Time-Interleaved ADCs, Past and Future

Ken Poulton
Agilent Laboratories

Outline

- Some History
- Reasons to Interleave
- Issues with Interleaving (with examples)
- Futures

Some History of Interleaving

- 1980 4-way 100 MSa/s 4 slices/chip CCD (T/H)
  Corcoran & Hornak
- 1980 4-way 4 MSa/s 4 slices/chip CMOS
  Hodges & Black
- 1987 4-way 1 GSa/s 1 slices/chip bipolar + GaAs
  Poulton, Corcoran & Hornak
- 1991 16-way 8 GSa/s 2 slices/chip bipolar
  Rush & Byrne
- 1997 4-way 8 GSa/s 2 slices/chip bipolar
  Poulton, Knudsen, Kerley, Kang, Tani, Cornish, VanGrouw
- 2002 32-way 4 GSa/s 32 slices/chip CMOS
  Poulton, Neff, Muto, Liu, Burstein, Heshami
- 2003 80-way 20 GSa/s 80 slices/chip CMOS
  Poulton, Neff, Setterberg, Wuppermann, Kopley, Jewett, Pernillo, Tan, Montijo
- 2008 16-way 24 GSa/s 64 slices/chip CMOS
  Schvan, Bach, Falt, Flemk, Gibbins, Greshishchev, Hamida, Pollex, Sitch, Wang, Wolczanski

- Trends:
  - Only modest on-chip interleaving until 2002
  - Massively-parallel interleaving introduced in 2002

Why Do Time Interleaving?

- Highest possible sample rate
  - When you max out the single-ADC sample rate in available technologies, interleaving is the way to further increase sample rate
- Reduce cost
  - Interleaving has moved multi-GSa/s ADCs from bipolar to CMOS
- Optimize power
  - CMOS ADCs are generally much more power-efficient
  - Can choose sample rate to optimize power efficiency

Why Not?

- Interleaving Errors
  - Calibration
- Extra complexity in HW, possibly SW and system

Issues and Example Solutions

- I'll describe issues for the design of interleaved ADCs, focusing mostly on on-chip interleaving
- I'll describe some example solutions, mostly from our own work

Core ADC Slice

- Objectives
  - Total power optimization – there is an optimum slice sample rate for a given technology ADC architecture
  - Area
  - Metastable error rate tends to limit sample rate of multi-stage ADCs
  - Calibration – more calibration can allow smaller-area devices with more mismatch

- Two examples:

Year	1997	2002
Process	25 GHz f_T Bipolar	0.35 um CMOS
Sample Rate	4 GSa/s, 7 bits	4 GSa/s, 8 bits
Slices	2x 2 GSa/s	32x 125 MSa/s
Slice Topology	Folded Flash, little calibration	Pipeline w/ per-stage calibration
Slice Power	3 W	50 mW
Chip Power	13 W	4.5 W

