

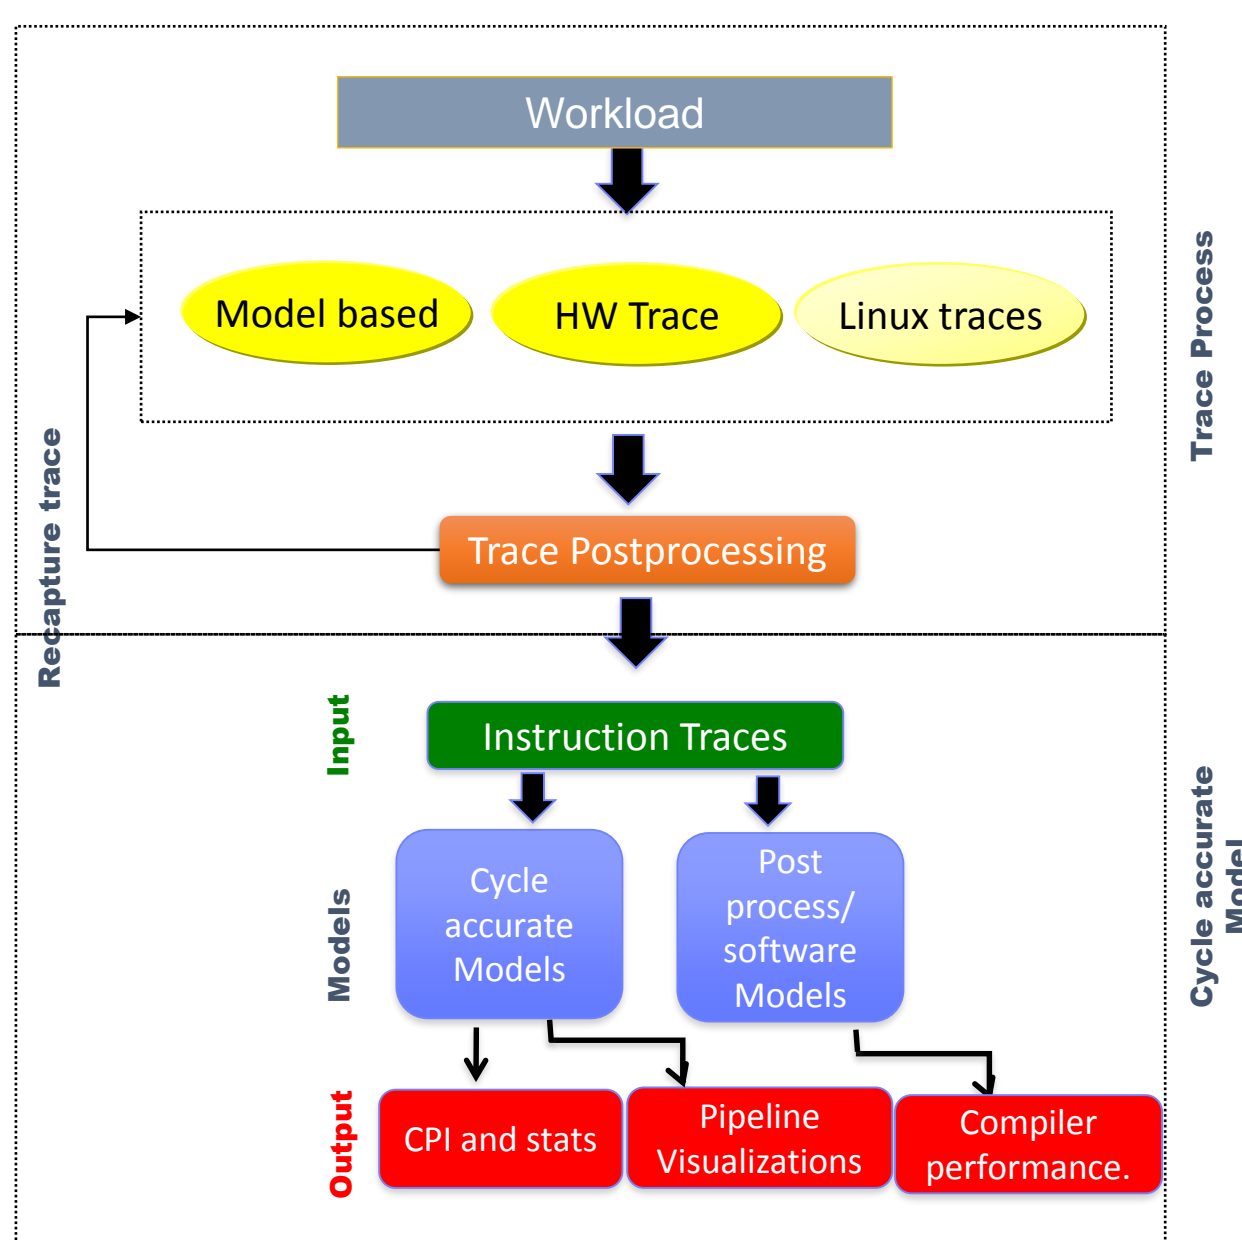
Instruction Tracing and dynamic codegen analysis to identify unique llvm performance issues.

