Extrapolation For Bin16

In support of comment #R2-191

Amir Rubin, Nvidia

Current Extrapolation for bin 16

Based on linear model of log10 of Hm(k), Hm(k) = portion of bin k from total #FEC blocks

- Quotation from 180.9.14 Receiver sensitivity, page 492 :

  "NOTE—In order to predict whether a receiver meets the BLER requirement in a short test time, extrapolation of the measured histogram (see 174A.9.3) to Hm(16) can be performed. One way of doing this is using a linear fit of log10(Hm(k)) for k = (n – 2) to n, where these values of k are the three highest bins having a count of 10 or more. If the slope of the histogram over the measured bins is increasing, it might be an indication of not meeting the BLER requirement. To verify that the slope is not increasing, the values of (log10(Hm(k)) – log10(Hm(4))) / (k – 4) should be less than (log10(Hm(4)) – log10(Hm(2))) / 2 for each k greater than 4."
- However, the linear model of log10 of portion of uncorrelated errors is:
  - Biased towards false failures
  - With excessive scatter over random scenarios

The model is biased up (1)

- The extent of bias is meaningful towards false failures of marginally passing links
- The table below shows the model bias vs. binomial distribution values, **for bin16 and not BLER**

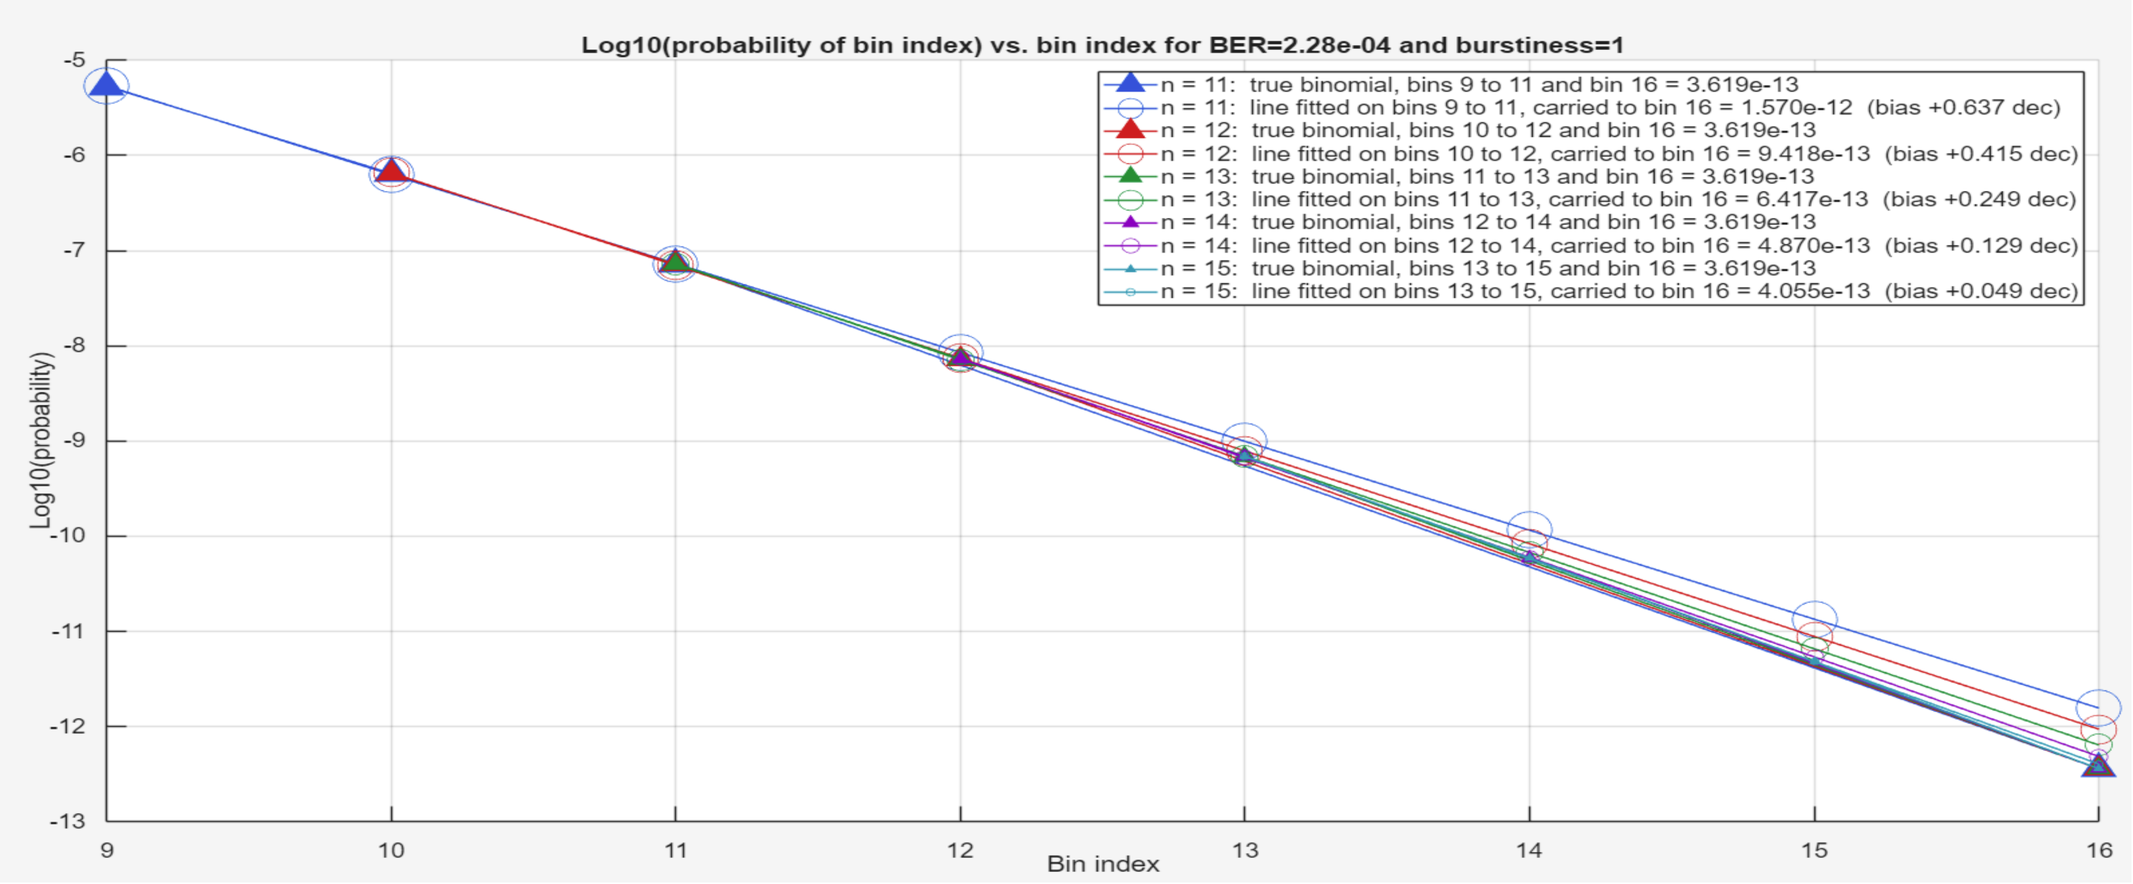

The model is biased up (2)

Summary table

- The extent of bias is meaningful towards false failures of marginally passing links
- The table below shows the model bias vs. binomial distribution values, **for bin16 and not BLER**

Max bin >=10	Min duration with BER = 2.28e-4	Extrapolated Bin 16	Model bias: Log10 $\left(\frac{Extrapolated\ bin\ 16}{Target\ bin\ 16}\right)$=Log10 $\left(\frac{Extrapolated\ bin\ 16}{3.619e-13^{*}}\right)$
11	>=4[Sec], Hm(11)=11	1.57e-12	0.637
12	>=36[Sec], Hm(12)=10	9.42e-13	0.415
13	>=6:17[Min:Sec], Hm(13)=10	6.42e-13	0.245
14	>=1:13[Hours:Min], Hm(14)=10	4.87e-13	0.129
15	>=14:57[Hours:Min], Hm(15)=10	4.05e-13	0.049

*Target bin16 with uncorrelated BER 2.28e-4 = 3.619e-13

The Monte Carlo Analysis (1)

- Draw 10,000 cases of Histogram bins Hm(k), per 3 durations X 3 BER
  - Durations: 60[Sec], 10[Min] and 60[Min], BER values: 2.28e-4 (target), 2.0e-4 and 2.177e-4 (BLER that is ½ vs. BLER with 2.28e-4)
- Per case execute the 180.9.14 log10-linear extrapolation and the proposed binomial extrapolation

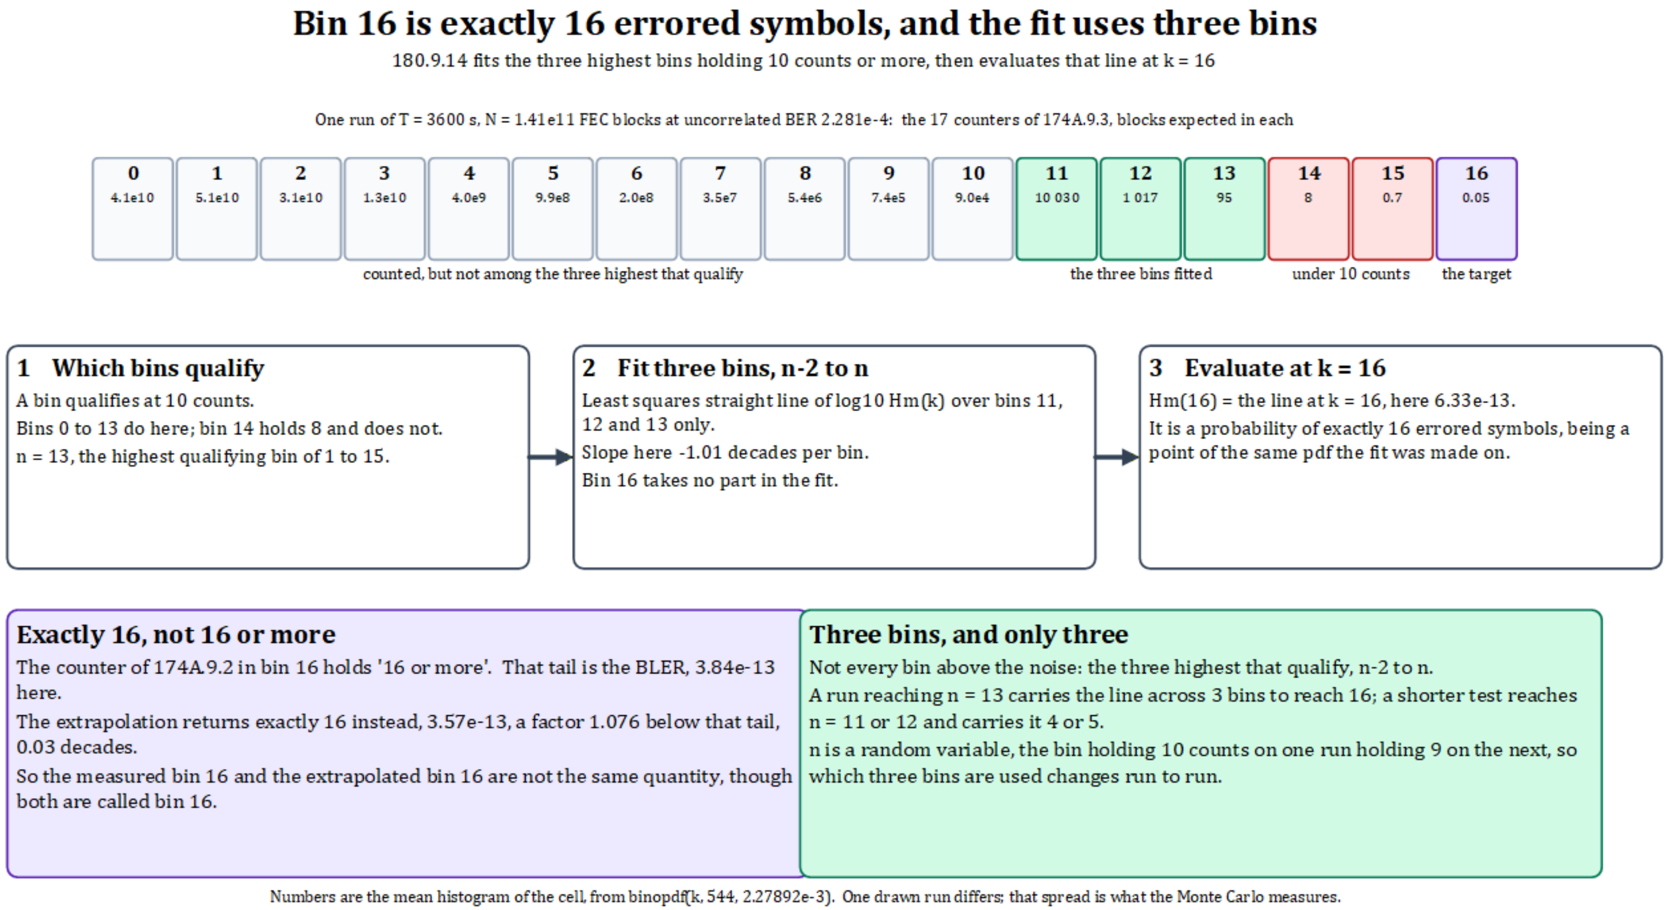

The Monte Carlo Analysis (2)

