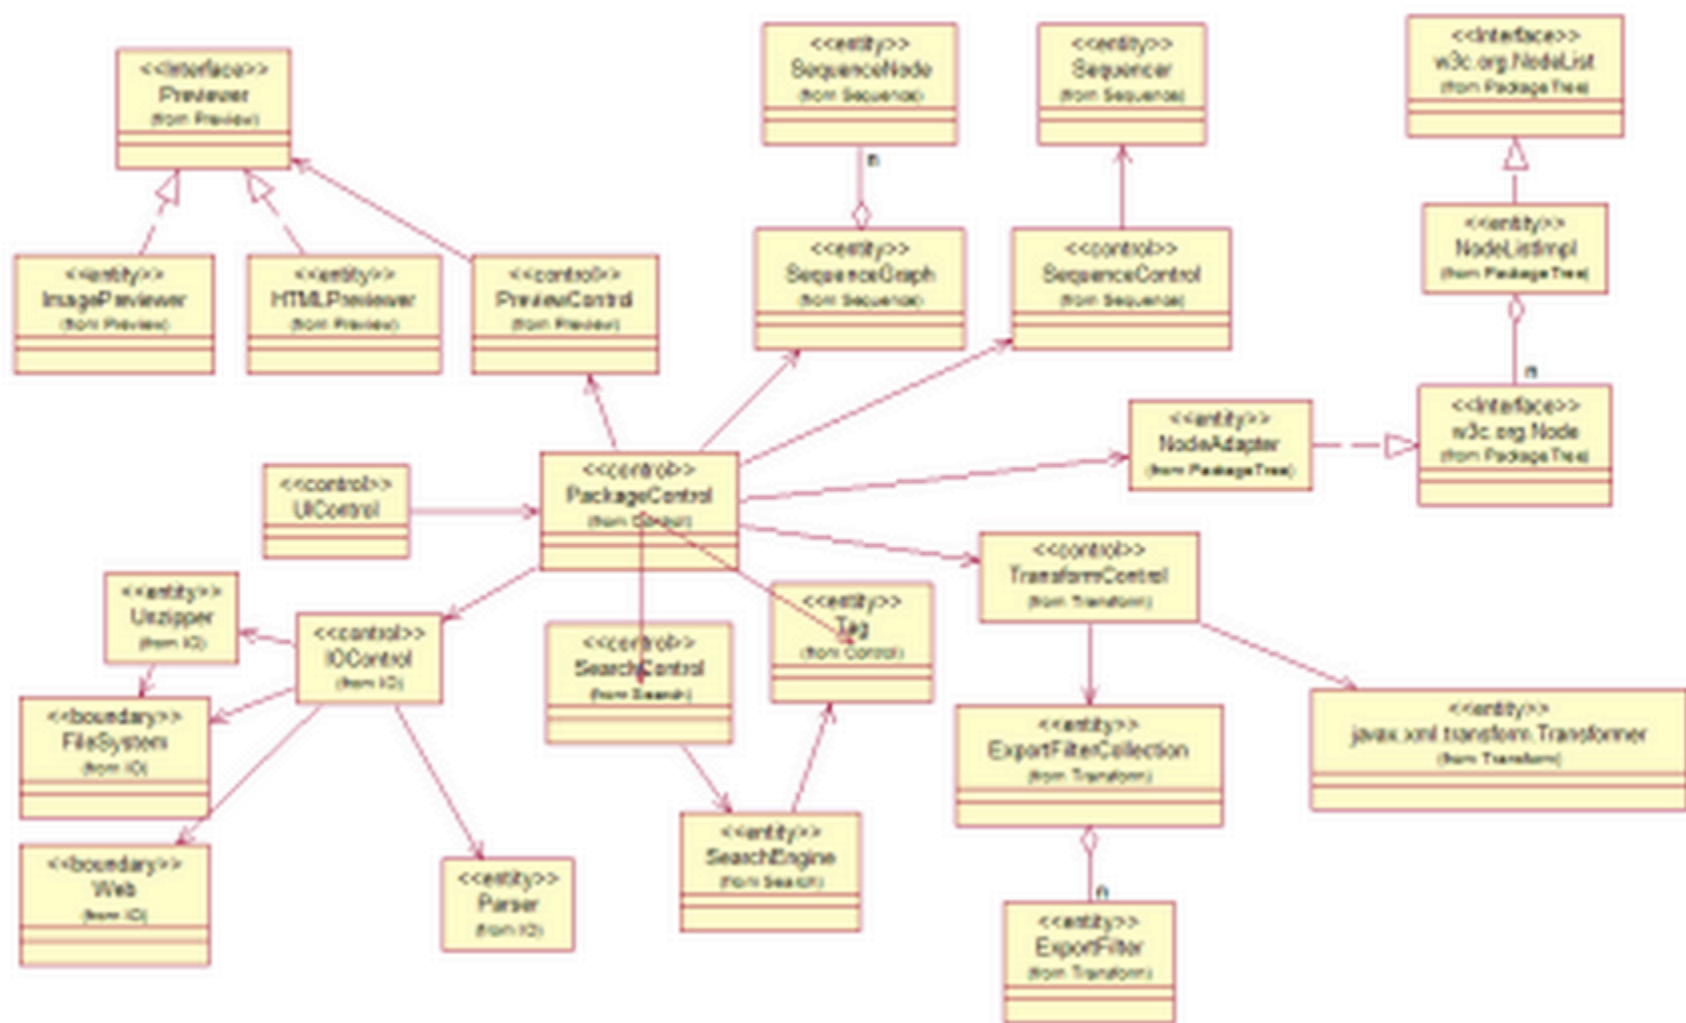
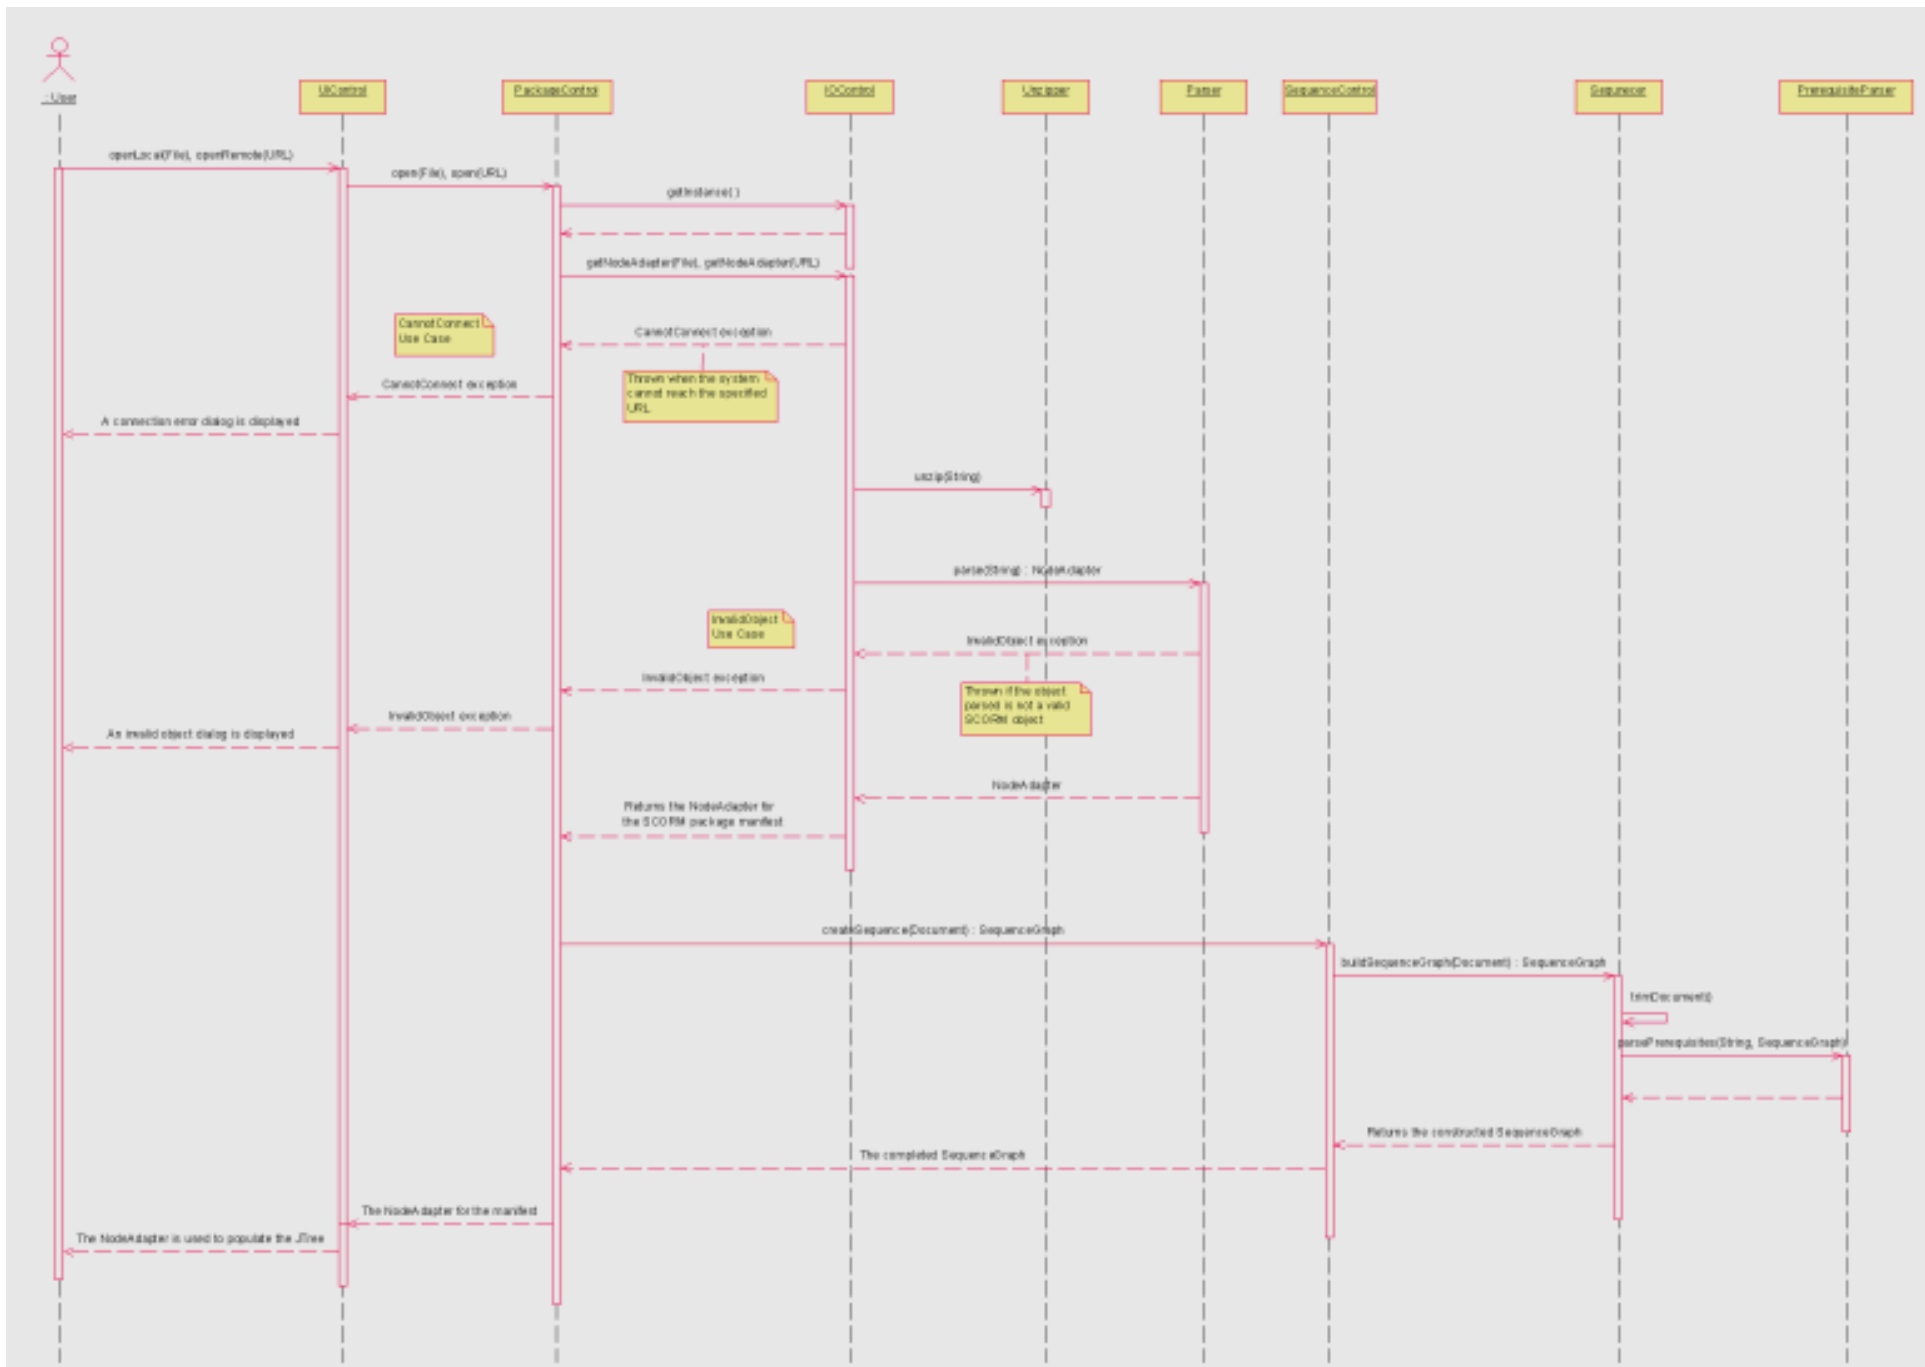


