Technical Report

UCAM-CL-TR-829
ISSN 1476-2986

Number 829

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Microelectronic security measures

Philip Christopher Paul

February 2013

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

This technical report is based on a dissertation submitted January 2009 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Pembroke College.

Some figures in this document are best viewed in colour. If you received a black-and-white copy, please consult the online version if necessary.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Abstract

In this dissertation I propose the concept of tamper protection grids for microelectronic security devices made from organic electronic materials. As security devices have become ubiquitous in recent years, they are becoming targets for criminal activity. One general attack route to breach the security is to carry out physical attack after depackaging a device. Commercial security devices use a metal wire mesh within the chip to protect against these attacks. However, as a microchip is physically robust, the mesh is not affected by depackaging.

As a better way of protecting security devices against attacks requiring the chip package to be removed, I investigate a protection grid that is vulnerable to damage if the packaging is tampered with. The protection grid is connected directly to standard bond pads on the microchip, to allow direct electronic measurements, saving the need for complex sensor structures. That way, a security device can monitor the package for integrity, and initiate countermeasures if required.

The feasibility of organic tamper protection grids was evaluated. To establish the viability of the concept, a fabrication method for these devices was developed, the sensitivity to depackaging was assessed, and practical implementation issues were evolved. Inkjet printing was chosen as fabrication route, as devices can be produced at low cost while preserving flexibility of layout. A solution to the problem of adverse surface interaction was found to ensure good print quality on the hydrophobic chip surface. Standard contacts between chip and grid are non-linear and degrade between measurements, however it was shown that stable ohmic contacts are possible using a silver buffer layer. The sensitivity of the grid to reported depackaging methods was tested, and improvements to the structure were found to maximise damage to the grid upon tampering with the package. Practical issues such as measurement stability with temperature and age were evaluated, as well as a first prototype to assess the achievable measurement accuracy. The evaluation of these practical issues shows directions for future work that can develop organic protection grids beyond the proof of concept.

Apart from the previously mentioned invasive attacks, there is a second category of attacks, non-invasive attacks, that do not require the removal of the chip packaging. The most prominent non-invasive attack is power analysis in which the power consumption of a device is used as oracle to reveal the secret key of a security device. Logic gates were designed and fabricated with data-independent power consumption in each clock cycle. However, it is shown that this is not sufficient to protect the secret key.

Despite balancing the discharged capacitances in each clock cycle, the power consumed still depends on input the data. While the overall charge consumed in each clock cycle matches to a few percent, differences within a clock cycle can easily be measured. It was shown that the dominant cause for this imbalance is early propagation, which can be mitigated by ensuring that evaluation in a gate only takes place after all inputs are present. The second major source of imbalance are mismatched discharge paths in logic gates, which result in data-dependent evaluation times of a gate. This source of imbalance is not as trivial to remove, as it conflicts with balancing the discharged capacitances in each clock cycle.

In memory of David Moore, 1951 - 2006.

I would like to thank my late supervisor Dr David Moore for supporting my PhD application, organising the Epson studentship, and his valuable feedback and inspiration during my first year of PhD studies.

I am very grateful to Dr Simon Moore for adopting me as his student, as well as his excellent support and his time and patience to advise in a less familiar field.

I would like to thank Seiko Epson for their generous financial support, studentship grant and use of the CRLE laboratory to carry out experiments. Dr Simon Tam as my supervisor at Epson provided me with valuable guidance, feedback, and support for my technical questions. I am also indebted to Dr Masaya Ishida for his help which allowed me to continue my research despite the closure of the CRLE laboratory in 2007.

Prof. Richard Friend, Prof. Piero Migliorato and Dr Daping Chu have also been very supportive in letting me use their laboratories and equipment, and providing me with advice and help.

I would also like to thank Prof. David Harris for inspiring the work on the secure array multiplier, and for his patience, help, and submitting the multiplier chip for fabrication.

Contents

List of Figures

List of Acronyms

AES advanced encryption standard

ATM Automated Teller Machine, cash dispenser machine

cdf cumulative distribution function

CMOS complementary metal oxide semiconductor

DEMA differential electro-magnetic analysis

DES data encryption standard

DPA differential power analysis

DI de-ionised, pure water

DIL dual in-line

DMSO Dimethyl sulfoxide, $(CH_3)_2SO$

EM electro-magnetic

ESD electro-static discharge

EU European Union

F8T2 poly(9,9-dioctyl-fluorene-co-bithiophene)

FET field-effect transistor

FIB focused ion beam

FIPS federal information protection standard

FPGA field programmable gate array

HCl hydrochloric acid

IBM International Business Machines

IC integrated circuit

IPA Isopropanol, a secondary alcohol. Other names: propan-2-ol, C_3H_7OH

MEMS Micro Electro-Mechanical Systems, also known as 'microsystems'

MOSFET metal oxide semiconductor field-effect transistor

MOSIS Metal Oxide Semiconductor Implementation System

NED normalised energy deviation

NMOS n-type metal oxide semiconductor

NSD normalised standard deviation

OTFT organic thin-film transistor

P3HT poly(3-hexylthiophene)

PE poly(ethylene)

PECVD plasma-enhanced chemical vapour deposition

PEDOT poly(3,4-ethylenedioxythiophene)

PMOS p-type metal oxide semiconductor

PSS poly(styrenesulfonate)

PVP poly(vinyl phenol)

RAL Rutherford Appleton Laboratory, Didcot, OX11 0QX

RC resistance capacitance

RFID radio-frequency identification

ROM read only memory

RSA Rivest, Shamir, Adlemann public key cipher

SEM scanning electron microscope

SIM subscriber identification module

SOI silicon on insulator

SPA simple power analysis

SPICE simulator with integrated circuit emphasis

TFT thin-film transistor

TPM trusted platform module

UV ultra-violet light

VLSI very large scale integrated circuit

Chapter One

1 Introduction

1.1 Motivation

"For generations, people have defined and protected their property and their privacy using locks, fences, signatures, seals, account books, and meters. These have been supported by a host of social constructs ranging from international treaties through national laws to manners and customs. This is changing, and quickly. Most records are now electronic, from bank accounts to registers of real property; and transactions are increasingly electronic, as shopping moves to the Internet. Just as important, but less obvious, are the many everyday systems that have been quietly automated" [Anderson, 2001b].

In analogy to physical locks and protection mechanisms, electronic record keeping and the automation of transactions require electronic security mechanisms to protect people's property and privacy. While the protection measures outside the electronic realm have been well established for a long time, electronic security mechanisms are still comparatively new. With the rapid progress in technology and the high complexity of electronic systems, security mechanisms to protect the systems are still evolving [Anderson, 1994a]. Security flaws are identified and removed, however new flaws are often introduced when features are added [Anderson, 1994a]. More importantly, in a distributed system, the liability for losses due to security failures may not lie with the party that is able to protect a system, leading to moral hazard [Anderson, 2001a; Anderson and Moore, 2006].

Protection and security mechanisms naturally carry a cost of implementation. If a security measure costs more than the value of fraud it will prevent, its implementation is uneconomical. Crime generally follows the same pattern of maximising benefit and minimising effort. If security is increased in one part of a system, criminals are likely to target less protected parts until these are also protected [Schneier, 1997]. The most likely security system to be implemented is one that is cost-efficient. Rather than maximising security, crime is deterred by making the system only slightly too difficult for a criminal to attack. Once the implemented security measures are defeated, the next round of improvements is made at a later stage.

In the European Union (EU) and some other countries, it has become economical for banks to replace magnetic strip cards with chip and pin smart cards (be it for security reasons or to diffuse liability), while in the USA where laws and liabilities are different, banks still rely on magnetic strip automated teller machine (ATM) cards [Anderson, 1994a,b]. The introduction of the 'chip and pin' system has caused ATM card fraud in the UK to shift to modes where either the card is not present (e.g. internet shopping) or where the fallback magnetic strip is

used instead of the chip [Burns and Weir, 2008].

Outside the banking system, smart cards, or more generally hardware security modules, have become widespread, as they are now cheap enough to the point of being disposable. Typical applications include subscriber identification module (SIM) cards in mobile phones, access control systems, radio-frequency identification (RFID) tags, billing systems for transportation networks, telephone cards, pre-payment electricity meters, and electronic passports and identification cards. Hardware security modules have also enabled new business models for providing software or entertainment content by means of pay-TV cards, hardware dongles to unlock software, and trusted computing platform modules (TPM) [Pearson, 2002]. In the most general form, these devices perform some form of encryption or decryption function, to either secure data that are transmitted or stored, or to provide a means of authentication, or a digital signature. To this end, a secret key is stored on the chip to render the cryptographic function unique. Attackers from unauthorised parties aim to retrieve this secret key and to duplicate or otherwise misuse the security module, or decrypt encrypted data.

With the rising popularity of hardware security modules and a multitude of applications, it is evident that the attention of criminals to security weaknesses is bound to increase. There are two examples of applications where security devices have been used for a number of years, namely telephone cards and pay-TV decoders. Attacks have been carried out on these devices since their introduction in the 1990s, therefore a number of attack methods have been established.

1.2 Background

Systems based on hardware security modules can be analysed by considering the trust boundaries between the parts of the system, which differ depending on the application. Apart from the parties carrying out a transaction, there may be separate parties that provide the security devices and the infrastructure to enable transactions. Each of these parties (as well as outsiders to the system) will have different incentives relating to the overall security system, depending on whether there is any benefit derived from cheating. The system and the security device need to be designed to be resilient against attacks by the parties that would benefit from cheating [Schneier and Shostack, 1999].

A wealth of protocols exist for transactions and cryptographic algorithms that have been proven secure under certain assumptions. A standard assumption is that there is only black-box access to secret information of honest parties [Gennaro et al., 2004]. For hardware security modules, processing of the secure transaction takes place inside the device, which may give the false impression that only black-box access is provided. Attacks on smart cards often aim to break the black-box assumption, and in turn break the security of the protocol.

Attackers have been classified into three groups according to knowledge and capability [Abraham et al., 1991]. Class I represents clever outsiders with insufficient knowledge of the security system and access to only moderately sophisticated equipment. Class II comprises knowledgeable insiders with good understanding of parts of the system and access to highly sophisticated tools and instruments. Class III includes funded organisations with substantial resources capable of in-depth analysis of the entire system, and thus the design of sophisticated attacks [Abraham et al., 1991]. The classification of attackers may be useful in the first instance when designing a security system. Depending on the trust boundaries between the parties and the application of the security system, the amount of value that can be extracted from breaches of trust can be quantified. The amount of value that can be extracted correlates with the category and capability of likely attackers. However, the presented classification of attackers is more transient in practice. The amount of knowledge about a system will increase with time, there-

fore becoming available to attackers with access to only moderately sophisticated equipment. Generally, the security of a system should not rely on available knowledge about it [Kerckhoffs, 1883]. Furthermore, the capability of attackers of a certain category also increases with time, as equipment becomes cheaper [Anderson and Kuhn, 1998] and more attacks are published (by attackers of a higher category). Once an attack has been published, it is possible to automate and allow less knowledgeable attackers to implement it. Similar to categorising the capability of adversaries, there are schemes for categorising security levels of hardware, such as federal information protection standard (FIPS) 140 or Common Criteria protection profiles. New attacks may be discovered which then reduce the device's security classification, making this metric transient as well.

Attacks on hardware security modules can also be categorised by attack route and method. That way, the classification is static with time, and not dependent on current hardware prices or availability of attack recipes. Grouping similar attacks also serves as a good starting point for developing countermeasures. The most general attack classification may differentiate between 'invasive', 'semi-invasive', or 'non-invasive' attacks. The criterion used to distinguish these categories is the extent to which the microchip is depackaged or physically tampered with. Each of these attacks can be 'active' or 'passive', which signifies whether the inputs to the chip and circuits are deliberately altered or not [Anderson et al., 2006].

1.2.1 Invasive and semi-invasive attacks

The most important characteristic for attacks in these two groups is the removal of the packaging of the hardware security module, which surrounds the microchip. The key difference between invasive and semi-invasive attacks is whether the passivation layer of the microchip is also (partially) removed. In an invasive attack this is the case, while in a semi-invasive attack, the die and passivation remain intact [Anderson et al., 2006].

The threat model associated with these attacks is primarily limited by the fact that depackaging or tampering is likely to leave a visible trace. For example, if the device is used as a means of identification of the card holder, an honest card holder will be alerted that an attack may have taken place if it leaves physical evidence. Conversely, a dishonest card holder may physically tamper with the module in order to impersonate someone else. In this case the operator of the terminal (e.g. a checkout clerk in a shop) may spot that the card has been tampered with. Invasive attacks are likely to be used if the card holder has a benefit from using a counterfeit or cloned card, e.g. in a pay-TV system [Schneier and Shostack, 1999].

The first step in both attack categories is to remove the packaging of the security device. If the device is mounted on a smart card, heating the card until it becomes flexible and the glue becomes soft will allow the chip to be removed from the card [Kömmerling and Kuhn, 1999]. The epoxy encapsulation of the chip may then be removed with pure nitric acid. Nitric acid does not etch away the aluminium bond pads, so long as it is anhydrous [Beck and Wilson, 1998; Kömmerling and Kuhn, 1999]. Local removal of the packaging epoxy is possible using acid in a jet etch system, or by oxygen or tetraflouromethane (CF_4) plasma etching. Specialised solvents are available for certain packaging materials, e.g. ethylenediamine to remove polyimides [Beck and Wilson, 1998] Depackaging may also be done physically rather than chemically. The obvious physical methods to remove the packaging material are drilling and milling [Beck and Wilson, 1998]. However, mechanical devices similar to vices are available to crack the epoxy packaging to reveal the die [Beck and Wilson, 1998]. Laser machining, water machining, or manual material removal techniques [Weingart, 2000] may also be used. To remove the chip passivation layer, dry etching [Beck and Wilson, 1998] or focused ion beams (FIB), or ultrasonically agitated probe needles may be used [Kömmerling and Kuhn, 1999].

There is a broad range of known invasive attacks. The most basic invasive attacks may be probes (either standard probe needles, or pico probes) on internal signal wires [Handschuh et al., 1999]. For a passive attack, the probes only record transmitted data, while for an active attack signals would be injected into the chip circuit. Alternatively, FIBs or lasers may be used to re-instate disabled circuitry, or alter the circuit or memory contents [Kömmerling and Kuhn, 1999]. Circuit modifications may also be used to disable parts of a chip, or to introduce (permanent) faults into computations. Differential fault analysis is a powerful cryptanalysis tool that can re-construct secret information from cryptography errors [Bar-El et al., 2006; Biham and Shamir, 1997]. Reconstruction of the chip's circuit diagram may be possible by dissecting the die layer by layer and analysing images taken of each layer [Blythe et al., 1993; Nohl et al., 2008]. The circuit diagram may then be used to aid and prepare other attacks, e.g. by locating regions of interest on the chip for electro-magnetic (EM) analysis, or locating wires to probe.

Invasive attacks may also be used to recover memory contents. For some read-only memory (ROM) types, the content may be recovered from the contact patterns after removal of the metal layers. Other ROMs store data by varying the dopant in a transistor, which becomes visible if a dopant-selective etchant is used [Kömmerling and Kuhn, 1999]. Re-writable memory types show data remanence effects, in particular if the same value is stored over long periods of time [Gutmann, 2001]. For long-term damage, the data may be visible as physical degradation (electromigration) under a scanning electron microscope (SEM) [Gutmann, 2001].

Semi-invasive attacks aim to probe internal data of the chip without the need to break the passivation layer. Radiation sources [Dyer et al., 2001] or lasers may be used to induce carriers into a transistor and make it conductive. That way, memory contents may be modified, or transient faults may be induced into computations to carry out differential fault analysis [Skorobogatov and Anderson, 2003]. EM pulses may be used in a similar way to induce faults or read memory contents, as well as re-construct the metal wiring layout of the upper metal layers [Samyde et al., 2002].

The analysis of EM emissions may also be carried out as a semi-invasive attack. The resolution of the EM probe depends directly on the distance to the die, therefore removal of the chip packaging increases the resolution. Positioning the probe above interesting locations on the chip is simplified if the layout is visible. [Gandolfi et al., 2001]

There are several approaches to defend against invasive and semi-invasive attacks. Devices may be made tamper-resistant, tamper-responding, or tamper-evident [Weingart, 2000]. The required grade of tamper protection is determined by the most likely attack method, which may be derived from the threat model.

In more sophisticated hardware security modules there is usually more margin for fitting tamper protection measures. For example, some military cipher machines go so far as to contain thermite charges to destroy them if they are tampered with [Anderson et al., 2006]. On the small end of the size scale, special coatings have been developed to prevent tampering and probing. Removal of this special conformal glue results in damage to the microchip. The coating technology is for military use and is classified [Anderson and Kuhn, 1996; Weingart, 2000]. Alternatively, hard barriers, silicon on insulator (SOI) technology and specially adapted chip topography may be used to prevent invasive attacks [Weingart, 2000]. An example of a hard barrier for smart cards is the Axalto 'Sishell' [Bonvalot and Leibenguth, 2003]. The actual security device is bonded to a second silicon protection block, and subsequently thinned down such that any attempt to separate the two silicon devices results in damage to the security device [Bonvalot and Leibenguth, 2003]. Another method to prevent tampering is to randomly distribute particles in the packaging material, and derive the cryptographic key based on the position of these particles [Koemmerling, 2000].

If attacks are likely to be carried out when the device is powered on (e.g. semi-invasive attacks, probing attacks), tamper resistance may not strictly be required, and a simpler tamper-responding approach may be sufficient. At the larger end of the size spectrum, the cryptographic key may be stored in battery backed memory surrounded by a flexible membrane containing conductive tracks. The battery of the module is used to continuously monitor the integrity of the membrane, and erase all data upon detection of abnormalities [Abraham et al., 1991; Anderson and Kuhn, 1996; Dyer et al., 2001; Gore; Weingart, 2000]. Other designs to detect tampering with circuit board-sized devices may include stressed glass circuit boards, piezo-electric sheets, light scattering pattern detection, ultrasonic waves, and radiation sensors [Weingart, 2000].

For single-chip devices such as smart cards, the cost of adding elaborate sensors or batteries is likely to exceed the small profit margins. These advanced features are not feasible for a device vendor to implement if the sale price of the device cannot be increased to cover the cost. Tamper detection techniques implemented on single chip devices may consist of a coating with high dielectic constant, and a capacitance sensor to detect its presence, or an optical sensor under an opaque coating [Anderson and Kuhn, 1996]. Fault-resistant or fault-detecting logic styles may be used to detect fault attacks [Moore et al., October 2003], however care must be taken not to introduce new vulnerabilities with these measures [Yen et al., 2002]. A common technique for intrusion detection is to reserve the topmost metal layer for a wire sensor mesh, and continuously monitoring the integrity of the mesh while the device is powered on [Anderson and Kuhn, 1996; Koeune and Standaert, 2005; Kömmerling and Kuhn, 1999]. However, the problem with these sensors is that they are located below the chip passivation, and are therefore not affected by depackaging. The lack of monitoring while the device is powered off, as well as the digital nature of the wire grid (connected or not) makes the grids tolerant to modifications. A FIB workstation can be used to cut fine holes through or between the wires, and then to fill the holes with metal to create a contact pad for probing as well as to re-route any severed wires [Anderson, 2001b; Kömmerling and Kuhn, 1999].

Tamper evident devices may be created by chemical or mechanical means, such as holographic tape, polished surfaces or bleeding paint in the packaging [Weingart, 2000]. However, tamper evidence is only applicable in certain cases. If traces of tampering are unlikely to be checked regularly, or the mechanism does not resist the wear and tear of everyday use of the device, then tamper evidence is not sufficient to protect a device.

1.2.2 Non-invasive attacks

In contrast to invasive and semi-invasive attacks, non-invasive attacks do not require physical tampering with the device. Generally, non-invasive attacks involve the monitoring of the externally accessible properties of the security device during operation. Passive attacks are carried out only by eavesdropping (and possibly sending valid, but chosen messages to the device), while active attacks send bogus inputs to the device. For these attacks, the term 'side-channel' has been established to describe the transmission (or leakage) of data through means not intended for communication. Examples of side-channels are the time duration [Kocher, 1996], power consumption [Kocher et al., 1999], and electromagnetic emission [Agrawal et al., 2003] of a computation. Other effective non-invasive attacks have been carried out by inserting glitches in the clock frequency and the power supply voltage [Anderson and Kuhn, 1998; Bar-El et al., 2006].

Even though several (improved) variants of the timing attack exist [Dhem et al., 2000; Handschuh and Heys, 1999; Kocher, 1996; Schindler et al., 2001; Schindler, 2000], they can be hampered using algorithmic measures. These aim to ensure constant run time, as well as using blinding factors to randomise the remaining run time variations [Kocher, 1996; Schindler, 2000].

In contrast to run time, power consumption is a fundamental property of the microchip circuit. Every time a logic gate switches, a small amount of power is consumed in the gate [Weste and Harris, 2004]. As the number of switching events in a circuit depends on the processed data, the power consumption of a microchip is data-dependent. Power analysis aims to recover the secret key of the cryptographic operation from the input current into the security device. The family of power analysis attacks was first introduced in 1999 by Paul Kocher [Kocher et al., 1999].

In addition to power analysis, EM emissions from microchips are directly related to the voltages and currents in a security device. These are termed direct EM emissions [Li et al., 2005a,b]. Appropriate probes to detect these emissions can be used to resolve the currents in parts of the circuitry [Gandolfi et al., 2001; Quisquater and Samyde, 2001]. The same mathematical techniques as used in power analysis may be used on direct EM emissions to recover the secret information. In that sense, direct EM emissions may be considered as spatially resolved power consumption.

Apart from direct EM emanations due to the flowing currents, coupling effects between signals in close proximity result in modulated emissions [Agrawal et al., 2003]. Similarly, modulation effects from RF signals injected into smart cards can be used to measure switching events as temporary changes in impedance between power and ground [Burnside et al., 2008]. These modulation effects may yield an even finer-grained picture of internal state and switching events in a security device.

The most basic method of power analysis (simple power analysis, SPA) attempts to interpret individual power consumption or EM emissions traces to recover the cryptographic key [Kocher et al., 1999]. Countermeasures against SPA include obscuring the power trace to prevent direct interpretation [Chevallier-Mames et al., 2004]. Differential power (DPA) and differential EM analysis (DEMA) methods are significantly more powerful than SPA, as they compare multiple power traces. Subtraction and averaging of traces allows unrelated parts of the curves to be removed, and even minute differences in data-dependent power consumption to be revealed. The relative amplitudes of the data-dependent power consumption from the cryptographic algorithm and the total power consumption represent a signal to noise ratio. The cipher and protocol may either be public information found in the data sheet, adhering to Kerckhoffs' principle [Kerckhoffs, 1883]. Or it may be reverse-engineered [Nohl et al., 2007]. In the power analysis attack, the cryptographic computation is retraced to generate a model of the internal state of the device. The power consumption traces are compared or correlated to the model of the internal state to reveal the secret key.

A wealth of mathematical methods and variants of the DPA attack have been established to increase the amount information extracted from power consumption traces, and thus reduce the number of traces required for a successful attack. Cipher-specific improvements can be made to the standard DPA attack [Bevan and Knudsen, 2003; Brier et al., 2004; Coron, 1999; Lemke et al., 2004; Wang et al., 1999].

One group of suggested countermeasures against differential power analysis focuses on hiding the signal by increasing variations (noise) in the power consumption traces. Apart from adding a noise source [Clavier et al., 2000], de-synchronisation of operations may prevent averaging of power traces. De-synchronisation may be achieved by random variations in clock frequency (clock gating) [Benini et al., 2003], or randomly inserting interrupts or no-op instructions [Clavier et al., 2000]. However, these measures do not completely prevent DPA, as timing may be re-synchronised by signal processing techniques [Akkar et al., 2000; Quisquater and Koeune, 2002] or by integration of the power traces [Clavier et al., 2000]. Another approach to hiding data is to carry out many operations in parallel, so that power traces of multiple operations overlap, making it difficult to observe the power traces of individual operations [Gebotys, 2004]. Generally, these measures increase the effort required for filtering out the additional noise

(more traces needed, as well as more pre-processing), but do not prevent a successful attack, as the signal remains present in the power consumption traces.

Another group of countermeasures against DPA focuses on methods to transform the cryptographic operation in a way that the sensitive computations are not carried out on the original data, but on obscured or masked data values. That way the internal state does not represent the values expected when re-tracing the cipher, and the power consumption traces do not correlate with the attacker's model. For example, for the data encryption standard (DES) cipher and the Rivest, Shamir, Adlemann public key cipher (RSA), the intermediate data may be split into several individual variables which are only combined later [Goubin and Patarin, 1999]. Masked algorithms have been also developed for other ciphers such as advanced encryption standard (AES) [Akkar and Giraud, 2001; Akkar and Goubin, 2003; Blomer et al., 2004; Oswald and Schramm, 2006; Oswald et al., 2005; Pramstaller et al., 2004; Trichina et al., 2005] and elliptic curve cryptography [Avanzi, 2003; Coron, 1999; Mamiya et al., 2004]. More universally, masking may be carried out on the circuit or logic gate level, thus being applicable to the entire chip rather than just the cryptographic algorithm [Popp and Mangard, 2005].

While masking is effective against standard (first-order) DPA, it may be broken using higher-order DPA where multiple points on the power consumption trace are used simultaneously [Messerges, 2000; Waddle and Wagner, 2004]. This DPA variant has been successfully used against several masking methods [Akkar et al., 2004; Coron and Goubin, 2000; Mangard et al., 2005; Oswald et al., 2006]. From an information-theoretical point of view, the most efficient version of DPA is the template attack [Chari et al., 2003; Fahn and Pearson, 1999; Rechberger and Oswald, 2005]. If a reference device is available, a device-specific template can be assembled which is then used as a matched filter for the attack. It has been found that, when using a template, information extraction from masked implementations of a cipher is as efficient as for non-masked implementations, therefore completely defeating masking [Oswald and Mangard, 2006].

Hiding the data-dependent power consumption in noise or masking it with random data have been shown to be ineffective, therefore it is clear that the only permanent solution to the problem must be to prevent the data dependency of the power consumption signal. De-coupling the chip from the power supply using capacitors allows the amount of charge drawn from the supply to be fixed. Any excess charge on the capacitor is dumped before re-charging [Shamir, 2000]. However, this method remains vulnerable to attacks using EM emissions, which can nevertheless resolve the actual power consumption of the chip. More recently, dynamic dual-rail logic styles have been proposed and evaluated [Moore et al., 2002; Sokolov et al., 2005, 2004; Tiri and Verbauwhede, 2004a; Tiri et al., 2002]. Dynamic logic is somewhat similar to implementing the suggested de-coupling capacitor method to the logic gate level. A standard complementary metal oxide semiconductor (CMOS) logic gate consists of two parts, a pull-up network made from p-type (PMOS) transistors, and a pull-down network made from n-type (NMOS) transistors. In dynamic logic, the pull-up part of the logic gate is replaced with a single clocked PMOS pre-charge transistor, which charges the (former) centre node whenever the clock is low. When the clock switches high, the PMOS transistor is turned off, and the NMOS branch of the logic gate either discharges the node, or not [Weste and Harris, 2004]. Dual-rail logic expands each logic gate by its complement. A zero is thus represented by '01' and a one by '10'. That way, one of the two halves of the logic gate remains charged and the other is discharged in every clock cycle. Dual-rail circuits are also suitable for self-timed designs without a central clock, therefore preventing clock glitch attacks, as well as for self-checking logic to identify attacks inserting faults into computations [Moore et al., 2002, October 2003]. To achieve data-independent power consumption, the parasitic capacitances of both halves of a logic gate must be equal [Tiri et al., 2002]. This extends to the wiring between logic gates, which requires additions to current

integrated circuit (IC) design tools [Kulikowski et al., 2006a; Tiri and Verbauwhede, 2004b]. A sample dynamic dual-rail IC was shown to have reduced data-dependent power consumption, but not to be perfectly balanced [Fournier et al., 2003]. Further pitfalls in processor design include resource sharing, optimisations and caches [Tiri, 2007], as well as early propagation effects in the circuit [Kulikowski et al., 2006b].

If the data dependency of the power consumption has not been completely eliminated, information leakage from the device occurs at a certain rate. If the information leakage rate of devices can be quantified, it is possible to implement a leakage-tolerant security protocol to change cryptographic key (or replace the device) after fewer cryptographic operations than are necessary for an attacker to extract the secret key [Chari et al., 1999; Kocher, 2005]. A framework needs to be established to quantify the information leakage rate [Coron et al., 2004; Macé et al., 2007; Standaert et al., 2006a,b]. Given the large number of DPA variants, quantification of leakage by number of cryptographic operations is difficult.

1.3 Thesis

So far, no robust countermeasure against neither invasive, nor non-invasive attacks has been developed for low-cost small scale security devices. This dissertation is a contribution to the development of more secure security devices, by evaluating proposals for protection measures against each of the two major smart card attack categories, invasive and non-invasive attacks. It is to be noted however, that in the present state of the art, no single technique allows provision of perfect security, even considering one particular attack [Koeune and Standaert, 2005].

As the majority of invasive and semi-invasive attacks are carried out while the device is powered on, a tamper responding protection measure may be sufficient to prevent these attacks from succeeding. Certified protection schemes exist for higher end devices of larger size (circuit boards). These schemes constantly monitor protection grids surrounding the security-sensitive circuitry. On the smaller, individual microchip scale it has been shown that the commonly used metal protection grids are not effective to prevent attacks. In this dissertation, an improved protection grid for microchip application is proposed, which addresses the issues associated with metal protection grids. It is estimated that fabrication cost of a device is not likely to increase, compared to devices devoting an entire metallisation layer to protection.

It has been shown that most countermeasures against non-invasive attacks can be defeated. Countermeasures aimed at decreasing the signal-to-noise ratio merely require a larger number of power consumption traces. Seemingly perfect countermeasures against power analysis attacks that randomise the processed data have been shown to be defeated by improved mathematical models. Therefore, the elimination of data-dependent power consumption entirely remains the only method to prevent power analysis attacks. However, no perfectly balanced logic gate exists so far. To better understand the cause for data-dependent power consumption, a dynamic dual-rail logic gate is designed and evaluated.

The protection measure against invasive and semi-invasive attacks is presented in Chapter 2. The fabrication of the proposed protection structure is described in Chapter 3. In Chapter 4, the security properties of the protection grids are evaluated. Chapter 5 presents the evaluation of environmental stability and prototype implementations. The balanced dual-rail dynamic logic gates against non-invasive attacks are discussed and evaluated in Chapter 6. Finally, the conclusions drawn from this work are presented in Chapter 7.

1.4 Publications

There are two patent applications pending for the tamper protection grids. They are filed under GB0717783.5/USP15943A and GB0718001.1/USP15968A.

This work has been published at the BLISS 2008 conference in Edinburgh. The proceedings paper is available under the following reference:
P. Paul; S. Moore; S. Tam, "Tamper Protection for Security Devices", ECSIS Symposium on Bio-inspired Learning and Intelligent Systems for Security, BLISS 2008, p. 92–96.

Journal publications are in preparation. The non-disclosure agreement with Epson and the patent applications prevented earlier publication, as did the closure of the Epson Cambridge laboratory.

Further publication:
P. Oikonomakos; P.C. Paul; S.W. Moore; S.W.-B. Tam; H. Ebihara, "Dynamic-logic PLA on low-temperature polysilicon TFT technology", IEE Electronics Letters, Volume 43, no 5, p. 23–24, 2007

Chapter Two

Tamper Protection Grids

2.1 Introduction

Most invasive and semi-invasive attacks aim to extract the internal state or the internal signals of a security device while it is powered on. There are only two invasive attacks that are carried out while the device is powered off, and both are destructive, in the sense that the security device is destroyed in the process. The two variants are the layer-by-layer reverse-engineering of the microchip circuit [Blythe et al., 1993], and the analysis of ROM contents [Kömmerling and Kuhn, 1999]. Tamper-resistant packaging would need to prevent the success of these two attacks, which is deemed impossible to achieve in practice [Smith and Weingart, 1998]. So long as the cryptographic key is not stored in an easily readable ROM memory, nor hard-wired into the circuit, these two attacks would not actually recover the key directly. To protect the cryptographic key it should therefore be sufficient to use a tamper-responding countermeasure. If a chip can detect tampering with the packaging, it can refuse to carry out any sensitive operations that may give access to the secret key.

Examples of current tamper-responding technology that have been used on commercial security devices are protection grids on the top metal layer of the chip, and/or light sensors. Both have been shown to be easily bypassed [Kömmerling and Kuhn, 1999]. For larger devices, tamper protection grids incorporated into the packaging have been more successful. The International Business Machines (IBM) 4758 cryptographic processor which utilises this protection measure has been certified under FIPS-140 [Weingart, 2000], a security standard for microelectronic devices. The main differences between the metal protection grids of a single chip and the protection grids for larger devices are the probability of damage and the grid resistance. Metal protection grids are not actually damaged in the depackaging process [Kömmerling and Kuhn, 1999], while the protection grids incorporated in the packaging are likely to be damaged when the device is tampered with. The integrity of the metal protection therefore does not guarantee that the packaging (which is the barrier against invasive attacks) is in place. Furthermore, a metal grid has a negligibly low resistance, therefore it is difficult to measure any resistance change from short circuits due to the small absolute differences in resistance. The IBM 4758 protection utilises higher resistance wiring in several layers which can be used to detect both open and (local) short circuits [Anderson, 2001b].

It is reported that the protection grid of the IBM 4758 cryptographic processor board consists of 'conductive organic lines on a polyester substrate' [Weingart, 2000]. Fabrication of these

conductive strips occurs by 'doping a urethane sheet' (a type of polyester) in tracks significantly wider than 0.1 mm [Anderson, 2001b]. Several layers are wrapped around the device and connected via a thin flexible cable [Smith and Weingart, 1998]. This fabrication method is too coarse for small size security devices such as smart cards (1 mm ^{2}), and needs to be scaled down.

Since the development of the IBM 4758, organic electronics have made significant advances [Burns et al., 2003; Dimitrakopoulos and Malenfant, 2002; Dimitrakopoulos and Mascaro, 2001; Horowitz, 2004; Kawase et al., 2005; Paul et al., 2003]. High-resolution patterning methods also exist [Sirringhaus et al., 2000], allowing organic grids to be scaled down. Therefore, organic protection layers are investigated as a tamper-responding protection measure for individual microchip devices.

2.1.1 Requirements

The following requirements for a tamper-responding protection grid have been identified:

1. **Sensitivity:** The grid must be damaged irreparably in an attack. Damage to the grid must be detectable by the security device, so that it can verify the integrity of the packaging.

2. **Entropy:** The properties of the protection grid must have sufficient entropy to be difficult for an attacker to determine and play back at a device.

3. **Economy:** Addition of the protection grid must have minimal impact on device cost.

4. **Durability:** The protection grid must be stable, with a lifetime similar to the device it protects.

The protection grid essentially either acts as a sensor that detects whether the (epoxy) packaging of the microchip is in place, or as a fragile key store. For cost-sensitive applications, separate components to supply security devices with energy when not in use (capacitors, batteries) are not affordable. This prevents continuous monitoring of the protection grid, as is done in larger devices [Weingart, 2000]. The integrity of the grid can therefore only be determined when the device is in operation. It may thus be possible for an attack to be carried out while the device is powered off. If the protection grid can be repaired by an attacker before powering on the device, the tampering will not be detected. Similarly, a problem exists if the properties of the protection grid can be read out. When the device is powered on again, these properties could be emulated to the device, defeating the tamper protection measure. Therefore it is important that the grid is designed in a way to prevent an attacker from repairing the protection grid, and from measuring the grid properties.

It is also important that the protection measure does not overly increase the fabrication cost, as any cost differences are multiplied by the large number of devices produced in a high-volume application. As metal protection grids have been used in commercial security devices, the cost of an additional metal layer may serve as a benchmark. The simplest way for a microchip to interface with external devices is by direct electrical connection through bond pads. Using suitable organic electronic materials, a protection layer can be designed to operate at a similar voltage range to the security device. Similar voltage ranges allow direct electrical connection between the chip and the protection grid via the bond pads. Furthermore, the electrical characteristics of the protection grid must not degrade to trigger a false alarm during the lifetime of the security device. Premature replacement obviously increases the cost of the security system, and a short lifetime may make the addition of protection grids infeasible.

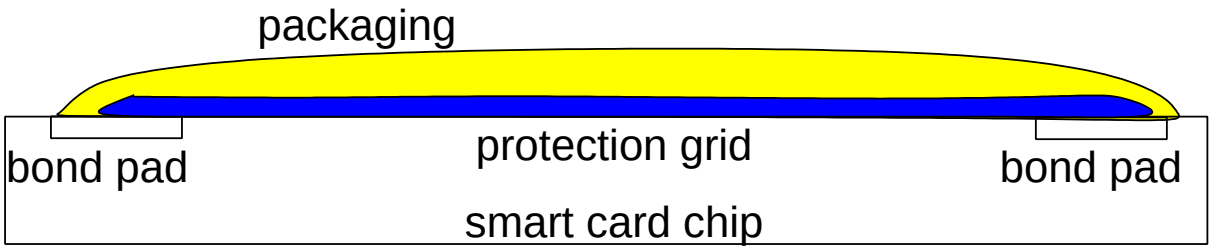

Figure 2.1: Side view of simple protection grid printed on top of a microchip

2.2 Passive protection grids

The most basic type of organic protection grid for microchips consists of resistor lines covering the device, similar to the protection grids used for the IBM 4758. This protection grid design will be termed 'passive protection grids' as it does not incorporate any active amplification or switching. If the resistor lines are integrated into the packaging, attempts to remove the packaging result in damage to the grid and a change in resistance.

The grid may be characterised by measuring the resistance directly (e.g. current for a fixed voltage), the resistance ratio(s) of multiple segments, the RC time constant, or possible non-linear characteristics such as Schottky barrier height of contacts or between different materials. To make spoofing (emulation of the protection grid) difficult for an attacker, the properties of the protection grid may be made unique for each device. Naturally occurring, or deliberate manufacturing variations may be utilised to alter the characteristics of the protection layer from one chip to the next. If the protection grid of one device is successfully characterised by an attacker, this knowledge is not useful for another device. However, care must be taken to design the grid such that characterisation is difficult, for example by ensuring that the contact pads connecting the grid to the chip cannot be probed without altering or destroying the grid itself.

2.2.1 Layout

The simplest form of a protection grid would be to use a single protection layer that is printed directly on the microchip, as illustrated in Figure 2.1. Fabrication of the protection grid (blue) would only involve one additional printing step before the packaging (yellow) is added. The layout pattern of the grid should be chosen so that security sensitive areas are sufficiently protected to prevent local depackaging. Using appropriate deposition and patterning technology, the pattern may also be varied dynamically between individual security devices, to obscure the location of the grid lines and increase the probability of damage. A potential weakness of printing directly onto the die may be lower sensitivity to depackaging, as the grid sits below the packaging material. This may allow the packaging to be thinned down without damaging the grid.

To guard against this problem, the protection grid may be sandwiched between two epoxy layers, as shown in Figure 2.2. The base layer (red) sits below the protection grid (blue), which is encapsulated by more packaging (yellow). This scheme should be designed to ensure that enough insulator material (red and yellow) is present below the protection grid to protect the chip from an attack, should an attacker manage to remove all packaging material above the protection grid. Topographical features may also be incorporated into the tamper protection grid layout, to make it harder to determine the amount of packaging that can be removed before damaging the protection grid. If the topographical layer and the encapsulation layer are made from similar or identical materials, then it is more difficult to stop the material removal or etch-

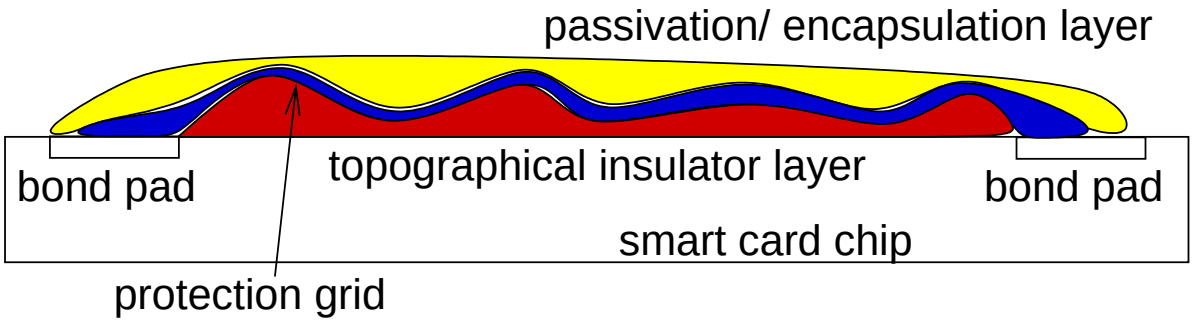

Figure 2.2: Side view protection grid with topographical features

ing at the boundary where the protection grid is located. The materials may also be chosen so that they preferentially dissolve in etchants that also damage the protection grid.

Similar to the IBM 4758 protection grid, multilayer protection structures may also be fabricated if an insulating layer (black) is deposited between the conductor layers (Figure 2.3). The two layers (blue and cyan) may be printed off set, so that there are no gaps between the grid lines through which a hole could be drilled into the packaging (omitted from drawing) without being detected. In Figure 2.3 the grid lines are drawn running in parallel, and they are connected to smaller, dedicated bond pads. This pattern is meant as an illustration rather than a definition of layout pattern. A multilayer structure may also consist of a single line wrapped back on itself, or broken into multiple segments. As the contact pads of a microchip are physically more robust, and allow better contact with probe needles than organic electronics, they are likely targets for an attacker to probe. Multilayer structures covering the contact pads may prevent successful probing if the higher layers are damaged by the probe needle.

2.2.2 Readout scheme

To be able to determine the integrity of the protection grid, repeatable measurements of grid characteristics must be possible. Generally, a measurement is carried out by quantifying the characteristic against a known (reference) value. As the security device may be operating in a hostile environment, it must be assumed that an attacker has control over all inputs to the device (e.g. clock frequency or supply voltage) as well as the ambient conditions (e.g. temperature). Therefore none of the inputs can be trusted as a measurement reference. Internal references may drift with supply voltage and temperature, thus are not automatically reliable. On-chip CMOS voltage references with minimal temperature drift [Giustolisi et al., 2003] may be used to guarantee a fixed supply voltage, and to detect too low a supply voltage. Alternatively, the temperature drift of on-chip reference and the protection grid characteristics may be compensated for, as it is possible to measure temperature with CMOS circuitry [Sanchez et al., 1997]. Differential measurement schemes may also be designed to avoid dependency on an external reference. Comparing parts of the protection grid with each other will cancel out variations due to the operating conditions. The differential schemes must be designed to detect all modes of damage, even if all branches are damaged by the same amount. Multilayer structures may be suitable for differential measurement schemes, as higher layers would generally suffer more damage than lower layers, as they would be exposed first.

The characteristic measurement value may be used in two ways. In the first method, the measurement may be used as part of the cryptography, following the particle distribution method proposed by Kömmerling [Koemmerling, 2000]. The cryptographic key may be derived from the properties of the protection layer, so that any change in properties makes the cryptographic key void. A schematic of this method is shown in Figure 2.4. For illustrative purposes the

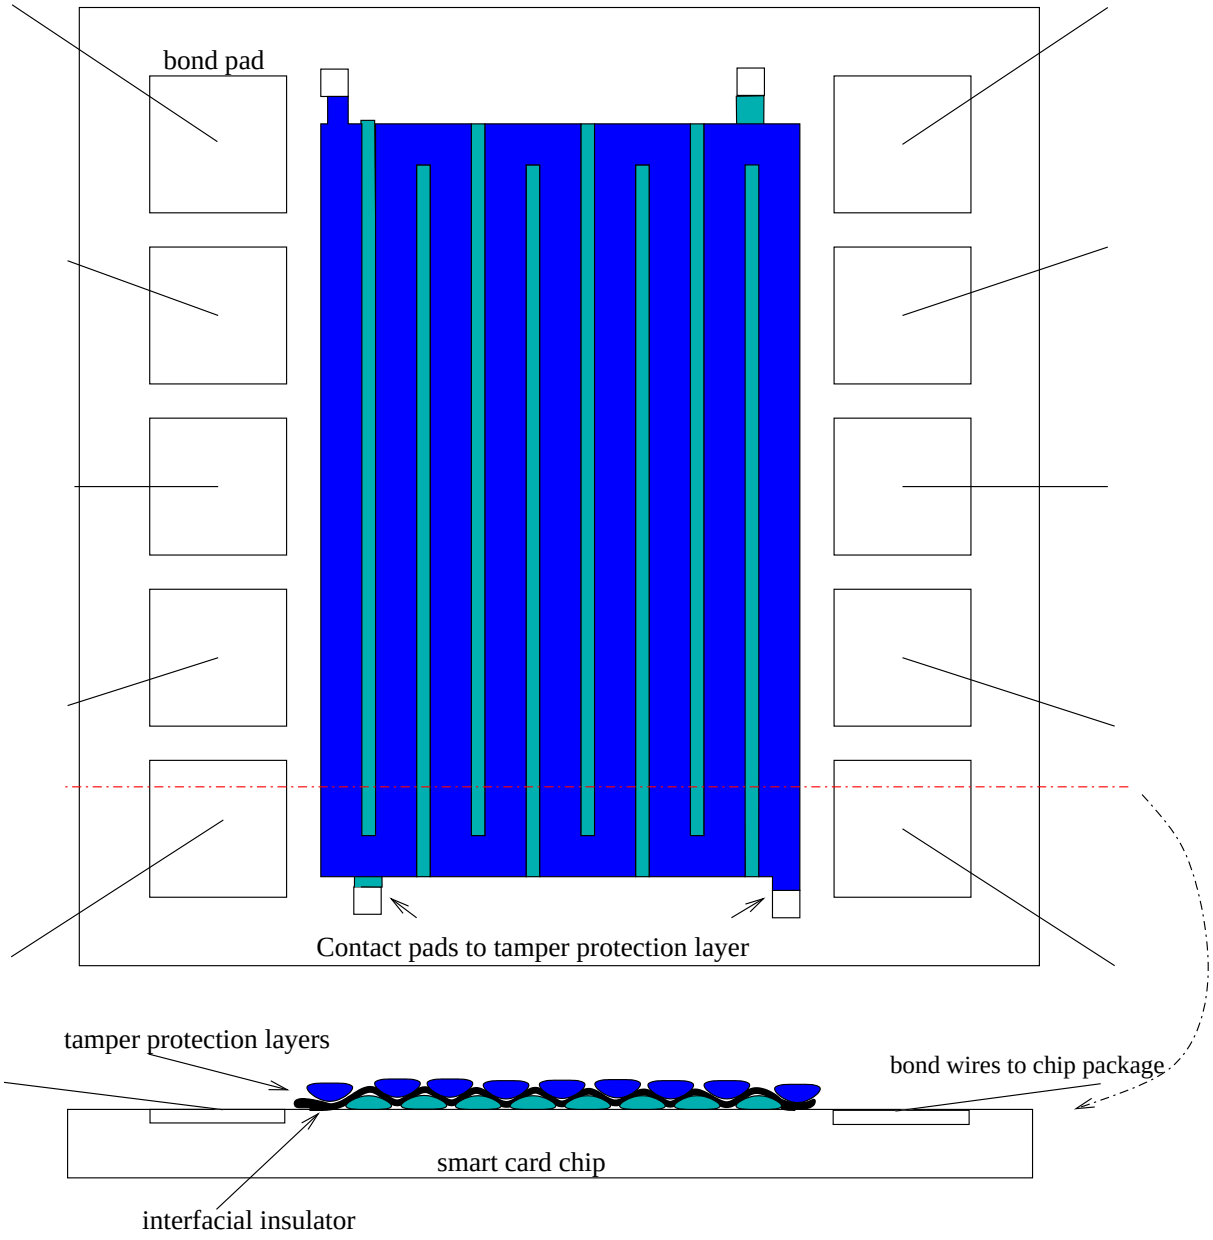

Figure 2.3: Top and side view of dual layer protection structure

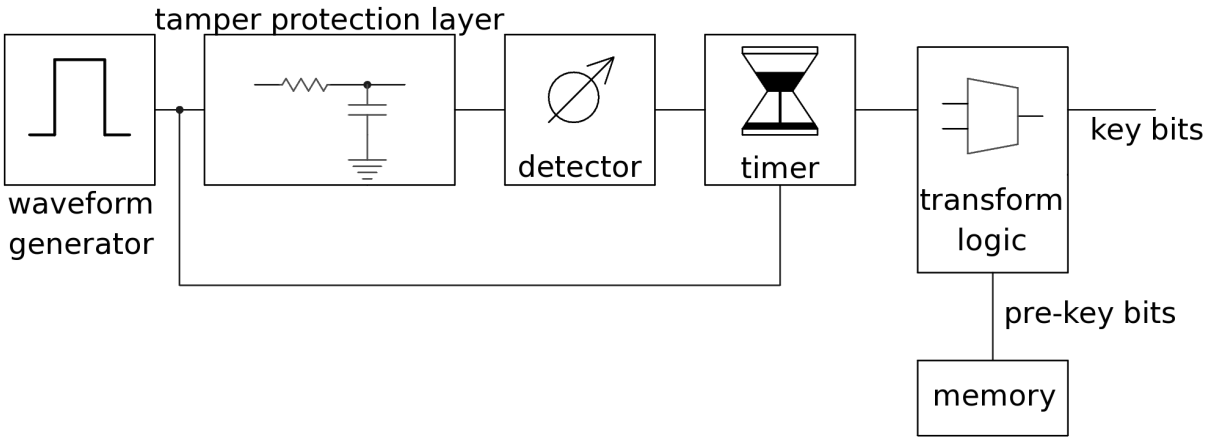

Figure 2.4: Block diagram of key generation principle

property of the tamper protection grid is shown as a resistance capacitance (RC) time delay, however other measurement methods may also be used. The cryptographic key may be derived directly from the signal by means of some transformation, e.g. using a hash function to distribute the entropy over all key bits. Without the transformation, the least significant bits show stronger variability compared to the more significant bits. However, if the resolution of the property measurement is too low, then the entropy of key values may be too small. The danger is that an attacker manipulates the protection grid properties to make the properties of the grid (and cryptographic key) match a second device, creating a clone of the device. To prevent this, the transformation function may also derive the key from a combination of the grid properties and (random) 'pre-key' bits which are taken from memory or derived from properties of the security device rather than the grid. However, a hash function also requires error correction before hashing, otherwise the cryptographic key cannot be recovered reliably. Alternatively, a low entropy of key values also opens the danger of a brute force key search.

In the second method, the random property may simply be compared to a reference value stored on the chip, as shown in Figure 2.5. As the properties of the protection layer vary (naturally or deliberately), the chip must initially perform a self-characterisation routine before first use to determine the reference value. This value could then be stored e.g. in non-volatile write-once memory directly on the chip [Paul, 2005]. Such a memory may consist of a pair of minimum-geometry transistors, one of which is damaged during data storage by subjecting it to a voltage beyond its safe operating envelope. A sense amplifier is used to detect which of the transistors is damaged [Paul, 2005]. If the properties of the grid and the reference values stored on the chip are mismatched, a self-characterisation phase may be triggered. Such a scheme in combination with write-once memory has the advantage that for the initial characterisation phase, the reference data will be stored correctly. For subsequent mismatches caused by tampering, the renewed data storage attempt will result in the write-once memory being corrupted and data will not be stored correctly [Paul, 2005]. This mechanism prevents device function, as the microchip will be stuck in an infinite loop comparing the measurement result with the value in memory, then attempting to store it, which yields corrupt data that does not match the measurement result.

The most suitable method of determining grid characteristics depends on the application and the design details of the security device. If certain features are already implemented in the security device, e.g. a reference oscillator, then a time delay read-out scheme is the most cost-efficient and obvious choice. Other constraints may be set by the overall device size, and availability of space for contact pads, as well as the CMOS fabrication technology. As the focus of

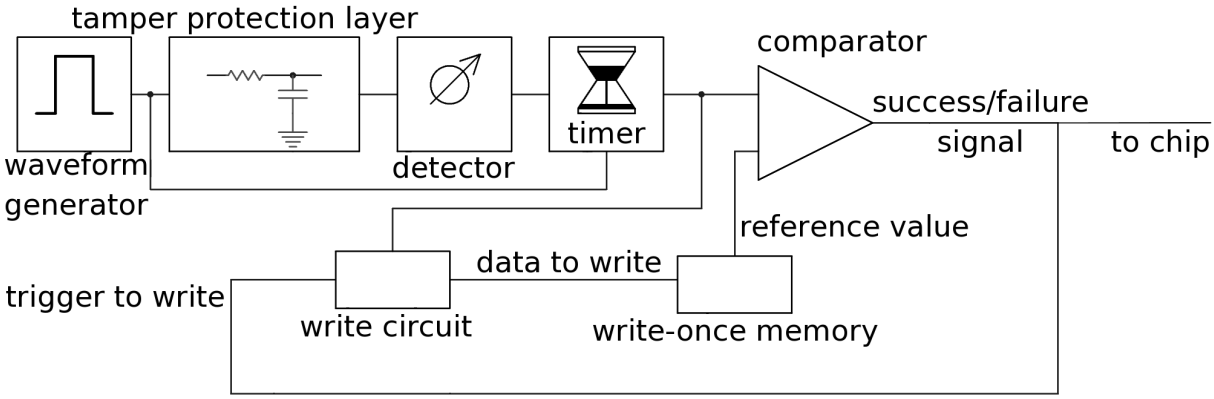

Figure 2.5: Block diagram of reference comparison principle

this dissertation is the the general evaluation of the suitability of organic electronics for tamper protection purposes, these schemes should only be considered as examples or suggestions.

2.3 Active protection grids

The inclusion of transistor circuits in the protection grid allows the mapping function between the contacts of the grid to be more complex than is achievable with resistive tracks. As these circuits may include amplification and switching, they are termed 'active protection grids'.

The current through a passive device depends on the voltage between the end points as well as the I-V-characteristics of the device. Passive devices are essentially combinations of two-terminal devices. Active circuits, in contrast, include multi-terminal devices such as organic transistors. As organic thin film transistors (OTFTs) are three-terminal devices, they can be combined such that the currents through the circuit depend on multiple voltages. As characterisation by an attacker requires more devices, the probability of damaging the grid in the process is increased. In addition, the layered structure of transistors compared to simple tracks allows more parameters to be varied to increase the range of possible unique grid characteristics. The larger range in circuit characteristics as well as the larger number of probes required for analysis makes characterisation and emulation more difficult for an attacker.

Organic TFTs may be fabricated either with the source and drain deposited on top of the semiconductor layer (top contact) or vice-versa (bottom contact) [Dimitrakopoulos and Mascaro, 2001]. If the semiconductor material covers the source and drain electrodes, direct probing of these contacts becomes more difficult to carry out without damaging the transistor. In order to probe the contacts, the semiconductor layer must be penetrated by the probe needle, which is likely to damage the transistor either physically, or by short-circuiting different contacts.

An active protection grid may also offer greater sensitivity to depackaging. While organic conductors have been reported to be environmentally stable [Elschner, 2002], only few environmentally stable organic semiconductors have been reported [Facchetti, 2007]. There also appears to be a tradeoff between chemical stability and mobility (performance) [Katz, 2004], though progress is being made to achieve both simultaneously [Tian et al., 2007]. The most important factors in environmental degradation of semiconductors are light, water, and air [Facchetti, 2007; Katz, 2004; Klauk, 2006]. Apart from chemical stress, organic semiconductors suffer from electrical stress, termed 'bias stress'. The transistor current decreases with prolonged transistor operation, which is a property of all disordered semiconductors, including amorphous silicon [Klauk, 2006]. Bias stress may exhibit a reversible and an irreversible component [Gomes et al.,

2006; Salleo et al., 2005; Sirringhaus, 2005]. Given these limitations, I conclude that transistor circuits may not be robust enough for application as a protection grid.

Multiple fabrication routes exist for organic semiconductors, such as thermal evaporation [Dimitrakopoulos and Malenfant, 2002], in-situ polymerisation [Chason et al., 2005], and dry thermal transfer [Katz, 2004], as well as various printing methods from solution [Klauk, 2006] including inkjet printing [Kawase et al., 2005]. The applicable fabrication method depends on the properties of the chosen semiconductors, such as solubility and thermal stability.

2.3.1 Circuit examples

In theory it is possible to create very complex circuits using (organic) transistors. However, the achievable complexity for active protection grids is limited in practice by the operating voltage, the switching speed, and the physical dimensions of the transistors.

The achievable transistor dimensions depend on several factors. The resolution of the printing method obviously sets a limit to what size can be fabricated reliably, in particular with respect to channel length and source/drain contact width [Burns et al., 2003; Sirringhaus et al., 2000]. The mobility in organic transistors is generally quite low. The best transistors achieve a mobility of the same order of magnitude as amorphous silicon [Facchetti, 2007], which is still three orders of magnitude less than typical single-crystalline silicon devices [Streetman and Banerjee, 2000]. To compensate for the lower mobility and achieve the required current levels, transistors need to be made wider. The contact resistance between the source/drain contacts and the semiconductor also reduces the current through the transistor. If the channel length is too short, the contact resistances dominate the transistor characteristics [Bürgi et al., 2003; Gundlach et al., 2006; Klauk et al., 2003; Lous et al., 1995; Luan and Neudeck, 1992; Necliudov et al., 2003; Street and Salleo, 2002; Zaumseil et al., 2003].

The reported operating voltages of organic transistors are of the order of 10s of volts [Chason et al., 2005; Dimitrakopoulos and Malenfant, 2002; Dimitrakopoulos and Mascaro, 2001; Katz, 2004], which is too high for typical CMOS chips without additions to extend the voltage range. Fortunately, progress is being made to lower the operating voltages of OTFTs [Liu et al., 2005]. As the traditional smart card operating voltage is around 5 V [Rankl and Effing, 2004], it may plausibly be assumed that an operating voltage for the protection grid below ca. 10 V is feasible.

As transistor circuits are suitable for both analogue and digital circuits, the protection grid could in theory also consist of logic gates or memory cells storing a unique key. The digital logic would need to be designed such that changes in transistor characteristics are detectable, even if the logical mapping is not altered. For example, changes in transistor characteristics impact the switching speed, which could be used to indicate damage to the protection grid.

Apart from monitoring transistor circuits for damage, it is also possible to implement chemical sensors using organic semiconductors [Weingart, 2000]. The sensors may be used to either sense a property of or chemical additive to the packaging (or a separate layer between packaging and sensor), or a property of the environment that is not present when the packaging is in place. If the chemical added to the packaging is sufficiently volatile, the molecule will evaporate when the packaging and encapsulation are broken, changing the sensor response. This type of grid is particularly useful against FIB and SEM attacks, as these typically require vacuum conditions that increase evaporation rates. However, the volatile material will also diffuse through the packaging. The diffusion rate of the material through the packaging must be low enough to support sufficient lifetime of the device.

Suitable circuits for active protection grids should allow detection of changes in transistor characteristics to ensure detection of tampering. To work around the size constraints, they would currently consist of only a few transistors until smaller organic devices are available. Given the typical high operating voltages of organic transistors, it may be required to use cir-

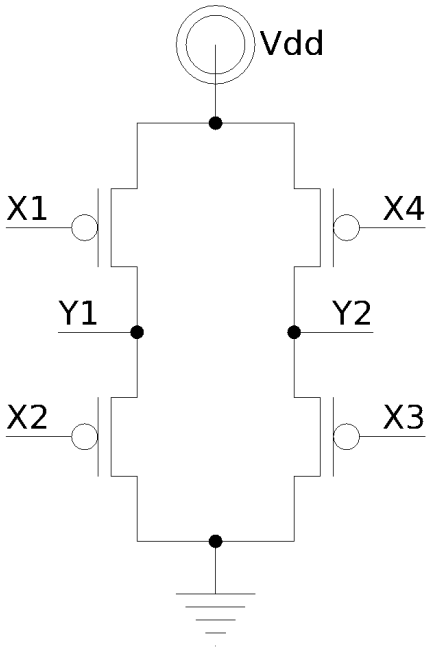

Figure 2.6: Schematic of four-transistor bridge

cuits that operate sub-threshold, i.e. below the threshold voltage necessary to turn the devices fully on.

Two circuits were found to be suitable for implementation of protection grids using current organic transistor technology. Both circuits are sufficiently simple, and are sensitive to changes in transistor characteristics. However, organic transistors may not be stable enough to achieve the required lifetime, which is estimated at several years.

2.3.1.1 Bridge implementation

An example of a simple active circuit is a four-transistor bridge configuration as shown in Figure 2.6. The gates of the four transistors (X1–X4) as well as the power and ground rails are the inputs to the circuit. The voltages of the two centre nodes (Y1, Y2) are used as a measure of integrity. If the transistor characteristics can be varied, the voltages of the centre nodes Y1 and Y2 will differ for the same input voltages, achieving a unique mapping between inputs and outputs for individual security devices.

The voltage at the centre node of each branch (Y1, Y2) is determined by the voltages at the transistor gates, and the supply voltage(s). The source of one of the two series-connected transistors is connected to a fixed voltage. For p-type transistors, and a positive supply voltage, the source of the top transistor is connected to the power supply. For n-type devices the source of the lower transistor would be connected to ground. If the (ideal) transistor is in saturation, the current through the channel is determined by the gate-source voltage ($V_{gs,1}$). The source of the second transistor is tied to the centre node, and the drain of the first transistor. As the voltage at that node is not fixed, it will adjust to set $V_{gs,2}$ for the second transistor such that the currents through both transistors are equal, as required by Kirchhoff's Current Law. Using the ideal metal oxide semiconductor field-effect transistor (MOSFET) equation for the saturation

region, and assuming equal threshold voltages V_T:

$$c_1 \quad (V(X1) \quad V_{dd} \quad V_T) \quad = \quad c_2 \quad (V(X2) \quad V(Y1) \quad V_T) \tag{2.1}$$

$$V(Y1) \quad = \quad V(X2) \quad \frac{c_1}{c_2} \quad (V(X1) \quad V_{dd}) + (1 \quad \frac{c_1}{c_2}) \quad V_T \tag{2.2}$$

For non-ideal transistors (many OTFTs) and for transistors in the linear operating region, the source-drain current is also determined by the source-drain voltage (V_{ds}), altering the simple relationship. In addition, it is expected that some leakage current is present in fabricated transistors

The on/off ratio of a transistor is defined as the ratio of minimum to maximum transistor currents for a given operating voltage range, typically $V_{gs} = 0$ V and $V_{gs} = V_{supply}$ [Katz, 1997; Kelsall et al., 2005; Schwoerer and Wolf, 2007]. The minimum current is the leakage current while the maximum transistor current is related to the constant of proportionality '*c*' in the equations above. The larger the on/off ratio of a transistor is, the larger is the dynamic range of the transistor current and the more stable the voltage at Y1/2 with respect to variations in the leakage currents. With a low on/off ratio of the transistors, it may be the case that a higher leakage current of one (badly fabricated) transistor exceeds the on-state current of the second transistor at the target voltages. Small transistor currents also make the bridge circuit more vulnerable to noise.

Similar to passive protection grids, the characteristics of the circuit could be compared to reference values stored on the chip, or used as a seed value to generate the cryptographic key. For the bridge implementation, the (absolute) voltages of the centre nodes for given input voltages (X1–X4) could be compared to reference values stored on chip. Alternatively, the input voltages when the bridge is balanced (voltages of Y1 and Y2 equal) could also be used as a metric for integrity. If the characteristics of individual transistors can be changed (variation in the constants of proportionality '*c*', the threshold voltage V_T and the leakage current), it is also possible to generate a unique sequence of balance points for different input voltages. This would enable a vector of data to be extracted from a single circuit rather than a single datum.

Using the balance point of the bridge as a measure of integrity, it is important to ensure that an attacker cannot simply short circuit the nodes together or apply a fixed voltage to both centre nodes. To prevent this simplistic attack, the bridge should be verified to go through an unbalanced phase before reaching the balance condition. As an example, randomised waveforms could be applied to the transistor gates (Figure 2.7). The input waveforms would reach the appropriate conditions for balancing the bridge from time to time, to verify the integrity of the circuit.

The advantage of a bridge circuit is that it also functions if transistors are not fully conducting in the sub-threshold region. However, with the reduced on/off current ratio it is also expected that the sensitivity of the circuit will be reduced.

2.3.1.2 Ring oscillator implementation

Ring oscillators are frequently added to CMOS microchip designs as test structures to monitor the variations of transistor characteristics for each fabrication run of a chip [Weste and Harris, 2004]. A ring oscillator consists of an odd number of inverters connected in series to form a ring. The frequency of the ring oscillator varies with layout geometry and transistor (semiconductor) characteristics. As the outputs of the ring are digital and the measured (analogue) characteristic is the frequency, a ring oscillator is a very convenient way to avoid precision analogue circuits for protection grid characterisation. Ring oscillators have been successfully implemented in

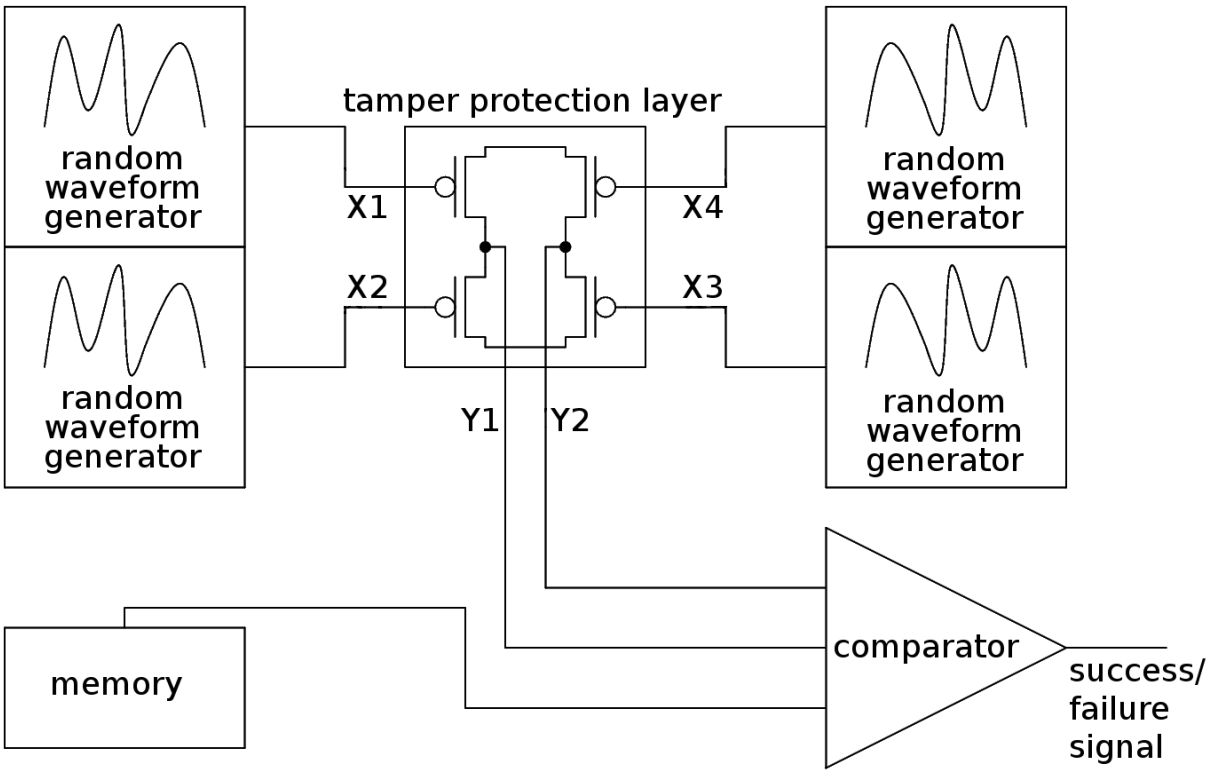

Figure 2.7: Block diagram of detection scheme for four-transistor bridge

organic transistor technology, therefore they may be suitable for protection grids [Baude et al., 2003].

When used as a protection structure, a timer would be used to determine the delay between two nodes of the ring oscillator, as shown in Figure 2.8. Altering transistor properties changes the frequency of the ring oscillator to give the protection grid a unique property. The transistor geometry, semiconductor properties, and number of inverter stages could all be varied for each individual security device, so long as the deposition patterns of the organic materials can be varied dynamically. To determine the timing characteristics using the CMOS security device, a stable on-chip timer must be implemented. As the switching speed of an inverter may change with temperature and supply voltage, it must be ensured that the supply voltage and environmental conditions are within specified bounds, using suitable references and sensors [Giustolisi et al., 2003; Sanchez et al., 1997].

