

# Hyperchaining Optimizations for an LLVM-Based Binary Translator on x86-64 and RISC-V Platforms

---

Jyun-Kai Lai and Wu Yang  
Department of Computer Science,  
National Yang Ming Chiao Tung University, Taiwan

# Outline

---

- Background
- Hyperchaining
- Experiments
- Conclusion

# Outline

---

- **Background**
- Hyperchaining
- Experiments
- Conclusion

# Binary translation

---

- Translate from guest (input) binary to host (output) binary and the translated program runs on a host machine
- **Rabbit** - an LLVM-based binary translator developed by our lab
  - Our work is based on Rabbit and translate from RISC-V binary to x86-64

# Architecture

---

- Frontend (Guest → LLVM IR)
  - Translate from guest binary to LLVM IR basic blocks
  - Incorporate one or more instructions into a LLVM IR basic block depend on SBT or DBT
- Backend (LLVM IR → Host)
  - Translate from LLVM IR basic blocks to binary chunks

# Translation

---

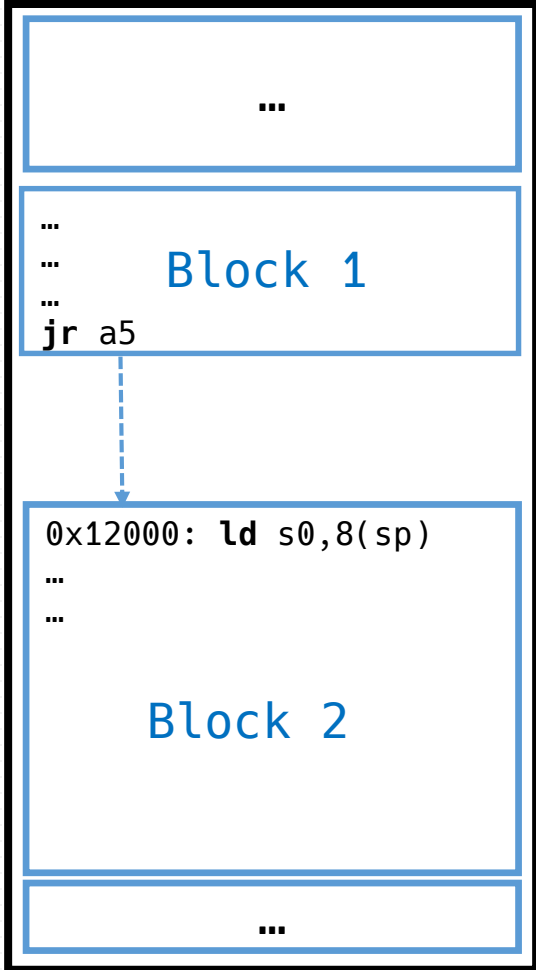
- Static Binary Translation (SBT)
  - Translate code blocks before run-time
- Dynamic Binary Translation (DBT)
  - Translate code blocks in run-time
- Hybrid Binary Translation (HBT; adopted by Rabbit)
  - SBT + DBT

# Translated code block

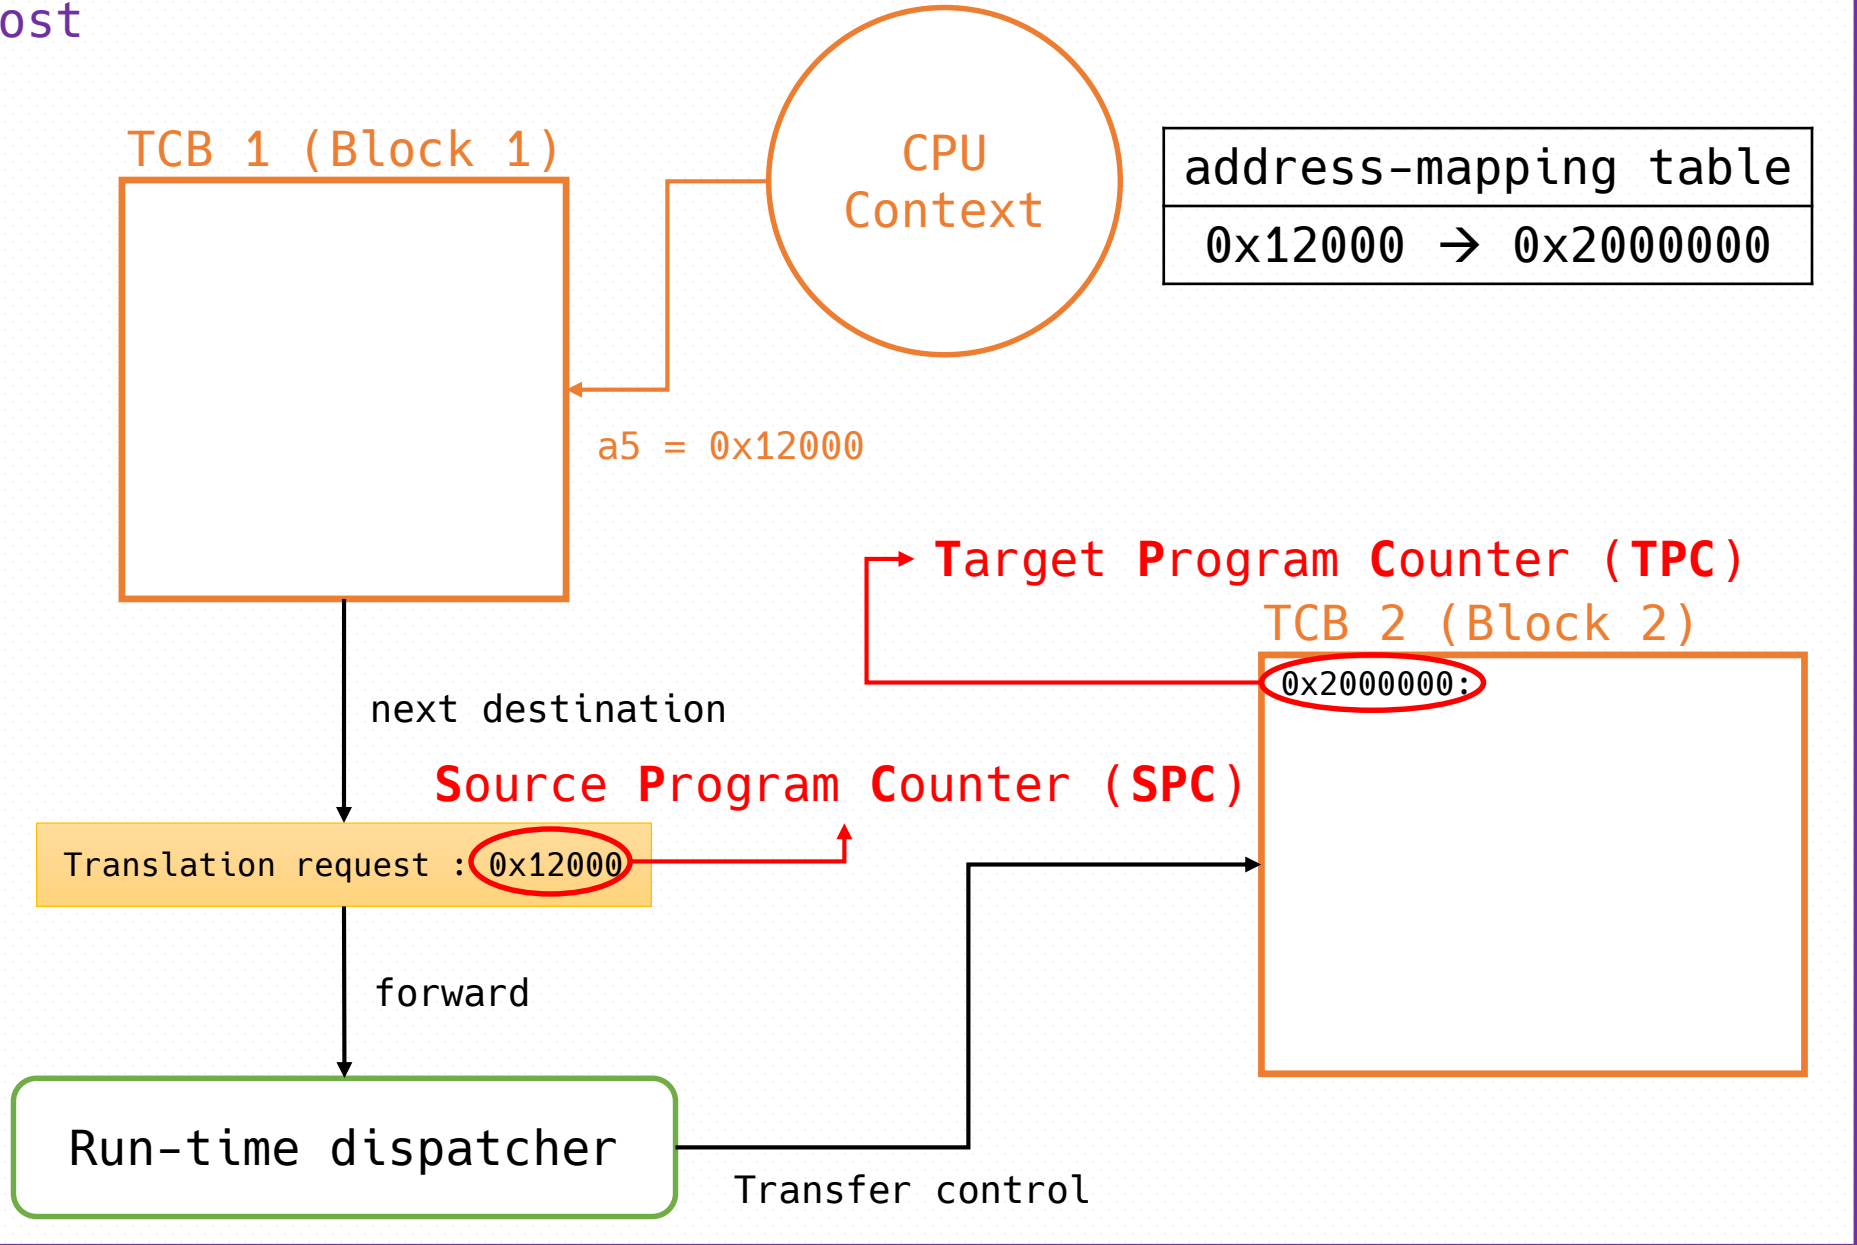
---

- Translation phase processes the client's instruction and translated into many code blocks, which we call **Translated Code Blocks (TCBs)**
- You may get many guest-to-host mapping relationship in the translation phase

