Strategies for Reducing Power Consumption and Increasing Reliability in IoT

Ricardo Reis^(✉)

Instituto de Informática – Universidade Federal do Rio Grande do Sul (UFRGS),
Caixa Postal 15.064, 91501-970 Porto Alegre, RS, Brazil
reis@inf.ufrgs.br

Abstract. The Internet of Things (IoT) demands new challenges in the design of computing and electronics components. One of the challenges is the power reduction of this expanding network of connected devices, where the majority is permanently connected. In a large set of applications, another significant issue is reliability, especially on critical areas as health and transport. This paper shows an overview of design strategies that we have developed to reduce power consumption and to increase reliability in circuits that are components of the IoT, as the reduction of the number of transistors in IoT devices, using optimisation techniques and the physical design of circuits tolerant to radiation effects.

Keywords: Internet-of-Things · Optimization · Physical design · Fault tolerance · Radiation effects · Nanoelectronics

1 Introduction

The growing number of connected devices in the Internet of Things (IoT) is one of the reasons for the ever increasing increase in the number of transistors produced annually in the world. Figure 1, based on (SIA 2005), shows the number of transistors manufactured annually in the world, year by year. This impressive growth is due to 3 main factors: the increasing number of transistors integrated into a chip, the growing number of products that include embedded chips and the increasing number of manufactured copies of each product. The manufacturing cost of a transistor is relatively cheap. In (The Economist 2010) a comparison is presented between the cost of a grain of rice and the cost of a transistor. The price of a rice grain can be equivalent to the manufacturing cost of more than 125,000 transistors. This would indicate that there is no need to optimise the number of transistors in a design, since the cost of them is relatively small. But the cost of energy required for the operation of a transistor is increasing a lot. We also have to consider that a high-power consumption can reduce the lifetime of a system, as well as increase the effects of variability that can cause an integrated system to malfunction and/or also reduce its useful life. With the increasing connection of electronic and computational devices on the Internet, that is, in the Internet of Things, power consumption problems tend to get worse, and a lot. How much Power Plants we will need to cope with the IoT/IoE (Internet of Everything) world? This is a major issue.

So, an essential keyword on the Internet of Things is **optimisation**, especially the optimisation of power consumption, which must be addressed at all levels of abstraction in the design flow of a computer or electronic system. The total power optimisation is a summation of the optimisation done at each level of design abstraction. So, sustainable computing requires optimisation at all design levels of a computer or electronic system design.

Fig. 1. Number of transistors produced annually in the world (adapted from SIA 2005)

2 Internet of Things

The term Internet of Things has already given rise to several other associated terms, such as the Internet of Health (IoH), Internet of People (IoP) and the Internet of Everything (IoE). In fact, the latter term becomes the most comprehensive, but each one of the others has some specific characteristics. When talking about the Internet of Health, which includes real-time monitoring of a person's clinical conditions, as well as chips injected in a person, the issue of reliability is a key one. And reliability is also related to power consumption in most cases. High power consumption can reduce the lifetime of a system. When it comes to the Internet of People, the issue of people's security and privacy is of great relevance. But in all cases, the importance of optimising energy consumption is growing more and more.

When considering optimisation, it means that integrated systems must increasingly be dedicated to the intended application to optimise the number of components, that means the number of transistors. Another important strategy for optimisation is the hardware and software codesign, where one can manage the compromise between performance, consumption, and reliability.

Devices connected to the Internet of Things (or the Internet of Everything), can have very different complexities. If it is analysed the complexity considering the number of components, we can find small devices with few transistors and large devices with billions of transistors. Of course, large devices will consume much more power, but we have to consider that most devices on the Internet of Things are devices with a low number of transistors. But, because they are found in large quantities, they can represent a total consumption more important than the consumption of the so-called large devices that are present in a lower number. Therefore, consumption optimisation must be performed on both large and small devices that are present in large quantities. Another aspect to consider is that some devices require the application of reliability techniques (such as those related to transport or health systems), which can increase the number of components, while other devices are not critical, such as a camera or video, where an error in viewing a pixel of an image does not cause significant problems.

Also, we can expect that many systems connected to the Internet of Everything will be Cyber Physical Systems (CPS), that are systems composed by different classes of components like electronic elements, mechanical elements, optical elements, physical sensors, chemical sensors, organic components, and many others. So, it is needed to obtain EDA tools to cope with the design of CPS composed of all these classes of devices.

