Bootstrap, Memory Management and Troubleshooting

(slides are based on Prof. Dr. Jian-Jia Chen and http://www.freertos.org)

Anas Toma

LS 12, TU Dortmund

February 01, 2018

Bootstrapping

- The entire boot process of a system, from application of power to performing its intended function
- Usually just called booting
- Procedure in PC (e.g., Linux)
  - Power on
  - BIOS gets control
  - BIOS initializes some hardware
  - BIOS loads bootloader
  - Bootloader loads operating system kernel
  - Kernel probes hardware
  - Kernel finds and mounts root filesystem
  - Kernel runs init
  - Init gets userspace up and running

How to Boot Embedded Systems?

- No BIOS; initial control passes to bootloader
- Limited options for storage
- Hardware usually remains the same
- Constraints on storage and memory: only need minimal root filesystem and userspace

Outline

Bootstrapping

Memory Management

Troubleshooting

Memory Management

- Subdividing memory to accommodate multiple tasks
- Memory needs to be allocated to ensure a reasonable supply of ready tasks to consume available processor time
- Memory allocation:
  - Programmers do not know where the program will be placed in memory when it is executed
  - While the program is executing, it may be swapped to external memory (disks etc.) and returned to main memory at a different location (for general OS)
  - Memory references must be translated in the code to actual physical memory address
  - This is implemented by calling malloc() in Linux and Unix
- Memory deallocation:
  - The unused memory space is given back to OS so that other tasks can use this part of memory again
  - This is implemented by calling free() in Linux and Unix

Memory Management in Embedded Systems

- The RTOS kernel has to allocate RAM each time a task, queue or semaphore is created. The malloc() and free() functions can sometimes be used for this purpose, but ...
  1. they are not always available on embedded systems,
  2. take up valuable code space,
  3. are not thread safe, and
  4. are not deterministic (the amount of time taken to execute the function will differ from call to call)
- An alternative scheme is required. One embedded / real time system can have very different RAM and timing requirements to another - so a single RAM allocation algorithm will only ever be appropriate for a subset of applications.
  - The memory allocation API is included in the RTOS portable layer
  - When the real time kernel in FreeRTOS requires RAM, instead of calling malloc() it makes a call to pvPortMalloc(). When RAM is being freed, instead of calling free() the real time kernel makes a call to vPortFree().

Default Schemes in FreeRTOS

- heap_1.c: This is the simplest scheme of all. It does not permit memory to be freed once it has been allocated, but despite this is suitable for a large number of applications.
  - Can be used if the application never deletes a task or queue (no calls to vTaskDelete() or vQueueDelete() are ever made).
  - Is always deterministic (always takes the same amount of time to return a block).
- heap_2.c: This scheme uses a best fit algorithm and, unlike scheme 1, allows previously allocated blocks to be freed. It does not combine adjacent free blocks into a single large block.
  - It can be used even when the application repeatedly calls vTaskCreate()/vTaskDelete() or vQueueCreate()/vQueueDelete()
  - It could possible result in memory fragmentation problems should your application create blocks of queues and tasks in an unpredictable order.
  - It is not deterministic - but is also not particularly inefficient.
- heap_3.c: This is just a wrapper for the standard malloc() and free() functions. It makes them thread safe.

Case 1: heap_1.c

```c
void *pvPortMalloc( size_t xWantedSize ){
void *pvReturn = NULL;
    /* Ensure that blocks are always aligned to the required number of bytes. */
    #if portBYTE_ALIGNMENT != 1
        if( xWantedSize & portBYTE_ALIGNMENT_MASK )
        {
            /* Byte alignment required. */
            xWantedSize += ( portBYTE_ALIGNMENT - ( xWantedSize & portBYTE_ALIGNMENT_MASK ) );
        }
    #endif
    vTaskSuspendAll();
    {
        /* Check there is enough room left for the allocation. */
        if( ( ( xNextFreeByte + xWantedSize ) < configTOTAL_HEAP_SIZE ) &&
            ( ( xNextFreeByte + xWantedSize ) > xNextFreeByte ) )/* Check for overflow. */
        {
            /* Return the next free byte then increment the index past this
            block. */
            pvReturn = &( xHeap.ucHeap[ xNextFreeByte ] );
            xNextFreeByte += xWantedSize;
        }
    }
    xTaskResumeAll();
    #if( configUSE_MALLOC_FAILED_HOOK == 1 )
    {
        if( pvReturn == NULL )
        {
            extern void vApplicationMallocFailedHook( void );
            vApplicationMallocFailedHook();
        }
    }
    #endif
    return pvReturn;
}
```

