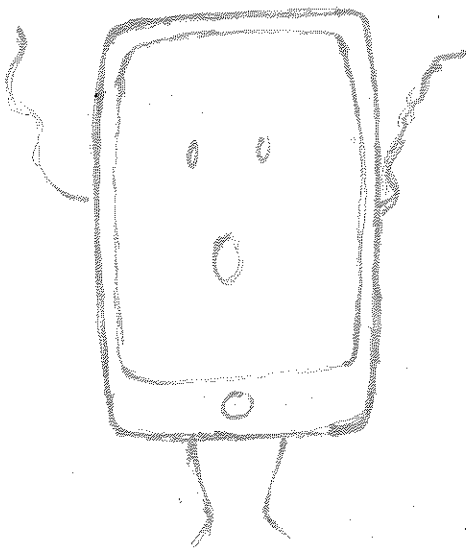

ERIK HAIGHT, CMPS 128 SPRING 2019

insync With NTP

BY:
ERIK
HAIGHT

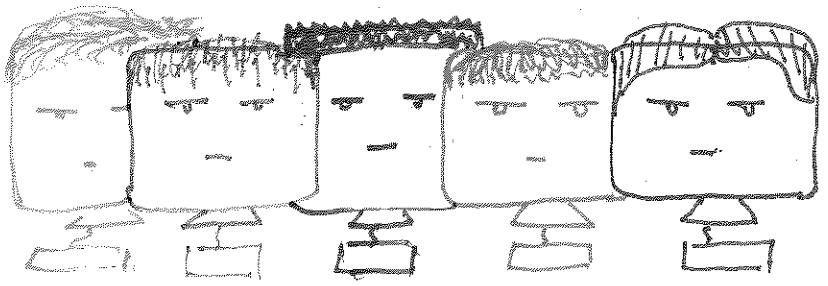

AN INTRODUCTION

HELLO MY NAME IS ERIK
AND I'M A COMPUTER
SCIENCE STUDENT AT UCSC!

I'VE BEEN TAKING A COURSE ON
DISTRIBUTED SYSTEMS THIS QUARTER.

A DISTRIBUTED SYSTEM
CAN BE THOUGHT OF AS
MANY COMPUTERS PRETENDING
TO BE A SINGLE COMPUTER.

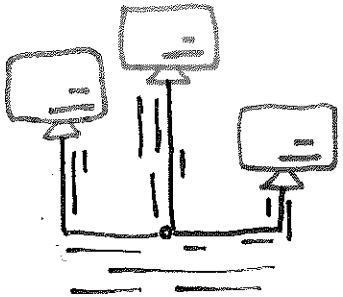

DISTRIBUTED SYSTEMS:

- ARE SO IMPORTANT!
- ARE FUN TO LEARN ABOUT!
- HAVE ENDLESS APPLICATIONS!
- STRUGGLE WITH ANSWERING

WHAT TIME IS IT!?

SO WHAT CLOCKS CAN YOU USE?
(FOR ORDERING EVENTS)

I MENTIONED IT BRIEFLY, BUT LOGICAL CLOCKS AND I ENCOURAGE ANYONE WHO HASN'T HEARD OF THEM TO CHECK OUT:

- LAMPORT CLOCKS
- VECTOR CLOCKS

OTHER TOPICS MENTIONED TO LOOK AT:

- DECENTRALIZED SYSTEMS
- FAULT-TOLERANCE
- LEAP SECONDS

RESOURCES I USED:

- http://support.ntp.org
- NTP DOCUMENTATION ON https://www.eecis.udel.edu
- NOTES FROM CMPS 128 SPRING 2019, LINDSEY KUPER
- https://en.wikipedia.org/wiki/NETWORK_TIME_PROTOCOL
- https://developers.google.com/time/smear

AND IF YOU DON'T BUY IT,

... LEAP SECONDS!

SOME SYSTEMS DEPEND ON PHYSICAL CLOCKS, AND NEGATIVE ELLAPSED TIMES CAN BE CATASTROPHIC... WHICH LEAP SECONDS CAN CAUSE.

THE STOCK MARKET SHUTS DOWN FOR THE HOUR SURROUNDING A LEAP SECOND.

IT'S SERIOUS BUSINESS.