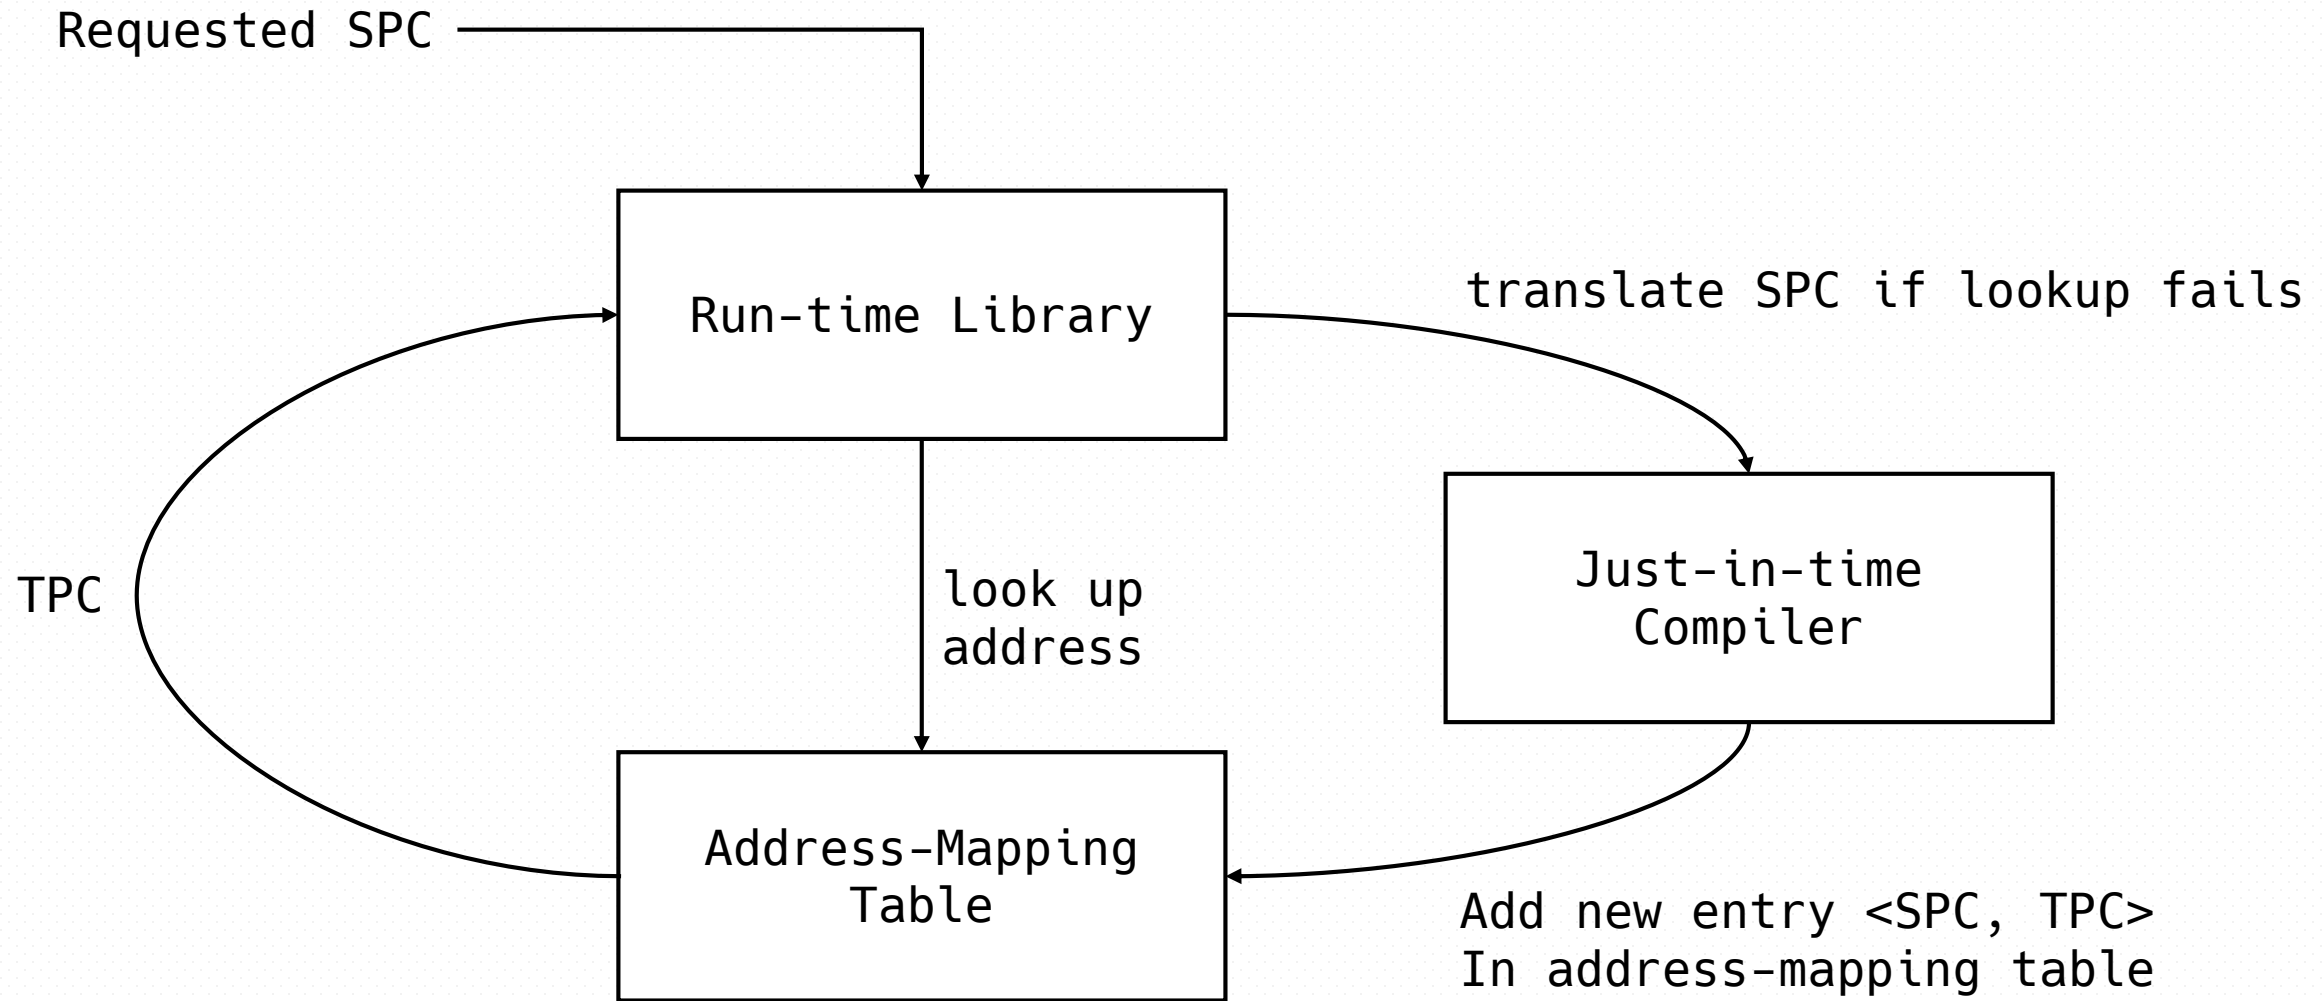
guest-binary



host

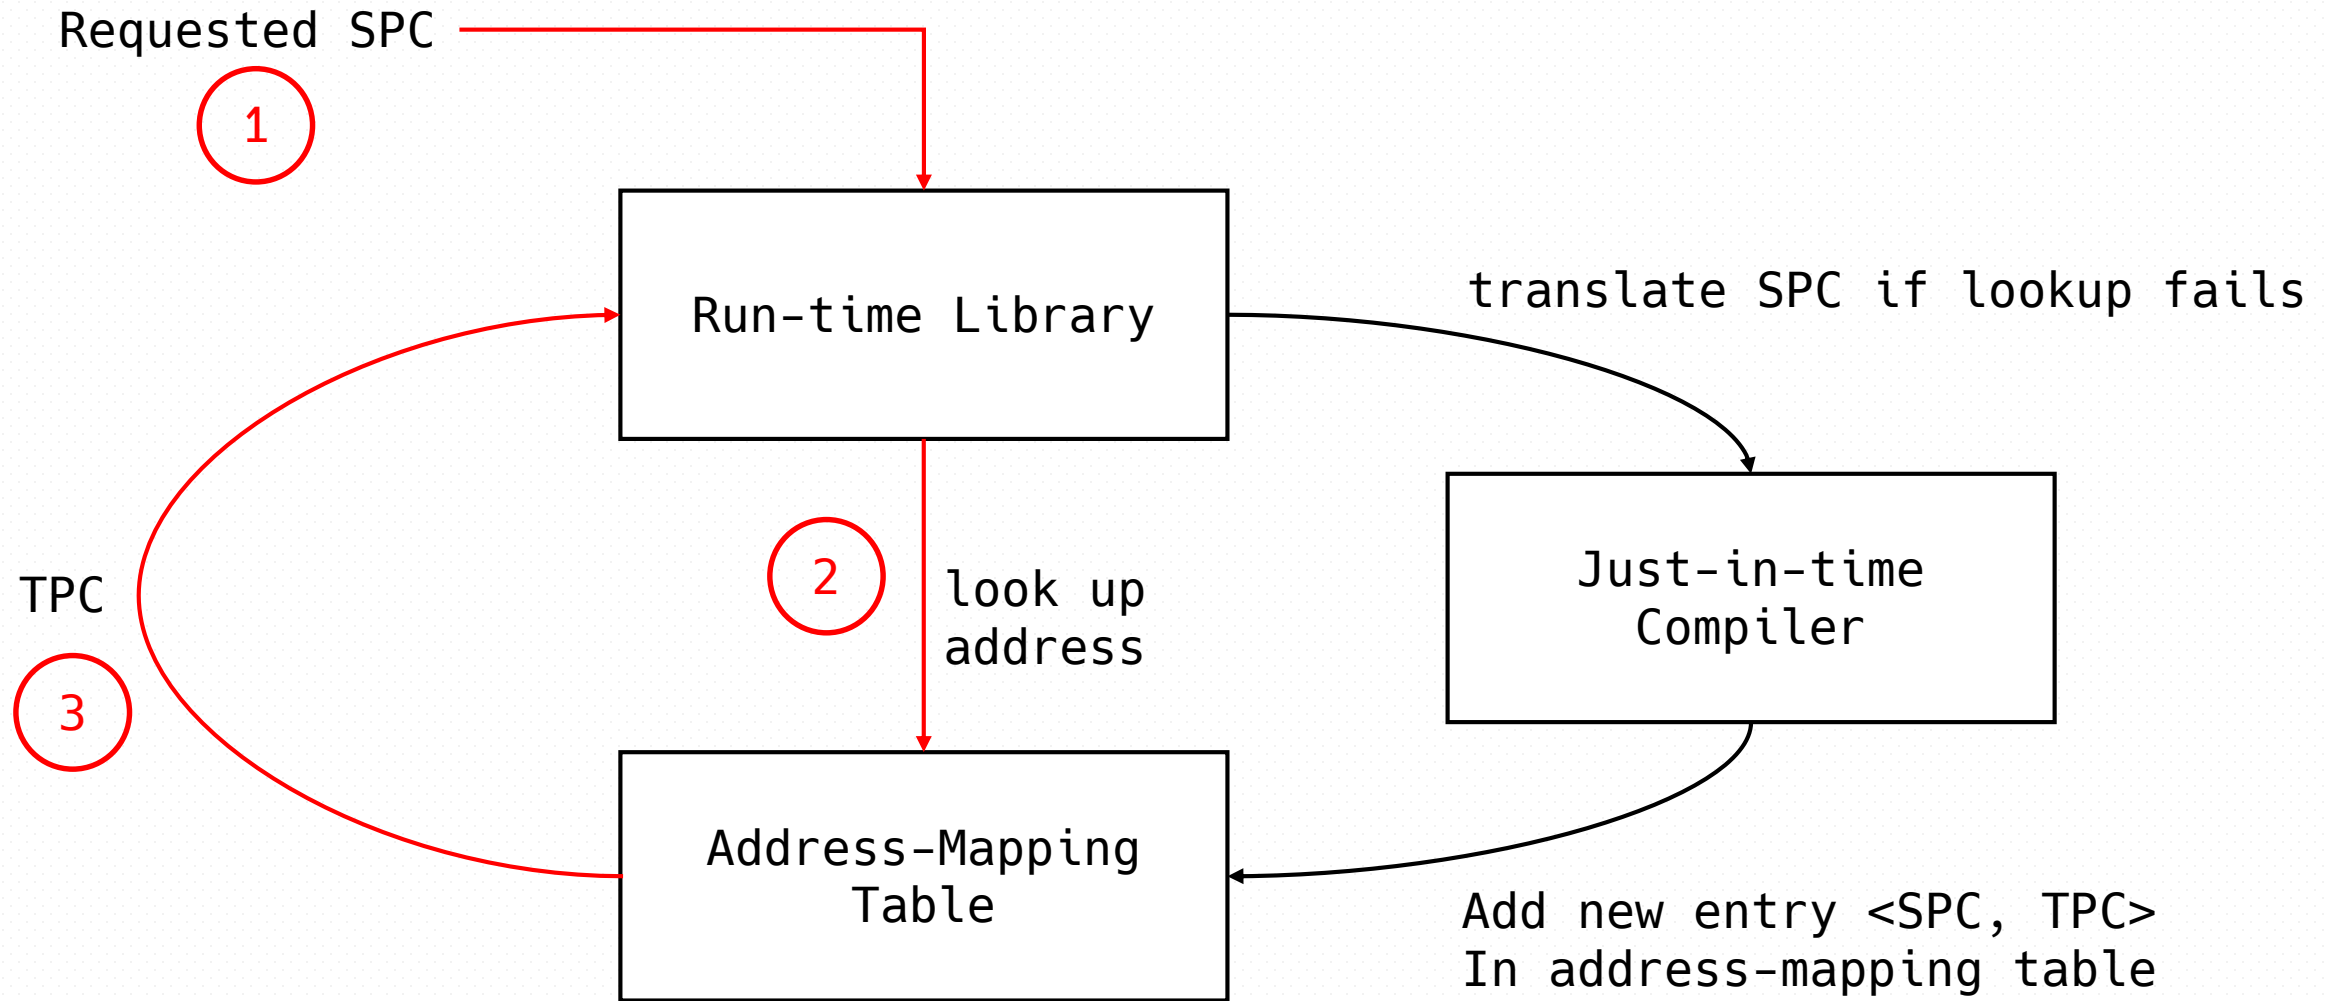






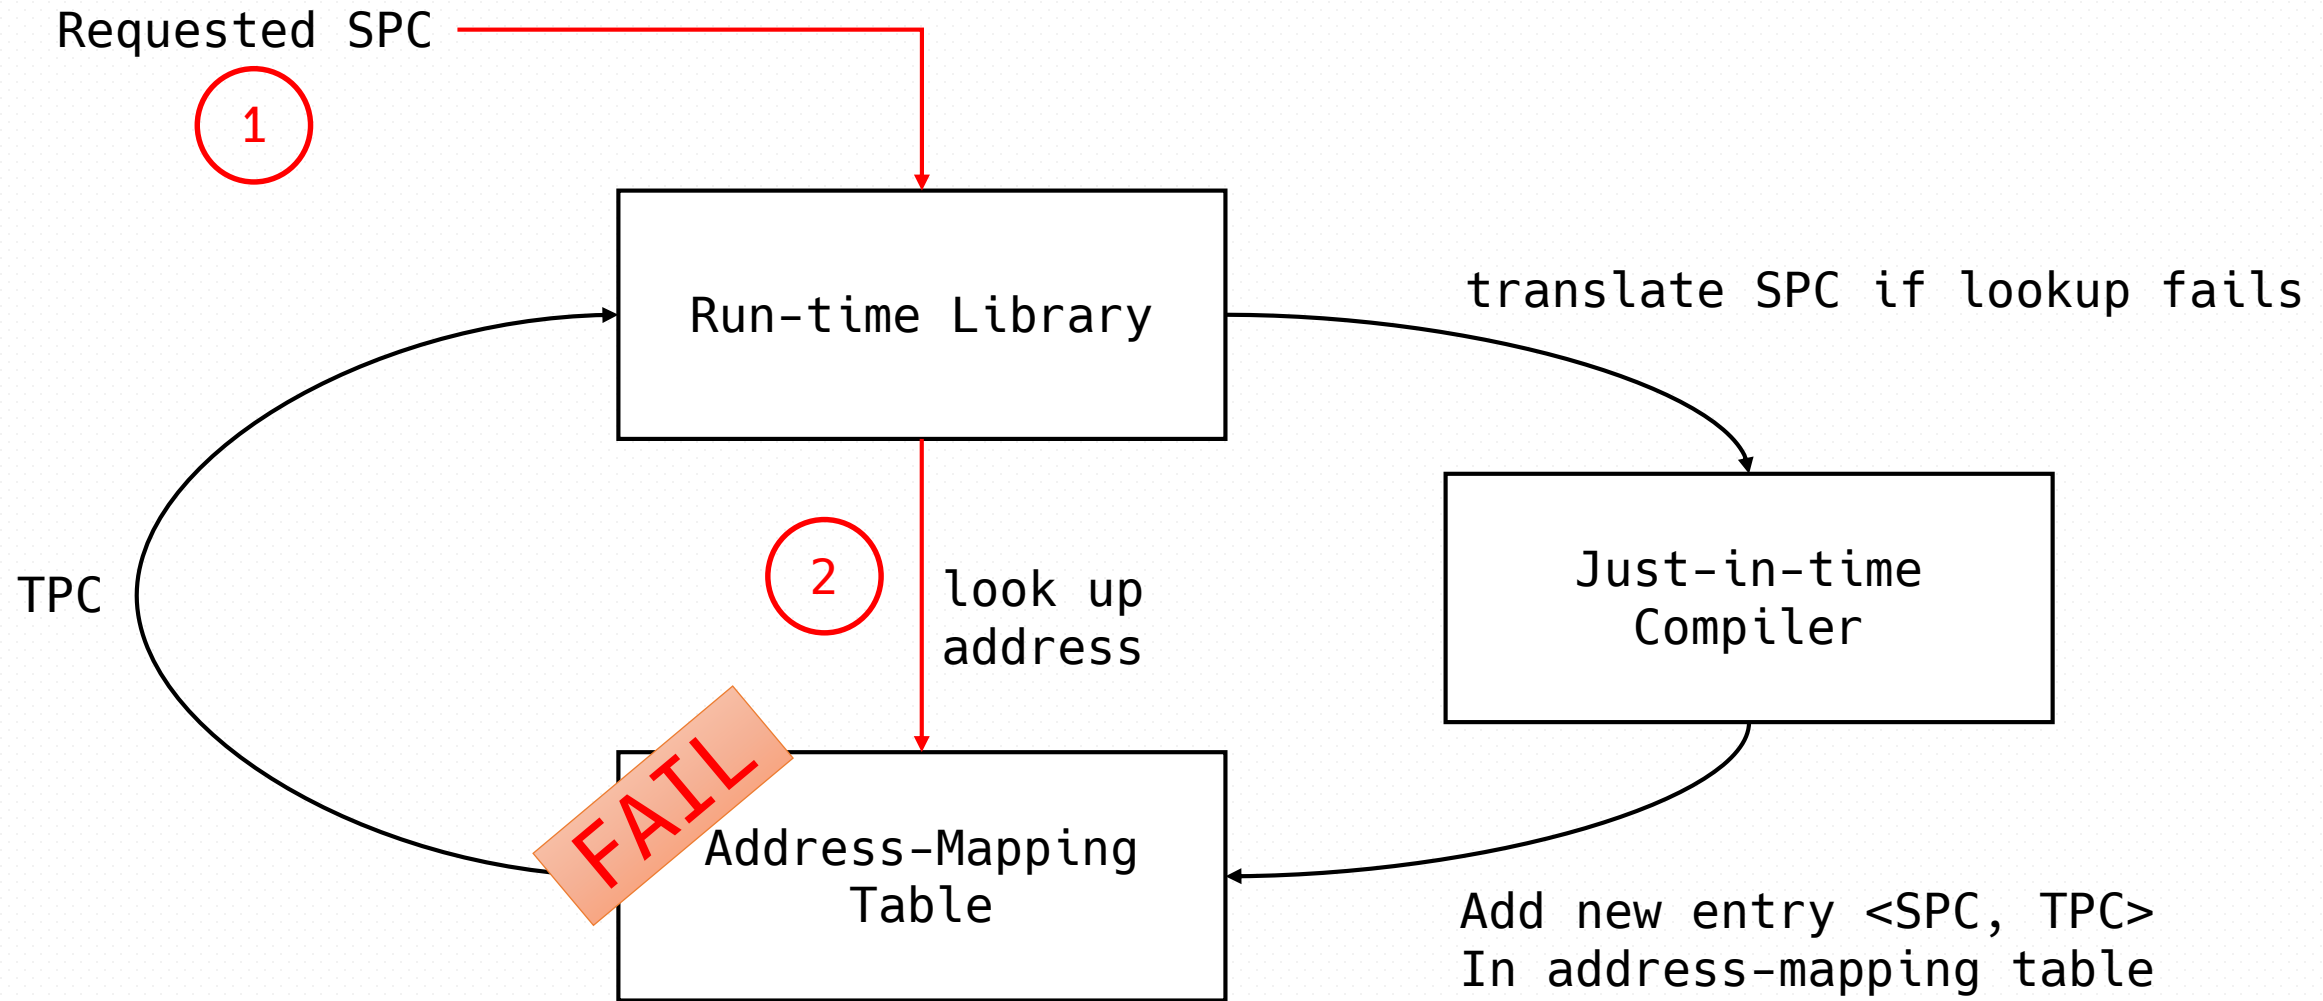
Possible Case 1 : the source block has been translated (from SBT or DBT)