Placement Algorithms

Different strategies may be taken as to how space is allocated to processes: (reference "Operating System Concepts" by Abraham Silberschatz, Peter B. Galvin (Author), and Greg Gagne.)

- First Fit: Allocate the first hole that is big enough. Searching may start either at the beginning of the set of holes or where the previous first-fit search ended.
- Best Fit: Allocate the smallest hole that is big enough. The entire list of holes must be searched unless it is sorted by size. This strategy produces the smallest leftover hole.
- Worst Fit: Allocate the largest hole. In contrast, this strategy aims to produce the largest leftover hole, which may be big enough to hold another process.
- Experiments have shown that both first fit and best fit are better than worst fit in terms of decreasing time and storage utilization. First fit is generally faster.

Case 2: heap_2.c

- It is more involving by using a best fit algorithm.
- We will look into the source code directly.

Memory Protection

- Using a Memory Protection Unit (MPU) can protect applications from a number of potential errors, ranging from undetected programming errors to errors introduced by system or hardware failures.
- In FreeRTOS: FreeRTOS-MPU
  - It can be used to protect the kernel itself from invalid execution by tasks and protect data from corruption.
  - It can also protect system peripherals from unintended modification by tasks and guarantee the detection of task stack overflows.

Creating Restricted Tasks in FreeRTOS

- The created task can run in either Privileged or User modes.
  - When Privileged mode it used the task will have access to the entire memory map
  - When User mode is used the task will have access to only its stack.
  - In both cases the MPU will not automatically catch stack overflows.
- If a task wants to use the MPU then the following additional information has to be provided:
  - The address of the task stack.
  - The start, size and access parameters for up to three user definable memory regions.
  - The memory regions allocated to a task can be changed using vTaskAllocateMPURegions().
- It is implemented in xTaskCreateRestricted() in task.h
- A Privileged mode task can call portSWITCH_TO_USER_MODE() to set itself into User mode. A task that is running in User mode cannot set itself into Privileged mode.

Outline

Outline

- Susceptible areas
  - Memory
  - Timeliness
  - Concurrency issues
- Tools and techniques
  - Code coverage tools
  - Unit tests
  - JTAG debugging

Memory issues - Stack

- Every thread in FreeRTOS has individual stack
- Stack requirement is often unpredictable
- Most common cause of spurious failures
- Particular high stack usage with library functions
  - printf, sprintf: better to write your own lightweight variant if the whole functionality is not required
- API Help:
  - uxTaskGetStackHighWaterMark(...): for testing phase
  - configCHECK_FOR_STACK_OVERFLOW: for runtime

Memory issues - Stack (2)

configCHECK_FOR_STACK_OVERFLOW

- **Calls a hooked function if a stack overflow is detected**
- **Checks made during the context switch**
- **Three options possible**
  - **=0 : no checks**
  - **=1 : checks the current value of stack pointer**
  - **=2 : checks the value of guard bytes between the stack spaces of different threads**
- **Makes context switch slower**

Memory Issues - Memory Allocation

- `malloc()` and `free()`: often not a good idea for embedded systems
- Dynamic memory allocation seldom used on safety critical parts of embedded systems due to:
  - **Sufficiency:** will a critical memory demand always be met
  - **Fragmentation:** what if all the available chunks are smaller than required chunks
  - **Garbage collection:** can this process be time-bounded
  - **Timeliness:** What is the upper bound on timeliness of fulfilling a memory request
- Static memory allocation
  - might be wasteful
  - inflexible
  - but less error prone

Timeliness

Most important concern in RT embedded systems.