2.3.1.3 Embedded gates

Some reported organic TFTs are deposited on a silicon wafer covered by a thin silicon dioxide insulator. The silicon wafer acts as gate electrode, and the oxide as gate insulator [Dimitrakopoulos and Malenfant, 2002; Dimitrakopoulos and Mascaro, 2001; Facchetti et al., 2005; Halik et al., 2003; Kymissis et al., 2001; Necliudov et al., 2000; Wang et al., 2005]. This configuration is known as the 'bottom gate' configuration, as the gate is the lowest layer. An adaptation of this design may be implemented in active protection grids, if the top metal layer of the security device is used as a gate and the chip passivation as gate dielectric. The semiconductor material and source and drain contacts would be deposited directly on the chip surface. The advantage of this transistor structure is the simplified structure of the protection grid achieved by eliminating separate gate electrodes and gate dielectric. As the gates do not require contact

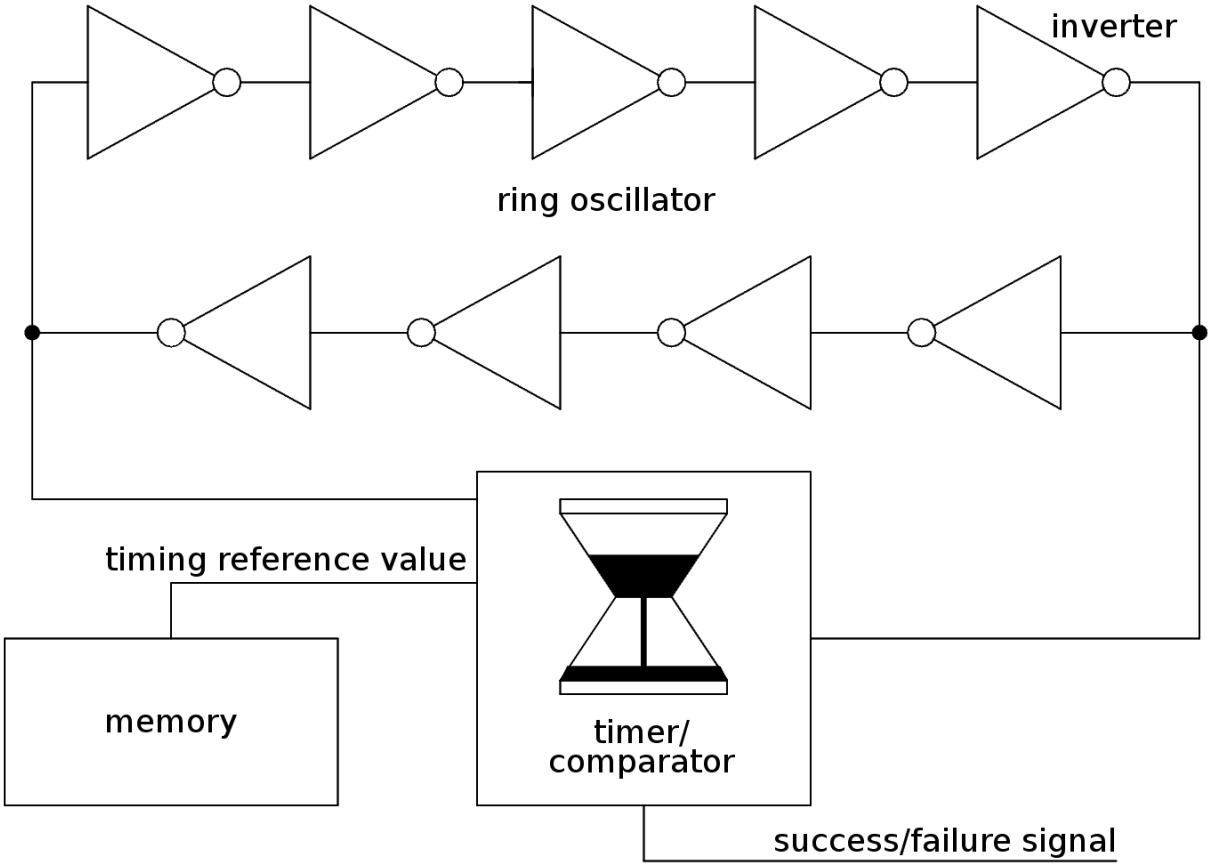

Figure 2.8: Function of ring oscillator embodiment

pads and pad drivers, I estimate that the consumed area on the chip is comparable. A second advantage of this scheme is the increased difficulty to probe the gate voltages of the transistors below the robust silicon nitride passivation of the microchip without damaging the protection grid. The disadvantages of this scheme are the reduced layout flexibility (fixed positions of the transistors) compared to all-polymer TFTs, and the position of the protection grid directly on top of the chip, rather than on a topographical interface layer.

2.4 Evaluation procedure

The evaluation of organic protection grids will focus on passive protection grids, as organic conductors are reported to be chemically stable and are thus most likely to meet the requirements. The stability of organic transistors is lacking to date. Therefore only simple transistor circuits are evaluated in place of full active protection grids. The measurements may be useful to derive requirements for transistor performance, should stable transistors become available.

For organic electronic materials to be feasible as a protection grid, requirements in terms of fabrication, security, and lifetime must be met. The passive protection grids are evaluated in three stages. First, a suitable organic conductor is chosen, and a reliable fabrication method is developed for the proposed structures. Once it is possible to make protection grids reliably, the security properties of the grids are evaluated. Independent of which readout scheme is chosen, tampering is detected by changes in the electrical properties of the protection grid lines. To complete the evaluation, the robustness of protection grids to environmental influences and ageing must be evaluated. A prototype grid for both passive and active protection grids is also assembled to test the properties and sensitivity readout schemes.

Chapter Three

Fabrication of Tamper Protection Grids

3.1 Introduction

While the currently used metal protection grids are simple to implement and do not require sophisticated sensors, they have the problem that they are encapsulated under an inert and robust encapsulation layer. They easily survive even aggressive depackaging methods [Anderson and Kuhn, 1996]. An alternative system to the metal grids is implemented in the IBM 4758 secure coprocessor module [Dyer et al., 2001]. It consists of a conductive grid on a polyurethane membrane surrounding the module, which is easily damaged in a depackaging attempt. While it allows the module to detect tampering, it requires a battery for monitoring the integrity of the membrane. Due to the cost and size constraints of small scale security devices, this concept is not feasible to implement on a microchip. So far, such simple yet responsive protection grid systems for microchips are lacking, prompting the proposal of organic electronics for this application.

Organic electronics have attractive properties for protection grids. These materials can be produced cheaply [Heeger, 2001], are stable [Heeger, 2001], and are suitable for low-cost additive fabrication methods which minimise waste [Burns et al., 2003; Chen et al., 2003]. Additive fabrication methods contrast with lithography, the standard method for patterning in CMOS device fabrication, where a large percentage of consumed materials are wasted [Franssila, 2004]. The grids should be given unique properties which are likely to be altered if the packaging of the device is tampered with (which is lacking in the IBM 4758). If the properties of the grid can be verified by the security device, the device will not have to constantly monitor the grid for anormalities. Verifying the grid before carrying out sensitive operations will be sufficient to protect the device against invasive attacks that require the chip to be powered on. This intermittent testing scheme thus saves the expense of a battery, which is critical for a small scale security device.

In this chapter, the fabrication of passive protection grids will be explored. The requirements for reliable fabrication of protection grids are predictable properties, a low defect rate, a low cost overhead, as well as dynamic pattern variation to ensure unique properties. The electrical properties should fall in a range that can be measured by a standard CMOS microchip. The required conductivity range depends on the chosen integrity metric (RC time delay, absolute resistance, or resistance ratios) and circuit for readout. While the choice of integrity metric and circuit are part of the design of the security device, a current of at least several μA at typ-

ical operating voltages will be required to carry out stable measurements. This current range corresponds to less than ca. 1 MΩ of resistance. In addition, good electrical properties require stable, ohmic contacts to the aluminium contact pads of the security device.

One of the most stable organic conductors is poly(3,4-ethylenedioxythiophene) (PEDOT) doped with poly(styrenesulfonate) (PSS) [Bantikassegn and Inganäs, 1997; Groenendaal et al., 2000; Heeger, 2001], therefore it was used to evaluate passive organic protection grids. It is commercially available as an aqueous dispersion from H.C. Starck [Starck, c]. For brevity, the PEDOT:PSS dispersion will be referred to as PEDOT in this dissertation. Even though other dopants exist for PEDOT [Bantikassegn and Inganäs, 1997], PSS is the most common dopant used in organic electronics [Groenendaal et al., 2000].

A versatile patterning method that is compatible with PEDOT is (drop on demand) inkjet printing [Burns et al., 2003; Chen et al., 2003]. The advantage of drop on demand inkjet printing is the low cost, and flexibility of deposition pattern. The low fabrication cost of inkjet printing compared to lithographic patterning of metals is partly due to the faster turnaround time (no vacuum necessary). In addition, only the required pattern is deposited, therefore no material is wasted. Inkjet compatible self-aligned patterning techniques have also been developed to yield gap sizes between adjacent lines of sub 100 nanometer dimension [Sele et al., 2005]. Another advantage of inkjet printing is the compatibility with organic transistor fabrication [Burns et al., 2003; de Gans et al., 2003; Kawase et al., 2005; Lim et al., 2006; Liu et al., 2005; Speakman et al., 2001; Stutzmann et al., 2003], which is required for active protection grids.

To develop a fabrication method for organic protection grids, several issues must be addressed. First, the parameters of the inkjet printing process need to be adjusted to achieve good print quality. In the next stage, the composition of the PEDOT ink and the conductivity range are evaluated. For a low defect rate, printing on top of a standard microchip and a selection of other substrate materials must be reliable. The electrical connection to the aluminium bond pads of standard CMOS microchips is examined and stabilised.

3.2 Printing setup

3.2.1 Printer

The printer used for fabrication is a custom-built inkjet printer located in a normal laboratory rather than a clean room. A detail photo of the print stage is shown in Figure 3.1. The print head is taken from an Epson Stylus Color II printer, and uses piezo technology to eject ink droplets [Epson]. The print head mechanically ejects ink drops, by means of electrically induced deformation of a piezo crystal. In comparison, a thermal inkjet print head boils the ink to eject the drops [Le, 1998]. The advantage of the mechanical ejection method is a larger range of printable inks, as there is no requirement on temperature stability and boiling point.

The print head is held fixed during printing while the sample is moved under the print head with the X–Y-stage. The Z-stage may be used to adjust the gap between print head and sample. The standard gap is ca. 1–2 mm to prevent the print head touching the sample, but small enough not to impact print quality. Ink is fed from the reservoir vial to the print head via a plastic tube. Drying time may be shortened by wafting the sample with a slow flow of dry nitrogen gas. High gas flow rates displace the ink droplets on the substrate, and thus degrade the print quality. There is a camera mounted below the X-Y stage for monitoring the printing process. The print head is controlled via a custom LabView program on a PC (not shown).

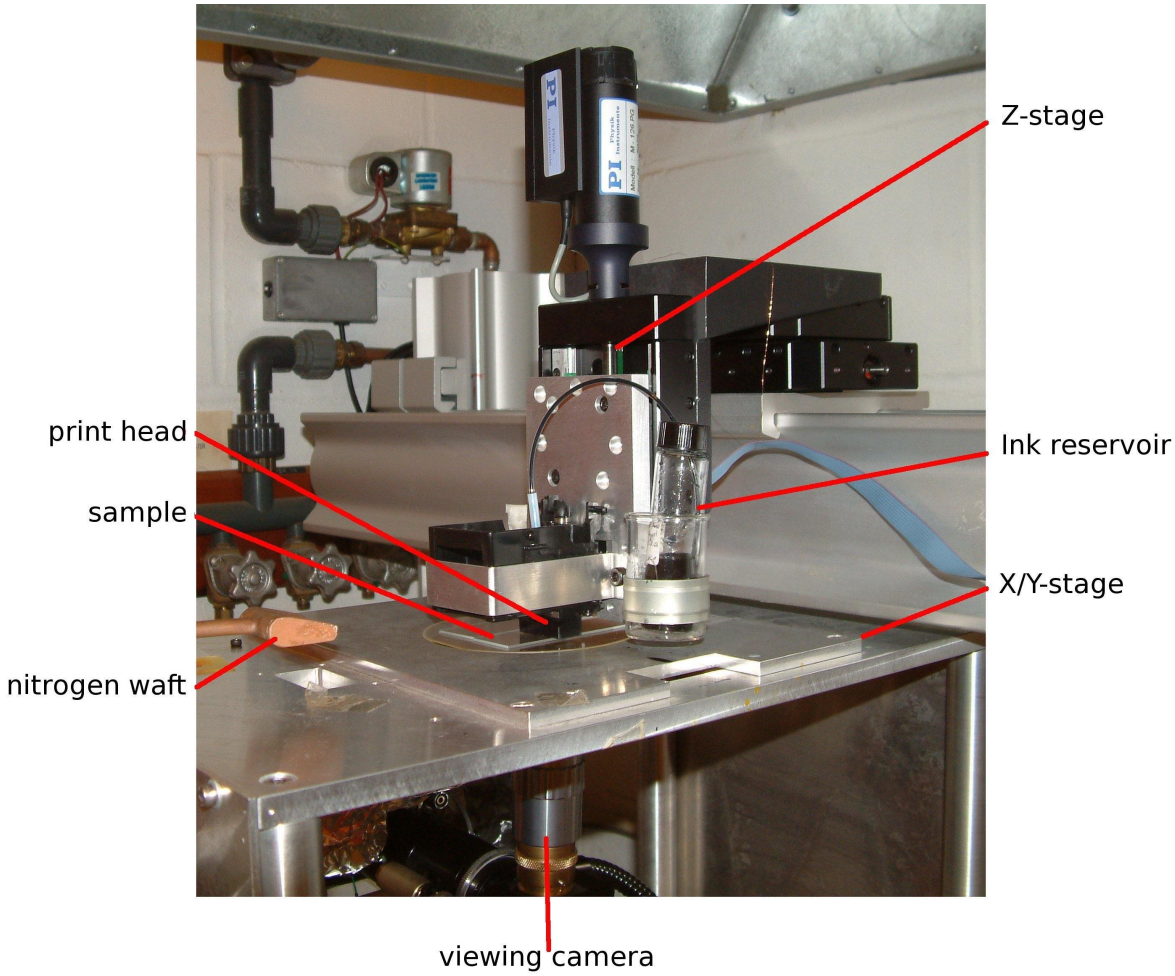

Figure 3.1: Inkjet printer used for fabrication

3.2.2 Printing preparation and procedure

Before running the printer, the print head must be rinsed thoroughly with de-ionised (DI) water. This removes old ink deposits and clears clogged nozzles. It is advisable to filter cleaning water and ink before use in the printer to remove clotted ink particles and dust. By running the printer with the rinsing water, it can be checked whether the print head has enough functioning nozzles available before loading the PEDOT ink. After the ink is loaded, drops of ink should be ejected continuously from the nozzles (idling mode) to prevent the ink drying and clogging the nozzles. The print head can be moved away from the samples for maintenance and during idling periods to prevent soiling of the sample. One or more functioning nozzles should be selected for printing, as well as the desired layout pattern.

The sample should be cleaned (at least in acetone, isopropanol (IPA) and DI water) before use, and placed alongside a blank/idling slide. For best printing results, one or two repetitions of the structure should be printed on a blank slide before printing the pattern on the sample slide. This allows the motion stages and air flow to reach their steady state during printing. The blank slide may also be helpful for adjusting the position of the monitoring camera to observe the printing process, without carrying out a print run.

3.3 Printing method and ink composition

Print quality and repeatability depend on several parameters. Before landing on the substrate, droplets may be deflected by statistical variations in the ejection angle from the nozzle, and also by the air flow between print head and substrate. Air motion is caused by a thin boundary layer of air adhering to the moving sample [Schlichting and Gersten, 2000], as well as the exhaust fan above the printer and the nitrogen waft (if used). The achievable positional accuracy of the ink droplets may also be limited by the motion and spreading of droplets on the substrate [Sirringhaus et al., 2000]. A smaller print gap reduces the impact of variations in ejection angle and air flow.

In addition to the positional accuracy of droplets, lower polymer concentrations in solutions influence the formation of satellite drops, which may degrade the printing results [de Gans et al., 2003]. Higher concentrations of polymer in the ink result in higher viscosity of the ink and faster drying, which is beneficial for print quality, but may lead to nozzle clogging if the ink in the nozzle dries between the ejection of droplets. If the viscosity of the ink is too high, the friction resisting flow through the nozzle capillary is larger than the force applied to the fluid by the piezo crystal, preventing drop ejection altogether.

To control the viscosity and drying properties, the aqueous PEDOT dispersion may be diluted with DI water. The PEDOT lines may also be made from several layers to reduce the effect of individual mis-placed drops. If the conductivity of PEDOT is found to be too low for a certain grid design, additives may be required to increase conductivity of the PEDOT ink. Dimethyl sulfoxide (DMSO) is a solvent recommended by the supplier for this purpose [Starck, a].

3.3.1 Sample space

For investigation of print quality and conductivity characteristics, a range of ink compositions and different print settings were trialled. To control viscosity, three dilutions of PEDOT with DI water were compared: neat PEDOT as supplied by H.C. Starck, 1:1 dilution with water and 1:2 dilution with water. For conductivity enhancement, the following compositions of PEDOT, water and DMSO were tested: 1:2:0.1 ,1:2:0.2 , 1:2:0.5, 1:2:1.

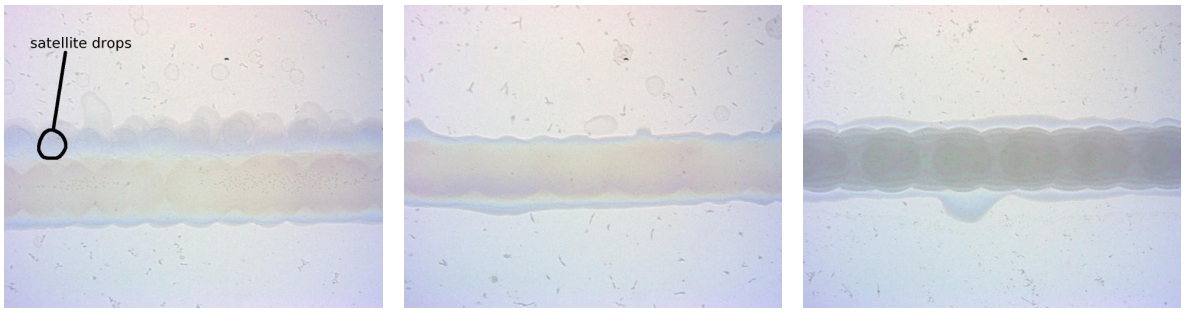

(a) 1:2 PEDOT:water diluted ink (b) 1:1 PEDOT:water diluted ink (c) neat PEDOT

Figure 3.2: Formation of satellite drops degrading print quality

3.3.2 Print quality

3.3.2.1 Drop size

The drop size as printed on glass substrate was measured to determine the default line width and suitable drop spacing. Sample geometry was measured using an Olympus BX-60 microscope with a magnification range of 50 –1000 and a calibrated camera and measurement software package. The drop diameter was measured to lie between 70 µm and 80 µm. The diameter of the visible satellite droplets for the 2:1 diluted PEDOT (Figure 3.2(a)) was measured to be 33 µm. For good overlap of droplets for a continuous line, a drop spacing of 40 µm was chosen.

3.3.2.2 Multi-pass printing

To form a continuous line, individual drops of ink are printed to overlap. Surface tension forces liquids to minimise their surface area [de Gennes et al., 2004], which on a flat surface is a dome shape. If overlapping drops are printed, the surface area is not minimal, therefore surface tension will exert a force to merge multiple drops to a larger single drop, obviously breaking the intended pattern. The extent of merging depends on the surface wetting properties, which determines how mobile freshly deposited ink drops are, as well as the drying time of the ink.

To avoid printing overlapping wet drops, the printer software was modified to allow fine grain control of the printing process. Apart from choosing the desired pattern, parameters such as drop spacing and number of passes per layer can be controlled. When printing lines in multiple passes (e.g. n passes), only every n^{th} drop of a pattern is printed in each pass. After a pass, the deposited ink is allowed to dry. In the next pass, the group of drops which has the furthest distance to the previously printed drops is printed. For the second pass this would correspond to around $\frac{n}{2}$. For example, in a four-pass print, drops 1,5,9…would be printed in the first pass, then 3,7,11…, and in the last two passes drops 2,6,10…and 4,8,12…

Figure 3.3 shows a microscope photo of printing 1:2 diluted PEDOT on glass, at a target drop spacing of 40 µm. The single pass print at the top of Figure 3.3 shows an irregular print result, as a result of fluid motion after deposition. In the two pass print, the drops printed in the same pass theoretically have a drop spacing of 80 µm (second from top), matching the previously measured drop diameter. This should leave individual drops with just enough space not to be in contact with each other. However, the positional and size variations of ink deposition and temporary expansion at impact are too large to prevent the wet drops from overlapping and merging. In contrast to the first two lines, the results for three (spacing 120 µm) and more passes are substantially better, showing no evidence of merging, and in turn a more regular print result.

The conclusion from Figure 3.3 is that the choice of number of passes is arbitrary above three for this specific drop spacing and diameter. More generally, the number of print passes should be chosen to result in a drop spacing larger than the drop diameter plus the positional variation of deposition. Four-pass printing is most symmetrical, in the sense that each drop reaching the surface will either have two or no neighbouring drops when reaching the surface. For some ink compositions or substrate materials it may be possible that the dry ink on the substrate has an influence on the surface area and energy of a freshly deposited drop, resulting in displacement of the drops. The symmetry of either having no or both neighbours helps to keep motion or spreading of drops equal in both directions, thus ensuring continuity of the printed line.

3.3.2.3 Effect of PEDOT dilution on print quality

The effect of polymer concentration and viscosity on the formation of satellite droplets predicted in literature [de Gans et al., 2003] is clearly visible for diluted PEDOT. Figure 3.2 shows photos of PEDOT lines printed under identical conditions, but with different grades of dilution. Figure 3.2(a) shows the satellite droplets resulting from 1:2 PEDOT:water dilution at the top of the line. It can be seen that these drops are significantly smaller than the main PEDOT drops forming the line (ca. 33 µm diameter). As the satellite droplets are lighter than the main drops and travel at a slower velocity [de Gans et al., 2003], the air flowing perpendicular to the direction of printing caused by the fume extraction pushed the drops next to the main PEDOT line. This effect increases the minimum distance lines can be printed next to each other without short circuiting. For higher polymer concentrations, no satellite droplets are visible next to the PEDOT line. The printing result for 1:1 diluted PEDOT, shown in Figure 3.2(b), is free from satellite droplets. Increasing the perpendicular air flow did not cause satellite droplets to appear in 1:1 diluted ink. Neat PEDOT ink with the highest polymer concentration also showed no satellite droplets on the substrate, as can be seen in Figure 3.2(c). However, neat PEDOT dries relatively quickly, clogging nozzles even during a print run.

3.3.2.4 Drying time

The drying time of the ink has an influence on the time that drops are mobile on a substrate, and potentially merge or travel away from their intended position. As the ink on the substrate must be allowed to dry between print passes and layers, the drying time of the ink also has a significant influence on the total printing time. Depending on the length of pause between print passes, the drying time also determines the potential for nozzle failure due to ink drying in the print head.

For pure diluted PEDOT the drying time was found to be of the order of several seconds. Print passes can be run continuously, with the setup time of the motion stage and print pattern data between passes being sufficient to allow drying of the most recently deposited ink drops. When DMSO was added to the PEDOT, the additional drying time significantly increased between passes (added to the setup times) beyond 20 seconds. The likely cause for this is the significantly higher boiling point of DMSO at 189 °C [Aldrich, accessed 2^{nd} Sept. 2008] which reduces the evaporation rate.

As the relative humidity of the surrounding atmosphere determines the evaporation rate of the droplets [Hu and Larson, 2002], drying time can be reduced by reducing the humidity around the sample. It was found that a very weak waft of nitrogen could reduce the drying time significantly. In order to minimise impact on print quality, the flow of nitrogen gas was dispersed around the sample to form a cloud of dry gas around the sample. Introducing the nitrogen drying aid reduced the required additional drying time back down to ca. 2 to 4 seconds depending on line length.

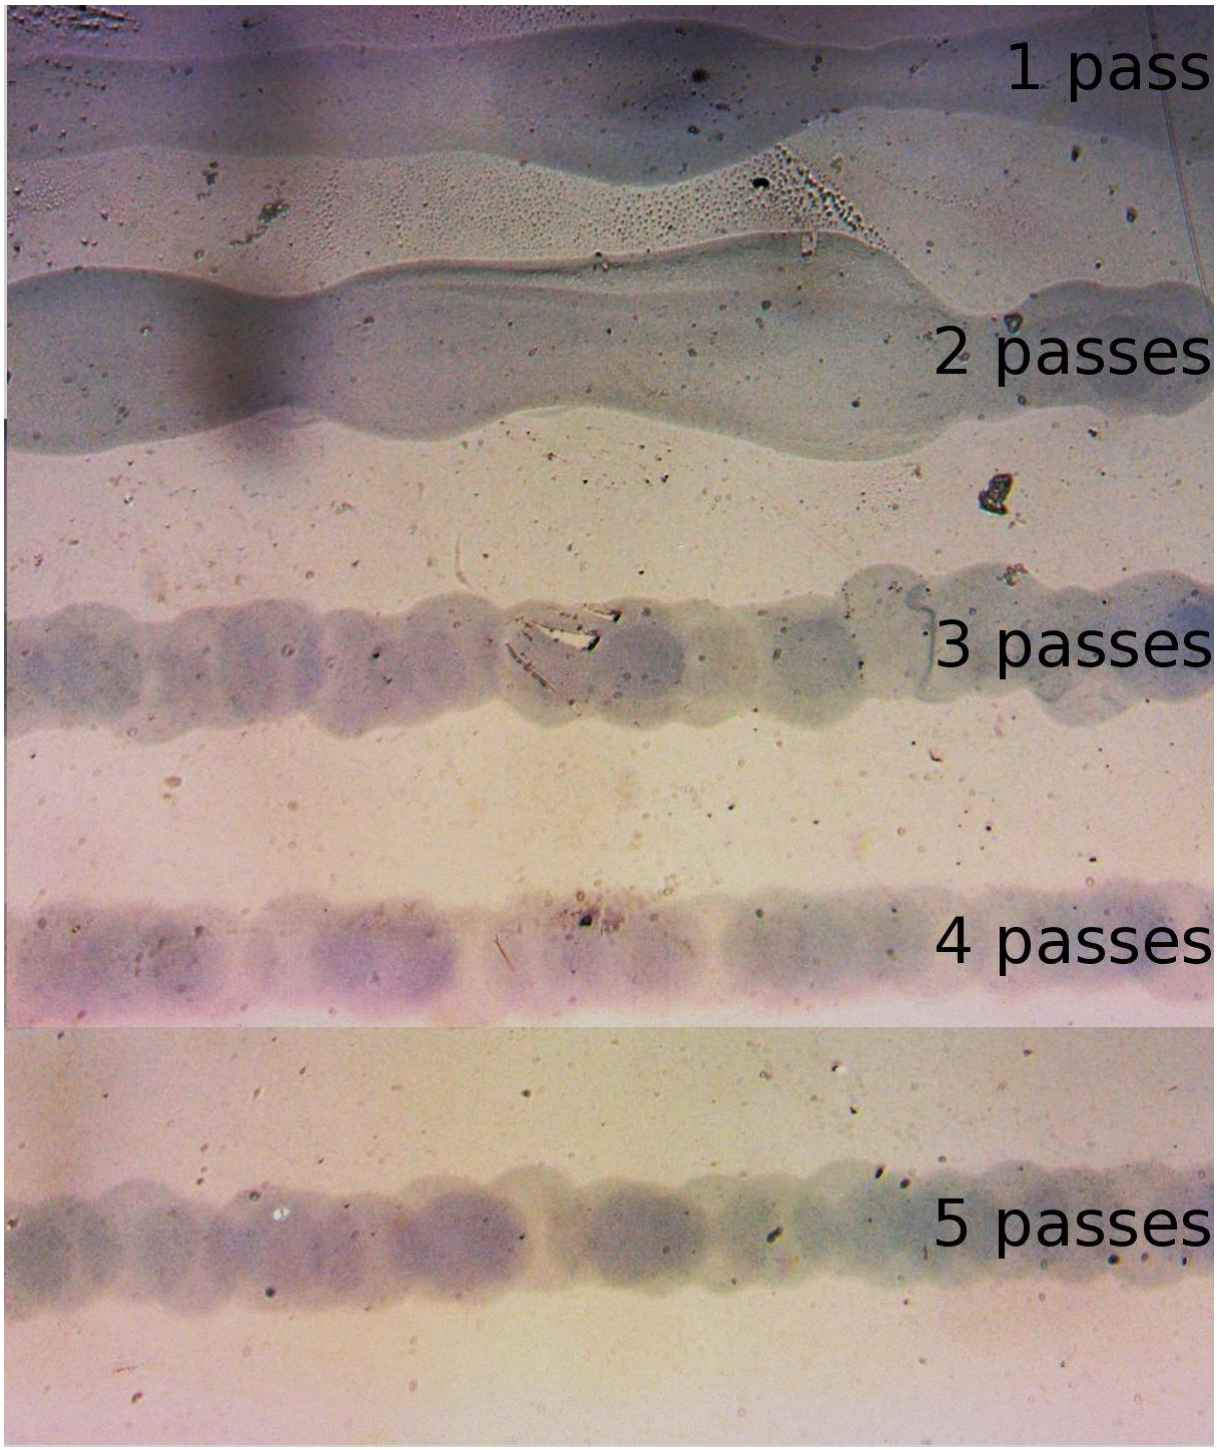

Figure 3.3: Printing of 1:2 PEDOT:Water in 1-5 passes

3.3.3 Electrical results

3.3.3.1 Electrical measurement setup and procedure

To determine the electrical characteristics, 1600 µm-long, resistive PEDOT lines were printed on clean glass microscope slides. The samples were transferred to a probe station where the lines were probed with tungsten probes. The I-V measurement was carried out as a four-point measurement with a HP4156A semiconductor parameter analyser. Four probe needles were brought into contact with the PEDOT line directly without contact pads. The voltage range chosen for the measurements was -5 V to 5 V, corresponding to the traditional smart card operating voltage [Rankl and Effing, 2004]. The measurements were carried out as a dual sweep, i.e. the voltage was scanned from -5 V to 5 V and back again in steps of 0.1 V. The dual sweep is used to identify problems such as bad, non-ohmic contacts as well as degradation or other changes in the sample, which typically result in hysteresis between the two directions of sweep. Measurements were carried out on ten samples each per configuration.

It was found that the thin layer of PEDOT offers little mechanical stability when in physical contact with tungsten probe needles. Good contact properties require a minimum cross-sectional area, which in turn requires the probe needles to penetrate the material rather than simply touch the surface. When lowering probes onto the substrate, the bending of the needle and the (non-orthogonal) angle of lowering move the sample, breaking the PEDOT line at the points where the other probes are already in contact with the sample. While the low physical robustness may be good in a tamper-proofing context, it makes lowering four probes onto the same resistor time-consuming to achieve. It was also found that PEDOT debris collecting on the probe needles introduced additional contact noise, therefore the debris had to be removed frequently. Comparing measurements on the same PEDOT line before and after cleaning the probes with a moist, lint-free wipe confirmed that debris degrades the electrical contact quality.

As PEDOT has a comparatively high resistance, it had to be evaluated to see if light had an influence on conductivity measurements. The fewer free charge carriers that are available in a material, the more the carriers generated by the impact of photons are relevant for electrical conduction. This effect is especially visible in semiconductors. If the energy transmitted by a photon is roughly equal or slightly larger than the band gap energy of a semiconductor, $h\nu \quad E_g$, then the photon may generate an electron-hole pair on impact. This increases the number of free carriers that can participate in conduction [Streetman and Banerjee, 2000]. The light senstivity in turn skews conductivity measurements.

The low physical robustness of PEDOT created problems when closing the probe station to carry out measurements in darkness. The unavoidable vibrations caused by the cover of the probe station were found to cause the probe needles to break the lines and contacts. As I-V measurements in light and dark conditions were found to be identical, measurements were carried out without closing the probe station cover to avoid breakage of the PEDOT lines.

3.3.3.2 Resistance of diluted PEDOT

The impact of DI water dilution on PEDOT conductivity was evaluated. The resistance per unit length (and layer) was determined for the three concentrations chosen earlier. As the drops of diluted PEDOT ink contain less conducting material per unit volume compared to the undiluted ink, the resistivity is expected to increase. The line width for all concentrations was measured to be nearly constant (ca. 75 µm for neat PEDOT, and ca. 79 µm for the diluted PEDOT, similar to the drop diameter). Therefore it is expected that the resistivity of the 1:1 diluted PEDOT is double the value of neat PEDOT, consistent with halving the amount of conducting material in the line same width line. For a 1:2 dilution, the solid PEDOT contents should be one third that

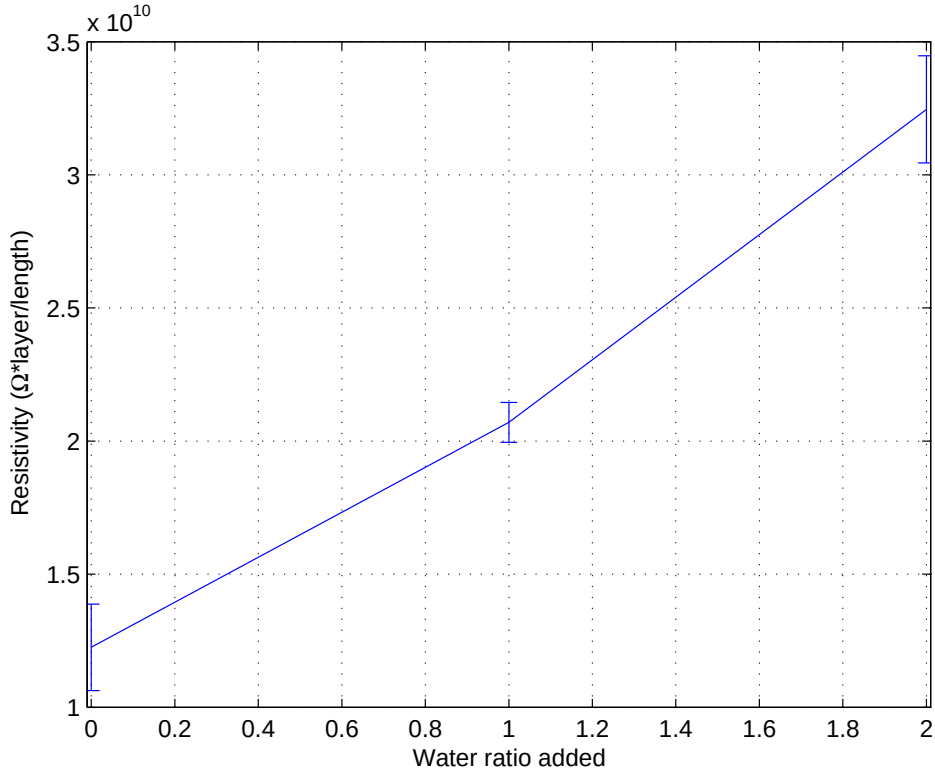

Figure 3.4: Resistance p.u. length per layer of PEDOT vs water dilution

of neat PEDOT ink, corresponding to a tripling of resistivity. As can be seen from Figure 3.4, the resistivity increases with addition of water, but at a slightly lower rate than predicted.

To achieve higher sample conductivity, it is possible to print multiple layers of PEDOT on top of each other. This also has the advantage of reducing the variation of the measured conductivity due to variations in the printing process, as e.g. caused by badly placed drops. The more layers that are printed in parallel, the smaller the impact of individual misplaced drops on the overall cross-sectional area. As the measured width of the PEDOT line did not increase significantly with additional layers, one can assume that additional layers increase the cross-sectional area linearly. Figure 3.5 shows a log-log plot of resistance per unit length against the number of layers for each of the three PEDOT dilutions. For identical resistors in parallel, theory predicts a resistance decrease per layer with a slope of -1 in a log-log plot. This corresponds to the inverse relationship of total resistance for parallel resistors described by Ohm's law. A linear fit of R vs. log(number of layers) for the graphs in Figure 3.5 gives slopes of -1.054 for 2:1 dilution, -1.013 for 1:1 dilution and -1.171 for neat PEDOT (taken from only three points), which corresponds to theory. This result confirms the validity of the assumption that additional layers add a constant cross-sectional area. Therefore conductivity and resistivity are expressed per layer and length.

3.3.3.3 Resistance of PEDOT with DMSO

As recommended by the PEDOT supplier, DMSO may be used as a conductivity-enhancing additive [Starck, a]. The conductivity of PEDOT/DMSO was determined for a fixed PEDOT:water ratio, with varying amounts of DMSO additive. The range of DMSO addition was varied between 10% and 100% of the amount of PEDOT by volume. The result is shown in Figure 3.6. The resistivity decreases by a factor of 43 compared to pure PEDOT when 10% DMSO is added, and

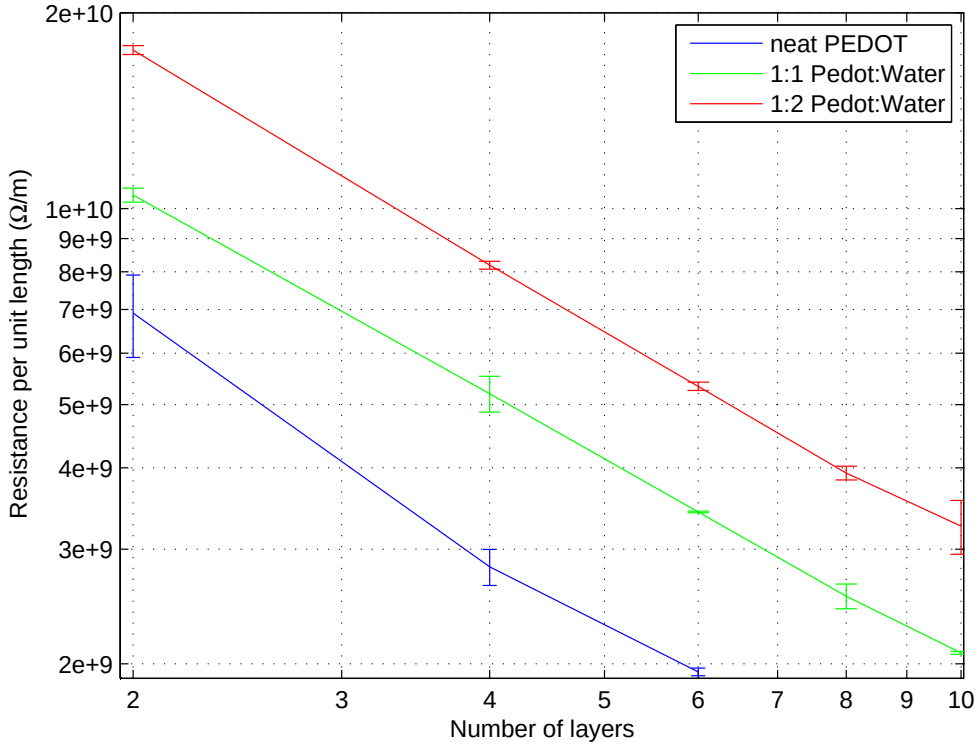

Figure 3.5: Resistance p.u. length of PEDOT lines vs number of layers

a factor of 182 for an addition of 20% DMSO. While further addition of DMSO further increases conductivity, the additional DMSO content has an adverse effect on printability, as can be seen in the optical microscope images in Figure 3.7 (taken for four-layer lines).

Additional DMSO increases the required inter-layer drying time, as well as the travel distance on the substrate and the surface roughness [Garnett and Ginley, 2005]. The increased surface roughness is also visible in Figure 3.7 as thin light or dark lines on the samples. While 1:2:0.5 was still printable at a slow rate (15s additional waiting time between layers with N_2 waft), 1:2:1 was found to be severely limited in terms of drying time and print quality and not practically usable for any more than this initial conductivity measurement. The lines for 1:2:1 PEDOT composition were frequently ($>$ 80% of lines) found to have gaps due to travelling drops, even for multi-layer lines, visible towards the top of Figure 3.7(d). A plausible explanation is that the slight reduction in surface area/energy due to a dry drop on the substrate is sufficient to cause wet drops to travel and deposit on top of the dry drop. It is also worth noting that the addition of DMSO increases the variance of the measured resistivity. This effect is most probably caused by the less uniform cross-sectional area. Visual inspection shows that surface roughness is increased and uniformity of the line width is reduced. These effects are visible when comparing Figure 3.7 to Figure 3.2.

3.4 Substrate materials

Apart from the drop ejection of the print head, the surface properties of the substrate material (e.g. wetability/degree of hydrophobicity) also determine the overall line quality and geometry of the inkjet printed PEDOT lines. Microchips are typically passivated with Silicon Oxide (SiO_2)

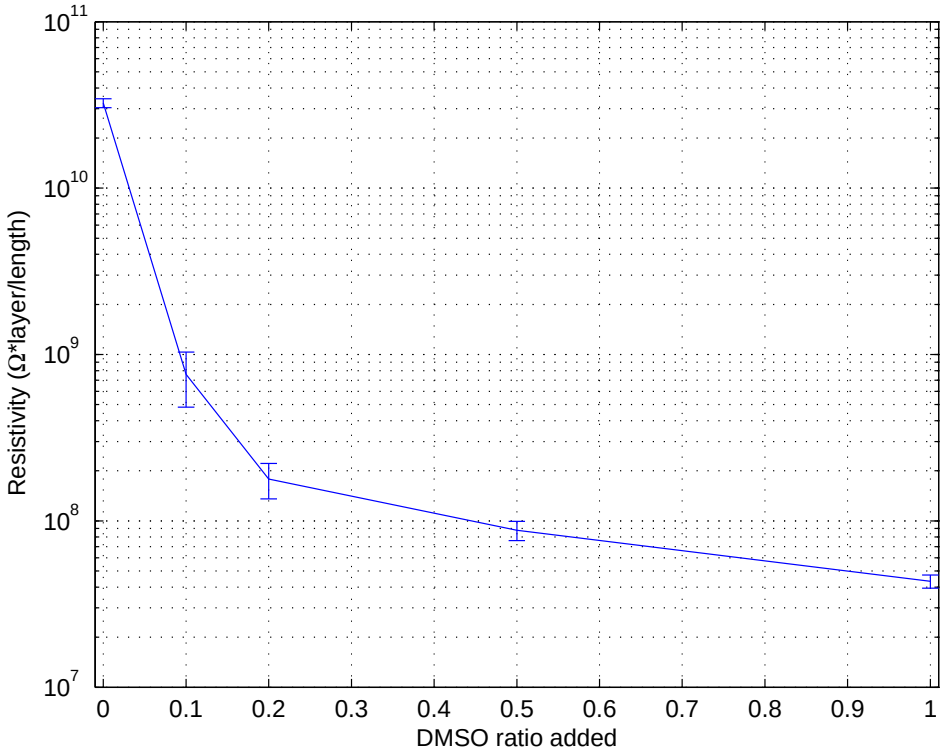

Figure 3.6: Resistance p.u. length per layer of PEDOT vs. added DMSO

and Silicon Nitride (Si_3O_4) [Van Zant, 2004]. By printing PEDOT lines on the sample microchip dies, it was found that the passivation of the chip surface is hydrophobic to the extent that inkjet printing of continuous PEDOT lines is not possible. Figure 3.8(a) shows the printing result on a microchip die. Continuous PEDOT lines are visible on the square aluminium contact pads of the microchip, however these lines are broken on the passivation layer in between the pads. This preliminary test reveals that printing of PEDOT onto a chip for the proposed protection grids is not possible without modifications.

There are two ways to solve this problem. The first possibility is the introduction of an interfacial layer between the silicon nitride surface and the PEDOT protection grids. This way, printing on the hydrophobic silicon nitride surface is simply avoided. This type of structure was also proposed in the context of a topographical surface profile. Several interfacial materials are tested: Araldite epoxy was chosen to emulate standard epoxy packaging material and adhesive [Anderson and Kuhn, 1996]. AZ5214E photoresist is patternable by ultra-violet light (UV) exposure. As the contact pads must remain clear for making contact to the protection grids as well as for the bonding/packaging process, patternability is advantageous. Photoresists have been proposed for structural material in Micro Electro-Mechanical Systems, also known as 'microsystems' (MEMS) products [Conradie and Moore, 2002; Lorenz et al., 1998], confirming their suitability for permanent features in addition to their standard short-term use in lithography. Poly(vinyl phenol) (PVP) is a spin-coatable insulator that may be used as gate insulator in OTFTs [Kawase et al., 2005; Stutzmann et al., 2003], and can be patterned using inkjet printing methods [Kawase et al., 2001].

A necessary requirement for good wetting is for the solid surface tension to be larger than the liquid surface tension [Berg, 1993], so the second possibility to allow printing on hydrophobic

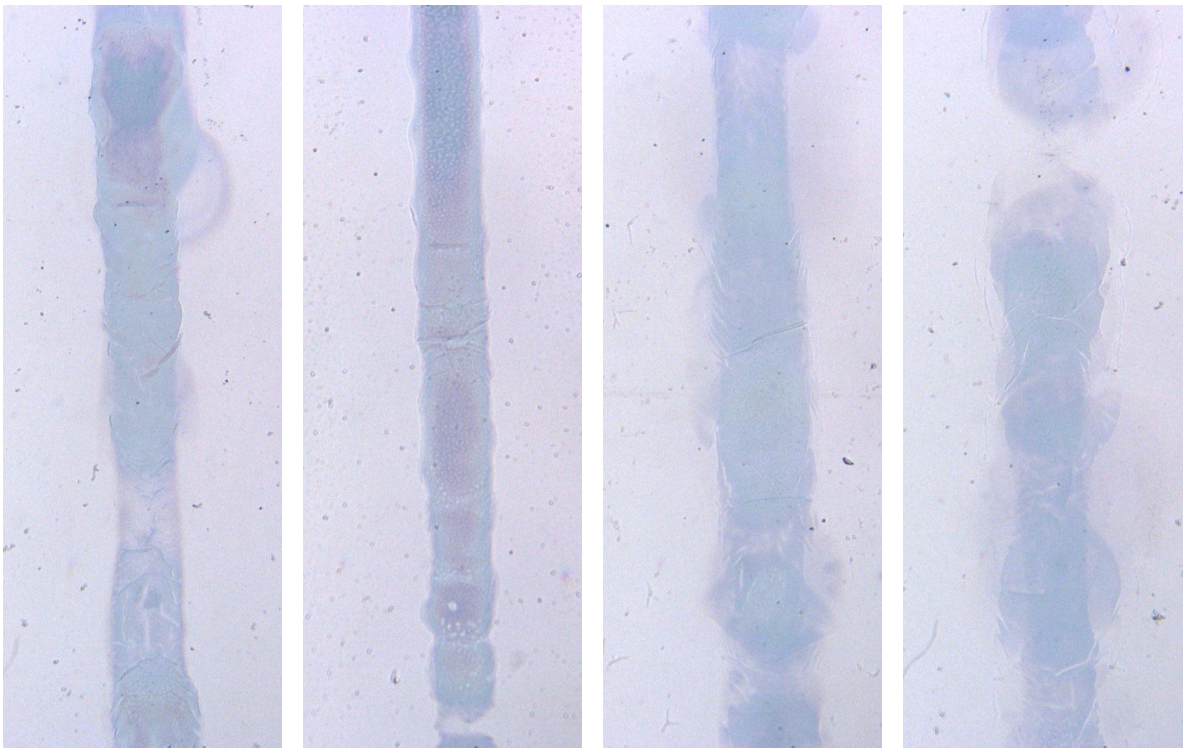

(a) 1:2:0.1 (b) 1:2:0.2 (c) 1:2:0.5 (d) 1:2:1

Figure 3.7: Print quality for lines with varying ratio of PEDOT:water:DMSO

surfaces (such as silicon nitride) is to reduce the surface tension of PEDOT by means of a surfactant [Ferri and Stebe, 2000]. Surfactants are molecules consisting of two parts, a hydrophobic portion and a hydrophilic group attached to it [Aldrich, accessed 5^{th} Sept. 2008]. In addition to lowering the surface tension, the addition of surfactant also reduces the surface roughness of PEDOT with added DMSO [Garnett and Ginley, 2005].

3.4.1 PEDOT printed on silicon nitride surface

A small number of microchip dies were available from another student's failed chip fabrication run, which were used to test the printability on the surface. However, these dies were unsuitable for electrical tests due to their small size and comparatively large aluminium contact pad area. Therefore the microchip surface was substituted with silicon nitride coated glass. Corning glass coated with 500 nm plasma-enhanced chemical vapour deposition (PECVD) silicon nitride was acquired from the Rutherford Appleton Laboratory (RAL).

Compared to the microchip surface, the PECVD silicon nitride coated samples were slightly less hydrophobic. While it was not possible to print directly onto a chip with 1:1 diluted PEDOT, printing succeeded on the substitute PECVD silicon nitride. Figure 3.8 shows photos of PEDOT printed on both types of silicon nitride sample. Figure 3.8(a) shows the failed printing attempt on the microchip surface. The printed lines on the aluminium pads (white squares) are clearly visible. Only patches of lines are visible between the aluminium pads. Figure 3.8(b) shows lines with 2,4,6,8,10 layers (top to bottom) printed on the PECVD silicon nitride. Visually comparing the lines to those printed on glass (e.g. Figure 3.2(b)), one can clearly see the tendency for droplets to travel on the substrate and combine to larger drops on silicon nitride.

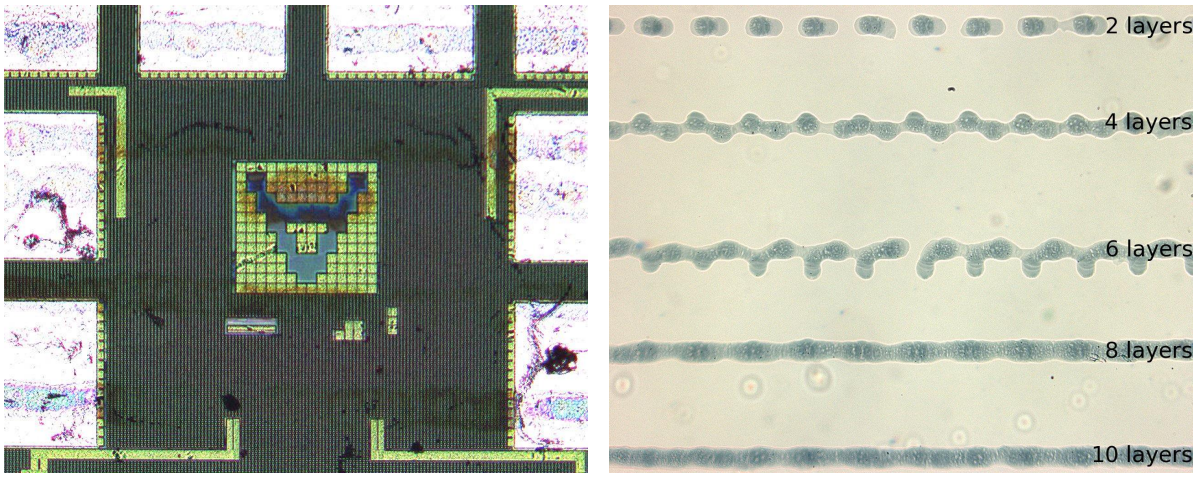

(a) Sample microchip (b) Silicon nitride coated glass substrate

Figure 3.8: Print quality of PEDOT on silicon nitride

For more diluted PEDOT and for PEDOT with DMSO additive, printing failed for the silicon nitride produced by RAL as well as for the microchip samples.

Figure 3.9 shows a log-log plot of the PEDOT resistivity vs. the number of printed layers of 1:1 diluted PEDOT printed on PECVD silicon nitride and on glass. Both samples have identical conductivity between 4 and 8 layers, showing that silicon nitride and glass have comparable influence on PEDOT. The data point for two layers is unreliable due to the large influence of printing faults on conductivity.

3.4.2 Print quality on interfacial layers

For qualitative comparison of the printing properties, PEDOT lines printed with 2 to 10 layers (in steps of 2 layers) are shown in Figure 3.10. All lines are printed with 1:1:0.2 PEDOT:water:DMSO ink. As a reference, Figure 3.10(a) shows lines printed on a glass substrate. For epoxy and PVP substrates (Figures 3.10(b) and 3.10(d)), the individual drops are clearly defined compared to the other substrates. The well-defined print is evidence of a reduced travelling distance of the drops on the substrate. The dark spots in the epoxy substrate are small air bubbles remaining from the mixing of the resin and hardener. The light striations on the PVP substrate are shallow surface features originating from the radial flow during spin coating. The print quality on AZ5214E photoresist (Figure 3.10(c)) is also satisfactory, though not as clearly defined as for PVP and epoxy.

3.4.3 Conductivity of PEDOT on epoxy interface layer

To test the influence of epoxy as substrate material on the conductivity of PEDOT, a glass slide was coated with a thin layer of epoxy. The edge of a second glass slide was used as a doctor blade to spread out a thin layer. Using this spreading method, it is possible to give the epoxy layer a deliberately uneven surface, and to thus create an irregular topographical surface (Figure 3.11).

The resistivity for pure PEDOT (1:1 diluted with water) and PEDOT with DMSO (1:1:0.2) is shown against numbers of layers in Figure 3.12. For pure PEDOT, the two curves coincide, showing that epoxy has no influence on the conductivity. For PEDOT with DMSO additive, a difference in resistivity of a factor of two was determined (Figure 3.12(b)). The precise cause

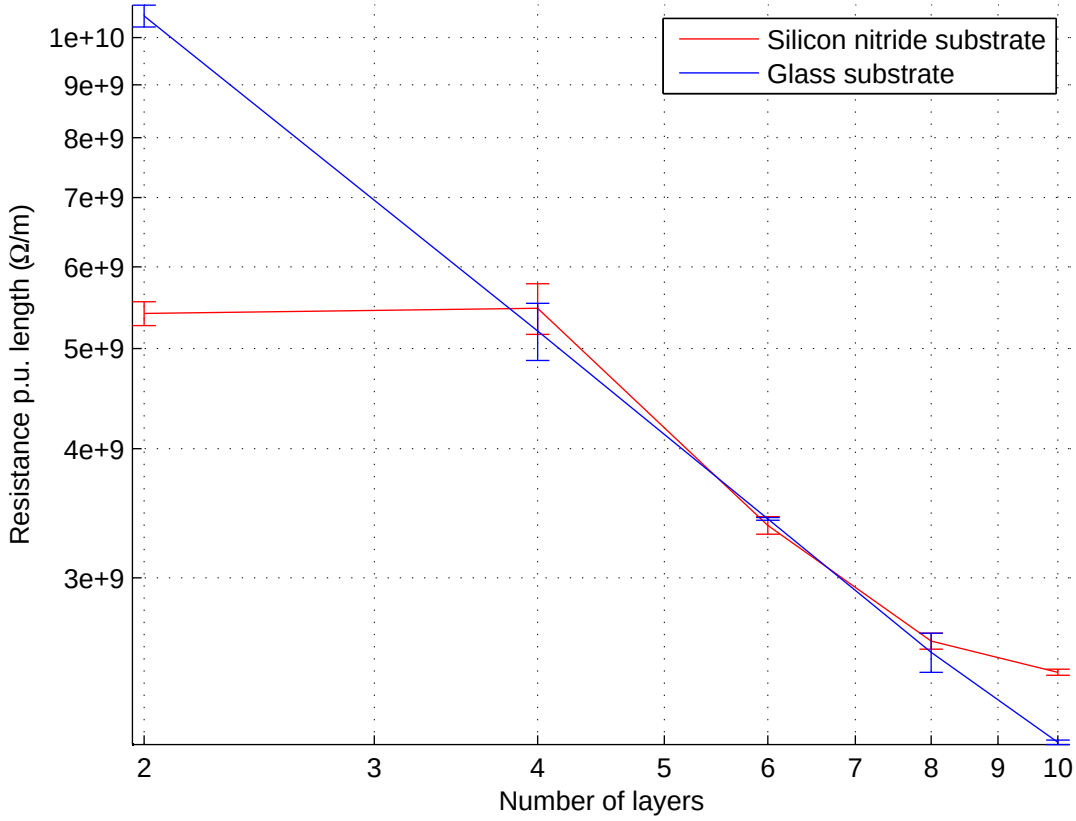

Figure 3.9: Resistivity vs. layers for silicon nitride substrate

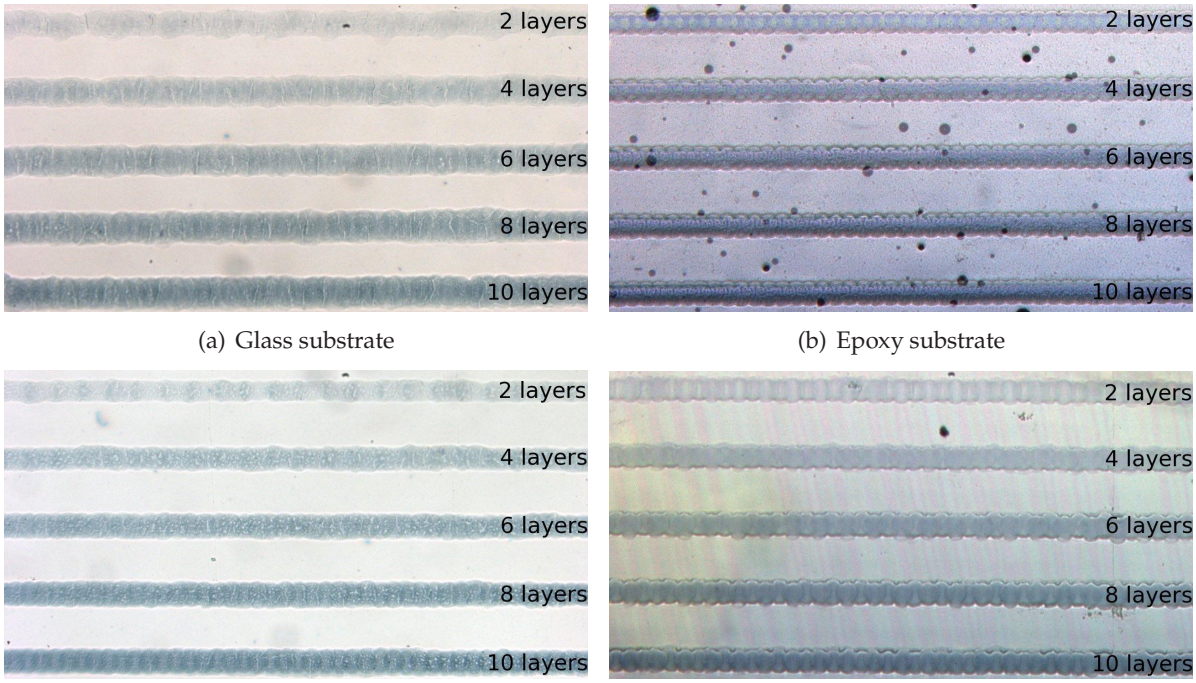

(a) Glass substrate

(b) Epoxy substrate

(c) Az5214e substrate

(d) PVP substrate

Figure 3.10: PEDOT lines on different substrate materials, using 1:1:0.2 PEDOT:water:DMSO ink

for this difference is not known, but may lie in the slightly different line geometry and print quality compared to glass (Figure 3.10). Another possibility is the interaction of DMSO with the Araldite epoxy, however this would require further investigation beyond the scope of this dissertation.

3.4.4 Conductivity of PEDOT on photoresist interface layer

The AZ5214E photoresist was spin-coated onto a glass slide at 4000 rpm for 30 s, then pre-baked for 50 s at 110 °C, and immediately developed with 1:4 diluted AZ400K developer. This recipe makes AZ5214E a positive photoresist, and development without prior UV exposure leaves a blanket layer in place. For patterning, a UV exposure step would be included with the desired pattern as holes or transparent areas in the mask.

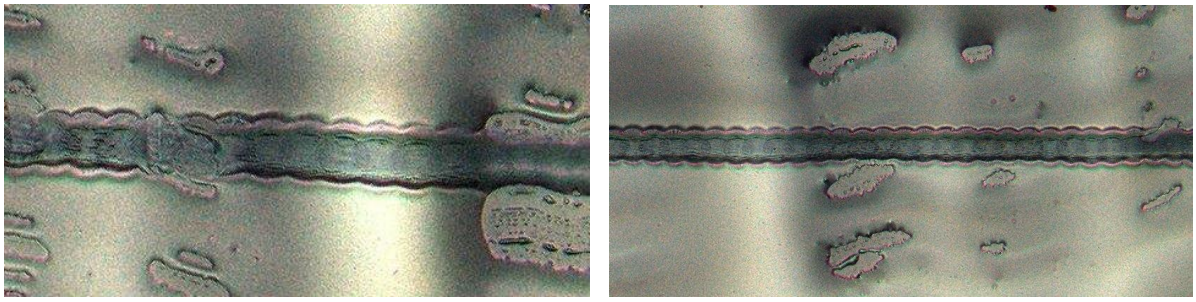

Figure 3.11: Topographic epoxy substrate

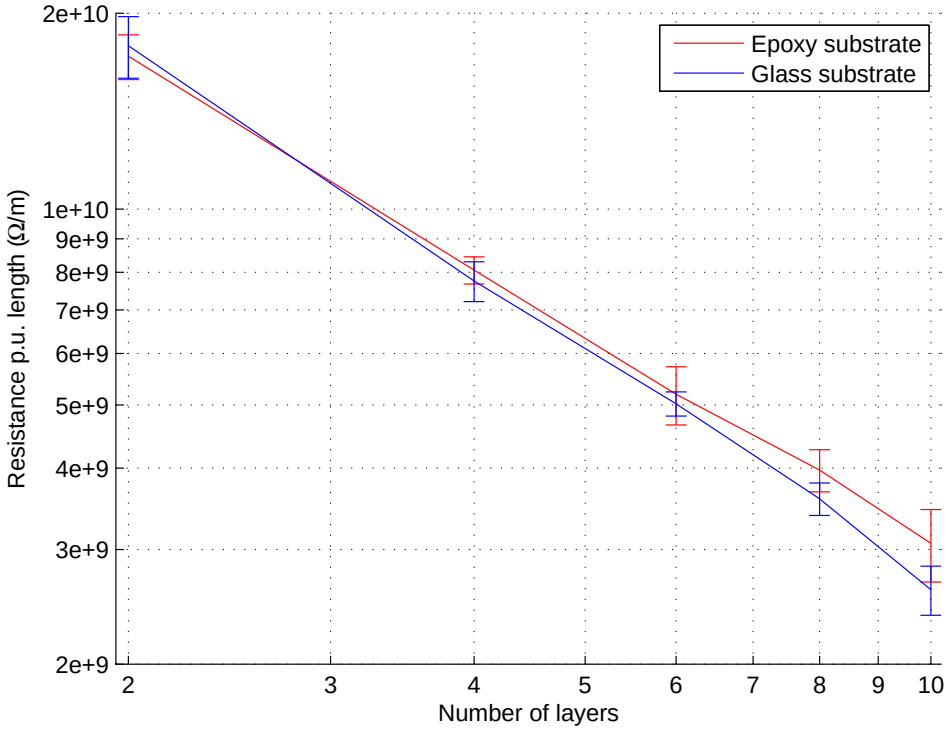

(a) Pure PEDOT ink

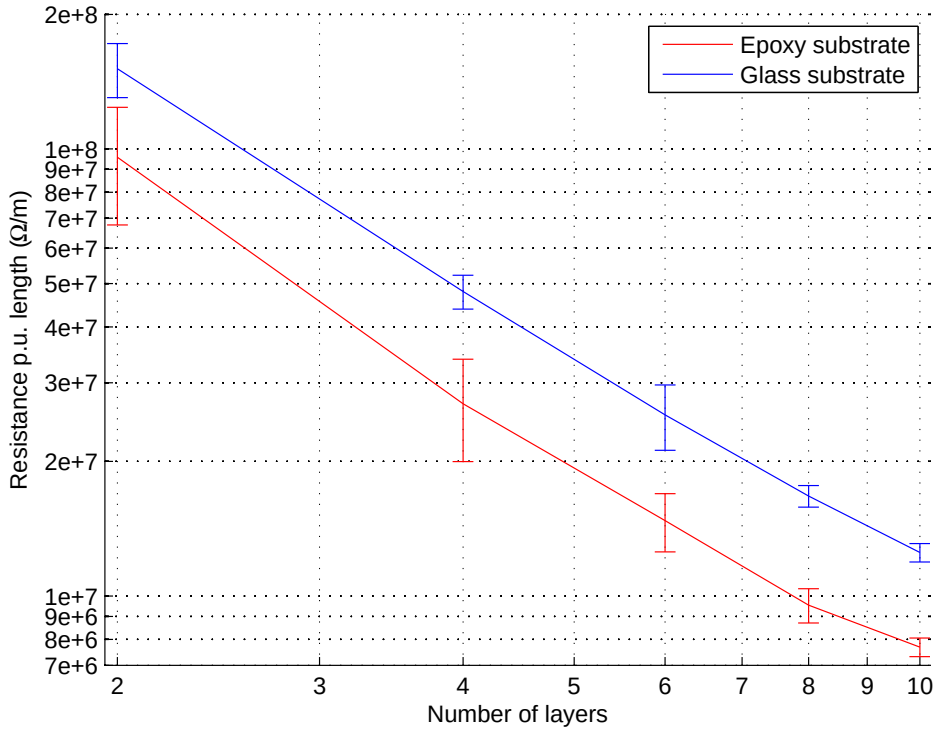

(b) PEDOT ink with DMSO additive

Figure 3.12: Resistivity vs. layers for epoxy interfacial layer

The measured resistivities of PEDOT on glass and photoresist are shown in Figure 3.13. The resistivities on photoresist are identical to the ones measured on glass for pure PEDOT (Figure 3.13(a)), showing that no interaction takes place between the two materials. However, the resistivity of PEDOT with DMSO (Figure 3.13(b)) is measured to be lower on epoxy compared to glass for lines with less than eight layers. The slope of the graph deviates from the expected $1/x$ relationship, which indicates that the simple per-layer resistivity model is not valid in this case.

3.4.5 Conductivity of PEDOT on PVP interface layer

PVP has been reported to be used as a gate insulator of OTFTs [Kawase et al., 2005; Stutzmann et al., 2003], therefore it is known to be compatible with organic semiconductors. Epoxy and photoresist may not necessarily be compatible. As PEDOT is a conductor, it may be printed on top of PVP as the gate electrode of OTFTs [Burns et al., 2003; Chason et al., 2005; Kawase et al., 2005; Liu et al., 2005; Stutzmann et al., 2003]. Thus, apart from using PVP as an interfacial layer, the interaction of PEDOT and PVP is also relevant for transistor characteristics.

PVP insulator was spin coated from isopropanol solution (70 mg/ml) at 2000 rpm for 30 s onto a glass slide. This results in a ca. 1 µm thick layer. After drying on a hot plate (80 °C) in air for several minutes, PEDOT lines were printed on the layer both with and without conductivity-enhancing DMSO additive. The same ink composition was used as for the previous measurements, 1:1 PEDOT:water and 1:1:0.2 pedot:water:DMSO.

The graph of resistivity vs. number of layers is shown in Figure 3.14 on PVP and on glass substrates. For pure PEDOT (Figure 3.14(a)) the resistivity matches well, similar to the previous measurements. However, for the samples printed with added DMSO, the resistance on PVP is significantly higher than that of PEDOT/DMSO printed on glass (Figure 3.14(b)). For four layers, the difference is a factor of ten, and for ten layers a factor of five. This difference cannot be explained by geometric differences alone, particularly given that the print quality on PVP is very regular (Figure 3.10(d)). The difference between the two ink compositions printed on PVP is the addition of DMSO. Given the relevance of PVP for OTFT fabrication, the interaction of DMSO with PVP was investigated. A simple solubility test (PVP powder immersed in DMSO) revealed that PVP has fast solubility in DMSO at room temperature. A powder generally has a larger surface area, thus greater solubility compared to a blanket surface coating. Single drops of DMSO from a pipette were therefore also applied to a dried PVP substrate layer. It was shown that DMSO dissolves PVP fast enough to remove the interfacial layer with only a few seconds of interaction, which is of the same time scale as the drying time of the PEDOT ink. While these tests aren't quantitative, it can nevertheless be concluded that re-dissolution of PVP will have taken place. The higher resistivity of the PEDOT lines makes it plausible that re-dissolved PVP may have been incorporated in the PEDOT line.

