

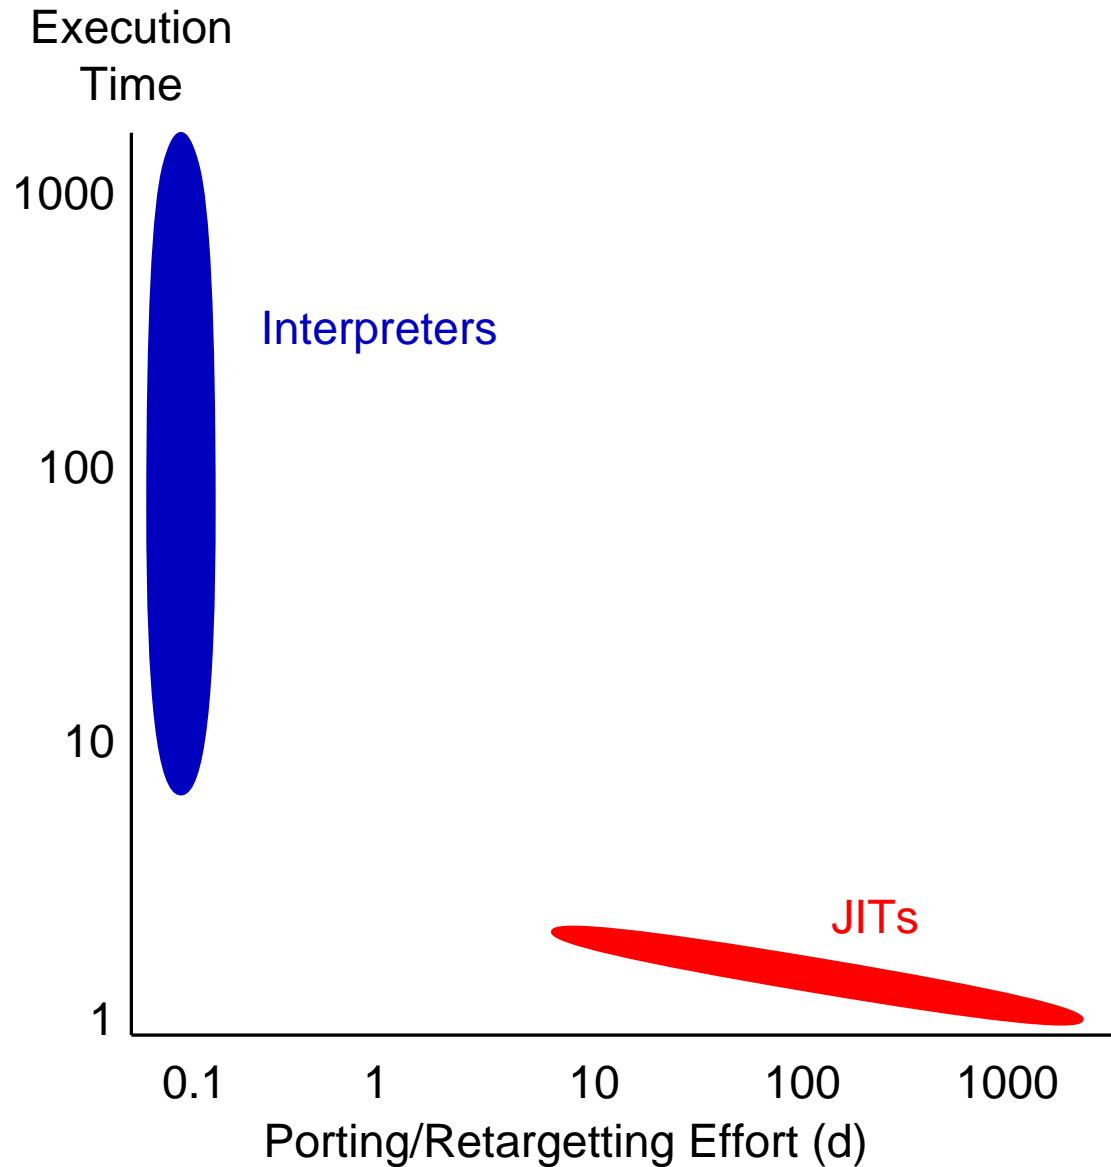
Superinstructions and Replication in the Cacao JVM interpreter

M. Anton Ertl

Christian Thalinger
TU Wien

Andreas Krall


Why interpreters?



