



SOLARFLARE
COMMUNICATIONS

How SolarFlare Communications Broke
the 10Gbps on UTP Barrier



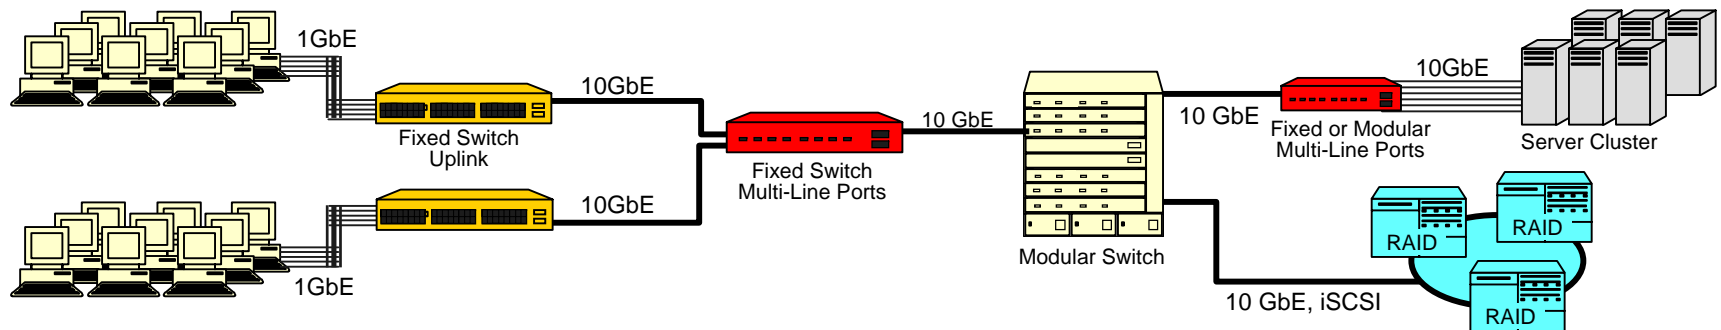
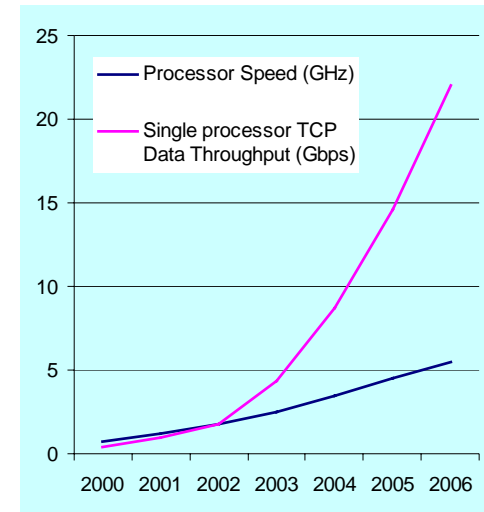
SolarFlare Communications

- Headquarters: Irvine, CA
- Product Focus: 10Gbps UTP Ethernet chip set
- Fabless business model with foundry CMOS process
- Key IP: DSP algorithm deriving its roots from NASA deep space communications
- Technology first demonstrated in March 2004



10G Ethernet – The Underlying Need

- Server Clusters and Distributed Computing
- High-End Computer I/O bottlenecks
- Storage over Internet Protocol – iSCSI
- Aggregation of 1000BASE-T on Clients
- What's missing? – Cheap, big pipes



Source: Sun & Intel IEEE Contributions, Dell 1000BASE-T White Paper, Forrester Research



Why UTP?

It's not *just* about running on installed cable

Optical Fiber and Modules

- Terminations are very expensive and require skilled labor to install
- Optical modules necessarily involve the mechanical assembly of many technologies: VCSELs, PIN diodes, laser drivers, transimpedance amps and SerDes chips.
- Optical modules are, by construction, one per port
- Optical modules are designed to operate at only one data rate and only purchased in applications where that data rate is needed


UTP and 10GBASE-T PHYs

- RJ45 is cheap, plastic, installed in field by any IT manager
- UTP PHYs are implemented in vanilla CMOS and are on a roadmap to single chip integration – lower COGs by construction
- UTP PHYs are capable of multi-port-on-a-chip implementations as lithography progresses – dramatically lowering price per port
- UTP PHYs are rate adaptive – making them attractive for PC LOM adoption where they are sold as “future proofing”



What's wrong with CX-4?

