Gabriele Proietti Mattia

# Advanced Operating Systems and Virtualization

[6] Time Management



Department of Computer, Control and Management Engineering "A. Ruberti", Sapienza University of Rome

gpm.name · proiettimattia@diag.uniroma1.it

A.Y. 2020/2021 · V1

### Outline

- 1. Introduction
- 2. Timekeeping Architecture
  - 1. Low-resolution Timers
  - 2. Generic Time Subsystem
- 3. Watchdogs

6.1

6. Time Management

# Introduction



Advanced Operating Systems and Virtualization

# **Time keeping**

A computer would be useless if programs would not have the possibility to keep track of time passing. This fundamental facility is handled by the kernel in two ways:

- 1. keeping the current time and date so they can be returned to user programs (e.g. time(), gettimeofday(), ...)
- maintaining timers, mechanisms that are able to notify the kernel or a user program (e.g. setitimer() and alarm()) that a certain interval of time has elapsed.

Time measurements are performed by hardware circuits based on oscillators.

# **Clock and Timer Circuits**

In the 8o86 architecture the CPU interacts with clock circuits that keep track of the time of day and timer circuits, programmed by the kernel so that they issue interrupts at a desired time. We have:

- **Real Time Clock** (RTC), which is independent of all the other chips and keeps track of the time of day, ticking even if the PC is off. The RTC is capable of issuing periodic interrupts on IRQ8 and can be programmed and used as an alarm. Linux uses the RTC only for time and date (/dev/rtc).
- **Timestamp Counter** (TSC), it is a counter incremented from an external oscillator connected at the CLK pin of the CPU, it can be read with the asm instruction rdtsc.
- **Programmable Interval Timer** (PIT), its role is similar to an alarm clock of a microwave oven, when timer ends it does not ring a bell but it issues an interrupt. The PIT can be programmed also for ticking at periodic intervals, these ticks are as a beat time for all the kernel activities, as a metronome.
- **CPU Local Timer** (in the LAPIC) is similar to PIT but local to the processor.
- High Precision Event Timer (HPET), provides a number of programmable timers .

6.2

6. Time Management

# **Timekeeping Architecture**

DIAG

Advanced Operating Systems and Virtualization

# **Keeping time in the Linux kernel**

Linux must carry on several time-related activities, for instance, the kernel periodically:

- updates the time elapsed since the startup
- updates time and date
- determines, for every cpu, how long the current process has been running and preempts it if it exceeded its time slot
- update resource usage statistics
- checks whether the interval of time associated with each software timer has elapsed

The *Linux timekeeping architecture* is a set of data structures and functions related to the flow of time. There are some differences in time keeping between a multi and a single processor architecture:

- in a single-processor system all time-keeping activities are triggered by interrupts raised by a global timer
- in a multi-processor system general activities (e.g. software timers) are triggered by the global timer but CPU-specific activities (e.g. monitoring the execution time of the current process) are triggered by the local APIC timer

### **Kernel Timers**

The entire kernel timing subsystem wraps around timers. Timers can be divided into:

- classic timers available since the first version of the kernel, their resolution is usually 4ms and they are called *low-resolution* timers or *timer wheel* timers
- **high-resolution** timers, they have a precision of nanoseconds, useful especially for media-oriented applications

Independently of the resolution, the kernel distinguishes between:

- **Time-outs** that represent events that are bound to happen after some time, e.g. expect a packet within the next 10 seconds. In general, resolution is not important for this kind of timers;
- **Timers** that are used for implementing temporal sequences, e.g. a sound card driver could issue data in small and periodic time intervals. These timers usually requires high resolution;

### **Kernel Timers**



Hardware clock chips

Figure 15-2: Overview of the components that build up the timing subsystem.

Mauerer, Wolfgang. Professional Linux kernel architecture. John Wiley & Sons, 2010.

### **Clock Events & Ticks**

The Clock Events are the foundation of periodic events. High Resolution Timers are based on the Clock Events abstraction, whereas the low-resolution mechanism can come or from a low-resolution clock device or from the high resolution subsystem. Two important tasks for which low-resolution timers assume the responsibility are:

- handling the global jiffies counter
- perform per-process accounting

#### **Configuration Options**

In general a periodic tick is enable for the whole life of the kernel but this would not allow the system to go in sleep mode (e.g. laptop), for this reason the kernel allows to configure a dynamic tick which does not require a periodic signal (*tickless* system).

On right the possible configurations.

High-res	High-res
Dynamic ticks	Periodic ticks
Low-res	Low-res
Dynamic ticks	Periodic ticks

Figure 15-3: Possible timekeeping configurations that arise because of high- and low-resolution timers and dynamic/periodic ticks.

Mauerer, Wolfgang. Professional Linux kernel architecture. John Wiley & Sons, 2010.



