The Darknet and the Future of Content Distribution

Peter Biddle, Paul England, Marcus Peinado, and Bryan Willman

Microsoft Corporation[1]

Abstract

We investigate the *darknet* – a collection of networks and technologies used to share digital content. The darknet is not a separate physical network but an application and protocol layer riding on existing networks. Examples of darknets are peer-to-peer file sharing, CD and DVD copying, and key or password sharing on email and newsgroups. The last few years have seen vast increases in the darknet's aggregate bandwidth, reliability, usability, size of shared library, and availability of search engines. In this paper we categorize and analyze existing and future darknets, from both the technical and legal perspectives. We speculate that there will be short-term impediments to the effectiveness of the darknet as a distribution mechanism, but ultimately the darknet-genie will not be put back into the bottle. In view of this hypothesis, we examine the relevance of content protection and content distribution architectures.

1 Introduction

People have always copied things. In the past, most items of value were physical objects. Patent law and economies of scale meant that small scale copying of physical objects was usually uneconomic, and large-scale copying (if it infringed) was stoppable using policemen and courts. Today, things of value are increasingly less tangible: often they are just bits and bytes or can be accurately represented as bits and bytes. The widespread deployment of packet-switched networks and the huge advances in computers and codec-technologies has made it feasible (and indeed attractive) to deliver such digital works over the Internet. This presents great opportunities and great challenges. The opportunity is low-cost delivery of personalized, desirable high-quality content. The challenge is that such content can be distributed illegally. Copyright law governs the legality of copying and distribution of such valuable data, but copyright protection is increasingly strained in a world of programmable computers and high-speed networks.

For example, consider the staggering burst of creativity by authors of computer programs that are designed to share audio files. This was first popularized by Napster, but today several popular applications and services offer similar capabilities. CD-writers have become mainstream, and DVD-writers may well follow suit. Hence, even in the absence of network connectivity, the opportunity for low-cost, large-scale file sharing exists.

1.1 The Darknet

Throughout this paper, we will call the shared items (e.g. software programs, songs, movies, books, etc.) *objects*. The persons who copy objects will be called *users* of the darknet, and the computers used to share objects will be called *hosts*.

---

[1] Statements in this paper represent the opinions of the authors and not necessarily the position of Microsoft Corporation.

The idea of the darknet is based upon three assumptions:

1. Any widely distributed object will be available to a fraction of users in a form that permits copying.
2. Users will copy objects if it is possible and interesting to do so.
3. Users are connected by high-bandwidth channels.

The *darknet* is the distribution network that emerges from the injection of objects according to assumption 1 and the distribution of those objects according to assumptions 2 and 3.

One implication of the first assumption is that any content protection system will leak popular or interesting content into the darknet, because some fraction of users--possibly experts–will overcome any copy prevention mechanism or because the object will enter the darknet before copy protection occurs.

The term “widely distributed” is intended to capture the notion of mass market distribution of objects to thousands or millions of practically anonymous users. This is in contrast to the protection of military, industrial, or personal secrets, which are typically not widely distributed and are not the focus of this paper.

Like other networks, the darknet can be modeled as a directed graph with labeled edges. The graph has one vertex for each user/host. For any pair of vertices (u,v), there is a directed edge from u to v if objects can be copied from u to v. The edge labels can be used to model relevant information about the physical network and may include information such as bandwidth, delay, availability, etc. The vertices are characterized by their object library, object requests made to other vertices, and object requests satisfied.

To operate effectively, the darknet has a small number of technological and infrastructure requirements, which are similar to those of legal content distribution networks. These infrastructure requirements are:

1. facilities for injecting new objects into the darknet (input)
2. a distribution network that carries copies of objects to users (transmission)
3. ubiquitous rendering devices, which allow users to consume objects (output)
4. a search mechanism to enable users to find objects (database)
5. storage that allows the darknet to retain objects for extended periods of time. Functionally, this is mostly a caching mechanism that reduces the load and exposure of nodes that inject objects.

The dramatic rise in the efficiency of the darknet can be traced back to the general technological improvements in these infrastructure areas. At the same time, most attempts to fight the darknet can be viewed as efforts to deprive it of one or more of the infrastructure items. Legal action has traditionally targeted search engines and, to a lesser extent, the distribution network. As we will describe later in the paper, this has been partially successful. The drive for legislation on mandatory watermarking aims to deprive the darknet of rendering devices. We will argue that watermarking approaches are technically flawed and unlikely to have any material impact on the darknet. Finally, most content protection systems are meant to prevent or delay the injection of new objects into the darknet. Based on our first assumption, no such system constitutes an impenetrable barrier, and we will discuss the merits of some popular systems.

We see no technical impediments to the darknet becoming increasingly efficient (measured by aggregate library size and available bandwidth). However, the darknet, in all its transport-layer embodiments, is under legal attack. In this paper, we speculate on the technical and legal future of the darknet, concentrating particularly, but not exclusively, on peer-to-peer networks.

The rest of this paper is structured as follows. Section 2 analyzes different manifestations of the darknet with respect to their robustness to attacks on the infrastructure requirements described above and speculates on the future development of the darknet. Section 3 describes content protection mechanisms, their probable effect on the darknet, and the impact of the darknet upon them. In sections 4 and 5, we speculate on the scenarios in which the darknet will be effective, and how businesses may need to behave to compete effectively with it.

2 The Evolution of the Darknet

We classify the different manifestations of the darknet that have come into existence in recent years with respect to the five infrastructure requirements described and analyze weaknesses and points of attack.

As a system, the darknet is subject to a variety of attacks. Legal action continues to be the most powerful challenge to the darknet. However, the darknet is also subject to a variety of other common threats (e.g. viruses, spamming) that, in the past, have lead to minor disruptions of the darknet, but could be considerably more damaging.

