# Perspectives on Network and Computer Timekeeping

WSTS 2024 / San Diego, CA / 2024-05

Kevin B. Stanton, Ph.D.
Stanton Consulting LLC
kevin.b.stanton@gmail.com
Linkedin

# End-to-end Time Distribution

Local

Global

Galactic

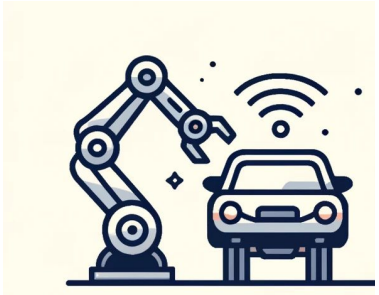

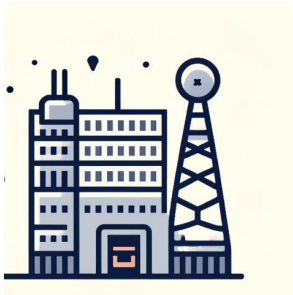

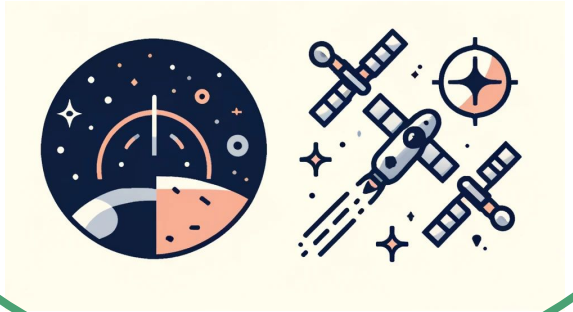

Not to scale

# Timing in the Last Centimeters
# ⇒ *The Need*

# Timing in the Last 2.54 Centimeters (“the last inch problem”)

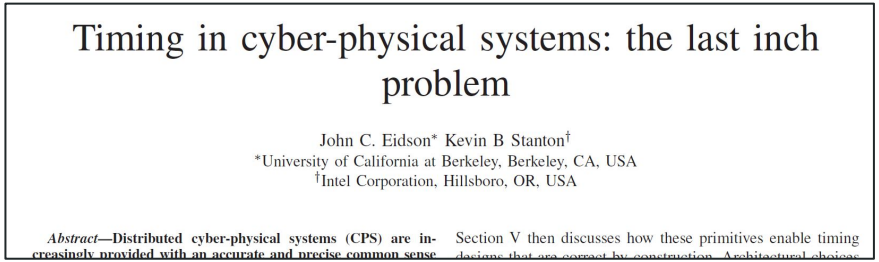

Timing in cyber-physical systems: the last inch problem

John C. Eidson* Kevin B Stanton†
*University of California at Berkeley, Berkeley, CA, USA
†Intel Corporation, Hillsboro, OR, USA

*Abstract*—Distributed cyber-physical systems (CPS) are increasingly provided with an accurate and precise common sense Section V then discusses how these primitives enable timing


“In either case, close attention must be paid to path latency between the PHY and the PTP clock in the microprocessor. In particular, **substantial degradation of accuracy** can occur when transferring PTP time from a **network interface card to the microprocessor** over a serial load/store interface such as PCI Express.”

“However, a recent **PCI Express enhancement** for in-band time transfer over PCI called **Precision Time Measurement (PTM)** promises to transfer time to the microprocessor with **accuracy in the nanoseconds**[35].”