3.4.6 Surfactant to relieve surface tension

As a high surface tension of PEDOT ink may be the cause for wetability problems, the effect of the reduction of ink surface tension on the printability was examined. A suitable surfactant for addition to PEDOT dispersion is Surfynol 104E produced by Air Products. The surfactant is listed by the supplier of PEDOT as one of the recommended surfactants for PEDOT processing [Starck, a]. An ink composition of PEDOT:Water:DMSO:Surfynol 1:1:0.1:0.01 was used for the test. This ink composition was chosen as a compromise between printability (without surfactant) and conductivity. 1:1 diluted PEDOT was shown to be barely printable on PECVD silicon nitride by itself. As DMSO is the culprit for degradation of print quality, only a small amount was added. The addition of 10% DMSO was determined as the minimal amount necessary to

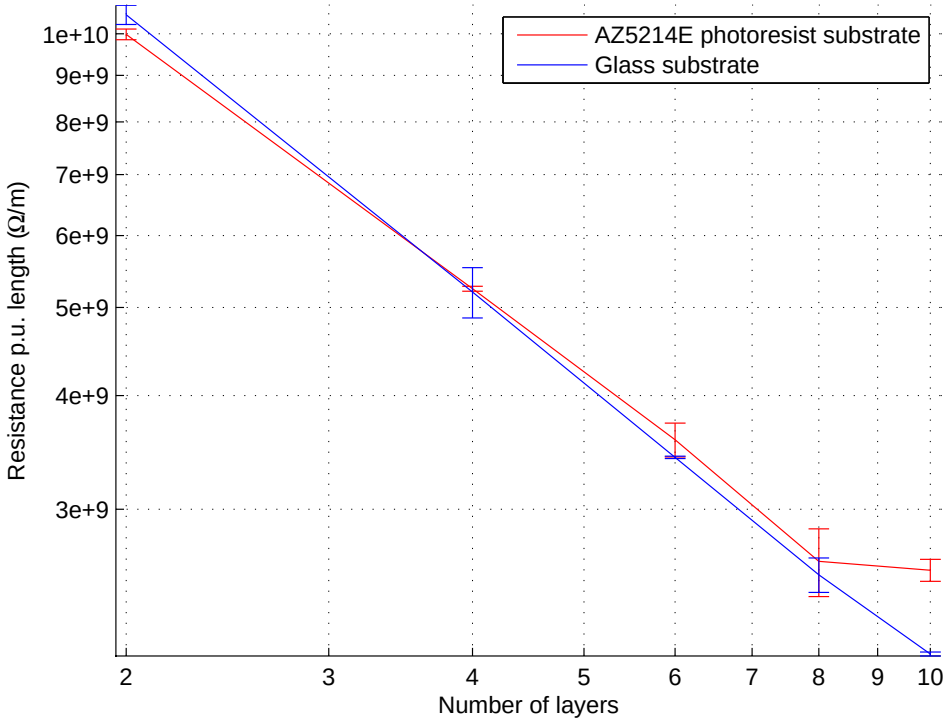

(a) Pure PEDOT ink

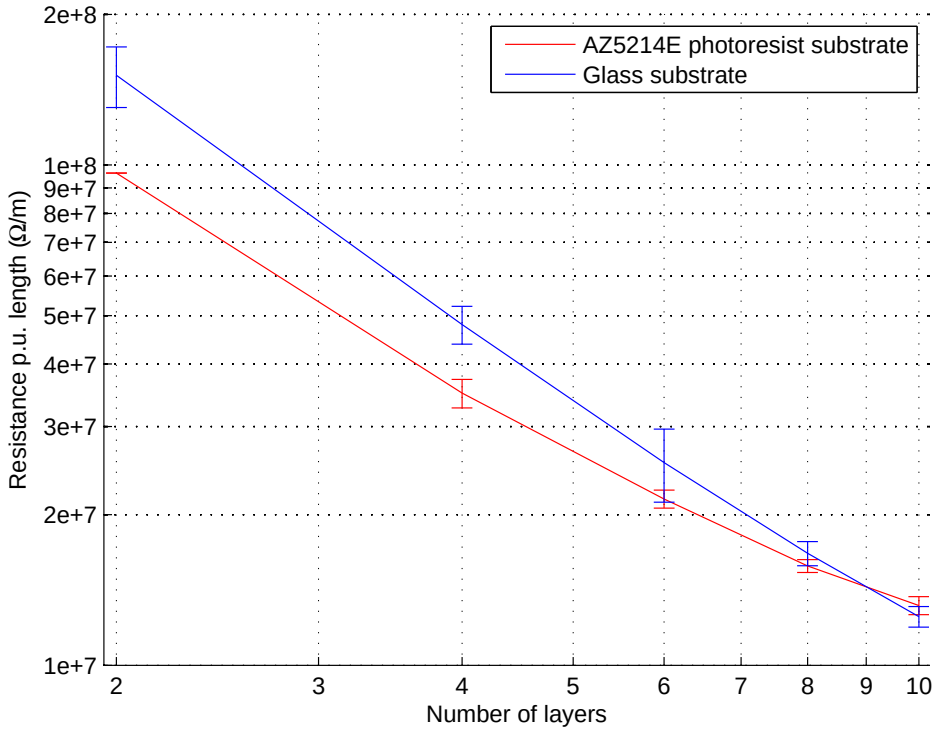

(b) PEDOT ink with DMSO additive

Figure 3.13: Resistivity vs. layers for photoresist interfacial layer

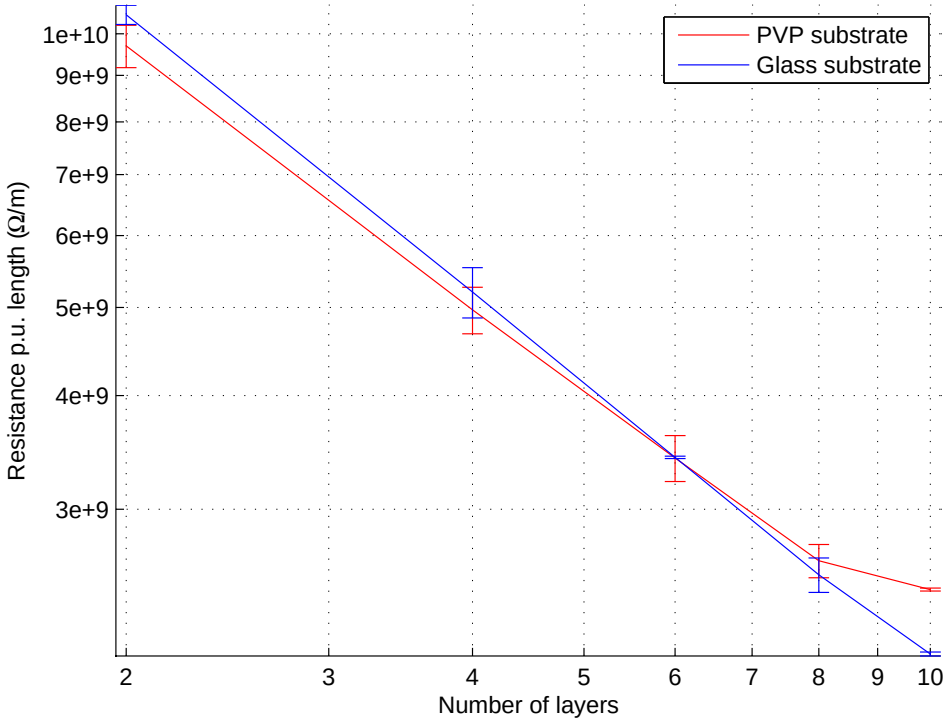

(a) Pure PEDOT ink

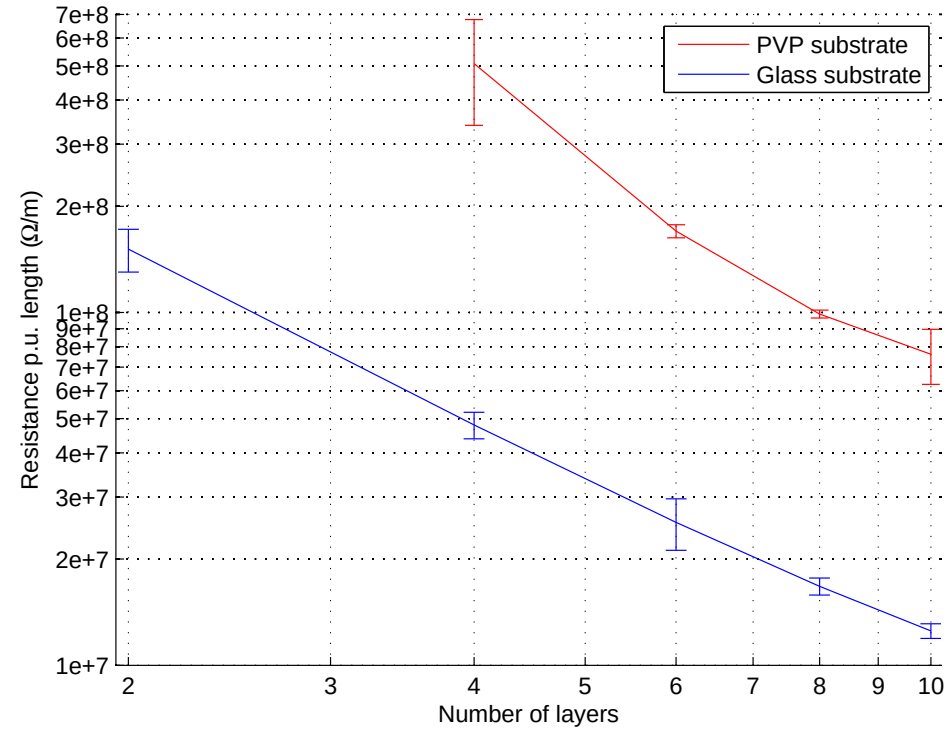

(b) PEDOT ink with DMSO additive

Figure 3.14: Resistivity vs. layers for PVP layer

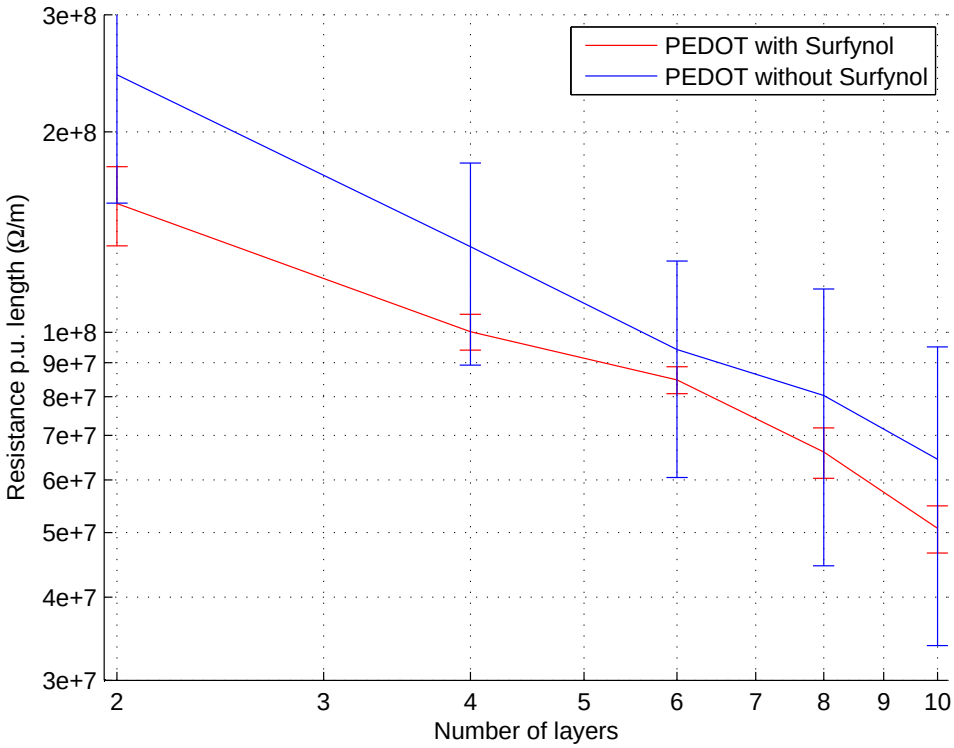

Figure 3.15: Resistivity vs. layers for PEDOT/DMSO on glass

achieve a significant increase in conductivity (Section 3.3.3.3). The results of the electrical characterisation of PEDOT with and without added Surfynol on glass are shown in Figure 3.15. It can be seen that the conductivity of both compositions are of the same order of magnitude, with a lower resistivity of PEDOT with Surfynol. This conductivity-enhancing effect was also observed by Garnett et al. [Garnett and Ginley, 2005].

While the addition of surfactant allows PEDOT to be printed directly on silicon nitride, the goal of reliable fabrication is not yet achieved. After printing several samples or layers, the substrate became flooded with PEDOT ink. As this problem occurred repeatedly, an investigation was carried out.

The ink is held in the nozzles of the print head by capillary forces until expunged by the piezo crystal [Le, 1998]. The capillary forces are created by the surface tension [de Gennes et al., 2004], therefore the reduction in surface tension also reduces the forces retaining the ink in the nozzles. For the Epson Stylus Color II print head used in this print setup, the nozzles must therefore be too large to prevent the ink from seeping through and collecting on the nozzle plate. The image series in Figure 3.16 shows the progress of ink seepage during normal idling (ink ejection to prevent nozzle clogging) operation. Initially, a small drop collects at one of the printing nozzles, marked '1' in Figure 3.16(a). After a short time (of the order of 1 minute) the drop grows to cover more nozzles ('1' in Figure 3.16(b)). The lighter streaks marked '2' in Figures 3.16(a) and 3.16(b) are correctly expunged ink drops from the print head. Finally, the drop on the nozzle plate expands to cover all nozzles (Figure 3.16(c)). After the drop reaches a critical size, the pressure difference caused by the drop is large enough to suck the ink out of the nozzles without the help of the piezo crystal, sustaining the seepage even when the printer is turned off. For comparison, Figure 3.16(d) shows a water drop of similar size to the large drop

in Figure 3.16(c) emanating from the print head. The water drop was generated by raising the ink reservoir to artificially increase the liquid pressure in the print head and force ink out of the nozzles. The higher surface tension of the pure water is visible by comparing the drop spread on the nozzle plate. A dry print head next to the drop is shown for size comparison (labelled '3'). A lowering of the ink reservoir could reduce the seepage of the ink with surfactant, but not entirely prevent it, as the reduction in ink pressure eventually prevents ink from re-filling the print head.

3.4.6.1 Application of surfactant to substrate surface

To avoid the problem of ink seepage, the surfactant was used on the substrate rather than as an additive to the ink. A 4% aqueous solution of Surfynol 104E surfactant was applied to the substrates and left to dry in place. After the Surfynol treatment, resistor lines were printed with 1:1:0.2 PEDOT:water:DMSO solution. Figure 3.17 shows microscope images of PEDOT lines printed for 2,4,6,8,10 layers (top to bottom) with and without Surfynol on different substrate materials. It is evident that printability is vastly enhanced, as printing of continuous lines is possible on both PECVD silicon nitride (Figure 3.17(a) vs. Figure 3.17(b)) and the surface of one of the sample microchips (Figure 3.17(c) vs. Figure 3.17(d)). Both reference images on silicon nitride were taken using PEDOT ink without DMSO addition, as printing with DMSO did not result in a useful pattern. Compared to lines printed without Surfynol (Figure 3.17(e)), PEDOT lines printed on glass with a Surfynol coating were wider (Figure 3.17(f)). The lines printed on Surfynol were found to be 100 µm to 110 µm wide compared the 70 µm to 80 µm width measured for samples without. Visual inspection of the glass/Surfynol samples shows formation of a more narrowly defined centre region of 65 µm to 70 µm width for 8 and 10 layers, which is close to the original print width. A higher magnification detail of 10 PEDOT layers is shown in Figure 3.17(g). A plausible explanation for the formation of the narrow centre region may be the incorporation of the surfactant in the PEDOT in the lower layers. As the layers are dry when adding the next layer, less Surfynol reaches the additional PEDOT layers. The electrical behaviour of PEDOT lines printed on Surfynol-coated glass is shown in Figure 3.18. While the conductivity is of the same order of magnitude for blank and Surfynol-coated glass, the strong deviation from the $1/x$ model shows that each additional layer has different properties, through a lower Surfynol content and different geometry.

3.5 Contacts between Aluminium and PEDOT

While the substrate material and surface properties are relevant for line print quality and thus line integrity, the protection grids must also interface with CMOS microchips. An important factor for successful integration are the contacts between the CMOS chip contact pads and the organic conductor.

In security devices made from standard microchip technology, the metal layers and contact pads are made from aluminium [Van Zant, 2004]. The problem with aluminium is that it oxidises rapidly. It almost instantaneously forms a thin insulating aluminium oxide (Al_2O_3) layer on its surface. The oxide layer is less of a problem with standard wire-bonding, as the oxide is physically broken by deformation of the contact pad during the bonding process [Servais and Brandenburg, 1991]. Inkjet deposition is carried out without force, thus without contact pad deformation or physical oxide removal. The electrical contacts between aluminium and the protection grid are therefore expected to be of low quality. The addition of an oxide removal step (e.g. chemical removal with phosphoric acid (H_3PO_4) [El-Kareh, 1995]) before printing is

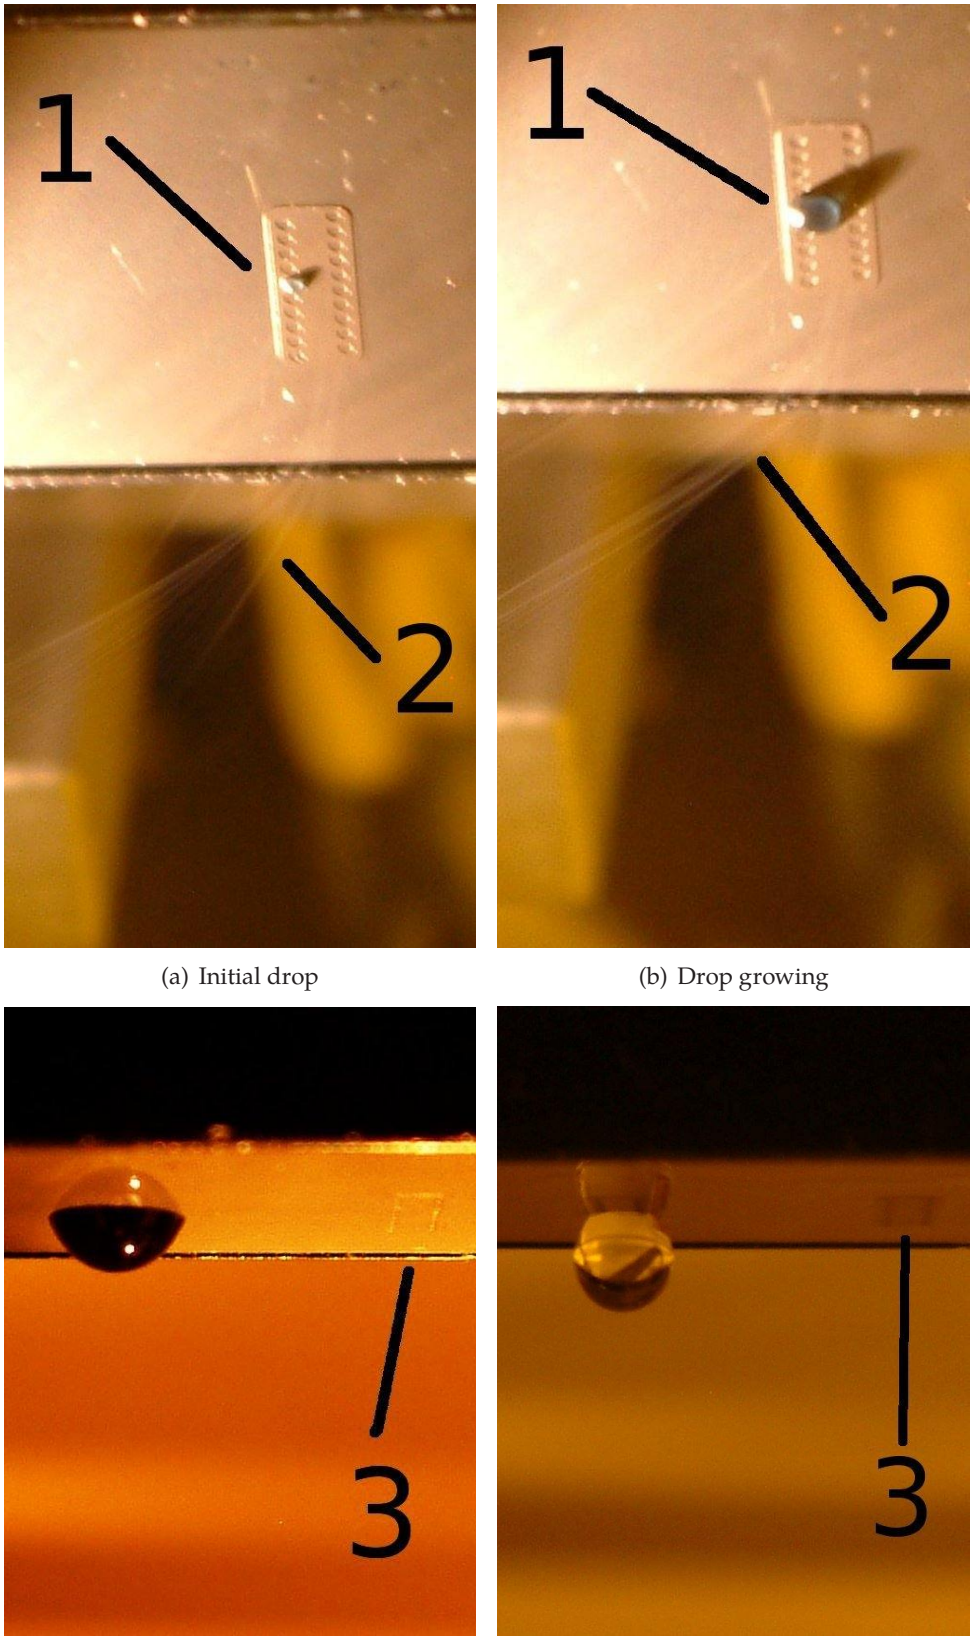

(a) Initial drop

(b) Drop growing

(c) Covered print head

(d) Pure water

Figure 3.16: Surfynol additive causing seepage

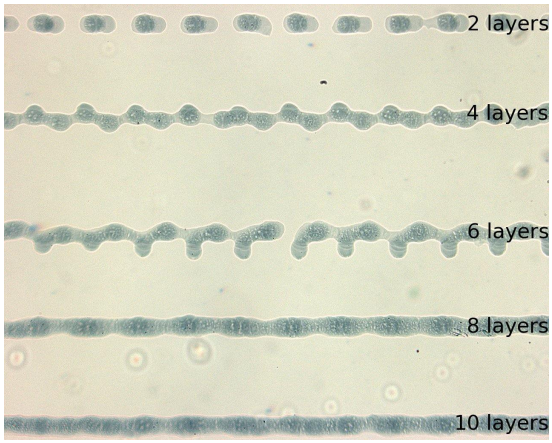

(a) Silicon nitride coated glass, no DMSO, cf. Figure 3.8(b))

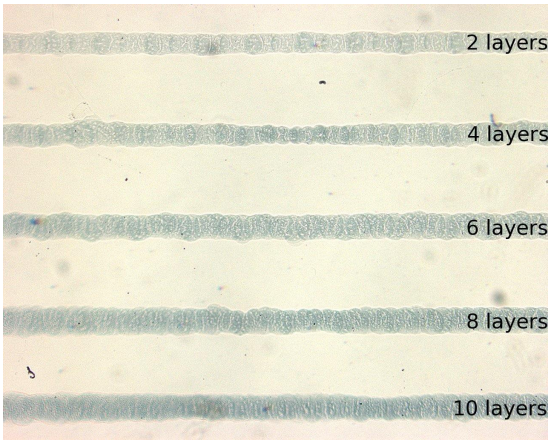

(b) Silicon nitride coated glass, DMSO, Surfynol

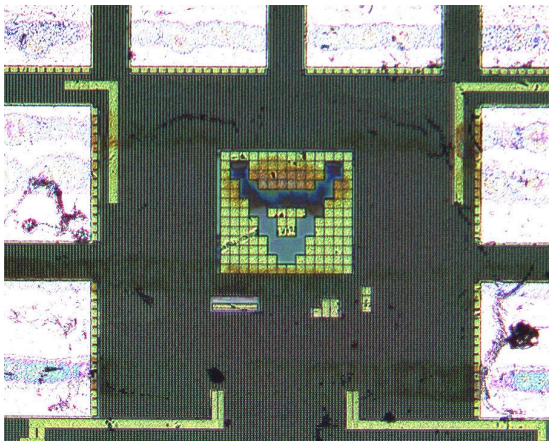

(c) Sample microchip, no DMSO (cf. Fig. 3.8(a))

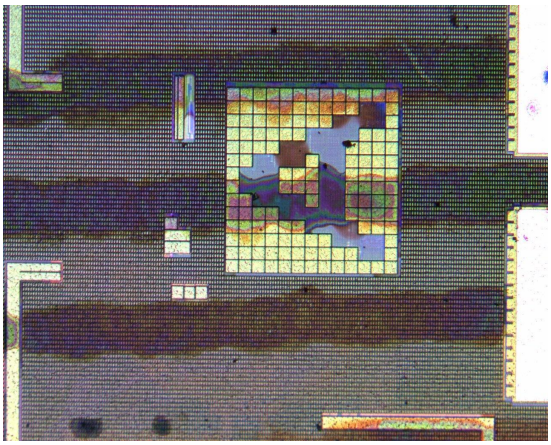

(d) Sample microchip, DMSO, Surfynol

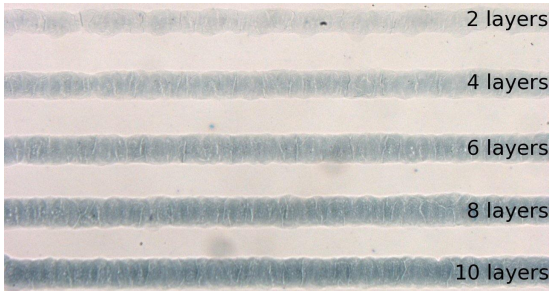

(e) Glass substrate (cf. Fig.3.10(a))

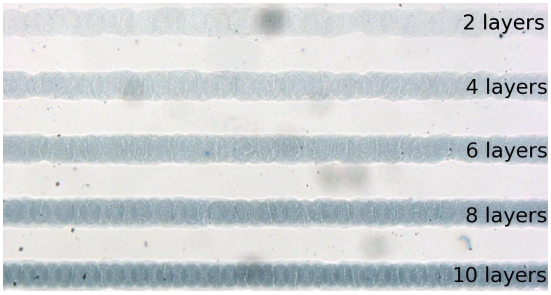

(f) Glass substrate, DMSO, Surfynol

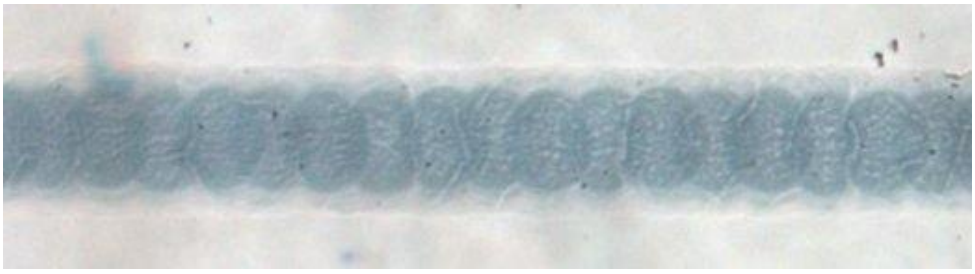

(g) Glass substrate, DMSO, Surfynol (detail of 10 layers)

Figure 3.17: Print quality of PEDOT with/without Surfynol coating

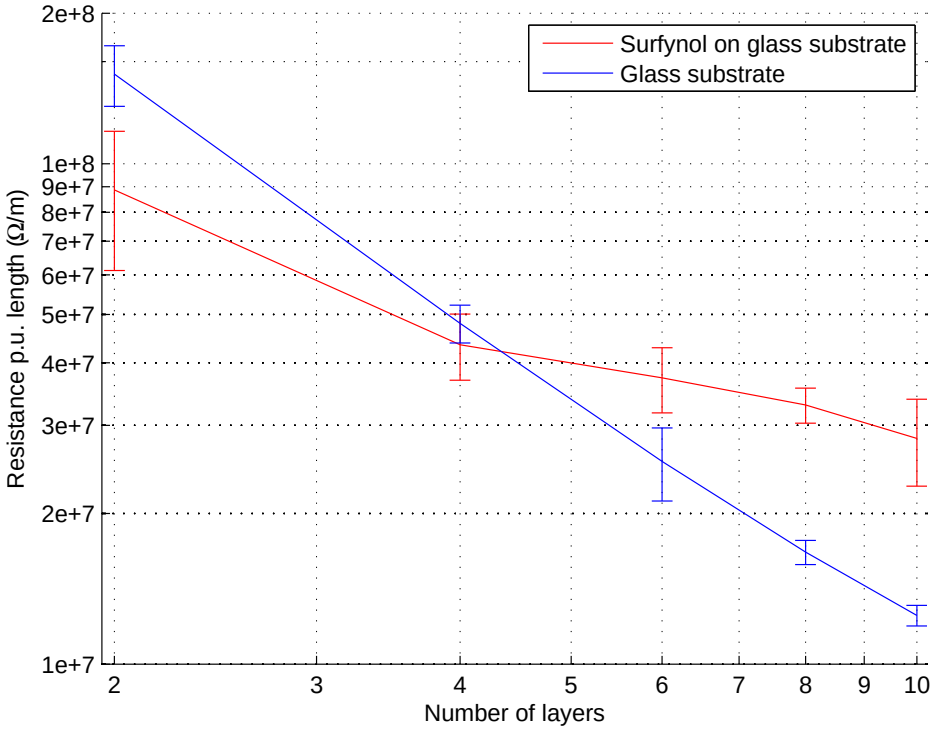

Figure 3.18: Resistivity vs. layers for PEDOT/DMSO printed on substrate coated with Surfynol

not a viable option, as the PEDOT solution is both aqueous and acidic [Starck, b,c]. The water and acid allow the re-formation of the oxide layer before the lines are dry.

As reported in literature, the I-V properties of PEDOT-aluminium contacts are non-linear. The contacts are either reported as rectifying Schottky type [Liang et al., 2002; Turut and Koleli, 1992] or non-ohmic and symmetrical, consistent with an insulator layer in series with a layer depleted of charge carriers [Bantikassegn and Inganäs, 1997]. Neither characteristic is desirable. For a PEDOT line printed between two aluminium contacts, a rectifying contact means that one of the two contacts is always reverse-biased. Current flow will be inhibited significantly until the reverse break-down voltage across the Schottky contact is reached. This voltage is reported at 5.5 V [Liang et al., 2002], 0.5 V higher than the standard operating voltage of a typical security device [Rankl and Effing, 2004]. This estimate also ignores any voltage drop at the second contact, and any I R drop along the line. The situation is only slightly better for symmetric properties, as the voltage drop at each contact is still likely to be significant (several Volts) in order for sufficient current to flow.

The contacts reported in literature [Bantikassegn and Inganäs, 1997] were fabricated by evaporation of aluminium on top of PEDOT, therefore avoiding the native oxide layer present in pre-deposited aluminium. Nevertheless, a layer of aluminium oxide was identified, as well as a voltage-dependent semiconducting layer [Bantikassegn and Inganäs, 1997; Jönsson et al., 2003]. This reiterates the fact that oxide removal prior to printing does not help the contact properties. Jönsson et al. reported that aluminium reacts with the $SO_3^- H^+$ groups of the PSS dopant, essentially removing the dopant from the PEDOT at the interface [Jönsson et al., 2003]. Therefore, an alternative method must be found to allow ohmic contacts to be made between microchip bond pads and the PEDOT protection grid.

Experiments were carried out to determine properties of PEDOT-aluminium contacts where

the aluminium is already coated by the native oxide layer. The non-functioning microchip samples were found to be unsuitable for the evaluation, as the contact pads include electro-static discharge (ESD) protection circuitry, which leaks current and thus skews the measurement results. Aluminium contacts were therefore deposited on a glass substrate. The properties of the aluminium/PEDOT contacts were measured on the fabricated aluminium pads. Silver paint does not react with PEDOT, therefore it is introduced as a method for contact enhancement.

3.5.1 Fabrication of aluminium contacts

Rectangular aluminium contacts were deposited on glass using a lift-off lithography process. Lift-off lithography was chosen over etching, as it does not require hazardous chemicals, which acidic aluminium etchant does [Microchemicals, 2007]. Lift-off lithography is compatible with other materials (e.g. gold or indium tin oxide). The method is more universal, in case other contact materials or patterns are required. The image reversal photoresist AZ5214E was used for patterning. An image reversal photoresist may be used both as positive or negative resist. If it is used as a negative resist, an additional image reversal step is required (baking and flood exposure). AZ5214E is well suited for lift-off lithography as the side walls have an overhanging/ negative slope after image reversal [Clariant]. The lift-off lithography recipe based on the AZ5214E photoresist datasheet [Clariant] was adapted for use with the available Canon mask aligner:

Ultrasonic substrate cleaning: acetone, IPA, DI water

30 minutes de-hydration bake at 125 °C

Spin coating of AZ5214E photoresist, 4000 rpm, 30 s

Pre-bake 105 °C, 1 minute

Exposure with mask, 3.5 Light Integer (LI) units (one light integer corresponds to ca. 1.5 s)

Image reversal bake to convert positive into negative resist (critical in terms of temperature and time), calibrated to 120°C on the digitally controlled hot plate, 2 minutes

Flood exposure, minimum 10 LI

Development, 1:4 AZ400K developer:DI water, ca. 45 s

The lithography was designed to use a negative photoresist, such as AZ5214E after image-reversal [Clariant]. The patterning results in windows in the photoresist layer where aluminium contacts are required. After patterning the photoresist, the samples were loaded into the thermal evaporator. The evaporation chamber was evacuated to less than $4 \quad 10^{-6}$ bars. A thin adhesion layer of chromium was thermally evaporated onto the substrates [Franssila, 2004], before deposition of the aluminium layer. The unwanted aluminium was removed by immersion in acetone under ultrasonic agitation. This dissolves the cured photoresist, lifting away the aluminium above it. To prevent re-deposition of aluminium, the samples were rinsed in acetone and IPA. The resulting contacts were found to be of correct geometry and without pin holes. For long, narrow wires, the adhesion of the metal to the glass substrate is a limiting factor. The lift-off procedure frequently removed or deformed parts of the thin connecting wires as well.

3.5.2 Evaluation of contact properties

3.5.2.1 Two-point measurement

Standard 1:1:0.2 PEDOT:water:DMSO lines with 6 layers and a drop spacing of 40 µm were printed between two as-deposited aluminium contacts. The aluminium contacts were placed between 1000 µm and 2000 µm apart. The contact area between aluminium and PEDOT for each sample was approximately 500 µm long and 80 µm wide. The electrical properties of the resulting aluminium-PEDOT-aluminium structure were determined with a HP4156A semiconductor parameter analyser.

An initial measurement to confirm good contact between the probe needles and aluminium pads was carried out by placing two probes at opposite ends of an aluminium contact pad without PEDOT, and measuring the electrical resistance. As expected, the resistance between the two probes was measured to be negligible ($<1\ \Omega$).

The probes were then placed on the aluminium pads at either end of a PEDOT line. Care was taken to ensure the probes only made contact to the aluminium pads, not the PEDOT lines. The I-V characteristics were determined as a two-point measurement both in light and dark conditions between -5 V and 5 V as a double sweep, the same conditions as used for the PEDOT characterisation. Measurements were then repeated several times for each sample, both in short succession (every few minutes), and also with longer gaps of more than 24 h.

Figure 3.19 shows typical I-V curves taken in six measurement cycle repetitions on a sample with low contact noise. Figure 3.19(b) shows a detail of the same curves, magnified to the region between -2 V and -5 V. The non-linearity of the curves is clearly visible, as well as degradation of conductivity. As the examined samples are symmetrical (Al-PEDOT-Al), the measured I-V characteristics are automatically also symmetrical, therefore no information can be extracted as to which contact model (rectifying or not) is valid. Other samples showed significant noise during measurement. An example for noisy I-V curves is included in Figure 3.21(b).

Regardless of contact model, the dominant feature of the I-V curves is the degradation between subsequent measurements. It can be seen from the detail plot (Figure 3.19(b)) that degradation takes place during the voltage sweeps of ca. 10 s duration. Taking the current at the lowest voltage (-5 V) as a benchmark, it can be seen that the conductivity was reduced by 80% within six measurement repetitions. The exact amount of degradation after each measurement cycle varied between samples, but was present in all samples. Pure PEDOT and PEDOT with added DMSO both showed the same degradation and hysteresis behaviour, with the only difference being that PEDOT with added DMSO generally showed more contact noise than pure PEDOT. No recovery was seen when re-measuring the samples more than 24h later. The degradation was also found to be independent of age of aluminium contacts (which determines the initial thickness of the oxide layer before PEDOT deposition), and of light or darkness conditions during measurement and sample storage.

Generally, non-linear contacts may not be ideal, but they are less of a problem than the permanent degradation of the contacts, which prevents repeatability, a key requirement for protection grids. If each measurement gives different results, it is impossible to determine the integrity of the grid. Recovery times of more than a minute or so severely limit the use of protection grids, as integrity tests would not be able to be repeated in the recovery period. For short recovery times, it would be feasible to delay the integrity test after power on, to ensure that the grid has recovered even if the card is used several times in short succession.

An alternative interpretation of the degradation effect may be a charge trapping model applied to the PEDOT-aluminium interface [Streetman and Banerjee, 2000]. With sufficiently deep traps, recovery would also take place slowly [Streetman and Banerjee, 2000]. However, the recovery of conductivity would be accelerated under illumination [Streetman and Banerjee, 2000],

which has not been observed here. A permanent change in the contacts is therefore most likely, such as an electro-chemical reaction at the interface [Jönsson et al., 2003; Liu et al., 2004].

3.5.3 Contact enhancement with silver conductive paint

Silver paint has been reported to make good contact with PEDOT [Hamedi et al., 2007; Mabrook et al., 2005]. Standard silver paint (manufactured by Electrolube) was therefore applied on top of the previously tested PEDOT lines with aluminium contacts. The conductive paint was deposited to overlap the aluminium contact pads and the PEDOT lines, as shown in Figure 3.20. In this experiment, the paint was manually deposited using an applicator. The standard silver paint is not compatible with the available Epson inkjet print head due to the high viscosity and low surface tension of the standard silver composition (solvents are ethanol, acetone, and ethyl acetate [Electrolube]). The combination of the large particle size of the silver (compared to PEDOT, individual silver particles are visible) and fast drying time, also increases the probability of blocking the inkjet nozzles. However, using nano-particulate silver, it is reported to be possible to make an inkjet printable silver preparation [Lee et al., 2005; Perelaer et al., 2008; van Osch et al., 2008].

3.5.3.1 Two-point measurement

A two-point measurement was carried out using the same setup and parameters as previously described in Section 3.5.2.1. The probes were lowered onto uncovered parts of the aluminium pads, as shown in the cross section of the contacts in Figure 3.20. Care was taken to avoid inadvertent contact to the silver paint and PEDOT line. Initially, three I-V curves were taken for the same sample before silver application, after which two I-V curves were taken. Figure 3.21(a) shows the full result of typical I-V measurement of the silver-enhanced contacts. The near horizontal graphs are the measurements before silver application, while the remaining two show the enhanced conductivity after silver application. Figure 3.21(b) shows a magnification of the noisy and severely degrading I-V curves of the PEDOT-aluminium contacts before silver application in more detail. While the extent of degradation for aluminium contacts varied between samples, all contacts were restored to high conductivity with silver paint.

From the curves it is evident that the contact properties of the aluminium-silver-PEDOT contacts are vastly improved compared to the previous aluminium-PEDOT contacts. The contacts are ohmic (linear and through the origin), remain stable under repeated stress and show no signs of hysteresis. Therefore they fulfill the requirements for reliable operation of the protection grid. Given the significantly lower conductivity of the aluminium-PEDOT contact compared to the aluminium-silver-PEDOT contact, it can be assumed that practically all current flows from the aluminium to the silver and into the PEDOT. Probing the samples on the silver pads rather than the aluminium did not change the measured I-V properties. This leads to the conclusion that the contact resistance between silver and aluminium is negligible. In future experiments, silver pads are painted on PEDOT lines for probing without the need for aluminium contacts. To determine whether the PEDOT-silver contact introduced a significant ohmic resistance to the system, a four-point measurement was carried out.

3.5.3.2 Four-point measurement

The contact resistance is difficult to quantify with precision, as the resistance of the PEDOT lines is significantly higher than the expected contact resistance value. The top view of a four-point measurement setup is shown in Figure 3.22. In a standard four-point measurement, the contact resistance is estimated by measuring the current through the outer injection probes (labelled

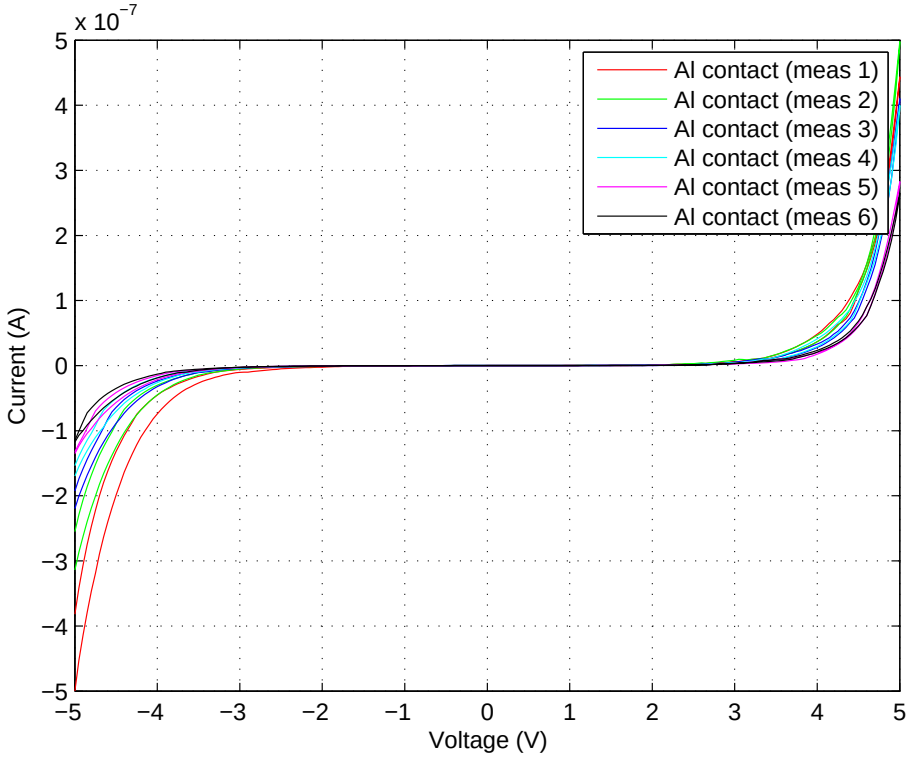

(a) Complete I-V curves

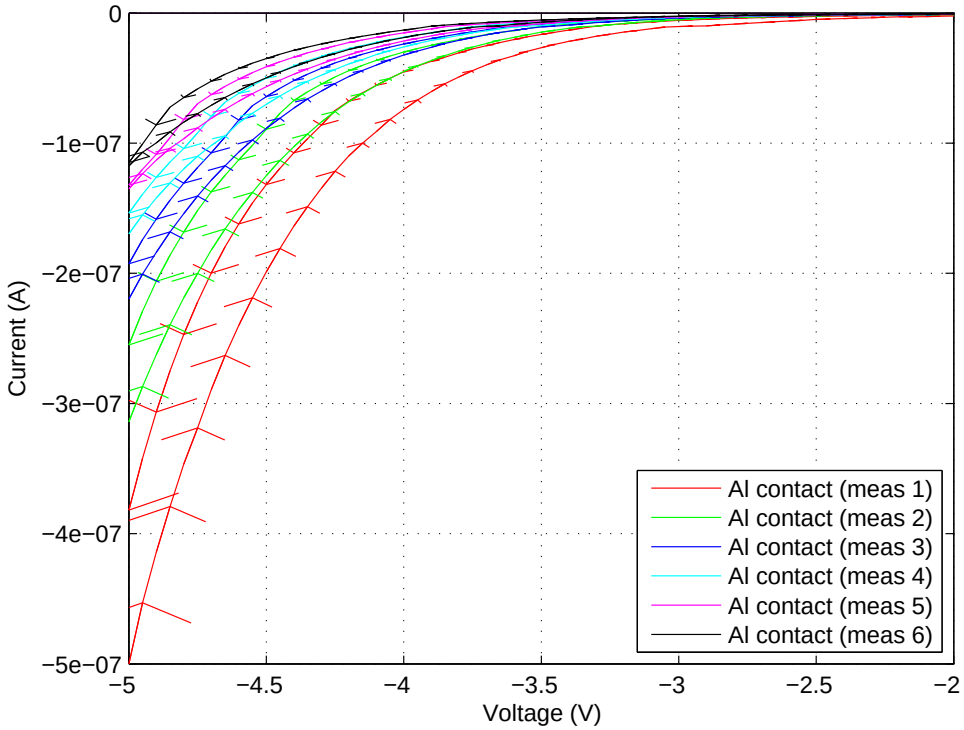

(b) Detail between -2 and -5V

Figure 3.19: Repeated I-V curves on single sample with aluminium contacts

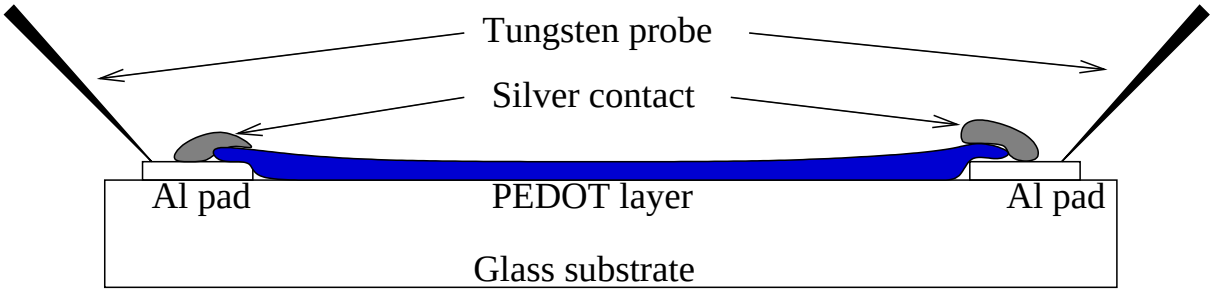

Figure 3.20: Side view of silver-enhanced contact

1 and 4), and the voltage drop between the two inner probes (labelled 2 and 3). The length between the inner probes gives an estimate of the resistance per unit length, allowing the line resistance between the outer and inner contacts to be subtracted. However, in the setup shown in Figure 3.22, the inner contacts are also silver padded. As silver has a low resistance per unit length, it is likely that the current flows through the silver contacts rather than the PEDOT lines, introducing further PEDOT-silver junctions and skewing the measurement. Using only tungsten probe needles for contacts 2 and 3 may also result in a distorted measurement. The mechanical fragility of PEDOT (reported in Section3.3.3.1) makes damage to the line due to the probe needle very likely. Scrapes in the PEDOT line result in an altered cross-sectional area, invalidating the measurement of resistance per unit length.

If the arrangement is reversed, the voltage is measured on contacts 1 and 4, and current is injected through contacts 2 and 3. In this setup, a small initial current would flow to charge the parasitic capacitances of line between inner and outer contacts, the contact pads, probes, and the measurement instrument. After charging these parasitics, no current flows between inner and outer contacts. Therefore additional junctions and changes in line geometry cease to play a role, resulting in a more accurate measurement result.

Similar to previous measurements, the voltage of the inner current injecting contacts was swept from -5 V to 5 V and back in steps of 0.1 V. The resistance of the contacts was determined for 12 samples. The contact resistance was found to lie in a range of 1 Ω to 115 Ω. Given that the voltage drop between the current injecting and measurement probes was of the order of several mV at absolute voltages of 5 V, measurement noise becomes significant. As the instrument is digital with a resolution of 16 bits per value, quantisation noise also plays a role, as well as any measurement offsets between inputs to the parameter analyser. In addition to the measurement accuracy for such small voltage offsets, it may also be possible that the resistance across the silver contact itself plays a role (probe position). Given the typical resistance of a PEDOT line of the order of 10 kΩ to several MΩ, contact resistances between 1 Ω and 100 Ω can be considered negligible. This result is further supported by the practically identical resistivity values measured with a four-point measurement and a two-point measurement with silver contact pads.

3.6 Epoxy encapsulation of PEDOT

While it has been previously established that a fully cured epoxy substrate has no effect on pure PEDOT, and a resistance-lowering effect on PEDOT with added DMSO, the effect of uncured epoxy resin must also be examined if lines are to be fully encapsulated. Encapsulation of PEDOT lines on different substrate materials was experimentally evaluated using Araldite Rapid Epoxy. A total of five independent measurements were carried out.

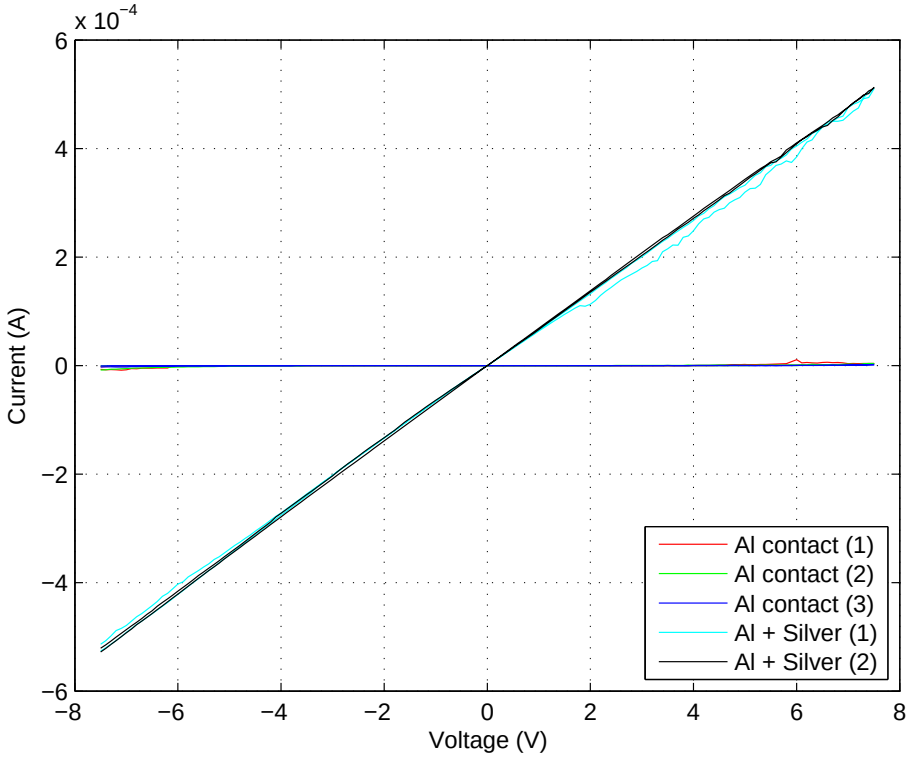

(a) Full I-V curves

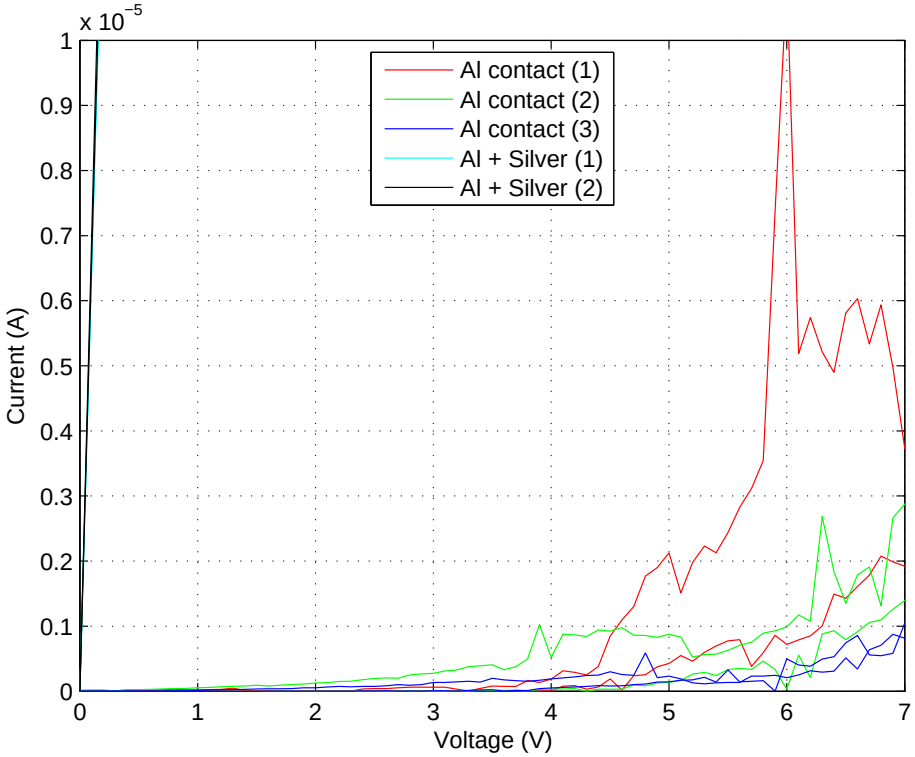

(b) Detail $0 - 10^{-5}$ A for positive voltages

Figure 3.21: Contact enhancement by application of silver conductive paint

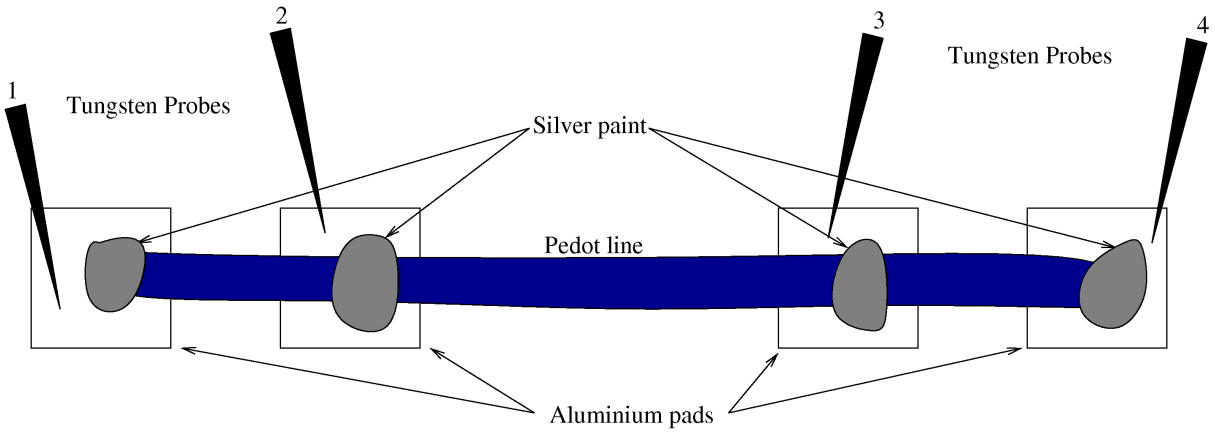

Figure 3.22: Top view of four-point measurement

date	sample	average change ratio
04^{th} July 2008	DMSO/Glass substrate	1.47
04^{th} July 2008	DMSO/Round glass substrates	1.85
04^{th} July 2008	DMSO/PVP substrate	2.33
04^{th} July 2008	DMSO/Epoxy substrate	1.49
25^{th} June 2008	DMSO/Glass substrate	5.04
25^{th} June 2008	pure/Glass substrate	5.85
25^{th} Jan 2008	DMSO/Glass substrate	1.59
12^{th} March 2008	DMSO/Glass substrate	0.51
12^{th} March 2008	pure/Glass substrate	0.94
18^{th} March 2008	DMSO/Glass substrate	1.01

Table 3.1: Resistance per unit length for epoxy coated and bare PEDOT samples

For the measurements carried out for the selection of substrates on 4^{th} July, the resistances were measured before and after application of epoxy to the same lines. For all other measurements, a group of randomly selected lines was covered with epoxy and the average resistances of the two groups were compared. The results for all experiments are shown in Table 3.1. All samples from the same day were printed and coated together. From the 04^{th} July experiments, it can be seen that the resistance change for lines on a PVP substrate is slightly higher than for other substrate materials, but the remaining samples are consistent. A comparison across experiments however shows great variation in resistance change. The cause for this variation is unknown, as many factors may have played a role, such as resin/hardener ratio, amount of water and DMSO remaining in the PEDOT line, ambient temperature, ambient humidity etc. As it is unlikely that Araldite is used as packaging material for a commercial security device, and other materials may have different effects, no attempt was made to identify the cause for the resistance change.

3.7 Conclusions

PEDOT composition For a good compromise between print quality, resistance per layer, and printing reliability, a 1:1 PEDOT:water dilution is recommended for this print head. Depending on the required conductivity, an addition of up to 20% DMSO is possible without excessive

degradation of printability. Some designs require high resistivity, such as RC delay lines, while other designs may require lower resistance for larger bias currents.

Surface properties For printing protection grids on top of standard CMOS microchips, both the interface layer method and using surfactant are viable options. As long as the interface layer is chemically inert and not dissolved by the ink composition, the resistivity of the printed PEDOT conductor lines is not changed. Altering the surface tension with surfactants is also possible, though in order not to affect the droplet ejection process it is better to coat the substrate with surfactant rather than adding it to the ink.

Contacts As predicted, it was found that the contacts between aluminium and PEDOT are problematic. The contact properties can be enhanced and stabilised with standard silver paint. While the silver paint was not inkjet printed in these experiments, inkjet printing has been reported in literature [Lee et al., 2005; Perelaer et al., 2008; van Osch et al., 2008].

It can be concluded that the fabrication of passive organic protection grids is feasible. The grid layout may be chosen flexibly, as the inkjet fabrication method does not require layout masks. Furthermore, inkjet printing is carried out in ambient laboratory conditions without the need for high vacuum conditions, therefore it is estimated that fabrication cost is less than that of a metal protection grid. It has been shown that with the appropriate printer settings and ink composition, PEDOT lines can be printed with predictable properties and a low defect rate. With the addition of DMSO, the conductivity range can be tuned to suit different readout schemes.

Chapter Four

Security of Tamper Protection Grids

4.1 Introduction

In the previous chapter, PEDOT was shown to have suitable properties as a conductor for organic tamper protection grids. It can be inkjet-printed onto a variety of materials with a low failure rate. The conductivity range can be adjusted to suit different protection grid designs. As it has been shown that fabrication is possible, the security properties of the PEDOT grids are the next step.

The evaluation of security properties is inherently difficult, as overall security is only as good as its weakest part. In the first instance, the security evaluation of this protection grids concept focuses on reliably achieving changes in characteristics of the protection grid. A full implementation of such a protection grid also requires all other aspects of the protection scheme (e.g. the readout circuit) to be secure, which is beyond the scope of this dissertation.

The depackaging methods that may be used by an attacker can vary widely, as will the attacker's skill level. Information about novel depackaging methods is not usually published by attackers. The security evaluation is therefore limited to known methods derived from microchip failure analysis. These depackaging methods do not by default consider the presence of protection grids, and thus may be too coarse for protected devices. If vulnerable protection grids become more common, there will be more incentive for attackers to develop custom depackaging methods targeted at defeating these. Any security evaluation is therefore only valid until new attack methods become known.

Apart from the difficulty in defining attack methods, there is also a large variety of different chip packages. Depending on the packaging type and material, different depackaging procedures may be used. The implementation details of the security device therefore influence the detection sensitivity, as the package may limit the choice of grid layout and size. The readout scheme and achievable measurement accuracy is also dependent on the security device and its layout. The implementation in a security device may be vulnerable to characterisation by non-invasive means. In particular, a time-delay measurement scheme may leave visible traces in the power consumption curve, which would reveal the delay time requirements to an attacker. These implementation specific vulnerabilities are material for further investigation, beyond the general evaluation of suitability of organic electronics as protection grids.

Prototype grid lines are tested against a variety of published depackaging methods, to assess and, if necessary, improve the sensitivity. Depackaging methods can be grouped into the broad

categories of physical and chemical methods. Physical depackaging methods generally involve mechanical means of material removal, as well as thermal means such as (local) evaporation. Mechanical methods such as polishing, grinding, or carving aim to remove the packaging material in layers, thus could stop just before damaging the protection grid. Other physical means such as drilling and thermal evaporation aim to create local holes in the packaging, which may be too small to detect by the protection grid. More global methods such as cracking the package open with a specialised tool [Beck and Wilson, 1998] may possibly remove the grid intact, so that it can be analysed separately. The most commonly reported chemical is fuming nitric acid which dissolves organic materials but spares silicon, silicon nitride, and aluminium [Anderson and Kuhn, 1996]. To rinse away the nitric acid, organic solvents such as acetone are used [Anderson and Kuhn, 1996]. For certain packaging materials, there are also specialised solvents [Beck and Wilson, 1998]. Chemical etching methods generally dissolve material with similar chemical properties at a certain rate, and spare other, dissimilar materials. If the protection grid material is inert to the chemical etchant, there is a danger of revealing the grid without damaging it.

To evaluate the general security properties of PEDOT protection grids, the sensitivity to depackaging and methods against probing must be considered. For high sensitivity, the resistance change of the lines must be measurable if the packaging has been tampered with. Preferably, the PEDOT lines should be cut completely. Probing of the PEDOT lines themselves must result in damage to the PEDOT lines, so that an attacker cannot read out the characteristics of the lines correctly.

4.2 Sensitivity to mechanical depackaging

Mechanical depackaging methods vary significantly depending on the skill of the attacker, and the tools or machines used. To be able to evaluate the sensitivity of PEDOT lines to these depackaging methods, a generalisation must be found.

For mechanical removal of the packaging, a force must be applied to the package in order to remove material. When cutting, drilling, or milling the package, shear forces are applied, as well as a compressive force to ensure the blade or grinding sheet penetrates the surface. Drilling is similar, except that the forces are applied in a circular motion rather than linear. When removing the security device out of a plastic carrier card [Kömmerling and Kuhn, 1999], the bending from carrier card exerts a tensile force on the mounting glue. To split the package apart [Beck and Wilson, 1998], compressive forces are applied to the sides of the device until the (brittle) packaging material fractures.

Generally, forces applied to ductile packaging material will result in deformation, which may in turn result in damage to the PEDOT line. The amount of deformation will vary, depending on the chosen packaging material and the depackaging method (and skill). Brittle packaging materials hardly deform prior to fracture, thus are not covered by this generalisation.

4.2.1 Resistance vs. strain

The most general measure of sensitivity to depackaging is the resistance change with strain. Even though the amount of deformation and size of affected region may vary for different mechanical depackaging methods and packaging materials, the resistance-strain relationship is a useful measure. If the deformation can be estimated, the total resistance change of a deformed grid line may be approximated by summing the local changes in resistance.

Under the assumption of a constant volume and constant resistivity, the resistance change with strain is expected to follow a quadratic law, if the material is strained in the direction of

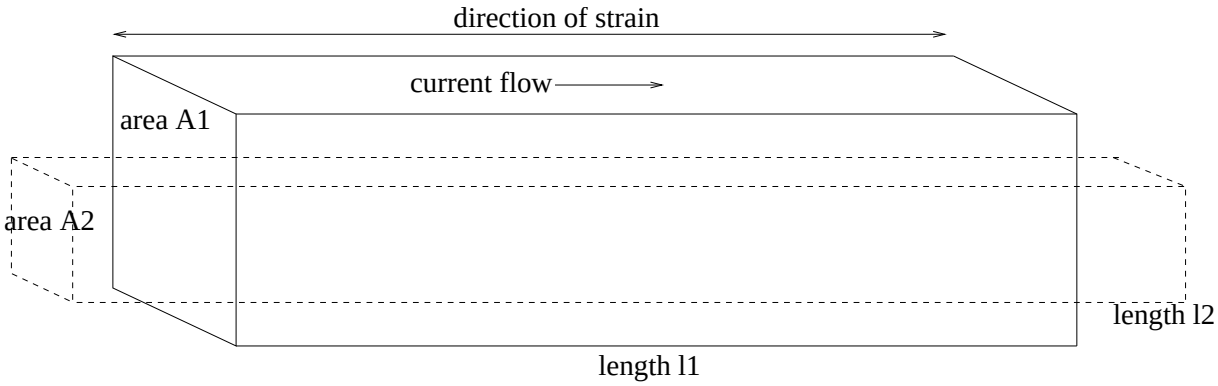

Figure 4.1: Wire strained in direction of current flow (inaccurate proportions)

current flow. Using the lengths and areas labelled in Figure 4.1, and the total volume V, the resistances R_1 (initial) and R_2 (final), as well as the material resistivity ρ:

$$V = A_1 \ l_1 = A_2 \ l_2 \tag{4.1}$$

$$A_2 = A_1 \ \frac{l_1}{l_2} \tag{4.2}$$

$$R_1 = \rho \ \frac{l_1}{A_1} \tag{4.3}$$

$$R_2 = \rho \ \frac{l_2}{A_2} \tag{4.4}$$

$$R_2 = \rho \ \frac{l_2}{A_1 \ \frac{l_1}{l_2}} = \rho \ \frac{l_2^2}{A_1 \ l_1} = R_1 \ \frac{l_2^2}{l_1^2} \tag{4.5}$$

Setup and procedure Standard PEDOT lines (1:1:0.2 PEDOT:water:DMSO, 40 µm drop spacing, 6 layers) were printed on ductile poly(ethylene) (PE) strips which allow large amounts of strain to be applied without cracking or yielding. The polyethylene strips were mounted on manual tensometers to apply linear strain. The resistance of the PEDOT lines was determined at intervals using a Fluke multimeter. The contact pads of the resistor lines were made from PEDOT, and the multimeter leads were coated with silver paint to achieve good contact. Silver contacts directly on the PE substrate were found to delaminate, as the dried silver paint is brittle and adhesion between the two materials is low.

Results The measured, normalised resistance-strain curve is shown on a log-log scale in Figure 4.2. For small strain (less than ca. 9%), the resistance change follows a square law, as predicted by the simplified theoretical derivation. Above the 9-10% threshold, cracks become visible in the contact pads, and delamination of the pads from the substrate begins. This coincides with the much more rapid increase in resistance, as well as the much larger error margins. Delamination takes place when the stress for a given strain exceeds the adhesion between the substrate and the PEDOT. As contact pads were significantly thicker than the PEDOT resistors, a larger stress was needed to strain the thicker material. Therefore it is plausible that the contacts delaminated while the lines remained attached to the PE substrate.

Discussion As PEDOT is a flexible material, the resistance change vs. strain is initially quite small before cracks are formed. Interpreting the results in a security context, the most problem-

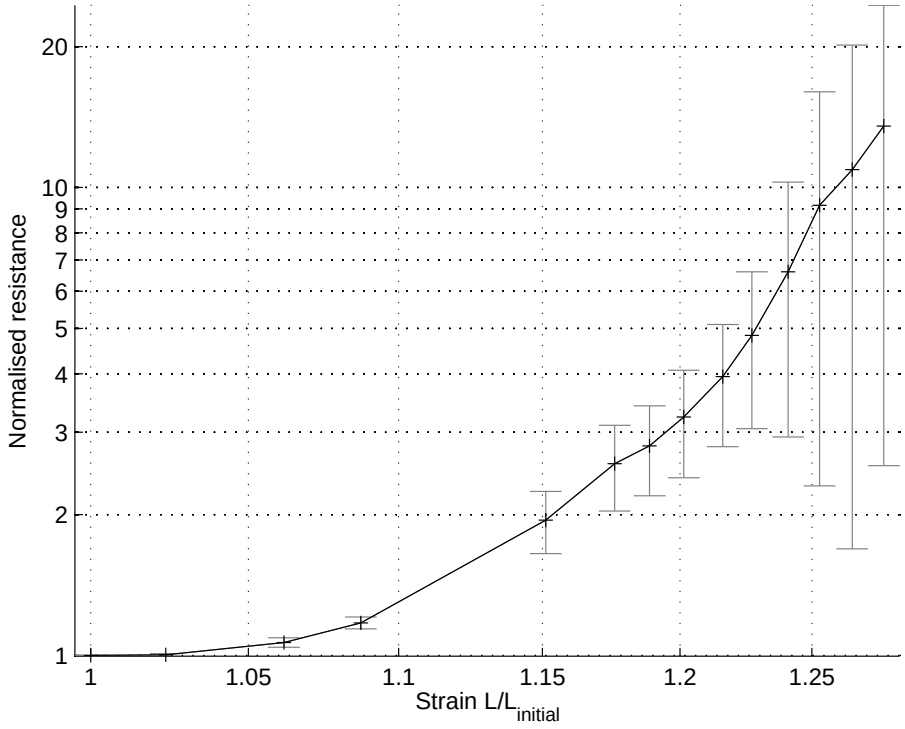

Figure 4.2: Normalised resistance vs. strain for 10 samples

atic types of damage to detect are either those localised to a small part of the protection grid, or small amounts of strain over a large area. It is evident that measurement accuracy is limited by the readout circuit and by contingency for environmental and degradation effects. Therefore a certain amount of physical deformation is tolerated by the security device. In the worst case it may be possible to penetrate the material below the grid with a probe needle without detection. To evaluate the extent of this potential problem, a further experiment is necessary to quantify the resistance change due to probe needle penetration.

4.2.2 Probing test

Setup and procedure Standard, 2 mm long PEDOT lines were printed on an epoxy substrate to represent a protection grid printed on a (topographical) base layer. The electrical resistance was measured with a HP4156A semiconductor parameter analyser before and after driving a tungsten probe needle into the epoxy below the PEDOT.

Results and Discussion The results of this evaluation are presented by means of two examples that illustrate what is probably the worst-case outcome. Figure 4.3 shows microscope images of two samples subjected to probing. For the first sample (Figure 4.3(a), detail in Figure 4.3(b)) the probe needle was placed immediately next to the PEDOT line. Despite the deformation, which is visible as wrinkles in the line, the resistance increased only by 1% from 60.3 kΩ to 61.0 kΩ. For the second sample (Figure 4.3(c), detail in Figure 4.3(d)), the probe needle was pushed below the PEDOT line, as is evident from the lighter area in the image. This more vigorous treatment resulted in the doubling of the sample resistance, from 550 kΩ to 1 MΩ. A small crack

is visible in the enlarged image of the PEDOT line (Figure 4.3(d)), which is the likely cause of the resistance change. However, it should be evident that a more delicate probing technique and thinner probe needles may result in undetected probing attacks.

