Darknets, DRM, and Trusted Computing: Economic Incentives for Platform Providers

Alessandro Acquisti*
Carnegie Mellon University

PRE-CONFERENCE, PRELIMINARY DRAFT
PLEASE DO NOT LINK/CIRCULATE/QUOTE
FINAL DRAFT TO BE PRESENTED AT TPRC 2004

Abstract

We discuss the incentives for platform and content providers to enforce digital rights management (DRM) through trusted computing (TC) initiatives in markets where consumers can choose between different platforms. Under what conditions does it make sense for platform vendors to second content providers' requests to protect the latter's content? Under what conditions will the market converge towards one dominant platform? What would the consequences of a trusted computing platform monopoly be for consumer welfare? In our preliminary analysis we discuss how platform providers' optimal decisions rely on a series of factor including the network effects associated with consumer-generated content. Even more than widely popular, high-demand content, the aggregate impact of low-demand individual content and the individual costs of platform adoption can determine the success or failure of trusted computing initiatives.

Keywords: Trusted Computing, Economics, Privacy, Security.

*Email: acquisti@andrew.cmu.edu

1 Introduction

In recent years, content providers and the computer industry have joined forces to protect intellectual property rights on digital information. Music publishers and movie studios have experimented with various digital rights management (DRM) initiatives in order to control access, use, and dissemination of their products. Software and hardware providers have developed various technological protections of such DRM. Those protections have been routinely broken, and the property rights they were supposed to defend have been bypassed by peer-to-peer file sharing networks. Biddle, England, Peinado, and Willman (2003) have coined the term "darknets" to refer to such networks, concluding that there exist no technical impediments to their growing in convenience and efficiency.

Advancements in computer security research such as Arbaugh, Farber, and Smith (1997), however, have recently made the prospect of building difficult to break DRM systems actually plausible. In 1999, Intel, Microsoft, HP, Compaq, and IBM founded the Trusted Computing Platform Alliance (TCPA), with the goal of creating open industry standards for a trusted computing subsystem to be added to personal computers. "Trusted computing" (TC) refers to standards (and the combination of hardware and software built upon them) aimed at increasing computer security of personal computers by letting certain components verify the trustworthiness of others before interacting with them, or allowing them to run on a machine. Although DRM is not explicitly mentioned in the TCPA specifications (Trusted Computing Group [2003]), one possible and natural applications of trusted computing is the enforcement of digital rights management (Anderson [2003a,b]). TC can be used as a building block for an architecture of protected computing environments in which every application, file, or communication needs approval in order to run on a system.

In order for DRM to be deployed effectively through trusted computing,

however, both software and hardware platform providers need to modify their products according to the TCPA specifications. Although the TCPA has grown to include most of the dominant players in the computer industry, almost ubiquitous adoption is instrumental to the full success of the initiative. Yet, compliance is costly. Recently, Microsoft announced a delay in the deployment of its TC initiatives due to the concerns of enterprise and independent application developers associated with the costs of making their products compatible with the specifications (Evers [2004]).

Because of vendors' costs and users' reaction, historically most hardware based DRM systems (from key lock-floppies to dongles) have been competed away in the marketplace. Under the opposite pulls of network effects (associated with non-DRM systems, which are preferred by users) and rights control (associated with DRM systems and preferred by content providers), most DMR-platforms have failed to gain widespread adoption, their producers eventually releasing their products in non-DRM forms. Whether things will go differently under a TC architecture may depend on the interaction of two dynamics. On the one hand, consumers have started using encryption, authentication, and social networks to create rings of trustd peers to exchange files out of the legal and technical reach of DRM enforcement.[1] The efficiency of these networks would keep consumers away from TC systems that may make those networks impossible. On the other hand, the size and strength of the alliance between key members of the computer industry may leave consumers no choice.

According to its proponents, TC would guarantee more information security and privacy than we have now. According to its detractors, TC would make it possible to transfer significant control of a computer away from its user and owner, to platform and content vendors. Artificially created lock-

[1]Some have started applying the term "darknets" in this particular sense: see Boutin (2004). Some have also noted that trusted computing itself may be used to trade copyrighted material untraceably: see Schechter, Greenstadt, and Smith (2003).

ins, incompatibilities, barriers to entry, and other obstacles to parties outside the TC alliance are described not only as possible, but as the actual business goals of the TC alliance (Green [2002]).

In this paper, we study the incentives for platform providers to produce TC compliant, DRM enforcing products in a market in which consumers can choose between "trusted" computing platforms and "open" platforms that do not enforce TC or DRM and on which darknets can thrive. Under what conditions does it make sense for platform providers to second content providers' requests to protect the latter's content? Under what conditions will the market converge towards one dominant platform? What would the consequences of a trusted computing platform monopoly be for consumer welfare?

The answers we seek are not based on the letter of current TC initiatives and TCPA statements. Instead, we focus on the likely business developments that will follow the economic incentives *given* the technological possibilities implied - though not unavoidable - in the TC specifications.

We discuss the interaction of content users, content providers, and platform providers. Content providers seek the help of platform providers to make copying content exceedingly costly, and use TC systems for DRM enforcement. Platform providers compare the payoffs from producing TC platforms (control, revenues from content providers, revenues from consumers, etc.) versus the payoffs from non-TC platforms (reduced costs, possibly increased network externalities, etc.). Content consumers compare the utility from non-TC platforms (control on their systems, free digital content - which may carry legal liabilities, and so on), with the utility from TC platforms (loss of control, paid digital content - without liabilities, and so on).

We show that, in addition to obvious factors (the magnitude of liabilities associated with copyright infringement, the usability of darknets, and the market power associated with controlling a dominant platform), the plat-

form providers' optimal decisions rely on the network effects of different types of content. Widely popular digital content attracts the attention of several users. Because of users' preferences and interests, that content's DRM protection is likely to be broken, and the content itself is likely to be eventually released on the darknet, and reproduced there. However, individually low demand content is unlikely to attract enough attention to justify the initial cost of breaking the protection, and is not likely to be reproduced on the darknet. Aggregating these individual inconveniences across many consumers, the non-TC platform may lose appeal compared to the TC platform.

In addition, two contrasting forces are at play: the individual costs of adoption of new platforms, and the providers' costs of compliance to the new platform. The interplay of network externalities and adoption dynamics can determine the success or failure of trusted computing initiatives.