In this section we consider the potential impact of legal developments on the darknet. Most of our analysis focuses on system robustness, rather than on detailed legal questions. We regard legal questions only with respect to their possible effect: the failure of certain nodes or links (vertices and edges of the graph defined above). In this sense, we are investigating a well known problem in distributed systems.

2.1 Early Small-Worlds Networks

Prior to the mid 1990s, copying was organized around groups of friends and acquaintances. The copied objects were music on cassette tapes and computer programs. The rendering devices were widely-available tape players and the computers of the time – see Fig. 1. Content injection was trivial, since most objects were either not copy protected or, if they were equipped with copy protection mechanisms, the mechanisms were easily defeated. The distribution network was a “sneaker net” of floppy disks and tapes (storage), which were handed in person between members of a group or were sent by postal mail. The bandwidth of this network – albeit small by today’s standards – was sufficient for the objects of the time. The main limitation of the sneaker net with its mechanical transport layer was latency. It could take days or weeks to obtain a copy of an object. Another serious limitation of these networks was the lack of a sophisticated search engine.

There were limited attempts to prosecute individuals who were trying to sell copyrighted objects they had obtained from the darknet (commercial piracy). However, the darknet as a whole was never under significant legal threat. Reasons may have included its limited commercial impact and the protection from legal surveillance afforded by sharing amongst friends.

The sizes of object libraries available on such networks are strongly influenced by the interconnections between the networks. For example, schoolchildren may copy content from their “family network” to their “school network” and thereby increase the size of the darknet object library available to each. Such networks have been studied extensively and are classified as “interconnected small-worlds networks.” [24] There are several popular examples of the characteristics of such systems. For example, most people have a social group of a few score of people. Each of these people has a group of friends that partly overlap with their friends’ friends, and also introduces more people. It is estimated that, on

average, each person is connected to every other person in the world by a chain of about six people from which arises the term “six degrees of separation”.

These findings are remarkably broadly applicable (e.g. [20][~~20~~,[3][~~3~~]). The chains are on average so short because certain super-peers have many links. In our example, some people are gregarious and have lots of friends from different social or geographical circles..

We suspect that these findings have implications for sharing on darknets, and we will return to this point when we discuss the darknets of the future later in this paper.

The small-worlds darknet continues to exist. However, a number of technological advances have given rise to new forms of the darknet that have superseded the small-worlds for some object types (e.g. audio).

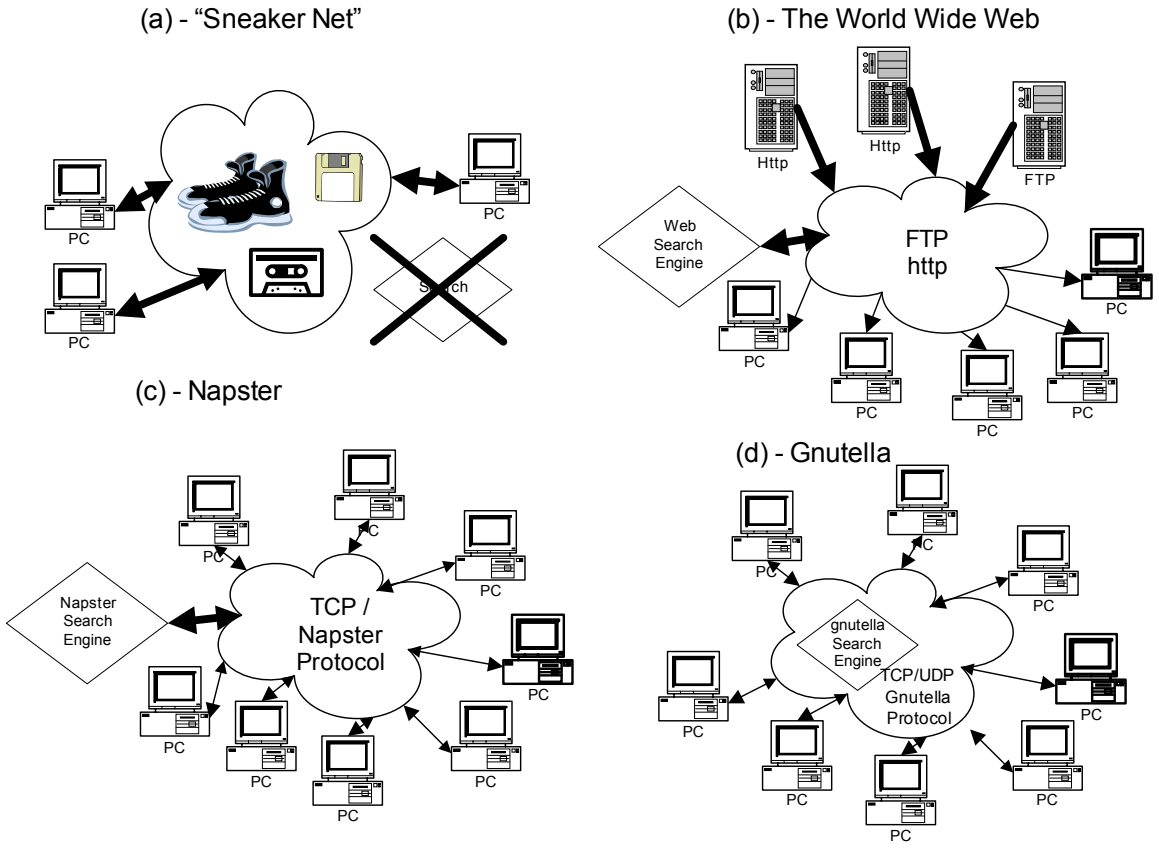

Figure 1: Historical evolution of the Darknet. We highlight the location of the search engine (if present) and the effective bandwidth (thicker lines represent higher bandwidth). Network latencies are not shown, but are much longer for the sneakernet than for the IP-based networks.

