

Kernel Probes for MIPS, ARM and PPC32

Sony India Software Centre Sony Corporation

Madhvesh Sulibhavi (madhvesh.s@ap.sony.com)

Lubna Badkar (lubna@ap.sony.com)

Tim Bird (tim.bird@am.sony.com) Presenting

Contents

- Kernel Probes (Kprobes) Overview
- Kprobes Support for MIPS
- Kprobes Support for ARM
- Kprobes Support for PPC32 (BookE)
- Overhead Measurement Details
- Source Patches
- References

Kprobes Overview

- Kprobes is simple and lightweight mechanism to collect debugging information dynamically.
- Linux main line kernel 2.6.16 (and later) includes support for X86, PPC64, Sparc, IA64, x86_64
- Kprobes support for MIPS, ARM and PPC32 arch is implemented by Sony for vanilla kernel 2.6.16-39.

Kprobes Overview (Contd)

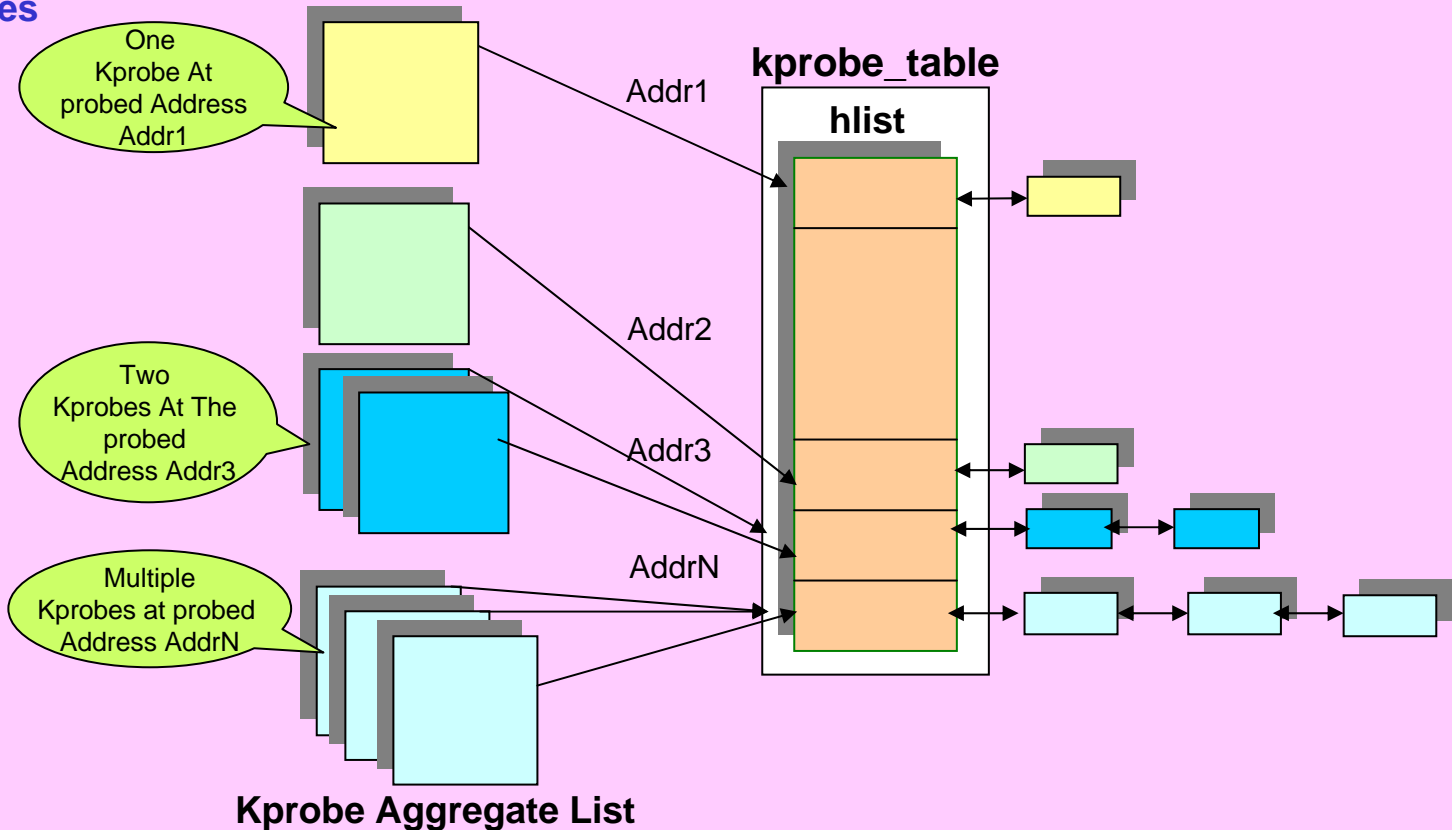
- Kprobe mainly contains following three handlers
- Pre-Handler
 - Executed before the probed instruction
 - It can be used to dump the register contents before executing the probed instruction
- Post-Handler
 - Executed after the probed instruction
 - It can be used to dump the register contents after executing the probed instructions
- Fault-Handler
 - Executed when some fault occurs in the pre handler or post handler or in the instruction being debugged

Kprobes Overview (Contd)

Includes:

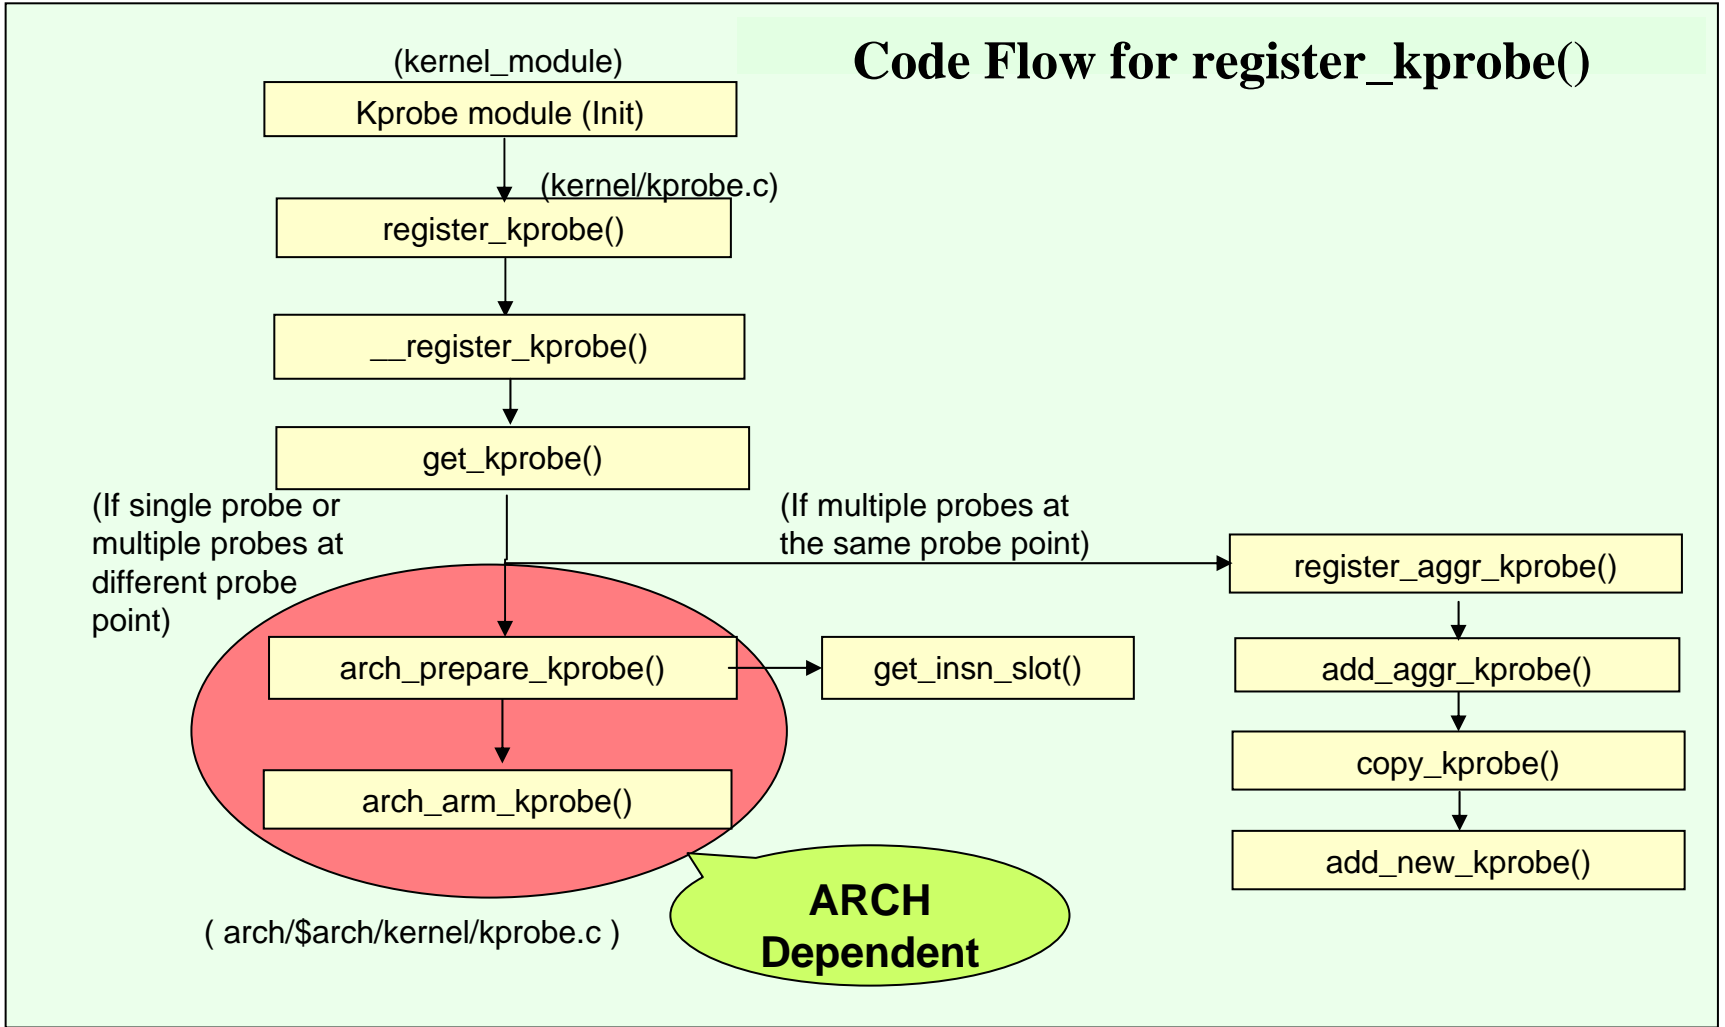
- Initializing the kprobes
- Register kprobes
- Deregister kprobes

Kprobes Management (Arch Independent)



