FP6-2002-IST-1-002131

PICMOS

Photonic Interconnect Layer on CMOS by wafer-scale integration

STReP - Specific Targeted Research Project

IST – Information Society Technologies

D1.3. Updated models and specifications based on data from WP3-5

Contributors	
CNRS-LEOM	**I. O'Connor**

Date of preparation : 30 March 2007

Project Co-ordinator: D. Van Thourhout *Revision*: 1.0

Dissemination level: RE (restricted)

Table of contents

1. Introduction

The aim of deliverable D1.3 is to complete optical-electrical link comparison based on updated electrical data and a new "optimistic" set of optical device characteristic data.

In previous work we completed the development of behavioural models enabling complete photonic link simulation (M1.2), as well as software for electrical and optical link synthesis and analysis (M1.3). Comparisons of electrical and optical link performance were carried out. This work was reported in deliverable D1.2, and represents the main result for WP1.

The goal was then to work towards D1.3 by defining an "optimistic" set of optical device characteristics and repeating the electrical-optical comparisons. Also, shortly after the completion of D1.2, an improved version of all of the PTM transistor models (from 130nm down to 32nm) was released at http://www.eas.asu.edu/~ptm/. In contrast to the previous models, which dated from three different releases, the present model set is said to exhibit a much better performance scaling across technology nodes. An additional goal was therefore to integrate these models into the synthesis toolsets.

In collaboration with WP3-5 we defined the "optimistic" device characteristic set (Table 1) to carry out further analysis of optical interconnect performance metrics. We also integrated the new PTM transistor models into both optical and electrical link synthesis toolsets and repeated the complete set of electrical interconnect evaluation experiments using these new models. The resulting interconnect performance scales more regularly.

Symbol	Description	Old value	New value	Units
E_s	Source total efficiency	0.1	**0.3**	mW/mA
A	Source/detector area	100x100	**10x10**	μm^2
I_t	Source threshold current	1.5e-3	**150e-6**	A
R_d	Detector responsivity	0.5	**0.7**	mA/mW
C_d	Detector capacitance	100e-15	**10e-15**	F

Table 1 Definition of device parameter variation

2. Analysis

To carry out the analyses with the new set of device data, we used the same optical link sizing process as described in D1.2, having previously changed the XML specification file according to the new specification set described in the following section.

2.1. Specifications

The following table (Figure 1) shows the sets of specifications used for analysis and interface circuit sizing. The same three predictive technologies as in D1.2 were

considered for this analysis: BPT models for 65nm, 45nm and 32nm gate length technology nodes. The updated data for the photonic devices is shown in bold type.

Specifications				BPT65	BPT45	BPT45
		Perfomance name	**units**	**Model value**	**Model value**	**Model value**
link	OPPLink	BER	bit/s	1.00E-18	1.00E-18	1.00E-18
		ITRS max frequency	bit/s	2.98E+09	5.20E+09	1.10E+10
		Length	m	{2500um,20000um}	{2500um,20000um}	{2500um,20000um}
		Activity rate		1	1	1
		Ambient temperature	℃	70	70	70
		Vdd (CMOS)	V	1.2	1.1	1
		Perfomance name	**units**	**Model value**	**Model value**	**Model value**
interface circuits	Driver	Vmodulation	V	1.2	1.1	1
		Vbias	V	0.7	0.7	0.7
		Vdd	V	1.2	1.1	1
		ModulationSpeed	bit/s	2.98E+09	5.20E+09	1.10E+10
		Ibias	A	1.81E-03	1.81E-03	1.81E-03
	Receiver	Vdd	V	1.2	1.1	1
		Speed	bits/s	2.98E+09	5.20E+09	1.10E+10
		Perfomance name	**units**	**Model value**	**Model value**	**Model value**
active devices	Source	TotalEfficiency	mW/mA	**0.3**	**0.3**	**0.3**
		Area	m2	**1.00E-10**	**1.00E-10**	**1.00E-10**
		Bandwidth	Hz	1.00E+10	1.00E+10	1.00E+10
		ThresholdCurrent	A	**1.50E-04**	**1.50E-04**	**1.50E-04**
	Detector	IDark	A	1.00E-18	1.00E-18	1.00E-18
		INoise	A	1.00E-15	1.00E-15	1.00E-15
		Bandwidth	Hz	2.00E+10	2.00E+10	2.00E+10
		Area	m2	**1.00E-10**	**1.00E-10**	**1.00E-10**
		TotalEfficiency	mA/mW	**0.7**	**0.7**	**0.7**
		Capacitance	F	**1.00E-14**	**1.00E-14**	**1.00E-14**
		Perfomance name	**units**	**SOI**	**SOI**	**SOI**
waveguide	InputCoupling	percentLoss (1)	%	0	0	0
	Strip	claddingIndex		1.46	1.46	1.46
		guideIndex		3.45	3.45	3.45
		wavelength	m	1.55e-6	1.55e-6	1.55e-6
		height	m	2.20E-07	2.20E-07	2.20E-07
		width	m	5.00E-07	5.00E-07	5.00E-07
		Loss	dB/cm	2.669	2.669	2.669
		Bend excess loss / 90°turn	dB	0.0269	0.0269	0.0269
		Bend radius	m	2.00E-06	2.00E-06	2.00E-06
		Pitch @ 20dB crosstalk	m	1.10E-06	1.10E-06	1.10E-06
		Line delay (n/ceff)	ps/mm	13.3	13.3	13.3
	OutputCoupling	percentLoss (1)	%	0	0	0
				Si3N4	**Si3N4**	**Si3N4**
waveguide	InputCoupling	percentLoss (1)	%	0	0	0
	Strip	claddingIndex		1.46	1.46	1.46
		guideIndex		2	2	2
		wavelength	m	1.55e-6	1.55e-6	1.55e-6
		height	m	8.00E-07	8.00E-07	8.00E-07
		width	m	4.00E-07	4.00E-07	4.00E-07
		Loss	dB/cm	1.5	1.5	1.5
		Bend excess loss / 90°turn	dB	1.52	1.52	1.52
		Bend radius	m	1.00E-05	1.00E-05	1.00E-05
		Pitch @ 20dB crosstalk	m	2.80E-06	2.80E-06	2.80E-06
		Line delay (n/ceff)	ps/mm	6.7	6.7	6.7
	OutputCoupling	percentLoss (1)	%	0	0	0

