ROSELINE: Enabling Robust, Secure and Efficient Knowledge of Time Across the System Stack

Award # CNS-1329755 (UCLA), CNS-1329644 (CMU), CNS-1329644 (UCSD), and CNS-1329650 (UCSB)

Type: Frontier; Start Date: June 2014

PI: Mani Srivastava (UCLA)

Collaborators: Rajesh Gupta (UCSD), Joao Hespanha (UCSB), Sudhakar Pamarti (UCLA), Raj Rajkumar (CMU), Anthony Rowe (CMU)

Managing Uncertainty in Time

To a software developer, time is a simple number….

clock_gettime( ) ➜ 1375599686.9

qot( ) ➜

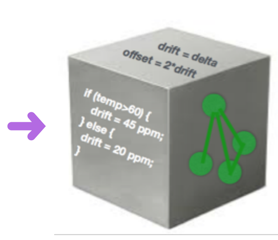

"QoT" to manage uncertainty

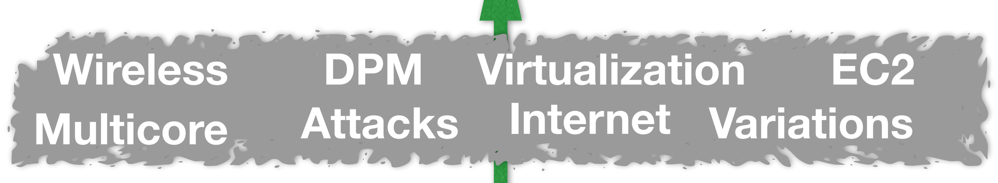

Uncertainty Amplified
Little Visibility
No Control

But reality is messy

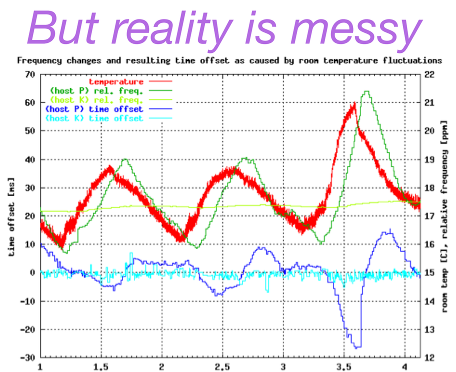

ROSELINE seeks to transform the uncertainty in time in computing system into a rich structure called "Quality of Time" (QoT) that is systematically measured, propagated and controlled throughout the system.