Architecture	Mono	Cacao
Alpha	interp.	JIT
AMD64	JIT	JIT
ARM	JIT	JIT
HP-PA	interp.	
IA32	JIT	JIT
IA64	JIT	
MIPS		JIT
MIPS64		JIT
PowerPC	JIT	JIT
PowerPC64		interp.
s390	JIT	
s390x	JIT	
SPARC		
SPARC64		

Threaded Code

VM Code



VM instruction routines

Machine code for imul
Dispatch next instruction

Machine code for iadd
Dispatch next instruction


Dynamic Superinstructions

data segment

dyn. superinst code

Machine code for iload
Machine code for iload
Machine code for isub
Machine code for istore
...
Dispatch next

data segment
threaded VM Code



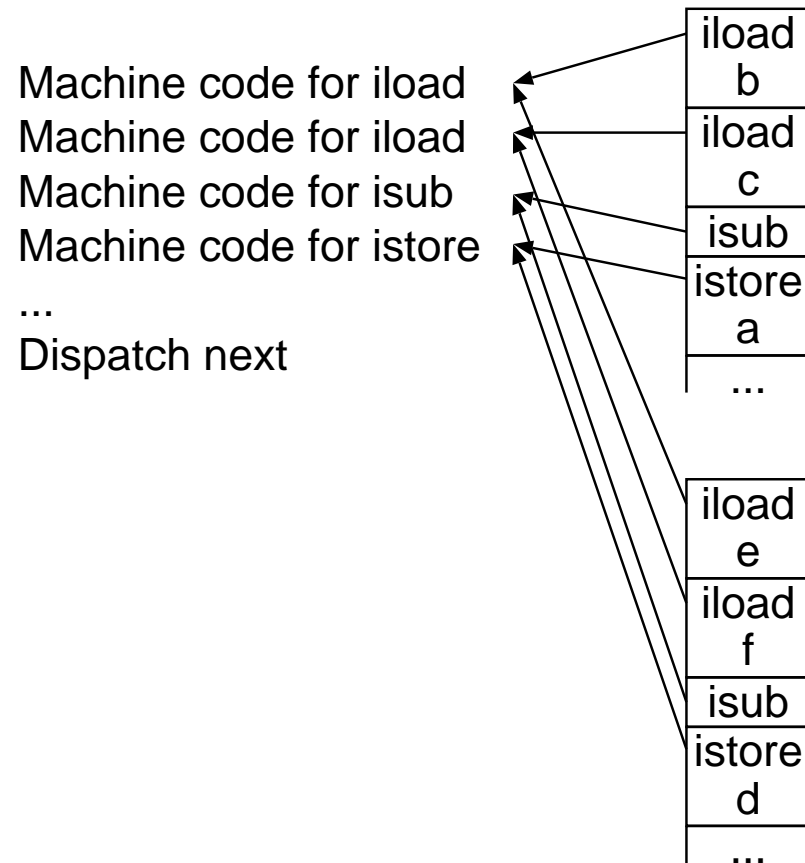
code segment

VM routine template

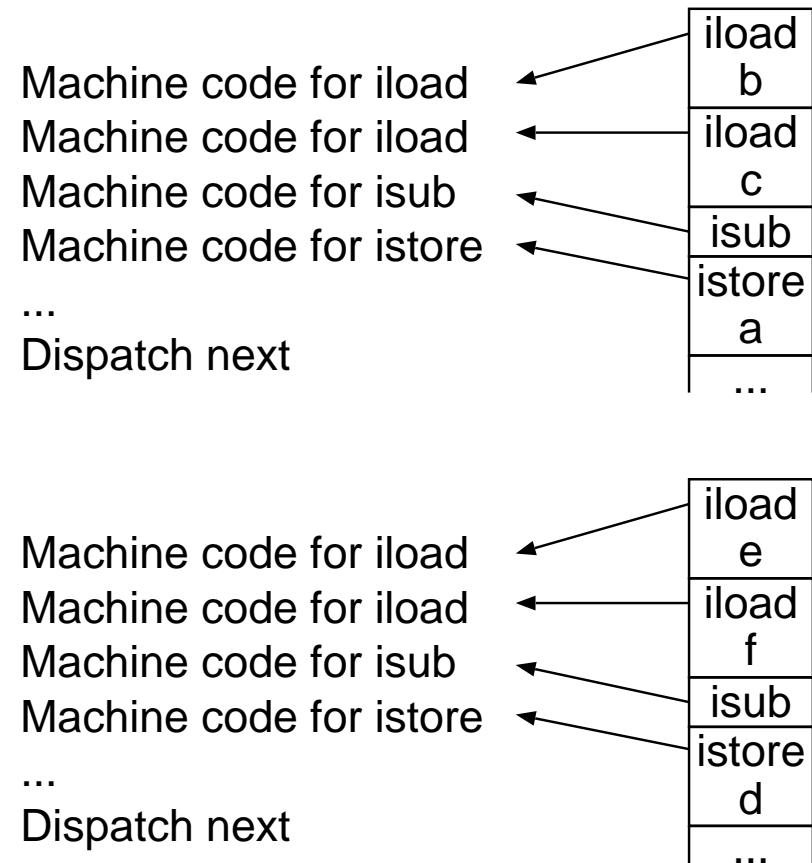
