

---

# Project 2 Roadmap

# Background – Context Switching

---

- One processor and multiple threads running concurrently – How?!!
- Give each thread a small time quantum to run.
- When this quantum expires, or the thread blocks, **context-switch** to a different thread.
  1. Where should I save the thread context during a context-switch?
  2. What should this context consist of?

# Background – Kernel Stack

---

- User process is a kernel thread with `USER_CONTEXT` structure.
- Store the current context (state) before context switching.
- Where is the kernel stack?

```
struct Kernel_Thread {  
    unsigned long esp; // Stack pointer (absolute)  
    void* stackPage; //The beginning of the stack  
    .....  
};
```

- *esp* points at the end of the stack (stack grows down from higher to lower address)

# Background – User Processes

---

- Two stacks: kernel stack and user stack.
- User Stack (store local variables)
- Start\_User\_Thread:

set up the *kernel* stack to look as if the thread had previously been running and then context-switched to the ready queue.

# Background – Context Information

User Stack Location

Interrupt\_State

- The items at the top are pushed first.
- Program Counter → EIP
- User stack pointer points to the end of the DS.
- Stack grows down from higher address to lower address.

Stack Data Selector (data selector)
Stack Pointer (end of data memory)
Eflags
Text Selector (code selector)
Program Counter (entry addr)
Error Code (0)
Interrupt Number (0)
EAX (0)
EBX (0)
ECX (0)
EDX (0)
ESI (Argument Block address)
EDI (0)
EBP (0)
DS (data selector)
ES (data selector)
FS (data selector)
GS (data selector)

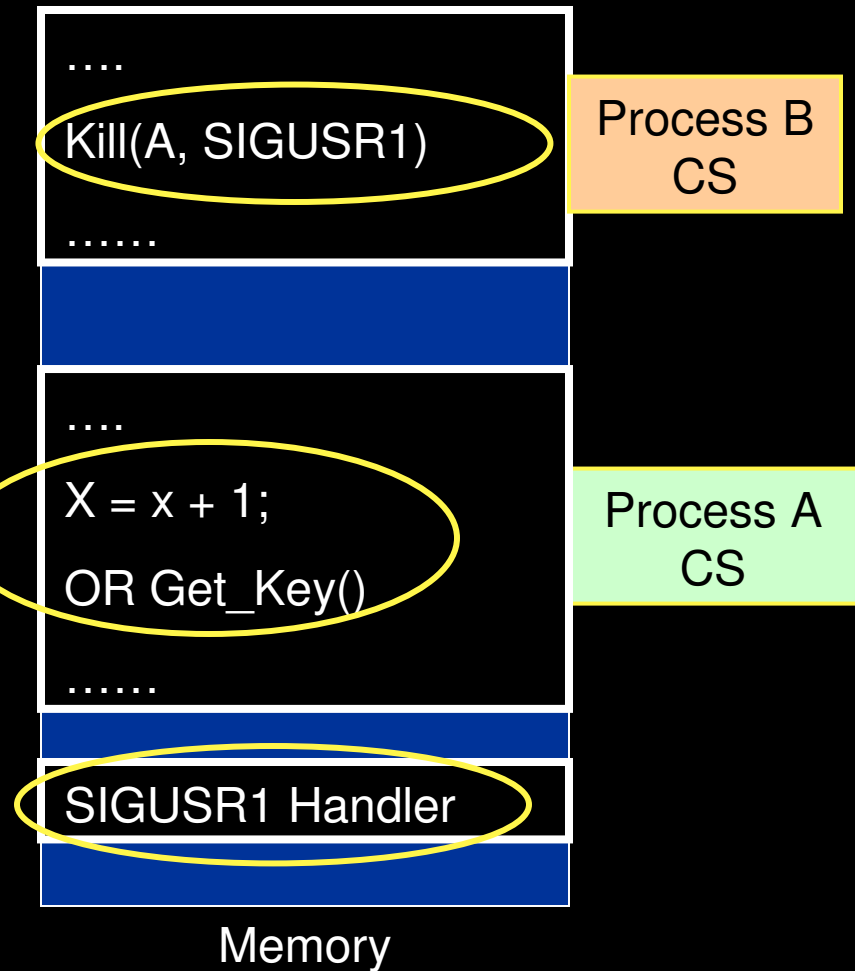
# Project 2: Signals

---

- Signals are to user processes what interrupts are to the kernel .
- Process temporarily stop what it is doing, and is instead redirected to the **signal handler**.
- When the handler completes, the process goes back to what it was doing (unless another signal is pending!)

# Signals

1. Process A is executing then either finishes quantum OR blocked
1. Process B is now executing and sends a signal to A.
1. Process A is executing again. However, control is transferred to SIGUSR1 handler.
1. SIGUSR1 handler finishes. Then control transfers to Process A again (if no other signal pending).