# Software Visualization

Presented by Sam Davis





# More than Just UML!

- UML is about **static structure** of software
- In terms of **abstractions** like
  - Procedures
  - Objects
  - Files
  - Packages
- But...

# Software is Dynamic!

- **Abstractions** are for developers
- Users care about **behaviour**
- Visualize behaviour of software at **run time**
  - Find errors
  - Find performance bottlenecks

What can we visualize?

# Test Results

- Hundreds, maybe thousands of tests
- For each test:
  - Purpose
  - Result (pass or fail)
    - Could be per-configuration or per-version
  - Relevant parts of the code

# Detailed Execution Data

- Could be for **many** executions
- Dynamic events as opposed to summary data



# Summary Data: Examples

- Total running time
- Number of times a method was called
- Amount of time CPU was idle

# Dynamic Events: Examples

- Memory allocation
- System calls
- Cache misses
- Page faults
- Pipeline flushes
- Process scheduling
- Completion of disk reads or writes
- Message receipt
- Application phases

# Really Detailed Execution Data

- Logging virtual machines can capture *everything*
  - Enough data to *replay* program execution and recreate the entire machine state at any point in time
  - Allows “time-traveling”
  - For long running systems, data could span months
- Uses:
  - Debugging
  - Understanding attacks

# Strata\_Various: Multi\_Layer Visualization of Dynamics in Software System Behavior

Doug Kimelman, Bryan Rosenburg, Tova  
Roth

Proc. Fifth IEEE Conf. Visualization '94, IEEE  
Computer Society Press, Los Alamitos, Calif.,  
1994, pp. 172–178.

# Strata\_Various

- Trace-driven program visualization
- Trace: sequence of <time, event> pairs
- Events captured from all layers:
  - Hardware
  - Operating System
  - Application
- Replay execution history
- Coordinate navigation of event views

# Strata\_Various: Main Argument

- Debugging and tuning requires simultaneously analyzing behaviour at **multiple layers** of the system

