

Improving Attribution of Performance Measurements for Optimized Code

John Mellor-Crummey and Mark Krentel

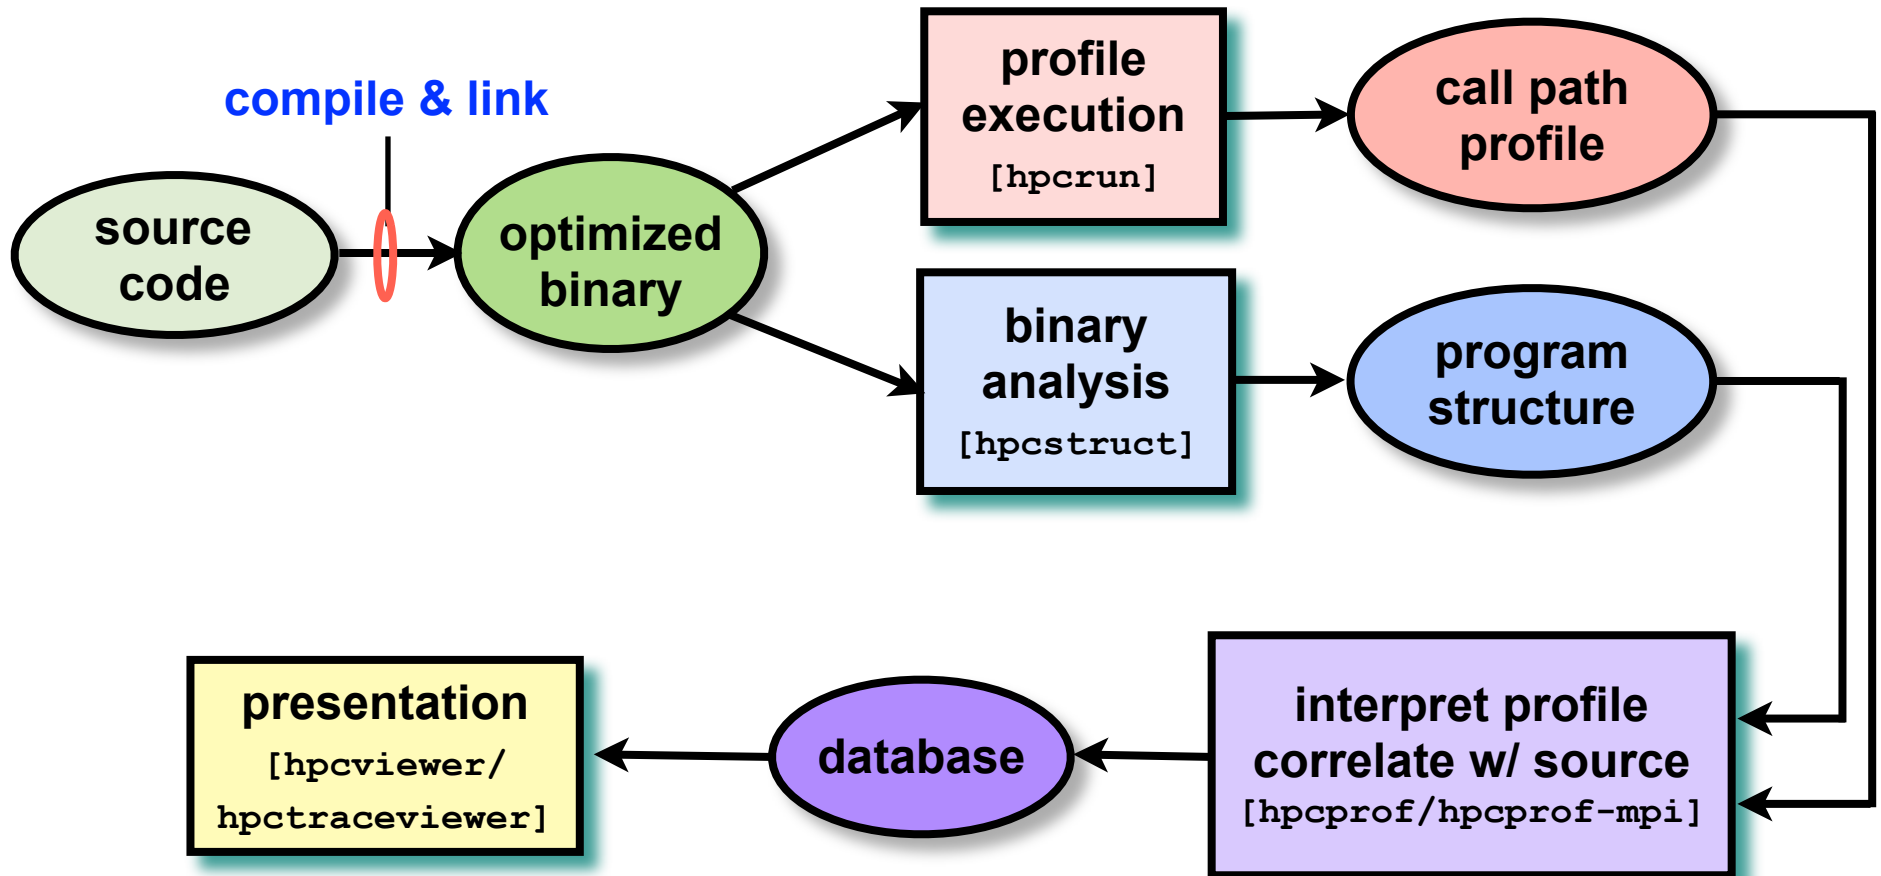
Department of Computer Science
Rice University

