Towards An Ownership Layer for the Internet

Trent McConaghy and David Holtzman

Version 1.03, June 24, 2015

ABSTRACT

One long-standing drawback to the Internet is the hidden 'artist penalty.' The very strengths of the Internet make it difficult for creators of digital content to be fairly compensated for their work. This paper describes our implementation and proposes a fix that makes it easy to control one's intellectual property by constituting a new "ownership layer" on top of the existing Internet. The approach has two pieces: **a registry with easy, secure legals;** and **visibility into usage / provenance** of the content. The legals formalize existing copyright rights of digital objects, making them easy and fluid for a creator or collector to use, transfer or modify. The bitcoin blockchain is used to securely record ownership transactions that are impossible to later repudiate or manipulate. Internet-scale media search provides visibility into usage of the media by crawling the web, applying machine learning to identify similar or identical media, and subsequently reporting their existence and location to the registered owners. Taken together, these pieces constitute "ownership processing" – a simple tractable solution to make ownership actions of digital property universally accessible.

1 INTRODUCTION

On the Internet, it's difficult for creators of digital content to get compensated fairly for their work. Digital artifacts have the unique attributes of lossless replication and near-zero transmission costs. This has led to the vast proliferation of websites designed by people who can't paint, sing, photograph or write, but can still express themselves by copying the works of others into a different framework of meaning. The downside is that the originators don't often get attribution and rarely receive compensation. Frightened publishers, fearful of losing their revenue, have reacted by creating onerous, unworkable and consumer unfriendly Digital Rights Management (DRM) systems. DRM technology is not really effective, which is one reason why the music industry has effectively stopped using them. But there remains a compelling need for a workable solution to the ownership and attribution problem, especially as higher value items such as digital art, videos and even 3D printer templates become readily accessible online. We believe that a workable solution for everyday license management will result in an unprecedented explosion of creativity on the Internet, as artists feel that they can safely show, sell and license their work without fear of being robbed.

The online attribution problem traces back to the 1989 design of the WWW, which has only unidirectional links and therefore no built-in attribution, let alone ownership. The "original hypertext" vision from the 60s actually anticipated this issue; but the proposed solution (Xanadu) was too complex to deploy properly.

A *redesign* of the Internet[1] to handle ownership would be wildly impractical. But, we can *retrofit* it, i.e. add an overlay that addresses title issues, without requiring the cooperation of millions upon millions of websites. There are two complementary components: management (easy secure disposition of rights) and transparency (visibility in the usage of the property).

1. **Registry with easy, secure legals**:
    - *Easy legals* via a creator- and consumer-friendly Terms Of Service (TOS). These formalize existing copyright rights that are traditionally difficult for the layperson (creator, consumer) to leverage. Actions include: registering a work, transferring ownership, licenses, loans and rentals. Electronic transfer of money between peers, using services such as Paypal, enable full seller-to-buyer transactions.
    - *Time-stamping evidence* of ownership actions onto a *trusted ledger (registry)*. We use the bitcoin blockchain [1], a distributed database where anyone can add information, but no one can delete an action once it's committed, and no one owns the overall database. These time stamps can be used as evidence in court to resolve ownership disputes. This architecture means that ascribe itself is not the custodian of the licensing. Since the owner always has access to the tamper-proof blockchain showing the chain of custody and rights, ascribe will never be the weak link in the security chain. Using a blockchain as a trust-minimization scheme in this manner has already been used as evidence in a court of law [2].

2. **Visibility into usage** of the digital property by **auto-discovering** where it's used online and recording changes to ownership (**provenance**).
    - Tracking: crawl the web and show creators where their works are being used. Tracking leverages elements of machine learning and big data.
    - Screening: whenever a new item is registered, owners of similar IP that have already registered are notified.
    - Provenance: recording the history of ownership, of loans, and other transfer transactions. This web-crawl aspect is an a *posteriori* approach to provenance discovery, and the registry an *a priori* approach.