Types of Interleaving Errors

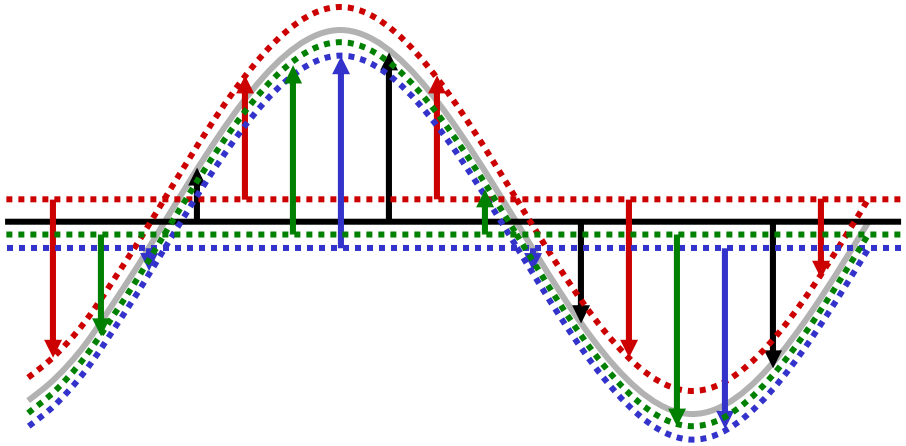

Offset Error
want <0.3% offsets for 8 bits

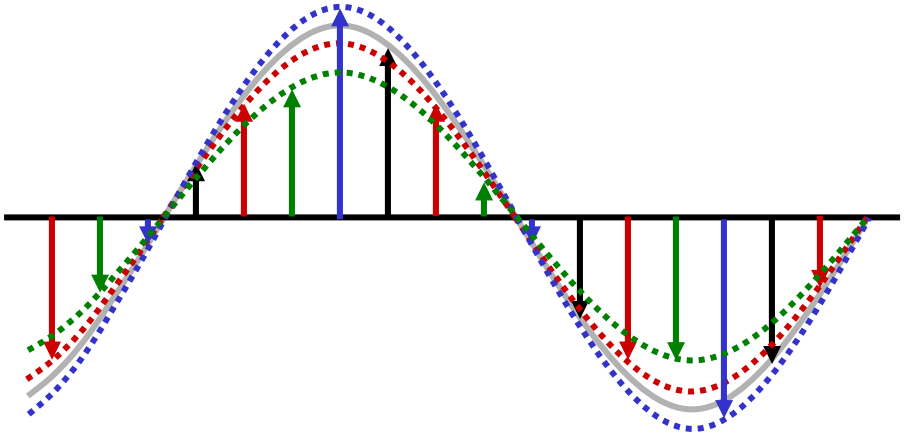

Gain Error
want <0.3% gain matching

- Errors due to mismatched paths
- Reconstructed waveform shows distortion.
- Spurs in the spectrum.

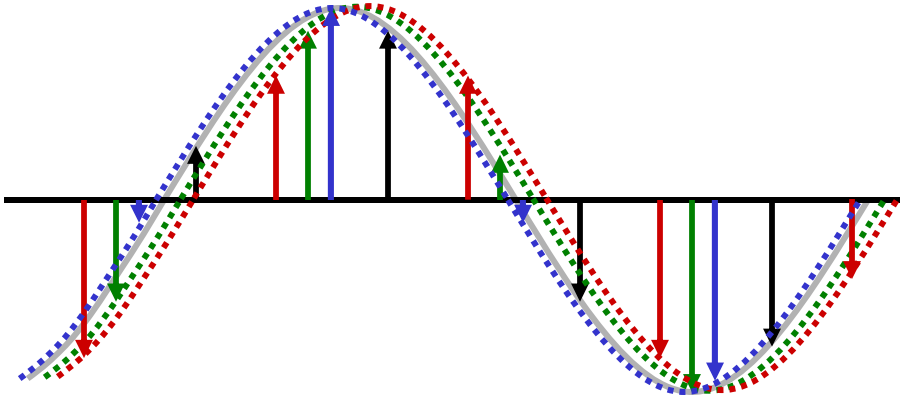

Sampling Instant Error
want <1 ps timing errors
for Fin of 1 GHz

Calibration Choices

- Calibration Approaches:
  - No Cal: Accurate by design
    - + Simplest in operation
    - -- Stringent matching requirements -> more area, higher power
    - -- Risk: any unforeseen mismatch needs a mask turn to fix
  - Analog Cal: Calibrate with adjustable offset, gain and delay DACs or trims
    - + Circuit sizes limited by SNR, not matching => much reduced power
    - + Maximize signal range valid in all slices
    - + Offset and gain adjusts can be low power
    - -- Delay adjusts add jitter and power
    - -- Risk: more analog circuits to design
  - Digital Cal: Calibrate with DSP
    - + Circuit sizes limited by SNR, not matching => much reduced power
    - + Fewer analog circuits than analog cal
    - -- Adds more quantization noise
    - -- Every adjust adds active data path power
      - +- Eventually, power of digital corrections will be lower than analog
    - -- More digital supply noise

