
*10GBASE-T:
TxFE solutions, dpSNR vs
length of precoding response,
and PMA training*

**IEEE P802.3an Task Force
Austin, May 18-20, 2005**

Gottfried Ungerboeck

Contents

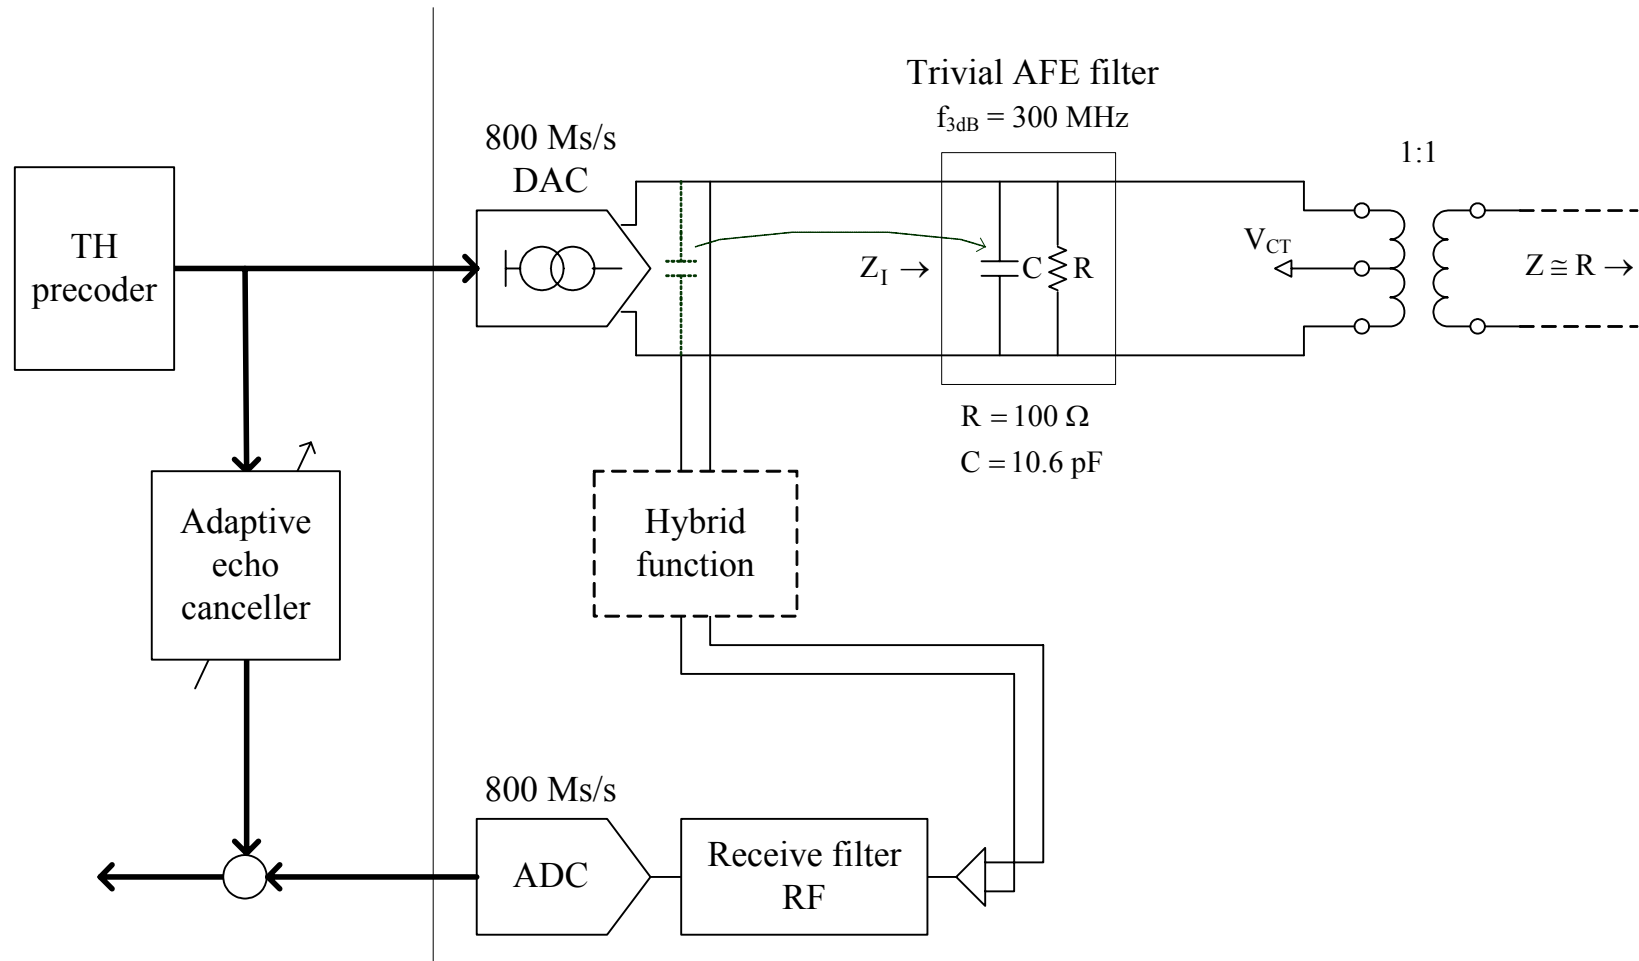
- Study of transmit front-end solutions
 - “Simple”: no digital filtering, 800 Ms/s DAC, simple R//C signal smoothing ($f_{3\text{dB}} = 300 \text{ MHz}$), 1:1 transformer → **transmit PSD depends on inaccurate analog components, no designed spectral nulls at dc and $1/2T$, poor return loss.**
 - “Baseline”: no digital filtering, 800 Ms/s DAC, signal smoothing by RLC front-end filter ($f_{3\text{dB}} = 300 \text{ MHz}$) with constant output impedance, 1:1 transformer
→ **transmit PSD depends on inaccurate analog components, no designed spectral nulls at dc and $1/2T$, good return loss.**
 - “Oversampled”: digital filtering and interpolation, 1600 Ms/s DAC, simple R//C signal smoothing ($f_{3\text{dB}} = 1 \text{ GHz}$), 1:1 transformer → **transmit PSD exhibits well defined shape with spectral nulls at dc and $1/2T$, good return loss.**