2 DRM, Darknets, and Trusted Computing

2.1 DRM

Over the last few years, advancements in digital compression, bandwidth, and communication technologies have made it possible for all sort of information goods to be cheaply digitized, reproduced, and distributed. One after the other various applications and protocols have made it possible to share files (including copyrighted material) online between parties that do not know each other. The control that creators or producers of content and copyrighted material have on the use and dissemination of their works has been eroded. An intense debate over the need and means (or lack thereof) for protection of copyrighted material has developed.

In an effort to contain the growth of illegal copies of digital content, content providers have considered various strategies to manage access, use,

and dissemination of their products. Alternatively, those strategies have relied on technology, legal provisions, or economic incentive. Dhamija and Wallenberg (2003) summarize them as follows: some strategies depend on making information goods "rival" in the economic sense, in that sharing one's information product with others is made impossible; other strategies focus on making the information goods "excludable," in that their use can be monitored (through, say, watermarking) and therefore regulated; yet other strategies accept the public good nature of information goods (naturally non rival and non excludable), and therefore rely on indirect ways to finance the creative activity of content producers.

Most institutional content providers (such as the movie and the music industries) have relied on combinations of the first and the second strategies. Making the good rival, in particular, has been attempted through technology, by building digital rights management (DRM) systems that move back the control of each copy of a digital information good from its buyer to its creator. Several DRM technologies (such as Apple FairPlay or RealNetworks Helix) have been brought into the market with the intent of distributing content in a protected form (such as encryption or "flags") that only compliant devices (such as MP3 players) can operate. That protection may limit the ability to access, modify, or disseminate the content.

2.2 Darknets

No DRM technology has, so far, resisted determined attacks. All DRM systems aim at resisting BOBE (break-once, break everywhere) attacks, so that, even if one DRM client is "broken," the same vulnerability cannot be used for all other DRM clients of the same type. However, "[m]ost commercial DRM-systems have BOBE-exploits" (Biddle et al [2003]), and all commercial protections for copyright material have been routinely broken soon after they were released. Much more successful have been file sharing,

peer-to-peer networking technologies and applications - such as Napster, Gnutella, Morpheus, or Kazaa - through which both uncopyrighted and copyrighted files are disseminated. Some have referred to these networks as "darknets" (Biddle et al [2003]): the darknet "is the distribution network that emerges from the injection of [copyrighted objects, assuming that] any widely distributed object will be available to a fraction of users in a form that permits copying." The distribution of those objects through the network relies on the assumption that its users will copy objects if it is possible and interesting to do so. Biddle et al (2003) conclude that there are no technical impediments to these file sharing networks growing in convenience, bandwidth, and efficiency. "In the presence of an infinitely efficient darknet [...] even sophisticated DRM systems are inherently ineffective."

2.3 Trusted Computing

Content providers such as the RIAA have reacted to the weaknesses of DRM technologies by adopting an aggressive legal posture - including bringing to court users detected disseminating copyrighted material. At the same time, however, new research in computer science (Arbaugh, Farber, and Smith [1997]) has made the prospect of building difficult to break DRM technologies actually plausible. The Trusted Computing Platform Alliance (TCPA) was founded in 1999 with the goal of creating open industry standards for a trusted computing (TC) subsystem to be added to personal computers. Since 1999, the number of software and hardware platform providers members of the alliance has vastly increased. It now includes the overwhelming majority of the large players in the computer industry (CPU manufacturers like Intel, AMD, and Motorola; BIOS vendors such as AMI and National Semiconductor; application vendors such as Microsoft and Adobe; and systems vendors such as HP, IBM, Dell, Gateway, Fujitsu, Samsung, and Toshiba). The fruits of this alliance (whose name has meanwhile changed to

Trusted Computing Group [TCG]) have already started appearing in computer products. According to Safford (2002a), IBM has been shipping a predecessor of TC in NetVista desktop and ThinkPad notebook computers since 2000.

According to its proponents, TC's goal is increased security in (personal) computers. The TCPA main specifications (Trusted Computing Group [2003]) describe a chip with particular security functions: public key (encryption) functions, trusted boot functions, and management functions (see Safford [2002a]). The public key functions ensure that public/private encryption key pair generation, encryption, decryption, signature and verification can be operated on the chip itself (which is more secure and less vulnerable than on software). The trusted boot functions (see Arbaugh, Farber, and Smith [1997]) imply that, upon booting, any data that has been sealed on that computer under a specific "Platform Configuration Register" will be unsealed only if the Platform Configuration Register is verified to have the same value displayed at the time of sealing. In other words, the chip checks the configuration of the machine; if the configuration is altered from that stored in the register (for example, because somebody is trying to boot a different system, or a virus has modified the operating system), the sealed data remain unaccessible. TC therefore makes it difficult for external parties (but also the owner and user of a computer) to tamper with certain aspects of the computer configuration, and allows for "remote" attestation to other parties that the computer has not, in fact, been tampered with, and therefore can be "trusted."

2.3.1 NGSCB

However, trusted by whom and for what goals? To answer this question one has to look for more than TC specifications' explicit statements, and deduce their implications.

The TCPA chip functionalities can be used as building blocks for a set of features hard to implement on today's personal computers. Soon after founding TCPA, Microsoft expressed interest in a combination of software and hardware, Palladium, that would build on the TC functionalities in order to extend them. The Palladium project was first announced as a trusted computing component for Windows operating systems. After the criticisms it received from security researchers and computer activists (see Anderson [2003c] and Stallman [2002]), the project was renamed Next-Generation Secure Computing Base (NGSCB), and scheduled to be part of Microsoft's next operating system, Longhorn. In May 2004 Microsoft announced that NGSCB would not have been part of the initial release of Longhorn (see Evers [2004]), although it denied that its development would have been interrupted.

NGSCB is meant to combine the hardware TC system with a software componenet, the Nexus. The Nexus acts as a manager for applications trying to interact with the functionalities provided by the TC hardware. NGSCB, among other things, would help isolate a "trusted" (that is, in this context, difficult to tamper with or access without required authorization) memory space, and create trusted paths from keyboards and mouses to monitors and printers. Applications could run in protected spaces, so that a certain application could be prevented from reading or writing another application's data.

