Toward Nanosecond-Accurate Timestamping on Low-Cost SDRs

Kevin Chu, Logan Byard, Ali Abedi
University of Wisconsin – Madison

(I) Motivation

Nanosecond-level timestamping on low-cost SDRs is critical for enabling precise synchronization, localization, and cross-platform signal analysis in distributed wireless systems. Many emerging applications, including spectrum monitoring, satellite interference attribution, coherent sensing, and multi-receiver measurements, require accurate knowledge of when each IQ sample was captured. Bringing nanosecond timing support to accessible SDR platforms can make these capabilities available beyond expensive lab-grade radios, enabling scalable and reproducible experiments in real-world deployments.

(II) Common Misunderstanding

A common misconception is that once an SDR is connected to a GPSDO, each received sample is automatically timestamped at the exact nanosecond the analog signal arrives at the RF front end. In reality, the signal experiences non-negligible and often hidden delays as it passes through the RF chain, ADC, digital filters, FPGA logic, and buffering before a timestamp is associated with the resulting samples. These internal delays can be much larger than a nanosecond; for example, on the USRP B210, the group delay of the AD9361 transceiver alone can introduce microsecond-scale timing error if it is not accounted for.

(III) Experiment Setup

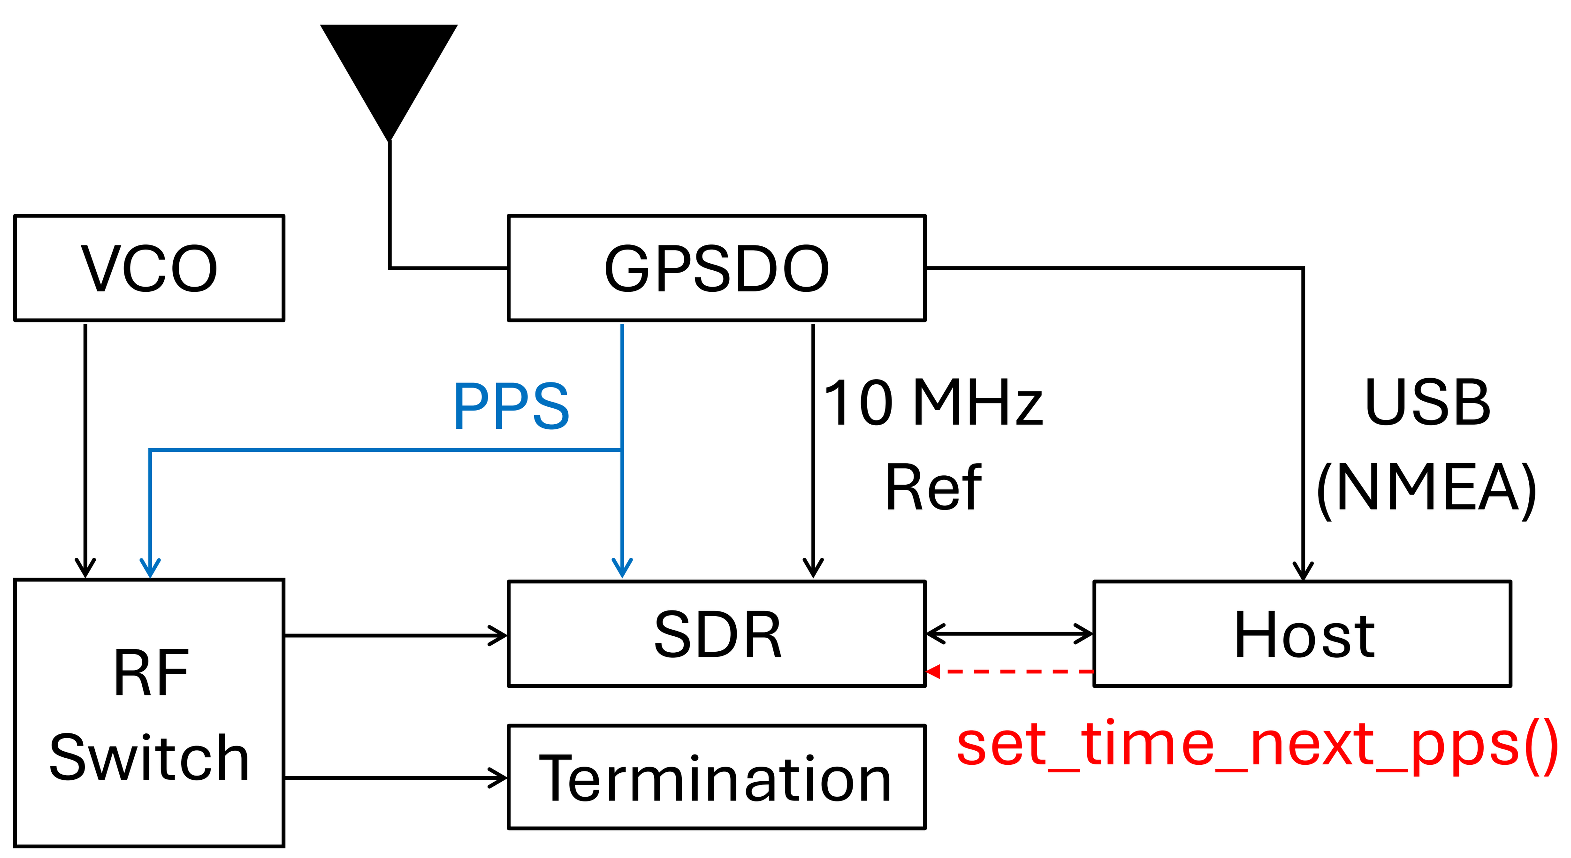

(IV) Methodology