Program Visualization 1.0 (C) IBM

Commands Configuration Suggest Help

Time 13.007041

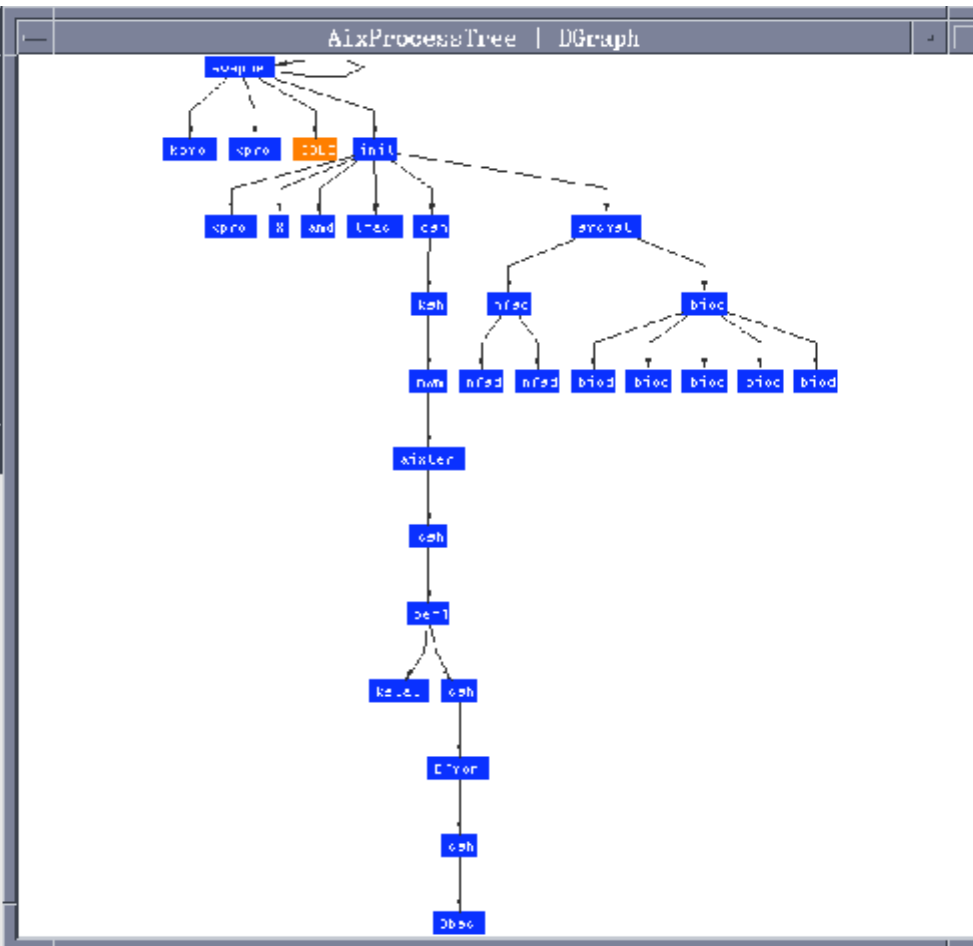
Speed Factor 1.00000

Log Wait

Trace File Position




Trace File D.pv

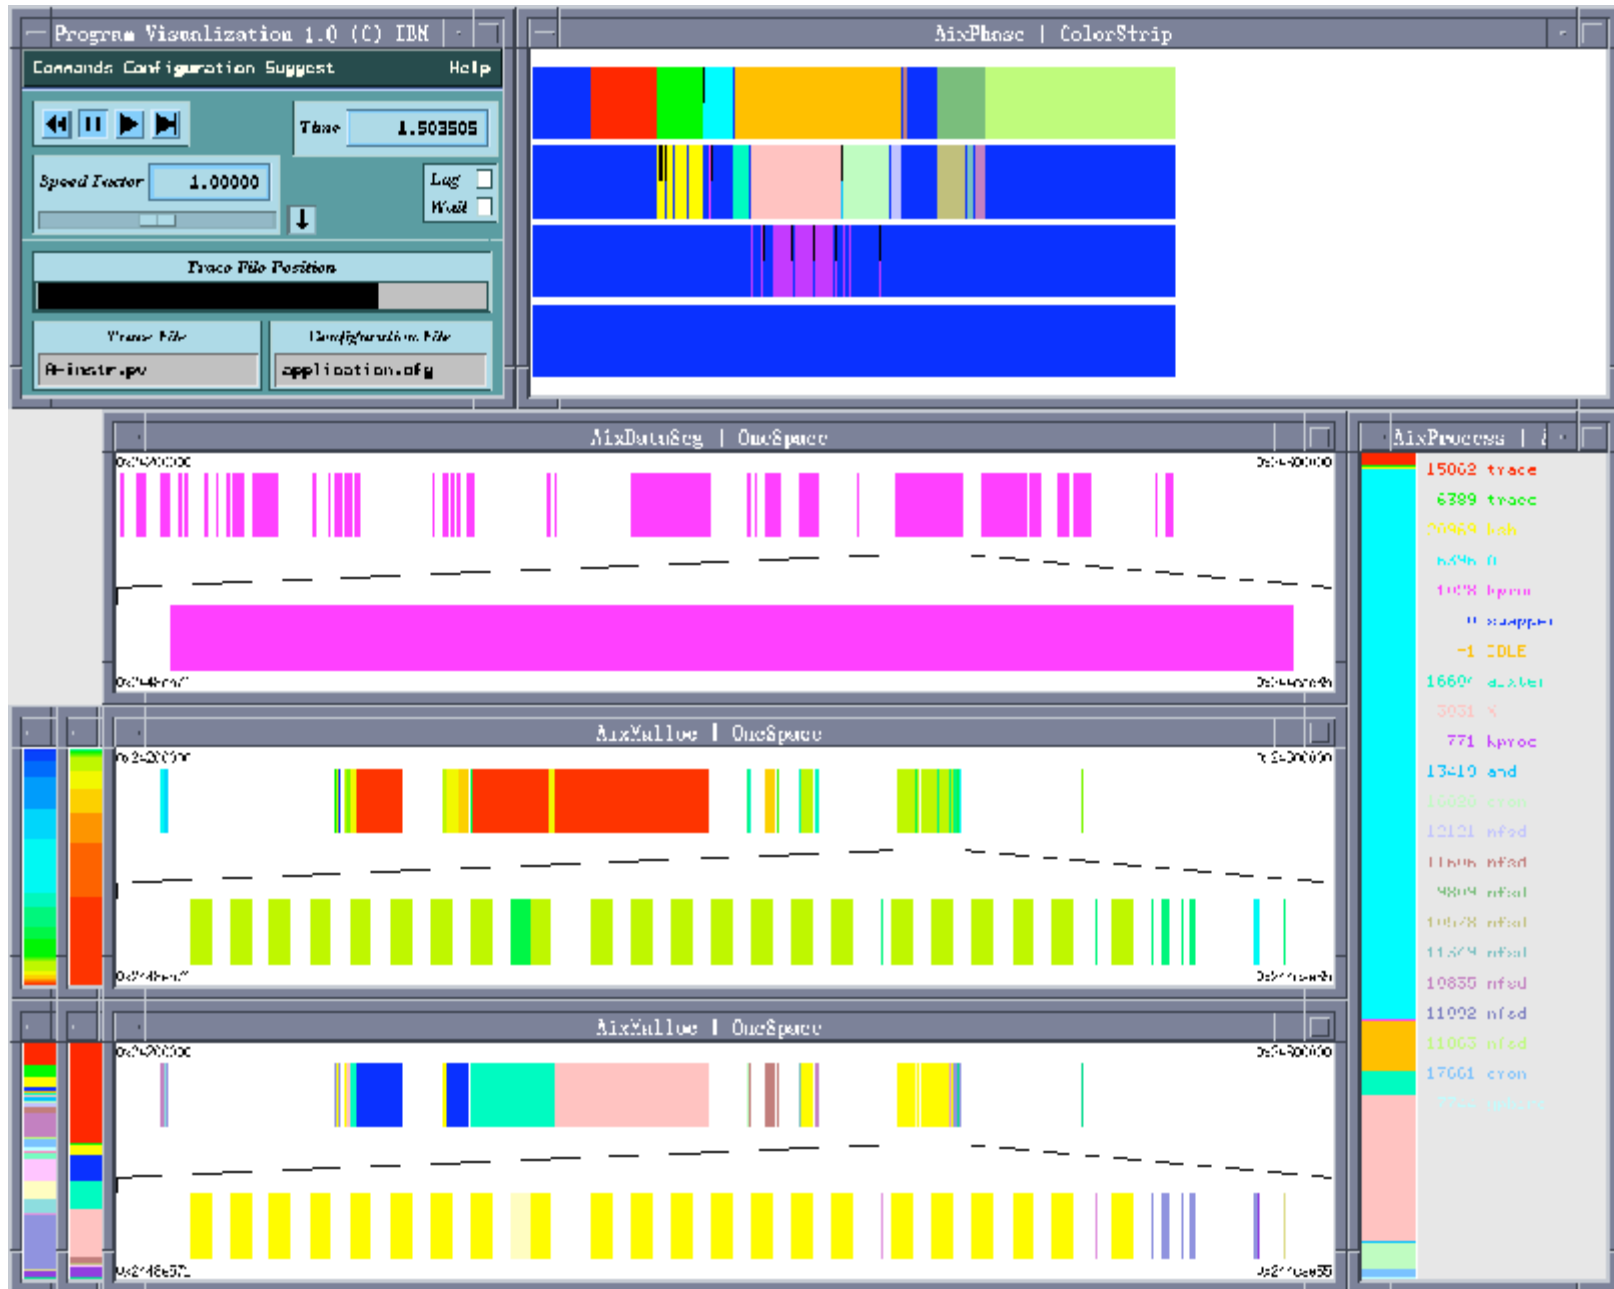
Configuration File system.cfg



AixProcess |

1441:	trace
1407:	trace
13313:	ef
14323:	pc-1
13371:	pc-ee
-1:	ILLE
3513:	r-fvc
7433:	Llvc
7720:	flvc
13474:	csh
4134:	erc
13333:	ksccc
3713:	bioc
771:	kp-cc
13323:	unarc
13330:	unarc
13331:	unarc
1416:	unarc
0:	skapper
13332:	unarc
13334:	unarc
13335:	unarc
12330:	slxcem
11332:	x









# Strata\_Various: Critique

- Examples demonstrate usefulness
- Fundamentally, a good idea
  - Increasing importance as multi-core machines become standard
- Many windows
  - Titles not meaningful
  - Virtual reality cop-out
- Dubious claim that tracing does not alter behaviour

# SeeSoft

- Zoomed out view of source code
  - Lines of code displayed as thin horizontal lines
  - Preserve indentation, length
  - Can colour lines according to data
- Link with readable view of code
- Allows tying data to source code

Stephen G. Eick, Joseph L. Steffen and Eric E. Sumner, Jr. "SeeSoft – A Tool for Visualizing Line-Oriented Software Statistics." *IEEE Transactions on Software Engineering*, 18(11):957-968, November 1992.

# SeeSoft Example



# Visually Encoding Program Test Information to Find Faults in Software (Tarantula)

James Eagan, Mary Jen Harrold, James A.  
Jones, and John Stasko, Proc. InfoVis 2001  
pp. 33-36.

# Tarantula

- Extends SeeSoft idea
- Defines colour mapping for LOC based on test results
- Goal: use test results to identify broken code

# Tarantula

- Input:
  - For each test:
    - Test number
    - Result (pass or fail)
    - Test coverage (list of line numbers)

# Tarantula: Discrete Colour Mapping

- Based on user tests
- Black background
- Colour each line
  - Red if executed by failed tests
  - Green if executed by passed tests
  - Yellow if executed by both



# Tarantula: Continuous Colour Mapping

- Extend discrete colour mapping by
  - Interpolating between red and green
  - Adjusting brightness according to number of tests
- Possibilities:
  - Number of passed or failed tests
  - Ratio of passed to failed tests
  - Ratio of % passed to % failed

# Tarantula: Continuous Colour Mapping

- For each line  $L$ 
  - Let  $p$  and  $f$  be the percentages of passed and failed tests that executed  $L$
  - If  $p = f = 0$ , colour  $L$  grey
  - Else, colour  $L$  according to
    - Hue:  $p / ( p + f )$ , where 0 is red and 1 is green
    - Brightness:  $\max( p, f )$

Tarantula Bug Finder

File

Default
  Discrete
  Continuous
  Passes
  Fails
  Mixed

Line: 6862

Test:

```


if (error != 0)
    *pqdim_unit_ptr = 0;

```

Line 6862

Executions: 66 / 300  
 Passed: 63 / 297  
 Failed: 3 / 3

Color Legend



# Tarantula: Critique

- Visualizing test results could be useful, this is a first step
- Future work: does colouring help to find broken code?
- Colouring: simple idea made complex
- Tests identified only by number
  - Better: name tests
  - Better still: can we visualize the *meaning* of tests?

# Visualization of Program- Execution Data for Deployed Software (Gammatella)

Alessandro Orso, James Jones, and Mary  
Jean Harrold.

Proc. of the ACM Symp. on Software  
Visualization, San Diego, CA, June 2003,  
pages 67--76.