Possible Case 2 : the source block has never been translated



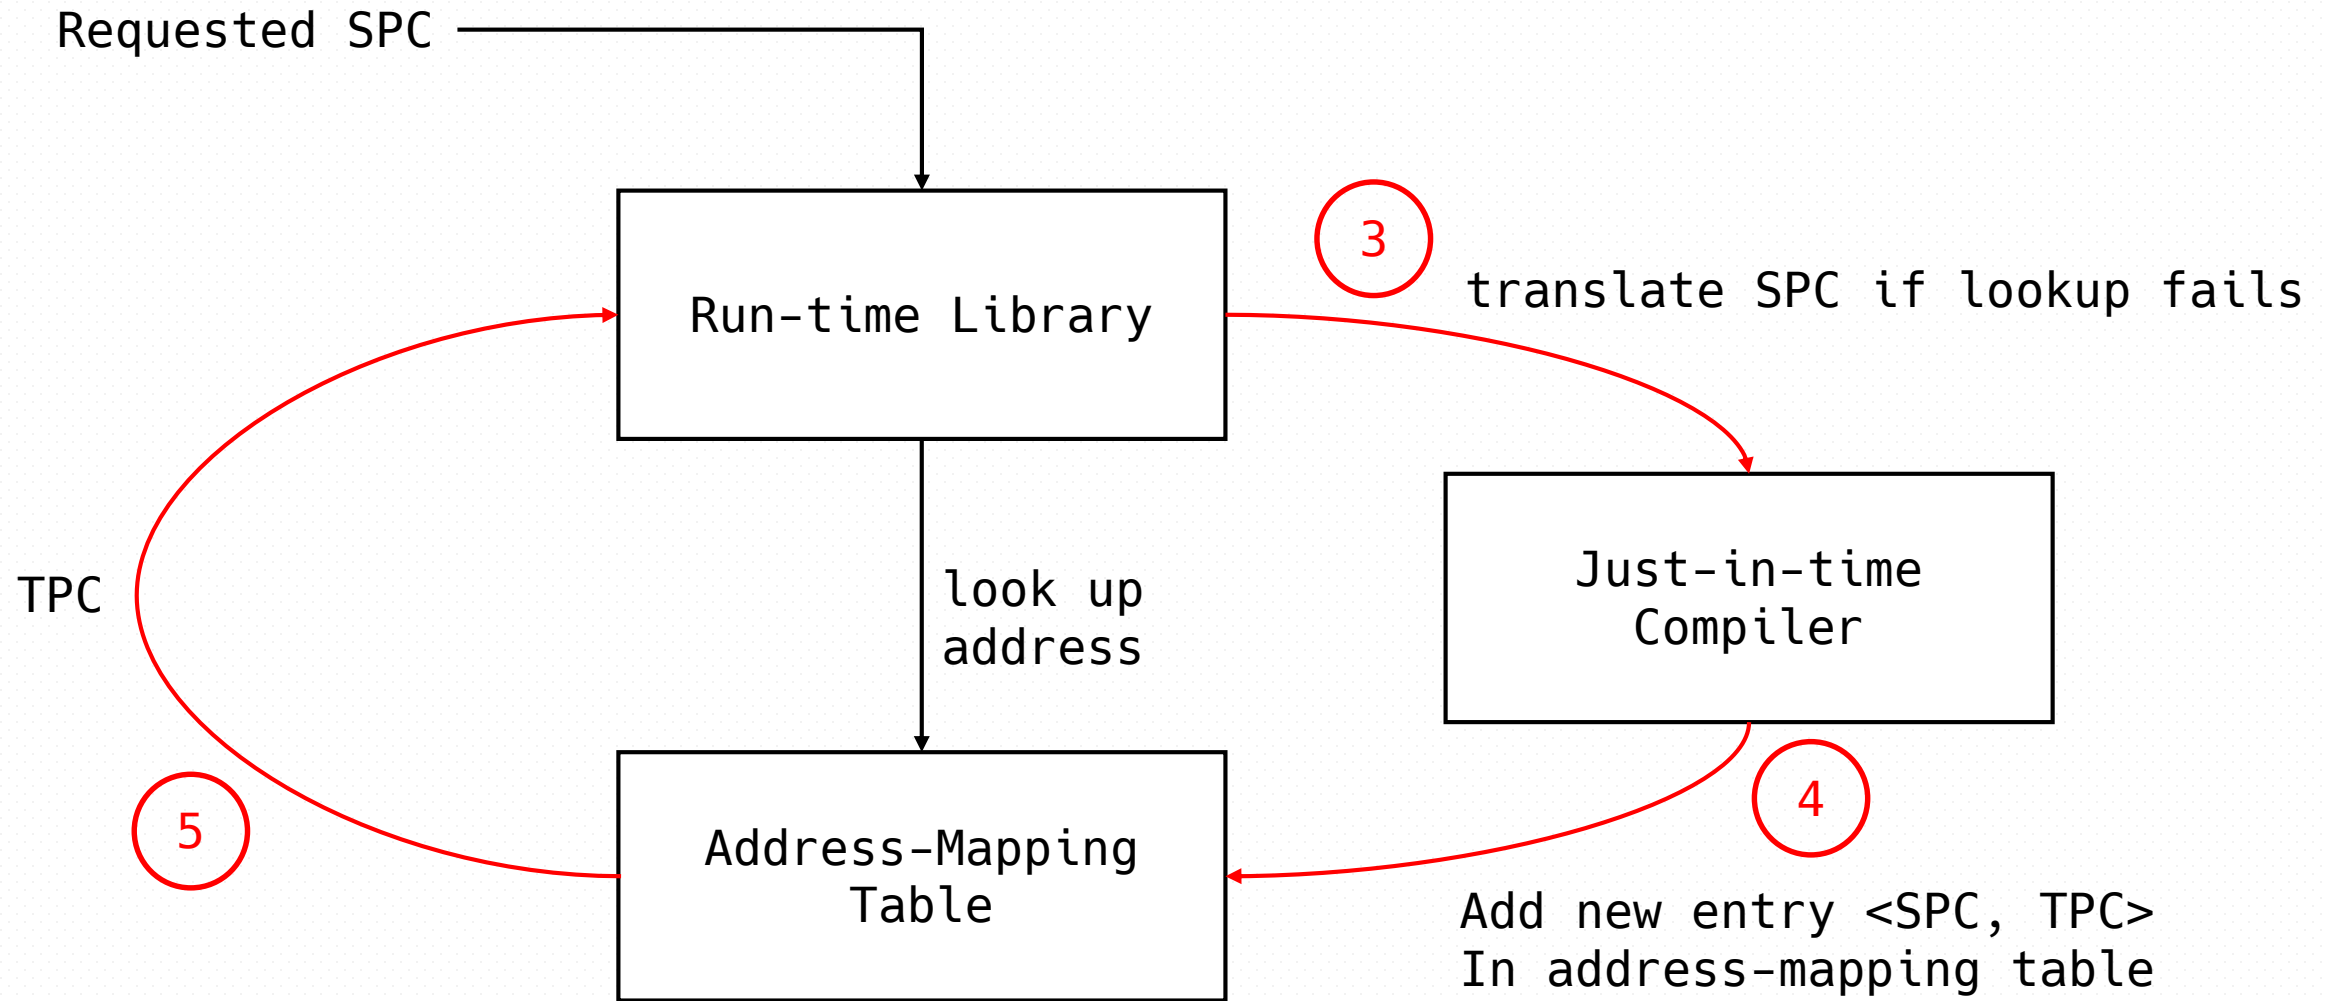
Possible Case 1 : the source block has been translated (from SBT or DBT)

Possible Case 2 : the source block has never been translated



Possible Case 1 : the source block has been translated (from SBT or DBT)

Possible Case 2 : the source block has never been translated



Possible Case 1 : the source block has been translated (from SBT or DBT)

Possible Case 2 : the source block has never been translated

# Instruction Terminator

---

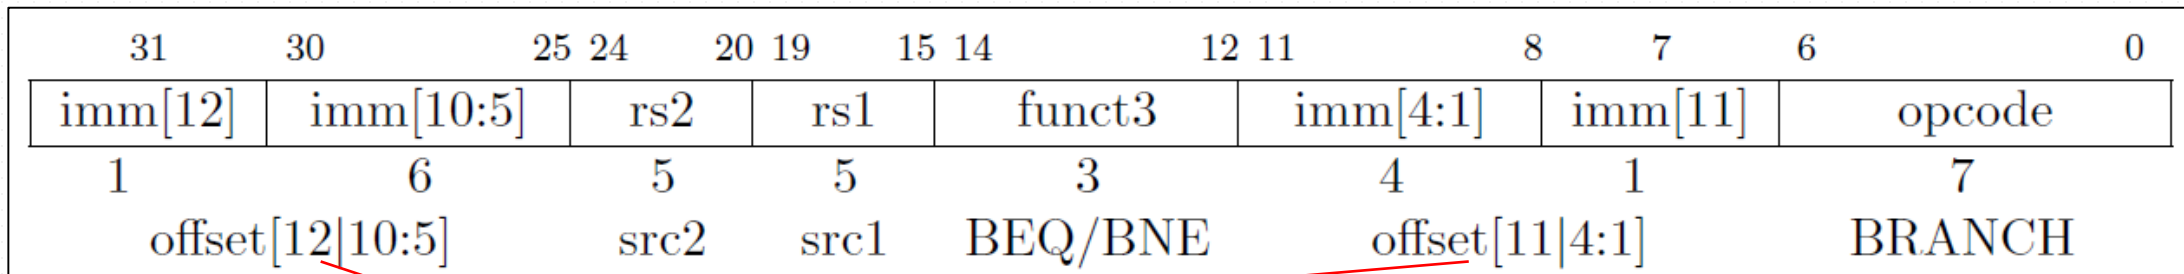
- Every basic block ends with an instruction terminator
- Can be classified into two groups
  - Direct Branches
  - Indirect Branches

instruction terminator

```
...  
...  
...  
jr a5
```

# Direct Branches

- Conditional branches, unconditional branches, direct function calls
- Single destination
- It is possible to know the destination in compilation time




Target  $\leftarrow$  \$PC + **offset** (PC-relative addressing)

From <https://github.com/riscv/riscv-isa-manual/releases/download/draft-20210612-98f3349/riscv-spec.pdf>

# Indirect Branches

---

- May have more than one destination
- The destination may be stored in a register or memory
- It is difficult to predict the destinations without context information



The value of `a5` ?

# Optimizations in Rabbit

---

- Direct Branch Link Optimization
  - In compilation time
- Direct Function Return Optimization
  - In compilation time
- Block Chaining Optimization
  - In run-time



# Direct Branch Link

---

## ■ Steps

1. Analyze LLVM basic blocks which end with direct branches
2. Try to find the SPC of successor block as much as possible and modify the former block to directly jump to the next one in compilation time

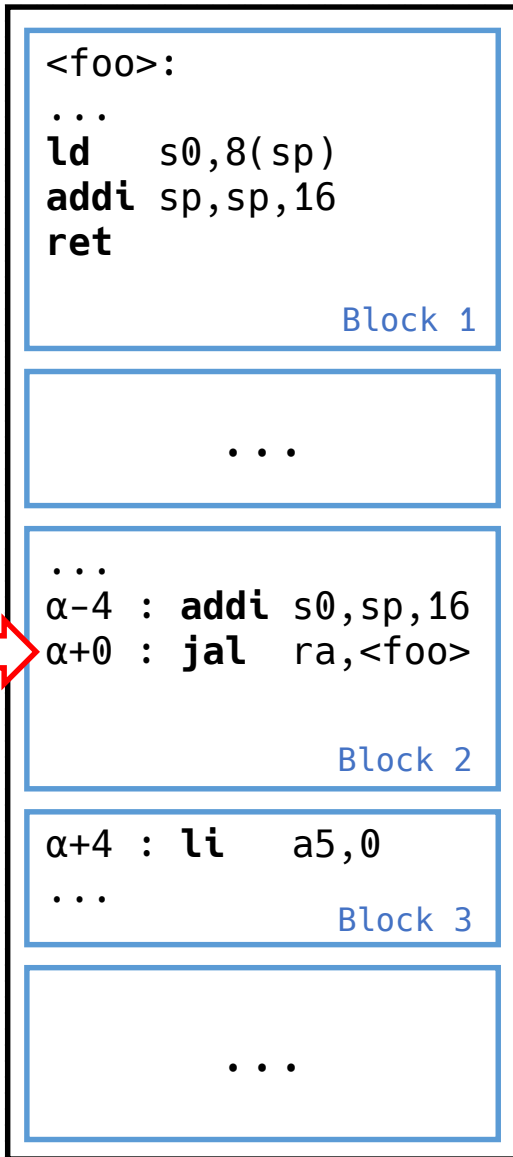
## ■ Mitigate run-time overhead and improve performance

# Direct Function Return

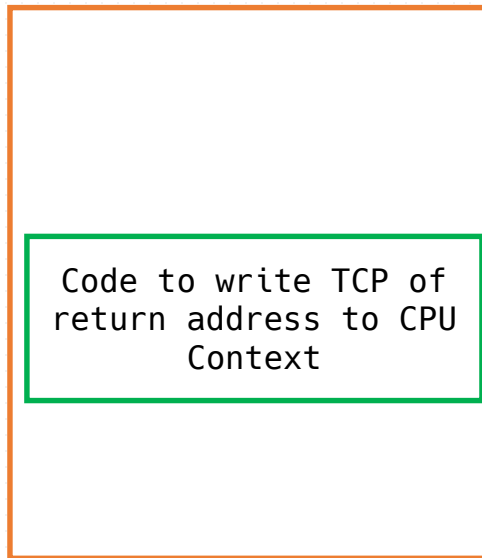
---

- In the normal situation, the return address of a function call should be the address of next instruction
- This optimization exploits that fact to tell the translator the TPC of return address in compilation time
  - No need to lookup return address in run-time

# Guest-binary



## TCB (Block 2)

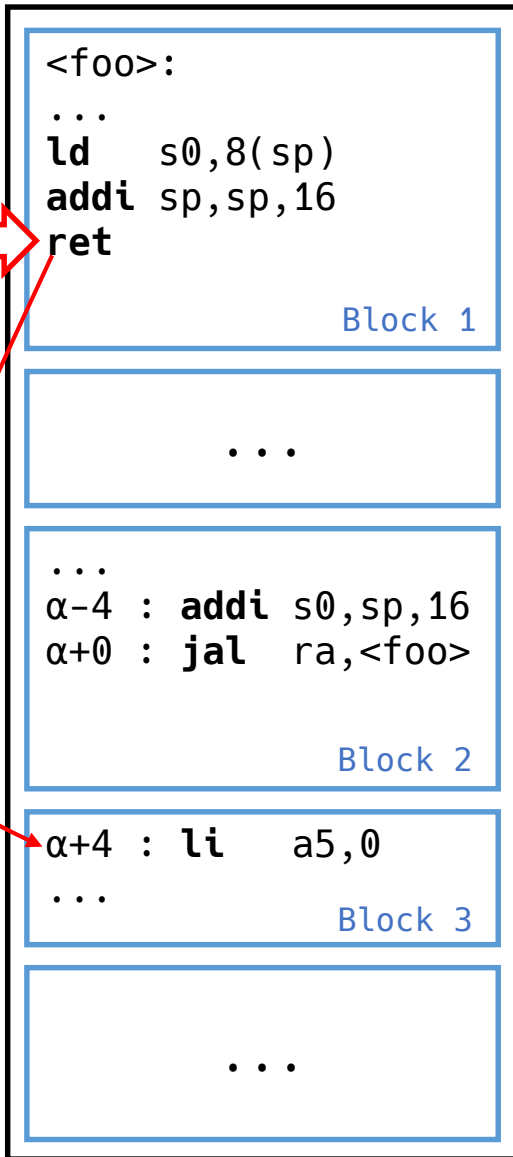


CPU Context	
register	value
...	...
ra	$\alpha+4 \rightarrow 0x4028000$
...	...

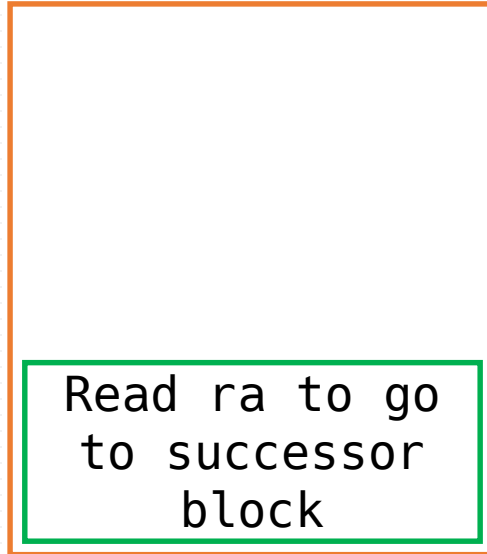
Replaced by TPC

Address-Mapping Table	
SPC	TPC
...	...
$\alpha+4$	$0x4028000$
...	...


# Guest-binary



# TCB (Block 1)



# TCB (Block 3)



CPU Context	
register	value
...	...
ra	0x4028000
...	...

👍 Do not need to do address lookup

# Direct Function Return

---

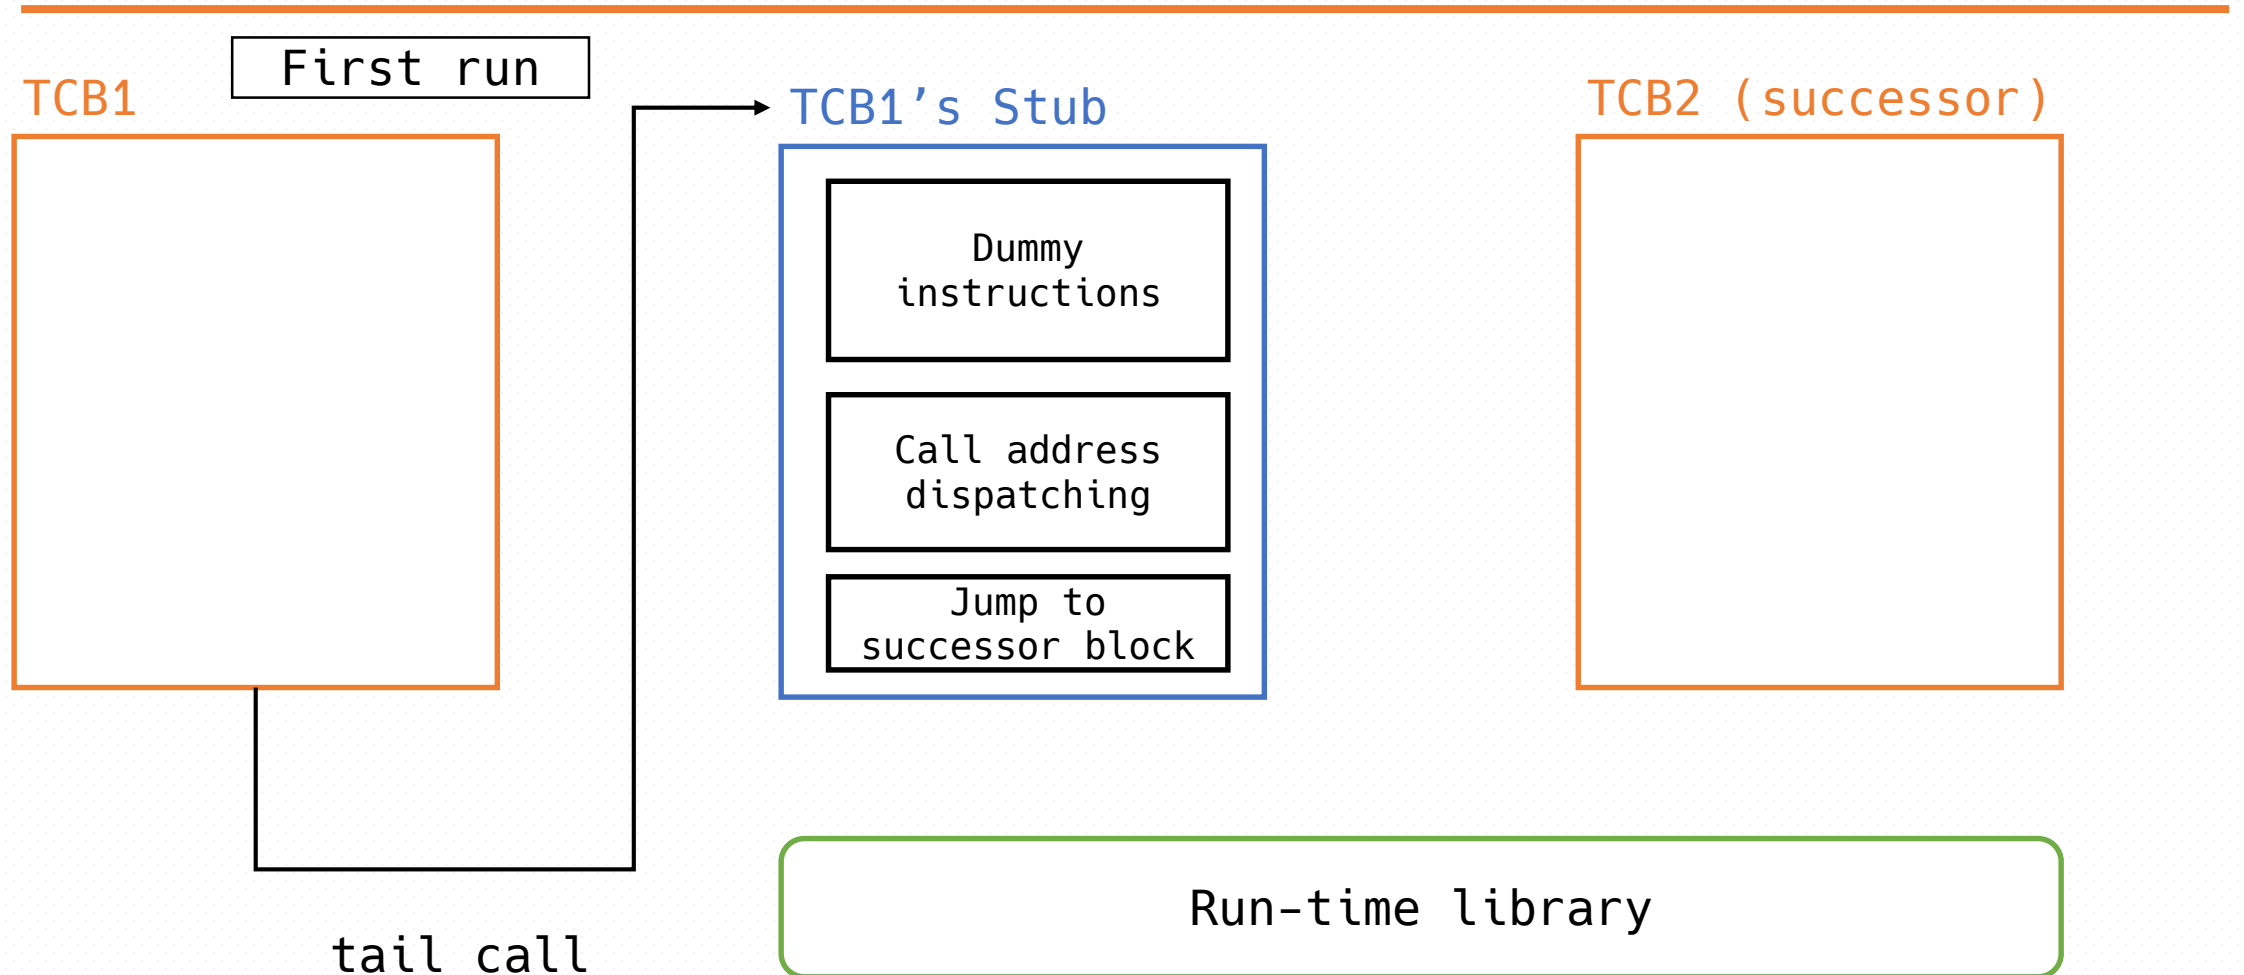
- Caveat : We assume the link register will not be modified in the subroutine
- Otherwise, the program will crash because of corrupt link register

