

Time and Clocks

Some of this material is based on [?, Kopetz, Ch. 3].
NOTE: *These notes are not complete. More to come.*

Global Time and Events

Global time is the exact time reference of an omniscient external observer, represented by a real half-line $\text{TIME} = [0, \infty)$. The the origin 0 is fixed at a standard point in the remote past, before any moment of interest occurs. The unit of measure for TIME is the second (s), but we are often interested in fractional units, such as *milliseconds* ($ms = 10^{-3}s$), *microseconds* ($\mu s = 10^{-6}s$), *nanoseconds* ($ns = 10^{-9}s$).

Conceptually, an *event* e is an observable condition to which we can associate an instant in TIME . Events have no duration, or more precisely, their *observation* is instantaneous. Later on (but not in this chapter) the notion of an event is generalized to include *compound events* whose conditions include prior events. However, even compound events are associated with instantaneous observation. Reserve variables t, t_5, t_k , etc., to denote instants (or points) in TIME . The function mapping events to their observation instants is named T . For any event e , $T\langle e \rangle \in \text{TIME}$ is that point. It is assumed that all events are distinguishable in TIME , so that T is injective. In other words,

$$\text{if } T\langle e_1 \rangle = T\langle e_2 \rangle \text{ then } e_1 = e_2$$

Clocks

Clocks are devices that measure discrete time, usually in a periodic way. A clock C is represented by an sequence of events called *ticks*:

$$C = \{C^0, C^1, C^2, \dots\}$$

Ticks are strictly ordered in time, so that for all $i < j$,

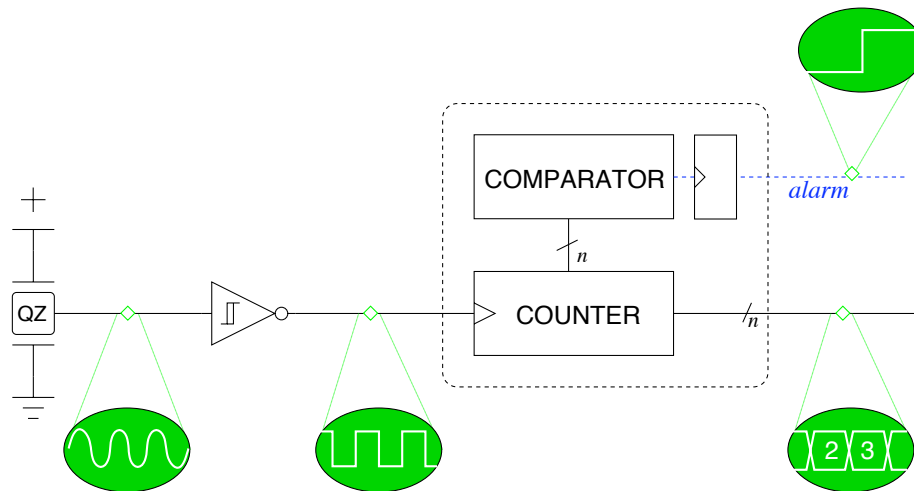
$$T\langle C^i \rangle < T\langle C^j \rangle$$

Just as the variable t is overloaded to denote both real-time instants and event observation times ($T\langle e \rangle$), let us define $C(e) \in \mathbb{N}$ to be the *index of the tick* at which e is observed by clock C . It should be clear that for any event e , $T\langle e \rangle \leq T\langle C^{C(e)} \rangle$; that is, no event can be observed in discrete time before it actually happens.

However, unlike the case in TIME , distinct events can be co-incident in “clock time;” we can have $C(e_1) = C(e_2)$ even if $e_1 \neq e_2$.

Physical Clocks

Physical clocks usually contain an *oscillator* of some kind¹ that generates a continuous sinusoidal voltage wave-form. The wave-form passes through a voltage-thresholding device that generates a digital wave form whose transitions (from low to high voltage, say) are the clock's ticks. This signal forms the *hardware clocking* hardware signal that synchronizes all digital components in its "domain." In effect, each device becomes an observer of the hardware-clock's ticks.



