EFM Copper (EFMC) Tutorial

June 2004

EFM Copper (EFMC) Tutorial

Introduction

While the demand for Ethernet services continues to grow, the challenge of delivering these services to end users remains a significant roadblock in the way to fulfill this demand. Until recently, fiber was the only way to deliver high speed Ethernet services. However, the limited coverage of the fiber network required an alternative solution to complement it and make Ethernet services ubiquitous. Twisted-pair copper wiring (plain old telephone line) dominates the local loop from the home or curb to the Central Office—the "first mile." Running Ethernet over this copper wiring is an ideal way to exploit the existing voice-grade copper infrastructure, within residential neighborhoods as well as business buildings. Using the existing voice wire infrastructure keeps deployment costs to a minimum: there is no need for new cabling inside or outside the residence or business and service providers enjoy new returns on their already amortized assets.

Hence, one of the solutions being touted in the IEEE Ethernet in the First Mile (EFM) standardization process is EFM over Copper, or EFMC. Running over existing Category 3 wire, EFM has set goals for a short reach option of at least 10 Mbps up to at least 750 meters, and a long reach option of at least 2 Mbps up to at least 2700 meters. While the EFM committee has set these objectives as a minimum, the standard does not limit systems to these rates, and in fact most EFMC systems available today support much higher rates. Additional mechanisms for bonding of multiple copper pairs allow an even higher throughput, providing a viable alternative for end-users served only by copper.

There are limited EFMC deployments today, but the EFM (802.3ah) standard will eliminate the proprietary nature of these early pre-standard implementations, while improving vendor interoperability for large public networks and denser deployments. As existing Ethernet PHYs are designed for engineered wiring, this is a new PHY level standard for the telephone line, leveraging field-proven DSL technology as its line code.

Finally, there are FCC requirements for spectrum compatibility and EMI – these requirements are not met by existing Ethernet PHY specifications. Similarly, existing DSL specifications are optimized for non-Ethernet protocols. 802.3ah addresses both of these issues.

EFMC: Based on DSL

The EFMC PHY uses DSL modulation techniques. This leverages years of work on DSL modulation development, and ensures spectral compatibility. Most importantly, there is a great increase in distance (with G.SHDSL.bis used for long reach covering over 5 km), and only one twisted pair is needed (as a minimum) for EFMC short reach or long reach (see Figure 1).

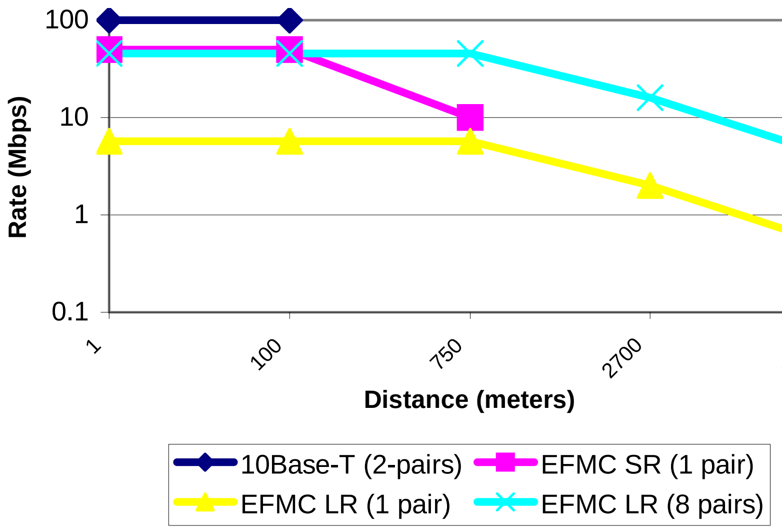

Figure 1: Distances and Bandwidth for10BASE-T, EFMC SR and EFMC LR

This is an evolutionary improvement over existing DSL, taking the existing DSL platform and simplifying it, while mandating interoperability. EFM simplifies the protocol layers and reduces configuration and provisioning options. All of this is in keeping with the 802.3 Ethernet tradition, which has always stressed interoperability. For example, while there were many non-interoperable DSL types, there are only two Ethernet port types.