# Block Chaining Optimization

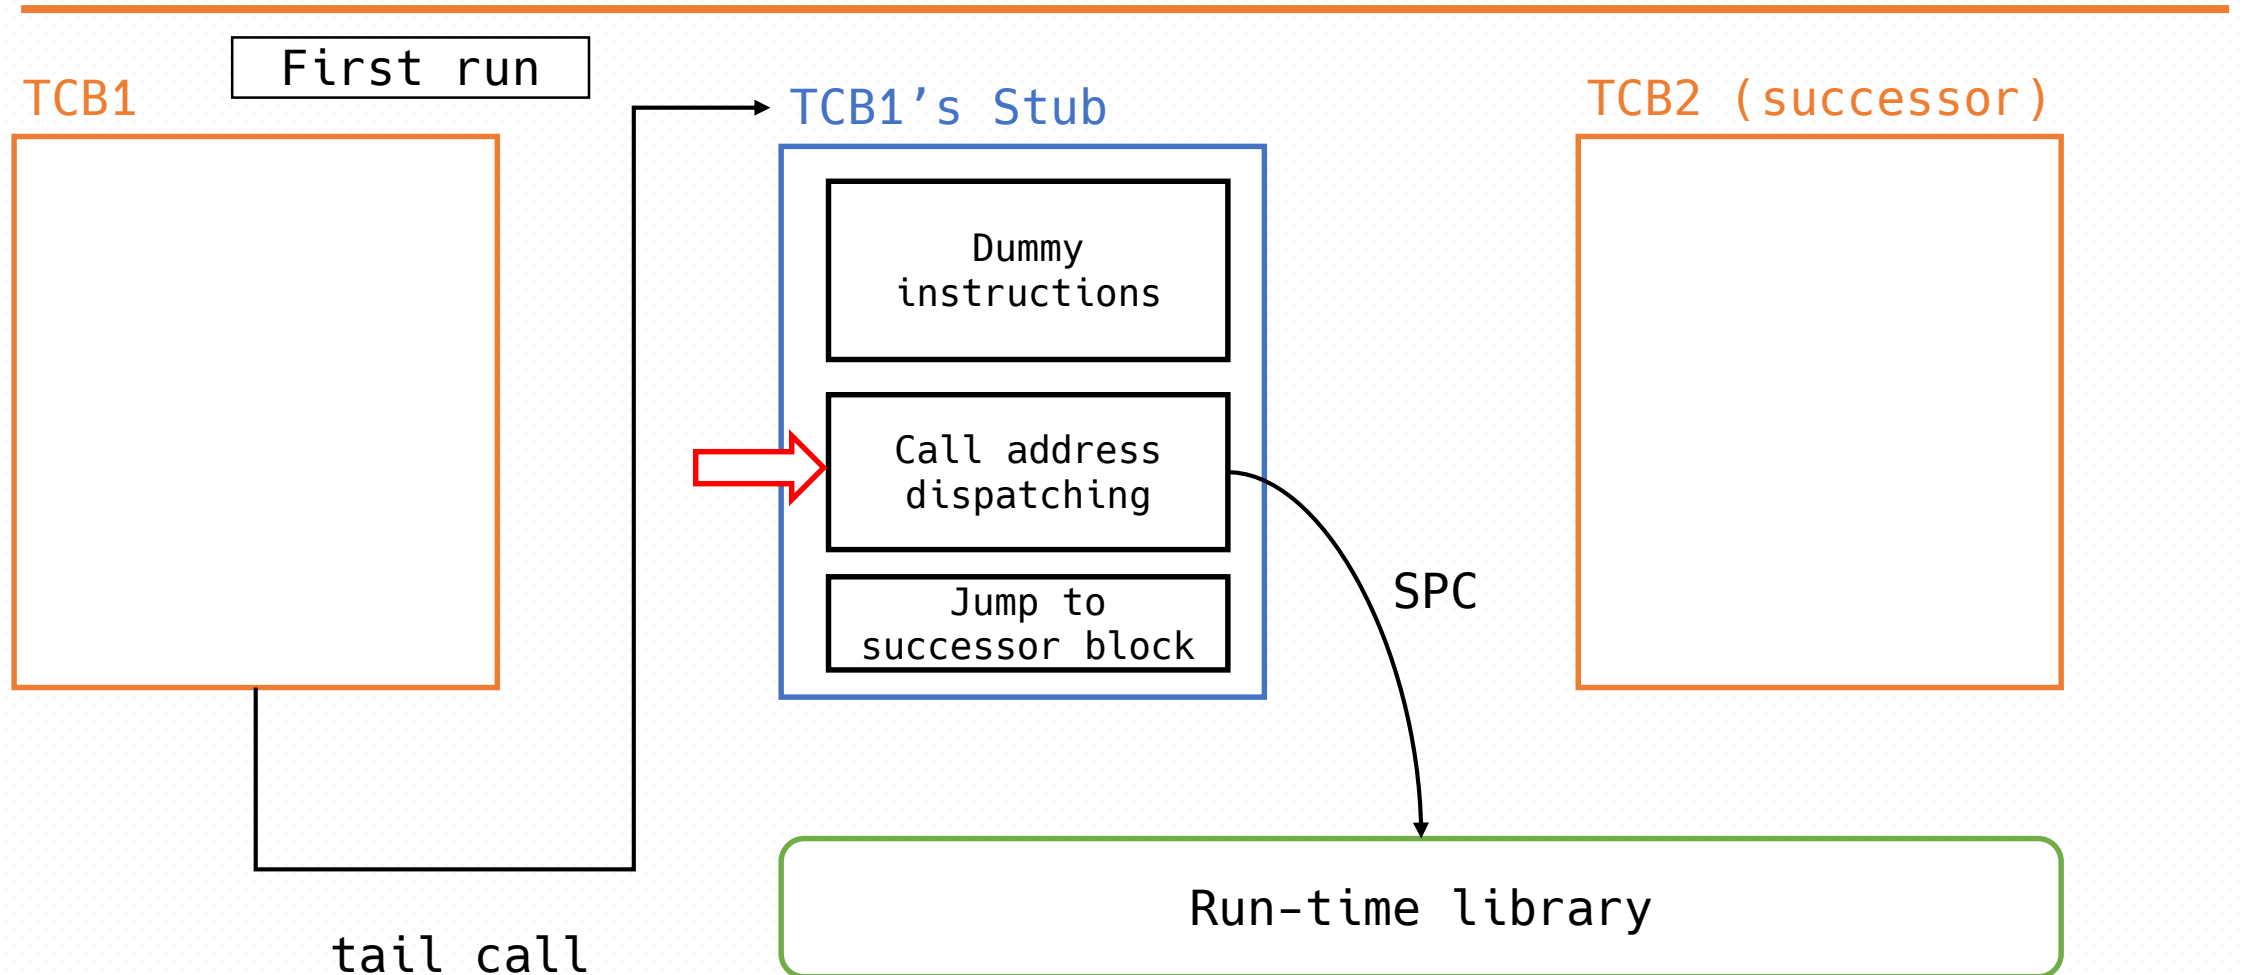
---

- Basically, block chaining optimization do the same thing as direct branch link optimization does but in run-time
- Adopt **self-modifying code**
  - Modify its own instructions in run-time
  - Patch the destination into the code