(1) coupling loss taken into account in source and detector total efficiency figures

Figure 1 Updated specification set for analysis

To follow the impact of these updated device parameters on system-level metrics, we developed an "impact tree" (Figure 2) showing the immediate impact of the device parameters on individual circuit-level parameters, which in turn impact on the overall system performance. The symbols used are summarized in Table 2.

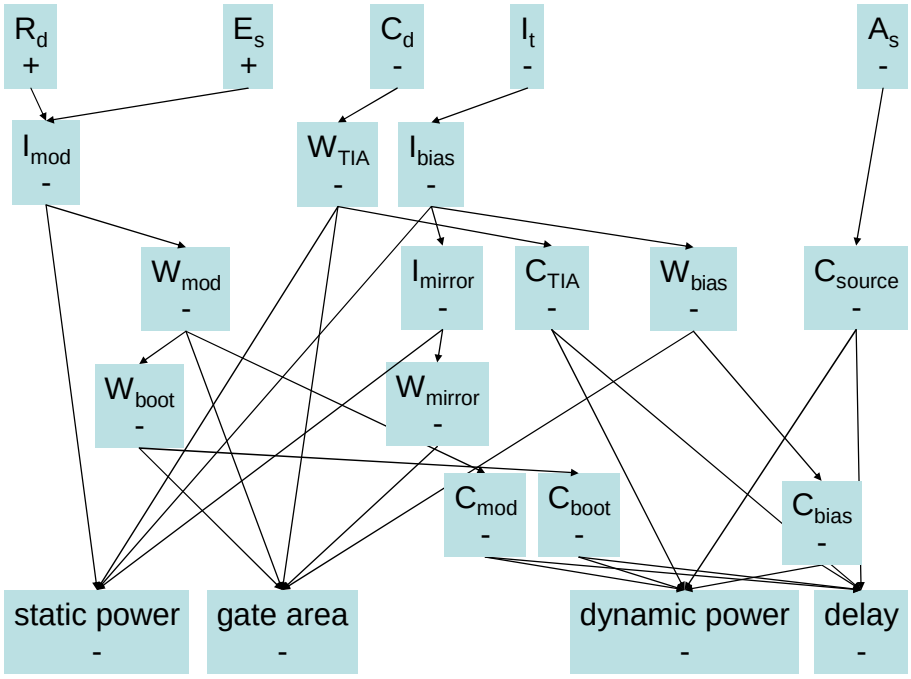

Figure 2 Impact tree showing routes from device-level parameters to performance metrics

Symbol	Description	Units
C_{bias}	Drain capacitance of driver bias transistor	F
C_{boot}	Gate capacitance of driver bootstrap transistors	F
C_{mod}	Drain capacitance of driver modulation transistor	F
C_{source}	Source parasitic capacitance	F
C_{TIA}	Gate capacitance of TIA circuit	F
I_{bias}	Source driver bias current	A
I_{mirror}	Mirror circuit quiescent current	A
I_{mod}	Source driver modulation current	A
W_{bias}	Width of driver bias transistor	µm
W_{boot}	Width of driver bootstrap transistors	µm
W_{mirror}	Width of mirror circuit transistors	µm
W_{mod}	Width of driver modulation transistor	µm
W_{TIA}	Width of TIA circuit transistors	µm

Table 2 Summary of symbols (device and circuit parameters) used in impact tree ()

2.2. Interconnect density and gate area analysis

Figure 3 shows a plot of the maximum achievable number of bi-directional links in a bundle for varying number of IP blocks (from 2 to 10) in a row using routing scheme 2 (described in D1.2). The die size was assumed to be 2cm.

