

## New Ways of Structuring Code

=====

C. Coroutines that call each other, with interrupts to handle hardware events and global variables used to communicate from ISR's to each coroutine, and between coroutines

## Advantages:

- \* Code is a bit less intertwined; each coroutine has its own while loop that only needs to focus on one task, keeps all its local variables since it has its own stack and doesn't have to return to the caller).
- \* No need for FLP code if we want the appearance of concurrency (but we do need to yield to other coroutines on a regular basis)
- \* Fairly simple to implement: use setjmp/longjmp which is ANSI C and a few lines of custom assembly

## Disadvantages:

- \* (Windows98) One coroutine, if it goes awry, can block other coroutines from running
- \* We need to yield to other coroutines on a regular basis, which is error-prone (depends on code author to remember to do this), non-maintainable, and still a bit of an FLP problem if responsivity is a design goal
- \* Each coroutine needs to know who to call next, which is error-prone and non-maintainable.
- \* Each coroutine needs its own stack, which can eat into available RAM.

D. Coroutines that call a scheduler, with interrupts to handle hardware events and global variables used to communicate from ISR's to each coroutine, and between coroutines

## Advantages:

- \* Code is even less intertwined, since each coroutine no longer has to worry about which coroutine to call next.
- \* Higher level of abstraction: each coroutine is oblivious to the details of coroutine switching, number of coroutines in the system. It just has to call 'yield()' once in a while. More abstraction --> lower complexity for the coroutine developer.
- \* All changes to the code regarding the number of coroutines, and what they are is done in one place (the scheduler).
- \* Scheduler can now explicitly implement prioritization, coroutine creation/termination.
- \* Fairly simple to implement: scheduler can be just another coroutine

## Disadvantages:

- \* One extra coroutine (hence stack and jump buffer) are needed.
- \* We need to yield to the scheduler on a regular basis, which is error-prone (depends on code author to remember to do this), non-maintainable, and still a bit of an FLP problem. Higher level of abstraction does not improve responsivity.
- \* Each coroutine needs its own stack, which can eat into available RAM.

Jun 29, 10 12:39

crsched.c

Page 1/3

```

1  #include <stdio.h>
2  #include <setjmp.h>
3  #include <stdlib.h>
4  #include "inc/hw_memmap.h"
5  #include "inc/hw_types.h"
6  #include "driverlib/gpio.h"
7  #include "driverlib/sysctl.h"
8  #include "driverlib/uart.h"
9  #include "rit128x96x4.h"
10 #include "scheduler.h"
11
12 #define STACK_SIZE 4096    // Amount of stack space for each thread
13
14 typedef struct {
15     int active;           // non-zero means thread is allowed to run
16     char *stack;         // pointer to TOP of stack (highest memory location)
17     jmp_buf state;       // saved state for longjmp()
18 } threadStruct_t;
19
20 // thread_t is a pointer to function with no parameters and
21 // no return value...i.e., a user-space thread.
22 typedef void (*thread_t)(void);
23
24 // These are the external user-space threads. In this program, we create
25 // the threads statically by placing their function addresses in
26 // threadTable[]. A more realistic kernel will allow dynamic creation
27 // and termination of threads.
28 extern void thread1(void);
29 extern void thread2(void);
30
31 static thread_t threadTable[] = {
32     thread1,
33     thread2
34 };
35 #define NUM_THREADS (sizeof(threadTable)/sizeof(threadTable[0]))
36
37 // These static global variables are used in scheduler(), in
38 // the yield() function, and in threadStarter()
39 static jmp_buf scheduler_buf; // saves the state of the scheduler
40 static threadStruct_t threads[NUM_THREADS]; // the thread table
41 unsigned currThread; // The currently active thread
42
43 // This function is called from within user thread context. It executes
44 // a jump back to the scheduler. When the scheduler returns here, it acts
45 // like a standard function return back to the caller of yield().
46 void yield(void)
47 {
48     if (setjmp(threads[currThread].state) == 0) {
49         // yield() called from the thread, jump to scheduler context
50         longjmp(scheduler_buf, 1);
51     } else {
52         // longjmp called from scheduler, return to thread context
53         return;
54     }
55 }
56
57 // This is the starting point for all threads. It runs in user thread
58 // context using the thread-specific stack. The address of this function
59 // is saved by createThread() in the LR field of the jump buffer so that
60 // the first time the scheduler() does a longjmp() to the thread, we
61 // start here.
62 void threadStarter(void)
63 {

