Synchronism, Jitter, and Communications in Accelerator Controls

Larry Doolittle, LBNL
October 2007 LLRF Workshop, Knoxville

- **Concept:**
  - Run all circuitry in the accelerator off a single master oscillator
- **Advantages:**
  - Cycle counting gives perfect interval measurement
  - No need for crude clock disciplining
  - Synchronous stimulus-response response measurements are easy accelerator-wide
  - No added jitter anywhere
  - Traditional communication techniques work with unlimited packet length
- **Challenges:**
  - Justify costs and reliability penalty of an additional cable for some chassis
  - Managing divider state
  - Establishing absolute time at each power-on
  - Leveraging commodity communication standards

"[We] don't know what's going on [with the reflected power leading edge] because [the data] acquisition isn't synchronized with the RF pulse"

– Tom Powers, this workshop

Outline

The lore of crossing clock domains

Clock coherence does not imply equality

Implications for bunch (train) spacing in pulsed linacs

Fine frequency adjustments for storage rings

Communication across synchronized clock domains

Establishing absolute time at power-on

The problem with computers

Implications for system simulation

Final challenge

The lore of crossing clock domains

Events

Gray codes

FIFOs

Metastability - can be managed but not eliminated

Inherently adds time indeterminism (jitter) of at least one output clock cycle

The simplest and most reliable FPGA designs use pure synchronous logic in a single clock domain

Clock coherence does not imply equality

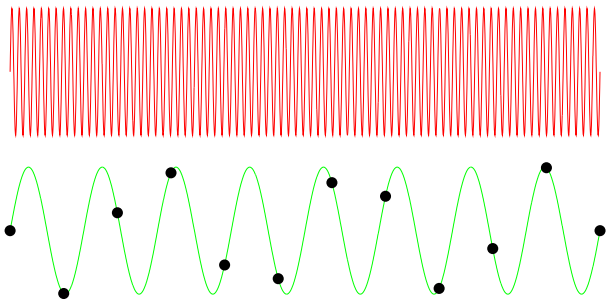

Example from “500 MHz” light source

$$\mathrm{RF} = 500\ \mathrm{MHz}$$
$$\mathrm{LO} = 550\ \mathrm{MHz}$$
$$f_\mathrm{S} = \mathrm{LO} \div 8 = 68.75\ \mathrm{MHz}$$
$$\mathrm{IF} = 50\ \mathrm{MHz}$$

Transverse feedback digital logic (FPGA) is clocked at 250 MHz (two samples per clock), one tick represents two RF buckets.

LLRF digital logic (FPGA) is clocked at 68.75 MHz, one tick represents $7\frac{3}{11}$ buckets.

Good thing an FPGA can implement a mod 11 counter!

Implications for bunch (train) spacing in pulsed linacs

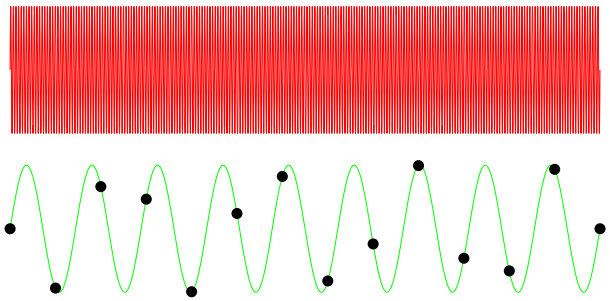

Example from ILC

$\mathrm{RF} = 1300$ MHz

$\mathrm{LO} = 1352$ MHz

$f_{\mathrm{S}} = \mathrm{LO} \div 18 = 75.111$ MHz

$\mathrm{IF} = 52$ MHz

Beam pulses are necessarily synchronous with the 1300 MHz bucket clock.

They should also be synchronized with the sample clock, or the pulse will "look" different in a way that will needlessly confuse the pulse-learning processes.

Conditions satisfied every 225 buckets = 173.1 ns.

5 Hz beam pulses should therefore arrive in multiples ($577778 \pm 1\%$) of this time.

Timing system must schedule pulses according to this plan, and broadcast the actual number that keeps the next pulse approximately line synchronous.

Constraints from other subsystems will multiply this granularity restriction.

Fine frequency adjustments for storage rings

Ring clocks traditionally tunable for optics reasons

1. analog or digital single-point tune, everything else follows
   - \+ simple, effective
   - \- difficult to get clock time (TAI) by counting machine cycles
   - \- can't use crystal VCXO for derived PLLs
2. distribute fixed frequency, digitally slew phase (DDS) at point of use
   - \+ phase angle of long reference lines doesn't change
   - \- frequency changes have to be precisely synchronized
   - \- absolute phase is difficult to restore after node power cycle

For choice 1, the DDS that creates the ring clock presumably runs from a high accuracy low jitter clock, which can be counted to get TAI. Messages need to flow from this DDS system to the rest of the machine giving the information needed to convert cycle count to TAI seconds. In many cases, *e.g.,* correlating waveforms between acquisition systems, it is both simpler and more precise to use cycle count instead of TAI seconds.