- Offline: verification using static analysis
  - WCET tools: aiT, Chronos et al.
  - Scheduling policy analysis using Real-time calculus
- Online: In system verification using tracing
- FreeRTOS allows tracing:
  - **Context switch** time, reason
  - **Queue** create, send, receive, peek, delete
  - **Mutex** create, give, failed
  - **Semaphore** create, give, failed
  - **Task** create, delay, resume, priority set, delete
- More can also be added
- Heisenberg bugs: Instrumenting the code changes the behavior of the code

Timeliness - Tracing Utility from FreeRTOS

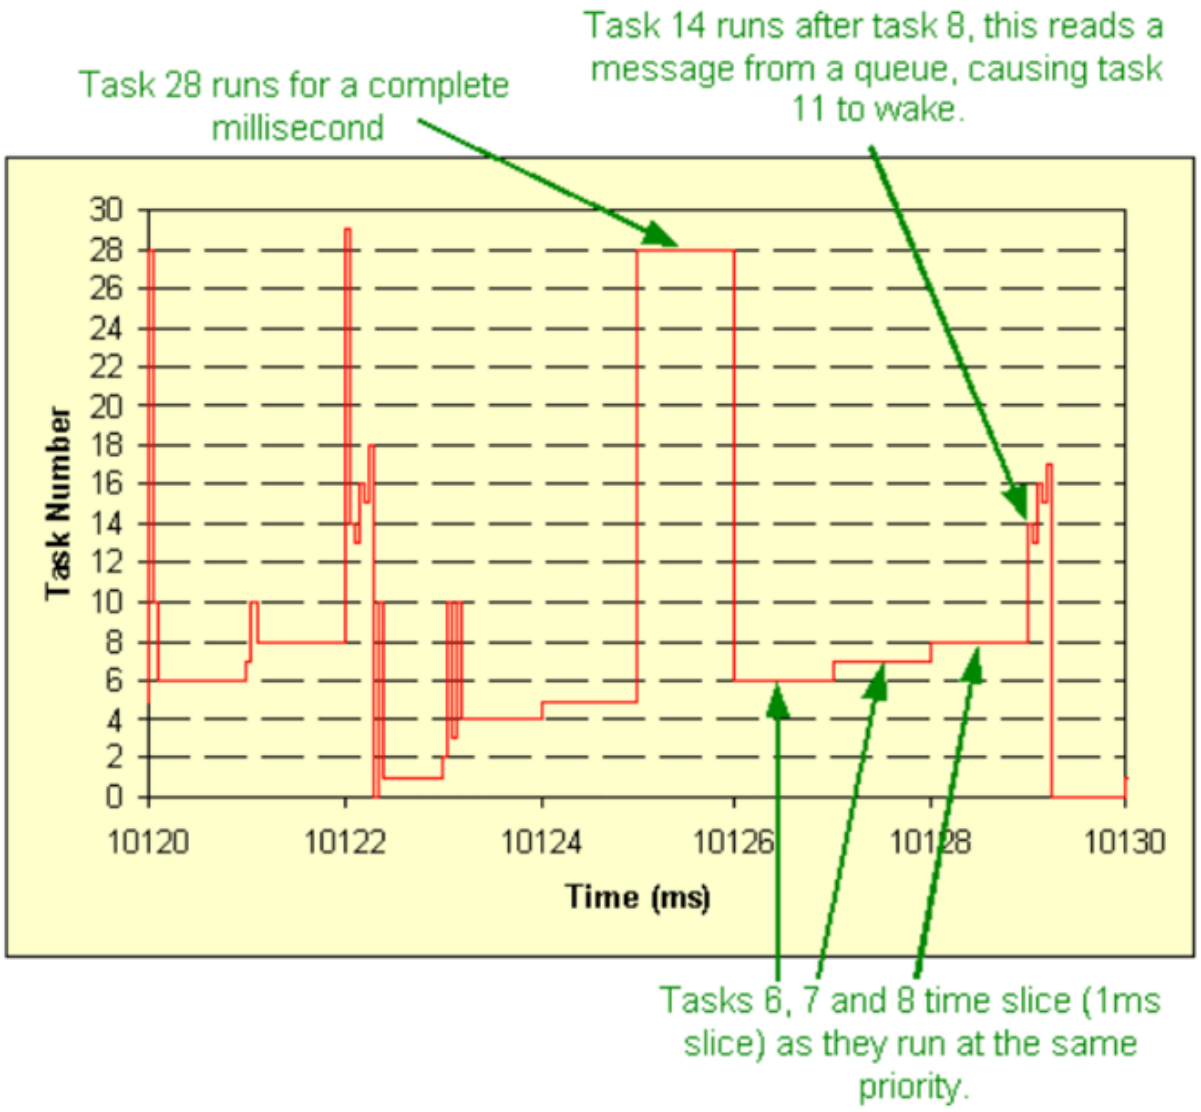

Concurrency issues

- Race conditions
  - The output is dependent on the sequencing or timing of the input
  - Resource Access should be carefully planned
  - Priorities inverted
- Interrupt priorites can cause problems
  - Nested interrupts can result in:
    - deadlocks
    - Interrupt misses

Useful Tools

- The right tool: provides the right level of abstraction
- Example: developing an NIC card

	Development Stage	Debugging Tools
1.	Hardware development	Oscilloscope, Logic analyzer
2.	Driver development	Data monitor, stack browser
3.	Protocol development	Wireshark, network analyzer

- Possible tools for memory issues:
  - valgrind, nmon, Intel parallel inspector
  - Tools only for Linux/x86

