

56Gs/s ADC Enabling 100GbE

Ian Dedic, Chief Engineer, Fujitsu Microelectronics Europe

OFC2010 Invited Paper, Digital Transmission Systems

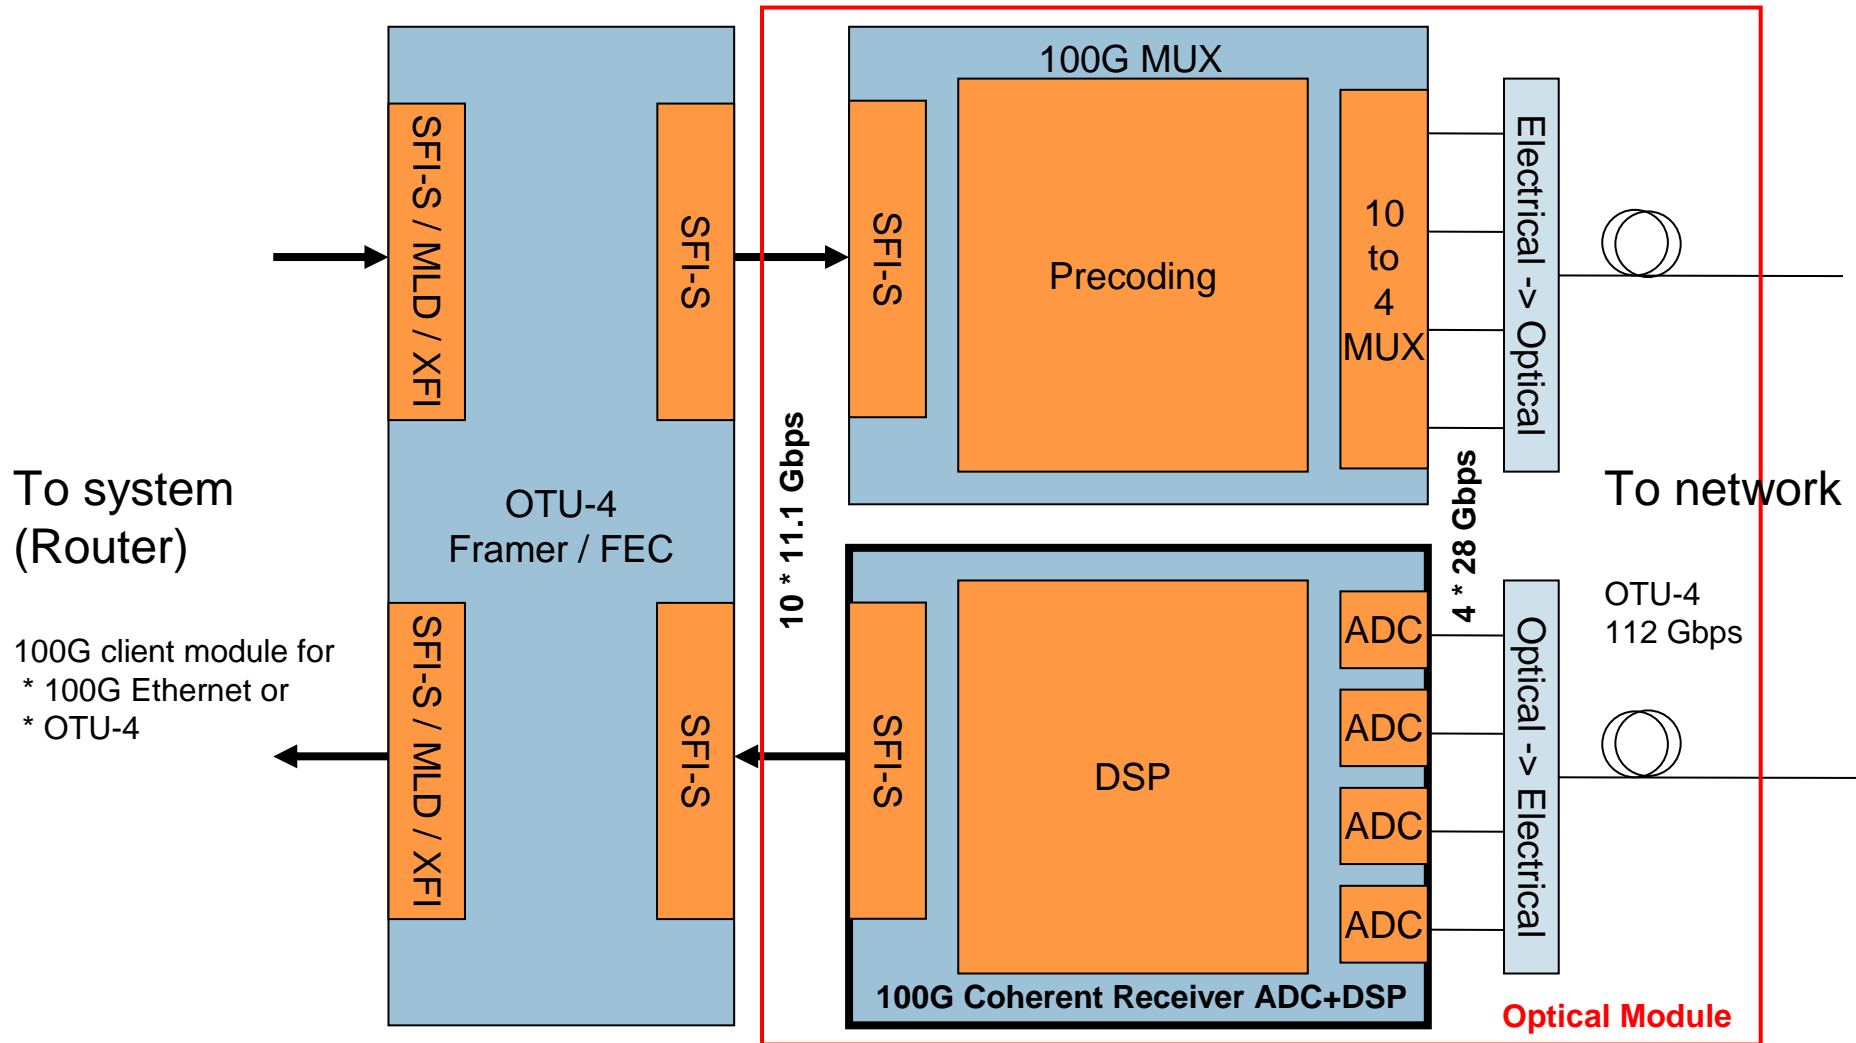
56Gs/s ADC : not just an ADC design problem



