IEEE Photonics Journal
An IEEE Photonics Society Publication

Energy Consumption and Energy Density in Optical and Electronic Signal Processing

Volume 3, Number 5, October 2011

Rodney S. Tucker, Fellow, IEEE
Kerry Hinton

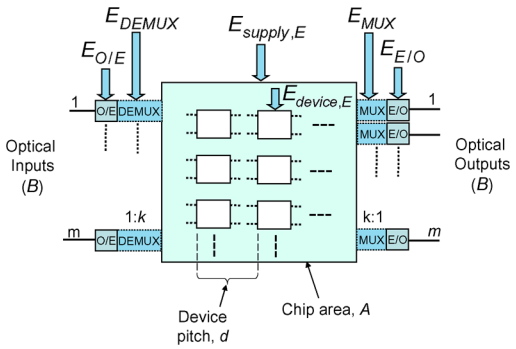

DOI: 10.1109/JPHOT.2011.2166254
1943-0655/$26.00 ©2011 IEEE

Energy Consumption and Energy Density in Optical and Electronic Signal Processing

Rodney S. Tucker, *Fellow, IEEE*, and Kerry Hinton

Institute for a Broadband-Enabled Society (IBES) and Centre for Energy-Efficient Telecommunications (CEET), Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne, Vic. 3010, Australia

DOI: 10.1109/JPHOT.2011.2166254
1943-0655/$26.00 ©2011 IEEE

Manuscript received July 19, 2011; revised August 16, 2011; accepted August 21, 2011. Date of publication August 30, 2011; date of current version September 23, 2011. Corresponding author: R. Tucker (e-mail: r.tucker@unimelb.edu.au).

Abstract: We compare the energy consumption of digital optical and digital electronic signal processing circuits, including the contributions to energy consumption of the optical to electrical (O/E) converters and electrical to optical (E/O) converters, and the demultiplexers (DEMUXs) and multiplexers (MUXs) required for electronic circuits to process high-speed optical signals. This paper focuses on three key practical considerations, namely, energy consumption, energy density, and the complexity of processing. We show that optical signal processing is potentially competitive with electronics in very high-speed circuits that provide only limited processing, i.e., when only a small number of processing operations are performed on each bit of data. However, in applications that require anything more than limited processing, electronics provides better energy efficiency and occupies a smaller footprint. In these applications, electronics is likely to remain the technology of choice. More attention needs to be paid to energy consumption issues in the research and development of new digital optical technologies.

Index Terms: Optical signal processing, electronic signal processing, energy consumption, energy efficiency.

1. Introduction

Today's society is heavily reliant on a sophisticated and ubiquitous cyber infrastructure underpinned by advanced digital signal processing technologies. In communications networks, signal processing functions range from computing-intensive processes, such as video coding and decoding, as well as forward error correction (FEC), to relatively simple functions, such as signal waveform reshaping and regeneration. The majority of this signal processing today is carried out by electronics, most of it in the form of complementary metal–oxide–semiconductor (CMOS) integrated circuits (ICs).

How does *optical* signal processing fit into the picture, and what is the future potential for digital optical signal processing (sometimes referred to as all-optical signal processing) to replace some electronic signal processing? Nonlinear optical components that can be used for signal processing include semiconductor optical amplifiers (SOAs) [1], periodically polled Lithium Niobate (PPLN) [2], highly nonlinear fibers (HNLFs) [3], and Silicon nanowires [4]. There have been many advances in nonlinear digital optical signal processing over the past 30 years (see, for example [5]). However, despite these advances, digital optical signal processing is not used in commercial telecommunications infrastructure.

There are several reasons why system designers prefer electronic signal processing over optical signal processing: First, electronic signal processing provides outstanding performance and

throughput (i.e., processing power and capacity for very high aggregated data rates). Second, electronic devices are very small. The feature size in state-of-the-art CMOS transistors is on the order of 22 nm—almost two orders of magnitude smaller than an optical wavelength in the communications band and at least three orders of magnitude smaller than practical optical devices. Third, electronic devices can be monolithically integrated with a very high device density. Consequently, electronic devices are many orders of magnitude less expensive than active optical components. Fourth, nearly all proposals for optical logic devices do not satisfy a variety of practical criteria needed for optical logic functionality [6]. Finally, and perhaps most importantly, digital electronic devices consume a relatively small amount of energy.

The importance of low energy consumption cannot be over emphasized. Energy consumption and the associated thermal issues have been a driving consideration behind much of the progress made by the electronic IC industry over the past half century. In comparison, the optical signal processing research community has paid scant attention to the question of energy consumption. A key objective of the present paper is to highlight the importance of energy consumption as a measure of performance in optical signal processing and to provide a benchmark for comparing energy consumption in optical and electronic circuits. We argue that energy consumption issues are critical in the development of new digital optical technologies.

In this paper, we compare signal processing in two domains: electronic and optical. If the signal to be processed is in one domain and the processing is carried out in the other domain, then optical to electrical (O/E) converters and electrical to optical (E/O) converters are required to convert the signal between domains. These conversions add to the overall energy consumption and this energy penalty can clearly be avoided if there is no change of domain. However, if the energy consumption of signal processing in one domain is larger than in the other domain, it may be beneficial to change domains.

A legitimate comparison between optical and electronic processing and switching must take into account the performance limitations imposed and the energy consumed by ancillary components such as the O/E and E/O converters needed to change domains [7], [8]. In addition, if the bit rate of the data to be processed is high, demultiplexers (DEMUXs) and multiplexers (MUXs) may be needed to reduce the bit rate to a level that the electronics can handle. A comparison of energy consumption between technologies also requires that *all* sources of energy consumption are considered, including energy consumed by interconnect wires in CMOS circuits, power-supply energies, and losses in optical circuits.

