RETROSPECTIVE: Printed Microprocessors

Nathaniel Bleier
nbleier3@illinois.edu
UIUC

Muhammad Husnain Mubarik
mubarik3@illinois.edu
UIUC

Farhan Rasheed
farhan.rasheed@kit.edu
KIT

Jasmin Aghassi-Hagmann
jasmin.aghassi@kit.edu
KIT

Mehdi B Tahoori
mehdi.tahoori@kit.edu
KIT

Rakesh Kumar
rakeshk@illinois.edu
UIUC

I. BACKGROUND

This research paper is the confluence of three disciplines: 1) the development of solution-based functional materials and printing fabrication processes, 2) electronic design automation, device modeling, process design kits and logic design, and 3) computer architecture.

Following the discovery of conducting polymers in the late 1970s, the idea of utilizing solution processible materials for the development of printed and flexible electronics emerged. However, it took several more decades to achieve devices with high field-effect mobility capable of operating at sub-volt levels, thereby making them suitable for applications in the Internet of Things and disposable electronics.

Since 2015, the research teams of Jasmin Aghassi-Hagmann and Mehdi Tahoori have been engaged in a collaborative effort within the framework of a state-funded graduate school called MERAGEM. Their focus has been on the fabrication, modeling, and design of printed inorganic, mainly metal oxide based electronic materials and circuits. Over time, several advancements have been made, resulting in the development of predictive device models, an electronic design automation (EDA) flow specifically tailored for basic printed standard cells, as well as parametric-cells (p-cells) for physical design. These achievements have enabled the elevation of the design abstraction level and facilitated the exploration of larger designs based on electronic devices for low-cost and low-power applications.

In parallel, the research group of Rakesh Kumar had been looking at ultra-low-power processors for quite some time when it started becoming apparent to them that there were a vast number of applications that had not seen much penetration of computing. Bandages were not smart, nor were beer bottles or food packaging. Disposable sensors and smart patches were nowhere to be seen either. A common thread across these applications was that they had stringent cost requirements. A quick lookaround at the prices of silicon-based electronics made it clear that silicon-based chips were still too expensive to support these applications. Furthermore, several of these applications had conformity requirements. E.g., a bandage must be able to bend to the contours of the body. Silicon-based chips couldn't bend.

Printed and flexible electronics were being proposed as a way to target applications with ultra-low-cost and conformality requirements.

Several printing technologies as well as many interesting materials ranging from organic polymers, to inorganic metal oxides and 2D materials were discussed in the research community most of them being solution processible and hence cost and material resource efficient. Depending on the temperature demand of the fabrication process some of these materials could be printed on plastic foils, and hence further reduce cost. Besides, the flexible and sometimes even stretchable substrates allow for great conformity.

Upon further research, it became apparent that previous studies on printed devices primarily focused on applications like displays and OLEDs, which were unsuitable for battery-powered usage due to their high voltages and low mobility. Although some progress had been made with carbon nanotube-based field-effect transistors, the expensive processing steps limited their viability. However, a new generation of printed devices, including electrolyte-gated Field Effect Transistors (EGFETs), with higher field-effect mobility and lower supply voltages - the focus of the teams at KIT - seemed more promising. The three groups stared collaborating in March 2019, which was the beginning of a longer fruitful collaboration.

II. THE STUDY

We decided to focus on printed microprocessors. It was not an obvious choice since previous printed circuits were mostly rudimentary. We reasoned that design, verification, and test costs may dominate overall costs of printed circuits once manufacturing became cheap due to printing. A programmable microprocessor will amortize these costs over larger volume, while supporting programmability for applications that may still need it.

We also had to choose target applications. Husnain Mubarik assembled a list of applications we considered would benefit from printed technologies. These applications were characterized by their ultra-low-cost and conformality requirements, but also by their ultra-low-power requirements (since it was clear the applications needed to be battery powered). Application set included sensing and monitoring applications, smart bandage, timer, Point-of-sale (POS) computation, etc. These applications also had relaxed requirements in terms of precision, sample rate, and duty cycle.

In order to evaluate printed microprocessors, we needed a Process Design Kit (PDK). However, there did not exist a synthesis and physical design ready standard cell library for any low voltage printed device technology. Farhan Rasheed developed two such libraries - one for EGFET and another for carbon nanotube thin-film transistor (CNTTFT) technology.

Once we had the libraries, Nathaniel (Nate) Bleier and Husnain evaluated some existing ultra-low-power microprocessors in printed technologies. These microprocessors were chosen as they were open source and had low gate count.

