Verifying the Authenticity of Chip Designs with the DesignTag System

Tom Kean, David McLaren and Carol Marsh
Algotronix Ltd.
Edinburgh, United Kingdom
tom@algotronix.com, david@algotronix.co.uk, carol@algotronix.co.uk

***Abstract*—This paper introduces DesignTag – a novel, patented, 'security tag' technology which can be used to verify the authenticity of semiconductor devices. The tag takes the form of a small digital circuit which is added to the chip design and communicates through the package with an external sensor. Falsely marked 'ghost' chips are present in the supply chain and cause economic damage to reputable semiconductor companies. They can also constitute a safety hazard in critical applications and act as a vector for inserting malicious 'Trojan' functionality into a secure system such as banking or government communications. DesignTag can also be used to address related threats such as copying of chip designs and Intellectual Property cores and unlicensed use of CAD tools.**

Keywords-counterfeit chips; IP tagging; security

I. INTRODUCTION

There are many electronic systems which require a high level of confidence that the chips being used within them are authentic. These include safety critical applications such as aerospace, automotive and medical, where failure of a device can have serious consequences. As well as compromising the end application, falsely marked chips threaten the revenue and reputation of reputable semiconductor companies whose name is misused.

There are two primary motivations for creating falsely labeled chips and injecting them into the supply chain. The first is straightforward economic gain from selling the falsely labeled chips at a higher price than they would otherwise command. Falsely marked chips may be cheaper copies of a brand name product, test failures recovered from scrap bins, chips recycled from scrapped equipment or even empty packages. More subtly, a 'bona-fide' chip may have speed or temperature grade information falsely marked to increase the product value [1]. In the case of memories 'near equivalent' devices with the same pinout may be remarked as a more valuable product [2].

The second motivation for creating falsely labeled chips is more malicious: falsely marked chips can be used to insert 'Trojan Horse' functionality into a secure system. In this case a device is created which appears to operate identically to a genuine device, in order to intentionally disable a system at some point in the future, or to alter the system behavior in some way for their own purposes. For example, cryptographic keys might be intentionally leaked or financial information stolen or modified. Standard microprocessors which are also available as Intellectual Property (IP) cores are a particularly attractive target for this form of attack. An attacker can purchase the standard microprocessor as an IP core, make a few changes to insert their Trojan Horse and then manufacture a chip which can be passed off as a standard microcontroller. This attack was recently demonstrated on the Leon processor [3]. Memory chips and modules are another potential vector for inserting Trojan functionality into a system – a less sophisticated version of this threat has already occurred with a Chinese contract manufacturer being blamed for the presence of a Windows virus on some newly manufactured video iPods [4].

When carried out with a modicum of care, false labeling of chips can be nearly impossible to detect. For example, a chip with a commercial temperature grade falsely labeled as a military component may perform perfectly until subjected to extreme cold. Chips recycled from scrapped equipment which then have their manufacturing date codes tampered could pass initial test but cause a higher than expected failure rate in deployed equipment containing them several years in the future. In many cases, the only technique for proving that a chip was falsely labeled would be to remove it from the equipment, depackage and examine it under a microscope. Because of the difficulty of detecting the crime it is likely that the reports of falsely labeled chips in the supply chain represent the tip of the iceberg.

The wholesale shift of electronics manufacturing to Asian countries, and China in particular, has exacerbated the problem in several ways. Firstly, contract manufacturers typically operating on wafer thin margins are highly motivated to secure the lowest possible prices for the components they use. With a 2% profit margin a reduction of 1% on the bill of materials is extremely attractive. Secondly, legal protections on intellectual property are less developed in China and there is a perception that it is fair to 'catch up' with the west by copying products. As one reverse engineering company claims on its website, "It will be because of the companies like ours that the gap in technological advantages between made in China and made abroad will narrow down" [5]. Product copying is of concern in this context because a Trojan Horse attack involves both product copying and false labeling. Thirdly, inserting malicious functionality into communications and computing systems is a known method for intelligence gathering or compromising of vital information infrastructure. During the cold war 'Eastern Block' countries were potentially vulnerable to intelligence gathering because of their dependence on western information

technology – but today the tables have turned and most information technology is manufactured in China.