2.3.2 DRM and Trusted Computing

The Trusted Computing Group maintains that TC can give users more secure local data storage and can lower the risks of identity theft and data losses from software and physical attacks. Organizations will gain increased ability to deploy secure systems and platform providers will enjoy the ability to develop secure systems based upon open standards.

However, other features that TC also makes *possible* have attracted intense criticism. NGSCB has been taken as example and proof of the risks associated with TC.[2] On top of a TC chip, it is conceivable to build a trusted architecture of hardware and operating system (NGSCB) and applications that transfer control of a personal computer from its owner and user to the writers of the components that run on it (see Anderson [2003c]). This means that a platform and content vendor may control how the consumer can use its products. This ability could limit competition and harm consumers.

For example, while Safford (2002a) notes that TCPA, Palladium, and DRM are not the same system, Green (2002), Anderson (2003a,b,c), and Felten (2003a,b) note that TC systems will make it possible to have stronger, difficult to break DRM technologies: "[s]ome portions of the trusted computing research agenda have roots in DRM, and Microsoft has announced a DRM technology (Microsoft Rights Management Services) that it says will make use of NGSCB" writes Schoen (2002).[3]

How would this happen? Building on the chain of trust, encryption, and attestation implicit in the TC specifications, the chip vendor could make sure that the machine can only boot with a certain known status and operating system. The operating system provider may be able to control which applications may be run on its platform. Application vendors may tie their products in ways that make interoperability with or by other applications difficult or impossible. Content providers may be able to sell information goods that can be played but never copied, or that can be decrypted only

[2]See see Arbaugh (2002), Green (2002), Stallman (2002), Ross (2003a,b,c), Felten (2003a,b), Schoen (2002). For some defense of TC, see Anonymous (2004), Safford (2002a,b), Lampson, Manferdelli, Peinado, and Willman (2003), Sadeghi and Stuble (2003).

[3]According to Fisher [2002], Intel's "LaGrande" TC initiative "will not, in and of it self, contain DRM technologies" but "technologies such as Palladium would be able to interface directly with [Intel technology] to effect strict controls on such things as CD copying, software installations, and who knows what else."

on one machine or type of platform. A web server may refuse to show a certain page to a certain browser.[4] Stronger access control on confidential documents may be deployed - including remote censorship, automatic document destruction, and mandatory access control conditional to arbitrarily set parameters (Anderson [2003a,b,c]). Unlicensed software or data may be made impossible to run on other machines, or could be remotely traced and deleted. The information goods protected by this form of control may not simply be copyrighted entertainment, but applications, or just individual users' own files, such as emails. An email written with a certain application may, in theory, be opened only by that same application.

Still, TC initiatives may render the above scenarios technically feasible, but not necessarily economically desirable for content and platform providers. Will vendors have the incentives to push their technological capabilities towards these directions? And would consumers accept those restrictions?

From a technology standpoint, TC makes hardware-based cryptographic support possible "for proofs that a potential receiver's machine is running an approved software stack. By making such proofs prerequisites for the transfer of sensitive data, owners of these data can ensure that only authorized applications will be run and only authorized actions will be taken by users." (Bergemann, Feigenbaum, Shenkerz, and Smith [2004]). From an economic standpoint, because this support is hardware-based, it is also more robust (i.e. costly to break) and easy to implement (i.e. efficient) than solutions not relying on TC. However, any computer security ultimately can be broken, given enough resources (such as time, money, knowledge, or incentives; see Schneier [2001]). TC systems will be vulnerable too - to physical attacks (see Greenstadt and Raymond [2004]), because of design flaws and bugs, or

[4] Already in absence of TC, "[i]n a widely publicized case, MSN, the Microsoft Network, briefly refused to serve web pages to non-Microsoft browsers." Schoen [2002].

to ingenious attackers who may bypass TC protection (for example, recording a TC and DRM protected MP3 file through an external microphone - see Anderson [2003c]).

In economic terms, TC may not make it impossible, but rather *more costly* to perform certain actions (copy protected files, use unlicensed software, etc.) than on existing computer platforms. And consumers may still be able to choose between trusted computing platforms (that may either limit their ability to perform certain actions or increase their costs) and platforms that do not enforce TC or DRM. This paper discusses what are the likely business developments of TC indicatives, given the economic incentives and the frontier of technological possibilities implied in the TC architecture.

3 Related Literature

The promises and risks of trusted computing have generated a vigorous debate. Economic consideration have motivated TCPA members and play a central role in that debate.

Green (2002) argues that the business objectives of TCPA include preventing the use of unlicensed software, enforcing DRM, preventing CD ripping and DivX creation, and enable information flow control.

Felten (2003a,b) highlights the importance of interoperability and network economic effects in the development of TC systems.

Schechter, Greenstadt, and Smith (2003) discuss the economics of copying digital content, and point out the possible use of TC to create darknets where content is illegally distributed out of legal and technological reach of the content providers.

Bechtold (2003) discusses legal and economic aspects of DRM and TC, and brings attention to the Sony Aibo case - which can be taken as example of the control that platform providers like to extend on both content

providers and consumers.[5]

Anderson (2003a,b,c) offers the most comprehensive analysis so far of the economic motivations and implications of TC. Anderson (2003b) concludes that the value to corporate and government users of TC is unclear, since new powers could be granted to those users, but also new risks and liability may be faced. According to Anderson, also the content industry - an obvious candidate to benefit from TC-enforced DRM - may not really receive advantage from limiting the diffusion of its products through controls and constraints. Anderson concludes that hardware vendors and software vendors have most to gain. Through an interplay of network externalities and technological lock-ins, Microsoft could "[invest] in equipping the operating system platform [...] with TC mechanisms in order to reap a reward through higher fee income from its applications." The likely consequences would be that TC would centralize economic power, favoring large companies over small ones.

Bergemann, Feigenbaum, Shenkerz, and Smith (2004) present a position paper with the first attempt to formally represent TC markets. Basing their approach on Caillaud and Jullien (2004) and Rochet and Tirole (2004), they present the interaction between content providers, platform providers, and consumers as a form of bilateral competition. The content provider needs to choose amongst competing platforms for its product, but also satisfy the largest possible network of users. The authors do not actually solve the model, but use it to raise several interesting questions - such as what is the

[5]Aibo is a robotic dog produced by Sony. Sony took legal action against a programmer who had written programs to enhance Aibo's behavior. The case "exemplifies how DRM systems can be employed to control the use of and access to technology platforms. Essentially, Aibo is a platform on top of which software applications can be built and run. If such a platform is protected by a DRM system, the platform owner can control who is able to build applications on top of the platform. This can prevent unaffiliated software developers from developing applications for the platform" - see Bechtold (2003).

likely effect of varying amounts of competing platform providers and their governance structures; and how do adoption decisions by content providers and consumers depend on the distribution of consumers' valuations of different products.