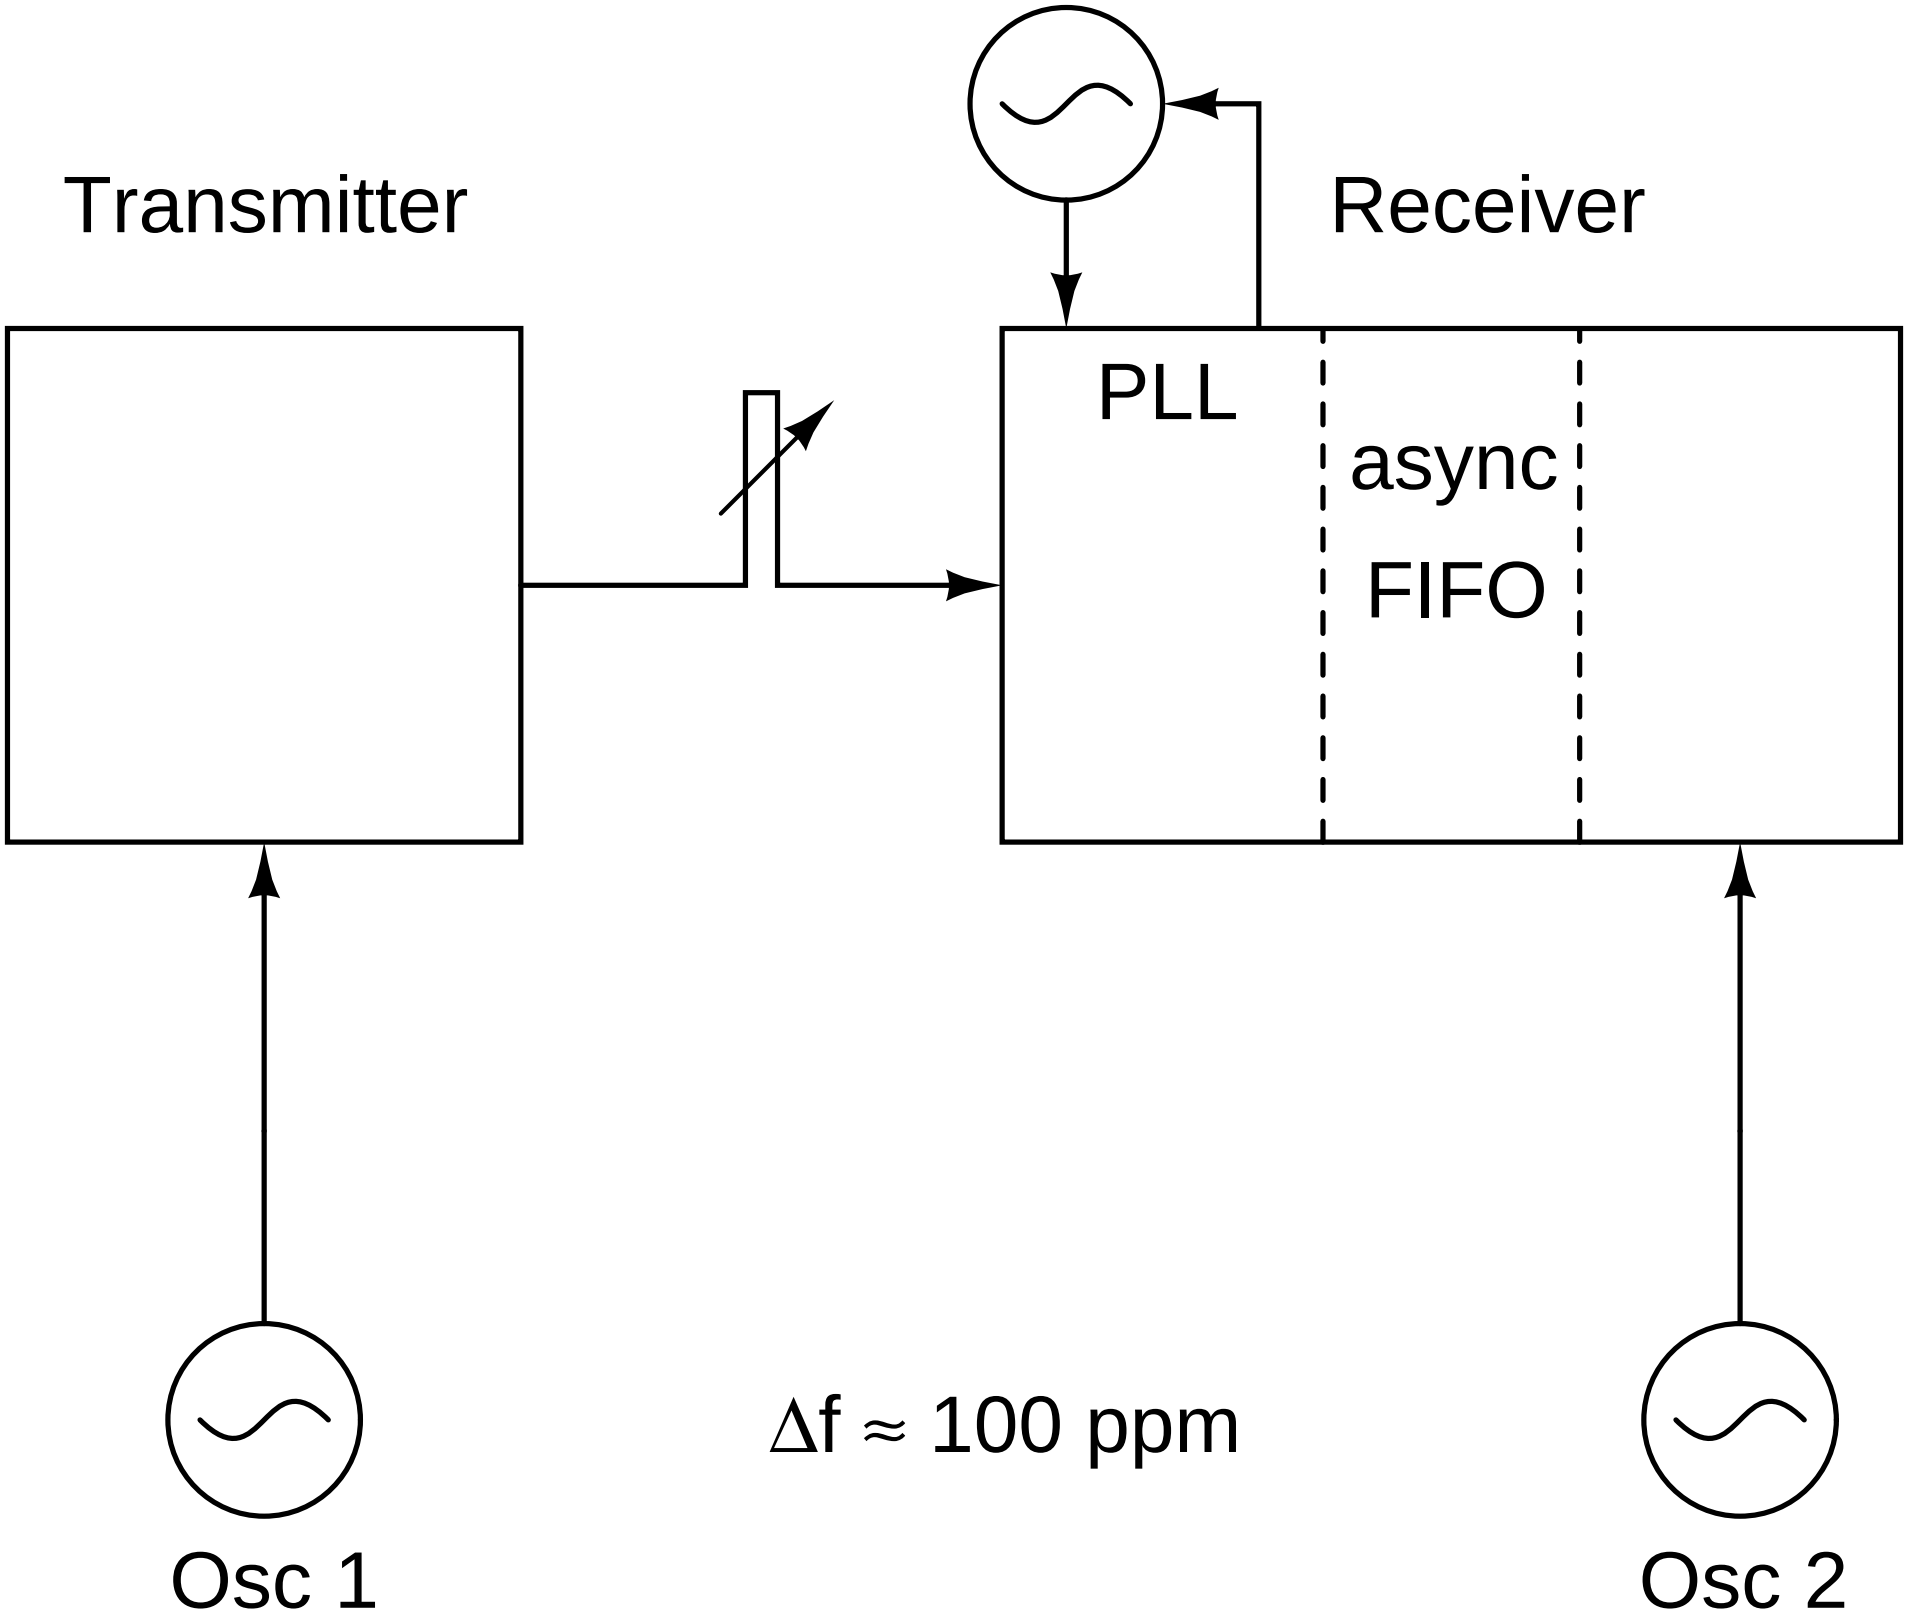
Transmitter
Receiver
PLL
async
FIFO
Δf ≈ 100 ppm
Osc 1
Osc 2