More Calibration Choices

- On-chip vs off-chip calibration (coefficient computation)
  - Off Chip:
    - + lower chip cost
    - + lower design risk
    - -- much more HW and/or SW complexity for the ADC user
  - On Chip
    - + More generally applicable
    - -- Higher design cost and risk
- Foreground calibration vs. Background calibration
  - Foreground
    - + Independent of the input signal
    - -- ADC user needs to provide cal signals
    - -- Dead time during calibration
  - Background
    - + Least complexity for the ADC user
    - -- Places some requirements on the input signal

Example of Calibration Choices – 32-way & 80-way CMOS ADCs

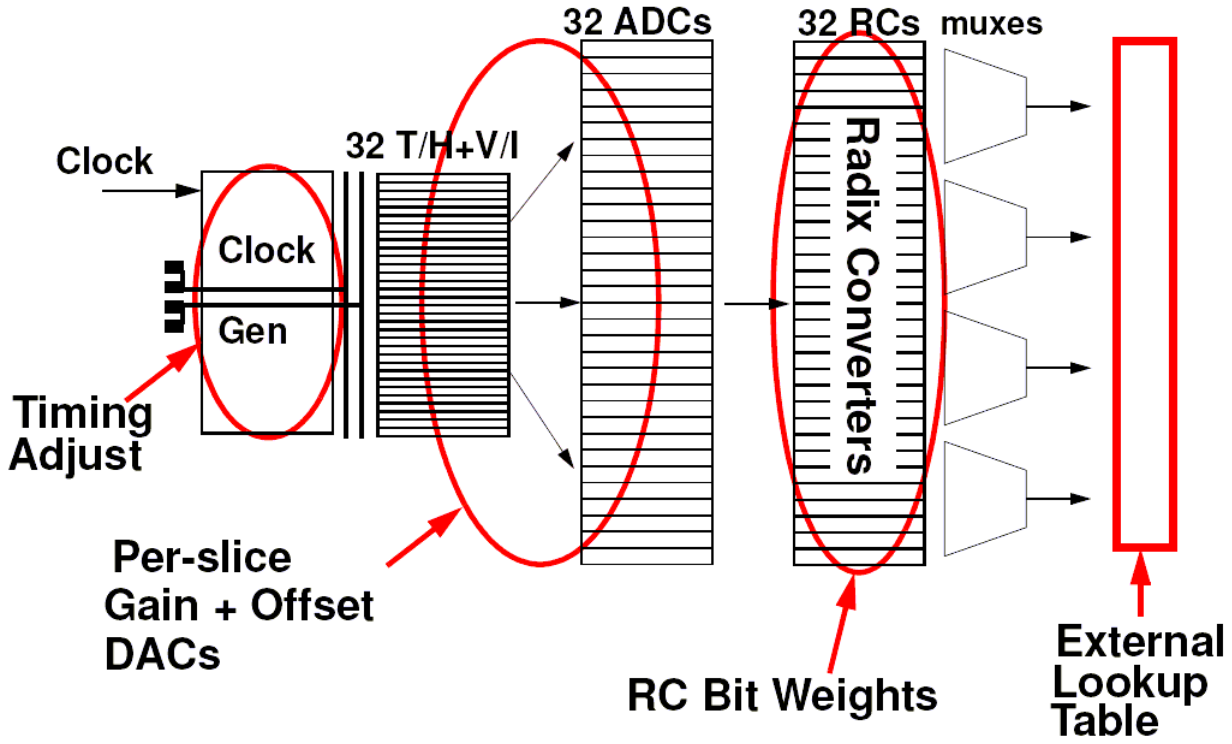

- Oscilloscope application
  - -> Foreground calibration
  - -> Off-chip calibration
- Non-binary-radix pipeline
  - -> Digital gain and offset for fine cal
- Maximize tolerated offset
  - -> Analog cal for coarse gain, offset
- 0.35, 0.18 um process
  - -> Adjustable delays for timing cal
  - -> External lookup for linearity correction