In our paper we present a less general model than the one by Bergemann et al (2004), and we use it to answer a more specific set of questions. We study the incentives for platform providers to produce trusted computing, DRM compliant products in a market in which consumers can choose between trusted platforms and platforms that do not enforce TC or DRM, and we examine the possible consequences of those decisions for competition and welfare.

4 Economic Incentives for Platform Providers

We study the incentives for platform providers to produce trusted computing, DRM compliant products in a market in which consumers can choose between trusted platforms and platforms that do not enforce TC or DRM. We focus on the business developments that are likely to follow the economic incentives, given the technological possibilities implied in the TC architecture. We discuss the interaction of content users, content providers, and platform providers. We start from a generic setup - involving two competing platform - and then specialize it for the trusted computing case, specifying the roles of "platform" and "content" providers.

4.1 Setup

Consider two competing platform technologies, A and B. There are, as in Caillaud and Jullien (2004), Rochet and Tirole (2004), an Bergemann et al (2004), three sets of players: N_I consumers i, N_G content providers g, and N_L platform providers l. Imagine a repeated 3 period game. In each period,

first the platform providers have to adopt either technology A or B, and produce platforms for content and consumers under that technology. Then, each content provider chooses to distribute content for either platform A or B, or both. Finally, each consumer adopts one platform, and purchases the desired content available on that platform. The game repeats thereafter - although budget constraints and the earlier investments associated with a certain platform may prevent each type of player from switching platforms.

4.1.1 Consumers

We analyze this game, as usual, going backwards. Consumers receive utility from consuming content goods under either platform. The utility to a given consumer i from a certain content good g comes from a known random distribution F: $u_i^g \sim F^g()$. However, we imagine that consumers' utilities are also affected by their preferences for the platform on which they consume the content good. Consider for simplicity that such preferences are distributed uniformly on a line between A and B. This means that if the same content good were available under both platforms, a consumer lying closer to the A end-point of the line (or at the end-point itself) would prefer that good on platform A (or only value it there); a consumer lying closer to the B end-point of the line (or at the end-point itself) would prefer that good on platform B (or only value it there); and a consumer lying exactly half-way the A and B end-points would be indifferent between using the good under either platform. Let us arbitrarily assign the value of 0 to the A endpoint, and the value of 1 to the B endpoint, and represent consumer i's preference for A or B with $t_i = [0, 1]$. Hence when $t_i < 1/2$, consumer i prefers platform A. When $t_i > 1/2$, consumer i prefers platform B. when $t_i = 1/2$, she is indifferent. Then the utility of content g to consumer i under platform A and B will be respectively $u_i^{g_A} = (1 - t)u_i^g$ and $u_i^{g_B} = tu_i^g$.

For consumer i on platform A, then, net utility is given by:

$$\sum_{g=1}^{N_{G_A}+N_{I_A}} x_{ig}(u_i^{g_A} - p_{g_A}) + \sum_{g=1}^{N_{G_B}++N_{I_B}} x_{ig}(u_i^{g_B} - d_{g_B}) - p_A \qquad (1)$$

N_{G_A} is the number of content providers releasing content for platform A (each content provider is assumed to produce one good, which can be sold without capacity constraints). In addition, consumers may also be producers of content - emails, text, various documents and projects in digital form exchanged among peers. Having other consumers using the same platform and producing compatible content provides both direct and indirect externalities that increases the value of the platform. This is captured by the sum: $\sum_{g=1}^{N_{G_A}+N_{I_A}}(u_i^{g_A})$. x_{ig} is a dummy variable with values 0 or 1, representing the decision of consumer i (not) to purchase good g. p_{g_A} is the price the customer pays for the good, while d_{g_B} is the cost she has to incur to use on platform A a good originally produced for platform B. This could include costs such as the cost of copying the good, the legal risks associated to copying, and so on. p_A is the price the consumer must pay to use platform A. Replace A with B in the equation and discussion above to derive the net utility for consumer i using platform B:

$$\sum_{g=1}^{N_{G_B}+N_{I_B}} x_{ig}(u_i^{g_B} - p_{g_B}) + \sum_{g=1}^{N_{G_A}++N_{I_A}} x_{ig}(u_i^{g_A} - d_{g_A}) - p_B \qquad (2)$$

Each consumer compares the net value from 1 to the net value from 2. In the first round, she adopts a platform. (Although we are not explicitly modelling it here, we discuss below how budget constraints will make such decision sticky and will impact the feasibility of deploying a successful TC platform.) At the time the consumer needs to make a decision, prices and costs $p_{g_B}, p_{g_A}, d_{g_B}, d_{g_A}, p_B$, and p_A are known. We also assume that each consumer knows the size of each network, and is not strategic in her decision process - that is, she acts according to what she knows about the current

platform membership, and does not consider how the strategic behavior of her peers at that period may affect the size of the network the following periods.

4.1.2 Content Providers

Each content provider produces one good - for either platform A or B, or both. We assume that each provider produces one unique good. However, competing providers can produce similar goods - so the content provider behaves as a monopolistic competitor. Provider g's profits will then result from:

$$p_{g_X} D_{g_X}(p_{g_X}, N_{I_X}, N_{G_X}) - c_{g_X} \tag{3}$$

where X can be A, B, or both. p_{g_X} is the equilibrium price for good g on platform X. Demand for that good, D_{g_X}, is function of its price, the size of the market N_{I_X}, and the number of competitors N_{G_X}. The cost of producing the first unit of good g under platform X, c_{g_X} is fixed - we assume that the marginal cost of reproduction is zero. In equilibrium, profits will be driven to zero. At the time content providers need to make their decision, the price of the platform (that impacts consumers' decision to adopt a platform) is known. Since all information in Equations 1 and 2 is either publicly known or can be calculated from random functions which are common knowledge, and since content providers also set p_{g_A} and p_{g_B} optimally, the consumers' optimal decision can be calculated and be inserted in Equation 3. Based on this, content providers will decide for which platforms they will distribute their goods.