- Draw 10,000 cases of Histogram bins Hm(k), per 3 durations X 3 BER
  - Durations: 60[Sec], 10[Min] and 60[Min], BER values: 2.281e-4 (target), 2.0e-4 and 2.177e-4 (BLER that is ½ vs. BLER with 2.28e-4)
- Per case execute the 180.9.14 log10-linear extrapolation and the proposed binomial extrapolation

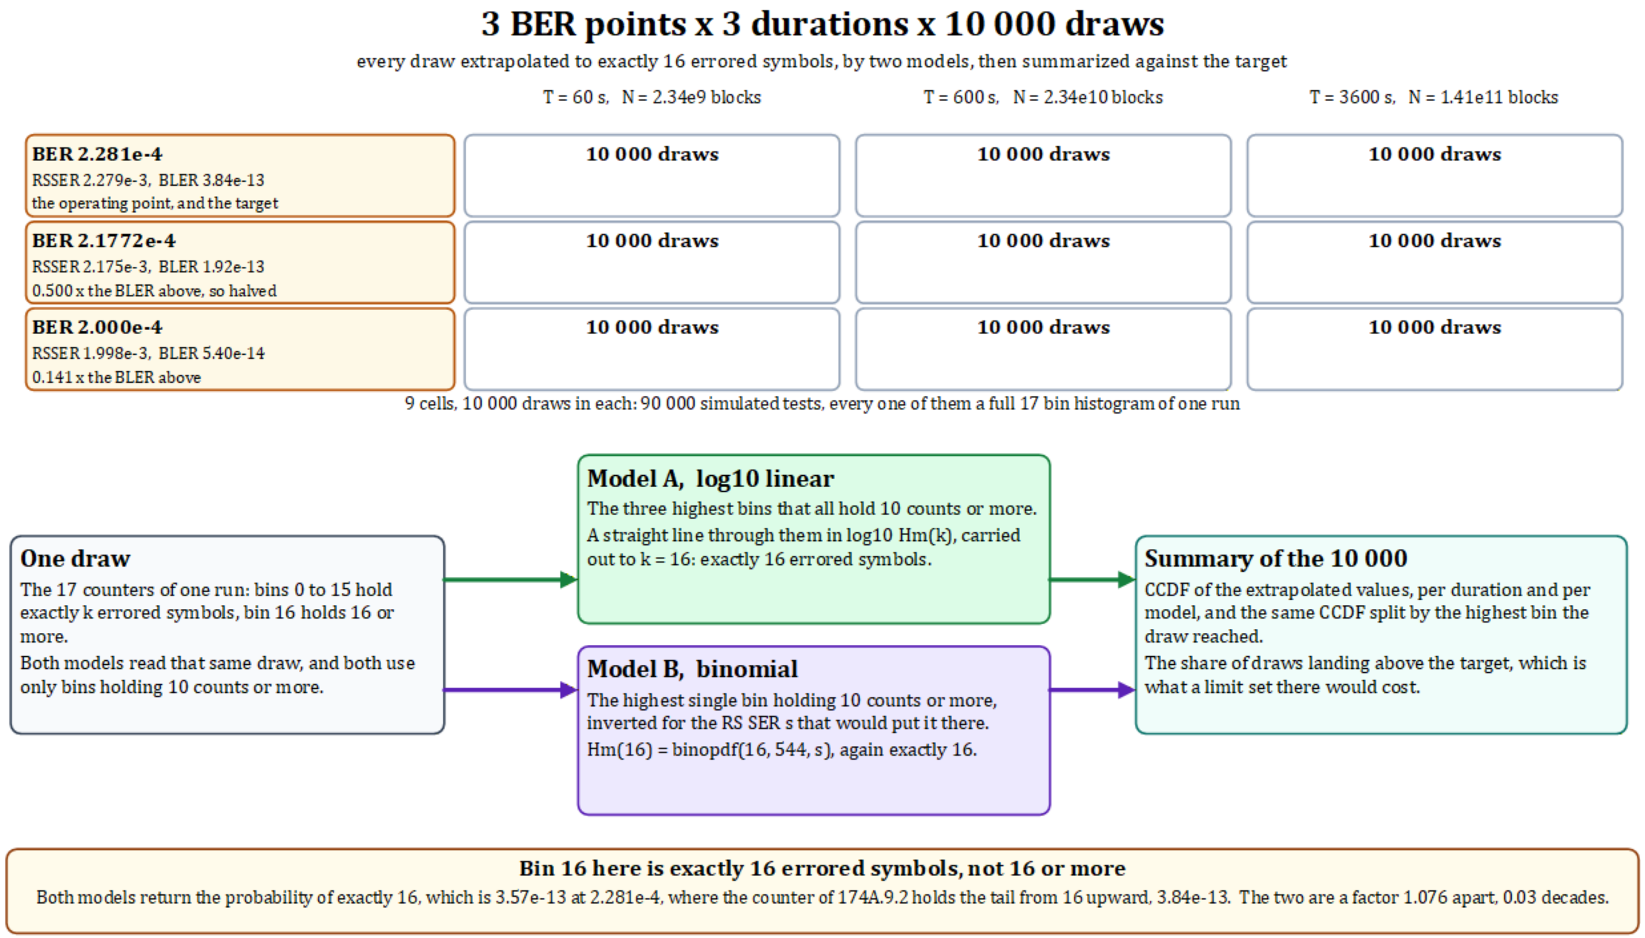

180.9.14 log10 linear model distribution (1)

Target BER=2.28e-4, leading to BLER=3.84e-13, vs. BER=2.0e-4 leading to BLER=5.42e-14 (-85.9%)

- Wide scatter over random scenarios, based on Monte Carlo draws
- Probability of more than 80% for failure of marginal links (uncorrelated BER of 2.281e-4), that should pass with 50%
    - For duration of 60[Sec], there’s ~18% for false failure for links with uncorrelated BER of 2e-4, that should always pass
- The figures below show the CCDF (Complementary CDF = probability of X value and above/worse) of current extrapolation scheme for 1X links

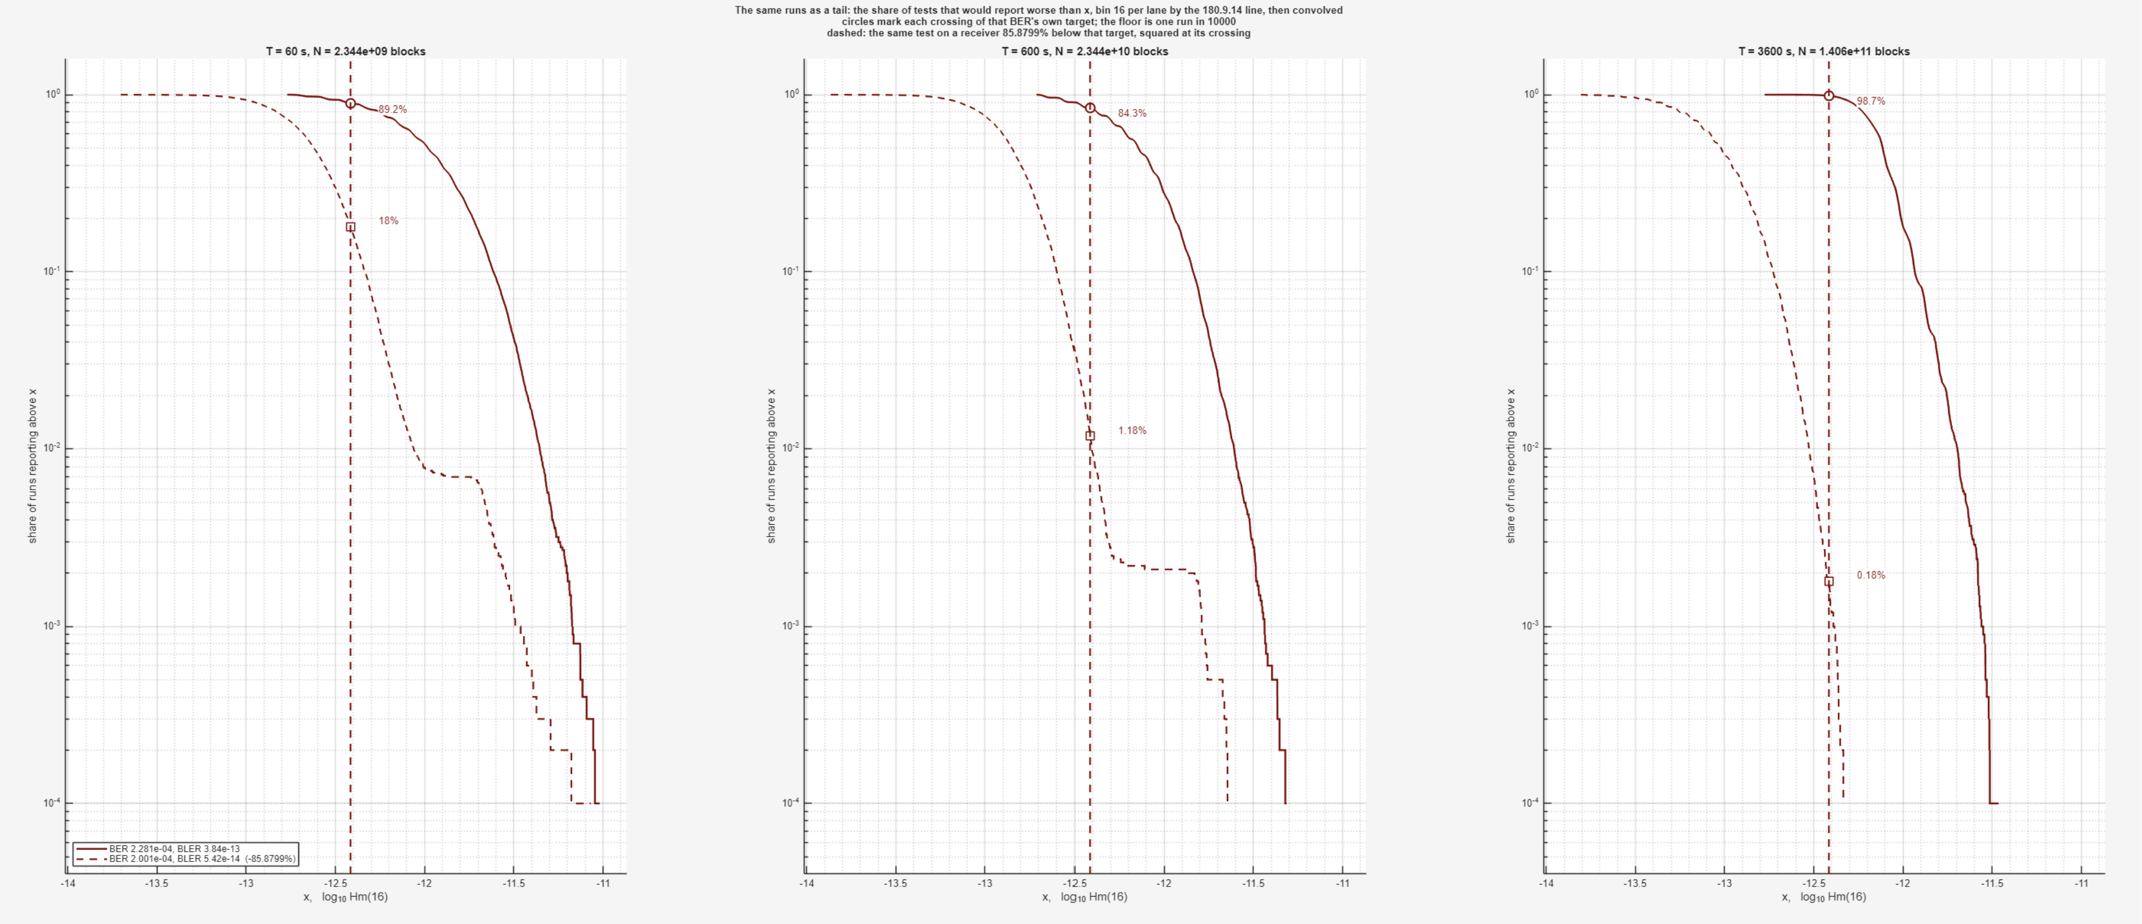

180.9.14 log10 linear model distribution (2)