```

Jun 29, 10 12:39

crsched.c

Page 2/3

```

64 // Call the entry point for this thread. The next line returns
65 // only when the thread exits.
66 (*(threadTable[currThread]))();
67
68 // Do thread-specific cleanup tasks. Currently, this just means marking
69 // the thread as inactive. Do NOT free the stack here because we're
70 // still using it! Remember, this function runs in user thread context.
71 threads[currThread].active = 0;
72
73 // This yield returns to the scheduler and never returns back since
74 // the scheduler identifies the thread as inactive.
75 yield();
76 }
77
78 // This function is implemented in assembly language. It sets up the
79 // initial jump-buffer (as would setjmp()) but with our own values
80 // for the stack (passed to createThread()) and LR (always set to
81 // threadStarter() for each thread).
82 extern void createThread(jmp_buf buf, char *stack);
83
84 // This is the "main loop" of the program.
85 void scheduler(void)
86 {
87     unsigned i;
88
89     currThread = -1;
90
91     do {
92         // It's kinda inefficient to call setjmp() every time through this
93         // loop, huh? I'm sure your code will be better.
94         if (setjmp(scheduler_buf)==0) {
95
96             // We saved the state of the scheduler, now find the next
97             // runnable thread in round-robin fashion. The 'i' variable
98             // keeps track of how many runnable threads there are. If we
99             // make a pass through threads[] and all threads are inactive,
100            // then 'i' will become 0 and we can exit the entire program.
101            i = NUM_THREADS;
102            do {
103                // Round-robin scheduler
104                if (++currThread == NUM_THREADS) {
105                    currThread = 0;
106                }
107
108                if (threads[currThread].active) {
109                    longjmp(threads[currThread].state, 1);
110                } else {
111                    i--;
112                }
113            } while (i > 0);
114
115            // No active threads left. Leave the scheduler, hence the program.
116            return;
117
118        } else {
119            // yield() returns here. Did the thread that just yielded to us exit? If
120            // so, clean up its entry in the thread table.
121            if (! threads[currThread].active) {
122                free(threads[currThread].stack - STACK_SIZE);
123            }
124        }
125    } while (1);
126 }

```

Jun 29, 10 12:39

crsched.c

Page 3/3

```

127
128 void main(void)
129 {
130     unsigned i;
131
132     // Set the clocking to run directly from the crystal.
133     SysCtlClockSet(SYSCTL_SYSDIV_1 | SYSCTL_USE_OSC | SYSCTL_OSC_MAIN |
134                   SYSCTL_XTAL_8MHZ);
135
136     // Initialize the OLED display and write status.
137     RIT128x96x4Init(1000000);
138     RIT128x96x4StringDraw("Scheduler Demo",          20,  0, 15);
139
140     // Enable the peripherals used by this example.
141     SysCtlPeripheralEnable(SYSCTL_PERIPH_UART0);
142     SysCtlPeripheralEnable(SYSCTL_PERIPH_GPIOA);
143
144     // Set GPIO A0 and A1 as UART pins.
145     GPIOPinTypeUART(GPIO_PORTA_BASE, GPIO_PIN_0 | GPIO_PIN_1);
146
147     // Configure the UART for 115,200, 8-N-1 operation.
148     UARTConfigSetExpClk(UART0_BASE, SysCtlClockGet(), 115200,
149                        (UART_CONFIG_WLEN_8 | UART_CONFIG_STOP_ONE |
150                         UART_CONFIG_PAR_NONE));
151
152     // Create all the threads and allocate a stack for each one
153     for (i=0; i < NUM_THREADS; i++) {
154         // Mark thread as runnable
155         threads[i].active = 1;
156
157         // Allocate stack
158         threads[i].stack = (char *)malloc(STACK_SIZE) + STACK_SIZE;
159         if (threads[i].stack == 0) {
160             iprintf("Out of memory\r\n");
161             exit(1);
162         }
163
164         // After createThread() executes, we can execute a longjmp()
165         // to threads[i].state and the thread will begin execution
166         // at threadStarter() with its own stack.
167         createThread(threads[i].state, threads[i].stack);
168     }
169
170     // Start running coroutines
171     scheduler();
172
173     // If scheduler() returns, all coroutines are inactive and we return
174     // from main() hence exit() should be called implicitly (according to
175     // ANSI C). However, TI's startup_gcc.c code (ResetISR) does not
176     // call exit() so we do it manually.
177     exit(0);
178 }
179
180 /*
181  * Compile with:
182  * ${CC} -o crsched.elf -I${STELLARISWARE} -L${STELLARISWARE}/driverlib/gcc
183  *      -Tlinkscript.x -Wl,-Map,crsched.map -Wl,--entry,ResetISR
184  *      crsched.c create.S threads.c startup_gcc.c syscalls.c rit128x96x4.c
185  *      -ldriver
186  */
187 // vim: expandtab ts=2 sw=2 cindent

```

Jun 29, 10 12:45

threads.c

Page 1/1

```
1  #include <stdio.h>
2  #include "scheduler.h"
3
4  // These are the user-space threads. Note that they are completely oblivious
5  // to the technical concerns of the scheduler. The only interface to the
6  // scheduler is the single function yield() and the global variable
7  // currThread which indicates the number of the thread currently
8  // running.
9
10 void thread1(void)
11 {
12     unsigned count;
13
14     for (count = 0; count < 10; count++) {
15         iprintf("In thread %u -- pass %d\r\n", currThread, count);
16         yield();
17     }
18 }
19
20 void thread2(void)
21 {
22     unsigned count;
23
24     for (count=0; count < 5; count++) {
25         iprintf("In thread %u -- pass %d\r\n", currThread, count);
26         yield();
27     }
28 }
```

