

Real Time and Clocks

Ken Birman

Cornell University. CS5410 Fall 2008.



Real time and clocks

- Lamport showed that if we care about event ordering, the best option is to use logical clocks
- But suppose we care about real time?
 - How well can clocks be synchronized?
 - To what extent can the operating system help us build applications that are sensitive to time?



Introducing “wall clock time”

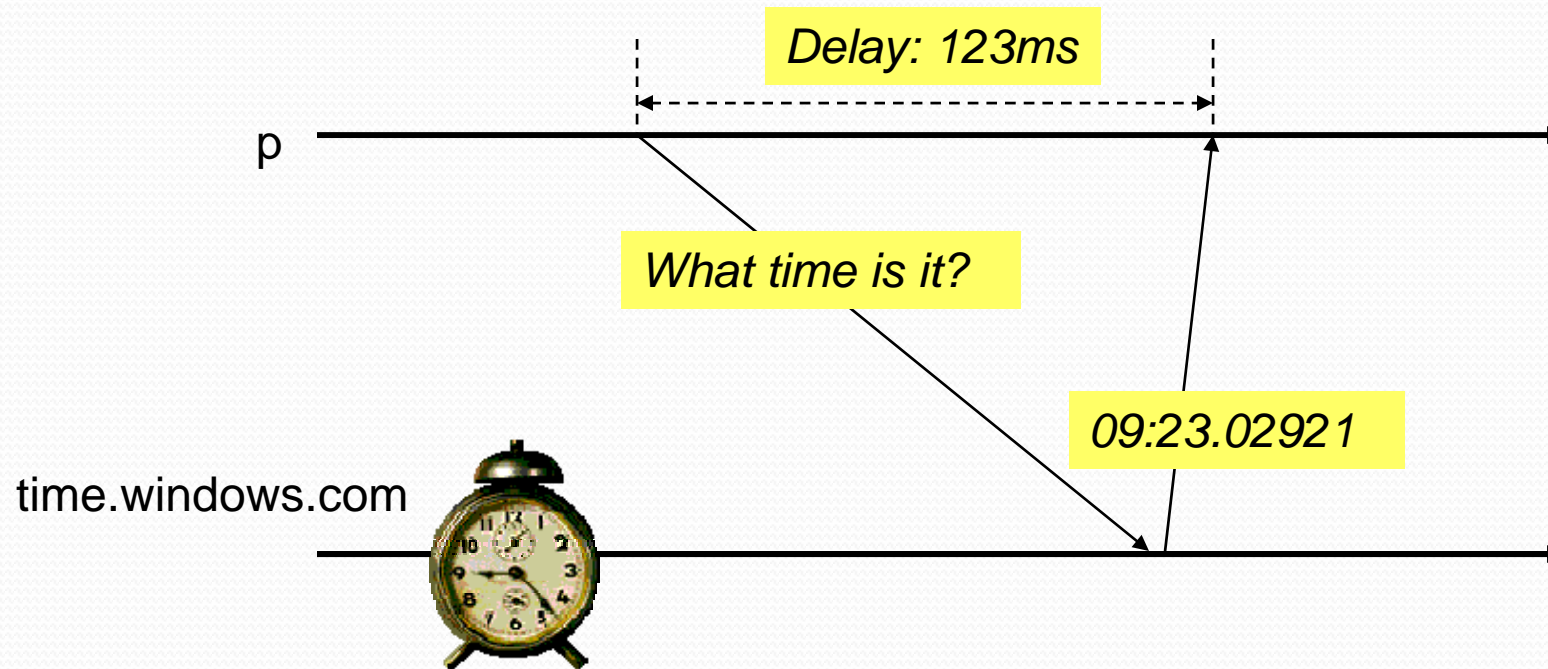
- There are several options
 - “Extend” a logical clock or vector clock with the clock time and use it to break ties
 - Makes meaningful statements like “B and D were concurrent, although B occurred first”
 - But unless clocks are closely synchronized such statements could be erroneous!
 - We use a clock synchronization algorithm to reconcile differences between clocks on various computers in the network



Synchronizing clocks

- Without help, clocks will often differ by many milliseconds
 - Problem is that when a machine downloads time from a network clock it can't be sure what the delay was
 - This is because the “uplink” and “downlink” delays are often very different in a network
- Outright failures of clocks are rare...