Target BER=2.281e-4, leading to BLER=3.84e-13, vs. BER=2.177e-4 leading to BLER=1.92e-13 (-50%)

- High probability of false failure for links with BLER that is 50% lower than target (1.92e-13), that should always pass
- The figures below show the CCDF (Complementary CDF = probability of X value and above/worse) of current extrapolation scheme for 1X links

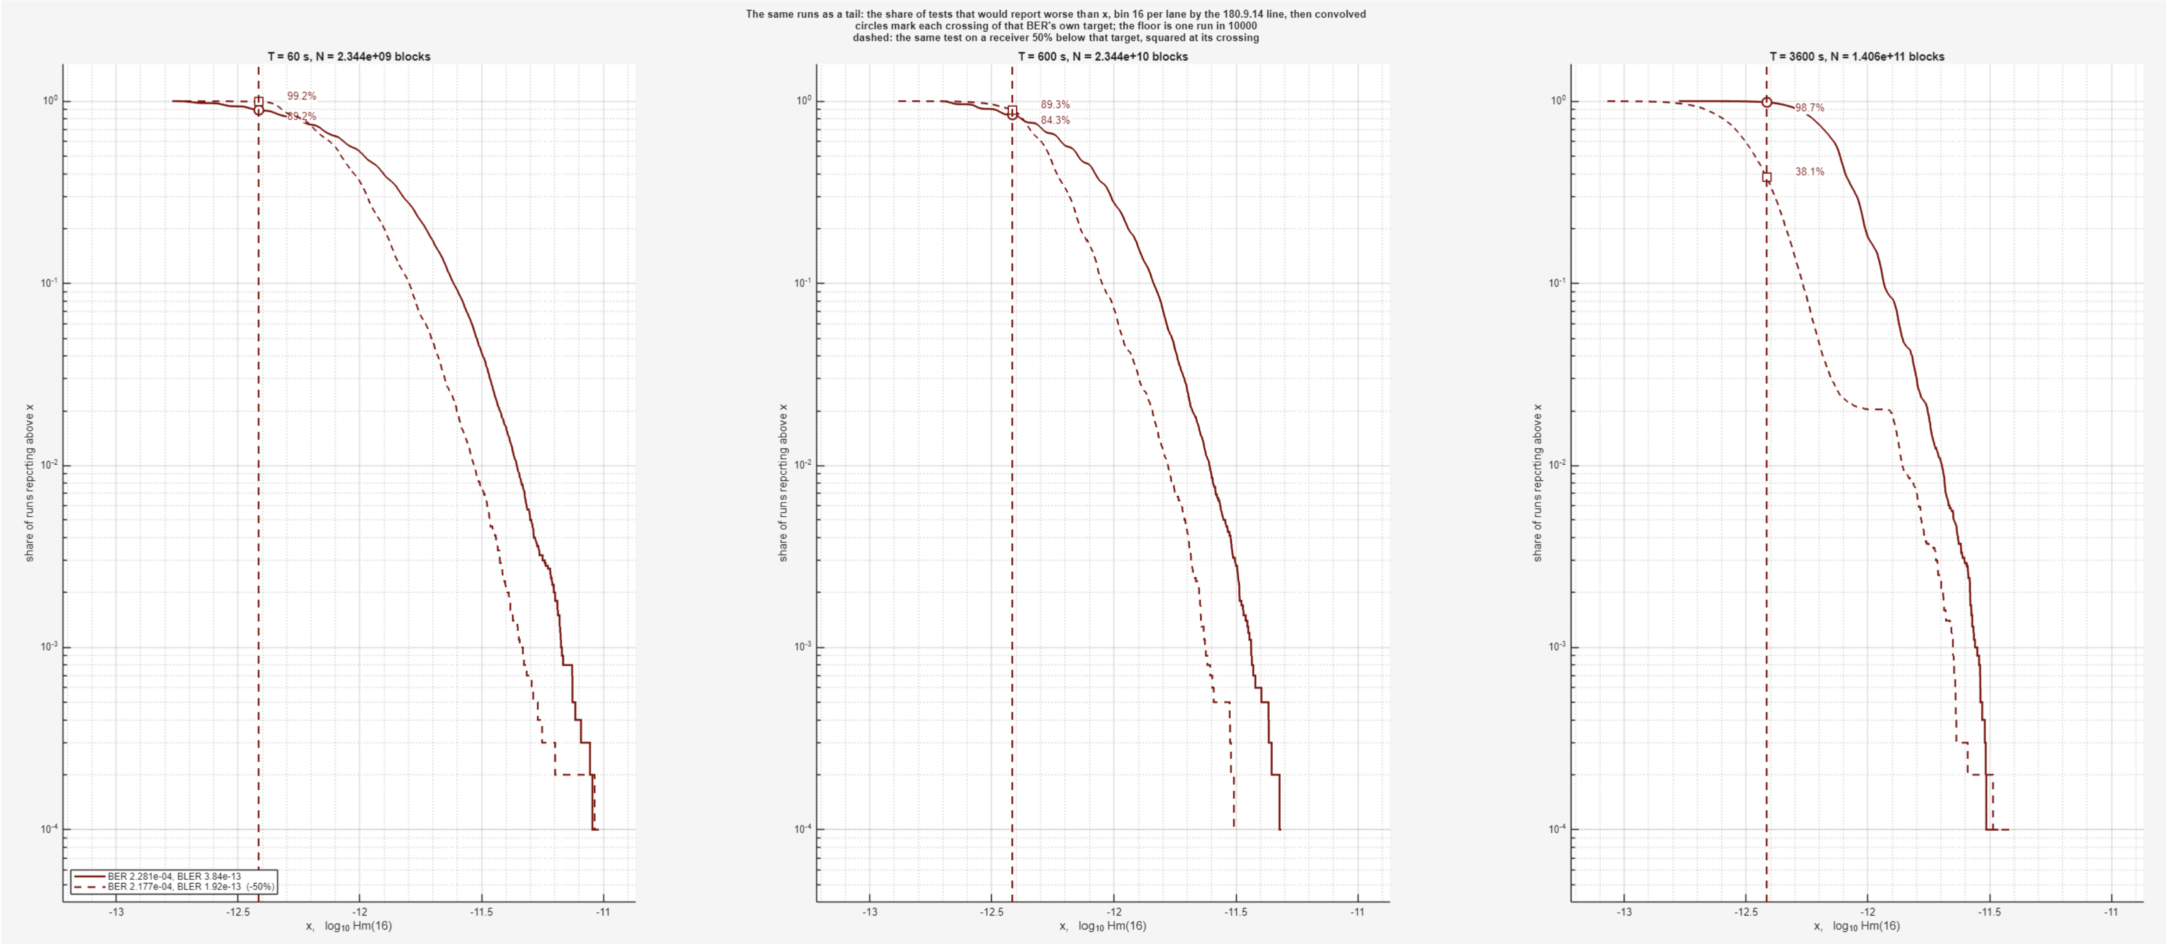

Proposed Remedy – Model

Use directly the binomial distribution function which has no bias

- Equivalent RSSER(k): The RSSER for which the portion of histogram k fits the binomial distribution function

$$\mathrm{Hmax(k)} = \mathrm{H(k)}$$

$$\frac{\mathrm{n!}}{\mathrm{k!\,(n-k)!}}\mathrm{RSSER}^{\mathrm{k}}(1-\mathrm{RSSER})^{\mathrm{n-k}} = \frac{\text{Number of FEC blocks with k RS symbol errors}}{\text{Total Number of FEC blocks}}$$

- Proposed extrapolation:
  - Find Equivalent RSSER for max nonzero histogram with count of 10 or more
  - Extrapolated bin 16 = $\sum_{\mathrm{k}=16}^{544} \frac{544!}{\mathrm{k!}(544-\mathrm{k})!} \text{Equivalent RSSER}^{\mathrm{k}}(1-\text{Equivalent RSSER})^{544-\mathrm{k}}$

- Alternatively see if the following is satisfied for max nonzero histogram k with count of 10 or more

$$\frac{\mathrm{n!}}{\mathrm{k!\,(n-k)!}}\text{Target RSSER}^{\mathrm{k}}(1-\text{Target RSSER})^{\mathrm{n-k}} = \frac{\text{Number of FEC blocks with k RS symbol errors}}{\text{Total Number of FEC blocks}}$$

Proposed remedy binomial model distribution (1)

- Narrow scatter over random scenarios, based on Monte Carlo draws
- Reduces the probability for failure of marginal links (uncorrelated BER of 2.218e-4), that should pass with 50%
- Secures that links with uncorrelated BER of 2e-4 will pass, equivalent to 0.09[dB]/0.045[dBo] margin
- The figures below show the CCDF (Complementary CDF = probability of X value and above/worse) of the proposed extrapolation scheme for 1X links