# Project Requirements

---

1. Add the code to handle signals.
2. System calls.
3. Background processes are NOT detached.

**Look for TODO macro**



# Supported Signals

---

1. SIGKILL: treated as Sys\_Kill of project1.
2. SIGUSR1 & SIGUSR2
3. SIGCHLD
  - Background processes are NOT detached any more (refCount = 2).
  - Sent to a parent when the background child dies.
  - Default handler = reap the child

# System Calls

---

- **sys\_signal**: register a signal handler
- **sys\_sigaction**: initialize signal handling for a process
- **sys\_kill**: send a signal
- **sys\_sigreturn**: indicate completion of signal handler
- **sys\_waitpid**: wait for any child process to die

# Sys\_Signal

---

- Register a signal handler for a process
  - state->ebx - pointer to handler function
  - state->ecx - signal number
  - Returns: 0 on success or error code (< 0) on error
- Calling Sys\_Signal with a handler to SIGKILL should result in an error.
- Initial handler for SIGCHLD (reaps all zombie) is  
Def\_Child\_Handler
- Two predefined handlers:
  - SIG\_IGN, SIG\_DFL (check include/libc/signal.h)
  - Used #define to set a fake address
  - Should be handled directly from kernel
- Example: Signal(SIGUSR1,SIG\_IGN);

# Sys\_RegDeliver

---

- Register trampoline function:
  - calls Sys\_ReturnSignal
- Signals cannot be delivered until this is registered.
  - state->ebx - pointer to Return\_Signal function
  - Returns: 0 on success or error code (< 0) on error

# Sys\_Kill

---

- Send a signal to a process
  - state->ebx - pid of process
  - state->ecx - signal number
  - Returns: 0 on success or error code (< 0) on error

# Sys\_ReturnSignal

---

- Complete signal handling for this process.
  - No Parameters
  - Returns: 0 on success or error code ( $< 0$ ) on error
- Called by a process immediately after it has handled a signal.

# Sys\_WaitNoPID

---

- Reap a child process that has died
  - state->ebx - pointer to status of process reaped
  - Returns: pid of reaped process on success, -1 on error.

# Signals Golden Rules

---

- Any user process stores THREE pointers to handler functions corresponding to (SIGUSR1, SIGUSR2, SIGCHLD).
- These pointers could be NULL if there is no registered handler.
- Any process also stores THREE pointers to the Ign\_Handler, Def\_Handler, Signal\_Return
- If there no handler registered, the default handler will be executed.
- Signal handling is **non-reentrant**.



# Signals Delivery

---

src/**geekos**/signal.c

1. Send\_Signal
2. Check\_Pending\_Signal
3. Set\_Handler
4. Setup\_Frame
5. Complete\_Handler

# Check\_Pending\_Signal

---

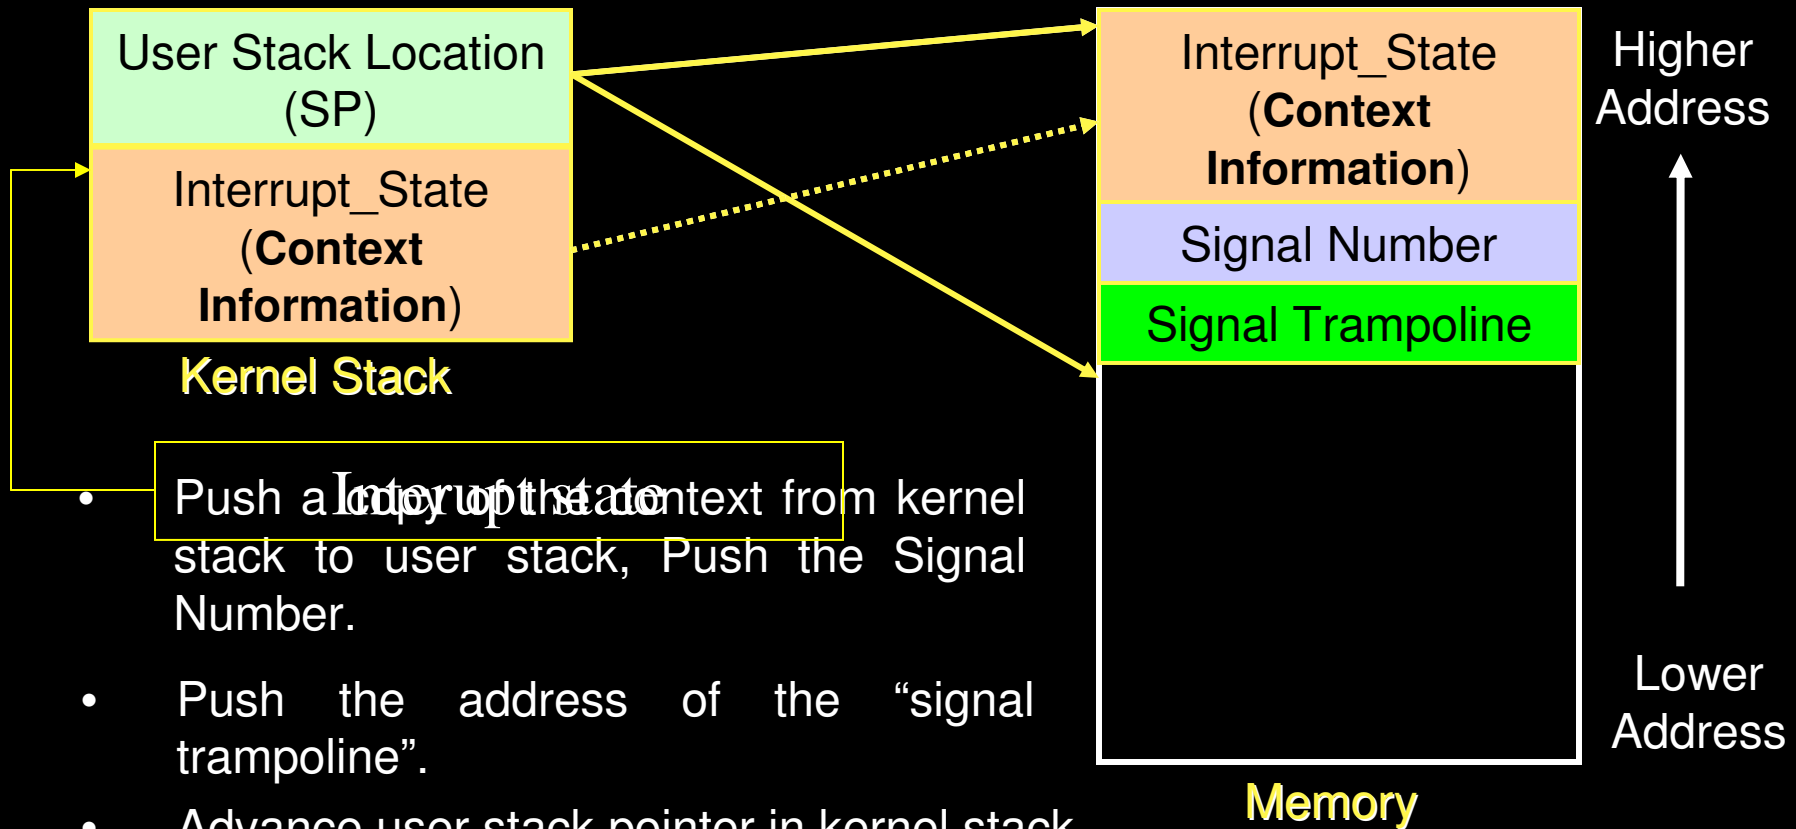
1. A signal is pending for that user process.
2. The process is about to start executing in user space.