The purpose of this paper is to compare the energy consumption of digital optical and digital electronic signal processing circuits. We consider the device footprint and the energy density per unit chip area and show how the energy density sets a limitation on the achievable density of integration on a chip. In our comparison, we take into account the limitations imposed by the O/E and E/O converters as well as MUXs and DEMUXs. We provide projections of device and circuit performance out to the year 2020, based on published trends in device technologies such as the International Technology Roadmap for Semiconductors (ITRS) [9].

We argue that optical signal processing may be competitive with electronics in simple high-speed circuits with limited processing power. However, unless there are many orders of magnitude of improvement in the energy efficiency of digital optical devices, electronics is likely to remain the technology of choice for the vast majority of digital signal processing and switching applications.

2. Energy Model

In this section, we develop an energy model of optical and electronic signal processing circuits. The model is based on an earlier model [8] but includes an important clock-frequency-dependent term in the model of nonlinear optical devices that was not included in the earlier model. To ensure a common basis for comparison between optical and electronic circuits, all of the signals processing circuits considered here have optical inputs and optical outputs. This is a mode of operation that is suited to optical signal processing and takes advantage of the ability of optical signal processing circuits to operate directly on optical data in communications systems.

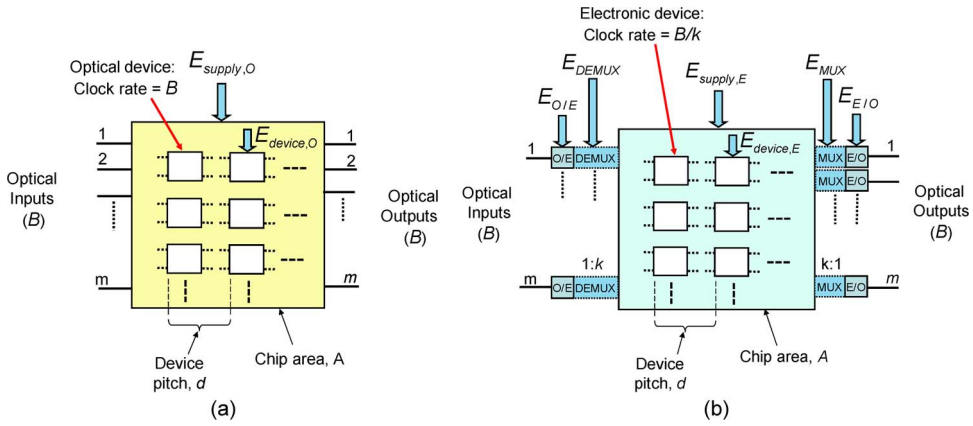

Fig. 1. (a) Optical signal processing circuit. (b) Electronic signal processing circuit.

2.1. Signal Processing Circuits

Fig. 1(a) shows an optical signal processing circuit, and Fig. 1(b) shows an electronic signal processing circuit. Both circuits in Fig. 1 have m optical input ports and m optical output ports. The data rate on each of the input ports in both circuits is B and the aggregate data rate across all input ports is $B_{aggregate} = mB$. In both circuits, the signal processing function is carried out by an interconnected array of digital devices [optical devices in Fig. 1(a) and electronic devices in Fig. 1(b)], shown as rectangles in Fig. 1. These devices are basic signal processing elements such as a logic gates or flip flops that form the building blocks of the overall circuit. In general, these devices have multiple inputs and multiple outputs, as shown in Fig. 1. A fundamentally important requirement on the characteristics of digital devices—whether optical or electronic—is that the input and output signal levels satisfy the requirements of logic level restoration, cascadability, and fan-out capabilities [6], [10], [11].

The chip area of both circuits (excluding O/E/O converters and (DE)MUXs) in Fig. 1 is A, and the spacing between devices (device pitch) is d. Assuming a uniform distribution of devices across the chip, the number of N_{device} devices on each chip is given by $N_{device} = A/d^2$. The supply energy to the signal processing chips in Fig. 1 per bit period of the input data (i.e., per $1/B$ s) is $E_{supply,O}$ and $E_{supply,E}$, and the energy consumed by each device per device operation is $E_{device,O}$ and $E_{device,E}$ for the optical and electronic circuits, respectively. The supply powers $P_{supply,O}$ and $P_{supply,E}$ for the optical and electronic circuits in Fig. 1(a) and (b) are simply the supply energies multiplied by the bit rate, i.e., $P_{supply,O} = BE_{supply,O}$, and $P_{supply,E} = BE_{supply,E}$.

A key differentiating factor between the two circuits in Fig. 1 is that the electronic signal processing circuit in Fig. 1(b) requires O/E conversion at the input ports and E/O conversion at the output ports. Another difference is that the clock frequency (i.e., the processing frequency) f_O of optical circuit in Fig. 1(a) is equal to the bit rate B of the incoming data and the clock period τ_O of the optical circuit is therefore $\tau_O = 1/f_O = 1/B$. In contrast, for electronic circuits, the input bit rate B may be higher than the maximum achievable speed capability of the electronic devices. It is therefore necessary to demultiplex the incoming data to a lower bit rate that matches the clock frequency capabilities of the electronic devices. Similarly, if the output bit rate is larger than the processing speed of the electronics, MUXs are required at the output ports. In Fig. 1(b), the multiplexing and demultiplexing ratio is k. Therefore, clock frequency of the electronic circuit is $f_E = B/k$ and the clock period of the electronic circuit is $\tau_E = 1/f_E = k/B$.

The total energy consumption per bit of all O/E and E/O converters in Fig. 1(b) is $E_{O/E/O} = E_{O/E} + E_{E/O}$, where $E_{O/E}$ and $E_{E/O}$ are the total energies per bit in the O/E converters and E/O converters, respectively. The total energy consumption of all MUXs and DEMUXs per input bit period $(1/B)$ is $E_{(DE)MUX} = E_{DEMUX} + E_{MUX}$, where E_{DEMUX} and E_{MUX} are the total energy per bit of the DEMUX and MUX, respectively.