The observed maximum frequencies of EGFET and CNTTFT microprocessors were below 25 Hz and 45 KHz, respectively. Although these frequencies met the performance requirements of many applications, there were still issues with sample rate and duty cycle requirements in EGFET technology. Additionally, the high energy consumption of the microprocessors posed concerns, as it resulted in unacceptably short lifetimes when using printed batteries. Furthermore, the area values, especially in EGFET technology, were excessively large, emphasizing the need for smaller and more energy-efficient microprocessor designs. It was evident that more compact and efficient designs were necessary to address these challenges.

One opportunity was ISA design from the ground up. ISAs for the existing processors we studied were either register-register ISAs or stack-based ISAs. However, D flipflops in printed technologies are very expensive (due to implementation in resistor-transistor logic). This makes a register-register ISA expensive. Similarly, stack-based ISAs are expensive as they require an expensive RAM-based stack. Unlike RAMs, crosspoint-based ROMs are cheap in a printed technology. So, a Harvard organization is cheaper than a Von Neuman organization.

With above in mind, Nate designed TP-ISA – Tiny Printed ISA targeting printed microprocessor cores. TP-ISA was a Harvard-based memory-memory ISA designed to produce low gate count microprocessors. The number and complexity of instruction were chosen to reduce program size and energy. With an 8-bit program counter, 24-bit instructions, and two 8-bit operands, TP-ISA supported up to 256 words of data memory.

The ISA served as a basis for a design space exploration of printed microprocessor architectures over multiple parameters-datawidths, pipeline depth, etc. The exploration immediately showed the benefits of TP-ISA. The largest TP-ISA cores during the exploration were smaller and lower power than the smallest equivalent pre-existing cores. The best cores outperform pre-existing cores by significant amounts in terms of area and power. We also found that the best printed microprocessor cores are single-stage pipelines.

The benefits of TP-ISA led us to more carefully look at unique opportunities in ISA design that printed technologies afford us. Previous approaches on application or domain-specific customization of hardware added custom instructions to the base ISA. However, the low fabrication cost of printed electronics meant that ISA changes could be made at finer granularity. We played with changing the number and size of operands and registers based on the program size and needs. This *program-specific ISA* would allow for most efficient execution. We think that program-specific microprocessor architecture is a promising area of study for technologies that have low fabrication and prototyping cost.

III. LOOKING BACK AND FORWARD

Looking back, the paper has several messages that are important and should be relevant going forward. First, there needs to be a clear acknowledgment that a vast number of application domains have not seen penetration of computing. Efforts to address these domains would be useful. Second, there is value to targeting metrics beyond the conventional ones (performance, power, security, reliability, etc.). Monetary cost and conformality, for example, are two interesting metrics that the paper focuses on. Toxicity, porosity, biodegradability, etc., also come to mind. Third, a large number of applications have extremely relaxed performance requirements - much less stringent than even the ones that were targeted by works on near and sub-threshold computing. Computer architectures for such applications may be a fruitful area of study. Fourth, the paper showed that unique architectural opportunities exist when using printed technologies.

The work had several limitations. First, it did not present any prototypes. So, it wasn't clear if it was feasible to fabricate printed or flexible processors at high yield. Second, the targeted applications require a microprocessor to be used in conjunction with a sensor, a power supply, and a communication system. it wasn't clear if it was possible to build the full system in a printed technology and what the corresponding benefits and challenges would be. Third, the targeted applications would benefit greatly from a non-volatile memory - such memory did not exist in the two technologies that were studied.

There has been followup work since the publication of the paper to address some of the limitations. Rakesh's group worked with Pragmatic to build one of the earliest programmable flexible microprocessors. These microprocessors - FlexiCores - were designed from the ground up for high yield and energy efficiency. Hundreds of FlexiCore chips were fabricated. We saw good yield establishing the feasbility of low cost flexible microprocessors. Jasmin's team managed to have fully printed process for both passive and active components, and together with Mehdi's group they designed and fabricated fully-printed artificial neurons. Her team also managed to have printed memristors which could potentially be used for printed neuromophic systems. We have also explored the benefits of printed technologies in context of machine learning, prototyping bespoke decision trees and neural networks. The other area of work, led by Mehdi's team, is analog and mixed-signal printed computing.

The ultra-low-cost objective of our work presents challenges, such as the need to sell billions of devices for profitability. Additionally, environmental impact is a concern, as short-lived electronics require biodegradable or recyclable substrates to ensure sustainability in high-volume applications. These challenges will persist and require attention from future researchers in printed and flexible electronics.

Overall, the materials and devices that support printed and flexible electronics have reached sufficient maturity that it is a good time for computer architects to get involved. Even the problems are interesting. The paper made a good case for it.