# Gammatella

- Collection and storage of program-execution data
- Visualization of data about **many executions**

# Gammatella: Executions

- Code coverage and profiling data
- Execution properties
  - OS
  - Java version
  - Etc.
- Filters
  - Boolean predicate logic
- Summarizers



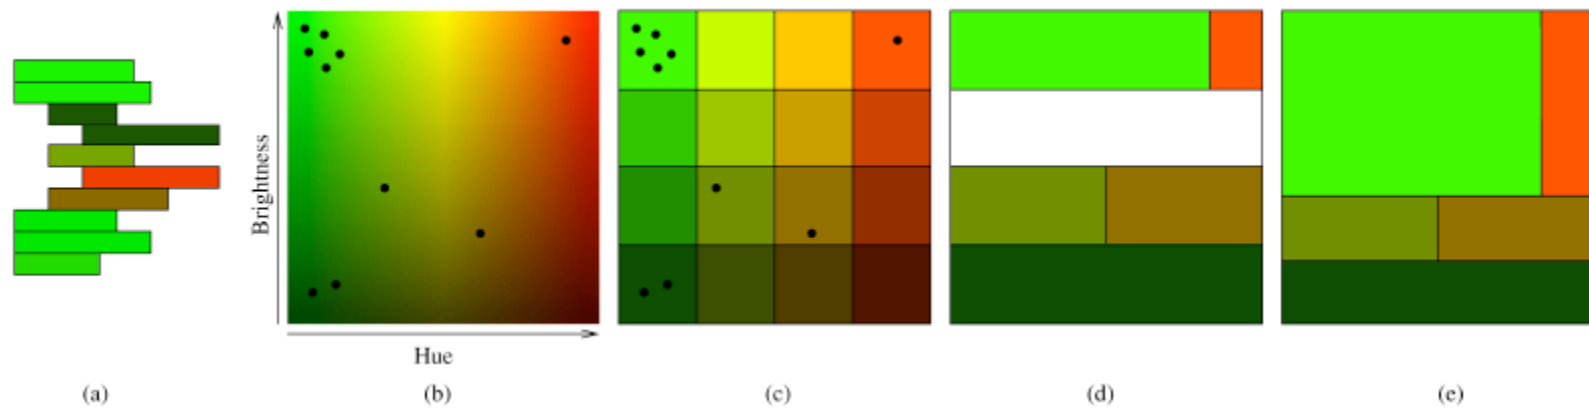
# Gammatella: Coloured, Tri-Level Representation

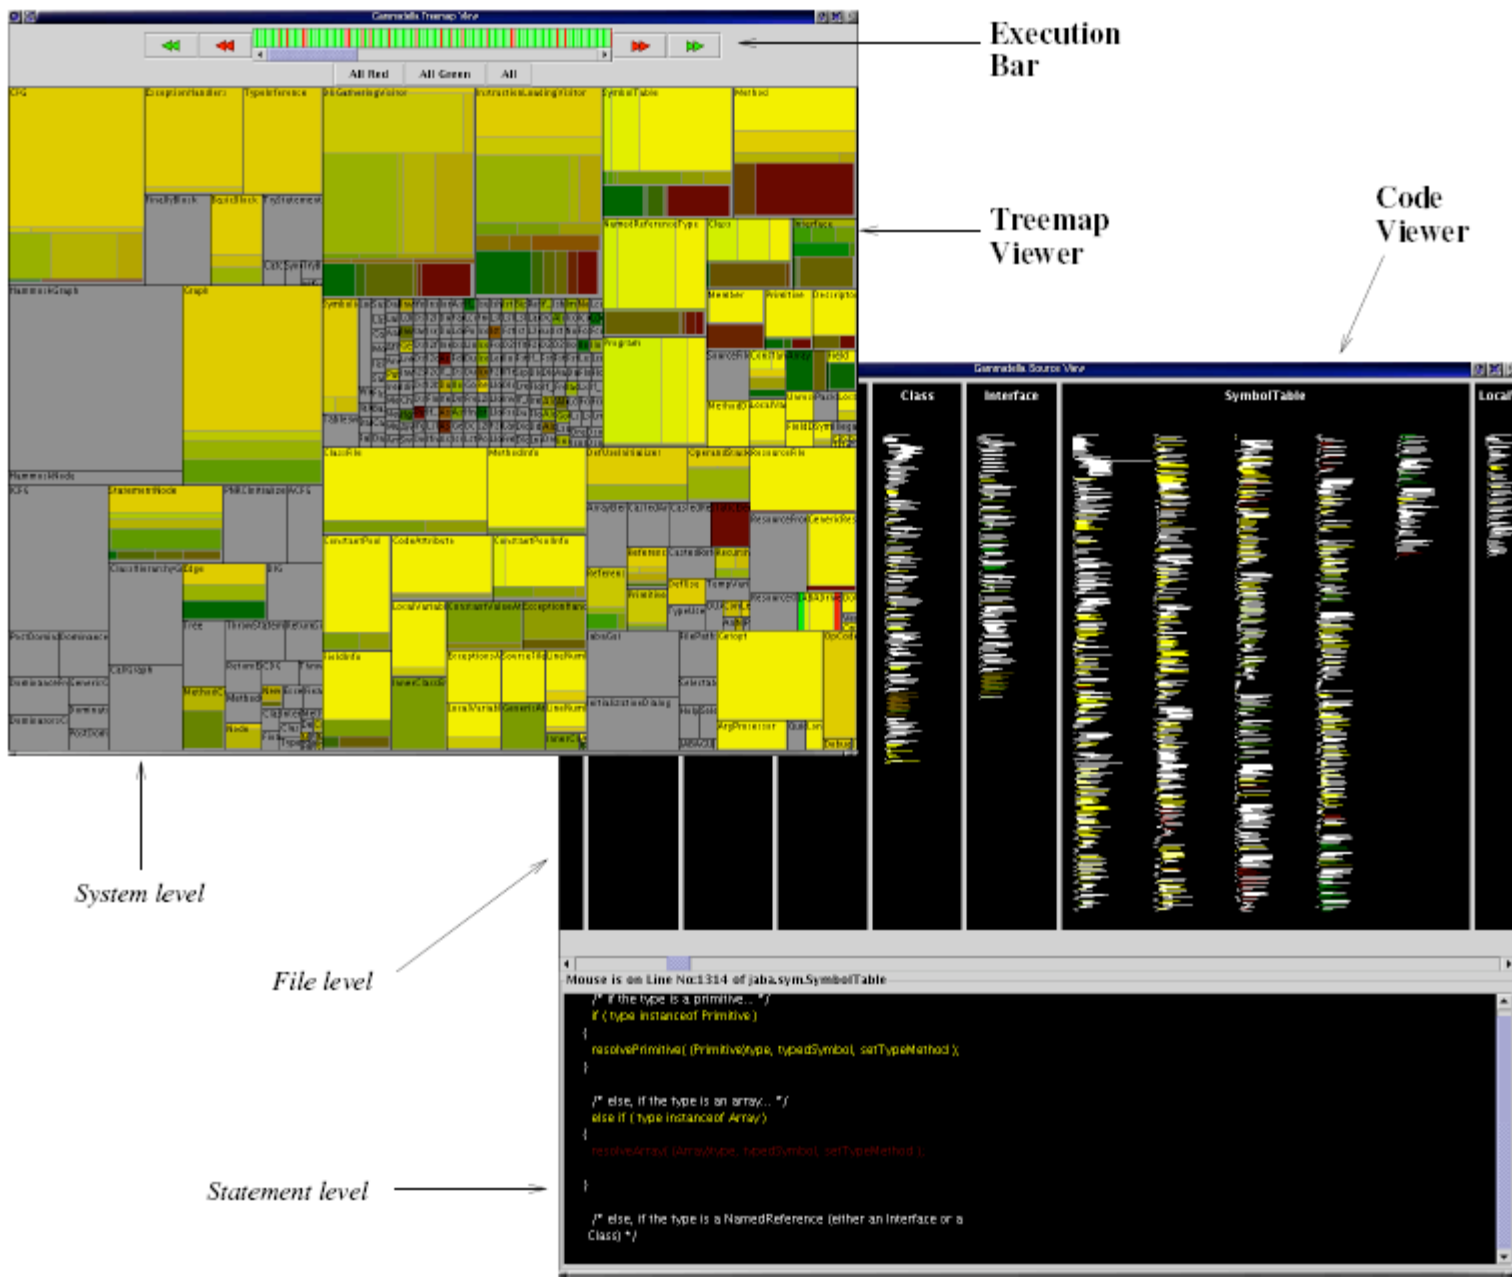
- System level
  - Treemap of package/class hierarchy
- File level:
  - SeeSoft-like view of code
- Statement level:
  - Source code (coloured text)
- Colours based on exceptions
  - Other colourings possible, e.g. profiling data



# One Level Treemap

- Layout algorithm for treemap of depth 1
  - Preserves relative placement of colours





# Gammatella: Critique

- Complete system – not just a visualization
- Effectively links code to structure
- Trial usage discovered useful but high-level information
  - Mainly relied on system view
  - Would be nice to see examples using file and statement level views