Background on DSL

DSL uses twisted pair access loops to transmit wideband digital signals. The various DSL flavors and their characteristics are specified in the following table.

DSL Type	**Characteristics**
SHDSL and Extended SHDSL	Symmetric, T1 carriage, no POTS overlay, medium to long loops
ADSL	Asymmetric, POTS overlay, medium to long loops
VDSL	Symmetric and asymmetric, short loops, high speed; operates at up to 12 MHz bandwidth

Table 1: DSL Types and Characteristics

There are two broad categories of DSL modulation techniques, both of which are commonly used in various DSL standards. These are:

- *Discrete Multitone Modulation (DMT):* Large number of narrowband, orthogonal, modulated carriers
- *Quadrature Amplitude Modulation (QAM):* Single wideband, modulated carrier

The two main regulatory issues with DSL are loop unbundling and spectral compatibility. Loop unbundling occurs because loops in a binder are often operated by different telephone companies. Crosstalk from pairs operated by one company can effect performance on pairs operated by another.

Spectral limits and deployment guidelines need to ensure the fair use of binder resources. ANSI T1.417 is the American National Standard for spectral compatibility. It requires compatibility with widely deployed “basis systems.” Other countries have issued their own specific spectral requirements.

A brief timeline showing the evolution of DSL is given in *Figure 2.*

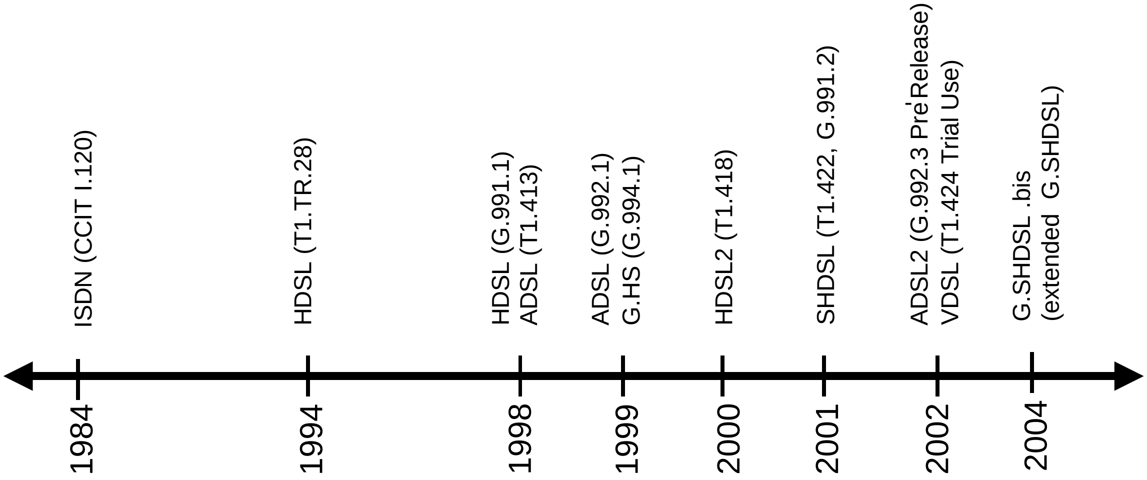

Figure 2: DSL Timeline

As *Figure 3* shows, the current typical DSL protocol stack is an outdated collection of options supporting PPP and ATM sublayers. It was built to accommodate services that were never deployed, which results in additional costs for needless provisioning, configuration, and maintenance.