To evaluate SDR timestamping accuracy, we build a PPS-aligned RF testbed using a USRP B210, an external GPSDO, and a controlled RF switch. The GPSDO provides both a 10 MHz reference and a 1 PPS signal to the SDR, while also sending NMEA time messages to the host. The host uses this information to configure the SDR clock with set_time_next_pps, aligning the SDR time counter to UTC at the next PPS edge. To test the actual timestamp accuracy, we feed a VCO-generated sine wave through an RF switch that is also triggered by the PPS. When the PPS rising edge occurs, the switch connects the tone to the SDR input. By detecting the rising edge in the received sample energy and comparing its reported timestamp to the exact start of the UTC second, we estimate the SDR's end-to-end timestamp offset.

(V) Results

Our measurements show that the received PPS-aligned RF edge appears with a microsecond-scale timestamp. We believe this offset is primarily caused by latency in the SDR receive signal chain, especially the group delay introduced by the AD9361 transceiver and its internal digital filtering. The results support the signal-chain delay hypothesis: as the sampling rate increases from 20 to 50 MSps, the measured deviation decreases. On the USRP B210, changing the sampling rate can also change the internal clocking and filtering configuration of the AD9361, which changes the effective signal-processing delay.

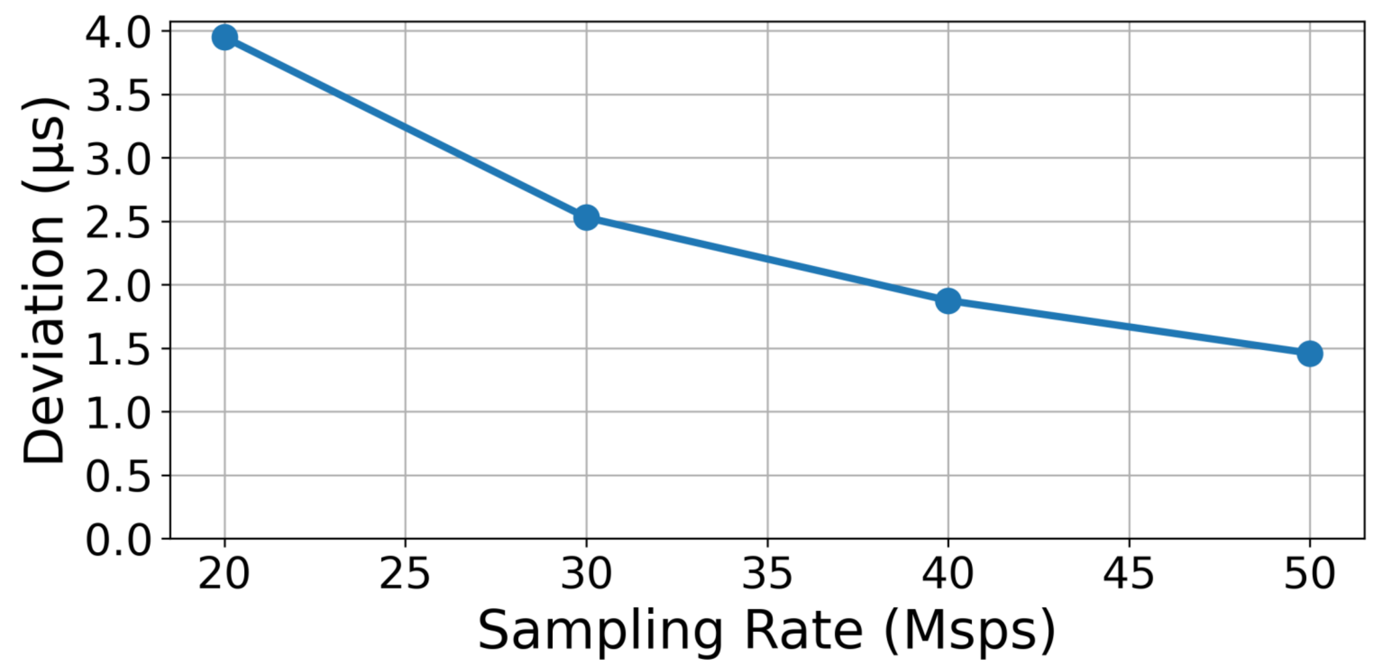

These results show that GPSDO synchronization should not be treated as a complete timing solution without calibration. Even when multiple SDRs share a GPSDO, their received samples may still have different effective timing offsets because each SDR has its own RF front-end, ADC, filtering, and FPGA signal-processing delay. This is especially important when combining measurements from different SDR models, which may not be truly time-aligned unless their end-to-end receive-path delays are considered.