The rest of this paper is organized as follows. Section 2 describes symptoms of the root problem; and section 3 covers Internet history where these issues were (quite amazingly) actually anticipated and (even more amazingly) designed for. Section 4 describes the ascribe approach. Section 5 provides examples of how ascribe is being used. Section 6 describes a vision of what the future could be like. Section 7 summarizes the paper.

[1] It's not really clear what the "internet" actually is. For purposes of this paper, we mean the World Wide Web (WWW) and content offered up by the websites that can are accessible using the http protocol

2 SYMPTOMS

The overall issue is that ownership of digital property on the web is demonstrably hard. Many people describe it as broken. This section explores some of the symptoms of this breakage.

2.1 STOLEN IMAGES

Trent really, *really* loves robots. He once had “robot” written on his birthday cake rather than his name. He grew up building Transformers out of Lego. He can relate to Chappie [5] and TARS [6] and maybe even YLIP [7]. He even does a Google image search for something robot-related probably once a week. Furry robots, robot music, robot marmots, whatever. When he does a Google image search for “robot”, he gets something Figure 1 (minus the crossed-out part). Many robots, hooray! But he wants to use the images fairly, so he filters the resulting list to just the images that he’s allowed to freely use. Alas, only *one* of those images remains.

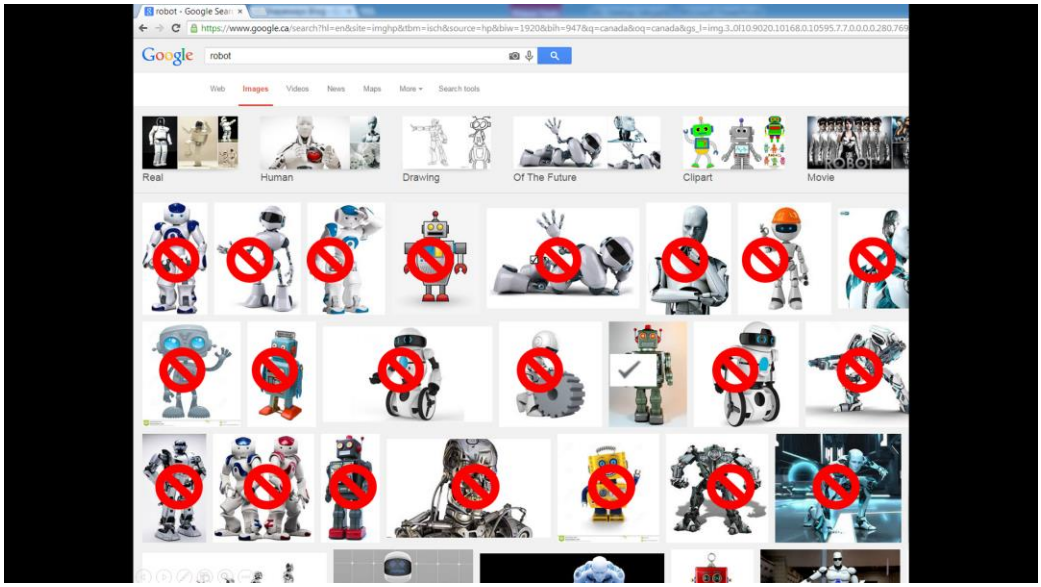

Figure 1: When doing a Google image search for "robot", only a fraction may be used fairly.

And if he wanted to purchase an image in the results of that search, how would he do it? There’s no clear way. He could go to a stock photo site, but it only holds a fraction of the total number of images.

What actually happens is that people grab whatever image they see, and just use it. 85% of images are actually ‘stolen’ this way [8]. If the actual rights holder is an individual, he probably won’t know about the infringement. And if he did know about it, it’s a giant legal pain to take action about it. Corporations and organizations that create content for money (e.g. pro photographers, media sites) *do* take action. They publish the images (e.g. as part of a news story), and then when a license violation occurs, they send the infringer a notice asking for payment (e.g. $350). It’s like a supermarket waiting for shoppers to stuff apples in their pockets and leave the store, then search passerbys’ pockets for the stolen fruit and hit them with a legal notice demanding payment. What a backwards business! So much easier to walk into a supermarket, get your apples, go to the checkout, and pay.