<http://hpctoolkit.org>

Motivation

- Modern software uses abstractions to manage complexity
 - procedures
 - classes
 - parameterized templates for algorithms and data structures
- Programmers rely on optimizing compilers to transform abstractions for efficient execution
 - compose algorithm and data structure templates
 - e.g., C++ Standard Template Library (STL), Boost, ...
 - inline procedures
 - transform loop nests
- Understanding the performance of modern software requires measuring the performance of optimized code and relating measurements back to the program source code

HPCToolkit Workflow




Call Path Profiling

Measure and attribute costs in context


sample timer or hardware counter overflows

gather calling context using stack unwinding

Call path sample



Calling context tree




Overhead proportional to sampling frequency...
...not call frequency


Understanding Optimized Code can be Difficult

- Structure of code is radically different after template instantiation, function inlining, and loop transformations
 - functions contain code from multiple files and functions

CCT unoptimized code



CCT optimized code



- Control flow graph structure is often rather complex
 - more than simple loops

Starting Point for This Work

Nathan Tallent, John Mellor-Crummey, and Michael Fagan. Binary analysis for measurement and attribution of program performance. PLDI '09. ACM, New York, NY, 441-452

- Binary analysis for call stack unwinding of unmodified optimized code
 - need to determine return address
 - parent's value for frame pointer register
- Binary analysis for attribution of performance to optimized code
 - identified inlined code as code from different source file
 - reported only one level of inlining
 - enclosing context
 - a single source line mapping for each generated instruction

An Example: small.cpp

```
using namespace std;
vector <int> v;
inline static void addToVector(int i) {
    v.push_back(i);
}
void do_work(int num) {
    v.clear();
    for (int i = 0; i < num; i++) {
        addToVector(i);
    }
}
int main(int argc, char **argv) {
    int len = 1000;
    int num, k;
    if (argc < 2 || sscanf(argv[1], "%d", &num) < 1) {
        num = 20;
    }
    num *= len;
    for (k = 0; k < num; k++) {
        do_work(len);
    }
    return 0;
}
```