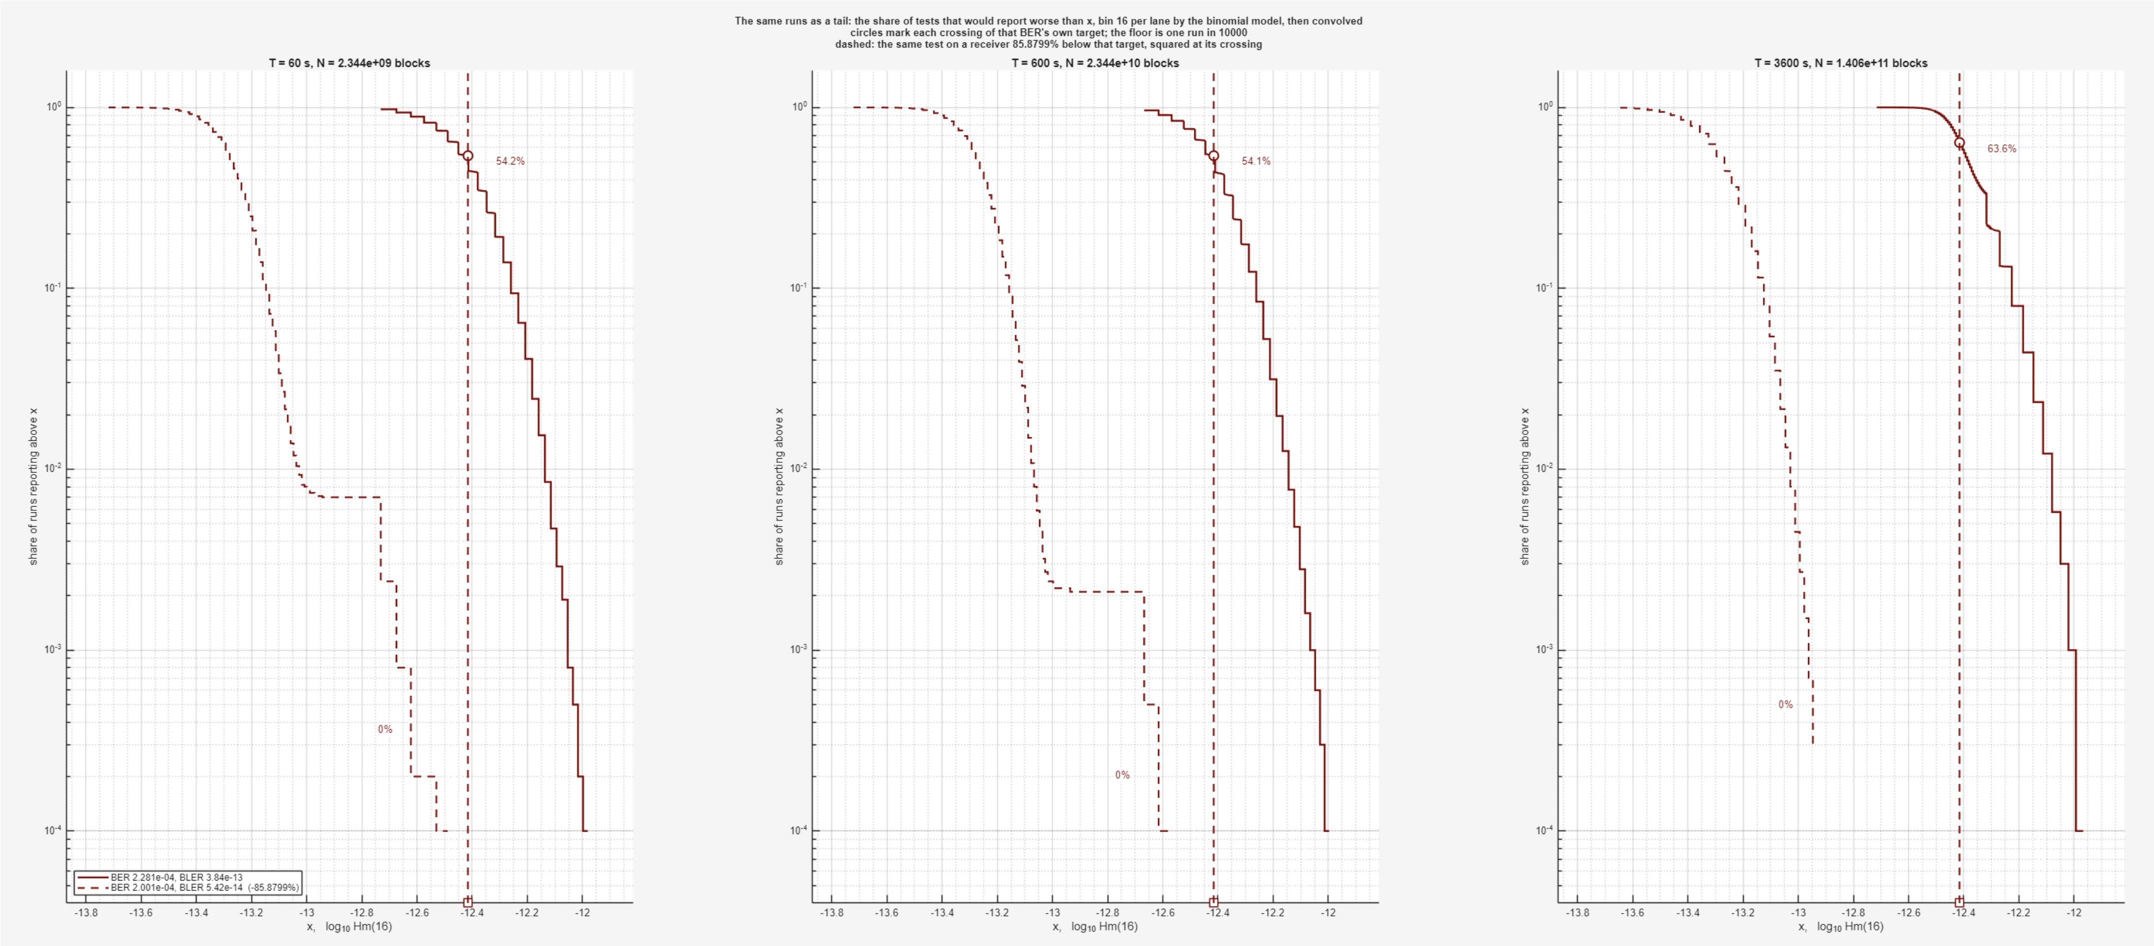

Proposed remedy binomial model distribution (2)

- Very meaningful drop of the probability of false failure for links with BLER that is 50% lower (1.92e-13), equivalent to 0.032[dB]/0.016[dBo] margin
- The figures below show the CCDF (Complementary CDF = probability of X value and above/worse) of the proposed extrapolation scheme for 1X links

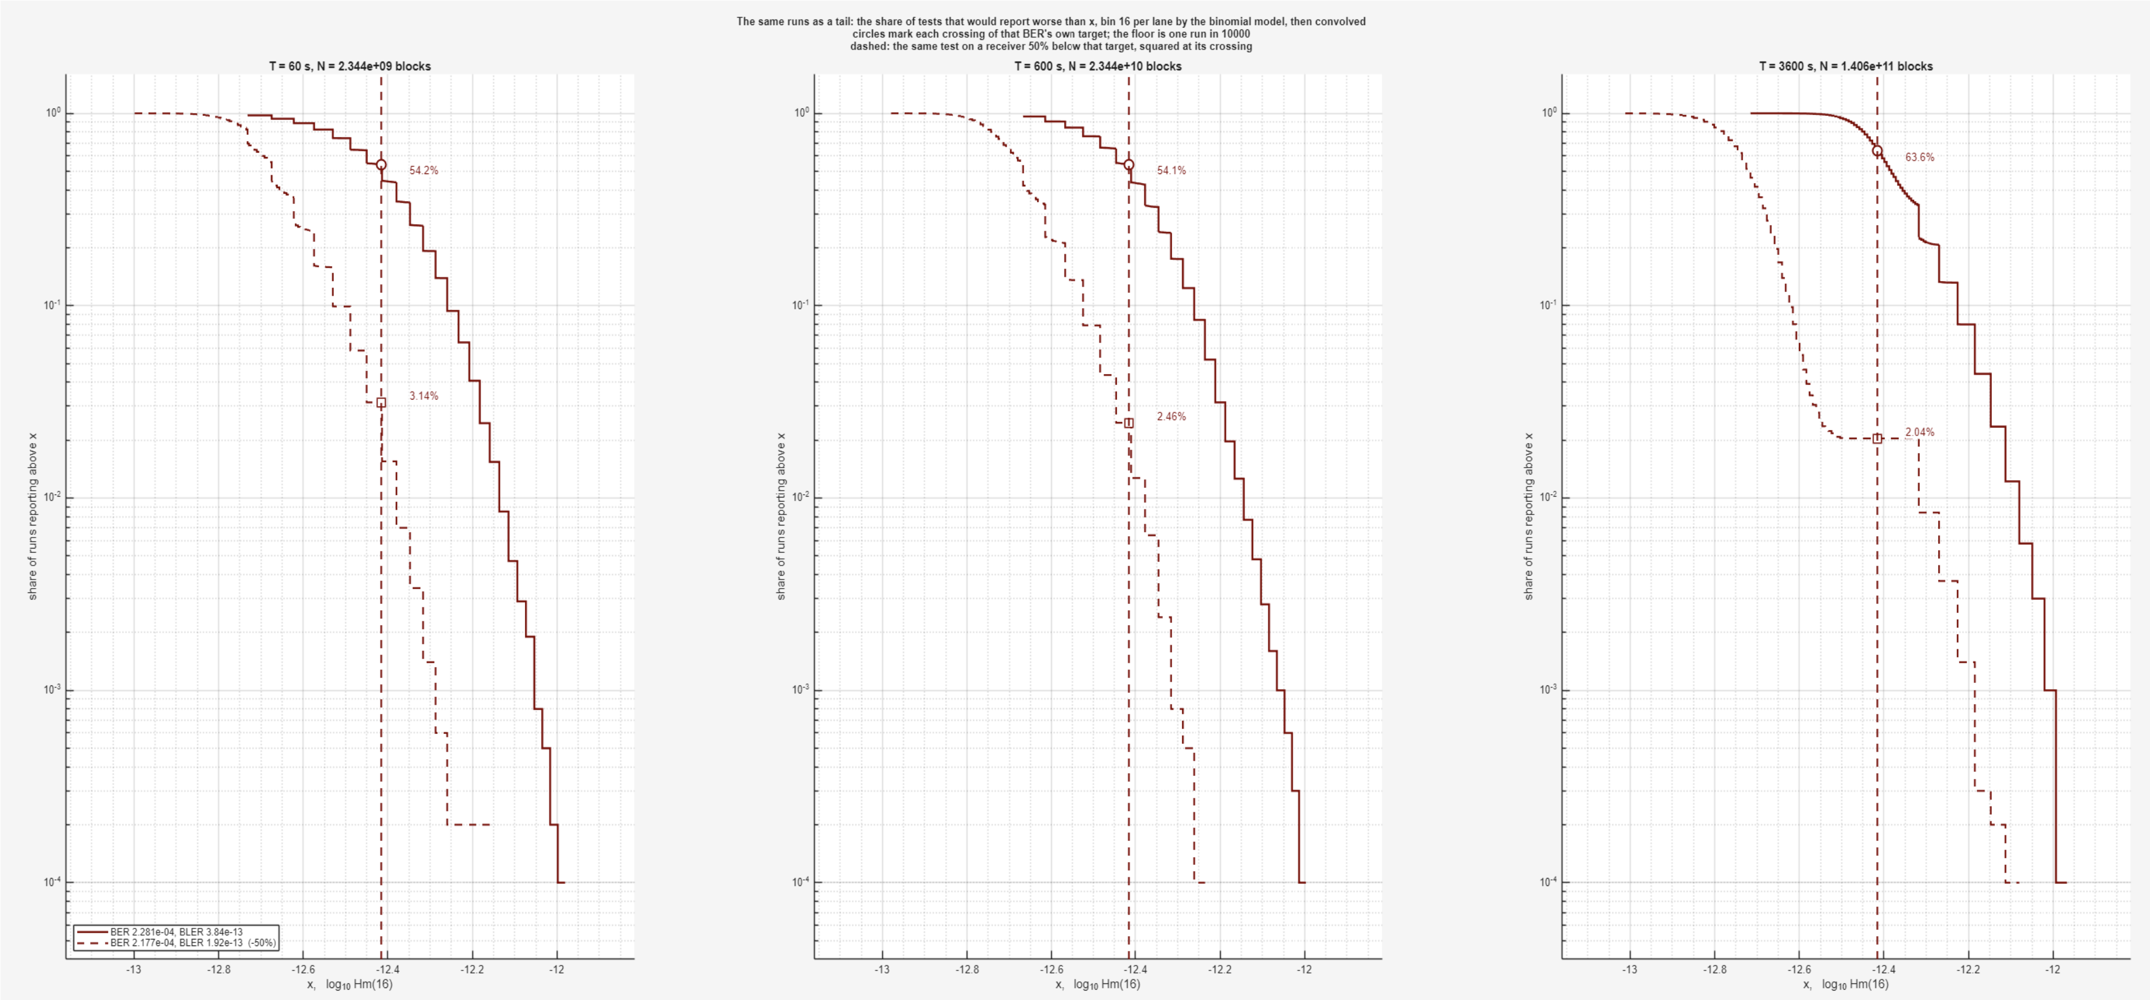

Proposed Remedy – Text

Replace the paragraph quoted from “180.9.14 Receiver sensitivity” with the following:

- NOTE—In order to predict whether a receiver meets the BLER requirement in a short test time, extrapolation of the measured histogram (see 174A.9.3) to Hm(16) can be performed. This can be done by using binomial distribution model for uncorrelated errors as follows: For the highest bin having count of 10 or more, find which RS SER in equation 174A-5 would lead to that count from the total amount of FEC blocks. Then estimate Hm(16) by accumulating equation 174A-5 with the found RS SER over bins 16 and above.

  Alternatively, plug the target RS SER for the scenario into equation 174A-5 and see if the count of the highest bin having count of 10 or more is not higher than that value.

Backup Slides

180.9.14 log10 linear model distribution (3)

Explanation of step in the CCDF

- The step in the CCDF (Complementary CDF = probability of X value and above/worse) for duration of 60[Sec] and for 600[Sec] results from having 2 values of max nonzero bins, where each of them has quite different bias.