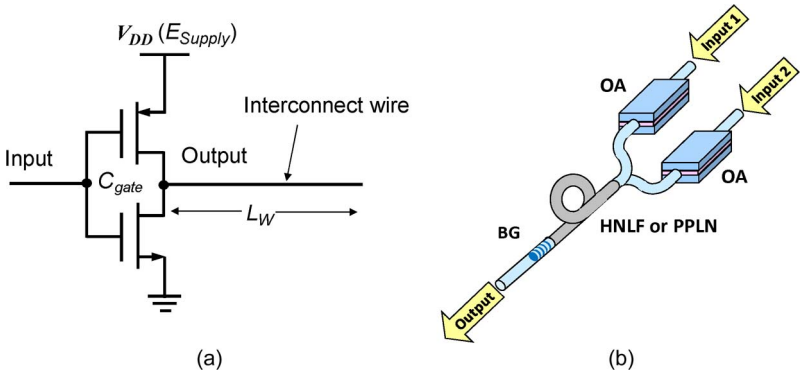

Fig. 2. (a) CMOS inverter circuit. (b) Schematic of a digital optical signal processing device providing switching by new frequency generation. PPLN = periodically polled Lithium Niobate. HNLF = highly nonlinear. OA = optical amplifier. BG = Bragg grating.

It is important to point out that while the clock frequency of the electronic signal processing chip in Fig. 1(b) is reduced from the line rate B to B/k, the O/E and E/O converters and DEMUX and MUX circuits must operate at B. The speed capabilities of these components ultimately limit the achievable input bit rate B in Fig. 1(b). One approach to increasing the effective B for electronic signal processing would be to use optical time division demultiplexing and multiplexing [12] at the optical input and output ports in Fig. 1(b). Another approach would be to replace the m optical inputs and outputs with mk inputs and outputs, each at a bit rate of B/k [6]. Both of these approaches are beyond the scope of this paper.

2.2. Signal Processing Devices

In this section, we develop simple models for the optical devices in Fig. 1(a) and the electronic devices in Fig. 1(b). As explained earlier, the term "device" in this paper refers to a building-block signal processing element such as a logic gate. A building-block device typically comprises a number of basic circuit elements. For example, Fig. 2(a) shows the circuit of a CMOS logic-level inverter connected to a wire that connects this device to the input of another device or devices in the circuit. This building-block device operates with well-defined input and output voltage levels and provides the necessary logic level restoration and fan-out that enables it to serve as a building block in a variety of signal processing circuits.

Fig. 2(b) is an example of a digital optical signal processing device. This particular example of a digital optical device provides switching by new frequency generation. The new frequency generation could be by 3-wave mixing in PPLN or 4-wave mixing in HNLF. In Fig. 2(b), the two inputs are boosted to relatively high power levels using optical amplifiers (OAs) and are used to generate a new (low-power) output. The unwanted frequencies are filtered out by Bragg grating (BG). To provide logic level restoration and a fan-out of 2, the gains of the OAs are set so that the net gain between inputs 1 and 2 and the output is 3 dB. For larger fan-out, the gain will need to be higher. This may require significant power output from the OAs, depending upon the efficiency of the process that generates the new frequency.

As shown in both Fig. 2(a) and (b), building-block devices for digital circuits—both electronic and optical—generally require a number of active devices (e.g., transistors, amplifiers, and nonlinear components). Fig. 3 is a general schematic of a nonlinear device. In Fig. 3, the device is shown with two inputs and two outputs, but more or fewer inputs and/or outputs are possible. The key component of the device in Fig. 3 is one or more nonlinear elements that carry out the signal processing operation. Also included are OAs at the inputs and/or outputs for signal-level restoration and fan-out.

A key parameter in the analysis of digital signal processing devices is the quantum of energy consumed by a device for each digital logic operation performed by that device. This quantum of

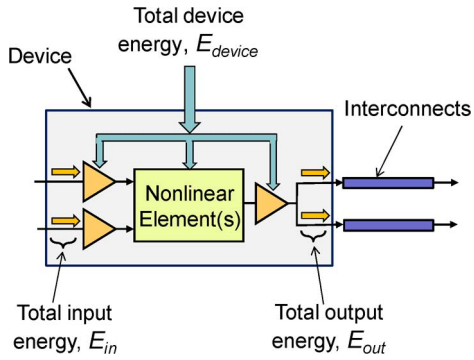

Fig. 3. Optical or electronic signal processing device with two input and two outputs. The device includes input and output amplifiers or buffers, and a nonlinear element.

energy is E_{device}, as defined earlier. As shown in Fig. 3, this energy includes the supply energy to all of the ancillary active components in the device. The total signal input energy per operation on all inputs to the device in Fig. 2 is E_{in}, and the total signal output energy on all outputs is E_{out}. This output energy includes the energy consumed by the interconnects between devices. In general, the interconnect energy is small in optical devices because the losses in optical waveguides are small. However, interconnect energy generally dominates in high-speed electronic circuits [13].

A common problem in the optics literature is that the signal input energy of an optical device is often confused with the "switching energy" of the device. For example, some authors use the term *switching energy* to refer to energy per bit at Input 1 or Input 2 in Fig. 2(b) or at the input to the HNLF or PPLN. Because $E_{in} \ll E_{device}$ [11], this typically leads to a gross underestimation of the energy consumption of the device and provides a misleading picture of the device energy consumption. Unfortunately, many papers do not provide sufficient information for the reader to estimate the total energy consumption, i.e., E_{device} in Fig. 3.

2.3. Device Energy Models

In this section, we present the device energy models used in our analysis. The CMOS device energy $E_{device,E}$ per logic operation is modeled as follows:

$$E_{device,E} = E_{ED}\frac{d}{d_{ref}} + E_{ES}\frac{f_{ref}}{\gamma f_E}. \tag{1}$$

The first term on the right-hand side of (1) is the dynamic energy of the device. It represents the energy consumed in charging and discharging the capacitance of the transistors and the interconnect wires. In general, the capacitance of the interconnect wires dominates over the capacitance of the transistors. Therefore, the dynamic energy supplied to each device per operation is proportional to the length of the interconnect wires L_W [see Fig. 2(a)]. For a given circuit, this wire length scales with the device pitch d. In (1), E_{ED} is the device dynamic energy for $d = d_{ref}$, and d_{ref} is a reference device pitch.