- 100G coherent receiver
- Why single-chip CMOS?
- So what is so difficult?
- CHAIS ADC
- DSP and integration
- Package and PCB
- Testing
- Examples
- Future challenges



100G Coherent Optical Transponder



100G coherent receiver requirements

- 4-channel 56Gs/s ADC, minimum 6b resolution
 - DP-QPSK receiver with H/V channels, I/Q conversion (OIF standard)
 - 2x oversampling minimizes impact on performance
 - 56Gs/s needed, >60Gs/s with soft FEC
 - In early 2009, predicted availability for such an ADC was in 2013
 - Major obstacle to target rollout date for 100G networks

- DSP to remove transmission channel imperfections and recover data
 - Dispersion, cross-polarization, time-varying channel, clock recovery
 - ~12 TeraOperations/sec (TOPS) for 40Gb/s → ~30 TOPS for 100Gb/s ?

- Power consumption and thermal constraints
 - OIF target for complete coherent transponder ~70W
 - Severe challenge for both ADC and DSP design
 - Total power for ADC+DSP needs to be < 50W
 - Environmental restrictions make problem even more difficult than it first appears

Why single-chip CMOS for 100G?

- Massive data bandwidth between ADC and DSP
 - 4-channel 6-8b 56Gs/s ADC means 1.3-1.8Tb/s of data at interface
 - Getting this from one chip to another costs power and chip area
 - 10G SERDES link ~100mW/channel → 3-4W per ADC

- Critical performance factor is power efficiency, not just speed
 - Discrete ADC dissipating >10W each are difficult to use
 - Total power dissipation is too high (>100W)
 - Skew management/calibration nightmare (especially over temperature/lifetime)
 - Higher intrinsic manufacturing cost and lower yield for a hybrid

- Single-chip CMOS solution is the “Holy Grail”
 - Integrate on ASIC with >50M gates (limited by power dissipation)
 - Leverage CMOS technology advances to drive down power and cost
 - Single-chip allows integration of calibration and channel deskew/matching
 - ADC and DAC get faster and lower power at the same rate as digital – ideally ☺

ADC+DSP -- so what is so difficult ?



- ADC is the biggest circuit design problem
 - Ultra-high speed, low noise and jitter, low power consumption – all same time
 - Conventional techniques cannot easily deliver required performance

- Digital-analogue noise coupling
 - Sampler/clock jitter <100fs on same chip as DSP with >100A current spikes

- Wide bandwidth (>15GHz), good S11 (up to >30GHz), low theta-jc
 - Sampler, package, PCB design all very challenging

- DSP design is out-of-the-ordinary (tens of TeraOPS)
 - Extremely power-efficient → use massive parallelism, not GHz clocks

- Test
 - Performance verification challenges limits of test equipment
 - Need at-speed performance verification in production, not just functional testing

The ADC problem

- Wideband low-noise sampler + demultiplexer + interleaved ADC array
 - Smaller CMOS geometries → higher speed → worse mismatch and noise

- Single 56Gs/s track/hold very difficult due to extreme speed
 - <9ps to acquire, <9ps to transfer to following interleaved T/H stages