Contents

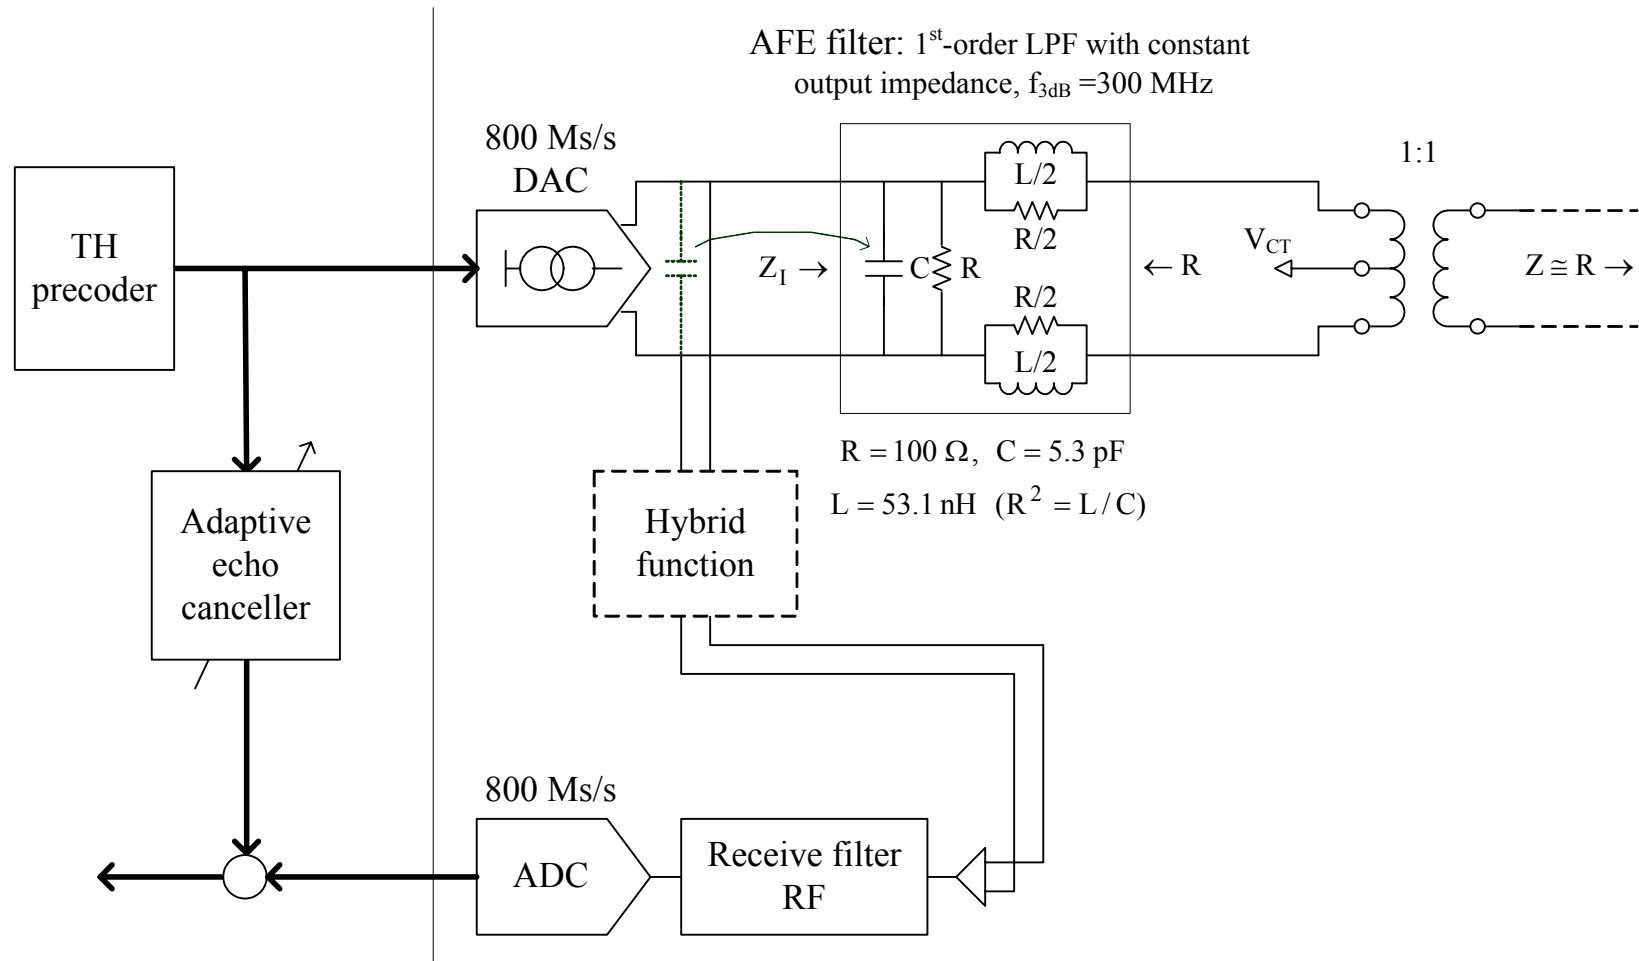
- **Decision-point SNR vs. length of precoding response**
 - Analysis: finite-length precoding response + infinite-length FFE optimized in MMSE sense.
 - Results obtained for “baseline” and “oversampled” transmit front-end, showing advantages of “oversampled” solution.
 - For worst case link characteristics, a programmable FIR precoding response of length $L = 32$ is found to be sufficient; $L = 16$ leads to small, but noticeable performance degradation.
- **PMA training issues**
 - Proposed main state diagram and state designations
 - Generation of PMA training sequences: proposal for concise and unambiguous description
 - Functional description of PMA training states.