Arguably, much of the theft happens because it’s (currently) inconvenient to do it any other way.

2.2 Fun with Rootkits

Remember the Sony rootkit issue? Let me summarize a typical user experience:

- Step 1 – buy a Ricky Martin "Life" CD (you know it!)
- Step 2 – put it in your computer to rip or play it
- Step 3 – a rootkit is installed – takes control of your computer at the most fundamental level
- Step 4 – You're now exposed to serious security holes in the rootkit, so hackers can steal your stuff. Just because you wanted to enjoy Ricky Martin.

Rootkits are small, usually auto-executable programs that when run, take over the host computer, and perform actions as if they were 'root' or the administrator, usually unbeknownst to the machine's owner. Rootkits are generally considered a hacker tool, but Sony installed them on music CDs to try to spot people who were pirating their music by preventing the installation of CD copying software and keeping any other software from accessing the Sony CD. The rootkit software installed automatically and had no uninstall feature, and also contained serious security flaws that could allow hackers to access any computer with a rootkit installed.

Alas, it wasn't just Ricky Martin, (you might be thinking that your good taste has kept you safe.). But they also got Our Lady Peace, a beloved Canadian rock band. And Celine Dion, Neil Diamond, Natasha Bedingfield, and more. In fact, Sony shipped 4.7 million CDs with their rootkit software, and 2.1 million of those were sold before the scandal forced a recall. That's a lot of rootkits. Even the Department of Homeland Security was concerned with the security vulnerabilities the rootkits created. It issued a statement saying media companies need to remember: "It's very important to remember that it's your intellectual property — it's not your computer."

DRM – the idea of locking down a file with a key – just doesn't work. Either there are unintended consequences as in the Sony example or they're so easy to crack, that they're effectively useless. DVDs are another great example: millions were spent developing a DRM for DVDs. Then it got hacked, within days of the release of DVDs. The studios said that it "must be a big organization to pull that off". Nope, it was a kid. A young hacker named DVD Jon [9]. Once the cat's out of the bag, it's out for good. All DRMs can be broken and it takes a fraction of the resources required to make it to break it. The proof is in the torrents.[2]

2.3 "Not Available in Your Country"

Trent is an English speaker and lives in Germany. It's 2015. Can he go online and find new release movies to watch in English, legally? Nope: iTunes Germany is in German, and iTunes USA detects that he's located in Germany and blocks him. Amazon Germany is in German, not English, even if you're an English speaker. Netflix works but it doesn't have new release movies, so the illustration in figure 3 is an

[2] Torrents are the files that are used in peer-to-peer file sharing, and are a common way to pirate movies, book or songs

all-too familiar sight.

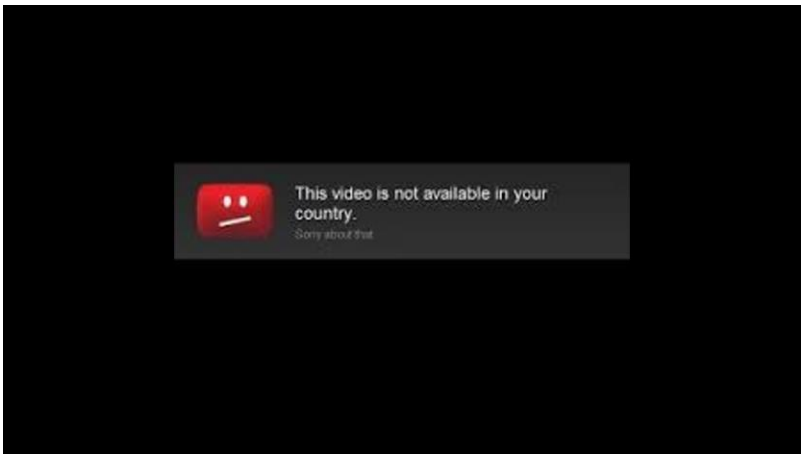

Figure 2: A common site on YouTube (www.youtube.com). Content is restricted to countries or regions.