CS register  $\neq$  KERNEL\_CS

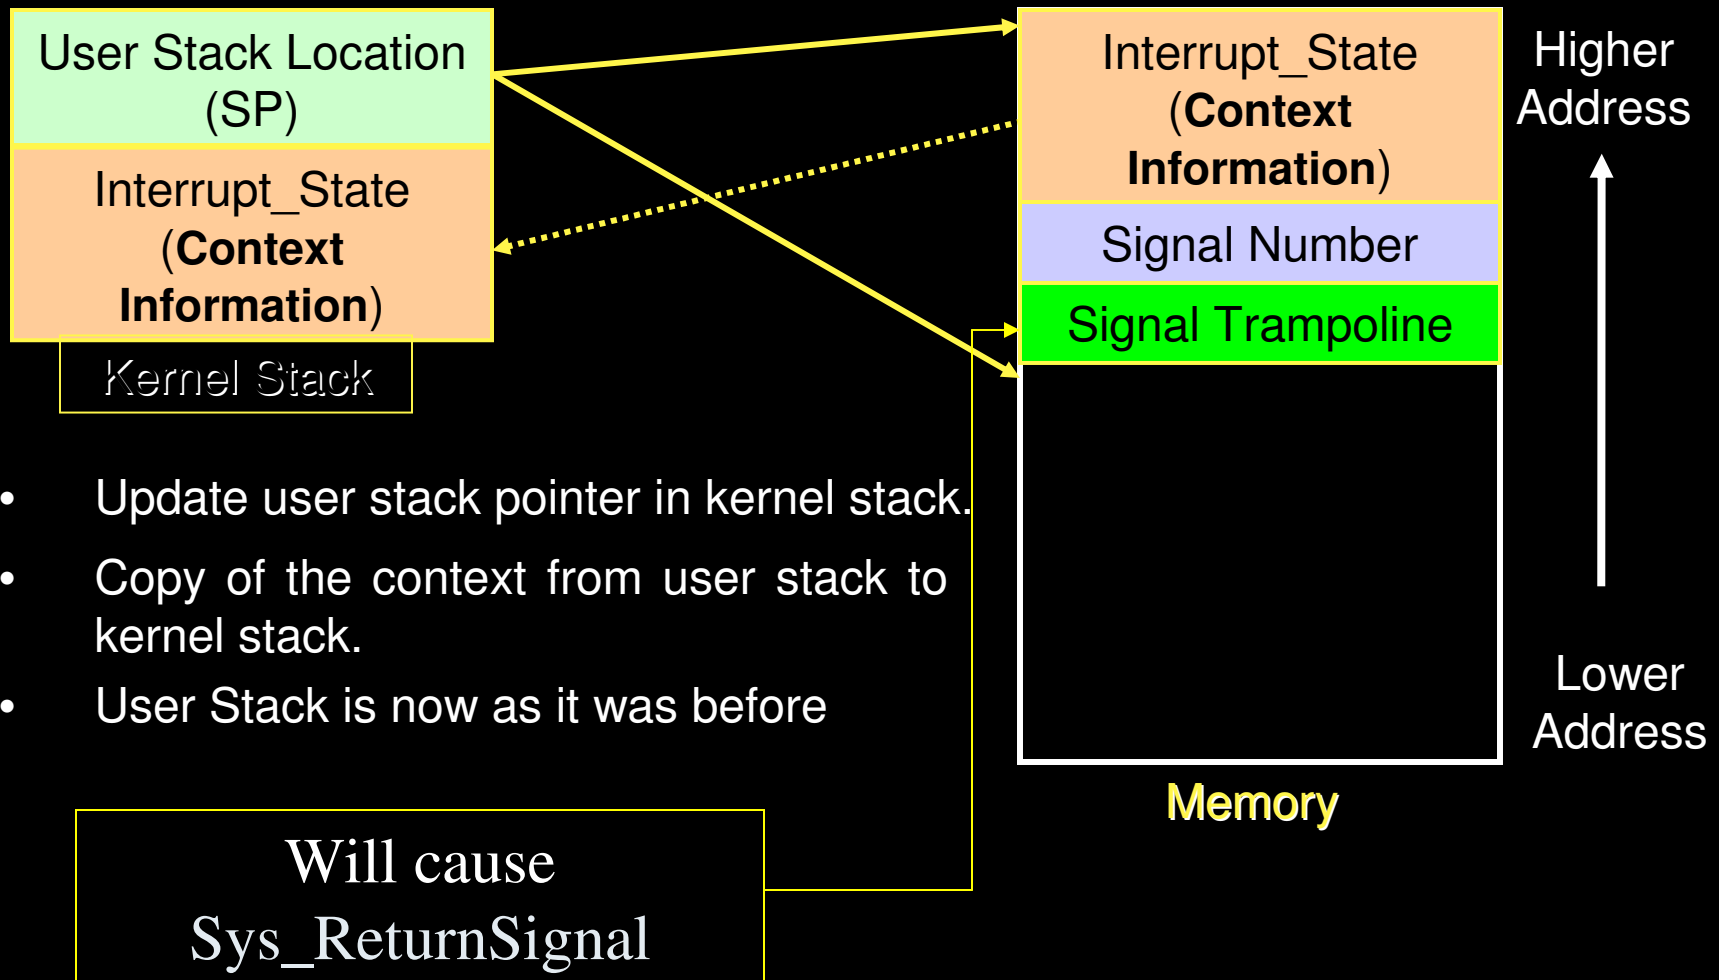
(see include/geekos/defs.h)