MOST SYSTEMS REPEAT A LEAP SECOND LIKE: 1 → 2 → 3 → 3 → 4

— DANGEROUS —

SOME SYSTEMS LIKE GOOGLE'S "SMEER" THE LEAP SECOND:

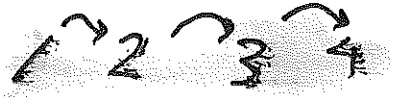

SECONDS ARE "LONGER" BUT AVOID LEAP SECOND FAULTS

ALL IN ALL: TIME IS TRICKY!

NOW I HOPE YOU LEARNED A BIT ON DISTRIBUTED SYSTEMS, NTP, AND WHY PHYSICAL CLOCKS AREN'T ALWAYS SAFE FOR ORDERINGS.

IT TURNS OUT THAT DISTRIBUTED SYSTEMS CAN'T REALLY RELY ON PHYSICAL CLOCKS AND INSTEAD USE LOGICAL CLOCKS TO REASON ABOUT ORDERINGS OF EVENTS.

BUT WE CAN'T JUST DITCH PHYSICAL CLOCKS ... AS MUCH AS WE MIGHT SOMETIMES WANT TO.

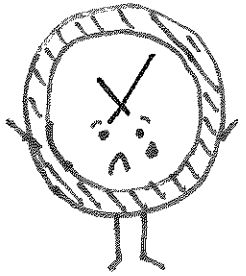

PEOPLE NEED TO KNOW WHAT TIME IT IS, WHEN AN EMAIL WAS SENT, OR ANYTHING ELSE RELATED TO DATA WITH A TIMESTAMP.

SO... WHAT CAN WE DO?

NETWORK TIME PROTOCOL
(OR NTP)

AND THAT'S WHAT THIS ZINE IS ABOUT!

NTP SHOWCASES SOME CORE CONCEPTS OF DISTRIBUTED SYSTEMS, AND I THINK NTP IS SWEET.

I HOPE BY THE END OF THIS, YOU DO TOO!

Some Definitions

BEFORE GETTING INTO DETAILS OF NTP, WE SHOULD FORMALLY DEFINE SOME TERMS.

TIME CONSUMER

CLIENT COMPUTER THAT WANTS TO SYNCHRONIZE IT'S LOCAL SYSTEM TIME WITH A TIME PROVIDER.

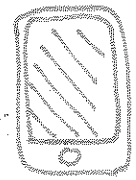

TIME PROVIDER

SERVER COMPUTER THAT SENDS IT'S LOCAL SYSTEM TIME TO TIME CONSUMERS.

STEPPING

LARGE TIME ADJUSTMENTS OVER A SHORT PERIOD OF TIME.

SLEWING

SMALL TIME ADJUSTMENTS OVER A LONG PERIOD OF TIME.

WHEN WE TALK ABOUT NTP, WE WILL BE MOSTLY TALKING ABOUT NTP DAEMON (NTPD), AS NTPD IS CONTINUOUSLY SYNCING.

Clocks And Time

SO WHERE DOES NTP STAND WITH OUR INITIAL PROBLEM OF TIME?

UNFORTUNATELY, IT DOESN'T REMOVE THE PROBLEM, BUT JUST REDUCES ODDS OF ERRORS OCCURING.

INACCURACY, IS STILL INACCURACY.

ON A **REALLY UNLUCKY** POLL FOR TIME, A DIFFERENCE BETWEEN A LATER POLL & AN EARLIER ONE...

CAN CAUSE **NEGATIVE ELLAPSED TIMES**

AND YOU JUST **CAN'T** ASSUME AN ORDERING OF EVENTS BASED ON A PHYSICAL TIMESTAMP.

THE IDEA IS THAT THE LOWER THE NUMBER OF THE STRATUM (CLOSER TO ATOMIC CLOCKS), THE MORE ACCURATE THE TIME IS.

STRATUM 1 DEVICES DIRECTLY CONNECT TO STRATUM 0 DEVICES WITH ~10ms ACCURACY. (ITS CRAZY)

EACH ADDITIONAL STRATUM IS ANOTHER $\frac{1}{2}$ - 100 MS ACCURACY DECREASE. EACH!

GREAT THINGS ABOUT THIS STRUCTURE:

- WORRIED ABOUT THE GAME OF TELEPHONE? ASK STRATUM 2 DEVICES, IT'S NOT NEARLY AS BAD ☺
- SOME DEVICES OVERLOADED? SERVER POOLS DISTRIBUTE WORKLOAD ☺
- SERVER CRASH? STRATUM 16 RESERVED FOR THIS!
  * THIS MAKES NTP SYSTEMS FAULT-TOLERANT *
    → THEY TOLERATE SOME SYSTEM FAILURES

THE PROTOCOL (NTPD)

THE NETWORK TIME PROTOCOL DAEMON WORKS AS FOLLOWS:

TIME CONSUMER SENDS A TIME SYNCHRONIZATION REQUEST TO PRODUCER.