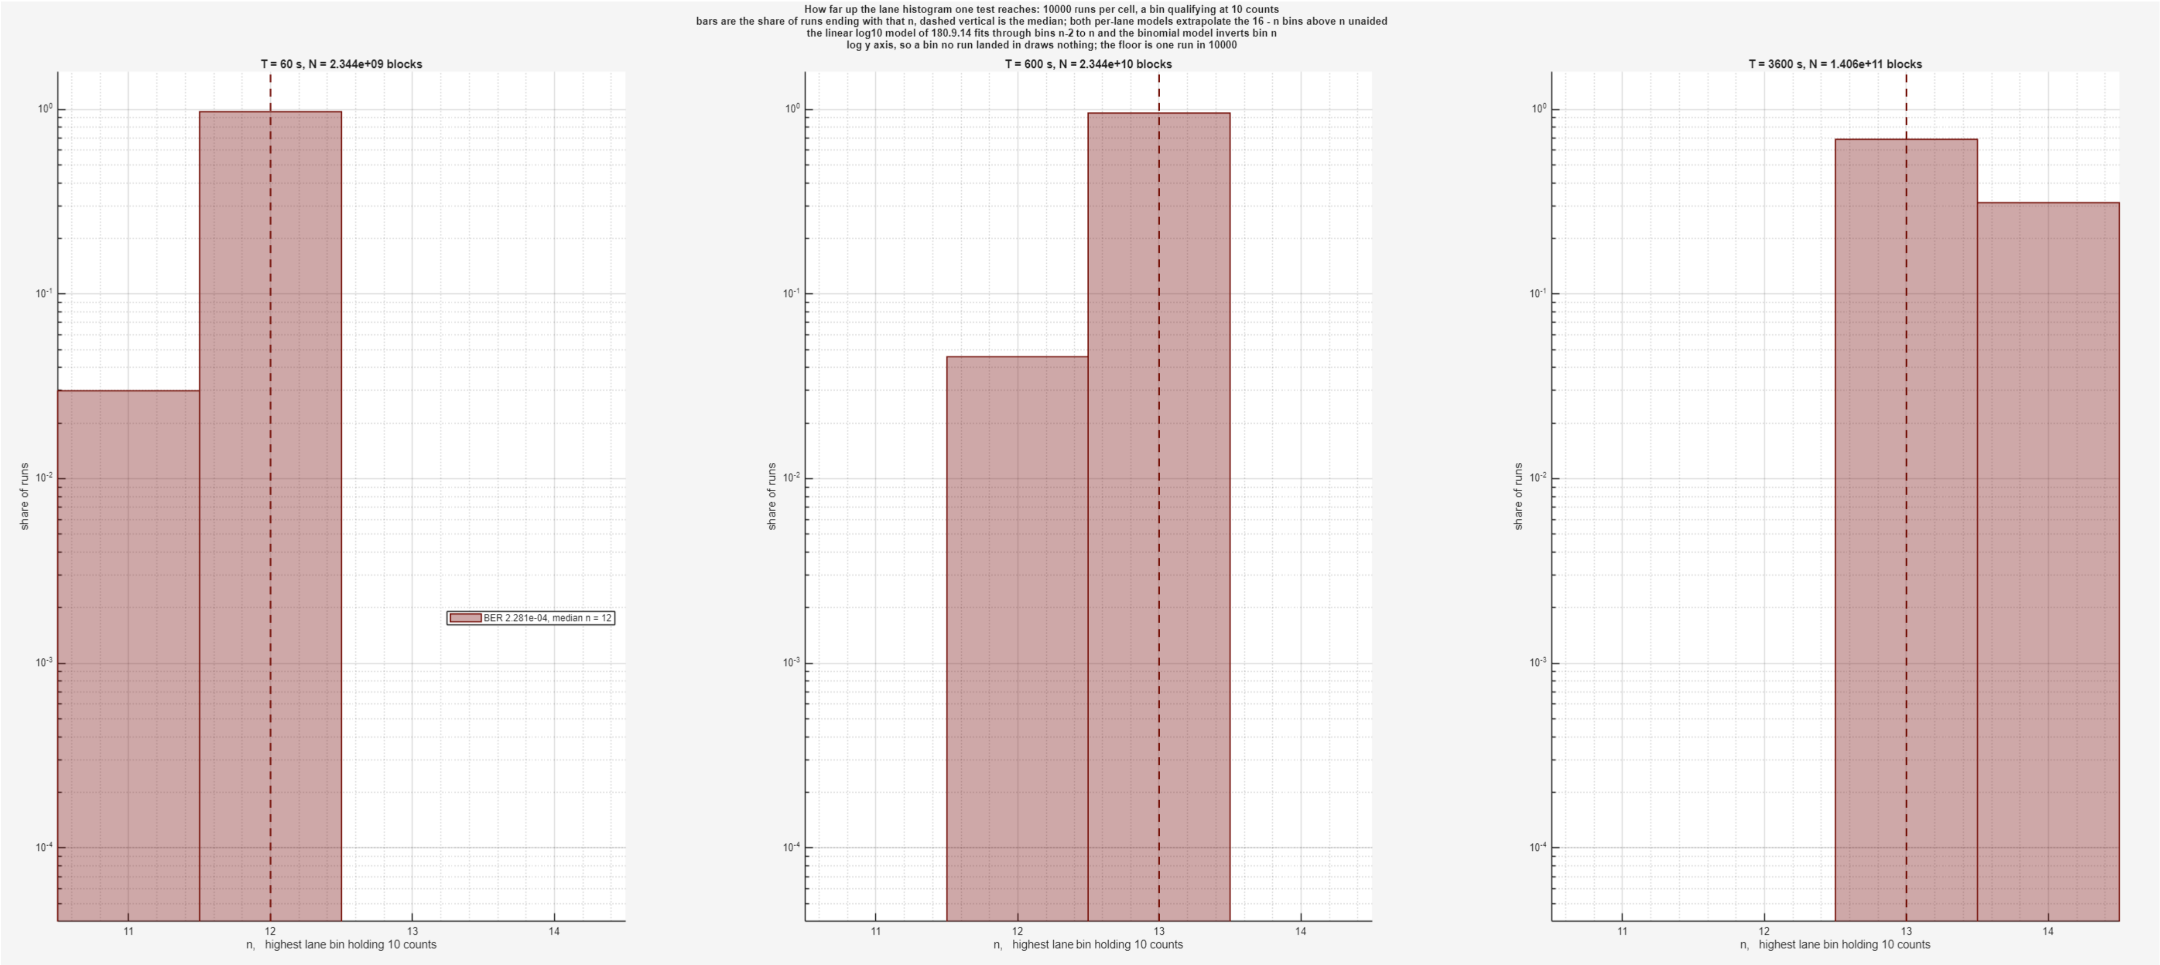

180.9.14 log10 linear model distribution (4)

Explanation of step in the CCDF

- The step in the CCDF (Complementary CDF = probability of X value and above/worse) for duration of 60[Sec] and for 600[Sec] results from having dichotomic partitioning of the max nonzero bin being used, where each of the 2 max nonzero bin has quite different bias.

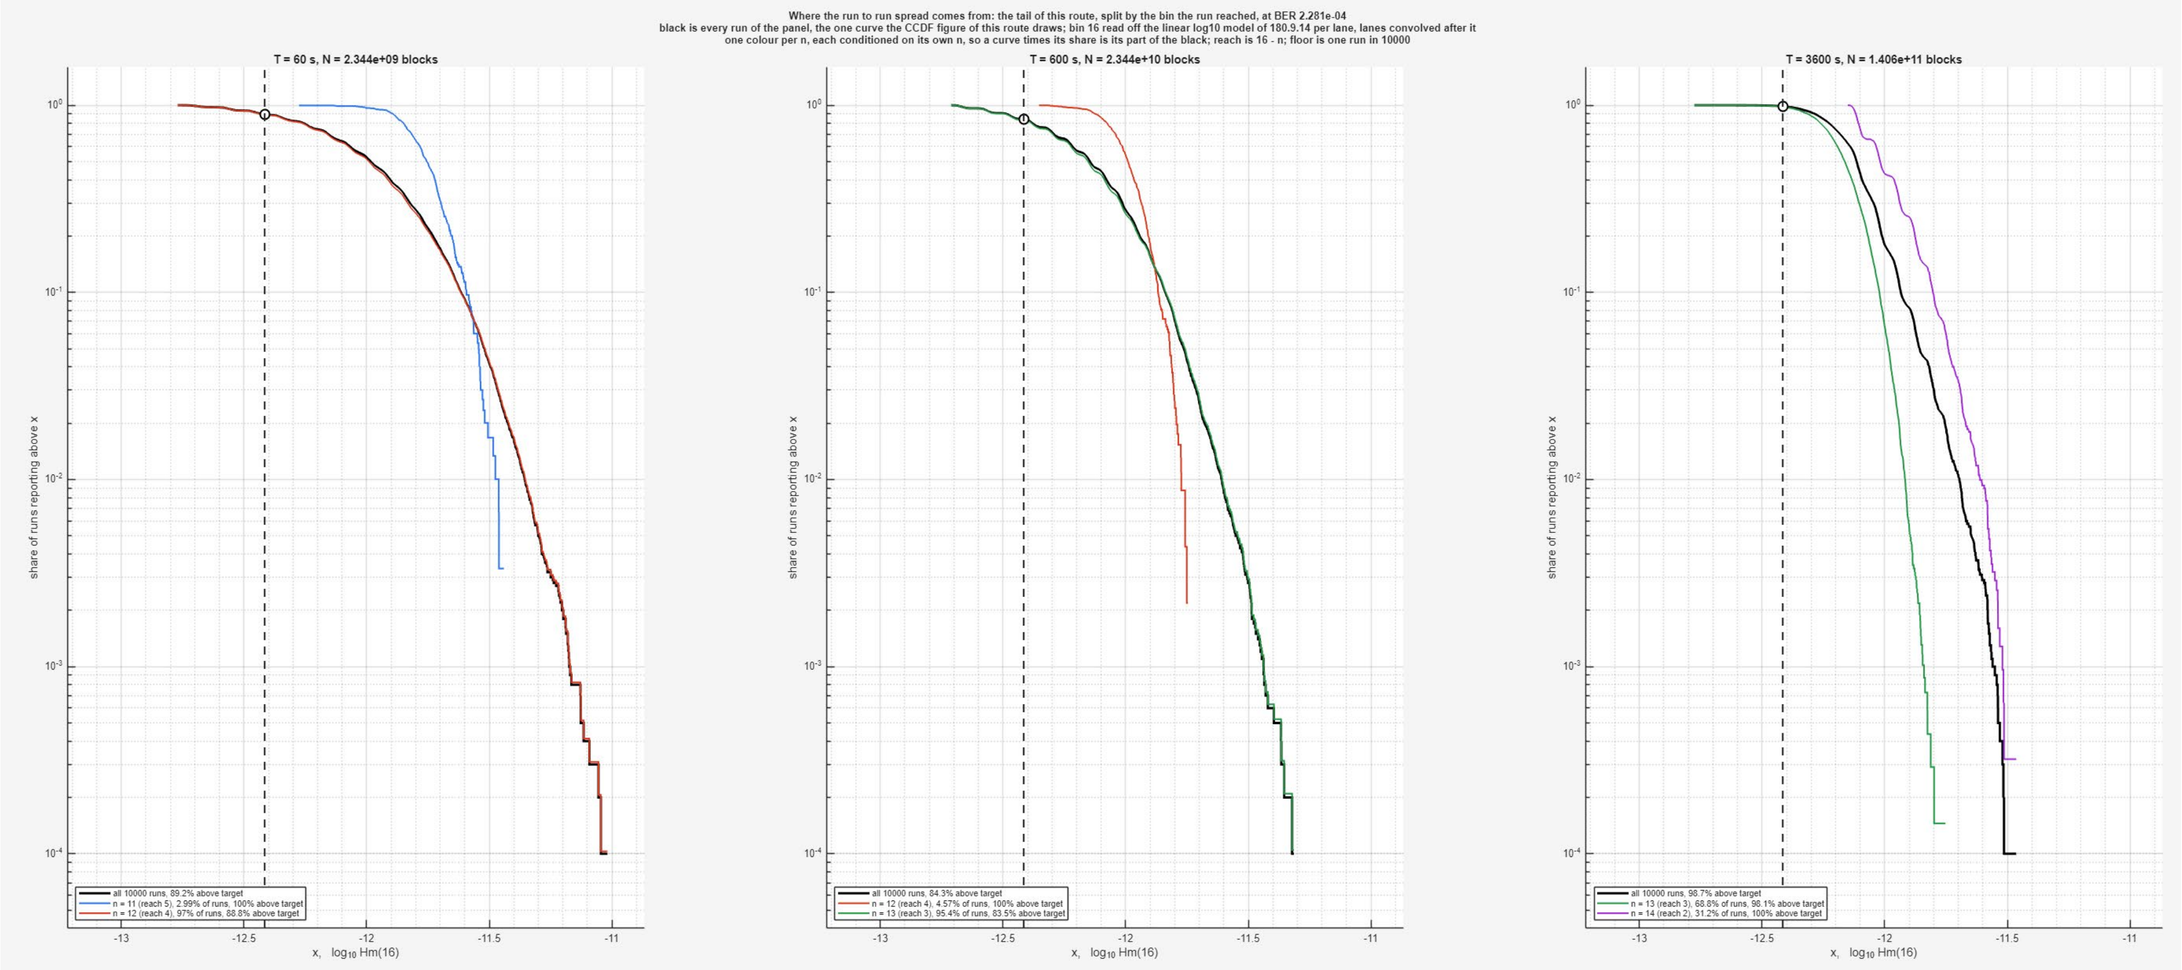

The Monte Carlo Analysis for P-lanes (1)

- Same flow as for 1X, only preceded by convolving the P-lanes bins values

p lanes: convolve first, extrapolate after

the limit is on the codeword, so there is nothing to extrapolate to until the lanes are one again