4.2.3 Brittle substrate layer

While a soft substrate material deforms under an applied force, it has been shown that the resistance change caused in the PEDOT grid is insufficient to prevent probing through the substrate material. As the PEDOT is flexible, the deformation initially follows the shape of the substrate, only starting to crack after a threshold of deformation is reached. While the flexible substrate material has been shown to facilitate deformation of the lines, additional measures are necessary to ensure that the PEDOT lines are broken rather than strained. Low strain-to-failure is a characteristic of brittle materials, therefore a multi-layer structure of the flexible epoxy substrate, a thin brittle polymer and a PEDOT resistor on top will be evaluated. In this sandwich structure, the flexible epoxy material will deform under the applied force, and cause the brittle layer to crack. The cracked layer is expected to then break the PEDOT line. A suitable brittle polymer is AZ5214E photoresist, which can be spin-coated in thin layers.

Similar to the previous experiment, the result of this investigation will be illustrated by means of a typical example structure. As the depackaging and probing technique and even the choice of materials for a production device are not fixed and defined, it is impossible to precisely quantify the amount of damage that will be caused by an attacker. The results should be interpreted as a proof of principle rather than a precise proof of security.

Results A microscope image of a probed sample structure is shown in Figure 4.4. The sample was probed a total of three times, as can be seen from the three areas of damage. The resistance change after the first, light probing attempt (bottom left) was negligible, from 56 kΩ to 58 kΩ. As the second probing attempt resulted in a small crack running through the PEDOT resistor, the resistance change was more pronounced. The resistance changed sixfold from 58 kΩ to 359 kΩ. The third probing attempt, which was of roughly the same force as the probing carried out in the previous experiment (Figure 4.3), resulted in the larger crack through the PEDOT line and a complete loss of conduction.

4.2.4 Mechanical depackaging tests

After the improvement in sensitivity to probing by including a brittle base layer below the PEDOT line, two further tests were carried out to determine the sensitivity to depackaging attempts.

Setup and procedure PEDOT lines printed on a soft/brittle substrate were covered with a thin layer of soft epoxy. Two different depackaging tests were carried out. For the first test, an attempt was made to remove the top epoxy coating by shaving off thin layers with a sharp blade. In the second test, a tensile force was applied to the top and bottom parts of a structure, to pull apart the layers. As the PEDOT layer is placed on a separate layer in the middle, it sits directly at a boundary of packaging layers. If the adhesion between layers is less than the yield strength of the packaging materials, then the package may split to reveal the grid.

Results An image of a sample taken from each experiment is shown in Figure 4.5. The left hand image, Figure 4.5(a), shows a still-covered PEDOT line (vertical dark shadow) after a package thinning attempt. The cracks running through the PEDOT line are visible as thin black lines,

(a) Probe needle pushed next to PEDOT line

(b) Detail of deformation area

(c) Probe needle tunnel below PEDOT line

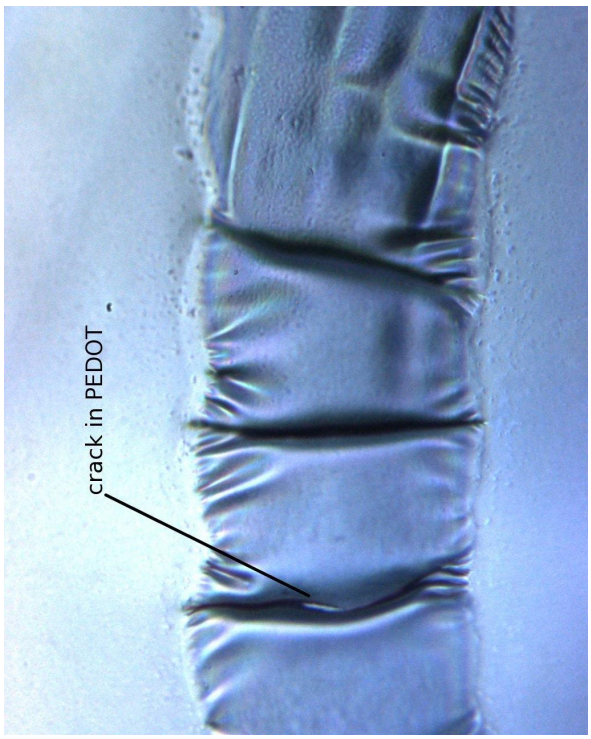

(d) Detail of deformation area

Figure 4.3: Probing below PEDOT lines on soft epoxy

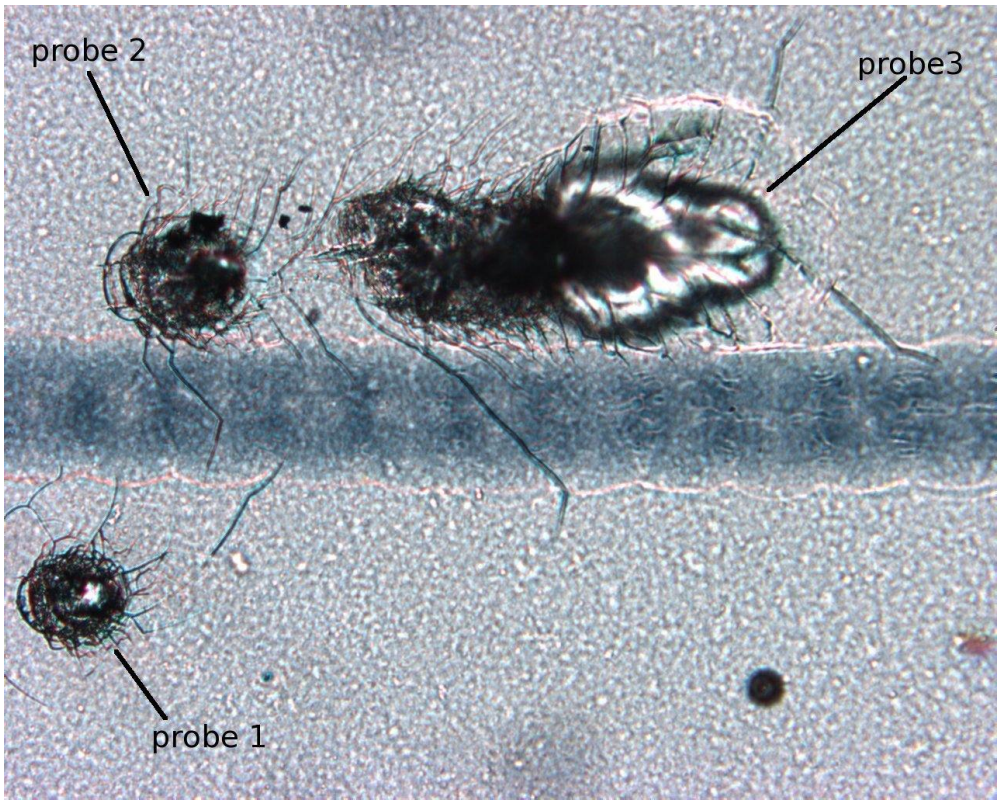

Figure 4.4: PEDOT on brittle sandwich substrate

and are orthogonal to the cutting direction. Depending on the direction, the cutting resulted in different amounts of compressive and shear forces on the material. The cracks shown in Figure 4.5(a) resulted in complete loss of conduction. Only a small force was applied to the sample to cause this damage, which did not remove significant amounts of material. None of the samples in this experiment were successfully depackaged with the PEDOT line intact.

In the second experiment, a tensile force, orthogonal to the substrate plane was applied to the epoxy material, succeeding in splitting it at the brittle centre layer. The sample in Figure 4.5(b) was only partially covered with epoxy, and the boundary between covered and bare parts is clearly visible from the damage boundary. While the brittle layer did not actually break the sample in this test, it can be seen from the image that the PEDOT line is destroyed. In other samples subjected to the tensile test, the entire PEDOT line was lifted off the substrate material. The adhesion of PEDOT to epoxy seems to be higher than to the brittle layer, causing the weak point and split to lie between the brittle material and the PEDOT. However, as the epoxy was deformed by pulling the packaging apart, the resistance of the samples that were removed in one piece changed by a factor of at least two, which will prevent an attacker from successfully analysing the removed grid lines.

4.2.5 Discussion

A brittle layer on the soft substrate epoxy greatly enhances the sensitivity of PEDOT lines to depackaging. The qualitative comparison of Figures 4.3 and 4.4 shows that the flexible substrate flows back into its previous shape after removal of the probe needle, whereas the cracks and damage remain with the brittle substrate. It can also be seen that the brittle material has cracked significantly in the immediate surroundings of where the probe made contact. If the

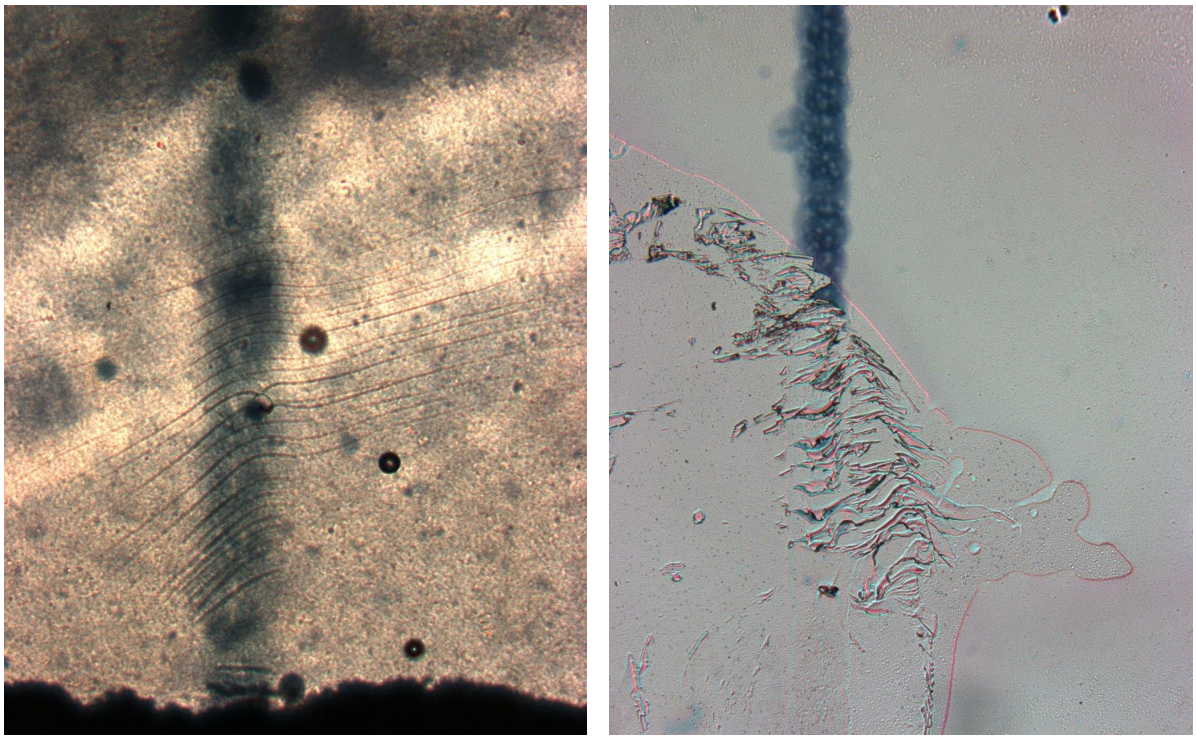

(a) Shaving away epoxy encapsulation

(b) Pulling apart multilayer packaging

Figure 4.5: PEDOT resistor lines after depackaging

substrate material and PEDOT resistor are cracked around the contact, then the measurement of properties is skewed or prevented altogether. This property of the substrate is useful, as it was described in Section 3.3.3.1 that a minimum amount of force and penetration is necessary for a probe needle to make good contact with PEDOT. If the PEDOT line is covered in epoxy when probing, it is also necessary to penetrate the epoxy material, which will require a larger force, likely to cause cracks.

For the depackaging test, three different types of force were applied to the sandwich structure to evaluate their influence on the integrity of the PEDOT line. While these tests do not directly represent the entire spectrum of mechanical depackaging methods, general conclusions can be drawn from the results of the test. The carving test was probably carried out more delicately than is achievable by automatic machinery. There is no guarantee that methods requiring less force applied to the packaging material cannot be found. However, as the brittle layer is cracked by the deformation rather than the applied force, the material properties of the flexible layers may be adjusted to increase sensitivity. A balance must be struck with handling robustness, possibly in the form of a hard shell covering the sandwich layer. From the pulling-apart test it can be concluded that the package splitting to reveal the grid is not necessarily a security problem. If the grid breaks in a seal-sticker like fashion, or the epoxy is sufficiently deformed to change the resistance of the grid, the properties of the PEDOT lines are irrecoverable.

It can be concluded that the ductile-brittle-grid-ductile sandwich structure shows good sensitivity to mechanical depackaging.

4.3 Sensitivity to acid depackaging

4.3.1 Fuming Nitric acid

Disolving in pure (fuming) nitric acid is the most widely reported method of chemical depackaging. It etches organic material without significantly damaging silicon, silicon nitride or the aluminium bond pads [Anderson and Kuhn, 1996; Beck and Wilson, 1998; Kömmerling and Kuhn, 1999]. Therefore, pure nitric acid was chosen to evaluate the sensitivity of PEDOT to acid exposure.

Setup and procedure Standard PEDOT resistor lines were printed on both epoxy and glass substrates. To test fully encapsulated grids, some samples were also covered with Araldite epoxy. One to two drops of room temperature nitric acid were administered to a sample using a pipette. After ca. 30s of exposure the sample was carefully dipped in DI water to remove the acid. Subsequently, the sample was blow-dried in a stream of nitrogen gas. As a control sample, a standard polycarbonate packaged microchip (Intel 440BX) was immersed in room temperature fuming nitric acid for several minutes.

Results Fuming nitric acid rapidly removes all traces of the PEDOT resistors both on glass and on the epoxy substrate. Microscope images of PEDOT lines partially exposed to fuming nitric acid are shown in Figure 4.6. In Figure 4.6(a) the PEDOT line was printed on a glass substrate. The effect of nitric acid on a PEDOT line is visible from the remnants of the line in the affected area.

For the polycarbonate reference, the etch rate at room temperature was significantly lower than for the Araldite epoxy and PEDOT lines. After several minutes of etching, the writing on the chip was still visible, indicating that only the surface of the polycarbonate packaging had been affected. Figure 4.6(b) illustrates the reaction of fuming nitric acid with Araldite epoxy. From the black spots in the image it is evident that the reaction releases a gas, which also causes physical motion in the epoxy layer. As the reaction rate was relatively high, and irregular due to the bubbles spreading the acid, it would be very difficult to terminate any acid depackaging of epoxy at precise depths. Apart from direct reaction with the acid, the PEDOT lines were also deformed and damaged by the flow of the surrounding dissolved epoxy.

4.3.2 70% nitric acid

As the pure nitric acid dissolved the PEDOT instantly, a lower concentration of nitric acid was also tested to evaluate the effect of dilution on PEDOT. Concentrated nitric acid (70%) contains water, and thus attacks aluminium [Kömmerling and Kuhn, 1999]. While it may not be used directly as a depackaging chemical, the concentration of the initially pure acid will drop as it consumes the packaging material. It is therefore necessary to evaluate lower concentrations of acid to determine whether they cause damage to PEDOT.

Setup and procedure As in the previous experiment, PEDOT lines were printed on glass and epoxy substrates. Each sample consisted of several PEDOT lines, of which one was initially padded with silver contact pads to determine the initial resistivity of the PEDOT. After completely covering the lines in a small amount of nitric acid for ca. 30s, the samples were carefully rinsed in DI water and blown dry. After acid immersion, the remaining lines were padded, and the resistivity was determined using the HP4156A. The reason for comparing the resistivities of two different lines on the same sample is that the silver contacts may be affected by the

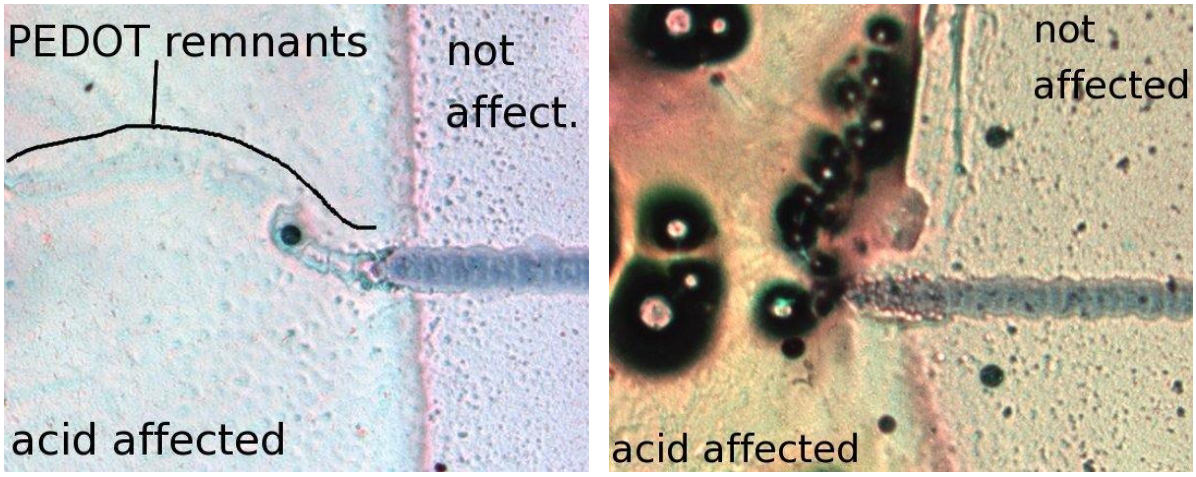

(a) PEDOT line on glass subjected to fuming nitric acid (b) PEDOT line on epoxy subjected to fuming nitric acid

Figure 4.6: Fuming nitric acid test

nitric acid, distorting the result. The resistivity of 12 PEDOT lines was measured before acid exposure, and for 16 PEDOT lines post-exposure.

Results The 70% concentrated nitric acid does not visibly etch away the PEDOT lines, and only has a small effect on the epoxy substrate (white discolouration, indicative of etching at the surface). It was found that the I-V characteristics of the initially padded samples measured after acid exposure were less linear (and more noisy) compared to the freshly padded samples. This indicates that the acid indeed affects the silver contacts, which may be an interesting result for a protection grid system, but in this context represents a distortion of the measurement. While there was no visual degradation of the PEDOT lines, it was found that the resistivity of PEDOT decreased by a factor of 2.5, from an average resistivity of $2.5 \quad 10^{7}\Omega/\mathrm{m}$ layers down to $9.97 \quad 10^{6}\Omega/\mathrm{m}$ layers.

4.3.3 Discussion

The rapid etching of PEDOT by fuming nitric acid is the best possible sensitivity that can be achieved. For the less agressive concentrated nitric acid, the considerable change in resistance should also be sufficient for detection (larger than the error margin). The lower resistance after acid exposure is not unique to nitric acid, but was also reported for diluted hydrochloric acid (HCl) treatment of PEDOT, where the resistivity was reported to have decreased by a factor of 45 [Nguyen et al., 2004]. A photoelectron spectroscopic evaluation of PEDOT:PSS reported the presence of poly(sodium 4-styrenesulfonate) in pristine PEDOT:PSS, which was converted to free PSS under the influence of HCl [Greczynski et al., December 2001], increasing the amount of dopant present. As the decrease in resistance is not a unique feature of nitric acid, it is concluded that acid depackaging can be detected by a PEDOT grid.

4.4 Sensitivity to solvents

Depending on the packaging material, a range of different chemicals beyond acids may be used for the depackaging process. The choice of packaging material is made individually for each security device, and depends on the application as well as the cost constraints. The solubility

of the packaging in solvents, and relevant solvent combinations may thus vary. For a general overview over sensitivity to solvents, the influence of the three most common laboratory solvents on PEDOT conductivity are evaluated. The evaluated solvents are DI water, isopropanol and acetone, which are very mild solvents. If PEDOT is sensitive to these solvents, or if sensitivity can be achieved, then it is likely that similar results would be found for other solvents.

4.4.1 Setup and procedure

Standard 6-layer PEDOT/DMSO lines were printed on the four substrate materials evaluated during PEDOT fabrication: Clean glass (as a prototype for a ceramic substrate), Araldite epoxy (as a prototype for an epoxy base layer), PVP (due to its use as OTFT gate insulator), and AZ5214E photoresist (brittle layer, easy patterning). The electrical resistance of each sample was determined before and after solvent exposure. Each sample comprised a set of individual PEDOT lines of ca. 2cm length with silver contacts placed at the ends. For direct interaction of the solvents with PEDOT, the lines were not encapsulated.

For the water test, a drop of water was placed in a glass petri dish, and the sample was placed face-down in the water. This method ensured uniform wetting compared to placing individual drops on the substrates, which contracted to larger drops on the hydrophobic surface. Water has a low evaporation rate in air at room temperature, compared to acetone and IPA. For a one-minute exposure cycle, the water-exposed samples were blow-dried in a stream of nitrogen.

For isopropanol and acetone, the surface tension is significantly lower, allowing the solvent to spread over the substrate. The evaporation rate of both solvents in air at room temperature is also sufficiently fast, and did not require nitrogen to dry the samples. One to two drops of solvent were placed on the PEDOT lines using a pipette. The samples were allowed to dry naturally, which also minimises the solvent flow over the sample.

4.4.2 De-ionised water

PEDOT ink is only available as an aqueous dispersion [Starck, b], therefore the PEDOT particles have already been in contact with water during deposition. While it is unlikely that a chemical reaction will take place, it may be possible that PEDOT particles are re-dispersed, or the PSS dopant is dissolved and removed with the water flow. It was also demonstrated in Section 5.2.2, that the humidity contents of the PEDOT lines reduces the resistivity, which may temporarily alter the measured properties until the lines are fully dry. The result of this experiment is also useful to estimate the effect of rinsing away the nitric acid in the previous experiment.

Results The strongest sensitivity to water was found using a glass substrate. As soon as a drop of water was placed on the sample with a pipette, the flow of the spreading drop was sufficient to remove and break the PEDOT lines. Using the immersion procedure to cover the entire surface also removed the PEDOT lines from the substrate. A microscope image of formerly parallel PEDOT lines subjected to water is shown in Figure 4.7(a). Folding of the PEDOT lines is visible in the centre of the image. However, it is also evident that the lines did not re-disperse and disintegrate, but are undercut by the water and floated away whole or in segments. The primary reason for line breakage is not re-dispersion, but lack of adhesion to the surface. A similar test with a silicon nitride substrate gave the same result.

The printed PEDOT lines showed greater adhesion to remaining substrate materials. Correspondingly, the resistors are significantly less vulnerable to damage when flooded with water. Both substrates initially contained six samples each, however three of the samples on the epoxy substrate were partially washed away upon water exposure. No lines printed on PVP

were washed away, which is plausible as the DMSO partially dissolves PVP, facilitating the integration of sample and substrate. Figure 4.8 summarises the resistance change of PEDOT with exposure time in water for samples printed on PVP (blue) and epoxy (black). Samples printed on PVP show a lower damage rate compared to the samples on epoxy. The reason for this difference may also lie in the integration of PVP into the PEDOT line, which may aid in binding the line together. However, a doubling of resistance in two to four minutes of exposure may not necessarily be sufficient for reliable detection.

Samples printed on AZ5214E photoresist were found to be destroyed within the first minute of exposure. Figure 4.7(b) shows the extensive cracking of the photoresist in the vicinity of the PEDOT line (with one island of the photoresist floated away). The reason for the cracking of the photoresist substrate material is unknown, though it provides a mechanism for sensitisation.

4.4.3 Acetone

Acetone is a polar solvent commonly used in the laboratory for cleaning and de-greasing purposes. It is also used to rinse away the nitric acid from depackaging [Anderson and Kuhn, 1996; Beck and Wilson, 1998; Kömmerling and Kuhn, 1999], therefore it is important to determine whether acetone has an effect on the conductivity of PEDOT.

Results Acetone has only a very small effect on the conductivity of PEDOT, if any. The graph of normalised resistance vs. acetone exposure is shown in Figure 4.9. Compared to the gradual doubling of resistance under the influence of water, the resistance of PEDOT changed by less than 10% on average in the same amount of time in acetone. The silver contacts were damaged by the acetone, as the silver particles were re-dispersed and spread over the sample, contributing to measurement variation. The samples printed on the epoxy substrate in particular were unusable after the third acetone exposure cycle.

For samples printed on AZ5214E photoresist and on PVP, the lines were destroyed after the first contact with acetone. Microscope photos of typical PEDOT samples are shown in Figure 4.10. Acetone is a solvent for both substrate materials, therefore the lines are undermined and break as they are moved with the flow of material. The effect is similar to the previous experiment with exposure of PEDOT to DI water on a glass or silicon nitride substrate, in the sense that the lines are physically broken. The main difference is that the entire substrate layer is removed in this experiment, therefore the adhesion of PEDOT to the substrate does not play a role. This mechanism is also utilised advantage of in lift-off lithography (Section 3.5.1), where a pre-patterned photoresist layer is covered with a blanket layer of material to be patterned. When the photoresist is removed, only the material deposited directly on the substrate (without photoresist) remains. The solvent for lift-off with AZ5214E is acetone.

4.4.4 Isopropanol

Isopropanol is another common laboratory solvent. As it is the solvent used for spin-coating PVP, similar lift-off behaviour is expected for the PVP substrate.

Results The lift-off effect was confirmed for PVP substrates, as can be seen in Figure 4.12(a). Figure 4.11 shows the normalised resistance vs. exposure time to isopropanol for glass and epoxy substrates. Similar to acetone, PEDOT has very low sensitivity to isopropanol for samples on glass and epoxy. The increase in resistance after 4 minutes of exposure is around 20% on glass, which is probably too low for detection. Interestingly, all samples printed on epoxy were broken after the third exposure to isopropanol. While no damage was immediately visible (Figure 4.12(b)), tiny, ca. $1\mu m$ wide microcracks were found on inspection at higher 1000

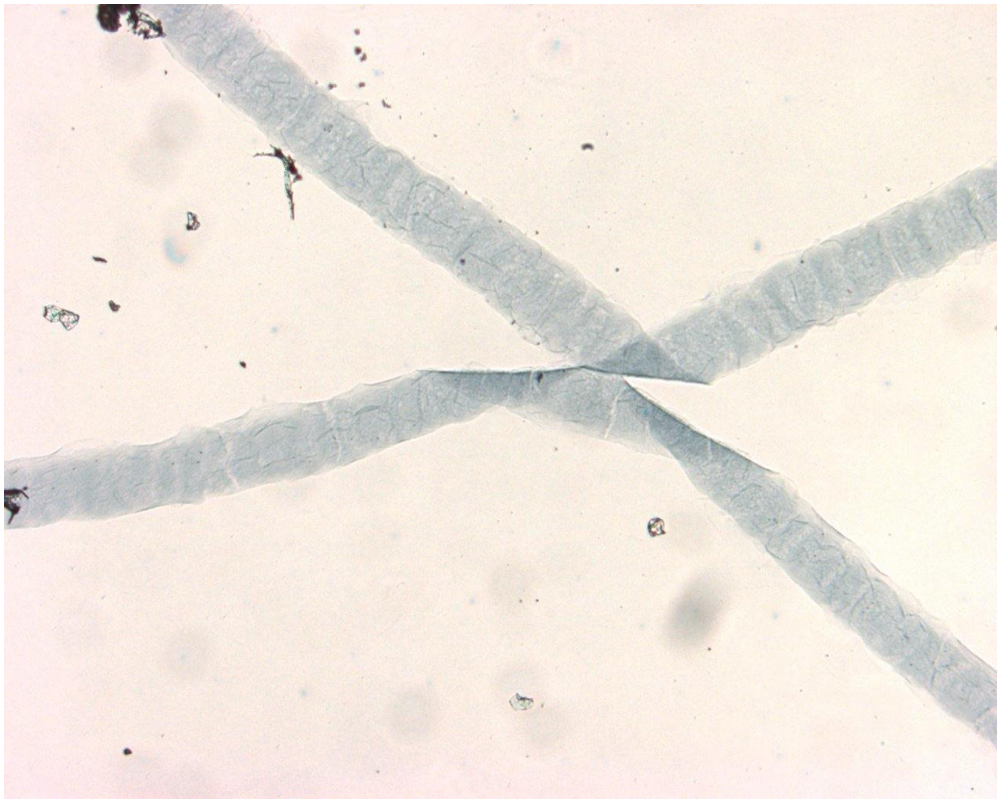

(a) PEDOT line on glass

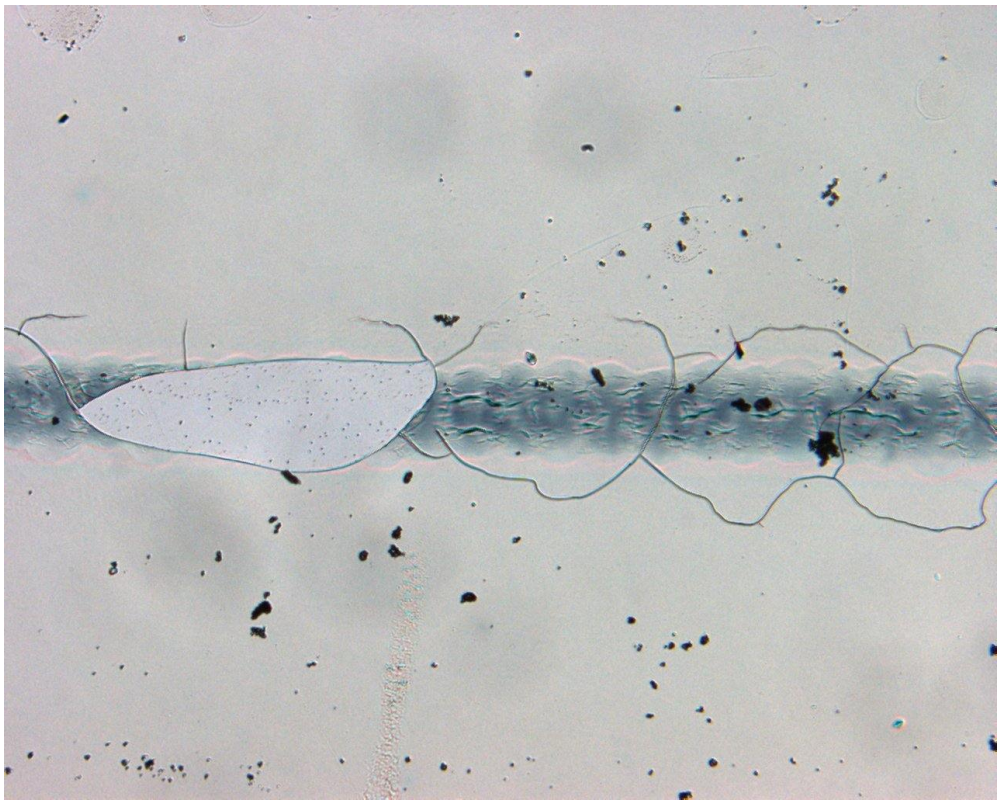

(b) PEDOT line on AZ5214E photoresist

Figure 4.7: PEDOT lines after water exposure

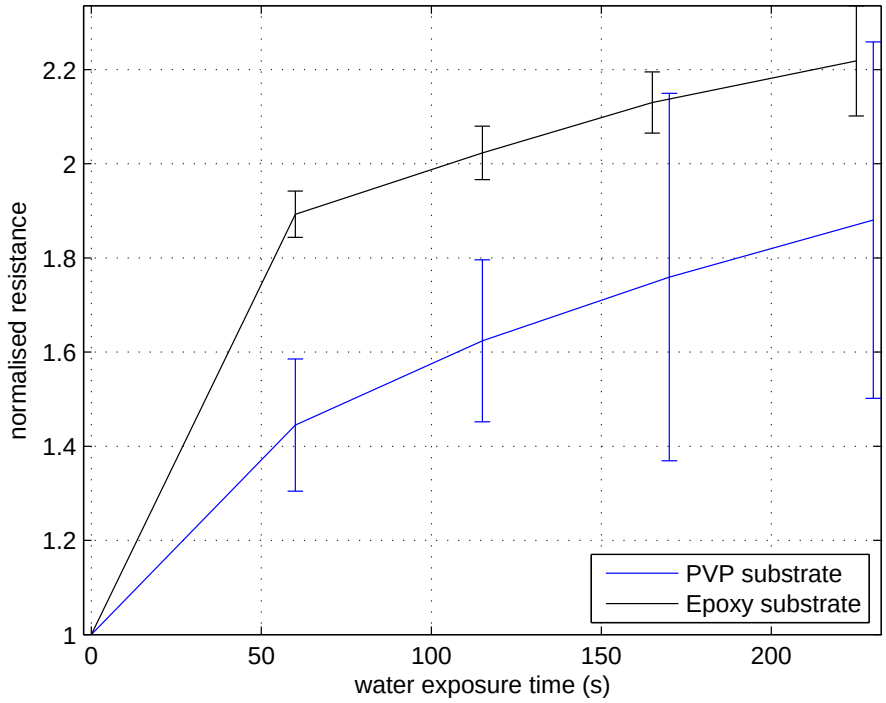

Figure 4.8: Normalised resistance of PEDOT vs. exposure to water for various substrate materials

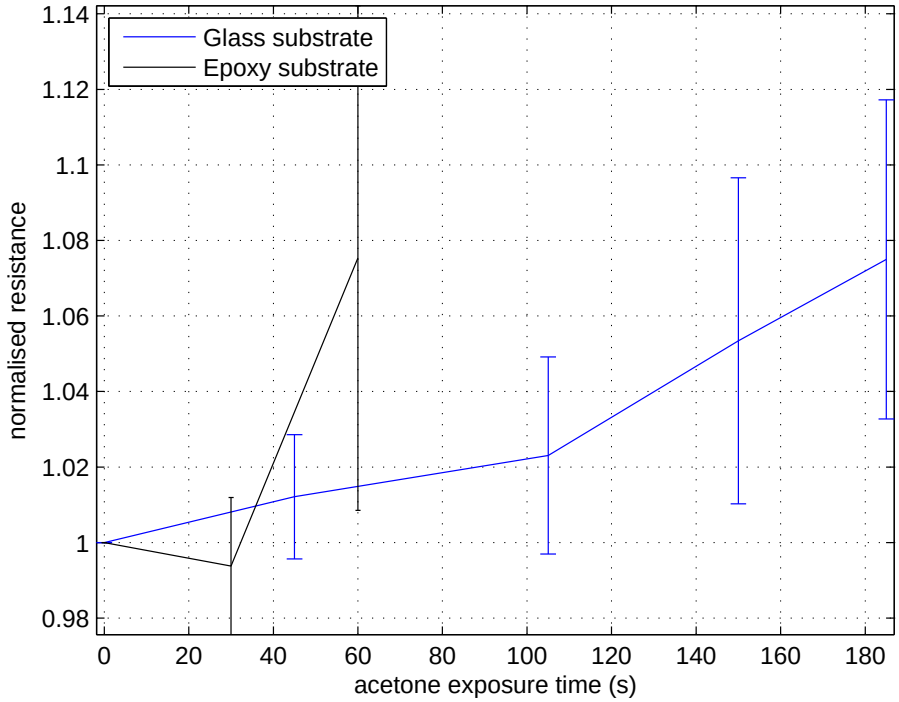

Figure 4.9: Normalised resistance of PEDOT vs. exposure to acetone for various substrate materials

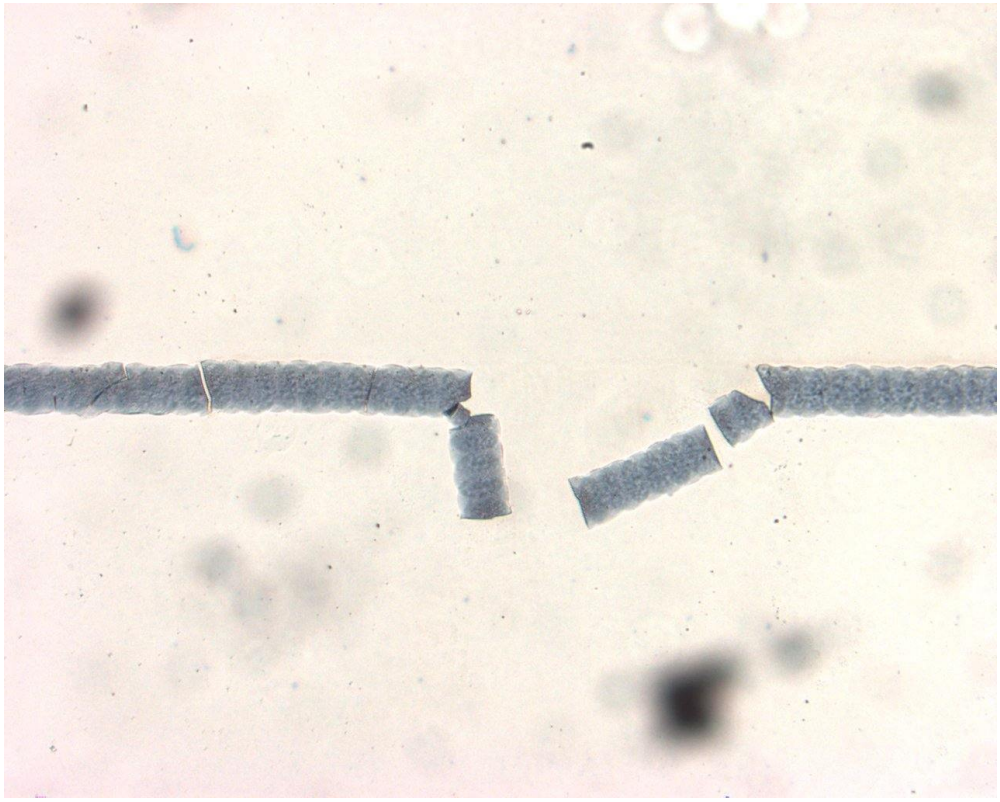

(a) PEDOT line on PVP

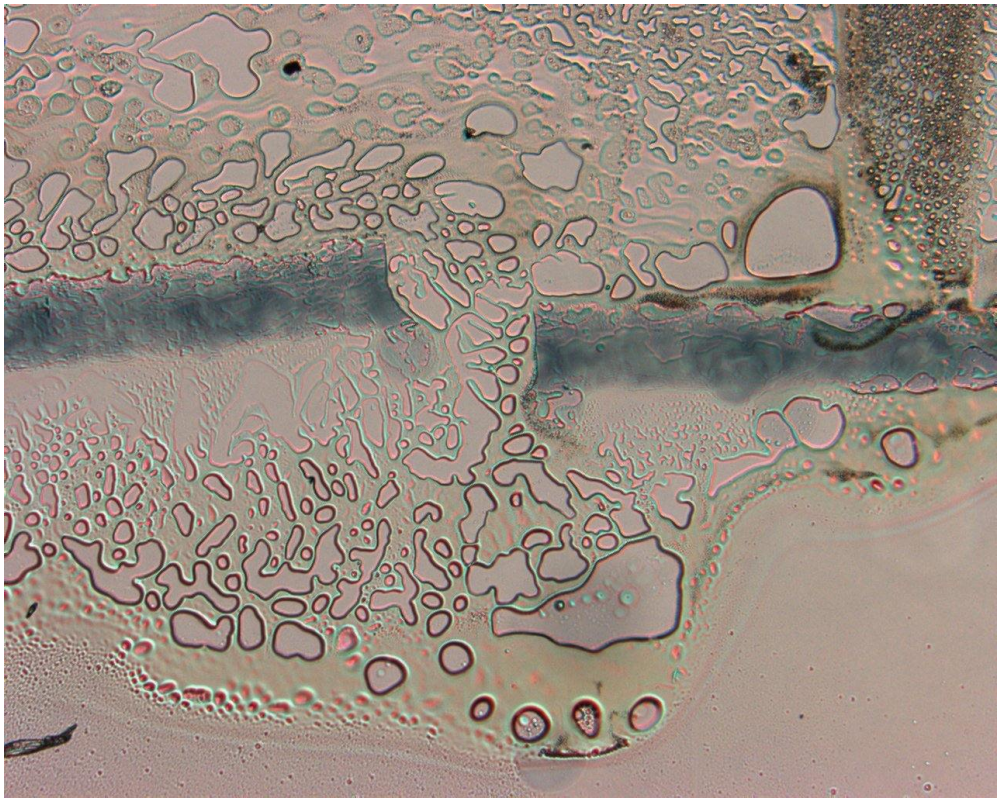

(b) PEDOT line on AZ5214E photoresist

Figure 4.10: PEDOT lines after acetone exposure

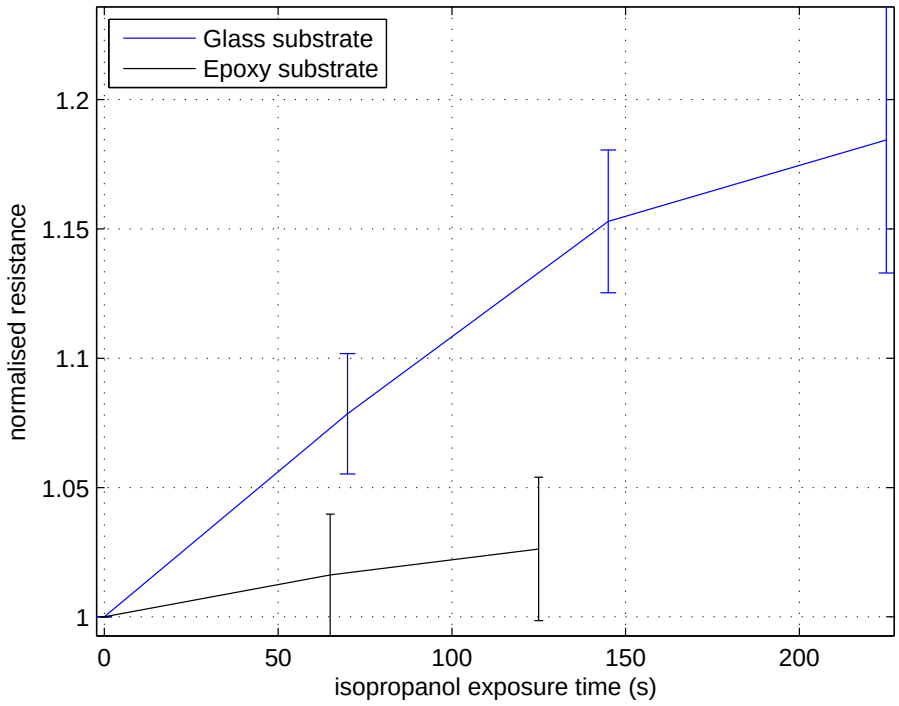

Figure 4.11: Normalised resistance of PEDOT vs. exposure to isopropanol for various substrate materials

magnification (Figure 4.12(c)). Similar microcracks were also found for samples on AZ5214E photoresist, directly after the first exposure (Figures 4.12(d) and 4.12(e)).

The cause for these cracks is unknown, so far. Such localised, narrow cracks could possibly be explained by local substrate swelling due to absorption of solvent. The swelling of the substrate may have strained the lines causing the cracks to form. After drying, the substrate would shrink back into its original shape, closing the cracks enough to hide them from visual inspection. The electrical contact is lost permanently, as a small gap remains. This model would however require significant substrate swelling, as it was shown in Section 4.2.1 that more than 10% strain is required to form cracks in PEDOT. As no definite cause for the formation of cracks can be identified, this damage method cannot be exploited to deliberately increase the sensitivity of the protection grid.

4.4.5 Discussion

While the solvents used for these experiments may not be exhaustive, some general guidelines may be derived. The sensitivity of PEDOT to the three example solvents is significantly lower than for acid. The highest sensitivity was found for water, which is unsurprising, as the PEDOT is delivered as an aqueous dispersion. It is to be noted that any resistance change is most likely the result of physical damage to the line, rather than a chemical reaction with PEDOT.

The physical damage mechanism may also be used deliberately to sensitise PEDOT to certain solvents. Using an appropriate substrate material that is soluble in the target solvent will cause the PEDOT lines to be undermined and broken upon exposure. If PEDOT is inert to the solvent used to dissolve the encapsulation layer, sensitisation may be required. In this case,

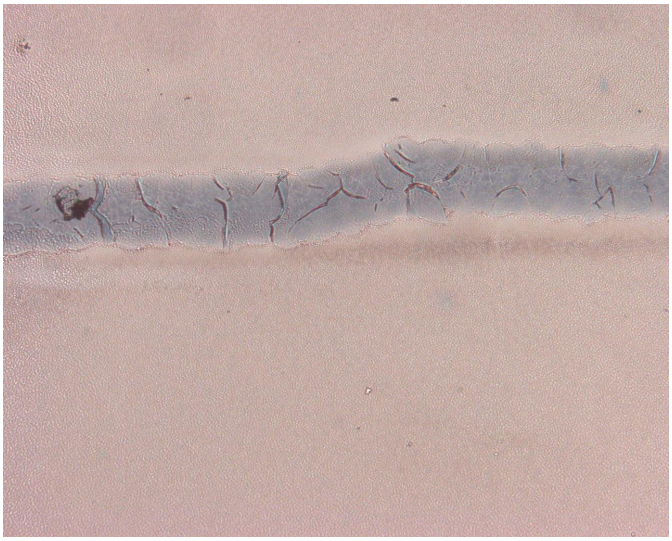

(a) PEDOT line on PVP

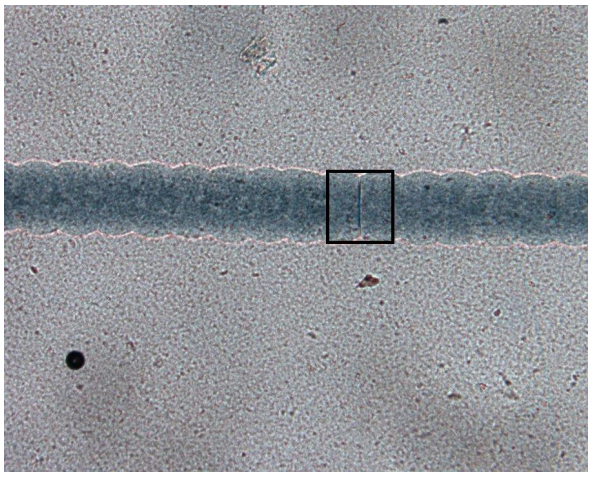

(b) PEDOT line on epoxy

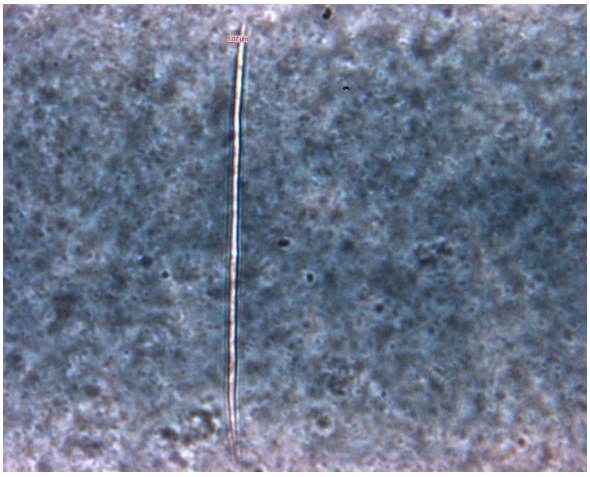

(c) PEDOT line on epoxy (detail)

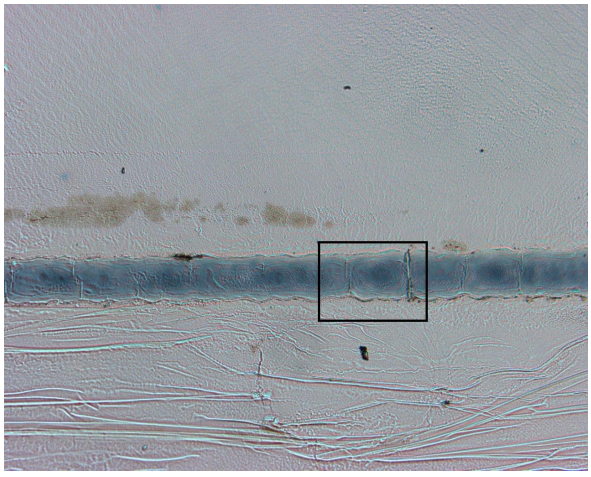

(d) PEDOT line on AZ5214E photoresist

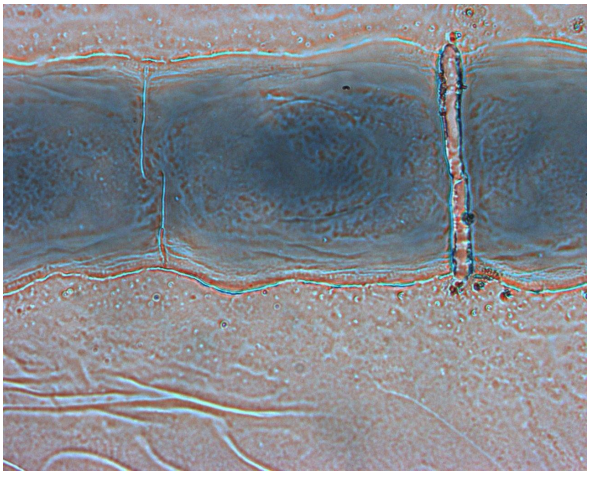

(e) PEDOT line on AZ5214E photoresist (detail)

Figure 4.12: PEDOT lines after isopropanol exposure

the base layer must also consist of a material with similar solubility as the top encapsulation to ensure that the PEDOT lines are broken physically upon exposure. In conclusion, chemical depackaging can be prevented reliably by the appropriate choice of packaging layers.

4.5 Laser depackaging

The final depackaging method in the security evaluation is laser ablation, as this method is independent from mechanical or chemical means. It is not covered by the guidelines derived in the previous sections.

4.5.1 Setup

PEDOT lines were printed on glass and epoxy substrate using the standard parameters. Several lines were encapsulated with Araldite epoxy. A New Wave laser ablation tool was used to test the response of the PEDOT resistors. The tool comprises a class 3b laser mounted in the camera port of a microscope. The microscope is used to focus the laser onto the sample, and allows the progress of irradiation to be monitored. The laser can be tuned to one of two frequencies, UV and green. A PC-controlled X-Y-stage allows positional control and automated cutting. The square spot size of the laser was around 50 µm to 60 µm. The resistance of the PEDOT lines pre- and post-exposure was determined with a Fluke multimeter.

4.5.2 Procedure

Two tests were carried out with the laser setup. Initially, the response of uncovered PEDOT lines to laser radiation was examined. As the PEDOT lines were evaporated with clearly defined edges, a test was carried out to estimate whether the remaining PEDOT was also affected by the radiation. Parts of the PEDOT lines were ablated, and the resistance change of the line was determined before and after the partial removal. The measured resistance change was compared to an estimate based on the altered geometry. The geometry of the ablated part was determined using the calibrated microscope, and the probably inaccurate assumption of a constant line thickness. To reduce the error caused by the assumption of a constant cross-section, material was removed from the centre to the edge of the line (estimated). Removing half the line relaxes the assumption about the line thickness to require only a symmetrical cross-section rather than a constant thickness.

In a second experiment, the epoxy-covered lines were examined to evaluate whether the epoxy encapsulation may be removed without damaging the PEDOT resistors. The highest intensity is in the focal point of the microscope, therefore the cutting depth of the laser ablation tool is limited. The focus of the laser was therefore also varied between the PEDOT line and the surface of the epoxy encapsulation.

4.5.3 Results

Cutting of the non-covered PEDOT lines was possible at high accuracy using both UV and green wavelengths. The resistance of the lines before and after cutting was found to be consistent with the change in line width (cross-sectional area), both for green and UV. For example, one sample changed resistance from 71.8 kΩ to 82.9 kΩ (UV), with an estimated final resistance of 79.4 kΩ, and another from 70.4 kΩ to 78.7 kΩ (green), estimated at 75.9 kΩ. As both lines were cut slightly beyond the centre, the estimate of the removed cross section was too low, underestimating the resistance change. The line is expected to be thickest towards the centre, from the hemispherical geometry of the sessile drop drying on the substrate [Hu and Larson, 2002]. The resistance

change is therefore deemed to be consistent within calculation tolerance, and an indication of only local damage to PEDOT lines under laser radiation.

For the epoxy covered lines it was found that the dark blue PEDOT was removed preferentially over the transparent epoxy. Figure 4.13(a) shows a single shot of the green laser on a PEDOT line completely surrounded by epoxy (encapsulation and substrate). The boundary of the laser can be seen from the cut-out area in the dark blue PEDOT line. The laser aperture was square, therefore the actual laser spot extends to below the PEDOT line. However, the damage to the epoxy coating is centered on the PEDOT, which leads to the conclusion that the evaporation of PEDOT may have caused the epoxy to burst rather than to evaporate. It can also be seen that the substrate material was also not affected by the laser from the lack of boundary edge between the irradiated areas and the non-affected areas. The same observation was made for non-encapsulated PEDOT, where the lines were removed without damage to the epoxy substrate.

The PEDOT lines printed on glass could be evaporated through the epoxy coating (UV and green). Figure 4.13(b) shows a sample PEDOT line with three laser shots. For the rightmost shot, it can be seen the PEDOT was removed without breaking the epoxy encapsulation. For the other shots, the encapsulation was thinner, hence the epoxy burst open. Depending on the thickness of the epoxy, focusing the laser on the surface of the epoxy also caused damage to the PEDOT lines if the intensity of the laser at the level of the PEDOT was still sufficiently high (the depth depends on the microscope optics). As the epoxy is transparent, most of the laser energy is transmitted rather than absorbed into the material. Damage to the epoxy is initially localised at imperfections (bubbles) and dust particles, until debris makes the epoxy opaque.

4.5.4 Discussion

The experiments show that very accurate cuts can be made into PEDOT using lasers. The second point is that transparent encapsulation materials transmit the radiation, thus allowing PEDOT to be removed preferentially. From an attacker's point of view, it may be possible to cut small enough holes into PEDOT grids to avoid changing the resistance above the detection threshold. However, for a transparent packaging material and a multi-layer grid it may be more difficult to drill a hole through the packaging without detection. Apart from the difficulty of coupling the laser into the epoxy, the bursting effect caused by the evaporation of the buried PEDOT may cause more severe damage to higher lying grid lines.

In the design phase of a security device, a balance must be struck between transparent material for difficult deposition of laser energy, or opaque material to obscure the location of the grid lines. Given that a transparent epoxy gives an attacker information about the location of the grid lines and the microchip, the material will most likely have to be opaque.

4.6 Conclusions

Mechanical sensitivity The sensitivity of PEDOT lines to mechanical depackaging depends very much on the substrate material. If the substrate material is ductile or flexible, then some mechanical deformation is tolerated by the PEDOT without significant resistance change. For protection layers this means that the grid may not detect localised depackaging. However, on a brittle substrate, damage is enhanced, since lines are broken as soon as the substrate material is cracked. The best protection was found to be a sandwich structure with the PEDOT line on a brittle layer between two soft layers. A thin brittle layer is also useful to make probing more difficult, as the force of a probe needle is likely to crack the brittle material before good contact can be made with the PEDOT.

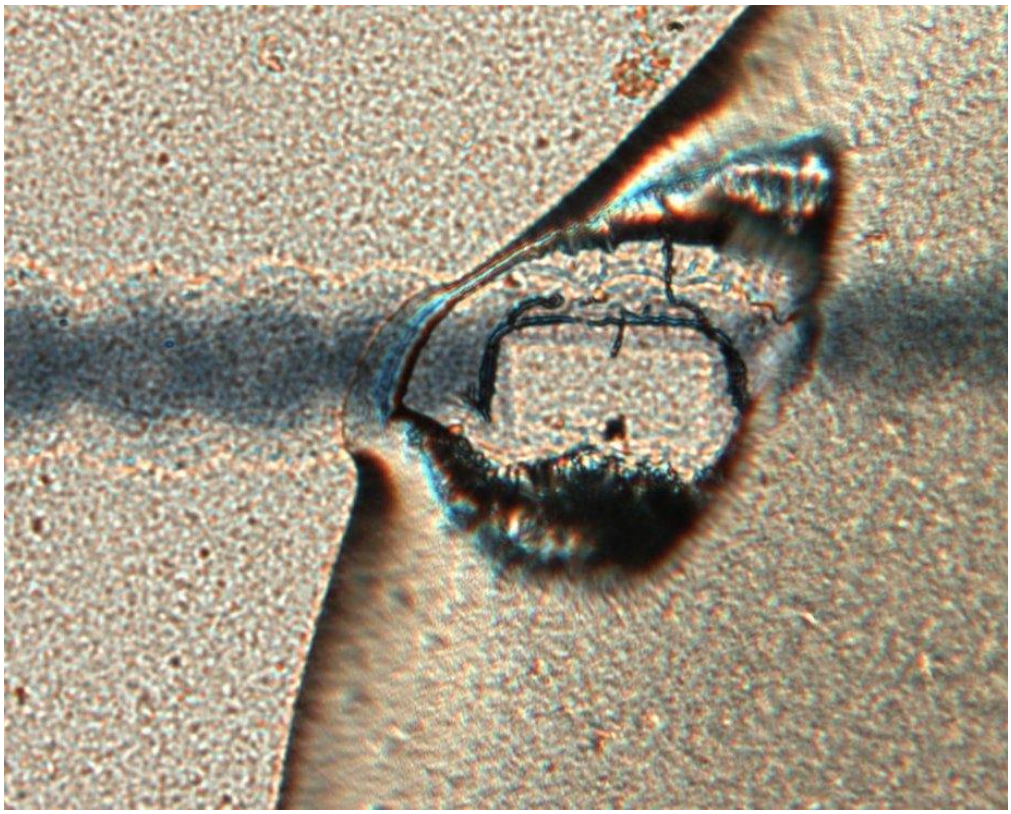

(a) Epoxy substrate

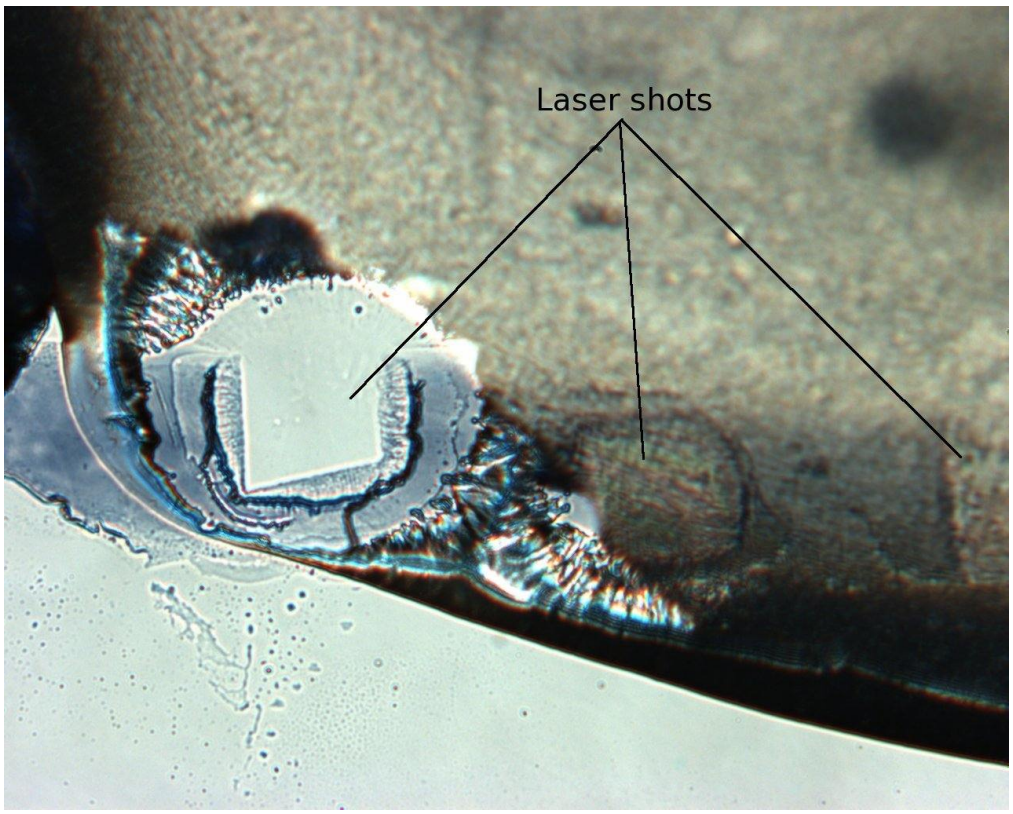

(b) Glass substrate

Figure 4.13: Laser depackaging test

Chemical sensitivity PEDOT is sensitive to the most commonly reported depackaging chemical, nitric acid. Pure nitric acid instantly removes the PEDOT, while concentrated nitric acid increases the conductivity of PEDOT measurably. PEDOT is not sensitive to the tested solvents: acetone, isopropanol and water. However, sensitivity to any solvent may be achieved if the substrate material of the PEDOT line is deformed by the solvent, breaking the PEDOT line physically.

Laser sensitivity PEDOT is easily ablated by laser radiation. It is possible to cut holes into the PEDOT lines with sharp definition, and without damaging the surrounding material. If coated by a transparent material, the PEDOT may be removed through the packaging. If the transparent packaging material is sufficiently thin, the evaporation of the buried PEDOT line causes a larger hole in the package than the diameter of the laser beam. For multi-layer protection grids the evaporation of a lower layer might that way destroy the higher layers.

In summary, it has been shown that the electrical properties of PEDOT lines change when subjected to various depackaging methods. Given a suitable readout scheme, depackaging of a microchip can be detected. If the PEDOT lines are not sensitive by default, sensitisation methods are available to enhance or create sensitivity.

Chapter Five

Prototype Tamper Protection Grids

5.1 Introduction

In the previous two chapters, it was shown that PEDOT protection grids can be fabricated reliably, and that they are either intrinsically sensitive to the removal of the packaging, or can be made sensitive by the choice of substrate material. In this chapter, some of the issues concerning the transfer from proof of concept to practical implementation are highlighted. These experiments are intended to establish the directions for future work. The properties that are required in a protection grid system are reliability, sufficient lifetime and sensitive readout method.

For reliable operation of the protection grids, the changes in properties due to ambient conditions such as temperature should either be minimal, or at least repeatable. If the variation of properties is sufficiently small, then an increased permitted measurement range can be defined to include all variations. However, this will reduce the sensitivity to depackaging. Alternatively, changes in properties could also be compensated for. This may be achieved using either differential measurement schemes, or by measurement of the environmental conditions and calculation of expected change in properties.

The lifetime of the protection grid should be at least as long as that of the security device, which requires long-term device stability. One reason why PEDOT was chosen for protection grids is its superior stability over most other stable organic conductors [Bantikassegn and Inganäs, 1997; Groenendaal et al., 2000; Heeger, 2001]. Nevertheless, the ageing of PEDOT resistors is examined to evaluate the rate of resistance change, and thus the expected lifetime.

Finally, prototype circuits for both passive and active protection grid types were assembled to test two of the proposed read-out schemes. For the passive protection grid prototype, a low-cost PIC microcontroller was used to measure the RC time delay along a PEDOT resistor. Apart from testing the integration of organic electronics with CMOS microchips, this evaluation allowed the measurement accuracy of a microcontroller-based read-out scheme to be determined. The evaluation of the active protection grid prototype was carried out to test the electrical properties of current organic TFTs and their performance in a protection grid circuit.

5.2 Effect of temperature

The evaluation of the properties of PEDOT has so far been carried out exclusively at temperatures in the range of ca. 20 °C to 25 °C. As security devices may operate across a wider tem-

perature range, possible changes in resistance of PEDOT due to temperature variation need to be considered. If the security device can measure the ambient temperature, a correction factor could be applied to the measured protection grid properties. Therefore the amount of change in PEDOT characteristics with temperature is less important than the repeatability. Silicon transistors and diodes exhibit variation of characteristics with temperature [Streetman and Banerjee, 2000], therefore it should be possible to use the variations to measure the temperature on the chip. A temperature sensor like the one developed by Sanchez et al. [Sanchez et al., 1997] should be suitable for this application.

The temperature dependency of resistance was determined for four PEDOT configurations. PEDOT with and without added DMSO were examined, as un-encapsulated and epoxy covered samples. As the addition of DMSO has a strong effect on conductivity, the temperature characteristics may also be affected. Similarly, the encapsulation with Araldite epoxy may also the high-temperature behaviour, by shielding the PEDOT from ambient oxygen or water, or by chemical interaction (Section 3.6).

5.2.1 Setup and procedure

A continuous resistance measurement of the PEDOT samples was carried out using Fluke F87 multimeters. To test the calibration of the multimeters, resistances measured with the multimeters and the HP4156A semiconductor parameter analyser were compared and found to coincide. As two multimeters were available for measurement, each glass slide was printed with two PEDOT samples. The samples were printed in the standard configuration of six layers, 40 µm drop spacing and either 1:1:0.2 PEDOT:water:DMSO or 1:1 PEDOT:water. Each sample was connected to a multimeter using copper/nickel alloy wires. The wires were attached to the samples by embedding them in the silver paint contacts. The resistances of the samples were read from the multimeters at intervals of 5 °C. The measurements are normalised to the resistance value read at 30 °C, to remove variation due to line geometry, epoxy degradation, and possible resistance changes from previous thermal cycles.

Measurement of the sample temperature was found to be challenging. The two available methods to measure the temperature in the used setup were the in-built thermometer of the hot plate, or an external thermocouple. The in-built thermometer of the hot plate measures the plate temperature, which will differ from the temperature of the sample on top of the glass slide. It was found that the reading of the in-built thermometer at room temperature was offset by ca. 10 °C.

The alternative method for temperature measurement is the use of an external thermocouple in physical contact with either the top of the sample or the plate surface. While the initial temperature readings were found to be more accurate (determined by comparison of several thermometers), this method gave less consistent and fluctuating readings at high temperatures. As the thermocouple needs to be at the same temperature as the hotplate or sample, the thermal contact and heat transport influence the measurement. For a thermocouple in direct contact with the hotplate, it was found that at high temperatures the reading changed with contact pressure. Pushing the thermocouple onto the hot plate using an insulator increased the temperature reading, which is evidence of bad thermal contact between the hot plate and the thermocouple. For a thermocouple placed on top of the 1 mm thick glass substrate, the temperature reading was lower than when the thermocouple was in contact with the hot plate. This is due to the significantly lower heat conductivity of glass (coefficient of thermal conductivity ca. 1.1 W/(m K)) compared to the metal hotplate (typically several 100 W/(m K)). Cooling effects from air flow over the thermocouple or heat transport through the connecting wires may be sufficient to skew the temperature reading.

While the calibration of the in-built thermometer may be debatable, it gave the most repeatable readings. As the thermometer is built into the hot plate, the thermal contact and heat transport properties were consistent throughout all measurement cycles. At this proof-of-concept stage, the approximate magnitude of change, and the repeatability of the R-T properties of PEDOT are evaluated. Therefore, the absolute precision of temperature measurements is less important than consistency temperature readings. For better consistency of temperature readings, the hot plate was covered with a glass bell jar during measurements. The heat transport through the sample substrate is low, as glass is an insulator. Measuring the temperature of the hot plate may not accurately reflect the sample temperature, as the air flow from the ventilation system in the laboratory cools the top surface of the sample. The bell jar ensures repeatable conditions by preventing external air flow.

Inaccurate temperature readings essentially result in a scaling of the X-axis. If accurate temperature measurements were available, the data points on the resulting R-T figures could be moved along the axis, however the total amount of resistance change would still be the same. In the final design of a protection grid system, temperature compensation must be re-calibrated to match the exact grid configuration, materials combination, and the on-chip temperature sensor.

5.2.2 PEDOT/DMSO not covered

Figure 5.1(a) shows the normalised resistivity vs. temperature taken in air for PEDOT/ DMSO samples (three repetitions per sample). To illustrate the wide variation in behaviour, all traces were plotted rather than plotting the average trace. Each line colour corresponds to measurements on a single sample. Following same colour traces, it can be seen that the R-T behaviour differs even between measurement runs on a single sample. As all samples were printed using the same batch of PEDOT, and under identical conditions, the R-T curves would be expected to coincide. The normalisation should also have removed line length and geometry influences, and even if not, the curves taken on the same sample would have been identical. The R-T curves flatten out at different levels at high temperatures, therefore inaccurate temperature measurements can also be excluded as the source of variation.

No systematic pattern could be seen in the differing behaviours of the samples, therefore the thermal history of the samples is unlikely to be the cause. The only remaining influence on the sample may be the ambient atmospheric conditions. To remove any possible influence of the weather and atmosphere, the bell jar covering the hot plate was repeatedly evacuated and flooded with nitrogen gas after loading a sample. It was found that during the initial evacuation of the bell jar, the resistances of the samples increased. Several cycles of evacuation and flooding with nitrogen were carried out to flush out all air. To prevent air re-entering the bell jar, nitrogen was continuously blown into the bell jar to raise the pressure in the bell jar to just above ambient. The increased resistance of the PEDOT lines settled down after ca. 5 to 10 minutes in nitrogen, to an average value 17.5% higher than the initial resistance. Removing the bell jar resulted in an immediate drop in resistance, back to the initial value.