System Software Mechanisms for QoT

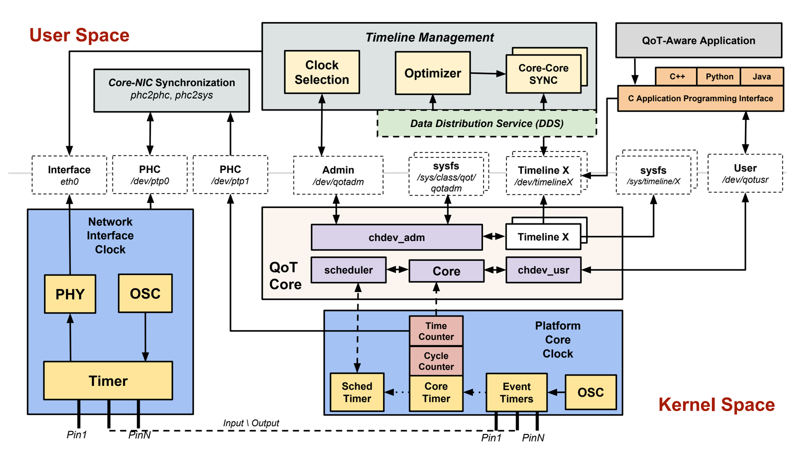

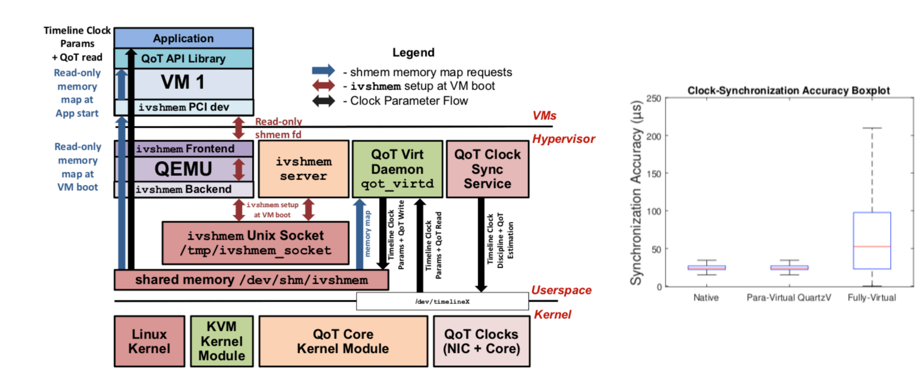

QuartzV: Bringing Quality of Time to Virtual Machines

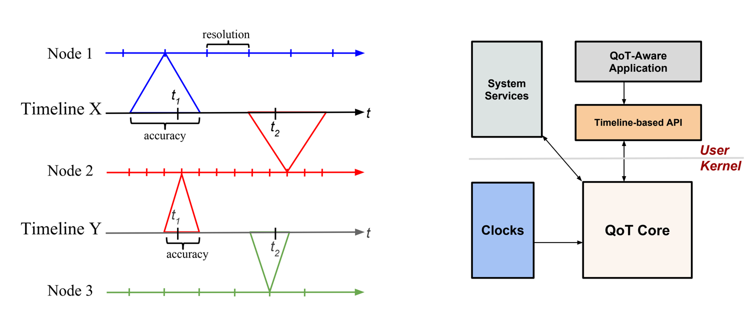

Timeline OS Abstraction

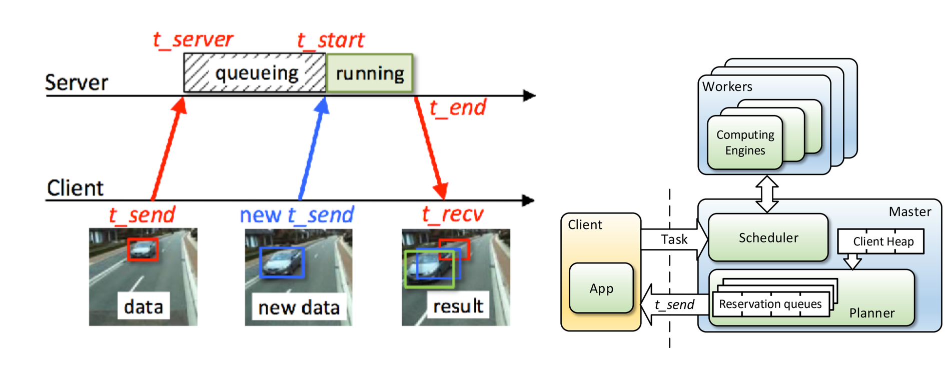

ATOM: Bringing Quality of Time across Edge-Cloud Boundary

QoT Challenges in Smartphone Apps

Shared notion of time is needed to fuse data across devices or to coordinate actions.

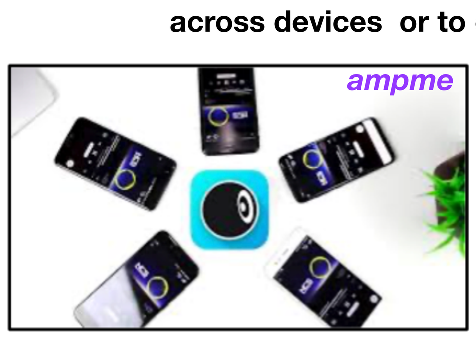

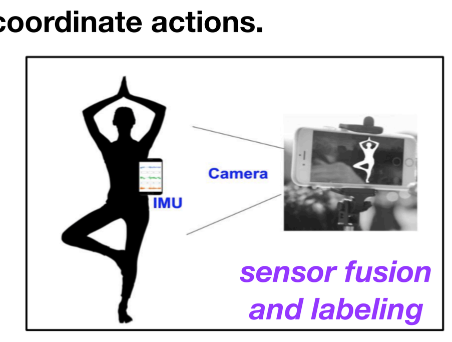

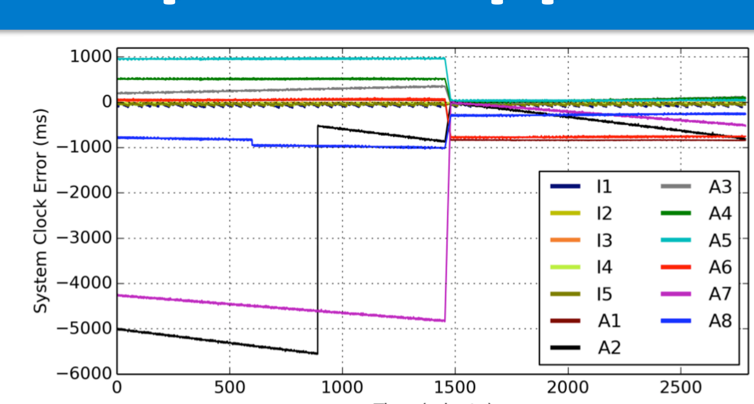

Observed system clock errors before & after a restart event at minute 1460 across five iOS devices (I1-I5) and eight Android devices (A1-A8) compared to an NTP baseline.