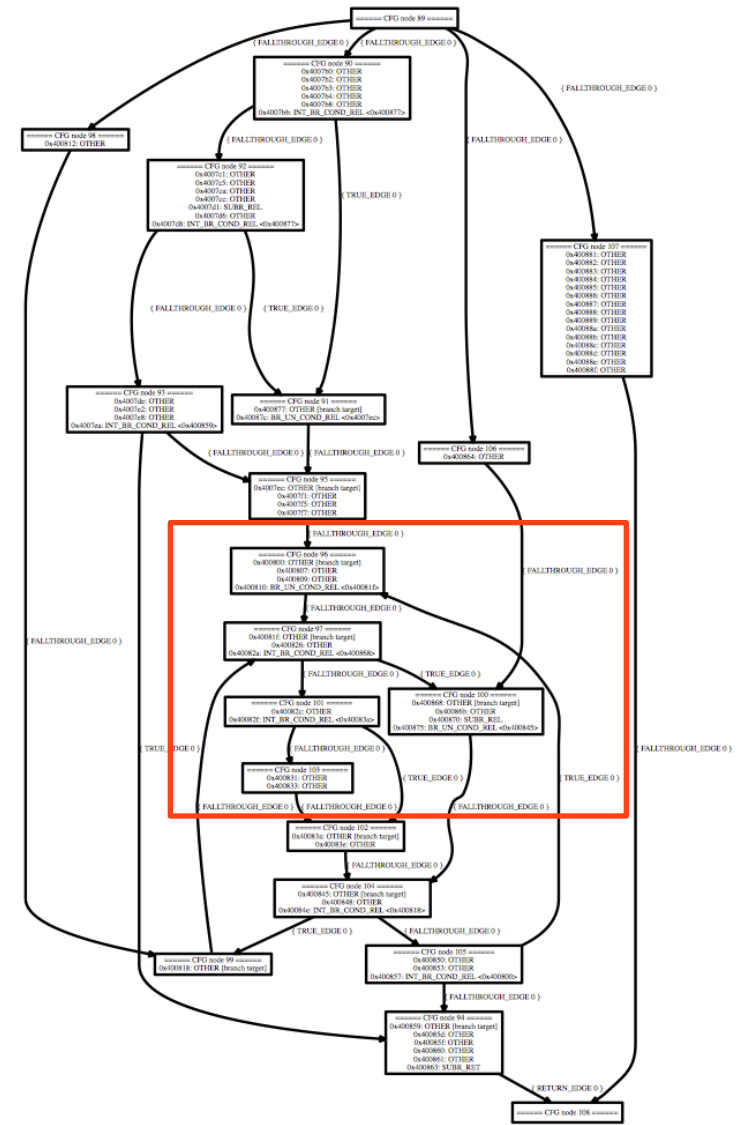
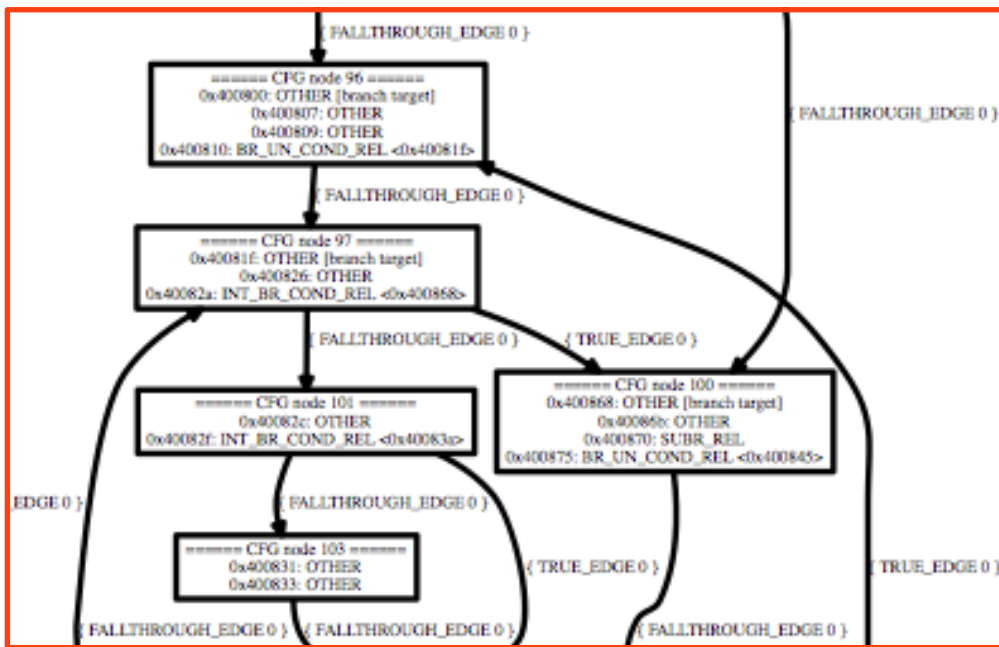
Generated Code for small.cpp (g++ 4.4.6)

91 lines of assembly code for main

- Multiple levels of inlining
- Inlines the following functions
 - dowork
 - addToVector
 - vector::push_back
 - __gnu_cxx::new_allocator
 - vector::clear
 - vector::_M_erase_at_end
- Only two function calls left
 - iterator in push_back
 - scanf