2.2 Central Internet Servers

By 1998, a new form of the darknet began to emerge from technological advances in several areas. The internet had become mainstream, and as such its protocols and infrastructure could now be relied upon by anyone seeking to connect users with a centralized service or with each other. The continuing fall in the price of storage together with advances in compression technology had also crossed the threshold at which storing large numbers of audio files was no longer an obstacle to mainstream users. Additionally, the power of computers had crossed the point at which they could be used as rendering devices for multimedia content. Finally, “CD ripping” became a trivial method for content injection.

The first embodiments of this new darknet were central internet servers with large collections of MP3 audio files. A fundamental change that came with these servers was the use of a new distribution network: The internet displaced the sneaker net – at least for audio content. This solved several problems of the old darknet. First, latency was reduced drastically.

Secondly, and more importantly, discovery of objects became much easier because of simple and powerful search mechanisms – most importantly the general-purpose world-wide-web search engine. The local view of the small world was replaced by a global view of the entire collection accessible by all users. The main characteristic of this form of the darknet was centralized storage and search – a simple architecture that mirrored mainstream internet servers.

Centralized or quasi-centralized distribution and service networks make sense for legal online commerce. Bandwidth and infrastructure costs tend to be low, and having customers visit a commerce site means the merchant can display adverts, collect profiles, and bill efficiently. Additionally, management, auditing, and accountability are much easier in a centralized model.

However, centralized schemes work poorly for *illegal* object distribution because large, central servers are large single points of failure: If the distributor is breaking the law, it is relatively easy to force him to stop. Early MP3 Web and FTP sites were commonly "hosted" by universities, corporations, and ISPs. Copyright-holders or their representatives sent "cease and desist" letters to these web-site operators and web-owners citing copyright infringement and in a few cases followed up with legal action [15]. The threats of legal action were successful attacks on those centralized networks, and MP3 web and FTP sites disappeared from the mainstream shortly after they appeared.

2.3 Peer-to-Peer Networks

The realization that centralized networks are not robust to attack (be it legal or technical) has spurred much of the innovation in peer-to-peer networking and file sharing technologies. In this section, we examine architectures that have evolved. Early systems were flawed because critical components remained centralized (Napster) or because of inefficiencies and lack of scalability of the protocol (gnutella) [17]. It should be noted that the problem of object location in a massively distributed, rapidly changing, heterogeneous system was new at the time peer-to-peer systems emerged. Efficient and highly scalable protocols have been proposed since then [9][~~9~~,[23]~~23~~].

2.3.1. Napster

Napster was the service that ignited peer-to-peer file sharing in 1999 [14]. There should be little doubt that a major portion of the massive (for the time) traffic on Napster was of copyrighted objects being transferred in a peer-to-peer model in violation of copyright law. Napster succeeded where central servers had failed by relying on the distributed storage of objects not under the control of Napster. This moved the injection, storage, network distribution, and consumption of objects to users.

However, Napster retained a centralized database[2] with a searchable index on the file name. The centralized database itself became a legal target [15]. Napster was first enjoined to deny certain queries (e.g. "Metallica") and then to police its network for all copyrighted content. As the size of the darknet indexed by Napster shrank, so did the number of users. This illustrates a general characteristic of darknets: there is positive

[2] Napster used a farm of weakly coupled databases with clients attaching to just one of the server hosts.

feedback between the size of the object library and aggregate bandwidth and the appeal of the network for its users.

2.3.2. Gnutella

The next technology that sparked public interest in peer-to-peer file sharing was Gnutella. In addition to distributed object storage, Gnutella uses a fully distributed database described more fully in [13]. Gnutella does not rely upon any centralized server or service – a peer just needs the IP address of one or a few participating peers to (in principle) reach any host on the Gnutella darknet. Second, Gnutella is not really "run" by anyone: it is an open protocol and anyone can write a Gnutella client application. Finally, Gnutella and its descendants go beyond sharing audio and have substantial non-infringing uses. This changes its legal standing markedly and puts it in a similar category to email. That is, email has substantial non-infringing use, and so email itself is not under legal threat even though it may be used to transfer copyrighted material unlawfully.

2.4 Robustness of Fully Distributed Darknets

Fully distributed peer-to-peer systems do not present the single points of failure that led to the demise of central MP3 servers and Napster. It is natural to ask how robust these systems are and what form potential attacks could take. We observe the following weaknesses in Gnutella-like systems:

- Free riding
- Lack of anonymity

2.4.1 Free Riding

Peer-to-peer systems are often thought of as fully decentralized networks with copies of objects uniformly distributed among the hosts. While this is possible in principle, in practice, it is not the case. Recent measurements of libraries shared by gnutella peers indicate that the majority of content is provided by a tiny fraction of the hosts [1]. In effect, although gnutella *appears* to be a peer-to-peer network of cooperating hosts, in actual fact it has evolved to effectively be another largely centralized system – see Fig. 2. *Free riding* (i.e. downloading objects without sharing them) by many gnutella users appears to be main cause of this development. Widespread free riding removes much of the power of network dynamics and may reduce a peer-to-peer network into a simple unidirectional distribution system from a small number of sources to a large number of destinations. Of course, if this is the case, then the vulnerabilities that we observed in centralized systems (e.g. FTP-servers) are present again. Free riding and the emergence of super-peers have several causes:

Peer-to-peer file sharing assumes that a significant fraction of users adhere to the somewhat post-capitalist idea of sacrificing their own resources for the "common good" of the network. Most free-riders do not seem to adopt this idea. For example, with 56 kbps modems still being the network connection for most users, allowing uploads constitutes a tangible bandwidth sacrifice. One approach is to make collaboration mandatory. For example, Freenet [6] clients are required to contribute some disk space. However, enforcing such requirements without a central infrastructure is difficult.

Existing infrastructure is another reason for the existence of super-peers. There are vast differences in the resources available to different types of hosts. For example, a T3 connection provides the combined bandwidth of about one thousand 56 kbps telephone connections.