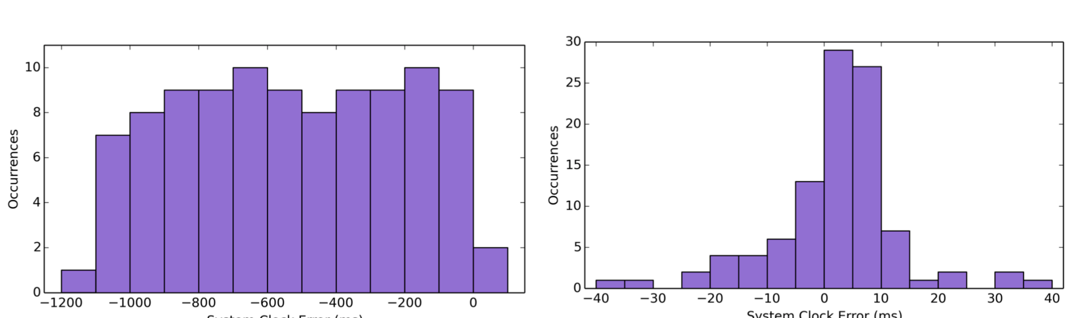

System clock error of an Android Nexus 5X device after triggering 100 independent NITZ timing updates (left), and 100 independent NTP timing updates (right).

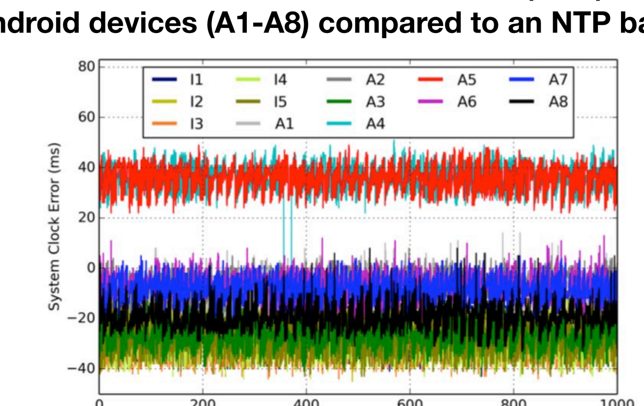

Error improvement with our GoodClock App-level Open-source NTP Library.