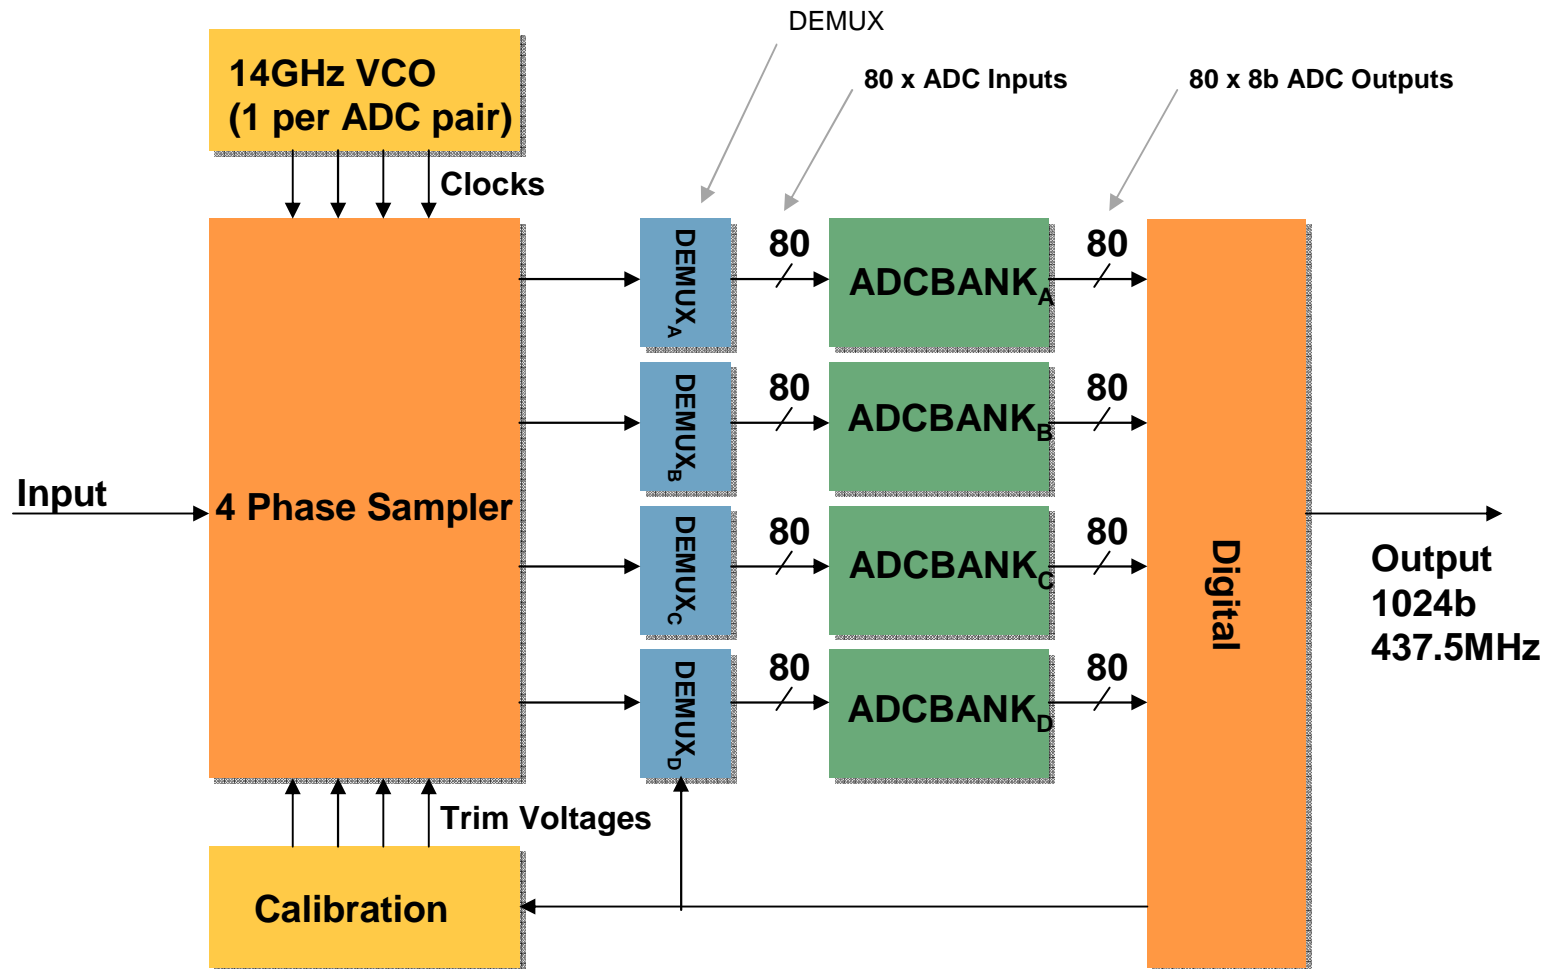
- Interleaved track/hold (e.g. 4-channel 14Gs/s) also very difficult
 - Signal/clock delays must match to <<1ps – how do you measure this?

- Noise, mismatch and power of cascaded circuits all adds up
 - Multiple sampling capacitors, buffers, switches, demultiplexers...

- Interleaved ADC back-end is not so difficult (but only in comparison!)
 - Design for best power and area efficiency rather than highest speed
 - Interleave as many as necessary to achieve required sampling rate
 - 8 x 175Ms/s 8b SAR ADCs fit underneath 1 solder bump → 45Gs/s per sq mm 😊

A 56Gs/s CMOS ADC solution

CHARGE-mode Interleaved Sampler (CHAIS)



[this page intentionally left blank]

- 56Gs/s with 8 bit resolution
 - Extra resolution allows some digital AGC after ADC instead of in OFE
 - Non-ideal INL/DNL has negligible impact on performance
 - 63Gs/s for 40nm ADC to allow for higher overhead soft decision FEC
 - Higher clock rate plus soft FEC needs 40nm to meet same power budget
- 16GHz bandwidth
 - Closely controlled, can be extended using digital equalization if required
- ENOB>5.7 for -6dBFS sinewave input
 - Similar power to 100G OFE output signal peaking at full-scale (PAR=9dB)
- ENOB almost constant with input frequency
 - On-chip PLL jitter is 30fs rms, THD < -40dBc at 15GHz
 - <0.2ENOB variation from 1GHz to 15GHz
- Power consumption 2W per ADC in 65nm, 9W for complete 4-channel RX
 - 5W for 63Gs/s 40nm 4-channel RX – power scales better than digital !

- Interleaved paths need accurate delay/gain/offset matching
 - 100fs inter-sample skew generates -40dBc distortion for 16GHz input