Figure 2 (The Connectivist 2014) shows an estimate of the number of devices connected to the Internet since 1992 when they were about 1 million devices. By 2020 when it is estimated that there will be more than 50 billion devices connected in the network, and there are currently around 35 billion connected devices. In (IHS Markit 2018) the number of devices connected to the network in 2018 is shown by industrial and commercial sectors, where almost half is in the area of communication. The significant growth in the number of connected devices to the Internet has naturally led to a considerable increase in the energy consumed in the Internet of Things. For how long will we have the energy to meet this growing demand? Therefore, it is necessary to use techniques to minimise the energy consumption of each connected device in the Internet of Things.

The Internet of Health (IoH) is a significant way to increase the life of human beings but also to improve life quality. Some of the examples of devices to be connected to the IoH are: Glasses that can advise eye correction; Toothbrush that can find cavities and breath issues; Razor that identify acne; Pacemakers that broadcast data to cardiologist; Underwearables that can provide early detection of cancer and other anomalies; Combs that can scan for fungus and hair loss; Earphones that does measurement of hearing, analysis of emotional level; Watches able to measure parameters like blood pressure, heart rate and others.

In critical areas such as the design of implanted devices (chips) in humans (Fig. 3), the reliability of the implanted systems is obviously critical. Some of the techniques used are based on the triplication of circuits and the temporal analysis of the propagation of a signal. Previously, the design of fault-tolerant circuits, to cope with radiation effects, was mainly in circuits that were sent to space. With the reduction of the value of the supply voltage of integrated circuits, nowadays the integrated circuits for use at ground level are also sensitive to errors caused by the radiation incident on the earth. Therefore, in critical areas such as implanted chips in humans, it is necessary to implement radiation effects tolerance techniques (Velazco et al. 2007). Also, critical

Fig. 2. Number of devices connected on the Internet (adapted from The Connectivist 2014)

systems used on the Internet of Health should be tolerant to any kind of noise (internal or external to the human body). They also must have a larger lifetime as possible, for obvious reasons and also should cope with environmental variability.

Fig. 3. The implantation of Chip Systems in humans demands reliability and ultra-low consumption

Also, there is the effect of “ageing”, that is, the ageing of the circuit, which is more eminent in nanometric technologies (Vasquez et al. 2012). One of the most important effects is known as NBTI (Negative Bias Temperature Instability) that alters the threshold voltage of the PMOS transistors, degrading the operation of the transistor. Another effect that causes failures in circuits throughout their life is the effect of electromigration, which can cause short circuits or rupture of connections (Fig. 4). In order to increase the lifetime of the chips, it is necessary to use physical design techniques that reduce the probability of electromigration (Posser et al. 2016, 2017).

Fig. 4. Exemplo of a Void (open circuit) and hillock (short circuit) (Geden 2011).

3 Electronic Design Automation (EDA) Tools

The use of EDA tools is essential for optimising energy consumption and increasing reliability, as the design flow has a large set of steps as well the number of components of a chip can reach billions of transistors. In Fig. 5 we can see the floorplan of an integrated circuit, where the hotter colours show regions (hot spots) with higher energy consumption, indicating that in some points there is a significant concentration of power consumption. One way to deal with the problem is to modify the placement of the logic cells in the circuit to distribute the cells with the highest energy consumption over the entire circuit area. But this must be done without compromising the area, wirelength and operating frequency specifications (much depends on the routing). Another way is to decrease the number of transistors, since the static consumption is related to the number of transistors (Reis 2011A).

Fig. 5. View of the consumption density on a chip

4 Power Consumption Reduction by Reducing the Number of Transistors

The reduction of the power consumption of a System on a Chip (SoC) is a function of a sum of techniques and strategies of design applied in different levels of abstraction in the design flow of an integrated system (Reis 2010). The summation of the gains is that it will set the total gain in power reduction. When we deal with the physical synthesis of a system on a chip, one technique is the optimisation of the number of components, that is, the number of transistors. In Fig. 6 (Reis 2011A) we can observe two solutions for the implementation of the same equation. The first solution makes use of 4 basic logic gates (3 NOR 2-input ports and one CMOS inverter), using a total of 14 transistors. The second solution makes use of only one logic gate, which performs the same function but with only 8 transistors. That is, the second solution, having a reduction in the number of transistors, will also have a proportionally smaller static power consumption. Furthermore, in the example of Fig. 6, we can see that the first solution also has 3 connections between the basic gates (and therefore even vias and contacts) that are eliminated in the second option with only one logic gate.