The two curves shown in Figure 5.1(b) show the resistance vs. temperature graphs for samples in nitrogen gas. The difference between the curves is the duration of nitrogen exposure before the measurement was taken. The blue graph represents 15 individual measurements with around 5 to 10 minutes exposure to nitrogen before carrying out the measurement. The red graph combines the change of resistance of 8 measurements after significantly longer exposure to nitrogen compared to the blue graph, with at least one hour of prior exposure (and a maximum of 16h exposure). For a nitrogen exposure time between ten minutes and one hour, the resistance would follow an intermediate curve.

The same sample could be made to follow either the red or the blue graph by varying the nitrogen exposure time. Two samples were exposed to nitrogen for several hours, and found to

follow the red curve on the first measurement. After cooling down, the samples were exposed to air, and the resistance reset to its initial value. After re-exposure to nitrogen for 10 minutes prior to measurement, the resistance change followed the blue curve. Conversely, a sample following the blue curve could be made to follow the red curve in the next measurement by long nitrogen exposure.

The central difference between ambient air and nitrogen atmosphere are the presence of moisture and oxygen in air, both of which are reduced or absent in the nitrogen atmosphere. Artificially increasing the moisture around the samples in an uncalibrated (qualitative) test in air reduced the resistance of the PEDOT/DMSO samples. After removal of the moisture source, the resistance increased back to its original value again. Therefore it is concluded that the PEDOT samples absorb moisture, which changes the resistivity. This conclusion is also consistent with the variance of the R-T measurement seen in ambient air. As the measurements were carried out over a period of two weeks, the humidity of the air, and correspondingly moisture in the sample, will have varied with the weather. The moisture absorbed in the sample evaporates during heating, therefore it is likely that there are two overlapping mechanisms changing the resistance. One is the resistance change due to the removal of water in the sample, the other the resistance change due to temperature variation. However, as the average resistance increases under nitrogen is 17.5%, and the average subsequent resistance decrease is 16% to 20%, the overall resistance change due to both components individually adds up to a decrease of 1% to 6% compared to the initial value. This value is lower than the decrease of at least 8% measured for the samples in air, therefore other effects (e.g. oxygen) must also play a role in the thermally induced resistance change of PEDOT.

Apart from the large measurement variation, it was found that for the measurements in air, the sample conductivity degraded by up to 5% after keeping the sample at maximum temperature for five minutes. The temperature was limited to a maximum of 130 °C to prevent excessive sample degradation at high temperatures. For samples measured in nitrogen, the hot plate was allowed to cool naturally under continued nitrogen flow after the five minute annealing time. This resulted in a much longer exposure to high temperatures. No degradation was found for these samples. Increasing the temperature further, continuous degradation was observed above ca. 160 °C in nitrogen. The reduced sample degradation in nitrogen was also observed by Li et al. [Li et al., 2004]. Similarly, Winter et al. [Winter et al., 1995] identified the degradation mechanism as partial oxidation of the sulfur in the PEDOT backbone, as well as the decomposition of the PSS dopant under the influence of high temperature and moisture.

5.2.3 Epoxy encapsulated PEDOT/DMSO

The samples for this experiment were encapsulated with Araldite Rapid epoxy in ambient air, including any atmospheric humidity and oxygen. Nitrogen exposure of encapsulated samples even for several hours before a measurement did not make any difference to the measured resistance, as the epoxy is a good barrier to diffusion.

The results of the R-T measurement for the epoxy-covered PEDOT/DMSO sample are shown in Figure 5.2. The measured characteristics separated out into four groups, which are represented by the four different coloured curves. The blue and green curves are taken from samples one day after encapsulation with epoxy. For these 'fresh' samples, the resistance change during the first temperature cycle (green) differed from subsequent cycles (blue), a phenomenon that was not observed for older samples examined more than a day after encapsulation. For the first temperature cycle (green), the resistance decrease accelerates above 70 °C to 80 °C, levelling off again around 110 °C. A permanent reduction in resistance was found in these samples after cooling down. The standard deviation (error bars) of the green curve increases at high temperatures due to different samples showing varying amounts of resistance change in their first

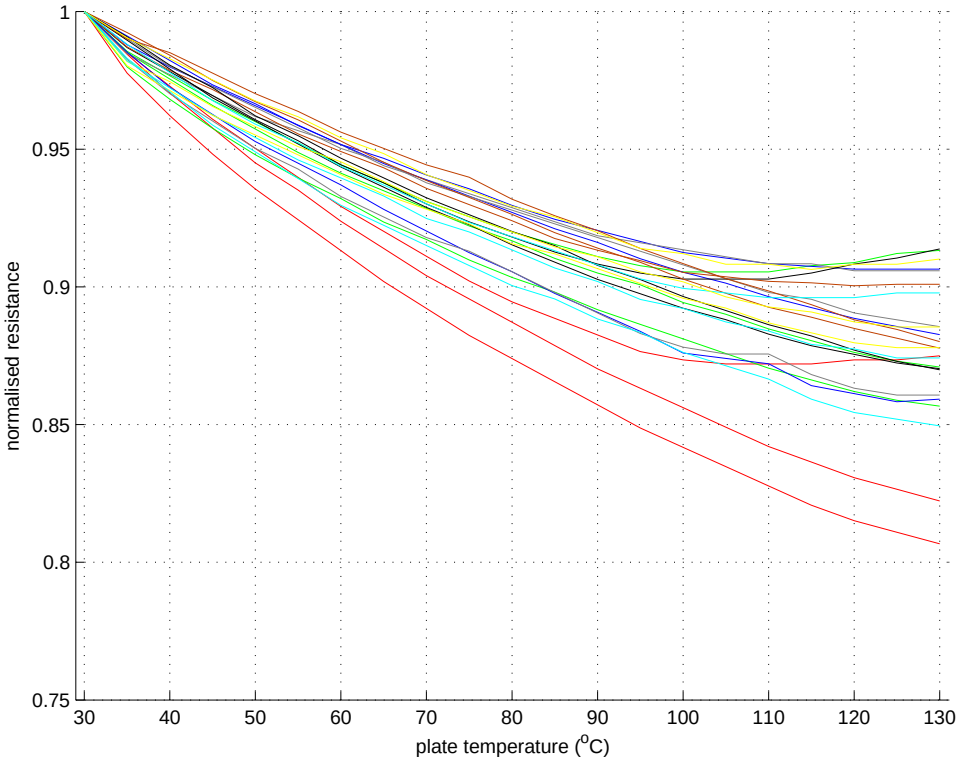

(a) PEDOT/DMSO in air

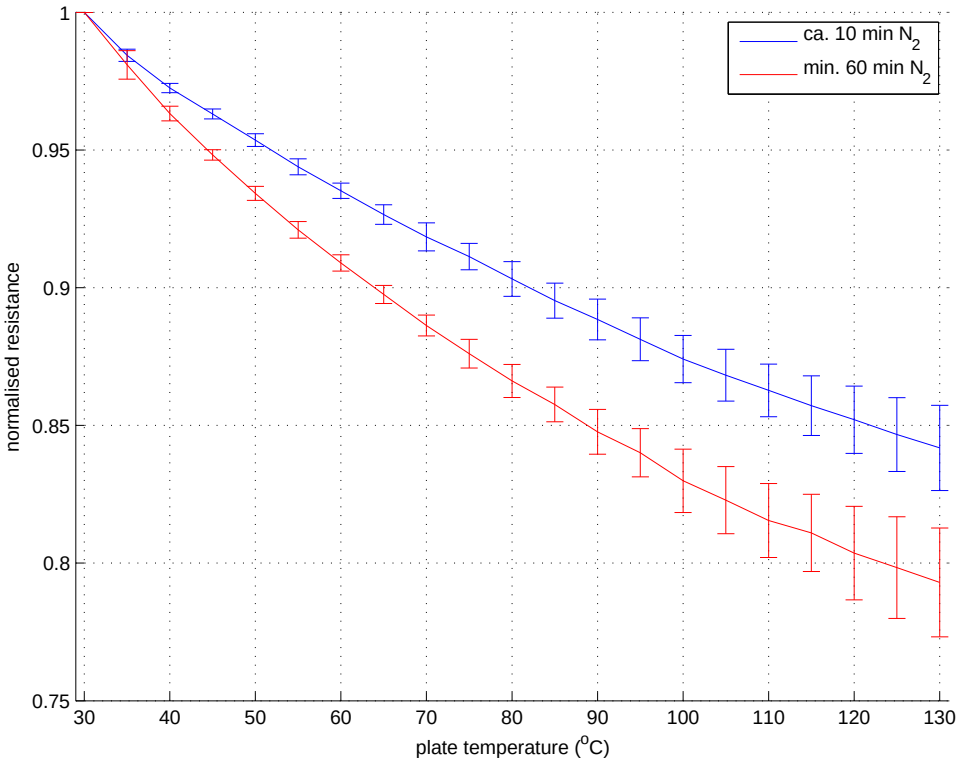

(b) PEDOT/DMSO in dry nitrogen

Figure 5.1: Normalised resistance vs. temperature for PEDOT/DMSO in air and in nitrogen

measurement cycle. After the first measurement cycle, all 'fresh' samples showed much closer R-T behaviour (blue curve). An obvious explanation for this phenomenon is that the epoxy was not fully cured before heating, and the elevated temperature accelerated and completed the curing process, possibly un-doing some of the increased resistance due to the epoxy coating (Section 3.6). It will be shown in Section 5.3 that the epoxy causes resistance fluctuations for two days after encapsulation, which is interpreted as the time for the epoxy to fully cure. The curing time is consistent with the accelerated resistance change seen only for one day old samples.

The remaining two curves in Figure 5.2 show the R-T behaviour of two day old samples (black), and samples greater than two days old (red). It can be seen that all curves, including the fresh epoxy, display a minimum at a certain temperature. The location of the minimum value is dependent on sample age. For fresh epoxy the minimum is found close to the maximum temperature of 110 °C, and decreases to lower temperatures with increasing age. For samples older than two days, degradation begins at temperatures as low as 65 °C. At temperatures of more than ca. 10 °C above the temperature of minimum resistance, the resistance increases continuously and irreversibly when the sample is held at a constant temperature. The low degradation temperature and fast rate of degradation were not seen for the non-encapsulated samples printed on glass. Therefore, the most likely cause for this behaviour is the Araldite epoxy, which may itself have started to degrade.

This low-temperature degradation may in fact be useful as an additional tamper-proofing property for smart card devices. Before the chip packaging is targeted, a smart card chip needs to be removed from the carrier card. A simple way to do this is to soften the glue by heating [Kömmerling and Kuhn, 1999]. If the temperatures are tuned so that the protection grid is affected by this heating to soften the glue, the protection grid will already be damaged before any properties can be measured. However, this security benefit is traded against robustness of the card against accidental heating (e.g. hot car in summer).

5.2.4 Pure PEDOT, not covered

The R-T curves for pure PEDOT taken in ambient air and nitrogen are shown in Figure 5.3. The behaviour of pure PEDOT differs significantly from PEDOT with DMSO. For measurements in air (Figure 5.3(a)), the resistance changes significantly more (-78%) than PEDOT/DMSO (Figure 5.1(a), -13% on average). The second difference between the two sample types is the behaviour in nitrogen. PEDOT/DMSO showed an average increase in resistance of 17.5%. In contrast, the resistance of pure PEDOT decreased by 66% when exposed to nitrogen. The R-T curve for pure PEDOT in nitrogen is shown in Figure 5.3(b). It can be seen that the decrease in resistance with temperature is close to linear, and shows little flattening at higher temperatures. Unlike PEDOT/DMSO, the R-T curve did not vary with exposure time. Adding the resistance changes from to nitrogen exposure and temperature change, a total resistance change of -82% was calculated for the sample in nitrogen. This value is in reasonable agreement with the resistance decrease of -78% determined from heating in air. The evaporation of humidity is consistent with the overall change in resistance, and the flattening off at high temperatures when the water is evaporated.

Interestingly, the R-T curve measured in air fluctuated less than for PEDOT/DMSO, despite the much larger resistance change exposing the samples to nitrogen. It may be possible that the conductivity reducing effect saturates at a comparatively low level of absorbed moisture. This moisture is removed either by heating, or by lowering the ambient humidity below the threshold for saturation (in nitrogen). The investigation of this effect is beyond the scope of this dissertation and research field.

Under the bell jar, the measurement accuracy of the resistance of the pure PEDOT samples was degraded by fluctuating multimeter readings. The measurement fluctuations disappeared

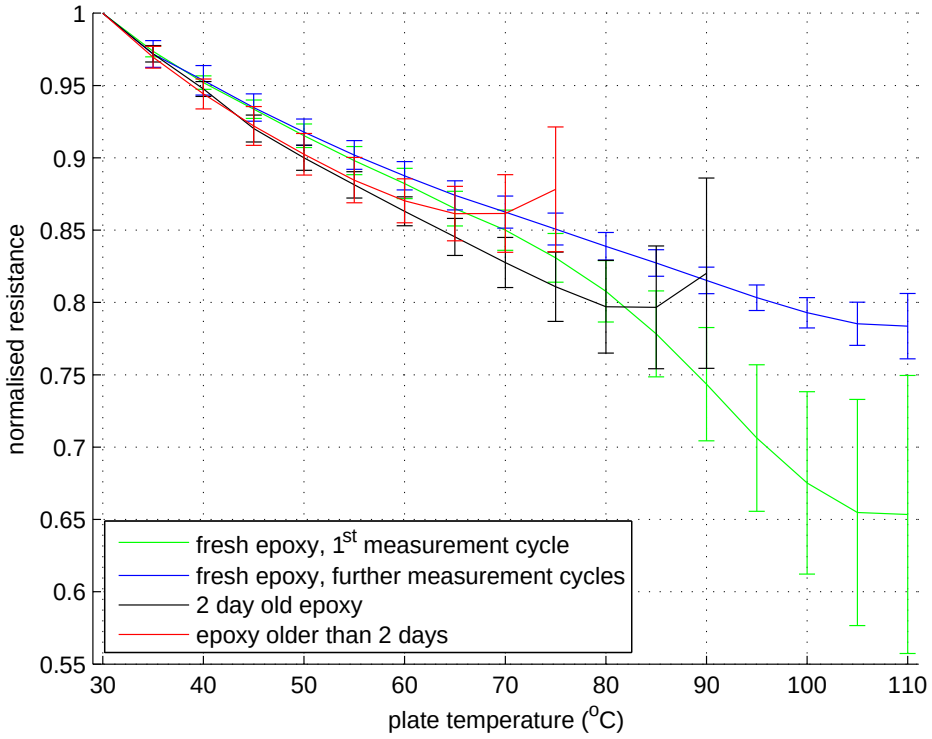

Figure 5.2: Normalised resistance vs. temperature for PEDOT/DMSO covered with epoxy

as soon as the bell jar was lifted. The variations were caused by the uninsulated wire leads making (weak) electrical contact with the hotplate under the pressure of the bell jar. For resistance measurements of the order of MΩ, the bias currents are very low due to the limited measurement voltage. Current leakage thus has a greater impact on measurement accuracy compared to the more robust bias currents of PEDOT/DMSO (kΩ).

5.2.5 Epoxy encapsulated pure PEDOT

The R-T curve for epoxy encapsulated pure PEDOT also shows different behaviour compared to encapsulated PEDOT/DMSO (Figure 5.4). The same Araldite Rapid epoxy was applied to the samples in ambient air, as to the PEDOT/DMSO samples. Regardless of age, all samples showed an accelerated reduction in resistance for the first heating cycle (blue), and repeatable characteristics thereafter (red). As before, the reduction in resistance after the first heating cycle was permanent. Sample degradation as seen in the PEDOT/DMSO samples was not observed, even though the epoxy was of the same type and age. It suggests like the degradation of epoxy or its decomposition products only have an effect on high-conductivity PEDOT. This differing behaviour may be explained by the degradation products of epoxy working against the conductivity enhancement from the DMSO. Alternatively, the decomposition products may simply have caused a fixed amount of resistance change, which in the 100 higher resistance of pure PEDOT will have fallen within the measurement tolerance. A detailed chemical analysis would be required to analyse the interaction, which is beyond the scope of this dissertation.

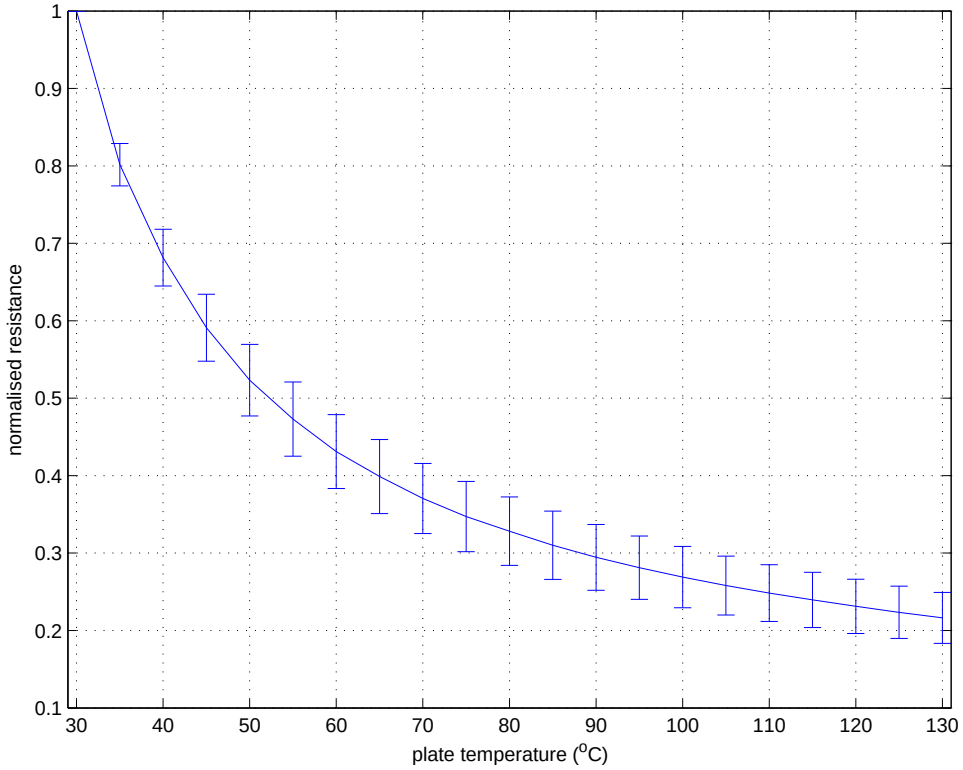

(a) Pure PEDOT in air

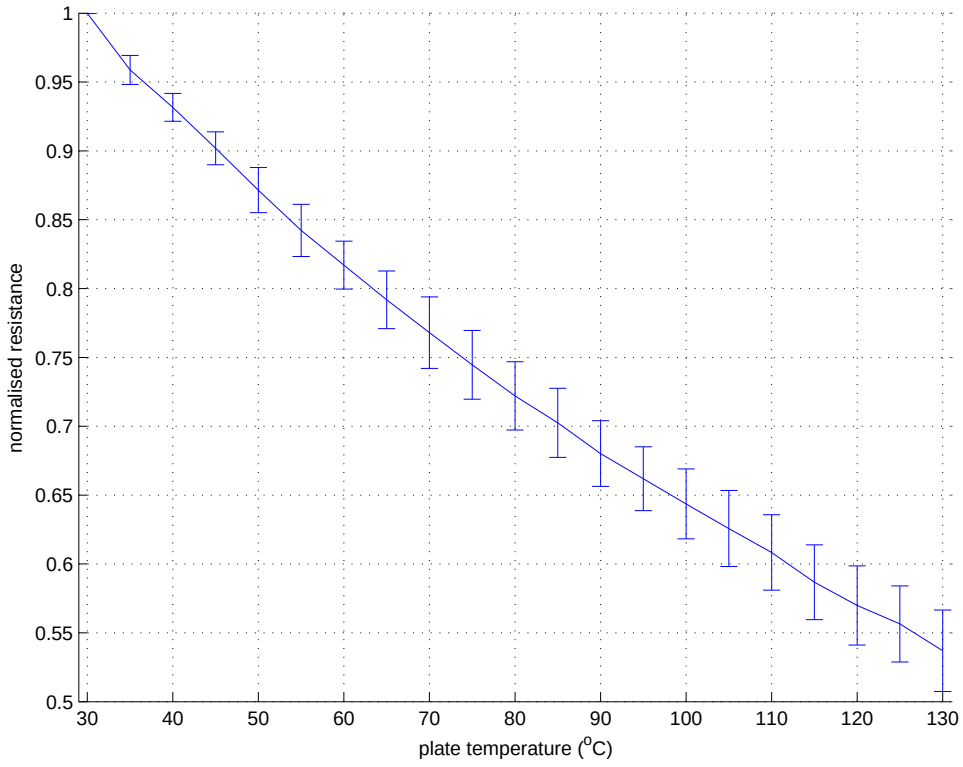

(b) Pure PEDOT in nitrogen

Figure 5.3: Normalised resistance vs. temperature for pure PEDOT in air and in nitrogen

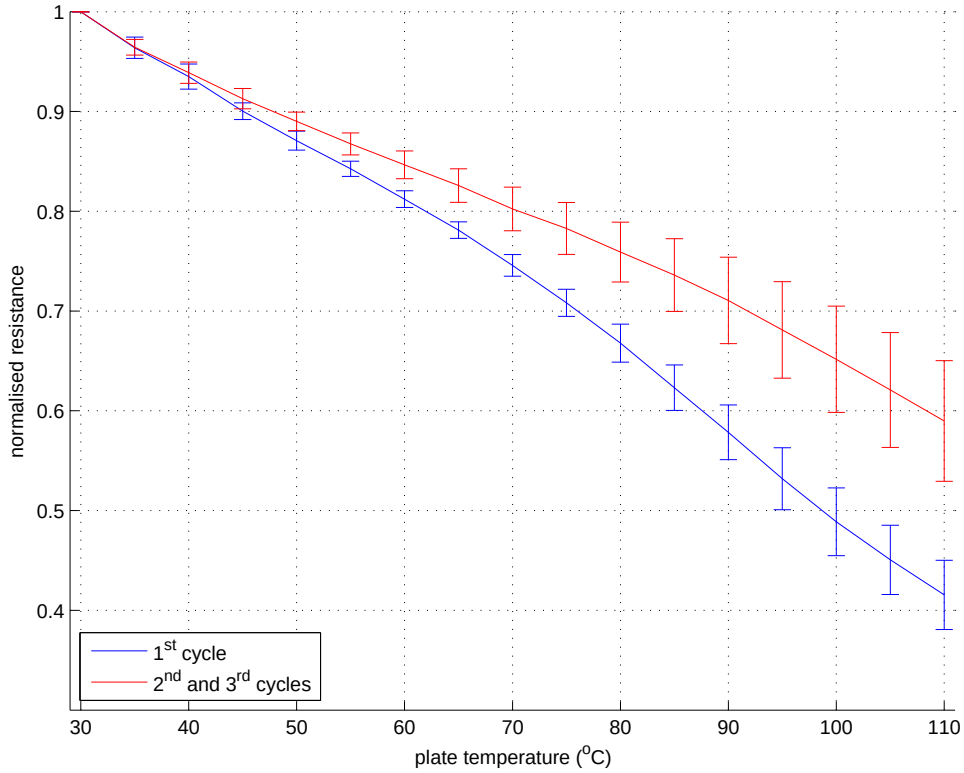

Figure 5.4: Normalised resistance vs. temperature for epoxy covered pure PEDOT

5.2.6 Conclusions

In summary, it has been shown that the resistance change of PEDOT at high temperature is not repeatable by default. For repeatable thermal behaviour, it is important that the PEDOT line is surrounded by an inert atmosphere, or encapsulation. Depending on design requirements, the PEDOT grid should either be encapsulated by a temperature-stable material, or by an epoxy that degrades at a certain maximum temperature.

5.3 Ageing behaviour

The proposed protection schemes rely on unique characteristics of the protection grid being measured, and compared to a reference value, to determine the integrity of the packaging. As changes in the properties of the grid are interpreted as tampering attempts, it is essential that the PEDOT characteristics must remain stable over time. To determine whether a protection grid fabricated from PEDOT permits a sufficient device lifetime for practical applications, the long-term stability of PEDOT lines were be examined.

5.3.1 Setup and procedure

Fifteen samples for each of the four sample types (1:1 PEDOT:water and 1:1:0.1 PEDOT:water:DMSO, epoxy encapsulated and bare) were used to determine the ageing behaviour. The I-V curves were taken over the course of 130 days with the HP4156A semiconductor parameter analyser. The samples were stored at ambient laboratory conditions (ca. 23 °C, light) throughout the measurement period.

5.3.2 Results

The resistance change with time for the two bare samples is shown in Figure 5.5. It can be seen that the resistance of both PEDOT with DMSO (Figure 5.5(a)) and pure PEDOT (Figure 5.5(b)) more than doubled over the measurement period. The resistance of PEDOT with DMSO increased by a factor ca. 2.65, while the resistance of pure PEDOT increased by a factor of 2.2. The ageing curve for pure PEDOT shows more stray points (local minima) compared to the PEDOT/DMSO curve. As the resistance change with temperature is much larger for pure PEDOT compared to PEDOT/DMSO, this is the most likely case. The temperature in the laboratory may have varied on different days, due to the air conditioning in the laboratory running or being turned off. A temperature change of 5 °C causes a resistance change of 20% (magnitude of the largest drop) for PEDOT in ambient air at a baseline temperature of 30 °C (Figure 5.3(a)). Extrapolating the R-T curve to 23 °C, a smaller temperature change (3–4 °C) is required to cause the observed fluctuations, which may well have been caused by the air conditioning unit. This result contrasts with the experiments of Xue et al., who found that the resistance change saturates for samples stored in air over a similar time period [Xue et al., 2005].

Figure 5.6 shows the resistance change for the epoxy encapsulated samples. The resistance of the encapsulated PEDOT/DMSO initially decreased for ca. three weeks, before stabilising and slowly increasing again (Figure 5.6(a)), at a slow rate compared to the bare samples. In the pure PEDOT samples, the resistance initially stayed roughly stable for the three week period, (Figure 5.6(b)), before also increasing at a slow rate. For the first two days after encapsulation, the PEDOT lines displayed significant variation in resistance, evident from the error bars in Figure 5.6. The normalisation reference was therefore moved to day three, when the resistances had stabilised. Comparing the decrease in resistance over the first three days to the resistance decrease in the temperature measurements (Section 5.2.3), it can be concluded that the elevated temperatures accelerate the curing process of the epoxy.

5.3.3 Discussion

These experiments show that the encapsulation with epoxy slows down the time-dependent conductivity degradation of PEDOT. While PEDOT has been named as one of the most stable organic conductors [Bantikassegn and Inganäs, 1997; Groenendaal et al., 2000; Heeger, 2001], degradation with time is clearly visible. Two presentations available from the PEDOT supplier confirm that PEDOT is stable in an inert atmosphere [Elschner, 2002], and that it degrades under heat, UV light, and oxygen, while the PSS dopant degrades with water [Starck, d].

The reduced degradation of the encapsulated samples shows that at least one component contributing to the degradation is blocked by the epoxy, most likely either oxygen or water. The samples were stored in light conditions, which contains a UV component. As it was possible to evaporate PEDOT through the (intact) epoxy with a UV laser (Section 4.5), it is evident that UV light can pass through epoxy. The UV light may contribute to the degradation of the encapsulated samples.

Furthermore, the epoxy does not completely cover the samples, as the ends of the PEDOT lines must be accessible for probing. This allows oxygen and moisture to diffuse into the PEDOT from the sides, at a lower rate than for bare lines. Some oxygen and moisture may also diffuse through the epoxy, but at a lower rate than through the open contacts. Apart from UV light and small amounts of oxygen and water, reaching the encapsulated PEDOT, the epoxy itself may also contribute to degradation. It was shown in Section 5.2 that the degradation of the epoxy at high temperature impacts the conductivity of PEDOT. Therefore it is possible that the Araldite epoxy may also have started to degrade, damaging the PEDOT resistors.

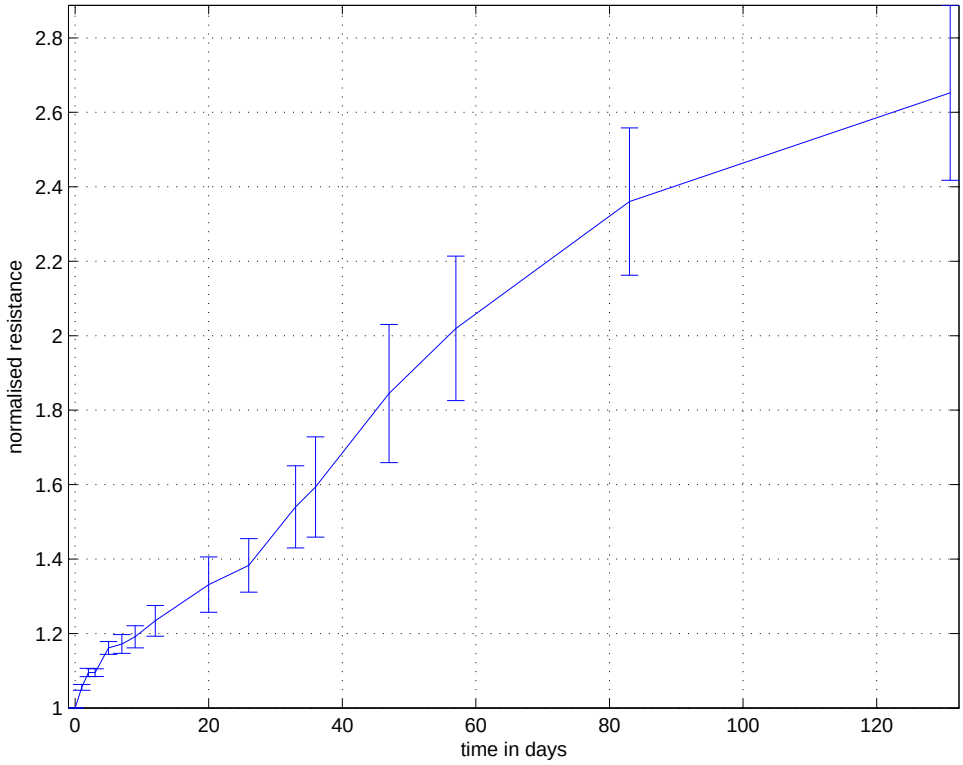

(a) PEDOT/DMSO, bare

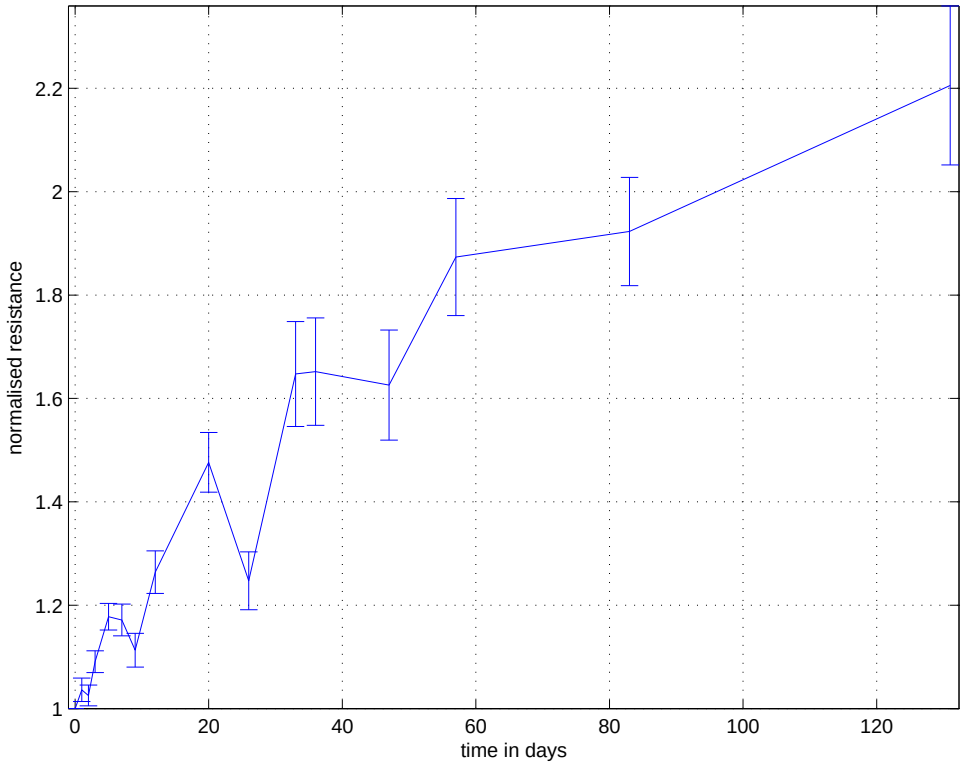

(b) Pure PEDOT, bare

Figure 5.5: Normalised resistance vs. time for bare samples

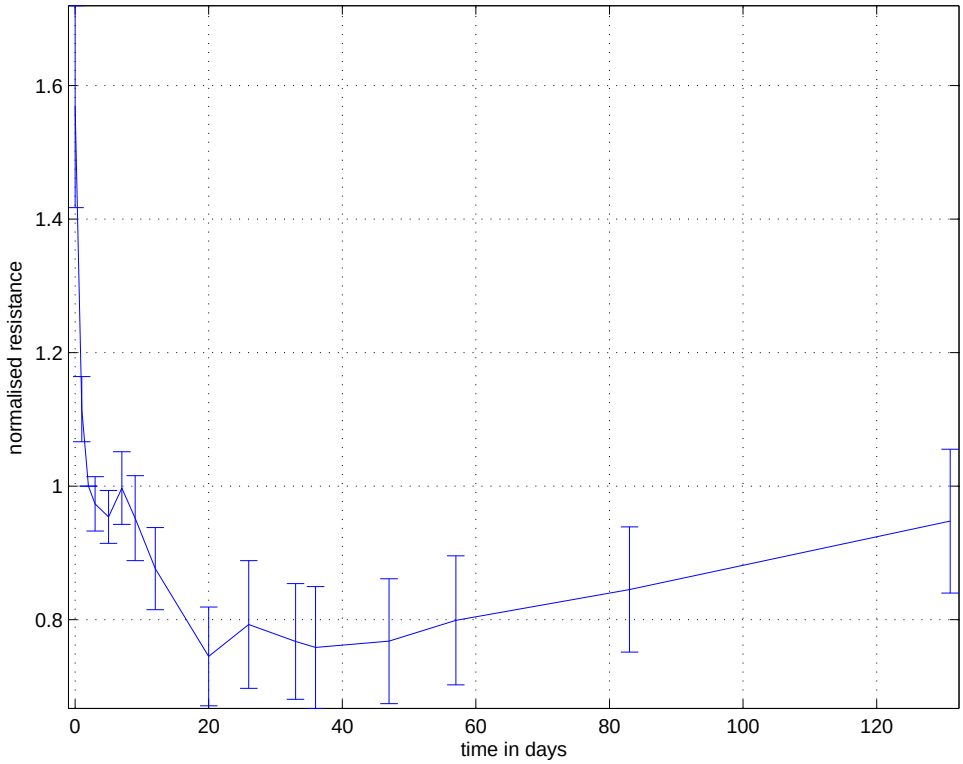

(a) PEDOT/DMSO, epoxy encapsulated

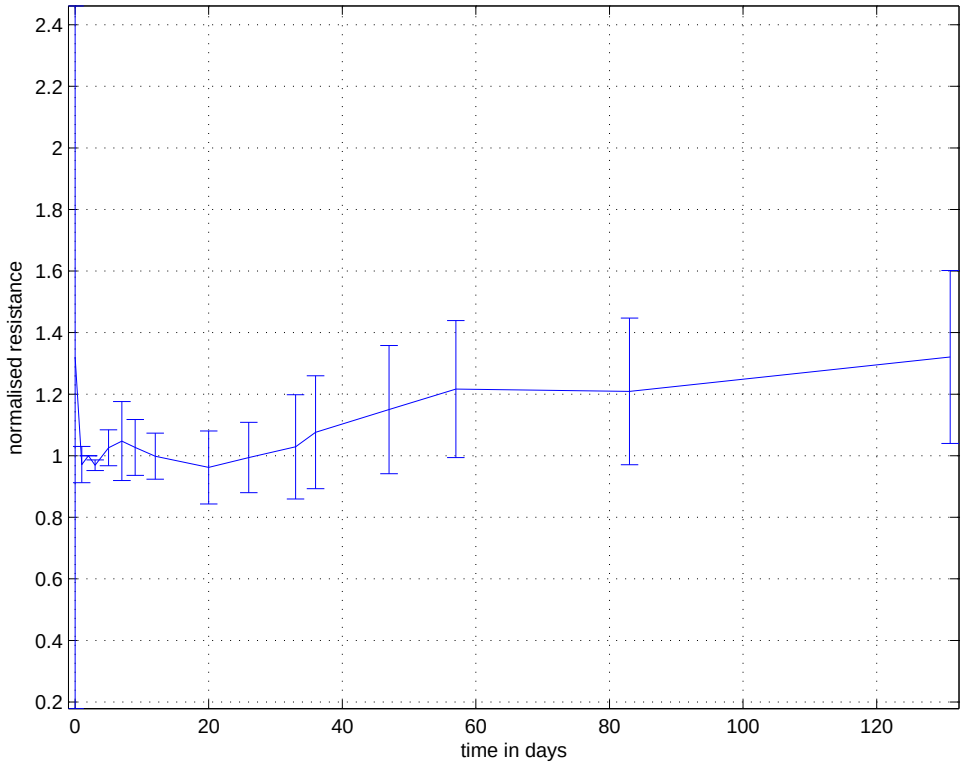

(b) Pure PEDOT, epoxy encapsulated

Figure 5.6: Normalised resistance vs. time for epoxy encapsulated samples

5.3.4 Conclusions

While it can be concluded that encapsulation of PEDOT decreases the ageing process, completely stopping PEDOT degradation has not been achieved. For a multi-year lifetime, a better encapsulation material must be found to prevent the slow degradation observed in this experiment. To eliminate degradation, the encapsulation material should completely enclose the PEDOT lines, be opaque to UV light, and have minimal diffusion rates for oxygen and water.

5.4 Evaluation of passive protection grid prototype

To evaluate the ability to measure the properties of the protection grid with a standard CMOS microchip, a prototype was assembled. A standard low-cost microcontroller was chosen to emulate the security device and to read out the properties of the PEDOT lines. Microcontrollers are similar to security devices, as they are also designed under stringent cost constraints. The Microchip PIC 16F689 microcontroller used for the prototype has a retail price of around $1, which may be used a benchmark value for the price of a security device.

The chosen readout scheme for the passive prototype is the RC delay line. This scheme measures the time delay of a signal between the input to the protection grid and the output from the protection grid. This scheme includes different aspects, that are also relevant for other readout schemes. Time delay measurements are used in the delay line, as well as for a ring oscillator active protection grid. To capture the signal at the end of the protection grid, an analogue voltage measurement is carried out. The voltage measurement at the end of the delay line is relevant to passive protection grids based on resistance ratios, as well as active bridge-type protection grids.

As it must be assumed that any input to the security device may have been manipulated, the reference clock for the time delay measurement must be generated internally on the microchip. The frequency stability of the internal CMOS oscillator has direct influence on the measurement accuracy of time intervals. Conveniently, the PIC microcontroller contains two internal oscillators, which may be used to evaluate the frequency stability of CMOS oscillators.

The prototype is evaluated in two steps. First, the frequency stability of the internal CMOS oscillators of the microcontroller is evaluated to give an indication how well a CMOS oscillator designed under cost constraints performs. In the second step, the sensitivity of the measurement scheme implemented using the microcontroller is evaluated. A completely severed protection grid would be trivial to detect, as no signal is transmitted. However, under the assumption that the packaging does not ensure complete destruction, the grid resistance would merely be altered. The sensitivity evaluation aims to quantify the detectable resistance change.

5.4.1 Theory

The measured time delay depends on the circuit properties, as well as the software of the microcontroller. As the microcontroller processes all instructions sequentially, there is an offset in the measured time delay. The timer is started before the output is switched, and the interrupt service routine also introduces a delay for stopping the timer. This time delay is constant for all devices, and it can simply be subtracted automatically.

As the resistance and capacitance of the PEDOT protection grid vary with printed pattern, the delay caused by charging the parasitic capacitances is unique to each device. Depending on the relative magnitudes of the capacitances of the grid line, and of the measurement circuit, different delay models are used.

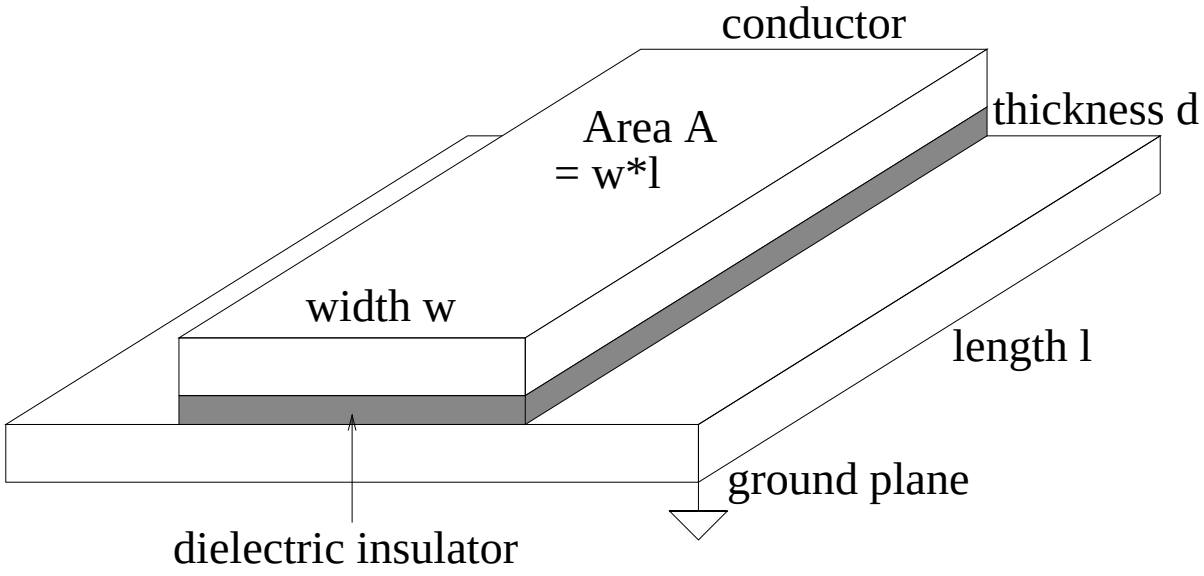

Figure 5.7: Diagram of capacitor

Generally, the parasitic capacitance per unit length of a grid line is calculated using the standard plate capacitor model, as drawn in Figure 5.7:

$$C = A \quad \frac{\epsilon}{d} \tag{5.1}$$

$$c = (w + 2d) \quad \frac{\epsilon_r \epsilon_0}{d} \tag{5.2}$$

Where the following symbols are used: area of a capacitor (A), thickness of dielectric (d), dielectric constant (ϵ), total capacitance (C) and capacitance per unit length (c). A common assumption for two plates of different sizes is that the fringe field extends both sides around the smaller plate by the thickness of the dielectric.

If the total capacitance is dominated by the input to the measurement circuit, then a simple lumped RC model may be used to calculate the total delay. In the prototype, the protection grid line is separate from the microcontroller, and connected through the circuit board, wires, and silver contacts. In addition, there is no ground plane in close proximity to the grid line, reducing the capacitance of the grid line. Therefore it is expected that the RC delay of the prototype will follow the lumped RC model.

A diagram of a lumped RC circuit with the symbols used for this derivation is shown in Figure 5.8. Summing currents at node X:

$$0 = \frac{V_x(t) \quad V_i(t)}{R_{grid}} \quad i_{leak}(t) \quad C_p \frac{dV_x(t)}{dt} \tag{5.3}$$

$$\frac{dV_x(t)}{dt} = \frac{V_x(t)}{R_{grid}C_p} \quad \frac{V_i(t)}{R_{grid}C_p} \quad \frac{i_{leak}(t)}{C_p} \tag{5.4}$$

$$V_x(t) = \exp\left(\frac{t}{R_{grid}C_p} \right) \left(const. \quad \int_0^t \left(\frac{V_i(t)}{R_{grid}C_p} \quad \frac{i_{leak}(t)}{C_p} \right) \exp\left(\frac{t}{R_{grid}C_p} \right) dt \right) \tag{5.5}$$

The response $V_x(t)$ depends significantly on the input function $V_i(t)$ and the leakage $i(t)$. For

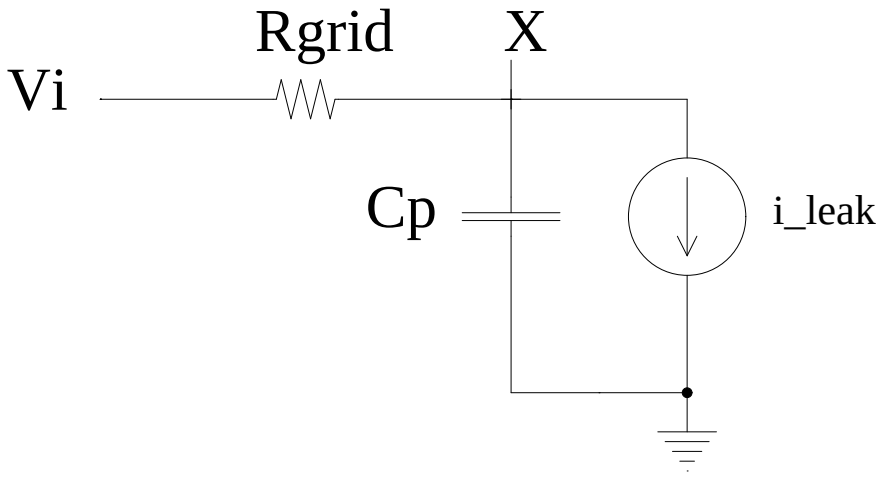

Figure 5.8: Schematic of lumped RC delay line

negligible leakage, an initial voltage $V_x(0) = 0$, and a step input $V_i(t) = V_{in}$ at $t = 0$, this expression is simplified to the well-known exponential approach curve:

$$V_x(t) = V_{in} \quad \left(1 \quad \exp \left(\quad \frac{t}{R_{grid}C_p} \right) \right) \tag{5.6}$$

If the parasitic capacitance of the grid line cannot be neglected, a distributed RC delay model must be used. The standard model is the Elmore delay model [Elmore, 1948], in which the the delay line is split into discrete RC-segments. In the Elmore delay model, the overall delay time is computed by summing the product of the capacitance of each segment with the resistance to the signal source. While this model works to compute the delay time constant, it does not fully capture the wave form expected at the output of the circuit. For a constant resistance and capacitance per unit length (which may not be the case for topographical base layer), it is possible to define the RC circuit as a 1-D diffusion equation [Rabaey et al., 2003]. The solution to the diffusion equation for a step response of an open-ended distributed RC line is computed in literature [Kahng and Muddu, 1994]. The computation of the output waveform for an arbitrary input waveform and/or non-constant capacitance per unit length require a convolution of the impulse response with the input waveform. The derivation of the output waveform is therefore best done as numerical simulation using the simulator with integrated circuit emphasis (SPICE).

A depackaging attempt is likely to cause the electrical properties of the PEDOT line to change. If the line geometry is not affected, the total resistance of the line will change, but the capacitance of the system will remain constant. This will result in a corresponding change in RC time constant, and hence be detectable by the circuit.

5.4.2 Prototype design

A PIC microcontroller is used to emulate the security device. For simplicity, the time delay along the grid line is measured by switching a digital output pin from zero to one (0 V to V_{dd}) as a step function. Using a digital input to the microcontroller as receiver was found to switch at a very low input voltage, resulting in a delay time which was too short for evaluation. Instead, the analogue comparator of the microcontroller is used to define a threshold voltage at which to measure the delay time. The voltage level of $\frac{15}{24} \quad V_{dd}$ was chosen arbitrarily, as it corresponds to a setting of '11111' in the control register.

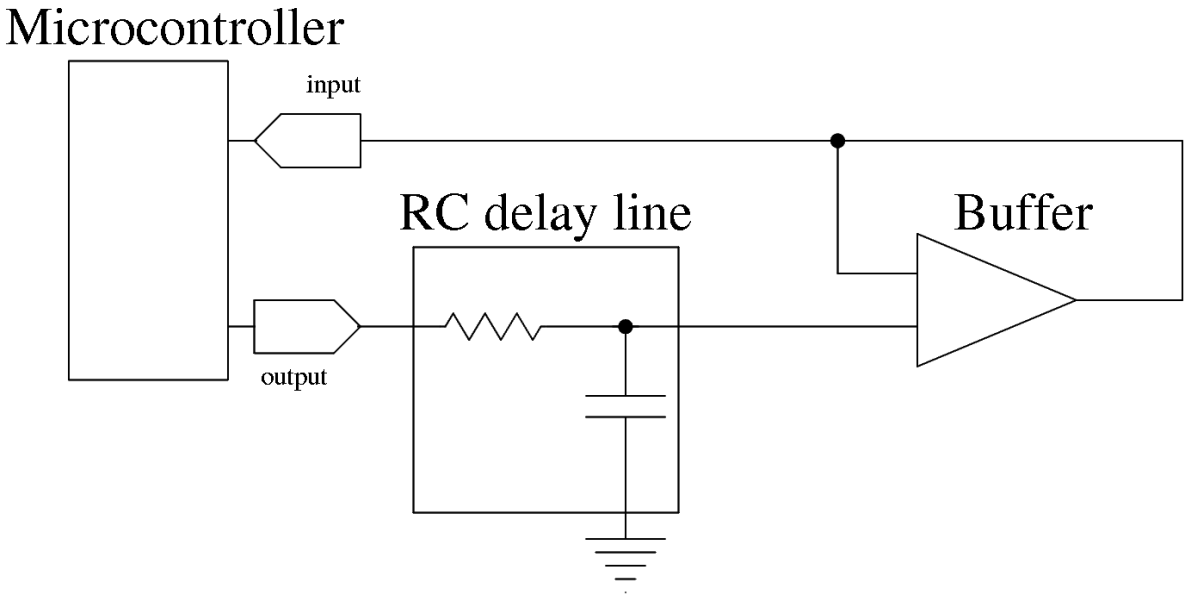

Figure 5.9: Schematic of passive protection grid prototype

The maximum input impedance into the analogue comparator of the microcontroller is specified at 10 kΩ, which is too low for PEDOT grid lines. Therefore, a CMOS op-amp was used as a voltage follower buffer. To minimise the leakage currents, a CMOS op-amp with a very high input impedance of the order of TΩs was chosen as the buffer. The final schematic of the passive delay line prototype is shown in Figure 5.9.

As the maximum operating frequency of the PIC microcontroller is 20 MHz, the resolution of the delay measurement is limited to steps of 50 ns. To reduce the noise added by the discrete time intervals, the rise time should span a sufficiently large number of clock cycles. If the delay time spans e.g. 100 clock cycles, the time resolution and quantisation error is 1%. Using the PIC microcontroller, the rise time of the grid must be slower than ca. 5 µs. This in turn requires that the resistance of the grid line is appropriately large, trading off delay time and feeble bias conditions at low currents. To run the microcontroller at the full frequency, an external 20 MHz quartz crystal was used to generate the clock frequency rather than the slower internal oscillator. As quartz crystals have very high frequency stability, the measurement accuracy of the delay time consists of two components: the CMOS oscillator frequency stability, and the accuracy of the delay measurement itself at a fixed oscillator frequency.

5.4.3 Measurements

5.4.3.1 Stability of CMOS oscillators

Setup and procedure Measurements were carried out for both internal oscillators of the microcontroller. An output pin was toggled to determine the oscillator frequency at different voltages and temperatures. The frequency was measured using a counter timer. The supply voltage to the microcontroller was varied using a laboratory power supply and a voltmeter was used to measure the voltage across the power and ground pins. The temperature was varied by means of ice spray and a hot air blower cooling or heating the top of the device. The temperature was measured using a thermocouple placed below the device. Before determining the oscillator frequency, the temperature was held stable for at least 30 s to allow the core of the device to reach the ambient temperature.

Resistance R MΩ	Mean delay μ clock cycles	standard deviation σ clock cycles	μ / R clock cycles/MΩ
14.92	358	12.39	23.99
16.81	410	12.00	24.39
19.82	475	12.65	23.96
21.7	483	13.31	22.25

Table 5.1: Delay line measurement results

Results The frequency-voltage diagram is shown in Figure 5.10(a). It can be seen that the frequency of both oscillators drops significantly at voltages below the specified minimum supply voltage of 2 V. Between 2 V and 6 V (detail in Figure 5.10(b)), the internal RC oscillator frequency is found to be very stable (within 0.5%). The high frequency internal oscillator frequency varies between +1% and -3.5%.

The frequency-temperature plot for both oscillators is shown in Figure 5.11. The frequencies of the internal oscillators vary by less than 1% (0.4% for the RC osillator) within the measured temperature range. This figure is in line with the accuracy quoted in the PIC 16F689 data sheet. This variation is small compared to the frequency variation with voltage and the change of PEDOT resistivity with temperature.

5.4.3.2 Delay line tests

Procedure Initially, the time delay of the microcontroller setup without a PEDOT protection grid was measured. The output of the microcontroller was short-circuited to the input of the voltage follower buffer. For the main measurement, a 5 mm, 6-layer PEDOT line printed on an epoxy substrate was wired between the two terminals. The resistance of the protection grid line was measured using a multimeter. One thousand time-delay measurements were taken for each measurement round. After a measurement round, the protection grid and substrate were deformed using a needle to effect a change in resistance.

Results The short-circuit delay time was determined to be a constant 29 clock cycles including the op-amp buffer and 26 clock cycles without. The results of four measurement cycles of an example delay line are shown in Figure 5.12. The results are also tabulated in Table 5.1. As can be seen from the last column, the μ/R ratio is near constant, showing good correlation with the lumped RC model. Measurements on other samples also showed a constant μ/R ratio. A different sample was characterised with 1000 and 10000 observations to test the independence of the measurements, and determine whether the PEDOT line characteristics might change with large numbers of measurements in a short period of time. For the 1000 observation measurement set $\mu = 289.70$, $\sigma = 14.56$ and for 10000 measurements of the same sample $\mu = 289.16$, $\sigma = 14.81$. As these values are practically identical, it can be concluded that the resistance measurement is a stationary process, and individual measurements are independent from each other.

5.4.4 Discussion

The RC oscillator is more stable than the high frequency internal oscillator, however the generated clock frequency is also much lower (31 kHz). For a security device to measure delay times or frequencies with high resolution, a faster oscillator than the in-built slow RC oscillator is required. As standard RC oscillators are limited to a few MHz frequency [Bala and Nandy,

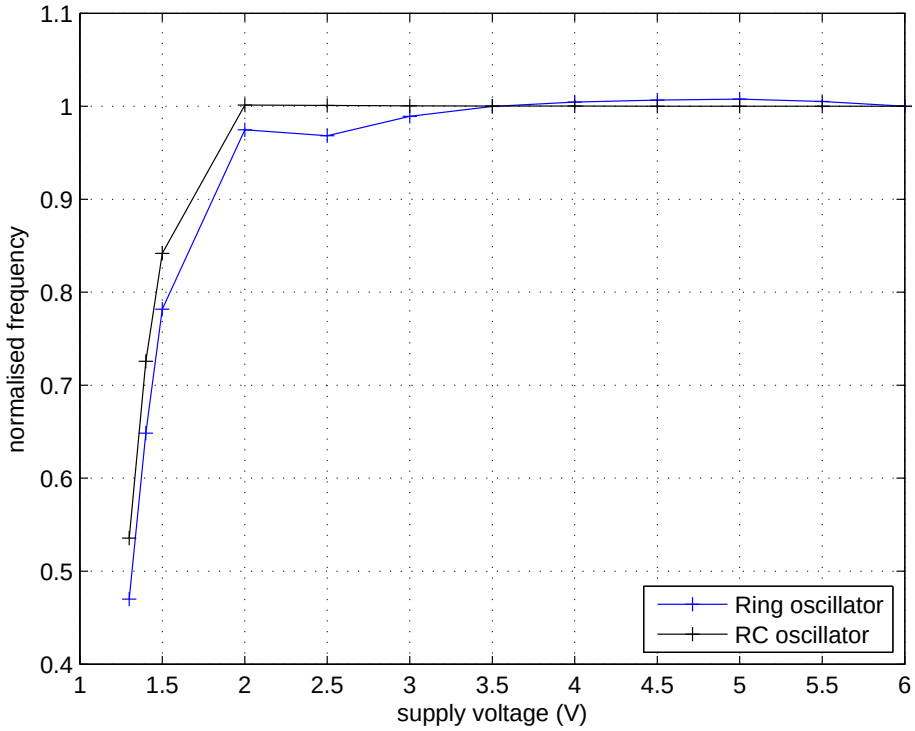

(a) Full measurement range

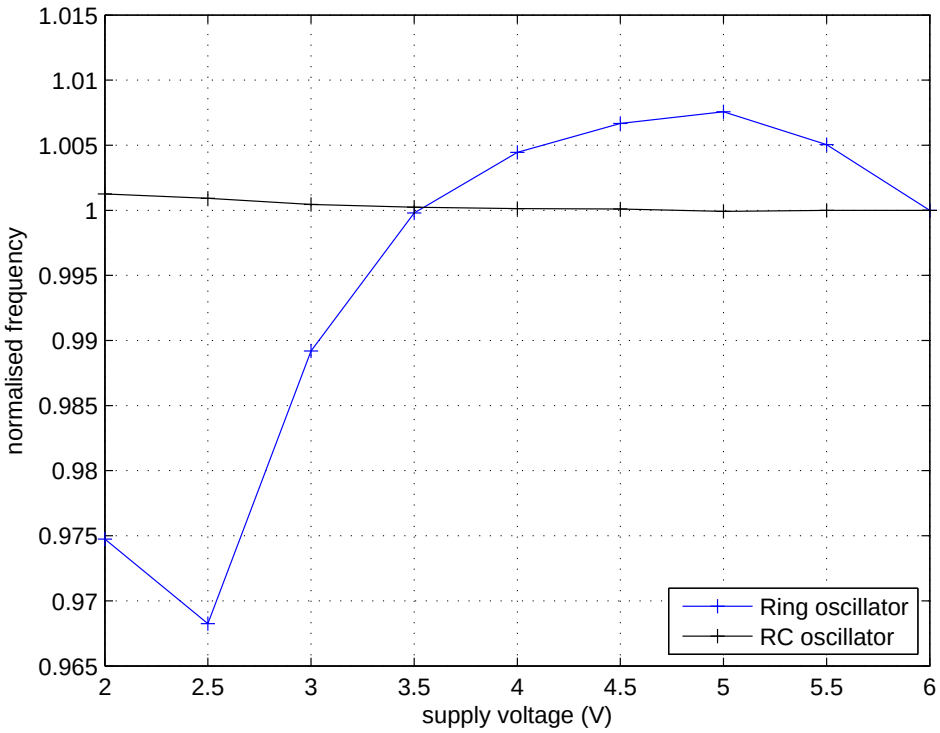

(b) Detail of 2 V to 6 V range

Figure 5.10: Frequency vs. supply voltage for PIC16F689 internal oscillators

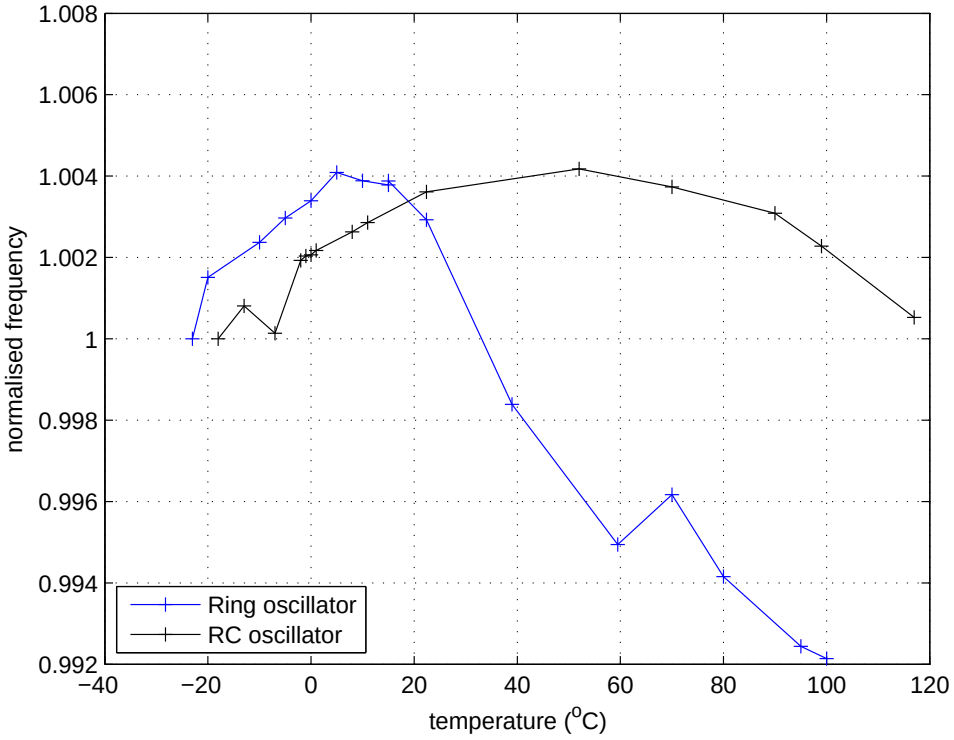

Figure 5.11: Frequency vs. temperature of PIC16F689 internal oscillators

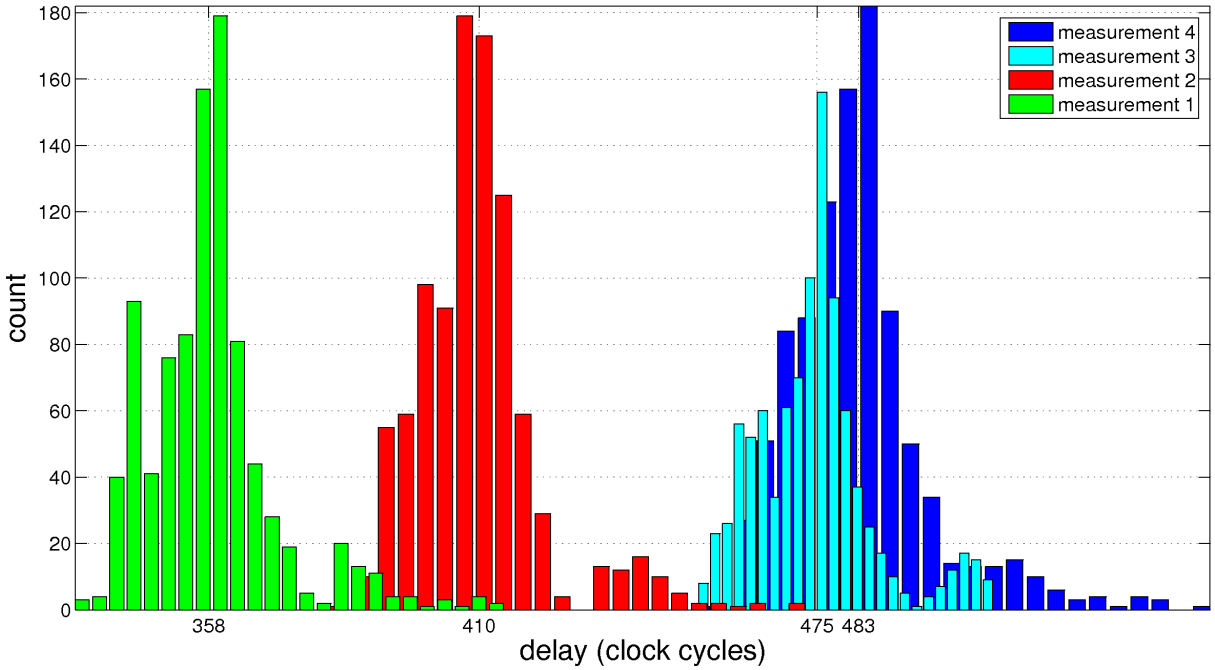

Figure 5.12: Delay line test with four measurement sets after damaging the PEDOT line

2005], a ring oscillator or phase-locked loop may be necessary to measure time intervals with sufficient resolution, however these may suffer from reduced frequency stability.

From the measurement results it can be seen that a variation in supply voltage generally causes a larger drift in frequency than the ambient temperature. A stable voltage reference [Giustolisi et al., 2003] may be used to ensure a constant supply voltage. If the supply voltage is stable, the frequency variation due to ambient temperature lies within 1%, or 4 to 5 clock cycles in this example. As the resistivity of PEDOT has a much higher variation with temperature (Section 5.2), the temperature-dependent frequency drift of the oscillator is negligible.

The delay measurements taken with the microcontroller do not show any variation with number of measurements, confirming the absence of any systematic measurement variation. It can be seen from Figure 5.12 that the delay measurements follow a Gaussian distribution. A statistical test may be used to verify whether a single, new measurement of the delay time is likely to have originated from the distribution of the undamaged protection grid. This integrity test corresponds to an individual measurement x_i of the protection grid delay time being tested against the null hypothesis $H_0 : x_i \quad S(\mu, \sigma)$. Using the individual delay time, and the known mean and standard deviation of the intact protection grid, the *z*-score may be computed:

$$z = \frac{x_i \quad \mu}{\sigma} \tag{5.7}$$

The *z*-score is equivalent to normalising the measurement data in units of standard deviations of the standard normal distribution [Lane, 1999]. The *z*-score is used as a metric in a (two-sided) test with a defined false rejection ratio. The threshold values outside of which the computed *z*-score is rejected are computed using the cumulative distribution function (cdf) of the standard normal distribution. As the security devices should have a low probability of premature failure (inadvertent self-destruct), it must be ensured that the probability of incorrectly assuming the grid is damaged (confidence) is sufficiently small, e.g. 0.1%. For a single measurement x_i, using a two-sided test, a confidence of 0.1% corresponds to a range of 3.3 standard deviations around the mean. In the example of the protection grid with a standard deviation of $\sigma \quad 12$ clock cycles, this corresponds to a deviation of ca. 40 clock cycles from the reference mean, or around 10% of the mean value.

To increase the resolution of the integrity test, multiple measurements may be taken each time the protection grid is tested for integrity. With *N* observations (measurements) $x_1, \ldots, x_n$, the computation of the *z*-score is changed by the square root of *N* [Lane, 1999]:

$$z = \frac{\bar{x} \quad \mu}{\sigma} \qquad \overline{N} \tag{5.8}$$

This *z*-score is again used in conjunction with the thresholds computed from the inverse cdf as before. Inserting different values for *N*, the measured deviation from the mean at which the protection grid should be labelled as defective can be computed. The result of this calculation is shown in Figure 5.13 against the number of measurements *N*. For this plot, a standard deviation of 12.39 clock cycles was assumed (first row of Table 5.1).

The previous evaluation computes the difference in the mean of the measurement and reference value required to achieve a probability of 0.1% of falsely characterising a protection grid as damaged. What is not captured by this calculation is the probability of actually detecting a damaged protection grid. If only a small number of observations or measurements is carried out, the observation mean may not reflect the mean of the underlying protection grid. In other words, as there is a spread in the measurement values, it may be possible that a damaged protection grid generates a 'valid' result. The probability of correctly rejecting a null hypothesis ("protection grid is intact") is described by the power of a test [Lane, 1999]. Power may thus be

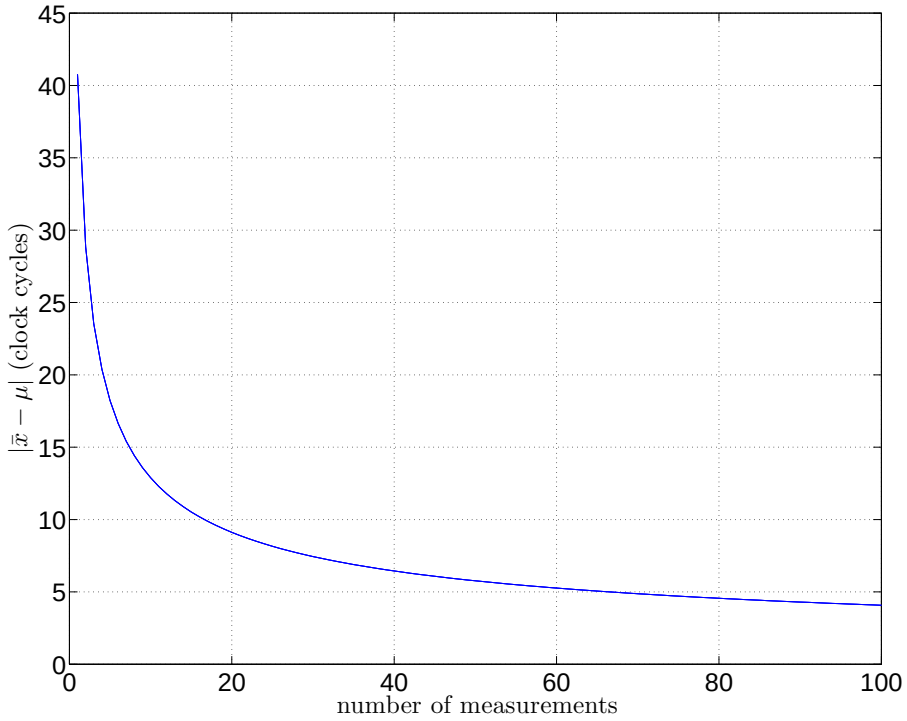

Figure 5.13: Rejection level vs. number of measurements for 0.1% confidence

traded against statistical significance (confidence level) [Cohen, 1988].