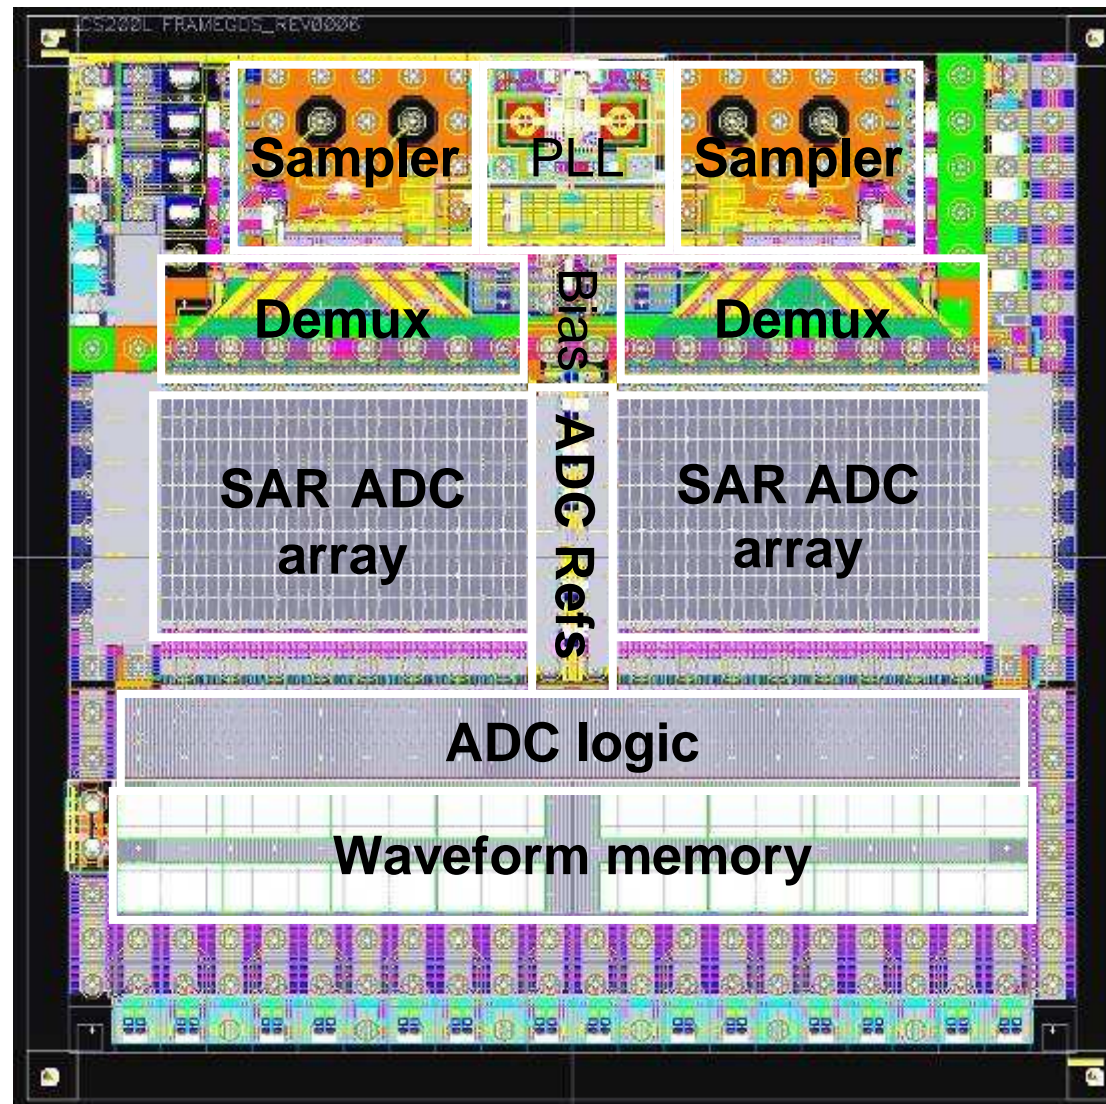
- Clock and signal mismatches get worse as process shrinks
 - Smaller gates (for higher speed/lower power) have worse matching

- Can't take ADC offline to recalibrate (*excuse me while I turn the Internet off...*)

- How to measure and correct these errors during operation?
 - Measurement accuracy problem (*Quis custodiet ipsos custodes?*)
 - Algorithm complexity (FFTs at these rates just for calibration are not desirable)

- CHAIS architecture can be calibrated using simple analysis of output data
 - Average errors calculated in background during operation (no FFTs, DSP, filters...)
 - Analogue trim coefficients calculated (low CPU load) to drive trim DACs
 - Interval between calibrations 0.1~10 seconds (just to track temperature changes)

Dual ADC layout (4mm x 4mm test chip)



The DSP problem

- Digital design tools (and designers) *really* don't like this type of DSP
 - The tools synthesize circuits, then worry about how to connect them up
 - 90% of power dissipation in interconnect, not gates
 - Massive data buses (4k bits at ADC outputs) → massive interconnect problem

- Partitioning into usable size blocks may be more difficult than it appears
 - Tools don't like doing flat designs with tens of millions of gates (turn-around time)
 - “OK, lets split that big DSP block into two and add some pipelining”
 - “Well, about this 16k bit wide data bus you've just introduced...” ☹

- Better system/architecture tools for this type of design are needed
 - Should really optimize the data flow, then shovel the circuits in underneath...
 - Designers' brains (and system-level design tools) don't always think this way ☹

- Very high average gate activity has several consequences – see later...

Noise coupling – round up the usual suspects...