Sample Timing Solutions: Single Samplers

+ Relaxes slice clock accuracy requirements

-- Requires a full-speed clock and sampler

- Two topologies:

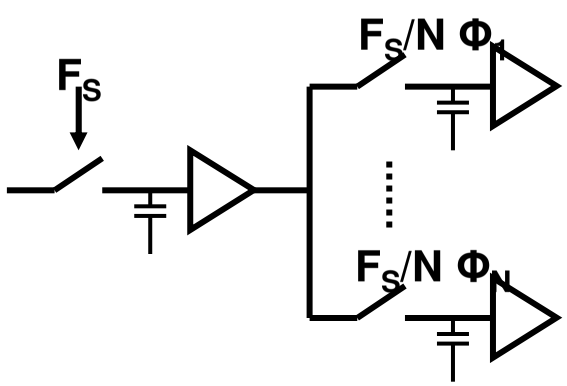

Two-Rank T/H system

+ best T/H BW

Corcoran, et al, ISSCC 1987

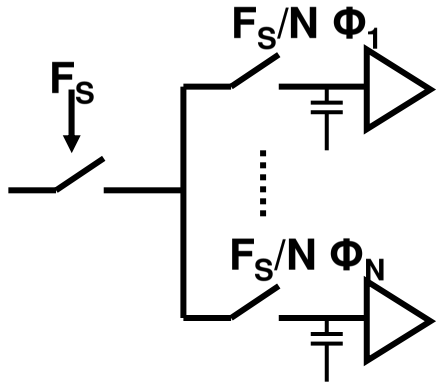

Two-Switch T/H system

+ lower noise and power

Gupta, et al, ISSCC 2007

Sample Timing Solutions: Interleaved Samplers

- Interleaved Samplers:
  - + Nothing needs to run at F_S
  - + No critical settling between T/Hs
  - -- Large C_{IN}
  - -- Hard to make many equal clock and V_{IN} paths
  - -- Timing correction usually required

- Timing Correction Methods:
  - Adjustable delays in clock paths
    - + Corrects the actual sampling instants
    - -- Drift
    - -- Delays add jitter and power
    - *Poulton, et al, ISSCC 2002*

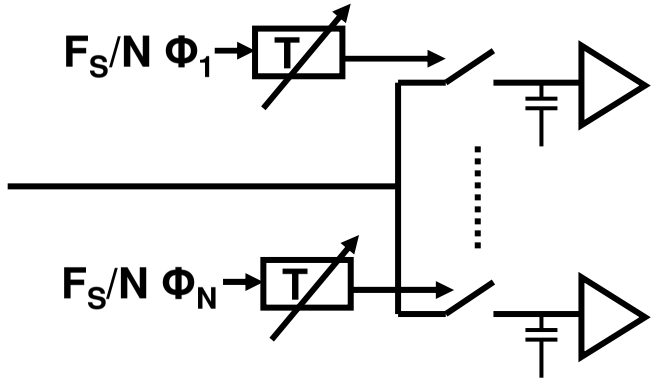

  - Delay adjustment by DSP after ADC
    - + Shorter timing paths to drift
    - -- Timing adjust filters cause gain vs. Fin errors near Nyquist
    - -- Power and digital noise generation
    - *Jamal, et al, ISSCC 2002*

Accuracy of Clock Generation

- Single Front-end Sampler
  - One full-accuracy clock (0.1-1 ps rms jitter)
  - N reduced-accuracy clocks (1-10 ps rms jitter and alignment)
    - Second rank requires fairly accurate clocks due to analog settling effects
- Interleaved front-end samplers
  - N Full-accuracy clocks (0.1-1 ps rms jitter **and** alignment)
    - Generating lots of very-low-jitter clocks can cost a lot of power
    - Simple DLL for N clocks, constant jitter spec:
      - **N stages, with power scaling as N to maintain jitter**
      - **=> Power scales as N^2**
- *For more detail, see: Robert Neff, "Clock Challenges for GHz ADCs" In the Clock Synthesis Design Forum (F7) on Thursday*

