

A 0.38 pJ/bit 1.24 nW Chip-to-Chip Serial Link for Ultra-Low Power Systems

Christopher J. Lukas and Benton H. Calhoun

Department of Electrical Engineering
University of Virginia
Charlottesville, USA
{cj14hd, bcalhoun} @virginia.edu

Abstract— As energy-constrained systems continue to reduce their power consumption, finding an optimal point of operation for the principle components in the energy budget becomes increasingly important. With energy dominant system components like communication circuits, it is important to consider both energy-per-bit and power in the context of the system's use cases. In this paper, we propose optimization of chip-to-chip links considering both energy-per-cycle and energy-per-bit to find the optimal operating voltage and activity factor while minimizing wasted energy and power. A fabricated 130 nm chip was used to verify this finding and resulted in an energy-per-bit of 0.38 pJ/bit and power of 1.24 nW.

Keywords—activity factor; chip-to-chip serial link; energy-per-cycle; energy-per-bit; Internet of Things; low throughput I/O; SPI; Ultra-low power

I. INTRODUCTION

Energy optimization in circuits has become an important field of research due to the demand of ultra-low power (ULP) systems that contribute to the internet of things (IoT). These circuits generally do not have the high throughput requirements of high-power, high-performance chips [1] and must be optimized using different techniques. These emerging technologies have a need for optimally energy efficient I/O for communication with off-chip nodes such as sensing devices, radios, and body area networks. There remains a need to address optimization of low throughput communication.

Current state-of-the-art energy efficient communication circuits [11][12] use the energy-per-bit metric as a key comparison point. However, optimizing energy-per-bit alone is a problem for ULP systems if the solution is increasing frequency and activity factor (probability of switching in a given clock cycle). Because IoT systems have both power and energy constraints, these changes can result in consumption levels above the power budget, making it an incomplete solution. In addition to this, circuits optimized for high frequency operation are designed with the assumption that a large amount of data is waiting to be transmitted. This is often not the case in the low-energy, low-performance domain.

Because of the need to amortize static and leakage energy, current research for chip-to-chip I/O is pushing into higher frequencies to reduce energy-per-bit. Fig. 1 shows that the lowest E/b designs use very high frequencies. Solutions for

lower data rate communications are lacking. For example, the I²C standard (up to 5 MHz) assumes a termination resistor of approximately 5k Ω that consumes significant static power which is amortized by increasing frequency. SPI has a similar typical max frequency but avoids the static current of I²C. Both standards assume a load capacitance of 5pF. Overall, the figure shows the need for lower energy I/O circuits at lower frequencies.

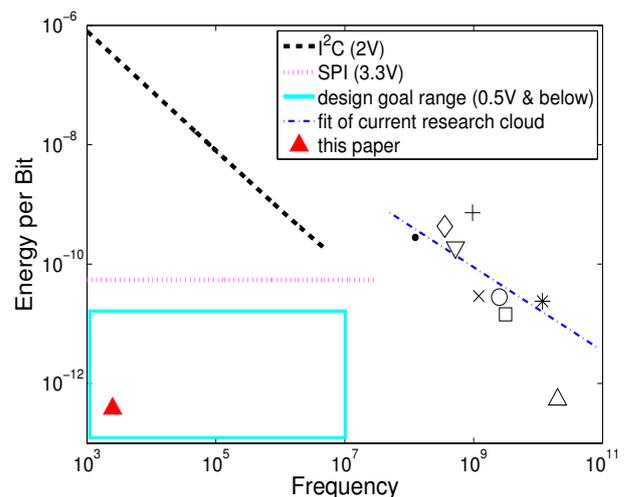


Fig. 1. De facto standards and state of the art communication circuits [1-9].

Energy constrained digital circuits minimize the energy-per-operation metric by determining the optimal operating voltage [13]. However, in transmission circuits where there is added control over the activity factor, and in turn, throughput of the circuit, determining the optimal operating voltage using this method is no longer enough. Activity factor provides a new knob that allows further reduction in energy-per-cycle.

Considering both metrics will allow us to create a communication circuit more suitable for ULP systems. This new design comes at the cost of lower throughput. However, IoT or body sensor nodes often will be transmitting data with a low throughput requirement, so reducing power at the cost of throughput is an acceptable tradeoff.

