Call Paths for Pin Tools

Milind Chabbi, Xu Liu, and John Mellor-Crummey Department of Computer Science Rice University

> CGO'14, Orlando, FL February 17, 2014

What is a Call Path ?



Chain of function calls that led to the current point in the program. (a.k.a Calling Context / Call Stack / Backtrace / Activation Record)



What is a Call Path ?

Performance Analysis Tools



Chain of function calls that led to the current point in the program. (a.k.a Calling Context / Call Stack / Backtrace / Activation Record)



Fine-grained monitoring tools

Correctness

• Performance



Fine-grained monitoring tools

- Correctness
 - Data race detection

• Performance



Fine-grained monitoring tools

- Correctness
 - Data race detection

Attribute each conflicting access to its full call path

Performance



Fine-grained monitoring tools

- Correctness
 - Data race detection
 - Taint analysis
 - Array out of bound detection
- Performance

Fine-grained monitoring tools

- Correctness
 - Data race detection
 - Taint analysis
 - Array out of bound detection
- Performance
 - Reuse-distance analysis
 [Liu et al. ISPASS'13]

Attribute distance between "use" and "reuse" to references in full context

Fine-grained monitoring tools

- Correctness
 - Data race detection
 - Taint analysis
 - Array out of bound detection
- Performance
 - Reuse-distance analysis
 [Liu et al. ISPASS'13]



Attribute distance between "use" and "reuse" to references in full context

Fine-grained monitoring tools

- Correctness
 - Data race detection
 - Taint analysis
 - Array out of bound detection
- Performance
 - Reuse-distance analysis
 - Cache simulation
 - False sharing detection
 - Redundancy detection (e.g. dead writes)

Fine-grained monitoring tools

- Correctness
 - Data race detection
 - Taint analysis
 - Array out of bound detection
- Performance
 - Reuse-distance analysis
 - Cache simulation
 - False sharing detection
 - Redundancy detection (e.g. dead writes)
- Other tools
 - Debugging, testing, resiliency, replay, etc.

State-of-the-art in Collecting Ubiquitous Call Paths



"It will slow down execution by a factor of several thousand compared to native execution -- I'd guess -so you'll wind up with something that is unusably slow on anything except the smallest problems."



"If you tried to invoke Thread::getCallStack on every memory access there would be very serious performance problems ... your program would probably never reach main."



No support for collecting calling contexts



Milind Chabbi

Call Paths for Pin Tools

State-of-the-art in Collecting Ubiquitous Call Paths



"It will slow down execution by a factor of several thousand compared to native execution -- I'd guess -so you'll wind up with something that is unusably slow on anything except the smallest problems."



"If you tried to invoke Thread::getCallStack on every memory access there would be very serious performance problems ... your program would probably never reach main."



No support for collecting calling contexts

We built one ourselves—CCTLib

Milind Chabbi

Roadmap

CCTLib

- Ubiquitous call path collection
- Attributing costs to data objects
- Evaluation
- Conclusions

Roadmap

CCTLib



Ubiquitous call path collection

- Attributing costs to data objects
- Evaluation
- Conclusions



1 Overhead

2

3

RICE

1 Overhead

2 **Overhead**



RICE

1 Overhead

2 **Overhead**

3 Overhead



1 Overhead

(Space)

2 **Overhead**

3 Overhead



1 Overhead (Space)

2 Overhead (Time)

3 Overhead

| RICE |
|-------------|
|-------------|



2 Overhead (Time)

3 Overhead (Parallel scaling)

<u>Problem:</u> Deluge of call paths

Problem: Deluge of call paths





Problem: Deluge of call paths





Problem: Deluge of call paths





<u>Problem:</u> Deluge of call paths



Instruction stream

<u>Solution</u>

- Call paths share common prefix
- Store call paths as a calling context tree (CCT)
- One CCT per thread



Problem: Unwinding overhead



Problem: Unwinding overhead



RICE

Problem: Unwinding overhead

Solution: Reverse the process. Eagerly build a replica/shadow stack on-the-fly.