Machine code for isub
Dispatch next
Machine code for iload
Dispatch next
Machine code for istore
Dispatch next

Replication

No Replication



Replication



- + Increases BTB prediction accuracy
- + Simpler
- Increases code size

JVM and .NET problems

- Quickening
- Potential exception-throwing instructions
- How much benefit?

Quickable Instructions

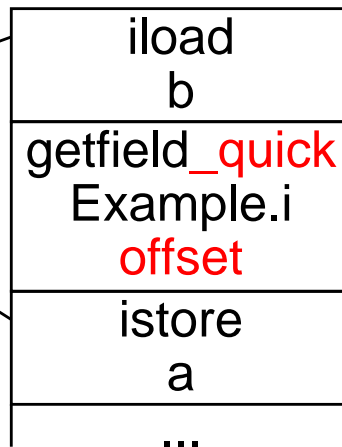
ACONST	INVOKESPECIAL
ARRAYCHECKCAST	INVOKESTATIC
CHECKCAST	INVOKEVIRTUAL
GETFIELD_CELL	MULTIANEWARRAY
GETFIELD_INT	NATIVECALL
GETFIELD_LONG	PUTFIELD_CELL
GETSTATIC_CELL	PUTFIELD_INT
GETSTATIC_INT	PUTFIELD_LONG
GETSTATIC_LONG	PUTSTATIC_CELL
INSTANCEOF	PUTSTATIC_INT
INVOKEINTERFACE	PUTSTATIC_LONG