Example: Clock Generation for 32-way 4-GSa/s ADC

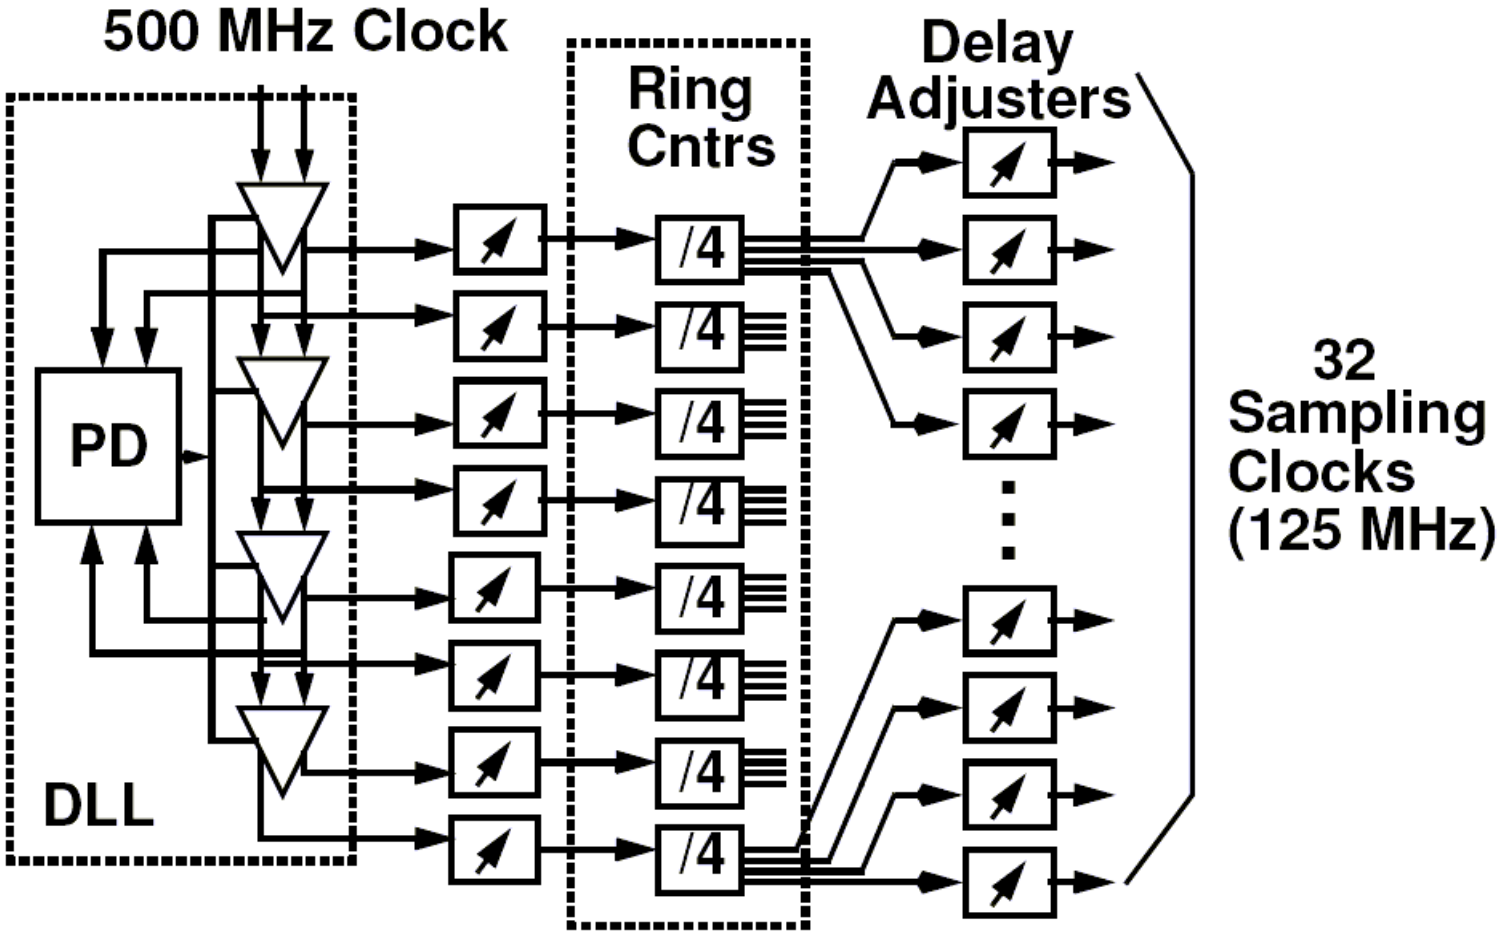

4 GSa/s: ~ 1 ps thermal jitter

Timing misalignments: before cal: 10 ps rms, after cal: 0.8 ps rms

Clock and Signal Distribution

- Multiple samplers requires distributing clocks
- Tradeoff between need for proximity of the samplers and physical size of the ADC slices
- Both clock and analog signal distribution can take a lot of area and power