Biplob Mishra

INTRODUCTION

Performance analysis of the machine code generated by a compiler can be carried out in different ways and can also be based on application in question. Common methods use some form of profiling on a running program which generally provides the statistical information about certain data and events. While this method does give important insights to a performance problem, some of the issues are more clearly understood when the compiled applications is actually run and the dynamic instructions of hot code execution paths are traced and analyzed in a small execution window. Trace records contain instructions and data, memory addresses and other information which provide complete visibility into the workings of an application.

TRACE OVERVIEW



```
68 00000001013e00 0 385050 addi r5,r21,80
69 00000001013e04 0 38505e addi r25,r21,84
70 00000001013e08 0 f18170 ne r1,r14
71 00000001013e0c 0 f81038 std r3,216(r1)
72 00000001013e10 0 38701e addi r3,r21,28
73 00000001013e14 0 f8103b std r3,208(r1)
74 00000001013e18 0 f8e103 ld r3,488(r1)
75 00000001013e1c 0 388103 ld r0,r(17)
76 00000001013e20 0 288000 cmovt cr1,r,8
77 00000001013e24 0 e3e1f9 ld r5,504(r1)
78 00000001013e28 0 832000 ld r6,r(15)
79 00000001013e2c 0 7c31a9 cmpld r1,r,4,133
80 00000001013e30 0 00 00 00 400000 bpe r0,(r17)
81 00000001013e34 0 00 00 00 288000 cmovl cr1,r,4,133
82 00000001013e38 0 00 00 00 418000 cmovl cr1,r,4,133
83 00000001013e3c 0 788020 cl_rldi r3,r,4,32
84 00000001013e40 0 795154 rdictr r3,r3,2,81
85 00000001013e44 0 7c31a9 lsw r3,r,3,8
86 00000001013e48 0 7c31a9 add r3,r,8
87 00000001013e4c 0 7c31a9 mctr r3
88 00000001013e50 1 00 00 4e90d0 bctr r3,r,112(r19)
89 00000001013e54 0 80c238 lsw r4,-15656(r2)
90 00000001013e58 0 8a6000 lsw r3,r,15)
91 00000001013e5c 0 704123 cmov r4,r3
```

Sample trace with data addresses

Trace collection

```
00000001013e04 0 addi r25,r21,84
00000001013e08 0 ne r1,r14
00000001013e0c 0 addi r3,r21,28
00000001013e10 0 std r3,208(r1)
00000001013e14 0 ld r3,488(r1)
00000001013e18 0 cmpld cr1,r,4,133
00000001013e1c 0 ld r0,r(17)
00000001013e20 0 cmovt cr1,r,8)
00000001013e24 0 cmovl cr1,r,4,133
00000001013e28 0 cmovl cr1,r,4,133
00000001013e2c 0 cl_rldi r3,r,4,32
00000001013e30 0 rdictr r3,r3,2,81
00000001013e34 0 lsw r3,r,3,8
00000001013e38 0 add r3,r,8
00000001013e3c 0 mctr r3
00000001013e40 0 bctr r3,r,112(r19)
00000001013e44 0 lsw r4,-15656(r2)
00000001013e48 0 lsw r3,r,15)
00000001013e4c 0 lsw r4,133
00000001013e50 0 cmpld r1,r,4,133
00000001013e54 0 ne r1,r14
00000001013e58 0 cmpld cr1,r,4,133
00000001013e5c 0 cmovl cr1,r,4,133
00000001013e60 0 cmovl cr1,r,4,133
00000001013e64 0 cmovl cr1,r,4,133
00000001013e68 0 cmovl cr1,r,4,133
00000001013e6c 0 cmovl cr1,r,4,133
00000001013e70 0 cmovl cr1,r,4,133
00000001013e74 0 cmovl cr1,r,4,133
00000001013e78 0 cmovl cr1,r,4,133
00000001013e7c 0 cmovl cr1,r,4,133
00000001013e80 0 cmovl cr1,r,4,133
00000001013e84 0 cmovl cr1,r,4,133
00000001013e88 0 cmovl cr1,r,4,133
00000001013e8c 0 cmovl cr1,r,4,133
00000001013e90 0 cmovl cr1,r,4,133
00000001013e94 0 cmovl cr1,r,4,133
00000001013e98 0 cmovl cr1,r,4,133
00000001013ea0 0 cmovl cr1,r,4,133
00000001013ea4 0 cmovl cr1,r,4,133
00000001013ea8 0 cmovl cr1,r,4,133
00000001013eac 0 cmovl cr1,r,4,133
00000001013eaf 0 cmovl cr1,r,4,133
00000001013eb0 0 cmovl cr1,r,4,133
```

-> Trace is the dynamic sequence of executed instructions as it occurred when the program was executed.