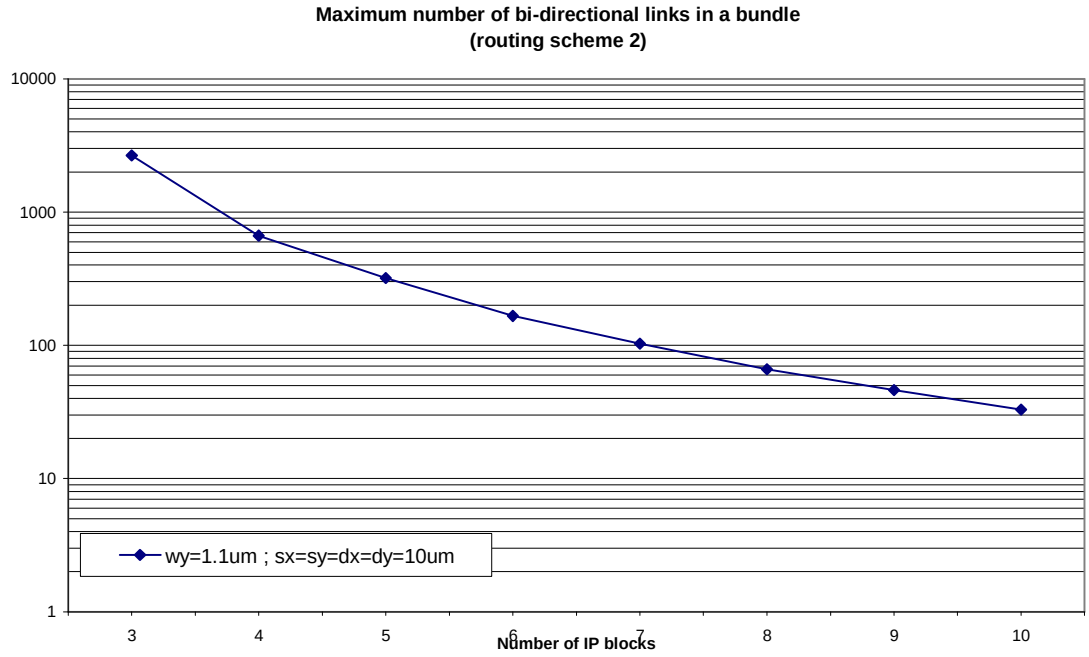

Figure 3 Number of optical links per bundle for varying number of IP blocks (routing scheme 2, pitch = 1.1µm)

Here, in contrast to the figure shown in D1.2 (using 100µm device dimensions), the bundle sizes remain sufficiently large. The number of links per bundle is greater than 32 for all considered values of IP blocks in a row, greater than 64 up to 8 IP blocks, and greater than 128 up to 6 IP blocks.

The optical link sizing method of D1.2 was applied according to the specifications for the BPT 65nm, 45nm and 32nm technologies. Figure 4 shows the results in terms of gate area (i.e. transistor channel dimensions only), extracted as the sum of all transistor gate channel areas W·L. The previous results from D1.2 are indicated in dotted lines.

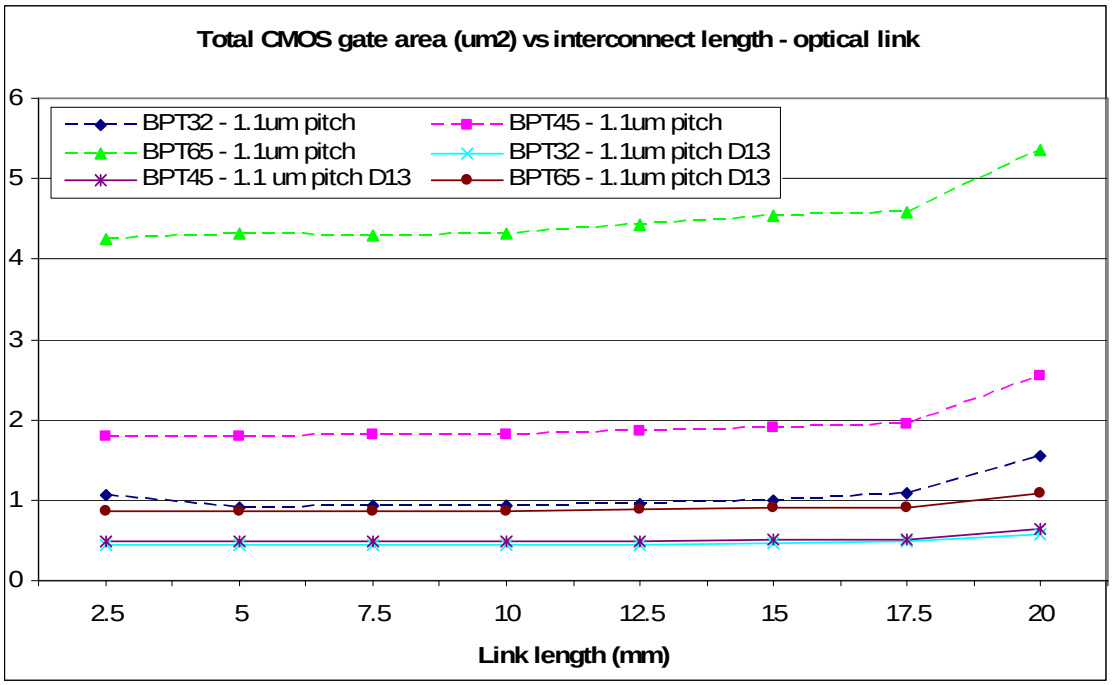

Figure 4 Total CMOS gate area (µm^2) for varying optical interconnect length and technologies

The new specification set reduces total CMOS gate area by a factor of between 2 (BPT 32nm) to 4 (BPT65nm) with respect to the previous analyses.