The true scale of the problem is hard to determine because of the difficulty of detection and the natural reluctance of affected parties to speak out: however press reports and publicly available information show it is both real and serious. In January 2008, two individuals faced federal charges in the USA for selling counterfeit goods to a number of high profile customers, including the Air Force, Marine Corps, Department of Energy and defense contractors such as Lockheed Martin [6]. Chinese and Lithuanian companies that provide a range of product copying services to Chinese manufacturers advertise openly on the internet [5],[7],[8]. English language product listing websites advertise a huge range of Chinese manufactured equipment to distributors in the west – including blatant copies of brand name equipment such as Nokia cellular phones [9],[10]. The US Department of Defense has identified the possibility of trojan functionality in chips used in military systems but manufactured abroad as an important concern [11].

To reduce the risks while retaining the cost benefits of outsourced electronics manufacturing, a simple, easy to use, technical means of verifying whether a chip is genuine is highly desirable.

Algotronix' DesignTag™ product presented here provides a method for uniquely identifying intellectual property within a chip, enabling the user to be confident that the chip or IP block is genuine. An individual IP block or an entire chip design is tagged with a code unique to that design, and this code is used to identify the design after placing a sensor in contact with the chip package. Unlike ink markings on the chip package the information from the active tag is only available to authorized users and is very difficult to tamper with.

DesignTag is an active circuit present in the finished chip which can be detected during operation, not an optical identification code on special mask layers or a watermark in the design source code [12]. Moreover, communication with the DesignTag is through a sensor placed on the chip package, not through signal pins. With modern chip packages such as Ball Grid Arrays it is difficult to directly probe signal pins on the package. Sensing the tag through the package means that accessing DesignTag is independent of the Printed Circuit Board layout and software within the system. This has two crucial advantages – firstly, a chip will be designed into many different systems and it would be cumbersome if access to the tag depended on factors specific to the system it was designed into. Secondly, in many usage scenarios the system containing the chip cannot be trusted – if the tag code was reported through system software it would be easy to falsify.

Verifying the provenance of an integrated circuit depends on checking that the correct tag is present within a chip which is suspected of being a fake. The inverse situation, where a chip has a tag signal from a reputable company but is not marked as such can indicate a 'pirated' design where the intellectual property has been copied without permission. Detecting this situation is of particular value to Field Programmable Gate Array (FPGA) designers (leading low cost FPGA families do not have bitstream encryption and bitstreams can be easily copied), IP core vendors and CAD tool companies. For example, a CAD tool company might program an educational or evaluation copy of their tools to insert a tag into the design being processed as a means of detecting if the tools were used to create a commercial product.

II. Description of DesignTag System

A. Overview

The DesignTag system consists of four components: a small 'tag' circuit incorporated within the IP core or chip design to be protected, hardware to collect data from the tagged chip, software which processes the collected data, and a database containing tag codes and design information. DesignTag scanners are low cost and can be widely deployed. In one usage scenario a semiconductor company requests a DesignTag scan of one of their chips which a customer claims is faulty before issuing a Returned Material Authorization. As well as confirming authenticity DesignTag can also provide a status code from on-chip self test circuits. In another use case a company which has contracted manufacturing out to a low cost supplier carries out a DesignTag scan of chips in a randomly selected sample of the products to determine if the bill of materials has been followed exactly. In a third use case an importer could use a DesignTag scan to identify whether there is stolen intellectual property in a 'no-name' electronics product they are considering distributing.

B. Communications Channel

In the DesignTag application the goal is to intentionally create a covert 'side channel' between the tag circuit on the chip and an external sensor. Side channels have been studied in the context of cryptographic hardware, in which there is a need to protect against so-called 'side channel attacks', where operation of the cryptographic circuit causes unintentional signals, for example fluctuations in power supply voltages or electromagnetic radiation which can be detected off-chip and provide information about the cryptographic key. Unlike side channel attacks in cryptography, this is an intentional transmission and the modulation scheme is within the control of the tag designer making the detection task considerably easier.

