



# Pin: building customized program analysis tools with dynamic instrumentation

C.-K. Luk & a bunch of other Intel guys

*presented by Bjoern Doebel*

## Static

GCC, (SP)Lint, Coverity, ...

- Compile-time
- Heavy-weight
- No environmental information



## Dynamic

Valgrind, Pin, DynamoRIO, ...

- Runtime
- Trade-off overhead vs. realistic observations
- Path coverage issues

---

## Source-level

Lint, Coverity

- Exact information
- Problem: 3<sup>rd</sup> party tools

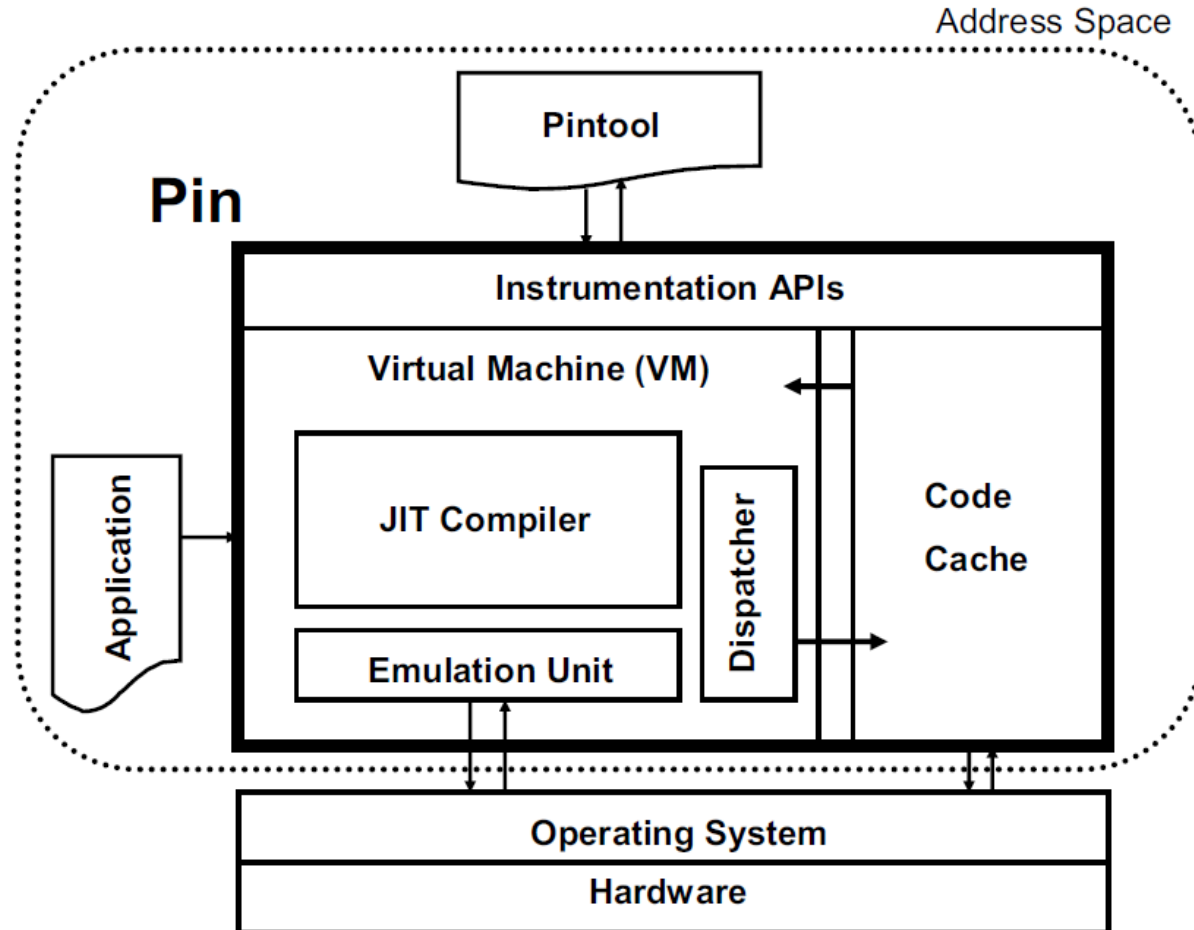


## Binary analysis

Bitblaze/Vine, Valgrind, Pin, ...