6. Time Management 2. Timekeeping Architecture

# **Low-Resolution Timers**



Advanced Operating Systems and Virtualization



### **Jiffies**

The jiffies variable is a counter that stores the number of elapsed ticks since the system was started. It is increased by one when a timer interrupts occurs.

In the 8086 the jiffies variable is of 32bit in size so it wraps in about 50 days, this because a long long (64bit) could not allow an atomic increment. This is true even today, with kernel 5.11.

jiffies and jiffies\_64 matches in their less significant bits and therefore the must point to the same memory location or same register. For achieving this, the two variables are declared separately but the linker merges them, in 64bit machine they are the same but in 32bit jiffies is the upper (or lower depending on endianness) of jiffies\_64.

https://elixir.bootlin.com/linux/v5.11/source/include/linux/jiffies.h#L74

Timers allow a generic function to be activated at a later time, they can be dynamically created and destroyed and they can be assigned to the kernel itself or to a process. For this reason, we need an efficient way for managing timers in the kernel. Timers are associated with deferrable functions Linux **does not guarantee** that activation takes place at exact time

#### APIs

- void init\_timer(struct timer\_list \*timer);
- void setup\_timer(struct timer\_list \*timer, void (\*function)(unsigned long),
   unsigned long data);
- int mod\_timer(struct timer\_list \*timer, unsigned long expires);
- void del\_timer(struct timer\_list \*timer);
- int timer\_pending(const struct timer\_list \*timer);

Timers are prone to race conditions (e.g., if resources are released). They should be deleted before releasing the resources.

#### struct timer\_list

A Timer is represented by the struct timer\_list.

11	struct <b>t</b> .	<pre>imer_list {</pre>	
12		/*	
13		* All fields that c	change during normal runtime grouped to the
14		* same cacheline	
15		*/	
16		struct <b>hlist_node</b>	entry;
17		unsigned long	expires;
18		void	<pre>(*function)(struct timer_list *);</pre>
19		u32	flags;
20			
21	#ifdef C	ONFIG_LOCKDEP	
22		struct lockdep_map	lockdep_map;
23	#endif		
24	};		

https://elixir.bootlin.com/linux/v5.11/source/include/linux/timer.h#L11

In jiffies

#### **The Timer Wheel**

Earlier versions of the kernel used a single timer list sorted according to the expiration time. This was significantly unreliable and inefficient. Newer versions of the kernel introduced the so-called **Timer Wheel** a nested structure for efficiently retrieving timers. For understanding the principle of operation we will consider **simplified** example.



Mauerer, Wolfgang. Professional Linux kernel architecture. John Wiley & Sons, 2010.

#### The Timer Wheel

Suppose that each bucket has a counter that stores the number of an array position.



The tv1 counter initially points to 0, then it is incremented at each tick and every timer executed When 255 is reached the tv1 counter is reset to o and the timer of the first position of tv2 bucket replenish the bucket tv1, the counter of tv2 is increased by 1 When 63 is reached the tv2 counter is reset to 0 and the timer of the first position of tv3 bucket replenish the bucket tv2, the counter of tv3 is increased by 1

To determine which timers have expired, the kernel need not scan through an enormous list of timers but can limit itself to checking a single array position in the first group

6.2.2

6. Time Management

2. Timekeeping Architecture

# **Generic Time Subsystem**

DIAG

Advanced Operating Systems and Virtualization

### Components

There are three main abstractions for providing time in the kernel:

- **Clock Sources** (struct clocksource) that are the backbox of time management. Essentially each clock provides a monotonically increasing counter with read only access;
- **Clock event devices** (struct clock\_event\_device) add the possibility of equipping clocks with events that occurs in the future.
- **Tick Devices** (struct tick\_device) extend the clock event sources to provide a continuous stream of tick events that happen periodically. A dynamic tick may stop a tick device if requested.

There are then two kinds of clocks:

- a **global clock** that is responsible for providing the periodic tick for updating jiffies (old PIT)
- **one local clock** for each CPU for performing process accounting, profiling and high-resolution timers





Figure 15-7: Overview of the generic time subsystem.

# Clocksource

V5.11

The clocksource struct represents an abstraction of a source of time. The assigned function read allows to retrieve the current cycle value from the source.