-> Trace provides complete history of what happens in a processor when the an application runs.

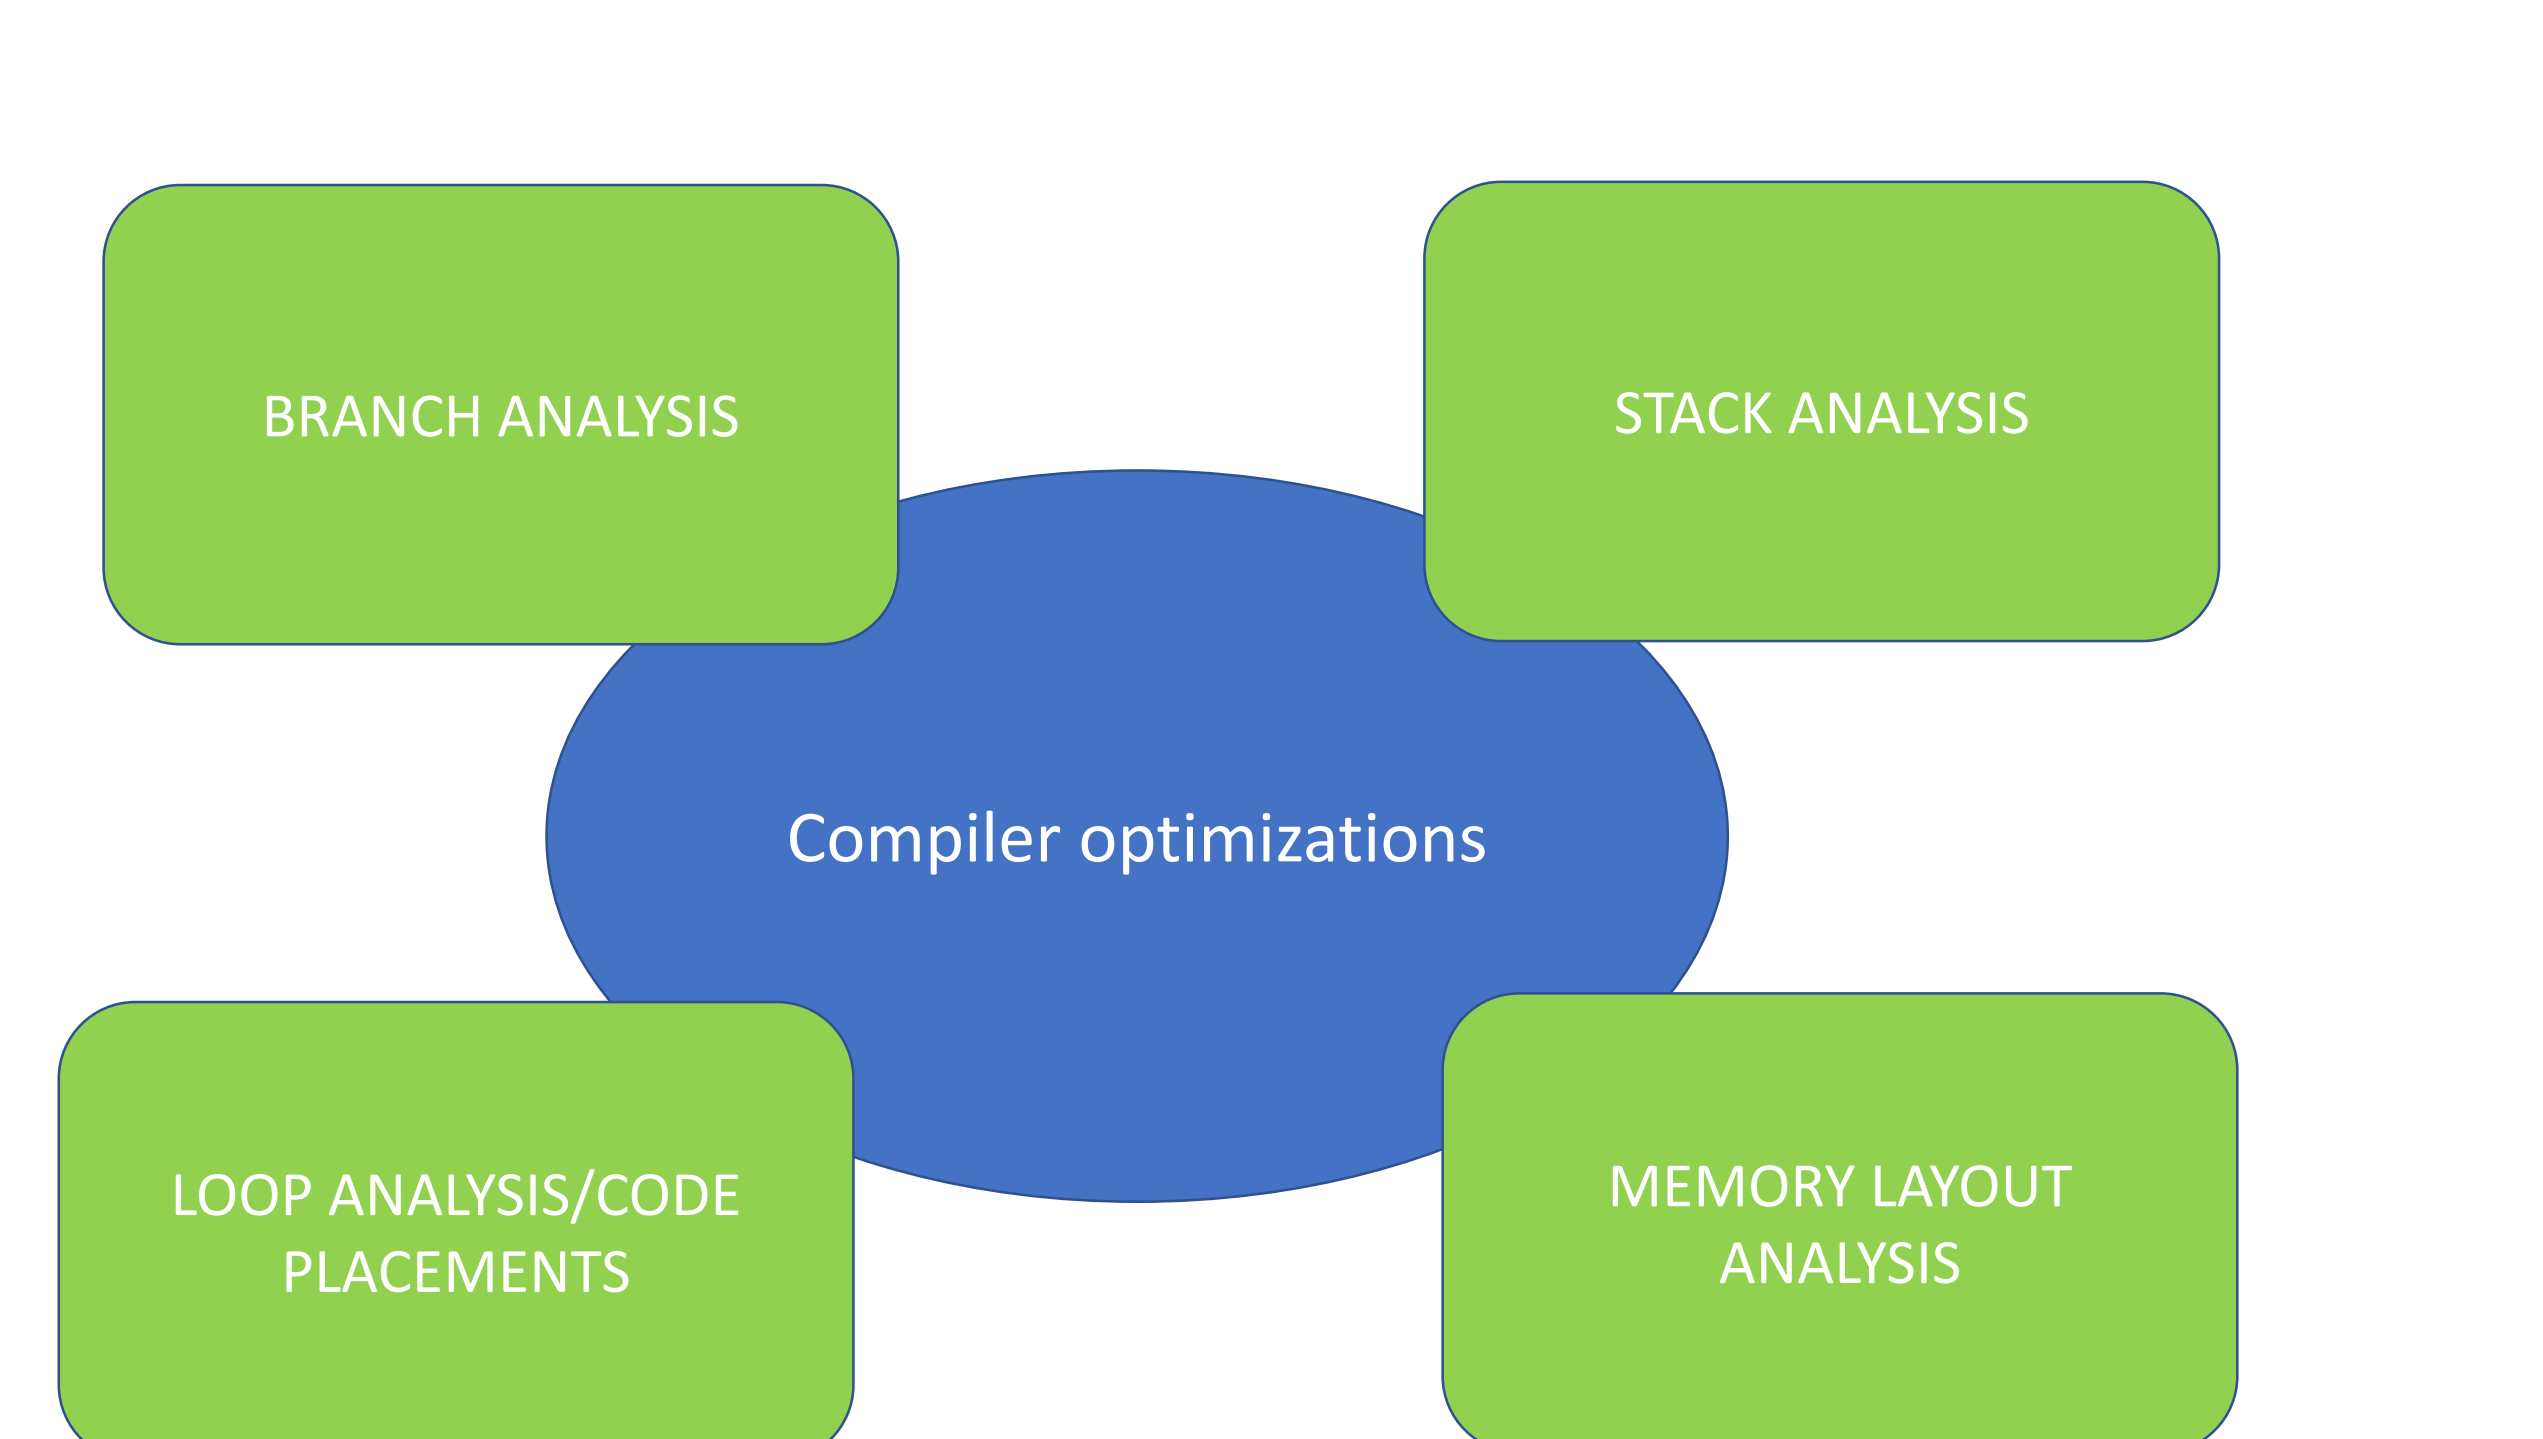
-> Multiple events as occurred on Instruction/code block execution can also be Superimposed on the dynamic Execution sequence.