- Short reach
 - Standardized to 15m; Some solutions up to 30m
 - The problem: Sweet Spot of Data Center applications is between 40m and 70m
 - Because of patch panel connections
- Expensive Cabling
 - \$200 for 15m link with connectors (compare to \$15 for same link on Cat6)
- Thick cumbersome cable and large connectors
- Can not be terminated in the field



15m CX-4: \$200

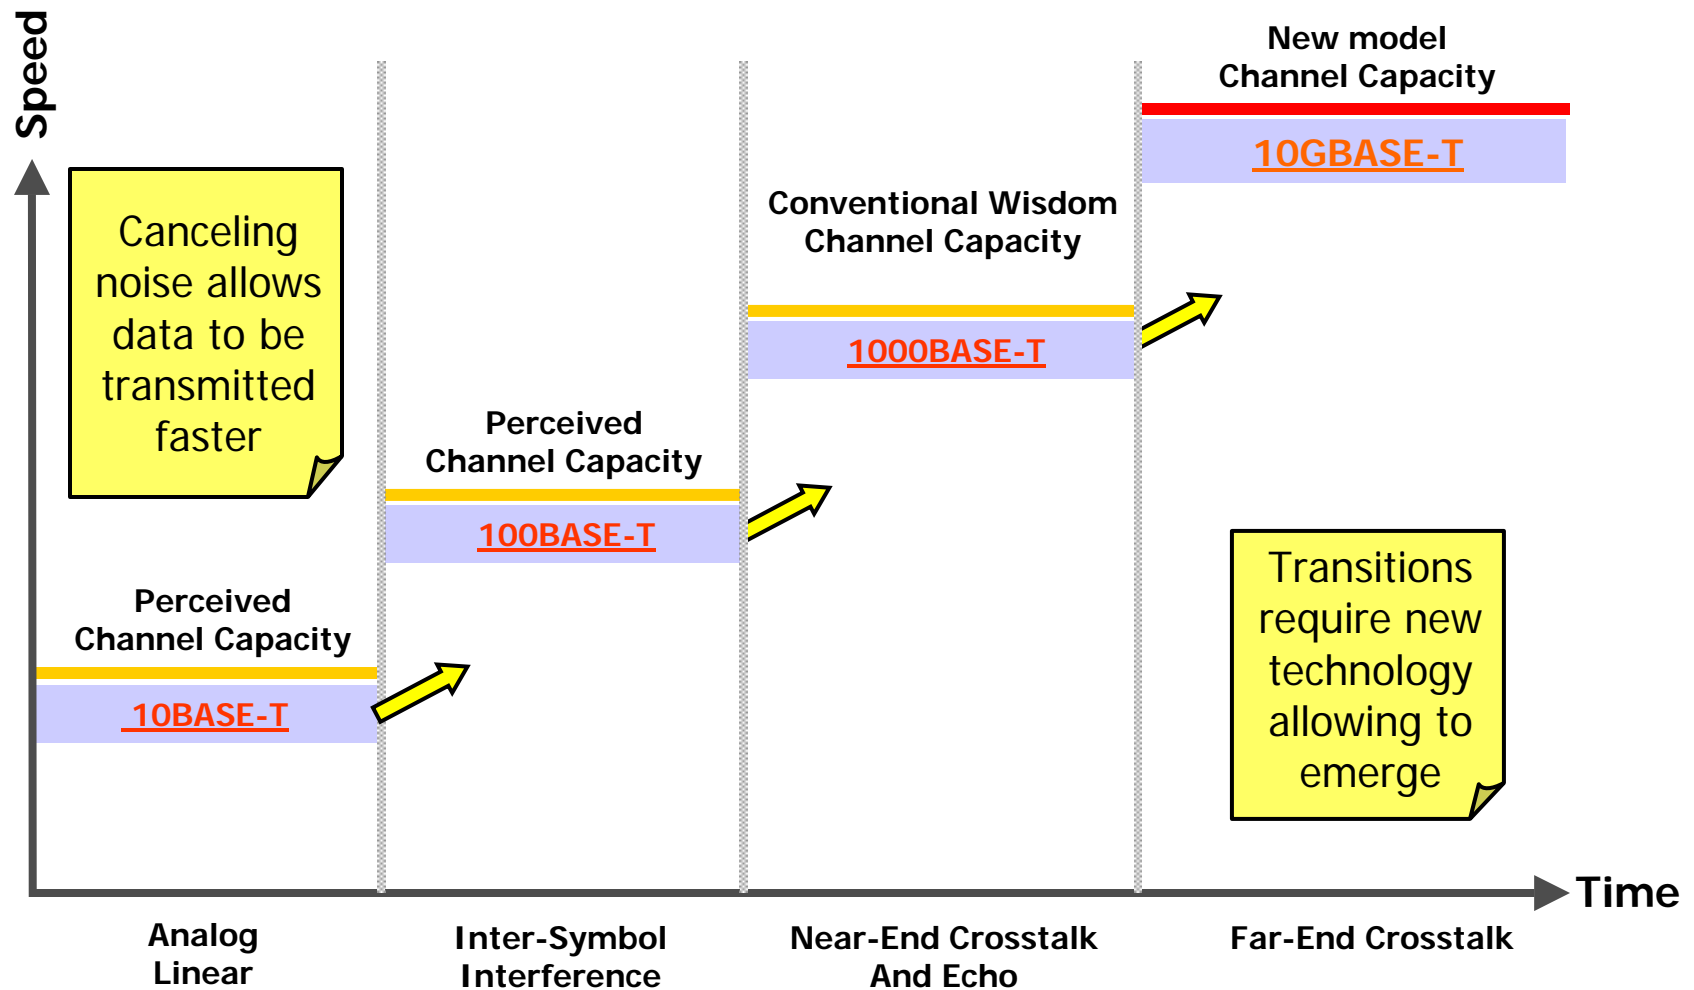


15m Cat6: \$15

Source: computercablestore.com



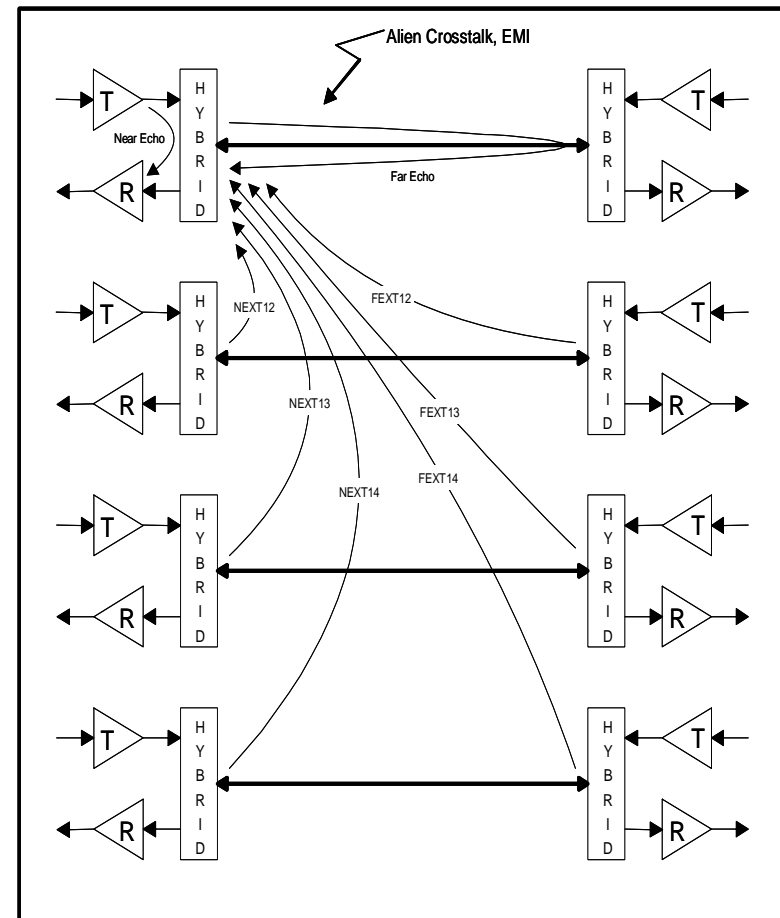
Technology Behind Ethernet Evolution





Why is 10G on UTP so difficult?

- Very small receive signals swamped by...
- Many sources of noise:
 - Far-End Echo
 - Near-End Crosstalk (NEXT)
 - Far-End Crosstalk (FEXT)
 - Inter-symbol interference (ISI)
 - Electromagnetic Interference (EMI)





Design Philosophy

- Evolutionary build on 1000BASE-T
 - Revolutionary complexity/performance enhancers
- Need to enhance treatment of media impairments
 - NEXT mitigation (less noise)
 - FEXT mitigation (less noise)
 - EMI friendly/Spectral efficiency (more bits/ baud)
- Best complexity / performance tradeoff while being Standards compatible

