Digital Complexity in DSL: An Extrapolated Historical Overview

Michael Timmers, Koen Hooghe, Mamoun Guenach, and Jochen Maes
Alcatel-Lucent Bell Labs
Copernicuslaan 50
B-2018 Antwerp, Belgium
Email: {michael.timmers,mamoun.guenach,jochen.maes}@ieee.org,koen.hooghe@alcatel-lucent.com

***Abstract*—Digital subscriber line (DSL) technology for copper twisted pair access networks has been evolving to meet the ever-growing demand for higher data rates. This evolution has gone hand in hand with the roll out of fiber deep in the access network. The most recent technology is vectored VDSL2, able to offer an aggregate data rate of 200 Mb/s on a single copper pair. The next step is to reach 500 to 1000 Mb/s over even shorter copper loops up to a few hundred meters. Such a DSL deployment is an enabler for the cost-effective continuation of the fiber roll-out closer to the end-user. In this paper, a reality check is performed on the digital complexity of a next-generation DSL (ΩDSL) transceiver. By taking into account Moore's law, it is shown that the time is right for this next-generation DSL.**

Keywords*—*Digital Subscriber Line, complexity, analysis

I. Introduction

After their first introduction in the early 1990s, wireline broadband networks, which includes fiber, coaxial cable and twisted pair, has evolved substantially. Despite the inherent attenuation of copper which limits the capacity, transmission over this medium remains attractive as it is abundantly present throughout the world due to historical telephone deployment. Hence, broadband over copper offers substantial deployment cost savings as compared to fiber-to-the-home (FTTH). Indeed, while FTTH has been technologically viable since 1988 [1], digital subscriber line (DSL) remains the predominant broadband access technology for the residential market [2]. However, as the access network remains the bottleneck in the end-to-end connection and due to the continuing demand for ever higher data rates, copper is being replaced by fiber step-by-step. Due to typical branched topologies, the cost per user of fiber deployment increases substantially when moving closer to the user. This is why different operators have expressed enthusiasm with recent technologies, such as phantom mode and vectoring, which hold the promise of delivering more than 300 Mb/s [3]. The success of vectoring and phantom mode transmission triggered interest in a next-generation broadband DSL, ΩDSL, beyond vectored VDSL2 to deliver 500 Mb/s to 1 Gb/s over relatively short loops, i.e., below 400 m [4]. Standardization of such an ΩDSL has been started in the project G.fast. However, due to the competition from other access technologies, the DSL capacity increase must remain cost-effective and, hence, low-complex. Indeed, as a healthy cost difference between FTTH and Fiber-To-The-Curb (FTTC) needs to stay in place, ΩDSL designs need to carefully evaluate the complexity of the used scheme.

In this paper, we focus on the digital complexity of the underlying scheme. We discuss different generations of DSL technologies based on discrete multi-tone modulation (DMT), using the methodology presented in [5]. We compare the complexity of different DSL flavors by introducing a time scaling based on Moore's Law [6]. Finally, we extrapolate these results to evaluate DMT-based proposals for a ΩDSL. We show that an evolutionary path is in line with previous complexity increases between generations. A DMT-based ΩDSL, which can leverage on proven technology, is, hence, recommended.

This paper is structured as follows. First, we give an overview of discrete multi-tone modulation in Section II. Afterwards, we present challenges and opportunities for the next-generation DSL in Section III. Then, in Section IV, we introduce the reference methodology for complexity analysis. We analyze the digital complexity of the different DSL flavors in Section V and draw conclusions in Section VI.

II. Overview of discrete multi-tone modulation

In this section, we give a short overview of the DMT modulation scheme and discuss its strengths and weaknesses.

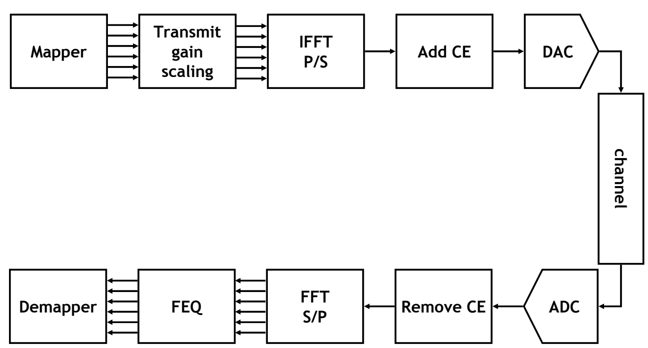

Fig. 1. Discrete multi-tone modulation splits the frequency-selective channel into a number of independent narrowband flat channels, which allows to adapt to the channel in a computationally efficient way.