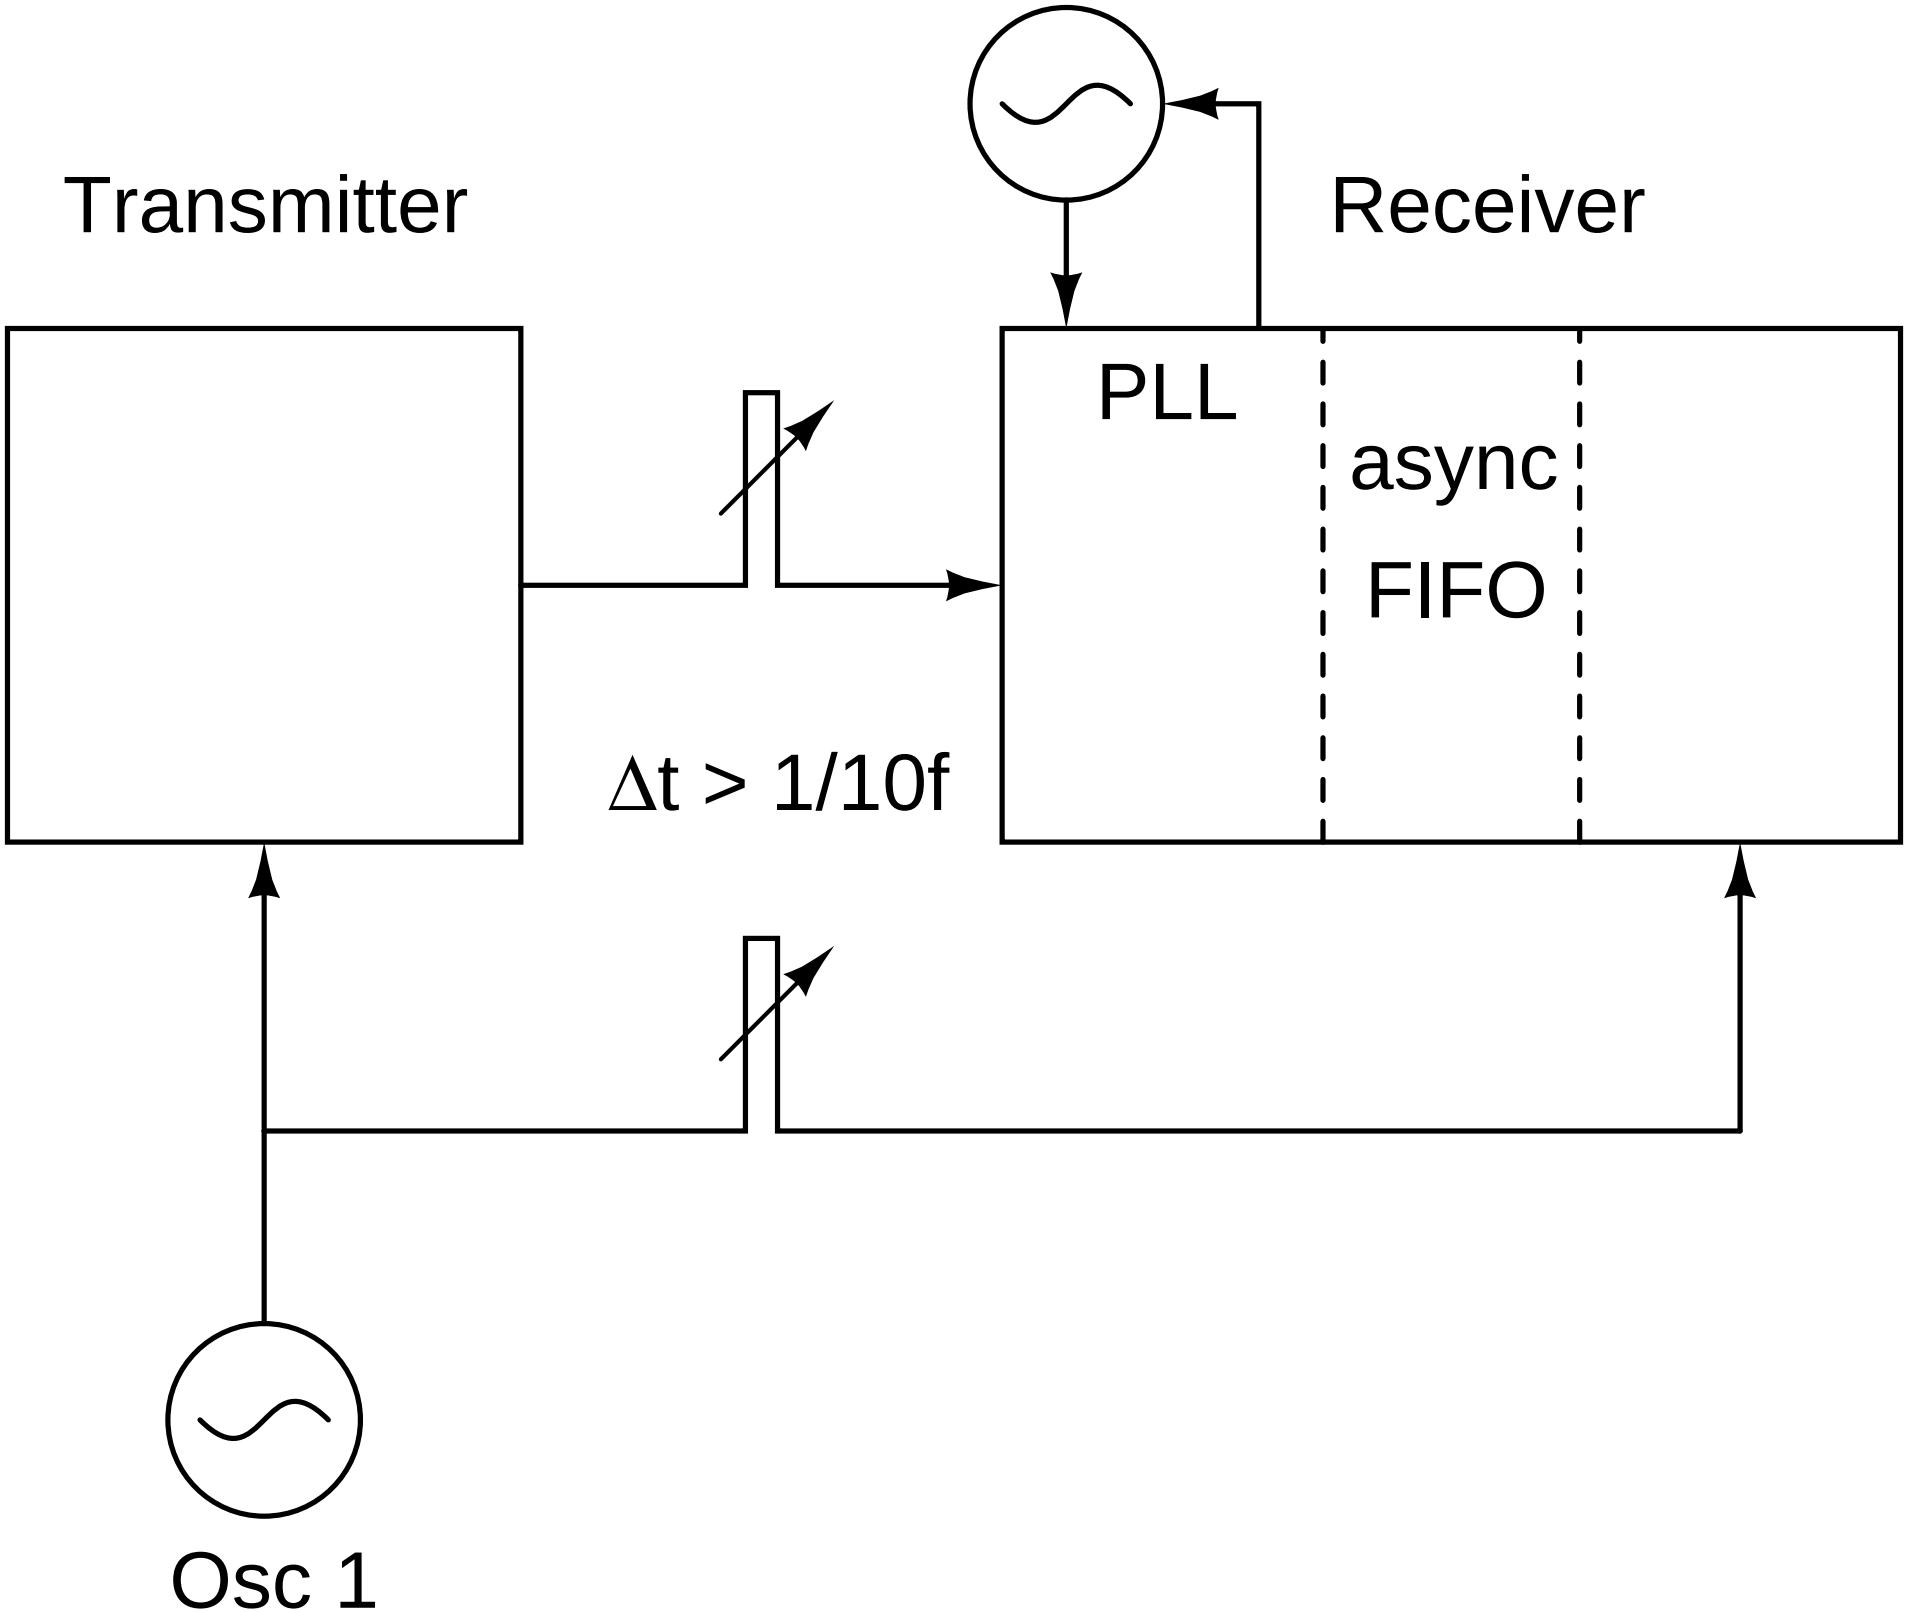
Transmitter
Receiver
PLL
async
FIFO
Δt > 1/10f
Osc 1