https://ptolemy.berkeley.edu/projects/chess/pubs/1156.html

# We Knew that PTP Would both Solve ***and Create*** Problems

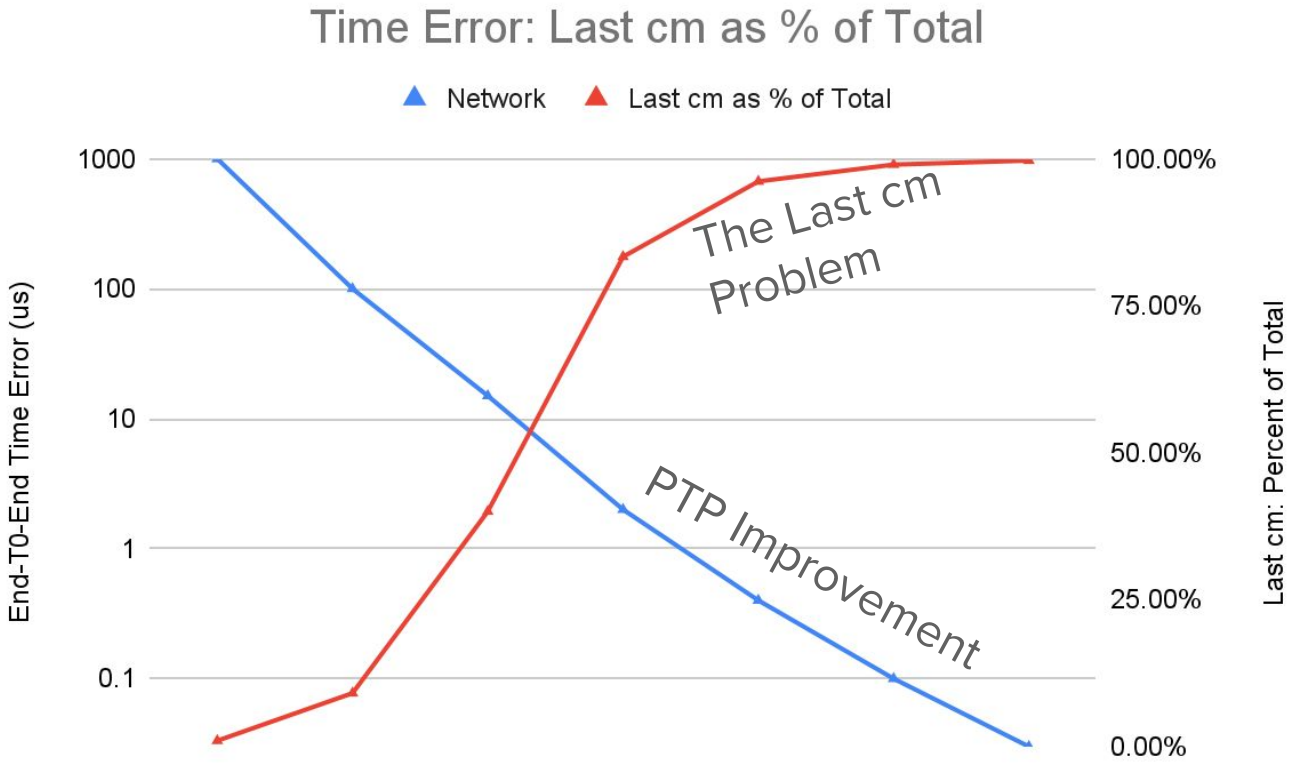

# Time Transfer To Software (Within the Compute System)

## The Way We've Always Done It (Software-Based)

1. Read CPU Counter
2. Read PTP Counter in PCIe NIC

   *Wait (for the long & variable delay)*
3. Read CPU Counter

⇒ What could possibly go wrong?