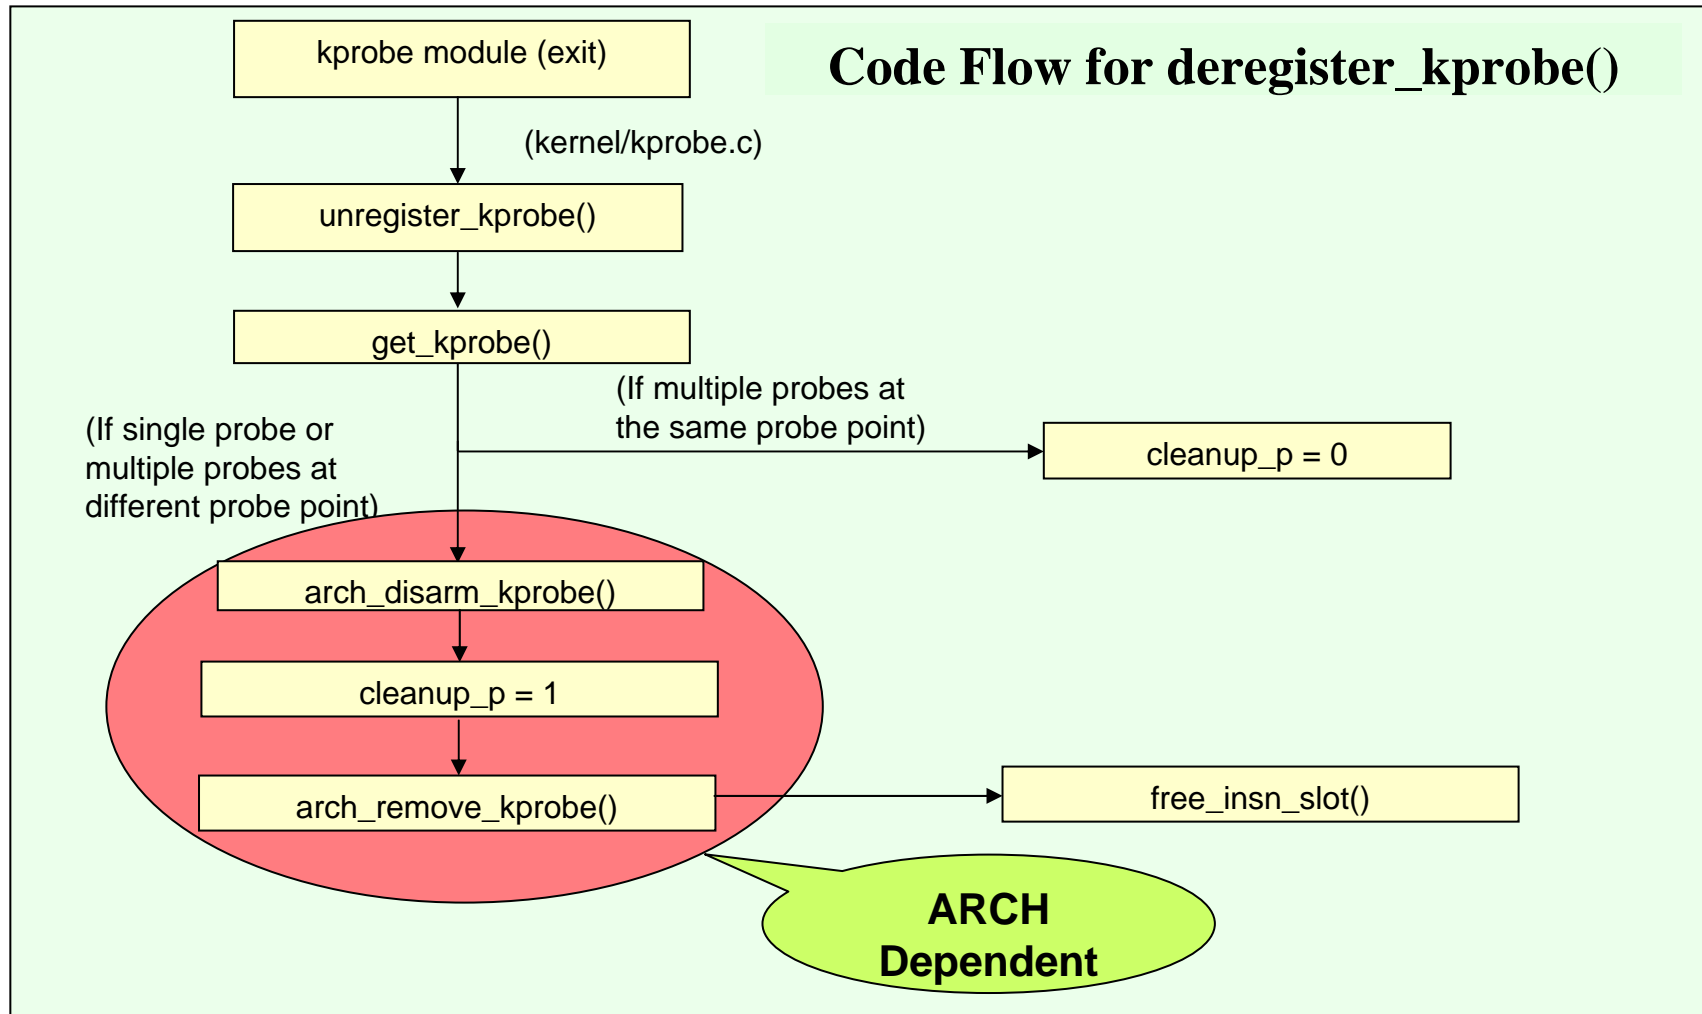
Kprobes Overview (Contd)

Code Flow for register_kprobe()



Kprobes Overview (Contd)

Code Flow for deregister_kprobe()



MIPS Arch Kprobes Contents

- Target Environment Details
- Kprobes Design
- Sample Example Output
- Sample Example Disassembly
- Limitations

MIPS Target Environment Details

- **Target Board**
 - Toshiba RBHMA4400 (TX4937)
- **OS**
 - Linux 2.6.16.39
- **Configuration**
 - `CONFIG_PREEMPT=y`
 - `CONFIG_KALLSYMS=y`
 - `CONFIG_KALLSYMSALL=y`
 - `CONFIG_DEBUG_KERNEL=y`



Hardware Features:
TX4937 @ 300Mhz.
System Clock @ 66Mhz
USB 1.1 Host Controller

MIPS Kprobes Design

- MIPS contains “break” instruction using which break point can be achieved.
- MIPS does not have hardware single step support and hence required alternate mechanism.
- Single step feature is implemented in software using “break” instruction with different code value
- MIPS Linux Kernel supports the following different break codes and two of these are used to implement Kprobes

MIPS Kprobes Design (Contd)

Opcode		Use
BRK_USERBP	0	User bp (used by debuggers)
BRK_KERNELBP	1	Break in the kernel
BRK_ABORT	2	Sometimes used by abort(3) to SIGIOT
BRK_BD_TAKEN	3	For bd slot emulation - not implemented
BRK_BD_NOTTAKEN	4	For bd slot emulation - not implemented
BRK_SSTEPBP	5	User bp (used by debuggers)
BRK_OVERFLOW	6	Overflow check
BRK_DIVZERO	7	Divide by zero check
BRK_RANGE	8	Range error check
BRK_STACKOVERFLOW	9	For Ada stack checking
BRK_NORLD	10	No rld found - not used by Linux/MIPS
_BRK_THREADBP	11	For threads, user bp (used by debuggers)
BRK_BUG	512	Used by BUG()
BRK_KDB	513	Used in KDB_ENTER()
BRK_MULOVF	1023	Multiply overflow


Linux Kernel Supported
Break Codes in
asm-mips/break.h

This is Used for Kprobe
to Implement BreakPoint

This is Used for Kprobe
to Implement Single Step
Operation

MIPS Kprobes Design (Contd)

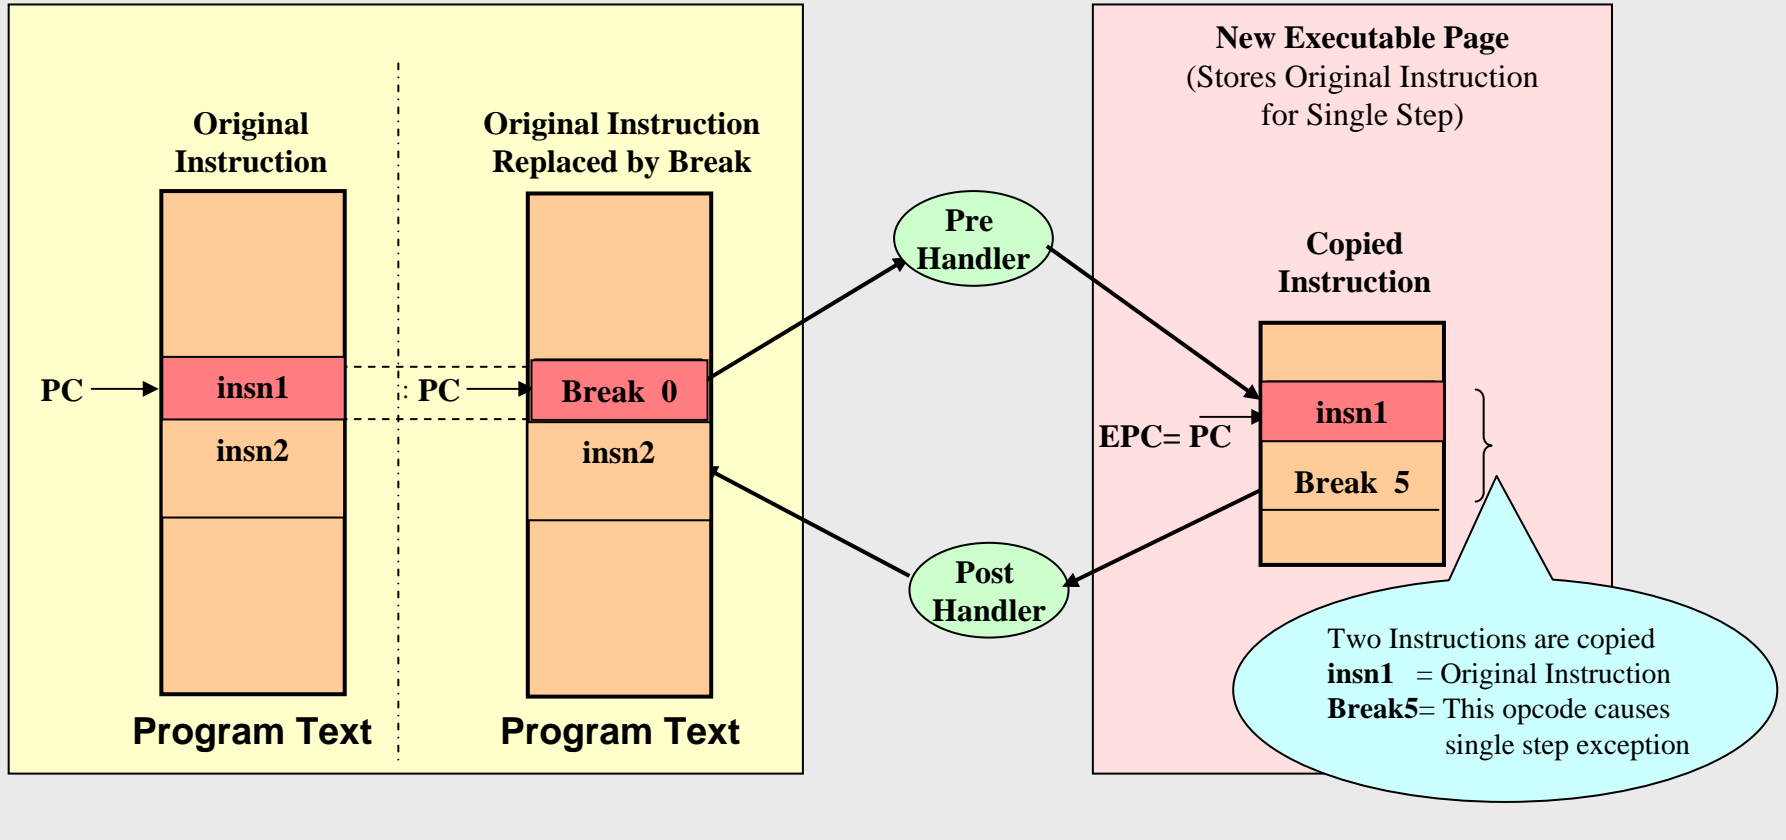
- Break instruction format is



- When exception is generated, Exception Program Counter (EPC) register holds the instruction which raised the exception. Using the EPC value, the “code” field is extracted to check whether the exception is of BREAK0 or BREAK5 type
- By using different values for **CODE** field, the break point and single step operations are achieved

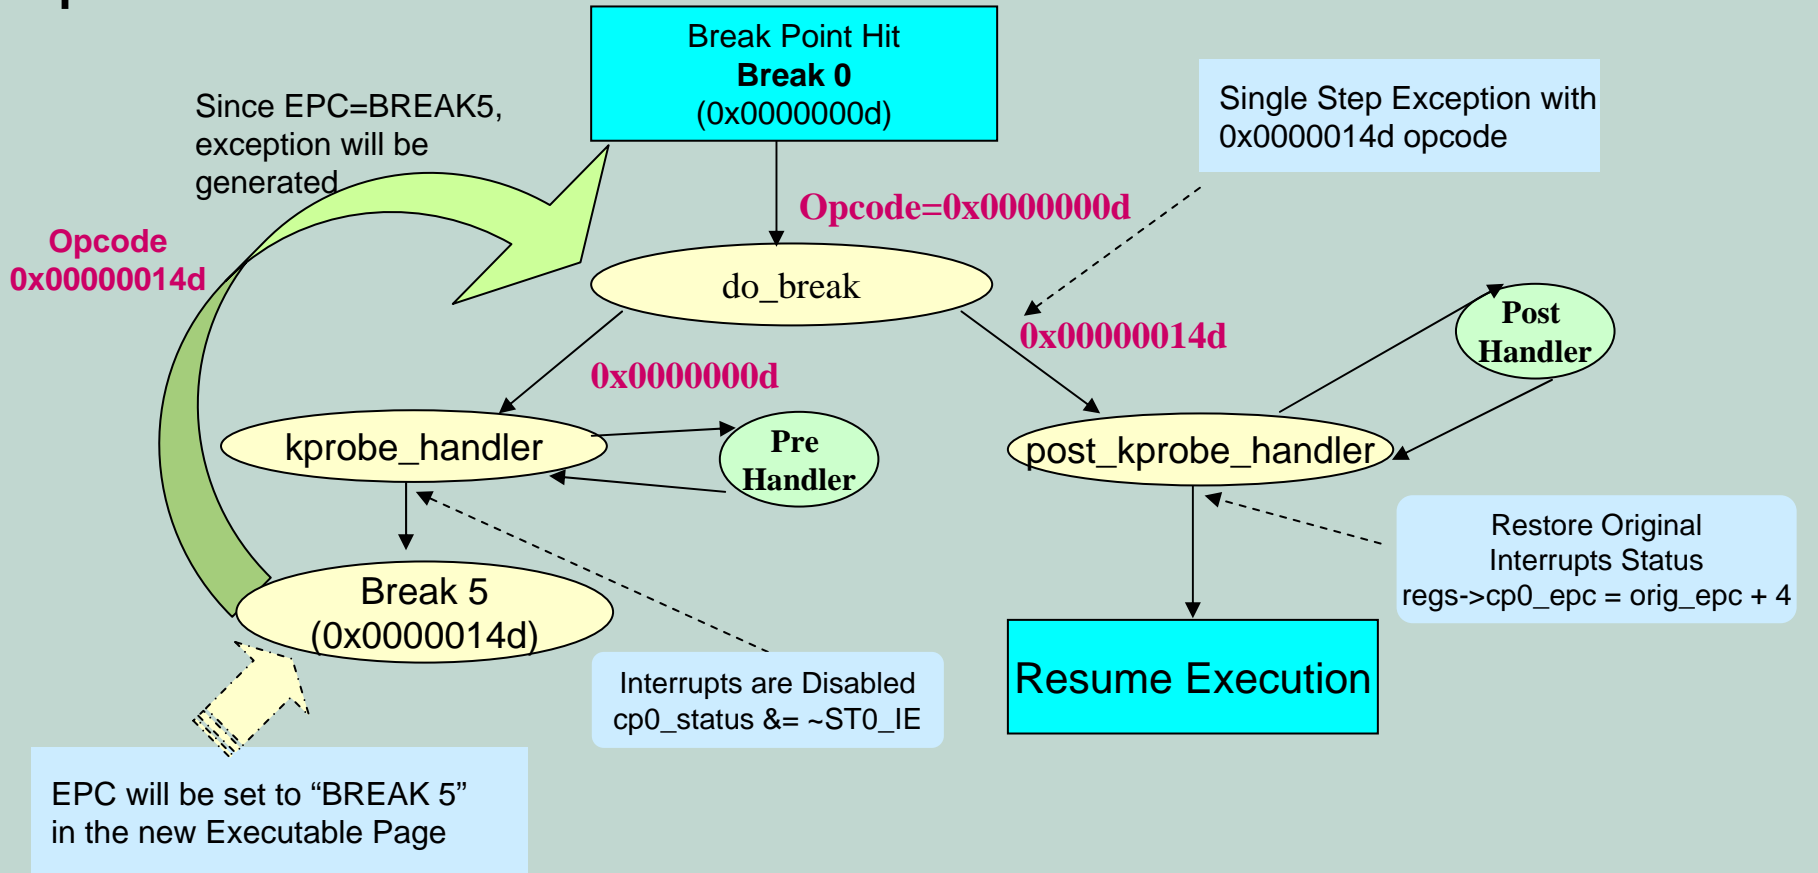
MIPS Kprobes Design (Contd)