A statistical test to quantify the difference of means is the 't-test', if the measurement data are normally distributed [Lane, 1999]. Different versions of the test exist for different types of data. A variant exists for two measurement sets of unequal sample size, if the mean and standard deviation of one measurement set is known and the standard deviation of the second measurement can be estimated. As the standard deviation of the measurement sets was found to be constant for a sample, this variant of the test may be used. This test is used to compute the difference in means of the underlying protection grid that can be detected for a given number of measurements, at a chosen confidence and power [Lane, 1999]. The graph in Figure 5.14 shows the achievable resolution vs. number of measurements for the same standard deviation used in the previous example. It assumes the same confidence level of 0.1%, as before. A statistical power of 80% was chosen, which is a commonly used value in scientific analyses [Cohen, 1988]. For small numbers of measurements (low N), distribution and standard deviation are badly defined, therefore the test becomes unreliable, hence a low detection accuracy.

Figures 5.13 and 5.14 therefore describe slightly different things. Figure 5.13 describes the difference in means between the reference value and measurement average that is required to determine that a device really is compromised. Figure 5.14 in turn describes the change in mean (or resistance) that is likely to be detected with a probability of 80% (and an error rate of 0.1%), if the metric of Figure 5.13 is applied. As the measurement mean approaches the sample mean for large numbers of measurements, the two figures converge.

As the standard deviations for all measurements of the same sample were nearly identical, it may be assumed that changes in resistance due to environmental conditions (temperature) will result in a shift of the mean. Therefore, compensation for temperature variation can be expressed as a linear combination of the measurement distribution of the time-delay, the mea-

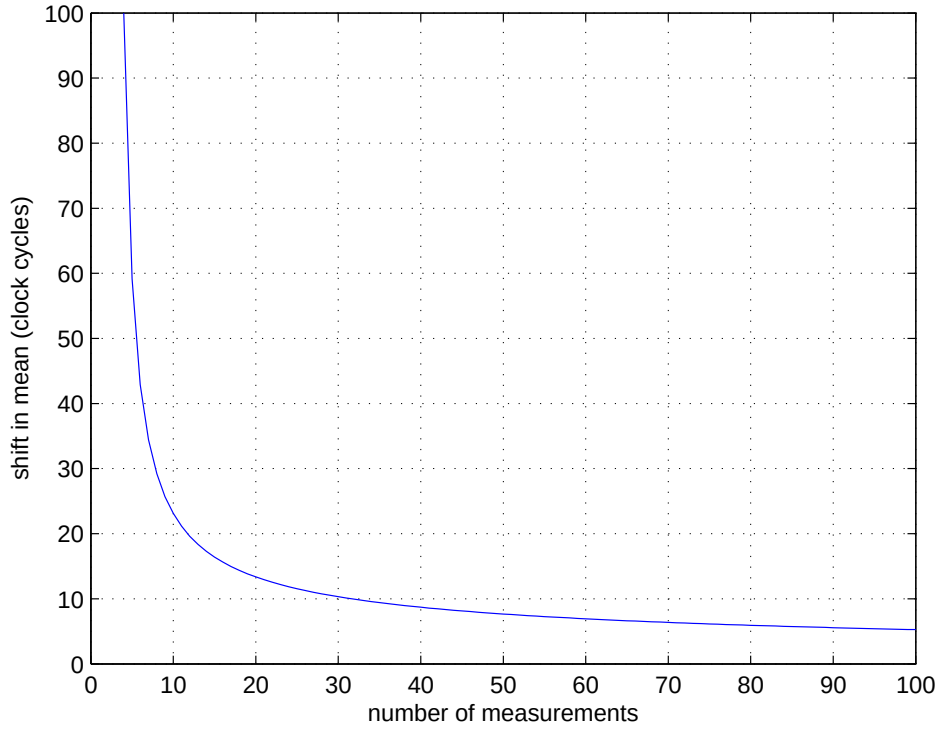

Figure 5.14: Resolution vs. number of measurements for 80% power and 0.1% confidence level

surement accuracy of the ambient conditions, and the uncertainty of resistance change at the measured environmental conditions.

5.4.5 Conclusions

The evaluation of the passive protection grid prototype has shown that it is possible to implement a time delay measurement circuit using a CMOS microcontroller. The measurement results show a spread in values, therefore statistical methods may be used to determine whether the protection grid has been damaged. The resolution of the test may be increased using multiple measurements and comparing the measurement set to the reference distribution. The standard deviations of the measurements were found to be constant over the examined resistance range, and to vary only between different samples. Resistance variations due to environmental conditions (e.g. temperature) are thus expected to affect only the measurement mean. Therefore, compensation for environmental influences essentially consists of the combination of the uncertainties of measuring the ambient conditions, and the uncertainty of the change in characteristics.

5.5 Evaluation of active prototype

To evaluate the sensitivity of active protection grids, the bridge circuit was selected as the simplest read-out scheme. A standard all-inkjet organic transistor was fabricated and characterised to evaluate current organic transistor technology. These circuits are simulated using the measured transistor characteristics, as it was not feasible to characterise fabricated organic TFT bridge circuits. The simulation results are compared to typical silicon transistors, which have

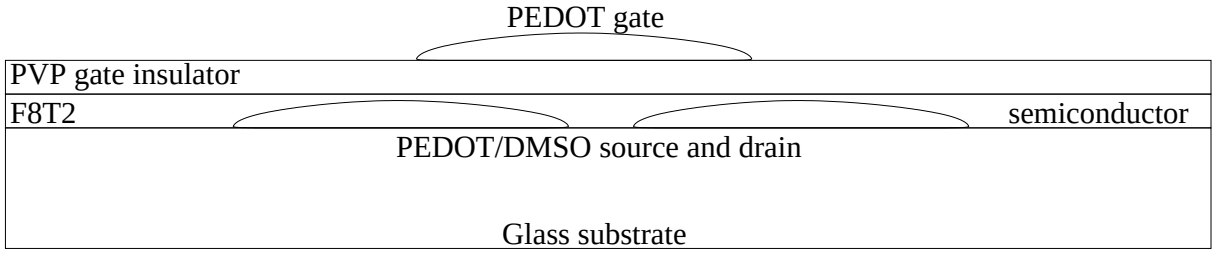

Figure 5.15: Structure of organic transistor

more ideal characteristics. The comparison is used as an aid to estimate what progress may be expected with future generations of organic transistors which are suitable for protection grid applications.

There are two semiconductors that are reported to be suitable for all-polymer organic TFTs. The first one is poly(9,9-dioctyl-fluorene-co-bithiophene) (F8T2), and the second is poly(3-hexyl-thiophene) (P3HT) [Kawase et al., 2005]. Both semiconductors have similar electrical properties, with P3HT achieving a slightly higher carrier mobility [Facchetti, 2007; Speakman et al., 2001]. The air stability properties of the polymers is reported to differ, with F8T2 being more stable than P3HT [Burns et al., 2003; Sirringhaus et al., 2000; Stutzmann et al., 2003]. The ionisation potential of the semiconductor (F8T2: 5.4 eV, P3HT: 4.9–5.1 eV) gives an indication of air stability, with a minimum potential of 5.2 eV being reported as stable [Kawase et al., 2005]. Due to its better stability, F8T2 was chosen for the prototype transistor.

5.5.1 Transistor design and fabrication

To reduce the exposure of the F8T2 semiconductor to oxygen and moisture, a top-gate transistor configuration was chosen, similar to one reported in literature [Rost et al., 2004]. The transistor structure and layers are shown in Figure 5.15.

The source and drain contacts were inkjet printed on a glass substrate using 1:1:0.2 water:PEDOT:DMSO ink. The print settings were identical to those of the PEDOT resistors of the passive protection grid. The transistor geometry was chosen to be similar to other reported F8T2 transistors [Burns et al., 2003; Kawase et al., 2005; Sirringhaus et al., 2000; Stutzmann et al., 2003]. The device width was 2 mm. A minimum nominal channel length of 50 µm was required to ensure that the source and drain contacts were not short-circuited by stray PEDOT drops. The actual channel length varied significantly, therefore no exact definition is possible for this transistor.

After printing the source and drain contacts, both contacts were padded with silver paint to allow probing through the insulator and semiconductor layers. To prevent the silver from influencing the transistor properties, it was deposited away from the transistor channel. Following the transistor recipe used by [Kawase et al., 2005], the F8T2 semiconductor was spin-coated from xylene solution followed by the PVP gate insulator from isopropanol solution. Finally, the transistor gate was printed on top of the PVP layer, using non-DMSO 1:1 PEDOT:water ink, as gates printed with PEDOT/DMSO ink dissolved enough PVP to short-circuit the transistor through the insulating layer. A microscope image of the organic TFT is shown in Figure 5.16. The gate contact was printed with sufficient width to cover the channels of all transistors in the same row on the substrate.

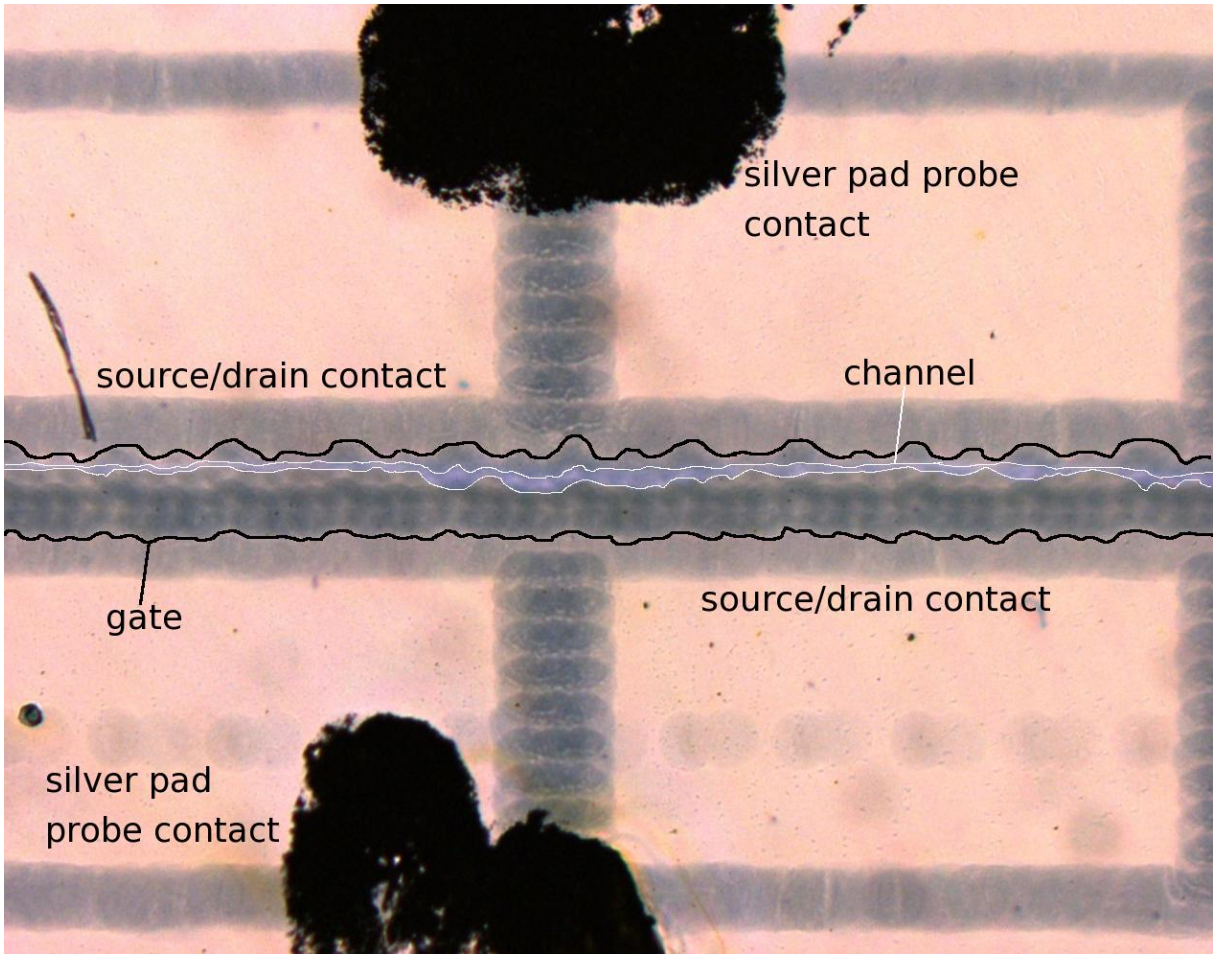

Figure 5.16: 50 magnified microscope image of F8T2 OTFT (white and black lines added to highlight edges)

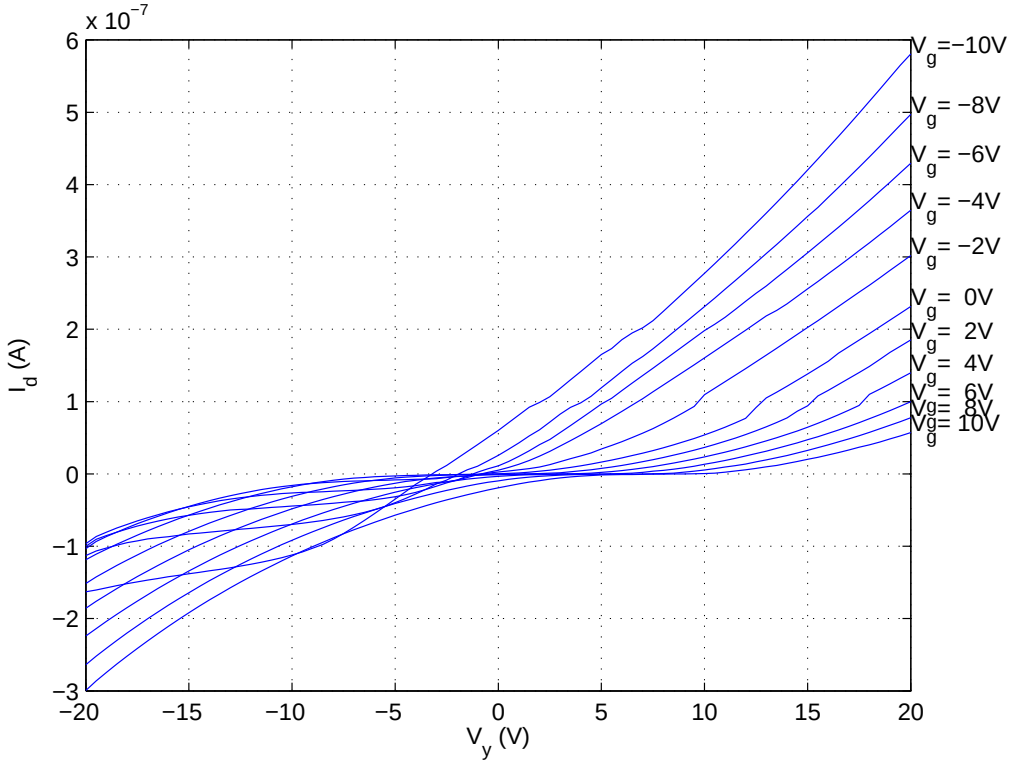

Figure 5.17: Measured characteristics of F8T2 OTFT

5.5.2 Transistor characteristics

The transistor characteristics were determined using a HP4156A semiconductor parameter analyser. To represent the target voltage range of 10 V, the examined source/drain voltage range was 20 V, at gate voltages between 10V.

The measured transistor current vs. source/drain voltage (at various gate voltages) of the best transistor is shown in Figure 5.17. The transistor shows p-channel field-effect transistor (FET)-like behaviour. Despite using non-DMSO ink, the transistor gate exhibits significant current leakage, altering the measured drain current. The I-V measurements were found not to be repeatable as the current varied between measurements. The off-state leakage current, in particular, increased with each measurement. While the general I-V characteristics of the transistors are similar to the values reported in literature (e.g. [Kawase et al., 2005]), the transistors degraded between measurements. Degradation was not observed in the transistors reported in literature, however these measurements were carried out in nitrogen rather than ambient air. With the observed degradation of characteristics between measurements, it is not feasible to measure the balance point of fabricated bridge circuits. Instead, a simulation of the bridge circuit is carried out using the measured transistor I-V curve from Figure 5.17. The problem of degradation in air may be avoided with appropriate encapsulation. The evaluation of compatible encapsulation methods, which provide protection from oxygen and moisture, is beyond the scope of this dissertation.

5.5.3 Simulated bridge

One half of the bridge circuit is examined using the load-line technique described in Appendix F of [Horowitz and Hill, 1989]. The load-line technique is derived from Kirchhoffs' Current and

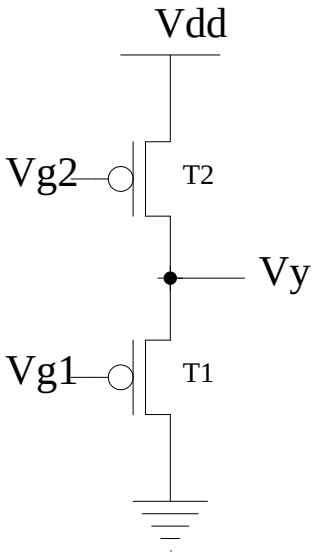

Figure 5.18: Schematic of half bridge circuit

Voltage Laws. It is a graphical method to analyse a series connection of two devices. For the example circuit in Figure 5.18, the I-V curve of the first device (T1) is plotted as usual. The I-V curve of the second device (T2) is plotted into the same figure, but for voltages of $V_{dd} \quad V_y$. The second plot is essentially reversed, with the origin at V_{dd}. The output voltage corresponds to the intersection of the two I-V curves, as required by Kirchhoffs' Current Law. Using the voltages and circuit shown in Figure 5.18, the following boundary conditions are used to construct a load-line diagram for the two transistors in one branch of the bridge:

Transistor 1:

$$V_{gs,1} = V_{g,1} \tag{5.9}$$
$$V_{ds,1} = V_y \tag{5.10}$$

Transistor 2:

$$V_{gs,2} = V_{g,2} \quad V_y \tag{5.11}$$
$$V_{ds,2} = V_{dd} \quad V_y \tag{5.12}$$

To illustrate the output voltage V_y for different input conditions, the drain current of T1 and T2 is plotted against V_y for different gate voltages (Figure 5.19). The supply voltage was chosen to be -10 V, which is defined to be the maximum feasible operating voltage to interface with standard low-cost silicon microchips. For comparison purposes, the same plot is constructed using the characteristics of a standard p-channel silicon MOSFET ($l = 0.5$ µm, $w = 2$ µm) in Figure 5.19(b).

The blue curves in Figure 5.19 correspond to T1. As the transistor has a fixed V_{gs} and a varying V_{ds}, the curves are identical to Figure 5.17. The red curves correspond to the second transistor, with V_{gs} and V_{ds} depending (linearly) on V_y. The red curve therefore represents the transistor current for a fixed offset between V_{gs} and V_{ds}. The intersection points of the blue and red curves mark the output voltage V_y for the given gate voltages, which are labelled at the end of each curve.

Comparing Figure 5.19(a) and 5.19(b), it can be seen that the F8T2 transistor is in its linear region ($V_{ds} < V_{gate} \quad V_{threshold}$) for most of the target voltage range, while the silicon transistor

reaches saturation. Significant leakage currents of the organic transistor are also visible compared to the silicon transistor. Both transistor types have in common that the drain currents show reduced dependency on the gate voltage in the linear region. This leads to low sensitivity of the output voltage V_y to the gate voltage V_{g1} of T1. This effect is especially visible in Figure 5.19(a), for gate voltages around -6 V to -8 V (and -10 V if it were not shifted to the left). As the protection grid scheme intends to map multiple inputs to a single output value, operation in the linear region is to be avoided.

Shifting both the supply voltage and the gate voltage of T2 (V_{g2}) by the same amount corresponds to a shift of the red curve along the x-axis. By relaxing the operating voltage constraint, better sensitivity of the output voltage can be obtained for the F8T2 transistors. Figure 5.20 shows the simulated characteristics obtained for an operating voltage range of 20 V. It can be seen that the the intersections of the red and blue curves have been shifted to a more saturated region of the T1 curves. This ensures that the output voltage V_y also depends on the gate voltage of T1, as with the silicon MOSFET. If the typical operating voltages of organic transistors can be reduced, then it should be possible to implement this read-out scheme using a standard low-cost silicon microchip.

Another difference between the OTFT and the silicon MOSFET are the current levels due to the reduced electron mobility in organic materials compared to silicon. Despite the OTFT having a very high width/length ratio, the current of the organic transistor is four orders of magnitude less than the current through the silicon transistor. The low current levels of the organic transistor of 10 nA to 100 nA are far more vulnerable to noise than the mA of the silicon transistors. In addition, any leakage currents or input currents to measurement circuits have a larger effect, the lower the transistor current. A higher transistor current would therefore be desirable for operating a protection grid.

5.5.4 Conclusions

From the measured transistor characteristics and the observed degradation, it is apparent that inkjet printed active protection grids are not yet feasible. However, if transistor characteristics are improved (lower operating voltages and higher currents), and more stable devices can be made, the bridge circuit may well be used to characterise transistors.

5.6 Conclusions

Temperature dependency The temperature dependent resistance of PEDOT was shown to be repeatable between samples, so long as the samples are heated in the same ambient conditions (humidity). If the atmosphere is not inert, degradation may occur at elevated temperatures. Depending on the required temperature range, the tested Araldite Rapid epoxy may either be useful as a tamper-prevention encapsulation, or be inferior to more temperature-stable epoxies.

Ageing Non-encapsulated PEDOT is not acceptable in terms of its ageing behaviour, neither for pure PEDOT nor PEDOT/DMSO. The epoxy encapsulated PEDOT is significantly better, but still shows some degradation for pure PEDOT, which should ideally be eliminated. A better encapsulation material would shield the epoxy from the three main causes of degradation (water, oxygen, and UV light).

Passive prototype The delay time of evaluated passive prototype may be described by a lumped rather than a distributed RC delay model. The resistance change of the grid line was found to be easily detectable as change in the measured delay time. The measurement values taken with the

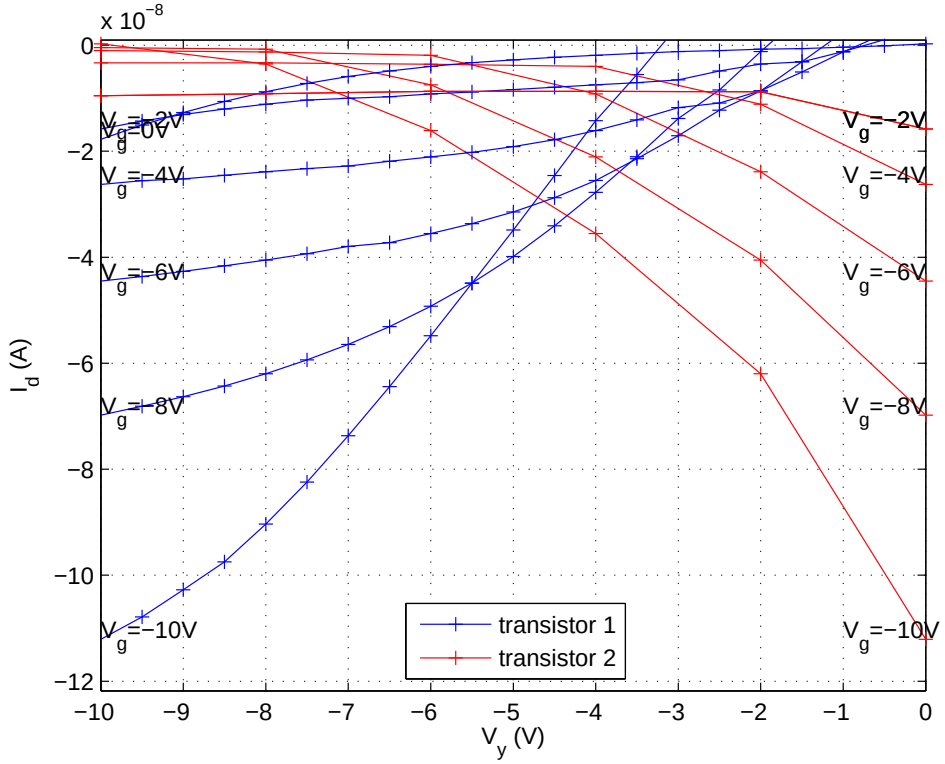

(a) F8T2 OTFT transistors

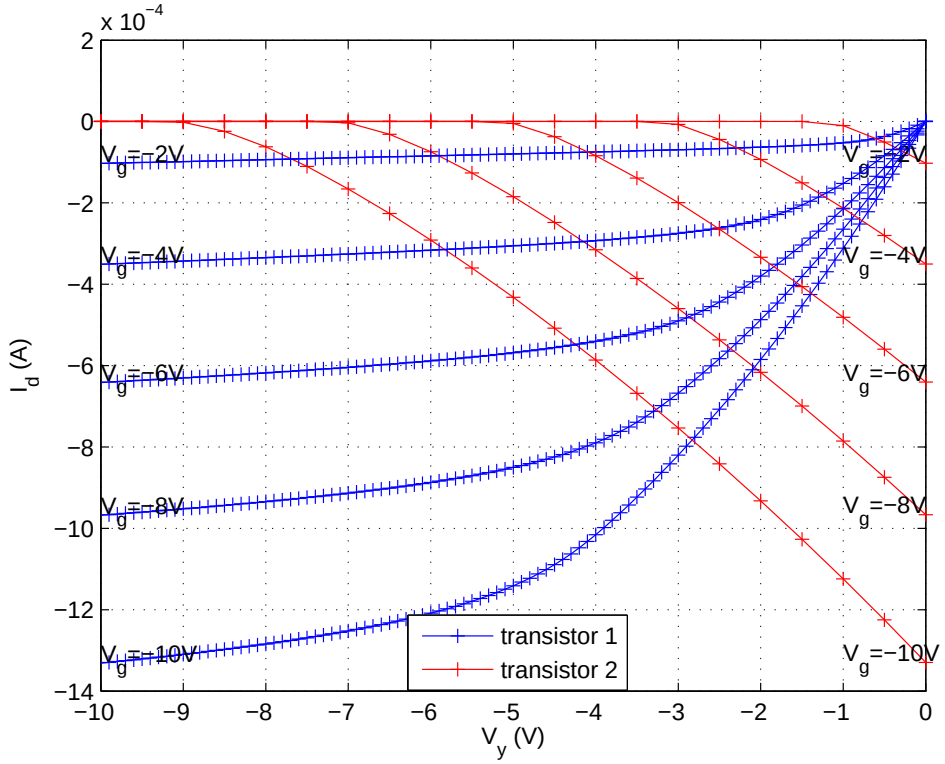

(b) Silicon MOSFET transistors

Figure 5.19: Simulation results of two transistor bridge

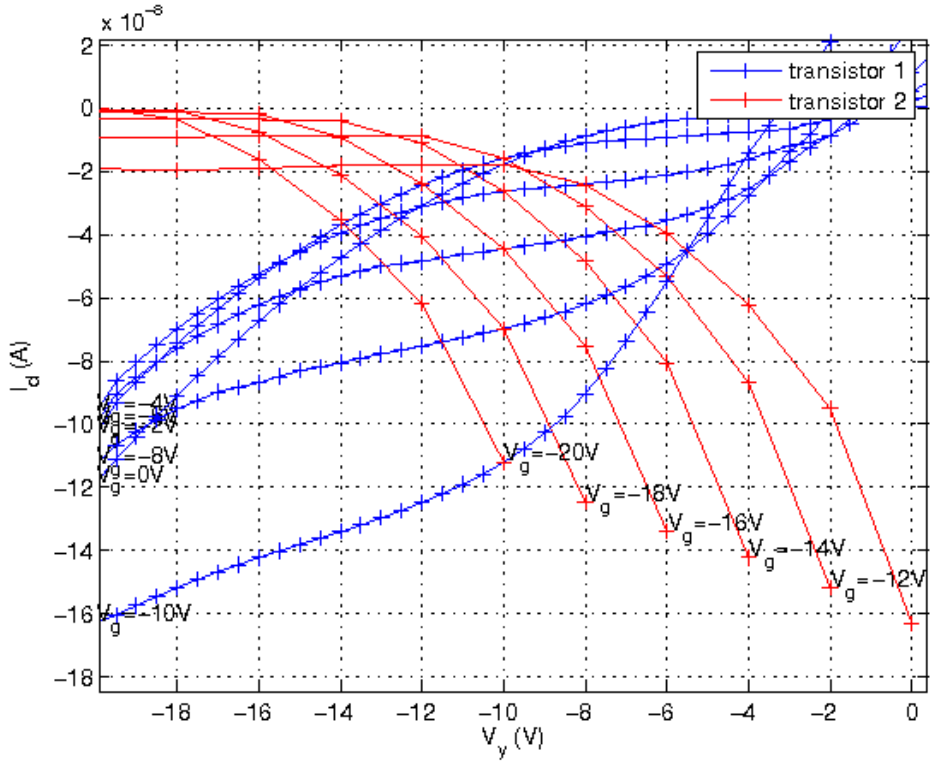

Figure 5.20: Simulation results of two transistor bridge at lower supply voltage

PIC microcontroller were distributed around a mean, rather than being a fixed value. Statistical tests and multiple measurements to determine integrity may be used to enhance the accuracy of the integrity measurement.

Active prototype The fabricated organic TFT test device shows significant degradation between measurements, as well as low current levels. Therefore, experimental evaluation of the bridge circuit is not feasible. A simulation showed that different output voltages can be achieved, though the sensitivity with respect to the gate voltage of the first transistor, and robustness against measurement noise are lacking. Using better transistors, the evaluated bridge circuit may be feasible for a protection grid.

In summary, the evaluation of the practical aspects of the protection grids has shown that changes to implementation details are required to assemble a robust protection grid. The tested Araldite epoxy has been shown to be inadequate both in terms of temperature stability, and in terms of preventing ageing of the protection grid. The passive prototype assembly showed that the RC delay measurement scheme can achieve good sensitivity to resistance changes of the protection grid. The active protection grid on the other hand is far from meeting requirements.

Chapter Six

Power Analysis Countermeasure

6.1 Introduction

A number of countermeasures against power analysis attacks have been published [Benini et al., 2003; Grassert and Timmermann, 2001; Kulikowski et al., 2006a; Macé et al., 2005; Plana et al., 2003; Sokolov et al., 2005, 2004; Tiri and Verbauwhede, 2004a; Tiri et al., 2002; Yu et al., 2003]. While reductions were made in the data-dependency of the power consumption, a small imbalance always remained. This remaining imbalance leaves the devices vulnerable to power analysis. A reduction in the data-dependency of the power consumption reduces the signal to noise ratio for an attacker. This eventually hides the data-dependent power consumption signal in noise, making different data indistinguishable in the first instance. Better mathematical models to enhance differential power analysis (see Section 1.2.2) help an attacker recover the secret data despite the reduced signal to noise ratio.

The large variety of power analysis methods hamper the definition of a precise quantitative measure of security, and what signal to noise ratio is required to recover secret data. However, on the lowest level, all power analysis variants involve some form of processing of the instantaneous power consumption of a security device. Therefore, a conservative guideline can be defined that any difference between power consumption traces that is visible in measurements must be assumed to leak data.

In this chapter, the problem of the remaining imbalance in nominally balanced circuits is evaluated. Balanced dynamic dual-rail logic gates were designed and analysed, both in simulation and as fabricated microchips. The work for this chapter was motivated by Professor David Harris of Harvey Mudd College, Claremont, California whilst he was on sabbatical in Cambridge. Prof. Harris proposed to modify dual-rail domino logic cells to achieve balance. Under the standard power consumption model, if equal capacitances are discharged for all inputs, power consumption is constant for all input data [Tiri et al., 2002]. Therefore, a perfectly balanced layout should result in data-independent power consumption.

As multiplication is relevant in modular exponentiation used in public key cryptography [Messerges et al., 1999], a multiplier was chosen as a test structure. An array multiplier is a tiled design, with each tile consisting of only three different logic gates [Weste and Harris, 2004]. As the structure is regular and simple, it can be laid out by hand, to give full control over layout geometry and parasitic capacitances. The small number of different components facilitates the analysis of power consumption.

The goal of this evaluation is not to develop a universal, new logic style to counter power analysis. Instead, the balanced dual-rail domino logic is used as a starting point to identify causes of data-dependency of the power consumption beyond discharging different sized parasitic capacitances. The findings are used to derive guidelines for designing better balanced logic gates. Therefore, the design of a full security device is not required, nor is it a problem that the fully manual layout method is labour intensive.

The evaluation is carried out in multiple stages. The three logic gates that form the multiplier cell are first modified from their original domino logic to achieve balance. These cells are assembled to design an 8-bit multiplier suitable for fabrication. To test the balance properties of the circuits, the multiplier circuit and the individual cells are simulated using the SPICE tool. The findings of the simulation are then compared to measurements carried out on fabricated multiplier chips obtained through the Metal Oxide Semiconductor Implementation System (MOSIS).

6.1.1 Evaluation procedure

The evaluation of the balanced dual-rail domino logic is carried out in simulation and measurement, as both have distinctive features. Simulation of logic gates allows the power consumption of logic gates to be evaluated at very fine-grained detail and temporal resolution. Currents through individual transistors or small wires on a microchip are typically not accessible to direct measurement, but can be estimated using the SPICE simulator. These wires may be embedded in the chip, and a measurement would require alterations in the layout to permit probing access. These alterations influence the power consumption behaviour, defeating the purpose of the measurement. The disadvantage of simulations is that they only reflect an idealised model of the circuit, which may not be accurate. Fabricated geometries or transistor characteristics have tolerances, which may vary from the ideal, skewing the results. In addition, simulations are restricted to the circuit, and do not include the connections to the chip packaging, power supply variations or substrate coupling.

The measured power supply current traces may be assumed to be low-pass filtered [Fournier et al., 2003] and noise-added versions of the simulated traces. The capacitances of the chip packaging and circuit board, as well as the measurement setup (typically resistor and oscilloscope probe) store charge, which may momentarily supply the device, smoothing out the measured current. In addition, noise is introduced into the measurement by the wires capturing external EM interference, mismatches in characteristic impedances causing signal reflections, an imperfect power supply, and the measurement noise in the oscilloscope itself. Measuring the current on a fabricated device also has the advantage that systematic transistor variations due to the layout are automatically included, and do not require a laborious monte-carlo simulation.

The details of the measurement setup vary between attackers, and with it the achievable measurement accuracy. If differences between power consumption traces can be detected using the experimental setup, then the results clearly demonstrate that the device is insecure. The converse, however, is not proof of security. The likelihood of detection outside the packaging is correlated with the variance of the simulated current traces. While metrics of the variances between traces have been defined [Gebotys, 2004; Tiri et al., 2002], using them in the context of evaluating single logic gates is misleading. The amount of imbalance becomes irrelevant if differences in the power consumption traces can be measured. In this case, a device is vulnerable to differential power analysis techniques, and therefore insecure. The magnitude of the imbalance only becomes relevant if the variation is close to indistinguishable. In this case, the magnitude of the imbalance correlates with the number of measurements required to compensate for the measurement noise swamping the data-dependent signal. Apart from variations in total charge consumed per clock cycle, there may also be a time shift in power consumption traces (data-dependent jitter). Current imbalances of short duration are more easily obscured by

the low-pass filtering effect of the packaging and measurement setup, while longer time shifts are more easily visible on the power supply pins of a device.

6.2 Design

6.2.1 Secure Array Multiplier

The array multiplier is a tiled array of carry-save adders, arranged in a regular layout as shown in Figure 6.1 (reproduced from [Weste and Harris, 2004]). An n-bit array multiplier consists of $n^2 + n$ tiles, as the bottom row of carry-propagate adders can be made from carry-save adders. Each carry-save adder tile comprises three logic gates: a 2-input AND, a 3-input XOR, and a 3-input majority gate [Weste and Harris, 2004], connected as shown in Figure 6.2.

6.2.2 Balanced dual-rail domino logic

Domino logic is a combination of dynamic logic gates and static (standard) inverters [Weste and Harris, 2004]. In dynamic logic gates, the pull-up/PMOS transistor network is replaced by a single clock-driven pre-charge transistor [Weste and Harris, 2004]. If the clock is low, the pre-charge transistor conducts current and charges the output node. When the clock rises, the pre-charge transistor ceases to conduct and the pull-down network of the logic gate can discharge the output node, if appropriate.

If any of the inputs to the NMOS logic transistors are high, the transistors are conductive. If a conducting path to ground is established during pre-charge, the N- and PMOS transistors will be in contention and waste power. If it cannot be guaranteed that all inputs to the logic gates are low during the pre-charge phase, a foot transistor (clocked NMOS transistor between the logic gate and ground) may be added [Nowka and Galambos, 1998]. A footed design is chosen to prevent accidental discharge. Another requirement for the input signals is monotonicity [Weste and Harris, 2004]. As the output node can only be discharged once during evaluation, the input signals must only rise monotonically from low to high, to ensure a correct result [Weste and Harris, 2004]. The output of a dynamic logic gate is high during pre-charge, and may switch to low during evaluation (monotonically falling). To meet the monotonicity requirements, the output nodes are inverted using a static inverter in domino logic [Weste and Harris, 2004].

A simple node counting technique was used to balance the logic gate design. In this technique, the nodes in a logic gate are grouped according to the number of connected transistors. Under the assumption of a symmetric layout, the pull-down transistor network is designed so that the number of discharged transistor nodes groups is equal for all possible inputs to the logic gates. To ensure all nodes of the logic gate are charged during the pre-charge phase, additional PMOS transistors are connected to each node in the logic gate.

Examination of the three logic gates showed that the XOR gate is balanced by default, and all that is required is a symmetric layout of the gate. The AND and majority gates, however, required significant modification. To prevent cross-talk between signal wires, power or ground wires were interspersed between the differential wire pairs.

6.2.3 Gate design

6.2.3.1 XOR Gate

The schematic of the XOR gate, as it is shown in Figure 6.3, was suggested by Prof. David Harris. The output inverters of the domino logic gate are overlaid with a blue shading. The PMOS pre-charge transistors are highlighted in green, and the foot transistors are highlighted in red. The

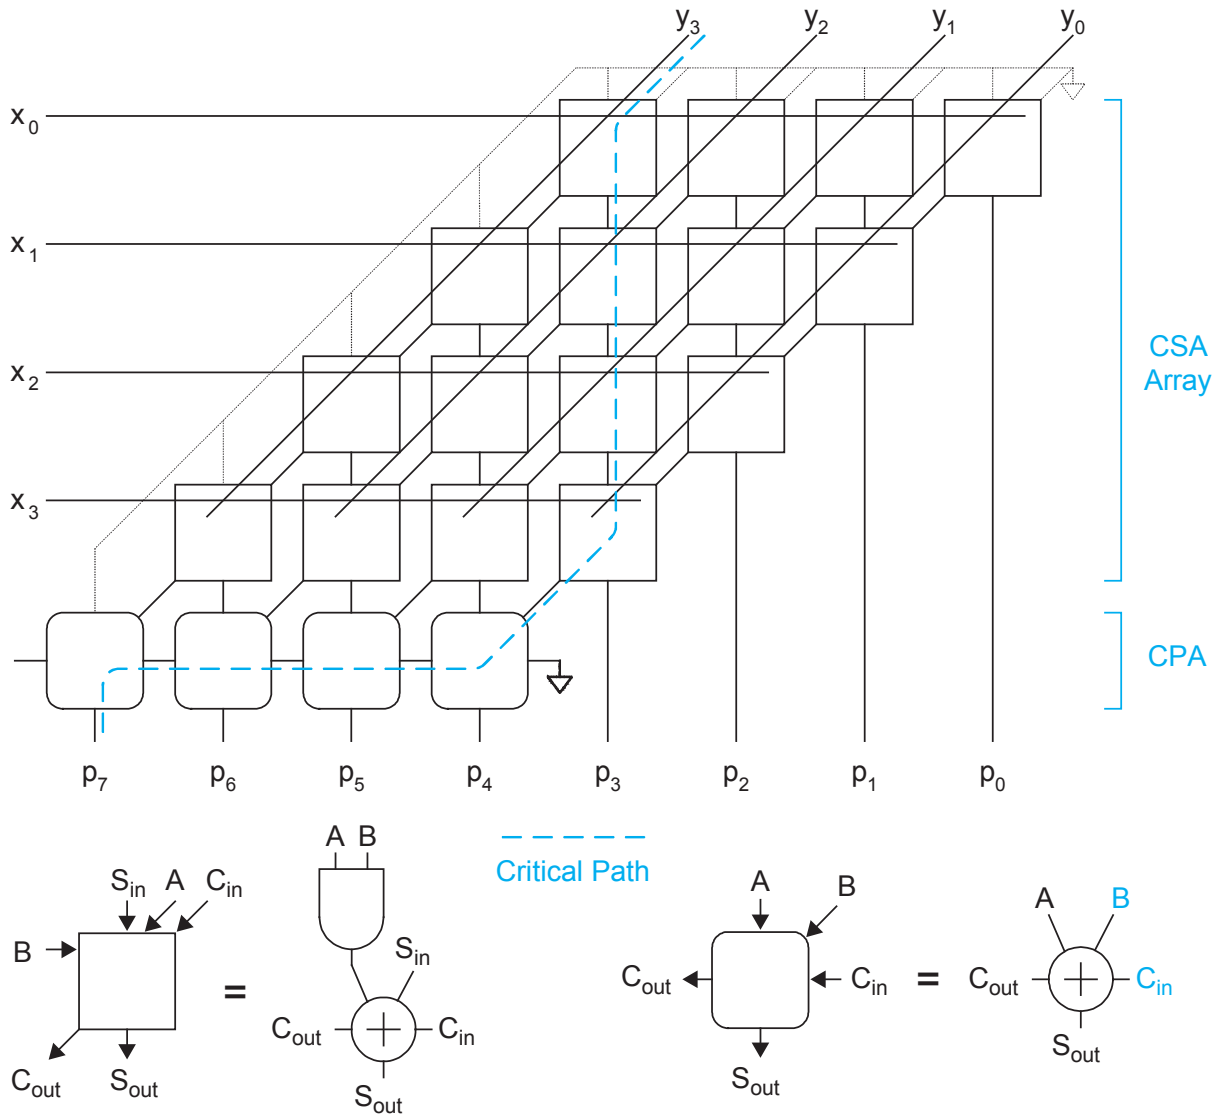

Figure 6.1: Schematic of the Array Multiplier[a]

[a]Weste/Harris, CMOS VLSI DESIGN: A CIRCUITS AND SYSTEMS PERSPECTIVE, ©2005 Pearson Education, Inc Reproduced by permission of Pearson Education, Inc.

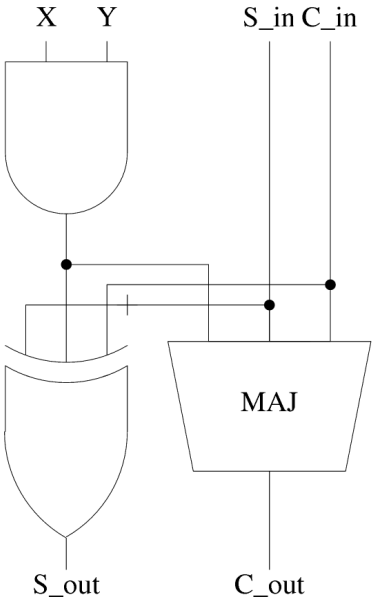

Figure 6.2: Array Multiplier Cell

remaining transistors are the evaluation transistors of the logic gate. The function of the logic transistors might be described as a set of multiplexers controlled by each of the inputs. During evaluation, the A-transistors connect either the left two or the right two B-transistors to ground. The B-transistor pair in turn selects which C-transistor pair is connected to ground. Finally the C-transistors choose which of the two inverted output nodes is discharged.

As each input to the logic gate (A, B, or C) essentially selects one of two nodes, the gate is balanced by default, so long as the layout is symmetric. Figure 6.4 shows my hand-crafted symmetric layout of the XOR gate.

The transistor spacing and layout was arranged to fit into the regular tiled layout, as well as the routing requirements for the output. The transistors are green, with different hatching to show the doping of the silicon (dotted for p-type transistors, diagonal hatching for n-type transistors). The metal layers are colour coded in blue (metal 1, lowest), magenta (metal 2, middle), and yellow (metal 3, top). The polysilicon gates of the transistors are pink, and the vias are filled black squares. For better orientation, the output inverters, pre-charge and foot transistors are highlighted in the same colours as the schematic. The evaluation transistors are emphasised in the black box.

6.2.3.2 AND Gate

The standard dynamic dual-rail AND gate is not symmetric by default. The 'true' branch (AND.h) comprises two series transistors, and the 'complement' branch (AND.l) two parallel transistors. The schematic of the balanced AND gate in Figure 6.5 shows the simple extension, which was also suggested by Prof. Harris. The highlighting scheme used in Figure 6.5 is identical to the one used for the XOR gate. To achieve balance, the logic transistors in each branch are doubled, so that there are four transistors in each branch. If the inputs are connected as shown in the schematic, one of the two output nodes is discharged, as well as two of the nodes between the series transistors, independent of input.

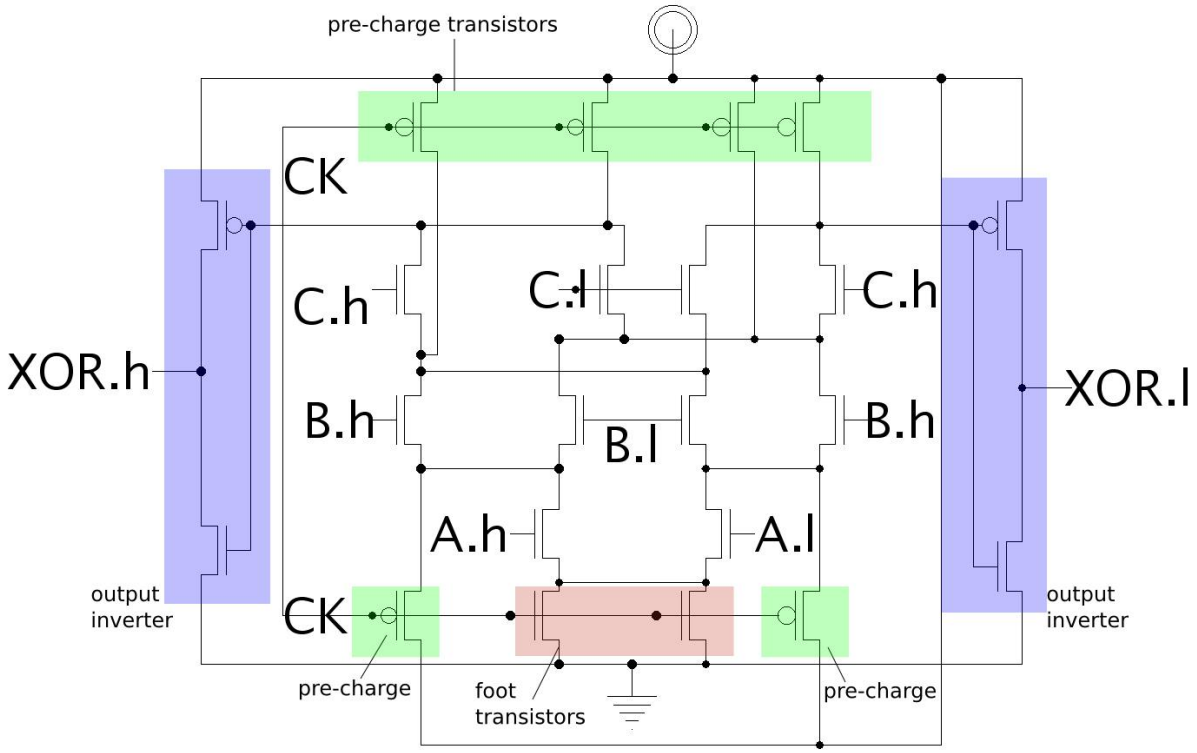

Figure 6.3: Schematic of XOR gate

I crafted the layout of the AND gate (Figure 6.6) to ensure the discharge of identical capacitances for all input cases. The layout is highligted in the same colours as the schematic (Figure 6.5). The eight evaluation transistors (inside the black rectangles) are arranged such that the input of each half of the gate is in the centre of the transistor pair closest to the (red) foot transistors. Correspondingly, the output of the dynamic logic part of the gate is in the centre of the transistor pair furthest from the foot transistors.

6.2.3.3 Majority Gate

The majority gate required substantial re-design to balance the discharged capacitances in each clock cycle. On evaluation, it was found that Prof. Harris' initially proposed design was not balanced. A new gate was designed to achieve balance with a minimum size footprint, but not adding transistors for all possible input combinations and permutations. This would require 144 transistors, as there are 8 possible input data, times 3 transistors in each combination, times 6 positional permutations per combination to ensure an even number of discharged nodes. As shown in Figure 6.7 (highlighted as before), the new design consists of a total of 18 evaluation transistors. A standard dual-rail majority gate requires at least 10 evaluation transistors, therefore the size penalty is less than twice the standard number of transistors.

Using the node numbering from Figure 6.7, the discharge status for all nodes and inputs is shown in Table 6.1. It can be seen that, independent of input, the same numbers of nodes between transistors are discharged. Some of the nodes are discharged indirectly, marked by the dagger symbol ( ). Indirect discharge takes place when a node is discharged via the output node. In this case, the transistor between that node and the output is conducting, but transistors between the node and the foot transistor are not. The discharge path for that node thus goes through three series transistors rather than directly through a single transistor.

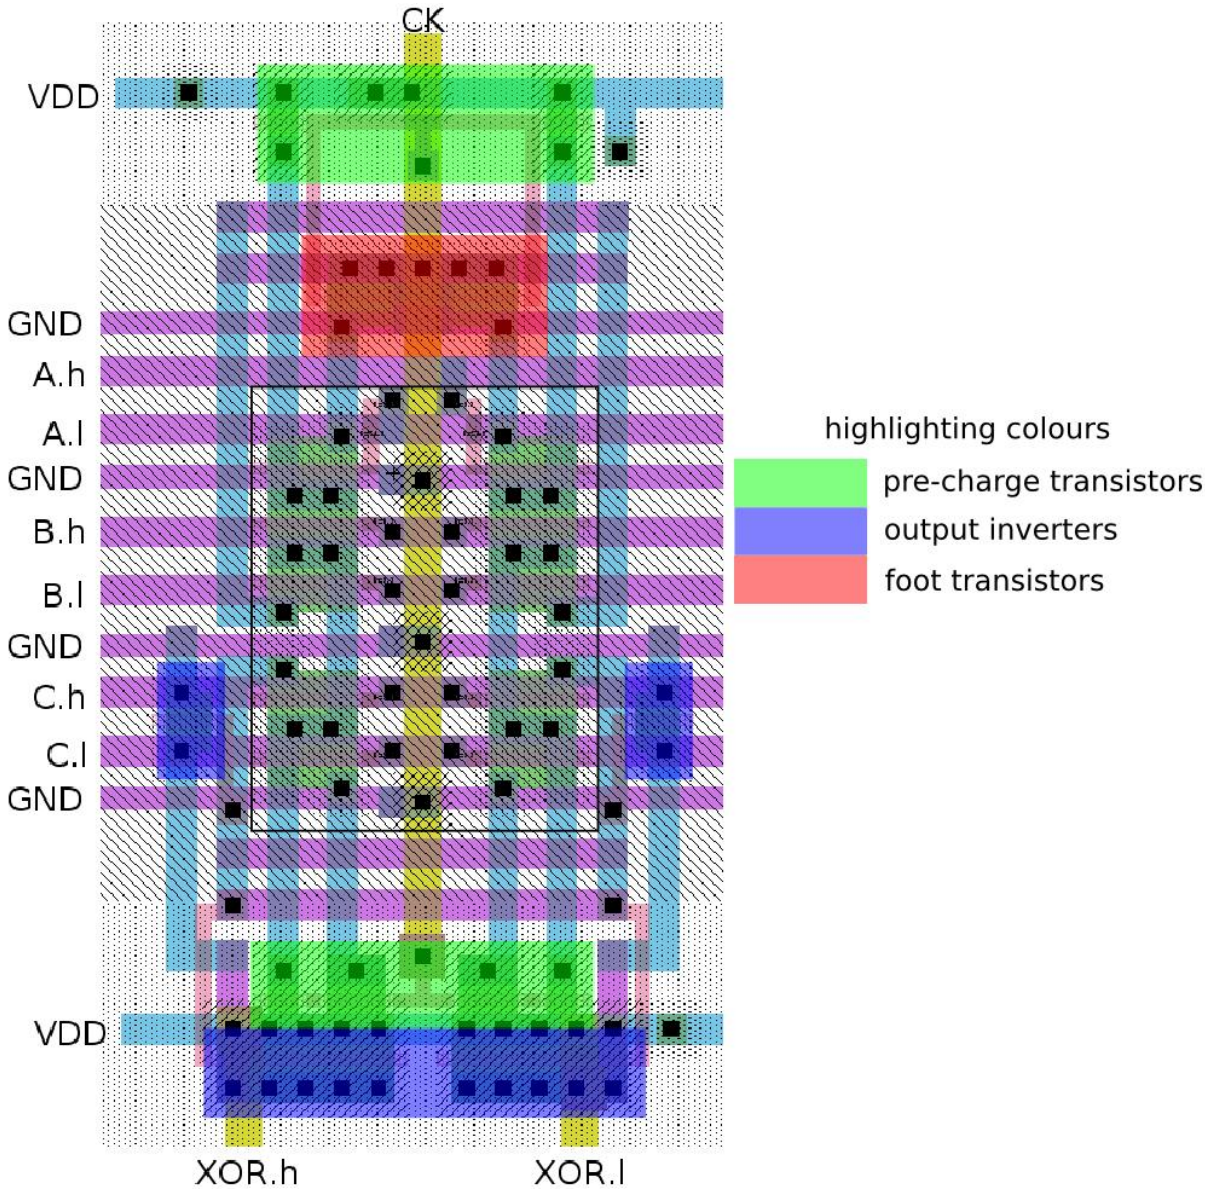

Figure 6.4: Layout of XOR gate

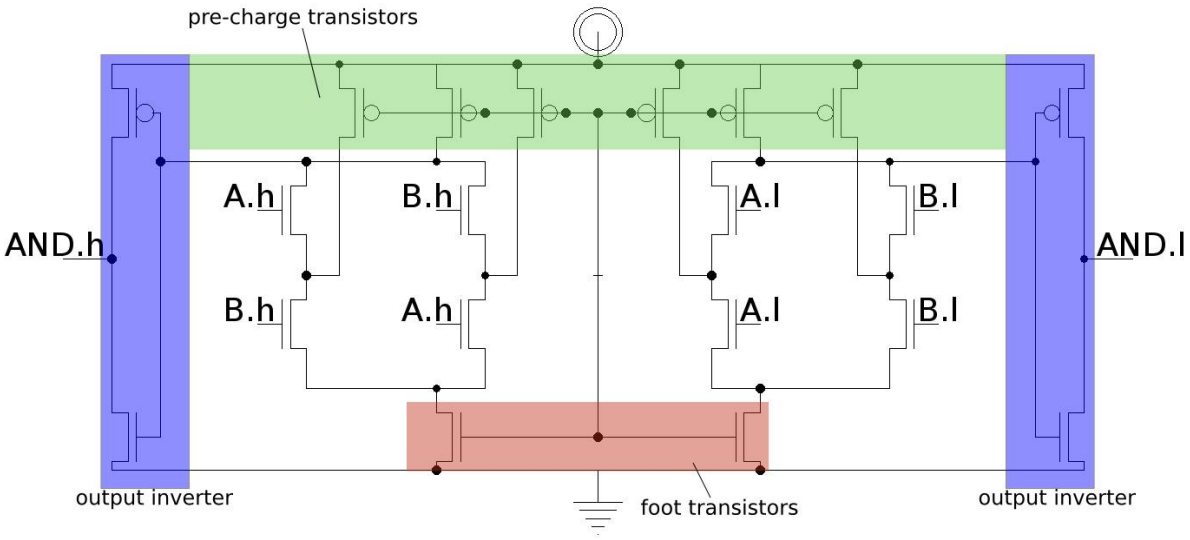

Figure 6.5: Schematic of AND gate

The layout of the majority gate is shown in Figure 6.8, using the same highlighting scheme as in the previous schematics. The six groups of three transistors are visible inside the black box. The symmetry axis of the layout runs horizontally through the centre of the gate, with the 'true' side on top, and the 'complement' below.

6.2.3.4 Carry-save adder cell

The layout of the assembled carry-save adder cell is shown in Figure 6.9. The individual logic gates are arranged in the following manner: the AND gate is placed above the side by side arrangement of the majority (left) and the XOR gates (right). The x-input runs horizontally, and the y-input vertically, corresponding to the top-level schematic of the array multiplier (Figure 6.1)[Weste and Harris, 2004]. The output of the majority gate represents the carry-out of the cell, and is routed vertically down. The output of the XOR gate is placed in the bottom right hand corner of the cell, so that it can be routed diagonally to the next cell.

6.2.4 Multiplier assembly and chip

The 8-bit multiplier design (layout shown in Figure 6.10) corresponds to the schematic in Figure 6.1. The tiled structure is visible, with an example carry-save adder cell framed with a black square. Input and output shift registers are added around the periphery for serial data transfer into and out of the multiplier. The carry-save adder cells of the first row and column (above and to the left of the highlighted carry-save adder cell) are simplified to AND gates, as the sum and carry inputs for the first row and column are always zero. Therefore the carry output is always zero, and the sum output equals the result of the AND gate.

6.3 Simulation of logic cells

Simulations were carried out to analyse the currents in the carry-save adder cell and the full array multiplier. The layout diagrams of the balanced dual-rail domino gates and circuits were processed with Mentor Graphics' Calibre parasitic extraction tool. Calibre generates a SPICE netlist that includes all layout parasitic resistances and coupling capacitances, to represent the

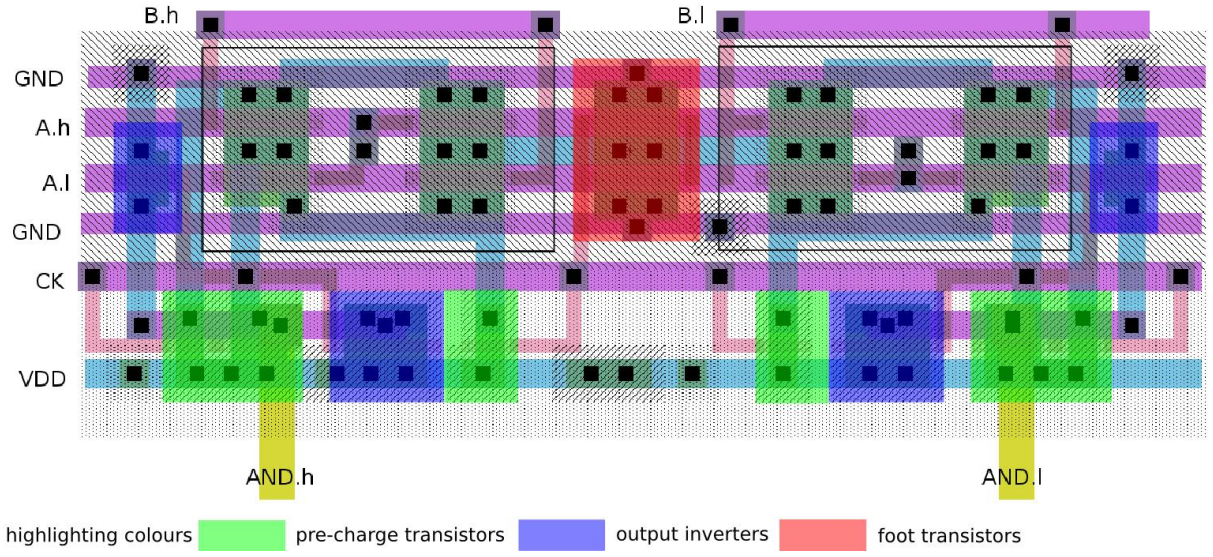

Figure 6.6: Layout of AND gate

circuit as it is laid out on the microchip. The transistor models for the fabrication run of the multiplier chip were supplied by MOSIS.

Evaluation of the secure multiplier design is carried out hierarchically, as a circuit can only be as secure as the components from which it is assembled. Simulation of smaller units also allows sources of data-dependency to be identified more easily. The carry-save adder cell is simulated and analysed in this section, before the full multiplier assembly is evaluated as part of the comparison between simulation and measurement in the next section.

6.3.1 Carry-save adder cell

The smallest design unit of the multiplier is the carry-save adder cell. The carry-save adder cells were laid out as a single, flat cell, in order to ensure routing requirements were satisfied for the multiplier and the cell area was minimised. While the carry-save adder cell is not strictly the lowest hierarchy level, the design is nevertheless simple enough to allow analysis of the entire unit of three logic gates. Within the adder cell, the only interaction between the logic gates is the output of the AND gate feeding into the majority and XOR gates.

6.3.1.1 Setup and procedure

As the carry-save adder cell has only four inputs, simulations were carried out for all input cases. The resolution of the simulation was set to 1 ps. Two clock cycles were simulated to confirm that the initial state of the circuit in the simulation run did not influence the results.

6.3.1.2 Results

The graph in Figure 6.11(a) shows the input currents to the multiplier assembly for all 16 different input cases over one clock cycle. A close-up of the input current to the cell at the evaluation clock edge is shown in Figure 6.11(b). The input data for each group of traces are labelled in the close-up graph. The order of the four input bits from left to right is x, y, sin, and cin.

The largest data-dependent imbalance in the power consumption traces of the carry-save adder is the timing variation of the negative current peak. If sin is equal to cin, then the current peak occurs approximately 0.5 ns earlier than if sin and cin are not equal.

The timing of the output from the three logic gates also varies with input data. The 50% to 50% delay between the clock edge and the output of the AND gate varies between 439 ps for a '0011' input (bit order x, y, sin, cin), and 516 ps for a '0100' input, a difference of 78 ps. The XOR gate shows a similar range in values, with a minimum delay of 877 ps ('0011') and a maximum delay of 962 ps ('0100'). The delay of the majority gate varies between 374 ps ('0000') and 888 ps ('0101'). Despite the timing variation, the total charge and energy consumed in one clock cycle vary only by 5% between minimum (49.88 fC) and maximum (52.31 fC).

6.3.1.3 Discussion

[Kulikowski et al., 2006b] identified the cause for the major timing imbalance as early propagation. If the inputs to a logic gate do not arrive simultaneously, the logic gate may set its output before all inputs have arrived. In the simulation, all inputs (x, y, sin, cin) were set simultaneously, together with the clock. As all logic gates are connected to the same clock signal, the output of the AND gate is delayed with respect to the inputs sin and cin. The standard majority gate evaluates as soon as sufficient information is present on the inputs to determine the result. The majority gate will therefore switch if two of the three inputs are present and equal. In the carry-save adder, the output of the AND gate is delayed compared to the other inputs. Therefore, if sin equals cin, early propagation of the majority gate takes place.

The timing of the majority gate output is evidence of early propagation, as the shortest delay time during early propagation is less than the delay time of the AND gate. Thus the majority gate switches its output before the result of the AND gate is set. The maximum delay time is roughly equal to the delay of the AND gate plus the early-propagation delay of the majority gate, explaining the larger timing variation over the other two logic gates.

The AND gate may also show early propagation. If one of the inputs is zero, the gate evaluates to zero. In this design, the AND gate is exclusively fed from the global x and y inputs, e.g. from a shift register or a data bus. Without the shift registers, buffering may be used to synchronise the inputs and prevent early propagation. The XOR gate is immune to early propagation, as the output always depends on all inputs.

To examine the smaller timing variations remaining in the evaluation current, simulations were carried out without the problem of early propagation.

6.3.2 Carry-save adder cell with dual clocks

To ensure the majority gate evaluates only after the AND gate output is set, the clock of the majority and XOR gates was delayed by 1 ns. This time delay corresponds to the duration of the current consumption of the AND gate, and the output of the AND gate reaching the full supply voltage. However, on a fabricated microchip, simple time delay assumptions are not guaranteed to be valid. Switching properties of logic gates vary with position on wafer, production run, operating temperature, and supply voltage, which may make the delay assumptions invalid. Inserting fixed time delays would also make the circuit more vulnerable to power supply glitch attacks [Anderson and Kuhn, 1996], which is contrary to robust security design.

The simple delayed clock approach is sufficient to carry out simulations to investigate the causes of the smaller data-dependent variation in the power supply current, despite it not being a universal solution to the early propagation problem.

6.3.2.1 Results

Figure 6.12(a) shows the simulation result for the delayed-clock version of the carry-save adder cell for all 16 possible input data. The evaluation clock edges are shown in blue (AND clock)

and red (MAJ/XOR clock) colour. Figure 6.12(b) shows a close-up of the evaluation current of the carry-save adder cell. To distinguish the different currents, the curves are labelled with the input data.

The time difference between evaluation current peaks is reduced using two clock signals. The maximum time delay in the current peak of the AND gate is 50 ps. For the other two gates the maximum time shift in the current peaks for different input data are 130 ps.

The 50% to 50% delay time of the individual logic gates was found to have the following range: For the AND gate, the minimum delay is 424 ps and the maximum delay 490 ps, for the same input cases as previously used. The majority gate has a delay between 308 ps and 456 ps, for inputs of '0000' and '1110' (bit order x, y, sin, cin). The delay of the XOR gate is between 445 ps ('1100') and 469 ps ('1101'). The charge consumption variation is unchanged from single-clock simulation, with a difference of 5% between minimum (47.12 fC) and maximum (49.13 fC).

6.3.2.2 Discussion

While the spread in evaluation time and position of the current peak of the majority gate is significantly reduced by eliminating early propagation, some timing variation remains. To determine its cause, the timings and input data were compared to the schematics of the individual gates (Figures 6.3, 6.5, and 6.7).

The first current peak is the evaluation of the AND gate. For all cases, the discharge currents start together, but then diverge. The faster evaluation corresponds to inputs x and y being equal, independent of the evaluation result (0 or 1). Comparing these inputs to the schematic (Figure 6.5), it can be seen that for equal x and y there are always two transistor groups discharging the dynamic logic gate. For non-equal x and y, only one transistor group discharges the dynamic logic gate. The dynamic logic gate capacitance and the resistance to ground from the transistors represents an RC discharge circuit with varying resistances. The variation in switching time of the dynamic logic part of the gate also determines the time taken for the static output inverter to switch. The slower the input to the inverter is switched, the longer the PMOS and NMOS transistors are conducting at the same time and consuming power. The independence of current consumption from output of the gate confirms the balance of the capacitances of the layout.

The second discharge peak in Figure 6.12(b) represents the combined current of the XOR and the majority gates. As the XOR gate is composed of essentially three multiplexors connected in series (Figure 6.3), there are always the same number of transistors connecting the output of the dynamic logic part of the gate to ground. This results in equal currents for all input cases. The origin of the timing variations of the current peaks can be traced to the majority gate. Similar to the AND gate, the majority gate has different numbers of transistors discharging the dynamic logic gate. From the transistor schematic (Figure 6.7) three cases can be distinguished: if sin is equal to cin and to the output of the AND gate, three transistor groups discharge the output of the dynamic logic gate. If sin and cin are equal, but differ from the AND output, there are two transistor groups conducting to ground, and if sin is not equal to cin, there is only one group conducting. The timing of the current peaks corresponds to these three cases, explaining the remaining time variation.

The problem of early evaluation can be mitigated by synchronisation of inputs, if all inputs are present before the evaluation phase. This may be achieved using asynchronous circuit design techniques. Asynchronous circuits have previously been proposed as protection measures for security devices, in the context of countermeasures against glitch attacks [Fournier et al., 2003; Moore et al., 2002, October 2003], or to prevent synchronisation of power consumption traces [Plana et al., 2003, 2002; Yu et al., 2003]. Adding synchronisation circuitry must not com-

promise security by revealing timing, or by having early propagation in the handshaking circuit. Alternatively, modifications to the gate design [Kulikowski et al., 2006b] may be suitable to some gate designs to alleviate problems from early propagation.

The cause of the smaller timing variation is not easily removed, if a logic gate is also required to preserve capacitive balance. For example, balanced capacitances of the AND gate were achieved by doubling the transistor groups on the 'true' side of the gate to match the two parallel transistors of the 'complement' side. As there is only one input case where the output of the AND gate is high, but three cases where the output is low (or the 'complement' is high), there is an inherent imbalance. Balancing the discharge path resistance may result in imbalanced capacitances. For example, the single combination of inputs to yield a '1' output calls for a single group of two series transistors to discharge the output node. On the 'complement' side, this requires three groups of two series transistors ($x\bar{y}$, $\bar{x}y$, $\bar{x}\bar{y}$). However, in this case the parasitic capacitances are not matched anymore. The insertion of dummy transistors to equalise capacitances and compensate for the imbalance may ensure balance in output node size. In addition, the number of discharged centre nodes between the series transistors must also be equalised, as well as effects from charging the gate capacitors of the transistors.

In the pre-charge phase, the static inverters are switched low, and all discharged capacitances of a logic gate are (p)re-charged. Inspecting the power supply current at the falling edge of the clock, it can be seen that the current for all input cases coincides. Therefore it can be concluded that the relevant capacitances of the gates are well-balanced. In summary, the timing variation in the evaluation phase dominates the overall data-dependency of the power consumption, and therefore needs to be addressed first.