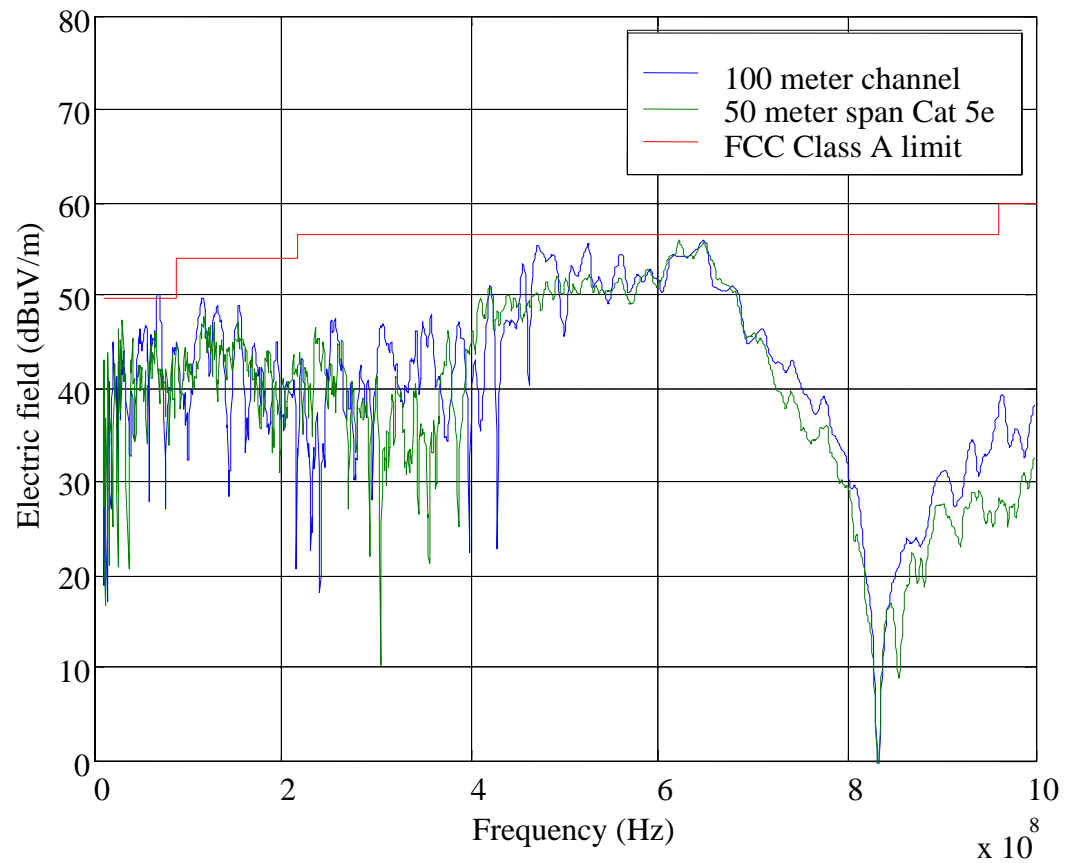


Comparison w/1000BASE-T

1000-BASE-T	10GbE Solution [UTP]
Multilevel coded PAM signaling (2-bits/symbol)	Multilevel coded PAM signaling (3-bits/symbol)
5-level with trellis code across pairs	10-level with trellis code across pairs
Full duplex echo-cancelled transmission	Full duplex echo-cancelled transmission
125 Mbaud, ~80 MHz used bandwidth	833 Mbaud, ~400 MHz used bandwidth
Moderate NEXT cancellation	High-Performance NEXT cancellation
No specified FEXT cancellation	High-Performance FEXT cancellation