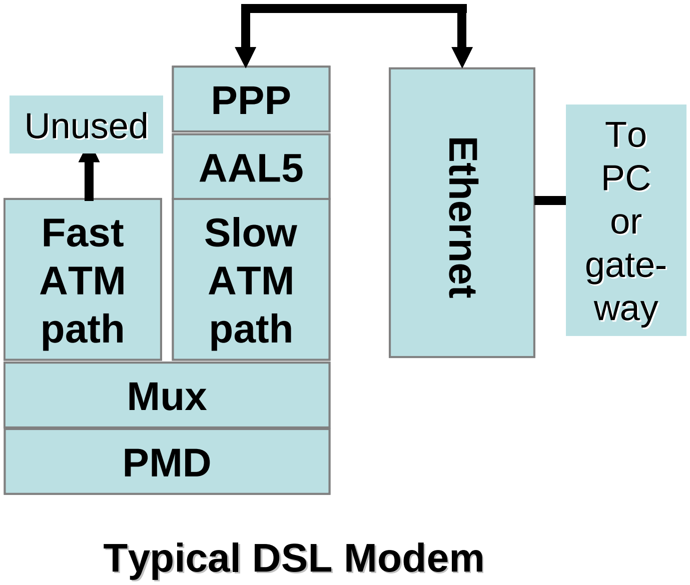

Figure 3: Typical Current DSL Protocol Stack

DSL Enhancements for EFMC

By contrast, a typical IP connection (whether it carries data, voice or video) begins and ends on Ethernet, so the flexibility of ATM is unutilized; supporting ATM just leads to unnecessary complexity. Additionally, the ATM "hardware-based" implementation that was planned to carry much higher traffic compared to the "software implementation" of IP routers, now finds an equal rival with hardware-based Ethernet switching modules that can carry traffic in similar rates at lower cost. As *Figure 4* shows, new DSL systems will strip out the intermediate sublayers and move to native Ethernet over DSL.

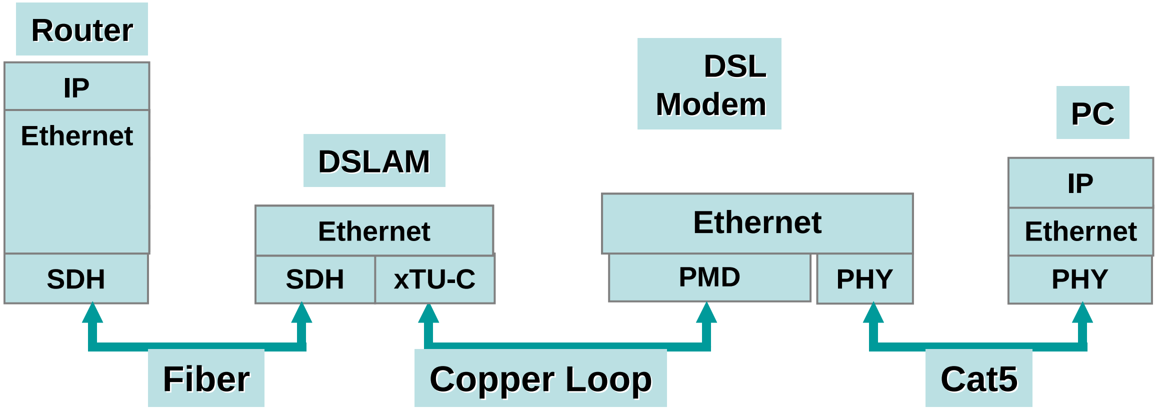

Figure 4: Newer DSL Systems Focus on Ethernet

With the EFMC architecture (*Figure 5*), new data link sublayers (closer to the PHY) are defined by existing DSL standards, and there are some new layers above those for rate matching and loop aggregation. Above that, you find the existing Ethernet layer.