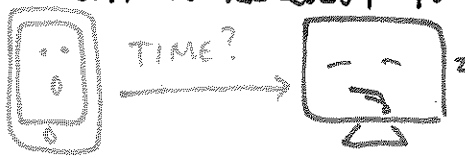

TIME PROVIDER SENDS BACK IT'S LOCAL SYSTEM TIME.

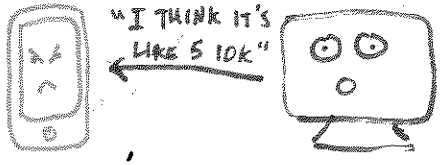

THE CONSUMER'S LOCAL TIME IS ADJUSTED:

IF IT'S PRETTY CLOSE (DIFFERENCE < 128ms)

SLEW

IF IT'S FAR OFF (128ms < DIFFERENCE < 17 MINS)

STEP

OTHERWISE, WE ARE REALLY FAR OFF (DIFF > 17MINS)

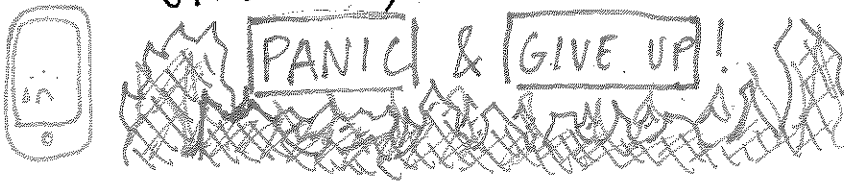

WHEN THE TIME DIFFERENCE BETWEEN THE CONSUMER AND PRODUCER IS OVER 17 MINUTES, IT'S CALLED:

INSANE TIME

(INSANE TIME)

AND WE GIVE UP ON EVEN ATTEMPTING TO SYNCHRONIZE TIME.

SO BEFORE WE USE NTPD, WE NEED TO GET SOMEWHAT CLOSE, BUT ~~HOW~~?

A MESSENGER BIRD?

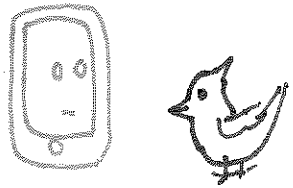

WE USE NTPDATE! SAME IDEA AS NTPD, BUT JUST TO GET THE "BALLPARK" NUMBER. IT'S A ONE-TIME SYNC, THEN NTPD NARROWS IN INDEFINITELY.

STRATA

TO TALK ABOUT WHAT TRUE TIME IS, WE NEED TO KNOW ABOUT STRATUM, OR PLURALIZED, STRATA.

IT LOOKS LIKE THIS:

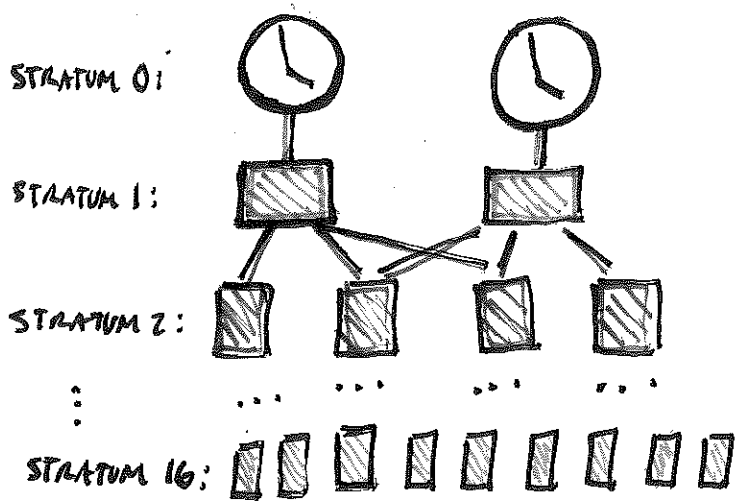

IT'S REALLY JUST A FANCY WORD FOR HIERARCHY

TRUE TIME COMES FROM STRATUM 0 DEVICES; GPS, ATOMIC CLOCKS, EXPENSIVE STUFF!