RICE

Problem: Unwinding overhead

Solution: Reverse the process. Eagerly build a replica/shadow stack on-the-fly.





Problem: Unwinding overhead

RICE

Solution: Reverse the process. Eagerly build a replica/shadow stack on-the-fly.



Problem: Unwinding overhead

RICE

Solution: Reverse the process. Eagerly build a replica/shadow stack on-the-fly.














Return to caller: Constant time update



Call Paths for Pin Tools

Milind Chabbi





Finding a callee from its caller involves a **lookup**



Finding a callee from its caller involves a **lookup**













Context Should Incorporate Instruction Pointer





Context Should Incorporate Instruction Pointer



Context Should Incorporate Instruction Pointer





A CCT node represents a *Pin trace*

- ◆ CCTLib maintains node → Pin trace mapping
- Each slot in a node represents an instruction in a *Pin trace*



- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





Milind Chabbi

- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





Milind Chabbi

- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





Milind Chabbi

- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query





- Problem: Mapping *IP* to *Slot* at runtime
 - Variable size x86 instructions
 - Non-sequential control flow
- Solution:
 - Pin's trace-instrumentation to hardwire *Slot#* as argument to context query routine for an IP
- **Result:**
 - Constant time to query



Roadmap

CCTLib

- Ubiquitous call path collection
- Attributing costs to data objects
- Evaluation
- Conclusions

```
int MyArray[SZ];
```

int main() {

if (...)

else

int * p;

Update(p);

```
int * Create() {
    return malloc(...);
```

```
void Update(int * ptr)
for( ... )
ptr[i]++;
```



- Associate each data access to its data object
- Data object

Dynamic allocation: Call path of allocation site

Static objects: Variable name

p = Create();

p = MyArray;

```
int MyArray[SZ];
int * Create() {
   return malloc(...);
void Update(int * ptr)
   for( ... )
      ptr[i]++;
int main() {
   int * p;
   if (....
      p = Create();
   else
      p = MyArray;
   Update(p);
```



- Associate each data access to its data object
- Data object
 - Dynamic allocation:
 Call path of allocation site
 - Static objects: Variable name

int MyArray[SZ];

int main() {

if (...)

else

int * p;

Update(p);

```
int * Create() {
    return malloc(...);
}
void Update(int * ptr)
    for( ... )
    ptr[i]++;
}
```



- Associate each data access to its data object
- Data object

Dynamic allocation: Call path of allocation site

Static objects: Variable name

p = Create();

p = MyArray;

int MyArray[SZ];

```
int * Create() {
   return malloc(...);
void Update(int * ptr)
   for( ... )
      ptr[i]++;
int main() {
   int * p;
   if (...)
      p = Create();
   else
```

```
p = MyArray;
Update(p);
```



- Associate each data access to its data object
- Data object
 - Dynamic allocation:
 Call path of allocation site
 - Static objects: Variable name

}

int MyArray[SZ];

```
int * Create() {
   return malloc(...);
}
void Update(int * ptr)
   for( ... )
   ptr[i]++;
}
```





- Associate each data access to its data object
- Data object

Dynamic allocation: Call path of allocation site

Static objects: Variable name

int MyArray[SZ];

int main() {

if (...)

else

int * p;

Update(p);

```
int * Create() {
    return malloc(...);
}
void Update(int * ptr)
    for( ... )
    ptr[i]++;
}
```



- Associate each data access to its data object
- Data object

Dynamic allocation: Call path of allocation site

Static objects: Variable name

}

p = Create();

p = MyArray;

int MyArray[SZ];

```
int * Create() {
   return malloc(...);
}
void Update(int * ptr)
   for( ... )
      ptr[i]++;
}
int main() {
   int * p;
   if (...)
```



- Associate each data access to its data object
- Data object

Dynamic allocation: Call path of allocation site

Static objects: Variable name

```
int MyArray[SZ];
int * Create() {
   return malloc(...);
void Update(int * ptr)
   for( ... )
     ptr[i]++;
int main() {
   int * p;
   if (...)
     p = Create();
   else
     p = MyArray;
  Update(p);
```



- Associate each data access to its data object
- Data object
 - Dynamic allocation:
 Call path of allocation site
 - Static objects: Variable name

Data-Centric Attribution

• How ?