Study of transmit front-end solutions

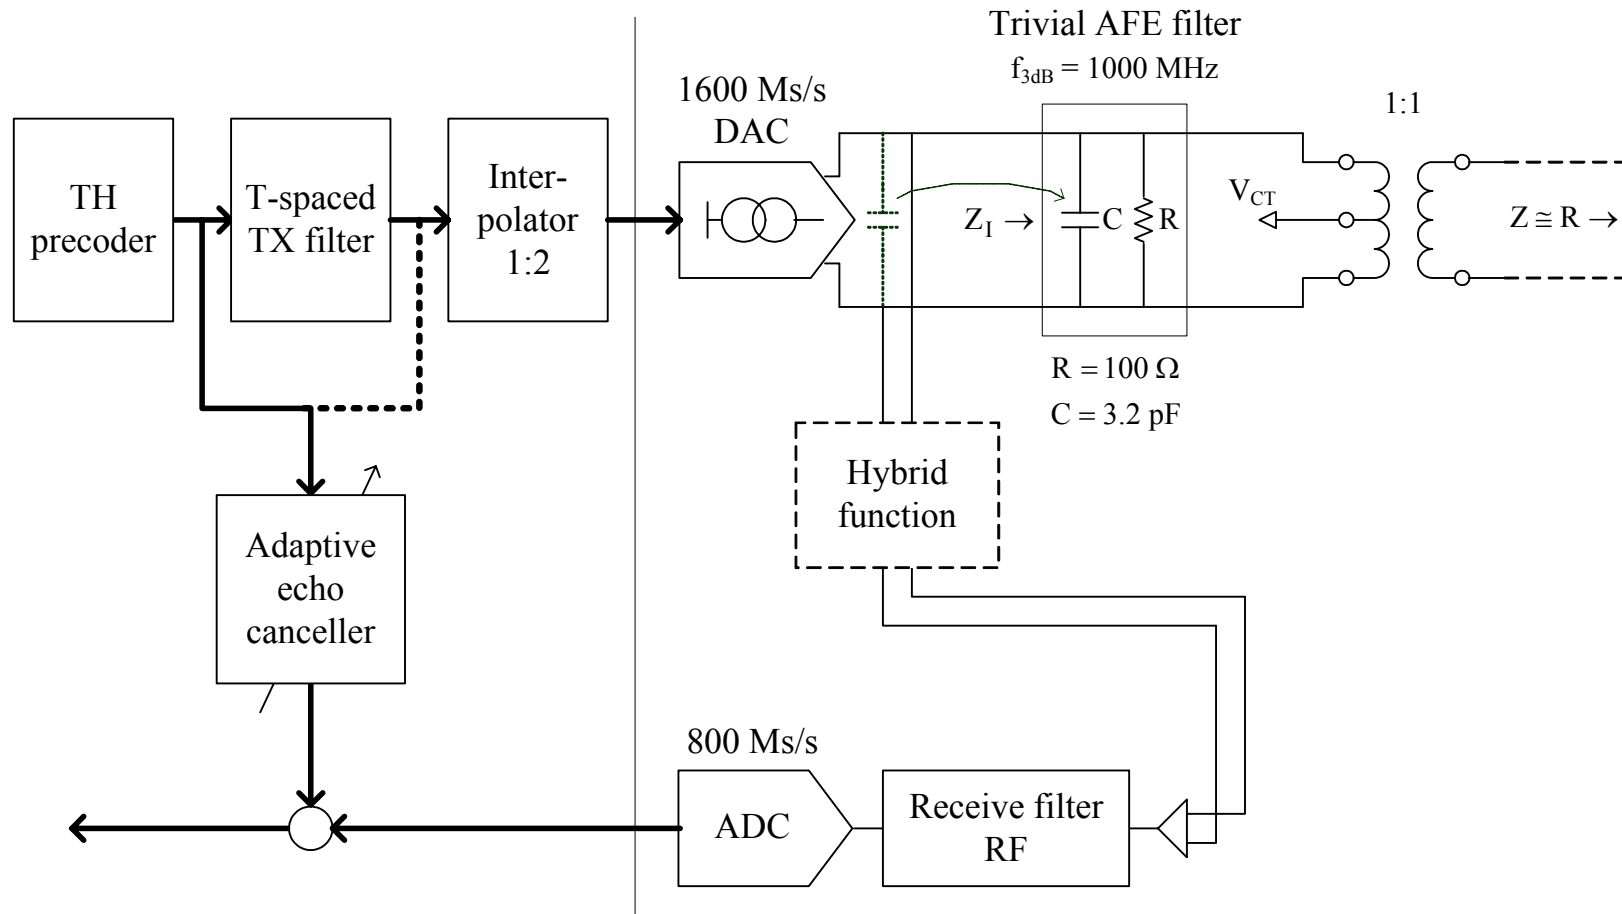
Transmitter front-end: “simple”




Transmitter front-end: “baseline”




Transmitter front-end: “oversampled”




Measurements of sample 10G quad-transformer



S_{21} : insertion loss PHY to Line (S_{12} is similar)




S_{22} : return loss Line to Line (S_{11} is similar)




Transmit PSD: “simple”


Perfectly implemented AFE filter + transformer pair 3




Transmit PSD: “simple”




Transmit PSD: “baseline”




Transmit PSD: “baseline”



Transmit PSD: “oversampled”




Transmit PSD: “oversampled”



MDI return loss


Perfectly implemented AFE filter + transformer pair 3



"Simple" solution does not meet MDI return loss spec: ruled out

Output voltages vs time: “baseline”

Simulated voltages at T/4-spaced time instances; rms and peak values determined in 1'000'000 T




Perfectly implemented AFE filter + transformer pair 3

$$P_T = 5 \text{ dBm into } 100 \Omega \Rightarrow 562 \text{ mV}_{\text{rms}}$$

Output voltages vs time: “oversampled”


Simulated voltages at T/4-spaced time instances; rms and peak values determined in 1'000'000 T




Perfectly implemented AFE filter + transformer pair 3

$$P_T = 5\text{ dBm into } 100\ \Omega \Rightarrow 562\text{ mV}_{\text{rms}}$$


Voltage distributions: “baseline”



Voltage distributions: “oversampled”



Single-ended voltage distributions at DAC output



Study of transmit-front-end solutions: summary

“Simple” solution ruled out: fails to meet MDI return loss spec.