## Changing the Rules with PTM Hardware-Based

1. Extend the CPU Counter to PCIe NIC
2. NIC "simultaneously" captures
   a. CPU Counter
   b. PTP Counter in PCIe NIC

⇒ *Synchronous Logic: "Simultaneous"*

# Timing in the Last Centimeters
# ⇒ *The Cure*

# Software Time versus PTP Time via PPS Outputs

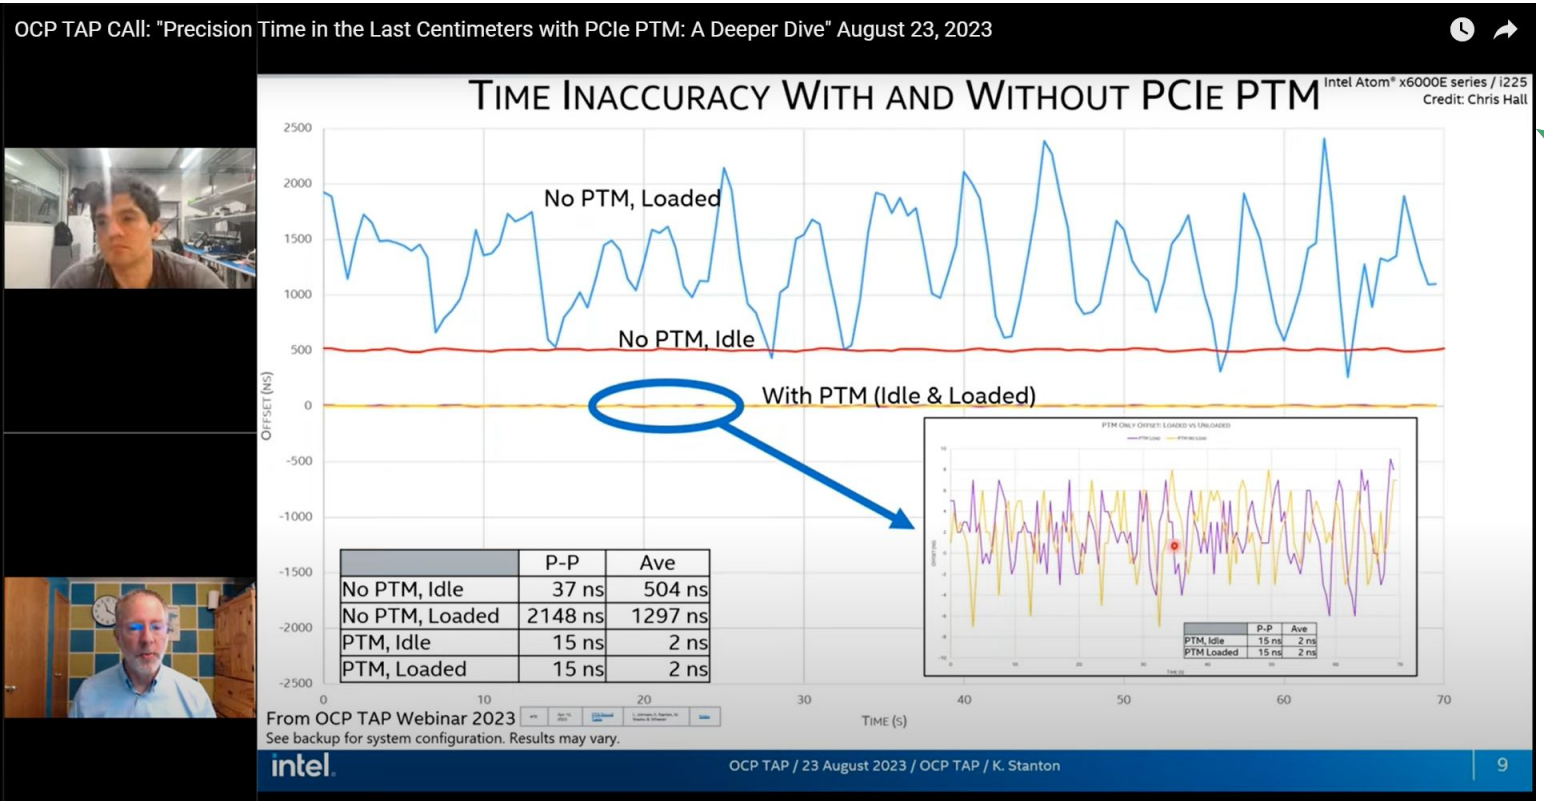


This Chris H. has a Poster at WSTS 2024. Check it out!