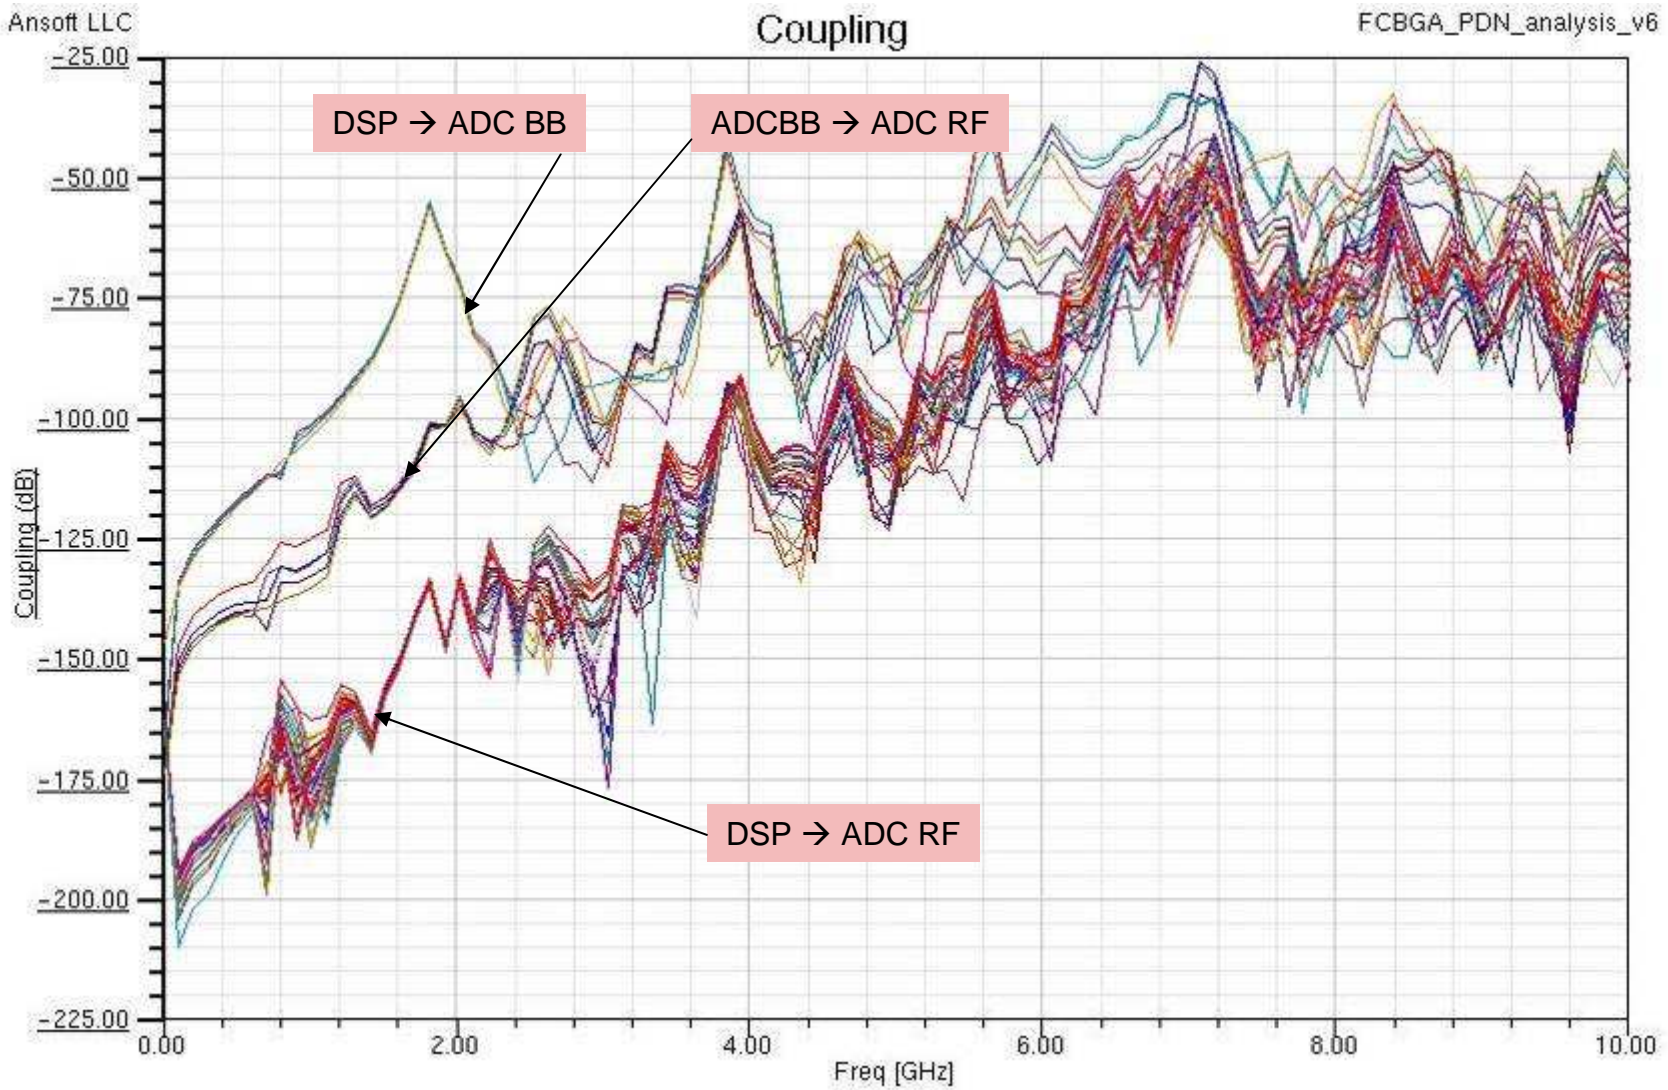


- Reduce aggressor (DSP logic) noise generation
 - Use intentional skew of clock timing within each block and between blocks
 - Reduces peak current and spreads out in time → >10x lower di/dt
 - Lots of on-chip (~300nF) and ultra-low-inductance (~4pH) in-package decoupling
 - Even with this can expect >100mVpp supply ripple (see later...)

- Increase victim (ADC analogue) immunity
 - Fewest possible noise/jitter sensitive circuits, all fully differential
 - Lots of on-chip (~100nF) and low-inductance in-package decoupling
 - For 100fs jitter and delay sensitivity of 1ns/V, need 100uV of supply noise

- Improve victim-aggressor isolation
 - Avoid low-resistance epi substrate (milliohm connection for substrate noise)
 - Build “nested walls” of isolation with most sensitive circuits in the middle
 - 100mV digital noise + 100uV analogue noise → 60dB isolation
 - Noise coupling is wideband (DC to GHz) not narrowband like RF – very challenging...