Kprobes Control Flow



MIPS Kprobes Design (Contd)

Kprobes Execution Flow



MIPS Kprobes Design (Contd)

- Kprobes Control Flow Details
 - When a particular instruction “insn1” is probed, the original instruction will be copied on the executable page and replaces the probed instruction by ‘break 0’ i.e 0x0000000d
 - ‘break 5’ (0x0000014d) will be copied on the executable page after the original probed instruction
 - A breakpoint exception occurs, immediately and unconditionally transferring control to the exception handler
 - When break 0 is hit, kprobes exception notifier will be called by the exception handler do_break() with the die value DIE_BREAK. This will in turn call kprobe_handler() which performs executing pre_handler() associated with kprobes

MIPS Kprobes Design (Contd)

- Kprobes Control Flow Details (Contd)
 - `kprobe_handler()` takes care of moving program counter to the copied instruction by calling `prepare_single_step()`.
 - After executing the probed instruction, the next instruction on the copied page 'break 5' will be executed
 - When break 5 is hit, again kprobes exception notifier will be called by the exception handler `do_break()` with the die value `DIE_SSTEPBP`. This will in turn call user defined `post_handler()`.
 - After the instruction debugging, the execution resumes to the next instruction "insn2"

MIPS Sample Example Output

Kernel Module With do_fork as Probe Point

Kprobe Module

Kp.addr=do_fork
 kp.pre_handler=before_hook
 kp.post_handler=after_hook

```
before_hook() {
    Stack_dump()
    Show_allregs()
}
```

```
after_hook() {
    Stack_dump()
    Show_allregs()
}
```

Stack dump

Registers

Stack dump

Registers

Call Trace:

```
[<c0020508>] before_hook+0x5c/0x94 [k_001]
[<80046bb8>] do_fork+0x0/0x274
[<8021025c>] kprobe_exceptions_notify+0x35c/0x410
[<8004eac4>] getnstimeofday+0x18/0x4c
[<80210360>] notifier_call_chain+0x38/0x70
[<8004eac4>] getnstimeofday+0x18/0x4c
[<8020fc8c>] do_break+0x94/0xc8
[<8006758c>] hrtimer_run_queues+0x40/0x1a4
[<80187608>] process_backlog+0x118/0x318
[<80023150>] handle_bp_int+0x24/0x34
[<80027b40>] sys_fork+0x0/0x2c
[<80027fcc>] _sys_fork+0x34/0x40
[<80046bb8>] do_fork+0x0/0x274
[<8002c850>] stack_done+0x20/0x40
```

Call Trace:

```
[<c0020000>] show_allregs+0x0/0x4ac [k_001]
[<c002059c>] after_hook+0x5c/0xa0 [k_001]
[<80050000>] insert_resource+0x70/0x148
[<8021024c>] kprobe_exceptions_notify+0x34c/0x410
[<80210360>] notifier_call_chain+0x38/0x70
[<c0020524>] before_hook+0x78/0x94 [k_001]
[<8020fc8c>] do_break+0x94/0xc8
[<8004eac4>] getnstimeofday+0x18/0x4c
[<80023150>] handle_bp_int+0x24/0x34
[<80027b40>] sys_fork+0x0/0x2c
[<80027fcc>] _sys_fork+0x34/0x40
[<80027fcc>] _sys_fork+0x34/0x40
[<8002c850>] stack_done+0x20/0x40
[<8002c850>] stack_done+0x20/0x40
```

Cpu 0

```
$ 0 : 00000000 10008401 00000002 2ac5e5f0
$ 4 : 00000012 7f9b0970 87e43f30 00000000
$ 8 : 00000000 80000010 80027b40 ffffffff
$12 : ffffffff 00000000 00000001 0000005d
$16 : 00000000 7f9b0a28 1000d598 00000000
$20 : 1002a468 00000000 00000000 00000000
$24 : 00000000 00000000
$28 : 87e42000 87e43f10 1002a0a8 80027fcc
Hi : 0000000c
Lo : 00000005
epc : 80046bb8 do_fork+0x0/0x274
ra : 80027fcc _sys_fork+0x34/0x40
Status: 10008403 KERNEL EXL IE
Cause : 10000024
PrId : 00002d30
```

Cpu 0

```
$ 0 : 00000000 10008401 00000002 2ac5e5f0
$ 4 : 00000012 7f9b0970 87e43f30 00000000
$ 8 : 00000000 80000010 80027b40 ffffffff
$12 : ffffffff 00000000 00000001 0000005d
$16 : 00000000 7f9b0a28 1000d598 00000000
$20 : 1002a468 00000000 00000000 00000000
$24 : 00000000 00000000
$28 : 87e42000 87e43eb8 1002a0a8 80027fcc
Hi : 0000000c
Lo : 00000005
epc : c00d0004 0xc00d0004
ra : 80027fcc _sys_fork+0x34/0x40
Status: 10008402 KERNEL EXL
Cause : 10000024
```

MIPS Example Disassembly

When kprobe is registered, original instruction is replaced with 'break' instruction '0000000d'

When probe point is hit, pre handler gets executed and 'pc' will be made to point to original instruction copied on executable page

00002b3c <do_fork>:

2b3c:	27bdffb0	addiu	sp,sp,-80
2b40:	afb40040	sw	s4,64(sp)
2b44:	afb3003c	sw	s3,60(sp)
2b48:	afb20038	sw	s2,56(sp)
2b4c:	afb10034	sw	s1,52(sp)
2b50:	afb00030	sw	s0,48(sp)
2b54:	afb0048	sw	ra,72(sp)
2b58:	afb50044	sw	s5,68(sp)
2b5c:	00808821	move	s1,a0
2b60:	00a08021	move	s0,a1
2b64:	00c09821	move	s3,a2
2b68:	0c000000	jal	0 <nr_processes>
2b6c:	00e0a021	move	s4,a3
2b70:	00409021	move	s2,v0
2b74:	06400032	bltz	s2,2c40 <do_fork+0x104>

kprobe

00002b3c <do_fork>:

2b3c:	0000000d	break	
2b40:	afb40040	sw	s4,64(sp)
2b44:	afb3003c	sw	s3,60(sp)
2b48:	afb20038	sw	s2,56(sp)
2b4c:	afb10034	sw	s1,52(sp)
2b50:	afb00030	sw	s0,48(sp)
2b54:	afb0048	sw	ra,72(sp)
2b58:	afb50044	sw	s5,68(sp)
2b5c:	00808821	move	s1,a0
2b60:	00a08021	move	s0,a1
2b64:	00c09821	move	s3,a2
2b68:	0c000000	jal	0 <nr_processes>
2b6c:	00e0a021	move	s4,a3
2b70:	00409021	move	s2,v0
2b74:	06400032	bltz	s2,2c40 <do_fork+0x104>

After original instruction execution, again 'break' instruction '00000014d' is executed. This causes single-step exception and post-handler gets executed. After this, 'pc' will be set to next instruction after the probed instruction

27bdffb0	addiu	sp,sp,-80
0000014d	break	

New Executable Page

MIPS Kprobes Limitations

- Only minimal testing is performed and needs to be tested using system tap test suites
- Not tested under SMP configuration
- Following instruction types are not supported
 - Jump register or jump and link register (ex: jr, jalr)
 - Jump And Link and Jump (ex: jal, j)
 - Branch instructions (ex: bltz, bgez, beq, bne, blez, bgtz)
 - coprocessor instructions (ex: cop1)

ARM Arch Kprobes Contents

- Target Environment Details
- Kprobes Design
- Sample Example Output
- Sample Example Disassembly
- Limitations

ARM Target Environment Details

- **Target Board**
 - OMAP Starter Kit (OSK 5912)
- **OS**
 - Linux 2.6.16.39
- **Configuration**
 - `CONFIG_PREEMPT=y`
 - `CONFIG_KALLSYMS=y`
 - `CONFIG_KALLSYMSALL=y`
 - `CONFIG_DEBUG_KERNEL=y`



Hardware Features:
OSK5912 @ 192 Mhz.
32 M Byte Flash ROM
USB Host Port

ARM Kprobes Design

- ARM arch does not have break or single step instructions
- Kprobes is implemented using one of the exception handling mechanism

ARM Supports Below Exception Types

Exception	Mode	Purpose
Fast Interrupt Request	FIQ	Fast interrupt handling
Interrupt Request	IRQ	Normal interrupt handling
SWI and RESET	SVC	Protected mode for OS
Pre-fetch or data abort	ABT	Memory protection handling
Undefined Instruction	UND	SW emulation of HW coprocessors

Undefined Exception Type is used for Kprobes

ARM Kprobes Design (Contd)

- Undefined instructions are those which the processor cannot decode in the absence of coprocessor. So this is used for kprobes implementation.
- Undefined Instruction opcode is as below



- If the “cond” field opcode[31:28] value = 1110, then it indicates the unconditional type.