Construct the CFG

- Parse the machine code in an executable
- Build a CFG at the level of basic blocks




Identify Loops

Directed Graph $G = (V, E)$

- Dominator
 - $x \text{ dom } y$ iff every execution path from entry to y goes through x
- Natural loop
 - defined by a back edge $y \rightarrow x$ where $x \text{ dom } y$
 - finds only single-entry loops
- Tarjan's algorithm finds single-entry, strongly-connected subgraphs
 - Robert Tarjan, "Depth-first search and linear graph algorithms," SIAM Journal on Computing 1(2):146–160, June 1972.
 - sketch
 - based on depth-first search
 - an SCC body includes nodes that reach a lower node than itself
 - loop head: node where lowest reachable is itself
 - complexity: $O(V + E)$

Coping with Irreducible Loops

- Problem: not all cycles are single-entry loops
 - multiple entry loop: irreducible
- Paul Havlak. Nesting of reducible and irreducible loops. ACM TOPLAS 19(4):557-567, 1997.
 - uses definitions of reducible and irreducible loops which allows arbitrary nesting of either kind of loop
 - loop nesting tree can depend on the depth-first spanning tree used to build it
 - header node representing a reducible loop in one version of loop nesting tree can represent an irreducible loop in another



Challenges to CFG Construction

- Compiler optimizations make it difficult to recover accurate CFGs
 - tail calls
 - functions that don't return, e.g., `exit`, `__cxa_throw`, `longjmp`, ...
 - calls to through PLT to dynamically-linked routines
 - calls to routines statically-linked in a load module
- No indication of these features in DWARF
 - recover this info by processing `/usr/include` and C++ ABI headers

Tail Call Example from LLNL's LULESH

Fragment of source code

```
if ( hgcoef > Real_t(0.) ) {  
    CalcFBHourglassForceForElems(determ,x8n,y8n,z8n,dvdx,dvdy,dvdz,hgcoef);  
}  
  
Release(&z8n) ;  
Release(&y8n) ;  
Release(&x8n) ;  
Release(&dvdz) ;  
Release(&dvdy) ;  
Release(&dvdx) ;  
  
return ;
```

Sketch of generated code (gcc 4.4.6 -O3)

```
if ( hgcoef > Real_t(0.) ) goto calc  
rel: free(&z8n)  
     free(&y8n)  
     free(&x8n)  
     free(&dvdz)  
     free(&dvdy)  
     push &dvdx  
     jmp free  
calc:inlined code for CalcFBHourglassForceForElems  
     goto rel
```

Non-returning Function Example from miniFE

- Non-returning functions occur frequently, even in scientific codes
 - casting associated with inlined C++ I/O helper routines

```
#ifndef _BASIC_IOS_H
...
__GLIBCXX_BEGIN_NAMESPACE(std)
template<typename _Facet>
inline const _Facet&
__check_facet(const _Facet* __f)
{
    if (!__f)
        __throw_bad_cast();
    return *__f;
}
...


```

Mapping Back to Program Structure

- For each instruction, identify its full provenance
 - use DWARF info to recover complete static call chains
 - recover a full inlined call chain for each machine instruction
- Integrate information about loops and inlining to assemble a representation of static structure
- Not as simple as it sounds
 - where do loops belong in an inlined call chain?