-> On right we have a "perf record" Profile of the same which contains cycle event imposed on a static profile.

-> The static profile though useful does not give complete detail of running program.

Instructions trace vs Static profile

WHERE CAN TRACES PLAY A ROLE



BRANCH ANALYSIS

```
00000001013e04 0 addi r25,r21,84
00000001013e08 0 ne r1,r14
00000001013e0c 0 addi r3,r21,28
00000001013e10 0 std r3,208(r1)
00000001013e14 0 ld r3,488(r1)
00000001013e18 0 cmpld cr3,r,3,0
00000001013e1c 0 ld r3,504(r1)
00000001013e20 0 cmpld r17,r6
00000001013e24 0 bpe 0000000101374b4 <S_regmatch+0x3f4>
00000001013e28 0 lbz r0,r(17)
00000001013e2c 0 cmplw cr1,r,4,133
00000001013e30 0 bgt cr1,0000000101374c0 <S_regmatch+0xa00>
00000001013e34 0 cl_rldi r3,r,4,32
00000001013e38 0 rdictr r3,r3,2,81
00000001013e3c 0 lsw r3,r,3,8
00000001013e40 0 add r3,r,8
00000001013e44 0 mctr r3
00000001013e48 0 bctr r3
00000001013e4c 0 lsw r4,-15656(r2)
00000001013e50 0 lsw r3,r,15)
00000001013e54 0 cmpld r1,r,4,133
00000001013e58 0 bgt- 00000001013a814 <S_regmatch+0x3d54>
00000001013e5c 0 ld r3,216(r1)
00000001013e60 0 lsw r4,r(10)
00000001013e64 0 lsw r3,r(3)
00000001013e68 0 cmplw cr3,r,4)
00000001013e6c 0 bgt- 000000010138280 <S_regmatch+0x17c0>
00000001013e70 0 stw r3,112(r19)
00000001013e74 0 stw r3,28(r21)
00000001013e78 0 lsw r3,116(r19)
00000001013e7c 0 mctr r3
00000001013e80 0 rdictr r3,r,4,133
00000001013e84 0 ld r3,40(r21)
00000001013e88 0 lsw r4,484(r1)
00000001013e8c 0 rlwimi r3,r,4,12,20,20
00000001013e90 0 cmpld r16,0
00000001013e94 0 mctr r3,4
00000001013e98 0 crands 4*cr4gt,4*cr4gt,0,0,11
00000001013ea0 0 beq- 00000001013a7f8 <S_regmatch+0x3d38>
00000001013ea4 0 cmpld r3,255
00000001013ea8 0 ble- 00000001013d0a4 <S_regmatch+0x45e4>
00000001013eac 0 ld r4,8(r25)
00000001013ead 0 rdictr r3,r3,1,62
00000001013eae 0 lhwx r5,r4,r3
00000001013eaf 0 ld r3,544(r1)
00000001013eb0 0 cmpld r3,0
00000001013eb4 0 bne- 00000001013c23c <S_regmatch+0x577c>
00000001013eb8 0 lsw r4,133
```