Exploiting Smatphone Peripherals for Better QoT

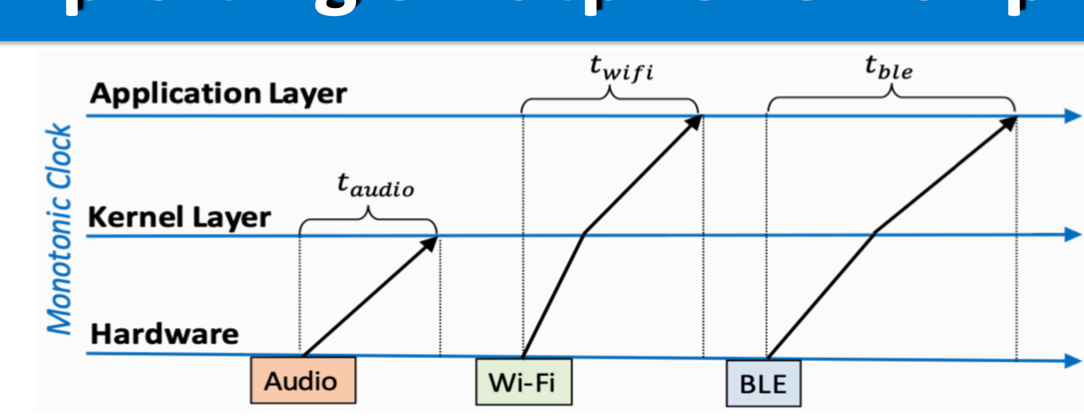

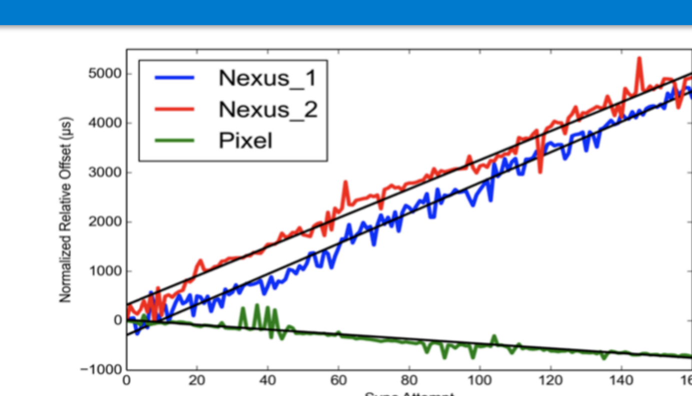

Timestamping events for audio, Wi-Fi and BLE peripherals and Android (left); and Drift in the relative clock offset for audio-based sync over a duration of 1 hour with respect to the fourth pixel phone serving as reference (right).

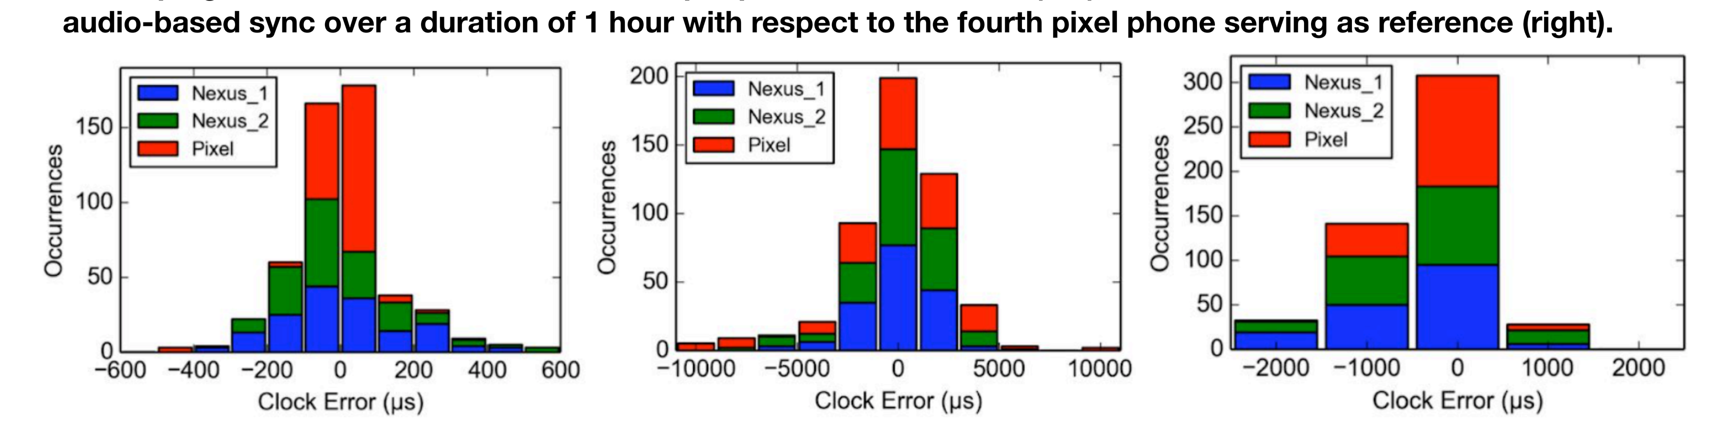

Sync offset variability in three Android smartphones with respect to a fourth reference device for (a) audio, (b) BLE, and (c) Wi-Fi implementations. Note: 86% of audio sync attempts fall within ±200µs. 85% of BLE sync attempts fall within ±3000µs. 95% of Wi-Fi sync attempts fall within ±1000µs.