To enable communications between an on-chip tag circuit and an external sensor, several potential side channels are available: in the cryptographic field, timing (changes in the delay between observable events), power supply voltage variations and electromagnetic emanations have all been extensively studied [13],[14],[15]. The authors have also studied and developed intellectual property on the use of these channels in security tags. However, for the initial DesignTag product it was decided to take the novel approach of signaling via a thermal channel, through changes in the chip package temperature.

Signaling with heat, while clearly possible in theory, has the important practical drawback that the data rate is low because there are physical limitations on how fast a chip package can heat up and how quickly a temperature measurement can be made. This makes it unsuitable for most applications, however in the context of a security tag only a

very small amount of information needs to be transferred and it would be quite acceptable to take a few minutes to do so – this is much faster than the alternative of extracting the chip from the system and sending it to a laboratory for analysis. It is also much faster than methods requiring electrical contact when the time taken to locate suitable probe-points on the circuit board and connect probes is taken into account. Unlike JTAG, DesignTag does not require additional package pins and is easily used during normal operation of the equipment rather than in a special test mode. Moreover, DesignTag can be deployed in updated versions of existing chip products without changing their pinout.

A practical advantage of signaling with heat in the DesignTag application is that it can be transmitted (and detected) using only digital circuits. Chip packages are also necessarily designed to transfer heat but may contain metallic components such as heat spreaders that would attenuate wireless signals. With a thermal signal it is immediately clear which chip is being measured, by contrast the source of a tag signal transferred using noise on the the power supply wiring is harder to localize with certainty.

C. Tag Circuit

The tag circuit must have a minimal impact on the design it protects in terms of power consumption and area. Large circuits will create an unacceptable area overhead. Structures typically found in radio frequency circuits such as antennas and inductors are undesirable in tags used for Intellectual Property protection since they are easy to visually identify and disable.

Each tag is a small digital circuit which generates a sequence of chip temperature changes corresponding to a unique binary code assigned to that particular tag. Fig. 1 gives an overview of the tag circuit. A heat source generates heat according to a sequence determined by the unique tag code output by the 'code generator' block, with timing information obtained from the 'timing' block. The circuit is protected by a 'Shell' providing anti-tamper and reverse engineering countermeasures. These countermeasures are not being disclosed publicly.

In the prototype tag circuit design, a number of parallel ring oscillators were used as the mechanism for heat generation: more subtle structures are also available. A binary '1' corresponds to turning the heat generation circuit on, i.e. enabling the ring oscillators which operate at a high frequency and generate heat. The heat generation circuit remains on for a length of time corresponding to the bit period. A binary '0' corresponds to turning the heat generation circuit off for one bit period. The temperature differences generated are small (less than 0.1°C on the chip package), and therefore the power consumed by the heat generation circuit is low.

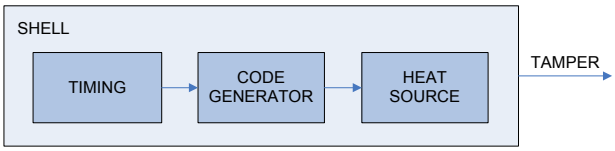

Figure 1. Tag Circuit

D. Tag Codes

Each DesignTag transmits a unique spread-spectrum 'pseudo-noise' temperature signal. The spread-spectrum codes are similar to those used by Code Division Multiple Access (CDMA) systems for wireless communications, in which each user is assigned a code sequence which is used to modulate ('spread') data transmitted by that user. When multiple CDMA users transmit over the same channel, a single user's data can be retrieved by correlating the received signal with the user's unique code, provided that the codes used have low cross-correlations. Equation (1) gives the cross-correlation between two discrete sequences, x and y, for sequence length N and offset k.