1. The process is not currently handling another signal.

# Setup\_Frame



# Complete\_Handler



---

# Project 2 Roadmap++

# Review

---

## Process A

```
main() {  
    for(1000)  
        Print "A"  
    Kill(B, SIGUSR1)  
}
```

## Process B

```
function handler() {  
    Print "HANDLING"  
}  
  
main() {  
    Signal(&handler, SIGUSR1)  
    for(;;)  
        Print "B"  
}
```

## Output

AABBAABBAABB..... HANDLING BBBBBBBBBBBBBBBB.....

1000 A'S

# System Calls

---

**sys\_signal**: register a signal handler

**sys\_kill**: send a signal

**sys\_sigaction**: initialize signal handling for a process

**sys\_waitpid**: wait for any child process to die

**sys\_sigreturn**: indicate completion of signal handler

# System Calls

---

**Sys\_Signal**: register a signal handler

**Sys\_Kill**: send a signal

} Referenced in user code

**Sys\_RegDeliver**: initialize signal handling for a process

**Sys\_WaitNoPID**: wait for any child process to die

**Sys\_ReturnSignal**: indicate completion of signal handler



# Review

---

Process A

```
main() {  
    for(1000)  
        Print "A"  
    Kill(B, SIGUSR1)  
}
```

Process B

```
function handler() {  
    Print "HANDLING"  
}  
  
main() {  
    Signal(&handler, SIGUSR1)  
    for(;;)  
        Print "B"  
}
```

Output

AABBAABBAABB..... HANDLING BBBB.....  
  
1000 A'S

# System Calls

---

**sys\_signal**: register a signal handler

**sys\_kill**: send a signal

**sys\_sigaction**: initialize signal handling for a process

**sys\_waitpid**: wait for any child process to die

**sys\_sigreturn**: indicate completion of signal handler

# System Calls

---

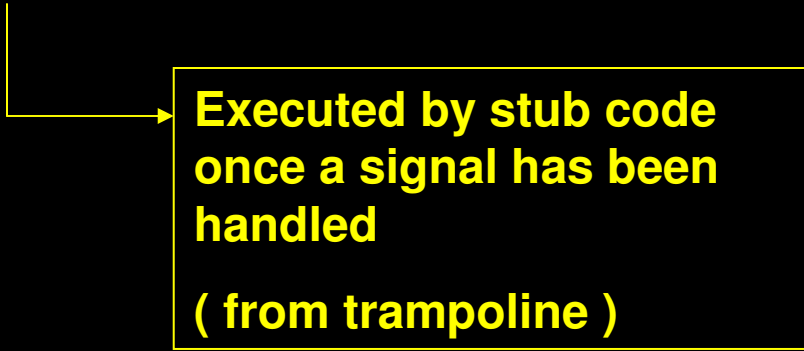
**Sys\_Signal**: register a signal handler

