10 Mb/s Single Twisted Pair Ethernet
10BASE-T1L and Clause 98

Steffen Graber

Pepperl+Fuchs

Content

- Single Twisted-Pair Auto-Negotiation

- 10BASE-T1L and Auto-Negotiation

- Matlab/Simulink Model

- Simulation Results

- Timing Values

- Compatibility (Data Rate)

- Compatibility (Signal Amplitude)

- Limitations

- Summary

Single Twisted-Pair Auto-Negotiation

- Clause 28 specifies standard auto-negotiation for 4-wire twisted-pair Ethernet using fast link pulses and is therefore not really useful for 2-wire auto-negotiation.
- Single twisted-pair auto-negotiation is described in Clause 98 for 100BASE-T1 and 1000BASE-T1.

- During the auto-negotiation process telegrams consisting of a start delimiter, a 48 bit payload field, a 16 bit CRC and a single bit end delimiter are transmitted in half-duplex.

- The data rate is 16.667 MBit/s using a DME encoding (a logic „1" is transmitted, if there is an edge between two clock transitions, while a logic „0" is transmitted with no edge between two clock transitions).
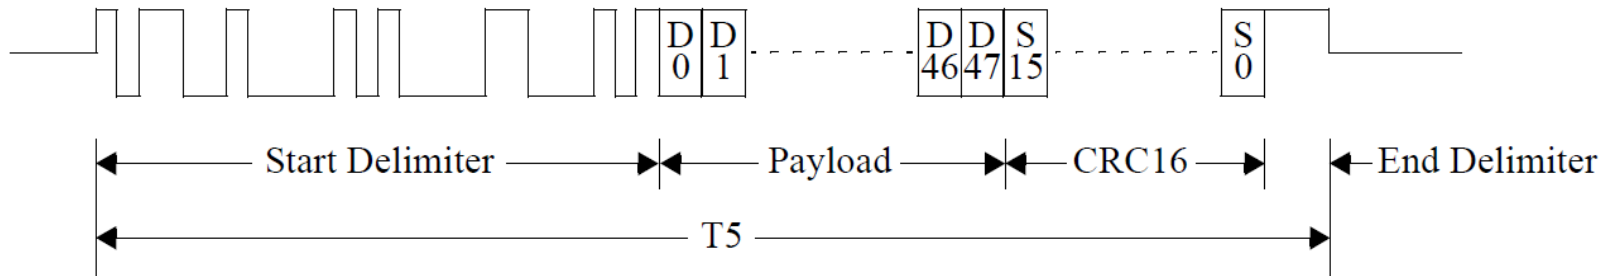

- For a long reach link segment the auto-negotiation communication speed needs to be reduced.
- 16.667 Mbit/s would cause an IL of 53 dB over a 1000 m link segment as specified for a 10BASE-T1L PHY, which is significantly too high.

- Due to the longer link segment length and the lower communication speed, the timer values controlling bit timing, frame timing, flow control and blanking time need to be adjusted.

10BASE-T1L and Auto-Negotiation

- The following slides show a Matlab/Simulink model and some simulations for a cable length of approx. 1500 m AWG18/1 cable.
- This is to add some margin to the intended 1000 m reach for the auto-negotiation process.
- The goal is to find a symbol rate for the auto-negotiation process, which allows an easy decoding without the need for a trained equalizer within the receive path of the PHY.

- The signal frequency for the auto-negotiation communication needs to be chosen so that it is working in conjunction with the power and signal coupling network but also in conjunction with the input filters of the PHY.

- A 200 kHz high pass filter at the input of the PHY has a significantly higher influence on the auto-negotiation communication than the power coupling network, as this is the high pass filter with the highest corner frequency within the system.

- Therefore there is need to find a baud rate, which fits to the requirements of the high-pass filter but also to the insertion loss requirements of the link segment and which can be easily oversampled by the integrated 7.5 or 15 MSPS/s ADC (not a hard requirement, but it could simplify the implementation).

- Assuming very short link segments, the 200 kHz high pass filter has the highest influence on the received signal, therefore a simulation for a short link segment is also provided.