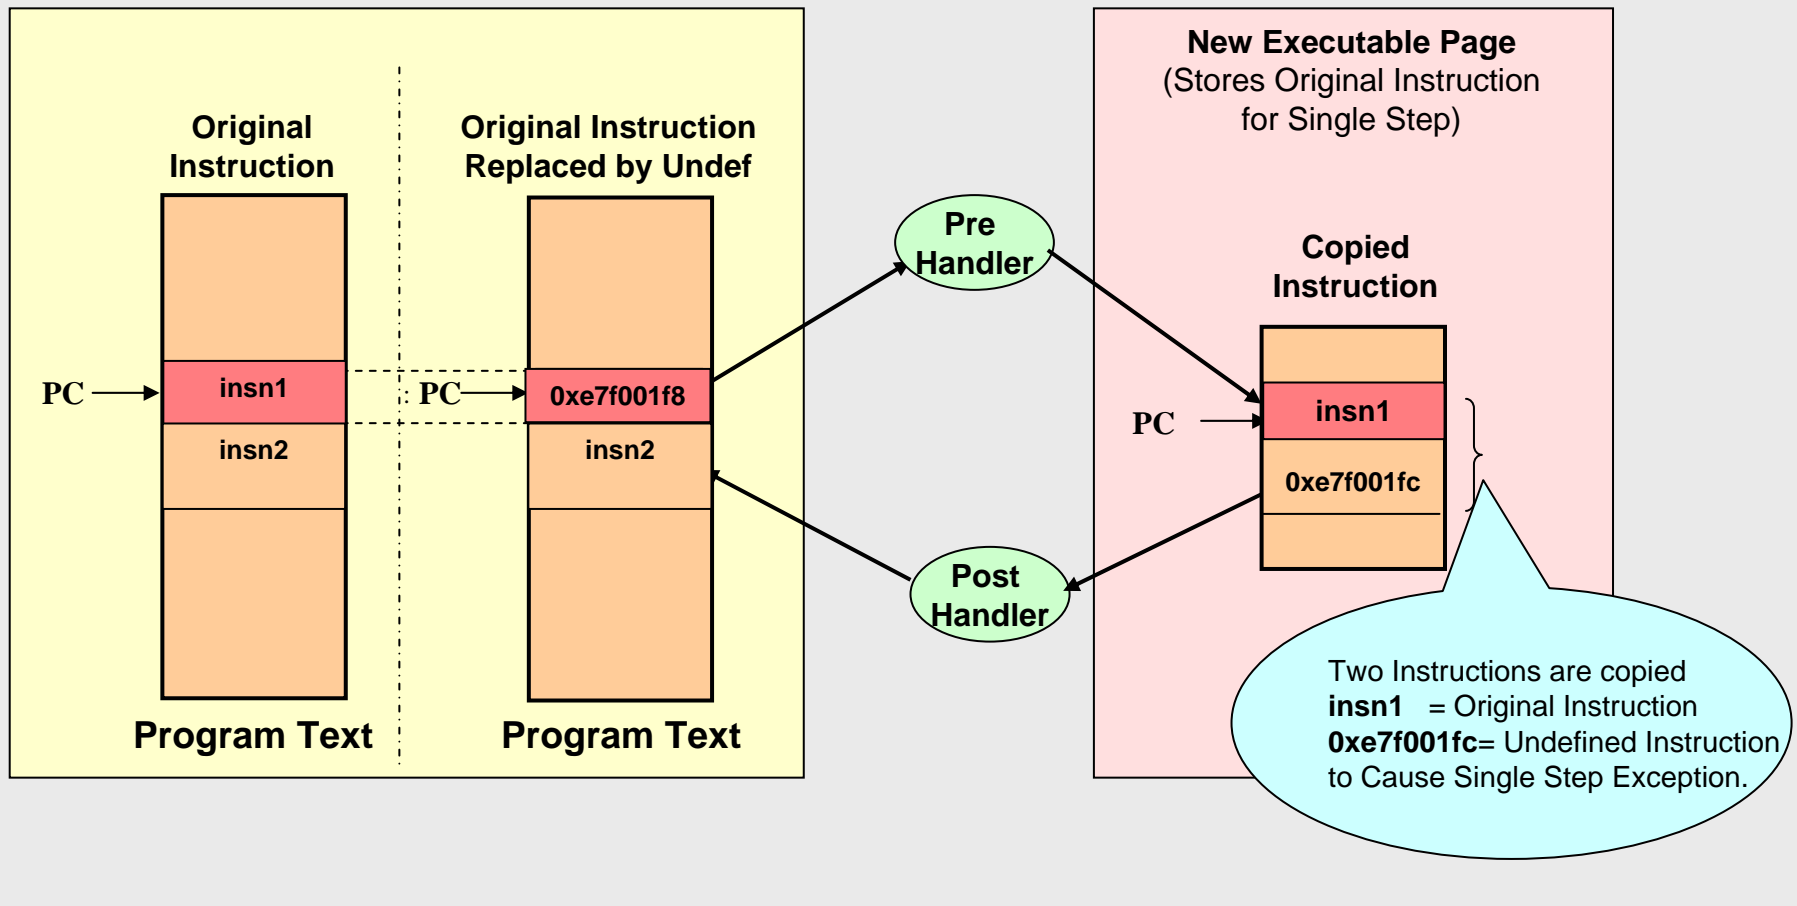


ARM Kprobes Design (Contd)

- When undefined exception is generated, Link Register (LR) will be updated with appropriate return address and this is used to implement single step operation and to resume the next instruction of the probe point
- In the exception handler, “die notifier” kernel feature is used to differentiate breakpoint and single step exceptions

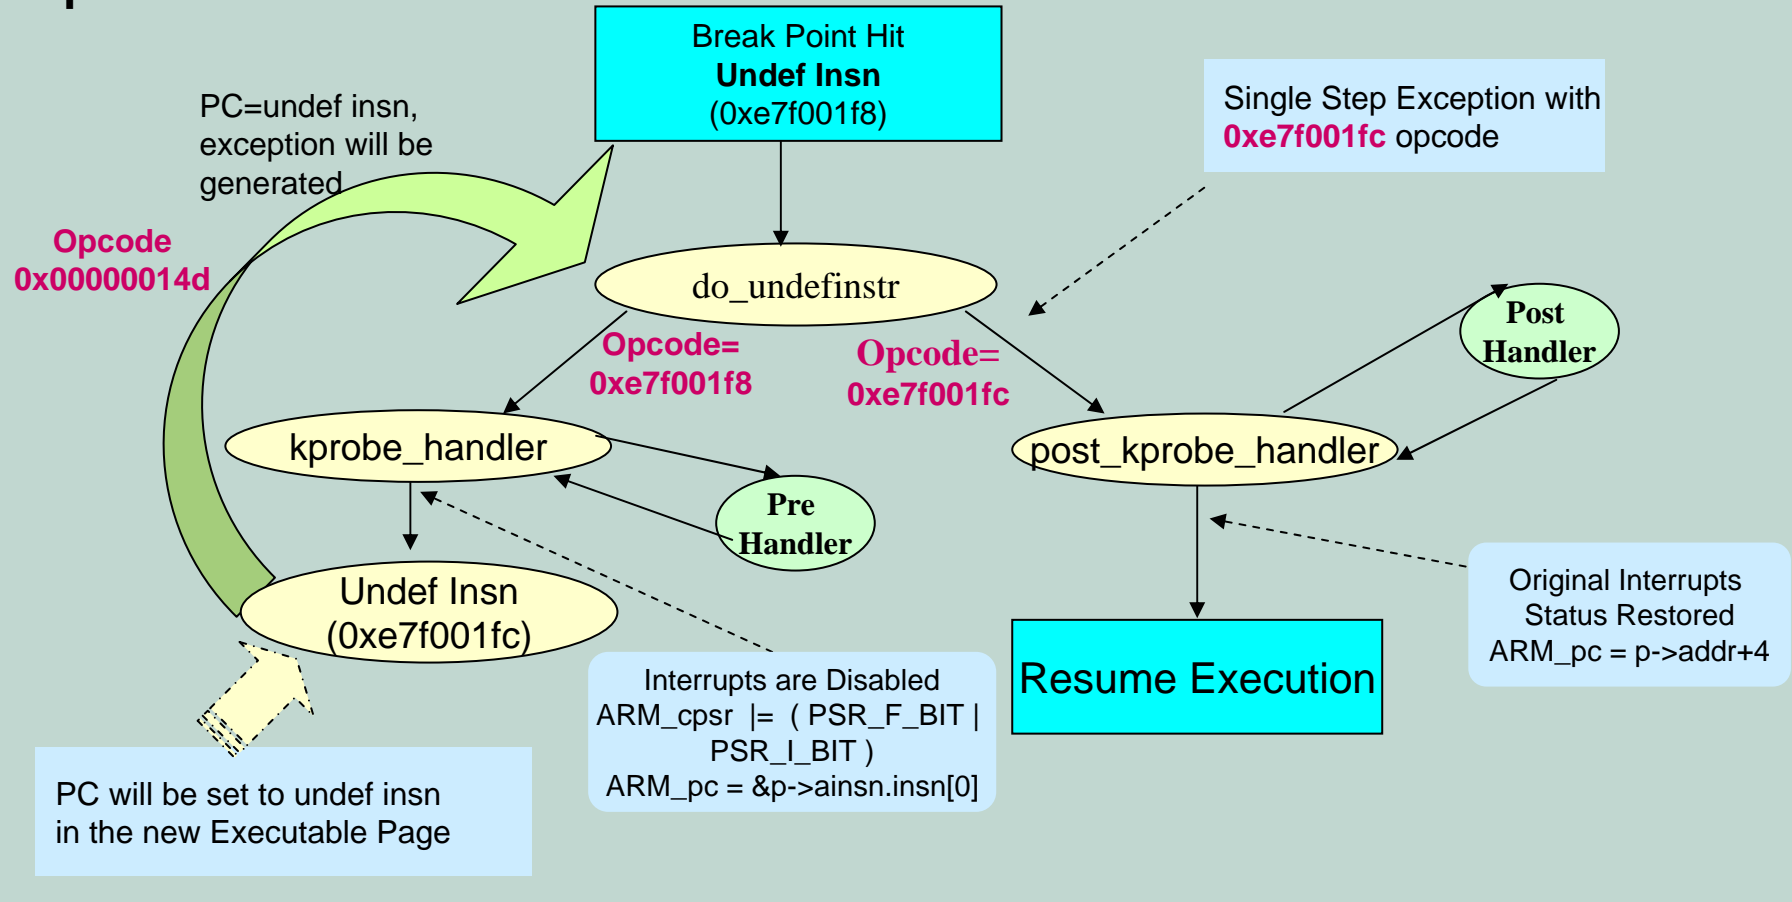
ARM Kprobes Design (Contd)

Kprobes Control Flow



ARM Kprobes Design (Contd)

Kprobes Execution Flow



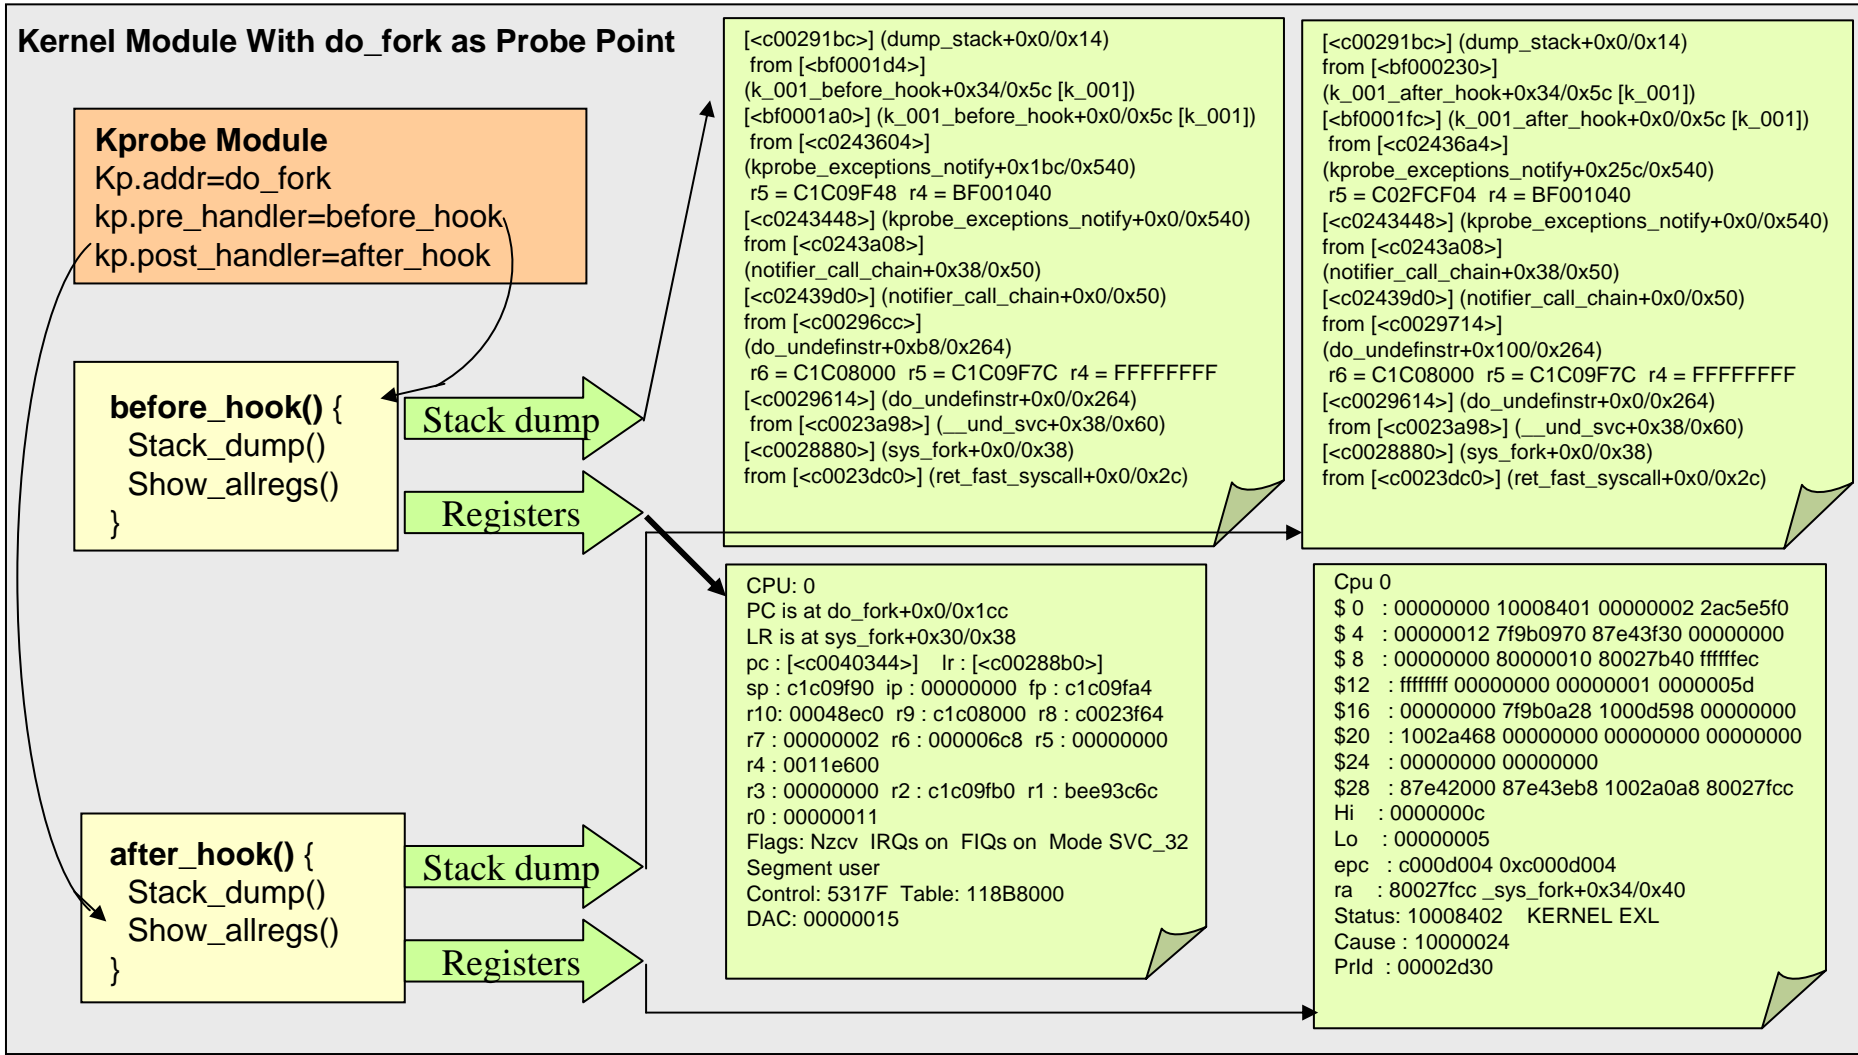
ARM Kprobes Design (Contd)