4.1.3 Platform Providers

The N_P platform providers p can produce either platform A or B. Their revenues come from the consumers' purchases of their platform. (For simplicity,

we do not consider possible additional revenues, such as royalties from the content providers. Rochet and Tirole [2004] note that platform providers in bilateral markets often use one side of the market as a loss-leader subsidized by the other side.) Since platform providers have the same information as the content providers, and can calculate their optimal strategies, they can also find the solution to the problem of which platform X will maximize the following:

$$(p_{l_X} - c_{l_X}) D_{l_X}(p_{l_X}, N_{I_X}, N_{L_X}) - c_X \tag{4}$$

c_{l_X} is the marginal cost of producing a unit of platform X. c_X is the fixed cost associated to that platform. We assume that platform providers are Cournot competitors.

Under general conditions, the model we are describing would be untractable without recurring to further simplifications. In a comparable setup, Caillaud and Jullien (2004), for example, consider only 2 intermediaries, and homogeneous consumers that only engage in a single transaction (rather than, as we imagine here, many potential ones). Rochet and Tirole (2004) first consider a monopoly platform benchmark, and then the duopoly platform case with symmetric equilibria and consumers engaging in single individual transactions. Still, if we consider symmetric providers, assume that $d_{g_B} = d_{g_A}$ and all other costs are symmetric between platforms, and imagine that half of the Cournot platform competitors may have already spent one penny towards the development of platform A technology (and the remaining half a penny towards the development of platform B technology), then we know that content providers will release their goods for both platforms, there will exist a market for both platform and, under the assumption of linear consumers' preferences for A and B, we will in fact have $N_{I_A} = N_{I_B}$.

4.2 TC vs. Darknet

Platforms A and B as presented above are undistinguishable. We specify the model by defining the TC platform as A and the non-TC platform as B. But what is really the difference?

The TC platform is a platform where platform producers can control what content the content providers can produce for their platform, and where content providers can enforce DRM, so that they can control how their content is used. In particular, they can make it it difficult to copy that content to others, and to transfer content produced for the TC platform A onto the platform B. Content providers can also make it difficult for content produced on the non-TC platform to be used on the TC platform.

The non-TC is simply a computer platform that does not adopt any form of trusted computing, and where DRM is therefore uneforceable (or simply ineffective, as on today's platforms) and darknets are possible. This does not mean that all content will be copied. Content (if any) created on the darknet can be cheaply reproduced, but not everybody on the B platform will copy. Content created on the TC platform may be introduced into the darknet only if the DRM protection is broken - which is costly. Hence content on the B platform may be similar to content on today's systems: a mix of open source content, consumer generated content (such as emails, documents, research, and so on), uncopyrighted material, content copied from the TC platform, or content by those providers that choose to bear the risk of having their products copied and to compete with the possibly lower prices of open source projects.[6]

Consumers' preferences are distributed along a line from A to B because

[6]Note however that we do not identify darknet with open source - and we use the term in the sense originally coined by Biddle et al (2003). Note also that we have been using the terms "platform" and "content" rather loosely. We explain why and further explain their meaning below.

some consumers may appreciate most the freedom and openness of a non TC system, while others may appreciate most the features and trust of the TC system.

Let us start from platform B. We will make some drastic simplifications to make the analysis tractable. First, we assume that the price consumers pay for the TC and non TC platforms is at least initially the same. Second, we assume that a symmetric equilibrium exists in the market for content goods, with a distribution of prices and utilities such that any given consumer is interested in a fraction s of content produced under any platform.

As discussed above, on the B platform (the darknet) content may be available - as open source project, academic research, non copyrighted content, and so on - at a very low price, p_{g_B}. Content from the TC platform is only available if copied - which is costly: d_{g_A}. We start considering the median consumer - the one indifferent between consuming goods under platform A and B. We ask under what conditions she will choose either platform. For this we must compare Equations 1 and 2, simplified after the above assumptions:

$$s(N_{G_B} + N_{I_B})(u_i^{g_B} - p_{g_B}) + s(N_{G_A} + N_{I_A})(u_i^{g_A} - d_{g_A}) - p_B \tag{5}$$

versus:

$$s(N_{G_A} + N_{I_A})(u_i^{g_A} - p_{g_A}) + s(N_{G_B} + N_{I_B})(u_i^{g_B} - d_{g_B}) - p_A \tag{6}$$

Simplifying, we find that the median consumer will prefer platform B if:

$$\frac{N_{G_A} + N_{I_A}}{N_{G_B} + N_{I_B}}(p_{g_A} - d_{g_A}) + d_{g_B} - p_{g_B} > 0 \tag{7}$$

This inequality depends, naturally, on the relative size of the networks on each platform and the relative prices and costs of obtaining content goods

under each platform (or copying them from one platform to the other). Low prices for content, p_{g_B}, will make the darknet more appealing. An high cost of copying content from A into B (d_{g_A}) must be compensated by a large (relative to A) network B. On the other side, for a platform provider to invest resources on the TC platform A, it should be the case that copying content into the darknet must be expensive compared to simply buying it on the legitimate platform, because this will tend to increase the size of the network of platform A. When p_{g_B} is low, as we assumed, it is also important to make it possible to use content from platform B on A - unless network A is much larger than B.

4.2.1 Why TC May Succeed and Why it May Fail: Low Demand Content, the Costs of Breaking TC, and Dynamics

The direct and indirect externalities implicit in Equation 7 may push the median consumer (and its marginal peers) towards one platform or the other. For example, even if content on B were freely or cheaply distributed, the prevalence of content only available on platform A may tilt the equilibrium in that platform's direction.

As for that, certain content providers may avoid releasing goods into the non-TC platform, for fear that their content could be pirated or sold at lower prices on the B platform. A content providers may use TC to make the diffusion of A content into B costly. This may mean, in practice, making it difficult to copy goods produced by the content providers on A, but also making it difficult for consumers in A to pass their own free content (for example, exchange emails) to consumers in the B platform. The content available on the non-TC platform would be content by any content provider that had decided to invest on the B market at the cost of having its products pirated, content produced by its users, open source content, and content copied from the TC platform. If copies are costly and network A is large

relatively to B, its advantage over network B could rapidly grow.