This elimination of connections is increasingly important because it decreases the number of connections to be implemented using the different metal layers. The decrease in the number of connections decreases the density of connections and, therefore, increases the routability of the circuit and also contributes to reduce the

average length of the connections, which implies in a reduction of the delay. In modern technologies, the delay in connections is so or more significant than the delay in the switching of logic gates. A greater spacing between the connections also contributes to an increase of reliability, due, for example, to the reduction of the possibility of electromigration, as already mentioned above.

Fig. 6. Two options for the implementation of the same function (Reis 2011A)

The reduction of the number of transistors depends on the use of efficient Electronic Design Automation (EDA) tools that transform the logical equations of a system so that in addition to mapping equations in CMOS gates, make optimum use of complex logic gates. In (Conceição et al. 2017) we present a tool to reduce the number of transistors in a circuit through the fusion of networks of transistors that present fanout equal to 1. Also, it is fundamental the use of an automatic synthesis tool that can perform the automatic layout of any logical function. There is no use to achieve a logical optimisation if it is necessary to map (transform) the equations according to the logic gates available in a traditional cell library [which have few functions, in general, no more than 100 functions], as is still done when using traditional EDA systems. This mapping step is called technology mapping, and it represents a step of deoptimization. With this aim, we have developed automatic layout synthesis tools such as ASTRAN (Ziesemer and Reis 2015) (Fig. 7), which allows automatic generation of the layout of any network of transistors (Reis 2011A, B).

Another technique to reduce consumption is through the sizing of the transistors. Modern integrated circuit manufacturing technologies show a significant increase in static power consumption that is often greater than dynamic power consumption. One way to mitigate power consumption, especially the static one, is to carry out a sizing of transistors to optimise power consumption (Posser 2011). In (Reimann et al. 2016) significant decreases in consumption are obtained through the use of automatic

transistor sizing tools. This is also called cell selection, where the cells are selected from a cell library. In this case, cell selection means the selection of cells with a specific size and V_{th} (threshold voltage). In traditional cell libraries, one function has in general 3 sizings (one for less area, one for less power, and one for less delay) and 3 V_{th} (threshold voltage).

Fig. 7. Transistor network layout generated automatically (Ziesemer and Reis 2015)

5 Reliability

As in the reduction of power consumption, in the design of critical systems, it is needed to use techniques to increase reliability at different levels of design abstraction. At the architectural level, a very applied method is the redundancy of modules, especially triple module redundancy (TMR) (Kastensmidt et al. 2006). Another is the temporal redundancy (Nicolaidis 1999) where a signal traverses two paths, one with higher delay and another one with less delay. The difference of delay must be longer than the duration of a transient. Comparing the signal after traversing the two paths indicates whether there has been a transient propagation or not. At the physical level, we can apply different techniques to reduce or avoid problems such as electromigration (Posser et al. 2016, 2017). In the example of Fig. 8, the position of the output pin in the centre (point 4) increases the lifetime of the circuit because it allows reducing the maximum density of current in the segments of the metal layer.

In (Velazco et al. 2007) it is presented a series of works aimed at mitigating the effects of radiation on integrated circuits. In (Kastensmidt et al. 2006; Neuberger et al. 2014; Gennaro et al. 2017; Aguiar et al. 2016; Lazzari et al. 2011) we present some of the results that our research group has obtained in the development of techniques aiming the design tolerant to faults due to transients, as the effects due to radiation.

Fig. 8. Changing current density with the change of position of the output pin (Posser et al. 2016, 2017)

6 Hardware Accelerators

The evolution of computer architectures, that today means, the evolution of microprocessor architectures has been very significant. In the 1970s, one marketing argument from microprocessor producers was the number of instructions that the microprocessor could execute as well as the clock frequency of the microprocessor. In the last decades, there has been a change of paradigm, discontinuing the race for the increase of the clock frequency, because the increment of the clock means an increase of the dynamic consumption. Instead, there was an increase in the number of cores (CPUs) aiming at increasing performance. Initially with homogeneous cores and later with heterogeneous cores.