- Kprobes Control Flow Details
 - When a particular instruction “insn1” is probed, the original instruction will be copied on the executable page and replaces the probed instruction by undefined instruction ‘0xe7f001f8’ i.e BREAKPOINT_INSTRUCTION
 - Undefined Instruction with opcode 0xe7f001fc (BREAKPOINT_INSTRUCTION_2) will be copied on the executable page
 - When “BREAKPOINT_INSTRUCTION” is hit, kprobes exception notifier will be called by the exception handler do_undefinstr() with the die value DIE_UNDEF_1. This will in turn call kprobe_handler() which executes pre_handler() associated with kprobes

ARM Kprobes Design (Contd)

- Kprobes Control Flow Details (Contd)
 - kprobe_handler() takes care of moving program counter to the copied instruction by calling prepare_single_step().
 - After executing the probed instruction, the next instruction on the copied page 'BREAKPOINT_INSTRUCTION_2' will be executed
 - When BREAKPOINT_INSTRUCTION_2 is hit, again kprobes exception notifier will be called by the exception handler do_break() with the die value DIE_UNDEF_2. This will in turn calls post_kprobe_handler ().
 - After the instruction debugging, the execution resumes to the next instruction "insn2"

ARM Sample Example Output



ARM Example Disassembly

When kprobe is registered, the original instruction is replaced with 'undefined' instruction '0xe7f001f8'.

When probe point is hit, kprobes executes pre-handler and pc is made to point original instruction copied on executable page

00001cac <do_fork>:

```

1cac: e1a0c00d  mov  ip, sp
1cb0: e92ddf0   stmdb sp!, {r4, r5, r6,
        r7, r8, sl, fp, ip, lr, pc}
1cb4: e24cb004  sub  fp, ip, #4 ; 0x4
1cb8: e24dd020  sub  sp, sp, #32 ; 0x20
1cbc: e1a05000  mov  r5, r0
1cc0: e1a04001  mov  r4, r1
1cc4: e1a08002  mov  r8, r2
1cc8: e1a0a003  mov  sl, r3
1ccc: ebfffffe  bl   0 <alloc_pidmap>
1cd0: e2506000  subs r6, r0, #0 ; 0x0
1cd4: b3e0000a  mvnlt r0, #10 ; 0xa
1cd8: ba000063  blt  1e6c <do_fork+0x1c0>
    
```

kprobe

00001cac <do_fork>:

```

1cac: e7f001f8  undefined
1cb0: e92ddf0   stmdb sp!, {r4, r5, r6,
        r7, r8, sl, fp, ip, lr, pc}
1cb4: e24cb004  sub  fp, ip, #4 ; 0x4
1cb8: e24dd020  sub  sp, sp, #32 ; 0x20
1cbc: e1a05000  mov  r5, r0
1cc0: e1a04001  mov  r4, r1
1cc4: e1a08002  mov  r8, r2
1cc8: e1a0a003  mov  sl, r3
1ccc: ebfffffe  bl   0 <alloc_pidmap>
1cd0: e2506000  subs r6, r0, #0 ; 0x0
1cd4: b3e0000a  mvnlt r0, #10 ; 0xa
1cd8: ba000063  blt  1e6c <do_fork+0x1c0>\
    
```

After original instruction execution, again 'undefined' instruction '0xe7f001fc' is executed. This causes single-step exception and post-handler will be executed. After this, pc will be moved to the next instruction after the probed instruction

```

e1a0c00d  mov  ip, sp
e7f001fc  undefined
    
```

Executable Page

ARM Kprobes Limitations


- Only minimal testing is performed and needs to be tested using system tap test suites
- Not tested under SMP configuration
- Following instruction types are not supported
 - Branch/link with exchange instruction type (ex: bx or blx)
 - Data processing instructions whose destination register is a PC (ex: AND, EOR, SUB, RSB, ADD, ADC, SBC, RSC, MOV, BIC, MVN)
 - load instructions which destination register is PC (ex:ldr)
 - Multiple load type instructions (ex: ldm)
 - Branch and branch with link type (Ex: b, bl)

PPC32 Arch Kprobes Contents

- Target Environment Details
- Kprobes Design
- Sample Example Output
- Sample Example Disassembly
- Limitations

PPC32 Target Environment Details

- **Target Board**
 - Ebony (PPC440gp)
- **OS**
 - Linux 2.6.16.39
- **Configuration**
 - CONFIG_PREEMPT=y
 - CONFIG_KALLSYMS=y
 - CONFIG_KALLSYMSALL=y
 - CONFIG_DEBUG_KERNEL=y



Hardware Features:


440GP Processor @ 500 Mhz
System Clock @ 66.66Mhz

PPC32 Kprobes Design

- Kprobe support for ppc64 is present in linux main line tree
- Kprobe support for PPC32 was posted recently
<http://patchwork.ozlabs.org/linuxppc/patch?id=9244>
- But the above patch addresses only for non BookE versions.
- Sony has provided kprobes support for PPC32 BookE versions
- PPC32 with BookE versions involved technical challenges while addressing single step exception handling. These issues are highlighted and also the solution to overcome these is discussed in detail.

PPC32 Kprobes Design (Contd)

- PPC32 Arch does have trap instructions 'tw'.
- Trap instruction format is



- There are different variations of trap instruction.
- This variation is mainly due to 'TO' field of the trap instruction
- In general trap instruction will be of the type 0x7c000008
- The trap used for kprobe is an unconditional trap "tw" i.e '0x7fe00008' for which TO =31 and RA=RB=0

PPC32 Kprobes Design (Contd)

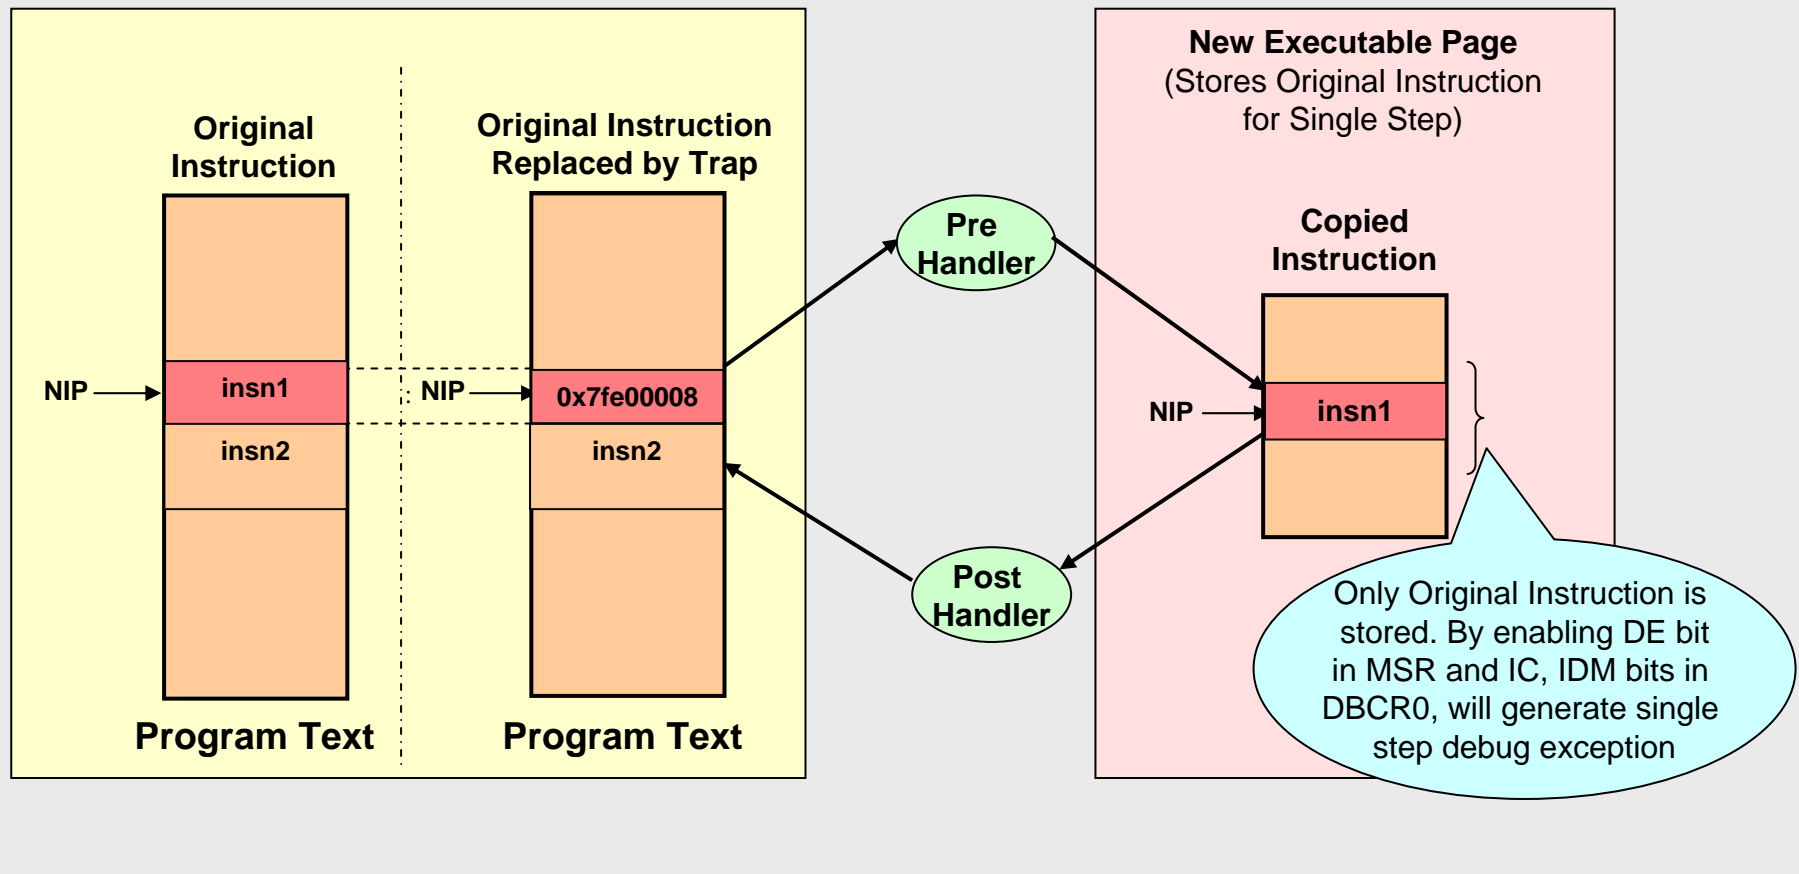
Different types of trap instruction are as follows

Code	Description	TO Encoding
twlt	Less than	16
Twle	Less than or equal	20
Twel	Equal	4
Twge	Greater than or equal	12
Twgt	Greater than	8
Twnl	Not less than	12
Twne	Not equal	24
Twng	Not greater than	20
Twllt	Logically less than	2
Twlle	Logically less than or equal	6
Twlge	Logically greater than or equal	5
Twlgt	Logically greater than	1
Twlnl	Logically not less than	5
Twlng	Logically not greater than	6
Tw	Unconditional	31