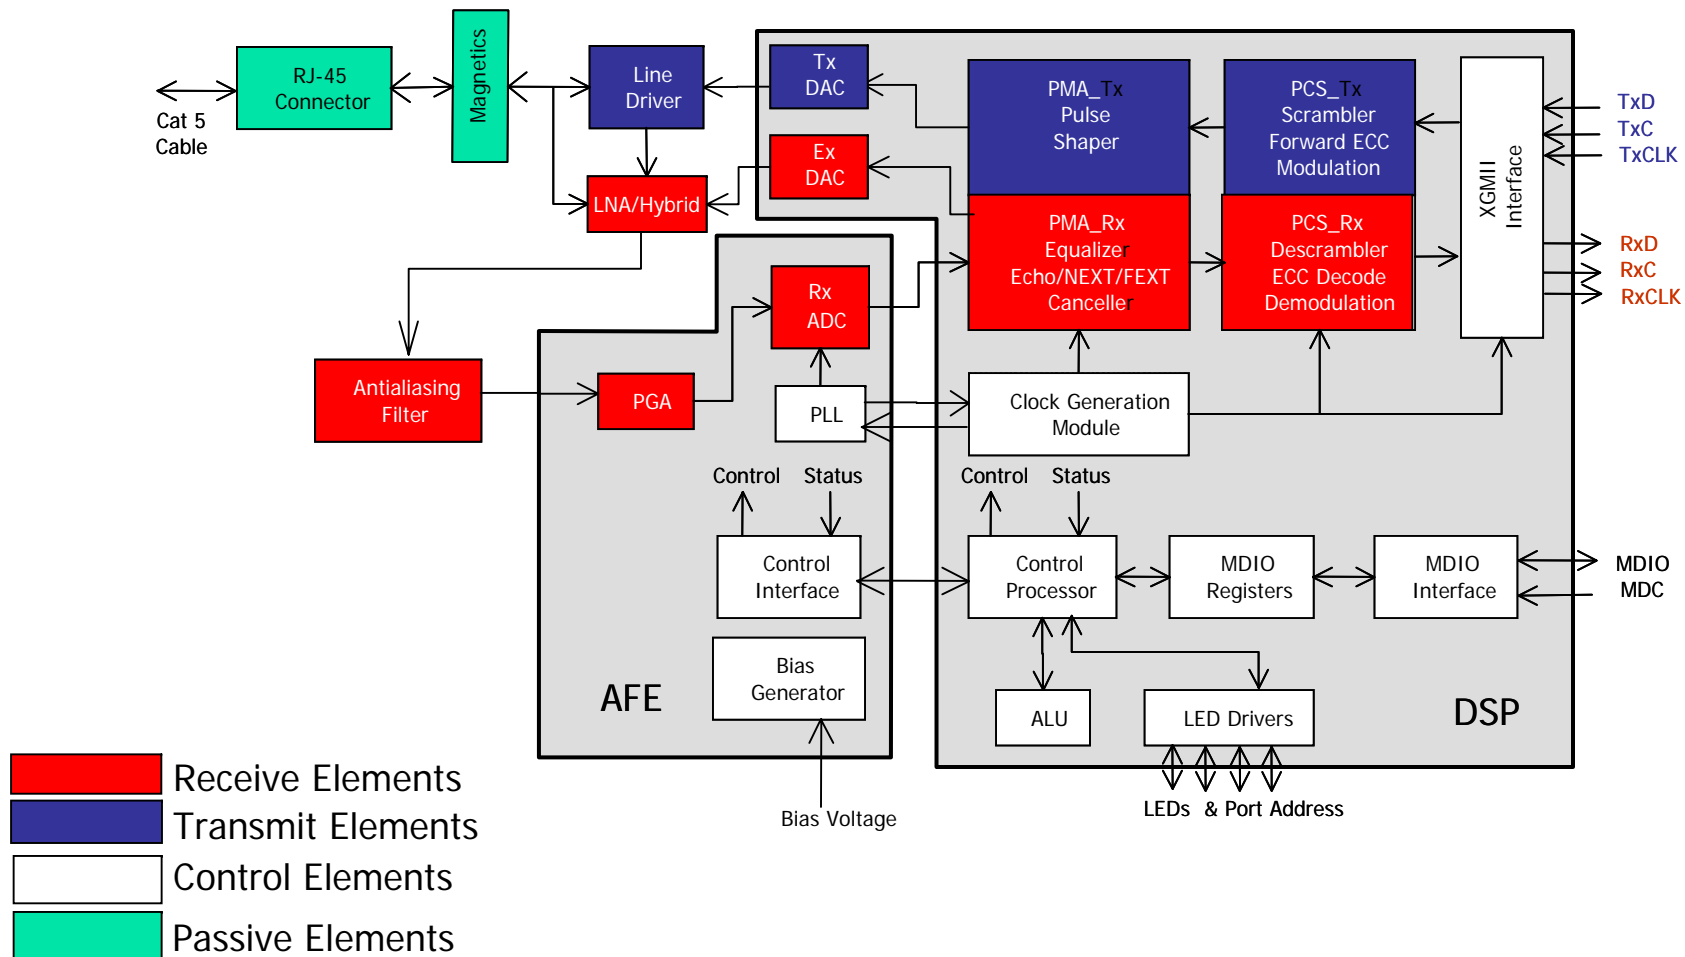


Measured Radiated Emissions






Chipset Block Diagram