The second term in (1) is the static energy of the device. It represents the energy associated with device leakage currents. This energy is inversely proportional to the effective clock frequency γf_E of the device, where f_E is the clock frequency of the electronic circuit, and γ is the activity factor. The average activity factor is the average number of digital operations that a device performs in each clock period. In (1), E_{ES} is the device dynamic energy for $\gamma f_E = f_{ref}$ and f_{ref} is a reference clock frequency. The static energy consumption of CMOS devices is becoming increasingly important as CMOS devices scale to small sizes [14]. However, in the analysis presented here, we consider circuits with effective clock frequencies γf_E of 1 GHz and higher, where static energy consumption is relatively small [14]. Therefore, we consider only the first term on the right-hand side of (1).

An important difference between digital CMOS devices and digital optical devices is that the dynamic energy per logic operation in a CMOS device is determined by the charging and

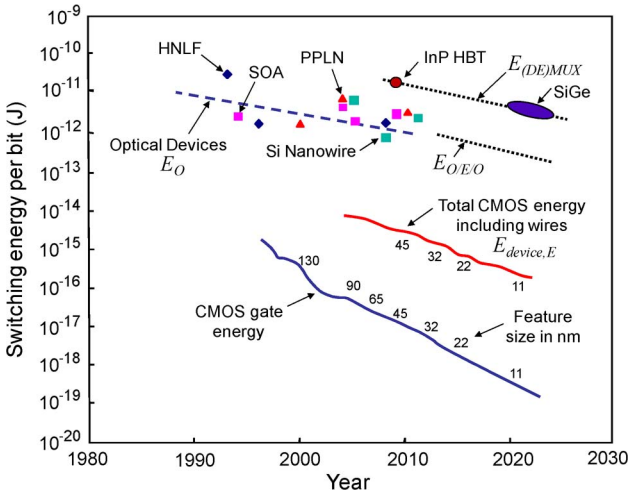

Fig. 4. Device energies per bit against time. The optical device energy per bit E_O is normalized to a reference clock frequency of $f_{ref} = 100$ GHz. HNLF = highly nonlinear fiber, SOA = semiconductor optical amplifier, PPLN = periodically polled Lithium Niobate.

discharging of a capacitance and is therefore independent of the clock frequency, but in digital optical devices, the energy per logic operation is determined by the (continuous) optical power required to activate an optical nonlinearity and the time taken to perform the logic operation. Therefore, in all-optical digital devices, in which all data and device control are carried out by optical signals, the nonlinear elements, amplifiers, and other active components require continuous supply power, regardless of whether that device is performing a digital logic operation or not. In addition, because $E_{in} \ll E_{device}$, the power consumption of nonlinear optical devices is approximately constant. This is unavoidable because of the need to provide ancillary active devices required for level restoration. Consequently, there is inverse relationship between the energy consumption per bit in optical devices and the clock frequency. We therefore model the device energy per clock period $E_{device,O}$ for digital optical devices as follows:

$$E_{device,O} = E_O \frac{f_{ref}}{\gamma\, f_O} \tag{2}$$

where E_O is the device energy when the effective clock frequency γf_O is equal to a reference clock frequency f_{ref}. Because interconnect losses are small in optical circuits, E_O is independent of the device pitch d.

2.4. Typical Device Energy Data

The calculations in this paper use energy data for typical optical and electronic devices. Fig. 4 shows some typical device energies as a function of time. The upper three curves (broken lines) in Fig. 4 provide estimates of E_O, $E_{(DE)MUX}$, and $E_{O/E/O}$. Some of the data for E_O in Fig. 4 were obtained from a recent study [11] of optical signal processing. In [11], $E_{device,O}$ was calculated for a variety of optical devices, using published experimental data, and (where possible) taking into account all the peripheral components, including OAs, optical pump circuits and drivers, etc. (see Fig. 3). In Fig. 4, we have included new data from some recent publications. As pointed out earlier, many papers do not provide details of all contribution to energy consumption. For example, while most papers document the signal levels at the inputs and outputs of the nonlinear element, many give no information on the supply energy for the input and/or output OAs in Fig. 3. In papers where energy data were missing, we made estimates of the total energy assuming OAs with power conversion efficiencies of 50%. Arguably, a conversion efficiency of 50% is somewhat optimistic, and therefore, the data should be viewed as being a lower limit on achievable energy.

The data presented in [11] are for $E_{device,O}$ and, therefore, do not take account of the clock frequency, which varies widely between different experiments published in the literature. To overcome this, we have renormalized the data presented in [11] from $E_{device,O}$ to E_O, using (2) with a reference clock frequency of $f_{ref} = 100$ GHz. This renormalization removes uncertainty about the influence of different clock frequencies. The broken line is a line of best fit to these data. The MUX and DEMUX data are taken from a survey of device capabilities presented in [15]. Data points are shown in Fig. 4 for an InP heterojunction bipolar transistor (HBT) DEMUX and a projection for future Silicon–Germanium (SiGe) high-speed MUX and DEMUX circuits [15]. The O/E/O data in Fig. 4 are based on estimates for long-reach WDM transmitters and receivers. Thus, these data are relevant to signal processing circuits used within long-reach transmission systems.

Some nonlinear devices reported in the literature use relatively long active regions in order to maximize the interaction between the optical signal field and the nonlinear medium. For example, recent experiments based on HNLF have used lengths of fiber as large as 1 km. Unfortunately, devices larger than a few centimeters or, at most, a few tens of centimeters, will have too large a footprint for practical applications and are also likely to suffer from difficulties with clock skew [13]. Therefore, the data in Fig. 4 do not include any devices that are larger than a few tens of centimeters.