In particular, widely popular digital content from platform A would still attract the attention of several users in B. Its protection would likely be broken, and the content would likely be released on the B platform. However, individually low demand content (such as emails between individual consumers on different platforms) will not attract enough attention to justify the cost of breaking each content's protection, and may not be made available to the other platform. Aggregating these individual inconveniences across many consumers, the non TC network loses value compared to the TC platform. As some consumers decide to switch sides, a self-reinforcing dynamics threatens the existence of the B platform: eventually, the economies of scale necessary to produce platform B will not be possible given the size of B's network, and the TC platform may become the only choice.

On the other side, the main reasons why the TC initiative may fail rely on: 1) the political or technical inability to avoid distribution of user-generated content from A to B; and 2) deployment dynamics. Individual switching costs from one platform to another will make consumers' adoption of A slow. At the same time, both content and platform providers face costs to comply to new platform specifications, which imply delays in the process through which DRM capabilities will be implemented from platform to content providers (Microsoft's recent decision to postpone the deployment of its TC component, NGSCB, into its new operating system, Longhorn, was due also to the requests of software developers who did not want to pay the costs of altering their products to make the compliant with the new architecture - see Evers [2004]).

The effectiveness of TC depends on how far-reaching will be its vertical (from platform to content) and horizontal (across consumers and providers) penetration in the market. Because of switching and compliance costs, such penetration will be slow, and the costs of bypassing TC may never become

as high as to force significant migrations into the TC platform. With an A network which is not dominant, the incentives for platform and content providers to invest in TC initiatives may decrease, and the incentives for A platform providers to deviate from the "alliance" and offer their products for network B may increase. In other words, the internal dynamics of adoption costs, implementation costs, and copy costs may cause TC to fail.

5 Consumer Welfare

If TC platforms succeeded in marginalizing non-TC platforms, to a point where the incentives to produce the non-TC platform could no longer justify the required investment, what would the implications be for consumer welfare?

First, significant adoption, transaction, learning, and switching costs for both vendors and consumers may be associated with moving to TC platforms. Second, a monopoly on computer standards would, in general, imply higher prices, loss of product differentiation and variety, and a reduction in social welfare.[7] On the other hand, TC may contribute to consumer welfare in terms of security, privacy, and innovation. We discuss this possibility below.[8]

[7] Technological lock-ins on content can make their upstream platform markets more competitive (see Varian [2002]), or, as in this case, can help providers propagate their monopolies. In addition, Economides (1996) notes that in competitions over standards and network externalities, a monopolist with capacity constraint will have an incentive to subside its own competitors. In the scenario we consider, however, the marginal cost of reproducing information goods can be so low that capacity constraints may not matter.

[8] The comparative social benefits associated with different forms of intellectual property protection (or lack thereof) have been already debated elsewhere, and we do no discuss them here. See for example Benkler (2001) and Bechtold (2003).

5.1 Security

TC's explicit goal is to improve the security of today's computers. By allowing a computer to boot only in a verifiable status, chances are reduced that an attacker may have modified (for example) the operating system, planting a worm, or virus, or other code with the intention of damaging the system or stealing its data. By using encryption, chances are reduced that sensitive information may be accessed by unauthorized parties. By using "attestation" between different modules (including modules in different machines), peers may be complete transactions with other parties they do not know and yet can be considered trustworthy.

The security literature, however, has raised significant doubts about the likelihood that TC could effectively decrease the vulnerability of a computer platform. Ultimately, any computer security can be attacked and broken with enough resources. TC systems will be vulnerable because of physical attacks, design flaws, or other attacks that simply bypass the TC defenses (see Anderson [2003c]). Anderson (2003a) concludes that there is no statistical difference in terms of security under a closed and an open system: any large, complex system will inevitably contain flaws and bugs, that will be exploited to attack tomorrow's TC systems similarly to today's systems.

TC security goals may depend on almost ubiquitous adoption of TC components (horizontally across different machines, and vertically across the different components of the same machine - from chip to operating systems to applications and peripherals: see for example Hendricks and Doorn [2004]). Adoption (and thereafter any required modification) of TC systems will be costly.[9] In face of these costs and historical clues that any complex software

[9] As an example of a costly scenario, consider the discovery of a BOBE vulnerability of one of the encryption protocols embodied in the TC chip. Today's operating systems can be in general be "patched" at low (users') costs. Tomorrow's TC system may require replacing hardware components on every affected machine, similarly to brick and mortar product recalls.

system will contain flaws that make it vulnerable, more evidence is needed to support the view that the security benefit-cost trade-off implied by TC is positive.

5.2 Privacy

Increased privacy is another benefit that TC may bring to consumers. Greenstadt and Raymond (2004), for example, present a novel TC-based and HIPAA system for medical privacy. Similarly to what content providers can do to protect their content on TC platforms, individual citizens may do thanks to TC features to protect their personal information.

There are, however, some difficulties.

First, security experts have claimed that the proposed TC specifications could be used, in fact, to decrease privacy and anonymity. The TC specifications (Trusted Computing Group [2003]) envision an process of anonymous authentication which involves a trusted third party. Since authentication implies revealing identifying information and the trusted third party may collude with others ito reveal the identity of a "trusted" computer, privacy may be compromised (see Arbaugh [2002]).

Second, as in the case of platform providers, economic incentives will ultimately determine how the technological possibilities of TC will be used - to protect personal information or to trade it more easily. In the absence of clear incentives to protect consumers' data (such as liabilities or consumers' actual demand for privacy), there is no reason to believe that organizations would freely implement TC features for their consumers' sake. In fact, lacking those external incentives, organizations may even decide to ignore TC-based protection of consumers' data if the value of that information is high enough: the costs of reproducing personal information by manually bypassing the TC protection will be lower than the costs of breaking larger

files such as media content.[10]

Privacy technologies already exist, but have met lukewarm reaction in the marketplace. Since TC will give control of data to whoever owns the keys under which the data is sealed, and since privacy and identity risks are often associated with abuses by authorized personnel, there is little evidence that TC would contribute to increased personal privacy.

5.3 Innovation: Trusted Computing and von Neumann Architecture - From EDVAC to TCP/IP

TC may increase innovation by making new products and services possible. There are, however, also some associated risks with the chain of control that TC may make possible.