	“Baseline” solution	“Oversampled” solution
Digital filters	none	$(1-D^2)/(1-0.75 D^2)$ + interpolator
DAC	800 Ms/s	1600 Ms/s
AFE filter	1-st order RLC LPF, $f_{3dB}=300$ MHz	Trivial R//C
rms and peak voltage at DAC output	higher rms, peak similar to “oversampled”	lower rms, peak similar to “baseline”
Excess bandwidth	substantial (→ sampling phase dependency in receiver)	sharp bandwidth limitation (EMI advantage)
Controlled spectral nulls	none	dc and 1/2T
Return loss	OK	OK
Transmit PSD shape	depends on analog components	digitally defined

Study of transmit-front-end solutions: conclusion

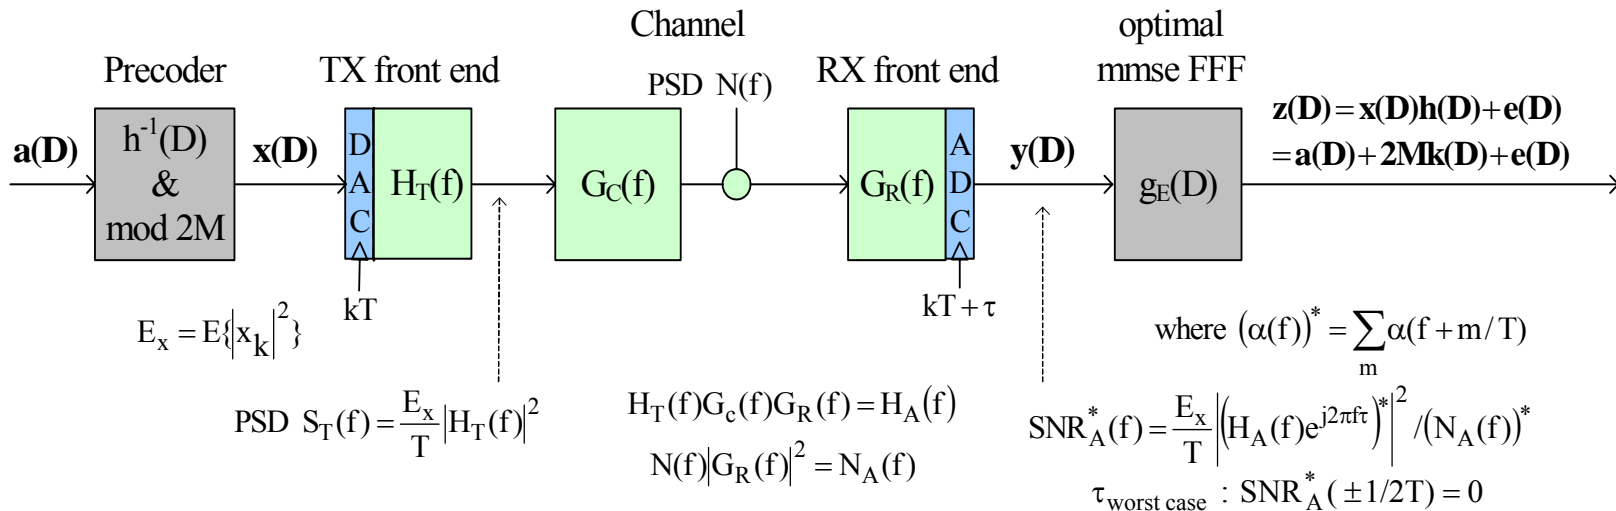
- Peak voltages at DAC output for “oversampled” and “baseline” solutions are similar; higher PAR of “oversampled” is compensated for by lower rms voltage.
- Cost of digital filtering and oversampling DAC outweighs disadvantages of “baseline solution”
 - o RLC AFE filter: two coils, concerns about balance, etc.
 - o PSD shape: substantial excess bandwidth, dependency on analog components, no controlled spectral nulls at dc and 1/2T
 - o hybrid function requires image impedances matching frequency-dependent input impedance of AFE filter.

Proposal: adopt well defined transmit PSD shape with sharp bandwidth limitation and spectral nulls at dc and 1/2T, as enabled by an “oversampled” transmit front-end solution.

$P_T = 5 \text{ dBm} \rightarrow 4 V_{ppd} \text{ at DAC output !!!}$

Decision-point SNR vs. length of precoding response

Optimum precoding response and decision-point SNR



Decision - point SNR for given $H_A(f)$, $N_A(f)$, τ , and precoding response $h(\mathbf{D})$


$$\text{SNR}_{\text{mmse}} = \left[T \int_{-1/2T}^{1/2T} \left| h(\mathbf{D} = e^{-j2\pi f T}) \right|^2 / (\text{SNR}_A^*(f) + 1) df \right]^{-1}, \quad h(\mathbf{D}) = 1 + \sum_{\ell=0}^L h_{\ell} \mathbf{D}^{\ell}$$

For given $\text{SNR}_A^*(f) + 1$, determine $(\arg) \max_{h_1, h_2, \dots, h_L} \text{SNR}_{\text{mmse}}$.