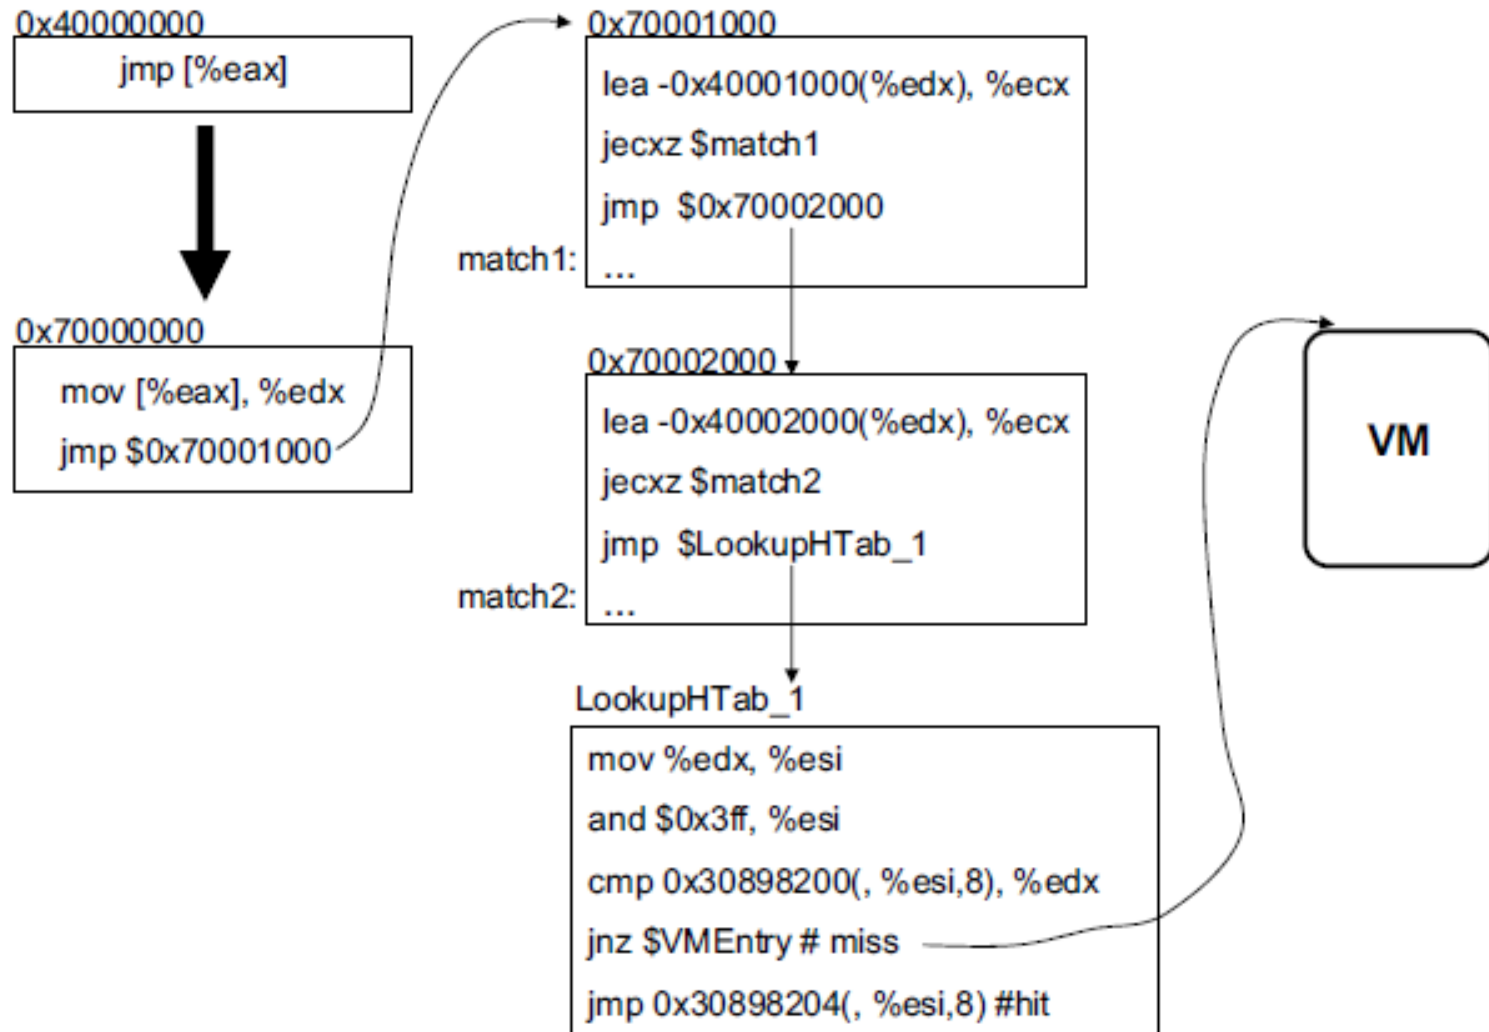
- Inexact w.r.t. source code
- Support for any kind of application / library