# Block Chaining Optimization

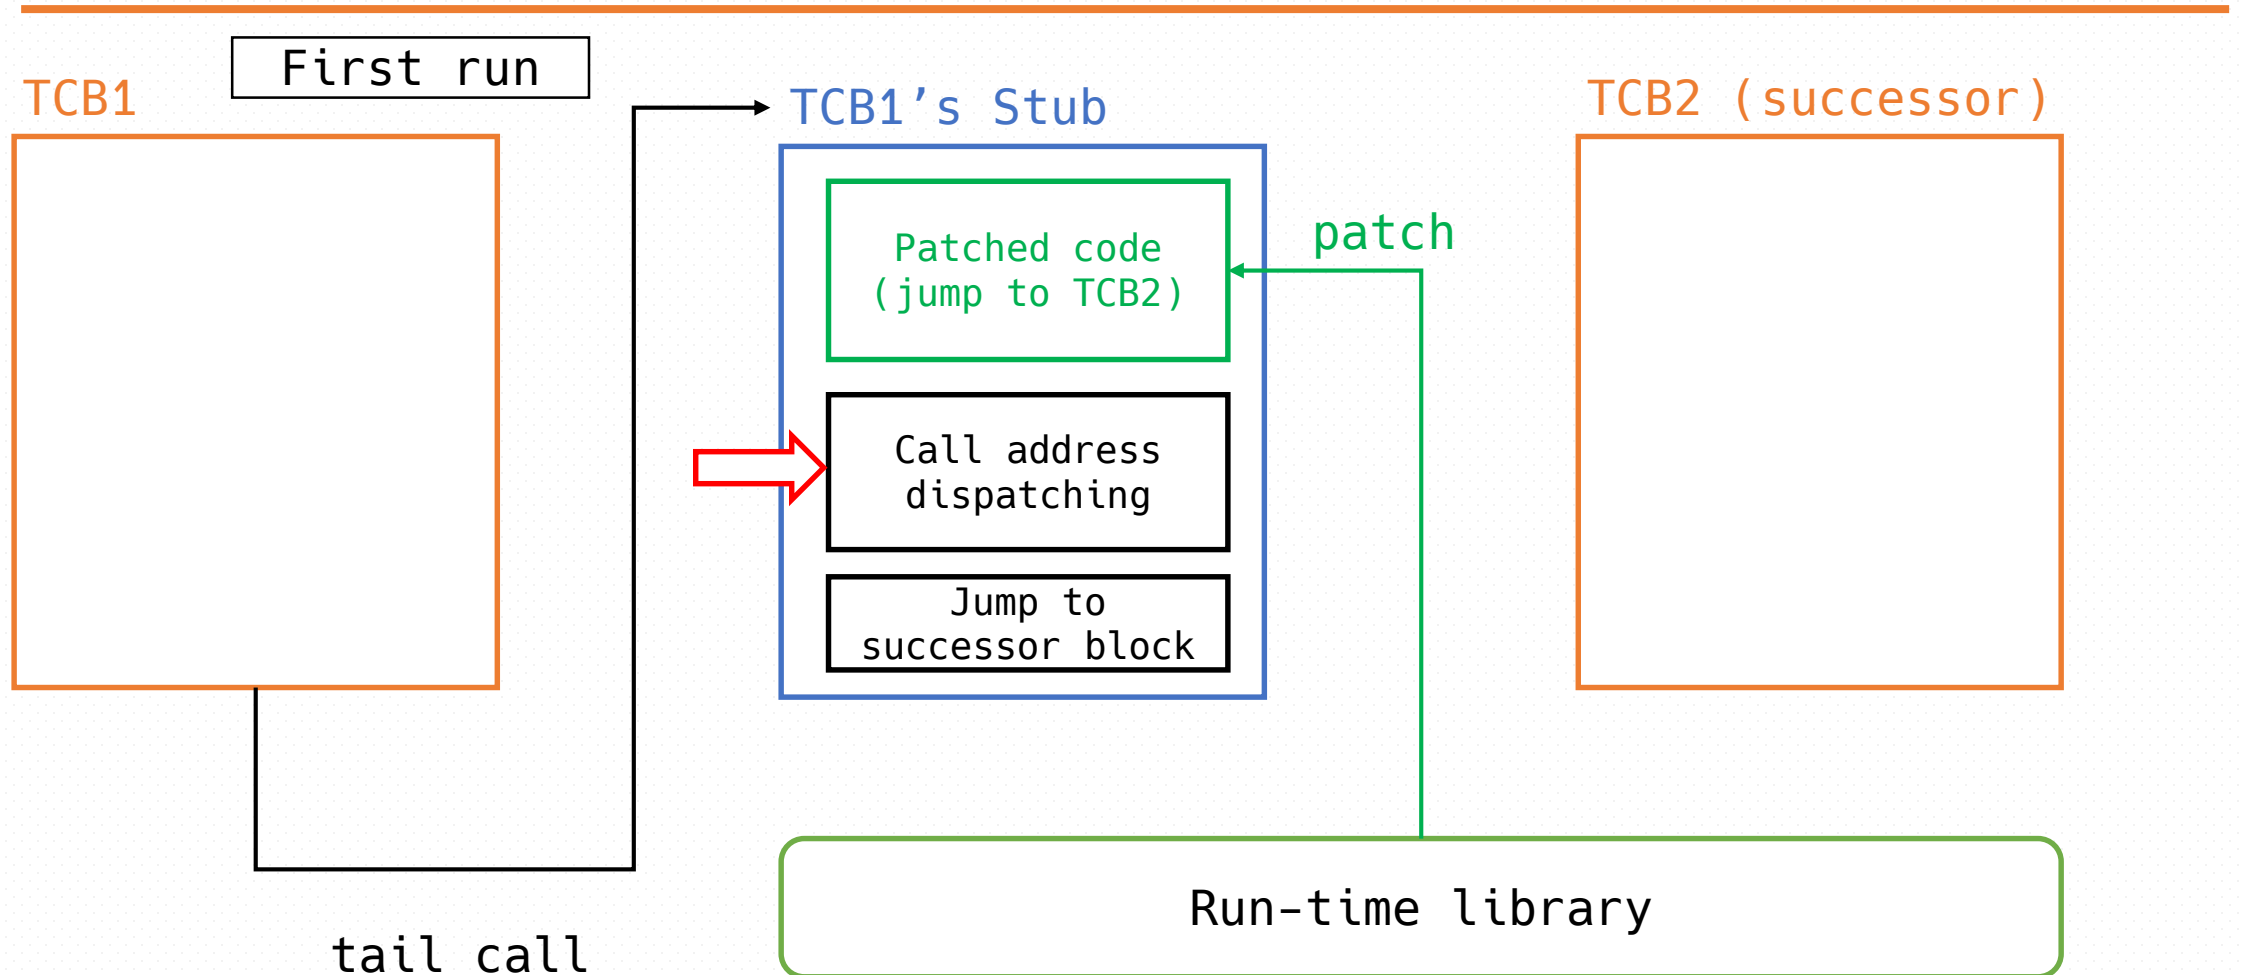


# Block Chaining Optimization

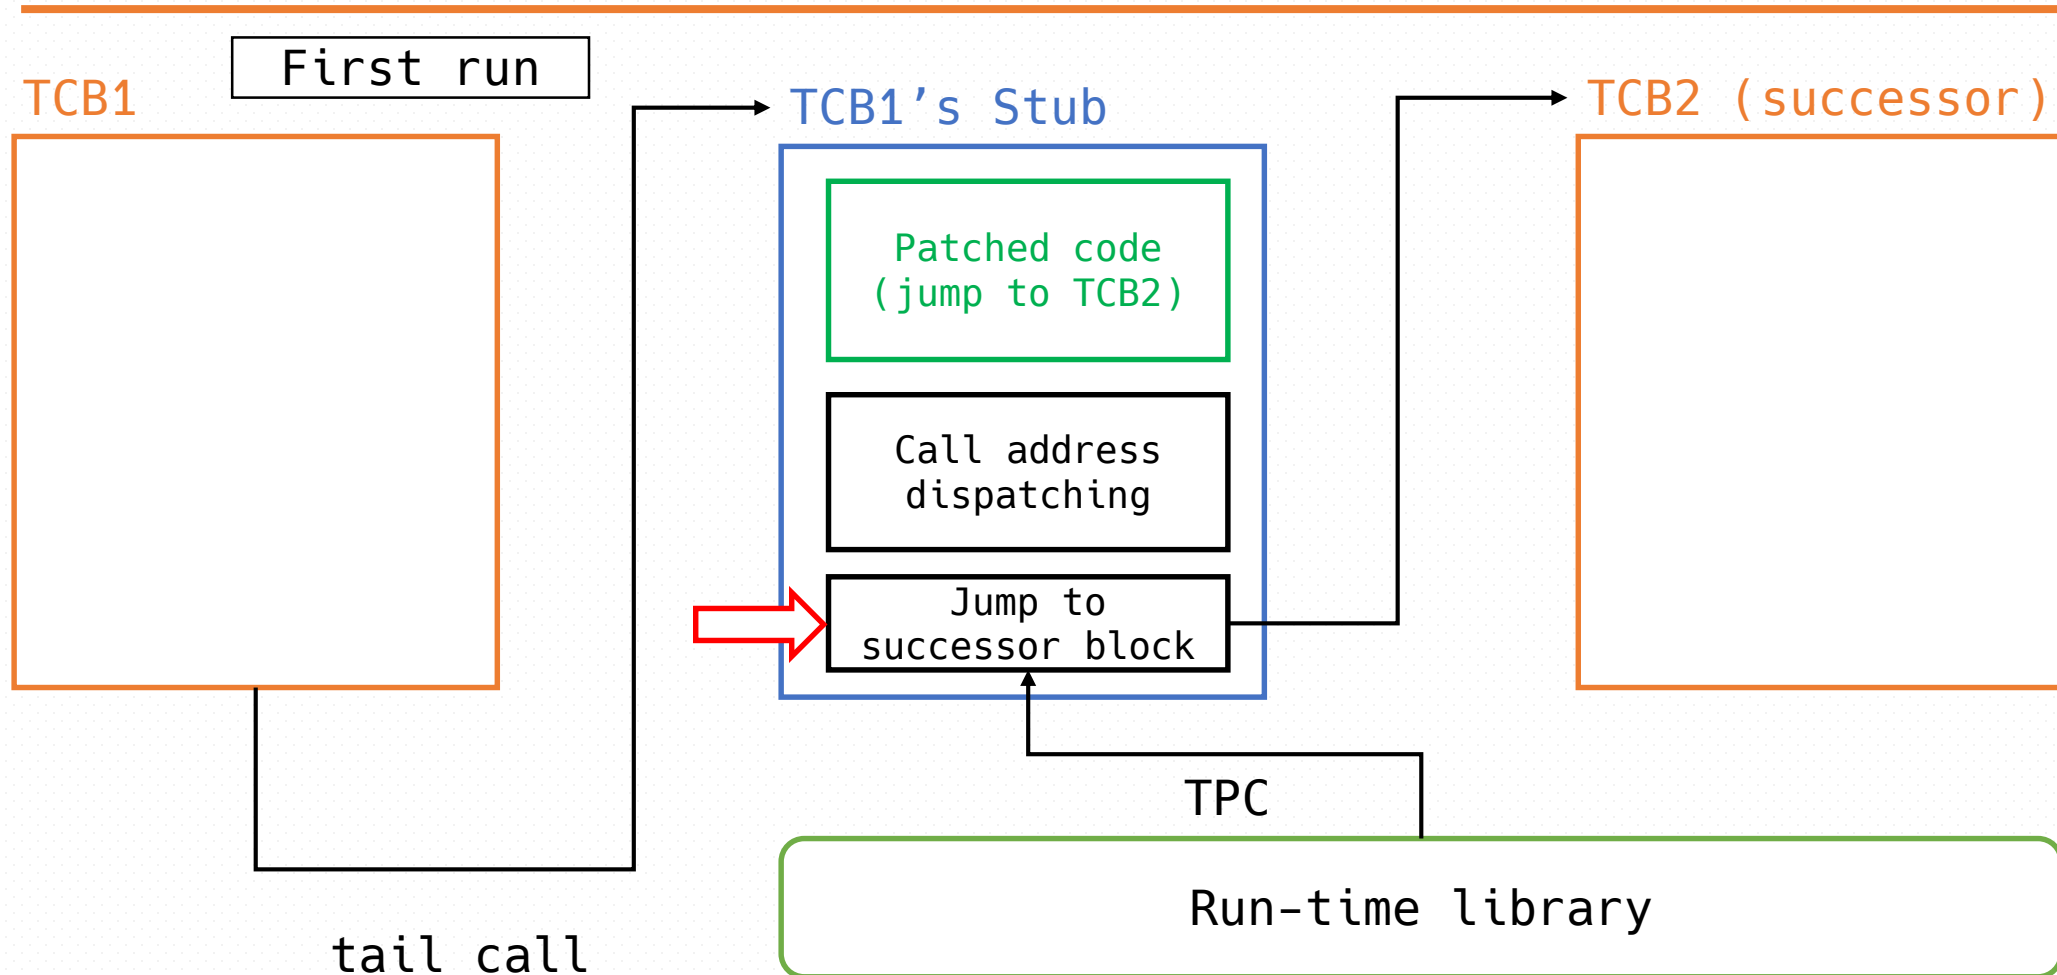




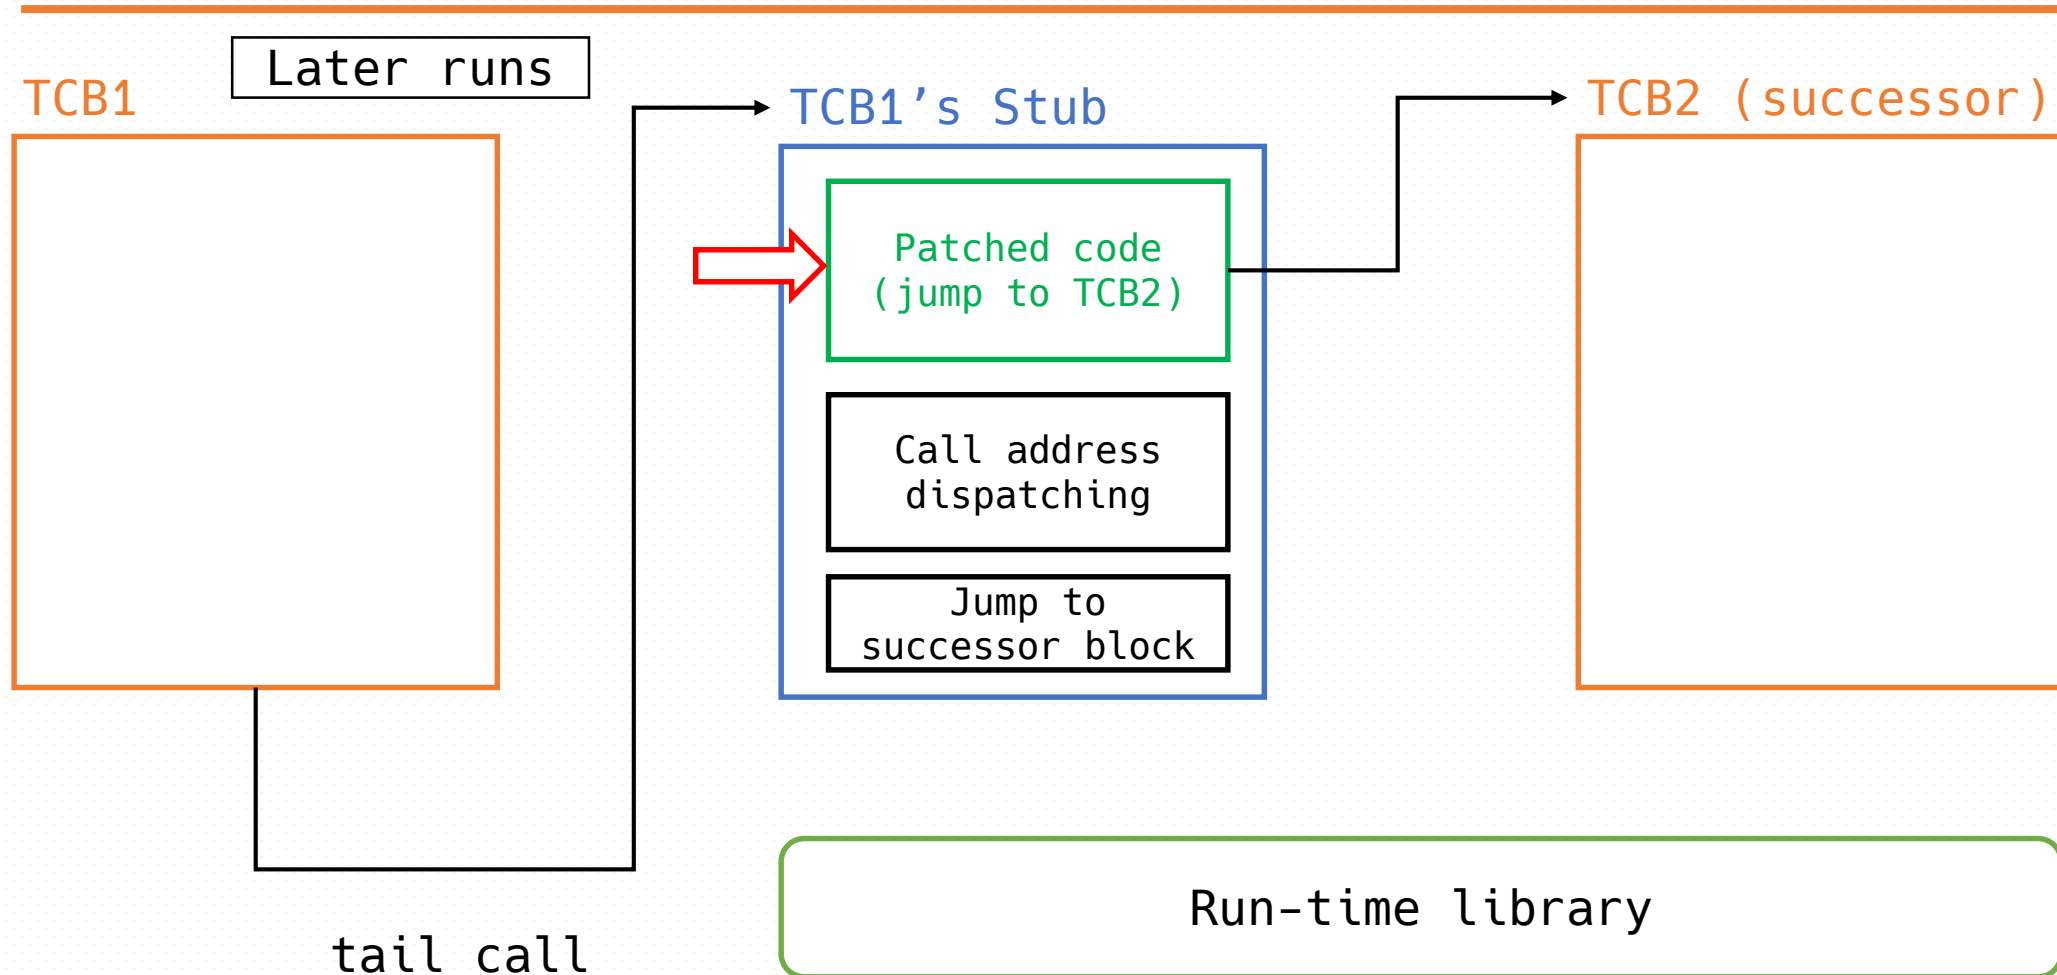
# Block Chaining Optimization



# Block Chaining Optimization



# Block Chaining Optimization



# Block Chaining Optimization

---

- Limitations
  - Only support chaining dtblocks to dtblocks
    - That means this optimization can't handle (stblocks → stblocks; stblocks → dtblocks and dtblocks → stblocks) chaining
  - Tail call optimization is involved that means we have an extra jump instruction. That is not efficient
  - Cannot handle indirect branch chaining
  - This optimization only supports old translator we had made before. Rabbit cannot use this optimization

# Outline

---

- Background
- **Hyperchaining**
- Experiments
- Conclusion

# Hyperchaining

---

- We develop **Hyperchaining** optimization, an improved version of block chaining optimization
- Be able to chain both stblocks and dtblocks
- Be able to handle indirect branches
- Design to work well with HBT and Rabbit

# Hyperchaining

---

- Depend on the implementation, hyperchaining optimization can be classified into two groups
  - **Platform-independent** (indep) hyperchaining
  - **Platform-dependent** (dep) hyperchaining

# Indep-Hyperchaining

---

- In indep-hyperchaining, we design a subroutine, **direct-branch chaining handler**, to chain blocks for direct branches
- Implemented in LLVM IR level
- This subroutine supports some functionalities of run-time dispatcher
  - Dynamic binary translation
  - Address lookup



# Indep-Hyperchaining

---

## ■ Steps

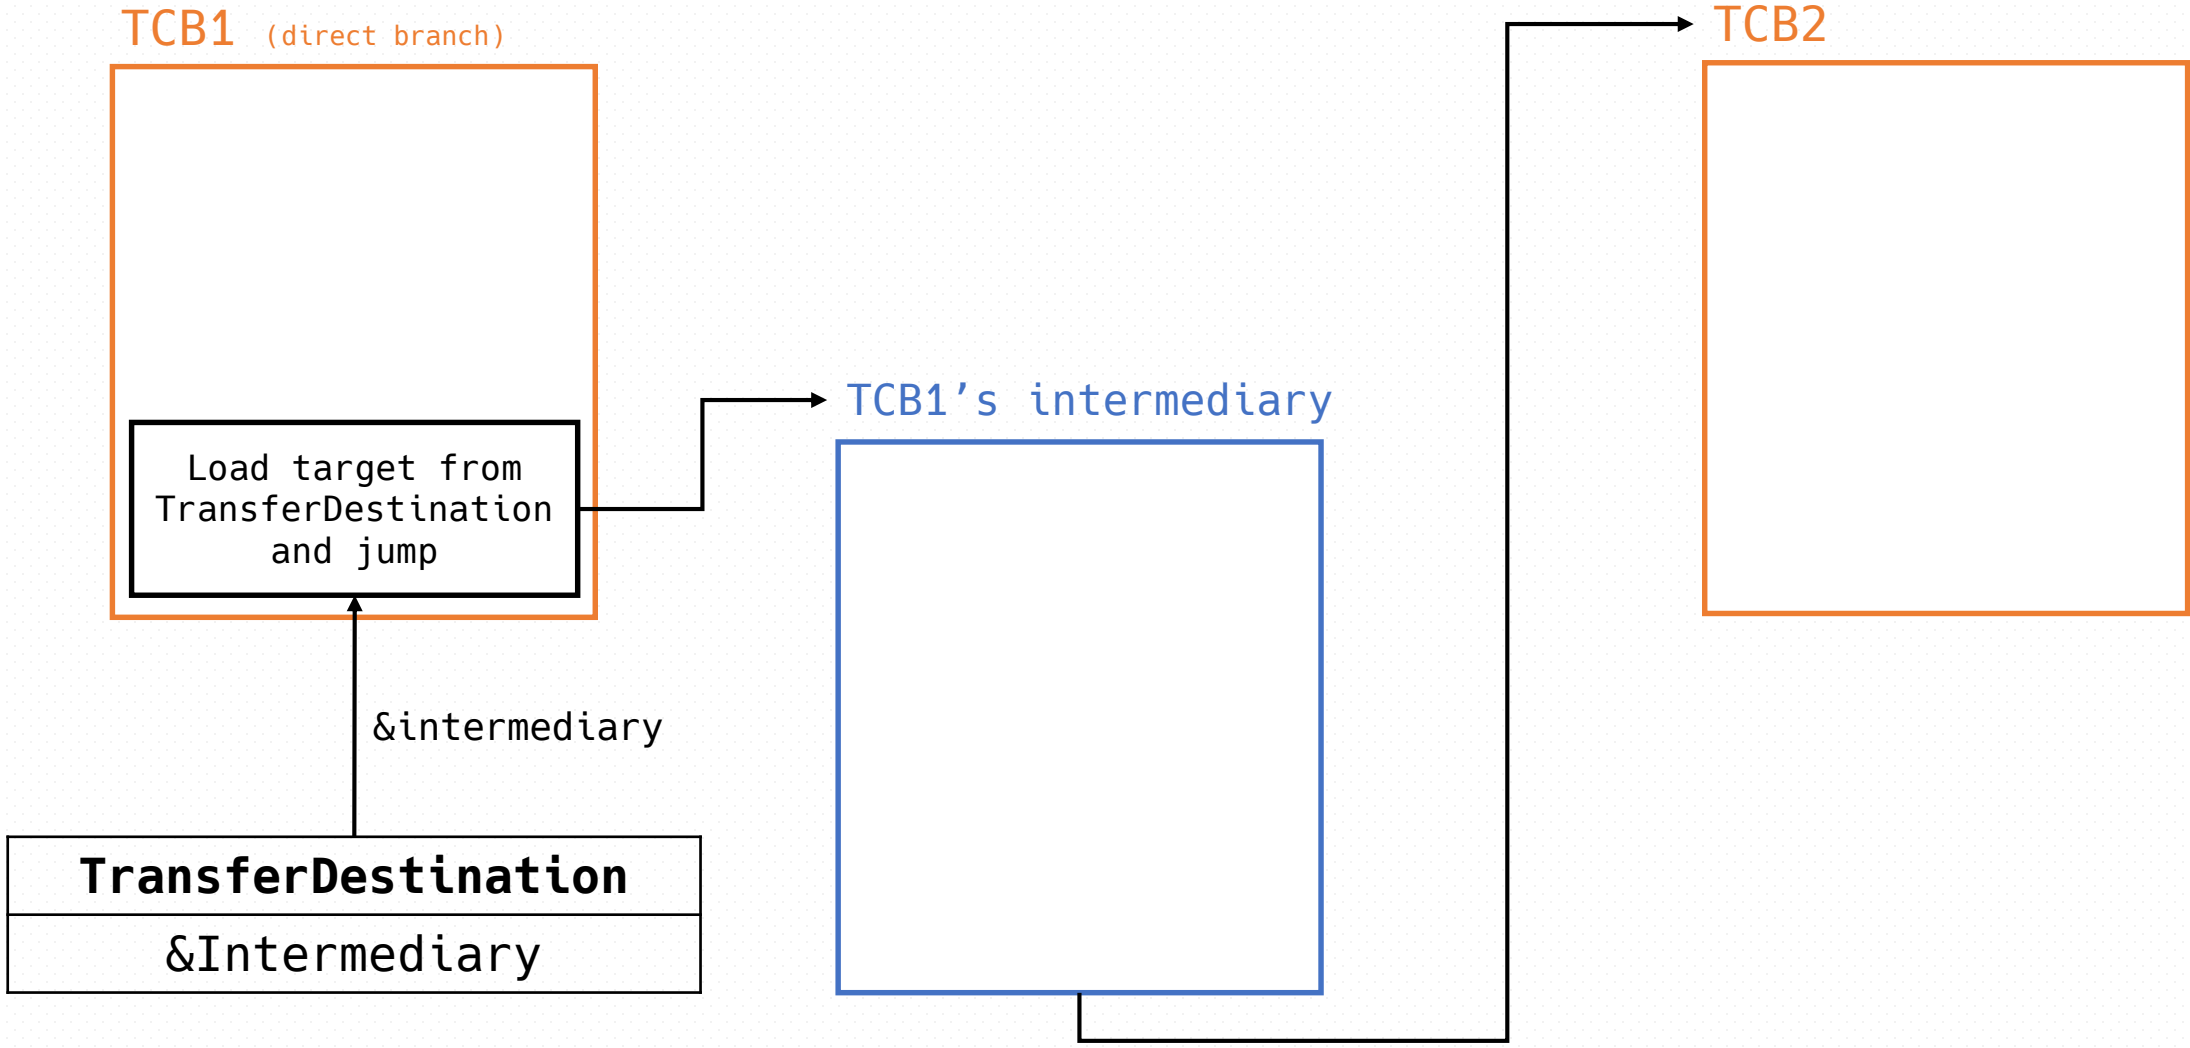
1. TCB loads address from a dedicated memory,