# Visualizing Application Behavior on Superscalar Processors

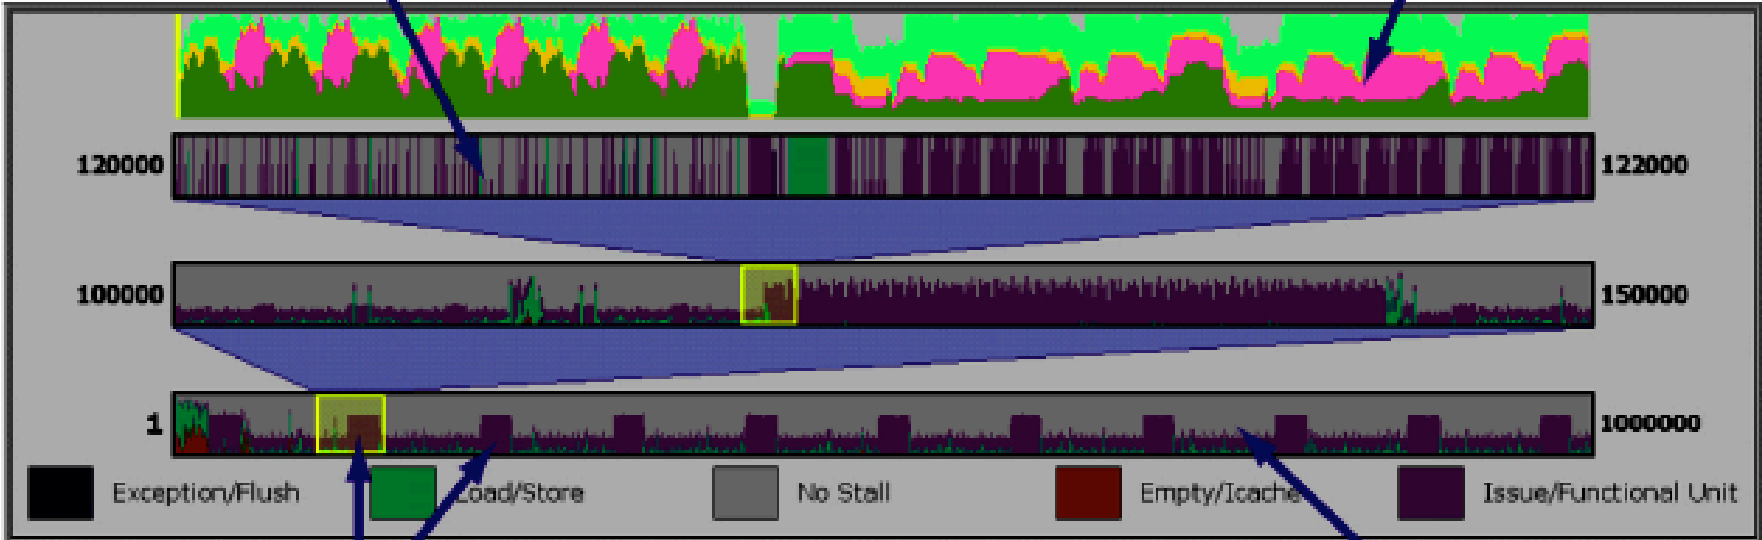
Chris Stolte, Robert Bosch, Pat  
Hanrahan, and Mendel Rosenblum  
Proc. InfoVis 1999

# Superscalar Processors: Quick Overview

- Pipeline
- Multiple Functional Units
  - Instruction-Level Parallelism (ILP)
- Instruction Reordering
- Branch Prediction and Speculation
- Reorder Buffer
  - Instructions wait to graduate (exit pipeline)

③ We are able to focus the area of interest to 2000 cycles -- few enough cycles that we can use animation for further investigation.

④ The instruction mix chart lets us see what types of instructions are in the pipeline during the time interval of interest.



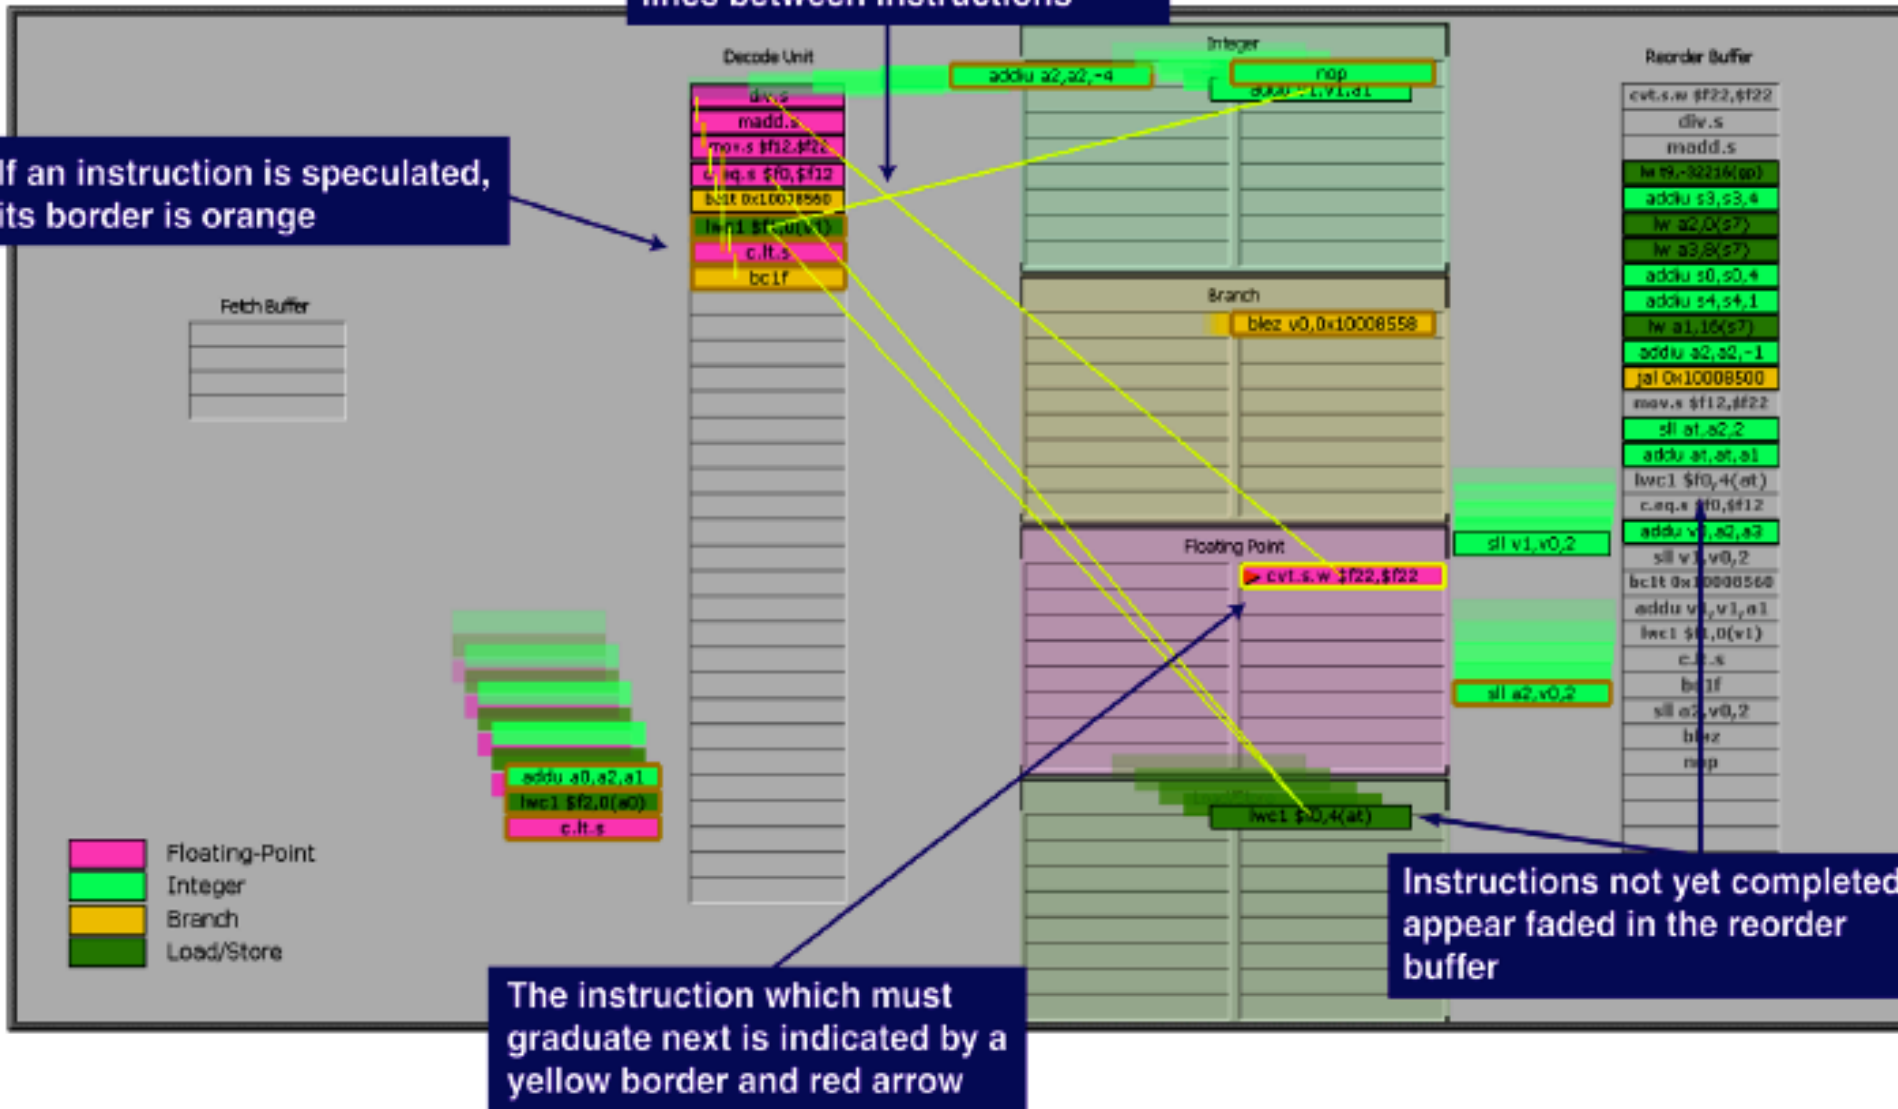
② There are periods of increased pipeline stall throughout the execution

① The overview displays stall and throughput information for the entire execution.