Synchronizing clocks




- Suppose *p* synchronizes with *time.windows.com* and notes that 123 ms elapsed while the protocol was running... what time is it now?



Synchronizing clocks

- Options?
 - P could guess that the delay was evenly split, but this is rarely the case in WAN settings (downlink speeds are higher)
 - P could ignore the delay
 - P could factor in only “known” delay
 - For example, suppose the link takes at least 25ms in each direction...

Synchronizing clocks



- Suppose *p* synchronizes with *time.windows.com* and notes that 123 ms elapsed while the protocol was running... what time is it now?



Synchronizing clocks


- In general can't do better than uncertainty in the link delay from the time source down to p
 - Take the measured delay
 - Subtract the “certain” component
 - We are left with the uncertainty
- Actual time can't get more accurate than this uncertainty!



What about GPS?

- GPS has a network of satellites that send out the time, with microsecond precision
- Each radio receiver captures several signals and compares the time of arrival
- This allows them to triangulate to determine position

GPS Triangulation





Issues in GPS triangulation

- Depends on very accurate model of satellite position
 - In practice, variations in gravity cause satellite to move while in orbit
- Assumes signal was received “directly”
 - Urban “canyons” with reflection an issue
- DOD encrypts low-order bits



GPS as a time source

- Need to estimate time for signals to transit through the atmosphere
 - This isn't hard because the orbit of the satellites is well known
 - Must correct for issues such as those just mentioned
- Accurate to $\pm 25\text{ms}$ without corrections
- Can achieve $\pm 1\text{ }\mu\text{s}$ accuracy with correction algorithm, if enough satellites are visible



Consequences?

- With a cheap GPS receiver, 25ms accuracy, which is large compared to time for exchanging messages
 - 10,000 msgs/second on modern platforms
 - ... hence .1ms “data rates”
 - Moreover, clocks on cheap machines have 10ms accuracy
- But with expensive GPS, we could timestamp as many as 100,000 msgs/second



Accuracy and Precision



- Accuracy is a measure of how close a clock is to “true” time
- Precision is a measure of how close a set of clocks are to one-another
 - Both are often expressed in terms of a window and a drift rate

Thought question

- We are building an anti-missile system

Radar tells the interceptor where it should be and what
time to get there

How do we want the radar and interceptor to be as accurate
as possible, or as precise as possible?





Thought question

- We want them to agree on the time but it isn't important whether they are accurate with respect to “true” time
 - “Precision” matters more than “accuracy”
 - Although for this, a GPS time source would be the way to go
 - Might achieve higher precision than we can with an “internal” synchronization protocol!



Real systems?

- Typically, some “master clock” owner periodically broadcasts the time
- Processes then update their clocks
 - But they can drift between updates
 - Hence we generally treat time as having fairly low accuracy
 - Often precision will be poor compared to message round-trip times



Clock synchronization

- To optimize for precision we can
 - Set all clocks from a GPS source or some other time “broadcast” source
 - Limited by uncertainty in downlink times
 - Or run a protocol between the machines
 - Many have been reported in the literature
 - Precision limited by uncertainty in message delays
 - Some can even overcome arbitrary failures in a subset of the machines!



Adjusting clocks: Not easy!

- Suppose the current time is 10:00.00pm
 - Now we discover we're wrong
 - It's actually 9:59.57pm!
- Options:
 - Set the clock back by 3 seconds...
 - But what will this do to timers?
 - Implies a need for a “global time warp”
 - Introduce an artificial time drift
 - E.g. make clock run slowly for a little while



Real systems

- Many adjust time “abruptly”
 - Time could seem to freeze for a while, until the clock is accurate (e.g. if it was fast)
 - Or might jump backwards or forwards with no warning to applications
- This causes many real systems to use relative time: “now + XYZ”
 - But measuring relative time is hard



Some advantages of real time

- Instant common knowledge
 - “At noon, switch from warmup mode to operational mode”
 - No messages are needed
 - Action can be more accurate than would be possible (due to speed of light) with message agreement protocols!