Matlab/Simulink Model
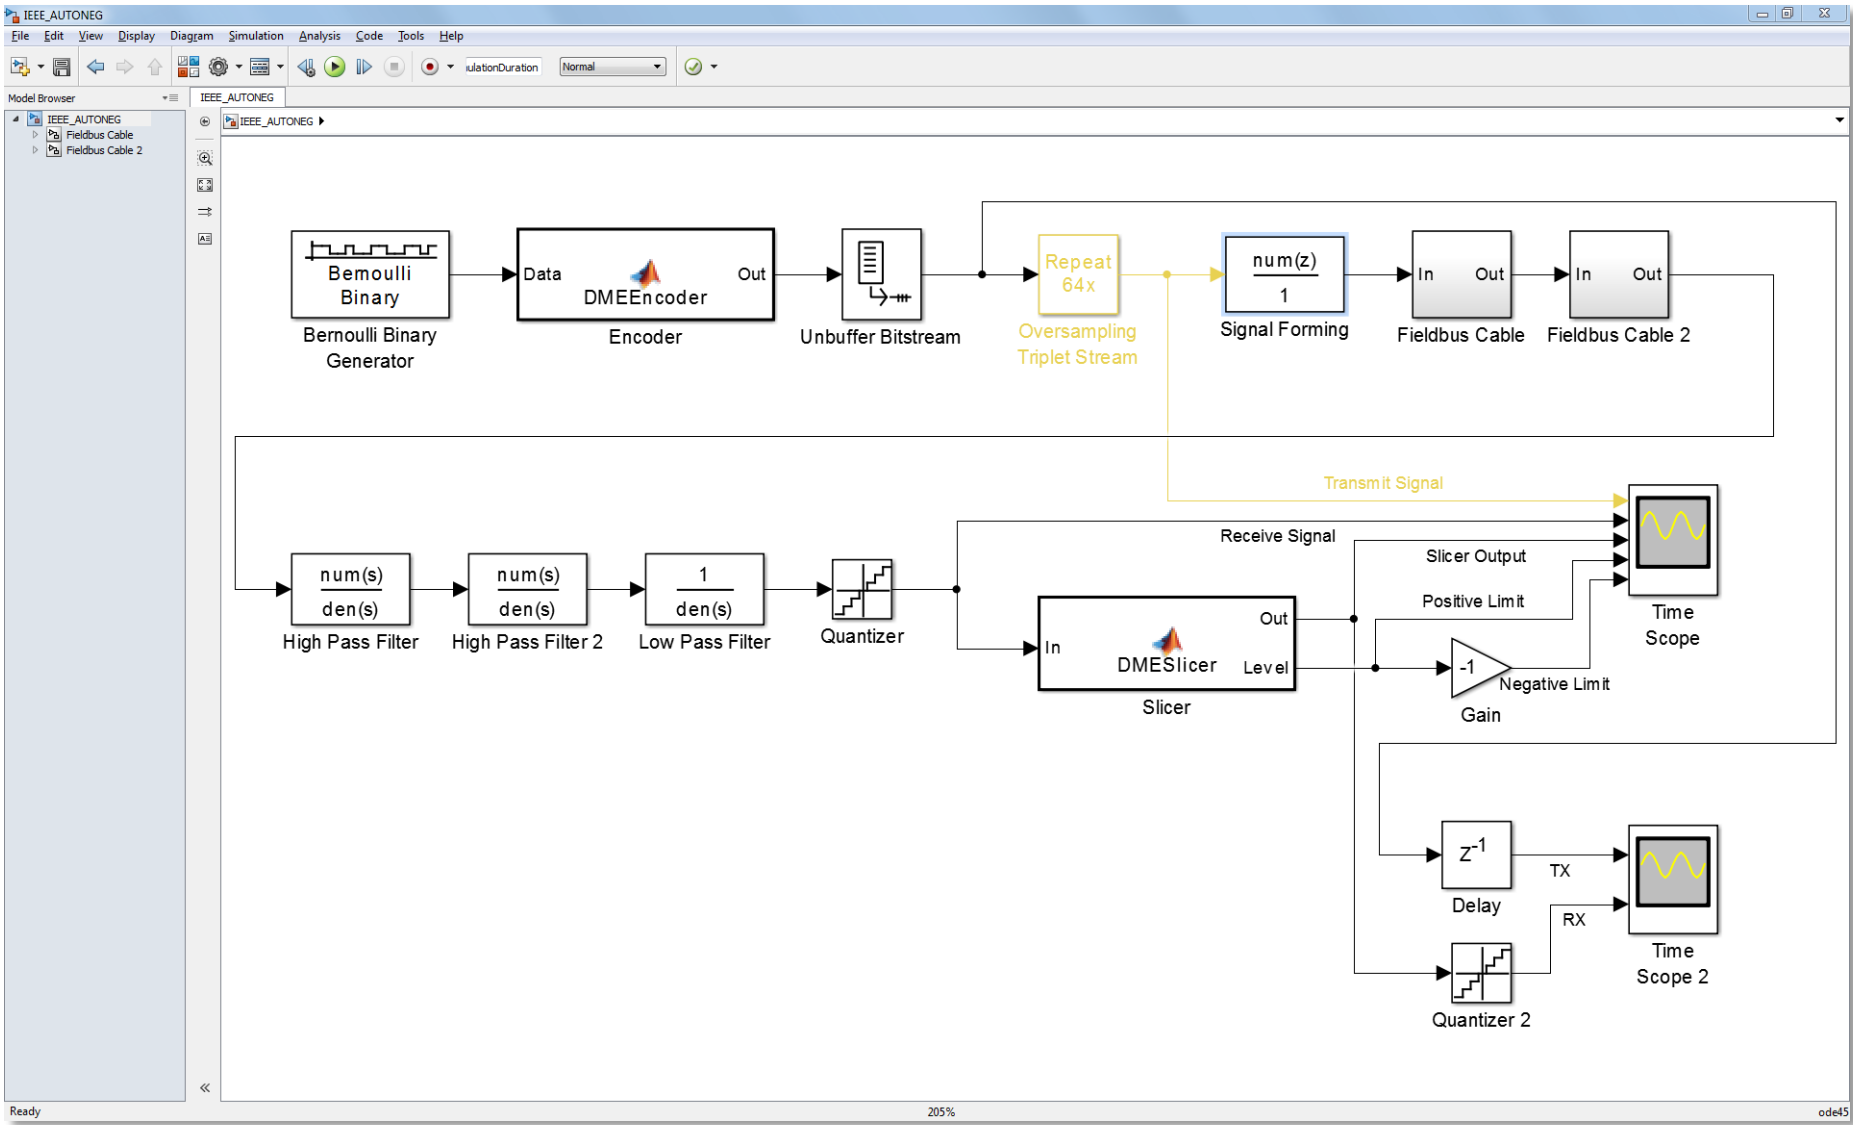
Simulation Results