**Sys\_Kill**: send a signal

**Sys\_RegDeliver**: initialize signal handling for a process

**Sys\_WaitNoPID**: wait for any child process to die

**Sys\_ReturnSignal**: indicate completion of signal handler



**Executed by stub code  
once a signal has been  
handled  
( from trampoline )**

# Helper Functions

---

- Send\_Signal
- Set\_Handler
- Check\_Pending\_Signal
- Setup\_Frame
- Complete\_Handler

# Review

---

## Process A

```
main() {  
    for(1000)  
        Print "A"  
    Kill(B, SIGUSR1)  
}
```

## Process B

```
function handler() {  
    Print "HANDLING"  
}  
  
main() {  
    Signal(&handler, SIGUSR1)  
    for(;;)  
        Print "B"  
}
```

# Overview

---

**A**

**B**

# Overview

---

**A**

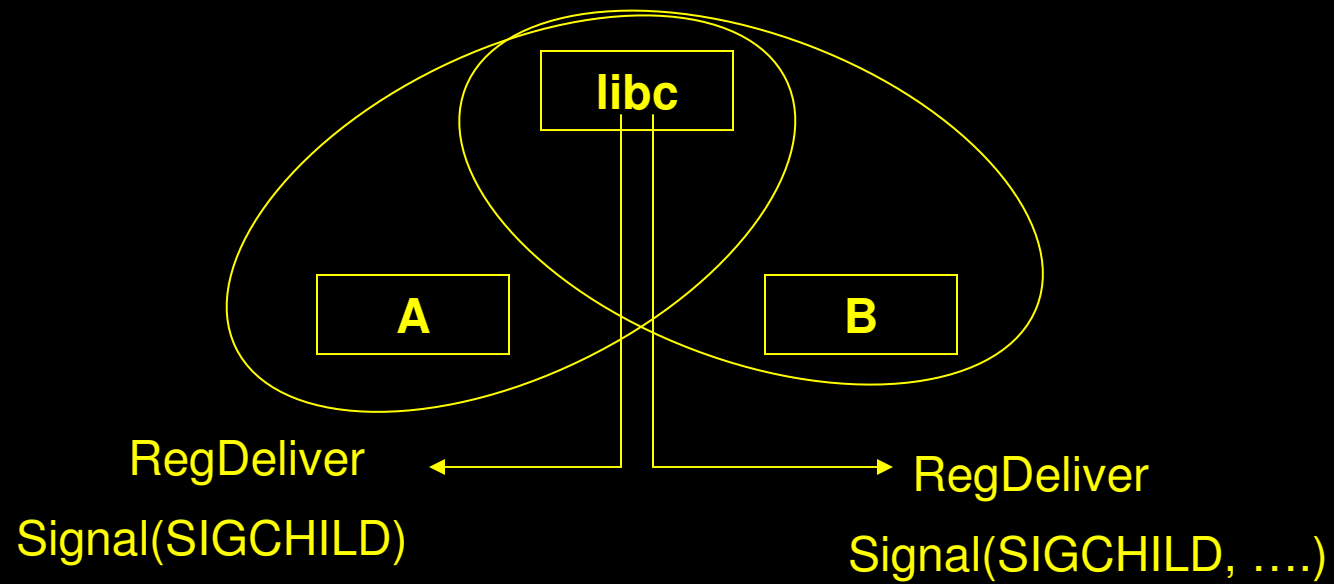
RegDeliver  
Signal(SIGCHILD)

**B**

RegDeliver  
Signal(SIGCHILD, ....)