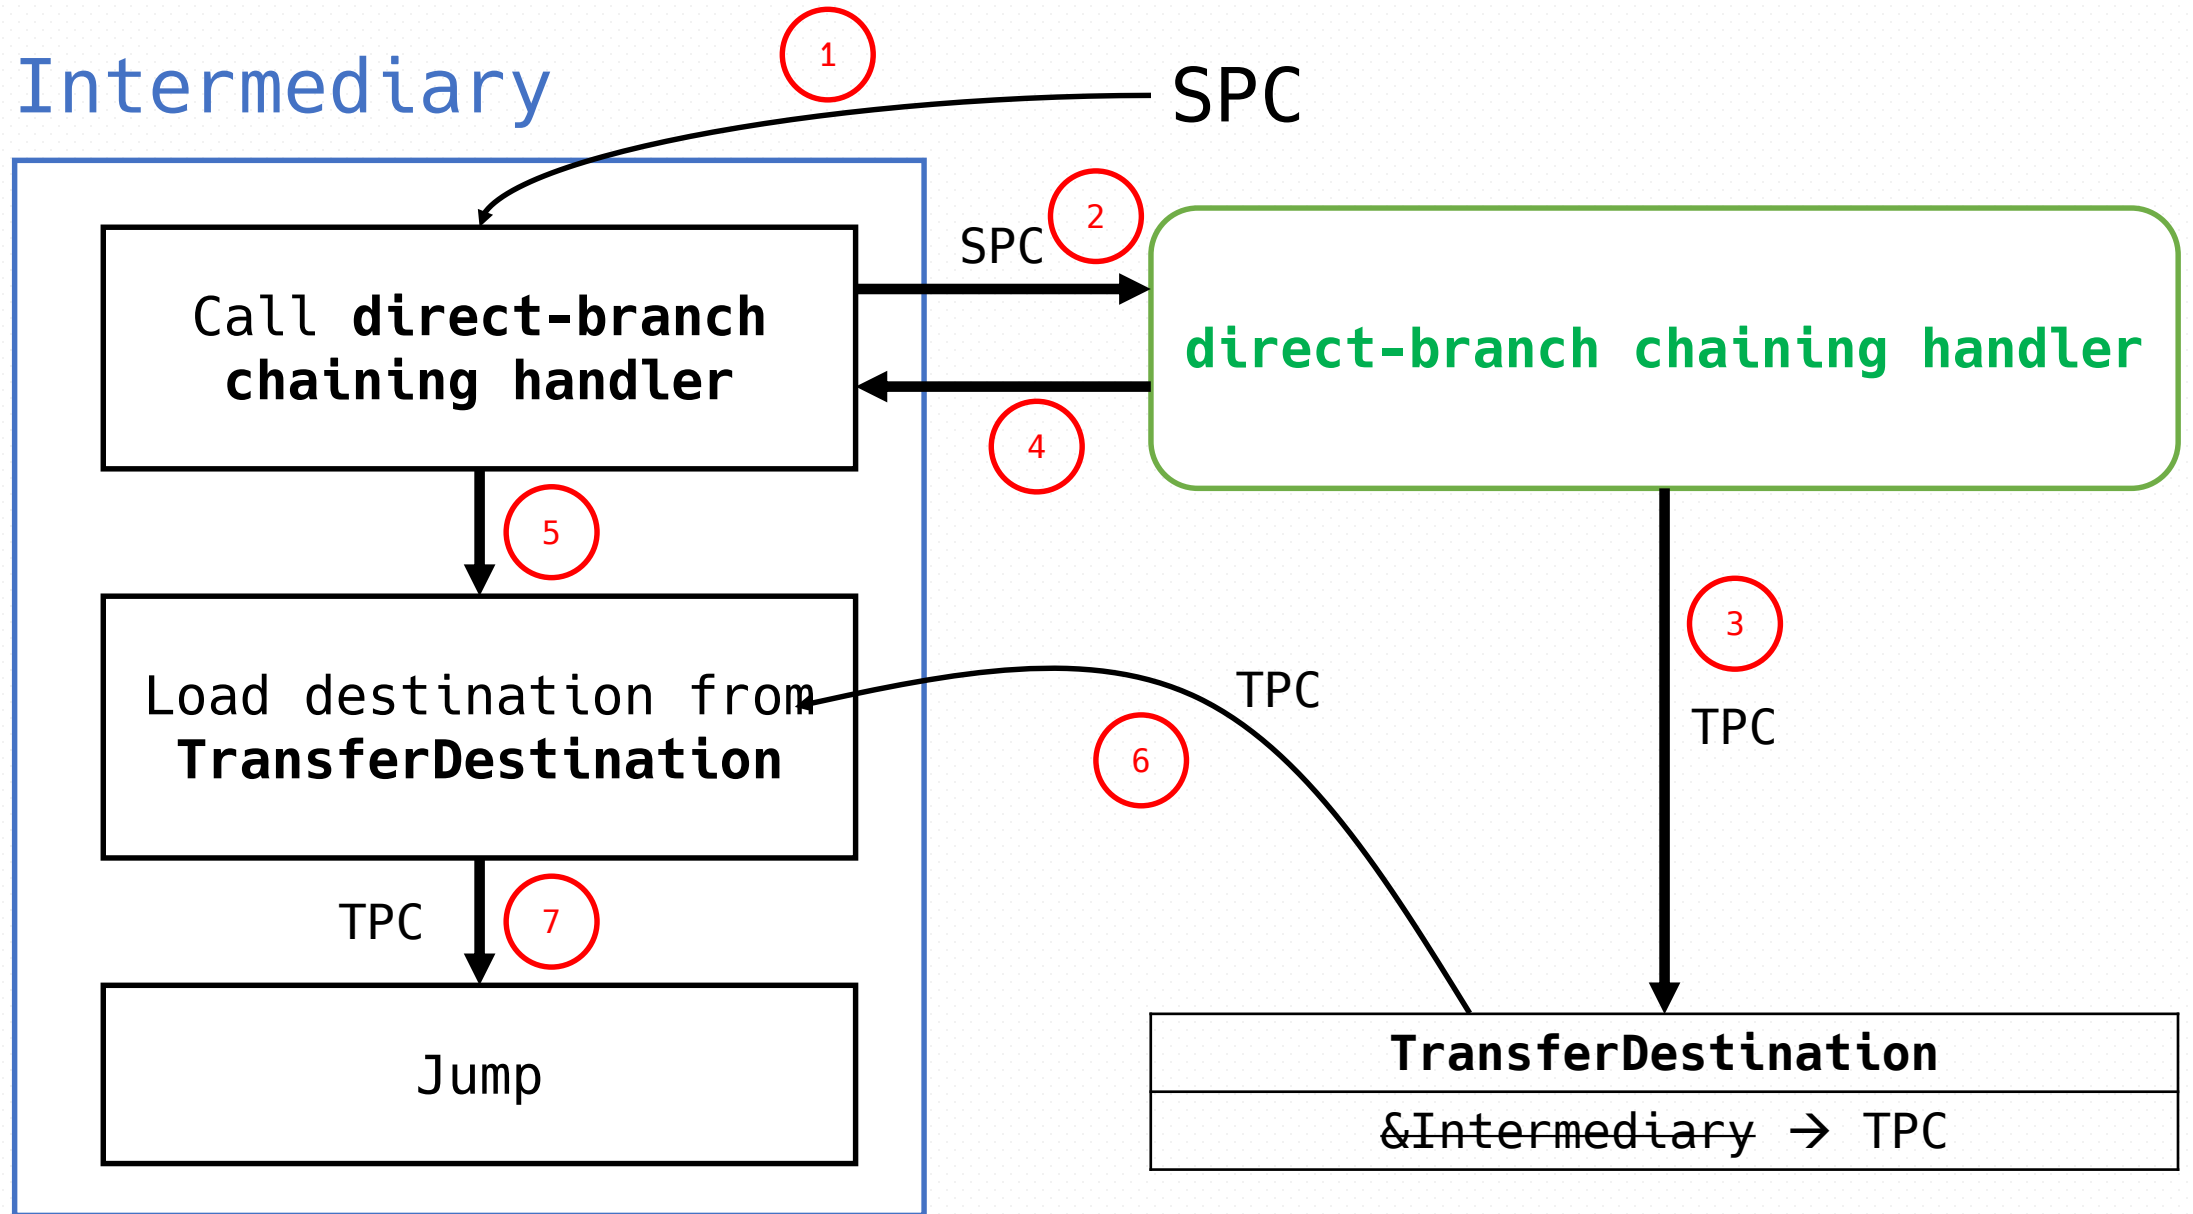
**TransferDestination**, for every TCB

■ TransferDestination is initialized with the address of  
**Intermediary**

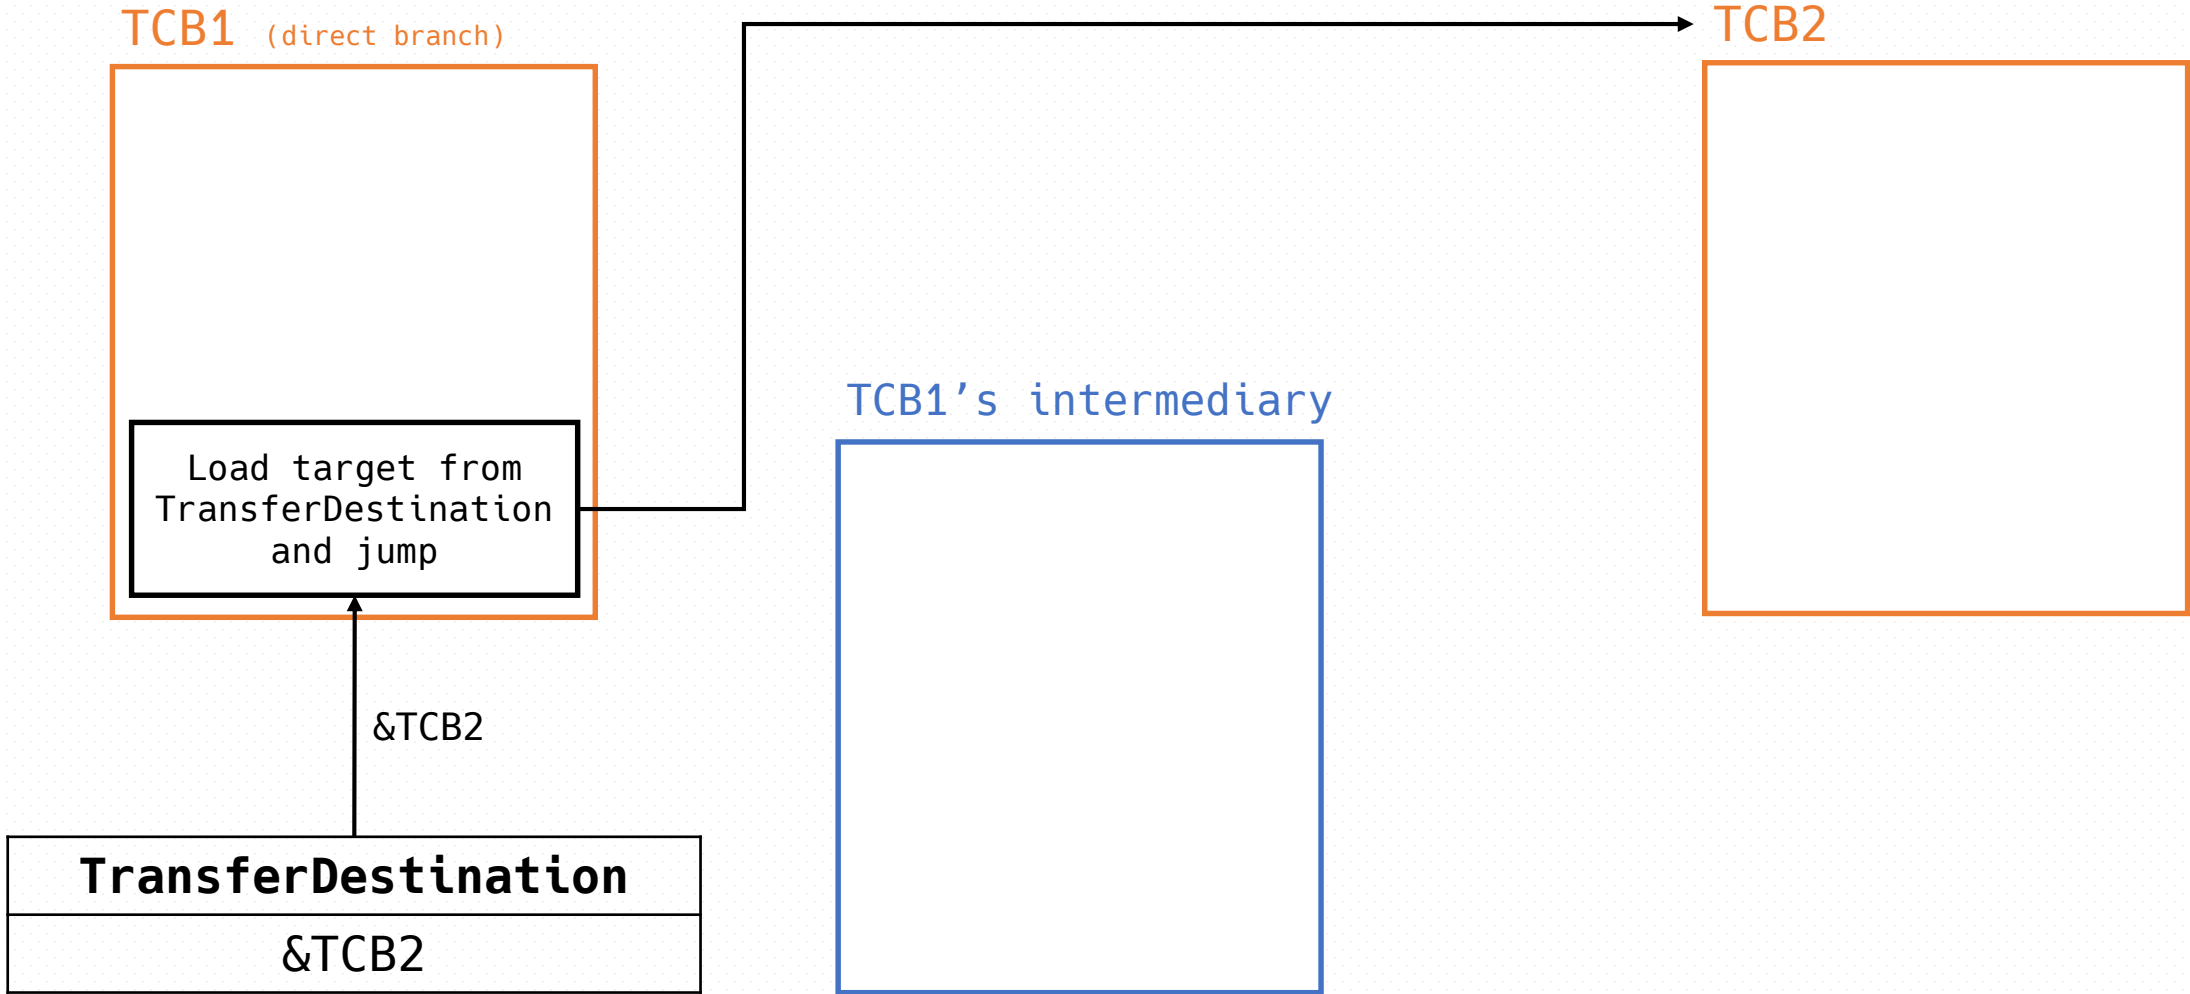
■ Intermediary calls direct-branch chaining handler to chain blocks and update TransferDestination with TPC