33	/**	
34	* struct clocksource -	hardware abstraction for a free running counter
35	* Provides mostly	state-free accessors to the underlying hardware.
36	* This is the str	ructure used for system time.
37	*	
38	* @read:	Returns a cycle value, passes clocksource as argument
39	* @mask:	Bitmask for two's complement
40	*	subtraction of non 64 bit counters
41	* @mult:	Cycle to nanosecond multiplier
42	* @shift:	Cycle to nanosecond divisor (power of two)
43	* @max_idle_ns:	Maximum idle time permitted by the clocksource (nsecs)
44	* @maxadj:	Maximum adjustment value to mult (~11%)
00	struct electrony	r
09	Struct CLOCKSOURCE	
90	u64	(*read)(struct clocksource *cs);
91	u64	mask;
92	u32	mult;
93	u32	shift;
94	u64	max idle ns;
95	u32	maxadj;
96	#ifdef CONFIG ARCH	CLOCKSOURCE DATA

# **Jiffies as Clocksource**

41 42	<pre>* The Jiffies based clocksour * denominator clock source wh</pre>	ce is the lowest common ich should function on
43	* all systems. It has the sam	e coarse resolution as
44	<pre>* the timer interrupt frequen</pre>	cy HZ and it suffers
45	* inaccuracies caused by miss	ed or lost timer
46	<pre>* interrupts and the inabilit</pre>	/ for the timer
47	* interrupt hardware to accur	atly tick at the
48	* requested HZ value. It is a	lso not recommended
49	* for "tick-less" systems.	
50	*/	
51	static struct clocksource cloc	<pre>ksource_jiffies = {</pre>
52	.name = "jif	fies",
53	.rating = 1, /	* lowest valid rating*/
54	read = jiff	ies_read,
55	.mask = CLOC	(SOURCE_MASK(32),
56	.mult = TICK	<pre>NSEC &lt;&lt; JIFFIES_SHIFT, /* details above */</pre>
57	.shift = JIFF	IES_SHIFT,
58	.max_cycles = 10,	
59	};	

https://elixir.bootlin.com/linux/v5.11/source/kernel/time/jiffies.c#L40

### **Timer Interrupt Management**

The handling of timer interrupts are handled according to the top/bottom half paradigm (using Task Queues, which have now been removed from the Kernel).

The top half executes the following actions:

- registers the bottom half
- increments jiffies
- checks whether the CPU scheduler needs to be activated, and in the positive case flags need\_resched (more on this later)

# **High Resolution Timers**

High resolution timers are based on the ktime\_t type (nanosecond scalar representation) rather than jiffies.

High resolution timers heavily depends on the architecture wrt the low-resolution ones. They:

- are **arranged** in a red-black tree
- they are independent of periodic ticks, they are not based on jiffies

Low-resolution timers **are based** on high-resolution ones.

97	/**			
98	* struct hrtimer - the basic hrtimer structure			
99	* @node: timerqueue node, which also manages node.expires,			
00	* the absolute expiry time in the hrtimers internal			
01	* representation. The time is related to the clock on			
02	* which the timer is based. Is setup by adding			
03	* slack to the _softexpires value. For non range timers			
04	* identical to softexpires.			
05	* @ softexpires: the absolute earliest expiry time of the hrtimer.			
06	* The time which was given as expiry time when the timer			
07	* was armed.			
08	* @function: timer expiry callback function			
09	* @base: pointer to the timer base (per cpu and per clock)			
10	* @state: state information (See bit values above)			
11	* @is_rel: Set if the timer was armed relative			
12	* @is_soft: Set if hrtimer will be expired in soft interrupt context.			
13	* @is_hard: Set if hrtimer will be expired in hard interrupt context			
14	* even on RT.			
15	*			
16	* The hrtimer structure must be initialized by hrtimer_init()			
17	*/			
18	struct hrtimer {			
19	struct timerqueue_node node;			
20	ktime_tsoftexpires;			
21	enum hrtimer_restart (*function)(struct hrtimer *);			
22	struct hrtimer_clock_base *base;			
23	u8 state;			
24	us is_rel;			
25	us is_soft;			
26	us is_nard;			
//				

6.3

6. Time Management

# Watchdogs



Advanced Operating Systems and Virtualization

# Watchdogs

A watchdog is a component that monitors a system for "normal" behaviour and if it fails, it performs a system reset to hopefully recover normal operation.

This is a last resort to maintain system availability or to allow sysadmins to remotely log after a restart and check what happened. In Linux, this is implemented in two parts:

- a kernel-level **module** which is able to perform a hard reset
- a **user-space** background **daemon** that refreshes the timer

At kernel level, this is implemented using a Non-Maskable Interrupt (NMI). The userspace daemon will notify the kernel watchdog module via the /dev/watchdog special device file that user space is still alive.

```
while (1) {
    ioctl(fd, WDIOC_KEEPALIVE, 0);
    sleep(10);
}
```

# Advanced Operating Systems and Virtualization

[6] Time Management

LECTURER Gabriele **Proietti Mattia** 

BASED ON WORK BY <u>http://www.ce.uniroma2.it/~pellegrini/</u>



gpm.name · proiettimattia@diag.uniroma1.it