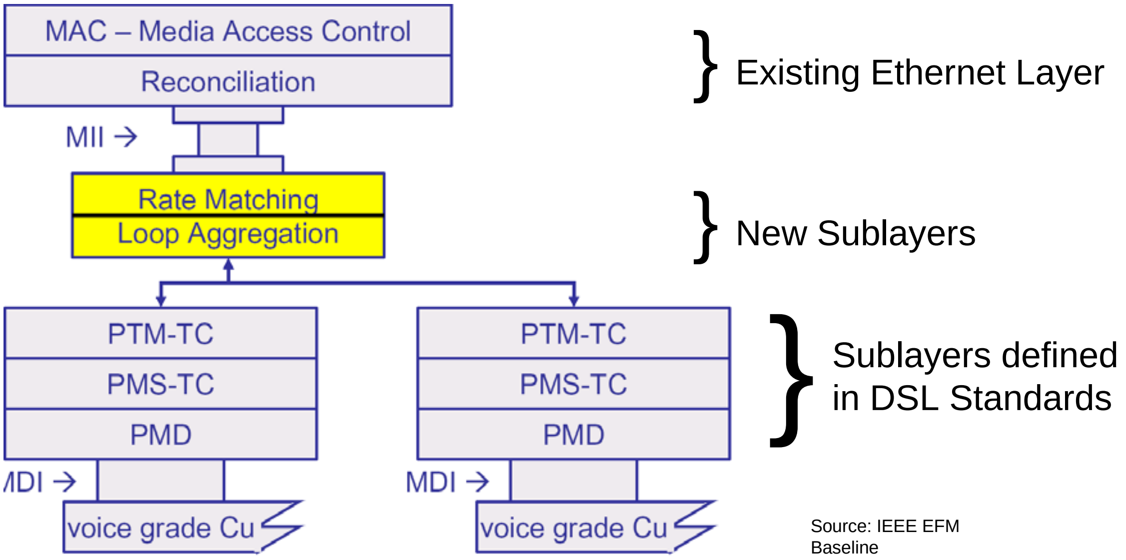

Source: IEEE EFM Baseline

Figure 5: EFMC Architecture

The EFMCu aggregation layer (at the transmitter) gets Ethernet packets through an MII interface after the IPG (Inter Packet Gap) and the Preamble have been removed. The optional aggregation layer (not required with a single modem) breaks the packet into variable length fragments. Each fragment is then forwarded to a specific modem's TC layer, where it is encapsulated with 64B/65B Framing and transmitted by a modem (PMA/PMD) onto the wire.

The receiver side decapsulates the fragments, reassembles the original Ethernet frames and restores the IPG and preamble. The EFM overhead is about 5%, depending on the packet size and the fragmentation algorithm (vendor specific).

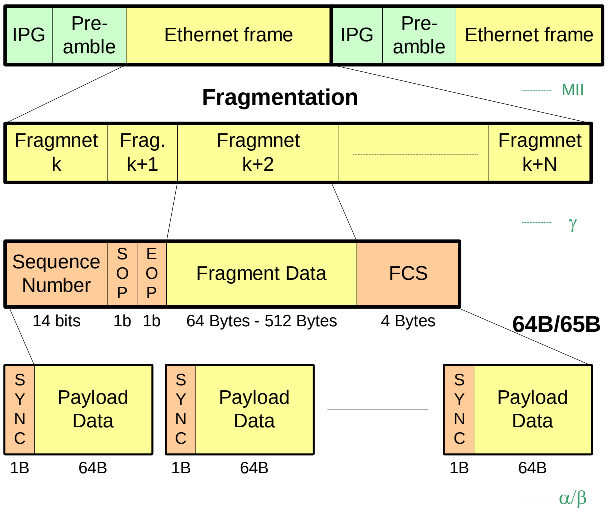

Figure 6: EFM Fragmentation and Framing

Copper Loop Issues and Solutions

Consider copper wiring in the local loop, between the central office and the subscriber. As *Figure 7* shows, there is a binder of 25 – 50 pairs going out to all subscribers.