First run





Second run



# Indep-Hyperchaining For Indirect Branches

---

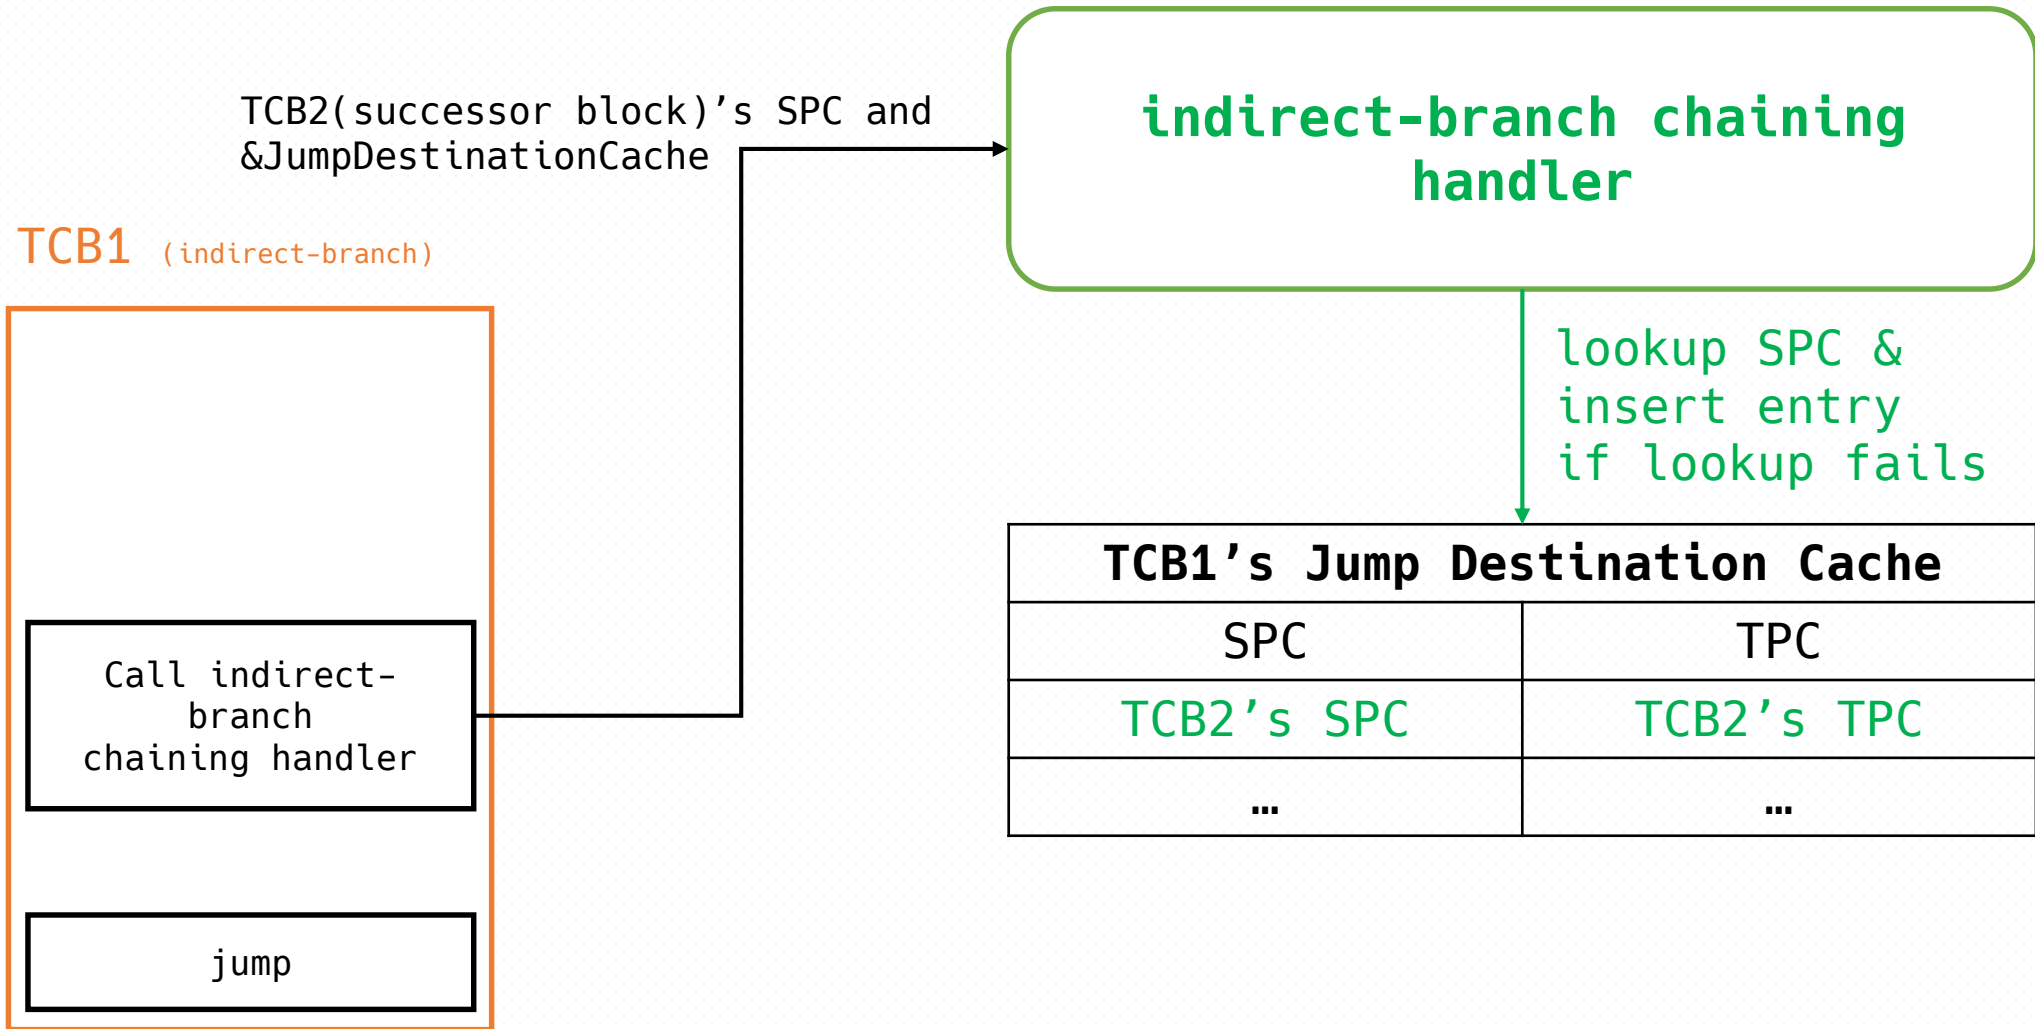
- For indirect branch, we have similar subroutine, **indirect-branch chaining handler**, to chain blocks with indirect branches
- Implemented in LLVM IR level

# Indep-Hyperchaining For Indirect Branches

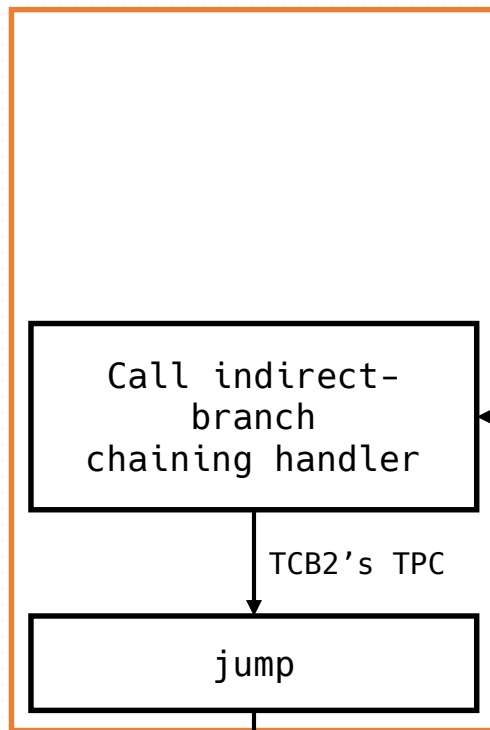
---

## ■ Steps

- TCB calls indirect-branch chaining handler and pass SPC to handler
- Chaining handler lookups SPC in **Jump Destination Cache** first. If nothing is found, do address lookup to acquire corresponding TPC and store address translation pair entry <SPC, TPC> to that cache



TCB1 (indirect-branch)



indirect-branch chaining handler

TCB2's TPC

TCB2's TPC

TCB1's Jump Destination Cache

SPC	TPC
TCB2's SPC	TCB2's TPC
...	...




# Dep-Hyperchaining For Direct Branches

---

- The reasons why we develop **dep-hyperchaining** (**platform-dependent hyperchaining**)
  - We observe some redundant operations in **Indep-hyperchaining**
    - For direct branches, TCBs need to read the same TPC of successor's block every time the block executes
    - For indirect branches, one possible case is single destination indirect branch which makes jump destination cache inefficient
  - The design of **dep-hyperchaining** is based on the **indep-hyperchaining**

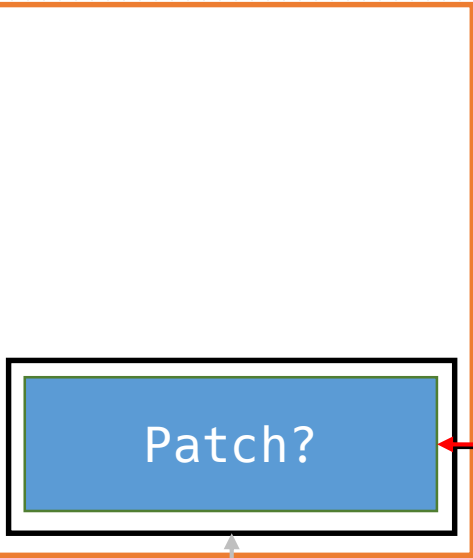
# Indep-hyperchaining for direct branches



How to simplify? Block chain optimization is a good idea but Rabbit doesn't support that.

# Indep-hyperchaining for direct branches

TCB1 (direct branch)



TCB2



&TCB2



How to simplify? ~~Block chain optimization is a good idea but Rabbit doesn't support that.~~ What if we directly patch to the TCB ?

# Challenges & Solutions

---

- Where is the instructions of load memory operations in TCBs
  - We need a anchor point to tell the run-time library the address of that instructions in TCBs
- How to patch the code. Is it practicable ?
  - We need to analyze possible generated code patterns

# Load Effective Address

---

## ■ LEA

- `lea` permits the program counter as its source operand

- `rip` minus the offset 7 (i.e. the length of `lea`) to get the address of `lea` instruction itself

- Store the address into the globally accessible **CPU context**

REX.W 8D ModRM F9 FF FF FF	<code>lea r64, QWORD PTR [rip - 0x7]</code>
----------------------------	---

# Store Address Into CPU Context

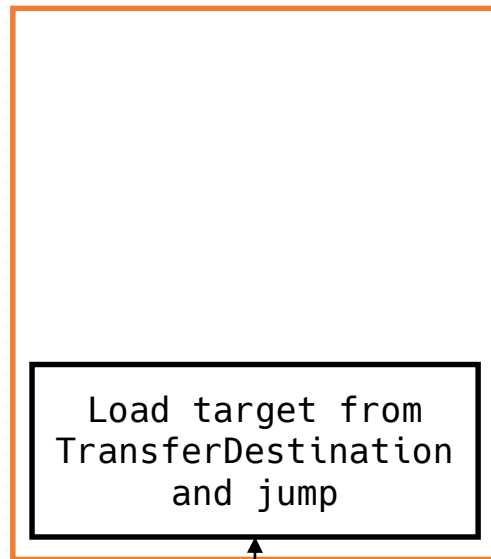
48 8D 05 F9 FF FF FF	<code>lea rax, QWORD PTR [rip-7]</code>
48 8B 0C 24	<code>mov rcx, QWORD PTR [rsp]</code>
48 89 41 0A	<code>mov QWORD PTR [rcx+10], rax</code>

CPU Context	
offset	field
...	...
0x10	anchor point

# Load TransferDestination's Address

48 B8 08 60 EA F7 FF 7F 00 00	<code>movabs rax, 0x7ffff7ea6008</code>
48 8B 00	<code>mov rax, QWORD PTR [rax]</code>

TCB (direct branch)



# Jump To Target

---

48 89 CF	<b>mov</b> rdi, rcx
FF E0	<b>jmp</b> rax

Every block contains only CPU context as a parameter to make code block be able to read and write simulated registers through CPU context.



# Code Pattern 1

---

- The aforementioned instruction sequences constitute code pattern 1 in our definition

48 8D 05 F9 FF FF FF	<b>lea</b> rax, <b>QWORD PTR</b> [rip-7]
48 8B 0C 24	<b>mov</b> rcx, <b>QWORD PTR</b> [rsp]
48 89 41 0A	<b>mov</b> <b>QWORD PTR</b> [rcx+10], rax
48 B8 08 60 EA F7 FF 7F 00 00	<b>movabs</b> rax, 0x7ffff7ea6008
48 8B 00	<b>mov</b> rax, <b>QWORD PTR</b> [rax]
48 89 CF	<b>mov</b> rdi, rcx
FF E0	<b>jmp</b> rax

# Then, How To Patch

---

- The intuitive way is patching the last instruction `“jmp r64”`
- Unfortunately, it is impracticable
  - Patch out-of-bound
  - Conditional branch block

# Jump Patch

---

- A minimal jump patch is a direct-jump patch
  - “JMP rel32” takes 5 bytes
  - “JMP rel8” is useless here because tiny displacement is rare to appear. We don't consider such case



Opcode	Instruction	Op/En	64-Bit Mode	Compat/Leg Mode	Description
EB <i>cb</i>	JMP <i>rel8</i>	D	Valid	Valid	Jump short, RIP = RIP + 8-bit displacement sign extended to 64-bits
E9 <i>cw</i>	JMP <i>rel16</i>	D	N.S.	Valid	Jump near, relative, displacement relative to next instruction. Not supported in 64-bit mode.
E9 <i>cd</i>	JMP <i>rel32</i>	D	Valid	Valid	Jump near, relative, RIP = RIP + 32-bit displacement sign extended to 64-bits

From Intel® 64 and IA-32 Architectures Software Developer's Manual

# Patch Out-of-bound

---

- A minimal patch, “`jmp rel32`”, takes 5 bytes
- However, last instruction “`jmp r64`” only provides 2 bytes space to patch



Patch 3 more bytes !!

# Conditional Branch Blocks

---

- “taken” and “not taken” blocks are contiguous in the run-time memory
  - There are no reserved bytes between two TCBs
  - Patch out-of-bound problem implies “not taken” block get overwritten at the first few bytes and becomes corrupt

# Patch

---


- For such concerns, we don't patch last instruction
- We pick "**movabs** rax, 0x7ffff7ea6008" as a proper instruction to patch
  - **movabs** offers 10 bytes space that is quite enough

# Patched TCB for code pattern 1

Before

48 8D 05 F9 FF FF FF	<code>lea rax, QWORD PTR [rip-7]</code>
48 8B 0C 24	<code>mov rcx, QWORD PTR [rsp]</code>
48 89 41 0A	<code>mov QWORD PTR [rcx+10], rax</code>
48 B8 08 60 EA F7 FF 7F 00 00	<code>movabs rax, 0x7ffff7ea6008</code>
48 8B 00	<code>mov rax, QWORD PTR [rax]</code>
48 89 CF	<code>mov rdi, rcx</code>
FF E0	<code>jmp rax</code>

After

48 8D 05 F9 FF FF FF	<code>lea rax, QWORD PTR [rip-7]</code>
48 8B 0C 24	<code>mov rcx, QWORD PTR [rsp]</code>
48 89 41 0A	<code>mov QWORD PTR [rcx+10], rax</code>
48 89 CF	<code>mov rdi, rcx</code>
E9 XX XX XX XX	<code>jmp &lt;TPC&gt; (successor block)</code> 
00 00	??? (don't care; unreachable)
48 8B 00	<code>mov rax, QWORD PTR [rax]</code>
48 89 CF	<code>mov rdi, rcx</code>
FF E0	<code>jmp rax</code>

# Indirect-Branch Patch

---

- If the displacement is too large to be encoded as a direct-jump patch, we use **indirect-branch patch**

<code>movabs rax, &lt;TPC&gt;</code>
<code>jmp rax</code>



# Patched TCB For Code Pattern 1

Before

48 8D 05 F9 FF FF FF	<b>lea</b> rax, <b>QWORD PTR</b> [rip-7]
48 8B 0C 24	<b>mov</b> rcx, <b>QWORD PTR</b> [rsp]
48 89 41 0A	<b>mov</b> <b>QWORD PTR</b> [rcx+10], rax
48 B8 08 60 EA F7 FF 7F 00 00	<b>movabs</b> rax, 0x7ffff7ea6008
48 8B 00	<b>mov</b> rax, <b>QWORD PTR</b> [rax]
48 89 CF	<b>mov</b> rdi, rcx
FF E0	<b>jmp</b> rax

After

48 8D 05 F9 FF FF FF	<b>lea</b> rax, <b>QWORD PTR</b> [rip-7]
48 8B 0C 24	<b>mov</b> rcx, <b>QWORD PTR</b> [rsp]
48 89 41 0A	<b>mov</b> <b>QWORD PTR</b> [rcx+10], rax
48 B8 XX XX XX XX XX XX XX XX	<b>movabs</b> rax, <TPC> (replaced by successor block)
0F 1F 00	<b>nop</b> <b>DWORD PTR</b> [rax](disable dereference)
48 89 CF	<b>mov</b> rdi, rcx
FF E0	<b>jmp</b> rax



# Code Pattern 2

---

- Similar to code pattern 1 but has less instructions
- **rdi** still holds **&CPUContext**

48 8D 05 F9 FF FF FF	<b>lea</b> rax, <b>QWORD PTR</b> [rip-7]
48 89 41 0A	<b>mov</b> <b>QWORD PTR</b> [rcx+10], rax
48 B8 08 60 EA F7 FF 7F 00 00	<b>movabs</b> rax, 0x7ffff7ea6008
48 8B 00	<b>mov</b> rax, <b>QWORD PTR</b> [rax]
FF E0	<b>jmp</b> rax


# Patched TCB For Code Pattern 2

---

**Before**

48 8D 05 F9 FF FF FF	<code>lea rax, QWORD PTR [rip-7]</code>
48 89 41 0A	<code>mov QWORD PTR [rcx+10], rax</code>
48 B8 08 60 EA F7 FF 7F 00 00	<code>movabs rax, 0x7ffff7ea6008</code>
48 8B 00	<code>mov rax, QWORD PTR [rax]</code>
FF E0	<code>jmp rax</code>

**After**


E9 XX XX XX XX	<code>jmp &lt;TPC&gt; (successor block)</code> 
FF FF	??
48 89 41 0A	<code>mov QWORD PTR [rcx+10], rax</code>
48 B8 08 60 EA F7 FF 7F 00 00	<code>movabs rax, 0x7ffff7ea6008</code>
48 8B 00	<code>mov rax, QWORD PTR [rax]</code>
FF E0	<code>jmp rax</code>

# Dep-Hyperchaining For Indirect Branches

---

- According to our study of SPEC CPU 2006 CINT benchmarks
  - 56% (avg.) of the indirect branches in the source binary are function-return instructions
  - 72.88% (avg.) of the indirect branches which are not function return have only a single jump destination
    - jump-destination cache is unnecessary in single destination case

SPEC CPU CINT 2006 (the percentage of indirect branches that have only one destination. Function returns are exclude in this calculation).



