RETROSPECTIVE: Printed Microprocessors

Rakesh Kumar
rakeshk@illinois.edu
UIUC

I. Background

By 2018, my research group had been looking at ultra-low-power processors for quite some time when it started becoming apparent to us 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 to us that silicon-based chips were still too expensive to support these applications. Furthermore, several of these applications had conformality 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 printed electronics technologies were based on maskless and additive processes that would reduce cost. Similarly, use of cheaper substrates such as plastic could also reduce cost. Use of flexible substrates enabled conformality.

The key challenge for the technology was that it was orders of magnitude worse than silicon-based electronics in terms of performance and density. Also, most circuits that had been built were tiny.

To better understand the technology, I travelled to IDTechex in November (still 2018), a large printed and flexible electronics trade show. I ended up having several interesting conversations at IDTechex and learned that printed and flexible electronics work had mostly focused on passive electronics. There were a small number of companies trying to build active electronics in printed and flexible electronics. I particularly enjoyed my meeting with Scott White. Scott was the CEO of Pragmatic Semiconductors (then Pragmatic Tech), a UK-based company building flexible chips. Scott made a compelling argument that there is value in building complex active circuits in flexible electronics.

After getting back to Champaign, we started looking carefully at research literature. We discovered that past research on printed devices had mostly focused on displays and OLEDs as target applications. These devices often had high voltages and poor mobility values making them unsuitable for battery powered applications. Some work based on carbon nanotube based field-effect transistors had started appearing, but these works relied on expensive processing steps.

However, some new printed devices had started emerging that had relative high mobilities and low supply voltages. Electrolyte-gated Field Effect Transistors (EGFETs), for example, had these properties.

During literature search, I discovered that Mehdi Tahoori and his colleagues at KIT were looking at EGFET-based active electronics. I sent Mehdi an email in March (2019) letting him know of our interest in exploring complex printed systems. Mehdi was excited to talk and that is how the collaboration on this work began.

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.

We observed that the maximum frequency of these microprocessors was less than 25 Hz for EGFET and less than 45 KHz for CNTTFT. However, these frequency values were high enough to meet many applications' performance requirements. Many applications still could not meet their sample rate and duty cycle requirements in EGFET technology. What was more worrisome, however, is that the energy consumption of the microprocessors was too high. Many applications would have unacceptably short lifetimes when using a printed battery.

Area values were also very high, especially in EGFET technology. EGFET microprocessors were at least 11 cm^2 big.

In fact, these numbers were conservative since they did not account for the cost of storing the programs. Since the chosen microprocessors were simple, their software was not dense and required large number of instructions, increasing area and energy overheads. It became clear to us that we needed much smaller, more energy-efficient microprocessor designs.

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 transistor-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. This was not surprising since DFFs are considerably more expensive than combinational cells for printed technologies.

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. For example, if the number of static instructions in a program is N, the program counter can be reduced from eight bits to $\lceil log2N \rceil$ bits.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. Unique performance and density tradeoffs, crosspoint-based instruction ROMs, program-specific ISAs, single-stage Harvard designs, memory-memory ISAs, etc., suggest interesting avenues for research.

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 is 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. In my own research group, we 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. Similarly, we have built and prototyped low-cost flexible encryption engines that may be needed for many applications, especially in the health domain. We have also explored the benefits of printed technologies in context of machine learning, prototyping bespoke decision trees and neural networks.

Our work targets ultra-low-cost. However, this brings its own challenges. An electronic device that costs a cent requires a sale of at least of 15 million units just to cover the salary of a single engineer make 150K per year. This illustrates the business challenge of printed and flexible electronics - needing to sell billions and trillions of devices for profitability. This is a challenge that will confound the entrepreneurs in this space for some time. Another challenge is environmental impact. The targeted applications require a short lifetime after which the electronics will be disposed. If the substrate is non-biodegradable or non-recyclable, the impact on environment can be severe, especially since the applications also require high volume. Any future work on printed or flexible electronics should also consider biodegradability or recylability.

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.