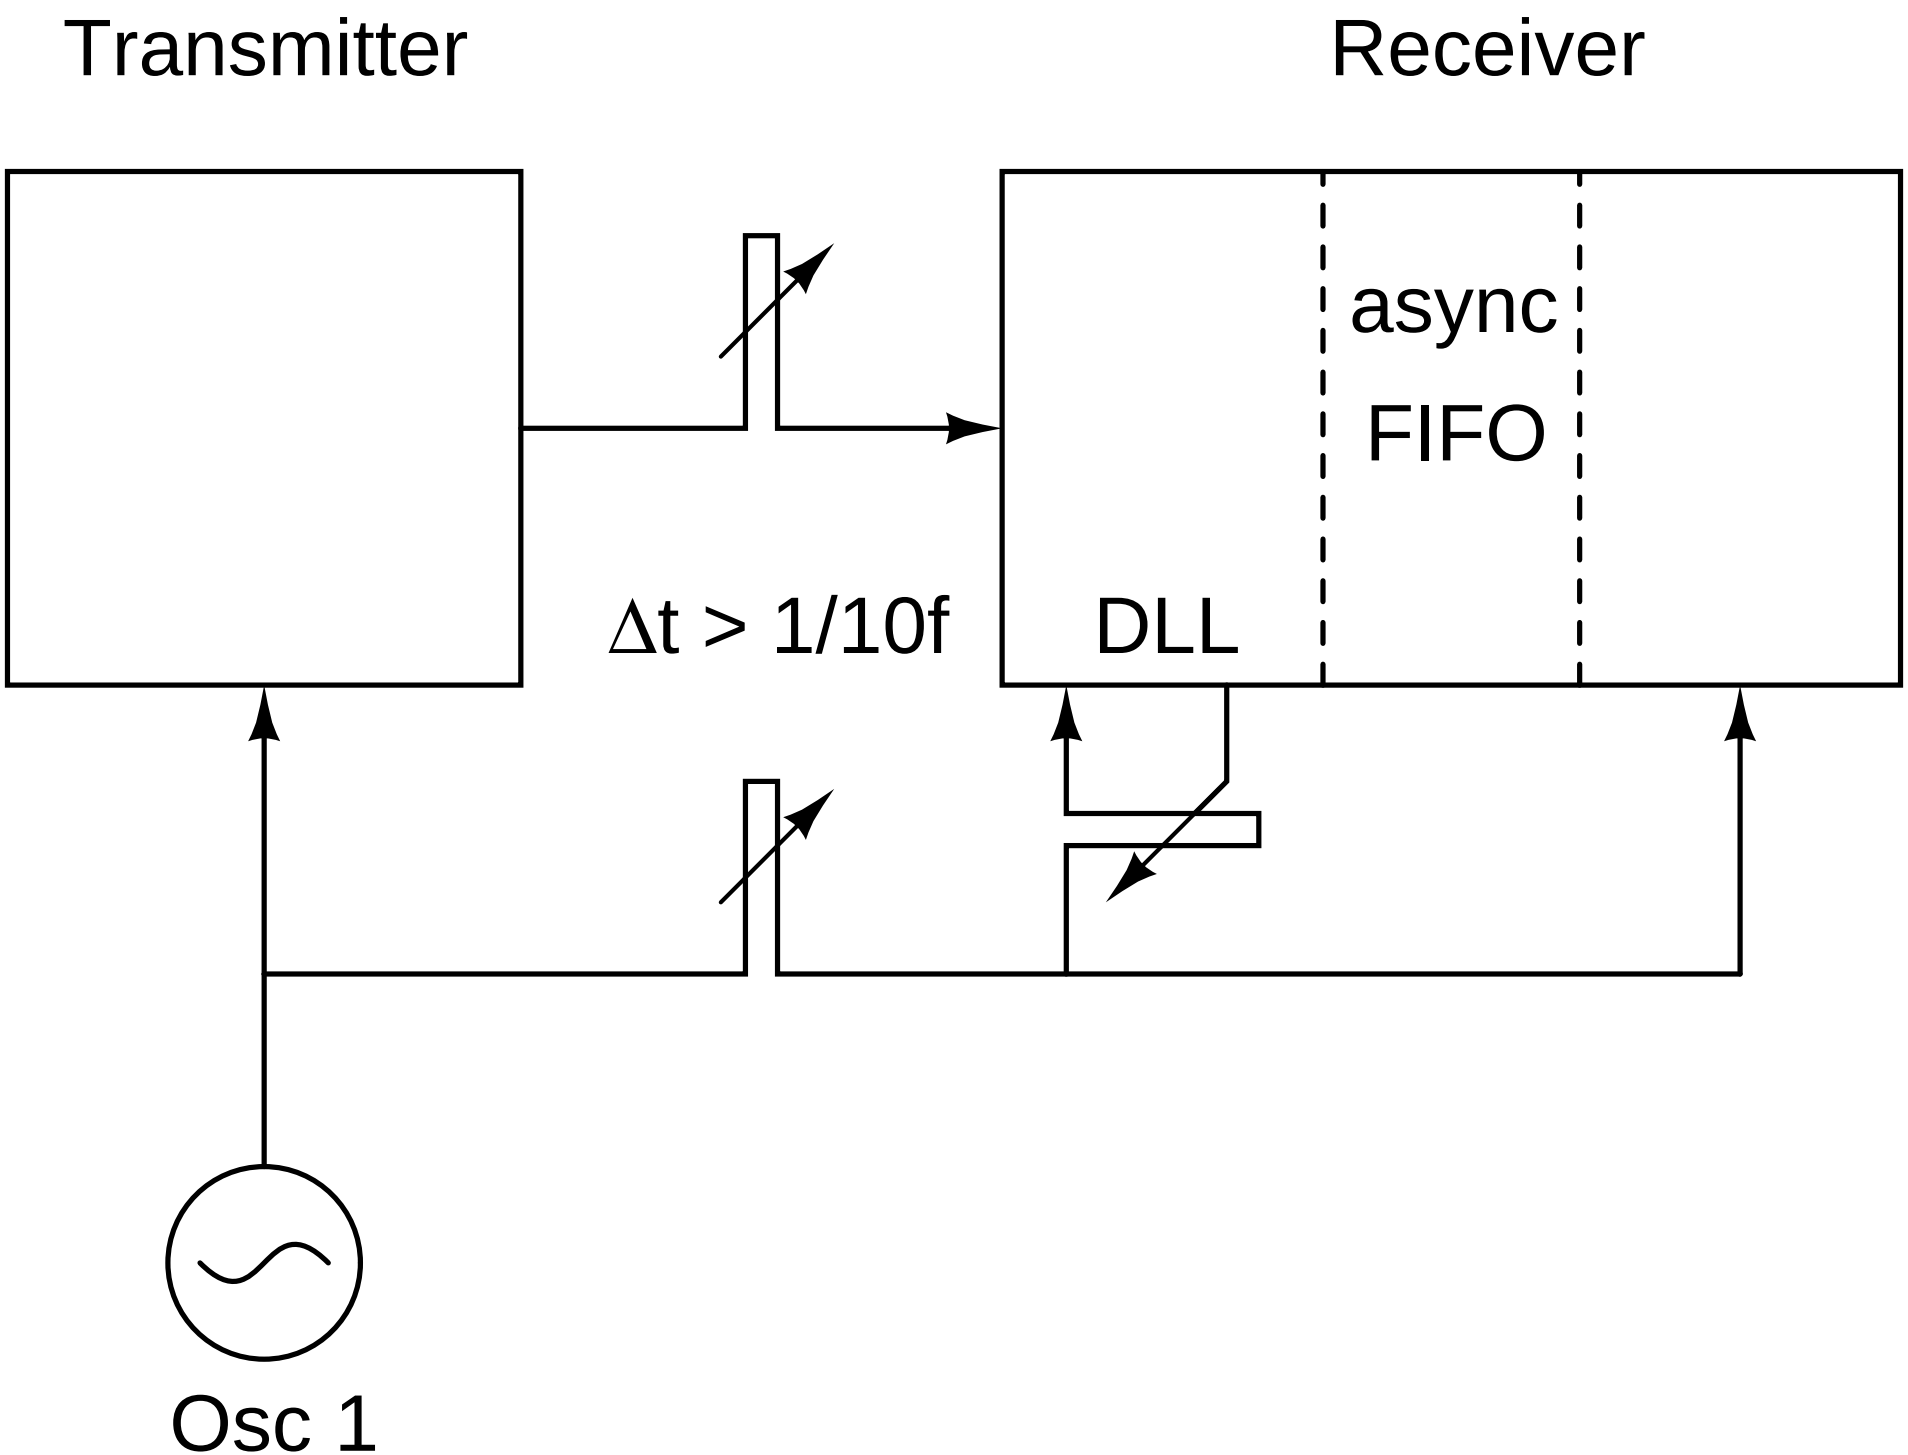
Transmitter
Receiver
async
FIFO
Δt > 1/10f
DLL
Osc 1