Dependencies appear as yellow lines between instructions

If an instruction is speculated, its border is orange



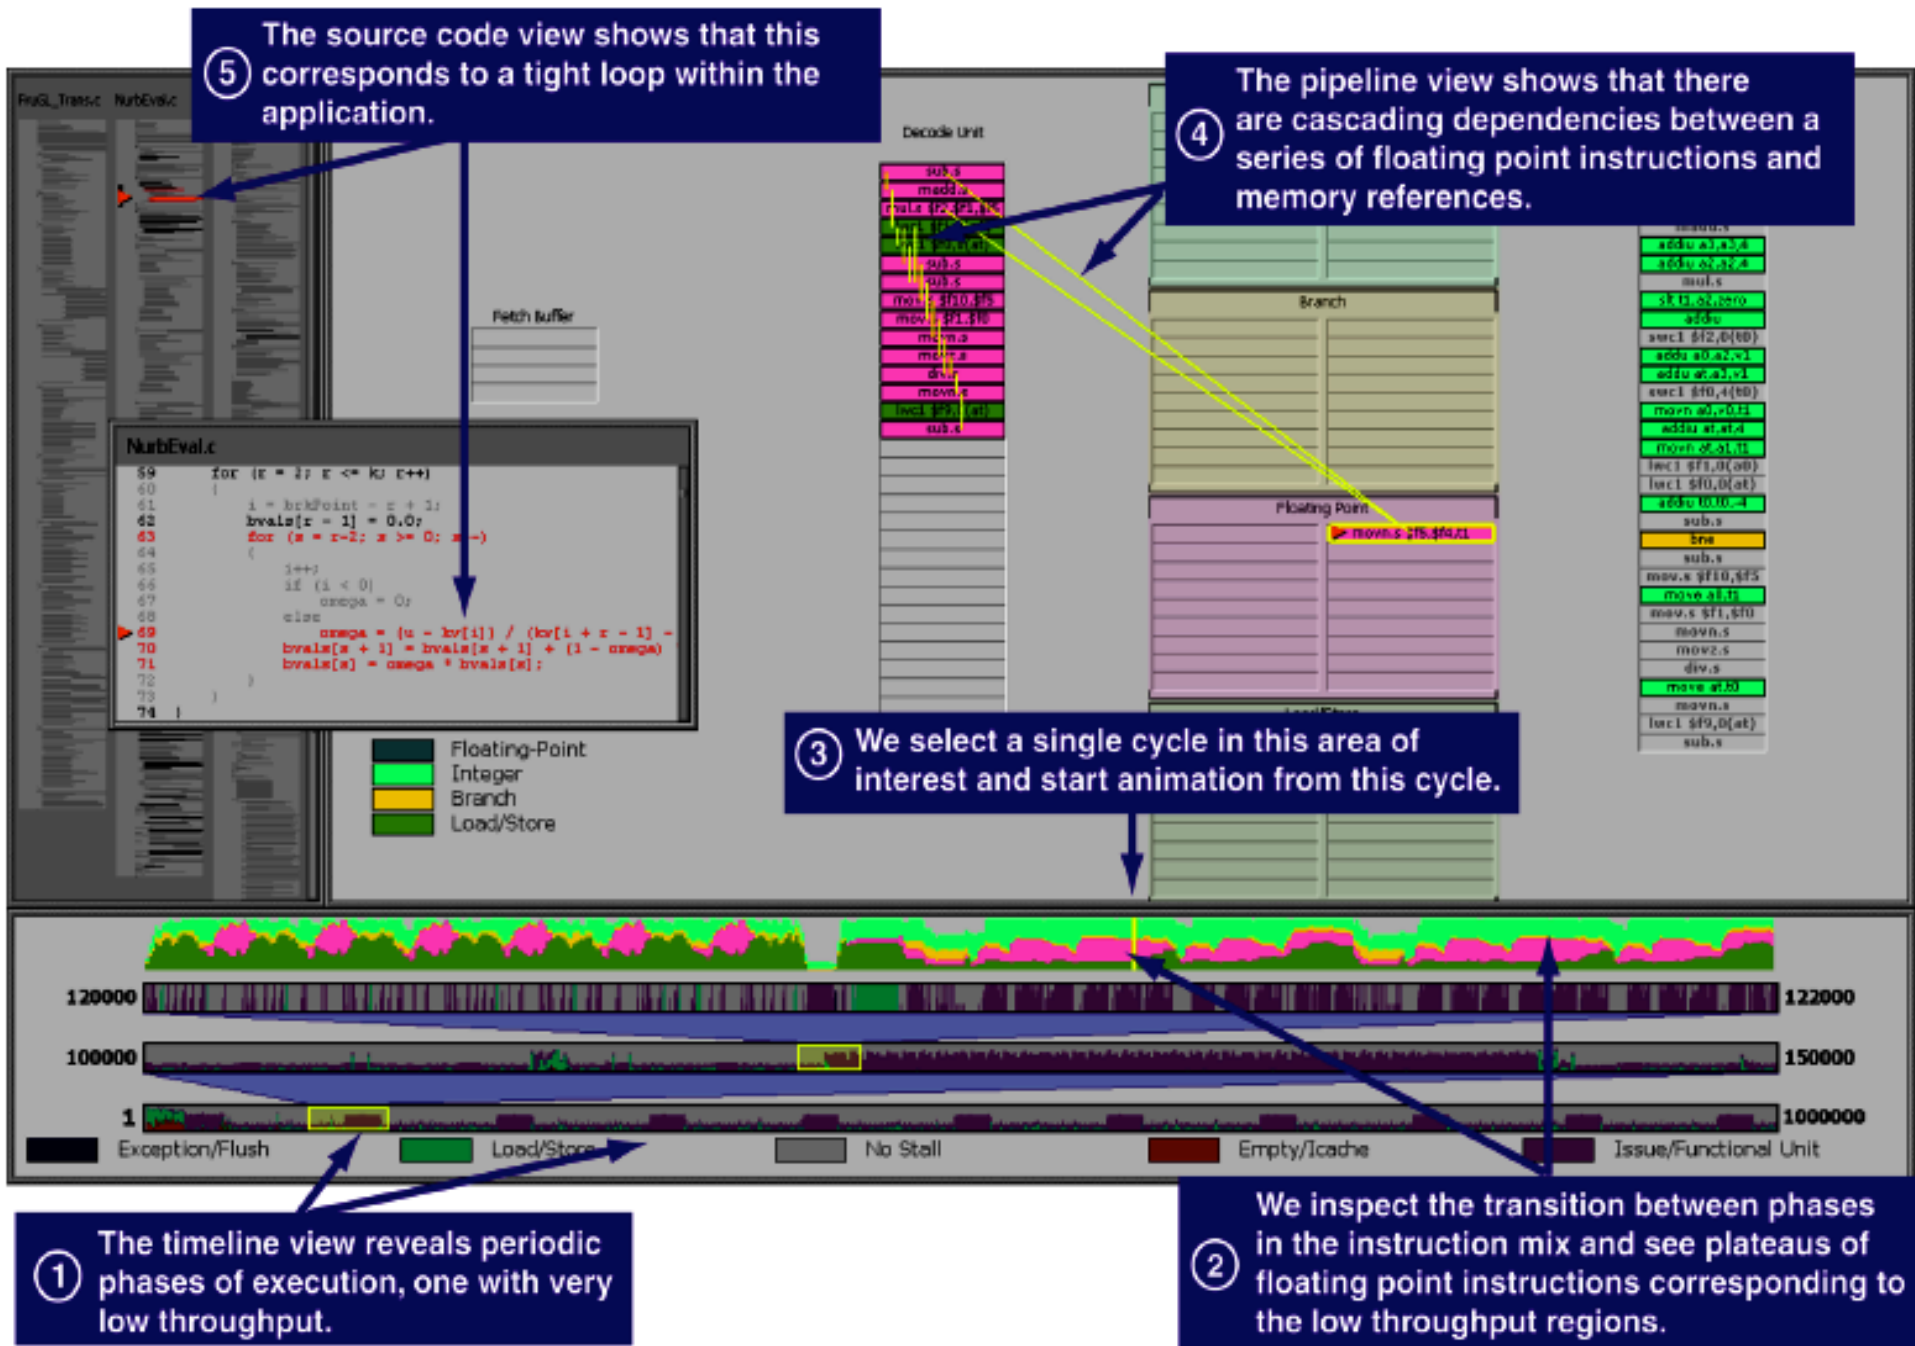
FruGL_Trans.c	NurbEval.c	TimeRun.c

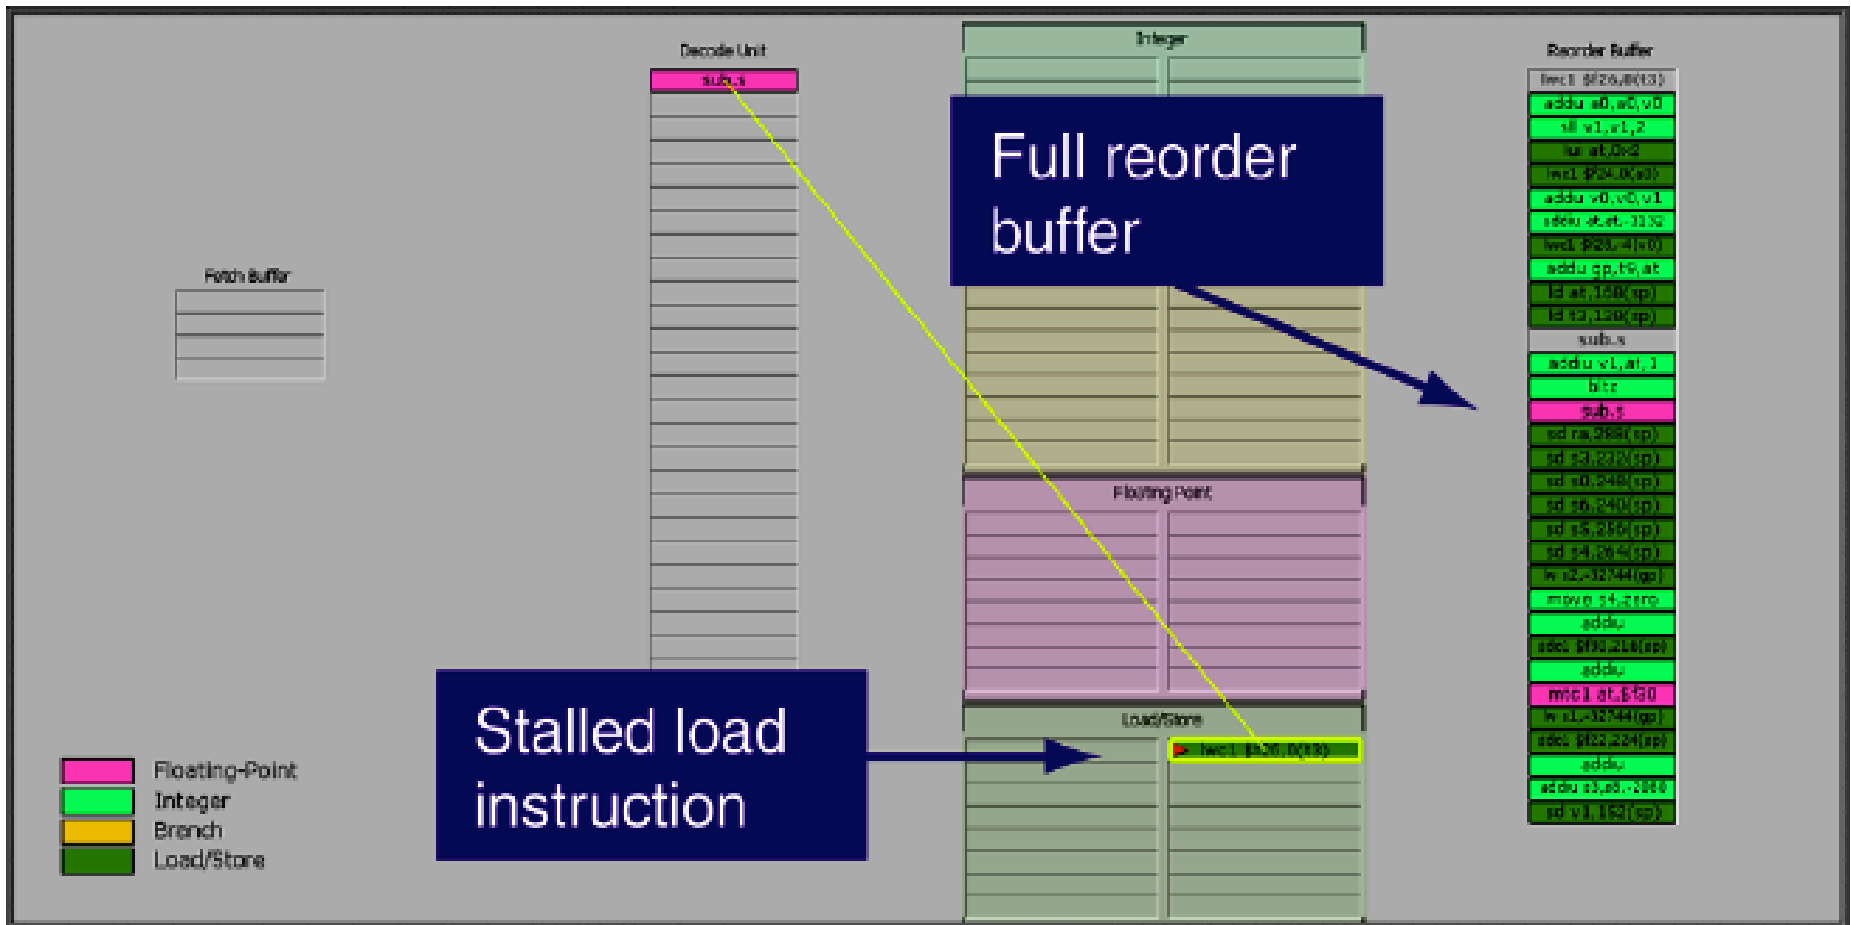
  

```

61     i = brkPoint - r + 1;
62     bvals[r - 1] = 0.0;
63     for (s = r-2; s >= 0; s--)
64     {
65         i++;
66         if (i < 0)
67             omega = 0;
68         else
69             omega = (u - kv[i]) / (kv[i + r - 1] - kv[i]);
70     bvals[s + 1] = bvals[s + 1] + (1 - omega) * bvals[s];
71     bvals[s] = omega * bvals[s];
72     }