6.3.3 Conclusions

Two main sources of imbalance have been identified in the simulated carry-save adder cells. The AND and the majority gates are vulnerable to early propagation, which is identified as the largest source of data-dependent power consumption. If the inputs to the gates arrive at different times, vulnerable logic gates set their outputs as soon as sufficient information is available to determine the result. If one input of the AND gate is low, then the output immediately switches to low. Similarly, in the majority gate early propagation takes place if the first two inputs are equal. Early propagation is manifested in the current traces by a temporal shift of the current peak with input data. To counter early propagation, all inputs must be set before the clock is asserted. This may be achieved using asynchronous circuit design techniques. Alternatively, modifications to the logic gates may be made to prevent early propagation, if it is possible without breaking the balance of the gate.

The second source of imbalance is the number of transistors discharging the gate capacitance. If the resistance of the discharge path varies, the RC time constant which determines the evaluation time is also affected. This effect results in a data-dependent time shift of the current peak, though the time shift is smaller than under early propagation.

As the total consumed charge is more or less constant for all input data, it can be concluded that the capacitive imbalances play a negligible role in simulation compared to the other two sources of imbalance.

6.4 Comparison of simulation and measurement

As the simulations showed that the overall charge consumed by a carry-save adder cell is independent of data, it is possible that the capacitances of the packaging and circuit board are sufficient to smooth out the temporal shift in current peaks. In order to evaluate the filtering

effect and noise addition of the physical measurement setup, simulation results are compared to power consumption measurements on the full chip.

The full array multiplier was chosen for this comparison, as its inputs and outputs are connected via shift registers. The input- and output shift registers were originally added to the multiplier to reduce the pin count required on the chip. Conveniently, they allow all data to be pre-set, and only the multiplier clock to be asserted to trigger an evaluation. If the shift registers on the output are not clocked, no output pad drivers are active, minimising any induced noise.

6.4.1 Corner cases

From the evaluation of the carry-save adder cell it is known that early propagation may occur in the majority gate. In the multiplier, the timing relationship between the inputs to the adder cells is different compared to the previous simulation, in which all inputs were asserted at the same time. An approximate timing diagram of the array multiplier is shown in Figure 6.13. The schematic is an 8-bit adaptation of Figure 6.1, except that the x and y inputs are not shown for clarity. The logic gates are marked as '&' for the internal AND gate, 'M' for the majority gate (cout), and 'X' for the XOR gate (sout). The time delays, marked in red, are abstracted to units of gate delay. This highlights the timing variation due to early propagation to be traced independently of the physical implementation and its specific gate delays. Cells that have outputs that may have early propagation are marked in blue.

The first row and column of the array multiplier consist only of AND gates. The majority gates of the second row are therefore only connected to AND gates (first row, and the internal AND gates of the cell), and a permanent ground connection. As the AND gates have the same delay, the majority gates of the second row will all evaluate at the same time, without early propagation. In the third row, the cell in the second column has two AND gate inputs, and the third input (cin) coming from the majority gate output of the second row. Therefore, the last input comes one gate delay after both AND gates have set their outputs, making the majority gate susceptible to early propagation. Given that the evaluation time for both the majority and XOR gates (without early propagation) was simulated to be of the order of 900 ps, it is assumed that both logic gates evaluate simultaneously. As the remaining gates in the third row are supplied by the majority and XOR outputs of the second row, they do not evaluate early.

In the second column of row four, the majority gate input from the previous row (cin) is the latest arriving input to the cell, as before. As its timing may vary due to early propagation, the XOR gate of this cell may now have data-dependent timing variation, passing on any early propagation from the majority gate of the previous row. As before, the remaining carry-save adder cells in the fourth row do not evaluate early if the two inputs sin and cin arrive simultaneously.

Timing variation of the XOR gate carries early propagation to the third column, as the sin and cin inputs may no longer be synchronised, and the AND gate result is already present after one gate delay. The majority gate in the third column of the fifth row may evaluate early, if the timing of the sin input to the cell is early, and its state is equal to the result of the AND gate. The timing of the sin input depends on the early evaluation status of the majority gate one column to the left and two rows up, as its timing is passed on through the XOR gate supplying the sin input.

The timing range of the remaining cells is computed analogously. The sum and carry bits propagate through the array multiplier row by row from top to bottom. If the XOR and majority gates have the same evaluation time, then the timing of the inputs to the cells in the next row will be synchronous, and early propagation is not a problem. Early propagation, however, feeds into the array multiplier from the first column, as it does not rely on data from the previous row and all AND gates evaluate immediately.

A pre-requisite for early propagation of the majority gate in columns after the second is early propagation in the previous column. Therefore, cases exhibiting early propagation can be identified. The x and y inputs feeding the AND gates in the first and second column are the most significant bits of y, y_7 and y_6, as well as all bits of x (following the naming convention of Figure 6.1). In order for a majority gate of a cell in the n^{th} row and second column to have the same inputs from both AND gates, $(x_n \quad y_6)$ must be equal to $(x_{n-1} \quad y_7)$. This is the case if either $y_7 = y_6 = 0$, or if $x = 0$. Alternatively, if the inputs are not zero, all bits of x should be equal to one, and $y_7 = y_6 = 1$. Early propagation properties in the third column depend on the status of early propagation of the previous column, as well as the appropriate sin and AND bits being equal.

For early propagation to occur in the carry propagate adder (the last row of the multiplier), both inputs from the previous row (sin and cin) must agree. The large number of 65536 possible input data make it infeasible to simulate and measure all possible cases, therefore analysis is limited to corner cases. Given that the largest sources of data dependent power consumption in the carry-propagate adder cell is from early propagation, corner cases can be defined based on the early propagation properties. A Matlab script was written to simulate the data flow through the multiplier and identify the corner cases where early propagation occurs in the multiplier and the carry-propagate adder.

The four evaluated cases are early propagation in the array and the carry-propagate adder, no early propagation in either, and early propagation in one part, but not the other. The 'fastest' case is an input of $0 \quad 0$, which has early propagation through both the array and the adder, as well as the shortest evaluation time in the AND gates. A slightly slower case with early propagation in the entire multiplier is $0 \quad 1111\ 1111$ (or vice versa), where the AND gates take slightly longer to evaluate. For early propagation to occur in the array, y_7 must be equal to y_6. The input of $1111\ 1111 \quad 0111\ 1111$ (in order of $x \quad y$) therefore does not have early propagation in the array, however it does have early propagation through the entire carry-propagate adder. Similarly, the case $1000\ 0000 \quad 1111\ 1111$ has early propagation in all multiplier rows, as the last AND gate of the first column feeds directly into the carry-propagate adder. However, there is no early propagation in the carry-propagate adder. The commutation of the previous value, $1111\ 1111 \quad 1000\ 0000$, on the other hand does not have early propagation in the multiplier nor in the carry-propagate adder.

6.4.2 Simulation setup

A similar SPICE simulation setup to the previous simulations was used. The transistor model used in the simulation was the same MOSIS-supplied model as before. While extracted circuit generated with the Calibre tool is very detailed, it comprised a large number of elements, requiring a correspondingly long simulation time. In the simulation of the individual carry-save adder cells this is tolerable, as the circuit itself is comparatively small. In the array multiplier this caused a simulation time of a day per input datum. As it was established in the simulation of the carry-save adder cells that the largest source of imbalance is early propagation rather than the capacitances of the circuit, a less complex circuit model was used. A reduced circuit model was generated with Electric, the very large scale integrated circuit (VLSI) design program that was used to design the multiplier layout. This model comprises all wire resistances but only the largest capacitances.

6.4.3 Simulation results of corner cases

The timings of the outputs of the array multiplier are shown in Figure 6.14. The curves are annotated with the output number, following the same naming convention (Figure 6.1) as before. As

the logic gates are dual-rail, one of the two rails is set high, depending on the evaluation result. The plotted graphs in each of the figures are the appropriate output rails switching high.

In Figure 6.14(a), the outputs are shown for an all-zero input. As the outputs of all AND gates are zero, all majority and XOR gates also evaluate to zero, causing early propagation in the entire structure. Comparing the timing of the outputs to the related case of 1111 1111 0, shown in Figure 6.14(b), small differences in output timing are visible. These can be seen by comparing he crossings of the output traces with the dotted grid lines. Figure 6.14(c) shows the output timing for the inputs of 1111 1111 0111 1111, in which the carry-propagate adder evaluates early, while the multiplier array does not. The other cases shown are inputs of 1111 1111 1000 0000 (Figure 6.14(d)), where there is no early propagation, and inputs of 1000 0000 1111 1111 (Figure 6.14(e)), where only the array shows early propagation.

The simulated power supply current for all five evaluated cases of the multiplier is shown in Figure 6.16 (moved to page 153 for comparison with the measured current curves). A full clock cycle is shown in Figure 6.16(a), and a magnification of the current in the evaluation phase is shown in Figure 6.16(b).

The total charge consumed during the evaluation phase lies in the range of 2.56 pC to 2.59 pC, a difference of 1% between maximum and minimum value. Integrating the current over a full clock cycle, the charge consumed varies between 6.36 pC to 6.48 pC (2% between minimum and maximum).

6.4.4 Measurement setup

The multiplier chip was fabricated using the AMIS C5 0.5 µm (5 V) CMOS technology with a single poly-silicon layer and three metal layers. The chip was mounted in a standard taped-lid dual in-line (DIL) package, which allowed photos of the die to be taken using a microscope (Figure 6.15). The regular layout structure is clearly visible, though the tracks in the top metal layer subjectively seem wider (smaller gaps) than in the layout design (Figure 6.10).

There are two of power supply pads on the chip, visible in the left column of contact pads in Figure 6.15. The top two pads (Figure 6.15) supply the input and output pad drivers of the chip, the shift registers, and other auxiliary circuits. The second supply is connected exclusively to the multiplier, so that the supply current can be measured more easily (bottom two pads in Figure 6.15). The ground connections are also separate, however there is some coupling between the two through the substrate.

The driver and measurement circuit were assembled on a standard strip board. To measure the current, 10 Ω surface mounted resistors were soldered in series with the power supply. The current was measured using a low-capacitance (0.1 pF) differential probe connected to a Tektronix TDS7254B oscilloscope. The resolution of the oscilloscope was set to the maximum sampling rate of 10 GHz or 100 ps per data point. On the power supply side of the differential probe, a 600 nF smoothing capacitor was added to assist in removing unwanted interference. For each input case, 100 curves were averaged to remove random noise.

A field programmable gate array (FPGA) chip was used to generate the appropriate waveforms to drive the array multiplier. The 16 input bits were clocked into the shift register, before the array multiplier clock was run for an appropriate number of repetitions. Finally, the result was clocked out of the output shift register, to verify correctness. To allow the ringing to settle between pre-charge and evaluation, a clock frequency of 4 MHz was chosen.

6.4.5 Measurement Results

The measured current through the power supply pins of the multiplier are shown in Figure 6.17. As before, the supply current of an entire clock cycle is plotted in the top figure (Figure 6.17(a)),

and a magnification of the evaluation phase below it (Figure 6.17(b)).

The charge consumed during the evaluate phase ranges from 35.1 pC to 36.5 pC (4% difference), while the charge consumed during a full clock cycle ranges from 94.1 pC to 99.5 pC (5% difference). These figures were determined by integration of the instantaneous current, thus are distorted by the ringing, noise, and cross-talk between the two power supply rails.

6.4.6 Discussion

The relative timings of the outputs largely follow theoretical prediction. In the cases with early evaluation through the carry-propagate adder, the outputs p_{10} to p_{15} are set in opposite order. For the two cases that do not have early propagation in the carry-save adder, the timing of the outputs is constant. While these two cases cannot be distinguished by the timing of the outputs, differences are visible in the power consumption traces (cyan and black curves in Figures 6.16 and 6.17).

A more detailed analysis of the output timing without early propagation in the array, but with early propagation in the carry-propagate adder (1111 1111 0111 1111, Figure 6.14(c)) reveals a more subtle timing effect which may also be caused by early propagation. According to the simple timing model presented in Figure 6.13, the outputs p_{10} to p_{15} should evaluate later than shown in the timing simulation, at around the same time as p_8. Only the first column of AND gates in the multiplier evaluates to zero, all other AND gates evaluate to one. The majority gates of the second column thus evaluate to zero, while the majority gates of the remaining six columns all evaluate to one. Dropping assumption of equal evaluation times of the majority and XOR gates, a slightly faster evaluation time of the majority gates (two transistors in series) compared to the XOR gates (three transistors in series, of the same dimension as the transistors in the majority gates) is assumed instead. The majority gates in the array may evaluate before the cin input is present, which is termed secondary early propagation. After the majority gate of one row has completed evaluation, two of the three inputs to the majority gates in the following row are equal, triggering early propagation. In the second column there is a zero input from the AND gates of the first column. In the remaining columns the internal AND gates evaluate to one, permitting the one result from the majority gates to propagate through without waiting for the sin input. The different order of output bits for the input 1111 1111 0111 1111 therefore is the result of different numbers of XOR gate evaluations required to compute output bits. A case with early propagation in the carry-propagate adder, but without this secondary early evaluation in the multiplier array was confirmed not to exist using the Matlab simulation script.

The simulated supply current of the array multiplier (Figure 6.16) correlates well with the timing of the outputs and the early evaluation status of the two parts. The large initial negative peak corresponds to the simultaneous evaluation of the AND gates. The slope after the first peak corresponds to the evaluation in the array. In the cases without early evaluation in the array (red and cyan), the current curve is more spread out compared to the cases with early evaluation. The two cases 0 0 and 1111 1111 0 with the slight timing difference in the AND gates also show a small time shift in the evaluation current. The evaluation of the carry-propagate adder starts around 8 ns for the cases without early propagation in the carry-propagate adder and takes around 4 ns.

While the measured current curve has significant oscillation (ringing) and noise, data-dependent differences are easily distinguishable above the noise, despite the crude DIL packaging and strip board measurement setup. Comparison of the data-dependent features in the simulated (Figure 6.16) and measured (Figure 6.17) supply current curves reveal common features, despite the filtering effect and noise of the measurement setup. Ringing oscillations are prominent in the measured curve. After the large negative current spike from the AND gate evaluation, the two cases without early propagation (red and cyan) draw less current than the other inputs. This

difference in currents between cases with and without early propagation also lasts around 7 ns in the measurement. Afterwards, starting at around 30 ns in the measurement curve, the two cases without early propagation in the carry-propagate adder draw more current compared to the cases with early propagation, and do this for approximately 4 ns, as in the simulation. The small time difference between the two cases with a zero output are not distinguishable in the measurement.

Comparing the charge consumption of the simulation and measurement of the array multiplier, significant differences between simulation and measurement are visible. The measured charge consumption is 15 times more than the simulated value, which may be due to the significant ringing in the measurement and the cross-talk between supply rails. In a separate experiment it was found that gate switching activity resulted in measurable oscillations on supply rails that were not feeding those logic gates. This cross-talk contributes to the measured power consumption, but is not present in the simulation. Despite the simulated difference in overall charge consumption of 2 % and a measured mismatch of 5 %, different inputs can clearly be distinguished.

6.4.7 Summary

Comparing simulation and measurement has shown that despite the crude packaging and measurement setup, the added capacitances on the power supply rails are not sufficient to conceal the data-dependent time shift in the power consumption traces from early propagation. The shorter time shift due to the mismatch in discharge resistance of the AND gate was not distinguishable in the power consumption traces.

6.5 Conclusions

Examination of balanced dual-rail domino logic circuits has shown that nominally balanced logic designs may still be vulnerable to power analysis attacks. It is not sufficient to ensure constant charge consumption per clock cycle. Early propagation may lead to timing differences, and thus create imbalance. To counter early propagation, modifications may be made to logic gates [Kulikowski et al., 2006b]. Alternatively, it must be ensured that all inputs are present before triggering the evaluation phase of the logic gate. The second source of data-dependent temporal shift in the instantaneous supply current is the resistance to ground discharging the dynamic logic gate.

A comparison of simulation and measured power consumption of the secure multiplier design revealed that the simulated mismatch is also present in the measurement. The cases could be distinguished despite nominally balanced charge consumption, smoothing effects of added capacitances in the packaging and measurement setup, and the addition of noise in the measured traces.

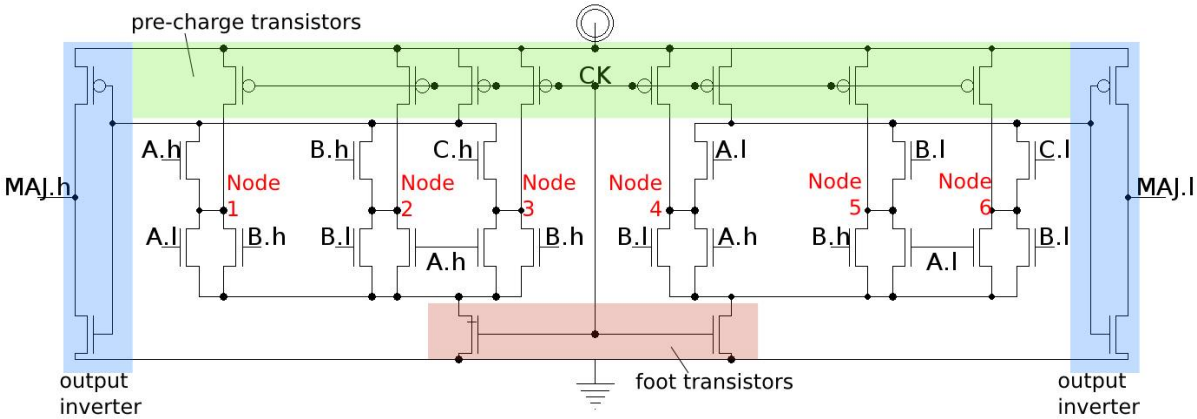

Figure 6.7: Schematic of majority gate

input (ABC)	node 1	node 2	node 3	node 4	node 5	node 6
000	✓	✓	–	✓	✓	✓
001	✓	✓	–	✓	✓	✓
010	✓	–	✓		✓	✓
011	✓		✓	–	✓	✓
100	–	✓	✓	✓		✓
101		✓	✓	✓	–	✓
110	✓	✓	✓	✓	✓	–
111	✓	✓	✓	✓	✓	–

Table 6.1: Discharged nodes for majority gate: direct (✓), no (–), indirect ( ) discharge

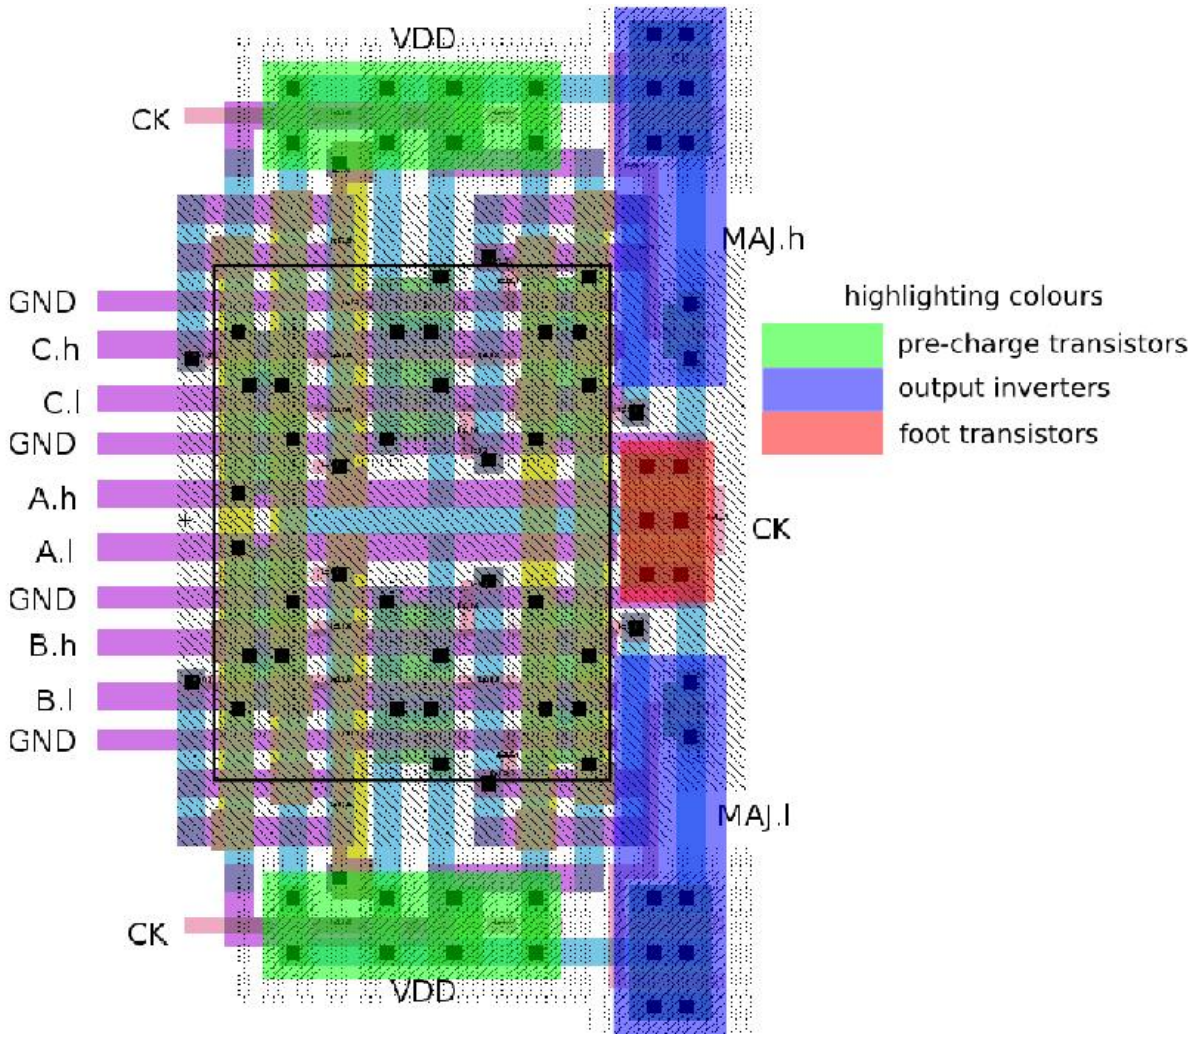

Figure 6.8: Layout of majority gate

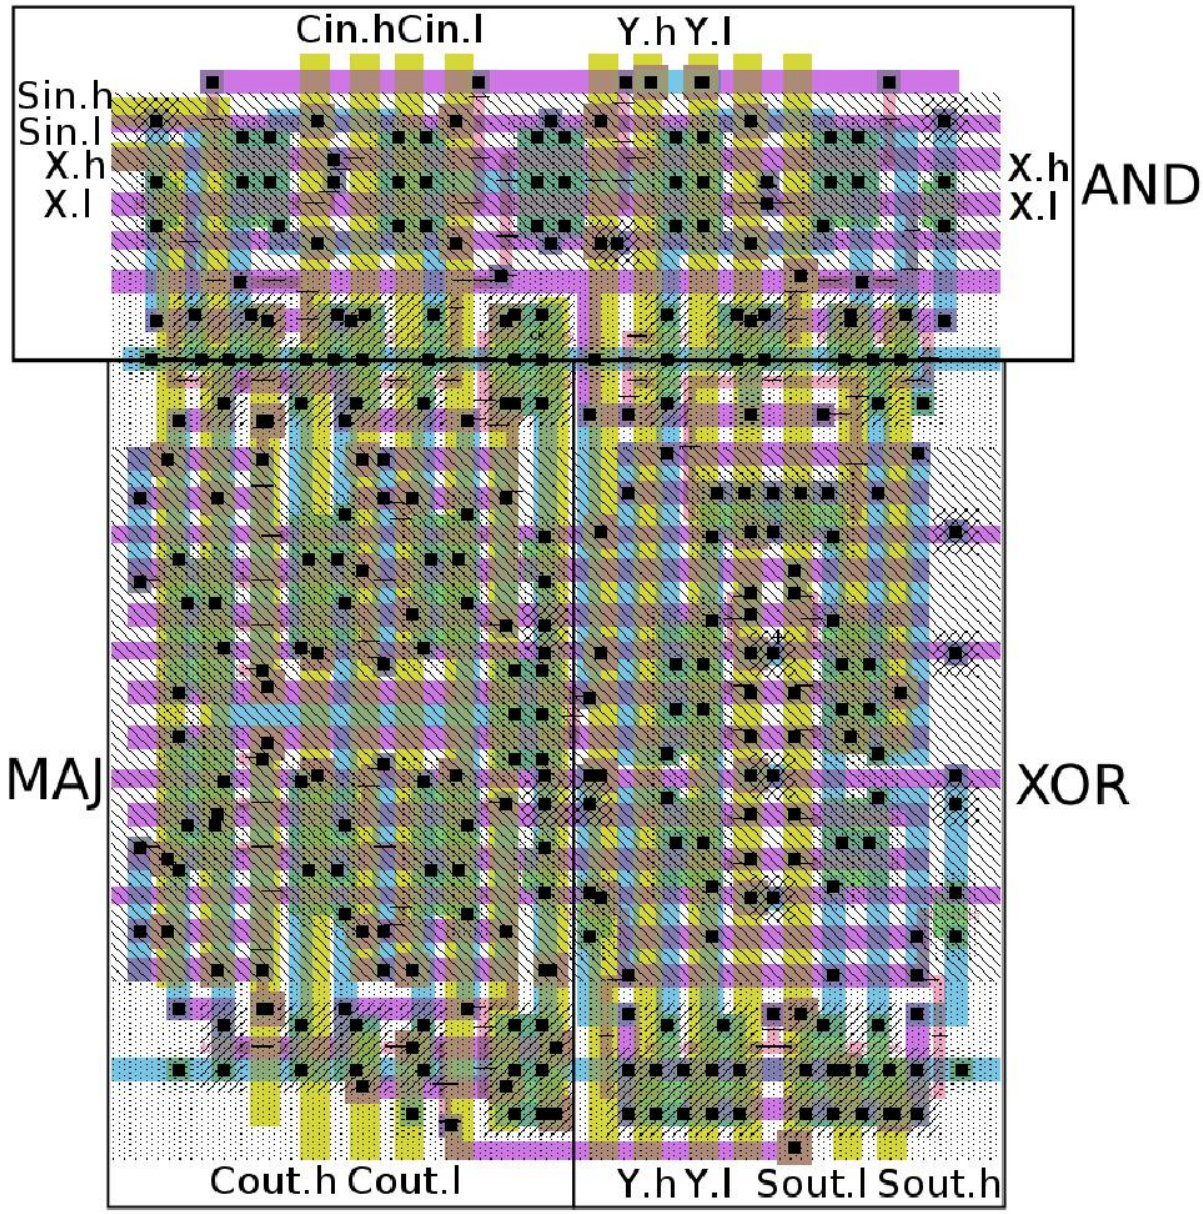

Figure 6.9: Layout of carry-save adder cell

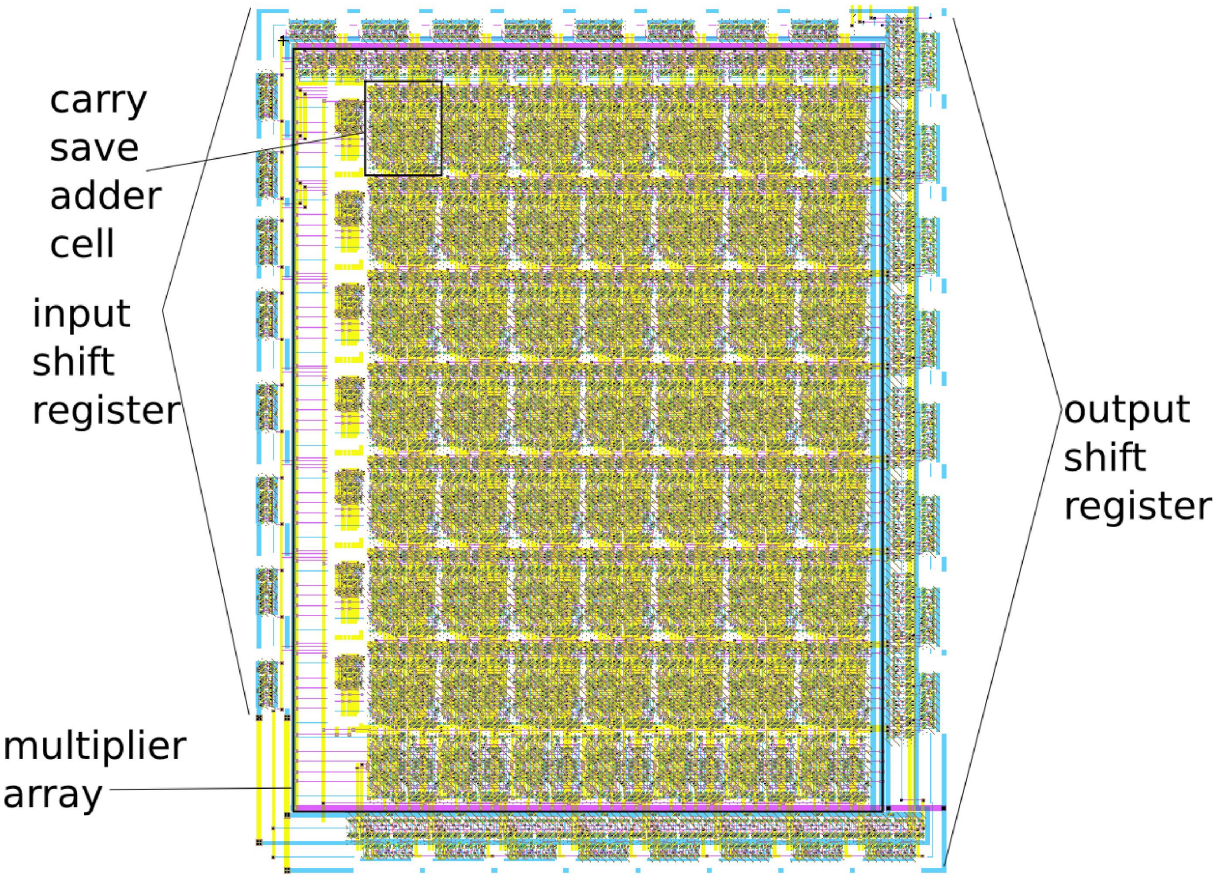

Figure 6.10: Layout of 8-bit multiplier

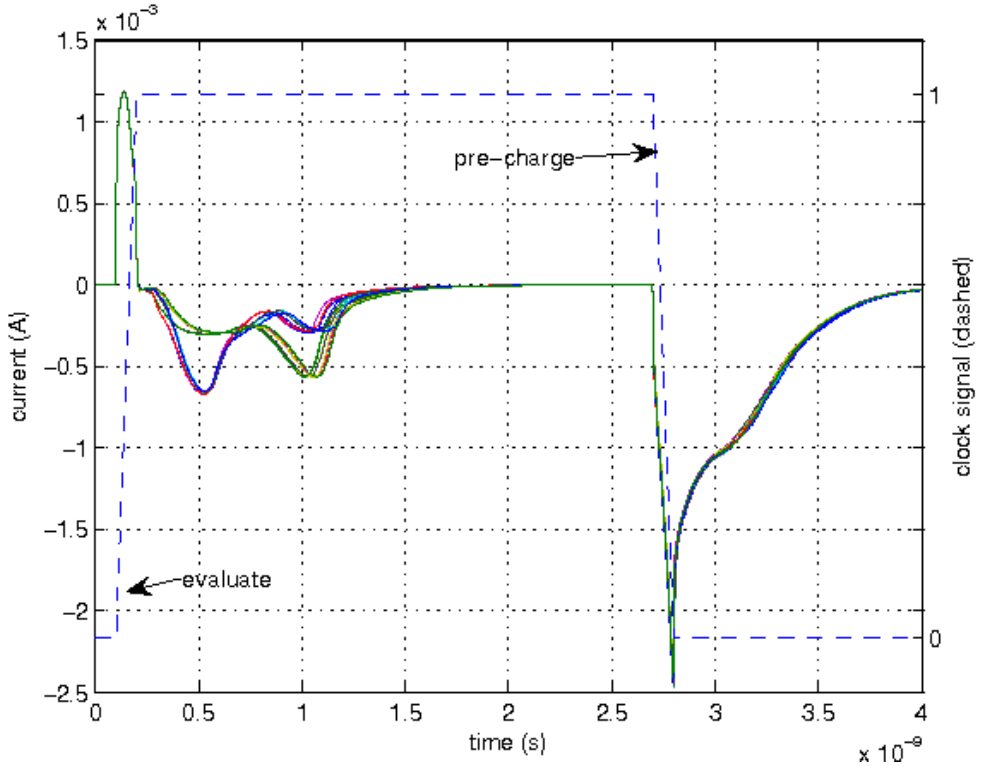

(a) Full cycle

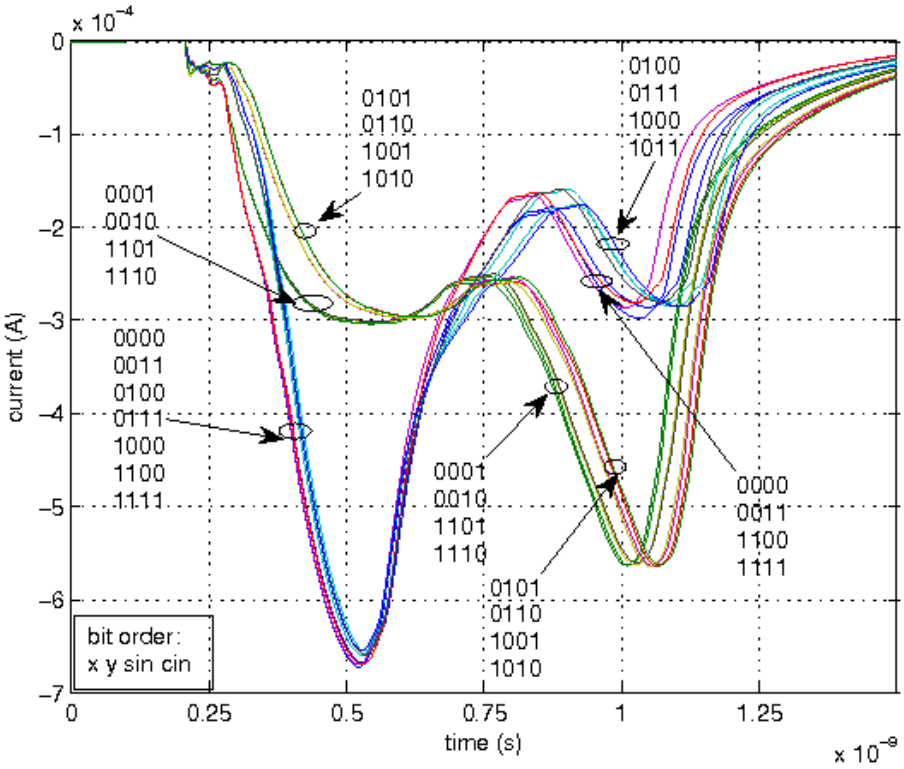

(b) Detail of evaluation clock edge

Figure 6.11: Input current to carry-save adder cell for all input cases

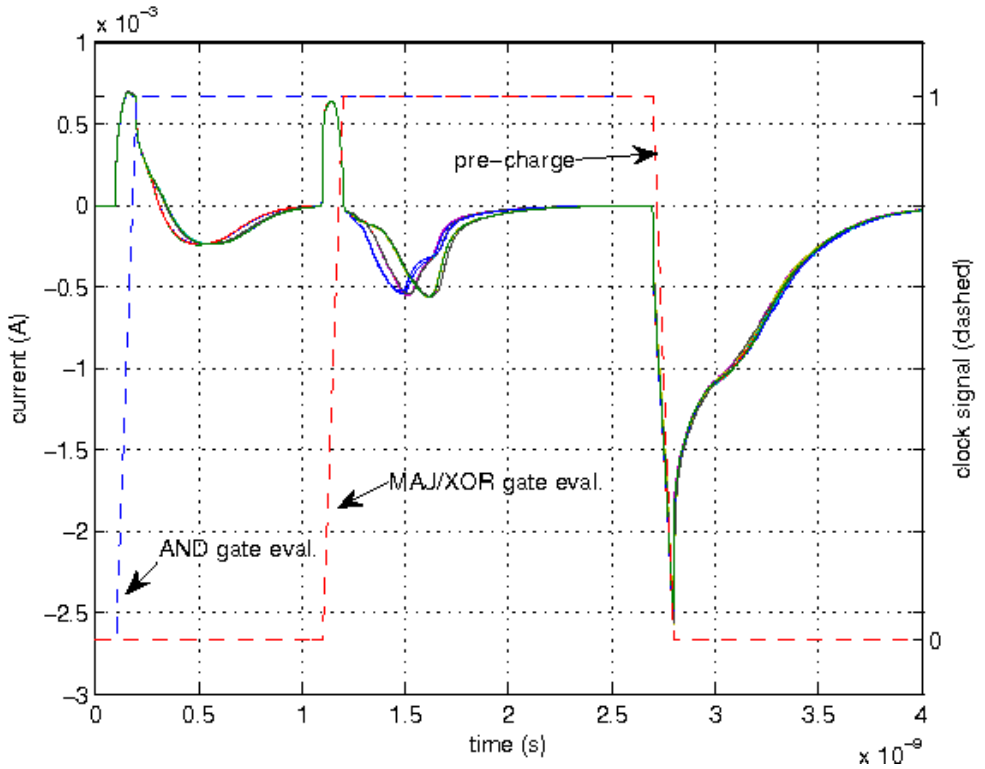

(a) Full cycle

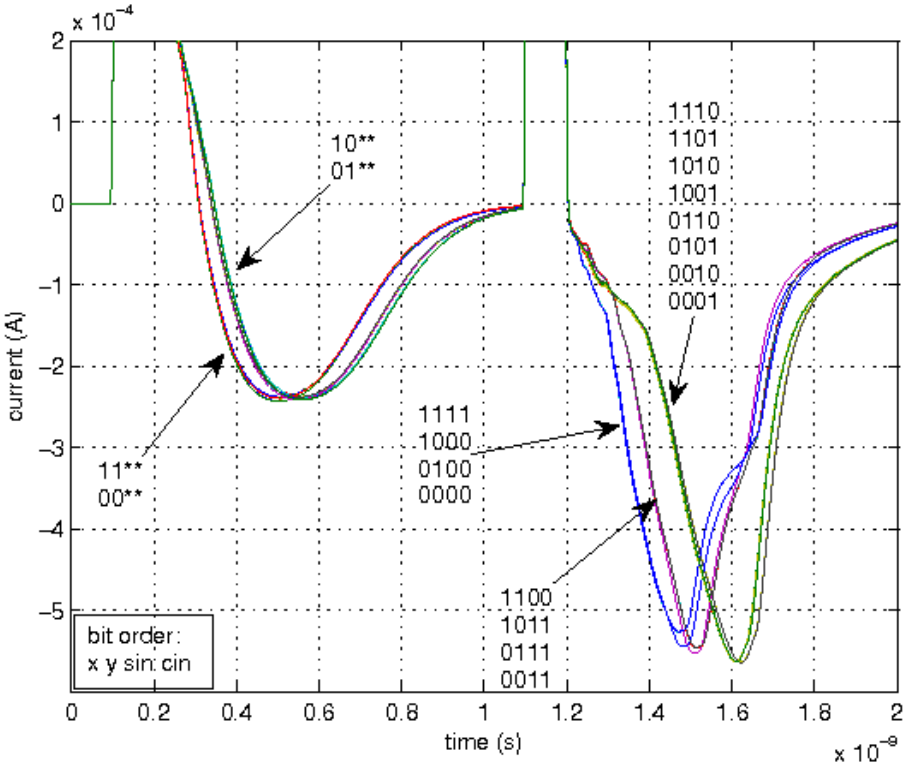

(b) Detail of evaluation clock edge

Figure 6.12: Input current to carry-save adder cell with delayed clock for all input cases

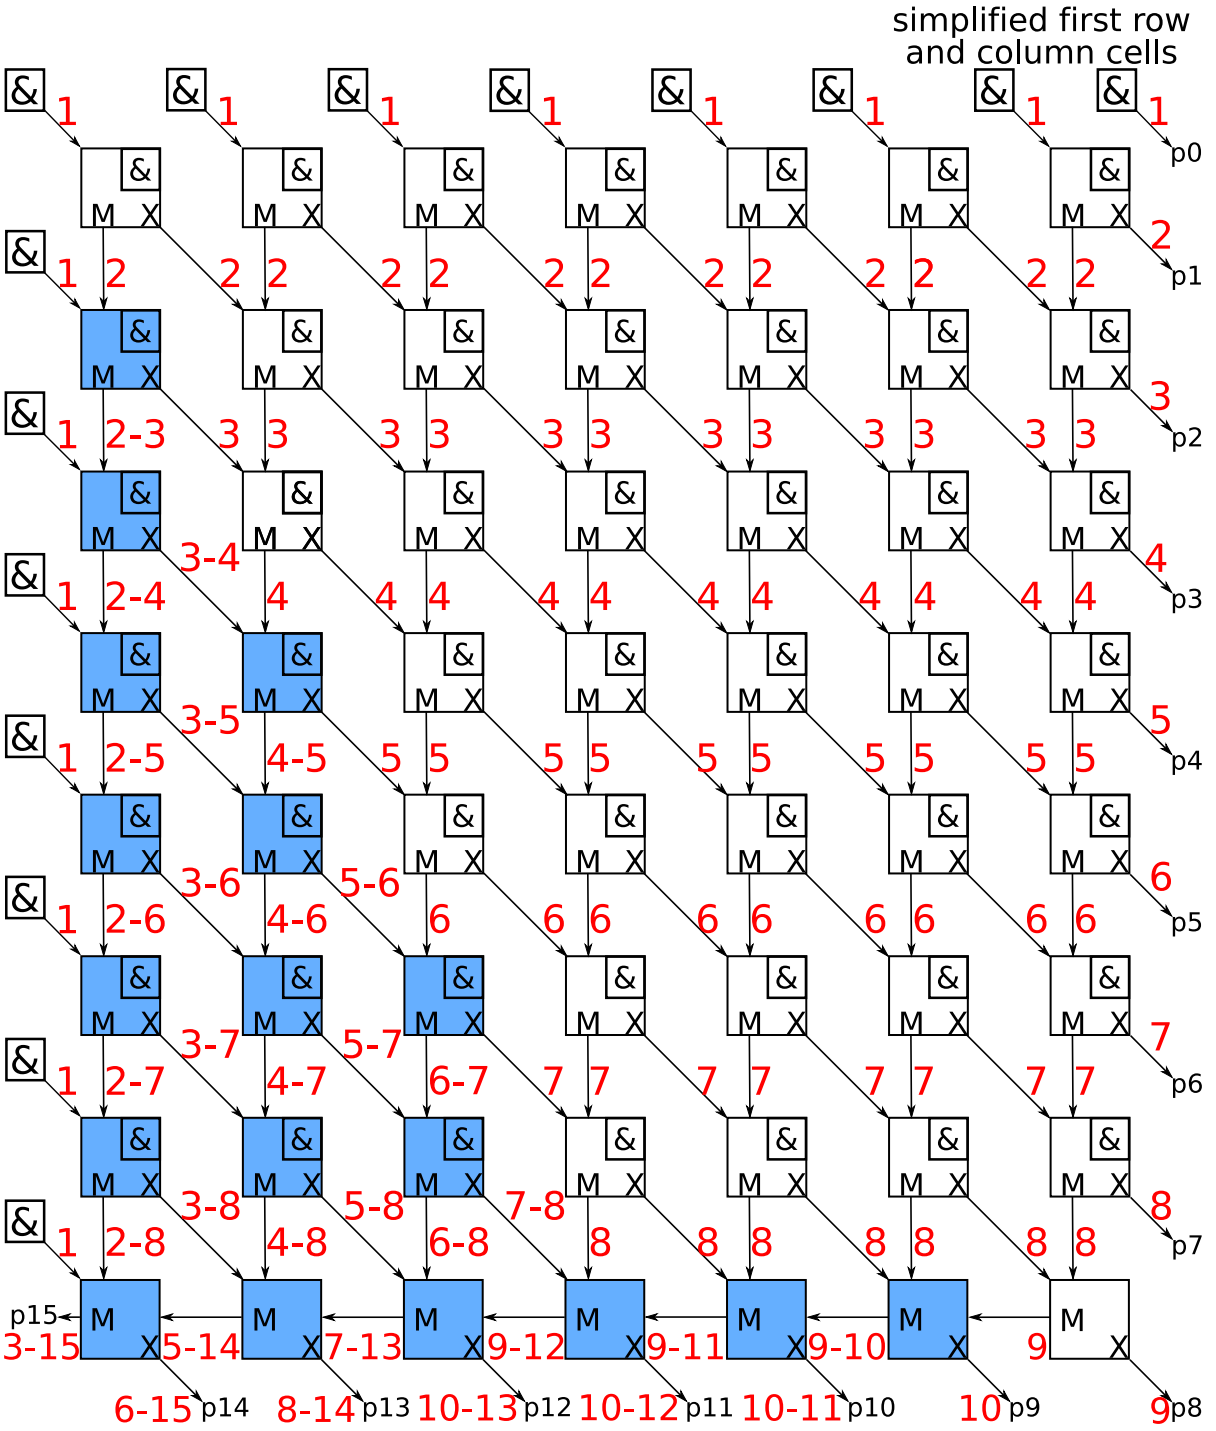

Figure 6.13: Timing diagram for 8-bit array multiplier

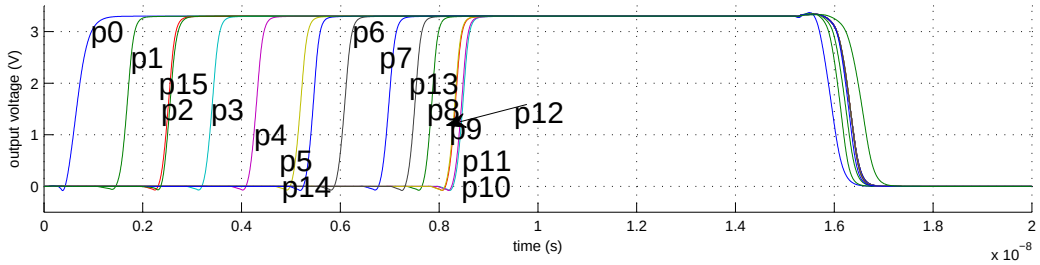

(a) 0×0 – all early

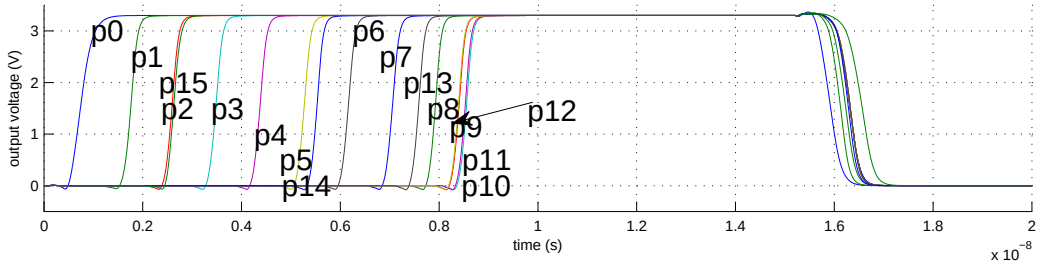

(b) 1111 1111 $\times$ 0– all early, slightly slower

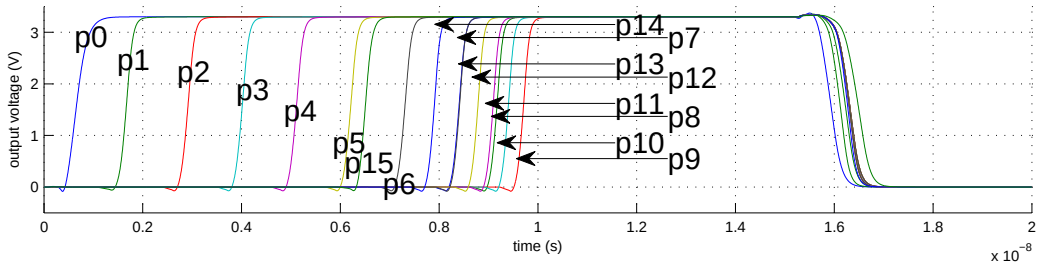

(c) 1111 1111 $\times$ 0111 1111 – carry-prop. adder early

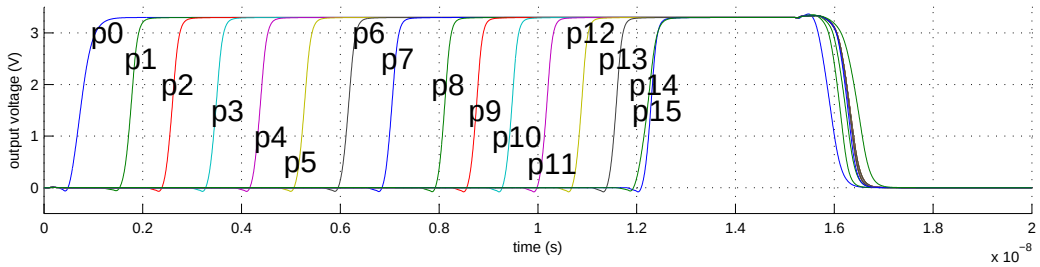

(d) 1111 1111 $\times$ 1000 0000 – no early propagation

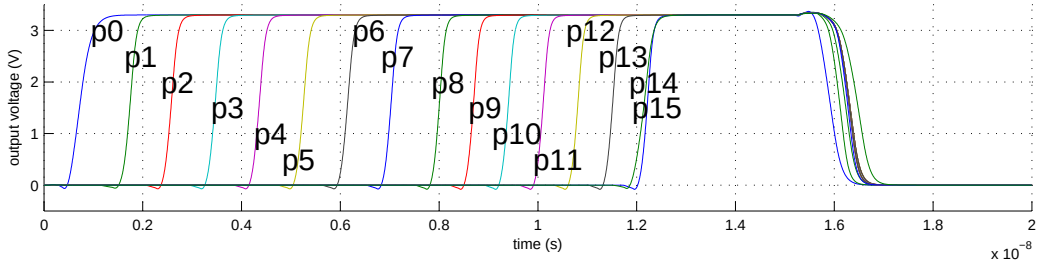

(e) 1000 0000 $\times$ 1111 1111 – array early

Figure 6.14: Timing of output data

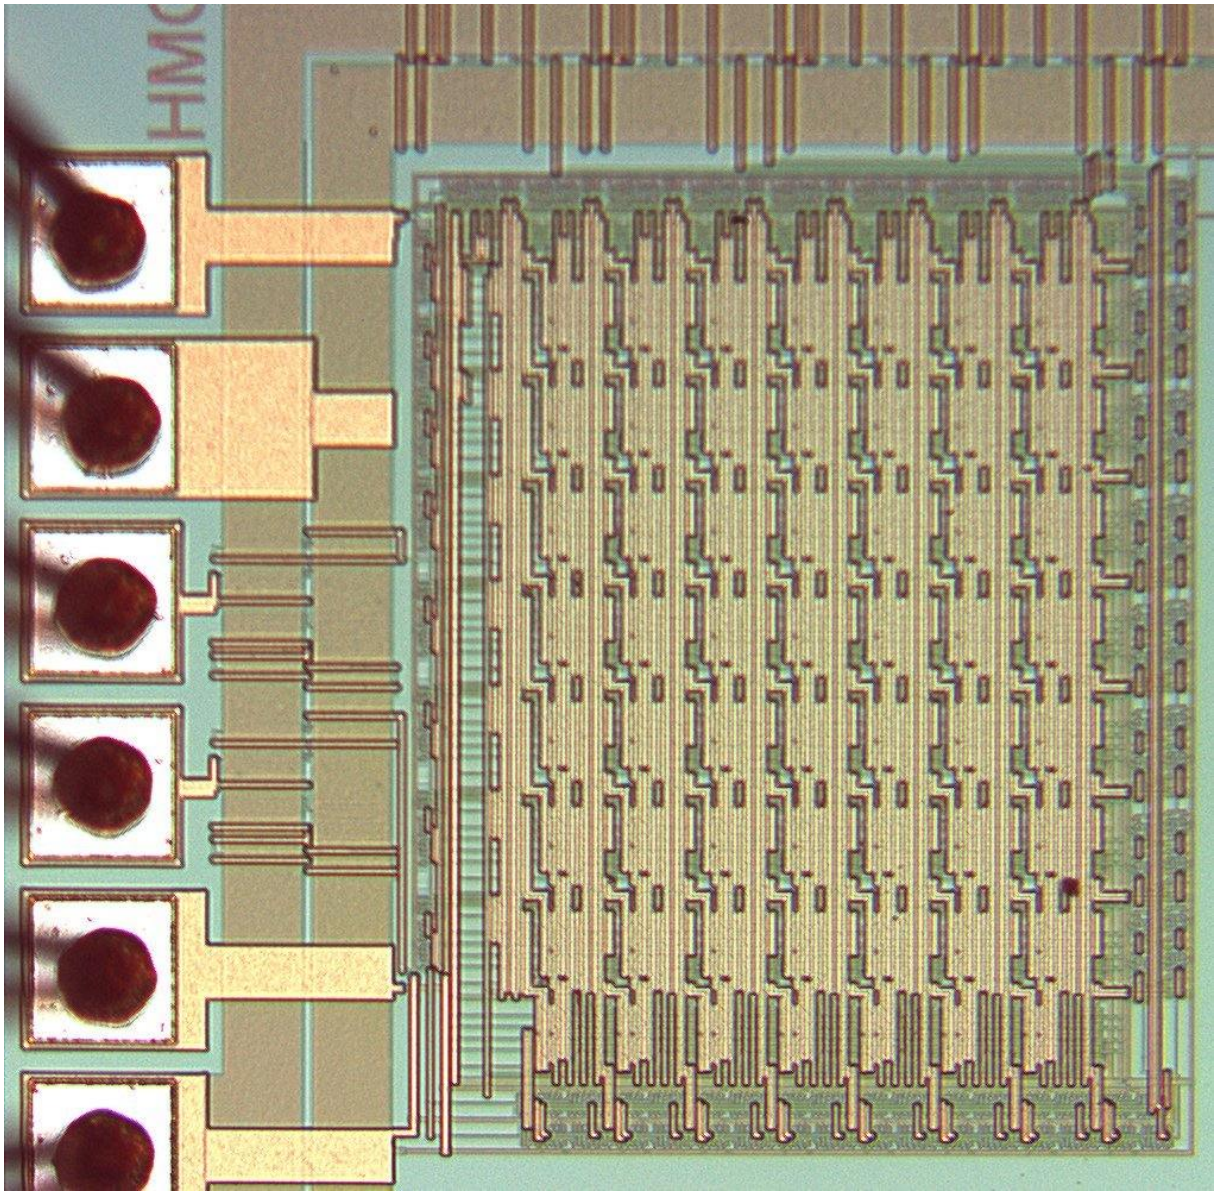

Figure 6.15: Microscope image of secure array multiplier

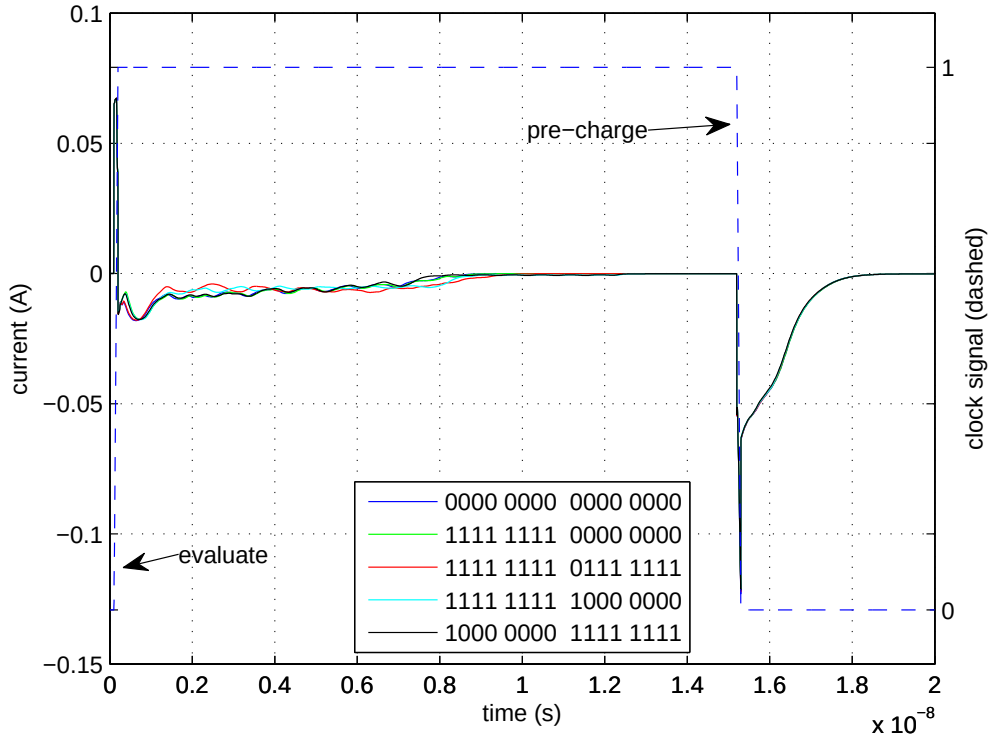

(a) Full cycle

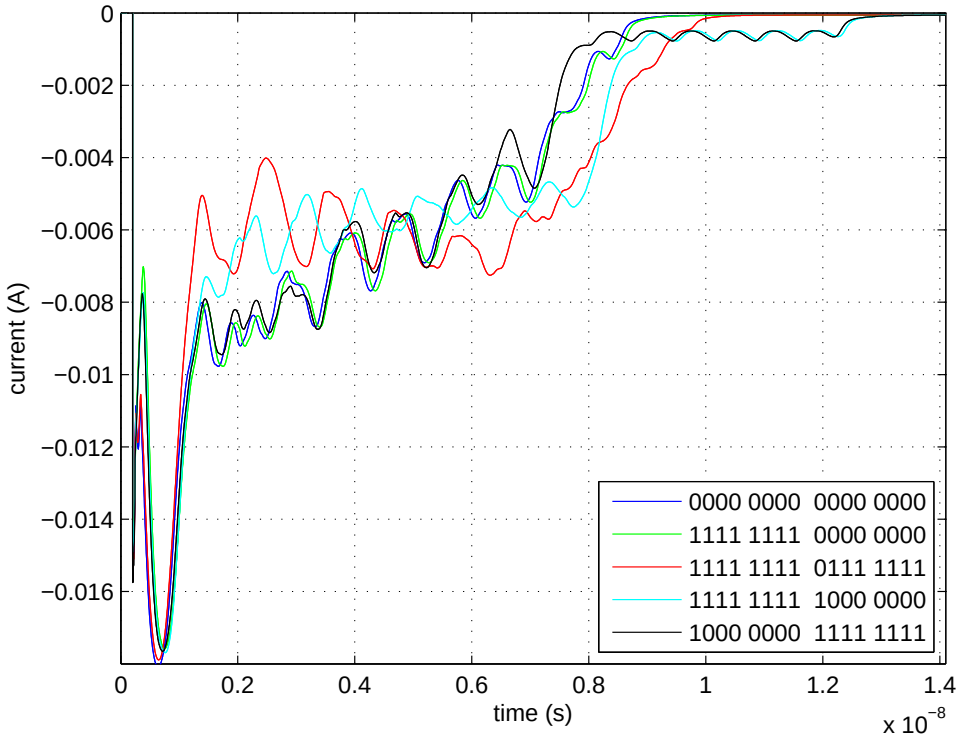

(b) Detail of evaluation clock edge

Figure 6.16: Simulated input current to full array multiplier

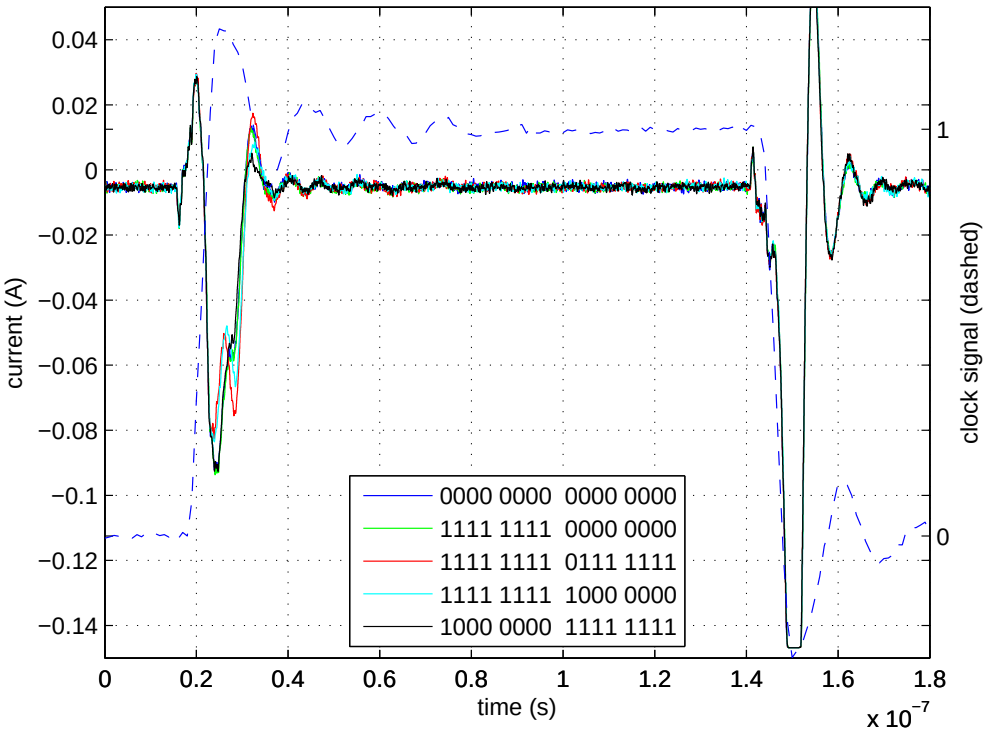

(a) Full cycle

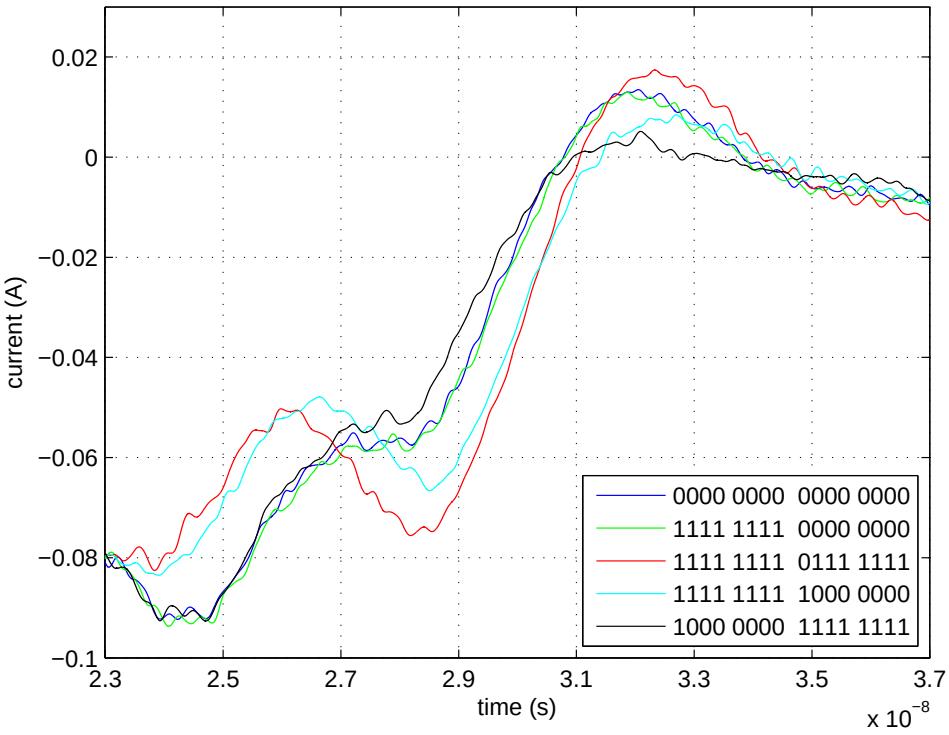

(b) Detail of evaluation clock edge

Figure 6.17: Measured input current to full array multiplier

Chapter Seven

7 Conclusions

Contributions were made to the protection of microelectronic security devices against the unauthorised extraction of secret information. Two patent applications were filed to cover the work on organic tamper protection grids. The expansion of the patent application from the UK to the US is evidence of Epson's commercial interest in this work. With more development effort to refine this proof of concept into a full protection system, organic protection grids will be a competitive application to Epson's inkjet-printed organic electronics technology. In contrast to metal protection grids, organic electronics allow the detection of depackaging by the security device and thus for it to take active countermeasures.

Organic electronics were proposed to be used in tamper protection grids of security devices. Practical issues were resolved with respect to fabrication, and the sensitivity of the system to depackaging was verified. To show directions for future work, practical implementation aspects were evaluted. To guard against non-invasive attacks, balanced logic gates were developed. An array multiplier was laid out and fabricated as a test structure for these secure logic gates. The design was simulated and measurements of the fabricated device were carried out to analyse power consumption balance and the origin of the data-dependency of the power consumption.

The balanced dual-rail domino logic gates were found to have measurable data-dependent power consumption. Simulations revealed that the origin of the data-dependency was not an imbalance in discharged capacitances, but were due to timing variations in the evaluation of the gates. Removing these sources of imbalance are likely to conflict with a balanced capacitance layout, or add significant size overhead.

7.1 Technical summary and future work

It was shown that simple organic protection grids can be fabricated using PEDOT:PSS as a conductive material, and that they are sensitive to standard depackaging methods. PEDOT is suitable for inkjet fabrication, which allows the grid pattern to be varied dynamically. This allows the grid properties to be made unique for each security device. Printing on standard microchips is possible despite the surface of microchips being hydrophobic, if a buffer layer is inserted, or if the surface is treated with surfactant. The unstable contacts between PEDOT and aluminium can be stabilised with silver paint. While the PEDOT is reported as the most stable organic conductor, ageing was still observed. Encapsulation of PEDOT reduces ageing significantly, however in the setup used, ageing was not completely eliminated. Further investigation of en-

capsulation materials must be carried out to ensure sufficient lifetime of the grid. The change of PEDOT resistance with temperature is sensitive to the humidity contents, thus varying with the weather. Encapsulation reduces the variability, however, after more than a day, the Araldite epoxy used for the tests degrades at temperatures below 100 °C, damaging the PEDOT grid.

Sensitivity to depackaging can be achieved by appropriate choice of packaging layers. PEDOT is sensitive to nitric acid, which is the most commonly used chemical to remove packaging material. If the protection grid is not sensitive to a particular depackaging solvent, sensitisation can be achieved by printing the protection grid on a base layer made from the same material as the remainder of the packaging. That way, the lines will be deformed and broken physically upon contact with the solvent. Mechanical depackaging was shown to be most effectively prevented by inserting a thin brittle layer between ductile base and cover layers. This sandwich packaging also prevents probing of the lines, as any physical force on the lines results in cracks forming in the substrate that break the lines. Laser depackaging may be hampered using a transparent encapsulation. The dark colour of the PEDOT absorbs more laser energy than the transparent packaging, thus the lines are destroyed before the packaging is removed.

While the focus of the security evaluation was to achieve maximum damage to the protection grids, the prototype evaluation focused on measuring the resistance and detecting changes. The evaluation of the prototype protection grid showed that there is a certain spread in measured delay times, therefore stochastic methods may be used to enhance the resolution of detectable resistance change of the lines. Organic transistors are not stable enough for repeatable measurements of properties. The voltages required for sufficient conduction of the transistors was also found to be too high for operation with a standard microcontroller. Until transistors with better properties become available, active protection grids are not feasible.

With the exception of the poor ageing of PEDOT, the evaluation indicates the feasibility of a passive protection grid scheme. However, the experiments were carried out on separate test structures printed on glass or polymer substrates. Future work to complete the development of passive organic protection grids must be concerned with fabrication of a full prototype. Packaging materials need to be evaluated to provide better encapsulation to minimise ageing, and mechanical robustness in everyday handling, while preserving sensitivity to depackaging. Furthermore, the vulnerability of the protection schemes against non-invasive methods has not been tested. Circuits which can reliably measure the properties of the protection grids without making the device vulnerable to power analysis of the protection grid circuitry need to be designed and tested. The evaluation of sensitivity to depackaging methods indicates that damage to the grid is likely if the removal of the package is attempted. However, a more robust security test would be to have an independent attack trial carried out on a full prototype.

The evaluation of the secure dual-rail domino logic has shown that it is not sufficient to simply balance the amount of charge consumed in a clock cycle. The timing of the evaluation or switching of a logic cell may vary due to the early propagation. This results in a time shift of the evaluation current peak, which was shown to be present in the measured power consumption traces. A secondary timing imbalance occurs if different numbers of transistors discharge the dynamic logic gate, for different input data. While a timing difference could not be distinguished in the measured multiplier circuit, this is not proof that these differences can be neglected. As the data-dependent evaluation time is the main source of imbalance, future research will need to focus on eliminating these differences, in order to achieve true data-independent power consumption and thus resistance against power analysis.