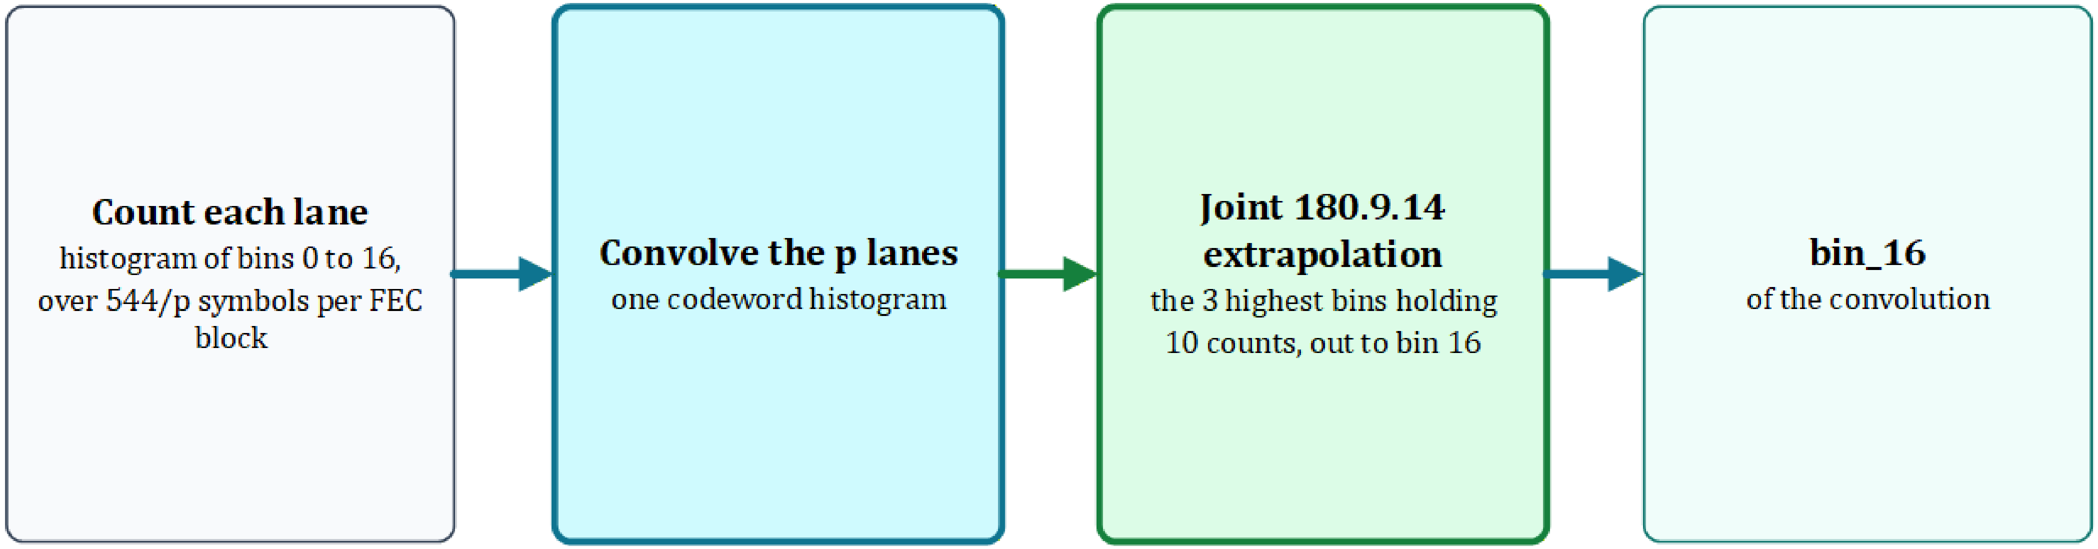

The Monte Carlo Analysis for P-lanes (1)

Numerical example – figure

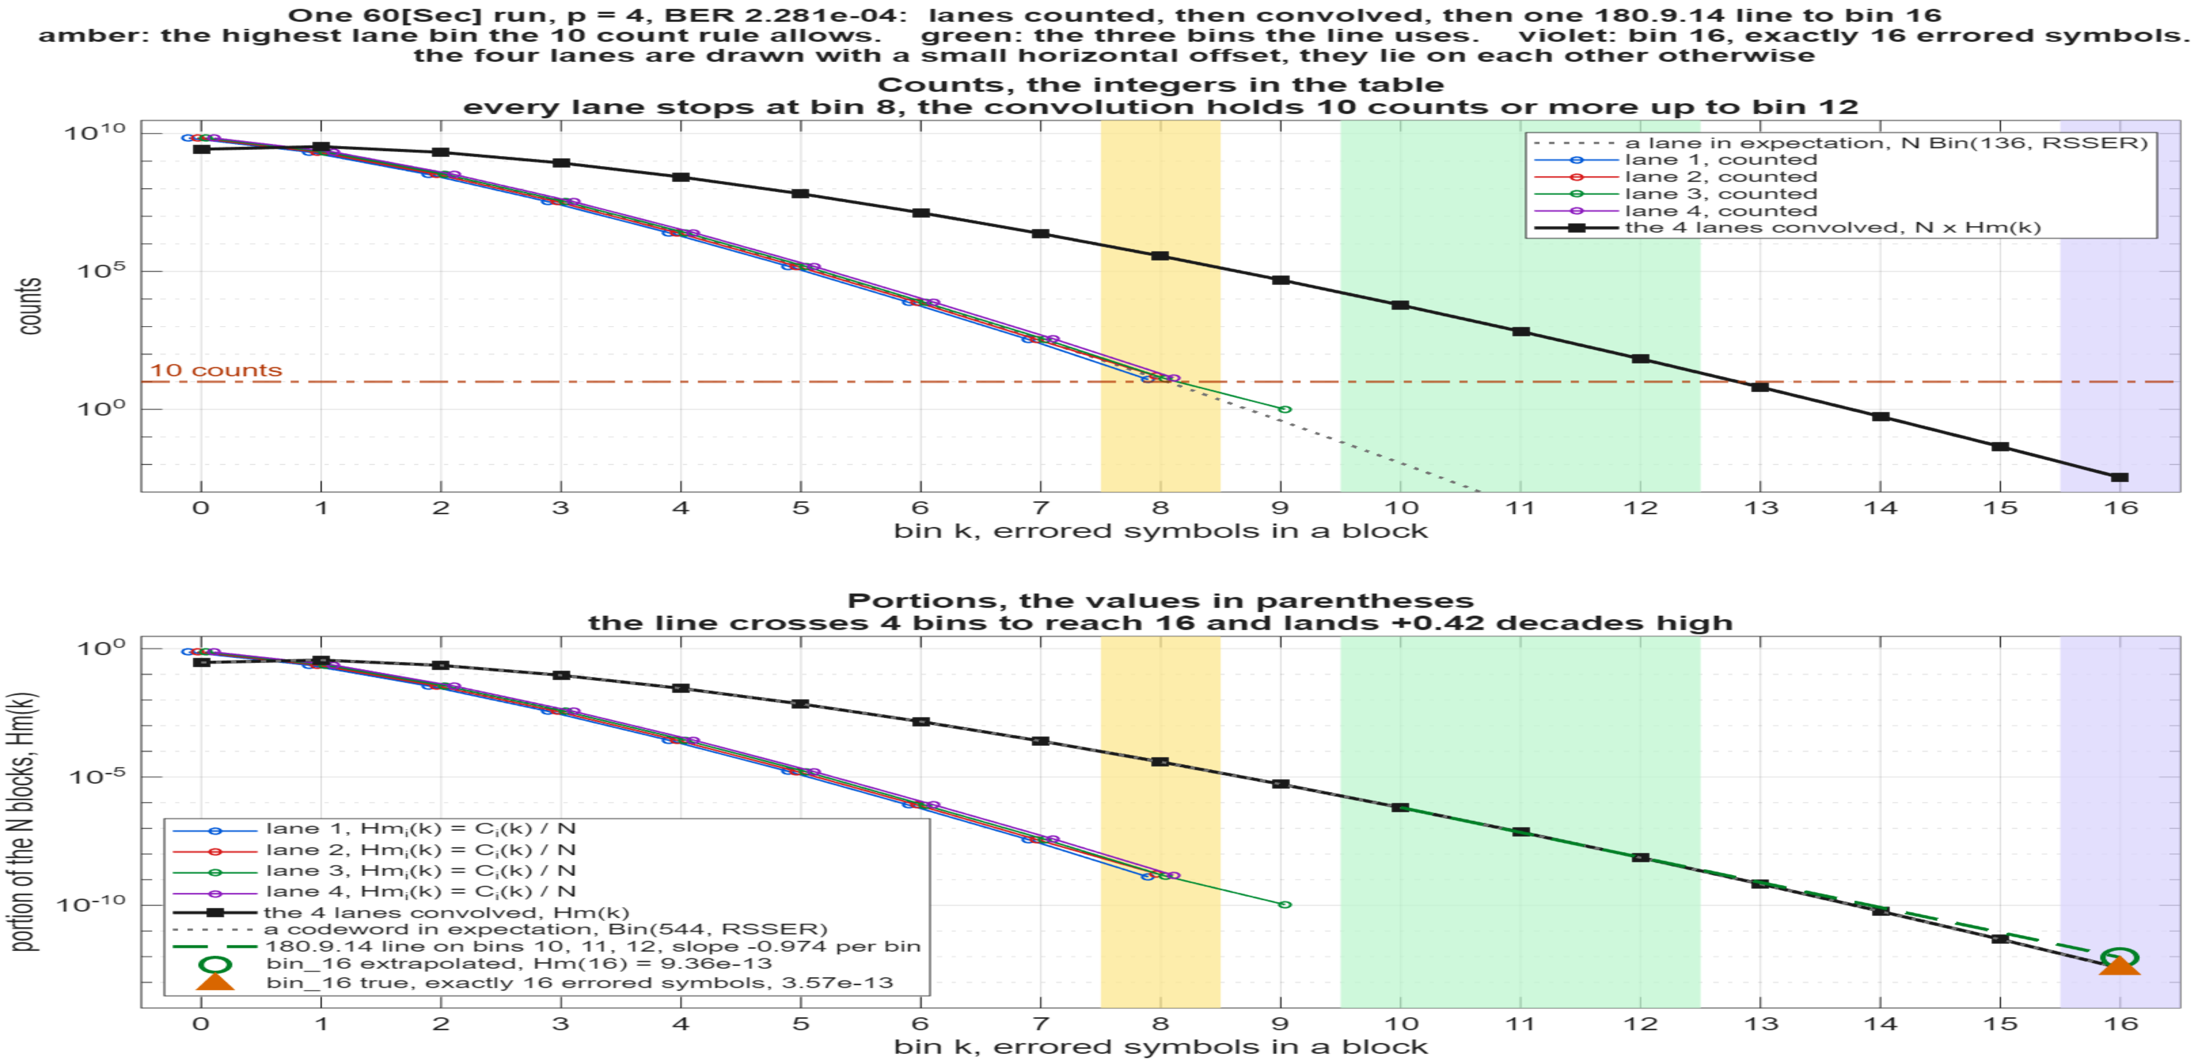

The Monte Carlo Analysis for P-lanes (2)

Numerical example - table

One 60[Sec] run, p = 4, BER 2.281e-4

the counts one draw put in every bin of every lane, the convolution of the four lanes, and the bin 16 the line gives

136 symbols per lane block, RSSER 2.279e-3, N = 9,375,000,000 blocks per lane

bin	lane 1	lane 2	lane 3	lane 4	convolution
0	6,874,207,395 (0.733249)	6,874,167,746 (0.733245)	6,874,141,030 (0.733242)	6,874,076,591 (0.733235)	2,709,952,832 (0.289062)
1	2,135,296,074 (0.227765)	2,135,302,482 (0.227766)	2,135,337,976 (0.227769)	2,135,366,363 (0.227772)	3,367,184,898 (0.359166)
2	329,194,545 (0.0351141)	329,225,893 (0.0351174)	329,217,711 (0.0351166)	329,255,323 (0.0351206)	2,088,080,383 (0.222729)
3	33,589,038 (0.00358283)	33,592,635 (0.00358321)	33,590,620 (0.003583)	33,586,918 (0.0035826)	861,666,835 (0.0919111)
4	2,550,531 (0.000272057)	2,549,290 (0.000271924)	2,551,121 (0.00027212)	2,552,887 (0.000272308)	266,191,160 (0.0283937)
5	154,425 (1.6472e-5)	153,885 (1.64144e-5)	153,756 (1.64006e-5)	153,807 (1.64061e-5)	65,665,575 (0.00700433)
6	7,638 (8.1472e-7)	7,711 (8.22507e-7)	7,425 (7.92e-7)	7,737 (8.2528e-7)	13,474,055 (0.00143723)
7	342 (3.648e-8)	342 (3.648e-8)	347 (3.70133e-8)	360 (3.84e-8)	2,365,435 (0.000252313)
8	12 (1.28e-9)	16 (1.70667e-9)	13 (1.38667e-9)	14 (1.49333e-9)	362,699 (3.86879e-5)
9	0 (0)	0 (0)	1 (1.06667e-10)	0 (0)	49,349 (5.26393e-6)
10	0 (0)	0 (0)	0 (0)	0 (0)	6,034 (6.43574e-7)
11	0 (0)	0 (0)	0 (0)	0 (0)	670 (7.14277e-8)
12	0 (0)	0 (0)	0 (0)	0 (0)	68 (7.25707e-9)
13	0 (0)	0 (0)	0 (0)	0 (0)	6 (6.7968e-10)
14	0 (0)	0 (0)	0 (0)	0 (0)	1 (5.90088e-11)
15	0 (0)	0 (0)	0 (0)	0 (0)	0 (4.77142e-12)
16	0 (0)	0 (0)	0 (0)	0 (0)	0 (3.60634e-13)