Folded spectral SNR function +1

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz
PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m

Fixed receive filter: 3rd-order BWF, $f_{3dB}=300$ MHz; worst-case sampling phase



Finite-length $h(D)$ + infinite FFE, MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
Fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase

TX front-end: “baseline”


L	SNR _{mmse}	$\sum_{\ell=0}^L h_{\ell} $
4	23.50 dB	6.17
6	23.93 dB	9.15
8	24.11 dB	12.00
12	24.27 dB	17.32
16	24.34 dB	22.04
24	24.40 dB	29.81
32	24.41 dB	35.72
48	24.43 dB	43.79
64	24.43 dB	48.84
96	24.43 dB	54.51
128	24.43 dB	57.59
192	24.43 dB	60.73

TX front-end: “oversampled”


L	SNR _{mmse}	$\sum_{\ell=0}^L h_{\ell} $
4	24.21 dB	5.49
6	24.45 dB	7.36
8	24.48 dB	8.22
12	24.49 dB	8.59
16	24.55 dB	10.74
24	24.67 dB	15.26
32	24.73 dB	17.95
48	24.75 dB	19.75
64	24.76 dB	20.08
96	24.76 dB	20.17
128	24.76 dB	20.52
192	24.76 dB	21.09

Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "baseline"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase




$$\sum_{\ell=0}^L |h_{\ell}| = 57.59$$




Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "oversampled"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase




$$\sum_{\ell=0}^L |h_{\ell}| = 20.52$$




Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "baseline"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase




$$\sum_{\ell=0}^L |h_{\ell}| = 35.72$$




Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "oversampled"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase




Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "baseline"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase




$$\sum_{\ell=0}^L |h_{\ell}| = 22.04$$



Finite-length $h(D)$ + infinite FFE , MMSE optimized

Class E cable, $P_T = 5$ dBm, $l = 100$ m, AWGN = -140 dBm/Hz, PS_ANEXT and PS_AFEXT from $P_T = 5$ dBm, $l = 100$ m
TX front-end: "oversampled"; fixed receive filter: 3rd-order BWF, $f_{3dB} = 300$ MHz; worst-case sampling phase