Trace32 - A nonfree FreeRTOS tool

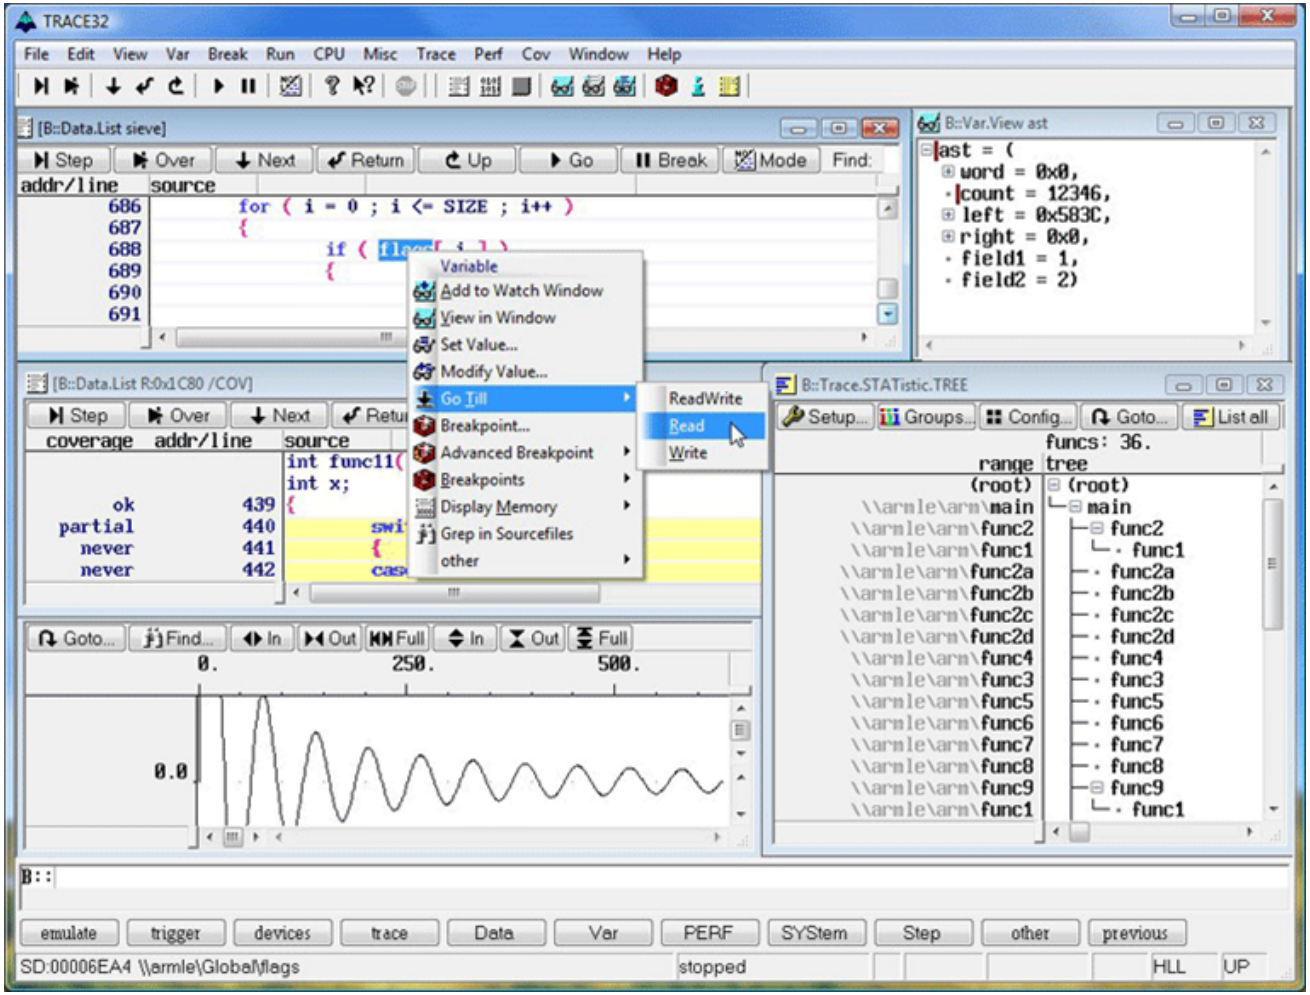

Unit Tests

- "Code a little, test a little" scheme
- Test the smallest possible units of code (function) in isolation from the complete application

- Saves time in integration
- Tools available for C
  - CUnit
  - Check

JTAG Debugging

- stands for Joint test action group
- Developed in 1980s by a consortia of over 200 member companies
- Main idea: have the test facilities / test points into the chips
- Standardized protocol using 5 pins
- The hardware implementation is normally available on chip, and can be accessed through serial/USB
- Allows single stepping while being in circuit, memory/register content reading + editing

JTAG via GDB and Eclipse

Debug - Hello_World/hello_world.c - Eclipse

File Edit Source Refactor Navigate Search Run Project Window Help

Debug

Step Over (F6)

Hello_World Automatic [C/C++ Remote Application]
gdbserver debugger (8/11/11 1:52 PM) (Suspended)
Thread [1] (Suspended)