$$r_{xy}(k) = \frac{1}{N}\sum_{i=0}^{N} x[i]y[i-k] \qquad (1)$$

Unlike CDMA, where data bits from a message are spread by the (short) spreading code and it is necessary to recover each data bit at the receiver, the aim in this application is to simply detect the presence of a tag signal and determine which spreading code has been used. It is therefore not necessary to use short codes with very low cross-correlations – long codes can be used and codes for two different tags may have quite high cross-correlations over particular code sections. If these codes are pseudo-random then over a long enough period the cross-correlations between codes (1), if stored as bipolar binary (+1, -1) values, will approach zero. These long codes allow the processing software to detect the tags. While it is not feasible (or necessary in this application) to generate a large number of codes with very low cross-correlations, the codes used are selected such that no two codes have particularly high cross-correlations over a specified range of code offsets (k) and sequence lengths (N). If two codes were similar, there is a risk that during detection a tag could be falsely identified as being in a chip, due to the presence of a tag with a similar code.

Each tag is uniquely identified by the processing software through its pseudo-noise spreading code. The tag circuit uses a linear feedback shift register (LFSR) or stream cipher to generate the spreading code based on a unique and secret 'tag code' which is stored within the tag. In order to detect test failures or chips with mislabeled speed grades the secret tag code is written to non-volatile memory within the tag following testing and speed grading. When the code register is not written the tag outputs a special 'test failure' code.

The amplitude of the temperature changes on the chip package caused by DesignTag is small and the signal is 'buried' in thermal noise from other circuits on the chip and environmental factors (for example drafts of air caused by people moving) – this makes determining whether a tag is present within a particular chip very difficult for anyone who does not know the tag code (and hence the spreading code).

E. Data Collection Hardware

The sensor used to measure chip temperature data is a basic thermocouple attached to the chip package with tape. To collect the data measured by the thermocouple, a low-cost data-

logging unit is used which connects directly with the software to process the data while sampling in real time. For evaluation and where tags are used infrequently, common test equipment such as a high precision digital multimeter with data logging capability can be used to collect the data, and this data can then be passed to the DesignTag software to check for the presence of tags.

F. Tag Detection Software

The signal processing performed by the software correlates data derived from the chip package temperature with each of the tag codes stored in a secure web-based database. As more samples are acquired, the correlation between the signal and a code used by a tag present in the chip will increase. Because of high levels of thermal noise, the magnitude of the correlation is likely to be small – but detectably higher than the correlation for non-matching codes which should tend to zero as the number of samples increases. The software can then display which tags have been detected and provide access to further information such as datasheets through the web-based database. IP owners can specify whether their tags are private, or are publically available and can therefore be detected by anyone with access to the public database. Fig. 2 shows data input to, and generated by, the signal processing software – Fig. 2A shows temperature data sampled by a thermocouple; Fig. 2B gives the data derived from the temperature data (the software is interested in whether temperature is rising or falling, not its absolute value). This processed data is then correlated, for a range of code offsets, with each code in the database – a section of one such code is given in Fig. 2C. Figs. 2D and 2E show the resulting correlation values for a non-matching code and matching code, respectively. A peak can clearly be seen in Fig. 2E, indicating a matching tag code. The magnitude of this peak will increase with time as more samples are collected, increasing confidence that the tag is present. Statistical techniques are used to determine a confidence value for the match and when the match reaches a specified threshold, detection of the corresponding tag is announced to the user.