# Overview

---



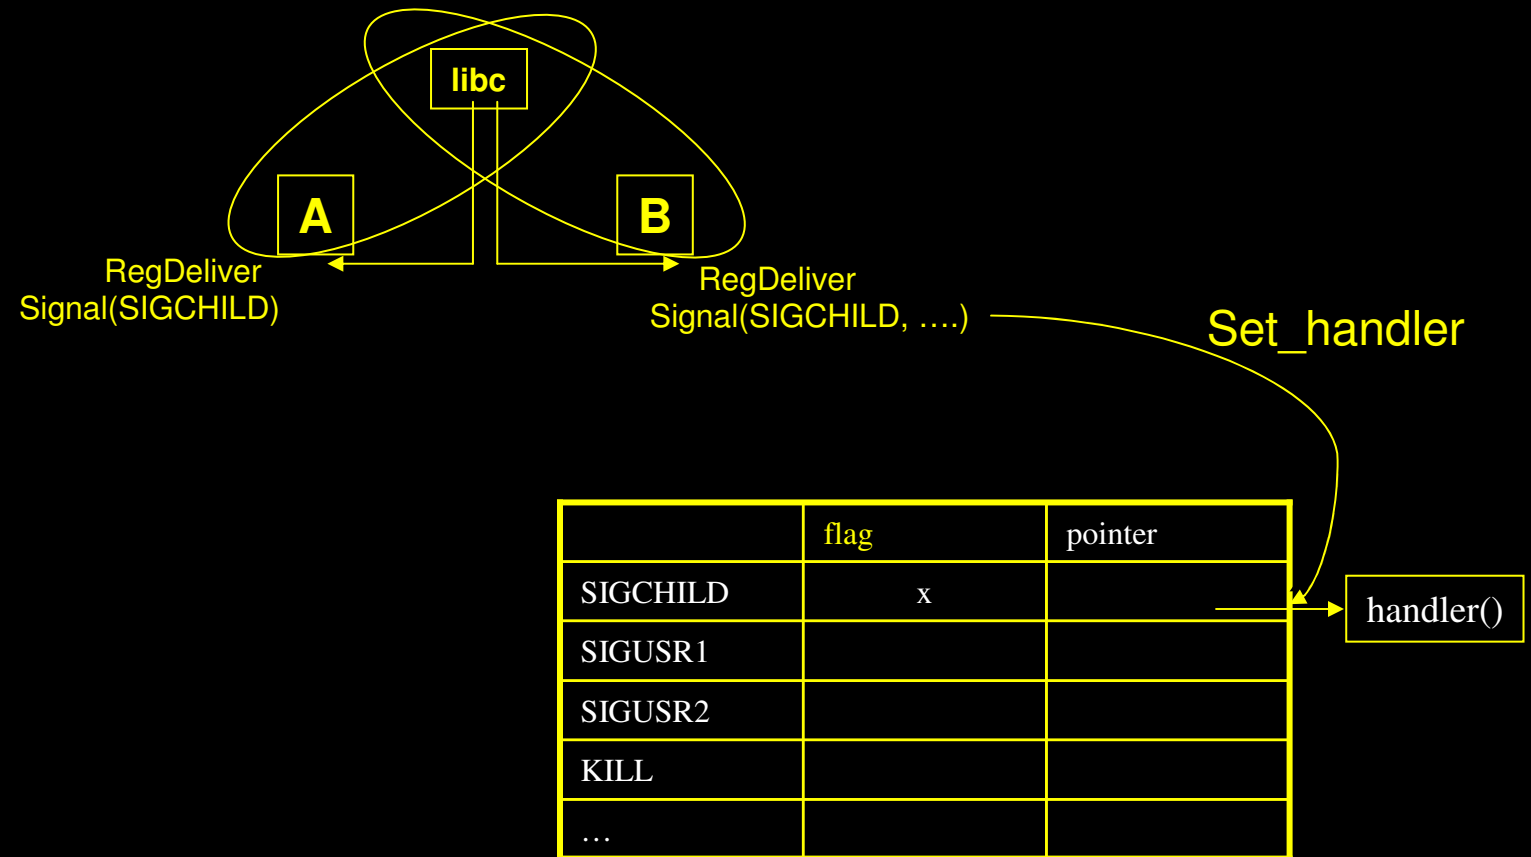


# Overview

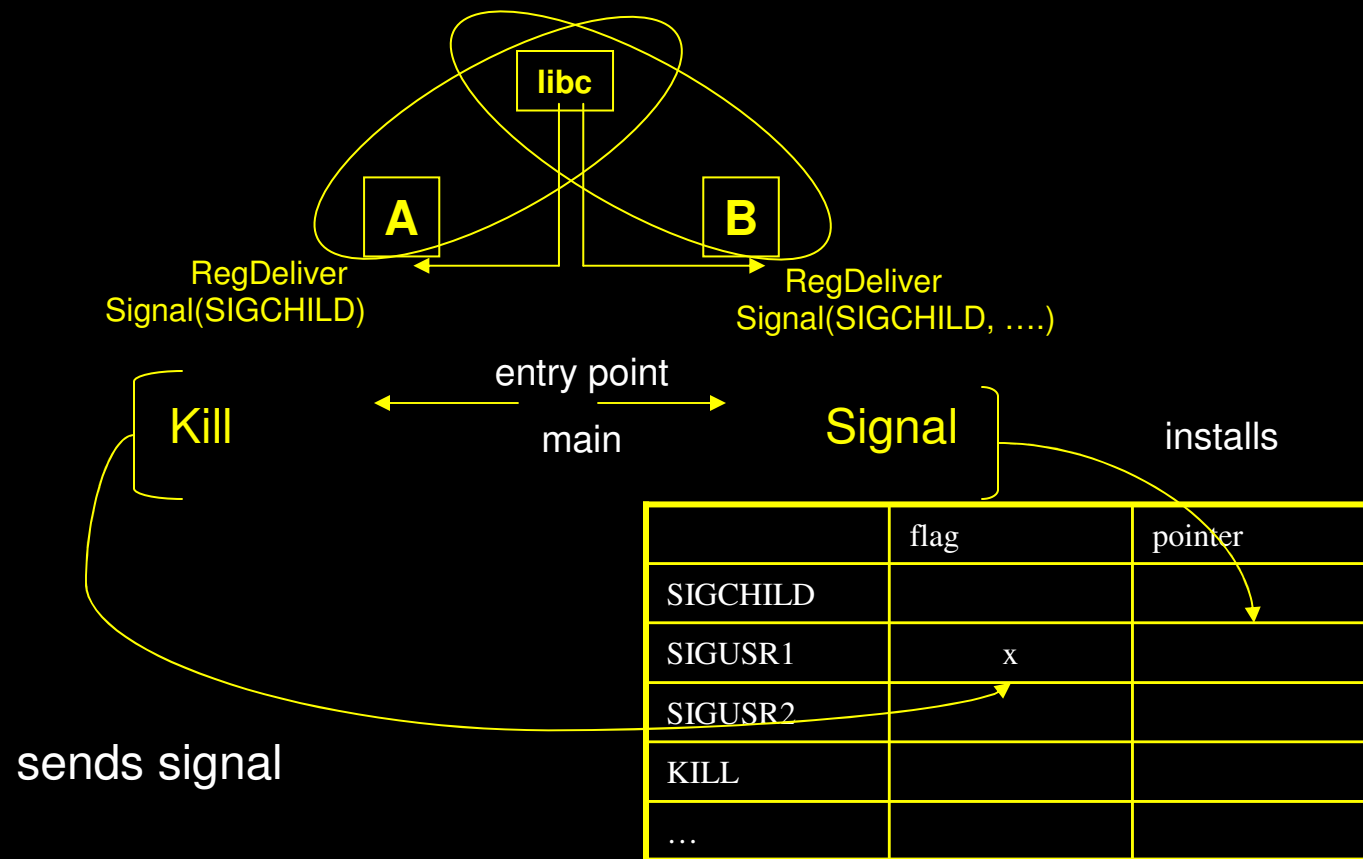


	flag	pointer
SIGCHILD		
SIGUSR1		
SIGUSR2		
KILL		
...		