73 }
74 )
75
76 /*
77  * Compute derivatives of the basis functions Bi,k(u)
78  */
79 static void
80 BasisDerivatives( float u, long brkPoint, float * kv, long k, float * dvals
81 {
82     register long s;

```






Memory stall

Dependencies exist between all of these instructions

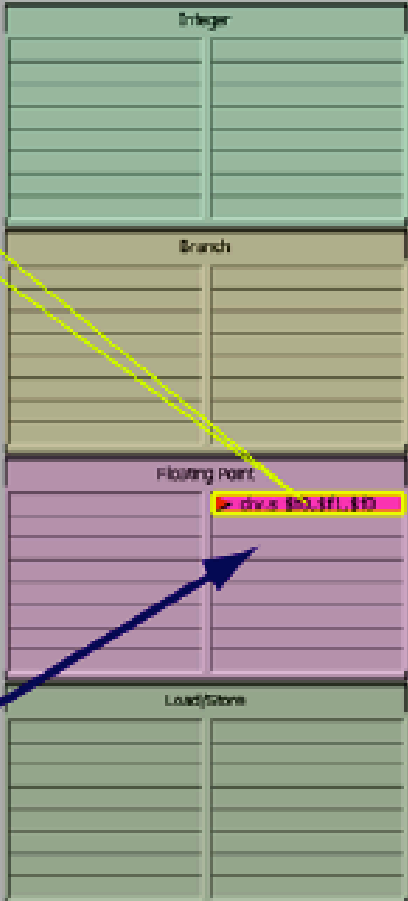


```
Decode Unit
mov.s $f0,$f0
mov.s $f0,$f0
mov.s $f0,$f0
mov.s $f1,$f1
mov.s $f1,$f1
mov.s $f2,$f2
srec.s $f0,$f0
srec.s $f0,$f0
lrec.s $f0,$f0
lrec.s $f0,$f0
sub.s $f0,$f0,$f0
sub.s $f0,$f0,$f0
mov.s $f0,$f0
mov.s $f0,$f0
mov.s $f0,$f0
div.s
mov.s
```