1 main() hello_world.c:5 0x00008394
Remote Shell
arm-linux-gdb (8/11/11 1:52 PM)
/home/user/workspace/Hello_World/hello_world (8/11/11 1:52 PM)

Variables | Breakpoints | Registers | Modules

Name	Value
(x)= i	2

hello_world.c | Makefile

```
#include <stdio.h>

int main (void) {
    int i = 1;
    for(i = 0; i < 100; i++) {
        printf("Hello World.\n");
    }
    return 0;
}
```

Outline | Disassembly | Enter location here

```
0000838c:  ldr r0, [pc, #44]      ; 0x83c0 <main+88>
00008390:  bl 0x8294 <puts>
 5           for(i = 0; i < 100; i++) {
00008394:  ldr r3, [r11, #-16]
00008398:  add r3, r3, #1  ; 0x1
0000839c:  str r3, [r11, #-16]
000083a0:  ldr r3, [r11, #-16]
000083a4:  cmp r3, #99     ; 0x63
000083a8:  ble 0x838c <main+36>
 8           return 0;
```

Tasks | Problems | Remote Systems | Terminals | Console | Memory

Hello_World Automatic [C/C++ Remote Application] Remote Shell

```
Remote debugging from host 172.21.10.65
Hello World.
Hello World.
Hello World.
```