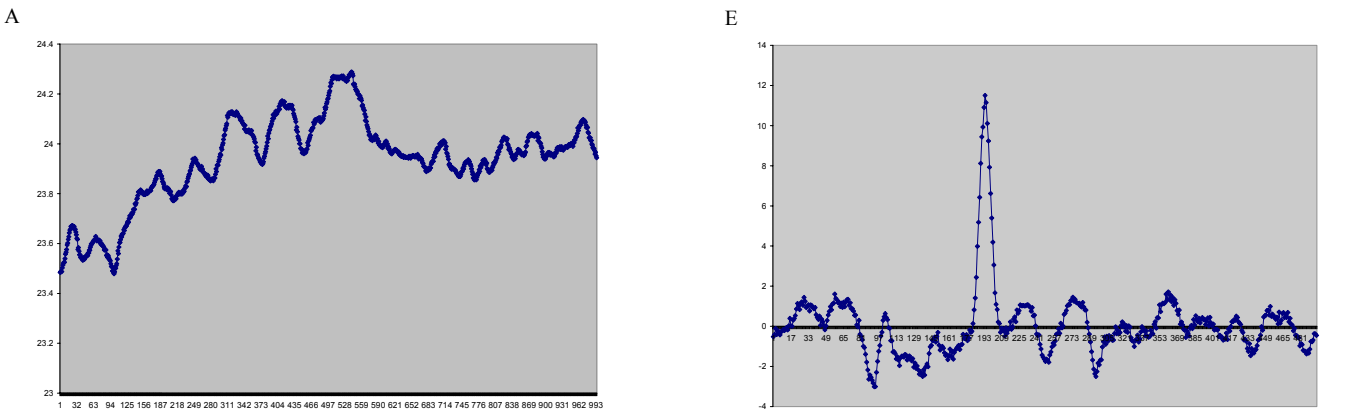

Figure 2. A) Temperature data sampled by sensor. B) Data derived from temperature data. C) Example tag code. D) Correlation data for non-matching code. E) Correlation data for matching code.

G. Tag Area and Power Requirements

Implementation details for the initial implementation of DesignTag on a Spartan 3A FPGA from Xilinx are presented in Table I. As can be seen the area of the circuit is very small, power consumption is low and including this tag would be a marginal cost when incorporated in a large IP core or complete chip design. The tag circuit can be configured to switch off after a fixed delay following power on, to reduce the the energy consumption impact on the protected design. In this case power must be cycled prior to detecting the tag but this is not usually a problem.

TABLE I. TAG IMPLEMENTATION DETAILS

Chip	Slices	RAM Blocks	Average Power Consumption
Spartan 3A	152	0	5mW

III. EXPERIMENTAL RESULTS

The experimental setup used for initial evaluation of DesignTag is shown in Fig. 3. A design to be tagged is implemented on a Xilinx Spartan XC3S700A chip on a Xilinx Spartan 3A evaluation board. The design used for this experiment is a Xilinx demonstration design provided with the evaluation board. It can be viewed as a typical System on Chip (SoC) design and makes use of several large IP blocks including a PicoBlaze soft core processor, VGA driver and audio driver to display various advertising messages on a VGA screen and via audio under the control of switches on the evaluation board. DesignTags were added to four of the main functional blocks of the design and the complete chip design was identified with an additional tag.

A Type-K thermocouple is attached to the top of the FPGA chip package, and temperature data is transferred to the tag detection software through an off-the-shelf data logging unit. The software runs on a laptop computer which connects to the data logging unit through a USB port. The software is then run for a specified period of time and any detected tags are displayed. The software searches through a database of 1000 different tag codes (including the five codes used by the on-chip tags).

The time taken to detect the five tags varies depending on which tag codes are used and temperature changes in the chip's environment, however the average time to successfully detect all five tags was found to be less than ten minutes. Tag detection takes place while the main circuit operates normally.

For the same design with a single tag, the time for detection was significantly reduced, at less than five minutes. The authors are continuing to optimize the DesignTag implementation and expect that they will be able to further improve the performance characteristics reported here.