DP-SNR vs length of precoding response: conclusion

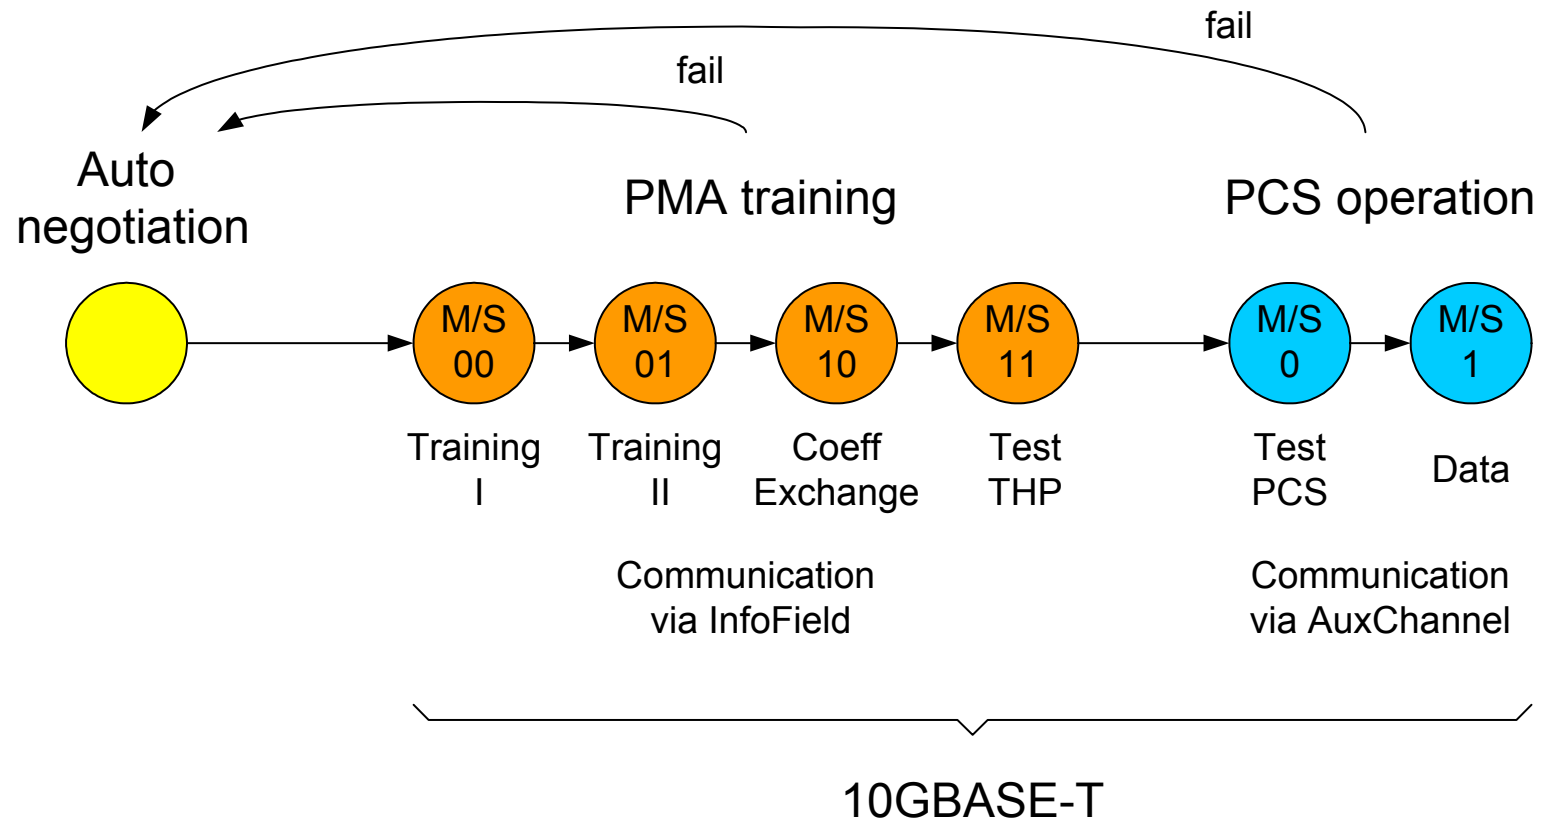
- Decision-point SNR is insensitive to length L of precoding response; a programmable FIR precoding response with $L = 32$ is adequate. $L = 16$ leads to small, but noticeable performance degradation. $L = 32$ provides headroom for dealing with non-smooth SNR(f). Coefficients >2 found.
- In addition, the results illustrate the benefits of the “oversampled” TX front-end:
 - o Higher decision-point SNR due to better PSD shape
 - o SNR performance always insensitive to sampling phase due to stricter bandwidth limitation and spectral null at $1/2T$
 - o Spectral null at dc reduces constellation expansion.

Proposal: adopt programmable FIR precoding with $L = 32$ and coefficient values in $[-4,+4]$. Forget fixed precoding responses.

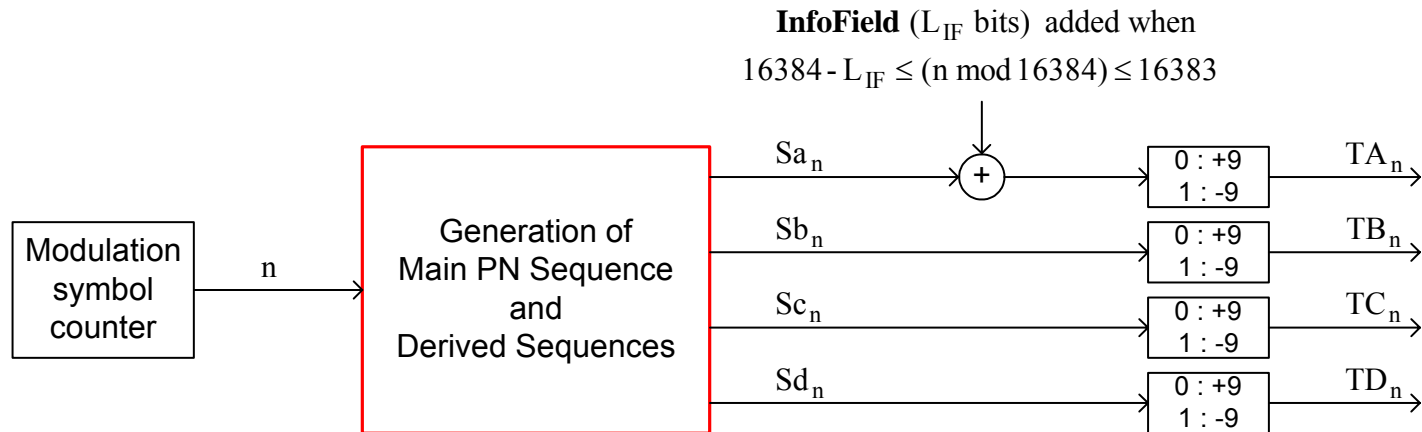
Same response for all pairs, or four individual responses?