Section II of this paper includes an explanation of energy-per-bit and energy-per-cycle and describes the importance of

each and explains the method used to find a point of operation where the benefits of optimizing both energy-per-bit and energy-per-cycle are captured. Section III will cover an example circuit that was fabricated for low-energy chip-to-chip communication, showing the quantifiable benefits of this optimization. Section IV shows the results of the fabricated chip and its comparison to existing state of the art. Finally, Section V concludes the paper.

II. ENERGY-PER-BIT AND ENERGY-PER-CYCLE

Energy-per-cycle is defined as the average energy consumed by a circuit during one clock cycle. This energy includes both sleeping and active cycles. Energy-per-bit (1) differs from energy-per-cycle (2) in that it has an opposite relationship with activity factor (α). Lowering activity factor in (2) results in a lower energy-per-cycle and therefore a lower average power at a cost of lower throughput. The same change in activity factor results in an increased energy-per-bit due to more time spent leaking for each bit transmitted. Equation (1) assumes transmission of pseudorandom data across the transmission line.

$$E_{bit} = \frac{CV_{DD}^2}{\frac{1}{\alpha} - \frac{2}{1-2\alpha}} + \frac{W_{eff}L_{DP}KC_gV_{DD}^2e^{-\frac{V_{DD}}{nV_{th}}}}{2\alpha} \quad (1)$$

$$E_{cycle} = \alpha E_{dyn} + E_{leak} = \alpha CV_{DD}^2 + I_{leak}V_{DD}t \quad (2)$$

In ULP sensor nodes, off-chip communication (wireless or wired) has often been the largest energy consumer when operating at the node's maximum data rate [10]. By reducing the activity factor, the energy consumption for a given cycle of these components will fall. This knob allows for optimization that is unique to communication circuits. In addition, by varying the operating voltage of a circuit, it is possible to calculate the minimum energy-per-cycle (3) for a given activity factor in order to find an optimal operating voltage (4) [13][14]. [13] assumes that the circuit is operating at its maximum frequency and calls the metric minimum energy-per-operation. Equation (5) takes into account that the throughput given by the maximum possible frequency for a voltage of operation may not be required. The optimal voltage (4) is a function of β , where n is the subthreshold slope and V_{th} is the thermal voltage. The Lambert W function is defined as $x = W(x)e^{W(x)}$ and for our purposes is constrained to the branch $W_{-1}(x)$. In (5), C is the total switching capacitance of the circuit, W_{eff} is the average transistor width relative to the characteristic inverter, L_{DP} is the logic depth, K is the delay fitting parameter, and C_g is the load capacitance. As α decreases, β approaches zero and V_{DDopt} increases, as shown in Fig. 2. This becomes a special case where increasing voltage will decrease the total E/b due to decreased active energy and less leakage over shorter cycle times.

$$E_{minat} \frac{\partial(\alpha CV_{DD}^2 + I_{leak}V_{DD}t)}{\partial V} = 0 \quad (3)$$

$$V_{DDopt} = nV_{th}(2 - \text{lambert}W(\beta)) \quad (4)$$

$$\beta = \frac{-2\alpha C}{W_{eff}L_{DP}KC_g}e^2 \geq \frac{-1}{e} \quad (5)$$

The minimum energy-per-cycle can be found for many activity factors to find a trend of minimum energy points. The final curve is the minimum energy-per-cycle versus as a function of activity factor. Solving analytically results in (6) [13].

$$E_{Copt} = [nV_{th}(2 - \text{lambert}W(\beta))]^2 * [\alpha C + W_{eff}L_{DP}KC_g e^{-(2 - \text{lambert}W(\beta))}] \quad (6)$$

If a power budget is defined for a communication circuit, this optimization results in a maximum activity factor and voltage that can be used to operate within a given system's budget. Both metrics are important to minimize in the design of an ultra-low power transmission line circuit and can be used together to find an optimal point of operation.

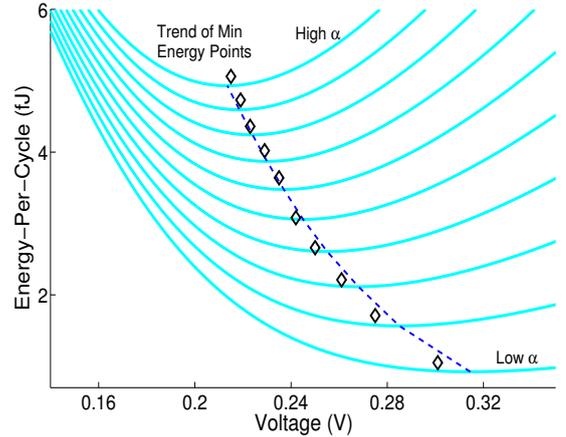


Fig. 2. Minimum energy points for a range of activity factors. Each activity factor has an associated optimal operating voltage, creating a new trend of minimum energy-per-cycle (6). Black diamonds show simulation values, while the dashed line shows the model.