- 312.5 kBit/s baud rate, 0 m (left) and 1500 m (right), due to the 200 kHz high pass filter significant overshoots can be seen.
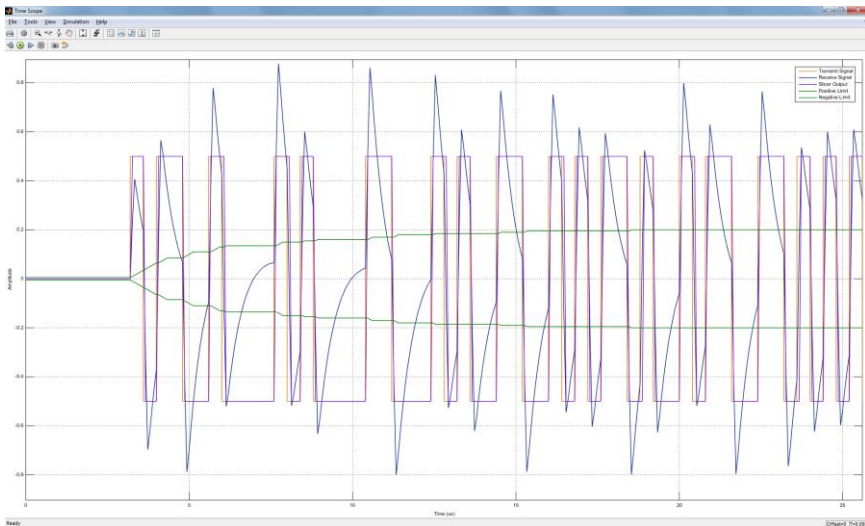
- 1250 kBit/s baud rate, 0 m (left) and 1500 m (right), for 1500 m there is a significant attenuation of the auto-negotiation signal.