not taken

indirect branches

taken branch

-> Trace can be used to derive critical information on branch intensive codes.

-> determine which branches are critical to performance.

-> for direct branches determine the % times it was taken vs non-taken.

-> above can be useful in optimal branch placements.

-> also evaluate the indirect branch call destination distribution.

CODE ANALYSIS

```
void quick_sort_(int *array, int l, int r)
{
  int partition_key, i, j;
  if (l <= r)
    return;
  i = l - 1;
  j = r;
  partition_key = array[l];
  while (1) {
    while (array[++i] < partition_key)
    while (partition_key < array[--j]) {
      if (i == j)
        break;
    }
    swap(array[i], array[j]);
    swap(array[l], array[i]);
    quick_sort_(array, l, i - 1);
    quick_sort_(array, i + 1, r);
  }
}

int tmp_array[10];
int size;
init();
for (i=0; i<10; i++) {
  // Let's read in an array and sort it
  read_array(tmp_array, size, i);
  quick_sort(tmp_array, size);
}
```

CALL NO	NUM_INSTR	
1		X
2		0.92X
3		X
4		0.90X
5		0.97X
6		0.93X
7		3X
8		0.92X
9		X
10		0.93X

-> Sample quicksort Program and the driver which calls it 10 times to sort an array.

-> Using statistical profiler it would be difficult to determine if a single iteration was slow.

-> Table derived from trace data tells us that call7 is the one of interest.

-> Events(Instruction) collected are as a ratio of event from 1st call.

STACK ANALYSIS

```
00000001013e04 <S_regmatch+0x344> 0 addi r25,r21,84
00000001013e08 <S_regmatch+0x348> 0 ne r1,r14
00000001013e0c <S_regmatch+0x34c> 0 addi r3,r21,28
00000001013e10 <S_regmatch+0x350> 0 std r3,216(r1)
00000001013e14 <S_regmatch+0x354> 0 ld r3,208(r1)
00000001013e18 <S_regmatch+0x358> 0 ld r3,488(r1)
00000001013e1c <S_regmatch+0x35c> 0 cmpld cr3,r,3,0
00000001013e20 <S_regmatch+0x360> 0 ld r3,504(r1)
00000001013e24 <S_regmatch+0x364> 0 cmpld r17,r6
00000001013e28 <S_regmatch+0x368> 0 bpe 0000000101374b4 <S_regmatch+0x3f4>
00000001013e2c <S_regmatch+0x370> 0 lbz r0,r(17)
00000001013e30 <S_regmatch+0x374> 0 cmplw cr1,r,4,133
00000001013e34 <S_regmatch+0x378> 0 bgt cr1,0000000101374c0 <S_regmatch+0xa00>
00000001013e38 <S_regmatch+0x37c> 0 cl_rldi r3,r,4,32
```