Source Code Attribution for Loops

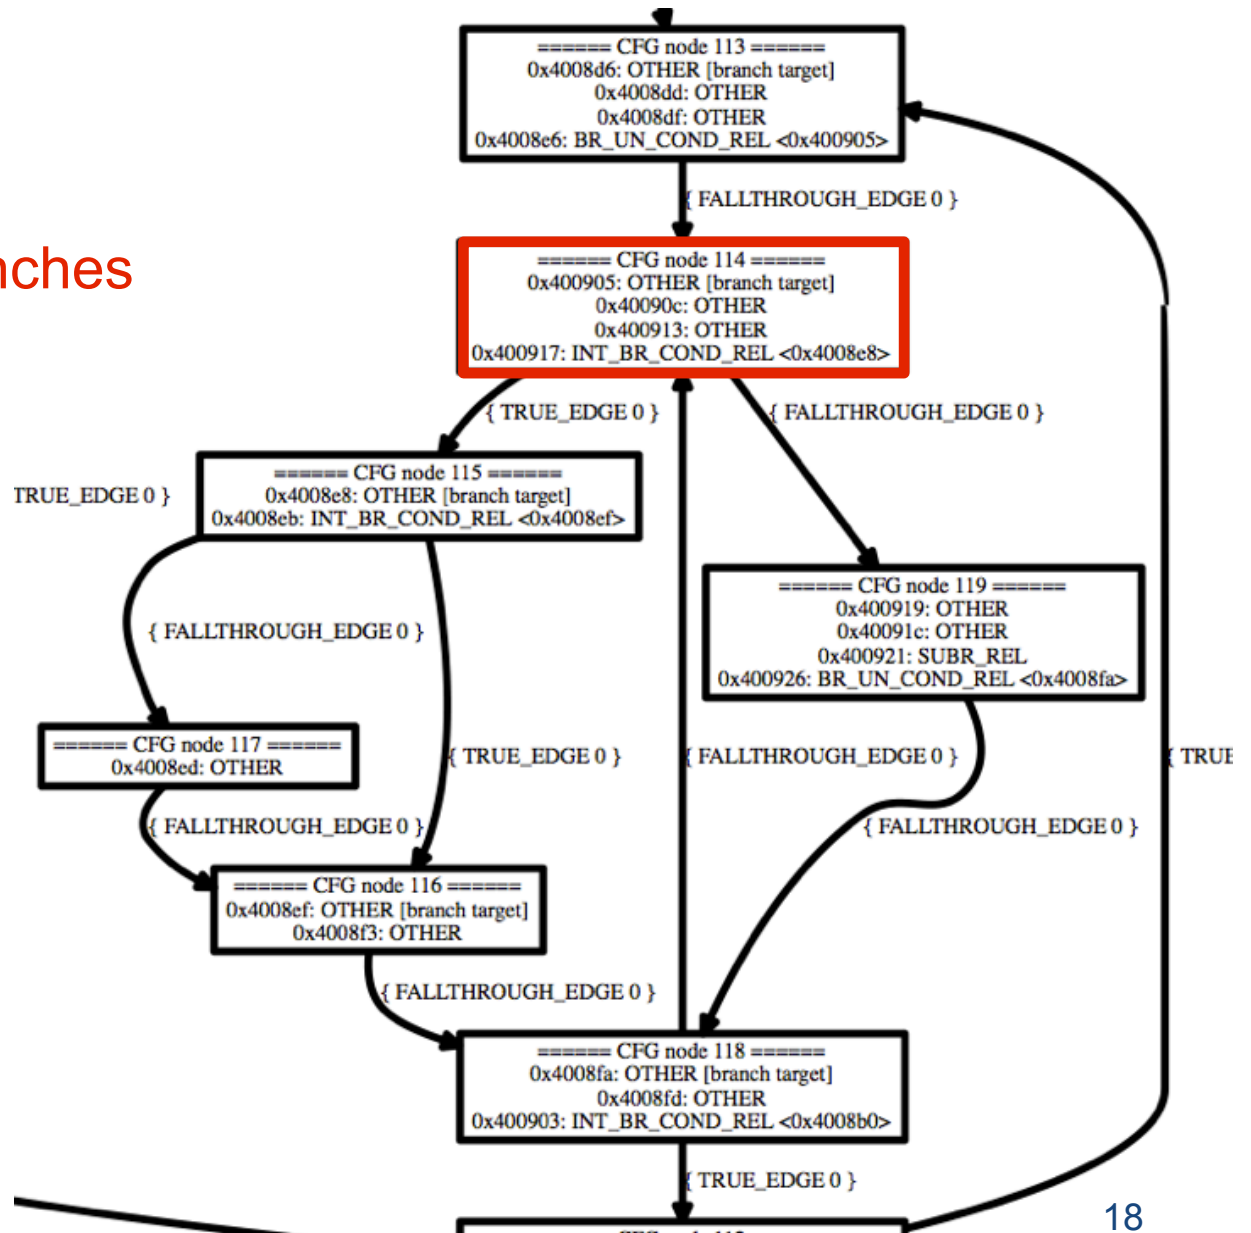
- Need to identify a source code position for each Interval and Irreducible interval
- What line number to use?
 - source line for first machine instruction in loop header?
 - source line for backward branch reaching loop header?
 - some complications ...
 - edges reaching loop header are not always backward branches



g++ 4.1.2

Detail of CFG for main (gcc 4.1.2)

Only fall through branches reach this header!