This is Used by Kprobes
to Implement Break Point
Exception

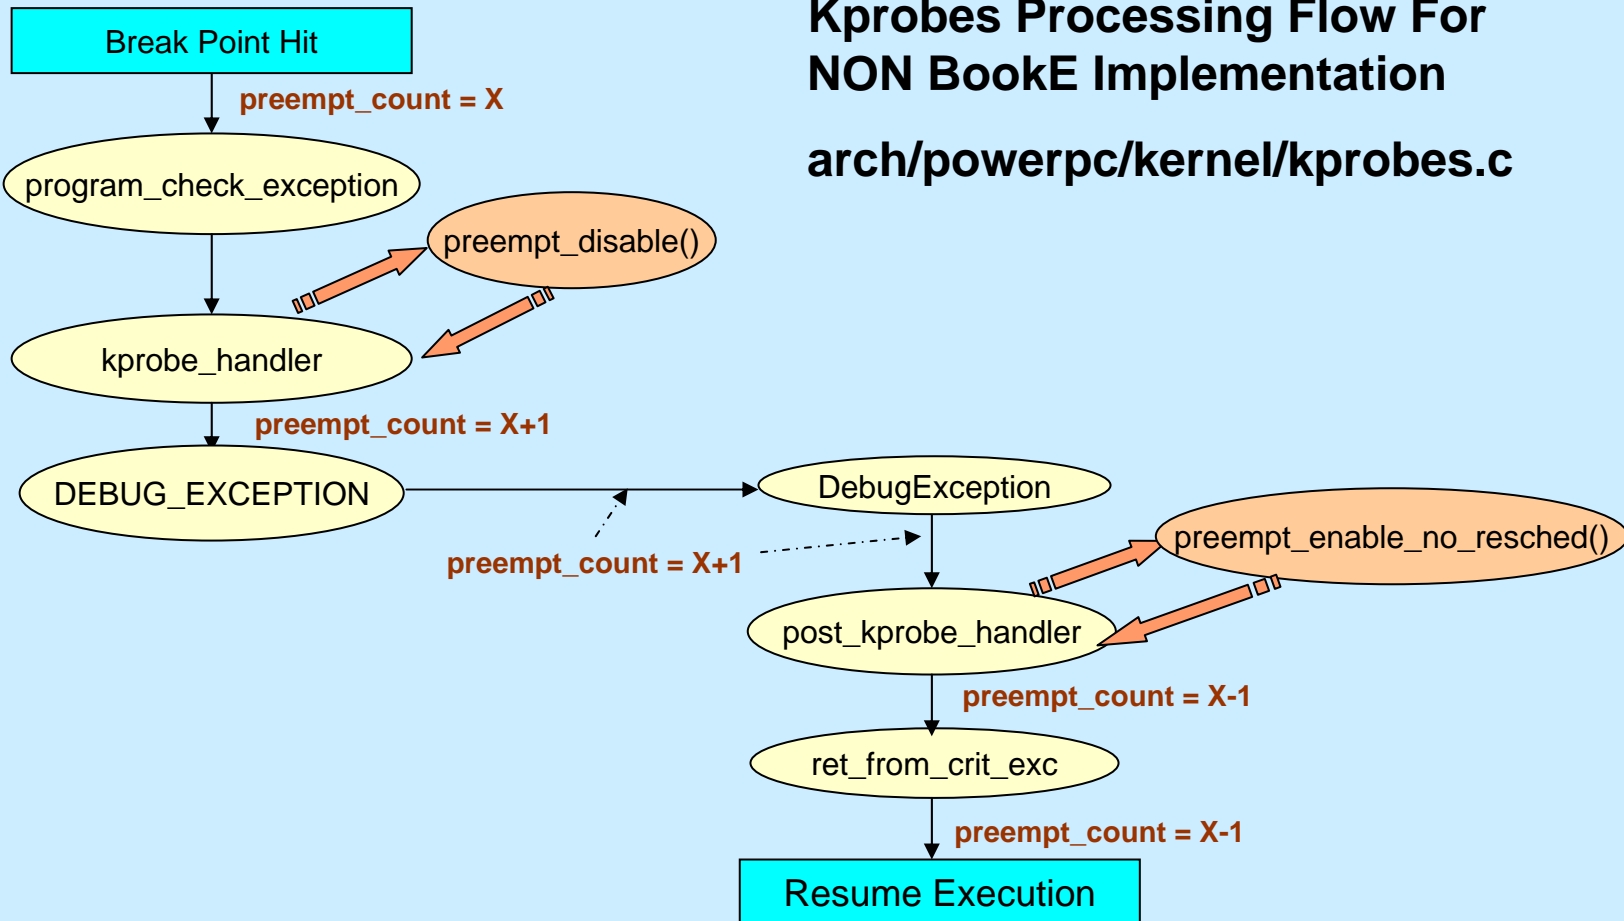
PPC32 Kprobes Design (Contd)

Kprobes Control Flow

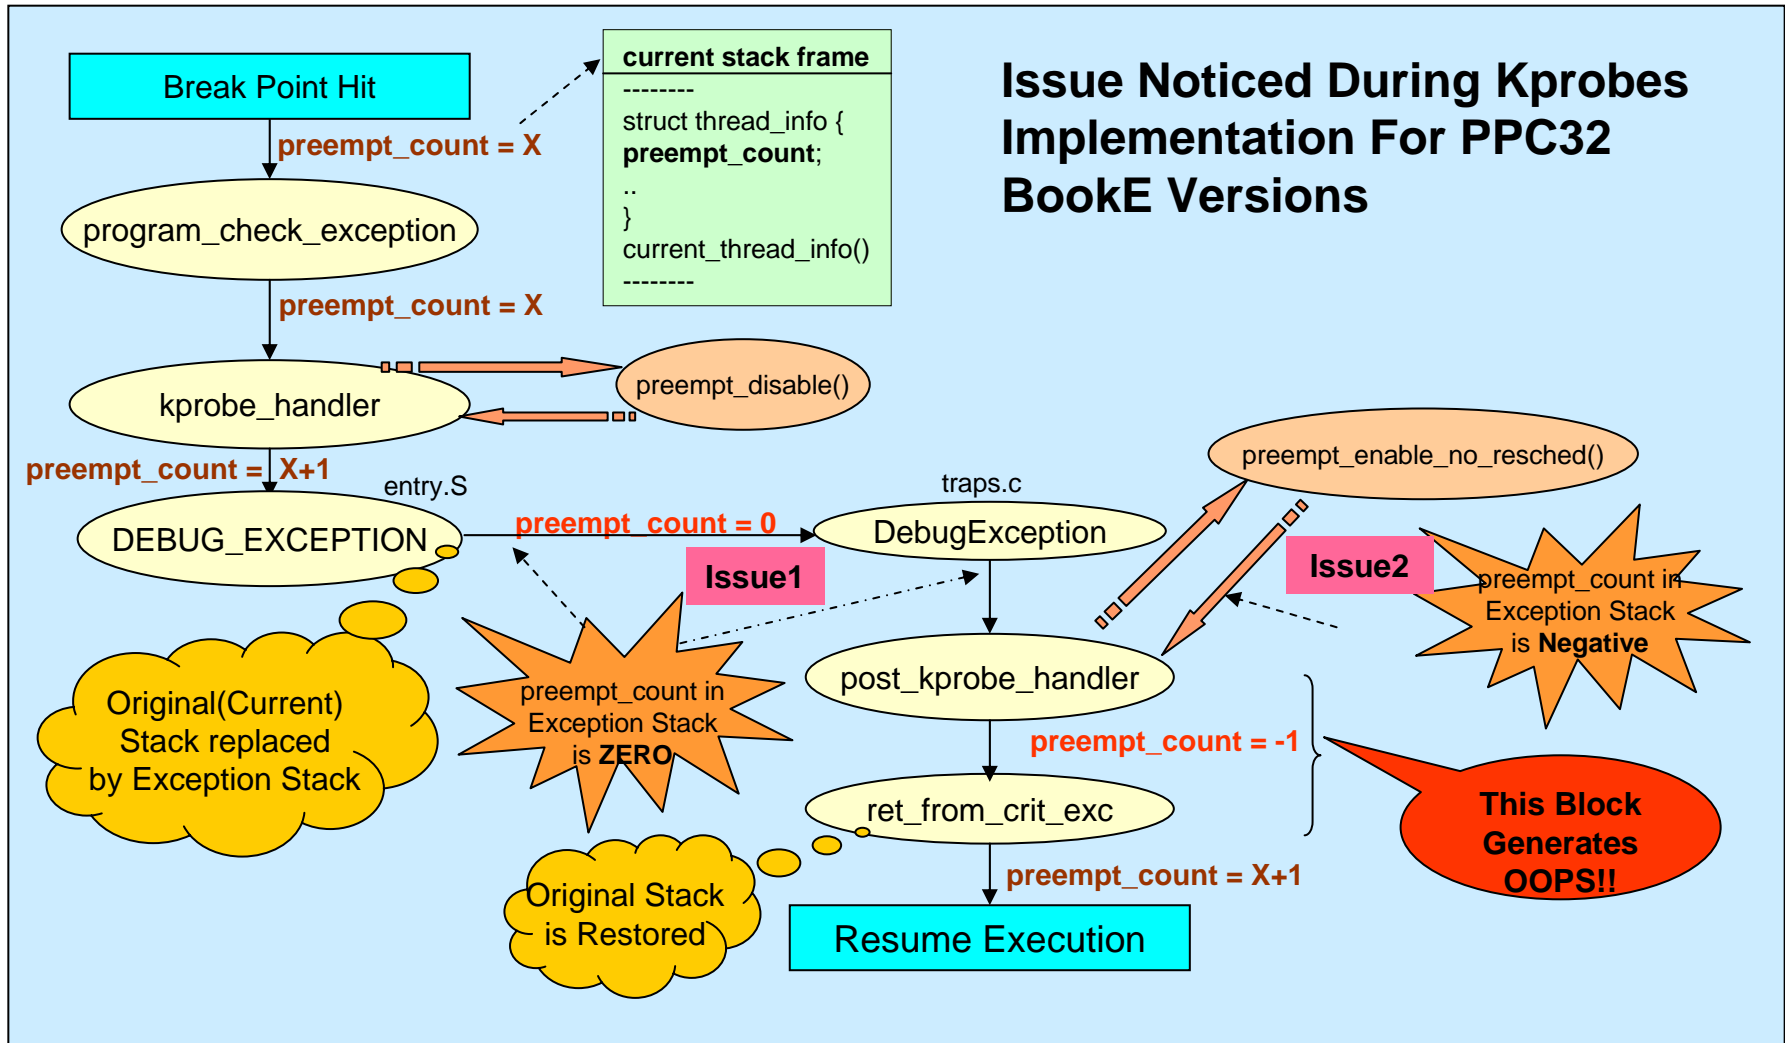


PPC32 Kprobes Design (Contd)

Kprobes Processing Flow For NON BookE Implementation arch/powerpc/kernel/kprobes.c