The device switching energy of a simple CMOS inverter is shown in the lower part of Fig. 4. Decreasing device feature sizes as a function of time are shown. The smallest commercial device at the time of writing is 22 nm, but projected energies out to feature sizes of 11 nm are included in Fig. 4, based on data in the ITRS [9].

The lower curve in the bottom part of Fig. 4 is the energy consumed by a CMOS inverter single gate, excluding the interconnect wires. The upper curve in the bottom part of Fig. 4 is an estimate of the total energy, including the energy consumed by the interconnect wires in a typical CMOS IC [16]. To estimate this total energy including the interconnect wires, we considered a commercial 32-nm two-core processor and estimated that the average total energy per transistor (including interconnect wires) per transition is around 3 fJ [16]. From this, we estimate that the average length of and interconnect wire length in a 32-nm processor is $L_W \sim 3$ μm. The data in Fig. 4 for other feature sizes were obtained by scaling the average interconnect wire length to the transistor feature size and using other parameters from the ITRS [9]. The rate of improvement of total CMOS energy including wires in Fig. 4 is approximately 25% per annum or a factor-of-10 improvement over 10 years. Note that the rate of improvement of total CMOS energy including wires is less rapid than for an isolated CMOS gate.

The total CMOS energy $E_{device,E}$ including the interconnect wires is more than two orders of magnitude larger than the energy of an isolated CMOS gate, indicating that the consumption of each device is dominated by the energy in the interconnect wires. This difference in energies is an upper limit as it applies to a fairly complex processor. In simpler circuits, the average interconnect lengths may be shorter, and this energy difference may be smaller. Our analysis of the circuit in Fig. 1(b) uses the $E_{device,E}$ in Fig. 4. From Fig. 4, $E_{device,E}$ (including interconnect wires) is 3 fJ in 2010 and 0.3 fJ in 2020. This translates into the E_E and d_{ref} figures given in Table 1.

3. Analysis

In this section, we compare the energy consumption of the circuits in Fig. 1 using estimated device energies for the years 2010 and 2020. Fig. 5 and Table 1 summarize the data used in the analysis. The optical device energy E_O (using $f_{ref} = 100$ GHz) is 1 pJ in 2010 and 200 fJ in 2020. This assumes significant improvements in optical device technology over the next 10 years.

As pointed out earlier, we have ignored leakage current (i.e., static energy consumption) in CMOS devices. Therefore, E_{ES} in Table 1 is set to zero. This is justified by analyses of CMOS devices including leakage [14] that confirm that at high clock frequencies, the static energy consumption in CMOS is smaller than the dynamic energy consumption.

The data for $E_{device,E}$, $E_{(DE)MUX}$, and $E_{O/E/O}$ in Fig. 5 and Table 1 are taken directly from Fig. 4. Note that the O/E/O energies in Fig. 5 and Table 1 apply to long-reach optical transmitters and

TABLE 1
Device energies

		2010	2020
Electronic	E_{ED}	2 fJ	0.2 fJ
	d_{ref}	3 μm	1.5 μm
	E_{ES}	~ 0	~ 0
Optical	E_O	1 pJ	200 fJ
	f_{ref}	100 GHz	100 GHz
$E_{O/E/O}$		1 pJ	500 fJ
$E_{(DE)MUX}$		20 pJ	8 pJ

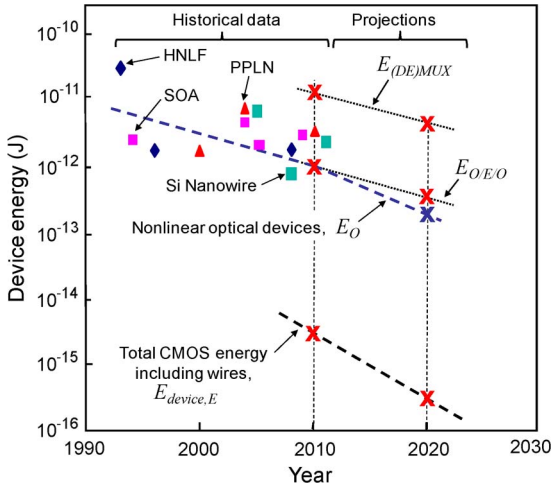

Fig. 5. Data points used in analysis, marked as "*X*".

receivers. These energies could be as much as an order of magnitude smaller if the transmitters and receivers are optimized for transmission over short distances [17].

3.1. Complexity of Signal Processing

One of the parameters in our analysis of energy consumption is the complexity of the signal processing carried out by the chip. We characterize the complexity of the signal processing in terms of the number N_{op} of digital operations performed on each bit of input data. The number of operations per bit of input data can vary widely, depending on the function being carried out by the circuit. For example, in an optical wavelength converter based on cross-phase modulation in SOAs [18], each incoming bit undergoes very little signal processing, and N_{op} would typically be on the order of one or two. For more complicated operations such as Internet protocol (IP) packet header recognition, more digital operations are needed, and N_{op} would be on the order of 100 or even larger for realistic size networks. In sophisticated signal processing circuits such as FEC [19], N_{op} would be more than 10^3 or 10^4.