The lanes reach bin 8

bin 8 holds 12, 16, 13, 14

bin 9 holds 0, 0, 1, 0

nothing above it

The convolution reaches 12

N x Hm(k) = 6034, 670, 68

in bins 10, 11 and 12

bin 13 gives 6, under 10

One 180.9.14 line

log10 Hm(k) on those bins

slope -0.974 per bin

four bins out to 16

bin_16, Hm(16) = 9.4e-13

exactly 16 errored symbols

N x Hm(16) = 0.009 counts

true 3.6e-13, +0.42 decades

Portions, Hm(k), in parentheses under every count, N x Hm(k)

Lane i is counted, and its portion of the N blocks is Hm_i(k) = C_i(k) / N, so lane 1 at k = 0 holds 6,874,207,395 / 9,375,000,000 = 0.733249.

The p = 4 lanes convolve into the codeword portion Hm(k) = SUM Hm_1(k1) Hm_2(k2) Hm_3(k3) Hm_4(k4) over k1 + k2 + k3 + k4 = k, read back as N x Hm(k).

k = 0, the one route, all four lanes empty: Hm(0) = 0.733249 x 0.733245 x 0.733242 x 0.733235 = 0.289062, N x Hm(0) = 2,709,952,832.

k = 1, the four routes, the single error in lane i: Hm(1) = 0.227765 x 0.733245 x 0.733242 x 0.733235 + the 3 alike = 0.359166, N x Hm(1) = 3,367,184,898.

green: the bins the line uses. violet: bin 16, exactly 16 errored symbols, never populated. amber: the highest lane bin the 10 count rule allows.

For p-lanes: 180.9.14 log10 linear model distribution (3)

Extremely narrowed

Target BER=2.28e-4, leading to BLER=3.84e-13, vs. BER=2.0e-4 leading to BLER=5.42e-14 (-85.9%)

- The figure below shows that for 4-lanes there is much narrower scatter vs. per single lane
- Reason is that 100% of the cases get to same highest bin post convolution

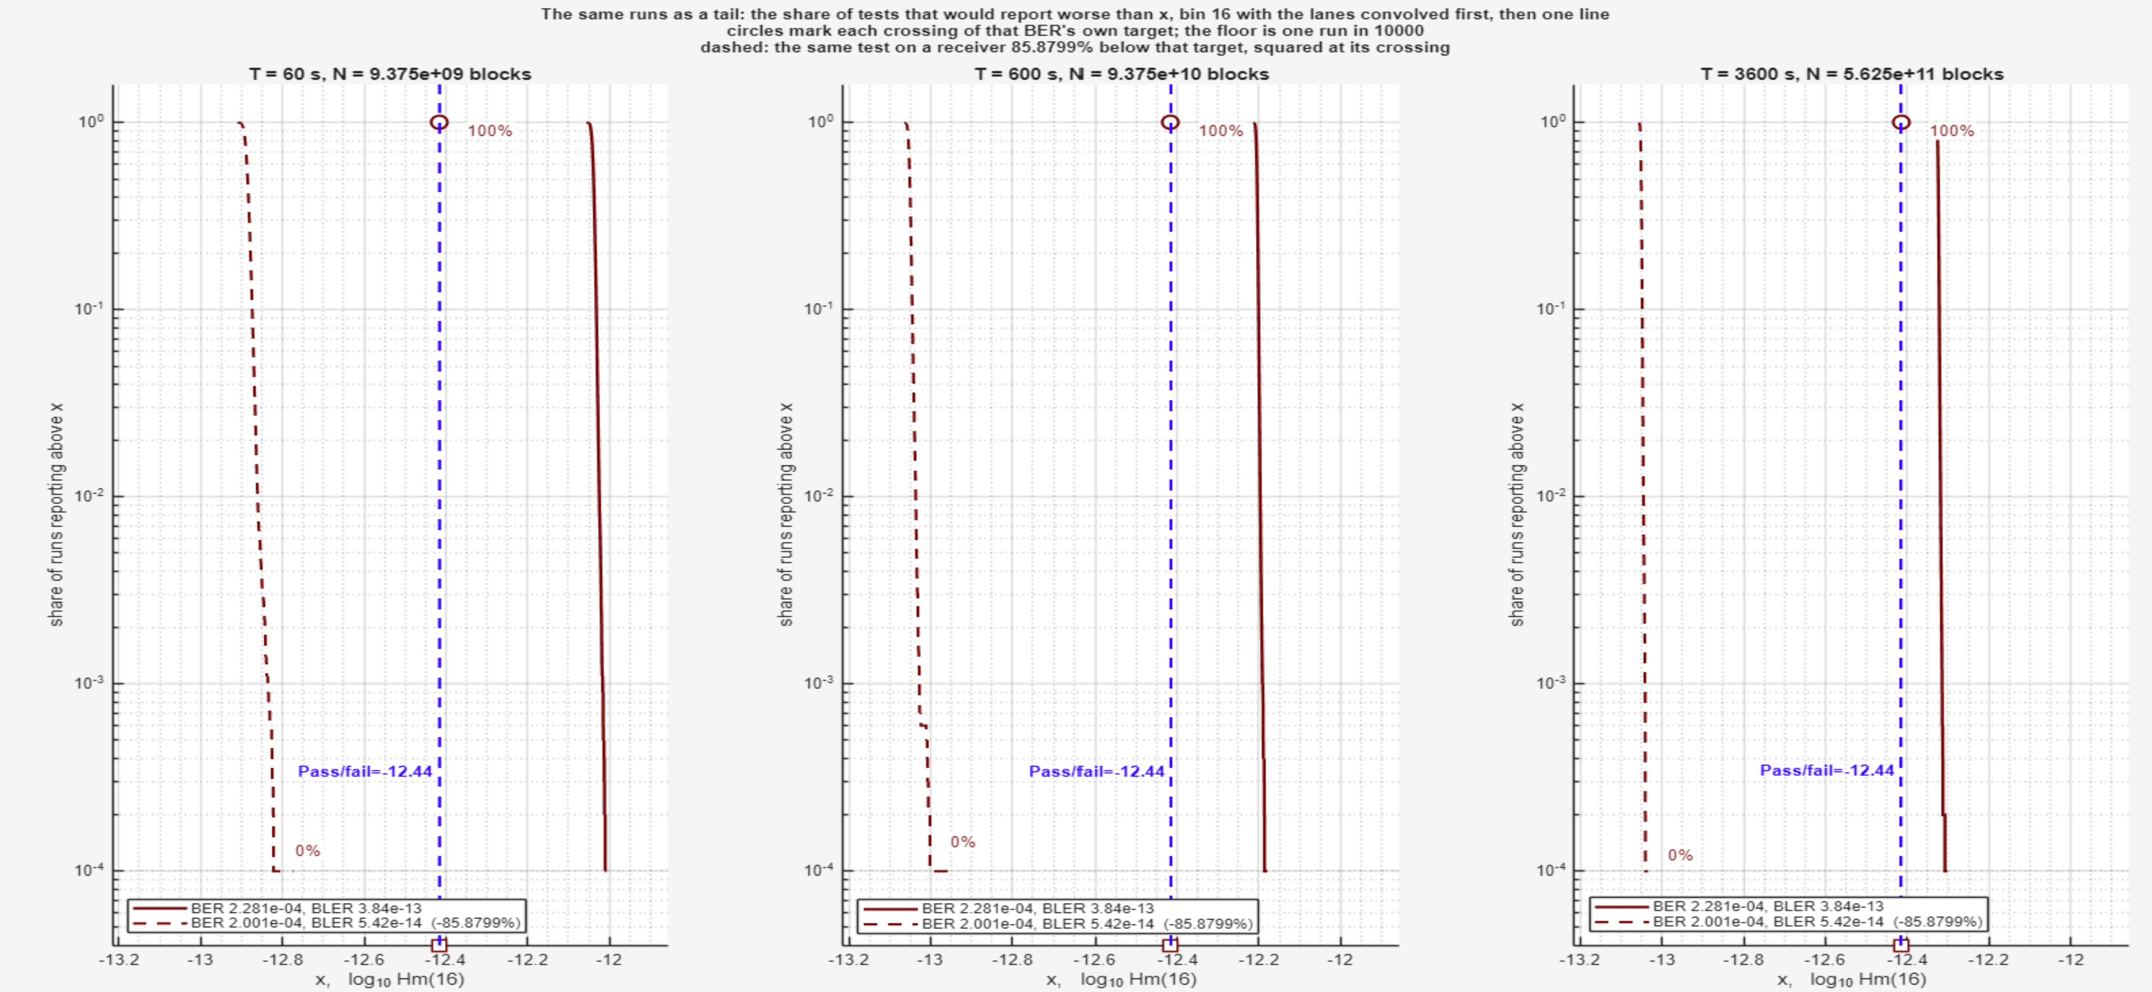

For p-lanes: 180.9.14 log10 linear model distribution (4)

Target BER=2.28e-4, leading to BLER=3.84e-13, vs. BER=2.0e-4 leading to BLER=5.42e-14 (-85.9%)

- The figure below shows that for 4-lanes there is much narrower scatter vs. per single lane
- Reason is that post convolution, the highest bin is fixed over the 10,000 draws