At the programming level, one is not so concerned with the clock-ticks themselves, but instead with how many ticks have transpired. Programs deal with *timers* which are basically digital counters that tally the ticks of the hardware clock. This counter can be programmed to generate an interrupt signal when the tally exceeds a pre-determined value. Every computer has one or more timers of this kind.

Of course, no physical clock keeps perfect time. True oscillators have specified frequency ranges and may differ from their frequency ratings within that range. True frequencies are also subject to fluctuations caused by electro-thermal conditions, and even the age. An oscillator's true frequency may *drift* (or deviate) by 0.00001% to 0.1% from its specified frequency, depending on its quality.

The reciprocal of a clock's frequency (cycles per second) is its *period* or *cycle time*. Technically, these terms refer to the continuous behavior of the oscillator. In discrete-time models, the comparative term *granularity* is also used to characterize the duration between ticks, and a specific cycle is called a *granule*. All these terms are interchangeable when the context is understood. A clock with a higher frequency, hence shorter cycle time, is said to be *fine-grained* relative to a lower-frequency, or *course-grained* clock.

¹The most common kind of oscillator is a quartz crystal, which vibrates in the presence of voltage, generating a fixed-period sine-wave.

The Reference Clock

The following kinds of questions are of interest:

- How accurate is a process’s estimate of global real time?
- How far apart in real time must two events be in order for a process to distinguish them?
- Can two processes observing the same event agree on when it occurred?
- If not, is there a best-case bound on their time estimates?
- Can two or more processes agree on what time it is?


To answer such questions, we need some kind of “ruler” providing a common reference to discrete time and to serve as a mathematical interface between digital time-keeping and global real time. The *reference clock*, called R , is a unique, abstract clock. R is observable externally but not internally, and it keeps perfect calendar time, relative to an International Time Standard. Each granule of R is strictly equal to all others. That is, R ’s frequency is *exactly* $f = \frac{1}{p}$ for some fixed p so that for all $i \in \mathbb{N}$,

$$T\langle R^{i+1} \rangle - T\langle R^i \rangle = p$$

R ’s is sufficiently fine-grained to make all digitization errors negligible, or equivalently, to make all distinct events separable. In current technologies, a frequency of femtoseconds (10^{15} cycles per second) is sufficient because digital *signal events*—transitions between valid voltage levels—require several femtoseconds.

Microticks and Macroticks

The reference clock’s miniscule granularity also allows us to mathematically measure and compare the presumably coarser-grained behavior of physical clocks. Toward this end, a coarser grained clock is superimposed on R , whose ticks represent the ideal behavior of a physical clock. Rather than postulating still another clock, we identify certain ticks of R as *macroticks*. The fine-grained ticks are called *microticks*. Unless the context requires it, the term *tick* means macrotick. So the picture of global time becomes:



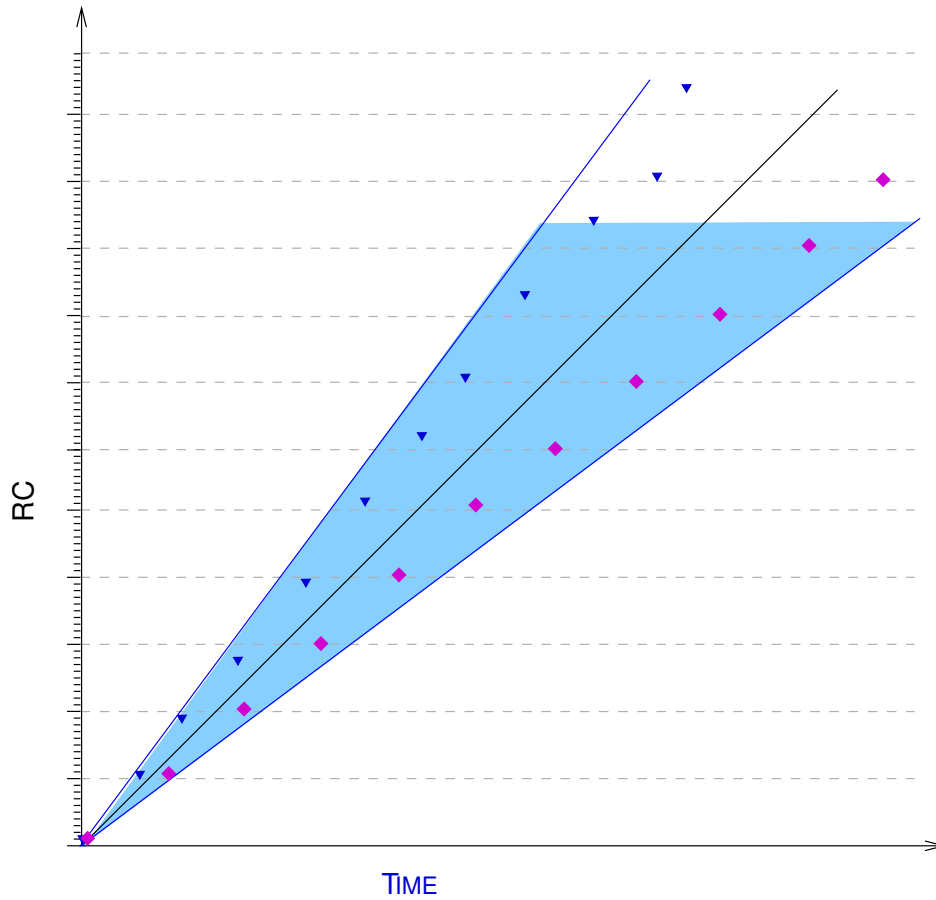


Figure 1: Clock drift

Notions of microticks and macroticks carries to physical clocks, as we shall see in Section , and beyond into levels of software abstraction. At software levels, it is not useful to reason at the physical clock's level of granularity. A single CPU instruction, for example, takes several hardware-clock ticks to interpret. So eventually, a notion of physical macroticks is needed. Similarly, at software levels, a macrotick may correspond to thousands or millions of physical clock ticks, depending on the granularity of software cycles.

Clock Drift

Figure 1 graphs ticks of two physical clocks against real time (horizontal axis) and reference-clock time (vertical axis). Both clocks deviate from the ideal

clock, whose ticks would lie along the diagonal line with slope 1. One is too fast and also fluctuates from tick to tick, when compared to the reference clock. The other is too slow, but appears to be more regular with respect to R . Clock *drift* is its deviation from a specified frequency as measured by R in microticks:

DEFINITION. *The incremental drift of a physical clock at tick i is*

$$\text{drift}(i, C) = \frac{R \langle C^{i+1} \rangle - R \langle C^i \rangle}{n}$$

where n is the nominal number of microticks successive ticks.

C 's incremental drift rate is its drift, normalized to an ideal clock whose incremental drift is 0.

$$\rho(i, C) = |\text{drift}(i, C) - 1|$$

The (accumulated) drift rate, $\rho(C)$, defines a bound on C 's drift over time.

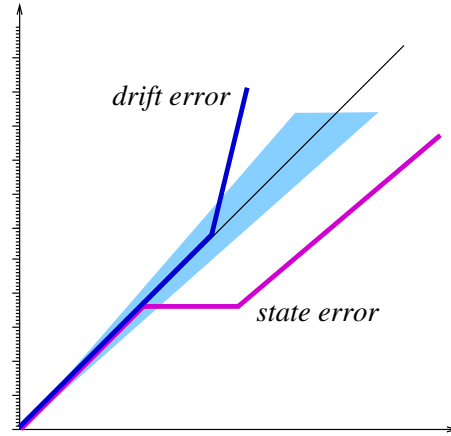
$$\rho(C) = \max\{\rho(i, C)\}_{i \in \mathbb{N}}$$

A physical clock's drift rate within specified environmental conditions is documented in its data sheet. As noted earlier, typical drift rates range from 10^{-2} to 10^{-7} , depending on the quality of the oscillator. Since no physical clock has drift rate 0, any *free running* clock will, over time, drift arbitrarily far from its nominal frequency; that is, that is, no bound can be placed on cumulative drift. Hence, physical clocks must occasionally be *adjusted* to correct for their drift.