# Design

---

- The code design need to consider situations for both single destination and multiple destinations
- Change the code behavior dynamically
  - Transform single destination → multiple destinations by patching code

# Solutions

---

- The code design need to consider situations for both single destination and multiple destinations
  - Single destination : **Code Cave**

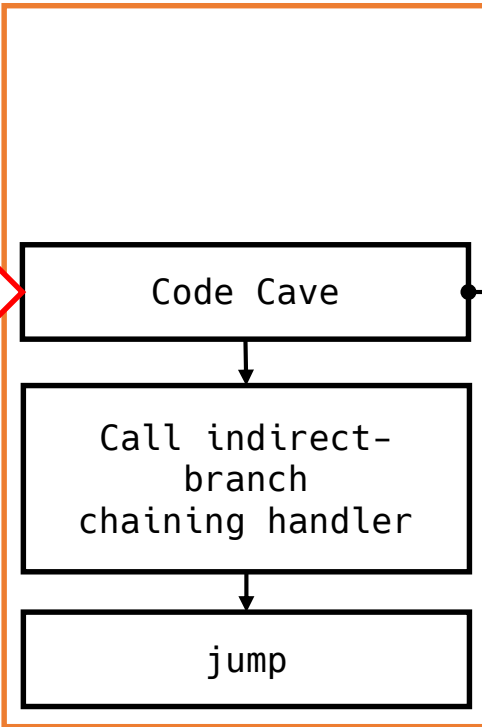
# Code Cave

---

- A **code cave** is one or several consecutive instructions in which some fields are intentionally filled with dummy values in the beginning
- Later these fields are patched with data related to the indirect branch, such as SPC, TPC, etc.



## TCB1 (indirect-branch)

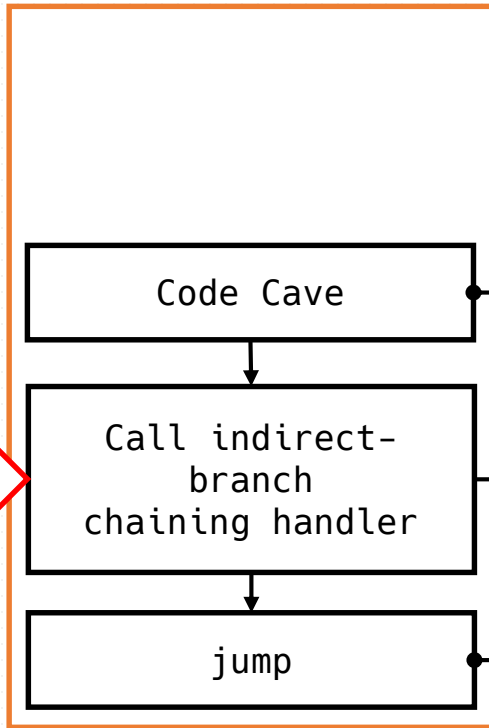


### Code Cave (initial state)

48 8D 15 00 00 00 00	<b>lea</b> rdx, <b>QWORD PTR</b> [rip]
48 81 FE FF FF FF 7F	<b>cmp</b> rsi, 0x7fffffff
48 0F 44 FF	<b>cmove</b> rdi, rdi
0F 1F 80 00 02 00 00	<b>nop dword ptr</b> [rax+0x200]
0F 1F 80 00 02 00 00	<b>nop dword ptr</b> [rax+0x200]

= nop  
= nop

# TCB1 (indirect-branch)



• Patched Code Cave (direct jump patch; α-patch)

48 8D 15 00 00 00 00	lea rdx, QWORD PTR [rip]
48 81 FE 00 35 01 00	cmp rsi, 0x13500
49 0F 44 FE	cmovbe rdi, r14
0F 84 XX XX XX XX	je 0x2000000
0F 1F 84 00 00 02 00 00	nop DWORD PTR [rax+rax*1+0x200]



SPC = 0x13500



Patch

```

mov rdi, r14
jmp rax
  
```


Analyze

```

0F 84 cd | JE rel32
  
```

From Intel® 64 and IA-32 Architectures Software Developer's Manual

# TCB1 (indirect-branch)



## Patched Code Cave (indirect jump patch; β-patch)

48 8D 15 00 00 00 00	lea rdx, QWORD PTR [rip]
48 81 FE 00 35 01 00	cmp rsi, 0x13500
49 0F 44 FE	cmovbe rdi, r14
75 0C	jne failure
48 B8 00 00 00 02 00 00 00 00	movabs rax, 0x20000000
FF E0	jmp rax
	failure:



SPC = 0x13500



Patch

```

mov rdi, r14
jmp rax
  
```

Analyze

```

75 cb | JNE rel8
  
```

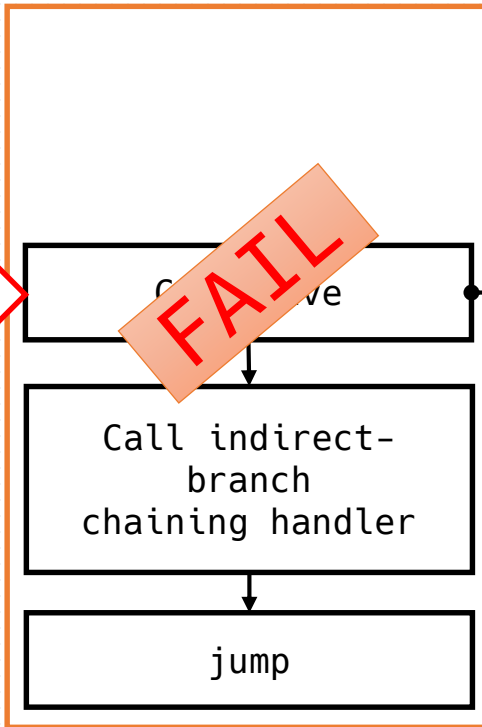
From Intel® 64 and IA-32 Architectures Software Developer's Manual

# Solutions

---

- The code design need to consider situations for both single destination and multiple destinations
  - Single destination : **Code Cave**
  - Multiple destinations : **Dynamic Address Mapping Table**
- Change the code behavior dynamically
  - **Repatch the code cave**

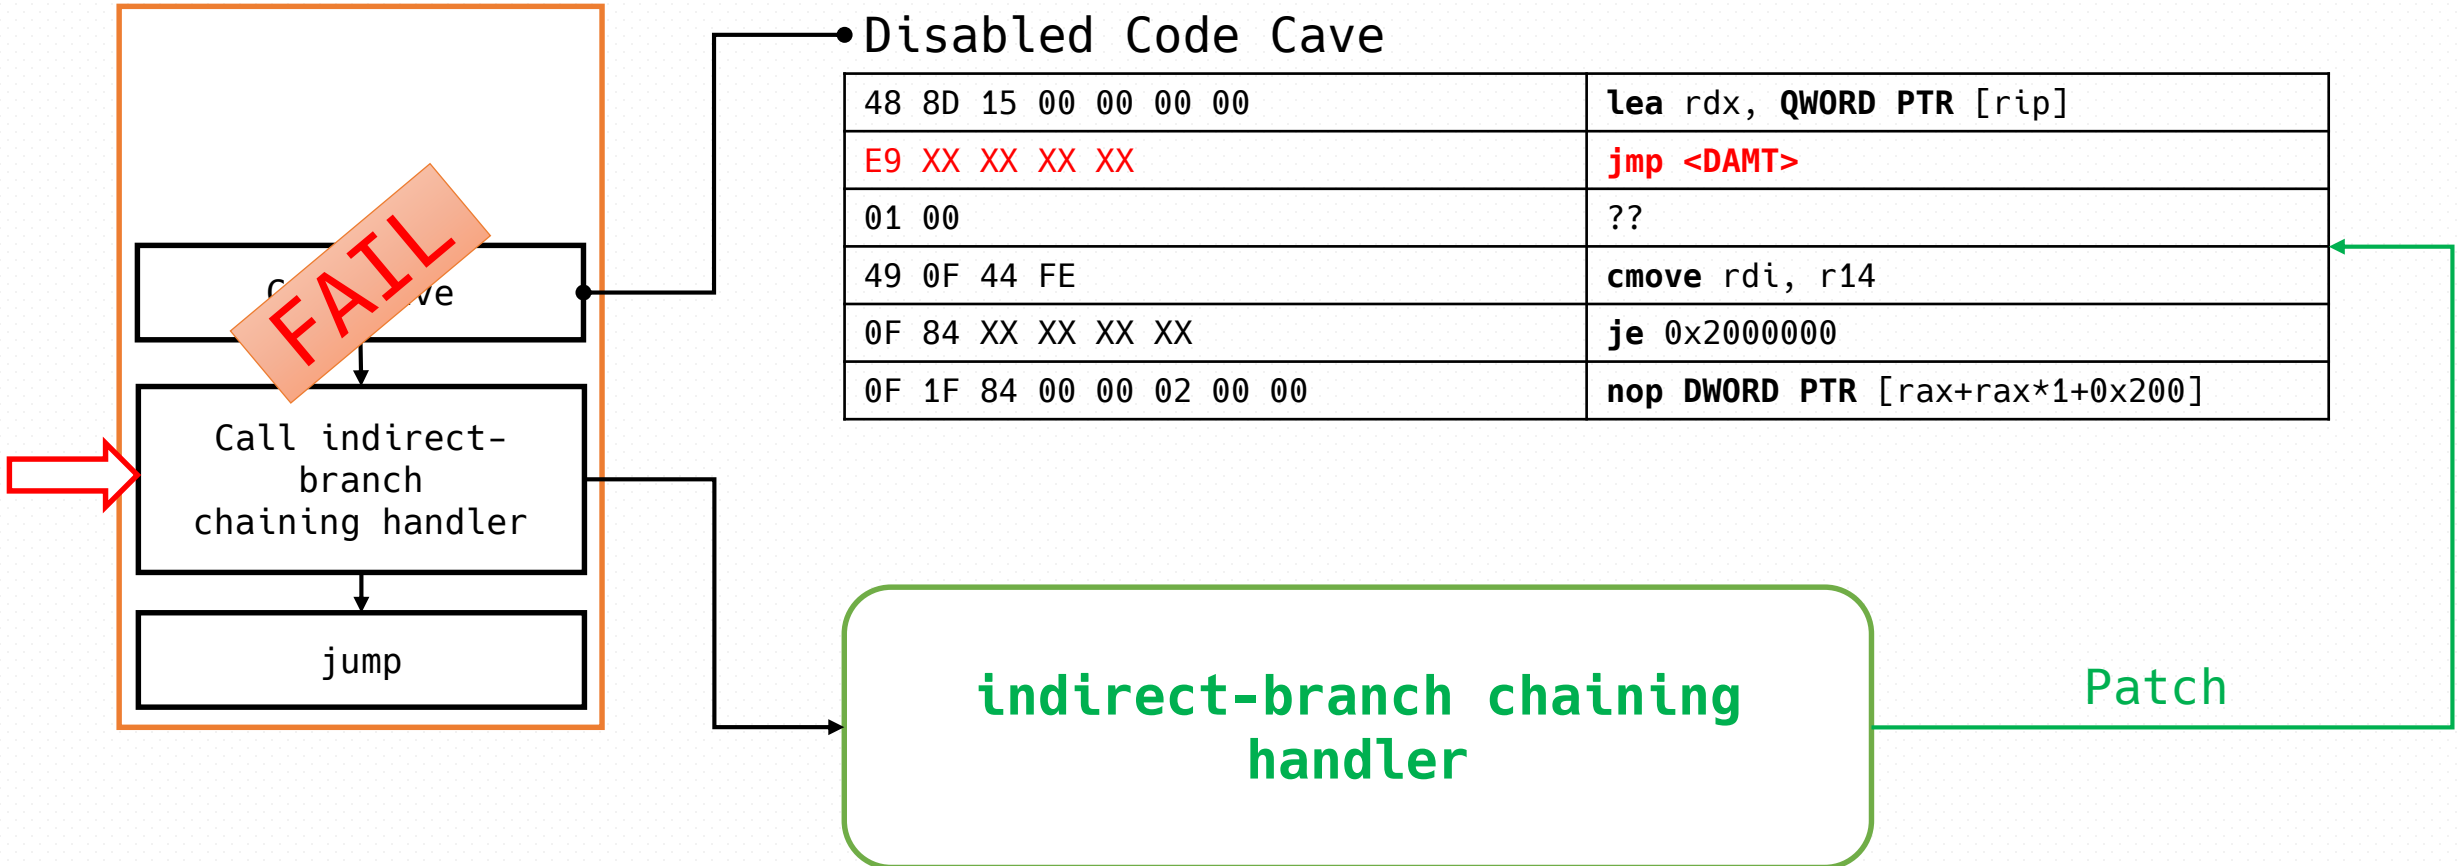
## TCB1 (indirect-branch)



### • Patched Code Cave (direct jump patch; $\alpha$ -patch)

48 8D 15 00 00 00 00	<b>lea</b> rdx, <b>QWORD PTR</b> [rip]
48 81 FE 00 35 01 00	<b>cmp</b> rsi, 0x13500
49 0F 44 FE	<b>cmov</b> e rdi, r14
0F 84 XX XX XX XX	<b>je</b> 0x2000000
0F 1F 84 00 00 02 00 00	<b>nop</b> <b>DWORD PTR</b> [rax+rax*1+0x200]


# TCB1 (indirect-branch)



We find that this indirect branch has more than one destination


# TCB1 (indirect-branch)

Rabbit allocates a block of memory for dynamic address mapping table



48 8D 15 00 00 00 00	lea rdx, QWORD PTR [rip]
E9 XX XX XX XX	jmp <DAMT>
01 00	??
49 0F 44 FE	cmove rdi, r14
0F 84 XX XX XX XX	je 0x20000000
0F 1F 84 00 00 02 00 00	nop DWORD PTR [rax+rax*1+0x200]

# TCB1 (indirect-branch)



TCB1's dynamic address mapping table

48 8D 15 00 00 00 00	lea rdx, QWORD PTR [rip]
E9 XX XX XX XX	jmp <DAMT>
01 00	??
49 0F 44 FE	cmovbe rdi, r14
0F 84 XX XX XX XX	je 0x2000000
0F 1F 84 00 00 02 00 00	nop DWORD PTR [rax+rax*1+0x200]
	lookup_fail_in_DAMT:

Lookup failed in DAMT



# Dynamic Address Mapping Table

---

```
DAMT:  
cmp rsi, 0x10000  
je tag1  
jmp <code cave>
```

```
tag1:  
mov rdi, r14  
movabs rax, 0x7ffff0000000  
jmp rax
```

Visited Destination	
SPC	TPC
0x10000	0x7ffff0000000

# Dynamic Address Mapping Table

```
DAMT:  
cmp rsi, 0x10000  
je tag1  
cmp rsi, 0x11000  
je tag2  
jmp <code cave>
```

```
tag1:  
mov rdi, r14  
movabs rax, 0x7ffff0000000  
jmp rax  
tag2:  
mov rdi, r14  
movabs rax, 0x7ffff1000000  
jmp rax
```

Visited Destination	
SPC	TPC
0x10000	0x7ffff0000000
0x11000	0x7ffff1000000

# Outline

---


- Background
- Hyperchaining
- **Experiments**
- Conclusion

# Experimental Setup


---

- CPU : Intel Core i7-7700 3.60GHz
- OS : Ubuntu 18.04.3 LTS with Linux kernel 4.15.0
- RAM : 16GB
- LLVM : 8.0.1
- Benchmark : SPEC CPU 2006 CINT


SPEC CPU CINT 2006 (platform-dependent hyperchaining for direct branches)




SPEC CPU CINT 2006 (platform-dependent hyperchaining for indirect branches)



SPEC CPU CINT 2006 (Memory usage for direct branch hyperchaining)



SPEC CPU CINT 2006 (Memory usage for indirect branch hyperchaining)





# Outline

---

- Background
- Hyperchaining
- Experiments
- **Conclusion**

# Conclusion

---

- This paper proposes platform-independent hyperchaining and platformdependent hyperchaining on the x86-64 platform
- The experimental results show that platform-dependent hyperchaining can reach 1.08x and 1.05x speedup compared to platform-independent hyperchaining

# Conclusion

---

- Memory usage of platform-dependent hyperchaining is 1.0089x and 1.019x times that of platform-independent hyperchaining for direct branches and indirect branches, respectively
- It has small memory overhead to adopt platform-dependent approach

# Future Work

---

- The approach of platform-dependent hyperchaining for indirect branches implies potential opportunities to be improved on performance
  - Schedule the first section of the SPC section to reduce and mitigate the miss penalty in run-time
  - Change the mechanism of searching in DAMT (Hash table ?)
  - Lack of a policy to disable platform-dependent hyperchaining when the dynamic address mapping table is full

# Credits

---

- The hack font - authors of Source Foundry, which is licensed under MIT  
<https://github.com/source-foundry/Hack>

THANK YOU FOR YOUR ATTENTION