First, Varian (2002) notes that TC may reduce user driven innovation, which is a significant contributor to technological developments in the computer industry (see Thomke and von Hippel [2002]).

Second, while the "concept of booting a machine into a known state is implicit in early PCs where the BIOS was in ROM and there was no hard drive in which a virus could hide" (Anderson [2003c]), TC may subvert the basic principles on which computer and Internet technologies have evolved since World War II.

During War World II, EDVAC became the first complete stored-program computer design to be conceived.[11] Like its immediate predecessor (ENIAC), EDVAC was a digital, electronic, general-purpose machine. Its electronic switches were able to manipulate their internal states to represent digital information according to schemes not constrained by the original design. Indeed, the machine's design was open ended to new data and new calcula-

[10] Shapiro and Varian (1998, ch. 1) discuss how ProCD *manually* transcribed phonebook directories to produce its own CD-ROM directories.

[11] Though not completed (see McCartney [1999] and Aspray, Burks [1987]; see also von Neumann [1945]).

tions. However, unlike ENIAC, EDVAC did not have to be programmed by connecting forests of cables and switches, nor did it need to receive its inputs by continuously and mechanically reading external tapes or punched cards. It could read the instructions from a tape and then store this 'program' in its electronic memory. EDVAC's stored program design offered "a new type of electronic memory [...], a store that could be both written into and read from at electronic speeds" (Aspray, Burks [1987], p. 4). The mercury retard lines used as memory-storage devices could keep in electronic form both data and instructions in binary fashion in variable-addresses. Their contents could be erased and modified again to store new information. EDVAC's design implied the capacity for storing programs internally within a computer's memory and "issuing instructions at electronic speeds which were comparable with those available in the rest of the machine" (Williams [1997], p. 300).

By storing both the instructions and the data inside the computer's memory, and by making that memory readable and modifiable by the machine itself, different programs could run simultaneously on a computer. The concept of 'program' was changed (cf. Reitwiesner [1997]). It became possible to create software components (like operating systems) that modify the memory of the computer which hosts them to control, load or run *other* programs. The stored-program computer's hardware structure of memory addresses became in other words a *tabula rasa*, an open space that gets a new structure and function as decided by a program (the operating system) running on it. In turn, through the operating system, the use of the hardware's functions is modified by the application.

What is the relevance of the stored-program architecture to trusted computing? In the stored program architecture, the operating system needs to be written to comply with the hardware specifications, and the application needs to be written to comply with the operating system specifications.

Beyond that, the hardware cannot decide what operating system will be loaded unto a given machine, and the operating system cannot know in advance what applications will run. Through the chain of attestations and trusts that TC may make possible, a chip could control which operating system can boot and which cannot, and the operating system could control which applications can run. TC makes it technically feasible to insert an *ex-post* validation of content, suggesting a potential reversal of the flow of control from the stored-program architecture.

One can only speculate about the potential consequences of these developments. The idea of the stored program architecture is that the designer of a certain system layer (say, the hardware) cannot and does not need to know what specific component may be later on designed and used on the next layer (say, an operating system; for example, one may purchase a Windows-based laptop, and then years later replace the operating system with a new version of Linux). All modern computer systems are based on this idea, and both the speed at which computer and Internet technologies have developed and the role of users in this process have been affected by it. For example, TCP/IP (a Internet 'transport protocol') permits host to host communication despite the heterogeneity of the host or network layers. This means that new capabilities can be added to the network without having to change the hardware of the participating host computers - thanks to the stored program architecture, new components can simply be added. Insomuch as the economic incentives may cause TC to be used to control and limit which components can be added to a system, the patterns of computer and Internet evolution may also be affected.

6 Limitations and Conclusions

The analysis presented in this paper is preliminary. In particular, the model highlighted here is an initial step in our ongoing research. We have discussed

the incentives for platform providers to produce trusted computing, DRM-enforcing equipment in a market in which content users can choose between trusted platforms and darknets. We have shown that platform providers' optimal decisions rely on a series of factor including the network effects associated with consumer-generated content. Even more than widely popular, high-demand content, the aggregate impact of low-demand individual content and the individual costs of platform adoption can determine the success or failure of trusted computing initiatives. In addition, the contrasting dynamics of individual costs of adoption of new platforms and providers' costs of compliance to new specifications may determine the success or failure of trusted computing initiatives.

Bibliography

R. J. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons, Inc., first edition, 2001.

R. J. Anderson, "Security in Open versus Closed Systems - The Dance of Boltzmann, Coase and Moore." `http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/toulouse.pdf`, 2003a.

R. J. Anderson. "Cryptography and Competition Policy Issues with Trusted Computing." In *PODC'03*, Boston, Massachusetts, July 13-16, 2003b.

R. J. Anderson. "'Trusted Computing Frequently Asked Questions: TC / TCG / LaGrande / NGSCB / Longhorn / Palladium / TCPA." `http://www.cl.cam.ac.uk/rja14/tcpa-faq.url`, 2003c.

Anonymous. "Interesting Uses of Trusted Computing." `http://invisiblog.com/1c801df4aee49232/article/0df117d5d9b32aea8bc23194ecc270ec`, 2004.

W. A. Arbaugh. "The TCPA - What's wrong; What's right and what to do about." Department of Computer Science and UMIACS, `http://www.cs.umd.edu/~waa/TCPA/TCPA-goodnbad.pdf`, 2002.

W. A. Arbaugh, D. J. Farber, and J. M. Smith. "A Secure and Reliable Bootstrap Architecture." In *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, 65-71, 1997.

Article 29 Data Protection Working Party. "Working Document on Trusted Computing Platforms and in particular on the work done by the Trusted Computing Group (TCG group)." 11816/03/EN WP 86, January 23, 2004.

W. Aspray. *John Von Neumann and the Origins of Modern Computing.* Cambridge MA: MIT Press, 1990.

W. Aspray and A. Burks (ed). *Papers of John Von Neumann on Computing and Computer Theory.* Cambridge MA: MIT Press, 1987.

S. Bechtold. "The Present and Future of Digital Rights Management Musings on Emerging Legal Problems." In *Digital Rights Management: Technological, Economic, Legal and Political Aspects*, ed. E. Becker, W. Buhse, D. Gnnewig, and N. Rump, Springer-Verlag, 2003.