Simulation Results


- 625 kBit/s baud rate, 0 m (left) and 1500 m (right), seem to be a good compromise.


- 652 kBit/s baud rate, 1000 m typical AWG18/1 cable (left) and decoded signal (right).

Simulation Results

- Overall a baud rate of **625 kBit/s** seems to be most suitable for a 10BASE-T1L link segment, even without an equalizer at the receive side, as the DME coding transmits the information over the edges within the communication signal and these edges have a relatively high frequency content, which goes through the high pass filter.

- As the baud rate works for link segments between 0 m and 1500 m (using typical AWG18/1 cables) there should be enough headroom when using this baud rate in combination with a 1000 m link segment.

- Most critical is the 200 kHz high pass filter, as the baud rate used during auto-negotiation is relatively low.
- Unfortunately the used start delimiter in Clause 98 is causing even lower signal frequencies, as there are several violations of the DME coding rules within the start delimiter.
- Additionally the start delimiter is not DC free, but this is, after processing the receive signal through a 200 kHz high pass filter, not really an issue.

- For an easy adoption of the slicer levels a short preamble (e.g. eight „1" DME symbols) before the start delimiter could make sense, but this would require an additional change in Clause 98.
- Alternatively the slicer levels could be adjusted during the first frame reception and could then be used for the subsequently received frames (in a worst case scenario the first received frame would just be dropped as the start delimiter is not being detected).

Timing Values

- Assuming a baud rate of e.g. 625 kBit/s this results in the following bit and page timings:

	Parameters	Min	Typ	Max	Units
T1	Transmit position spacing (period)	799.96	800	800.04	ns
T2	Clock transition to clock transition	1590	1600	1610	ns
T3	Clock transition to data transition (data = 1)	795	800	805	ns
T4a	+1 to -1 or -1 to +1 transitions in a DME page	79	---	143	---
T4b	0 to ±1 or ±1 to 0 transitions in a DME page	2	2	2	---
T5	DME page width	124793	124800	124807	ns

- Additionally the timer values of the state machines have to be adopted to reflect the longer link segment length (up to approx. 1500 m) and the longer bit times and therefore resulting longer auto-negotiation data packet length.
- A maximum length of 1500 m is chosen to have some margin, as e.g. an AWG14 cable is having a lower attenuation than a AWG18/1 cable thus in principle allowing a longer link segment than 1000 m when using such a cable.

Timing Values

- The following table shows possible timer values (need to be discussed in detail) for a 625 kBit/s communication:

Timer	Min	Max	Unit	Remarks
blind_timer	15000	15900	ns	7.5 times link segment length
break_link_timer	tbd.	tbd.		
clock_detect_max_timer	1680	2000	ns	5 – 25 % more than time T2
clock_detect_min_timer	1200	1520	ns	5 – 25 % less than time T2
data_detect_max_timer	880	1200	ns	10 – 50 % more than time T3
data_detect_min_timer	400	720	ns	10 – 50 % less than time T3
interval_timer	799.96	800.04	ns	800 ns ± 0.005 %
link_fail_inhibit_timer	tbd.	tbd.		
page_test_max_timer	128000	131200	ns	1600 ns instead of 60 ns bit time
receive_DME_timer	143040	147140	ns	1600 ns instead of 60 ns bit time and 7.5 times link segment length
rx_wait_timer	tbd.	tbd.	µs	Time to detect DME frames
silent_timer	15900	16800	ns	7.5 times link segment length

- backoff_timer:
  if T[4] bit is 1 then the duration is set as (143040 ns to 147140 ns) + (random integer from 0 to 15) x (15900 ns to 16800 ns)
  if T[4] bit is 0 then the duration is set as (151215 ns to 155315 ns) + (random integer from 0 to 15) x (15900 ns to 16800 ns)
  (adopted to reflect a bit time of 1600 ns instead of 60 ns and 7.5 times higher link segment length)

Compatibility (Data Rate)

- Due to the significantly lower communication speed of the 10BASE-T1L auto-negotiation process compared to a 100/1000BASE-T1 PHY, even if both versions use the same data encoding **they will not be compatible without additional measures**, especially when taking the coupling networks and receive high pass filters into account.