2.4.2 Lack of Anonymity

Users of gnutella who share objects they have stored are not anonymous. Current peer-to-peer networks permit the server endpoints to be determined, and if a peer-client can determine the IP address and affiliation of a peer, then so can a lawyer or government agency. This means that users who share copyrighted objects face some threat of legal action. This appears to be yet another explanation for free riding.

There are some possible technological workarounds to the absence of endpoint anonymity. We could imagine anonymizing routers, overseas routers, object fragmentation, or some other means to complicate the effort required by law-enforcement to determine the original source of the copyrighted bits. For example, Freenet tries to hide the identity of the hosts storing any given object by means of a variety of heuristics, including routing the object through intermediate hosts and providing mechanisms for easy migration of objects to other hosts. Similarly, Mnemosyne [10] tries to organize object storage, such that individual hosts may not know what objects are stored on them. It is conjectured in [10] that this may amount to common-carrier status for the host. A detailed analysis of the legal or technical robustness of these systems is beyond the scope of this paper.

2.4.3 Attacks

In light of these weaknesses, attacks on gnutella-style darknets focus on their object storage and search infrastructures. Because of the prevalence of super-peers, the gnutella darknet depends on a relatively small set of powerful hosts, and these hosts are promising targets for attackers.

Darknet hosts owned by corporations are typically easily removed. Often, these hosts are set up by individual employees without the knowledge of corporate management. Generally corporations respect intellectual property laws. This together with their reluctance to become targets of lawsuits, and their centralized network of hierarchical management makes it relatively easy to remove darknet hosts in the corporate domain.

While the structures at universities are typically less hierarchical and strict than those of corporations, ultimately, similar rules apply. If the .com and .edu T1 and T3 lines were pulled from under a darknet, the usefulness of the network would suffer drastically.

This would leave DSL, ISDN, and cable-modem users as the high-bandwidth servers of objects. We believe limiting hosts to this class would present a far less effective piracy network today from the perspective of acquisition because of the relative rarity of high-bandwidth consumer connections, and hence users would abandon this darknet. However, consumer broadband is becoming more popular, so in the long run it is probable that there will be adequate consumer bandwidth to support an effective consumer darknet.

The obvious next legal escalation is to bring direct or indirect (through the affiliation) challenges against users who share large libraries of copyrighted material. This is already happening and the legal threats or actions appear to be successful [7]. This requires the collaboration of ISPs in identifying their customers, which appears to be forthcoming due to requirements that the carrier must take to avoid liability[3] and, in some cases, because of corporate ties between ISPs and content providers. Once again, free riding makes this attack strategy far more tractable.

It is hard to predict further legal escalation, but we note that the DMCA (digital millennium copyright act) is a far-reaching (although not fully tested) example of a law that

[3] The Church of Scientology has been aggressive in pursuing ISPs that host its copyright material on newsgroups. The suit that appeared most likely to result in a clear finding, filed against Netcom, was settled out of court. Hence it is still not clear whether an ISP has a responsibility to police the users of its network.

is potentially quite powerful. We believe it probable that there will be a few more rounds of technical innovations to sidestep existing laws, followed by new laws, or new interpretations of old laws, in the next few years.

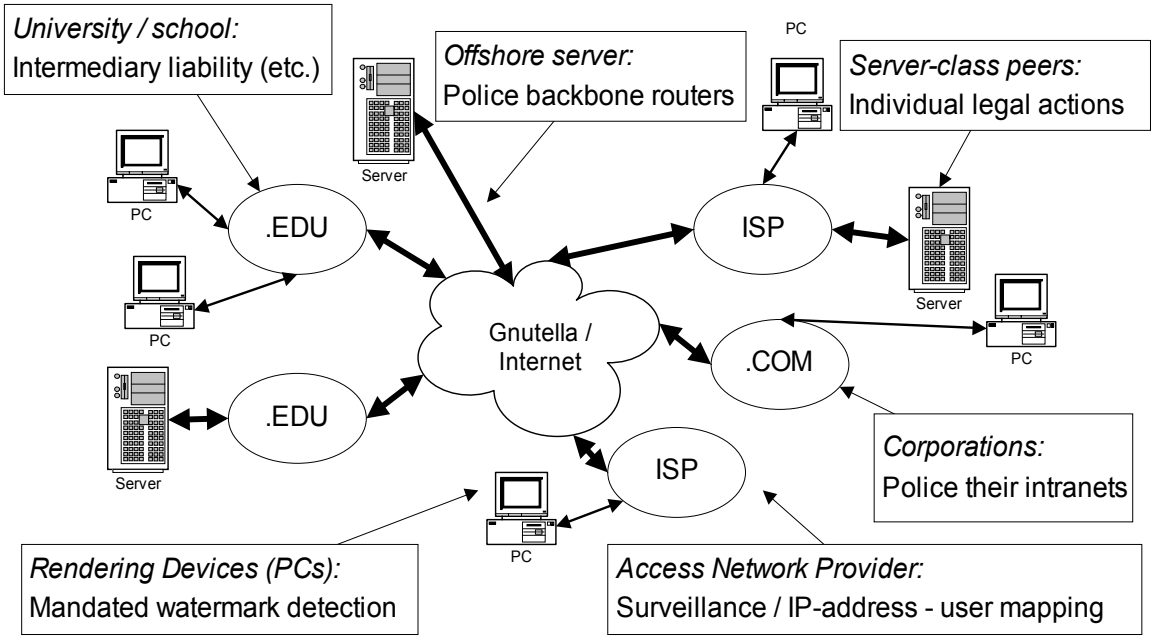

Figure 2: Policing the darknet. Gnutella-style networks appear hard to police because they are highly distributed, and there are thousands or millions of peers. Looking more closely there are several potential vulnerabilities.

2.4.4 Conclusions

All attacks we have identified exploit the lack of endpoint anonymity and are aided by the effects of free riding. We have seen effective legal measures on all peer-to-peer technologies that are used to provide effectively global access to copyrighted material. Centralized web servers were effectively closed down. Napster was effectively closed down. Gnutella and Kazaa are under threat because of free rider weaknesses and lack of endpoint anonymity.