The maximum activity factor in Fig. 2 corresponds to the minimum energy-per-operation from [13]. However, the minimum energy point in [13] may result in a higher throughput than needed for a given application. This energy can be further reduced through lowering activity factor. This produces a new curve with a new minimum energy point. As fewer cycles are spent transmitting data, the leakage energy dominates the curve and pushes the optimal point to the right, increasing its voltage.

III. MEASURED RESULTS

Using equations (1) and (6) we designed a circuit that would minimize energy-per-cycle while still retaining an acceptable throughput. We designed the circuit to have as low leakage as possible, reducing the minimum energy point.

Because of the low leakage, the minimum energy point at high α is below the circuit's minimum operating voltage. This means we want to reduce the voltage until it matches our throughput requirements.

We fabricated a single-ended SPI-compatible chip-to-chip link in a commercial 130 nm process with one pin for clock and one for data. While a differential transmission is preferred in higher throughput circuits, it requires a minimum of 2x energy-per-bit transition on the transmission side due to having two drivers and transmission lines, as well as a differential amplifier on the receiving end, which generally requires a constant current. This is not ideal for our design space, as we will not be amortizing this energy with higher throughput.

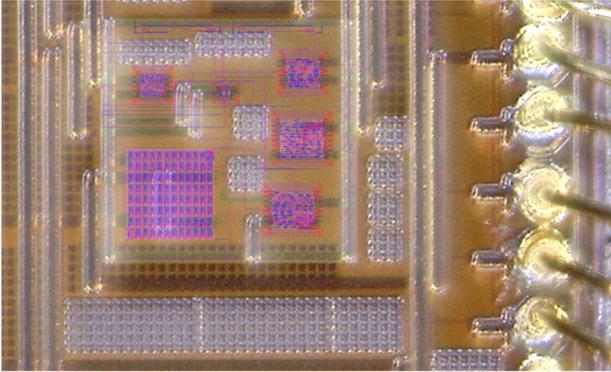


Fig. 3. Chip micrograph of the fabricated chip. Because of low visibility due to the metal fill layers above the circuit, the layout has been overlaid.

The transmission line is not terminated, as the circuit operates at a low enough frequency where termination is not required. This is a benefit of low throughput communication, as a termination resistor will result in a non-amortizable constant current. Higher voltages will result in faster transition slopes, creating more noise in the transmission line. This design decision limits the maximum voltage of the circuit. Impedance matching isn't necessary, as we are not worried about reflections at low frequencies or maximum power transfer. On the contrary, the highest power efficiency is desirable, so a low source resistance is chosen, but not so low as to be high leakage. Fig. 3 shows a chip micrograph of the fabricated circuit.

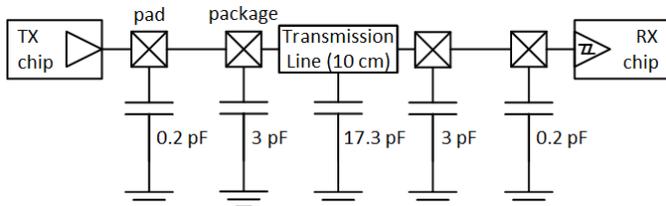


Fig. 4. Model of transmission line. Series inductors and resistors are ignored in low frequency operation. The clock and data lines were both designed to have these values to match delay.

The fabricated chip's block diagram (Fig. 4) shows the major sources of capacitance, which decrease the energy efficiency of the circuit. We chose to use a 10 cm trace between the transmitter and receiver to get a pessimistic capacitance on the PCB transmission line. In a non-test environment, the transmission line capacitance can be greatly

reduced by both optimizing the distance between the transmitter and receiver, and moving from a PGA package to one with lower capacitances.