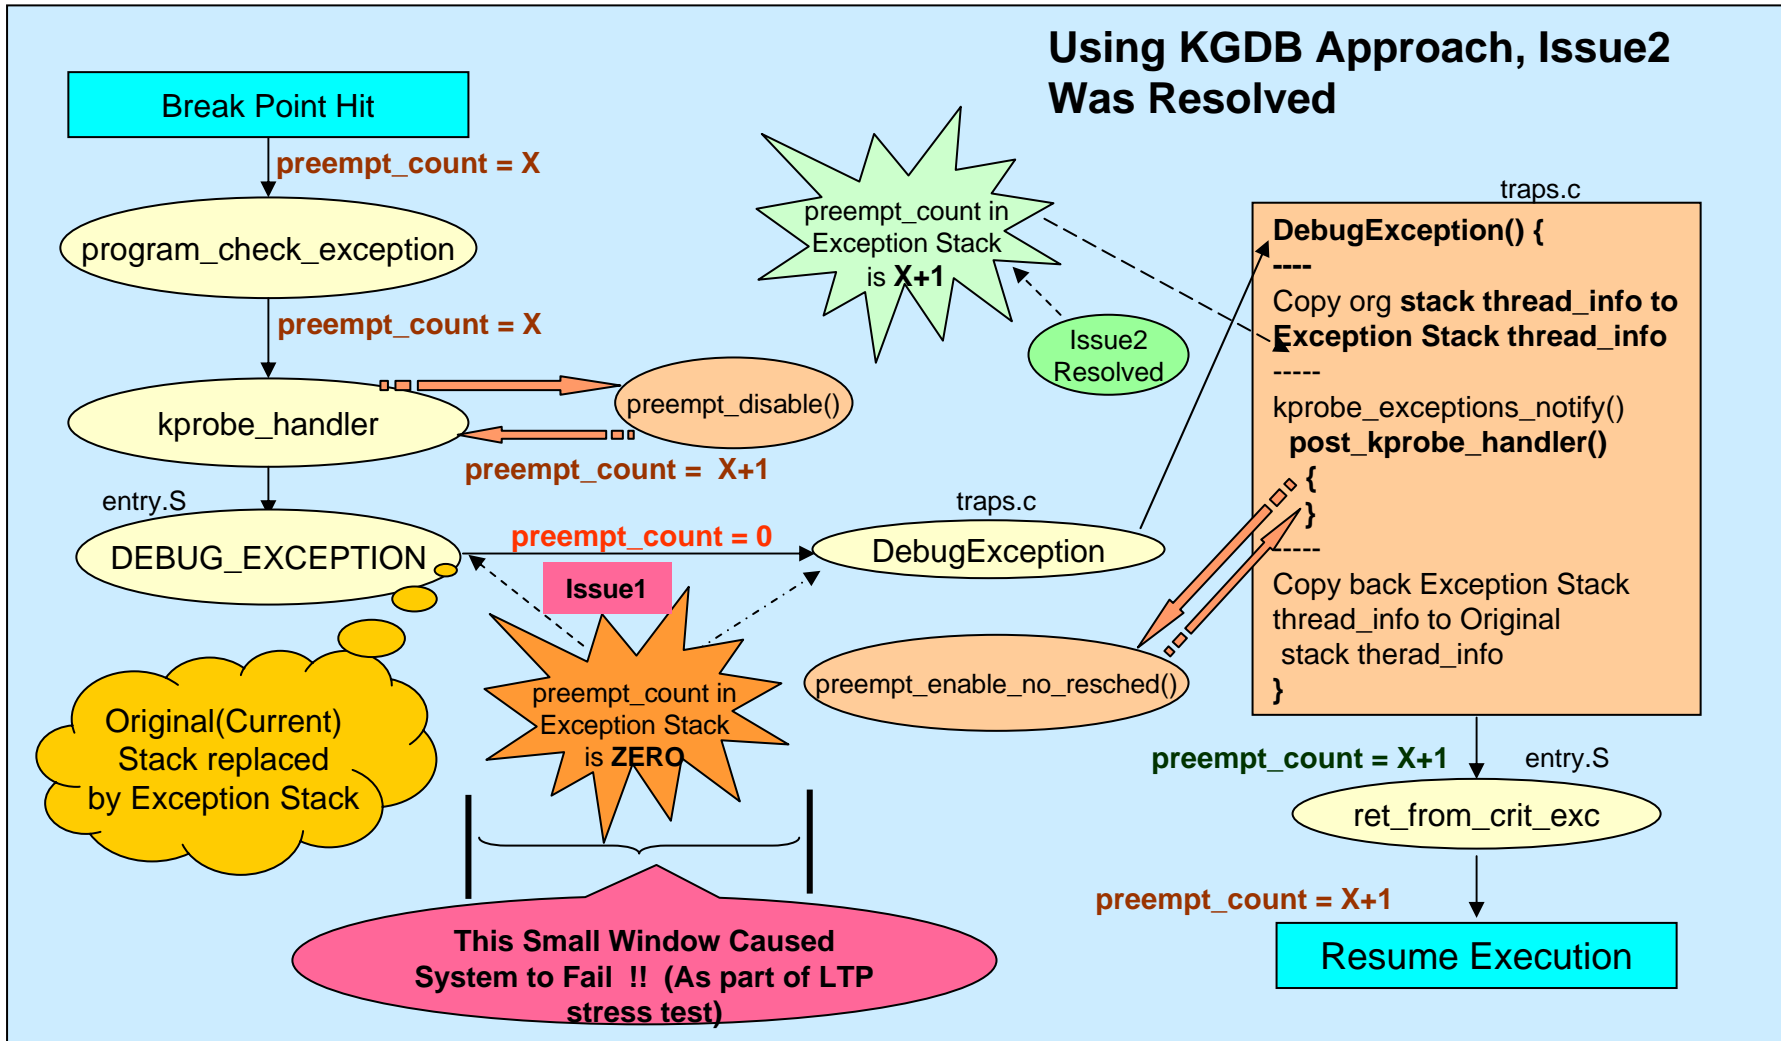


PPC32 Kprobes Design (Contd)



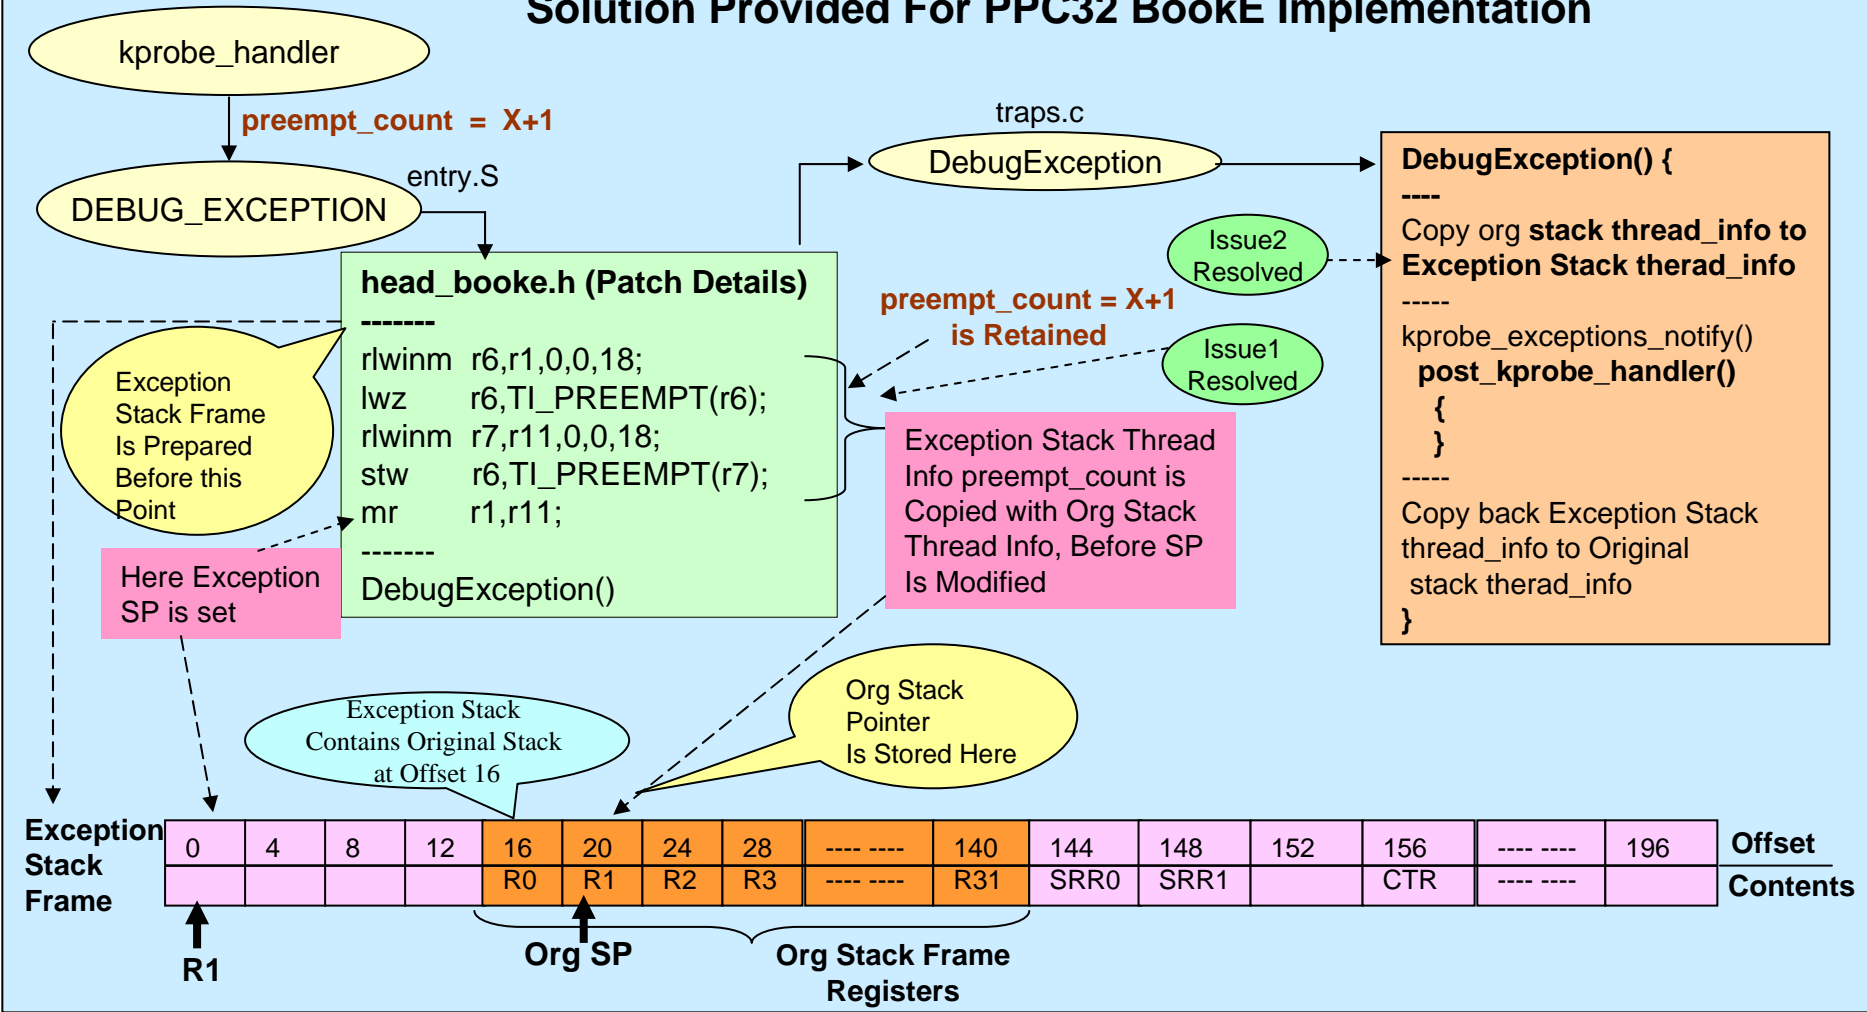
Issue Noticed During Kprobes Implementation For PPC32 BookE Versions

PPC32 Kprobes Design (Contd)



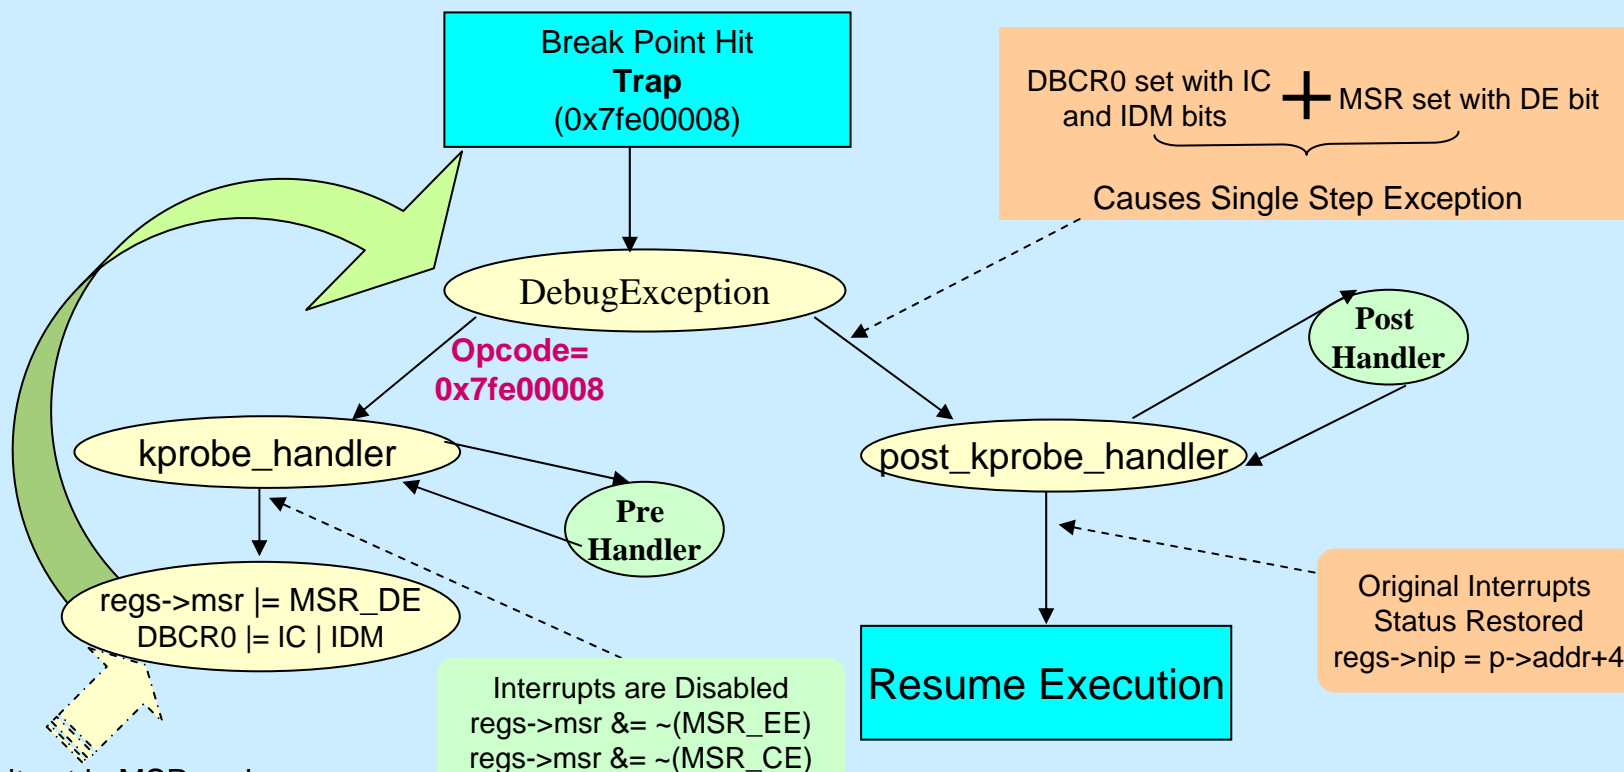
PPC32 Kprobes Design (Contd)