Some advantages of real time

- The outside world cares about time
 - Aircraft attitude control is a “real time” process
 - People and cars and planes move at speeds that are measured in time
 - Physical processes often involve coordinated actions in time



Disadvantages of real time

- On Monday, we saw that causal time is a better way to understand event relationships in actual systems
 - Real time can be deceptive
 - Causality can be tracked... and is closer to what really mattered!
- For example, a causal snapshot is “safe” but an instantaneous one might be confusing



Internal uses of time

- Most systems use time for expiration
 - Security credentials are only valid for a limited period, then keys are updated
 - IP addresses are “leased” and must be refreshed before they time out
 - DNS entries have a TTL value
 - Many file systems use time to figure out whether one file is fresher than another



The “endless rebuild problem”

- Suppose you run Make on a system that has a clock running slow
 - File xyz is “older” than xyz.cs, so we recompile xyz...
 - ... creating a new file, which we timestamp
 - ... and store
- The new one may STILL be “older” than xyz.cs!



Implications?

- In a robust distributed system, we may need trustworthy sources of time!
 - Time services that can't be corrupted and won't run slow or fast
 - Synchronization that really works
 - Algorithms that won't malfunction if clocks are off by some limited amount



Fault-tolerant clock sync


- Assume that we have 5 machines with GPS units
- Each senses the time independently
- Challenge: how to achieve optimal precision and accuracy?



Srikanth and Toueg

- You can't achieve both at once
 - To achieve the best precision you lose some accuracy, and vice versa
- Problem is ultimately similar to Byzantine Agreement
 - We looked at this once, assuming signatures
 - Similar approach can be used for clocks

Combining “sensor” inputs



- “Shout at 10:00.00”



Combining “sensor” inputs

- Basic approach
 - Assume that no more than k out of n fail
 - Depending on assumptions, k is usually bounded to be less than $n/3$
 - Discard outliers
 - Take mean of resulting values
- Attacking such a clock?
 - Try and be “as far away as possible” without getting discarded



How do real clocks fail?

- Bits can stick
 - This gives clocks that “jump around”
- The whole clock can get stuck, perhaps erratically
- Clock can miscount and hence drift (backwards) rapidly



Using real-time

- Consider using a real-time operating system, clock synchronization algorithm, and to design protocols that exploit time
- Example: MARS system uses pairs of redundant processors to perform actions fault-tolerantly and meet deadlines. Has been applied in process control systems. (Another example: Delta-4)



Using time with sensors


- Many distributed systems monitor something in the outside world
 - They use “sensors” to capture data such as temperature, video images, etc. Often data comes with build-in precision limits
 - Then label these with time
- We’ve seen that time comes with imprecision too
- How does this impact applications that “sense” things?



Time with sensors


- Suppose that an application tracks temperature in Ithaca
 - At 10:00am, 52 degree F
 - At noon, 68 degrees F
 - At 2:00pm, 74 degrees F
 - At 6:00 pm 58 degrees F
- And temperature is +/- 2 degrees

Temperatures



Do we really know the value?

- The 12pm value was really within a “bounding box”
 - The value was between, say, 63 and 67 with a “best estimate” of 65
 - But the *time* was also in a range of possible times
 - Perhaps, between 11:59 and 12:01
- So we should think of the sensor value as a box