- Record all <AddressRange, VariableName> tuples in a map
- Instrument all allocation/free routines and maintain
 <AddressRange, CallPath> tuples in the map
- At each memory access: search the map for the address

• Problems

- Searching the map on each access is expensive
- Map needs to be concurrent for threaded programs

Data-Centric Attribution using a Balanced Tree

- Observation:
 - Updates to the map are infrequent
 - Lookups in the maps are frequent
- Solution #1: sorted map
 - Keep <AddressRange, Object> in a balanced binary tree
 - Low memory cost—O(N)
 - Moderate lookup cost—O(log N)
 - Concurrent access is handled by a novel replicated tree data structure

Data-Centric Attribution using Shadow Memory

Solution #2: shadow memory

Application

CCTLib



Data-Centric Attribution using Shadow Memory

Solution #2: shadow memory

ObjA Application

CCTLib
• Solution #2: shadow memory







Solution #2: shadow memory







Solution #2: shadow memory



Solution #2: shadow memory



- For each memory cell, a shadow cell holds a handle for the memory cell's data object
 - Low lookup cost—O(1), high memory cost—
- $O(\sum_{i=1}^{N} sizeof(Obj(i)))$
- Shadow memory supports concurrent access
- CCTLib supports both solutions, clients can choose

Roadmap

CCTLib

- Ubiquitous call path collection
- Attributing costs to data objects
- Evaluation
- Conclusions



| Program | Running time in sec | |
|------------|------------------------|--|
| astr | 361 | |
| bzip2 | 161 | |
| gcc | 70 | |
| h264ref | 618 | |
| hmmer | 446 | |
| libquantum | 462 | |
| mcf | 320 | |
| omnetpp | 352 | |
| Xalan | 295 | |
| ROSE | 24 | |
| LAMMPS | 99 | |
| LULESH | 67 | |

Experimental setup:

- 2.2GHz Intel Sandy Bridge
- 128GB DDR3
- GNU 4.4.6 tool chain

Spec Int 2006 reference benchmark

| Program | Running time in sec | |
|------------|------------------------|--|
| astr | 361 | |
| bzip2 | 161 | |
| gcc | 70 | |
| h264ref | 618 | |
| hmmer | 446 | |
| libquantum | 462 | |
| mcf | 320 | |
| omnetpp | 352 | |
| Xalan | 295 | |
| ROSE | 24 | |
| LAMMPS | 99 | |
| LULESH | 67 | |

Experimental setup:

- 2.2GHz Intel Sandy Bridge
- 128GB DDR3
- GNU 4.4.6 tool chain

Spec Int 2006 reference benchmark

| Program | Running time in sec | |
|------------|------------------------|--|
| astr | 361 | |
| bzip2 | 161 | |
| gcc | 70 | |
| h264ref | 618 | |
| hmmer | 446 | |
| libquantum | 462 | |
| mcf | 320 | |
| omnetpp | 352 | |
| Xalan | 295 | |
| ROSE | 24 | |
| LAMMPS | 99 | |
| LULESH | 67 | |

Experimental setup:

- 2.2GHz Intel Sandy Bridge
- 128GB DDR3
- GNU 4.4.6 tool chain

Source-to-source compiler from LLNL

3M LOC compiling 70K LOC Deep call chains

Spec Int 2006 reference benchmark

| Program | Running time in sec |
|------------|------------------------|
| astr | 361 |
| bzip2 | 161 |
| gcc | 70 |
| h264ref | 618 |
| hmmer | 446 |
| libquantum | 462 |
| mcf | 320 |
| omnetpp | 352 |
| Xalan | 295 |
| ROSE | 24 |
| LAMMPS | 99 |
| LULESH | 67 |