Associating a Loop with a Source Line

Today's heuristic

- Priority scheme
 - back edge
 - backward branch closing natural loop
 - true branches from within the loop
 - fall through edges from within the loop
- If none of these has a source mapping, use the mapping for the loop header
- If the source mapping for the loop header is less deeply nested than the source of the edge targeting it, use that instead

Assembling the Source View

- Perform interval analysis of the CFG
- Recursively assemble the CCT for a procedure
 - for each interval
 - insert source code for all machine instructions inside into CCT
 - insert the call chain for the loop
 - never make the loop a child of any node inserted inside the loop
 - create copies of context where necessary
 - identify the least common ancestor between a loop and the calling context for machine instruction inside it
 - treat copies of contexts along respective paths as equivalent
 - take the path below the LCA and insert that inside the loop
- For each “alien” context in inlined code, record information about
 - call site
 - callee
- Gracefully handle case where no static call chain information available
 - simply indicate that inlined code came from the following source file and line
- Present this in hpcviewer’s source code view as if real call chains, but indicate when function is inlined

LULESH: Attribution for Optimized Code

- Present full calling context and loops, as if an unoptimized executable

The screenshot displays a debugger window for the file `lulesh.cc`. The source code is shown with line numbers 1362 through 1378. Line 1365, `CalcFBHourglassForceForElems(determ,x8n,y8n,z8n,dvdx,dvdy,dvdz,hgcoef);`, is highlighted. Below the code, the 'Calling Context View' is active, showing a tree of scopes. A table below the tree provides performance metrics for each scope.

Scope	WALLCLOCK (usec).[0,0] (I)	WALLCLOCK (usec).[0,0] (E)
Experiment Aggregate Metrics	3.02e+07 100 %	3.02e+07 100 %
main	3.02e+07 99.9%	
loop at lulesh.cc: 2871	3.02e+07 99.9%	
2873: LagrangeLeapFrog	3.02e+07 99.9%	
2647: [I] LagrangeNodal	2.00e+07 66.1%	1.40e+05 0.5%
1498: [I] CalcForceForNodes	1.95e+07 64.4%	8.02e+04 0.3%
1422: [I] CalcVolumeForceForElems	1.94e+07 64.1%	2.00e+04 0.1%
1403: CalcHourglassControlForElems	1.54e+07 50.8%	4.47e+06 14.8%
1365: [I] CalcFBHourglassForceForElems	7.79e+06 25.8%	6.65e+06 22.0%
loop at lulesh.cc: 1178	7.79e+06 25.8%	6.65e+06 22.0%
92: cbrt	1.12e+06 3.7%	8.19e+05 2.7%
lulesh.cc: 1201	7.79e+05 2.6%	7.79e+05 2.6%
lulesh.cc: 1205	6.99e+05 2.3%	6.99e+05 2.3%
lulesh.cc: 611	4.00e+05 1.3%	4.00e+05 1.3%

A vertical black bar on the left contains the text 'LULESH' in white, with arrows pointing to the corresponding scopes in the table.

miniFE with Non-returning Function Analysis

```
171 TICK();
172 p_ap_dot = matvec_and_dot(A, p, Ap);
173 TOCK(tMATVECDOT);
174 #else
175 TICK(); matvec(A, p, Ap); TOCK(tMATVEC);
176
177 TICK(); p_ap_dot = dot(Ap, p); TOCK(tDOT);
178 #endif
179
180 #ifdef MINIFE_DEBUG
181     os << "iter " << k << " p_ap_dot = " << p_ap_dot << endl;
```