A

Appendix One

Details of printing setup

A.1 Introduction

The printer used to fabricate the PEDOT lines is custom-built design based on an Epson piezo-electic print head. Piezo-electric print heads are particularly suitable for patterning organic electronics, as the drop ejection mechanism is purely mechanical. In contrast, HP/Canon inkjet printers use bubble-jet technology, which is based on boiling the ink to expunge the drops. This thermal cycling may have detrimental effects on the electronic material, as well as being problematic when if the boiling point of the solvent of the ink is different from the design value (e.g. when printing organic solvent-based ink, or DMSO/water mixtures).

When moving laboratories (due to the closure of the laboratory I was originally working in), I took over a disused printer that was partially disassembled and whose driver software was unfit for my purposes. Therefore I re-built the printer, making minor alterations to the setup, and re-wrote most of the software.

Depending how careful a print head is treated, it will last between a week of daily use to several months of use. The standard mode of failure is blocking of nozzles. Therefore the ink should be filtered before use to remove aggregated (PEDOT–) particles and dust. Most of the time, only a single nozzle is used during a print run, therefore the remaining nozzles dry out starting from the surface/nozzle plate. Wiping the nozzle plate and rinsing usually clears the nozzles again, but with time more and more nozzles will cease to function.

A.2 Top-level design

The top-level design of the printer was apparent from the state it was found in and not altered significantly apart from the addition of the nitrogen waft. A photo of the re-assembled printer is shown in Figure A.1. The print head is held fixed while the sample is moved below it using the x/y motion stages. The print head may be moved along the horizontal print head mounting beam in order to facilitate loading a sample or for servicing the head without interfering with a loaded sample, i.e. eject ink drops to prevent nozzle clogging when idle (termed idling), fill ink, or rinse the head using the nozzle cleaning suction cup syringe. The camera objective is placed below the printer to monitor the sample in real time (on the top PC monitor). The printer is mounted on an optical table to reduce vibrations and thus increase print quality. Similarly, the nitrogen waft aids drying of the ink to improve print quality. The fume hood is set to a

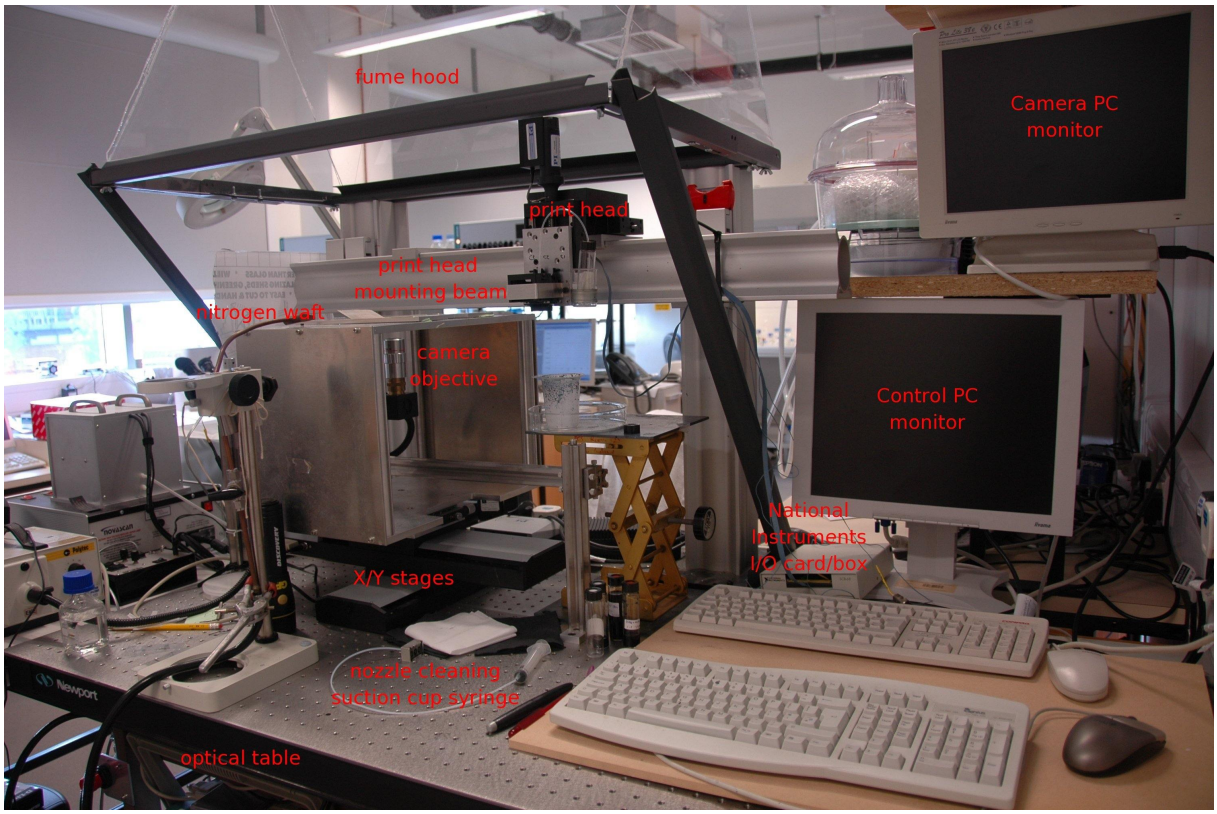

Figure A.1: Printer setup

moderate flow rate to extract fumes and stray droplets (during idling), but not to disturb the printing process. The printer setup is controlled from the control PC which also controls the National Instruments I/O card.

The schematic of the printer is shown in Figure A.2. The print setup comprises eight different modules. The control PC (running LabView) forms the core of the printer, controlling the other components. In order to control the other devices, it has dedicated I/O hardware for the PI motion stages, a GPIB card to control the data generator, and a National Instruments (NI) I/O card with breakout box.

To ensure accurate timing and position of the droplets, the PC controls a Sony/Tektronix DG2030 data generator via GPIB. The data generator is a device that plays back pre-programmed digital waveforms with accurate timing. If the velocity of the motion stage is uniform, the accurate data timing allows regular spacing of the ink drops. The time resolution of the data generator was set to 1 µs.

Similar functionality may also be achieved using FPGA or microcontroller circuits, so long as the output is clock-cycle accurate and the oscillator is sufficiently stable. The data generator output controls the trigger input of the arbitrary waveform generator, such that each square pulse is essentially converted into the piezo drive waveform. The drive pulse is then amplified using a separate piezo drive amplifier, as the print head does not have its own piezo driver. The NI I/O box controls the data communication with the print head, i.e. to set which nozzle to print from. The motion stages are standard devices that are delivered with LabView drivers. The only modification made to the motion stages is that the trigger line of the motion stage was also connected to the AND gate. The second input to the AND gate is an enable/disable line from the NI I/O box. The output of the AND gate may then be used to trigger and synchronise the data generator/print head and the motion stage.

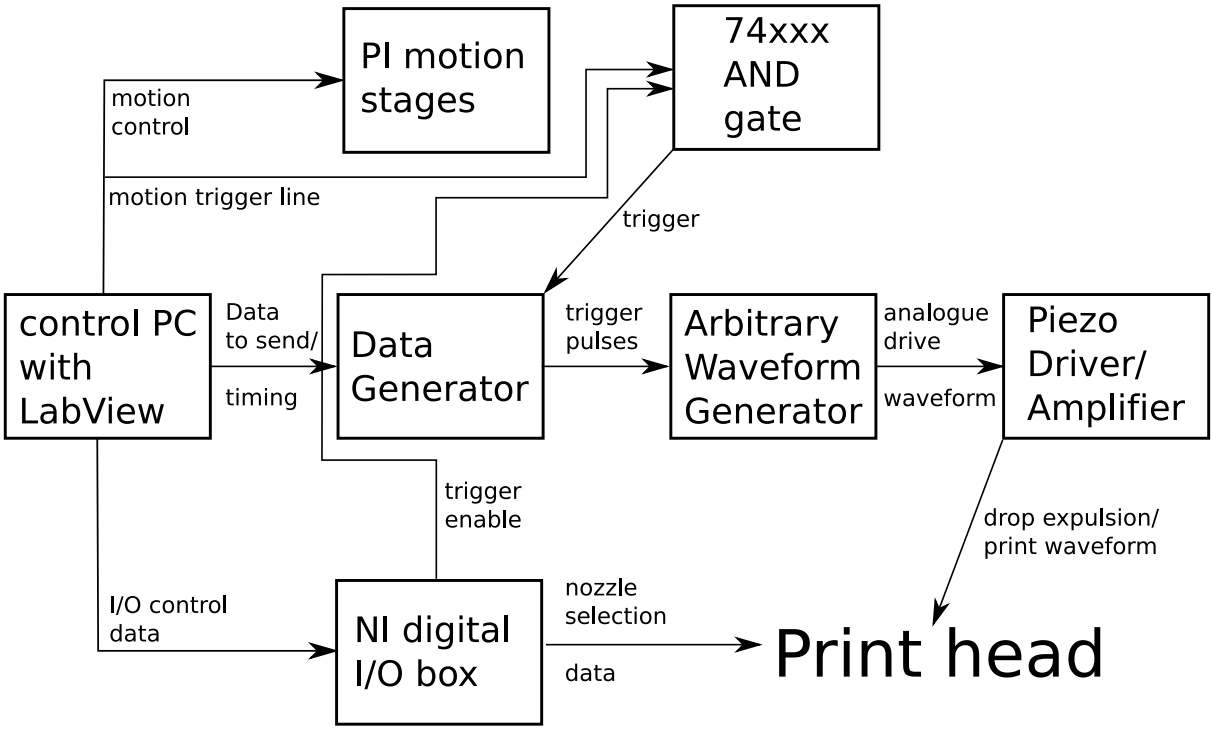

Figure A.2: Top-level schematic of the printer

A.3 Detailed design

A.3.1 Motion stages

The computer-controlled motion stages used were made by Physik Instrumente (PI). The stages may be programmed with the total travel distance and velocity. This allows the spacing of the drops to be calculated from the drop ejection frequency. As the motion stages need some distance to accelerate, the first pattern printed is printed incorrectly. Therefore, any pattern should be repeated, so that the first one can be discarded.

Only the X and Y motion stages were controlled with the PC. The Z-stage controlling the height of the head above the sample is operated manually, as it only needs adjustment when a new head is mounted, or a sample of different thickness.

A.3.2 Print head mounting

A detailed image of the print head mounting is shown in Figure A.3 The head assembly is held in place with a mounting bracket and two screws (black hexagonal). As the contacts to the circuit board of the print head is made via pads, spring-loadeded pins are used to make the electrical connection (transparent plate in centre). The screw holes of the mounting bracket (not shown) are actually u-shaped, to allow the position of the print head to be adjusted and align the pins with the contact pads.

Depending on the viscosity and surface tension of the ink, the liquid pressure may need to be adjusted to either enable reliable drop ejection, or prevent seeping. For this purpose, the ink bottle may be raised or lowered to achieve higher or lower pressure (one can unscrew the holder). Alternatively, the ink may be loaded into a syringe connected to the same type plastic tube to control the pressure manually. However, these methods are rather crude. Other printers,

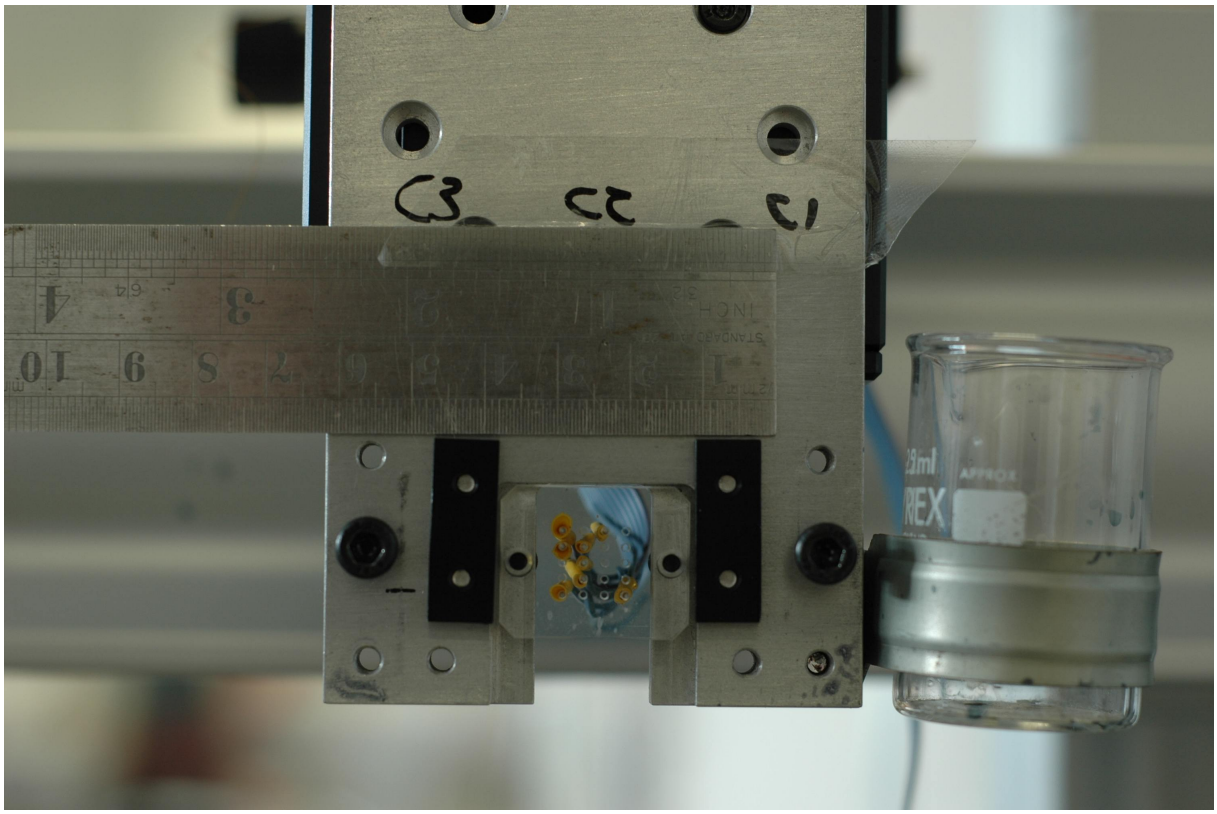

Figure A.3: Detail of the print head mounting plate, contacts, and ink vial holder

e.g. by 'Microdrop' have in-built ink pressure control facilities, which allow better pressure control.

A.3.3 Drop ejection

Drop ejection happens by a change of shape of the piezo crystal, depending on print head design either by a forward 'push' motion, or by bending/flexing of a beam [Le, 1998]. The waveform for drop ejection in general is a trapeze shape, but the details may vary. The waveform used for this print head is shown in Figure A.4. This waveform was pre-programmed into the function generator. The slopes of the trapeze shape, and the voltage levels are the two main parameters to control drop ejection. The waveform may be adjusted to work with ink of different viscosity and surface tension, which will require different levels of drive force.

The waveform shown in Figure A.4 to drive the piezo print head had been pre-programmed into the function generator, and served well for printing the PEDOT. Therefore there was no need to make changes to the waveform. If a drive waveform needs to be designed for a custom-built printer, the following articles may be of use: [Dong et al., 2006; Gan et al., 2009; Kim et al., 2006; Kwon and Kim, 2007]

The maximum frequency of drop ejection for inkjet print heads is specified in the printer's service manual (which may be found on the internet from third-party sites). For this printer, the repetition frequency was set to a maximum of 10kHz, but other printers allow higher frequencies (e.g. 14.4 kHz for the Epson Stylus Color 400). The relevant service manual also shows the connection diagram of the print head and the required supply voltages (5 V or 5 V/ 42 V).

As the output voltage of the function generator is low, a Trek 603 piezo amplifier is used to reach the appropriate voltage levels. The Trek 603 amplifier was used with a gain of 50, at

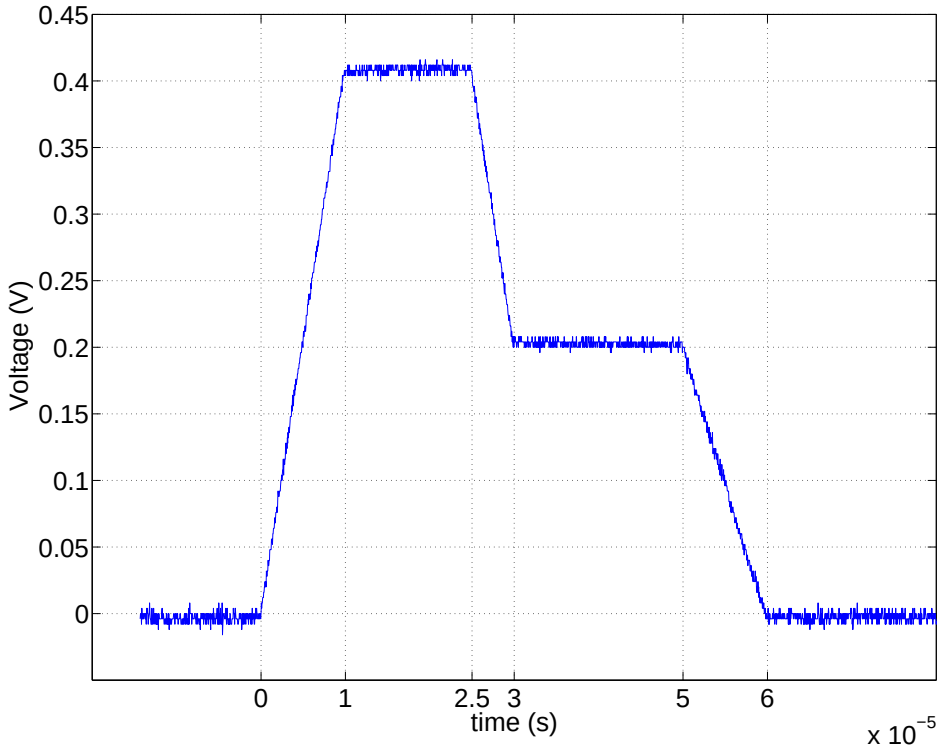

Figure A.4: Drop ejection waveform

an output range of 125 V and a "hi" load range. Some newer print heads have an in-built piezo drive amplifier, making the design of a custom-built printer easier. The required voltages thus depend very much on the print head design. Alternatively, if the cost of a dedicated piezo amplifier is too high for a custom built printer, Cirrus Logic/Apex offer integrated circuits to drive piezo actuators and print heads. There are white papers and application notes that can be downloaded from their web site[1].

Nozzles are selected by sending an n-bit data word to the nozzle control pin of the print head, where n is the number of nozzles on the print head. The print head control circuit essentially acts as a multiplexer to distribute the drive pulse to the individual piezo crystals. For some print heads, the nozzle selection is split by colour, with one line dedicated to each colour, or one line for black and a second for all colours. The details are also described in the relevant service manuals.

A.3.4 Software

The printer is controlled from a PC running a LabView program. First, the nozzle(s) for printing are selected. This may be done either by ejecting ink onto a blank slide and checking correct drop ejection using the live camera. Alternatively, a torch may be used to shine light onto the nozzles from behind, which makes the ejected ink droplets visible (especially against a dark background). Then, the pattern is chosen as an array of drops, with a defined drop spacing in

[1] http://www.cirrus.com/en/products/apex/documents.html
http://www.cirrus.com/en/support/design/whitepapers.html
http://www.cirrus.com/en/pubs/appNote/Apex_AN44U_1.pdf
http://www.cirrus.com/en/pubs/whitePaper/1105_Drive_PE_Actuators_with_Op_Amps.pdf

X and Y. The number and spacing of repetitions of the pattern in X and Y is also set, as well as the number of layers printed on top of each other. Further settings that have proven useful is printing a pattern in several runs (as described in the main dissertation), and selecting extra time for a print pattern to dry before printing the next set of drops).

Further control functions of the software is to set all nozzles to idle mode (eject ink when not in use), move the stage without printing, and translate selected lengths of simple lines to the appropriate drop pattern.

References

D. G. Abraham, G. M. Dolan, G. P. Double, and J. V. Stevens. Transaction security system. *IBM Syst. J.*, 30(2):206–229, 1991. ISSN 0018-8670.

Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, and Pankaj Rohatgi. The EM Side-Channel(s). In *CHES '02: Revised Papers from the 4th International Workshop on Cryptographic Hardware and Embedded Systems*, pages 29–45, London, UK, 2003. Springer-Verlag. ISBN 3-540-00409-2.

Mehdi-Laurent Akkar and Christophe Giraud. An Implementation of DES and AES, Secure against Some Attacks. In *CHES '01: Proceedings of the Third International Workshop on Cryptographic Hardware and Embedded Systems*, volume 2162/2001 of *Lecture Notes in Computer Science*, pages 309–318, London, UK, 2001. Springer Berlin / Heidelberg. ISBN 3-540-42521-7. doi: 10.1007/3-540-44709-1.

Mehdi-Laurent Akkar and Louis Goubin. A Generic Protection against High-Order Differential Power Analysis. In *Fast Software Encryption 10th International Workshop, FSE 2003, Lund, Sweden, February 24-26, 2003, Revised Papers*, volume 2887 of *Lecture Notes in Computer Science*, pages 192–205, 2003. doi: 10.1007/b93938.

Mehdi-Laurent Akkar, Régis Bevan, Paul Dischamp, and Didier Moyart. Power Analysis, What Is Now Possible.... In *ASIACRYPT '00: Proceedings of the 6th International Conference on the Theory and Application of Cryptology and Information Security*, pages 489–502, London, UK, 2000. Springer-Verlag. ISBN 3-540-41404-5.

Mehdi-Laurent Akkar, Régis Bévan, and Louis Goubin. Two Power Analysis Attacks against One-Mask Methods. In *Fast Software Encryption*, volume 3017/2004 of *Lecture Notes in Computer Science*, pages 332–347. Springer Berlin / Heidelberg, 2004. doi: 10.1007/b98177.

Sigma Aldrich. Properties of DMSO, http://www.sigmaaldrich.com/catalog/search/ProductDetail/FLUKA/41650, accessed 2^{nd} Sept. 2008.

Sigma Aldrich. Detergents / Surfactants, http://www.sigmaaldrich.com/Area_of_Interest/Biochemicals/BioUltra/Detergents_Surfactants.html?cm_mmc=wiki-_-social-_-surfactants-_-Surfactants, accessed 5^{th} Sept. 2008.

R. Anderson. Why information security is hard - an economic perspective. *Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual*, 0:358–365, Dec. 2001a.

R. Anderson and M. Kuhn. Tamper Resistance – a Cautionary Note. In *The Second USENIX Workshop on Electronic Commerce Proceedings*, pages 1–11, Oakland, California, 18–21 1996. ISBN 1-880446-83-9.

R. Anderson, M. Bond, J. Clulow, and S. Skorobogatov. Cryptographic Processors – a survey. *Proceedings of the IEEE*, 94(2):357–369, Feb. 2006.

R. J. Anderson and T. Moore. The Economics of Information Security. *Science*, 314:610–613, October 2006. doi: 10.1126/science.1130992.

Ross J. Anderson. Why cryptosystems fail. *Commun. ACM*, 37(11):32–40, 1994a. ISSN 0001-0782. doi: 10.1145/188280.188291.

Ross J. Anderson. Liability and Computer Security: Nine Principles. In *ESORICS '94: Proceedings of the Third European Symposium on Research in Computer Security*, pages 231–245, London, UK, 1994b. Springer-Verlag. ISBN 3-540-58618-0.

Ross J. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley, 2 edition, 2001b.

Ross J. Anderson and Markus G. Kuhn. Low Cost Attacks on Tamper Resistant Devices. In *Proceedings of the 5th International Workshop on Security Protocols*, pages 125–136, London, UK, 1998. Springer-Verlag. ISBN 3-540-64040-1.

Roberto M. Avanzi. Countermeasures against Differential Power Analysis for Hyperelliptic Curve Cryptosystems. In *Cryptographic Hardware and Embedded Systems CHES 2003*, volume 2779/2003 of *Lecture Notes in Computer Science*, pages 366–381. Springer Berlin / Heidelberg, 2003.

Falguni Bala and Tapas Nandy. Conventional RC oscillators, though offer inexpensiveProgrammable High Frequency RC Oscillator. In *VLSID '05: Proceedings of the 18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design*, pages 511–515, Washington, DC, USA, 2005. IEEE Computer Society. ISBN 0-7695-2264-5. doi: 10.1109/ICVD.2005.70.

W. Bantikassegn and O. Inganäs. Electronic properties of junctions between aluminum and doped poly(3,4-ethylenedioxythiophene). *Thin Solid Films*, 293(1-2):138–143, January 1997.

Hagai Bar-El, H. Choukri, D. Naccache, M. Tunstall, and C. Whelan. The sorcerer's apprentice guide to fault attacks. *Proceedings of the IEEE*, 94(2):370–382, Feb 2006.

P. F. Baude, D. A. Ender, M. A. Haase, T. W. Kelley, D. V. Muyres, and S. D. Theiss. Pentacene-based radio-frequency identification circuitry. *Applied Physics Letters*, 82(22):3964–3966, 2003.

Friedrich Beck and Stephen S. Wilson. *Integrated Circuit Failure Analysis*. John Wiley and Sons, 1998.

Luca Benini, Alberto Macii, Enrico Macii, Elvira Omerbegovic, Fabrizio Pro, and Massimo Poncino. Energy-aware design techniques for differential power analysis protection. In *DAC '03: Proceedings of the 40th conference on Design automation*, pages 36–41, New York, NY, USA, 2003. ACM. ISBN 1-58113-688-9. doi: 10.1145/775832.775845.

John C. Berg. *Wettability*. CRC Press, 1993. ISBN 0-8247-9046-4.

Régis Bevan and Erik Knudsen. Ways to Enhance Differential Power Analysis. In *Information Security and Cryptology - ICISC 2002: 5th International Conference, Seoul, Korea, November 28-29, 2002. Revised Papers*, volume 2587/2003 of *Lecture Notes in Computer Science*, pages 327–341, 2003.

Eli Biham and Adi Shamir. *Differential Fault Analysis of Secret Key Cryptosystems*, volume 1294 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, advances in cryptology – crypto '97: 17th annual international cryptology conference, santa barbara, california, usa, august 1997. proceedings edition, 1997.

Johannes Blomer, Jorge Guajardo Merchan, and Volker Krummel. Provably Secure Masking of AES. In *In SAC*, volume 3357/2005, pages 69–83. Springer Berlin / Heidelberg, 2004.

S. Blythe, B. Fraboni, S. Lall, H. Ahmed, and U. de Riu. Layout reconstruction of complex silicon chips. *Solid-State Circuits, IEEE Journal of*, 28(2):138–145, Feb 1993. doi: 10.1109/4.192045.

Béatrice Bonvalot and Joseph Leibenguth. Ultimate hardware security for smart cards. In *Minatec*, http://www.minatec.com/minatec2003/act_pdf/5_FRIDAY_BONVALOT.pdf, 2003.

Eric Brier, Christophe Clavier, and Francis Olivier. Correlation Power Analysis with a Leakage Model. In *Cryptographic Hardware and Embedded Systems CHES 2004*, volume 3156/2004 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2004. doi: 0.1007/b99451.

L. Bürgi, T. J. Richards, R. H. Friend, and H. Sirringhaus. Close look at charge carrier injection in polymer field-effect transistors. *Journal of Applied Physics*, 94(9):6129–6137, 2003.

Seamus E. Burns, Paul Cain, John Mills, Jizheng Wang, and Henning Sirringhaus. Inkjet Printing of Polymer Thin-Film Transistor Circuits. *MRS Bulletin*, 28,11:829–835, 2003.

Susan Burns and George R. S. Weir. Trends in Smartcard Fraud. In Hamid Jahankhani, Kenneth Revett, and Dominic Palmer-Brown, editors, *Global E-Security, 4th International Conference, ICGeS 2008, London, UK, June 23-25, 2008. Proceedings*, volume 12 1 of *Communications in Computer and Information Science*, pages 40–47. Springer Berlin Heidelberg, 2008. doi: 10.1007/978-3-540-69403-8_5.

Andrew Burnside, Ahmet Erdogan, and Tughrul Arslan. The Re-emission Side Channel. *Bio-inspired, Learning, and Intelligent Systems for Security, ECSIS Symposium on*, 0:154–159, 2008. doi: 10.1109/BLISS.2008.22.

Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards Sound Approaches to Counteract Power-Analysis Attacks. In *CRYPTO '99: Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, pages 398–412, London, UK, 1999. Springer-Verlag. ISBN 3-540-66347-9.

Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template Attacks. In *CHES '02: Revised Papers from the 4th International Workshop on Cryptographic Hardware and Embedded Systems*, pages 13–28, London, UK, 2003. Springer-Verlag. ISBN 3-540-00409-2.

M. Chason, Jr. Brazis, P.W., Jie Zhang, K. Kalyanasundaram, and D.R. Gamota. Printed organic semiconducting devices. *Proceedings of the IEEE*, 93(7):1348–1356, July 2005. doi: 10.1109/JPROC.2005.850306.

Bin Chen, Tianhong Cui, Yi Liu, and Kody Varahramyan. All-polymer RC filter circuits fabricated with inkjet printing technology. *Solid-State Electronics*, 47:841–847, 2003.

Benoit Chevallier-Mames, Mathieu Ciet, and Marc Joye. Low-Cost Solutions for Preventing Simple Side-Channel Analysis: Side-Channel Atomicity. *IEEE Transactions on Computers*, 53 (6):760–768, 2004. ISSN 0018-9340. doi: 10.1109/TC.2004.13.

Clariant. *AZ5214E Image Reversal Photoresist Product Data Sheet*. Clariant GmbH, Business Unit Electronic Materials, Rheingaustrasse 190, D-65203 Wiesbaden, Germany.

Christophe Clavier, Jean-Sebastien Coron, and Nora Dabbous. Differential Power Analysis in the Presence of Hardware Countermeasures. In *CHES '00: Proceedings of the Second International Workshop on Cryptographic Hardware and Embedded Systems*, pages 252–263, London, UK, 2000. Springer-Verlag. ISBN 3-540-41455-X.

Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates, 1988.

Ewan H. Conradie and David F. Moore. SU-8 thick photoresist processing as a functional material for MEMS applications. *Journal of Micromechanics and Microengineering*, 12(4):368–374, July 2002.

Jean-Sébastien Coron. Resistance against Differential Power Analysis for Elliptic Curve Cryptosystems. In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, pages 292–302, London, UK, 1999. Springer-Verlag. ISBN 3-540-66646-X.

Jean-Sébastien Coron and Louis Goubin. On Boolean and Arithmetic Masking against Differential Power Analysis. In *CHES '00: Proceedings of the Second International Workshop on Cryptographic Hardware and Embedded Systems*, volume 1965/2000 of *Lecture Notes in Computer Science*, pages 231–237, London, UK, 2000. Springer Berlin / Heidelberg. ISBN 3-540-41455-X. doi: 10.1007/3-540-44499-8.

Jean-Sébastien Coron, David Naccache, and Paul Kocher. Statistics and secret leakage. *Trans. on Embedded Computing Sys.*, 3(3):492–508, 2004. ISSN 1539-9087. doi: 10.1145/1015047.1015050.

B.-J. de Gans, P. C. Duineveld, and U. S. Schubert. Inkjet Printing of Polymers: State of the Art and Future Developments. *Advanced Materials*, 16:203 – 213, 2003. doi: 10.1002/adma.200300385.

Pierre-Gilles de Gennes, Françoise Brochard-Wyart, David Quéré, and Axel Reisinger. *Capillarity and Wetting Phenomena*. Springer, 2004.

Jean-François Dhem, François Koeune, Philippe-Alexandre Leroux, Pat rick Mestré, Jean-Jacques Quisquater, and Jean-Louis Willems. A Practical Implementation of the Timing Attack. In *CARDIS '98: Proceedings of the The International Conference on Smart Card Research and Applications*, pages 167–182, London, UK, 2000. Springer-Verlag. ISBN 3-540-67923-5.

Christos D. Dimitrakopoulos and Patrick R.L. Malenfant. Organic Thin Film Transistors for Large Area Electronics. *Advanced Materials*, 14(2):99–117, January 2002. doi: 10.1002/1521-4095(20020116)14:2 99::AID-ADMA99 3.0.CO;2-9.

Christos D. Dimitrakopoulos and D. J. Mascaro. Organic thin-film transistors: a review of recent advances. *IBM J. Res. & Dev.*, 45(1):11 – 27, January 2001.

Hongming Dong, Wallace W. Carr, and Jeffrey F. Morris. Visualization of drop-on-demand inkjet: Drop formation and deposition. *Review of Scientific Instruments*, 77(8):085101, 2006. doi: 10.1063/1.2234853. URL `http://link.aip.org/link/?RSI/77/085101/1`.

Joan G. Dyer, Mark Lindemann, Ronald Perez, Reiner Sailer, Leendert van Doorn, Sean W. Smith, and Steve Weingart. Building the IBM 4758 Secure Coprocessor. *Computer*, 34(10):57–66, 2001. ISSN 0018-9162. doi: 10.1109/2.955100.

Badih El-Kareh. *Fundamentals of Semiconductor Processing Technologies*. Springer, 1995.

Electrolube. *Silver Conductive Paint Materials Safety Data Sheet*. Electrolube, Kingsbury Park, Midland Road, Swadlincote, Derbyshire, DE11 0AN, UK http://www.electrolube.com/docs/-msds/044/044SCP.pdf.

W. C. Elmore. The Transient Response of Damped Linear Networks with Particular Regard to Wideband Amplifiers. *Journal of Applied Physics*, 19(1):55–63, 1948. doi: 10.1063/1.1697872.

Andreas Elschner. Baytron P as a Functional Layer in OLEDs – physical and electrical characterization. Baytron Micro Symposium 2002 in Cologne, Germany, H.C. Starck, 2002.

Epson. *Epson Stylus Color II Printer Specifications*. `http://files.support.\\epson.com/pdf/sc2\_\_\_/sc2\_\_\_sl.pdf` , Accessed 1^{st} Sept. 2008.

A. Facchetti, M.-H. Yoon, and T.J. Marks. Gate Dielectrics for Organic Field-Effect Transistors: New Opportunities for Organic Electronics. *Advanced Materials*, 17(14):1705–1725, 2005. doi: 10.1002/adma.200500517.

Antonio Facchetti. Semiconductors for organic transistors. *Materials Today*, 10(3):28–37, March 2007.

Paul N. Fahn and Peter K. Pearson. IPA: A New Class of Power Attacks. In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, pages 173–186, London, UK, 1999. Springer-Verlag. ISBN 3-540-66646-X.

James K. Ferri and Kathleen J. Stebe. Which surfactants reduce surface tension faster? A scaling argument for diffusion-controlled adsorption. *Advances in Colloid and Interface Science*, 85(1): 61–97, February 2000. doi: 10.1016/S0001-8686(99)00027-5.

Jacques J. A. Fournier, Simon Moore, Huiyun Li, Robert Mullins, and George Taylor. Security Evaluation of Asynchronous Circuits. In *In Proceedings of Cryptographic Hardware and Embedded Systems - CHES2003*, volume 2779/2003 of *Lecture Notes in Computer Science*, pages 137–151. Springer Berlin / Heidelberg, 2003.

Sami Franssila. *Introduction to Microfabrication*. John Wiley and Sons, 2004. ISBN 0470851058.

H Y Gan, Xuechuan Shan, T Eriksson, B K Lok, and Y C Lam. Reduction of droplet volume by controlling actuating waveforms in inkjet printing for micro-pattern formation. *Journal of Micromechanics and Microengineering*, 19(5):055010 (8pp), 2009. URL `http://stacks.iop.org/0960-1317/19/055010`.

Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic Analysis: Concrete Results. In *CHES '01: Proceedings of the Third International Workshop on Cryptographic Hardware and Embedded Systems*, pages 251–261, London, UK, 2001. Springer-Verlag. ISBN 3-540-42521-7.

Erik Garnett and David Ginley. Electrical and Morphological Properties of Inkjet Printed PEDOT/PSS Films. *Journal of Undergraduate Research*, V:24–29, 2005.

Catherine H. Gebotys. Design of secure cryptography against the threat of power-attacks in DSP-embedded processors. *Trans. on Embedded Computing Sys.*, 3(1):92–113, 2004. ISSN 1539-9087. doi: 10.1145/972627.972632.

Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. *Algorithmic Tamper-Proof (ATP) Security: Theoretical Foundations for Security Against Hardware Tampering*, volume 2951 of *Lecture Notes in Computer Science*. Springer, 2004. doi: 10.1007/b95566.

G. Giustolisi, G. Palumbo, M. Criscione, and F. Cutri. A low-voltage low-power voltage reference based on subthreshold MOSFETs. *Solid-State Circuits, IEEE Journal of*, 38(1):151–154, Jan 2003. ISSN 0018-9200. doi: 10.1109/JSSC.2002.806266.

H. L. Gomes, P. Stallinga, M. Colle, D. M. de Leeuw, and F. Biscarini. Electrical instabilities in organic semiconductors caused by trapped supercooled water. *Applied Physics Letters*, 88(8): 082101, 2006.

Gore. D3 anti-tamper sytem. http://www.gore.com/en_xx/products/electronic/specialty/-antitamper.html, Accessed 29^{th} Sept. 2008.

Louis Goubin and Jacques Patarin. DES and Differential Power Analysis (The "Duplication" Method). In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, volume 1717/1999 of *Lecture Notes in Computer Science*, pages 158–172, London, UK, 1999. Springer-Verlag. ISBN 3-540-66646-X.

F. Grassert and D. Timmermann. Dynamic single phase logic with self-timed stages for power reduction in pipeline circuit designs. In *Circuits and Systems, 2001. ISCAS 2001. The 2001 IEEE International Symposium on*, volume 4, pages 144–147, 6–9 May 2001. doi: 10.1109/ISCAS.2001.922191.

G. Greczynski, T. Kugler, M. Keil, W. Osikowicz, M. Fahlman, and W.R. Salaneck. Photoelectron spectroscopy of thin films of PEDOT-PSS conjugated polymer blend: a mini-review and some new results. *Journal of Electron Spectroscopy and Related Phenomena*, 121:1–17(17), December 2001. doi: 10.1016/S0368-2048(01)00323-1.

L. Groenendaal, F. Jonas, D. Freitag, H. Pielartzik, and J. R. Reynolds. Poly(3,4-ethylenedioxythiophene) and Its Derivatives: Past, Present, and Future. *Advanced Materials*, 12(7):481 – 494, 2000. doi: 10.1002/(SICI)1521-4095(200004)12:7 481::AID-ADMA481 3.0.CO;2-C.

D. J. Gundlach, L. Zhou, J. A. Nichols, T. N. Jackson, P. V. Necliudov, and M. S. Shur. An experimental study of contact effects in organic thin film transistors. *Journal of Applied Physics*, 100 (2):024509, 2006.

Peter Gutmann. Data Remanence in Semiconductor Devices. In *Proceedings of the 10th USENIX Security Symposium*, pages 39–54, 2001.

M. Halik, H. Klauk, U. Zschieschang, G. Schmid, S. Ponomarenko, S. Kirchmeyer, and W. Weber. Relationship Between Molecular Structure and Electrical Performance of Oligothiophene Organic Thin Film Transistors. *Advanced Materials*, 15(11):917–922, 2003. doi: 10.1002/adma.200304654.

Mahiar Hamedi, Robert Forchheimer, and Olle Inganas. Towards woven logic from organic electronic fibres. *Nature Materials*, 6:357–362, 2007. doi: 10.1038/nmat1884.

Helena Handschuh and Howard M. Heys. A Timing Attack on RC5. In *SAC '98: Proceedings of the Selected Areas in Cryptography*, pages 306–318, London, UK, 1999. Springer-Verlag. ISBN 3-540-65894-7.

Helena Handschuh, Pascal Paillier, and Jacques Stern. Probing Attacks on Tamper-Resistant Devices. In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, pages 303–315, London, UK, 1999. Springer-Verlag. ISBN 3-540-66646-X.

A.J. Heeger. Semiconducting and Metallic Polymers: The Fourth Generation of Polymeric Materials. *Journal of Physical Chemistry B*, 105(36):8475–8491, 2001. ISSN 1520-6106.

Gilles Horowitz. Organic thin film transistors: from theory to real device. *J. Mater. Res.*, 19:1946, 2004.

P. Horowitz and W. Hill. *The Art of Electronics*. Cambridge University Press, New York, NY, USA, 1989. ISBN 0-521-37095-7.

H. Hu and R.G. Larson. Evaporation of a Sessile Droplet on a Substrate. *Journal of Physical Chemistry B*, 106(6):1334–1344, 2002. ISSN 1520-6106. doi: 0.1021/jp0118322.

S K M Jönsson, W R Salaneck, and M Fahlman. X-ray photoelectron spectroscopy study of the metal/polymer contacts involving aluminum and poly[3,4-(ethylenedioxy)thiophene]-poly(styrenesulfonic acid) derivatives. *Journal of Materials Research*, 18(5):1219–1226, May 2003. doi: 10.1557/JMR.2003.0167.

A.B. Kahng and S. Muddu. Delay Analysis of VLSI Interconnections Using the Diffusion Equation Model. *Design Automation, 1994. 31st Conference on*, 0:563–569, June 1994. ISSN 0738-100X.

H. E. Katz. Organic molecular solids as thin film transistor semiconductors. *Journal of material chemistry*, 7(3):369–376, 1997. doi: 10.1039/a605274f.

H.E. Katz. Recent Advances in Semiconductor Performance and Printing Processes for Organic Transistor-Based Electronics. *Chemistry of Materials*, 16(23):4748–4756, 2004. ISSN 0897-4756.

T. Kawase, H. Sirringhaus, R. H. Friend, and T. Shimoda. Inkjet Printed Via-Hole Interconnections and Resistors for All-Polymer Transistor Circuits. *Advanced Materials*, 13(21):1601–1605, November 2001. doi: 10.1002/1521-4095(200111)13:21 1601::AID-ADMA1601 3.0.CO;2-X.

Takeo Kawase, Soichi Moriya, Christopher J. Newsome, and Tatsuya Shimoda. Inkjet Printing of Polymeric Field-Effect Transistors and Its Applications. *Japanese Journal of Applied Physics*, 44(6A):3649–3658, 2005. doi: 10.1143/JJAP.44.3649.

Robert W. Kelsall, Ian W. Hamley, and Mark Geoghegan. *Nanoscale Science and Technology*. John Wiley and Sons, 2005.

K. Kerckhoffs. La cryptographie militaire. *J Sci Militaires*, 9:5–38, 1883.

Youngjae Kim, Wonchul Sim, Changsung Park, Youngseuck Yoo, Jaewoo Joung, and Yongsoo Oh. The effects of driving waveform of piezoelectric industrial inkjet head for fine patterns. pages 826–831, Jan. 2006. doi: 10.1109/NEMS.2006.334905.

Hagen Klauk. *Organic Electronics*. Wiley-VCH, 2006.

Hagen Klauk, Günter Schmid, Wolfgang Radlik, Werner Weber, Lisong Zhou, Chris D. Sheraw, Jonathan A. Nichols, and Thomas N. Jackson. Contact resistance in organic thin film transistors. *Solid-State Electronics*, 47(2):297–301, February 2003. doi: 10.1016/S0038-1101(02)00210-1.

Paul Kocher. Design and Validation Strategies for Obtaining Assurance in Countermeasures to Power Analysis and Related Attacks. This paper was prepared for the NIST Physical Security Workshop, 26-29 Sept. 2005, September 2005.

Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *CRYPTO '96: Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology*, pages 104–113, London, UK, 1996. Springer-Verlag. ISBN 3-540-61512-1.

Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. *Lecture Notes In Computer Science*, 1666:388 – 397, 1999. ISSN 0302-9743.

Oliver Koemmerling. US patent 7005733: Anti-tamper encapsulation for an integrated circuit. US patent office, 2000.

François Koeune and François-Xavier Standaert. A Tutorial on Physical Security and Side-Channel Attacks. *Lecture Notes in Computer Science*, 3655:78–108, 2005. doi: 10.1007/11554578_3.

Oliver Kömmerling and Markus G. Kuhn. Design Principles for Tamper-Resistant Smartcard Processors. In *Proceedings of the USENIX Workshop on Smartcard Technology, Chicago, 10–11 May, 1999.*, pages 9–20, 1999.

Konrad Kulikowski, Alexander Smirnov, and A. Taubin. Automated Design of Cryptographic Devices Resistant to Multiple Side-Channel Attacks. In *Cryptographic Hardware and Embedded Systems CHES 2006*, Lecture Notes in Computer Science, pages 399–413. Springer Berlin / Heidelberg, 2006a. doi: 10.1007/11894063.

Konrad J. Kulikowski, Mark G. Karpovsky, and Alexander Taubin. Power Attacks on Secure Hardware Based on Early Propagation of Data. *IEEE International On-Line Testing Symposium/On-Line Testing Symposium, IEEE International*, 0:131–138, 2006b. doi: 10.1109/IOLTS.2006.49.

Kye-Si Kwon and Wousik Kim. A waveform design method for high-speed inkjet printing based on self-sensing measurement. *Sensors and Actuators A: Physical*, 140(1):75 – 83, 2007. ISSN 0924-4247. doi: DOI:10.1016/j.sna.2007.06.010. URL `http://www.sciencedirect.com/science/article/B6THG-4P06CD3-4/2/5a270a7615d0315bb67b9f78c92f47cc`.

I. Kymissis, C.D. Dimitrakopoulos, and S. Purushothaman. High-performance bottom electrode organic thin-film transistors. *Electron Devices, IEEE Transactions on*, 48(6):1060–1064, June 2001. doi: 10.1109/16.925226.

David M. Lane. *Hyperstat*. Atomic Dog Publishing, 2 edition, 1999. online version at http://davidmlane.com/hyperstat/.

Hue P. Le. Progress and Trends in Ink-jet Printing Technology. *Journal of Imaging Science and Technology*, 42(1):49–62, January/February 1998.

H.-H. Lee, K.-S. Chou, and K.-C. Huang. Inkjet printing of nanosized silver colloids. *Nanotechnology*, 16:2436–2441, October 2005. doi: 10.1088/0957-4484/16/10/074.

Kerstin Lemke, Kai Schramm, and Christof Paar. DPA on n-Bit Sized Boolean and Arithmetic Operations and Its Application to IDEA, RC6, and the HMAC-Construction. In *Cryptographic Hardware and Embedded Systems CHES 2004*, volume 3156 of *Lecture Notes in Computer Science*, pages 205–219. Springer Berlin / Heidelberg, 2004. doi: 10.1007/b99451.

Huiyun Li, A. Theodore Markettos, and Simon Moore. Security Evaluation Against Electromagnetic Analysis at Design Time. In *Cryptographic Hardware and Embedded Systems CHES*

2005, volume 3659/2005 of *Lecture Notes in Computer Science*, pages 280–292. Springer Berlin / Heidelberg, 2005a. doi: 10.1007/11545262.

Huiyun Li, A.T. Markettos, and S. Moore. Security evaluation against electromagnetic analysis at design time. *High-Level Design Validation and Test Workshop, 2005. Tenth IEEE International*, 0:211–218, Nov.-2 Dec. 2005b. ISSN 1552-6674. doi: 10.1109/HLDVT.2005.1568839.

Jun Li, Lianhua Fan, and C.P. Wong. Influence of thermal treatment on the conductivity of PEDT films. *Advanced Packaging Materials: Processes, Properties and Interfaces, 2004. Proceedings. 9th International Symposium on*, 0:204–207, 2004. doi: 10.1109/ISAPM.2004.1288014.

Guirong Liang, Tianhong Cui, and Kody Varahramyan. Fabrication and electrical characteristics of polymer-based Schottky diode. *Solid-State Electronics*, 47(4):691–694, 2002. doi: 10.1016/S0038-1101(02)00324-6.

Jung Ah Lim, Jeong Ho Cho, Yeong Don Park, Do Hwan Kim, Minkyu Hwang, and Kilwon Cho. Solvent effect of inkjet printed source/drain electrodes on electrical properties of polymer thin-film transistors. *Applied Physics Letters*, 88(8):082102, 2006.

Yi Liu, Kody Varahramyan, and Tianhong Cui. Low-Voltage All-Polymer Field-Effect Transistor Fabricated Using an Inkjet Printing Technique. *Macromolecular Rapid Communications*, 26(24): 1955–1959, 2005. doi: 10.1002/marc.200500493.

Zhengchun Liu, Yi Su, and Kody Varahramyan. Electrical Hysteresis of PEDOT/PSS-Metal Contact Devices. *Materials Research Society Symposium Proceedings*, 814:253–257, 2004.

H. Lorenz, M. Despont, N. Fahrni, J. Brugger, P. Vettiger, and P. Renaud. High-aspect-ratio, ultrathick, negative-tone near-UV photoresist and its applications for MEMS. *Sensors and actuators. A, Physical*, 64:33–39, January 1998. doi: 10.1016/S0924-4247(98)80055-1.

E.J. Lous, P.W.M. Blom, L.W. Molenkamp, and D.M. de Leeuw. Schottky contacts on a highly doped organic semiconductor. *Phys. Rev. B*, 51(23):17251–17254, Jun 1995. doi: 10.1103/PhysRevB.51.17251.

Shengwen Luan and Gerold W. Neudeck. An experimental study of the source/drain parasitic resistance effects in amorphous silicon thin film transistors. *Journal of Applied Physics*, 72(2): 766–772, 1992.

M. F. Mabrook, C. Pearson, and M. C. Petty. An inkjet-printed chemical fuse. *Applied Physics Letters*, 86(1):013507, 2005. doi: 10.1063/1.1846950.

François Macé, François-Xavier Standaert, Jean-Jacques Quisquater, and Jean-Didier Legat. A Design Methodology for Secured ICs Using Dynamic Current Mode Logic. In *Integrated Circuit and System Design*, volume 3728/2005 of *Lecture Notes in Computer Science*, pages 550–560. Springer Berlin / Heidelberg, 2005. doi: 10.1007/11556930.

François Macé, François-Xavier Standaert, and Jean-Jacques Quisquater. Information Theoretic Evaluation of Side-Channel Resistant Logic Styles. In *Cryptographic Hardware and Embedded Systems CHES 2007*, volume 4727/2007 of *Lecture Notes in Computer Science*, pages 427–442. Springer Berlin / Heidelberg, 2007. doi: 10.1007/978-3-540-74735-2.

Hideyo Mamiya, Atsuko Miyaji, and Hiroaki Morimoto. Efficient Countermeasures against RPA, DPA, and SPA. In *Cryptographic Hardware and Embedded Systems CHES 2004*, Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2004. doi: 10.1007/b99451.

Stefan Mangard, Norbert Pramstaller, and Maria Elisabeth Oswald. Successfully Attacking Masked AES Hardware Implementations. In Josyula R. Rao and Berk Sunar, editors, *Cryptographic Hardware and Embedded Systems CHES 2005*, Lecture Notes in Computer Science (LNCS), pages 157 – 171. Springer, 2005.

Thomas S. Messerges. Using Second-Order Power Analysis to Attack DPA Resistant Software. In *CHES '00: Proceedings of the Second International Workshop on Cryptographic Hardware and Embedded Systems*, volume 1965/2000 of *Lecture Notes in Computer Science*, pages 238–251, London, UK, 2000. Springer Berlin / Heidelberg. ISBN 3-540-41455-X. doi: 10.1007/3-540-44499-8.

Thomas S. Messerges, Ezzy A. Dabbish, and Robert H. Sloan. Power Analysis Attacks of Modular Exponentiation in Smartcards. In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, pages 144–157, London, UK, 1999. Springer-Verlag. ISBN 3-540-66646-X.

Microchemicals. *Aluminium Etching*, Feb 2007. http://www.microchemicals.net/technical-information/ , Accessed 3^{rd} July 2007.

S. Moore, R. Anderson, P. Cunningham, R. Mullins, and G. Taylor. Improving smart card security using self-timed circuits. In *Asynchronous Circuits and Systems, 2002. Proceedings. Eighth International Symposium on*, pages 211–218, 8-11 April 2002.

S. Moore, R. Anderson, R. Mullins, G. Taylor, and J.J.A. Fournier. Balanced self-checking asynchronous logic for smart card applications. *Microprocessors and Microsystems*, 27:421–430(10), October 2003. doi: 10.1016/S0141-9331(03)00092-9.

P. V. Necliudov, M. S. Shur, D. J. Gundlach, and T. N. Jackson. Modeling of organic thin film transistors of different designs. *Journal of Applied Physics*, 88(11):6594–6597, 2000.

Peter V. Necliudov, Michael S. Shur, David J. Gundlach, and Thomas N. Jackson. Contact resistance extraction in pentacene thin film transistors. *Solid-State Electronics*, 47(2):259–262, February 2003.

T. P. Nguyen, P. Le Rendu, P. D. Long, and S. A. De Vos. Chemical and thermal treatment of PEDOT:PSS thin films for use in organic light emitting diodes. *Surface and Coatings Technology*, 180-181:646–649, 2004. Proceedings of Symposium G on Protective Coatings and Thin Films-03, of the E-MRS 2003 Spring Conference,.

Karsten Nohl, Starbug, and Henryk Plötz. Mifare Security. In *Presentation at the 24th Congress of the Chaos Computer Club in Berlin*, December 2007.

Karsten Nohl, David Evans, Starbug, and Henryk Plötz. Reverse-Engineering a Cryptographic RFID Tag. In *Proceedings of the 17^{th} USENIX Security Symposium*, pages 185–193, 2008.

K.J. Nowka and T. Galambos. Circuit design techniques for a gigahertz integer microprocessor. In *Computer Design: VLSI in Computers and Processors, 1998. ICCD '98. Proceedings. International Conference on*, pages 11–16, Oct 1998. doi: 10.1109/ICCD.1998.727017.

Elisabeth Oswald and Stefan Mangard. Template Attacks on Masking—Resistance Is Futile. In *Topics in Cryptology - CT-RSA 2007*, volume 4377/2006 of *Lecture Notes in Computer Science*, pages 243–256. Springer Berlin / Heidelberg, 2006. doi: 10.1007/11967668_16.

Maria Elisabeth Oswald and Kai Schramm. An Efficient Masking Scheme for AES Software Implementations. In *WISA 2005*, Lecture Notes in Computer Science, pages 292 – 305. Springer, 2006.

Maria Elisabeth Oswald, Stefan Mangard, Norbert Pramstaller, and Vincent Rijmen. A Side-Channel Analysis Resistant Description of the AES S-box. In *Fast Software Encryption*, volume 3557/2005 of *Lecture Notes in Computer Science*, pages 413 – 423. Springer, 2005. doi: 10.1007/b137506.

Maria Elisabeth Oswald, Stefan Mangard, Christoph Herbst, and Stefan Tillich. Practical Second-Order DPA Attacks for Masked Smart Card Implementations of Block Ciphers. In *Topics in Cryptology - CT-RSA 2006*, Lecture Notes in Computer Science, pages 192 – 207. Springer, 2006. doi: 10.1007/11605805_13.

Kateri E. Paul, William S. Wong, Steven E. Ready, and Robert A. Street. Additive jet printing of polymer thin-film transistors. *Applied Physics Letters*, 83(10):2070–2072, 2003. doi: 10.1063/1.1609233.

P. Paul. Non- volatile memory for wireless applications. Master's thesis, Cambridge University Engineering Department, 2005.

Siani Pearson. Trusted Computing Platforms, the Next Security Solution. Technical Report HPL-2002-221, HP Laboratories, 2002.

Jolke Perelaer, Antonius W. M. de Laat, Chris E. Hendriks, and Ulrich S. Schubert. Inkjet-printed silver tracks: low temperature curing and thermal stability investigation. *Journal of Materials Chemistry*, 18:3209 – 3215, 2008. doi: 10.1039/b720032c.

L. A. Plana, P. A. Riocreux, W. J. Bainbridge, A. Bardsley, S. Temple, J. D. Garside, and Z. C. Yu. SPA–a secure Amulet core for smartcard applications. *Microprocessors and Microsystems*, 27(9): 431 – 446, 2003. ISSN 0141-9331. doi: 10.1016/S0141-9331(03)00093-0.

L.A. Plana, P.A. Riocreux, W.J. Bainbridge, A. Bardsley, J.D. Garside, and S. Temple. SPA - a synthesisable Amulet core for smartcard applications. In *Asynchronous Circuits and Systems, 2002. Proceedings. Eighth International Symposium on*, pages 201–210, April 2002.

Thomas Popp and Stefan Mangard. Masked Dual-Rail Pre-Charge Logic: DPA-Resistance without Routing Constraints. In Josyula R. Rao and Berk Sunar, editors, *Cryptographic Hardware and Embedded Systems CHES 2005, 7th International Workshop, Edinburgh, Scotland, August 29 - September 1, 2005, Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 172–186. Springer, 2005.

N. Pramstaller, F.K. Gurkaynak, S. Haene, H. Kaeslin, N. Felber, and W. Fichtner. Towards an AES crypto-chip resistant to differential power analysis. *Solid-State Circuits Conference, 2004. ESSCIRC 2004. Proceeding of the 30th European*, 0:307–310, Sept. 2004. doi: 10.1109/ESSCIR.2004.1356679.

Jean-Jacques Quisquater and François Koeune. Side-channel Attacks: State-of-the-Art. Technical report, CRYPTREC project deliverable, http://www.ipa.go.jp/security/enc/-CRYPTREC/fy15/doc/1047_Side_Channel_report.pdf, accessed 20^{th} Sept. 2008, 2002.

Jean-Jacques Quisquater and David Samyde. ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In *E-SMART '01: Proceedings of the International Conference on Research in Smart Cards*, pages 200–210, London, UK, 2001. Springer-Verlag. ISBN 3-540-42610-8.

Jan M. Rabaey, Anantha Chandrakasan, and Borivoje Nikolic. *Digital Integrated Circuits*. Prentice Hall, 2 edition, 2003.

Wolfgang Rankl and Wolfgang Effing. *Smart Card Handbook*. John Wiley and Sons, 2004.

Christian Rechberger and Elisabeth Oswald. Practical Template Attacks. In *Lecture Notes in Computer Science*, volume 3325 of *Lecture Notes in Computer Science*, pages 440–456, Berlin / Heidelberg, February 2005. Springer. doi: 10.1007/b103174.

Henning Rost, Jurgen Ficker, Juan Sanchez Alonso, Luc Leenders, and Iain McCulloch. Air-stable all-polymer field-effect transistors with organic electrodes. *Synthetic Metals*, 145(1): 83–85, August 2004. doi: 10.1016/j.synthmet.2004.04.008.

A. Salleo, F. Endicott, and R. A. Street. Reversible and irreversible trapping at room temperature in poly(thiophene) thin-film transistors. *Applied Physics Letters*, 86(26):263505, 2005.

David Samyde, Sergei Skorobogatov, Ross Anderson, and Jean-Jacques Quisquater. On a New Way to Read Data from Memory. In *SISW '02: Proceedings of the First International IEEE Security in Storage Workshop*, page 65, Washington, DC, USA, 2002. IEEE Computer Society. ISBN 0-7695-1888-5.

H. Sanchez, R. Philip, J. Alvarez, and G. Gerosa. A CMOS Temperature Sensor For PowerPC RISC Microprocessors. *VLSI Circuits, 1997. Digest of Technical Papers., 1997 Symposium on*, 0: 13–14, Jun 1997.

W. Schindler, F. Koeune, and J. Quisquater. Unleashing the Full Power of Timing Attack. Technical Report CG-2001/3, Universite Catholique de Louvain – Crypto Group, 2001.

Werner Schindler. A Timing Attack against RSA with the Chinese Remainder Theorem. In *Cryptographic Hardware and Embedded Systems CHES 2000*, volume 1965/2000 of *Lecture Notes in Computer Science*, pages 109–124. Springer Berlin / Heidelberg, 2000. doi: 10.1007/3-540-44499-8.

Hermann Schlichting and Klaus Gersten. *Boundary-layer theory*. Springer, 2000.

Bruce Schneier. Why Cryptography is Harder than it Looks. *Information Security Bulletin*, 2(2): 31–36, 1997.

Bruce Schneier and Adam Shostack. Breaking Up Is Hard to Do: Modeling Security Threats for Smart Cards. In *USENIX Workshop on Smartcard Technology*, pages 175–185, May 1999.

Markus Schwoerer and Hans Christoph Wolf. *Organic Molecular Solids*. Wiley-VCH, 2007.

Christoph Sele, Timothy von Werne, Richard Friend, and Henning Sirringhaus. Lithography-Free, Self-Aligned Inkjet Printing with Sub-Hundred-Nanometer Resolution. *Advanced Materials*, 17(8):997–1001, April 2005.

G. E. Servais and S. D. Brandenburg. Wire bonding - a closer look. In *Proceedings ASM International ISTFA 91*, pages 525–529, 1991.

Adi Shamir. Protecting Smart Cards from Passive Power Analysis with Detached Power Supplies. In *CHES '00: Proceedings of the Second International Workshop on Cryptographic Hardware and Embedded Systems*, pages 71–77, London, UK, 2000. Springer-Verlag. ISBN 3-540-41455-X.

H. Sirringhaus, T. Kawase, R. H. Friend, T. Shimoda, M. Inbasekaran, W. Wu, and E. P. Woo. High-Resolution Inkjet Printing of All-Polymer Transistor Circuits. *Science*, 290:2123–2126, December 2000.

Henning Sirringhaus. Device Physics of Solution-Processed Organic Field-Effect Transistors. *Advanced Materials*, 17(20):2411–2425, 2005. doi: 10.1002/adma.200501152.

Sergei P. Skorobogatov and Ross J. Anderson. Optical Fault Induction Attacks. *Cryptographic Hardware and Embedded Systems CHES 2002: 4th International Workshop Redwood Shores, CA, USA, August 13-15, 20. Revised Papers*, 2523:31 – 48, 2003.

Sean Smith and Steve Weingart. Building a High-Performance, Programmable Secure Coprocessor. Technical Report RC21102, IBM, 1998.

D. Sokolov, J. Murphy, A. Bystrov, and A. Yakovlev. Design and analysis of dual-rail circuits for security applications. *Computers, IEEE Transactions on*, 54(4):449–460, April 2005. doi: 10.1109/TC.2005.61.

Danil Sokolov, Julian Murphy, Alexandre V. Bystrov, and Alexandre Yakovlev. Improving the Security of Dual-Rail Circuits. In Marc Joye and Jean-Jacques Quisquater, editors, *CHES*, volume 3156 of *Lecture Notes in Computer Science*, pages 282–297. Springer, 2004. ISBN 3-540-22666-4.

Stuart P. Speakman, Gregor G. Rozenberg, Kim J. Clay, William I. Milne, Adelina Ille, Ian A. Gardner, Eric Bresler, and Joachim H. G. Steinke. High performance organic semiconducting thin films: Ink jet printed polythiophene [rr-P3HT]. *Organic Electronics*, 2(2):65–73, September 2001.

François-Xavier Standaert, Tal G. Malkin, and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. Technical Report 2006/139, Cryptology ePrint Archive, 2006a. http://eprint.iacr.org.

François-Xavier Standaert, E. Peeters, C. Archambeau, and J. J. Quisquater. Towards Security Limits in Side-Channel Attacks (With an Application to Block Ciphers). In *Cryptographic Hardware and Embedded Systems CHES 2006*, volume 4249/2006, pages 30–45. Springer Berlin / Heidelberg, 2006b. doi: 10.1007/11894063.

H.C. Starck. Clevios FAQs, http://www.clevios.com/index.php?page_id=972#anker20, a. accessed 2^{nd} Sept. 2008.

H.C. Starck. Clevios FAQs, http://www.clevios.com/index.php?page_id=972#anker17, b. accessed 2^{nd} Sept. 2008.

H.C. Starck. Clevios FAQs, http://www.clevios.com/index.php?page_id=972#anker4, c. accessed 2^{nd} Sept. 2008.

H.C. Starck. Clevios Chemical Properties Presentation, http://www.clevios.com/pages/985/-chemical_properties_clevios_20080314.pdf, d. Version 14^{th} March 2008, accessed 2^{nd} Sept. 2008.

R. A. Street and A. Salleo. Contact effects in polymer transistors. *Applied Physics Letters*, 81(15): 2887–2889, 2002.

Ben G. Streetman and Sanjay Banerjee. *Solid State Electronic Devices*. Prentice Hall, 5 edition, 2000. ISBN 0-13-026101-7.

Natalie Stutzmann, Richard H. Friend, and Henning Sirringhaus. Self-Aligned, Vertical-Channel, Polymer Field-Effect Transistors. *Science*, 299(5614):1881–1884, 2003. doi: 10.1126/science.1081279.

H. K. Tian, J. W. Shi, D. H. Yan, L. X. Wang, Y. H. Geng, and F. S. Wang. Naphthyl End-Capped Quarterthiophene: A Simple Organic Semiconductor with High Mobility and Air Stability. *advanced materials*, 18(16):2149 – 2152, 2007. doi: 10.1002/adma.200600178.

K. Tiri and I. Verbauwhede. Charge recycling sense amplifier based logic: securing low power security ICs against DPA [differential power analysis]. In *Solid-State Circuits Conference, 2004. ESSCIRC 2004. Proceeding of the 30th European*, pages 179–182, 21-23 Sept. 2004a. doi: 10.1109/ESSCIR.2004.1356647.

K. Tiri, M. Akmal, and I. Verbauwhede. A dynamic and differential CMOS logic with signal independent power consumption to withstand differential power analysis on smart cards. In *Solid-State Circuits Conference, 2002. ESSCIRC 2002. Proceedings of the 28th European*, pages 403–406, 24-26 Sept. 2002.

Kris Tiri. Side-channel attack pitfalls. In *DAC '07: Proceedings of the 44th annual conference on Design automation*, pages 15–20, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-627-1. doi: 10.1145/1278480.1278485.

Kris Tiri and Ingrid Verbauwhede. Place and Route for Secure Standard Cell Design. In *CARDIS*, pages 143–158, 2004b.

Elena Trichina, Tymur Korkishko, and Kyung Hee Lee. *Advanced Encryption Standard - AES*, volume 3373/2005, chapter Small Size, Low Power, Side Channel-Immune AES Coprocessor: Design and Synthesis Results, pages 113–127. Springer Berlin / Heidelberg, 2005. doi: 10.1007/b137765.

Abdulmecit Turut and Fatih Koleli. Semiconductive polymer-based Schottky diode. *Journal of Applied Physics*, 72(2):818–819, 1992. doi: 10.1063/1.351822.

Thijs H. J. van Osch, Jolke Perelaer, Antonius W. M. de Laat, and Ulrich S. Schubert. Inkjet Printing of Narrow Conductive Tracks on Untreated Polymeric Substrates. *Advanced materials*, 20(2):343–345, 2008. doi: 10.1002/adma.200701876.

Peter Van Zant. *Microchip Fabrication*. McGraw-Hill, 5 edition, 2004. ISBN 0-07-143241-8.

Jason Waddle and David Wagner. Towards Efficient Second-Order Power Analysis. In *Cryptographic Hardware and Embedded Systems CHES 2004*, volume 3156/2004 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin / Heidelberg, 2004. doi: 10.1007/b99451.

Chua-Chin Wang, Chi-Feng Wu, Rain-Ted Hwang, and Chia-Hsiung Kao. A low-power and high-speed dynamic PLA circuit configuration for single-clock CMOS. *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on [see also Circuits and Systems I: Regular Papers, IEEE Transactions on]*, 46(7):857–861, July 1999. doi: 10.1109/81.774233.

Liang Wang, Daniel Fine, Deepak Sharma, Luisa Torsi, and A. Dodabalapur. Nanoscale organic and polymeric field-effect transistors as chemical sensors. *Analytical and Bioanalytical Chemistry*, 384(2):310–321, 2005. doi: 10.1007/s00216-005-0150-2.

Steve H. Weingart. Physical Security Devices for Computer Subsystems: A Survey of Attacks and Defenses. *Lecture Notes in Computer Science*, 1965:302–317, January 2000.

N. Weste and D. Harris. *CMOS VLSI Design*. Addison Wesley, May 2004. ISBN 0321149017.

I. Winter, C. Reese, J. Hormes, G. Heywang, and F. Jonas. The thermal ageing of poly(3,4-ethylenedioxythiophene). An investigation by X-ray absorption and X-ray photoelectron spectroscopy. *Chemical Physics*, 194:207–213, May 1995. doi: 10.1016/0301-0104(95)00026-K.

F. Xue, Y. Su, and K. Varahramyan. Modified PEDOT-PSS Conducting Polymer as S/D Electrodes for Device Performance Enhancement of P3HT TFTs. *IEEE Transactions on Electron Devices*, 52:1982–1987, September 2005. doi: 10.1109/TED.2005.855062.

Sung-Ming Yen, Seungjoo Kim, Seongan Lim, and Sangjae Moon. A Countermeasure against One Physical Cryptanalysis May Benefit Another Attack. *Lecture Notes in Computer Science*, 2288:414–427, 2002.

Z.C. Yu, S.B. Furber, and L.A. Plana. An investigation into the security of self-timed circuits. In *Asynchronous Circuits and Systems, 2003. Proceedings. Ninth International Symposium on*, pages 206–215, 12-15 May 2003. doi: 10.1109/ASYNC.2003.1199180.

J. Zaumseil, K. W. Baldwin, and J. A. Rogers. Contact resistance in organic transistors that use source and drain electrodes formed by soft contact lamination. *Journal of Applied Physics*, 93 (1):6117–6124, 2003.