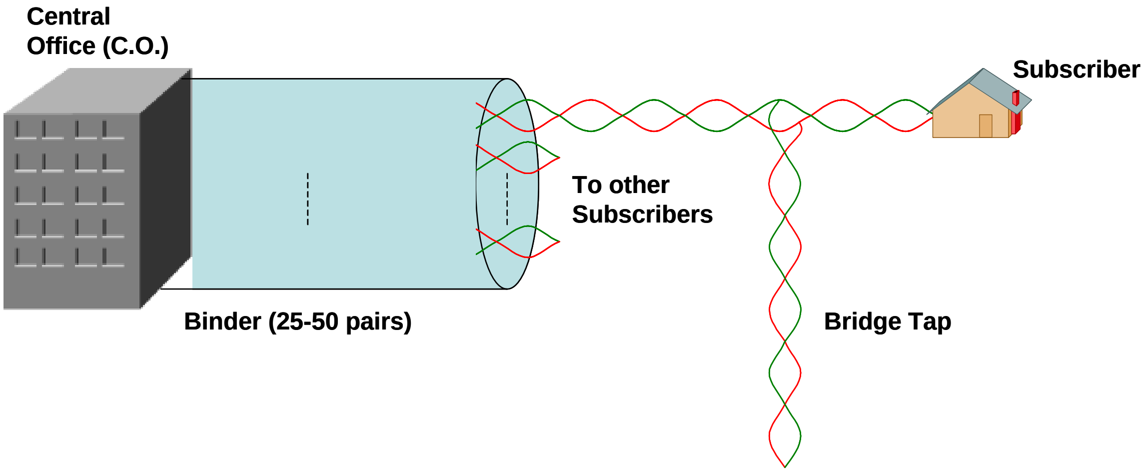

Figure 7: PSTN Loop Plant

Multiple pairs are wrapped tightly together in each binder, and binders fan out as they extend toward subscribers. “Bridge Taps” occur where stubs are left unconnected, and in-building wiring is also a factor to consider.

What are the transmission limitations of this architecture? Attenuation is one, because loss increases with frequency. Crosstalk is also a concern, because of the predominant impairment in the loop plant. There is interference from the same type of service on other pairs in the binder (self-crosstalk), or other types of service (alien-crosstalk). Finally, there is the issue of the POTS/IDSN overlay: either POTS (0-25 KHz) or ISDN (0-138 KHz) may be operating on the same pair.

Band Plan Definitions

To mitigate some of these limitations, regulators administer band plan definitions (*Figure 8*) to help endure the operation of different services in the same binder.

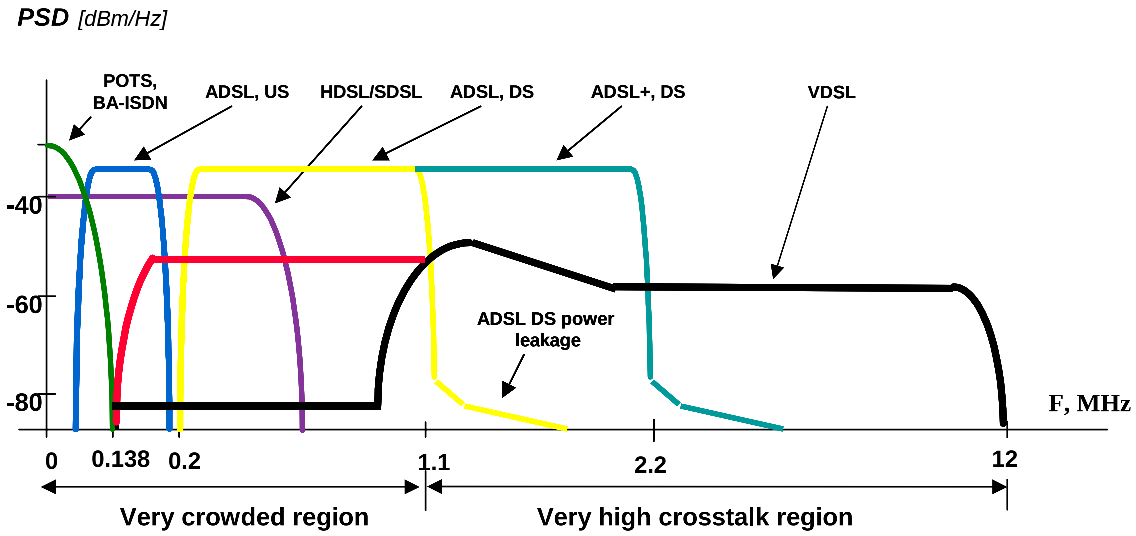

Figure 8: Band Plan Definitions

Coping with Crosstalk

There are two types of crosstalk problems: Far End Crosstalk (FEXT) and Near End Crosstalk (NEXT). These are shown in *Figure 9*.