Currently, we can find chips with multiple CPUs and several GPUs (as can be seen in Fig. 9 (Shao 2016) showing the floorplan of the A8 microprocessor (from Apple). In this same figure, it can be observed that about half of the area is occupied with hardware accelerators, which are modules dedicated to the execution of a specific function. For example, an encryption module placed next to the output/input pins and which will encode the output data and decode the received data. So, the execution of this function will be faster, because it is done by a dedicated module (that means smaller) and with only the needed number of components to perform that function. It also will consume less power.

A more important fact is that the use of hardware accelerators leads to greater energy efficiency (allowing more sustainable computing), mainly due to the reduction in the number of components used to perform a function. At any given time, only the hardware accelerators in use at that time are being powered. So, the hardware accelerators that are not in use are disconnected from the power supply. This strategy is also known as “Dark Silicon”. We can even predict architectures consisting essentially of hardware accelerators, with only one or two small CPUs to manage these hardware accelerators.

Fig. 9. Apple 8 floorplan with 29 hardware accelerators (AnandTech [2014](#); Shao [2016](#))

The introduction of an NPU in A11 is another element characterising the heterogeneity of the SoC (chip system). And we can expect increasingly heterogeneous architectures, with dedicated modules for different operations to be performed by a SoC. In Fig. 10 (Techinsights [2017](#)) the floorplan of the Apple A11 microprocessor is presented, where one of the modules is an NPU (Neural Processing Unit). NPU is mostly dedicated to facial recognition (Techinsights [2017](#)), processing machine learning tasks more efficiently, consuming less energy than CPUs do. The CPUs occupy about 15% of the area of the chip and 6 GPUs occupy about 20% of the area. Most of the area is filled with the hardware accelerators. That is, it is growing in the architecture of Apple microprocessors the use of hardware accelerators.

Fig. 10. Apple 11 floorplan with an NPU (Techninsights [2017](#))

7 Conclusions

To have sustainable computing, when the number of connected devices in the Internet of Things is fast increasing, it is fundamental the design of devices optimised regarding energy consumption. Most of the chips produced today use much more transistors than necessary to perform a function. So, there is a significant space for the optimisation of the number of components. In many devices related to critical applications, the application of techniques for fault tolerance is also fundamental, as nowadays circuits at ground level can have faults due to radiation effects. The reduction of power consumption must be treated at all design abstraction levels in a synthesis flow of integrated systems, from the specification of them in high-level languages to the physical synthesis. It was presented several works that were developed to reduce the power consumption and increase the reliability of integrated systems on a chip, and more

details are shown in the mentioned references. The keyword in the age of the Internet of Things is **optimisation**.

Acknowledgements. We thank CNPq, FINEP, Fapergs, and CAPES for financial support for the development of our team's work, as well as the master's and doctoral students of PGMICRO and PPGC and students of Scientific Initiation who have contributed to the research works that served as the basis for this paper.