# Overview



# Overview



# Helper Functions

---

Send\_Signal

Set\_Handler

Check\_Pending\_Signal

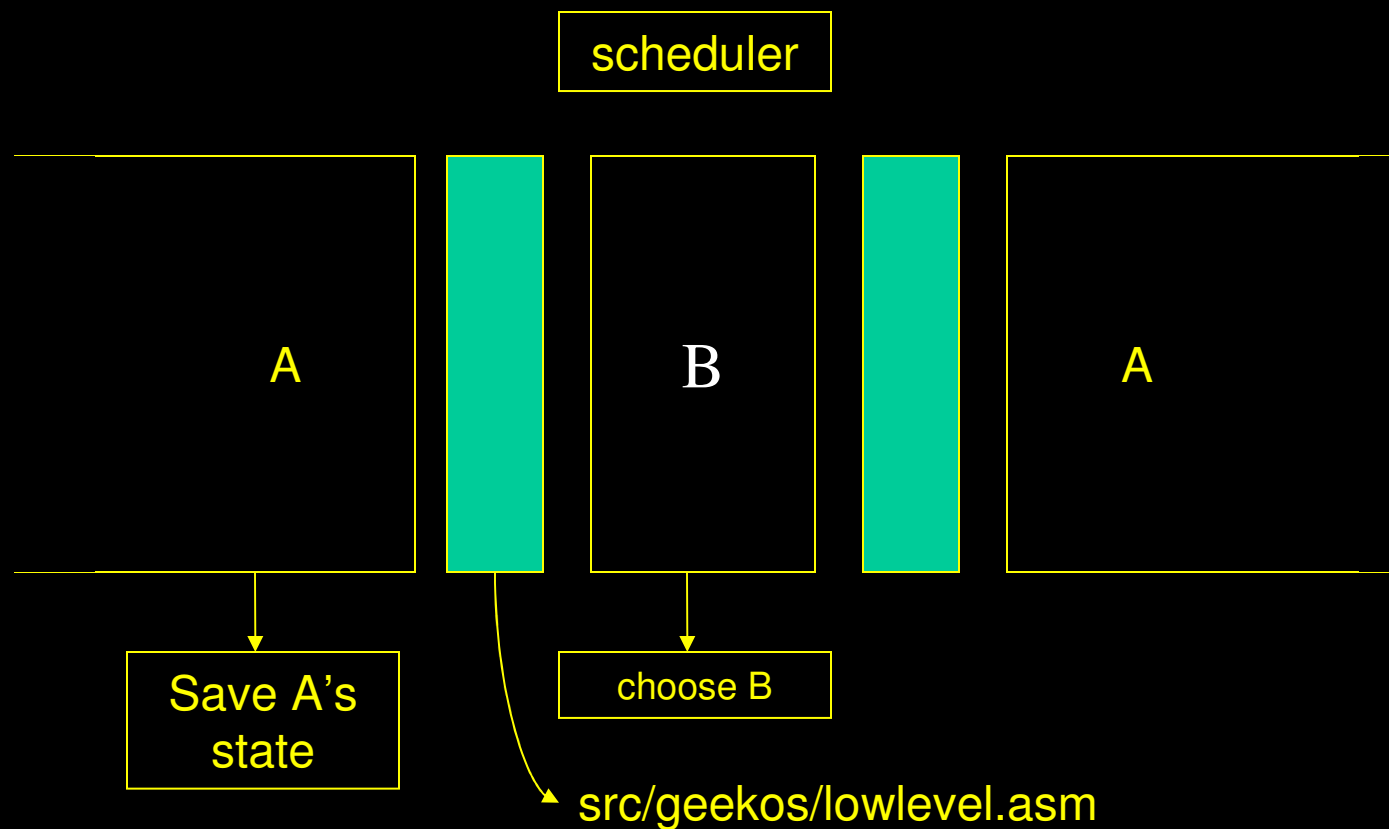
Setup\_Frame

Complete\_Handler

**Look at Scheduler**


# Scheduler: w/o signals

---



# Scheduler: w/ signals

---



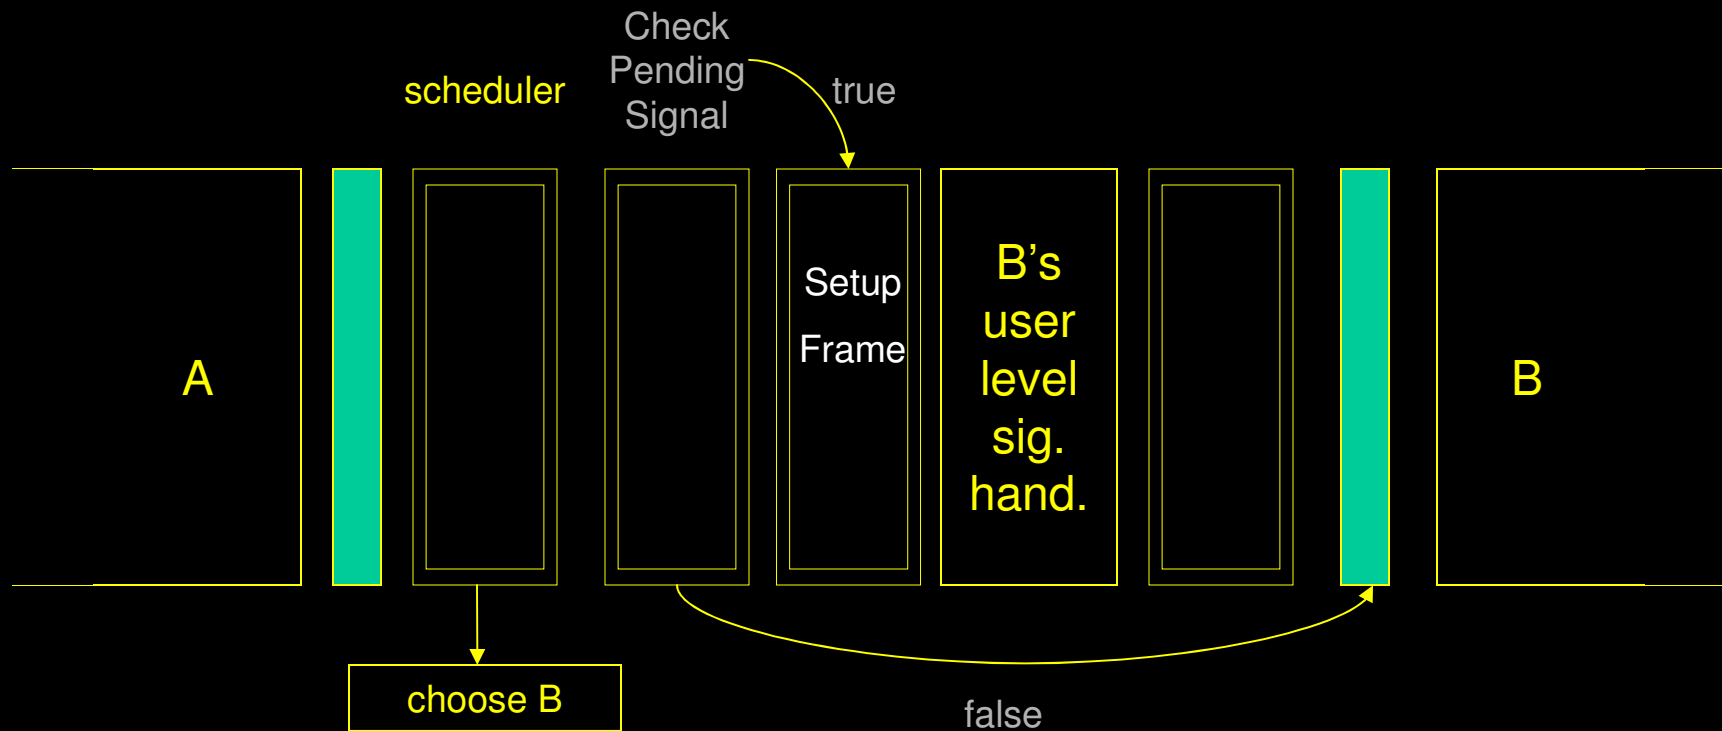
# Check Pending Signal

---

Boolean output

Determines whether to proceed with signal handling

# Scheduler: w/ signals





# Setup Frame

---

Sets up state to enable user-level handling code execution

# Setup Frame

---

Sets up state to enable user-level handling code execution

How are functions called?

# Function Calls

---

Parameter of return address is stored on the stack so  
when finished

- Pop off stack
- Continue execution

Setup Frame

- Enables user stack to keep:
  - Interrupt\_State Vector
  - Return address

# Storing Return Address

---

Want `complete_handler` to execute once user level handling done.

## Hack

- Place address of `return_signal` as return address on stack
- Now `return_signal` stored as function

# Scheduler: w/ signals

---