Lack of endpoint anonymity is a direct result of the globally accessible global object database, and it is the existence of the global database that most distinguishes the newer darknets from the earlier small worlds. At this point, it is hard to judge whether the darknet will be able to retain this global database in the long term, but it seems seems clear that legal setbacks to global-index peer-to-peer will continue to be severe.

However, should Gnutella-style systems become unviable as darknets, systems, such as Freenet or Mnemosyne might take their place. Peer-to-peer networking and file sharing does seem to be entering into the mainstream – both for illegal and legal uses. If we couple this with the rapid build-out of consumer broadband, the dropping price of storage, and the fact that personal computers are effectively establishing themselves as centers of home-entertainment, we suspect that peer-to-peer functionality will remain popular and become more widespread.

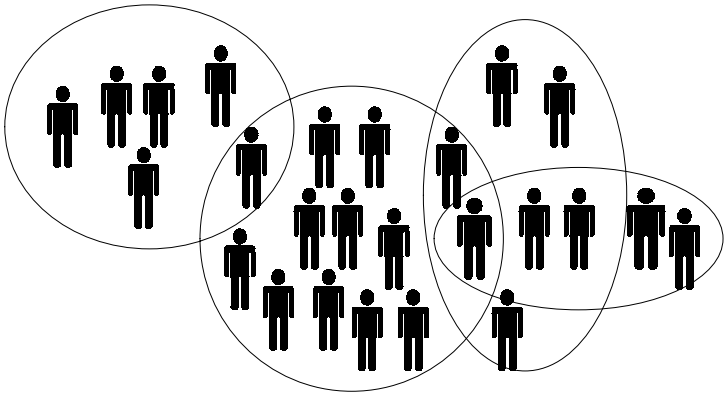

Figure 3: Interconnected small-worlds darknets. Threats of surveillance close global darknets. Darknets form around social groups, but use high bandwidth, low latency communications (the internet) and are supported by search engines. Custom applications, Instant-Messenger-style applic or simple shared file-systems host the darknet. People's social groups overlap so objects available in one darknet diffuse to others: in the termin used in this paper, each peer that is a member of more than one darknet is an introduction host for objects obtained from other darknets.

2.5 Small Worlds Networks Revisited

In this section we try to predict the evolution of the darknet should global peer-to-peer networks be effectively stopped by legal means. The globally accessible global database is the only infrastructure component of the darknet that can be disabled in this way. The other enabling technologies of the darknet (injection, distribution networks, rendering devices, storage) will not only remain available, but rapidly increase in power, based on general technological advances and the possible incorporation of cryptography. We stress that the networks described in this section (in most cases) provide poorer services than global network, and would only arise in the absence of a global database.

In the absence of a global database, small-worlds networks could again become the prevalent form of the darknet. However, these small-worlds will be more powerful than they were in the past. With the widespread availability of cheap CD and DVD readers and writers as well as large hard disks, the bandwidth of the sneaker net has increased dramatically, the cost of object storage has become negligible and object injection tools have become ubiquitous. Furthermore, the internet is available as a distribution mechanism that is adequate for audio for most users, and is becoming increasingly adequate for video and computer programs. In light of strong cryptography, it is hard to imagine how sharing could be observed and prosecuted as long as users do not share with strangers.

In concrete terms, students in dorms will establish darknets to share content in their social group. These darknets may be based on simple file sharing, DVD-copying, or may use special application programs or servers: for example, a chat or instant-messenger client enhanced to share content with members of your buddy-list. Each student will be a member of other darknets: for example, their family, various special interest groups, friends from high-school, and colleagues in part-time jobs (Fig. 3). If there are a few active super-peers - users that locate and share objects with zeal - then we can anticipate that content will rapidly diffuse between darknets, and relatively small darknets arranged around social groups will approach the aggregate libraries that are provided by the global darknets of today. Since the legal exposure of such sharing is quite limited, we believe that sharing amongst socially oriented groups will increase unabated.

Small-worlds networks suffer somewhat from the lack of a global database; each user can only see the objects stored by his small world neighbors. This raises a number of interesting questions about the network structure and object flow:

- What graph structure will the network have? For example, will it be connected? What will be the average distance between two nodes?
- Given a graph structure, how will objects propagate through the graph? In particular, what fraction of objects will be available at a given node? How long does it take for objects to propagate (diffuse) through the network?

Questions of this type have been studied in different contexts in a variety of fields (mathematics, computer science, economics, and physics). A number of empirical studies seek to establish structural properties of different types of small world networks, such as social networks [20] and the world-wide web [3]. These works conclude that the diameter of the examined networks is small, and observe further structural properties, such as a power law of the degree distribution [5], A number of authors seek to model these networks by means of random graphs, in order to perform more detailed mathematical analysis on the models [2][~~2~~,[8][~~8~~,[21][~~21~~,[22][~~22~~] and, in particular, study the possibility of efficient search under different random graph distributions [18][~~18~~,[19][~~19~~]. We will present a quantitative study of the structure and dynamics of small-worlds networks in an upcoming paper, but to summarize, small-worlds darknets can be extremely efficient for popular titles: very few peers are needed to satisfy requests for top-20 books, songs, movies or computer programs. If darknets are interconnected, we expect the effective introduction rate to be large. Finally, if darknet clients are enhanced to actively seek out new popular content, as opposed to the user-demand based schemes of today, small-worlds darknets will be very efficient.

3 Introducing Content into the Darknet

Our analysis and intuition have led us to believe that efficient darknets – in global or small-worlds form -- will remain a fact of life. In this section we examine rights-management technologies that are being deployed to limit the introduction rate or decrease the rate of diffusion of content into the darknet.

3.1 Conditional Access Systems

A conditional-access system is a simple form of rights-management system in which subscribers are given access to objects based (typically) on a service contract. Digital rights management systems often perform the same function, but typically impose restrictions on the use of objects after unlocking.