DSL modems use a particular modulation format referred to as discrete multi-tone modulation (DMT), which splits the frequency-selective channel into a number of independent narrowband flat channels. This allows to cope with severe channel conditions in a flexible and computationally efficient

 ISBN: 978-1-61208-142-7

way. For instance, the frequency-selectivity of the twisted pair channel can be dealt with by means of simple frequency domain equalization techniques. Moreover, narrowband interference can be handled by vacating (notching) the transmission over the corresponding subcarriers, complemented with time-domain windowing. Echo reflection can be tackled by duplexing two-way data transmission in non-overlapping upstream and downstream bands, which is referred to as frequency-division duplexing (FDD). DMT furthermore allows for a spectrally efficient design, as the amount of power and the number of bits transmitted on each subcarrier can be allocated in a flexible and optimal way, using so-called power and bit loading procedures.

In Fig. 1, the general blocks of DMT modulation are shown [5, 7]. In the transmitter, the input data bits are first mapped onto QAM constellations points for each of the n_c carriers. Each carrier has a complex transmit gain to control the carrier power. The frequency domain samples X_k are then converted into a time-domain symbol with an N-point inverse fast Fourier transform (IFFT), where N is twice the number of carriers, n_c:

$$x_n = \sum_{k=0}^{N-1} X_k e^{\frac{-j2\pi kn}{N}} \tag{1}$$

After the IFFT, a cyclic extension (CE) is added with a length ν that is larger than the time delay spread to combat inter-block interference. Windowing is further applied to reduce out-of-band leakage. The resulting signal then passes through a digital-to-analog converter (DAC) and is sent over the channel.

In the receiver, the received signal is digitized with an analog-to-digital converter (ADC). An FFT per DMT block of N samples is applied to switch back to the original frequency-domain symbols. The single-tap frequency domain equalizer inverts the channel on a per-carrier basis, followed by the demapper to recover the transmitted information bits.

III. THE NEXT-GENERATION DSL

In this section, we discuss the objectives for ΩDSL. Afterwards, we present a general overview on where the opportunites lie to achieve these objectives.

A. Objectives

To compete with other access network technologies, such as wireless, coax or fiber, ΩDSL should target a data rate of 1Gb/s aggregated over upstream (US) and downstream (DS). Although the residential market demand is still far below this limit, it is expected that the bandwidth consumption will continue its exponential growth and it is expected that the top line data rate demands will exceed 1Gb/s by 2030 [8]. Aside from this data rate, ΩDSL should be very flexible, both in the frequency as in the time domain to adapt to dynamic channel conditions and applications requirements. An important aspect is also energy efficiency: the energy consumption of ΩDSL should scale with the applied load, i.e., the actual throughput. As the DSL environment will become a lot more dynamic, it is also important to be robust by avoiding packet errors, retransmitting and keeping the downtime of the DSL communication as low as possible.

B. Approach

DMT remains one of the prominent candidates for achieving the objectives listed in Section III-A. Indeed, DMT is very flexible in the frequency domain and especially suited for spectral confinement, which is important when moving to higher bandwidths, where additional notching is required.

Furthermore, as any other multicarrier based modulation the frequency selectivity of the channel can easily be addressed by very basic single tap equalizers. It is also robust to near-end crosstalk and echo through frequency division duplexing. Far-end crosstalk can be avoided through precoding, which can be easily implemented in the generic DMT scheme. However, it is not so robust against transient noise. Therefore, we should rely on fast adaptation mechanisms like a fast bit swap and bit loading or adaptive coding techniques, such as FEC or ARQ. Luckily, due to the flexibility of DMT such fast adaptation is possible. However, due to the FFT/IFFT core, DSL is a complex modulation scheme. Hence, in this paper, we want to perform a reality-check and look if ΩDSL is feasible. In this section, we will first describe how the DMT scheme needs to be applied to deliver the high data rates that are required.