2.3. Delay analysis

The optical link sizing method of D1.2 was applied according to the specifications for the BPT 65nm, 45nm and 32nm technologies. As before, delay was calculated as :

$$\Delta t = t_{V_{out}=\frac{V_{pit_{\max}}+V_{out_{\min}}}{2},slope=+,5} - t_{V_{in}=\frac{V_{in_{\max}}+V_{in_{\min}}}{2},slope=+,5}$$

where the notation

$$t_{V=V_1,slope=\{+,-\},n}$$

signifies the time corresponding to the nth simulation point where the quantity V is equal to V_1 on a rising (+) or falling (-) slope. Also calculated as a point of reference was the intrinsic waveguide delay, using t_{TOF}[*]=13.3ps/mm for the Si/SiO_2 waveguides. Figure 5 shows the delay results for varying link lengths. The previous results from D1.2 are indicated in dotted lines.

[*] TOF: Time of Flight

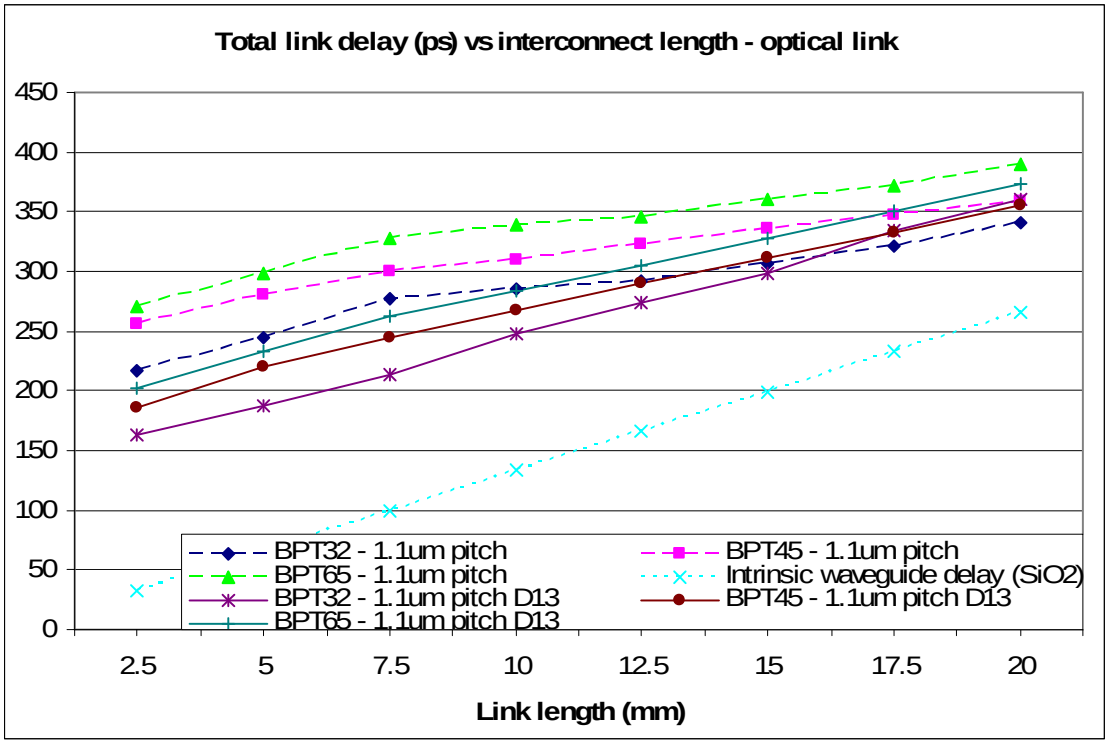

Figure 5 Total link delay (ps) for varying optical interconnect length and technologies

Significant improvement is not demonstrated here. At best a further 50ps delay reduction can be achieved at the shorter link lengths (where optical links are less likely to be used). Towards the longer links the higher driver modulation current I_m, required to compensate higher overall waveguide loss, is able to drive the source capacitance faster and tends towards a small contribution to overall delay. As such the delay is dominated at these lengths by parameters unchanged from the previous analysis (detector and source delay, waveguide delay, comparator delay) and the new delay figures match the previous figures.

2.4. Power analysis

The optical link sizing method of D1.2 was applied according to the specifications for the BPT 65nm, 45nm and 32nm technologies. As before, the average static power was extracted from transient simulations using:

$$\overline{P} = \frac{I_{source_0} + I_{source_1}}{2} \cdot V_p + \frac{I_{\det_0} + I_{\det_1}}{2} \cdot V_d + \frac{I_{cct_0} + I_{cct_1}}{2} \cdot V_{dd}$$

where I_{source}, I_{det} and I_{cct} represent the currents flowing through the source, detector and circuit voltage supplies of V_p, V_d and V_{dd} respectively. Figure 6 shows the average static power results for varying link lengths. The previous results from D1.2 are indicated in dotted lines.