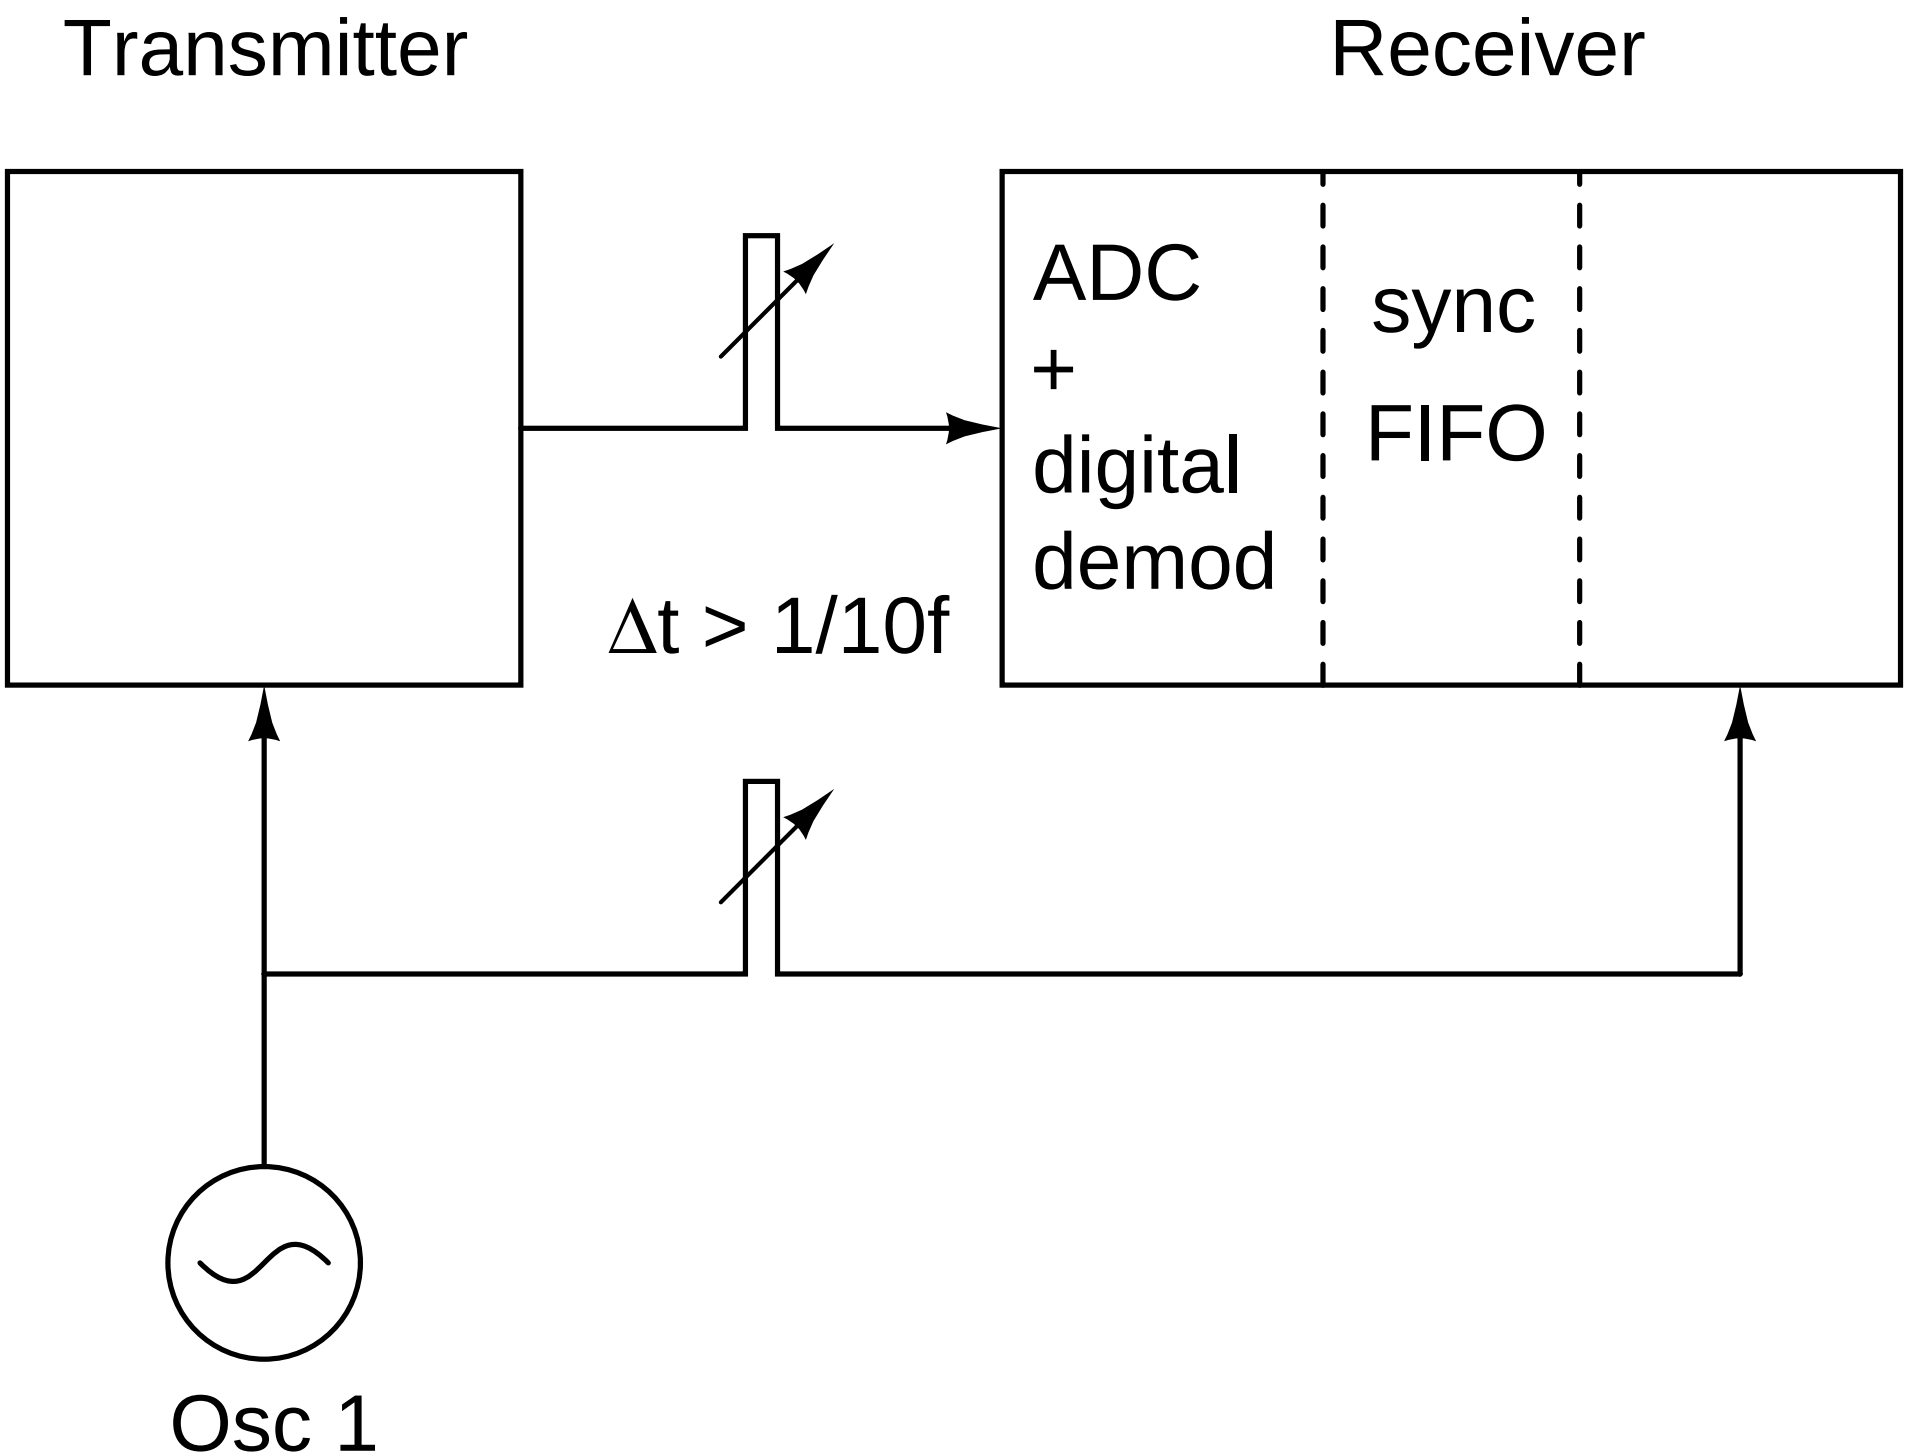
Transmitter
Receiver
ADC
+
digital
demod
sync
FIFO
Δt > 1/10f
Osc 1

Establishing absolute time at power-on

Network provides coarse time – microseconds, perhaps down to 20 ns

Looking for decent techniques to distribute fiducial accurate enough to find individual clock cycles, without adding yet another cable.

- PRN modulation of distributed clocks?
- Out of band signalling?

T. Rohlev wants to send all frequency and timing information, including the fiducial, over optical MGT links.

The problem with computers

Computers and software adds another layer of temporal indeterminism: Cache, interrupts, threads

It would take an incredible code coverage tool to evaluate all combinations of if statements and loop counts to see if DSP code "makes timing"

Implications for system simulation

Simulating and understanding a system with hundreds or thousands of incoherent crystal oscillators is statistical at best.

A believable simulation of a system with a known coherence relation between all the clocks is achieved in a time on the order of the LCM of the clock divisors.

Final Challenge

Actually make progress on this front, especially when it involves the larger accelerator community: diagnostics, timing, controls