We start our analysis from the Shannon-Hartley theorem [9]:

$$C = W \log_2 \left(1 + \mathrm{SNR}\right), \tag{2}$$

where C is the Shannon capacity of the communication channel and W is the analog bandwidth used. The signal-to-noise ratio at the receiver is denoted as SNR. This Shannon capacity is an upper limit on the throughput that can be achieved over a channel with no errors. However, a gap exists between the Shannon capacity and the practical data rate, R, of the channel. For DSL systems the following formula is generally used [12]:

$$R = \eta \Delta_c \sum_{k=0}^{N-1} \min \Big(\log_2 \big(1 + \frac{|H(k)|^2 P_t(k)}{\Gamma(\sigma_0^2(k) + I(k))}\big), b_{\max} \Big). \tag{3}$$

Here, η represents the efficiency, which takes into account the coding and CE overhead, Γ is the SNR gap, Δ_c the carrier spacing and $b_{\max}$ the bit cap. The channel attenuation is represented as $|H|^2$ and the transmit power is denoted P_t. A typical value for the transceiver noise power spectral density, σ_0^2, in state-of-the-art DSL designs is -135 to -140 dBm/Hz. The interference power I is the result of radio ingress, alien crosstalk and self-crosstalk.

To increase the capacity of a system, several methods can be applied. The most straightforward is to control the channel attenuation. This is done by pushing fiber deeper in the access network and, thus, shortening the looplengths of the copper channel. ΩDSL will consider looplengths up to 400 m and will be optimized for looplengths from 0 m to 200 m, which is typical for a FTTC or FTTB scenario (see Fig. 2).

Another parameter that can be optimized is the coding efficiency, for which today a typical range of 78% is used [12].

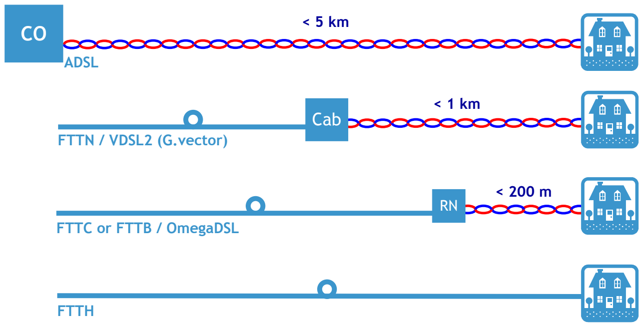

Fig. 2. Fiber is contiuously being deployed closer to the user, starting from an all-copper deployment from the CO (ADSL). An economically interesting next step exist in the combination of FTTB or FTTC with ΩDSL

TABLE I
PARAMETERS USED FOR FIGURES IN SECTION III

	Fig. 3	Fig. 4
Looplength	200 m	variable
Twisted pair	24 awg	24 awg
Noise power	-135 dBm/Hz	-135 dBm/Hz
Transmit power	-60 dBm/Hz	-60 dBm/Hz ($f <= 30$ Mhz)
		-76dBm/Hz ($f > 30$ MHz)
Γ	9.45 dB	9.45 dB
b_{max}	15 bits	15 bits
η	78%	78%
Notches	none	none

However, it is clear that the coding will need to provide the desired robustness and resiliency to counter the channel dynamics. Some components of the interference, such as self-crosstalk, can be cancelled today, e.g., in G.vector [10]. Another typical parameter to increase system capacity is the analog bandwidth, W. The widest profile of VDSL2 uses a 30 MHz bandwidth. Today, VDSL2 is limited at lower frequencies by the bit cap, $b_{\max}$, which is set to 15 bits per carrier. Also, here, headroom exist to remove this constraint in low-bandwidth systems, constrained by ADC technology.