- Example: 80-slice ADC:
  - samplers: 1.9 mm wide
  - ADCs: 9 mm
  - Memory 13 mm

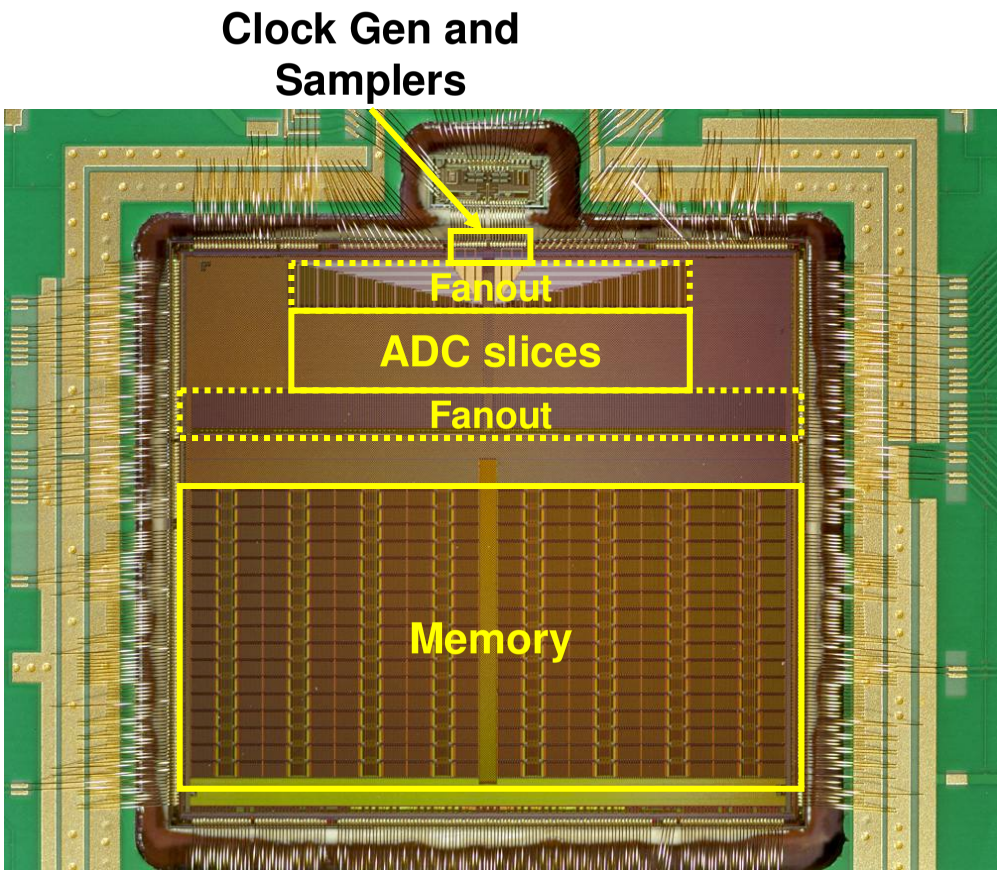

Overhead Factors for Interleaved ADCs

- Power: Even with a power-efficient ADC slice, you can still have a lot of chip power
  - Example: 32-way, 4-GSa/s ADC
    - 32 ADC slices 1.6 W
    - Samplers + clocks 0.9 W
    - Multiplexors and outputs 1.7 W
    - Total 4.2 W