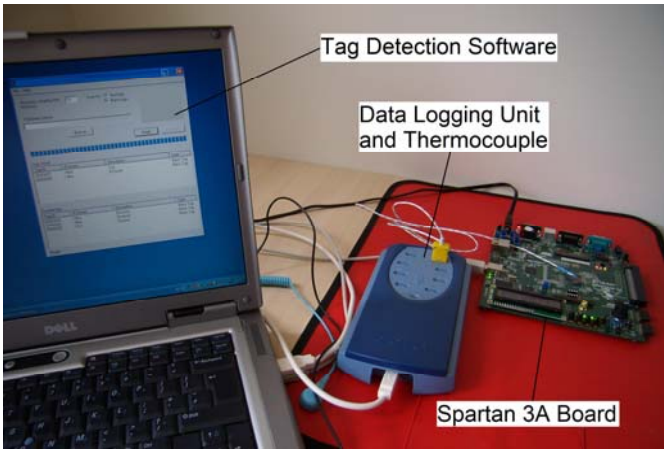

Figure 3. Experimental Setup

IV. SUMMARY

A novel 'active tag' technology has been proposed and developed. Using this technology, users of integrated circuits can determine whether the design they are using is genuine, or an inferior replica potentially with hidden, undesirable functionality. DesignTag can also be used to address Intellectual Property misuse scenarios such as overbuilding by licensed customers, bitstream piracy and misuse of CAD tool licenses. Other applications include enforcing export control and patent license restrictions, reporting status codes from on-chip test circuits and determining design version information for IP core vendors and FPGA users who cannot directly mark the chip package.

The proposed active tag is a very small and low power circuit which can be added to chip or IP core designs and detected using an external sensor. Tag detection is achieved using a thermal scheme allowing the tag to be built using only digital components and making it suitable for use with FPGAs as well as ASICs.

REFERENCES

[1] M. Singer. "Fake Chips Shadow AMD's New Alchemy". InternetNews.com. Jan. 3, 2005. [Online]. Available: http://www.internetnews.com/bus-news/article.php/3453541

[2] G.A. Quirk. "Under the Hood Special Report: Counterfeit parts, legitimate woes". EE Times. June 8, 2007. [Online]. Available: http://www.eetimes.com/news/latest/showArticle.jhtml?articleID=201201791

[3] R. McMillan. (2008, Apr. 16). "Researchers Uncover Undetectable Chip Hack". IDG News Service. [Online]. Available: http://www.techworld.com/security/news/index.cfm?newsid=11993

[4] Apple Inc. "Small Number of Video iPods Shipped With Windows Virus". Apple Inc. [Online]. Available: http://www.apple.com/support/windowsvirus

[5] DragonMen PCBWork Room. "Decipher Brief Introduction". DragonMen PCBWorkRoom. [Online]. Available: http://www.pcbwork.net/en/chip.asp

[6] G. Derene and J. Pappalardo. "Counterfeit Chinese Chips Raise Big Hacking and Terror Threats". Popular Mechanics. April 2008 [Online]. Available: http://www.popularmechanics.com/technology/industry/4253628.html

[7] Semiconductors Research. "Semiconductors Research". [Online]. Available: http://www.semiresearch.com/

[8] Beijing Diamond Institution of Computer Application. “Trade Leads”. [Online]. Available: http://en.pcbtn.com/Product/View_Trade_Leads.asp?id=42756

[9] Alibaba.com. “Alibaba Manufacturer Directory”. [Online] Available: www.alibaba.com

[10] MadeInChina.com. “China Manufacturing Manufacturers in China Product Sourcing Suppliers Export Import”. [Online] Available: www.madeinchina.com

[11] Defense Science Board. Defense Science Board Task Force on High Performance Microchip Supply. February 2005.

[12] A.B. Kahng et al. “Watermarking techniques for Intellectual Property protection”, 1998 ACM/IEEE Design Automation Conference Proceedings, June 1998, pp 776-781.

[13] Kocher Paul, “Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems”, Proceedings of Crypto’96, Springer-Verlag, August 1996, pp. 104–113.

[14] Kocher, Jaffe and Jun, “Differential power analysis”, Proceedings of Crypto’99, Springer-Verlag, 1999, pp. 388-397

[15] Gandolfi, Mourtel and Olivier, “Electromagnetic analysis: concrete results”, Proceedings of CHES’01 Springer-Verlag, 2001, pp, 251-261

[16] Algotronix Ltd, “Method and apparatus for secure configuration of a Field Programmable Gate Array”, US Patent Application US2001/0015919

[17] Algotronix Ltd, “Method of using a mask programmed key to securely configure a Field Programmable Gate Array”, US Patent Application US2001/0037458

[18] Algotronix Ltd, “Method of protecting Intellectual Property cores on Field Programmable Gate Array”, US Patent US2002/0199110