PMA training issues

Proposed main state diagram and state designations



Unambiguous generation of PMA training sequences



Main PN sequence

$n \bmod 16384 = 0$: $Scr_n[0:32] = 33 \text{ lsbs of } 0x15979A422$ (periodic initialization)

$n \bmod 16384 \neq 0$: $Scr_n[1:33] = Scr_{n-1}[0:32]$

$$Scr_n[0] = \begin{cases} Scr_n[20] \oplus Scr_n[33] & \text{if PMA_CONFIG} = \text{MASTER} \\ Scr_n[13] \oplus Scr_n[33] & \text{if PMA_CONFIG} = \text{SLAVE} \end{cases}$$

Derived sequences

$$Sa_n = \begin{cases} Scr_n[0] \oplus 1 & \text{if } n \bmod 256 = 0 \\ Scr_n[0] & \text{otherwise} \end{cases}$$

$$Sb_n = Scr_n[3] \oplus Scr_n[8]$$

$$Sc_n = Scr_n[6] \oplus Scr_n[16]$$

$$Sd_n = Scr_n[9] \oplus Scr_n[14] \oplus Scr_n[19] \oplus Scr_n[24]$$

← really needed?

Functional description of PMA training states

Master Training I (M00)

- Send 16K-periodic PMA training frames w/o THP
- Adjust echo/next cancellation
- Gradually increase TX power starting from minimum power
- Check for reception of PMA training frames from Slave

Relevant entries in transmitted InfoField

- | | |
|-------------------------------------|--------|
| - current state = 00 | 2 bit |
| - current TX power | 3 bit |
| - announced next TX power | 3 bit |
| - announce transition to next state | 1 bit |
| - transition counter | 12 bit |

Condition for transition to next PMA training state (M01)

Reception of PMA training frames from Slave detected

Do you prefer to say “increase TX power” or “decrease PBO”?

Functional description of PMA training states

Master Training II (M01)

- Send 16K-periodic PMA training frames w/o THP
- Refine echo/next cancellation adjustments
- Adjust receiver settings for DFE operation
- InfoField decoding

Relevant entries in transmitted InfoField

- | | |
|-------------------------------------|--------|
| - current state = 01 | 2 bit |
| - current TX power | 3 bit |
| - announced next TX power | 3 bit |
| - requested TX power | 3 bit |
| - decision-point MSE in dB | 6 bit |
| - announce transition to next state | 1 bit |
| - transition counter | 12 bit |

Relevant entries in received InfoField

Same as in transmitted InfoField with current state 01

Condition for transition to next PMA training state (M10)

Master and Slave decision-point MSE < MSE threshold

Functional description of PMA training states

Master Coeff Exchange (M10)

- Send 16K-periodic PMA training frames w/o THP
- Refine echo/next cancellation adjustments
- Refine receiver settings for DFE operation
- InfoField decoding

Relevant entries in transmitted InfoField

- | | |
|-------------------------------------|-----------------------------------|
| - current state = 10 | 2 bit |
| - coefficient index | 4 bit (3 bit) |
| - two coefficients | 16 bit (four coefficients 32 bit) |
| - all coefficients received | 1 bit |
| - announce transition to next state | 1 bit |
| - transition counter | 12 bit |

Relevant entries in received InfoField

Same as in transmitted InfoField with current state = 10

Condition for transition to next PMA training state (M11)

Master and Slave have received all 32 coefficients

Functional description of PMA training states

Master Test THP (M11)

- Send 16K-periodic PMA training frames with THP
- Refine echo/next cancellation adjustments
- Refine receiver settings for THP operation
- InfoField decoding

Relevant entries in transmitted InfoField

- | | |
|-------------------------------------|--------|
| - current state = 11 | 2 bit |
| - current TX power | 3 bit |
| - announced next TX power | 3 bit |
| - requested TX power | 3 bit |
| - decision-point MSE in dB | 6 bit |
| - announce transition to next state | 1 bit |
| - transition counter | 12 bit |

Relevant entries in received InfoField

Same as in transmitted InfoField with current state 11

Condition for transition to PCS operation

Master and Slave decision-point MSE < MSE threshold

Functional description of PMA training states

Slave functions in Training I (S00), Training II (S01), Coeff Exchange (S10), and Test THP (S11) are conceptually similar.