Conditional access systems such as cable, satellite TV, and satellite radio offer little or no protection against objects being introduced into the darknet from subscribing hosts. A conditional-access system customer has no access to channels or titles to which they are not entitled, and has essentially free use of channels that he has subscribed or paid for. This means that an investment of ~$100 (at time of writing) on an analog video-capture card is sufficient to obtain and share TV programs and movies. Some CA systems provide post-unlock protections but they are generally cheap and easy to circumvent.

Thus, conditional access systems provide a widely deployed, high-bandwidth source of video material for the darknet. In practice, the large size and low cost of CA-provided video content will limit the exploitation of the darknet for distributing video in the near-term.

The same can *not* be said of the use of the darknet to distribute conditional-access system broadcast keys. At some level, each head-end (satellite or cable TV head-end) uses an encryption key that must be made available to each customer (it is a broadcast),

and in the case of a satellite system this could be millions of homes. CA-system providers take measures to limit the usefulness of exploited session keys (for example, they are changed every few seconds), but if darknet latencies are low, or if encrypted broadcast data is cached, then the darknet could threaten CA-system revenues.

We observe that the exposure of the conditional access provider to losses due to piracy is proportional to the number of customers that share a session key. In this regard, cable-operators are in a safer position than satellite operators because a cable operator can narrowcast more cheaply.

3.2 DRM Systems

A classical-DRM system is one in which a client obtains content in protected (typically encrypted) form, with a license that specifies the uses to which the content may be put. Examples of licensing terms that are being explored by the industry are “play on these three hosts,” “play once,” “use computer program for one hour,” etc.

The license and the wrapped content are presented to the DRM system whose responsibility is to ensure that:

a) The client cannot remove the encryption from the file and send it to a peer,
b) The client cannot “clone” its DRM system to make it run on another host,
c) The client obeys the rules set out in the DRM license, and,
d) The client cannot separate the rules from the payload.

Advanced DRM systems may go further.

Some such technologies have been commercially very successful – the content scrambling system used in DVDs, and (broadly interpreted) the protection schemes used by conditional access system providers fall into this category, as do newer DRM systems that use the internet as a distribution channel and computers as rendering devices. These technologies are appealing because they promote the establishment of new businesses, and can reduce distribution costs. If costs and licensing terms are appealing to producers and consumers, then the vendor thrives. If the licensing terms are unappealing or inconvenient, the costs are too high, or competing systems exist, then the business will fail. The DivX “DVD” rental model failed on most or all of these metrics, but CSS-protected DVDs succeeded beyond the wildest expectations of the industry.

On personal computers, current DRM systems are software-only systems using a variety of tricks to make them hard to subvert. DRM enabled consumer electronics devices are also beginning to emerge.

In the absence of the darknet, the goal of such systems is to have comparable security to competing distribution systems – notably the CD and DVD – so that programmable computers can play an increasing role in home entertainment. We will speculate whether these strategies will be successful in the Sect. 5.

DRM systems strive to be BOBE (break-once, break everywhere)-resistant. That is, suppliers anticipate (and the assumptions of the darknet predict) that individual instances (clients) of all security-systems, whether based on hardware or software, will be subverted. If a client of a system is subverted, then all content protected by that DRM client can be unprotected. If the break can be applied to *any other* DRM client of that class so that all of those users can break their systems, then the DRM-scheme is BOBE-weak. If, on the other hand, knowledge gained breaking one client cannot be applied elsewhere, then the DRM system is BOBE-strong.

Most commercial DRM-systems have BOBE-exploits, and we note that the darknet applies to DRM-hacks as well. The CSS system is an exemplary BOBE-weak system. The knowledge and code that comprised the De-CSS exploit spread uncontrolled around

the world on web-sites, newsgroups, and even T-shirts, in spite of the fact that, in principle, the Digital Millennium Copyright Act makes it a crime to develop these exploits.

A final characteristic of existing DRM-systems is *renewability*. Vendors recognize the possibility of exploits, and build systems that can be field-updated.

It is hard to quantify the effectiveness of DRM-systems for restricting the introduction of content into the darknet from experience with existing systems. Existing DRM-systems typically provide protection for months to years; however, the content available to such systems has to date been of minimal interest, and the content that *is* protected is also available in unprotected form. The one system that was protecting valuable content (DVD video) was broken very soon after compression technology and increased storage capacities and bandwidth enabled the darknet to carry video content.

3.3 Software

The DRM-systems described above can be used to provide protection for software, in addition other objects (e.g. audio and video). Alternatively, copy protection systems for computer programs may embed the copy protection code in the software itself.

The most important copy-protection primitive for computer programs is for the software to be bound to a host in such a way that the program will not work on an unlicensed machine. Binding requires a machine ID: this can be a unique number on a machine (e.g. a network card MAC address), or can be provided by an external dongle.

For such schemes to be strong, two things must be true. First, the machine ID must not be "virtualizable." For instance, if it is trivial to modify a NIC driver to return an invalid MAC address, then the software-host binding is easily broken. Second, the code that performs the binding checks must not be easy to patch. A variety of technologies that revolve around software tamper-resistance can help here [4].

We believe that binding software to a host is a more tractable problem than protecting passive content, as the former only requires tamper resistance, while the latter also requires the ability to hide and manage secrets. However, we observe that all software copy-protection systems deployed thus far have bee broken. The definitions of BOBE-strong and BOBE-weak apply similarly to software. Furthermore, software is as much subject to the dynamics of the darknet as passive content.

4 Policing Hosts

If there are subverted hosts, then content will leak into the darknet. If the darknet is efficient, then content will be rapidly propagated to all interested peers. In the light of this, technologists are looking for alternative protection schemes. In this section we will evaluate watermarking and fingerprinting technologies.