Power Distribution Network simulation



On-chip decoupling and effect of noise



- Coherent receiver has unusually high gate activity levels
 - Clock frequency is relatively low (~500MHz) to maximize power efficiency
 - Large fraction of gates toggle every clock cycle
 - Power dissipation is relatively high (~50W) ← *this means “between 20W and 100W”* 😊
 - Power per mm² is higher than normal (because gate activity is so high)
 - If your digital gate count is too big to fit in the chip, your power will be far too high...

- All charge to sustain supply voltage on clock edge comes from on-chip
 - For 40A@1.2V, 500MHz clock, 200mVpp ripple → 400nF total capacitance
 - Capacitance of logic 100nF → 300nF on-chip decoupling is needed !
 - Same power at 1GHz would only need 100nF, at 2GHz logic capacitance is enough
 - CPU, GPU, FPGA, standard ASIC don't have this problem (lower activity levels)

- Measured results show negligible degradation in analogue performance
 - Compare noise floor with digital logic off and in test mode (~40W extra power)
 - **ENOB > 5.7, SINAD degrades by <1dB with 9GHz sinewave input** 😊

- 1mm pitch FCBGA, >1000 pins, 19 internal layers
 - Use similar package for test chips as typical ASIC to get same performance
 - Low-loss high-TCE LTCC (12ppm/C) for improved second-level reliability

- Multiple power/ground regions and shields for noise isolation

- Ultra-low-inductance internal decoupling for supplies and bias/reference
 - Multiple interleaved VDD/VSS planes connect chip to multi-terminal decouplers
 - Noise dealt with inside package → simplifies customer PCB design


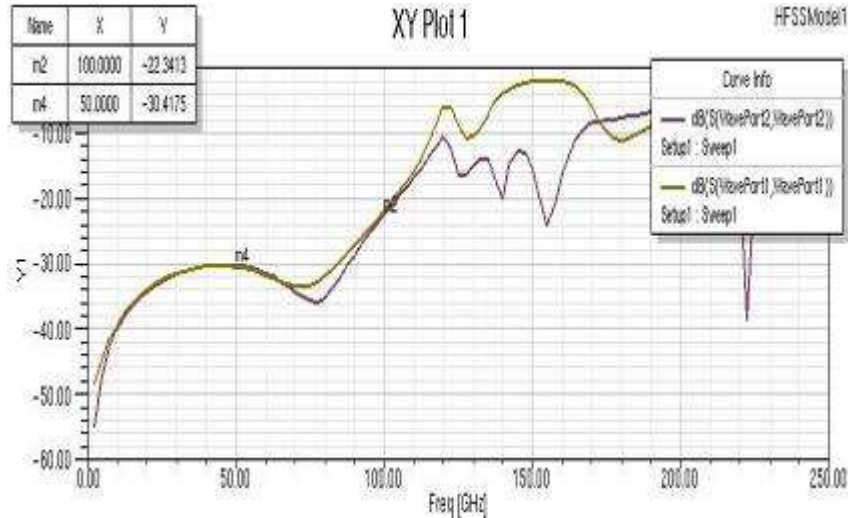
- Coaxial via and waveguide structures, <1dB loss at 20GHz
 - Ground planes removed above signal balls to reduce capacitance
 - Increased layer spacing (wider tracks) to reduce losses

- Optimized launch to G-S-G coplanar waveguide on low-loss PCB
 - Balls on row inside signal pins removed to reduce capacitance, grounds cut back
 - Outer PCB layers use MEW Megtron 6 (very low loss, lead-free multilayer compatible)


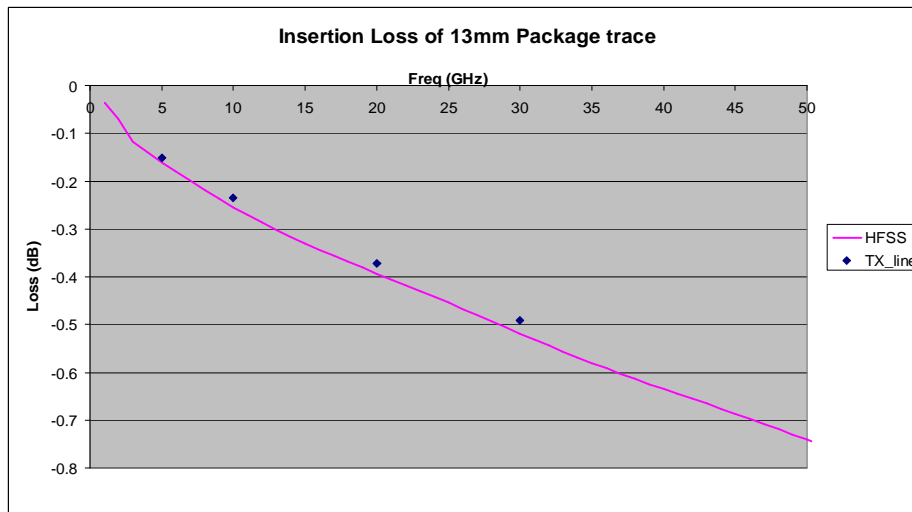
Chip-to-package transition : $S_{11} < -20\text{dB}$ to $\sim 100\text{GHz}$