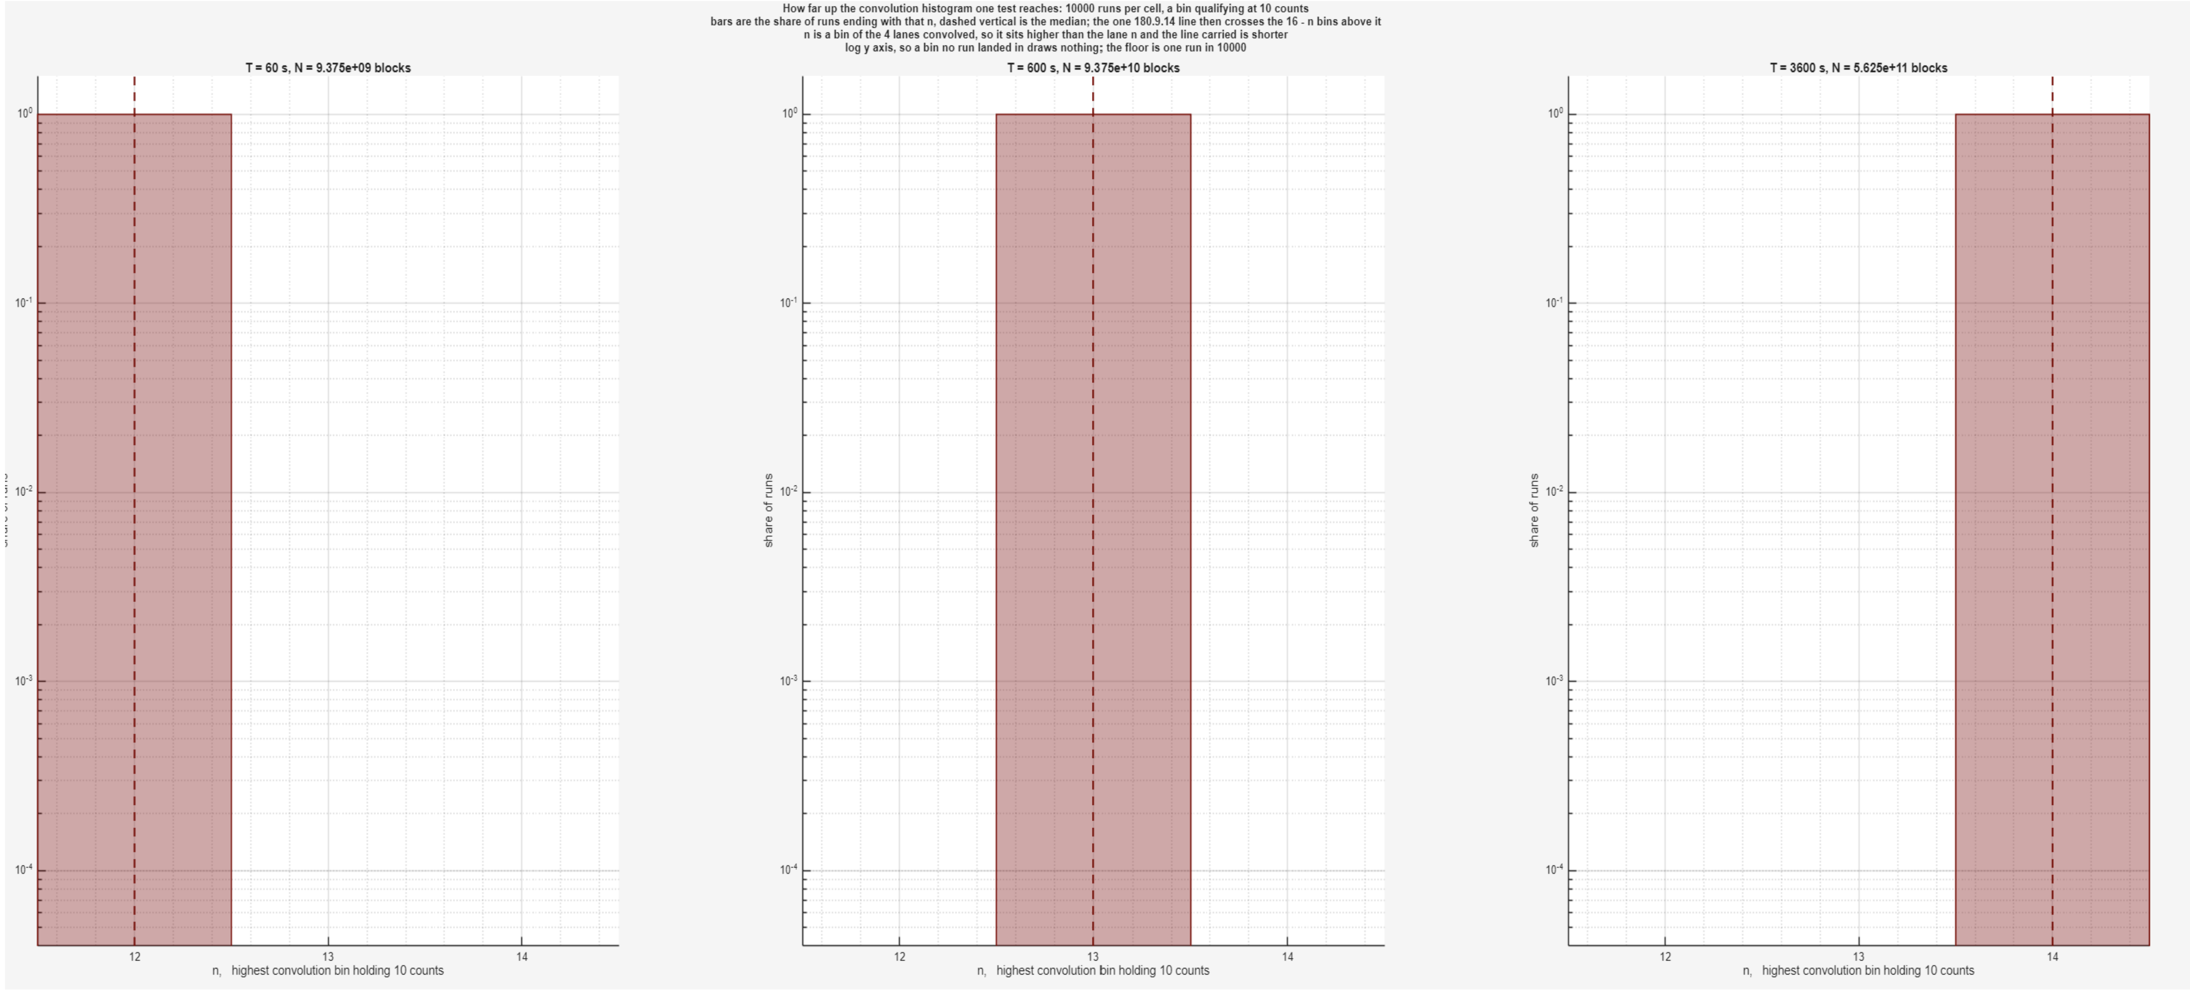

Correlation impact on model bias (1)

- As correlation of RS symbol errors increases the model bias is reduced
- With the used model* the correlation leading to negative bias (false pass) is only for failing links

Max bin >=10	Fixed Burstiness**=1.00		Fixed Burstiness**=1.05		Fixed Burstiness**=1.1	
	Extrapolated Bin 16	Model bias	Extrapolated Bin 16	Model bias	Extrapolated Bin 16	Model bias
11	1.57e-12	0.637	2.91e-13	-0.094	6.16e-14	-0.769
12	9.42e-13	0.415	2.95e-13	-0.088	1.02e-13	-0.551
13	6.42e-13	0.245	3.09e-13	-0.068	1.59e-13	-0.359
14	4.87e-13	0.129	3.27e-13	-0.043	2.28e-13	-0.200
15	4.05e-13	0.049	3.46e-13	-0.019	3.00e-13	-0.082

*Assuming Equivalent RSSER of 2.28e-3 for bin 16, and rising to this RSSER with fixed burstiness from the bins used for extrapolation
**Definition: Burstiness = Equivalent RSSER of bin (n) / Equivalent RSSER of bin (n – 1)

Correlation impact on model bias

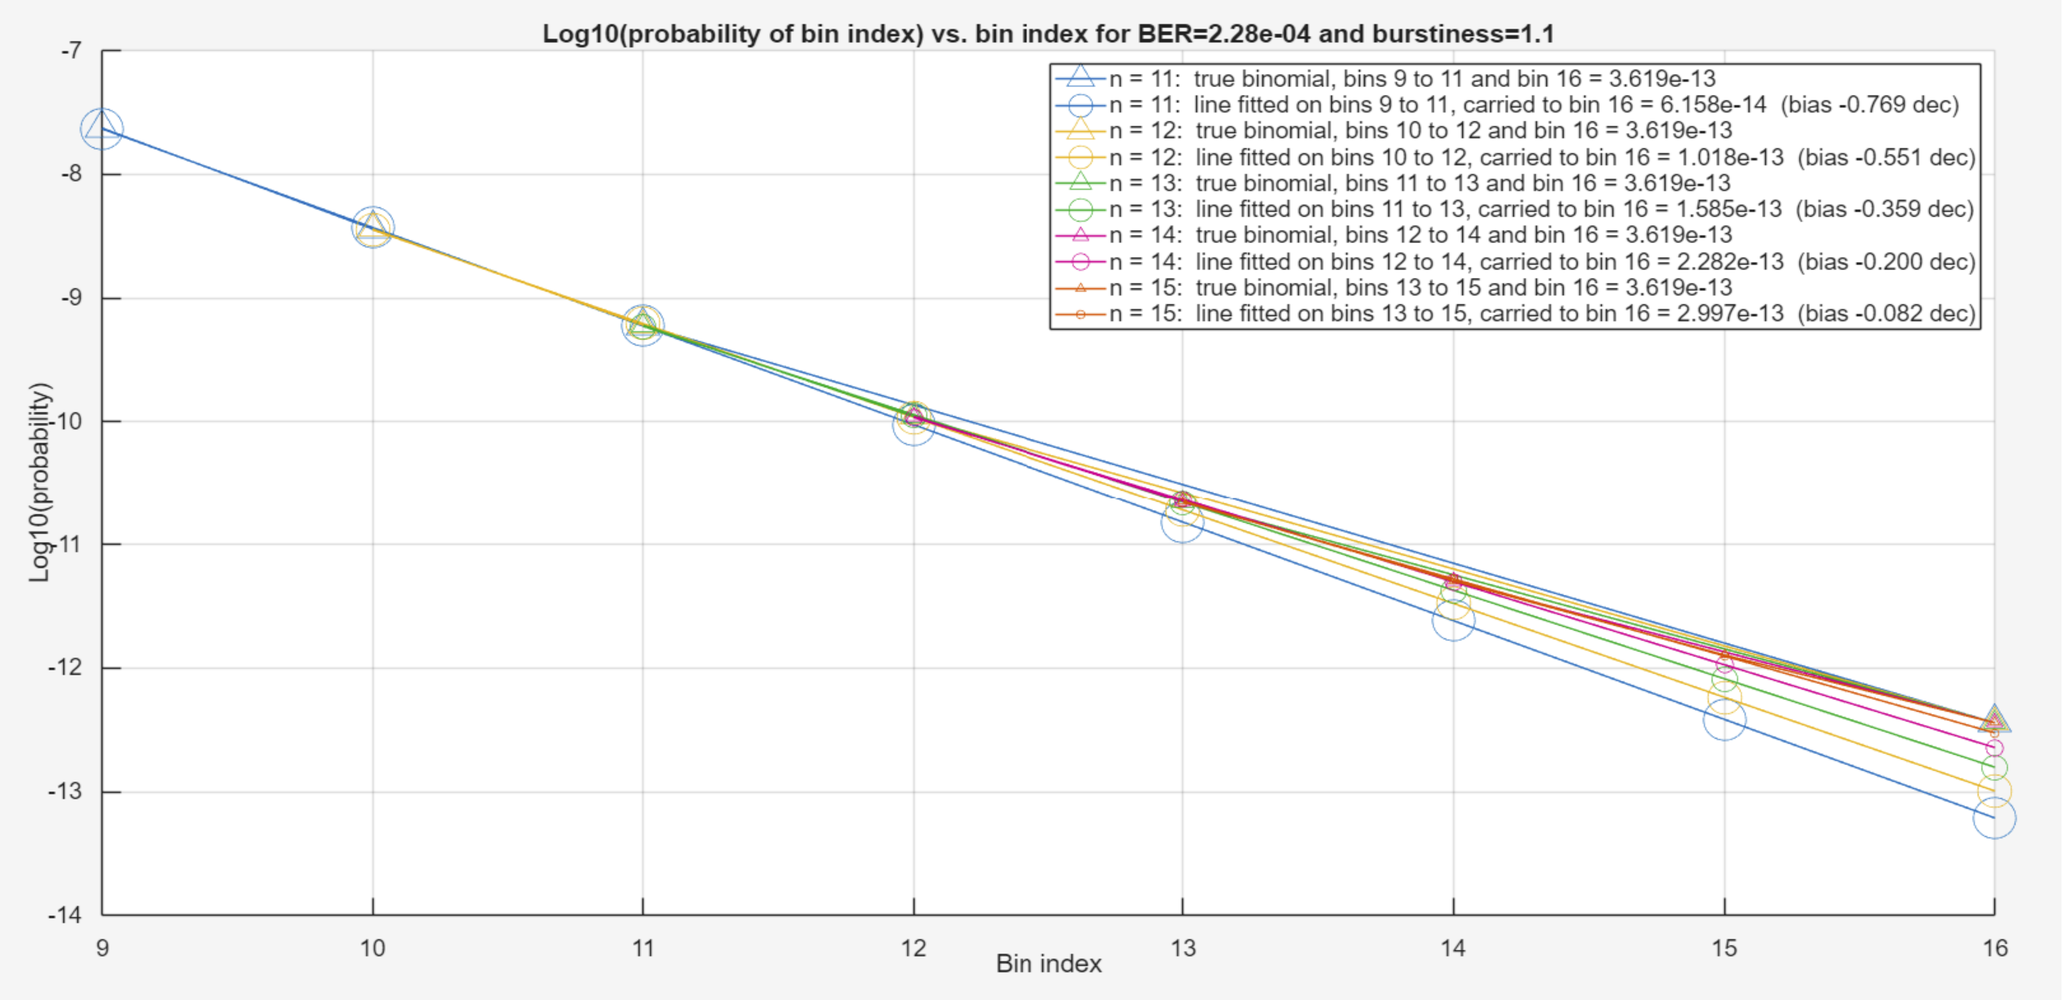

*Assuming Equivalent RSSER of 2.28e-3 for bin 16, rising to this RSSER with fixed burstiness from the lowest bin used for extrapolation
**Definition: Burstiness = Equivalent RSSER of bin (n) / Equivalent RSSER of bin (n – 1)