Simple Solution

data segment
dyn. superinst code

Machine code for iload
Dispatch next
Machine code for istore
...
Dispatch next


data segment
threaded VM Code



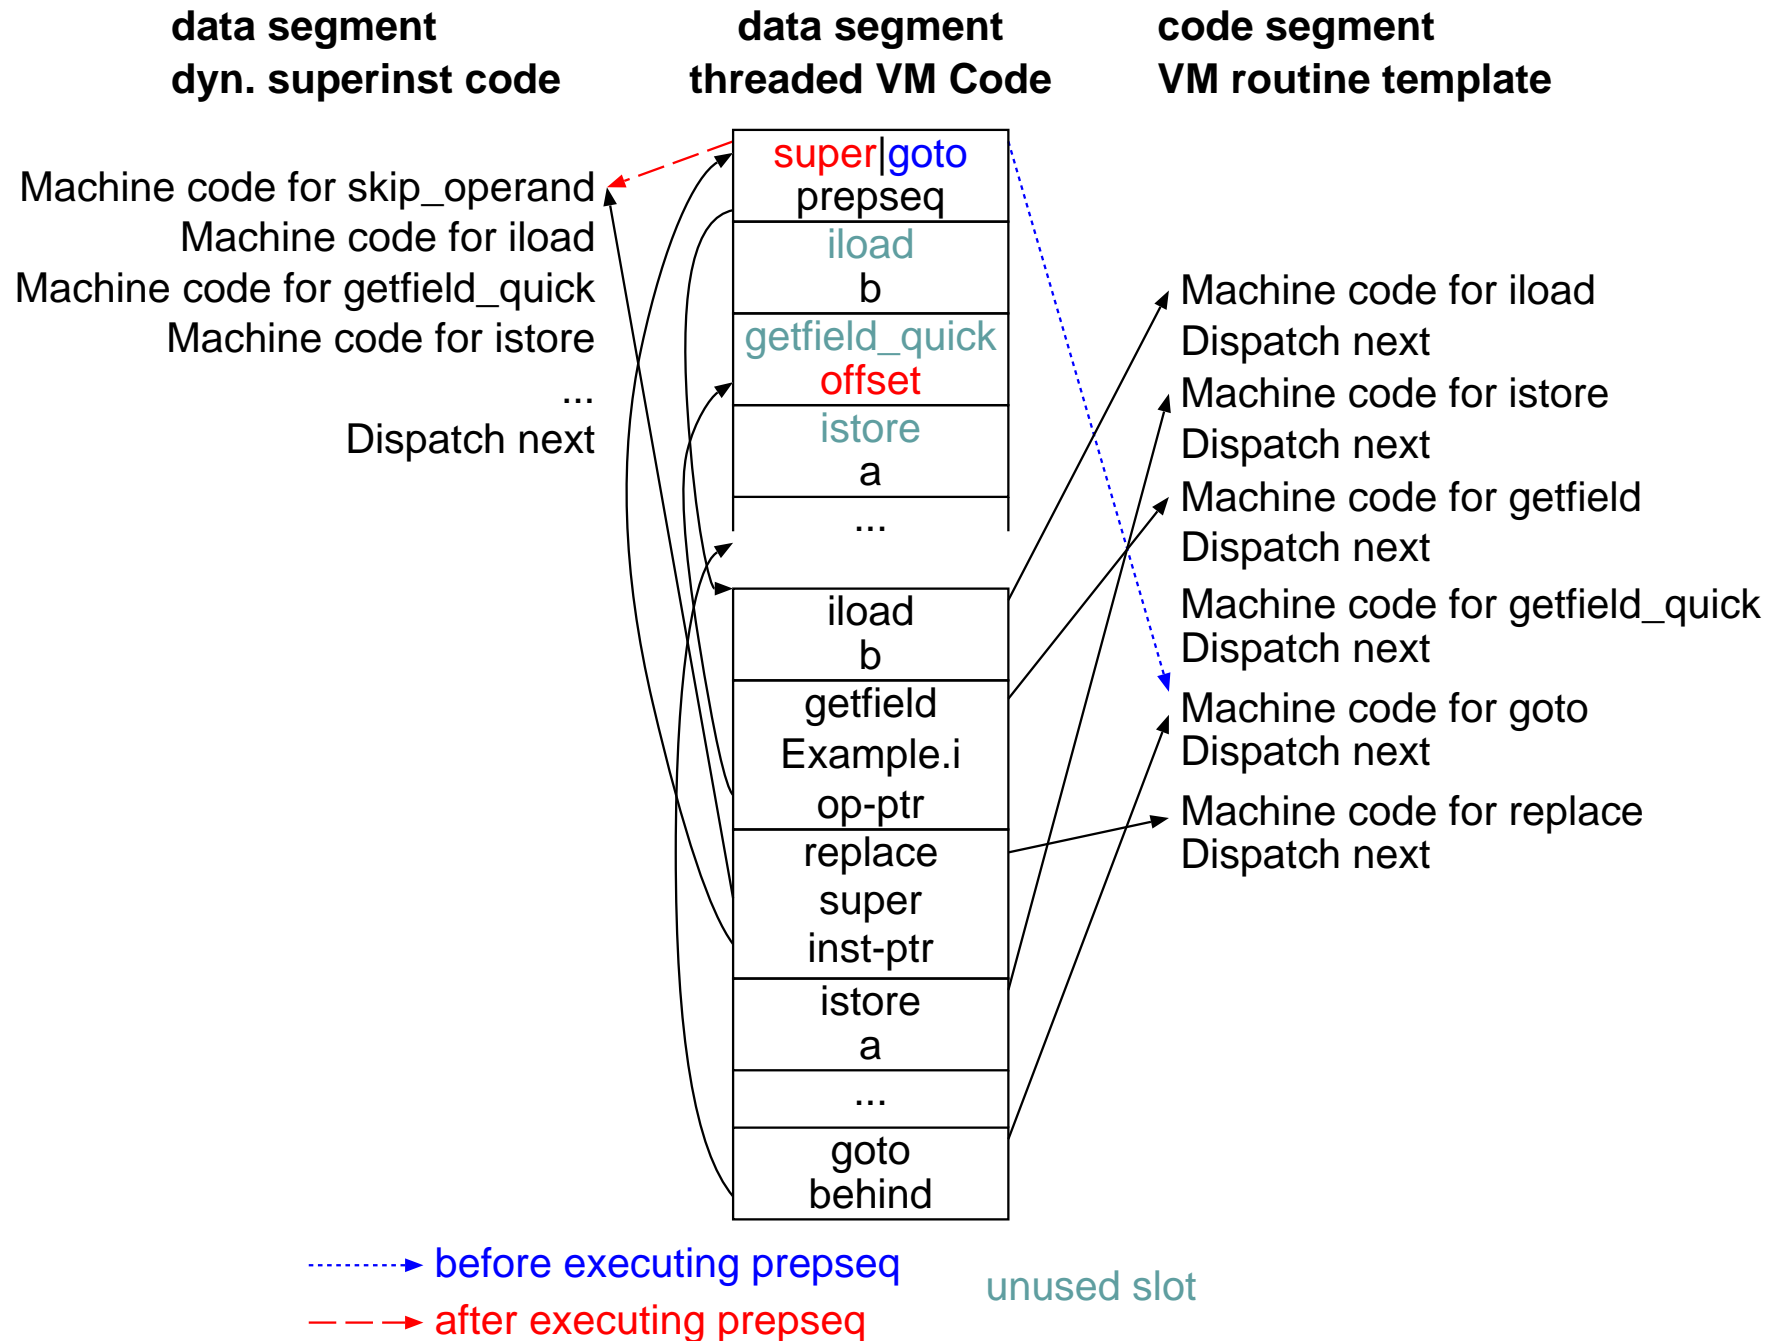
code segment
VM routine template

Machine code for iload
Dispatch next
Machine code for istore
Dispatch next
Machine code for getfield
Dispatch next
Machine code for getfield_quick
Dispatch next