In Fig. 3, we show the headroom for these different options for a 24 awg twisted pair with a looplength of 200 m for a transmit power of -60 dBm/Hz and a noise power of -135 dBm/Hz. Each block in the grid of this figure, represents 50 Mb/s (10 MHz multiplied by 5 (b/s)/Hz). Hence, we can immediately see that a 30 MHz bandwidth will not suffice, as it is only able to deliver 750 Mb/s Shannon capacity even on a null-loop (15 blocks in Fig. 3, i.e., 25 (b/s)/Hz over 30MHz). We also see that removing the bit cap, $b_{\max}$, only has limited benefit in the lower frequency range. The efficiency does have a large impact, but, as indicated above, coding overhead is required to deliver the desired robustness and resiliency. An interesting and valuable parameter is the SNR gap, Γ. Traditionally, this is used to cover the channel dynamics. However, when we leverage on the flexibility of the DMT modulation scheme, we can lower this gap. As indicated in Fig. 3, this increases the SNR with 6 dB or, equivalently, about 2 (b/s)/Hz.

In Fig. 4, we show the looplength distribution for three proposals for a DMT-based ΩDSL, using the parameters listed

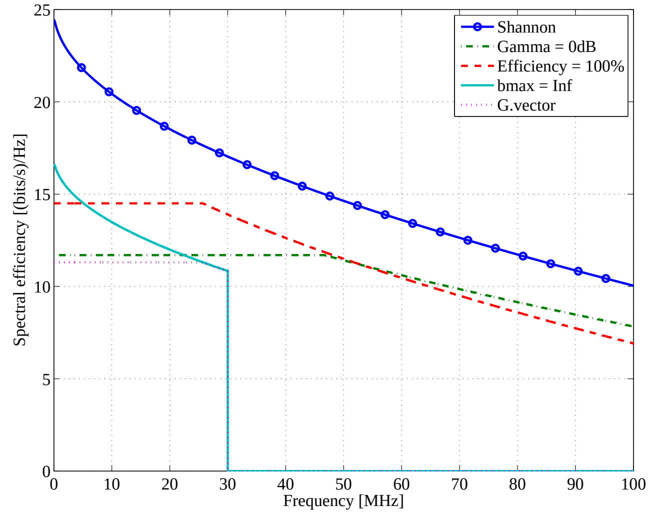

Fig. 3. The headroom of different options shows the relative contributions of each parameter optimization. To reach the 1Gb/s objective, increasing the bandwidth is a must.

in Table I. It shows that a 70 MHz profile cannot reach the desired 1 Gb/s for any looplength. However, a bandwidth of 140 MHz is in line with the objectives for ΩDSL. As a reference, we also included a 280 MHz profile, which is capable to deliver over 2Gb/s for the shortest looplengths.

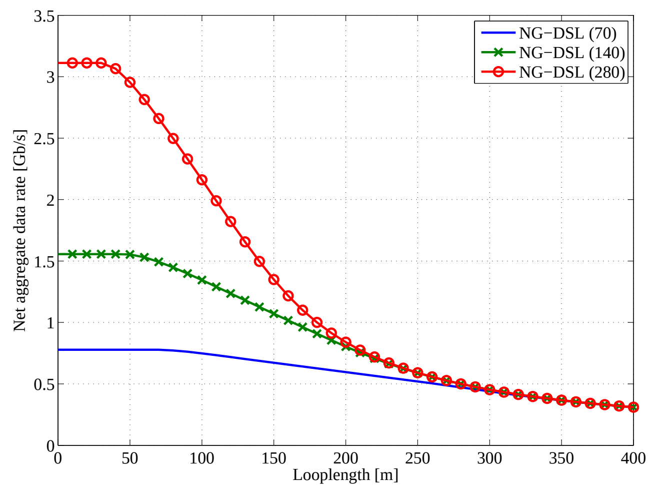

Fig. 4. Using parameters listed in Table I, we show that a 140 MHz profile will reach 1Gb/s for the target looplengths.

IV. METHODOLOGY FOR DMT COMPLEXITY ANALYSIS

In this section, we discuss the methodology for complexity analysis of DMT modulation schemes. We rely on the work, presented in [5], which we summarize below. In this paper, however, we rely on the number of multiply-accumulates per second (R_{MAC}) and the number of memory locations per second (R_{MEM}) for comparison of the different DSL technologies, rather than on the number of multiply-accumulates per bit as done in [5].