- Approach:
  - Disassemble binary (→ intermediate language)
  - Insert tool-specific instrumentation at IL level
  - Recompile into machine code
- Usually JIT-based
  - Disassemble & Resynthesize
  - Copy & Annotate
- Main focus of research: runtime optimizations

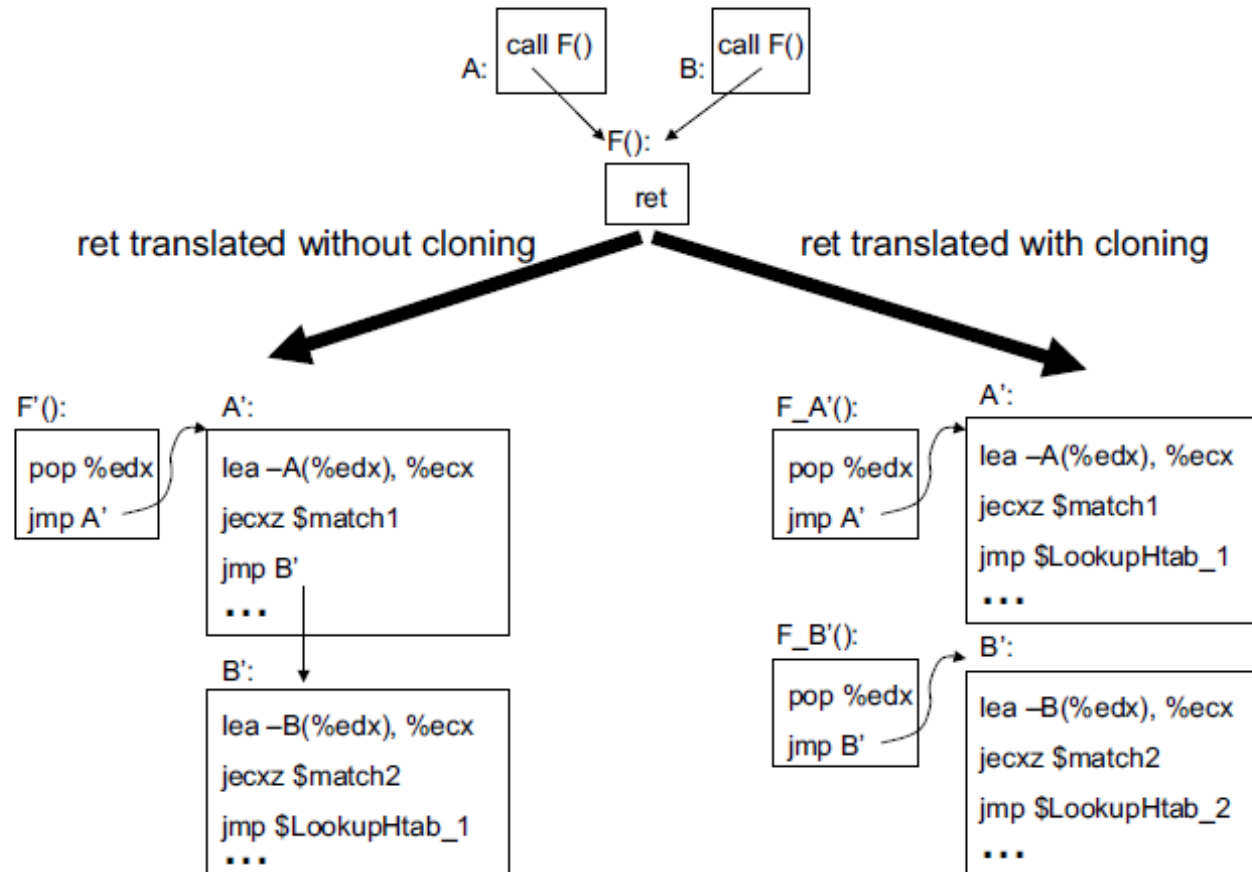


**Figure 2.** Pin's software architecture

- No intermediate language
- Trace-based recompilation
- Can attach to running program



## (b) Using cloning to help predict return targets



**Figure 3.** Compiling indirect jumps and returns

## (b) Valgrind's approach


### Trace 1

```
mov $1, %eax  
mov $2, %esi  
cmp %ecx, %edx  
mov %eax, EAX  
mov %esi, EBX  
jz t'
```

### Trace 2

**t':**

```
mov EAX, %eax  
mov EBX, %edi  
add $1, %eax  
sub $2, %edi  
...
```



**t:**

```
mov $1, %eax  
mov $2, %ebx  
cmp %ecx, %edx  
jz t  
...  
add $1, %eax  
sub $2, %ebx  
...
```



```

mov $1, %eax
mov $2, %ebx
cmp %ecx, %edx
jz t
...
t: add $1, %eax
   sub $2, %ebx
   ...

```

## (c) Pin (no reconciliation needed)

### Trace 1

```

mov $1, %eax
mov $2, %esi
cmp %ecx, %edx
jz t'

```

*Compile Trace 2 using the bindings:*

Virtual	Physical
%eax	%eax
%ebx	%esi
%ecx	%ecx
%edx	%edx

**t'**:

### Trace 2

```

add $1, %eax
sub $2, %esi
...

```

```

mov $1, %eax
mov $2, %ebx
cmp %ecx, %edx
jz t
...
t: add $1, %eax
   sub $2, %ebx
   ...