In the optical circuit in Fig. 1(a), the clock period is $\tau_{clock,O} = 1/B$, and the number of operations per input bit achievable on a chip containing $N_{device,O}$ devices, each of which is capable of one operation per clock period $\tau_{clock,O}$, is given by

$$N_{op} = \gamma N_{device,O}/m \tag{3}$$

where γ is the activity factor. Similarly, in the electronic circuit in Fig. 1(b), the clock period is $\tau_{clock,E} = k/B$, and the number of operations per input bit on a chip containing $N_{device,E}$ devices,

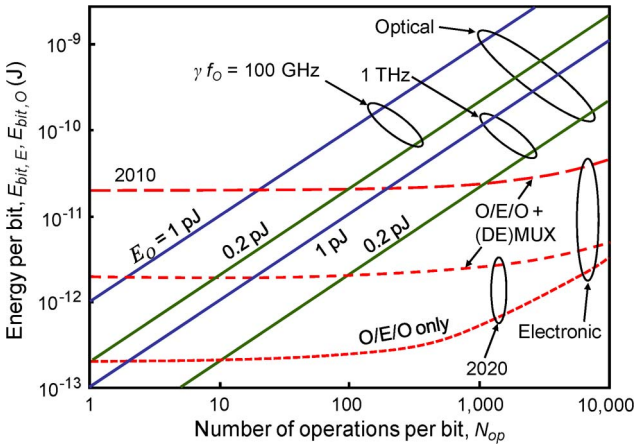

Fig. 6. Total energy per bit processed against number of operations per bit. The solid diagonal lines give $E_{bit,O}$ for optical devices with $\gamma f_O = 100$ GHz and 1 THz, as well as with $E_O = 1$ pJ and 0.2 pJ. The upper two broken lines give $E_{bit,E}$ in 2010 and 2020 for electronic circuits with DE(MUX) and O/E/O circuits at the inputs and outputs as shown in Fig. 1(b). The lowest broken curve gives $E_{bit,E}$ in 2020 for circuits with O/E/O converters only (i.e., without DE(MUX)s). This would apply in situations where the bit rate is low enough not to require the (DE)MUXs.

each of which is capable of one operation per clock period $\tau_{clock,E}$, is given by

$$N_{op} = \gamma N_{device,E}/mk. \tag{4}$$

3.2. Total Energy per Input Bit Processed

The total energy $E_{bit,O}$ per input bit consumed by the optical chip in Fig. 1(a) per input bit period or optical clock period $(1/B)$ is equal to the number of logic operations N_{op} per input but multiplied by the total energy $E_{device,O}$ per device per logic operation. Using (2), this gives

$$E_{bit,O} = N_{op}E_{device,O} = N_{op}E_O\frac{f_{ref}}{\gamma f_O}. \tag{5}$$

Similarly, for the electronic chip in the circuit in Fig. 1(b), the total energy consumed by the chip per input bit is $E_{bit,E} = N_{op}\, E_{device,E}$. Using (1) with $E_{ES} = 0$ and adding the energy consumed by the O/E/O converters and (DE)MUXs, the total energy per bit in the circuit in Fig. 1(b) is

$$E_{bit,E} = N_{op}E_{ED}\frac{d}{d_{ref}} + E_{O/E/O} + E_{(DE)MUX}. \tag{6}$$

Fig. 6 shows the total energy per bit $E_{bit,E}$ and $E_{bit,O}$ against the number of operations N_{op} per bit for optical and electronic circuits using the data in Table 1. The energy per bit in an optical circuit increases linearly with the number of digital operations per bit and decreases as the effective clock frequency γf_O increases and E_O decreases. In contrast, the energy per bit in an electronic circuit is largely independent of the number of operations per bit and is largely independent of the clock frequency. However, it is important to remember that the effective clock frequency of an electronic circuit is limited to around 100 GHz by the limited speed capabilities of the O/E/O converters and the (DE)MUXs. The energy per bit in the electronic circuit is dominated by the (DE)MUXs and, to a smaller extent, the O/E/O converters. The energy per bit in electronic circuits for 2010 includes both the O/E/O converters and the (DE)MUXs, but for 2020, we have drawn two curves—one with both the O/E/O converters and the (DE)MUXs and one with O/E/O converters only to allow for situations where the bit rate is low enough not to require (DE)MUXs. Fig. 6 shows that for circuits with low complexity (fewer than 10 operations per bit), optical and electronic circuits consume similar amounts of energy. However, for more than 100 operations per bit, electronic circuits are generally

more competitive from an energy consumption point of view. However, as stated earlier, optical circuits can, in principle, operate at higher speed.

3.3. Chip Power Density

In this section, we explore how the chip power density affects scaling properties of optical chips and compare the scaling properties of optical and electronic chips. The power density (i.e., the supply power per unit chip area) is a key parameter that limits the integration density and the clock speed in electronic ICs. In today's IC chips, thermal considerations limit the maximum allowable power density to less than about 100 W/cm^2 [9], [20]. The implications of power management on the scaling properties of electronic chips are reasonably well understood [20], but this topic has received little attention in the optics literature.

If the output signal energy from an electronic chip is small compared with the supply energy, all of the supply energy is converted to heat. In an optical circuit, some of the supply energy is converted to heat. However, in principle, some unused optical energy (e.g., unused spontaneous emission or unused nonlinear mixing products) could be removed from the chip, either using waveguides or by free space propagation, and dissipated off-chip. In the present analysis, we consider two power density limits for optical circuits: 100 W/cm^2—the same limit as for electronic circuits, and 1 kW/cm^2. This latter limit would apply for a circuit where 90% of the supply energy is dissipated off-chip in optical form.

In [11], it is shown that the input and output power to and from the chip at the signal input and output ports is small compared with the supply power. Therefore, these powers are ignored in the present analysis. In addition, we do not include the O/E and E/O converters and the (DE)MUXs in this part of the analysis.

The chip supply power for optical and electronic circuits (see Fig. 1) is given by

$$P_{supply,X} = B_{aggregate}\ E_{bit,X} \tag{7}$$

where, as before, $B_{aggregate} = mB$. The chip power density P_d for optical and the electronic chips is given by [13]

$$P_{d,X} = \frac{P_{supply,X}}{A} = \frac{mBE_{bit,X}}{d^2 N_{device}} = \frac{\gamma f_X E_{device,X}}{d^2} \tag{8}$$

where $P_{supply,X} = B E_{supply,X}$, and the term X in the subscripts in (7) and (8) can be either O or E.