- Design Overhead:
  - More optimization variables
  - More circuits that interact
  - More clocks
  - V_{IN} and clock distribution
  - Data collection and multiplexing
  - Calibration
  - Optional: Memory and/or DSP

Historical Data: Relative Design Time vs. Interleave Factor

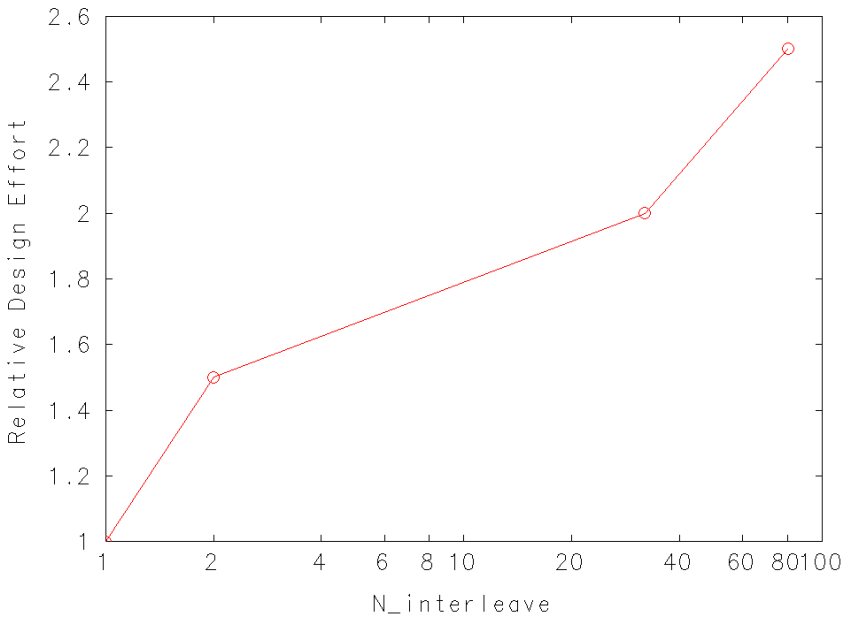

Where Does All That Data Go?

- Totally application-dependent
- Off Chip
  - Wide parallel buses (~1 Gb/s per pin) or faster serial ports
  - I/O power can easily be as much as the ADC power
- On-chip memory
  - Memory power can easily be as much as the ADC power
  - When you need more memory, moving to the next process is a huge redesign for the ADC
- On-chip DSP
  - A big win if the DSP reduces the amount of data that must be stored or transported off-chip
  - Otherwise, it can be a problem to tie together the ADC process and the DSP process due to different power scaling and development times.

Trends

- Much more power-efficient core slices
  - Power-Efficiency FOM has dropped by ~10x since 2002
  - Pipelines and SARs dominate
- More background cal
  - Making the transition from academia to industry
  - Partly enabled by targeting specific apps like communication channels with somewhat-known signals
- Interleaving is breaking out of “just for scopes”
  - Lots of papers at ISSCC
  - A startup whose major focus is interleaving existing ADCs
  - Massive interleaving for 12-25 Gb/s wireline links
  - Becoming “just another ADC topology”

What is the Future for Interleaved ADCs?

- When will all ADCs be interleaved?

  Never!

  - Design complexity
  - Overhead for power and design time

- When will time-interleaved ADCs become common?
  - It's happening right now

Acknowledgements

- John Corcoran and Robert Neff of Agilent Labs