Mitigawting Adverse QoT in ML-enabled Apps

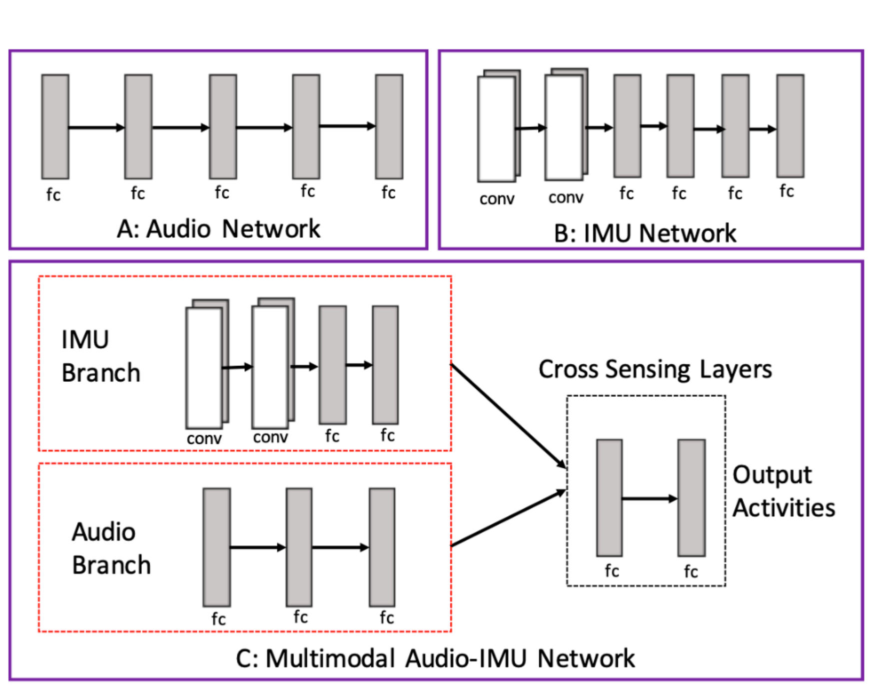

Multimodal Audio- IMU Deep Neural Network with separate branches for Audio and IMU data fused in cross sensing layers.

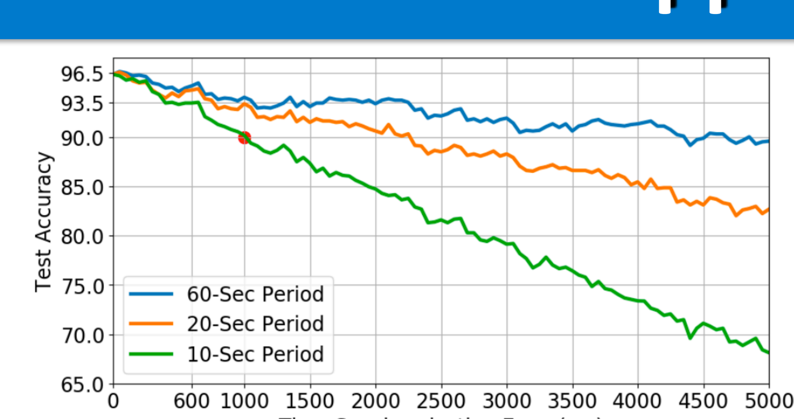

Variation in accuracy of Multimodal Audio-IMU network with increased timing errors between smartphones collecting audio and IMU data.

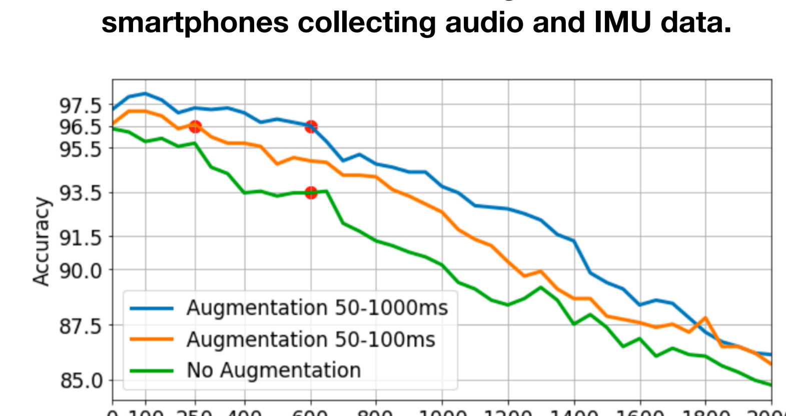

Comparison of Multimodal Audio-IMU Network test accuracy with different augmented training datasets.