From (8), the device energy for both optical and electronic devices can be written as

$$E_{device,X} = P_{d,X} \frac{d^2}{\gamma f_X}. \tag{9}$$

Fig. 7 shows $E_{device,X}$ from (9) plotted (broken lines) against the device pitch d, for $P_d = 100$ W/cm^2 and for $\gamma f_X = 1$ GHz, 100 GHz, and 10 THz. These broken lines in Fig. 7 give the upper bound on allowable device energy for a given γf_X and d. If the device energy for a particular device falls above one of these lines, then the chip power density will exceed 100 W/cm^2. Shown on the upper horizontal axis is the power consumed by each device, namely $d^2 P_d = \gamma f_x E_{device,X}$. Also plotted in Fig. 7 is E_{device} for an 11-nm CMOS device (i.e., $E_{device,E}$) from (1). At a device pitch of 10^{-7} m, a chip using these devices would be limited to an effective clock frequency γf_X of less than 1 GHz. However, at larger device pitches, the effective clock frequency γf_X can, in principle, increase to around 100 GHz, with a the effective clock frequency limited by the O/E/O converters and the (DE)MUXs.

For comparison with the 11-nm CMOS example, Fig. 7 shows vertical lines representing the constant power per device of optical devices with $E_O = 200$ fJ and 1 pJ. This indicates that the device pitch of optical devices is limited by thermal considerations to a minimum around 10^{-4} m. However, with a device pitch on this order, optical devices are capable, in principle, of operating at less than 100 W/cm^2 of dissipated power up to 10 THz and beyond.

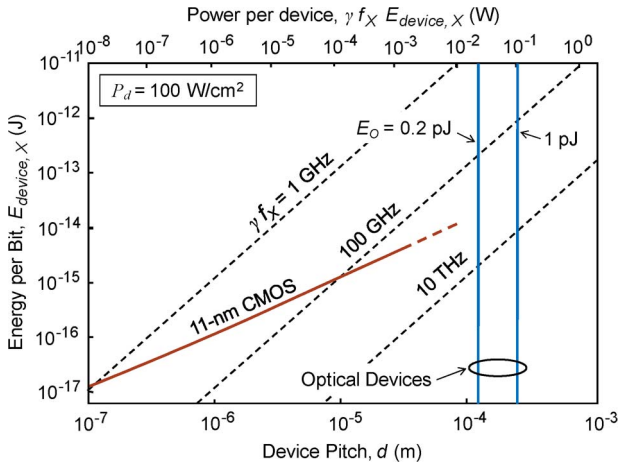

Fig. 7. Supply energy per bit against device pitch.

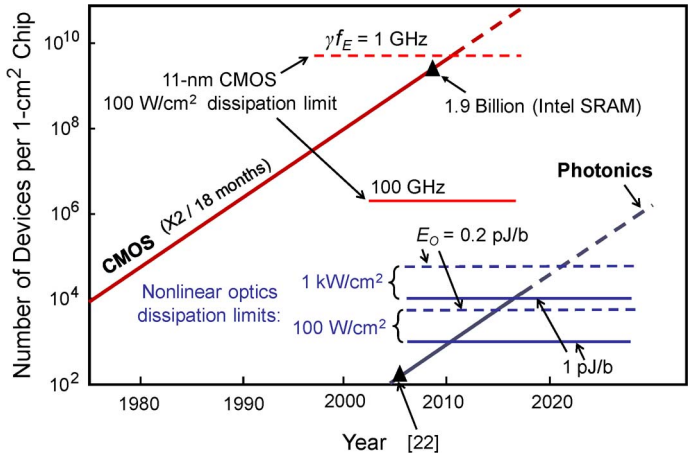

Fig. 8. Number of devices per 1-cm^2 chip.

3.4. Device Integration Density

As shown in Section 3.3 above, the number of devices that can be integrated on a chip is limited by dissipation on the chip. Fig. 8 provides an indication of historical trends in chip integration densities and relates this to energy limits for 11-nm CMOS and optical devices. The upper diagonal line traces improvements in the density of integration over time for CMOS ICs. For CMOS devices, the number of devices on a 1-cm^2 chip has grown at a rate of around a factor of two every 18 months, according to Moore's law [21]. There is no equivalent law for optical ICs; therefore, we have given the lower diagonal line in Fig. 8 the same slope as the CMOS curve. Trends to date indicate that it is unlikely optical ICs will match Moore's law, and we have therefore adopted this as a rather optimistic upper limit on the rate of improvement. This curve passes through a data point at an integration density of $\sim$100 cm^{-2} in 2005 [22]. Note that the circuits described in [22] are not specifically designed for digital signal processing, but [22] gives an indication of the practical level of optical device integration in the 2005 timeframe.

The upper two horizontal lines in Fig. 8 represent the dissipation-limited number of 11-nm CMOS devices per 1-cm^2 of chip area for circuits with effective clock frequencies of 1 GHz and 100 GHz. CMOS chips, with around 2 billion devices per 1-cm^2 of chip area, operate close to the thermal dissipation limit at 1 GHz. Dissipation limits for 11-nm CMOS restrict the device density to about 1.5 million per cm^2 if the clock frequency is increased to 100 GHz. The lower four horizontal lines

in Fig. 8 represent the dissipation-limited number of optical devices for the two optical device energies E_O used in Fig. 7 and for power density limits of 100 W/cm^2 and 1 kW/cm^2. While we recognize that there is no photonic equivalent to Moore's law, Fig. 8 suggests that there is some potential for future increases in the integration density of photonic ICs. However, it appears likely that photonic integration densities will be restricted by thermal limits to less than 10^4 cm^{-2}, i.e., around six orders of magnitude lower than the integration density of CMOS circuits.

4. Conclusion

Energy consumption is a key consideration in the development components, devices, and circuits for digital optical signal processing. We have shown that, from an energy consumption point of view, digital optical signal processing is potentially competitive with electronic signal processing if the signal to be processed is in optical form and if the signal processing function is simple—i.e., when there are only a few digital operations performed on each bit of data. For circuits requiring more operations on each bit of data, electronic signal processing uses less energy, even if the data to be processed is in optical form.