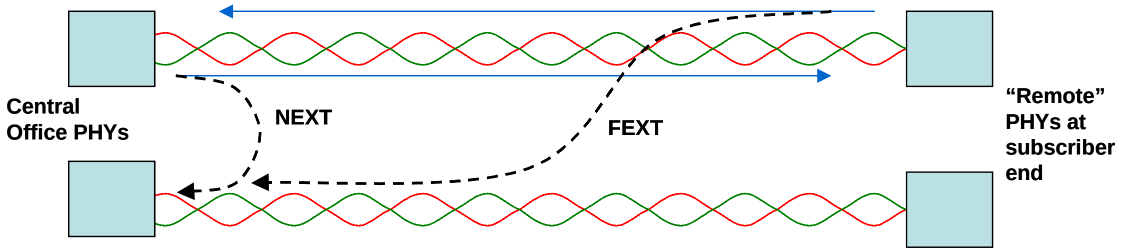

Figure 9: Far End and Near End Crosstalk

FEXT is caused by a transmitter operating on another pair in the binder, at the opposite end from the receiver. The crosstalk level is smoothed out by loop attenuation.

NEXT is caused by a transmitter operating on another pair in the binder, at the same end as the receiver. There is no loop attenuation; this is a higher level of crosstalk than FEXT. The impact on NEXT or FEXT on the transmission depends on the signal type that is transmitted. For VDSL systems that are DMT based, the transmit and receive signals are utilizing different frequency bands, and therefore an adjacent transmit signal has a very low impact on a received signal, making the NEXT much smaller so the FEXT becomes the dominant disturbance. On the other hand, with G.SHDSL systems that utilize the same frequencies for transmit and receive (and echo cancellation techniques to separate between them), the NEXT is the dominant factor that limits transmission capabilities.

Determining Channel Capacity

The example in *Figure 10* is for downstream VDSL (with Plan 998).

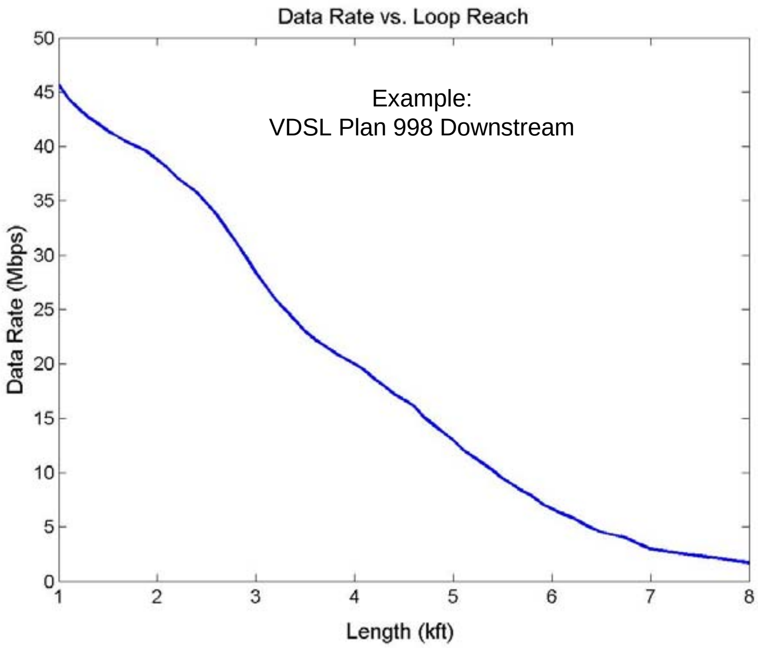

Feet	Meters
1,000.00	304.80
2,000.00	609.60
3,000.00	914.40
4,000.00	1,219.20
5,000.00	1,524.00
6,000.00	1,828.80
7,000.00	2,133.60
8,000.00	2,438.40