References

- Aguiar, Y., Zimpeck, A., Meinhardt, C., Reis, R.: Permanent and single event transient faults reliability evaluation EDA Tool. In: *Microelectronics Reliability*, September 2016, vol. 64, pp. 63–67. Elsevier B.V., Amsterdam (2016). ISSN 0026-2714
- AnandTech (2014). <https://www.anandtech.com/show/8562/chipworks-a8>
- Conceição, C., Moura, G., Pisoni, F., Reis, R.: A cell clustering technique to reduce transistor count. In: *24th IEEE International Conference on Electronics, Circuits and Systems – ICECS 2017*, Batumi, Georgia, 5–8 December 2017, pp. 186–189 (2017). <https://doi.org/10.1109/icecs.2017.8291996>
- Gennaro, R., Rosa, F., Oliveira, A., Kastensmidt, F., Ost, L., Reis, R.: Analyzing the impact of fault tolerance methods in ARM processors under soft errors running Linux and parallelization APIs. *IEEE Trans. Nucl. Sci.* **64**(8) (2017). <https://doi.org/10.1109/tns.2017.2706519>. ISSN 1558–1578
- Geden, B.: Understand and avoid electromigration (EM) & IR-drop in custom IP blocks. Synopsys (2011)
- Lazzari, C., Wirth, G., Kastensmidt, F., Anghel, L., Reis, R.: Asymmetric transistor sizing targeting radiation-hardened circuits. *J. Electr. Eng.* (2011A). <https://doi.org/10.1007/s00202-011-0212-8>. Accessed June 2011
- Kastensmidt, F., Carro, L., Reis, R.: *Fault-Tolerance Techniques for SRAM-Based FPGA*, pp. 1–183. Springer, New York (2006). <https://doi.org/10.1007/978-0-387-31069-5>. ISBN 0-387-31068-1
- Neuberger, G., Wirth, G., Reis, R.: *Protecting Chips Against Hold Time Violations Due to Variability*, pp. 1–107. Springer, New York (2014). <https://doi.org/10.1007/978-94-007-2427-3>. ISBN 978-94-007-2426-6
- Nicolaidis, M.: Time redundancy based soft-error tolerance to rescue nanometer technologies. In: *Proceedings of IEEE VLSI Test Symposium*, vol. 17, pp. 86–94. IEEE Computer Society (1999)
- Posser, G., Flach, G., Wilke, G., Reis, R.: Gate sizing minimizing delay and area. In: *ISVLSI 2011, IEEE Computer Society Annual Symposium on VLSI*, Chennai, India, 4–6 July 2011, pp. 315–316 (2011). <https://doi.org/10.1109/isvlsi.2011.92>. ISBN 978-0-7695-4447-2
- Posser, G., Mishra, V., Jain, P., Reis, R., Sapatnekar, S.: Cell-internal electromigration: analysis and pin placement based optimization. *IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst.* **35**(2), 220–231 (2016). <https://doi.org/10.1109/TCAD.2015.2456427>. ISSN: 0278-0070
- Posser, G., Sapatnekar, S., Reis, R.: *Electromigration Inside Logic Cells*, 118 p. Springer (2017). <https://doi.org/10.1007/978-3-319-48899-8>. ISBN: 978-3-319-48898-1
- Reimann, T., Sze, C., Reis, R.: Challenges of cell selection algorithms in industrial high performance microprocessor designs. In: *Integration*, vol. 52, pp. 347–354. Elsevier B.V., January 2016. <https://doi.org/10.1016/j.vlsi.2015.09.001>. ISSN 0167-9260

- Reis, R.: Redução de Consumo pela Otimização de Componentes. In: SEMISH 2010, Anais do 37º Seminário Integrado de Software e Hardware, Belo Horizonte, 21 a 22 de julho de 2010, pp. 371–379 (2010). ISSN 2175-2761
- Reis, R.: Design automation of transistor networks, a new challenge. In: IEEE International Symposium on Circuits and Systems, ISCAS 2011, Rio de Janeiro, Brasil, 15–19 May 2011, pp. 2485–2488. IEEE Press (2011A). <https://doi.org/10.1109/iscas.2011.5938108>. ISBN 978-1-4244-9472-9
- Reis, R.: Power consumption & reliability in NanoCMOS. In: IEEE NANO, 11th International Conference on Nanotechnology, Portland, USA, 15–19 August 2011 (invited talk), pp. 711–714 (2011B). <https://doi.org/10.1109/nano.2011.6144656>. ISBN 978-1-4577-1515-0
- The Connectivist (2014). <http://ow.ly/i/5vph6/original>
- The Economist (2010). Accessed 6 Sept 2010
- SIA: Semiconductor Industry Association, Rrebooting the IT Revolution (2015). <http://www.semiconductors.org/clientuploads/Resources/RITR%20WEB%20version%20FINAL.pdf>
- IHS Markit: IoT Trend Watch 2018 (2018). https://ihsmarkit.com/forms/thankyou.html?efid=t+m2jEyFYkYQYyoP3YvuHA==&&gasc_id=862037098&&gasc_label=scrXCLnM7m0Q6siGmwM
- Techinsights (2017). <http://techinsights.com/about-techinsights/overview/blog/apple-iphone-8-teardown/>
- Vazquez, J., et al.: Delay sensing for long-term variations and defects monitoring in safety-critical applications. *Analog Integr. Circ. Sig. Process.* **70**(2), 249–263 (2012). <https://doi.org/10.1007/s10470-011-9789-0>. ISSN 0925-1030
- Velazco, R., Fouillat, P., Reis, R.: Radiation Effects on Embedded Systems. Springer, New York (2007). ISBN 978-1-4020-5645-1
- Shao, Y.S.: Design and modeling of specialized architectures. Ph.D. thesis, Harvard, May 2016. <https://ysshao.github.io/papers/shao2016-dissertation.pdf>
- Ziesemer, A., Reis, R.: Physical design automation of transistors network. *Microelectron. Eng.* **148**, 122–128 (2015). <https://doi.org/10.1016/j.mee.2015.10.018>. ISSN 0167-9317

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.