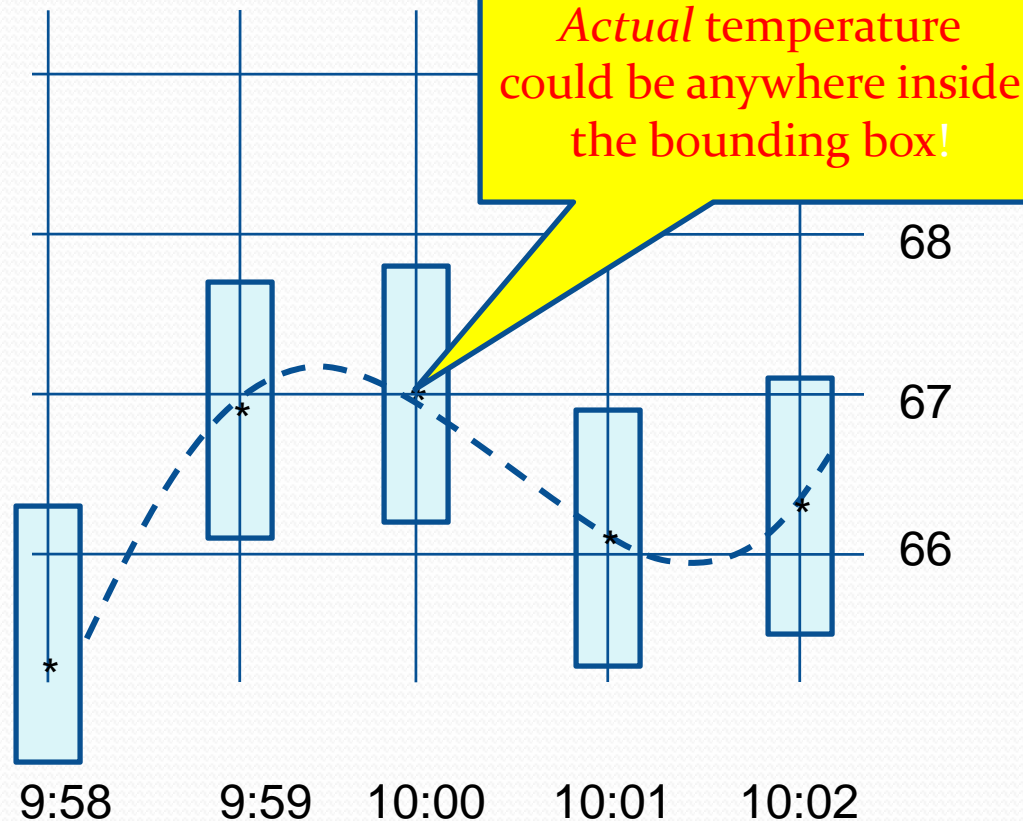


Does this matter?

- Suppose that we are supposed to *only activate the assembly line once all the furnaces have reached operating temperature*
- Or *vent the reactor vessel if the pressure goes over 100 lbs per square inch*
- How would we translate these rules to work with sensors that return values in “boxes”?

“Maybe” versus “Definitely”

- Suppose a sensor returns 67 ± 1 at $10:00 \pm 10$ secs



- Was it *definitely* 68 degrees? Or just “maybe”?




Wood and Marzullo

- Looked at issues of clock and sensor synchronization
- Developed fault-tolerance mechanisms for estimating data values and synchronizing clocks
- Showed how to deal with imprecision
 - You needed to tell them which behavior you wanted
 - Then they interpreted the question relative to the “bounding box” for the sensor

Overcoming errors

- If we have n readings and at most k are faulty, intersect boxes (excluding all possible





Back to our questions

- *Only activate the assembly line once all the furnaces have [definitely] reached operating temperature*
 - We want to know that the temperature is *definitely* high enough. *Entire bounding box* must be above the threshold temperature to be safe, since any point in the box is a “possibility” for the current temperature
- *Vent the reactor vessel if the pressure [may be] over 100 lbs per square inch*
 - Trigger vent if *any portion of the box* is over threshold, because (perhaps) the vessel has reached that pressure.



Other issues to consider

- Source of imprecision is often the operating system
 - Scheduling delays
 - Paging
 - Contention for resources (locking)
- To overcome these problems it can be helpful to use a *real time operating system* in addition to using clock synchronization or sensor synchronization protocols
- By reducing uncertainty these “shrink the box”



Summary

- On Monday we saw that events in a system are best understood in terms of the logical progression of time
- Now we've looked at real (clock) time, which is one form of sensor, and also other kinds of sensor inputs
- Imprecise measurements force us to think in terms of bounding boxes with values in the box
 - We can use this to overcome errors
 - And we can also interpret queries over sensors and time in ways that explicitly cope with imprecision