Figure 10: Sample Data Rate versus Loop Reach (Downstream VDSL)

To those of us who are used to normal DSL rates on the order of 384 Kbps, it's surprising to see, for instance, that 10 Mbps is theoretically possible at a greater than 5000 foot reach. The above performance chart assumes ten self-FEXT disturbers, no self-NEXT disturbers, and a background noise of -140 dBm/Hz.

EFMC Ports: Short Reach and Long Reach

The two EFMC port types are a short reach PHY, type 10PASS-TS and a long reach PHY, type 2BASE-TL. The bandwidth and distance capabilities of these options are shown in *Figure 11*.

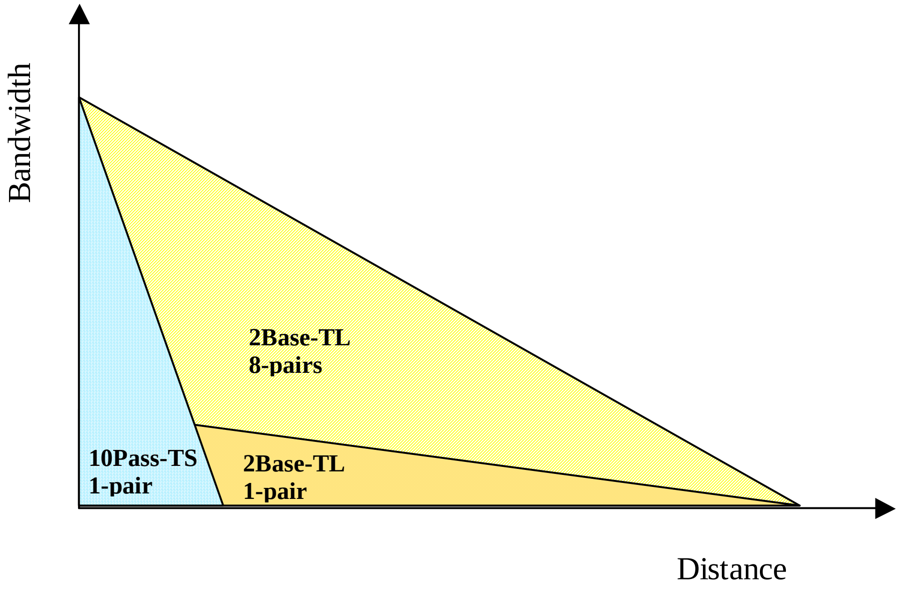

Figure 11: Short and Long Reach Options for EFMC

For the short reach PHY (EFMC SR), the band plan works as follows. Plan 997 is used in Europe, and Plan 998 is used in North America. The short reach PHY details are illustrated in Figure 12.