```

## (d) Pin (minimal reconciliation needed)

### Trace 1 (being compiled)

```

mov $1, %eax
mov $2, %esi
cmp %ecx, %edx
mov %esi, EBX
mov EBX, %edi
jz t'

```


### Trace 2 (previously compiled)

```

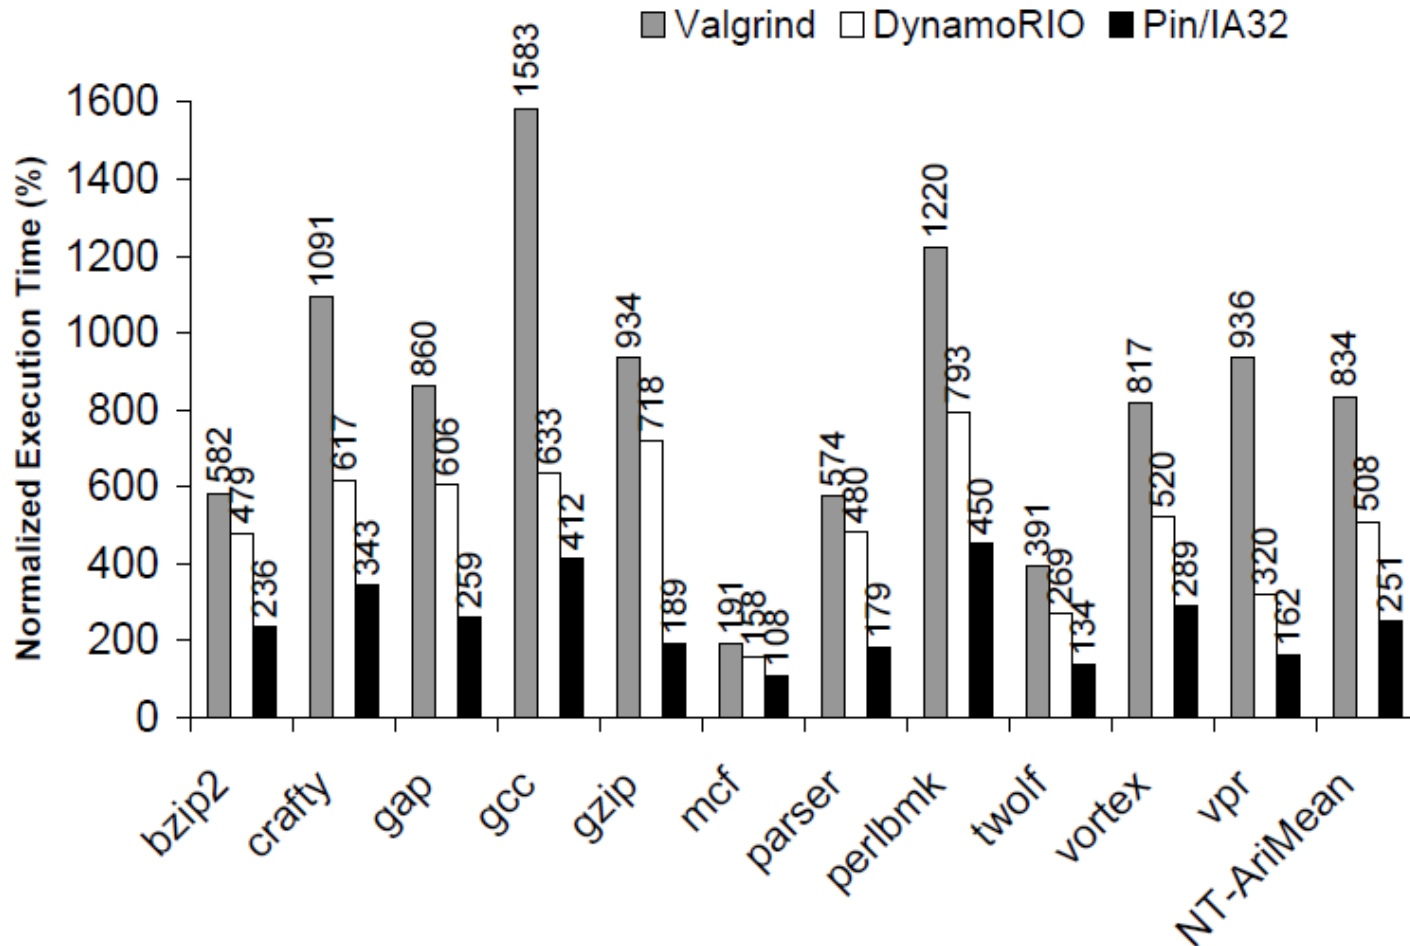
t': add $1, %eax
    sub $2, %edi
    ...

```

No need to recompile  
Trace 2, simply reconcile  
the bindings of virtual  
%ebx in Traces 1 and 2



(b) With basic-block counting



- Is there any advantage from having no intermediate language?
  - In fact, the compiler has some kind of intermediate representation, even if it is no language.
- Comparing Pin with Valgrind: How would it perform, if Pin were required to provide
  - Shadow values
  - Address space management
  - System call interception

- “[..] We implemented basic-block counting by modifying a tool in the Valgrind package named *lackey* [..]” [1]
  - “*Lackey* is a simple Valgrind tool that does various kinds of basic program measurement. It **adds quite a lot of simple instrumentation** to the program's code. It is primarily **intended to be of use as an example tool**, and consequently **emphasises clarity of implementation over performance.**” [2]

[1] The paper

[2] <http://valgrind.org/docs/manual/lk-manual.html>