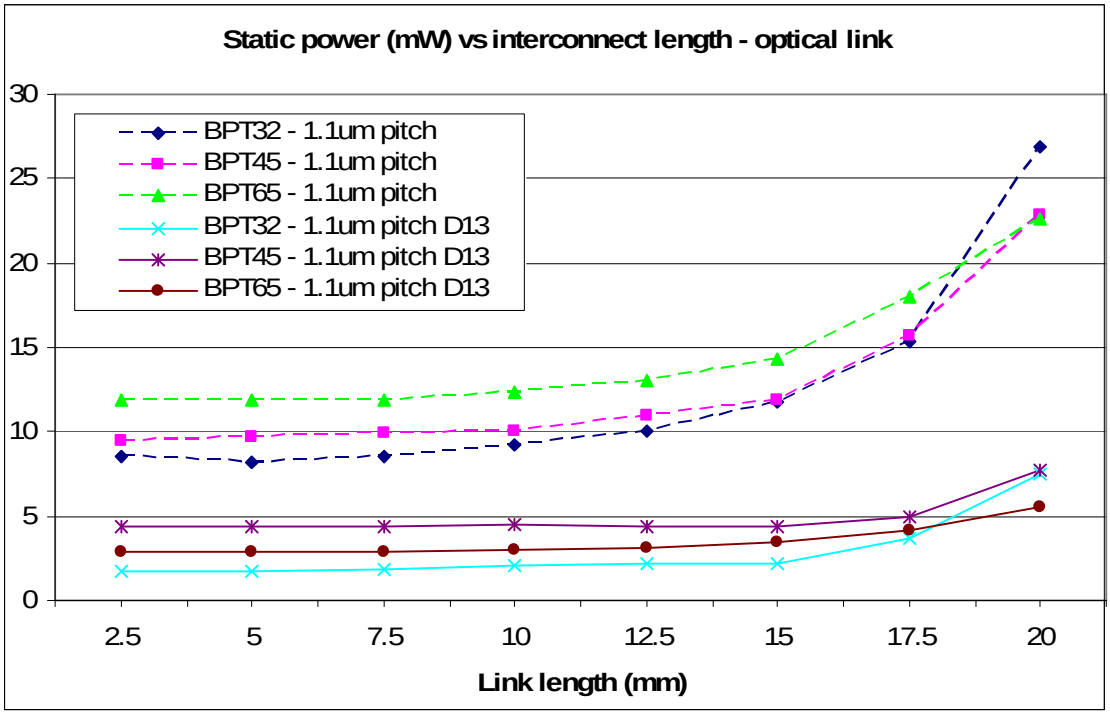

Figure 6 Average static power (in mW) vs interconnect length for BPT 65nm, 45nm and 32nm technologies

Significant reductions in static power are observed with the new specification set (a minimum factor of 2 for BPT45 and 3 for BPT32 at low link lengths; and a maximum factor of 4 for BPT65 and BPT32 at high link lengths). This result is due mainly to the lower source threshold current, but also to higher source efficiency and detector responsivity leading to lower modulation currents, and to lower detector capacitance which lowers the receiver circuit quiescent current.

The dynamic power is calculated from rising and falling edge transitions (the average switching energy extracted from simulations as the integral of supply currents in edge transitions). Figure 7 shows the dynamic power results for varying link lengths. The previous results from D1.2 are indicated in dotted lines.

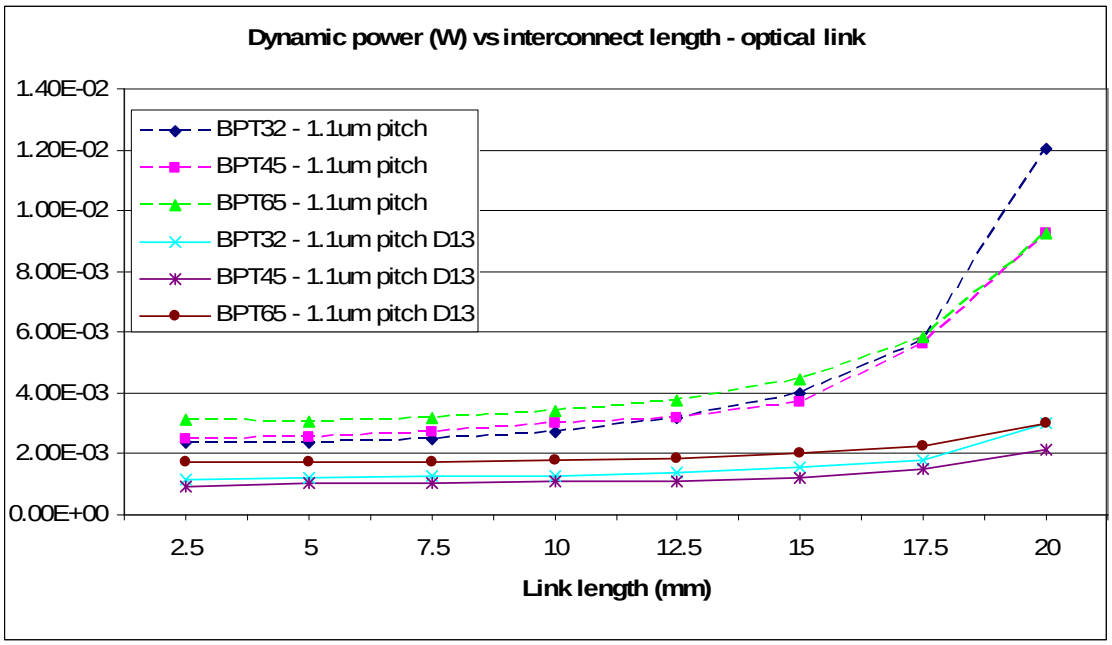

Figure 7 Average dynamic power (in W) vs interconnect length for BPT 65nm, 45nm and 32nm technologies

Here power reductions of the order of 2 to 4 are observed. This is attributed mainly to lower transistor capacitances due to lower bias current in the driver (due to the reduction in source threshold current), and to lower modulation currents in the driver (due to the increase in source efficiency and detector responsivity).