stack usage

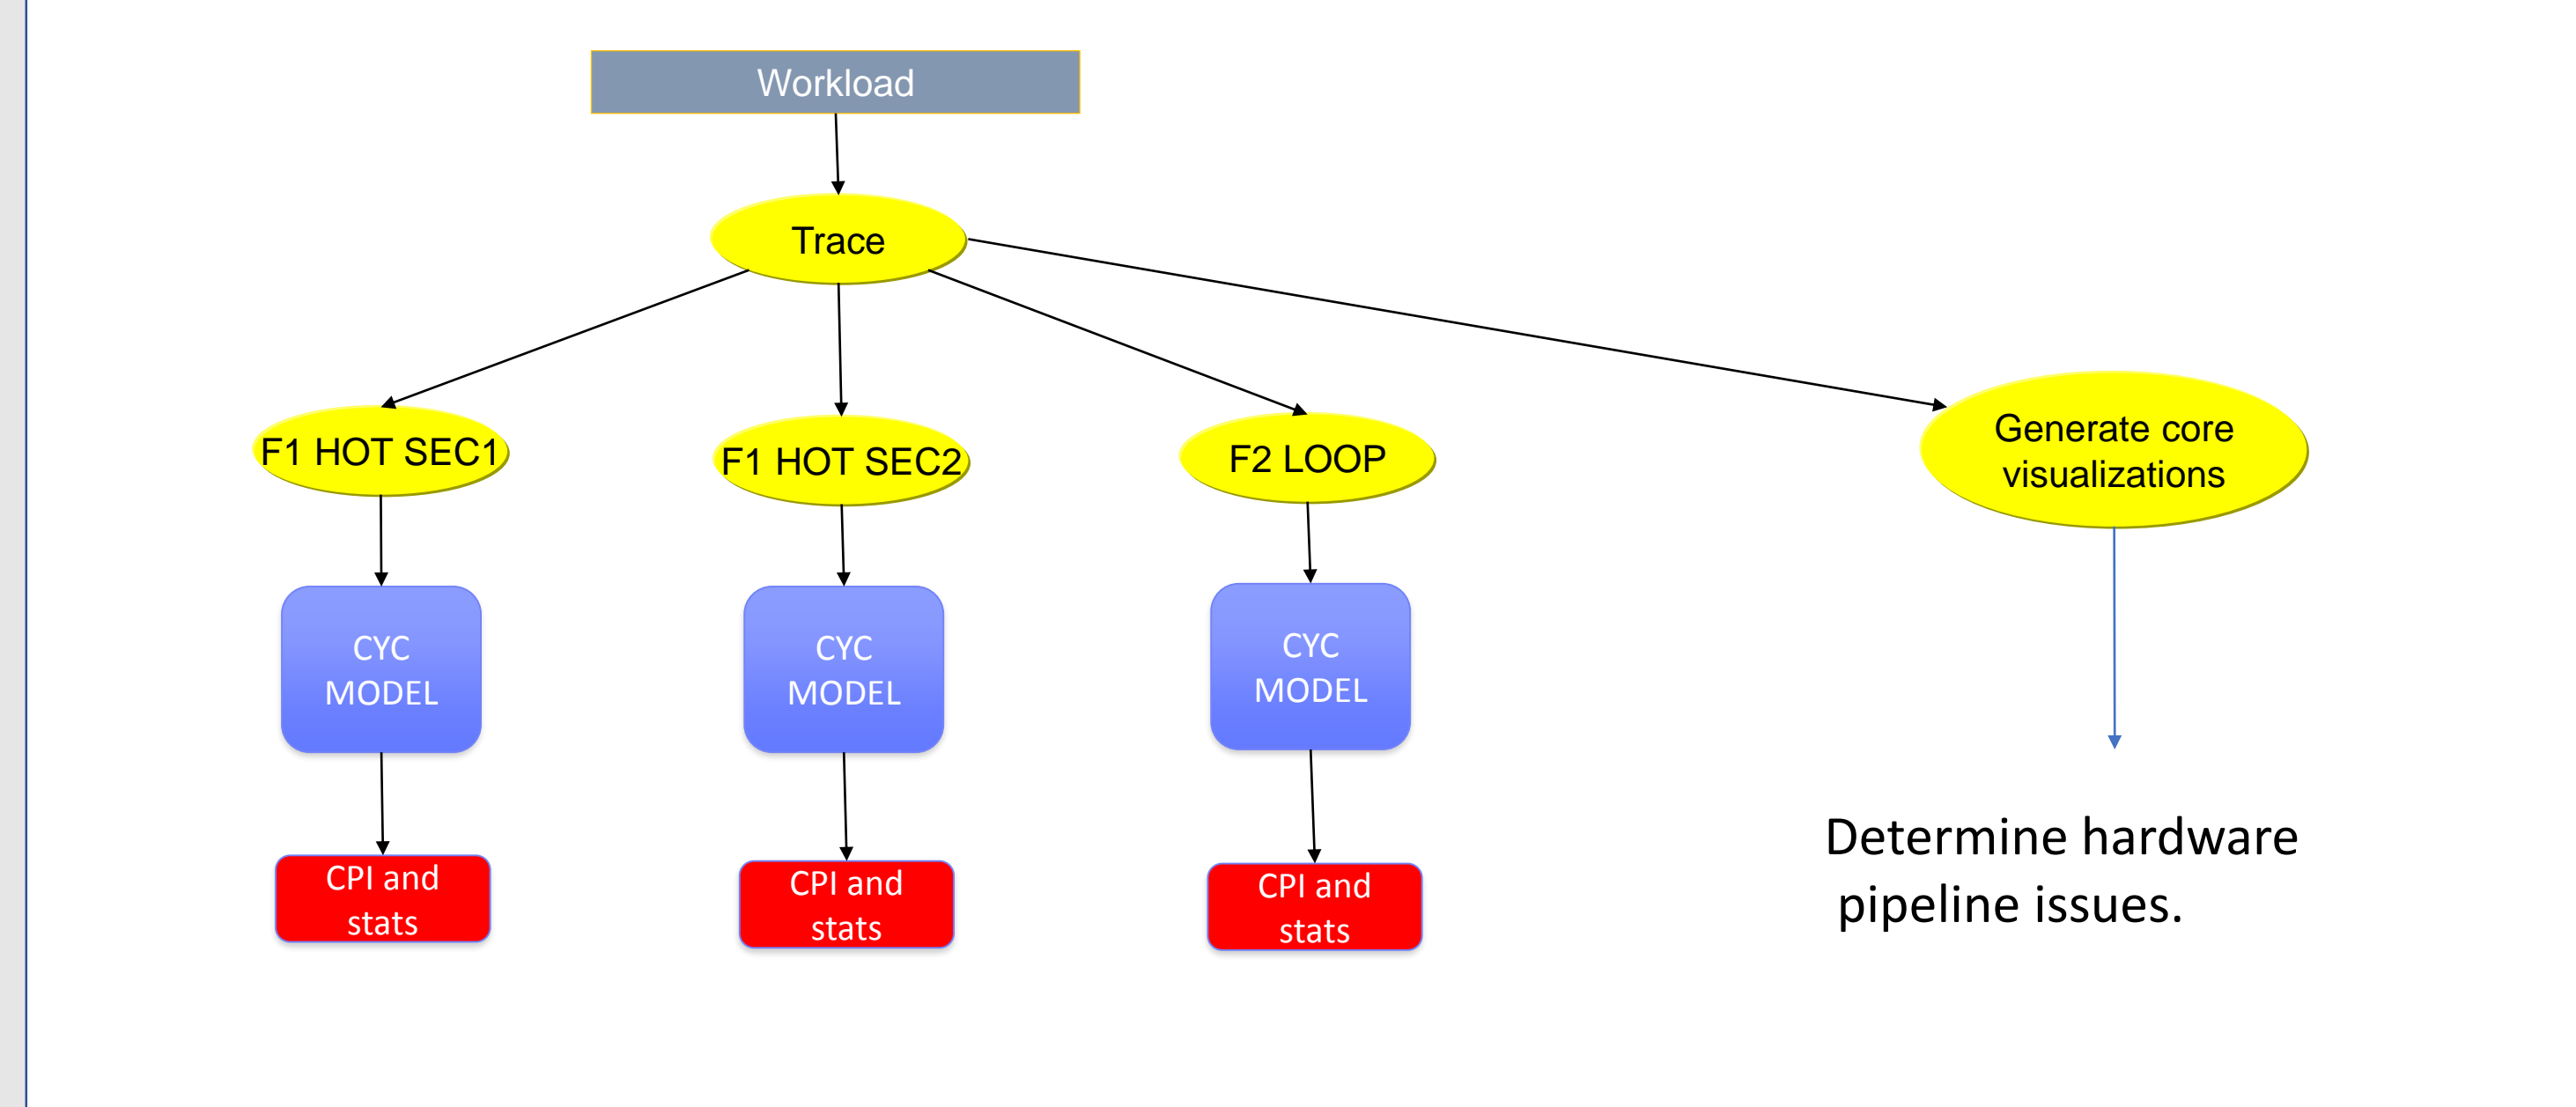
-> the image shows a push to and pop from the stack.

-> In a trace keep a count of stack accesses at different code levels.

-> with help of the instruction addresses which map to a particular function can be determined.

-> Generally useful in comparing two different compiler register usage and stack accesses.
-> Analyze what percentage of instructions and cycles are used as stack operations.
-> Stack usage in prologue/epilogue vs hot path.
-> Register spill analysis in critical paths.
*as in above example it can be useful to determine if stack usage shoots up in a particular call to the function/or a different execution path is taken.

BEYOND



Enable cycle and other event analysis at lower code levels, particularly the hot loops or code sections within a function. So the above quicksort example cycle impact can be evaluated per function call level. The traces can further be used to evaluate core pipeline visualizations and determine hardware issues which cannot be explained by just disassembly analysis.

Disclaimer: This poster is intended to represent the opinion of the author and not on any organization that the author is/was associated or affiliated with.