TABLE II
COMPLEXITY OF A DMT SYSTEM [5].

Operation Block	N_{MAC}	N_{MEM}
Transmit gain scaling	N	N
IFFT	$0.75N\log_2 N$	$1.5N$
Tx-windowing	2ν	2ν
Rx-windowing	2ν	2ν
FFT	$0.75N\log_2 N$	N
FEQ	$1.5N$	$2N$
Total	$2.5N + 1.5N\log_2 N + 4\nu$	$5.5N + 4\nu$

In Table II, the number of multiply-accumulates per symbol (N_{MAC}) and the number of memory locations per symbol (N_{MEM}) are shown for each block of the DMT modulation scheme. Below, we indicate the number of multiply-accumulates, $N_{\mathrm{MAC,b}}$, in (4) and memory locations, $N_{\mathrm{MEM,b}}$ in (5) to transmit a single bit. We also define the precision-scaled metrics, $NB_{\mathrm{MAC,b}}$ in (6) and $NB_{\mathrm{MEM,b}}$ in (7), respectively [5].

$$N_{\mathrm{MAC,b}} = \frac{N_{\mathrm{MAC}}}{b_{symb}}, \quad (4)$$

$$N_{\mathrm{MEM,b}} = \frac{N_{\mathrm{MEM}}}{b_{symb}}, \quad (5)$$

$$NB_{\mathrm{MAC,b}} = N_{\mathrm{MAC,b}} B_{\mathrm{MAC}}, \quad (6)$$

$$NB_{\mathrm{MEM,b}} = N_{\mathrm{MEM,b}} B_{\mathrm{MEM}}, \quad (7)$$

where B_{MAC} and B_{MEM} are the precisions of the multiply-accumulates (MAC) in bits and the number of bits per word in the memory, respectively. The comparison metrics, R_{MAC} and R_{MEM}, can readily be found as follows:

$$R_{\mathrm{MAC}} = N_{\mathrm{MAC,b}} R, \quad (8)$$

$$R_{\mathrm{MEM}} = N_{\mathrm{MEM,b}} R, \quad (9)$$

where R is the rate (throughput) of the DMT system in b/s. The total complexity of a DMT scheme can be found in Table II. Below, we summarize how the precisions, B_{MAC} and B_{MEM}, can be found. The signal-to-quantization-noise ratio (SQNR) at the ADC output is given by [11]:

$$\mathrm{SQNR} = 6.02B_{ADC} + 4.77 - \mathrm{PAPR[dB]}, \quad (10)$$

where B_{ADC} is the ADC precision, PAPR is the peak-to-average power ratio in dB of the underlying modulation scheme. The SQNR needs to configured so that the quantization noise only has limited impact on the bit-error-ratio (BER). Typically, the quantization noise should have an impact below 0.25 dB:

$$SQNR > \max_i SNR_i + 12.27[\mathrm{dB}], \quad (11)$$

assuming flat noise. The ADC precision can then be derived as:

$$B_{ADC} = \left\lceil \frac{1}{6.02}\left(\max(\mathrm{SNR}_i) + \mathrm{PAPR} + 7.5\right)\right\rceil. \quad (12)$$

The most complex blocks in DMT modulation are the FFT and IFFT blocks. The quantization noise, $\mathrm{SQNR_{FFT}}$, on these operations needs to be less than the ADC quantization noise. The $\mathrm{SQNR_{FFT}}$ can be expressed as [11]:

$$\mathrm{SQNR_{FFT}} = 6.02B_{\mathrm{FFT}} - 12.64 - 10\log N, \quad (13)$$

which leads to:

$$B_{\mathrm{FFT}} > \lceil B_{\mathrm{ADC}} + 1.67\log N - 0.17\mathrm{PAPR} + 2.8\rceil, \quad (14)$$

where B_{FFT} is the precision of the FFT in bits. Using Table II and equations (3) and (4-14), we can find the complexity metrics as:

$$R_{\mathrm{MAC}} = B_{\mathrm{FFT}}^2\left((2.5 + 1.5\log_2 N)(1-\alpha) + 4\alpha\right) f_s, \quad (15)$$

$$R_{\mathrm{MEM}} = B_{\mathrm{FFT}}\left(5.5(1-\alpha) + 4\alpha\right) f_s, \quad (16)$$

where α is the relative overhead due to the CE and f_s the sampling frequency.