Using static and dynamic power information, the total energy can be calculated (Figure 8). The previous results from D1.2 are indicated in dotted lines.

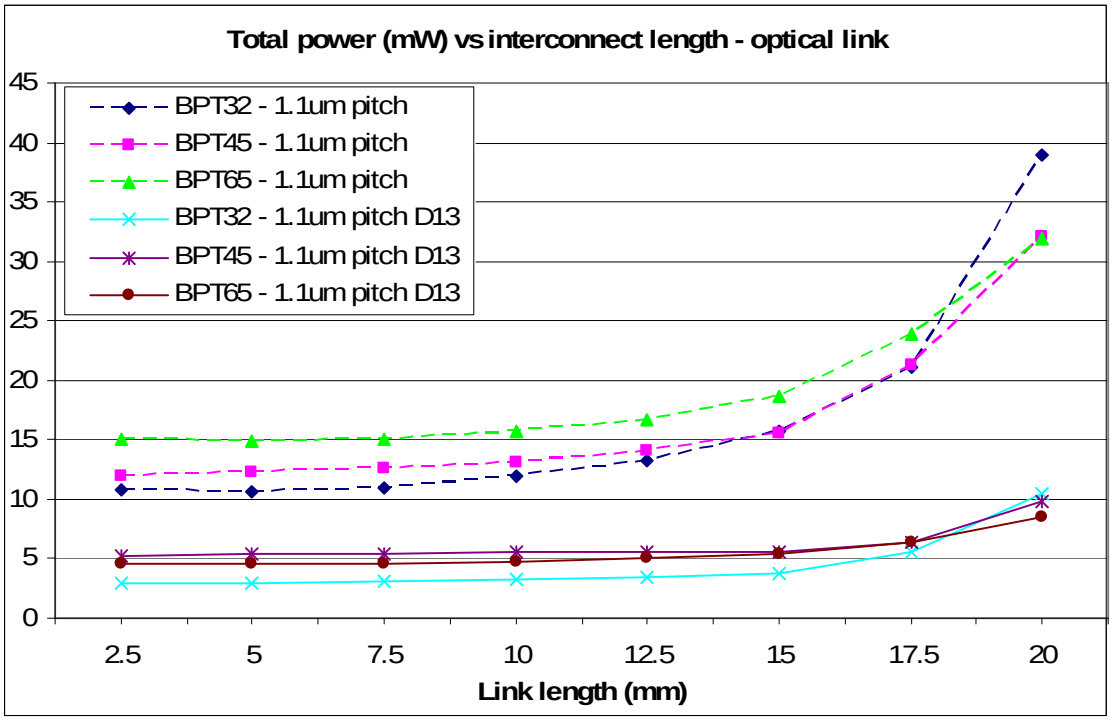

Figure 8 Total power (in mW) vs interconnect length for BPT 65nm, 45nm and 32nm technologies

Overall power is shown to be reduced by factors of between 2 and 4 with respect to previous analyses. The greatest reduction is achieved at higher link lengths, which is the expected context for the use of such optical links.

2.5. Conclusion and impact on project

These results show how device-level specifications impact on the overall performance of links. There is little particular impact for the project since this deliverable arrives at the end of the project; however the results will be disseminated to generate further interest in the viability of the optical links developed in the project. The results will be disseminated in particular through updated comparisons with electrical data (detailed in section 3).

3. Electrical-optical comparison

3.1. Introduction

For all electrical analyses, the previous data from D1.2 was used. The simulation conditions were:

- invertors with a 2/1 PMOS-to-NMOS ratio were used
- for each CMOS technology used, the maximal overall link input capacitance was restricted to that of a CMOS invertor with minimal gate length (defined as 2λ) and 60λ and 30λ for the PMOS and NMOS widths, respectively
- the minimal output drive strength was set to that of the same invertor

For all optical analyses, the data was taken from the results described in the previous section.

The comparison results that are presented in the following sections are between optical interconnect and 1.1µm pitch unshielded electrical interconnect. The results are presented in the form of reduction factors, calculated as P_e/P_o, where P_o represents the optical performance figure and P_e represents the electrical interconnect figure (where smaller performance figures for area, delay and power mean improved performance).