4.1 Watermarking

Watermarking embeds an "indelible" invisible mark in content. A plethora of schemes exist for audio/video and still image content and computer programs.

There are a variety of schemes for exploiting watermarks for content-protection. Consider a rendering device that locates and interprets watermarks. If a watermark is found then special action is taken. Two common actions are:

1) *Restrict behavior:* For example, a bus-adapter may refuse to pass content that has the "copy once" and "already copied once" bits set.

2) *Require a license to play*: For example, if a watermark is found indicating that content is rights-restricted then the renderer may demand a license indicating that the user is authorized to play the content.

Such systems were proposed for audio content – for example the secure digital music initiative (SDMI) [16], and are under consideration for video by the copy-protection technical working group (CPTWG) [12].

There are several reasons why it appears unlikely that such systems will ever become an effective anti-piracy technology. From a commercial point of view, building a watermark detector into a device renders it strictly less useful for consumers than a competing product that does not. This argues that watermarking schemes are unlikely to be widely deployed, unless mandated by legislation. The recently proposed Hollings bill is a step along these lines [11].

We contrast watermark-based policing with classical DRM: If a general-purpose device is equipped with a classical DRM-system, it can play all content acquired from the darknet, *and* have access to new content acquired through the DRM-channel. This is in stark distinction to reduction of functionality inherent in watermark-based policing.

Even if watermarking systems were mandated, this approach is likely to fail due to a variety of technical inadequacies. The first inadequacy concerns the robustness of the embedding layer. We are not aware of systems for which simple data transformations cannot strip the mark or make it unreadable. Marks can be made more robust, but in order to recover marks after adversarial manipulation, the reader must typically search a large phase space, and this quickly becomes untenable. In spite of the proliferation of proposed watermarking schemes, it remains doubtful whether robust embedding layers for the relevant content types can be found.

A second inadequacy lies in unrealistic assumptions about key management. Most watermarking schemes require widely deployed cryptographic keys. Standard watermarking schemes are based on the normal cryptographic principles of a public algorithm and secret keys. Most schemes use a shared-key between marker and detector. In practice, this means that all detectors need a private key, and, typically, share a single private key. It would be naïve to assume that these keys will remain secret for long in an adversarial environment. Once the key or keys are compromised, the darknet will propagate them efficiently, and the scheme collapses. There have been proposals for public-key watermarking systems. However, so far, this work does not seem practical and the corresponding schemes do not even begin to approach the robustness of the cryptographic systems whose name they borrow.

A final consideration bears on the location of mandatory watermark detectors in client devices. On open computing devices (e.g. personal computers), these detectors could, in principle, be placed in software or in hardware. Placing detectors in software would be largely meaningless, as circumvention of the detector would be as simple as replacing it by a different piece of software. This includes detectors placed in the operating system, all of whose components can be easily replaced, modified and propagated over the darknet.

Alternatively, the detectors could be placed in hardware (e.g. audio and video cards). In the presence of the problems described this would lead to untenable renewability problems --- the hardware would be ineffective within days of deployment. Consumers, on the other hand, expect the hardware to remain in use for many years. Finally, consumers themselves are likely to rebel against “footing the bill” for these ineffective content protection systems. It is virtually certain, that the darknet would be filled with a continuous supply of watermark removal tools, based on compromised keys and weaknesses in the embedding layer. Attempts to force the public to “update” their hardware would not only be intrusive, but impractical.

In summary, attempts to mandate content protection systems based on watermark detection at the consumer's machine suffer from commercial drawbacks and severe technical deficiencies. These schemes, which aim to provide content protection beyond DRM by attacking the darknet, are rendered entirely ineffective by the presence of even a moderately functional darknet.

4.2 Fingerprinting

Fingerprint schemes are based on similar technologies and concepts to watermarking schemes. However, whereas watermarking is designed to perform *a-priori* policing, fingerprinting is designed to provide *a-posteriori* forensics.

In the simplest case, fingerprinting is used for individual-sale content (as opposed to super-distribution or broadcast – although it can be applied there with some additional assumptions). When a client purchases an object, the supplier marks it with an individualized mark that identifies the purchaser. The purchaser is free to use the content, but if it appears on a darknet, a policeman can identify the source of the content and the offender can be prosecuted.

Fingerprinting suffers from fewer technical problems than watermarking. The main advantage is that no widespread key-distribution is needed – a publisher can use whatever secret or proprietary fingerprinting technology they choose, and is entirely responsible for the management of their own keys.

Fingerprinting has one problem that is not found in watermarking. Since each fingerprinted copy of a piece of media is different, if a user can obtain several different copies, he can launch collusion attacks (e.g. averaging). In general, such attacks are very damaging to the fingerprint payload.

It remains to be seen whether fingerprinting will act as a deterrent to theft. There is currently no legal precedent for media fingerprints being evidence of crime, and this case will probably be hard to make – after all, detection is a statistical process with false positives, and plenty of opportunity for deniability. However, we anticipate that there will be uneasiness in sharing a piece of content that may contain a person's identity, and that ultimately leaves that person's control.

Note also that with widely distributed watermarking detectors, it is easy to see whether you have successfully removed a watermark. There is no such assurance for determining whether a fingerprint has been successfully removed from an object because users are not necessarily knowledgeable about the fingerprint scheme or schemes in use. However, if it turns out that the deterrence of fingerprinting is small (i.e. everyone shares their media regardless of the presence of marks), there is probably no reasonable legal response. Finally, distribution schemes in which objects must be individualized will be expensive.

5 Conclusions

There seem to be no technical impediments to darknet-based peer-to-peer file sharing technologies growing in convenience, aggregate bandwidth and efficiency. The legal future of darknet-technologies is less certain, but we believe that, at least for some classes of user, and possibly for the population at large, efficient darknets will exist. The rest of this section will analyze the implications of the darknet from the point of view of individual technologies and of commerce in digital goods.