A reference point of operation will be used for a comparison to the methodology used in this paper. A typical operating voltage for ULP body sensor nodes is assumed (500 mV) [10]. The accepted methodology is to use the maximum working frequency for the circuit at its operating voltage to reduce energy per bit by amortizing leakage, so a frequency of 2.1 MHz is used as a reference. Because a bit is sent every cycle, the energy-per-bit will equal the energy-per-cycle. This operating point was found to have energy-per-bit and energy-per-cycle value of 2.60 pJ for the fabricated circuit while consuming 5.46 μ W of power.

When using our methodology, we found the minimum operating voltage to be 200 mV (3.25 kbps), below which the circuit became unreliable. In a situation where the throughput is not high enough and optimal energy-per-bit is preferred, the voltage can be raised until the requirement is satisfied. If optimal energy-per-cycle is preferred due to lower throughput requirements, the minimum energy-per-cycle curve can be followed by lowering activity factor only enough to satisfy the throughput requirements. A comparison of the above solutions is compared to the reference point of operation in Table 1.

TABLE I. A COMPARISON OF A TYPICAL SPI TRANSMISSION VERSES THE OPTIMIZED TRANSMISSION OF OUR FABRICATED CIRCUIT.

| | Reference | Optimized E/b |
|----------------|--------------|---------------|
| Voltage | 500 mV | 200 mV |
| Energy-Per-Bit | 2.60 pJ | 0.38 pJ |
| Power | 5.46 μ W | 1.24 nW |

IV. RESULTS

Using the optimal point of operation from earlier in the paper along with the ULP circuit design strategies, we were able to create a circuit that operated with 30% lower energy-per-bit than the state of the art but with almost seven orders of magnitude less power than that same design. The flexibility of the circuit allows it to operate into the above threshold range to \sim 700mV before reflections in the transmission line affect the reliability of the circuit.

The black filled square (leftmost) in Fig. 5 shows the point of minimum energy per bit. The curves show the flexibility of the fabricated chip. The black line shows increasing V_{DD} while retaining maximum frequency to increase throughput at the cost of energy per bit and power. The dashed blue lines show the use of the activity factor knob introduced in Section II to reduce power at the cost of energy per bit. The green lines show varying voltage and activity factor to retain constant throughput.

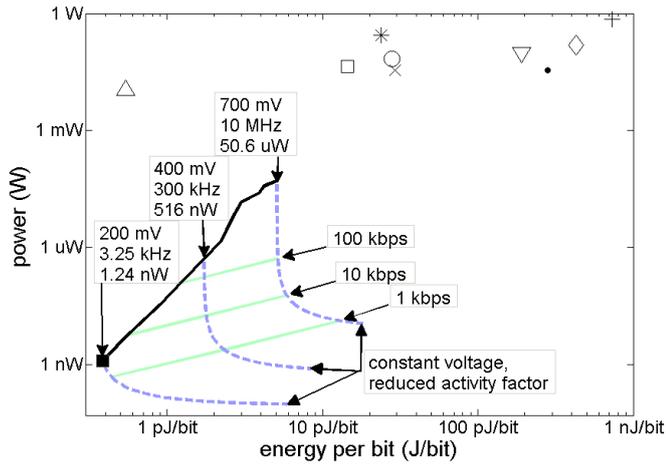


Fig. 5. A comparison of power and energy per bit with current state of the art.

Ideally a circuit will be at the bottom left of Fig. 5, having both a power consumption and energy per bit as close to zero as possible. Pushing to the lower left of the figure is limited by the physical constraints of the circuit, mainly its ability to reduce its operating voltage. However, operating this circuit at 126 mV is possible and results in reduced power consumption (~800 pW) at the cost of severely degraded throughput (60 bps). Once the minimum operating voltage is found it can be used to find the maximum throughput. The maximum throughput gives the minimum energy-per-bit due to amortization of leakage energy. If this throughput is above the system's requirement, the best course of action is to travel down the blue line in the figure by reducing activity factor until the throughput matches the budget. If the throughput is below the system's requirement the voltage can be increased by traveling up the black line, increasing the maximum frequency of the circuit until the necessary throughput is met.

Fig. 6 shows an eye diagram for the transmission line operating at its lowest energy-per-bit at around 200 mV peak-to-peak. Even at the low voltage, the receiver can clearly distinguish between high and low bits. The driver operates at 200 mV whereas the supporting logic operates at 400 mV.