Y. Benkler. "Coase's Penguin, or, Linux and the Nature of the Firm." Presented at the *Conference on the Public Domain*, Duke Law School, Noveber 9-11, 2001.

D. Bergemann, J. Feigenbaum, S. Shenkerz, and J. M. Smith. "Towards an Economic Analysis of Trusted Systems." Presented at *The third Annual Workshop on Economics and Information Security*, College Park, Maryland, May 29-30, 2004.

P. Biddle, P. England, M. Peinado, and B. Willman. "The Darknet and the Future of Content Protection." In *Digital Rights Management: Technological, Economic, Legal and Political Aspects*, ed. E. Becker, W. Buhse, D. Gnnewig, and N. Rump, Springer-Verlag, 2003.

P. Boutin. "See You on the Darknet Why we don't really want Internet security." Slate.com, January 28, 2004.

B. Caillaud and B. Jullien. "Chicken and Egg: Competition among Intermediation Service Providers." *Rand Journal of Economics*, 34(2), 309-328, 2004.

V. G. Cerf and R. E. Kahn. "A Protocol for Packet-Network Intercommunication." *IEEE Transactions on Communications*, May, 1974.

V. G. Cerf and J. B. Postel. "Specification of Internetwork Transmission Control protocol: TCP Version 3." Information Sciences Institute, University of Southern California, January, 1978.

R. Dhamija and F. Wallenberg. "A Framework for Evaluating Digital Rights Management Proposals." In *First International Mobile IPR Work-*

shop: Rights Management of Information Products on the Mobile Internet, Helsinki, Finland, 2003.

N. Economides. "The Economics of Networks." *International Journal of Industrial Organization*, 14(2), 1996.

P. England, B. Lampson, J. Manferdelli, M. Peinado, and B. Willman. "A trusted open platform." *IEEE Computer*, 55-62, July 2003.

J. Evers. "Microsoft revisits NGSCB security plan." Computerworld, May 06, 2004.

German Federal Government. "Federal Government's Comments on the TCG and NGSCB in the Field of Trusted Computing." `http://www.bsi.de/trustcomp/stellung/StellungnahmeTCG1_2a_e.pdf`, 2004.

E. W. Felten. "Understanding Trusted Computing - Will Its Benefits Outweigh Its Drawbacks?" IEEE Security and Privacy, 60-62, May-June, 2003a.

E. W. Felten. "A Skeptical View of DRM and Fair Use." Communications of the ACM, 46(4), 57, 2003b.

K. Fisher. "Intel joins the DRM fray." arstechnica.com, `http://arstechnica.com/news/posts/20020910-1425.html`, September 10, 2002.

C. Flick. "The Controversy over Trusted Computing." BSc Thesis, University of Sydney Unit for History and Philosophy of Science, June, 2004.

I. Fried. "Longhorn to put squeeze on gadgets." CNET News.com, September 9, 2004.

L. Green. "Trusted Computing Platform Alliance: The Mother(board) of all Big Brothers." Presented at *DEFCON 10*, `http://www.cypherpunks.to/TCPA_DEFCON_10.pdf`, 2002.

R. Greenstadt and J. F. Raymond. "Trusted Computing for Medical Privacy." Presented at the *PORTIA Workshop on Sensitive Data*, Stanford, 2004.

S. Haber, B. Horne, J. Pato, T. Sander, ad R. E. Tarjan. "If Piracy is the Problem, Is DRM the Answer?" Trusted Systems Laboratory HP Laboratories Cambridge, HPL-2003-110 May 27th , 2003.

K. Hafner and M. Lyon. *Where Wizards Stay up Late - The Origins of the Internet.* New York: Touchstone, 1996.

J. Hendricks and L. van Doorn. "Secure Bootstrap is Not Enough: Shoring up the Trusted Computing Base." *Proceedings of the Eleventh SIGOPS European Workshop*, Leuven, Belgium, September 2004.

P. Liang. "What to expect from Microsoft's NGSCB plan." Computerworld, August 19, 2004.

Lickelider. "Man-computer symbiosis." *IRE Transactions on Human Factors in Electronics,* 1(1), 4-11, 1960.

S. McCartney. *ENIAC.* New York: Walker and Company, 1999.

J. Naughton. *A Brief History of the Future.* London: Weidenfeld and Nicolson, 1999.

J. von Neumann. "First Draft of a Report on the EDVAC." In Aspray W., Burks A. (eds.), (1987), 17-82, 1945.

G. W. Reitwiesner. "The First Operating System for EDVAC." *IEEE Annals of the History of Computing*, 19(1), 1997.

J. Rochet and J. Tirole. "Platform Competition in Two-Sided Markets." Forthcoming, *Journal of the European Economic Association*, 2004.

A. R. Sadeghi and C. Stuble. "Bridging the Gap between TCPA/Palladium and Personal Security." Technical Report, Saarland University, 2003.

D. Safford. "Clarifying Misinformation on TCPA." Technical report, IBM Research. `http://www.research.ibm.com/gsal/tcpa/tcparebuttal.pdf`, 2002a.

D. Safford. "The Need for TCPA." Technical report, IBM Research, 2002b.

B. Schneier. "The Futility of Digital Copy Prevention" Crypto-Gram Newsletter, May 15, 2001.

C. Shapiro and H. R. Varian. *Information Rules*. Harvard Business School Press, 1998.

S. E. Schechter, R. A. Greenstadt, and M. D. Smith. "Trusted Computing, Peer-To-Peer Distribution, and the Economics of Pirated Entertainment." Presented at *The Second Annual Workshop on Economics and Information Security*, College Park, Maryland, May 29-30, 2003.

S. Schoen. "Trusted Computing: Promise and Risk." EFF, 2002.

R. R. Stallman. "Can you trust your computer?" In *Free Software, Free Society: The Selected Essays of Richard M. Stallman*, 2002.

S. Thomke and E. von Hippel. "Customers as Innovators: A New Way to Create Value." *Harvard Business Review*, 80(4), 74-81, 2002.

Trusted Computing Group. "TCG TPM Specification v1.2: Design Principles." Technical report, 2003.

H. R. Varian. "New Chips Can Keep a Tight Rein on Consumers."New York Times, July 4, 2002.

M. R. Williams. *A History of Computing Technology*. Los Alamitos, CA: IEEE Computer Society Press, 1997.