ADC Bump-Pkg Transition



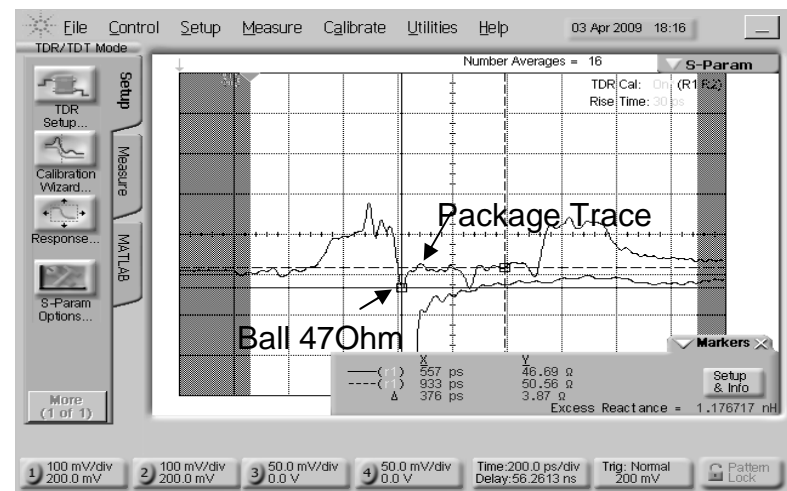
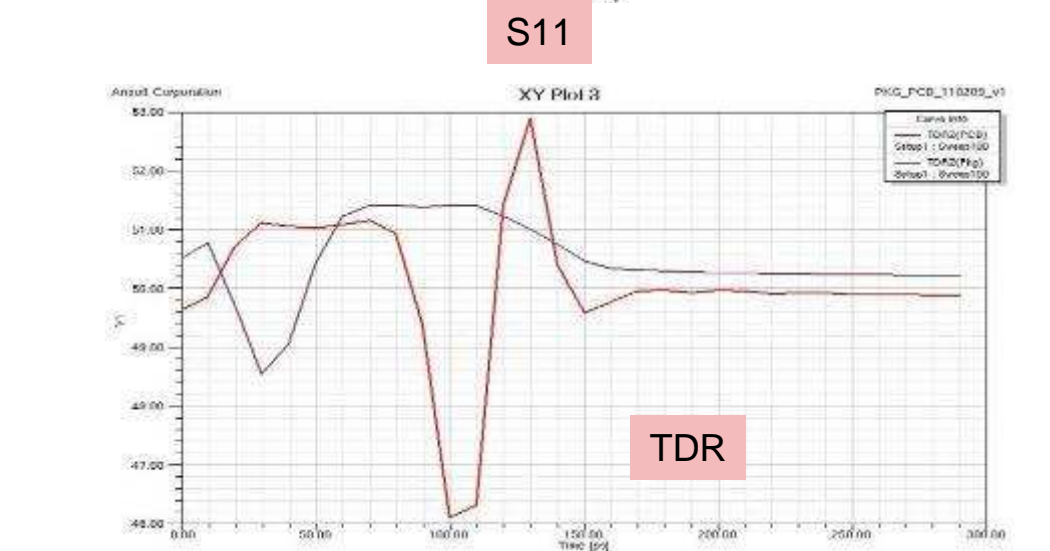
Signal line loss $< 0.02\text{dB}/\text{GHz}$



Package-to-PCB transition : S11 < -20dB to ~50GHz

ADC Pkg-PCB Transition

HFSS sim model




- Power dissipation seems not so bad compared to high-power GPU/CPU
 - But thermal environment is much more restrictive – no massive heatsink/fan

- Available airflow and heatsink size/aspect ratio are restricted
 - Incoming air typically 55C maximum (even higher if a fan fails)
 - Flat heatsinks have poor heat spreading performance (hot-spot over package)
 - Heatpipes (e.g. **NanoSpreader™**) may be needed to distribute heat

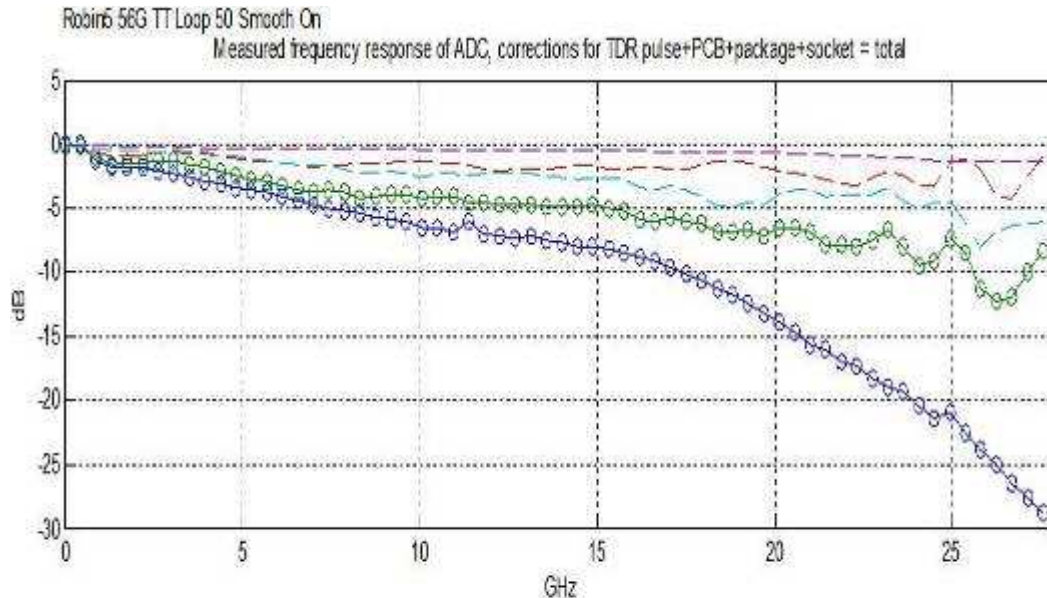
- Optical front-ends cannot tolerate high temperatures
 - Rated to ~75C → can't be placed too close to high power ADC+DSP
 - Low-loss PCB needed