Clock Failure Modes

Even an adjusted clock may still *fail*, however. Two distinct *modes* are commonly of concern. A *drift failure* occurs when a clock exceeds its specified drift rate. Drift failure is typically caused by environmental conditions exceeding the specified constraints, or when rate-adjustable clocks (discussed later) are mis-configured. A *state error* occurs when the clock's counter is over-written with an invalid value. Both of these scenarios, illustrated below, are possible effects

of software errors in clock adjustment.



Multiply Clocked Systems

Embedded systems usually contain multiple, separately clocked components. Coordinating component behavior requires that they agree on time at some level of granularity. More specifically, if two components observe the same event, they should agree as to when that event occurred. Similarly, if two components observe two events, they should agree on their temporal order.

Ordinarily², no physical clock can observe the ticks of any other clock, but the reference clock can observe their relative drift, or *offset*.

DEFINITION. *The offset between two physical clocks, C and D , at tick i , is the number of microticks between their corresponding i^{th} ticks:*

$$\text{offset}(i, C, D) = |R\langle C^i \rangle - R\langle D^i \rangle|$$

Given an ensemble (or set) of clocks, $\mathcal{C} = \{C_1, \dots, C_n\}$, the precision of \mathcal{C} at tick i is

$$\Pi(i, \mathcal{C}) = \max\{\text{offset}(i, C_j, C_k) \mid 1 \leq j, k \leq n\}$$

The precision Π of \mathcal{C} over an interval of interest, $[a, b]$, is

$$\Pi = \max \{ \Pi(i, \mathcal{C}) \mid a \leq i \leq b \}$$

Since precision is simply a measure related to relative drift, it grows without bound over time unless the clocks are internally synchronized occasionally. Clock synchronization is the subject of Section ??, later in this chapter.

²Important exceptions to this generality are discussed later in this book, but not in this chapter.

In most applications, the primary concern is precision; that is, how closely the physical clocks agree on the number of ticks that have occurred since time zero. Sometimes, it is also important to know how far the ensemble has drifted in its estimate of global real time.

DEFINITION. A clock's Accuracy at tick i is

$$\text{accuracy}(i, C) = \text{offset}(i, C, R)$$

The accuracy of an ensemble $\mathcal{C} = \{C_1, \dots, C_n\}$ is the maximum accuracy its individual clocks over an interval of interest.

When accuracy is a concern, all physical clocks must also be *externally synchronized* with the reference clock from time to time.

If an ensemble \mathcal{C} is externally synchronized to an accuracy of A then it is also internally synchronized to a precision of $\Pi = 2A$. An internally synchronized ensemble is not necessarily accurate, however.

Time Estimation

We have laid enough groundwork, now, to address the questions posed at the beginning of Section . In particular, let us consider events that are visible³ to the physical clocks and whether they can agree on the global time of such a common observation.

It should be clear that agreement cannot be closer than the precision of the ensemble, so let us impose a coarser granularity of that magnitude. Given an ensemble $\mathcal{C} = \{C_1, \dots, C_n\}$ with precision Π , choose a *resolution* $m \in \mathbb{N}$ such that for each $C_k \in \mathcal{C}$ and for all $i \in \mathbb{N}$,