- The use cases for a 100/1000BASE-T1 short reach PHY and a 10BASE-T1L long reach PHY are quite different, so initially it is not too likely that there will be a need to negotiate between the different standards.
- Nevertheless as 2-wire Ethernet evolves in the future there likely will be a requirement for some kind of universal PHY, which needs to negotiate with both short and long reach PHYs (and even if it is just for e.g. 10BASE-T1L and 10BASE-T1S) using different auto-negotiation speeds.

- Therefore, even if the two auto-negotiation speeds are not compatible, we need to think about measures how to negotiate between PHYs using different auto-negotiation data rates.

- One idea is shown below:

  - A 100/1000BASE-T1 PHY only supports the high speed auto-negotiation scheme, therefore two of these PHYs always negotiate at 16.667 MBit/s without causing an issue.
  - A 10BASE-T1L PHY only supports the low speed auto-negotiation scheme, therefore two of these PHYs always negotiate at e.g. 625 kBit/s without causing an issue.
  - More interesting are the universal 10/100BASE-T1L and 10/100/1000BASE-T1(L) PHYs, which also may try to negotiate with 100/1000BASE-T1, 10BASE-T1L, or future 100BASE-T1L only PHYs.

Compatibility (Data Rate)

- A PHY only supporting one auto-negotiation speed starts auto-negotiation communication immediately.
- The receiver of a universal PHY initially tries to decode high speed and low speed auto-negotiation frames (in parallel or sequential).
- To stay as compatible as possible most reasonable would be for a universal PHY to initially wait, if there are high speed or low speed auto-negotiation frames received from a far end link partner (as a PHY only supporting one auto-negotiation speed starts auto-negotiation communication immediately).
- Depending on the received frame type (high speed or low speed) the auto-negotiation speed of the local universal PHY can be adopted accordingly before starting the own auto-negotiation frame transmission.
- If there are no auto-negotiation frames received from the far end side, there are three possibilities:
    - there is no link partner,
    - the link partner transmits auto-negotiation frames at a fixed speed, which is not suitable for the link segment, or
    - there is also a universal PHY.
- For the first two cases there is no way to establish an auto-negotiation communication, therefore the third case is assumed.
- In such a case the local universal PHY should start the transmission of auto-negotiation frames after a defined period of time.
- As it is assumed that there is a universal PHY on both ends of the link segment supporting low speed auto-negotiation communication and there is no knowledge about the link segment, transmission should start in low speed auto-negotiation mode, as this would also allow communication over a long reach link segment.
- If there is still no response, after some time the universal PHY needs to stop transmission again and listen, if the far end PHY is now starting to transmit data (which e.g. has been disturbed by the wrong communication speed of the local PHY), and so on.

- Alternatively probably other algorithms like e.g. toggling the communication speed after some time in a pseudo random manner could be used.

- **The right algorithm mainly depends on what existing 100/1000BASE-T1 PHYs are able tolerate and what would disturb existing implementations (if there is need to stay compatible to these)**.

Compatibility (Signal Amplitude)

- A **long distance PHY** operates using a signal amplitude of **2.4 V_{pp}**, while a **powered intrinsically safe link segment** operates using a reduced signal amplitude of **1.0 V_{pp}**.

- One possibility would be to also **adopt the auto-negotiation signal amplitude** to the maximum allowed signal amplitude of the link segment (2.4 V_{pp} during normal operation, 1.0 V_{pp} when operating in low power mode, e.g. in powered intrinsically safe applications).

- Not reducing the signal amplitude for a powered intrinsically safe link segment will cause the auto-negotiation signal to get asymmetrically clipped by the clamping diodes within the intrinsically safe circuit.

- An other possibility would be to **always use a signal amplitude of 1.0 V_{pp}** for the auto-negotiation communication, independently from the operation mode (which reduces the SNR running over a long reach link segment).