AND STRATUM 0 / 1 DEVICES AREN'T PUBLICALLY AVAILABLE, (SECURITY REASONS) & ALSO TRAFFIC!

THIS COULD EASILY TURN INTO A GAME OF TELEPHONE, & IT IS ONE!

FOR ANYONE WHO HASN'T PLAYED TELEPHONE, IT'S A GAME WHERE A PERSON STARTS WITH A MESSAGE, PASSES IT ON TO SOMEONE, THEY PASS IT ON & SO ON. THE IDEA IS THAT THE FURTHER THE MESSAGE, THE LESS RELIABLE IT'S CONTENT GETS.

THEN WITH THE BUS STOP EXAMPLE, SOMEONE COULD TELL ME:

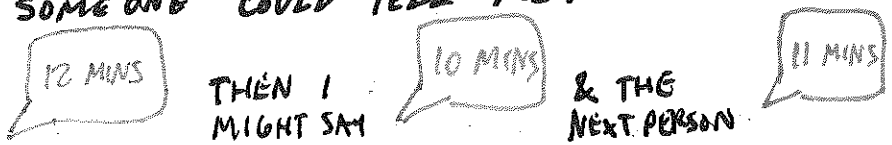

AND SO ON!

NTP IS A GAME OF TELEPHONE, WHERE THE "ORIGINAL TIME" OR TRUE TIME, BECOMES INCREASINGLY MORE WRONG THE FURTHER DOWN A CHAIN OF PRODUCERS IT IS.

QUESTIONS I HAD

* WHY DO WE NEED NTP?

NTP MEASURES AND CORRECTS FOR INCIDENTAL CLOCK FREQUENCY ERRORS. IT ALSO PREVENTS DRIFT.

* WHAT IS DRIFT?

SYSTEMS CALCULATE TIME BY COUNTING MOVEMENTS OF A SMALL CRYSTAL BASED ON THE CPU'S CLOCK FREQUENCY.

SYSTEMS CAN RUN AT DIFFERENT CLOCK FREQUENCY => THEY COUNT TIME DIFFERENTLY. OVER A PERIOD OF TIME, THESE SYSTEMS "DRIFT" APART AND GO OUT OF SYNC.

* IS IT REALLY THAT BIG OF A DEAL?

MAYBE NOT FOR A NOTIFICATION ON TWITTER, BUT IT'S VERY SERIOUS ON REAL TIME SYSTEMS. LIKE A SYSTEM THAT CONTROLS AIR TRAFFIC.

Sweet Things About The NTP Algorithm

1. It doesn't just copy the provider's time to the consumer's; it accounts for latency, drift, & more!

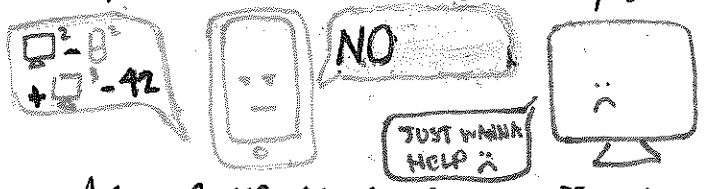

2. As consumer becomes more in sync with producer, NTPD makes less requests!
   Initially, once a minute ...
   Later, once every 17 minutes ☺

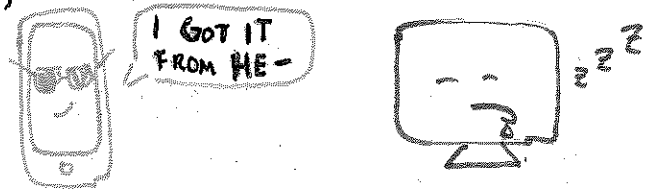

3. It's a decentralized algorithm, meaning systems can be both a time consumer and time provider.

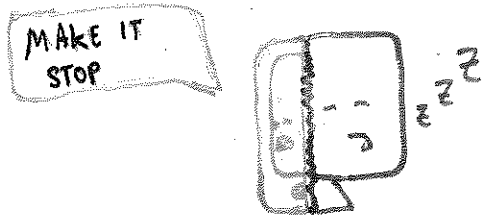

Decentralized Systems

A set of connected NTP servers characterizes a **decentralized distributed system**, or a distributed system with no single point where decisions are made.

Consider the following analogy:

Im walking on campus and I want to know:

I could ask anyone, suppose someone says:

Now I can tell anyone else, just like how a time consumer can be a time producer for another consumer.

But...