TABLE III
THE DIFFERENT DSL FLAVORS, CONSIDERED IN THIS ANALYSIS.

DSL flavor	Year of standardization by ITU	b_{max}	n_c	Δ_c[kHz]
ADSL	2001	15	256	4.3125
ADSL2	2002	15	256	4.3125
ADSL2+	2003	15	512	4.3125
VDSL DMT	2004	15	2782	4.3125
VDSL2-8	2006	15	2048	4.3125
VDSL2-12	2006	15	2782	4.3125
VDSL2-17	2008	15	4096	4.3125
VDSL2-30	2008	15	3478	8.625
G.hn	2010	12	2048	48.82
ΩDSL (70)	?	15	4096	17.25
ΩDSL (140)	?	15	8192	17.25
ΩDSL (280)	?	15	16384	17.25

V. RESULTS AND EXTRAPOLATION TO ΩDSL

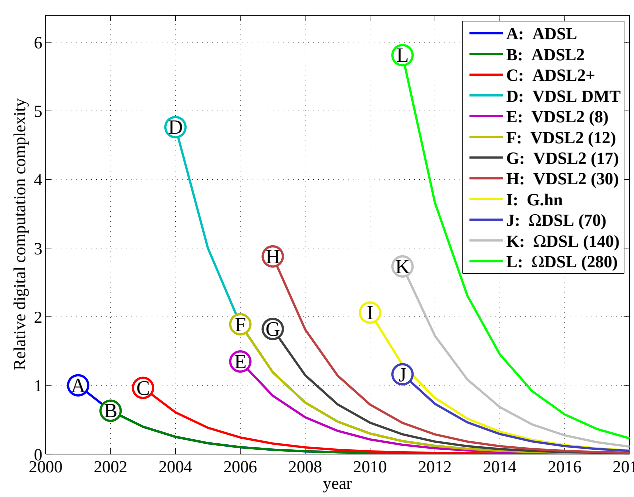

Fig. 5. Relative computational complexity of the different DSL flavors as a function of time.

In Table III, we present an overview on the different DMT-based DSL technologies. For ΩDSL, we include three proposals, all using a carrier spacing of 17.25 kHz, which is deemed appropriate for the looplengths we are considering

 ISBN: 978-1-61208-142-7

for this future technology [12]. However, the results qualitatively hold for other carrier spacing. In Fig. 5, we see the precision-scaled complexity, NB_{MAC}, relative to ADSL in its ITU standardization year of 2001. This relative complexity, $\overline{R_{MAC}(x,y)}$, metric of DSL technology x in year y is based on Moore's law that states that the transistor density on a chip doubles each 18 months [6]. While studies show a flattening of this law, we assume that it holds to facilitate our analysis:

$$\overline{R_{MAC}(x,y)} = \frac{R_{MAC}(x)}{R_{MAC}(ADSL)2^{\frac{2}{3}(y-2001)}}. \tag{17}$$

A few interesting things can be observed. First, ADSL seems to be a cheap technology. However, the technology for ADSL was already available in 1995, long before the ITU standardization, making it a complex technique at market introduction. The pressure on complexity is mainly related to an increased competitive market.

The innovation from ADSL2 with respect to ADSL were mainly targeted at coding, using the same modulation scheme. This is why they come out equivalently complex in this comparison. The main reason why ADSL2 did not become a commercial success is because it was standardized by ITU almost simultaneously as ADSL2plus and the complexity (cost) of ADSL2plus was acceptable to allow cost-effective and dense line cards and customer premise equipment. Furthermore, we see that all standards, when introduced, have a relative complexity metric in the range $[0.6, 1.6]$. The most prominent outlier is VDSL1 at a relative complexity metric of 4.5. This can indicate why VDSL1 only had moderate commercial success, because of the high complexity cost involved. Furthermore, after two years, the 12 MHz profile of VDSL2 was standardized, which uses the same bandwidth and number of carriers, rendering VDSL1 obsolete.