Solution Provided For PPC32 BookE Implementation



PPC32 Kprobes Design (Contd)

Working Kprobes PPC32 BookE Execution Flow



Since DE bit set in MSR and IC, IDM bits are set in DBCR0, Single Step Exception is generated

Interrupts are Disabled
`regs->msr &= ~(MSR_EE)`
`regs->msr &= ~(MSR_CE)`

DBCRO set with IC + MSR set with DE bit and IDM bits
 Causes Single Step Exception

Original Interrupts Status Restored
`regs->nip = p->addr+4`

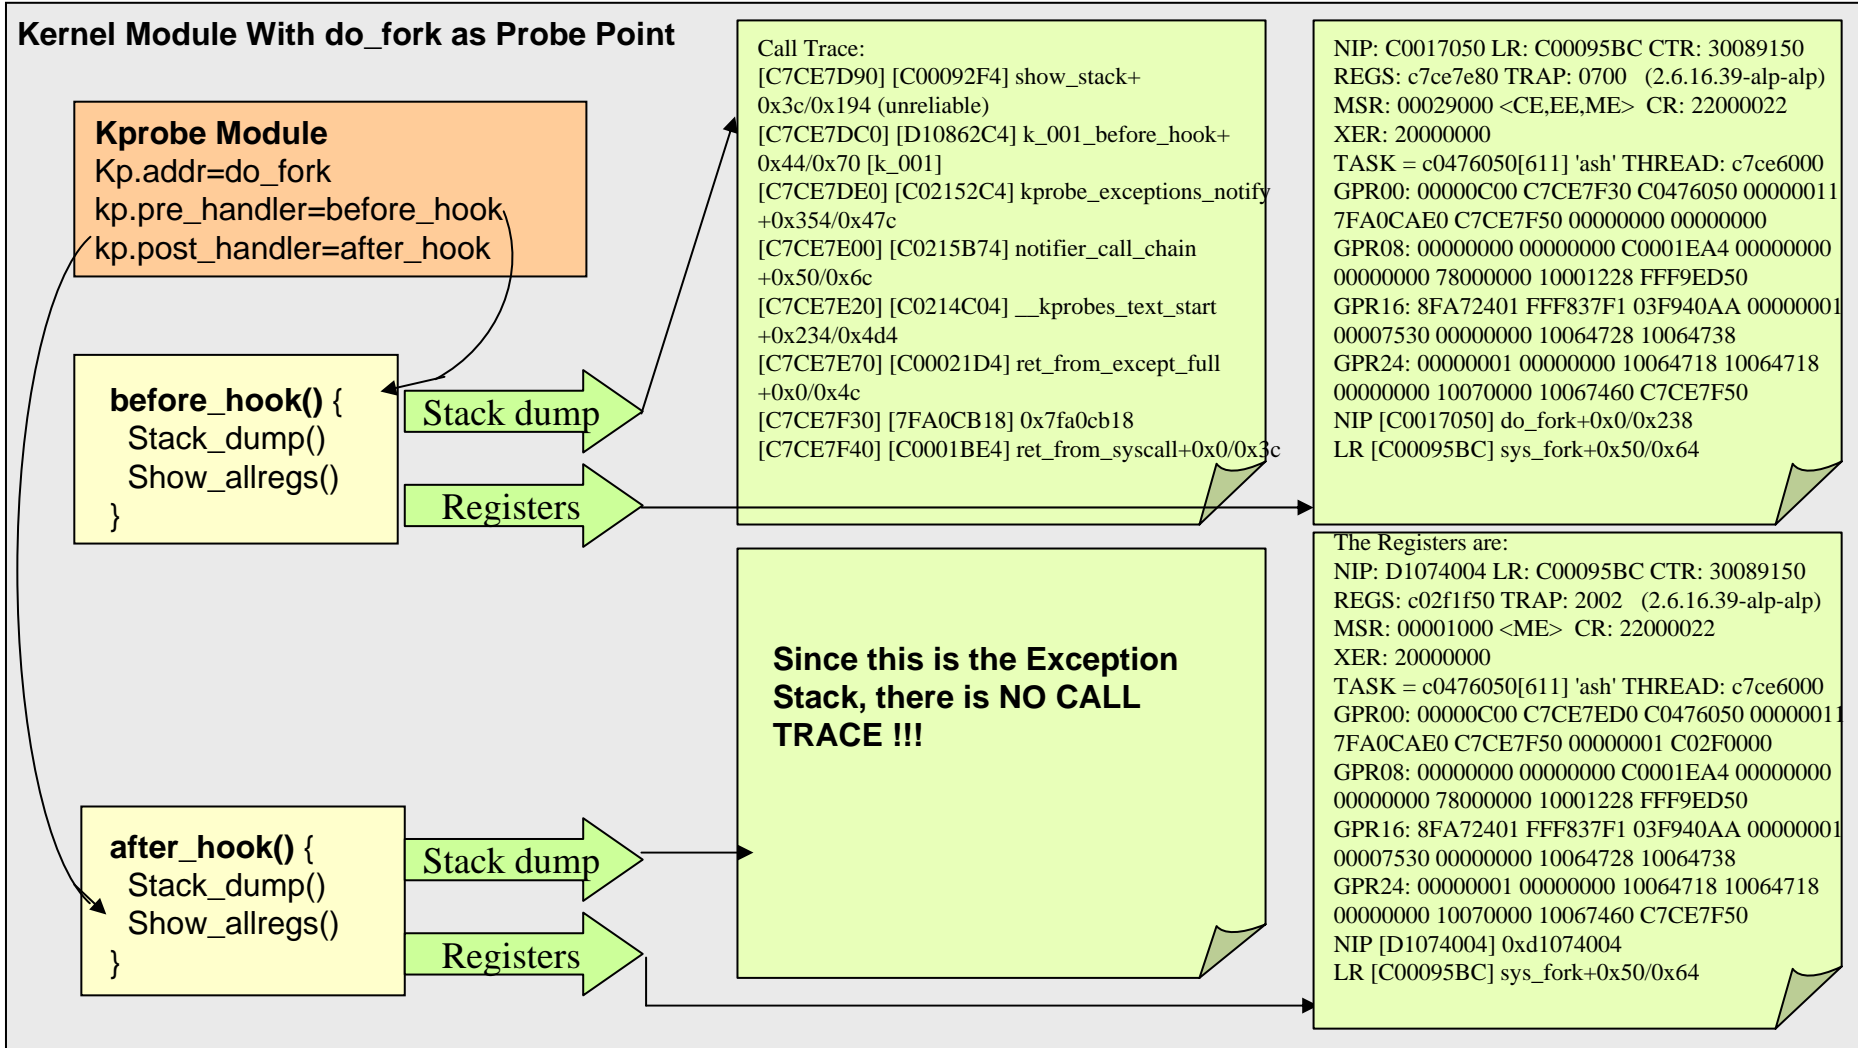
PPC32 Kprobes Design (Contd)

- Kprobes Control Flow Details
 - When a particular instruction “insn1” is probed, the original instruction will be copied on the executable page and replaces the probed instruction by ‘trap’ i.e 0x7fe00008
 - A breakpoint exception occurs, immediately and unconditionally transferring control to the exception handler
 - When ‘tw’ is hit, kprobes exception notifier will be called by the exception handler `program_check_exception ()` with the die value `DIE_BPT`. This will in turn call `kprobe_handler()` which performs executing `pre_handler()` associated with kprobes

PPC32 Kprobes Design (Contd)

- Kprobes Control Flow Details (Contd)
 - `kprobe_handler()` takes care of moving program counter to the copied instruction by calling `prepare_single_step()`. DE bit of MSR and IC, IDM bits of DBCR0 will also be set as part of `prepare_single_step()` and this will generate debug exception.
 - Since MSR[DE] bit is set, single step operation will be performed on the copied instruction in the executable page and kprobes exception notifier will be called by the exception handler `DebugException()` with the die value `DIE_SSTEP`. This will in turn call user defined `post_handler()`
 - After the instruction debugging, the execution resumes to the next instruction “`Insn2`”

PPC32 Sample Example Output



PPC32 Example Disassembly

When kprobe is registered, the original instruction is replaced with 'unconditional trap' instruction '0x7fe00008'.

When probe point is hit, kprobe executes pre-handler and DE bit in MSR register will be enabled and pc is made to point original instruction copied on executable page

```
00001fa4 <do_fork>:
1fa4: 94 21 ff c0  stwu  r1,-64(r1)
1fa8: 7d 80 00 26  mfcrl r12
1fac: 7c 08 02 a6  mflr  r0
1fb0: bf 01 00 20  stmw  r24,32(r1)
1fb4: 7c 7e 1b 78  mr    r30,r3
1fb8: 90 01 00 44  stw   r0,68(r1)
1fbc: 7c 9f 23 78  mr    r31,r4
1fc0: 91 81 00 1c  stw   r12,28(r1)
1fc4: 7c bc 2b 78  mr    r28,r5
1fc8: 7c db 33 78  mr    r27,r6
1fcc: 7c fa 3b 78  mr    r26,r7
1fd0: 7d 19 43 78  mr    r25,r8
1fd4: 48 00 00 01  bl    1fd4 <do_fork+0x30>
```

```
00001fa4 <do_fork>:
1fa4: 7f e0 00 08  trap
1fa8: 7d 80 00 26  mfcrl r12
1fac: 7c 08 02 a6  mflr  r0
1fb0: bf 01 00 20  stmw  r24,32(r1)
1fb4: 7c 7e 1b 78  mr    r30,r3
1fb8: 90 01 00 44  stw   r0,68(r1)
1fbc: 7c 9f 23 78  mr    r31,r4
1fc0: 91 81 00 1c  stw   r12,28(r1)
1fc4: 7c bc 2b 78  mr    r28,r5
1fc8: 7c db 33 78  mr    r27,r6
1fcc: 7c fa 3b 78  mr    r26,r7
1fd0: 7d 19 43 78  mr    r25,r8
1fd4: 48 00 00 01  bl    1fd4 <do_fork+0x30>
```

Kprobe

After original instruction execution on executable page, since DE bit is enabled in MSR register, it generates single step debug exception. During debug exception, the post-handler gets executed and pc will be moved to the next instruction after the probed instruction

```
94 21 ff c0  stwu  r1,-64(r1)
```

Executable Page

PPC32 Kprobes Limitations

- If post handler tries to dump stack contents, then call trace cannot be displayed because during post handler execution, the stack corresponds to Exception Stack.
- Only minimal testing is performed and needs to be tested using system tap test suites
- Not tested under SMP configuration

Overhead Measurement Details

Kprobes Overhead Measurement Data for MIPS, ARM and PPC32 Arch

Overhead Measurement Test Sample: <http://sourceware.org/ml/systemtap/2006-q1/msg00318.html>

Overhead	MIPS(usec/call)	ARM(usec/call)	PPC32(usec/call)
Time taken to execute gettimeofday	1.16	5.22	0.99
Time taken to execute gettimeofday when a kprobe is inserted at do_gettimeofday with only Pre-handler	2.98	11.85	3.26
pre_handler Overhead	1.81	6.62	2.27
Time taken to execute gettimeofday when a kprobe is inserted at do_gettimeofday with only Post-handler	2.99	12.04	3.27
post_handler Overhead	1.82	6.81	2.28
Time taken to execute gettimeofday when a kprobe is inserted at do_gettimeofday with both Pre and post handler	3.01	12.10	3.30
(pre+post)_handler Overhead	1.84	6.87	2.31

Source Patches

- MIPS patches can be extracted from below link under patches for 2.6.16
 - <http://tree.celinuxforum.org/CelfPubWiki/PatchArchive>
- ARM patches can be extracted from below link under patches for 2.6.16
 - <http://tree.celinuxforum.org/CelfPubWiki/PatchArchive>
- PPC32 patches will be uploaded very soon in CELF Patch Archive page

References

- <http://www-128.ibm.com/developerworks/library/l-kprobes.html?ca=dgr-Inxw42Kprobe>
- <http://www-users.cs.umn.edu/~boutcher/kprobes/>
- <http://tree.celinuxforum.org/pipermail/celinux-dev/2006-October/001335.html>
- <http://tree.celinuxforum.org/pipermail/celinux-dev/2007-January/001395.html>