It is often argued that optical signal processing will replace digital signal processing because of its high-speed capabilities. However, simplistic arguments based on speed alone often miss the critically important point that digital optical devices are generally very energy hungry. Unless there are many-orders-of-magnitude improvement in the energy efficiency of digital optical devices, electronics is likely to remain the technology of choice for the vast majority of digital signal processing and switching applications in telecommunications networks.

There is a large disparity between the research literature on CMOS signal processing and the research literature on optical signal processing. While energy considerations are primary drivers in CMOS R&D, energy considerations receive very little attention in the optical signal processing literature. Even in papers where energy is mentioned, full details are often missing, and incorrect interpretations of energy consumption in optical devices are not uncommon. We believe that energy consumption is a potentially significant barrier to the commercial exploitation of digital optical circuits. More attention needs to be paid to energy consumption issues in the research and development of new digital optical technologies.

References

[1] K. E. Stubkjaer, "Semiconductor optical amplifier-based all-optical gates for high-speed optical processing," *J. Sel. Topics Quantum Electron.*, vol. 6, no. 6, pp. 1428–1435, Nov./Dec. 2000.
[2] G. S. Kanter, P. Kumar, K. R. Parameswaran, and M. M. Fejer, "Wavelength-selective pulsed all-optical switching based on cascaded second-order nonlinearity in a periodically poled lithium-niobate waveguide," *IEEE Photon. Technol. Lett.*, vol. 13, no. 4, pp. 341–343, Apr. 2001.
[3] M. Asobe, T. Kananori, and K. Kubodera, "Applications of highly nonlinear chalcogenide glass fibers in ultrafast all-optical switches," *IEEE J. Quantum Electron.*, vol. 29, no. 8, pp. 2325–2333, Aug. 1993.
[4] S. Gao, X. Zhang, Z. Li, and S. He, "Polarization-independent wavelength conversion using an angled-polarization pump in a silicon nanowire waveguide," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 250–256, Jan./Feb. 2010.
[5] M. Vasilyev, Y. Su, and C. McKinstrie, "Special issue on nonlinear-optical signal processing?" *IEEE J. Sel. Topics Quantum Electron.*, vol. 14, no. 3, pp. 527–528, May/Jun. 2008.
[6] D. A. B. Miller, "Are optical transistors the logical next step?" *Nature Photon.*, vol. 4, no. 1, pp. 3–5, Jan. 2010.
[7] R. S. Tucker, K. Hinton, and G. Raskutti, "Energy consumption limits in high-speed optical and electronic signal processing," *Electron. Lett.*, vol. 43, no. 17, pp. 906–908, Aug. 2007.
[8] R. S. Tucker, "Optics versus electronics for high-speed switching and signal processing," in *Proc. Photon. Soc. Summer Top. Meeting Playa del Carmen*, 2010, pp. 111–112.
[9] *International Technology Roadmap for Semiconductors,* 2010, Edition (2005). [Online]. Available: http://public.itrs.net/2005
[10] D. Miller, "Optics for low-energy communications inside digital processors: Quantum detectors, sources, and modulators as efficient impedance converters," *Opt. Lett.*, vol. 14, no. 2, pp. 146–148, Jan. 1989.
[11] K. Hinton, G. Raskutti, P. M. Farrell, and R. S. Tucker, "Switching energy and device size limits on digital photonic signal processing technologies," *J. Sel. Topics Quantum Electron.*, vol. 14, no. 3, pp. 938–945, May/Jun. 2008.
[12] R. S. Tucker, G. Eisenstein, and S. K. Korotky, "Optical time-division multiplexing for very high bit-rate transmission," *J. Lightw. Technol.*, vol. 6, no. 11, pp. 1737–1749, Nov. 1988.
[13] R. S. Tucker, "Energy consumption in digital optical ICs with plasmon waveguide interconnects," *IEEE Photon. Technol. Lett.*, vol. 19, no. 24, pp. 2036–2038, Dec. 2007.

[14] D. J. Frank, W. Haensch, G. Shahidi, and O. H. Dokumaci, "Optimizing CMOS technology for maximum performance," *IBM J. Res. Develop.*, vol. 50, no. 4.5, pp. 419–431, Jul. 2006.
[15] R. S. Tucker, "Green optical communications—Part I: Energy limitations in transport," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 245–260, Mar./Apr. 2011.
[16] R. S. Tucker, "Green optical communications—Part II: Energy limitations in networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 261–274, Mar./Apr. 2011.
[17] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, Jul. 2009.
[18] M. L. Masanovic, V. Lal, J. A. Summers, J. S. Barton, E. J. Skogen, L. G. Rau, L. A. Coldren, and D. J. Blumenthal, "Widely tunable monolithically integrated all-optical wavelength converters in InP," *J. Lightw. Technol.*, vol. 23, no. 3, pp. 1350–1362, Mar. 2005.
[19] I. B. Djordjevic, M. Arabaci, and L. L. Minkov, "Next generation FEC for high-capacity communication in optical transport networks," *J. Lightw. Technol.*, vol. 27, no. 16, pp. 3518–3530, Aug. 2009.
[20] S.-C. Lin and K. Banerjee, "Cool chips: Opportunities and implications for power and thermal management," *IEEE Trans. Electron Devices*, vol. 55, no. 1, pp. 245–255, Jan. 2008.
[21] L. Chang, Y.-K. Choi, J. Kedzierski, N. Lindert, P. Xuan, J. Bokor, C. Hu, and T.-J. King, "Moore's law lives on," *IEEE Circuits Devices Mag.*, vol. 19, no. 1, pp. 35–42, Jan. 2003.
[22] R. Nagarajan, M. Kato, J. Pleumeekers, P. Evans, S. Corzine, S. Hurtt, A. Dentai, S. Murthy, M. Missey, R. Muthiah, R. A. Salvatore, C. Joyner, J. R. Schneider, M. Ziari, F. Kish, and D. Welch, "InP photonic integrated circuits," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 5, pp. 1113–1125, Sep./Oct. 2010.