- Nevertheless comparing the received auto-negotiation signal amplitude while transmitting with 1.0 V_{pp}, running over a 1000 m link segment, with the PAM-3 signal amplitude being received during normal PHY operation, the received auto-negotiation signal amplitude is still higher and during auto-negotiation there is no need to reach a BER of 10^{-9} or better.

- Further analysis/discussions are necessary, if a 1.0 V_{pp} auto-negotiation signal amplitude is suitable also for a long reach link segment.

Compatibility (Signal Amplitude)

- Theoretically universal devices transmitting auto-negotiation signals at 2.4 V_{pp} and 1.0 V_{pp} are possible.

- Such a device can start transmitting at a signal amplitude of 2.4 V_{pp} during the auto-negotiation process, and therefore is likely not understood by an intrinsically safe port, but as soon as it receives an auto-negotiation telegram from the far end side, which only shows a capability to communicate at an amplitude of 1.0 V_{pp}, it has to reduce its own transmit level to 1.0 V_{pp}.

10BASE-T1L PHY (reduced auto-negotiation signal amplitude)	1.0 V_{pp} auto-negotiation signal amplitude → ← 2.4 V_{pp} auto-negotiation signal amplitude	10BASE-T1L PHY (default auto-negotiation signal amplitude)

Switch to reduced signal amplitude, if the far end device only supports reduced amplitude.

10BASE-T1L PHY (reduced auto-negotiation signal amplitude)	1.0 V_{pp} auto-negotiation signal amplitude → ← 1.0 V_{pp} auto-negotiation signal amplitude	10BASE-T1L PHY (reduced auto-negotiation signal amplitude)

- Nevertheless it needs to be discussed if there is really a need for an automatic signal amplitude adoption during auto-negotiation, as this would require an analysis of the negotiated data by the auto-negotiation process itself within the state machines.
- As this adds quite some amount of additional complexity during adoption of Clause 98, it should be avoided, if possible.

Limitations

- Thinking about a future 10/100BASE-T1L PHY, the auto-negotiation procedure can **negotiate the features** of the used PHYs, but auto-negotiation **is not intended to check, if the link segment is suitable** for a 100 MBit/s communication speed or not.

- This is different compared to current 4-/8-wire 10/100/1000 MBit/s PHYs, which assume that there is a minimum quality of the link segment, which allows to just run at the highest negotiated speed (e.g. when using Cat. 5e cabling up to 100 m).

- Therefore for a long reach link segment a future 10/100BASE-T1L PHY needs to be forced to 10 MBit/s by default and a higher communication speed needs to be set by user interaction (e.g. during installation) to guarantee a reliable operation.

- Alternatively, in addition to the auto-negotiation process, a future 10/100BASE-T1L PHY may do a probing if the link segment is suitable for a 100 MBit/s communication or not.

- If the link segment is not suitable (e.g. because the insertion loss is too high for a 100 MBit/s communication), then a future 10/100BASE-T1L PHY has to fall back to a 10 MBit/s communication speed, even if both PHYs have the 100 MBit/s capability field set.

- The (optional) link segment probing will likely have to be part of a future 100BASE-T1L standard.

Summary

- With adoption of the communication speed and timing parameters Clause 98 auto-negotiation seems to be suitable for a 10BASE-T1L PHY.

- Based on the shown simulations a 625 kBit/s DME communication seems to be appropriate for a 10BASE-T1L link segment, testing in real life is still necessary, a detailed review of the required timings would also be necessary.

- If we agree that auto-negotiation between short and long reach PHYs is necessary, we need to implement an algorithm to negotiate at two different speeds, even if this makes the adoption of Clause 98 more complicated.

- There is need for further analysis/discussions, if an auto-negotiation signal amplitude of 1.0 V_{pp} is also suitable for a long reach link segment.

- As for several applications no auto-negotiation is necessary, also a forced setting of the relevant parameters (e.g. by bootstrap option or by MDIO registers or a similar implementation) is required, in this case the auto-negotiation procedure should be disabled.

- **This presentation just shows some principle ways how to implement Clause 98 for a 10BASE-T1L PHY, there is still a lot of work to do.**

Thank You