3.2. Results

3.2.1. Area

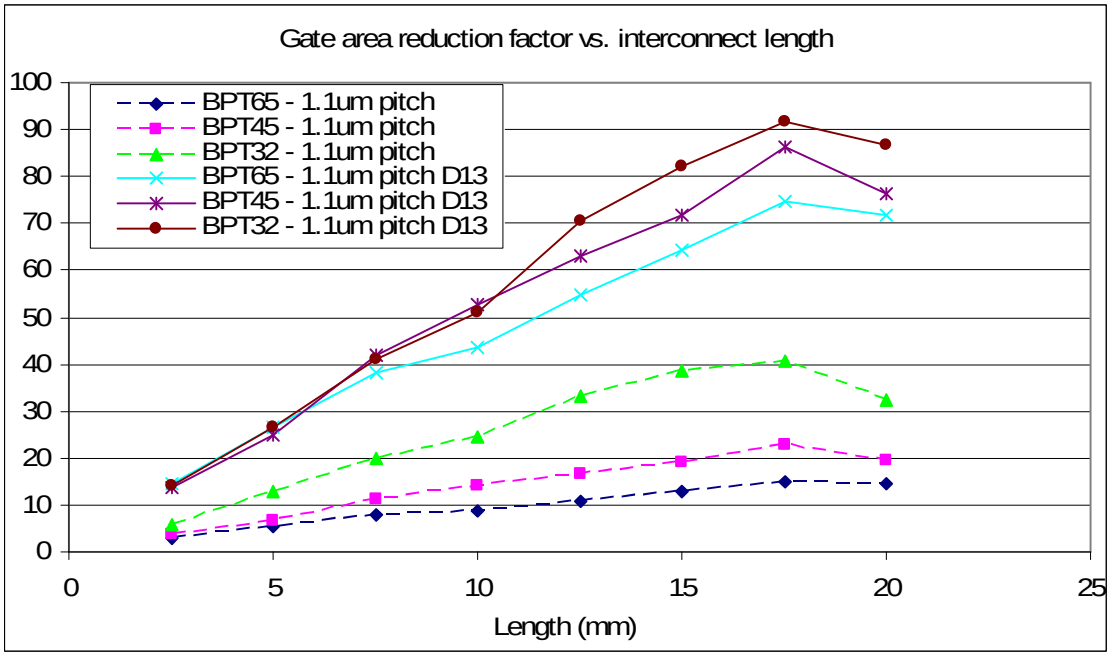

Figure 9 Gate area comparison for varying interconnect length and technologies at 1.1µm pitch: reduction factor

These figures show an impressive reduction in gate area in favour of optical interconnect (of the order of 60x-90x for link lengths above 1cm and for the two most advanced technology nodes). Optical interconnect will thus introduce a significantly lower area penalty for data routing functions at transistor-level. These figures are to be considered in the context of scenarios indicating the use of up to 25% of transistors on chip in electrical interconnect buffers. While only a part of this number is used for long links, it is still clear that the use of optical links will free up a large number of transistors for use in functions other than interconnect (data processing, memory …)

3.2.2. Delay

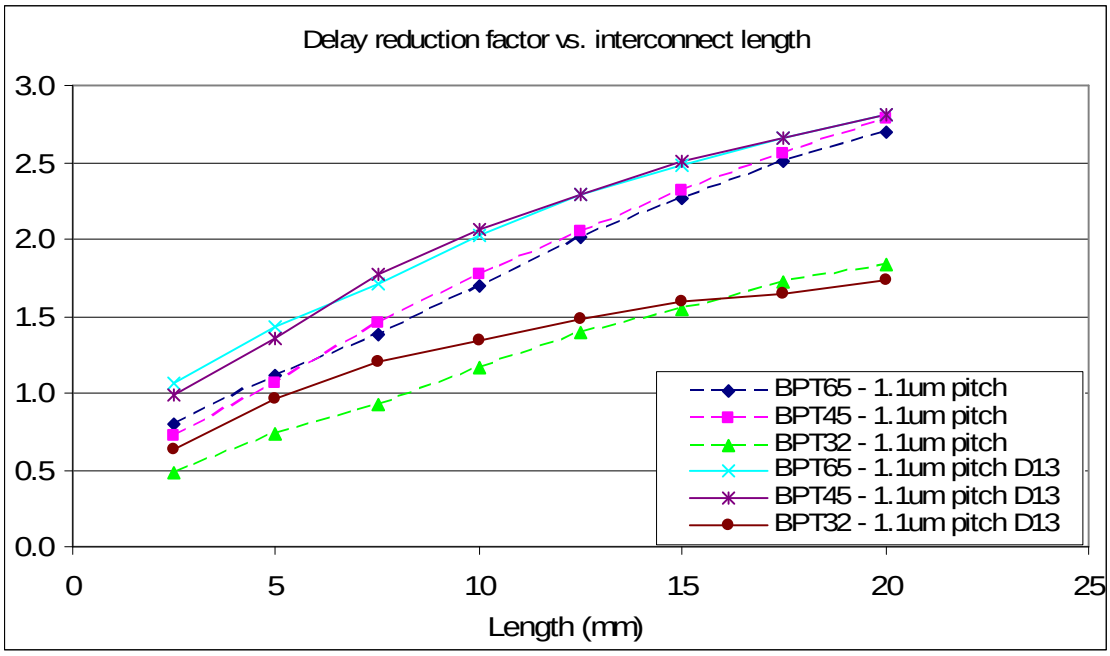

Figure 10 Delay comparison for varying interconnect length and technologies at 1.1µm pitch: reduction factor

As in D1.2, it is shown that optical interconnect will have a slight advantage over electrical interconnect in terms of delay for long interconnect lengths (above 10mm). The basic reason for this is that delay for optical interconnect does not depend as strongly on interconnect length as electrical interconnect, because no additional circuit stages are added – the increase stems from higher intrinsic waveguide delay only. However, the advantage decreases for more advanced technology nodes and indeed does not achieve any delay reduction (actually the opposite) for unscaled interconnect at 32nm gate length.