Instructions are being executed sequentially

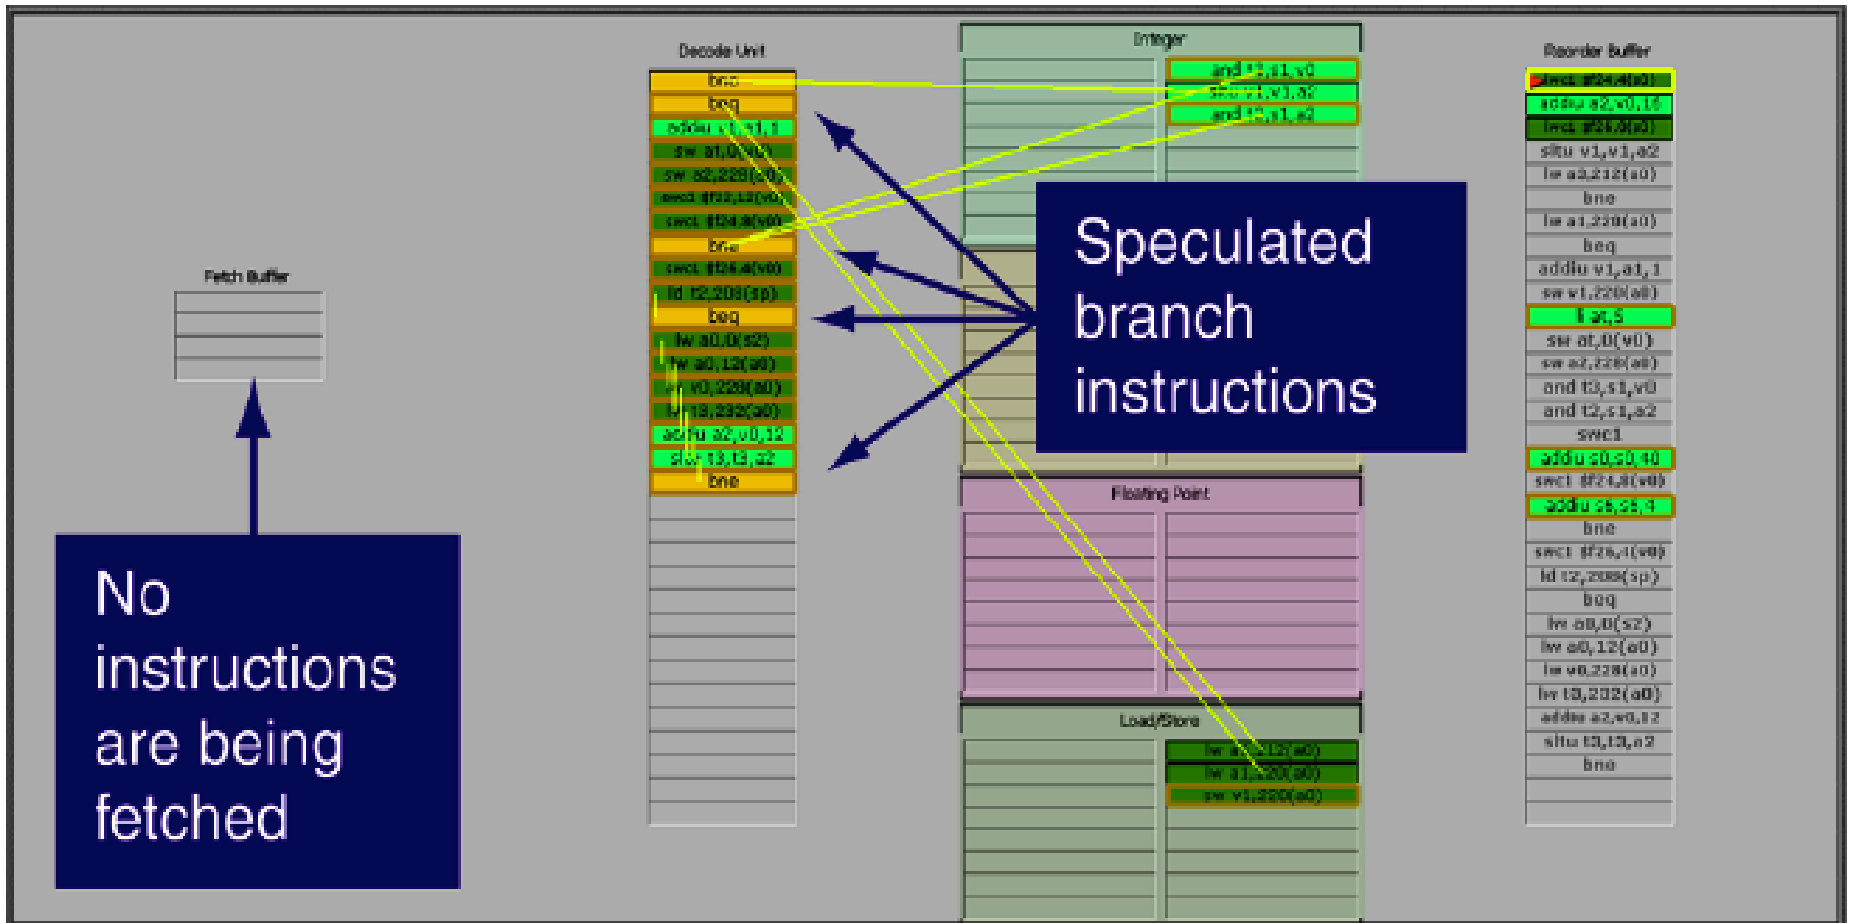


- Floating-Point
- Integer
- Branch
- Load/Store

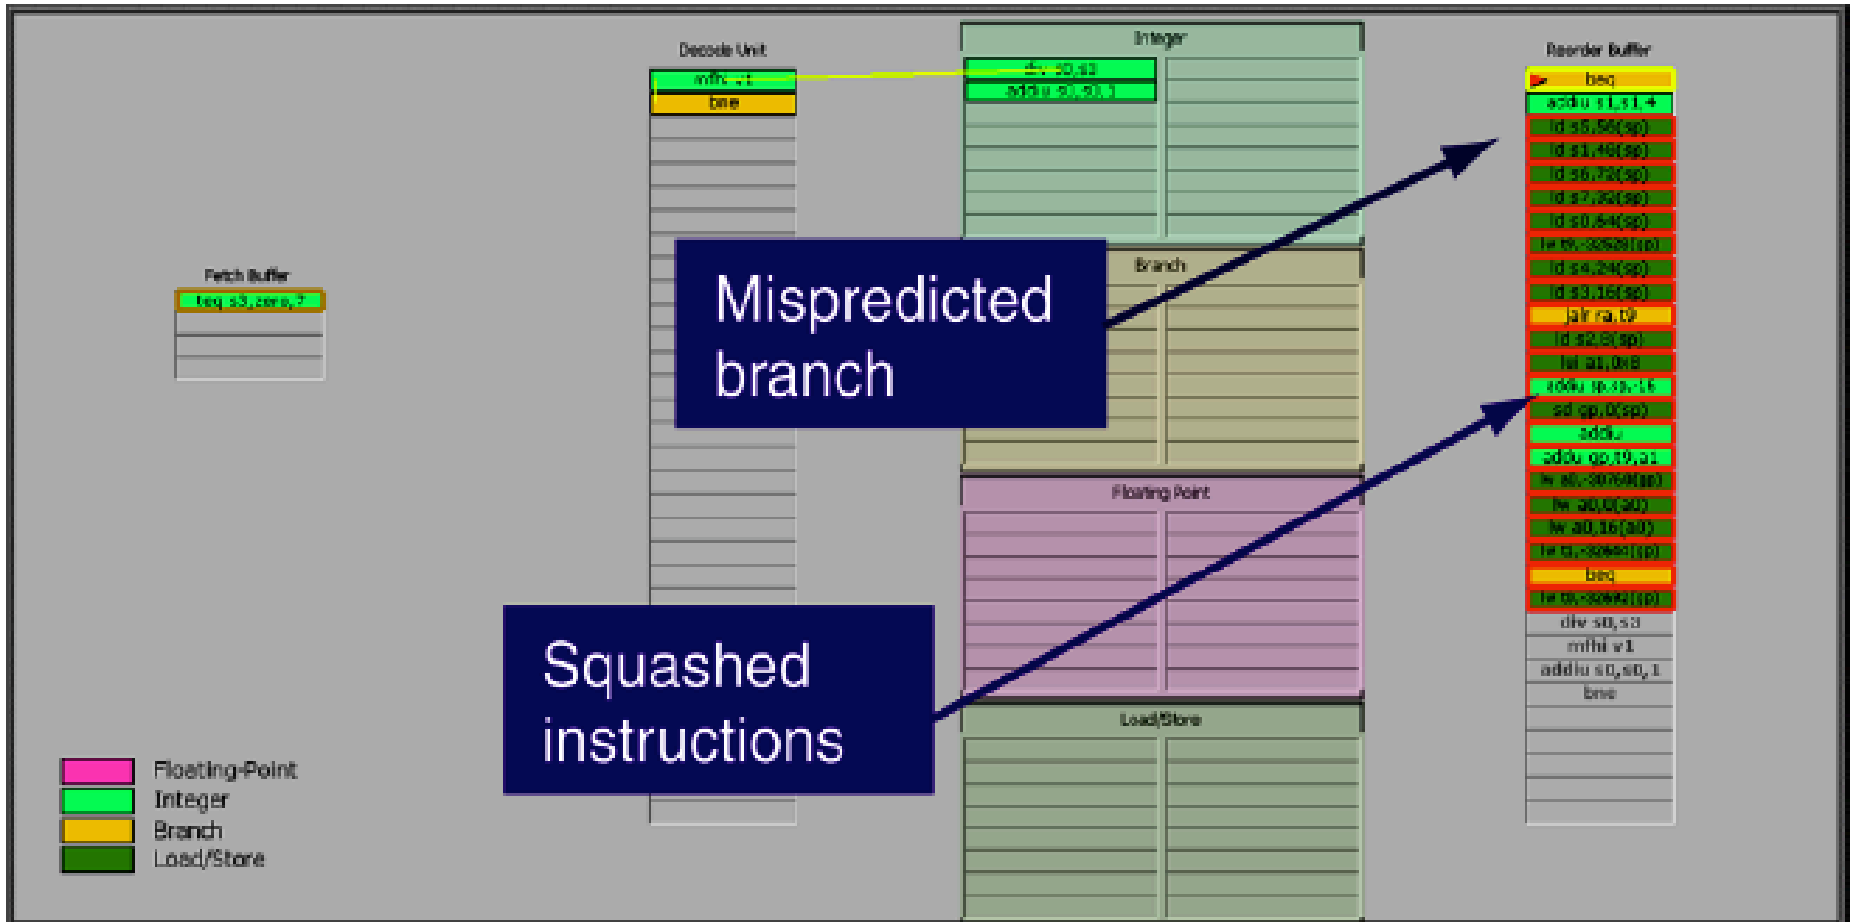


```
Reorder Buffer
div.s $f0,$f0,$f0
mov.s $f0,$f0
mov.s
lrec.s $f0,$f0
sub.s $f0,$f0,$f0
mov.s
mov.s
addu a3,a3,4
addu a2,a2,4
mul.s $f2,$f1,$f0
lt t1,a2,zero
addu
srec.s $f2,$f0,$f0
addu a0,a0,v1
addu a1,a3,v1
srec.s $f0,$f0,$f0
mov.s a2,v0,t1
addu a1,a1,4
mov.s a1,a1,t1
lrec.s $f1,$f0,$f0
lrec.s $f0,$f0,$f0
addu t0,t0,-4
sub.s $f0,$f0,$f0
bne
sub.s $f0,$f0,$f0
mov.s $f1,$f0
mov.s
mov.s
div.s
move a1,t0
mov.s
```

# Dependencies



Deep speculation



## Branch misprediction

# Critique

- Most code doesn't need this level of optimization, but
  - The visualization is effective, and would be useful for code that does
  - May reduce the expertise needed to perform low level optimization
- Might be effective as a teaching tool
- Bad color scheme: black/purple/brown
- Does it scale with processor complexity?



# Papers

- D. Kimelman, B. Rosenburg, and T. Roth, "Strata-Various: Multi-Layer Visualization of Dynamics in Software System Behavior," Proc. Fifth IEEE Conf. Visualization '94, IEEE Computer Society Press, Los Alamitos, Calif., 1994, pp. 172–178.
- James Eagan, Mary Jen Harrold, James A. Jones, and John Stasko, "Visually Encoding Program Test Information to Find Faults in Software." Proc. InfoVis 2001 pp. 33-36.

# Papers

- Alessandro Orso, James Jones, and Mary Jean Harrold. "Visualization of Program-Execution Data for Deployed Software." Proc. of the ACM Symp. on Software Visualization, San Diego, CA, June 2003, pages 67--76.
- Chris Stolte, Robert Bosch, Pat Hanrahan, and Mendel Rosenblum, "Visualizing Application Behavior on Superscalar Processors." Proc. InfoVis 1999