Scope	WALLCLOCK (usec).[0,0] (l)	WALLCLOCK (usec).[0,
Experiment Aggregate Metrics	1.37e+07 100 %	1.37e+07
main	1.37e+07 100 %	
147: int miniFE::driver<double, int, int>(Box const&, Box&, mi	1.37e+07 100 %	
289: void miniFE::cg_solve<miniFE::CSRMatrix<double, int,	8.72e+06 63.7%	
loop at cg_solve.hpp: 152	8.61e+06 62.9%	
loop at cg_solve.hpp: 161	8.59e+06 62.7%	
175: [l] miniFE::matvec_std<miniFE::CSRMatrix<do	7.03e+06 51.4%	7.03e+06
loop at SparseMatrix_functions.hpp: 570	7.03e+06 51.4%	2.60e+05
loop at SparseMatrix_functions.hpp: 573	6.77e+06 49.5%	6.77e+06
SparseMatrix_functions.hpp: 574	4.79e+06 34.9%	4.79e+06
SparseMatrix_functions.hpp: 573	1.99e+06 14.5%	1.99e+06
SparseMatrix_functions.hpp: 74	1.80e+05 1.3%	1.80e+05
SparseMatrix_functions.hpp: 577	5.99e+04 0.4%	5.99e+04
SparseMatrix_functions.hpp: 571	2.00e+04 0.1%	2.00e+04
210: [l] waxpby<miniFE::Vector<double, int, int>	3.80e+05 2.8%	3.80e+05
158: [l] waxpby<miniFE::Vector<double, int, int>	3.70e+05 2.7%	3.70e+05
209: [l] waxpby<miniFE::Vector<double, int, int>	2.80e+05 2.0%	2.80e+05
177: [l] dot<miniFE::Vector<double, int, int>	2.70e+05 2.0%	2.70e+05
156: [l] dot<miniFE::Vector<double, int, int>	2.40e+05 1.8%	2.40e+05

903

miniFE without Non-returning Function Analysis

```
basic_string.h  atomicity.h  SparseMatrix_function...  CSRMatrix.hpp  cg_solve.hpp
73 static inline _Atomic_word
74 __attribute__((__unused__))
75 __exchange_and_add_dispatch(_Atomic_word* __mem, int __val)
76 {
77 #ifdef __GTHREADS
78   if (__gthread_active_p())
79     return __exchange_and_add(__mem, __val);
80   else
81     return __exchange_and_add_single(__mem, __val);
82 #else
83   return __exchange_and_add_single(__mem, __val);
84 #endif
85 }
86
87 static inline void
88 __attribute__((__unused__))
```

bogus loop distorts
CFG for miniFE::driver

Calling Context View Callers View Flat View

Scope | WALLCLOCK (usec).[0,0] (l) | W

Scope	WALLCLOCK (usec).[0,0] (l)	W
Experiment Aggregate Metrics	1.37e+07	100 %
main	1.37e+07	100 %
147: int miniFE::driver<double, int, int>(Box const&, Box&, miniFE::Parameters&, YAM	1.37e+07	100 %
143: [l] compute_imbalance<int>	1.37e+07	99.8 %
88: [l] ~basic_string	1.37e+07	99.8 %
[l] std::string::Rep::M_dispose	1.37e+07	99.8 %
[l] __exchange_and_add_dispatch	1.37e+07	99.8 %
loop at atomicity.h: 78	1.37e+07	99.8 %
289: void miniFE::cg_solve (miniFE::CSRMatrix<double, int, int>, mi	8.72e+06	63.7 %
loop at atomicity.h: 64	4.46e+06	32.6 %
218: void miniFE::impose_dirichlet<miniFE::CSRMatrix<double, int,	4.80e+05	3.5 %
179: [l] ~CSRMatrix	3.20e+04	0.2 %

CFG

What's left?

- Technical issues
 - explore cases where embedding of loops in static call chains still isn't satisfactory
 - is there a better interpretation of the graph depending on depth first parse
 - can exhaustive analysis of a loop yield better results?
 - beyond just looking at loop header and incident edges
 - new 2007 flow graph analysis algorithm
 - better results?
 - better performance?
 - analysis speed for huge binaries?
- Community issues
 - lobby DWARF community to enhance standard with information about functions that don't return

Flowgraph Analysis References

- Robert Tarjan, “Depth-first search and linear graph algorithms,” *SIAM Journal on Computing* 1(2):146–160, June 1972.
- Paul Havlak. Nesting of reducible and irreducible loops. *ACM TOPLAS* 19(4): 557–567, July 1997.
- Tao Wei, Jian Mao, Wei Zou, and Yu Chen. A New Algorithm for Identifying Loops in Decompilation. *Static Analysis 14th International Symposium (SAS)*, LNCS 4634, pp. 170–183, 2007.