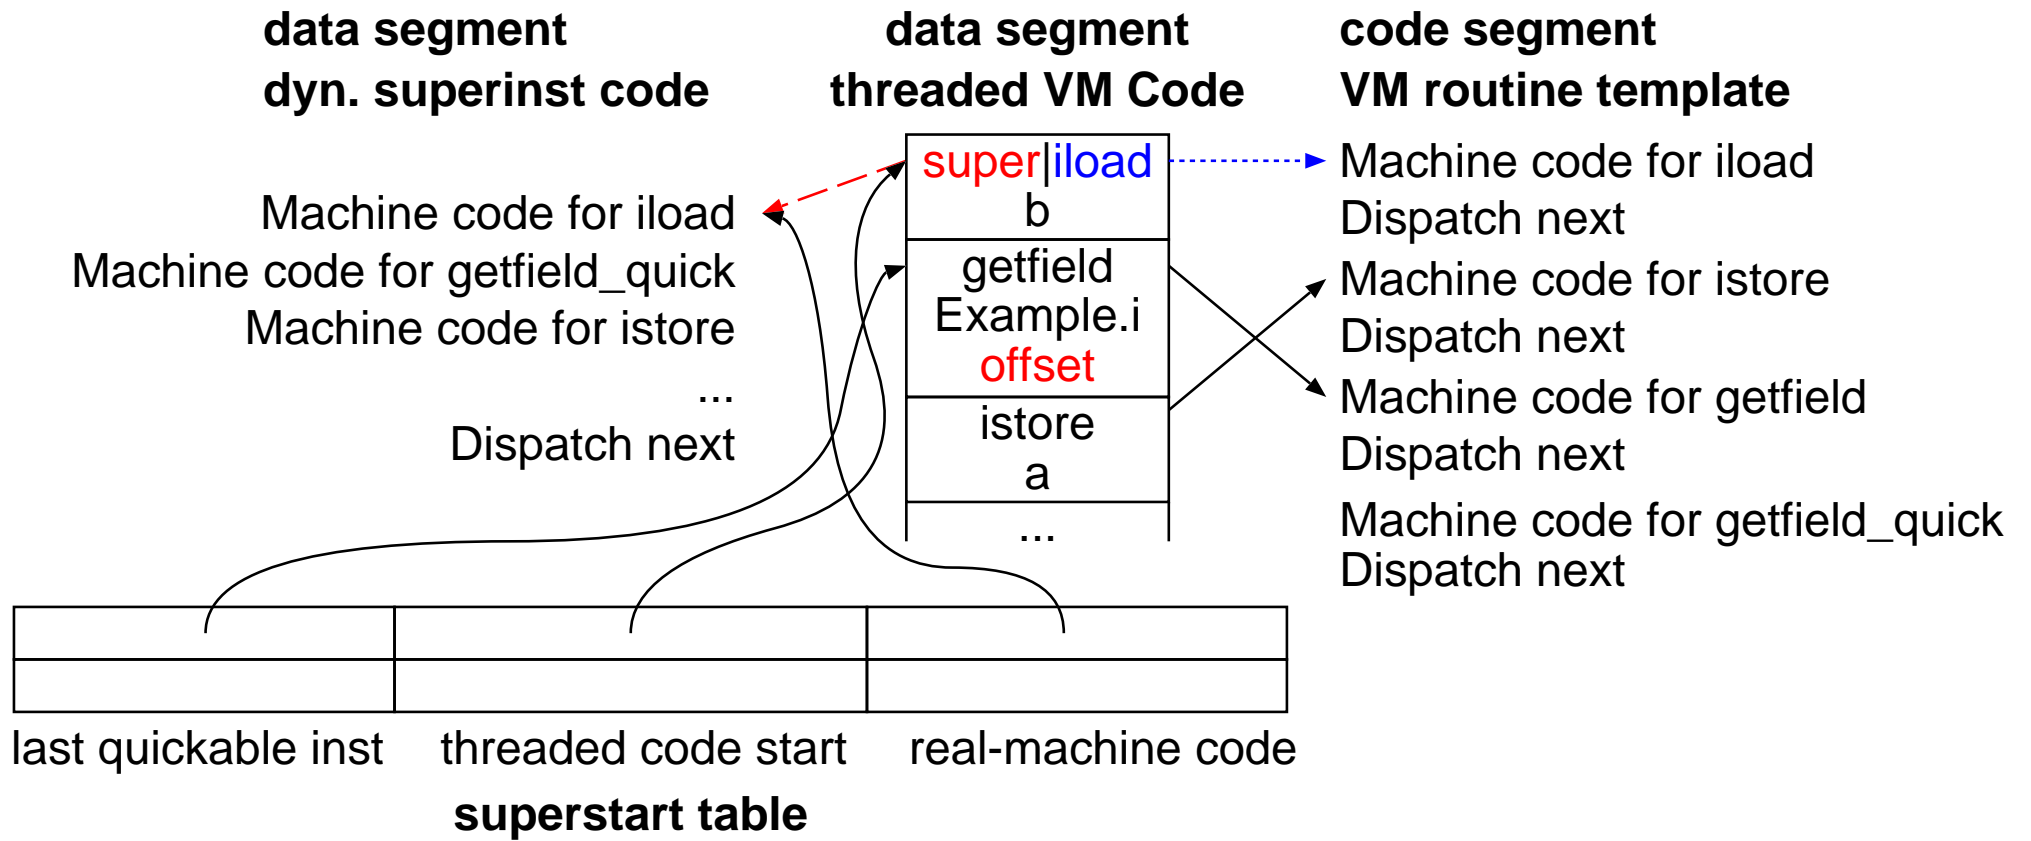
.....> before executing getfield
---> after executing getfield