The other outlier is the 30 MHz profile of VDSL2, which has lower deployment volumes than the 17 MHz profile. This is mainly due to the fact that the 30 MHz profile targets FTTB deployment, which takes a long time to establish. Counter examples are Japan and Korea. More interestingly, the relative complexity metric of the three ΩDSL proposals already fall in 2012 within this appropriate standardization range. This is in line with recent start of actvities at standardization bodies. A similar analysis can be performed for the relative precision-scaled memory complexity metric, $\overline{R_{MEM}(x,y)}$ (see Fig. 6). Given the fact that N is the main contributing factor to both complexity and memory, we draw similar conclusions for the memory complexity.

VI. Conclusions

In this paper, we have discussed the digital complexity of different DSL flavors. We have shown that successful adoption of DMT-based DSL technologies occurs in a certain complexity range, when corrected with Moore's Law. Recently, a next-generation DSL technology is being targeted for very short looplengths and very high data rates. We have shown that the opportunity window for standardization of ΩDSL has come, using an extrapolation of the digital complexity.

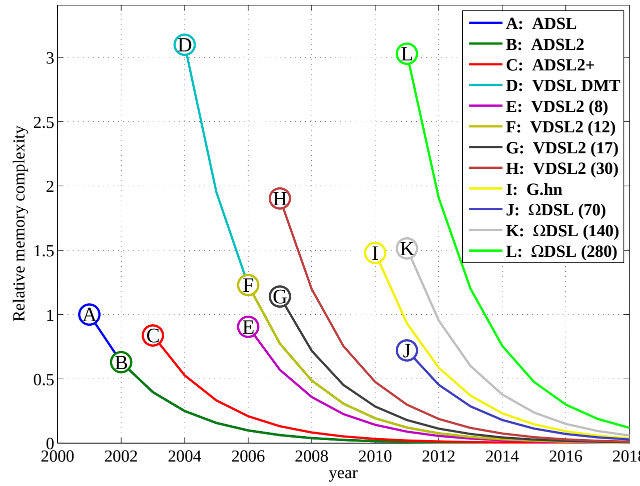

Fig. 6. Relative memory complexity of the different DSL flavors as a function of time.

Acknowledgments

This work was supported by the IWT under contract number 100334, Phanter. The authors would like to thank Paul Spruyt and Adriaan De Lind Van Wijngaarden for their careful review.

References

[1] J. Bourne, "Fiber to the home: practically a reality," in *Proc. of the IEEE International Conference on Communications (ICC)*, 1988, pp. 890–894.
[2] T. Mastrangelo, "Annual Market Outlook Report," *Broadband Trends*, July 2009.
[3] M. Peeters and S. Vanhastel, "The Copper Phantom," *OSP Magazine*, Jan. 2011.
[4] L. Humphrey, "Fibre to the DP," *Broadband Forum*, Dec. 2010.
[5] B. Shim and N. R. Shanbhag, "Complexity analysis of multicarrier and single-carrier systems for very high-speed digital subscriber line," *IEEE Transactions on Signal Processing*, vol. 51, no. 1, Jan. 2003.
[6] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, pp. 114–117, Apr. 1965.
[7] J. Bingham, "Multicarrier modulation for data transmission: An idea whose time has come," *IEEE Communications Magazine*, vol. 28, pp. 5–14, 1990.
[8] Zhone, "Choosing the right fttx architecture," Retrieved on January 20 2010 from http://www.zhone.com/solutions/docs/zti-wp-fttx_choices.pdf.
[9] C. E. Shannon, "A Mathematical Theory of Communication," vol. 27, pp. 379–423, July 1948.
[10] V. Oksman, "The ITU-T's new g.vector standard proliferates 100 mb/s dsl," *IEEE Communications Magazine*, vol. 48, no. 10, pp. 140–148, Oct. 2010.
[11] A. V. Oppenheim and R. W. Schafer, "Discrete-time signal processing," *Englewood Cliffs, NJ: Prentice-Hall*, pp. 637–641, 1989.
[12] J. Maes, M. Timmers, and M. Guenach, "Optimized modulation parameters for a next-generation dsl," *Technical Report*, Jan. 2011.

 ISBN: 978-1-61208-142-7