https://youtu.be/9OILFLV-Sfc?si=U9EX2iB4PTiGiIV3&t=899

# Last Centimeters Time-Sync…

# …is “Solved”

Link to Youtube on Previous Page

# Two-Way Transfer: Detail

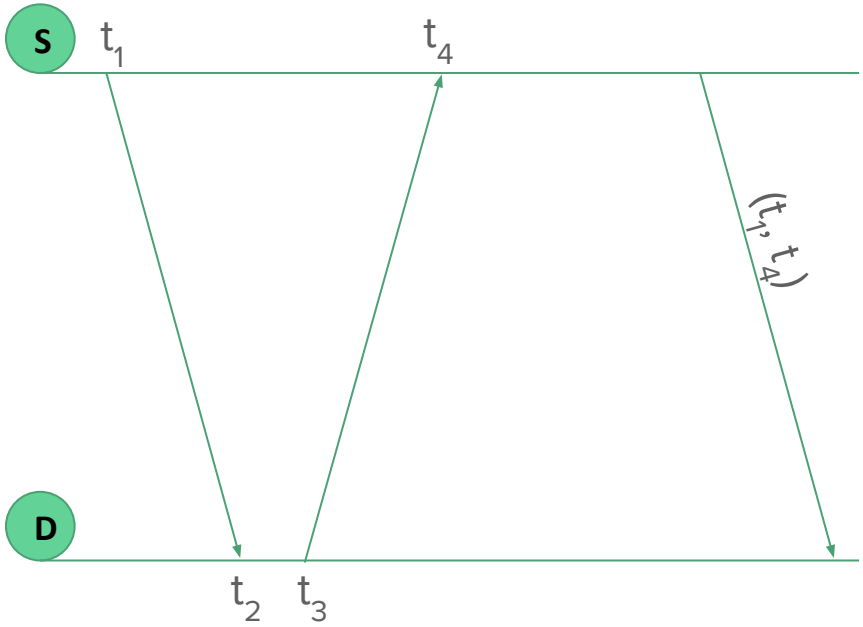


Time **S**ource

Time **D**estination

Link Delay = $[(t_4-t_1) - (t_3-t_2)] / 2$

Clock Offset = $[(t_2-t_1) - (t_4-t_3)] / 2$

Source time at $t_2$ = $t_1$+Link Delay

= $t_1 + [(t_4-t_1) - (t_3-t_2)] / 2$

Questions

1. Is the Delay Symmetrical
   a. Over Ethernet?
   b. Over PCIE tree
   c. Over SOC Fabric?

# Timing in the Last Centimeters
# ⇒ *Lessons Learned*

# The Path to PTM

## The Three Criteria

https://commons.wikimedia.org/wiki/File:Three_first_fingers.JPG

## IEEE 1588 / PTP

1588

## Ethernet NICs


## PCIe PTM Spec

## CPU and SoC Timekeeping Upgrade

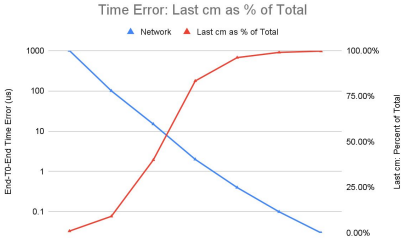

# PTM Readiness

List of Precision Time Measurement Readiness [edit]