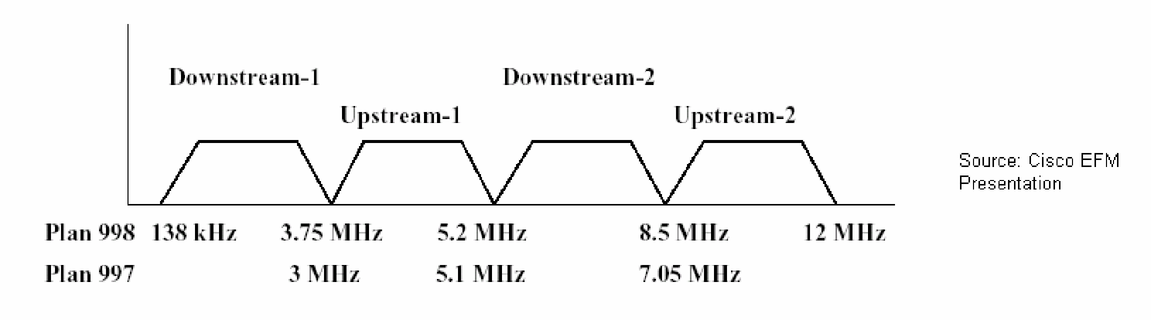

Source: Cisco EFM Presentation

Figure 12: EFMC Short reach PHY Details

The long reach PHY (EFMC LR) is based on the ITU-T G.991.2 standard for single-pair high-rate DSL; it is inherently symmetric and uses TC-PAM. An extended version of this line code called the G.SHDSL.bis is now being standardized by ITU-T and ANSI, and has been adopted by the EFM committee. This extended version allows bit rates of up to 5.7 Mbps symmetrical, while still complying with spectral compatibility requirements such as ANSI T1.417. The G.SHDSL.bis enables delivery of high speeds to long distances, covering almost all the customer base of the service providers.

Copper Loop Bonding

EFM has introduced an important capability to copper-based system – the ability to utilize more than one pair and carry far more bandwidth over the exiting copper infrastructure. The EFM Aggregation layer allows multiple pairs to be used as a single, high capacity link, providing a “fiber replacement” in places where fiber does not exist. With loop bonding capability, no customer will be left without high-speed business-class Ethernet service, making Ethernet service ubiquitous.

While some other bonding mechanisms were suggested in the past to perform this function, the EFM bonding has proven to be most efficient for delivering Ethernet traffic. Specifically designed to meet the requirements of copper transmission, with less-than-optimal predictability and variable rates and delays on the pairs, the EFM bonding is a natural bonding scheme choice for Ethernet services. While specific vendor enhancements to the transmission layer may allow additional performance gain, the EFM standard bonding allows for vendor interoperability and enables mass deployment.

When comparing the EFM bonding to alternative bonding schemes such as inverse multiplexing over ATM (IMA), or the M-pair G.SHDSL scheme, EFM proves in to be superior in all aspects. The following table provides a detailed comparison.

	EFM 802.3ah	**IMA**	**M-pair G.SHDSL**
Service Optimization	Ethernet	ATM	TDM
Network Integration	Ethernet/IP native interface	ATM networks	Not Defined (system feature)
Service conversion required	No	Yes	Yes
Overhead	~5%	~20% in mixed traffic ~40% in short frames	~20% in mixed traffic ~40% in short frames
Operates on pairs with different rates	Yes (can utilize the high and low rates)	No (uses lowest rate for all pairs)	No (uses lowest rate for all pairs)
Typical bonding delay	2-4ms	25-100ms	2-4ms
Noise Immunity	High-Med (depending on vendor implementation)	Low (long recovery time)	Low (single pair-loss drops the link)
Management complexity	Low (Ethernet management)	High (ATM management)	Not defined in standard (system implementation)

Table 2: Copper Bonding Scheme Comparison

Conclusion

We need a better solution than existing DSL to accommodate business and residential growth in bandwidth demand and quality requirements. The IEEE EFMC committee standardizes the efficient delivery of Ethernet packets directly over copper pairs at 10 Mbps and above in both directions. This “Ethernet-pure” solution provides a seamless integration into today’s and tomorrow’s networks.

EFM has set objectives to achieve 10 Mbps at 750m (EFMC SR) and 2 Mbps at 2,700m (EFMC LR). However, the standard does not limit implementation to these rates and existing products already exceed them, delivering higher throughput to longer distances. Additionally, the introduction of copper bonding into the standard allows delivery of even higher bandwidth to longer distances over multiple copper pairs, enabling a good alternative in places where fiber does not exist or is not economical to deploy.

By reducing service provider capital expenditures for implementation, EFMC is an easy, low-cost, and immediate solution for providing feature-rich, high-speed access and services to subscribers. This is an attractive access solution for both residential and business users, and can coexist with ADSL, VDSL, ISDN and PSTN in the same cables, bringing native Ethernet to the first mile over a twisted pair access network.