3.2.3. Power

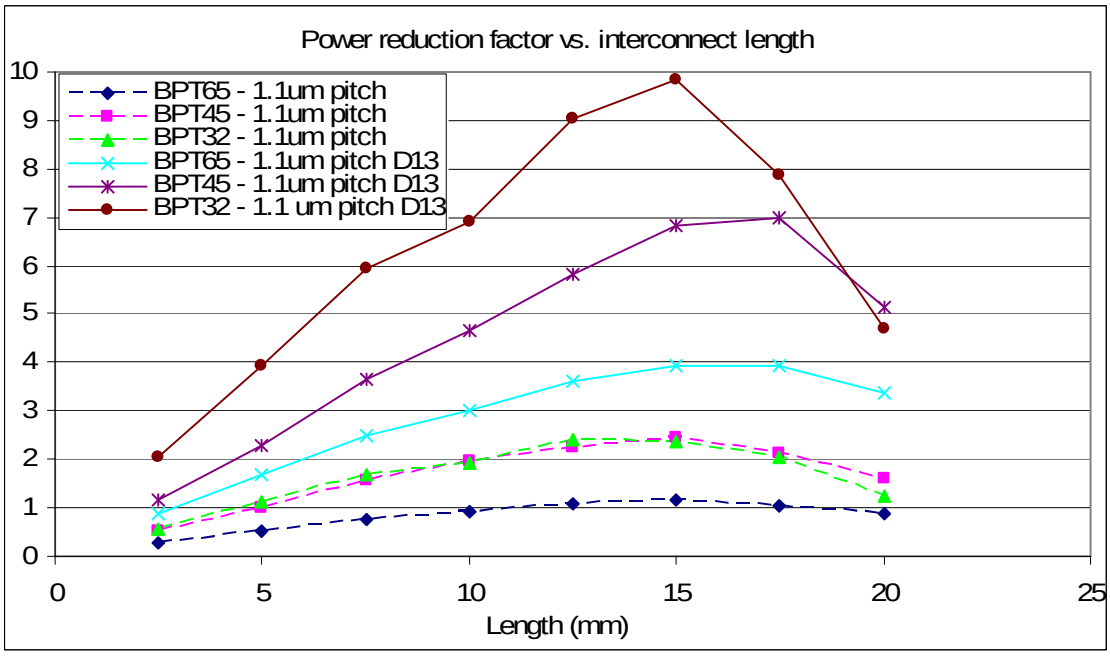

Figure 11 Average total power comparison for varying interconnect length and technologies at 1.1µm pitch: reduction factor

In contrast to the conclusions of the previous analyses in D1.2, power reduction can be considered here to be a major argument in favour of optical interconnect. In D1.2, the static power comparison was the weak point for optical interconnect, due to continuous biasing of the source (avoiding turn-on times to achieve the required bit rate) and of the receiver circuit (the circuit bandwidth is directly related to quiescent bias current). The reduction in source threshold current and in detector capacitance has a significant impact on both these factors, to the extent that static power in optical links, while still higher than that of electrical interconnect, is no longer dominant. It is likely that comparisons using technologies with transistor gate lengths below 32nm will further improve this comparison since on the electrical side the static power dissipation will increase with leakage current. We can thus consider that the optical device improvements, if feasible, constitute the main path to the solution to the following recommendation in D1.2 ("*Power, and particularly static power, is therefore a key performance metric to optimise during exploration of optical interconnect device specifications*").

Dynamic power in optical interconnect is further reduced with the smaller overall circuit transistors used in this analysis. The total power comparison shows power reduction factors between 5x-10x for link lengths above 10mm and for the two most advanced technology nodes.

3.3. Conclusion and impact on project

This investigation program shows how optical interconnect performance metrics compare to those of electrical interconnect performance metrics, using the "optimistic" set of

device characteristics. Optical interconnect compares favourably for longer link lengths and more advanced technology. This is particularly true for gate area, and also now for power. It is also true to some extent for delay. However, delay figures still show reduced advantage to optical interconnect for more advanced technology nodes. The figures for total power are also still non-monotonic with respect to interconnect length, implying that there exists an optimum length for which power reduction is maximised.

As stated in the previous section, there is little particular impact of these results for the project since this deliverable arrives at the end of the project. The results will be disseminated to generate further interest in the viability of the optical links developed in the project.

4. Conclusion

This report has summarised the application of the tools, methods and models developed in WP1 and described in D1.2 to regenerate data based on a new set of device specifications. The main results are reproduced here:

- for gate area, optical interconnect gives a reduction in gate area with respect to electrical interconnect of the order of 60x-90x for link lengths above 1cm and for the two most advanced technology nodes
- for delay, optical interconnect demonstrates a slight but not significant advantage over electrical interconnect for long interconnect lengths (above 10mm)
- for total power, optical interconnect gives power reduction factors of between 5x-10x for link lengths above 10mm and for the two most advanced technology nodes

5. Dissemination

5.1. Journals

I. O'Connor, F. Tissafi-Drissi, F. Gaffiot, J. Dambre, M. De Wilde, D. Stroobandt, J. Van Campenhout, D. Van Thourhout, "Systematic Simulation-Based Predictive Synthesis of Integrated Optical Interconnect," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, accepted for publication 2007

5.2. Invited conference papers

I. O'Connor, F. Tissafi-Drissi, D. Navarro, F. Mieyeville, F. Gaffiot, J. Dambre, M. De Wilde, D. Stroobandt, M. Brière, "Integrated optical interconnect for on-chip data transport," *International IEEE-NEWCAS Conference*, June 18-21, 2006, Gatineau, Canada