| | Class | Manufacturer | Product | Quanta | TGPIO | PTM Status |
|---|---|---|---|---|---|---|
| 1 | CPU | Intel | Core, Gen 11 (Tiger Lake) | 4ns | | Supported |
| 2 | CPU | Intel | Core, Gen 12 (Alder Lake) | 4ns | | Supported |
| 3 | CPU | Intel | Core, Gen 13 / Gen14 (Raptor Lake) | 4ns | | Supported |
| 4 | CPU | Intel | Xeon, Gen 4 (Sapphire Rapid) | 2ns | | Supported |
| 5 | NIC | Intel | I225 (Foxville) | 4ns | | Supported |
| 6 | NIC | Intel | I226 (Foxville) | 4ns | | Supported |
| 7 | FPGA | Intel | Intel Agilex | 4ns | | Supported |
| 8 | Microprocessor | Texas Instruments | AM64x / AM65x / AM68x / AM69x / DRA82x ARM SoCs | 4ns | | Supported |
| 9 | NIC | Nvidia | ConnectX-6 | 2ns | | Beta Tested |
| 10 | NIC | Nvidia | ConnectX-7 | 2ns | | Beta Tested |
| 11 | Time Card | OCP | Time Card 2 | 4ns | | Supported |
| 12 | Time Card | Safran | ART2 | 4ns | | In Progress |
| 13 | Time Card | Adtran | 5400 SyncModule | 4ns | | In Progress |
| 14 | NIC | Liquid-Markets-Solutions | UberNIC | 4ns | | Supported |
| 15 | NIC | Broadcom | P2100G | 4ns | | In Progress |
| 16 | Corundum | UCSD | Corundum | 4ns | | In Progress |
| 17 | Motherboard | ASRock | SP2C741D16-2T | 2ns | 2 exposed | Supported |
| 18 | Motherboard | ASRock | W790D8UD-1L1N2T/BCM | 2ns | 2 exposed | Supported |
| 19 | Motherboard | Supermicro | X13SEDW-F | 2ns | | Beta Testing |
| 20 | CPU | Intel | Core, Gen 14 (Raptor Lake Refresh) | 4ns | | Supported |
| 21 | Motherboard | Asus | W790SE | 2ns | 1 under BGA | Supported |
| 22 | Motherboard | Asus | Z790 Maximus Extreme | 2ns | 1 under BGA | Supported |
| 23 | Motherboard | Asus | Z690-I | 2ns | 1 under BGA | Supported |
| 24 | Motherboard | Asus | Z790-G WIFI | 4ns | 1 under BGA | Supported |
| 25 | Motherboard | Asus | Z690-Extreme | 4ns | 1 under BGA | Supported |
| 26 | Motherboard | Asus | W790-ACE | 2ns | 1 under BGA | Supported |
| 27 | Motherboard | Asus | Z790-I | 4ns | 1 under BGA | Supported |

https://www.opencompute.org/wiki/PTM_Readiness

# The Future of Time Synchronization
# ⇒ What's Next?

# Google’s Andrew Fikes on Time


https://www.youtube.com/watch?v=nvlt0dA7rsQ

“Everybody tells you [that] you ***can’t trust Time***. The clock on your server is ***absolutely*** something that you ***can’t trust***, and you develop this ***innate fear*** of that thing. *[Distrust of Time]* ***is Legendary…***

I can tell you that I ***completely trust our time system***, and it has made my life so much easier. Having a ***source of global ordering*** that I don’t have to reach out to talk to, that I can ***reference locally***, is just a ***huge huge huge thinking shift***.

… It took two or three versions [of the Google database, Spanner] to get over that [fear of relying on time]”

— Andrew Fikes (VP of Infrastrustructure @ Google, CoudNext 2019) (Emphasis Mine)

# The Future (Beyond Accuracy)

APIs Beyond “Now()”

Authenticated Time Source

Resilient Timing Sources / Paths

Explicit Time Accuracy / Jitter Bounds

Explicit Consistency / Availability Tradeoffs

Non-Newtonian Timekeeping

CISA
CYBER+INFRASTRUCTURE

DEFEND TODAY. SECURE TOMORROW.

## TIME – THE INVISIBLE UTILITY

### WHY IS TIME IMPORTANT?

Time is critical to certain services used within most organizations, yet many organizations are unaware of their dependence on time, the source of their time, or the existence of a world time standard.[1] As systems grow in complexity, becoming global and mobile, access to resilient, accurate, and precise time is a necessity in both the private and public sectors worldwide. Without accurate and resilient time, critical functions and services can become unreliable, inaccurate, or unavailable.