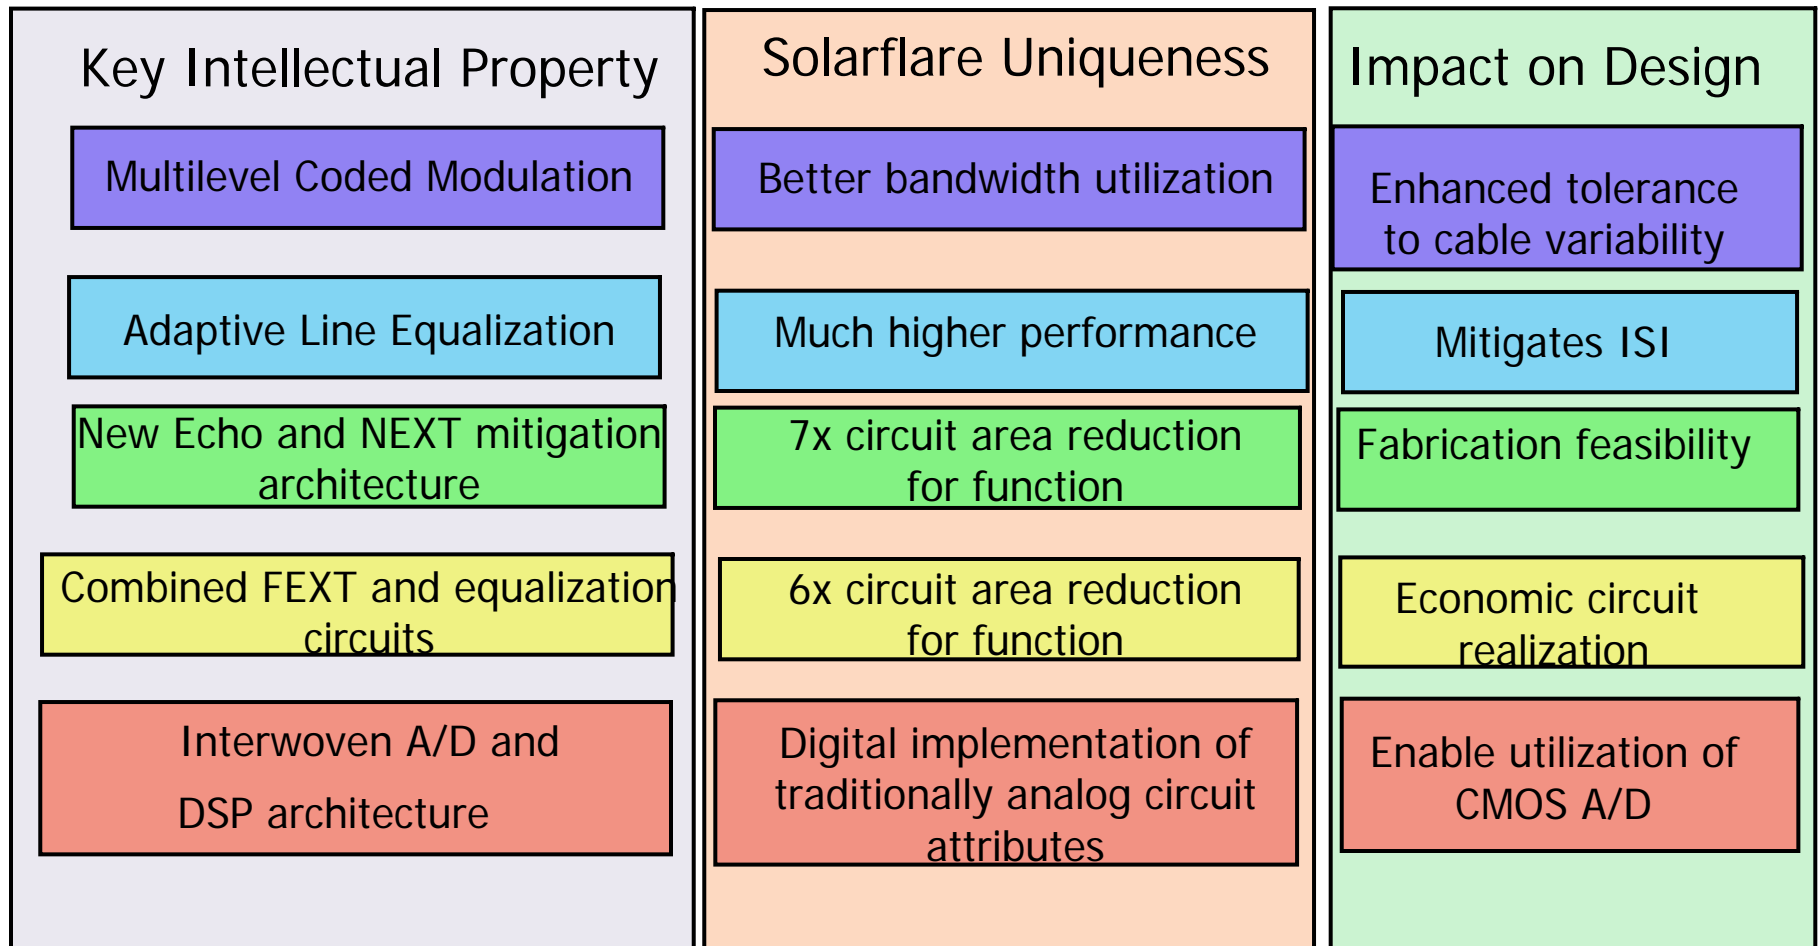


Noise Cancellation





SolarFlare IP Platform






Chipset Details

DSP:

- 0.13u CMOS
- 25 x 25 mm 575 pin BGA package
- 5.4 million gates + 2 Mbits memory



AFE:

- 0.18u CMOS
- 23 X 23 mm 473 pin BGA package
- 2.76 million transistors (0.7 million eq. Gates)

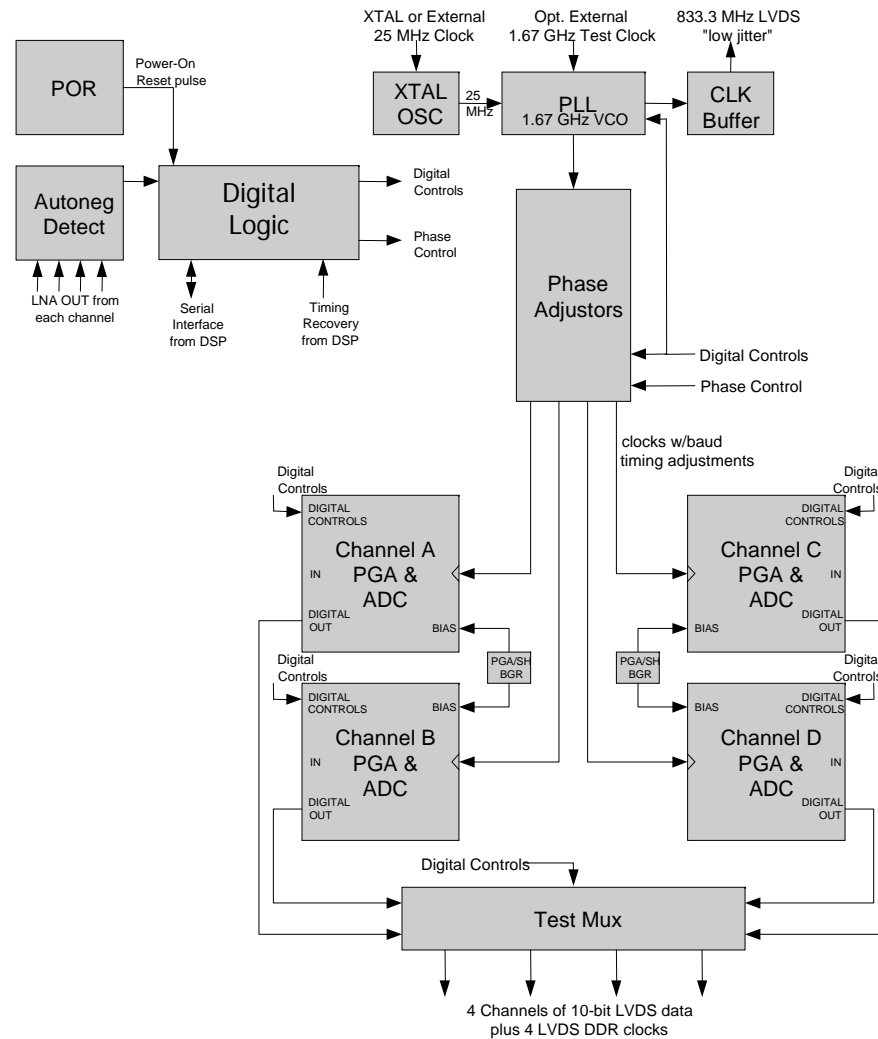


AFE Highlights

- Integrates:
 - Four 10-bit, 1Gbps analog-to-digital converters
 - Programmable gain amplifiers
 - PLL clock multiplication and distribution circuits
- 0.18 μ m CMOS technology
 - Low risk, low cost
- 2.7 Million Transistors




