8802-3/802.3 REVISION REQUEST 1118

DATE: 31st Oct, 2003
NAME: Eric Lynskey
COMPANY/AFFILIATION: UNH-IOL
E-MAIL: elynskey@iol.unh.edu

REQUESTED REVISION:
- STANDARD: IEEE Std. 802.3ae-2002
- CLAUSE NUMBER: 48.2.4.2
- CLAUSE TITLE: Idle ||I||

PROPOSED REVISION TEXT:

Within bullet item (d) remove the word uniform, allowing the sentence to read "...r is a randomly distributed number between 16 and 31..."

For details see attached draft.

RATIONALE FOR REVISION:

A bit that chooses between ||K|| and ||R|| is produced during every clock cycle. The counter that determines when it is time to send an ||A|| is only evaluated and reset every 16-31 clock cycles. When the A_CNT variable reaches zero, an ||A|| is sent and a new value is loaded into the counter. The four bits loaded into the counter can take on any value from zero through fifteen, and these are combined with an MSB that is always high, thus initializing the counter to 16-31. This counter will then start to count down, decrementing once for every clock cycle. While this counter is decrementing, the pseudo-random generator is continuing to shift through its 127-bit cycle, and ||R|| and ||K|| codes are being transmitted. Thus, the bits being fed to the A counter are not being fed continuously, but only after the expiration of the current A counter.

If a separate pseudo-random generator existed that only provided a new value upon the expiration of A_CNT, then the next sequence would get shifted into the counter and no possible combinations would be missed. If the pseudo-random generator generated its first four output states as W1, W2, W3, and W4 (where each Wn is 4-bits in length and takes on a decimal value between 0 and 15, inclusive) then the A counter would be initialized with W1+16. After A_CNT expired, which would be W1+16 clock cycles, the new value of the counter would be set to W2+16. This same process would then continue through all 127 different combinations: W4...W127.

However, this is not what happens within a device with a single random number generator. If there is a single pseudo-random generator, the A_CNT will only get initialized after the previous A_CNT has expired. Since the pseudo-random generator is continuously running to pick between ||R|| and ||K||, the A_CNT will not be initialized with state WX+1 if the

current state is WX, as was shown in the previous paragraph. If A_CNT is originally initialized with WX, then the next state would be W(WX+16). By not sampling the pseudo-random generator after every clock cycle, not every possible value output by the pseudo-random generator will be observed.

The plots on show simulations of what the ||A|| spacing would actually look like when the pseudo-random generator is sampled in this manner. The first plot shows the distribution of ||A|| spacing for the X^7+X^6+1 polynomial, and the second plot uses the X^7+X^3+1 polynomial. The x-axis shows the number of columns between consecutive ||A|| columns. The y-axis shows the number of times that ||A|| spacing existed. For each simulation, 4096 trials were done. It is clear that after several thousand trials, there are many spacing values that have not been used, thus making it impossible to have a uniformly distributed spacing.

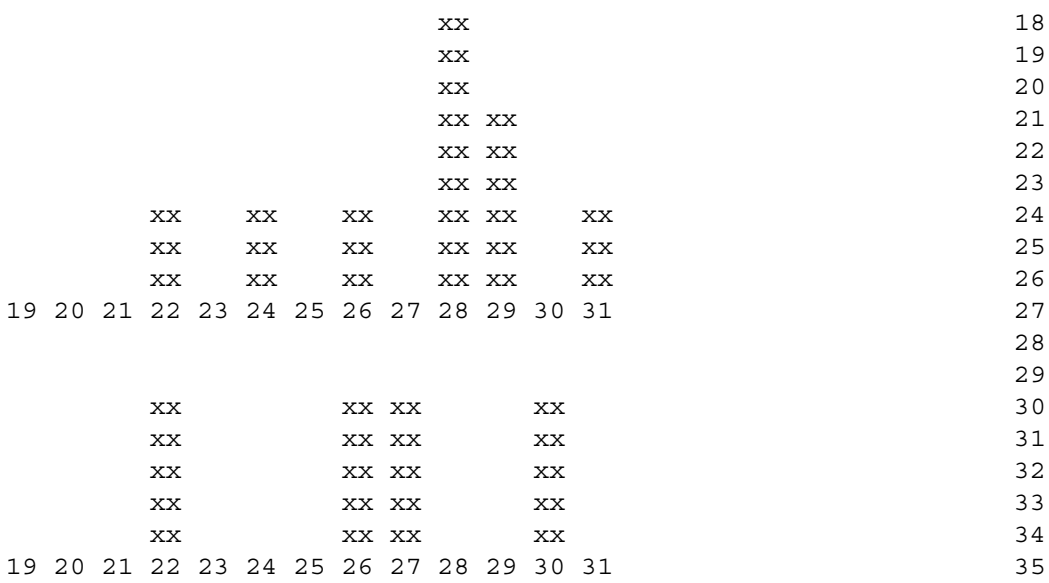

IMPACT ON EXISTING NETWORKS:

The standard currently defines that the A spacing be uniform randomly distributed between 16 and 31. A brief survey of 6 implementations of this PCS showed that 4 had a distribution similar to one of the two plots that are shown above. One had a uniform distribution, and one had a distribution that was neither uniform nor fitting the above plots. Removing the word "uniform" should allow for both uniform and non-uniform implementations to exist provided that they are based on one of the two defined polynomials.

Please attach supporting material, if any	
Submit to:- Bob Grow, Chair IEEE 802.3	
E-Mail: Bob.Grow@intel.com	
For official 802.3 use	
REV REQ NUMBER: 1118	
DATE RECEIVED: 31st Oct, 2003	
~~EDITORIAL~~/TECHNICAL	
ACCEPTED/DENIED	
BALLOT REQ'D YES/~~NO~~	
COMMENTS: 05-Jul-04 Ver: D1.2 Status: B	

For information about this Revision Request see -

http://www.ieee802.org/3/maint/requests/revision_history.html#REQ1118

48.2.4.2 Idle (||I||)

Change the item d) of subclause 48.2.4.2 as follows:

d) Each ||A|| is sent after r non-||A|| columns where r is a ~~uniform~~ randomly distributed number between 16 and 31, inclusive. The corresponding minimum spacing of 16 non-||A|| columns between two ||A|| columns provides a theoretical 85-bit deskew capability.