Clocks for IoT Energy & Complexity Reduction

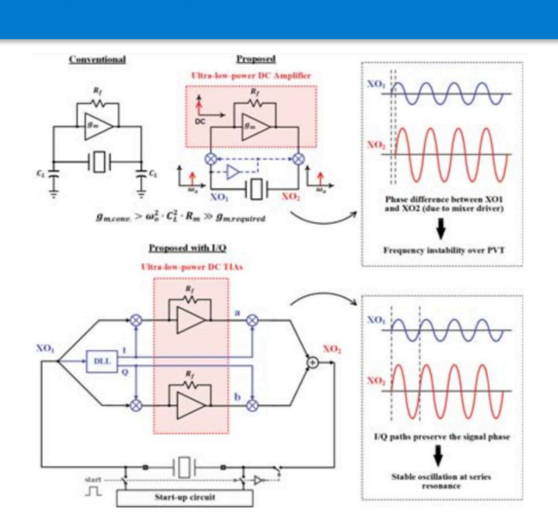

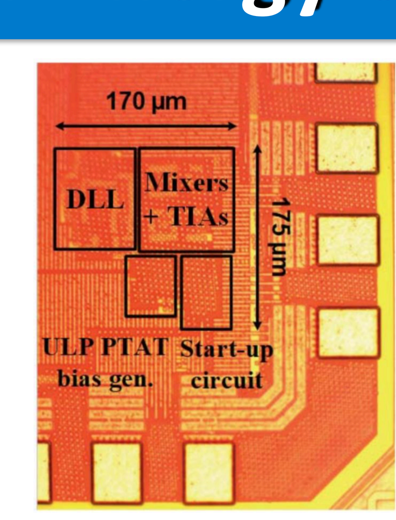

Conventional Pierce XO architecture vs our ultra low power XO (left); and. a die micrograph of the IC (right)

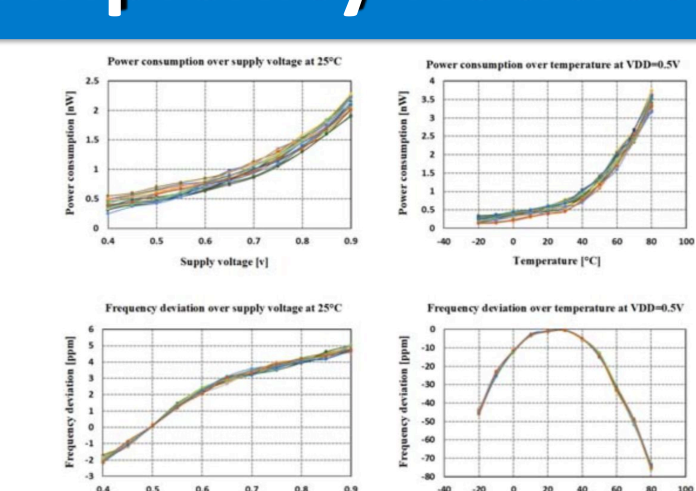

Power consumption and frequency variations over VDD (left) and temperature (right) measured over 20 dies.

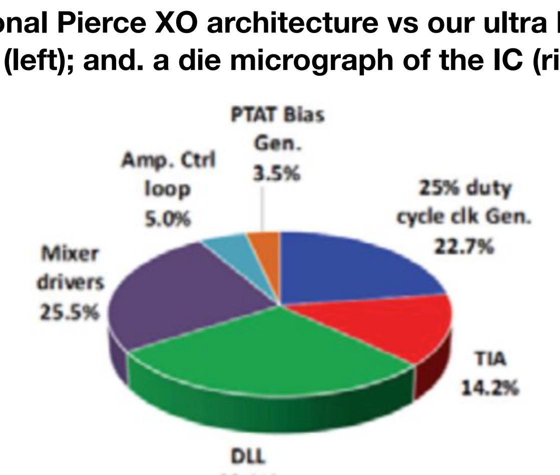

	This work*	VLSI'17 [1]	JSSC'16 [2]	ISSCC'14 [3]	JSSC'16 [4]
Technology (nm)	65	55	180	28	180
Area (mm^2)	0.027	0.16	0.3	0.03	0.062
Supply voltage (V)	0.5	0.5	1.2	0.2	0.3
Power at 25°C (nW)	0.55	1.7	5.58	5	1.5
Temperature stability (ppm)	80 (-20°C to 80°C)	109 (-20°C to 80°C)	133 (-20°C to 80°C)	50 (-20°C to 80°C)	150 (0°C to 80°C)
Line sensitivity (ppm/V)	13	6.7	30.3	85	7
Allan Deviation (ppb)	14	25	10	10	70
Calibration required	NO	YES	YES	NO	YES

*Averaged across 20 dies with worst-case power of 0.7nW.

Carnegie Mellon University

UCLA

UCSB

UCSD