AFE Block Diagram





AFE Die Layout



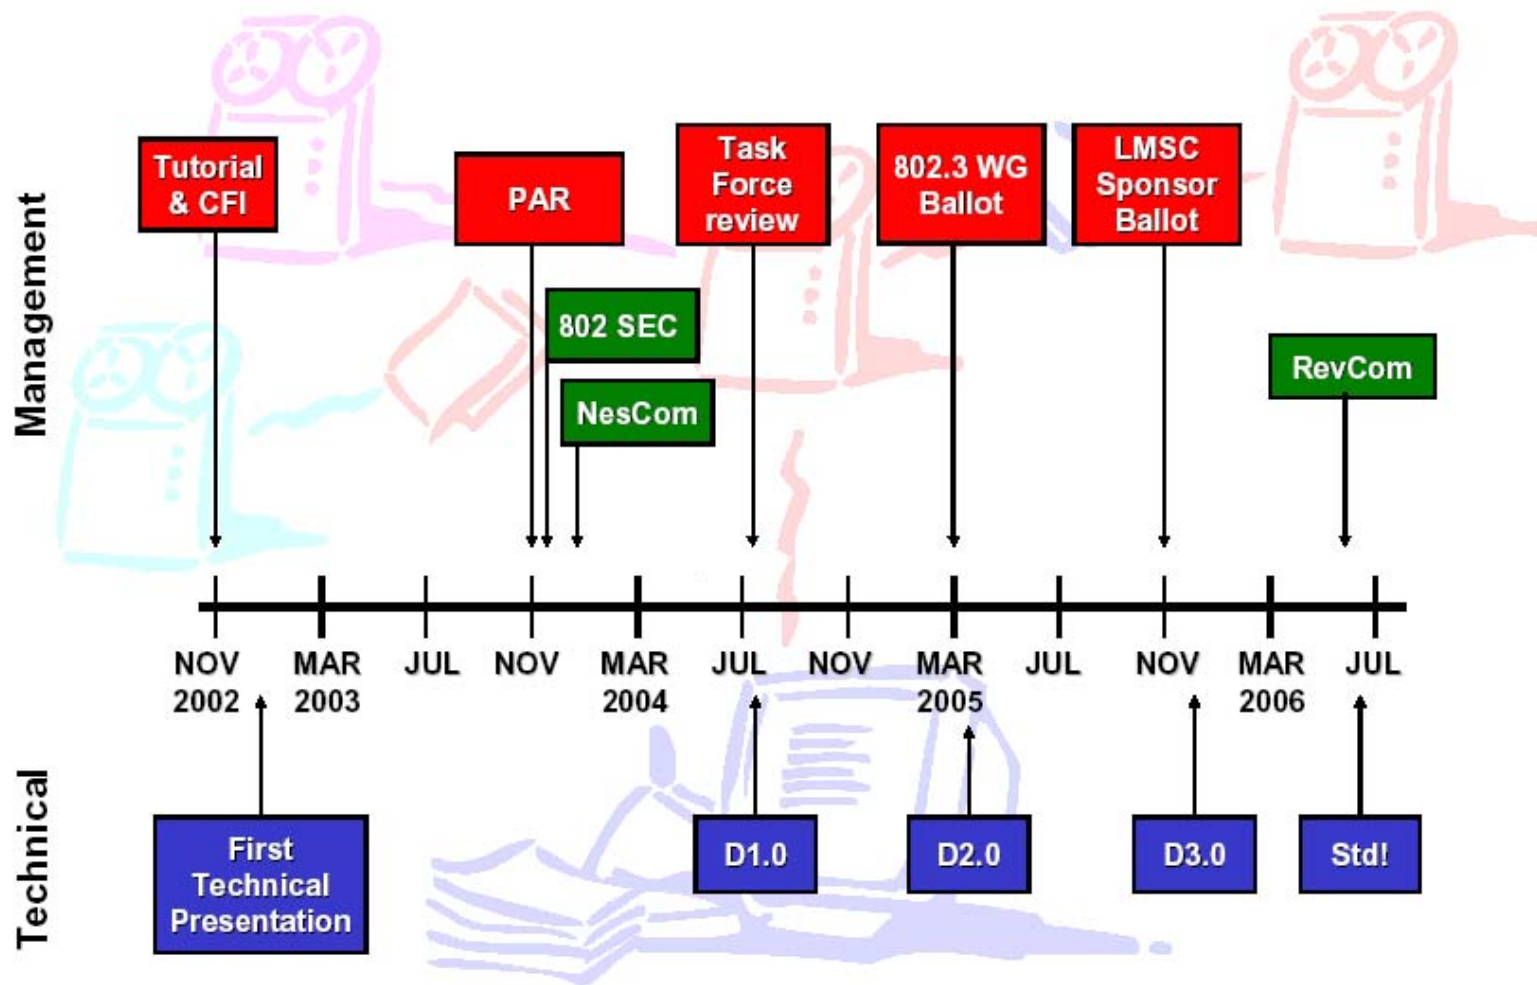


Key Challenges and Mitigation Strategies

- High input bandwidth (500MHz) needs
 - Custom sampling circuit with very short sampling aperture and high bandwidth PGA front end.
- Low jitter and low skew clock distribution
 - Self adjusting PLL based closed-loop architecture
- Low front-end noise and high linearity requirements
 - Careful partitioning of gain elements in signal path
- Channel to channel noise coupling
 - Dedicated power and ground systems per channel
 - Isolation guard rings and on-chip shielding structures
 - Extensive packaging/substrate modeling
 - Flip-chip packaging for bond-wire inductance elimination



IEEE Standard Timeline



Source: IEEE P802.3an agenda_1_0504.pdf



SOLARFLARE
COMMUNICATIONS