Experimental setup:

- 2.2GHz Intel Sandy Bridge
- 128GB DDR3
- GNU 4.4.6 tool chain

Source-to-source compiler from LLNL

3M LOC compiling 70K LOC Deep call chains

Molecular dynamics code 500K LOC

Deep call chains Multithreaded

Spec Int 2006 reference benchmark

| Program | Running time in sec | |
|------------|------------------------|--|
| astr | 361 | |
| bzip2 | 161 | |
| gcc | 70 | |
| h264ref | 618 | |
| hmmer | 446 | |
| libquantum | 462 | |
| mcf | 320 | |
| omnetpp | 352 | |
| Xalan | 295 | |
| ROSE | 24 | |
| LAMMPS | 99 | |
| LULESH | 67 | |

Experimental setup:

- 2.2GHz Intel Sandy Bridge
- 128GB DDR3
- GNU 4.4.6 tool chain

Source-to-source compiler from LLNL

3M LOC compiling 70K LOC Deep call chains

Molecular dynamics code 500K LOC Deep call chains Multithreaded

Hydrodynamics mini-app from LLNL

Frequent data allocation and de-allocations Memory bound Multithreaded, Poor scaling

| | Call path collection |
|--|----------------------|
| Time overhead | |
| relative to original | 30x |
| (Null Pin tool) | |
| Time overhead | |
| relative to simple | 1 7v |
| instruction | |
| counting Pin tool | |
| <u>Memory overhead</u> relative to original | 1.8x |
| program | |

| | Call path collection |
|--|----------------------|
| <u>Time overhead</u> relative to original program (Null Pin tool) | 30x |
| Time overhead relative to simple instruction counting Pin tool | 1.7x |
| Memory overhead relative to original program | 1.8x |

| | Call path collection | Data-centric attribution | |
|---|-------------------------|--------------------------|---------------|
| | | Balanced Tree | Shadow Memory |
| <u>Time overhead</u> relative to simple instruction counting Pin tool | 1.7x | 4.5x | 2.3x |
| Memory overhead relative to original program | 1.8x | 2.0x | 11 x |

| | Call path collection | Data-centric attribution | |
|---|-------------------------|--------------------------|---------------|
| | | Balanced Tree | Shadow Memory |
| <u>Time overhead</u> relative to simple instruction counting Pin tool | 1.7x | 4.5x | 2.3x |
| <u>Memory overhead</u> relative to original program | 1.8x | 2.0x | 11 x |



CCTLib Scales to Multiple Threads



Conclusions

- Many tools can benefit from attributing metrics to full calling contexts and/or data objects
- Ubiquitous calling context collection was previously considered prohibitively expensive
- Fine-grain attribution of metrics to calling contexts and data objects is practical
- Full-precision call path collection and data-centric attribution require only modest space and time overhead
 - Choice of algorithms and data structures was a key to success

Conclusions



- Many tools can benefit from attributing metrics to full calling contexts and/or data objects
- Ubiquitous calling context collection was previously considered prohibitively expensive
- Fine-grain attribution of metrics to calling contexts and data objects is practical
- Full-precision call path collection and data-centric attribution require only modest space and time overhead
 - Choice of algorithms and data structures was a key to success

Conclusions



- Many tools can benefit from attributing metrics to full calling contexts and/or data objects
- Ubiquitous calling context collection was previously considered prohibitively expensive
- Fine-grain attribution of metrics to calling contexts and data objects is practical
- Full-precision call path collection and data-centric attribution require only modest space and time overhead
 - Choice of algorithms and data structures was a key to success

http://code.google.com/p/cctlib/



Other Complications in Real Programs

- Complex control flow
 - Signal handling
 - Setjmp-Longjmp
 - C++ exceptions (try-catch)
- Thread creation and destruction
 - Maintaining parent-child relationships between threads
 - Scalability to large number of threads