Auto-Negotiation (AN) Overview

BRAD BOOTH, MICROSOFT
ARTHUR MARRIS, CADENCE

Scope

Common text found in Clause 28, 37 and 73

- Allows a device to advertise enhanced modes of operation it possesses to a device at the remote end of a link segment and to detect corresponding enhanced operational modes that the other device may be advertising
- Exchanges information between two devices that share a link segment and to automatically configure both devices to take maximum advantage of their abilities
- Does not test the link segment characteristics
- Permits management to disable or enable the Auto-Negotiation function
- Reject the use of operational modes that are not shared by both devices
- Allows management to select a specific operational mode

Outside the Scope of Auto-Negotiation

Management interaction with external devices

- Example, EEPROM inside an AOC or DAC module
- Cable assembly assumes a fully-passive system with no management capabilities
- AN sublayer cannot assume access to external EEPROM data
- Implementations may use management interaction between AN and external devices

Does not guarantee it will find a compatible mode of operation

- Management entity will have access to link partners advertised abilities after AN has finished

Does not guarantee a compatible mode of operation will result in a link being established or maintained

- AN completes before link training or PCS alignment and sync; therefore, has no indication if a link is possible

Where Are They Used?

Clause 28 AN

- BASE-T PHYs
- Mandatory at 1G+ data rates
- Sits between the MDI and PMA/PMD

Clause 37 AN

- Included as part of the Clause 36 PCS
- Only used in non-backplane 1000BASE-X PHYs

Clause 73 AN *(Most relevant to P802.3by)*

- Borrowed from Clause 28 as its baseline
- Modified to support backplane (≥ 1G) and copper cabling (> 1G)
- Mandatory to implement, optional to use
- Sits between the MDI and PMA/PMD

Clause 73 Overview

Multi-speed devices
- Can permit incremental speed upgrades in network architecture
- Can provide a “plug and play” feature

Uses low speed signalling that can work over low capacity channels
- Lane 0 used in multi-lane interfaces

Restarted by link failure or management request
- Either end of the link may restart AN

After the speed is negotiated
- Link training is performed
- Management communicates with AN and indicates the selected speed to the MAC

AN is always under the control of management
- Management may intervene to select a specific mode of operation or to modify advertised abilities
- If the advertised abilities of both devices are compatible, management intervention is not necessary

Management

Management is undefined in 802.3 but is considered pervasive

- Beyond the scope of the standard
- Left up to the implementer

Clause 22 and Clause 45

- Define Management Data I/O (MDIO) interfaces
- Define the register bit mappings and definitions
- Optional to implement the interfaces

Can assist in finding a compatible mode of operation

- Example, 40G port (capable of 4x10G) interfacing to four 10G ports
- 40G port management can see the exchanged information
- Restart auto-negotiation on each lane of the port to negotiate to 10G

AN Page Exchange

Link partners (LPA and LPB) exchange abilities and modes of operation via the exchange of base pages and, if requested, next pages

If requested, next pages are exchanged until both link partners have exhausted their next pages

- Next page messages are defined in Annex 28C
- Exchanges identifier tags, EEE parameters, operating parameters and vendor specific information

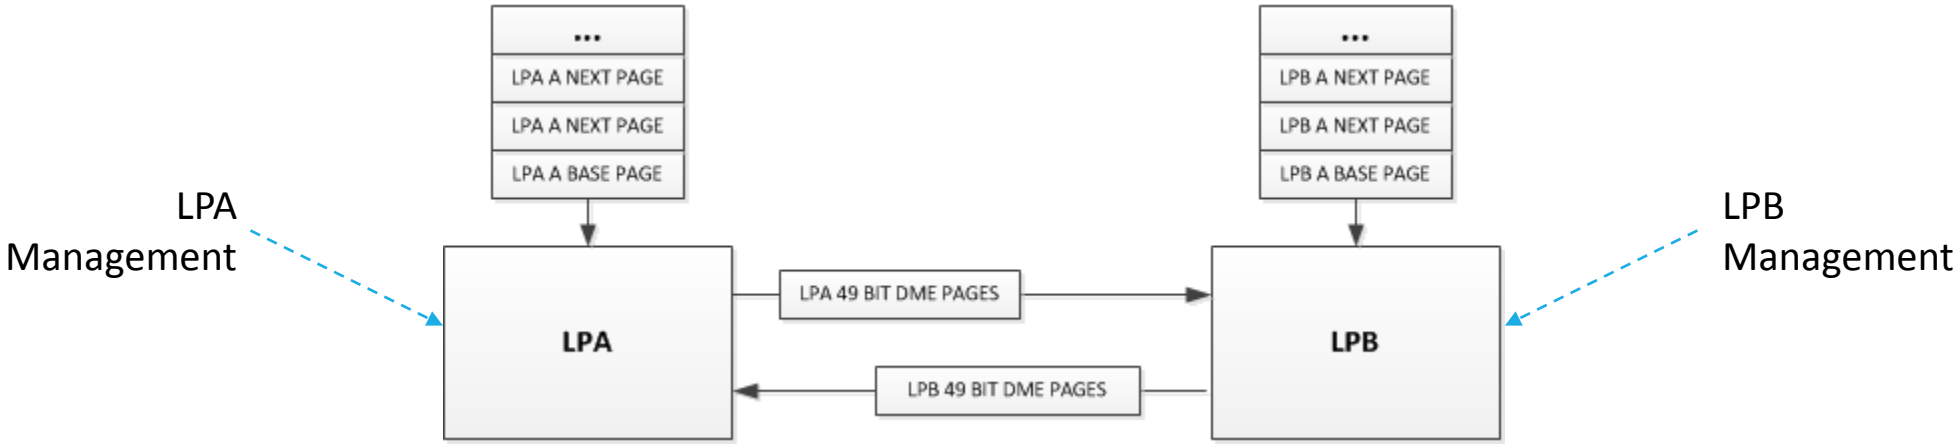

DME = Differential Manchester Encoding

Example 1 – Compatible Abilities

Using 1G/10G backplane as an example

Both link partners are equipped with AN and both are capable of 1000BASE-KX, 10GBASE-KX4 and 10GBASE-KR and FEC

LPA Management has LPA indicate 1000BASE-KX and 10GBASE-KR capabilities with FEC ability and FEC requested

LPB Management has LPB indicate 10GBASE-KX4 and 10GBASE-KR capabilities with FEC ability and no FEC requested

AN completes with the resolution of 10GBASE-KR with FEC enabled

Each LP Management reads the resolution from AN and waits for indication that the link is up before initiating the exchange of packets

Example 2 – Incompatible Advertised Abilities

Using 1G/10G backplane as an example

Both link partners are equipped with AN and both are capable of 1000BASE-KX, 10GBASE-KX4 and 10GBASE-KR and FEC

LPA Management has LPA indicate 1000BASE-KX

LPB Management has LPB indicate 10GBASE-KR capabilities with FEC ability and requested

AN completes with no highest common denominator resolution

Each LP Management reads the resolution from their respective AN and could do the following:

- Either LP Management could change its advertised abilities and restart AN
- Flag to user there is an incompatible mode and user can manually configure the operating mode

Example 3 – Incompatible Devices

Using 1G/10G backplane as an example

Both link partners are equipped with AN

LPA is only capable of 1000BASE-KX and LPA Management has it advertise that ability

LPB is only capable of 10GBASE-KR with FEC and LPB Management has it advertise those abilities

AN completes with no highest common denominator resolution

Each LP Management reads the resolution from their respective AN and would flag to user to incompatible mode

- User would be required to find compatible hardware