5.1 Technological Implications

DRM systems are limited to protecting the content they contain. Beyond our first assumption about the darknet, the darknet is not impacted by DRM systems. In light of our first assumption about the darknet, DRM design details, such as properties of the tamper-resistant software may be strictly less relevant than the question whether the current darknet has a global database. In the presence of an infinitely efficient darknet – which allows instantaneous transmission of objects to all interested users – even sophisticated DRM systems are inherently ineffective. On the other hand, if the darknet is made up of isolated small worlds, even BOBE-weak DRM systems are highly effective. The interesting cases arise between these two extremes – in the presence of a darknet, which is connected, but in which factors, such as latency, limited bandwidth or the absence of a global database limit the speed with which objects propagate through the darknet. It appears that quantitative studies of the effective "diffusion constant" of different kinds of darknets would be highly useful in elucidating the dynamics of DRM and the darknet.

Proposals for systems involving mandatory watermark detection in rendering devices try to impact the effectiveness of the darknet directly by trying to detect and eliminate objects that originated in the darknet. In addition to severe commercial and social problems, these schemes suffer from several technical deficiencies, which, in the presence of an effective darknet, lead to their complete collapse. We conclude that such schemes are doomed to failure.

5.2 Business in the Face of the Darknet

There is evidence that the darknet will continue to exist and provide low cost, high-quality service to a large group of consumers. This means that in many markets, the darknet will be a competitor to legal commerce. From the point of view of economic theory, this has profound implications for business strategy: for example, increased security (e.g. stronger DRM systems) may act as a *disincentive* to legal commerce. Consider an MP3 file sold on a web site: this costs money, but the purchased object is as useful as a version acquired from the darknet. However, a securely DRM-wrapped song is strictly *less* attractive: although the industry is striving for flexible licensing rules, customers *will* be restricted in their actions if the system is to provide meaningful security. This means that a vendor will probably make more money by selling unprotected objects than protected objects. In short, if you are competing with the darknet, you must compete on the darknet's own terms: that is convenience and low cost rather than additional security.

Certain industries have faced this (to a greater or lesser extent) in the past. Dongle-protected computer programs lost sales to unprotected programs, or hacked versions of the program. Users have also refused to upgrade to newer software versions that are copy protected.

There are many factors that influence the threat of the darknet to an industry. We see the darknet having most direct bearing on mass-market consumer IP-goods. Goods sold to corporations are less threatened because corporations mostly try to stay legal, and will police their own intranets for illicit activities. Additionally, the cost-per-bit, and the total size of the objects have a huge bearing on the competitiveness of today's darknets compared with legal trade. For example, today's peer-to-peer technologies provide excellent service quality for audio files, but users must be very determined or price-sensitive to download movies from a darknet, when the legal competition is a rental for a few dollars.

References

[1] E. Adar and B. A. Huberman, *Free Riding on Gnutella,* http://www.firstmonday.dk/issues/issue5_10/adar/index.html

[2] W. Aiello, F. Chung and L. Lu, *Random evolution in massive graphs*, In Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, pages 510—519, 2001.

[3] R. Albert, H. Jeong and A.-L. Barabási, *Diameter of the world-wide web*, Nature 401, pages 130—131, 1999.

[4] D. Aucsmith, *Tamper Resistant Software, An Implementation*, Information Hiding 1996, Proceedings: Springer 1998.

[5] A.-L. Barabási, R. Albert, *Emergence of scaling in random networks*, Science 286, pages 509—512, 1999.

[6] I. Clarke, O. Sandberg, B. Wiley and T. Hong, *Freenet: A distributed information storage and retrieval system*, International Workshop on Design Issues in Anonymity and Unobservability, 2000.

[7] R. Clarke, *A defendant class action lawsuit* http://www.kentlaw.edu/perritt/honorsscholars/clarke.html

[8] C. Cooper and A. Frieze, *A general model of web graphs*, Proceedings of ESA 2001, pages 500-511, 2001.

[9] F. Dabek, E. Brunskill, M. F. Kaashoek, D. Karger, R. Morris, I. Stoica and H. Balakrishnan, *Building peer-to-peer systems with Chord, a distributed lookup service*, In Proceedings of the Eighth IEEE Workshop on Hot Topics in Operating Systems (HotOS-VIII), pages 81—86, 2001.

[10] S. Hand and T. Roscoe, *Mnemosyne: peer-to-peer steganographic storage*, In Proceedings of the First International Workshop on Peer-to-Peer Systems, 2002.

[11] Senator Fritz Hollings, *Consumer Broadband and Digital Television Promotion Act.*

[12] http://www.cptwg.org

[13] http://www.gnutelladev.com/protocol/gnutella-protocol.html

[14] http://www.napster.com

[15] http://www.riaa.org

[16] http://www.sdmi.org

[17] M. Javanovic, F. Annextein and K. Berman, *Scalability Issues in Large Peer-to-Peer Networks - A Case Study of Gnutella,* ECECS Department, University of Cincinnati, Cincinnati, OH 45221

[18] J. Kleinberg, Navigation in a small world, Nature 406, 2000.

[19] J. Kleinberg, *Small-world phenomena and the dynamics of information*, Advances in Neural Information Processing Systems (NIPS) 14, 2001.

[20] S. Milgram, The small world problem, Psychology Today, vol. 2, pages 60—67, 1967.

[21] M .Newman, *Small worlds: the structure of social networks*, Santa Fe Institute, Technical Report 99-12-080, 1999.

[22] M. Newman, D. Watts and S. Strogatz, *Random graph models of social networks*, Proc. Natl. Acad. Sci. USA 99, pages 2566—2572, 2002.

[23] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, H. Balakrishnan, *CHORD: A scalable peer-to-peer lookup service for internet applications*, In Proceedings of the ACM SIGCOMM 2001 Conference *SIGCOMM-01*, pages 149—160, 2001.

[24] D. J. Watts and S. H. Strogatz, *Collective dynamics of small-world networks*, Nature, 393:440-442, June 1998.