$$R \langle C_k^{i+m} \rangle - R \langle C_k^i \rangle \leq \Pi$$

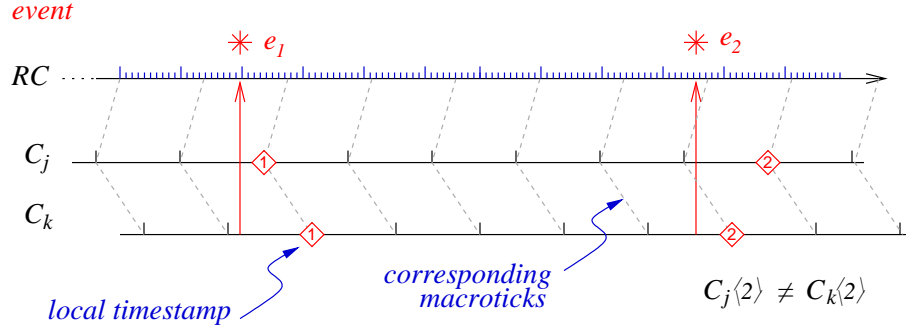
As with the reference clock, call the fine-grained ticks of a physical clock its *microticks* and call every m^{th} microtick a *macrotick*, or simply *tick* if the context allows. A *timestamp* is a local clock's estimate of global time, is measured in (macro)ticks, still denoted by $C \langle e \rangle$. This coarser *global granularity* of macroticks, we can assure that the local timestamps of two clocks observing the same event differ by at most one tick; for any event e

$$| C_j \langle e \rangle - C_k \langle e \rangle | \leq 1$$

This is as close as we can get to clock agreement with respect to a single event. Because physical clocks cannot be perfectly synchronized in dense TIME, for an event to occur between the i^{th} tick of one clock and the $i + 1^{\text{st}}$ tick of another, resulting in different timestamps. In the illustration below, two clocks agree on

³Of course, the clocks themselves don't observe or act on events; they are simply event generators, used by the observers to mark the time of their observations.


event e_1 and disagree on e_2 . It is of no help to make the granules larger. This is a fundamental disparity.



Event Ordering

It should be apparent that if two events, e_1 and e_2 , occur close enough in time, then a single clock cannot distinguish their order. Similarly, in the illustration below,

- (a) $C_j \langle e_1 \rangle = C_j \langle e_2 \rangle$, but $C_k \langle e_1 \rangle \neq C_k \langle e_2 \rangle$. Clock C_j sees events e_1 and e_2 occurring simultaneously but C_k sees them at different global times.
- (b) $\mathcal{C} \langle e_3 \rangle = \mathcal{C} \langle e_4 \rangle$, but $R \langle e_4 \rangle - R \langle e_3 \rangle > \Pi$. The *ensemble* $\{C_j, C_k\}$ sees events e_3 and e_4 simultaneous in global time, even though the actual interval between them exceeds the precision of the system.
- (c) $\mathcal{C} \langle e_6 \rangle < \mathcal{C} \langle e_5 \rangle$, but $R \langle e_5 \rangle < R \langle e_6 \rangle$. The *ensemble* $\{C_j, C_k\}$ sees events e_5 and e_6 in a different temporal order than they actually occur.



Scenarios (b) and (c), above, demonstrate that of the temporal order of two events can be correctly deduced provided the actual time-interval between them is at least two granules of local time. As a consequence of synchronization and digitization errors it is not (always) possible for a multi-clocked *system* to reconstruct event order if the timestamps differ by one tick, even though this is possible for a single local clock.

Measuring Time Intervals

As we have just seen, the synchronization and digitization errors associated with a single event sum to at most $2g$. If we wish to measure the time interval between starting and ending events, e_s and e_e , define *durations*

$$d_{act} = R\langle e_e \rangle - R\langle e_s \rangle, \text{ the actual interval from } e_s \text{ to } e_e$$

$$d_{obs} = \mathcal{C}\langle e_e \rangle - \mathcal{C}\langle e_s \rangle, \text{ the observed interval from } e_s \text{ to } e_e$$

Then if $g > \Pi$ is the global granularity of the system.

$$(d_{obs} - 2g) < d_{act} < (d_{obs} + 2g)$$

At best the actual duration is within a $4g$ time window of what can be observed.

π/Δ Precedence

So far, we have looked at external event discrimination from within a distributed system. Let us now turn this around and consider the external view of events generated from within such a system. Suppose three separately clocked nodes are each to *generate* events at regular intervals, at ticks 1, 5, and 9.