Jun 29, 10 12:46

## create.S

Page 1/1

```

1  /*
2  * Implement the thread creation task:
3  *
4  *   - initialize the jump buffer with appropriate values for
5  *     R13 (stack) and R14 (first address to jump to)
6  *   - all other registers are irrelevant upon thread creation
7  *
8  *   In the jump buffer, the R13 slot is set to the second parameter of this
9  *   function (the top-of-stack address, passed in R1). The R14 slot is set to
10 *   the address of the threadStarter() function.
11 *
12 *   The C prototype for this function call is:
13 *       createThread(threads[i].state, threads[i].stack)
14 *   thus:
15 *       R0 ←- state (a setjmp()-style jump buffer)
16 *       R1 ←- stack (address of top-of-stack)
17 */
18     .syntax unified
19     .text
20     .align 2
21     .thumb
22     .thumb_func
23     .type createThread,function
24     .global createThread
25 createThread:
26
27     /* Save registers in the jump buffer. Their values are
28     not important when the thread is first created. This line is the same as
29     the first two lines of setjmp(), except we don't save SP and LR since we
30     want to set these to our own values. Really, the only point of this
31     instruction is to advance R0 to the right location in the jump buffer for
32     pointing to SP (without having to do any math :-) */
33
34     stmea  r0!, { r4-r10, r11 }
35
36     /* Now we save SP and LR in that order. SP is the R1 parameter, and we have
37     * to get the address of threadStarter() into a higher register (so they are
38     * placed in the jump buffer in the right order). */
39
40     ldr    R2, .L0
41     stmea  R0!, { R1, R2 }    @ Store "SP" and "LR" for the new thread
42
43     bx    lr
44
45 .L0:
46     .word  threadStarter

```

## New Ways of Structuring Code

=====

D.

## Disadvantages:

- \* We need to yield to the scheduler on a regular basis, which is error-prone (depends on code author to remember to do this), non-maintainable, and still a bit of an FLP problem. Higher level of abstraction does not improve responsiveness.

E. Time-sliced threads, with interrupts to handle hardware events and global variables used to communicate from ISR's to each thread, and between threads

## Advantages:

- \* No need to explicitly yield to a scheduler...the scheduler forcibly interrupts (i.e., pre-empts) a thread when its "time slice" is up. This is a natural use for the SysTick timer.
- \* The code is minimally intertwined and there are only well-defined interfaces between threads<-->scheduler and thread<-->thread.
- \* If time slices are small enough, the system provides the "appearance of concurrency" and has very good latency and responsiveness.

## Disadvantages:

- \* Complex: interrupt-driven system is easy to get wrong, debugging is difficult.
- \* Each thread needs its own stack.
- \* The "concurrency problem". When program state is saved/restored at an explicit function call boundary (i.e., yield) then behavior is (mostly) deterministic. But when program state is saved/restored at any assembly-language instruction boundary, then behavior is non-deterministic and previously-true assumptions no longer hold.

Jul 03, 12 13:03

**threads\_preemptive.c**

Page 1/1

```
1  #include <stdio.h>
2  #include "scheduler.h"
3
4  // This is "the next level": threads that are pre-emptively interrupted
5  // in order to return control to a central scheduler, and are mostly
6  // entirely unconcerned with (a) other threads, and (b) the scheduler itself.
7  //
8  // NOTE: A thread can be interrupted AT ANY TIME, even in the middle of
9  // a function call, on ANY assembly language boundary.
10
11 void thread1(void)
12 {
13     unsigned count;
14
15     for (count = 0; count < 10; count++) {
16         iprintf("In thread %u -- pass %d\r\n", currThread, count);
17     }
18 }
19
20 void thread2(void)
21 {
22     unsigned count;
23
24     for (count=0; count < 5; count++) {
25         iprintf("In thread %u -- pass %d\r\n", currThread, count);
26     }
27 }
```



So what do you need to do?

=====

1. Program SysTick for periodic interrupts:
  - why are shorter periods better?
  - why are shorter periods worse?
2. On a SysTick interrupt, do a "forced setjmp", i.e., save the state of the processor as for any exception.
3. Determine which thread should run next (round-robin? priorities?)
4. Restore the state of the processor, not of the thread that was interrupted, but of the thread that is to run.

Things that need to change

=====

- \* threadStruct\_t: jmp\_buf is not enough to save exception state
 

```
typedef struct {
    int active;           // non-zero means thread is allowed to run
    char *stack;         // pointer to TOP of stack
    jmp_buf state;       // saved state for longjmp()
} threadStruct_t;
```
  - \* yield(): must simulate a SysTick exception, i.e., cause a context switch
- NOTE: Simply sitting in a loop waiting until the next SysTick interrupt is NOT OK!!!
- \* createThread: must simulate saved exception state, instead of simulating setjmp state
    - What should exception return address be?
    - What should R14 be?
    - What should R13 be?
    - What should R0-R12 be?
  - \* scheduler(): goes away...scheduling functions are done in the SysTick ISR
  - \* main(): how do you kickstart the process? change privilege levels? change to process stack?