- Cannot afford large temperature difference between die and package lid
 - Most of thermal resistance needs to be allocated to heatsink and system
 - Need very low θ_{jc} : $<0.2\text{C/W}$ keeps die $<10\text{C}$ hotter than package lid

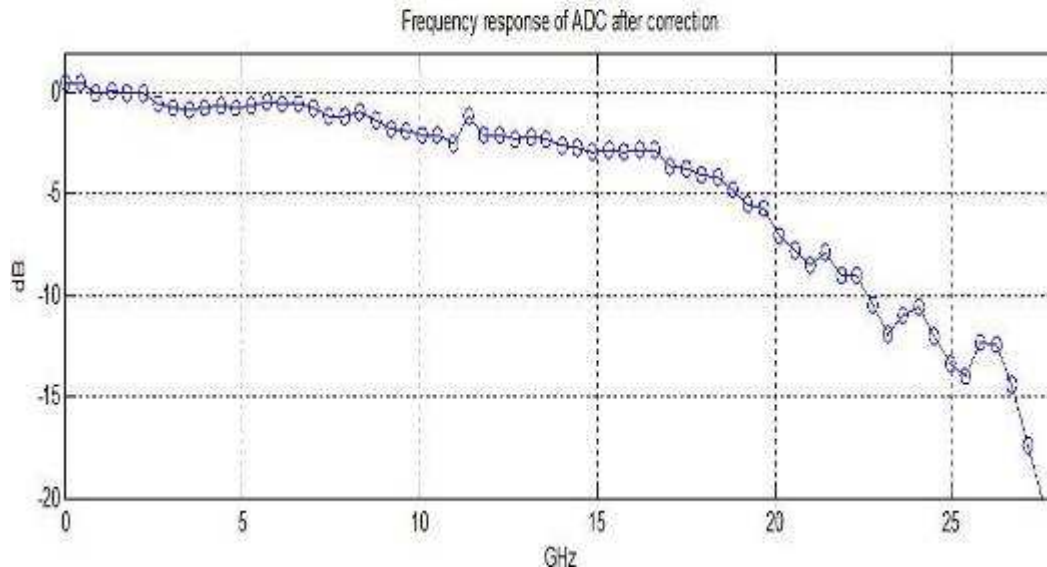
BATBOARD and ROBIN



Frequency Response (test setup and ADC)



- Frequency response of test setup
 - TDR step (measured)
 - Batboard PCB (measured)
 - ENIG not Ag finish (Ni is lossy!)
 - Socket (estimated -1dB @ 20 GHz)
- Test setup loss bigger than ADC loss!



- Corrected ADC frequency response
 - accurate measurements are not easy ☹
- ADC -3dB bandwidth 16GHz
 - close to simulation and specification


- Need proper performance verification, not just functionality
 - Increased confidence that chip actually meets design specifications
 - Make chip self-testing as far as possible and do at-speed performance tests

- Test ADC ENOB using sinewave input
 - Sampled data stored in on-chip RAMs then read out and analysed (ENOB)
 - Filtered low-noise SMT VCO signal source
 - Average amplitude and slew rate similar to 28Gb/s signal from OFE
 - Production test guarantees ENOB specification limit is met for real signal

- Test CEI-11G outputs by looping back into 56Gs/s ADC inputs
 - We have a “free” 4-channel 56Gs/s scope with ~100fs jitter, so use it 😊
 - Can equalize out test setup losses digitally after ADC (like RX equalizer in CDR)
 - 5 samples per bit gives complete waveform analysis on all TX channels
 - Full-speed measurement of eye opening and jitter possible in production


Example of 100G coherent receiver

- Architecture: Single CMOS die
- Technology: 65nm CMOS
- Interconnect: 12 layer metal
- Die size: 15 mm x 15 mm
- Package: FCBGA-1148



- Analogue macros: 4 channel 56 Gs/s ADC + multi channel CEI-11G TX
- ADC power : 9W (4 channels)

Example HiTCE LTCC substrate and chip

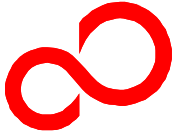


A “gold-plated” package solution...



- Advanced package technology originally developed for server CPU
 - It's nice to have in-house packaging ☺
- High-reliability Hi-TCE LTCC package
 - >2000 thermal cycles
- Metal TIM and gold-plated AlSiC lid
 - Very strong reliable chip-lid bond
 - All-metal path for heat transfer
- Very low thermal resistance
 - $\Theta_{jc} < 0.2\text{C/W}$ (JEDEC)
- Package structure and design changed for better noise isolation

- **What obstacles are there to progress beyond 100Gb/s?**
- **Sampler noise/bandwidth/interleave skew/clock jitter**
 - Can be solved using new CMOS techniques instead of exotic technology
 - CHAIS sampler/demux/ADC is capable of >100Gs/s even in 65nm
 - 112Gs/s or 126Gs/s for 400Gb/s 16QAM perfectly feasible, about 2W/ch in 40nm
- **Input bandwidth increase (30GHz?) and S11 improvement (up to 50GHz?)**
 - FBGA package changes may be needed to optimize design for very high frequencies
 - Smaller ball pitch conflicts with second-level reliability and PCB issues
- **Power consumption – this is really a DSP issue, not an ADC one**
 - ADC is 9W/4ch (56Gs/s 65nm), scales with technology (5W/4ch for 63Gs/s 40nm)
 - DSP power is several times ADC power, especially with more complex systems
 - Power increase (complexity) is outrunning power savings (process shrink)
- **The ADC is no longer the limiting factor ☺**



FUJITSU