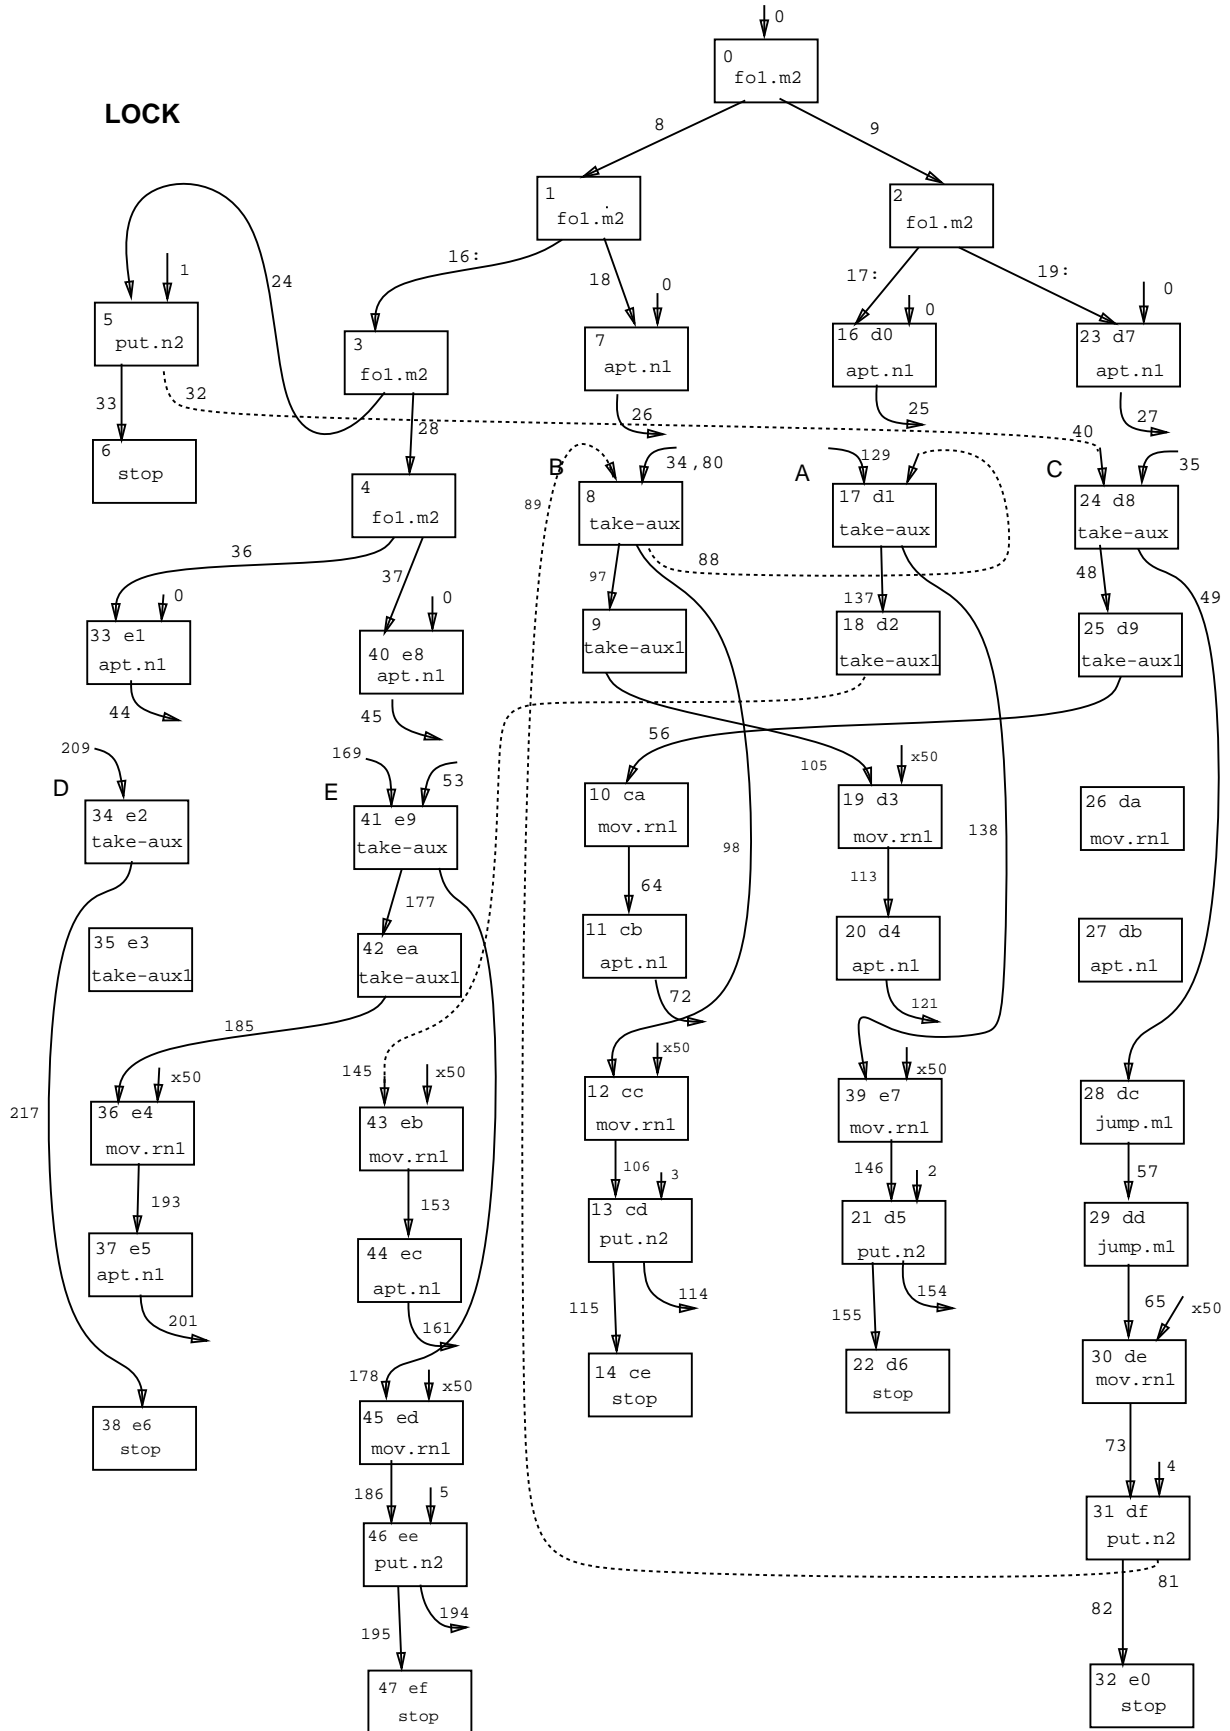




Figure 9: Multiple Deferring Locks

0



1



14

%%plmem

A rectangular box with a black border containing the assembly instruction '%%plmem'. An arrow points down and to the right to the top center of this box from the box above.