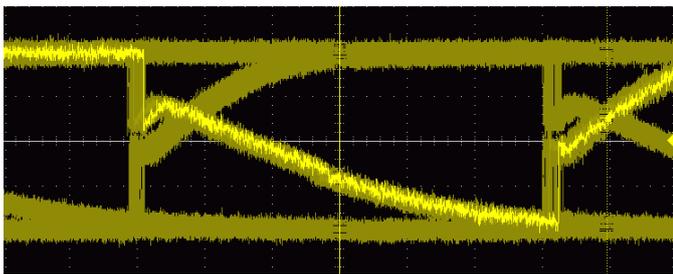


Fig. 6. Eye diagram of the transmission line operating at 200mV peak-to-peak.

Using this design in an ultra-low power SoC will reduce energy-per-bit spent communicating with other devices by over two orders of magnitude when compared to existing SPI (Fig. 1). Because this saved energy is often a large part of the system's budget, it can be used to implement other logic blocks

and stay within the same energy budget. Future work involves using the throughput of the serial link as a driving force in the design of surrounding circuits.

V. CONCLUSION

This paper describes a method for finding the minimum energy-per-cycle curve by varying activity factor, and then using this as an energy budgeting technique for designing a communication circuit with energy constraints. We then went on to relate this with the energy-per-bit metric and create a trade-off between the two in which an optimal point is found, resulting in an operating voltage and activity factor. A fabricated circuit was then used to show the benefits of using this optimization, having an energy-per-bit of 0.38 pJ/bit and power of 1.24 nW. This methodology will allow designers to have a more accurate energy model that can be used for designing communications circuits more suited for body sensor nodes and other ultra-low power systems.

VI. ACKNOWLEDGEMENTS

This work was funded in part by NSF (1035771) and the NSF NERC ASSIST Center (EEC-1160483). The authors thank Ben Boudaoud and John Poulton for their assistance.

VII. REFERENCES

- [1] Poulton, J.W. et al., "A 0.54pJ/b 20Gb/s ground-referenced single-ended short-haul serial link in 28nm CMOS for advanced packaging applications," ISSCC pp.404,405, 2013.
- [2] Rylov, S.; et al. "10+ Gb/s 90nm CMOS serial link demo in CBGA package," CICC pp.27,30, 2004.
- [3] Chen, L. et al., "A 90nm 1-4.25-Gb/s Multi Data Rate Receiver for High Speed Serial Links," ASSCC pp.391,394, 2006.
- [4] Boni, A., "1.2-Gb/s true PECL 100K compatible I/O interface in 0.35- μ m CMOS," JSSC pp.979,987, Jun 2001
- [5] Chih-Chien Hung et al., "A Sub-1V CMOS 2.5Gb/s Serial Link Transceiver Using 2X Oversampling," ASSCC pp.37,40, 2005.
- [6] Sungjoon Kim et al., "A 960-Mb/s/pin interface for skew-tolerant bus using low jitter PLL," JSSC pp.691,700, May 1997
- [7] Bin Guo et al., "A 125 Mbs CMOS all-digital data transceiver using synchronous uniform sampling," ISSCC pp.112,113, 1994.
- [8] Tamura, M. et al., "A 1V 357Mb/s-throughput transferjet™ SoC with embedded transceiver and digital baseband in 90nm CMOS," ISSCC pp.440,442, 2012.
- [9] Miyashita, D. et al., "A -70dBm-sensitivity 522Mbps 0.19nJ/bit-TX 0.43nJ/bit-RX transceiver for TransferJet™ SoC in 65nm CMOS," VLSIC pp.74,75, 2012.
- [10] Yanqing Zhang et al., "A Batteryless 19 uW MICS/ISM-Band Energy Harvesting Body Sensor Node SoC for ExG Applications," JSSC pp.199,213, Jan. 2013.
- [11] Yao-Hong Liu et al., "A 1.9nJ/b 2.4GHz multistandard (Bluetooth Low Energy/Zigbee/IEEE802.15.6) transceiver for personal/body-area networks," ISSCC pp.446,447 2013.
- [12] Roberts, N.E.; Wentzloff, D.D., "915MHz ultra low power receiver using sub-Vt active rectifiers," SubVT pp.1,3 2012.
- [13] Calhoun, B.H.; Wang, A.; Chandrakasan, A., "Modeling and sizing for minimum energy operation in subthreshold circuits," JSSC pp.1778,1786, Sept. 2005.
- [14] Markovic, D. et al., "Methods for true energy-performance optimization," JSSC pp.1282,1293, Aug. 2004.