SableVM's Sophisticated Solution



Cacao's Sophisticated Solution



-----> before executing `getfield`

-- --> after executing `getfield`

Potential Exception-Throwing Instructions

ILOAD	GETFIELD_CELL
LALOAD	GETFIELD_INT
AALOAD	GETFIELD_LONG
BALOAD	PUTFIELD_CELL
CALOAD	PUTFIELD_INT
SALOAD	PUTFIELD_LONG
IASTORE	INVOKEVIRTUAL
LASTORE	INVOKESPECIAL
BASTORE	INVOKEINTERFACE
CASTORE	ARRAYLENGTH
IDIV	CHECKNULL
IREM	

Problem and Solution

getfield_cell:

mov (%edi),%eax

add \$0x8,%edi

test %ebp,%ebp

je throw

add \$0x4,%edi

mov (%eax,%ebp,1),%ebp

jmp *-4(%edi)

getfield_cell:

mov (%edi),%eax

add \$0x8,%edi

test %ebp,%ebp

jne no_throw

jmp *0x2a0(%esp)

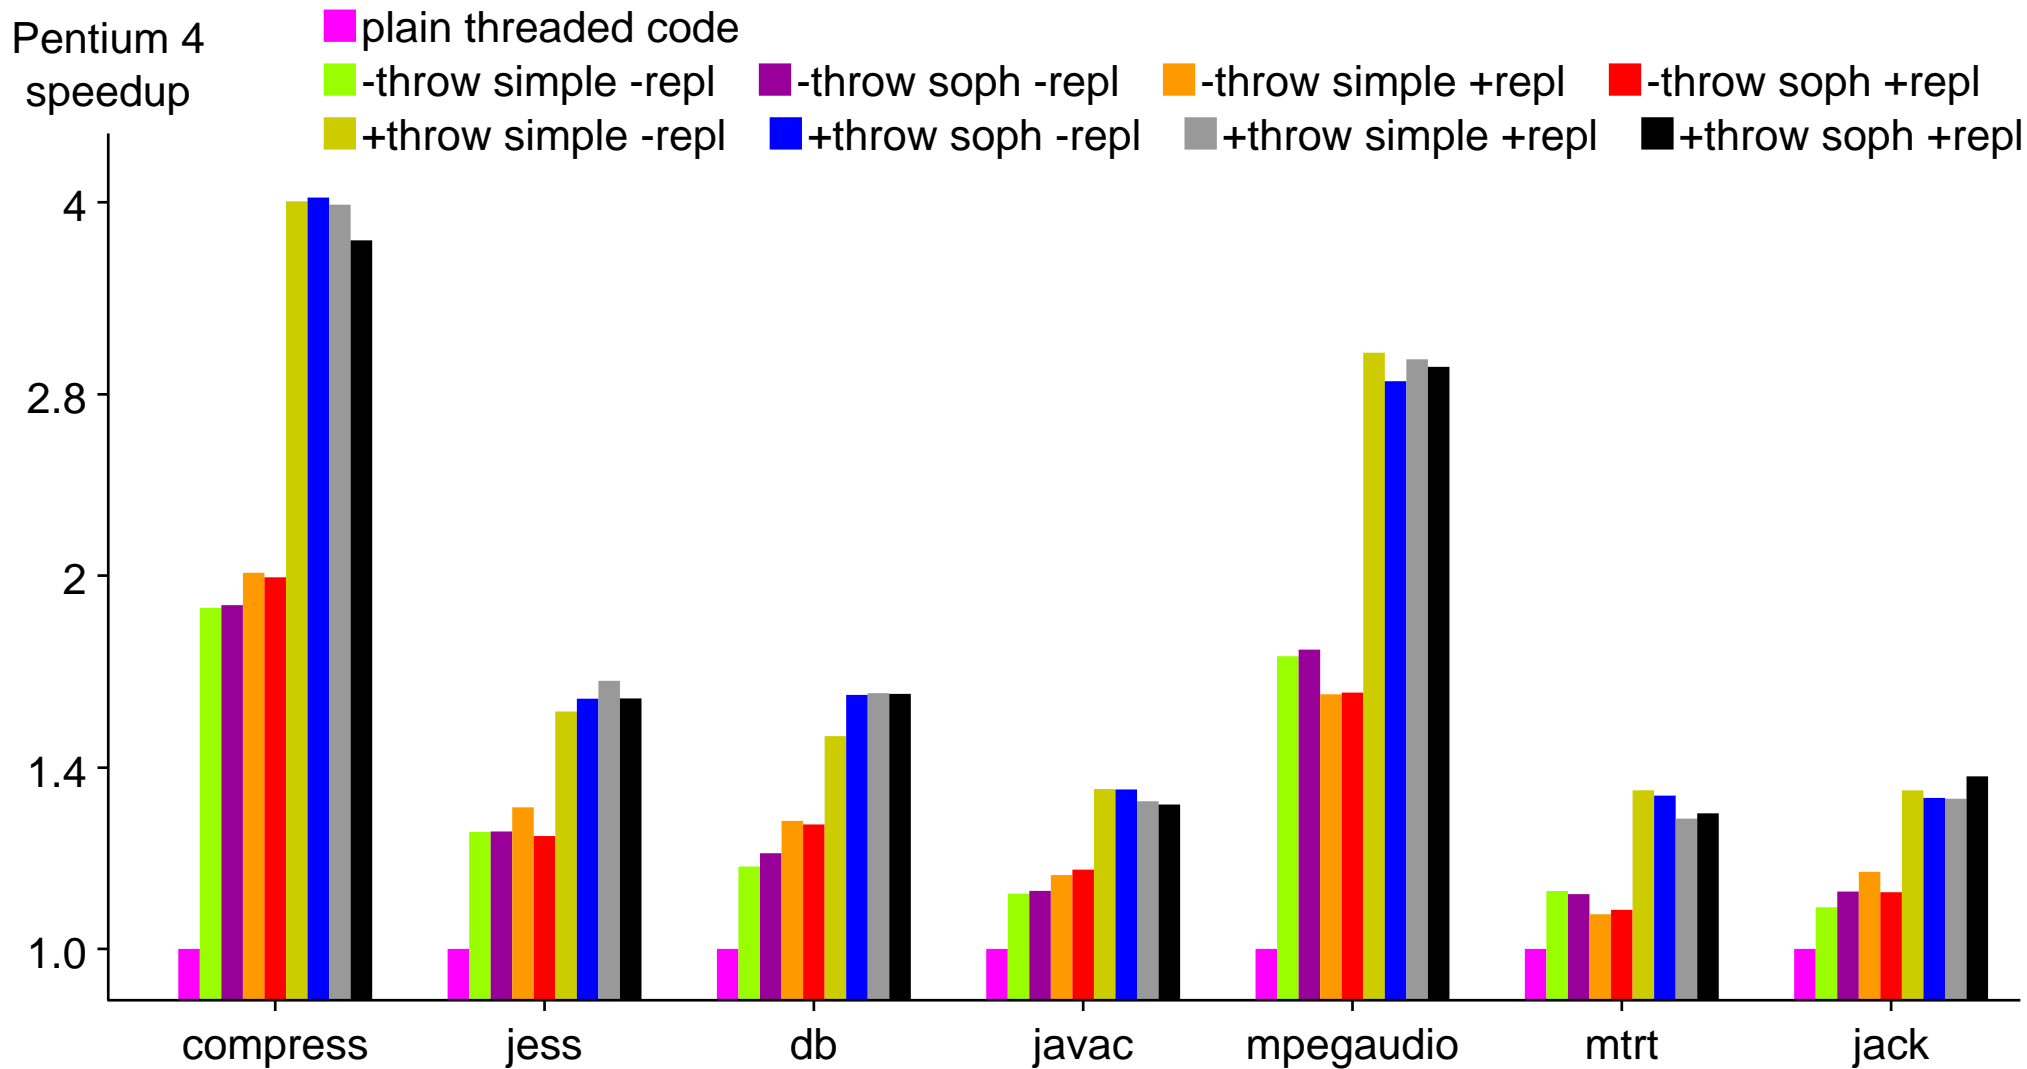
no_throw:

add \$0x4,%edi


mov (%eax,%ebp,1),%ebp

jmp *-4(%edi)

Speedup over plain threaded code



Speedup of various JVMs over Cacao with Superinstructions



Conclusion

- Superinstructions can provide big speedups
- Replication has little impact
- Quickening:
New sophisticated solution
but simple solution performs well in JIT setting
- Relocatability of throwing VM instructions:
Big performance impact
Solution: replace relative with indirect jumps