### SECTORS AND INDUSTRIES DEPENDENT ON TIME

| Communications | Transportation | Power Grid | Finance | Security | IT |
|---|---|---|---|---|---|
| Telecommunication<br>Cloud Operations<br>Internet of Things (IoT) | Aviation<br>Maritime<br>Pipelines<br>Rail | Frequency Monitoring<br>Multi-rate Billing<br>Fault Detection | Regulatory Requirements<br>ATM Networks | Cryptography<br>Access Control<br>Forensics<br>Surveillance | Smart Devices<br>Incident Investigations |

### WHY SHOULD YOU BE CONCERNED ABOUT TIME NOW?

GPS has become the *de facto* time standard for many commercial users because of its relatively low cost and ubiquitous availability. In 1997, the President's Commission on Critical Infrastructure Protection

https://www.cisa.gov/sites/default/files/2023-02/Technical-Level_Resilient_Timing_Overview-CISA_Fact_Sheet_508C.pdf

# Summary

1. The Last Centimeters Problem is “solved”

   Implement, adopt, and deploy PTM

2. Software applications need more information

   Today, they get only

```
now()
```

As we increase our TRUST in the TIME, let’s not forget to inform the application

# Thank You!

# Software Access to “Now”

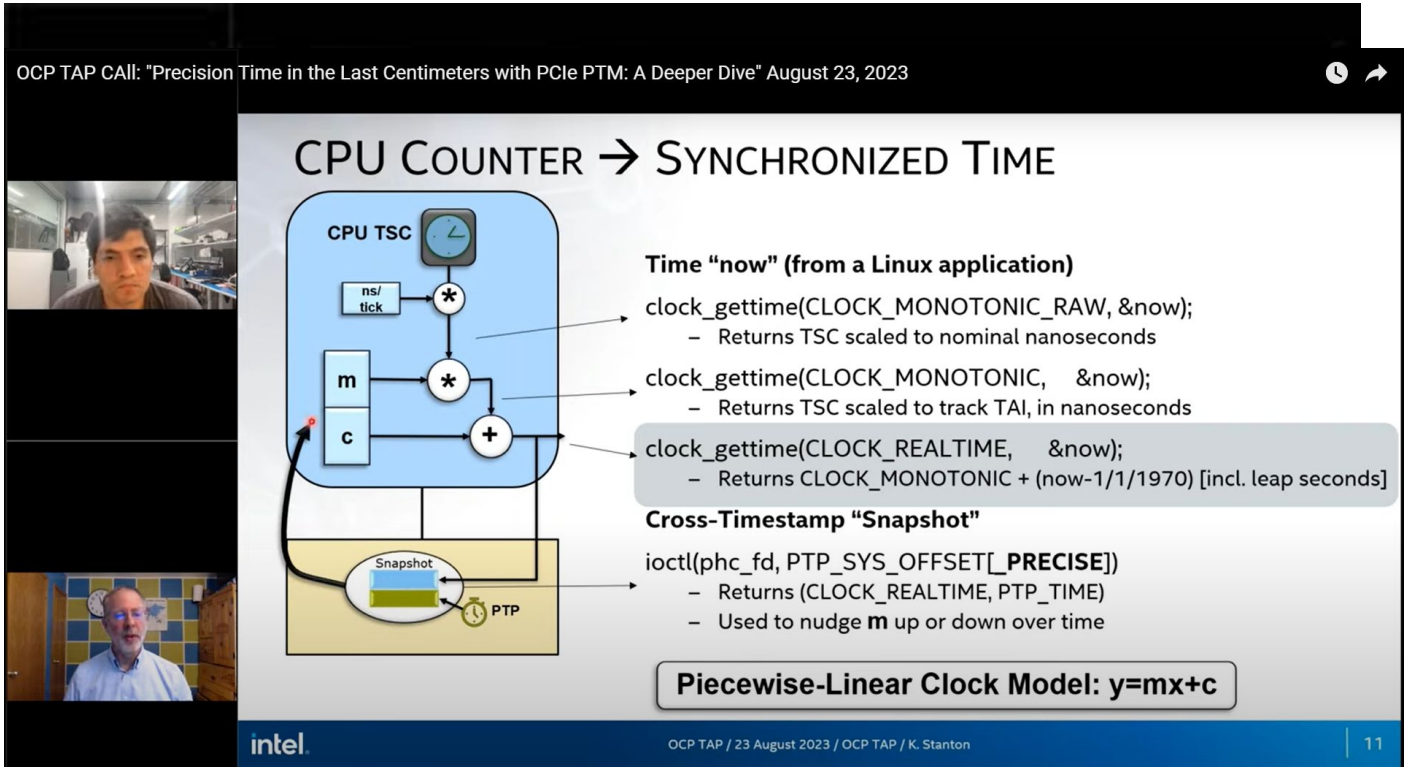


https://youtu.be/9OILFLV-Sfc?si=ohnZ9U07JT7eoYN6&t=1211

# PCIe PTM: How it works — between chips / chiplets

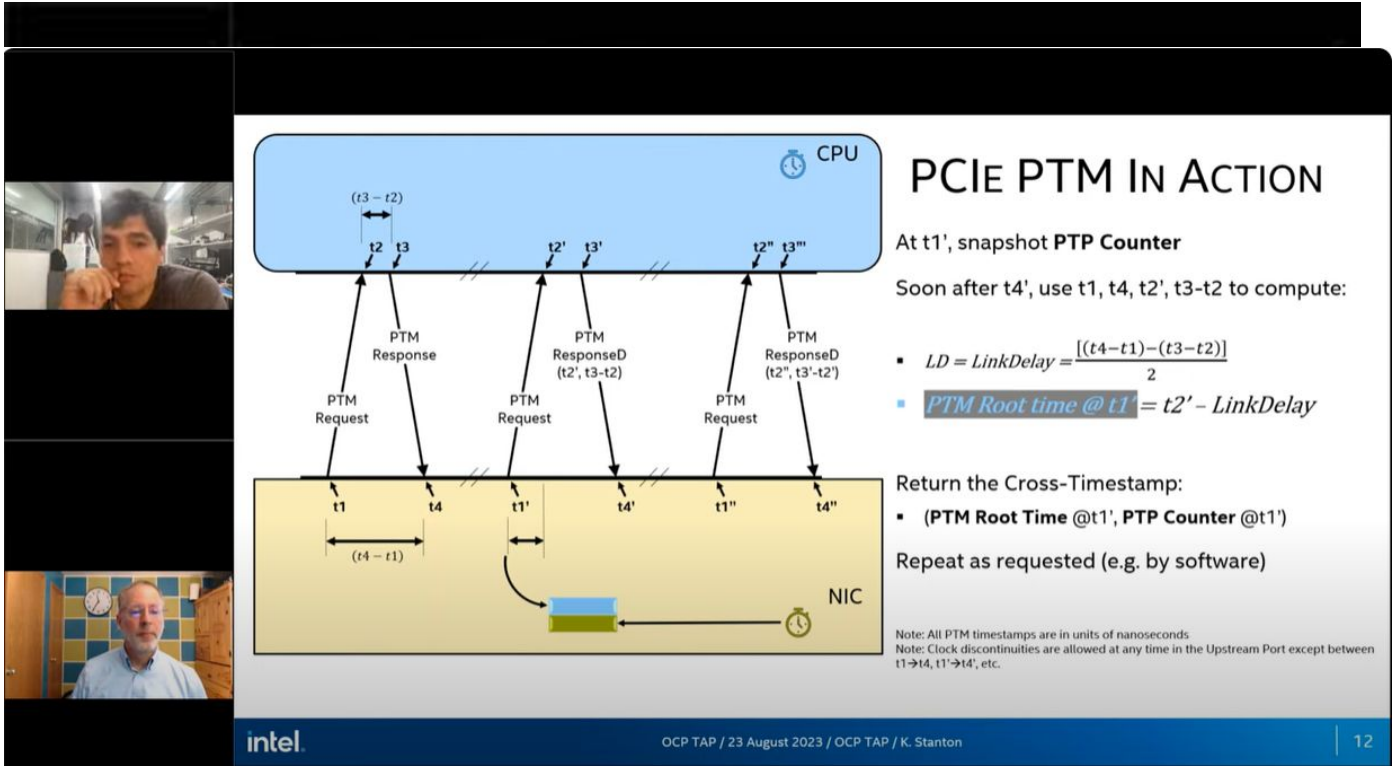


https://youtu.be/9OILFLV-Sfc?si=6n4kS-W7645JfwJd&t=1975