Region coding on discs are another example. This burns consumers for both movies and video games. DVDs typically allow you to switch regions five times and no more. This is a huge pain for people who maintain residences in different countries or travel internationally frequently.

Video games – even worse: Xbox and WII games are encoded to only work in the country of purchase. If you buy a game in Brussels, you cannot play it in Canada.

The region codes are explained as a means to force consumers to abide by the local copyright laws. Though it seems as if the main motivation is to try to prevent piracy. But the anti-theft logic doesn't work; pirates selling ripped DVDs everywhere from Richmond, BC to Khao San Road, Bangkok simply copy both the US and European versions or just remove the regional encoding. You can also buy players that play all discs, regardless of region, by simply ignoring the codes. It's another example where the consumer is penalized for fair use of their purchased content without having to resort to serpentine circumnavigation of the technology or just getting fed up and flat out pirating the material.

These are examples of a much larger problem: while the Internet offers global connectivity, it ends at the legals. The World Intellectual Property Office (WIPO) is a body of the United Nations (UN) that creates guidelines for intellectual property (IP). But it's up to each country to implement. The implementations can vary wildly. And they're very rarely consumer-centric.

For example, the USA's implementation is the dreaded DMCA (Digital Millennium Copyright Act), which is far more onerous than the WIPO guidelines, due to excessive lobbying by US media conglomerates.

The one thing all of these implementations have in common is that it's hard for someone without a legal education to know their rights and responsibilities. And even for those who do, it's often nearly impossible to exercise what few rights remain to them without an obscene amount of legal fees to fight a case that they're likely to lose anyway. For example, the DMCA allows a user to make a single backup copy of any digital media that they buy for their own use, but most DRMs stop consumers from performing this legal action.

"DRM tries to defy the physics of bits."

Bits can be copied, shared, stored, and viewed essentially for free. In trying to lock down bits, DRM tries to defy the physics of bits. But there are gaping holes. The most obvious (and least talked about) is the "analog hole": if I want my own copy of a piece of media, I simply press play on a device, then record with a physically second device. DRM hasn't worked, and consumers have paid the price. The Internet needs a system of ownership that acknowledges the physics of bits.

2.4 How To Collect Digital Art?

Collecting *physical* art is straightforward: an artist makes a painting, and consigns it to a gallery; the gallerist puts it for sale in an exhibition; the collector walks into the gallery, chooses a work, pays for it; the gallerist gives the physical painting to the collector to take home; the collector hangs it on his wall. The collector is said to 'own' the purchased painting. Possession of the physical artifact has always simplified the transfer of art as tritely stated in the old adage: "possession is 9-10ths of the law." This practice started before the Renaissance. The Mona Lisa is 500 years old – older than most nations. There's no question of who owns it. (There was however, a period where a Louvre janitor stole it, in order to "keep it safe.")

It's not so easy when it comes to digital art, when what you're buying is bits. How does a collector actually collect digital art in a way that it can be re-sold? How does one establish provenance of a digital file, if the file can be copied at will? One could try to lock up the file, DRM-style, but this is fighting the physics of bits and is horrible for nearly everybody. Unless you like rootkits.

This problem of digital art ownership has been called "the elephant in the room" of the art world:

> *"the question of ownership — and how you get someone to pay notoriously high art-market prices for something as relatively immaterial as Molly's webcam video or a 24-second YouTube clip — is still unsolved, and what the organizers of Paddles ON repeatedly called "**the elephant in the room**." But curator Lindsay Howard waxes poetic about giving digital artist both the recognition and the dollars they deserve." --The Verge, 2013 [2]*

The workaround used by digital artists is to try to induce physicality in some way. This has roots in photography, sculpture, and engravings. In photography, the owner of the IP kept the negative, and collectors would get physical prints. The photographer chooses a fixed number of editions in advance, in order to create scarcity. The limited-edition concept has roots long before photography: industrial-era bronze statues had molds from which multiple bronze "prints" were made (there are >20 copies of Rodin's "The Thinker" [13]); and prints from wood-block etchings go back centuries as well. Figure 3

illustrates.

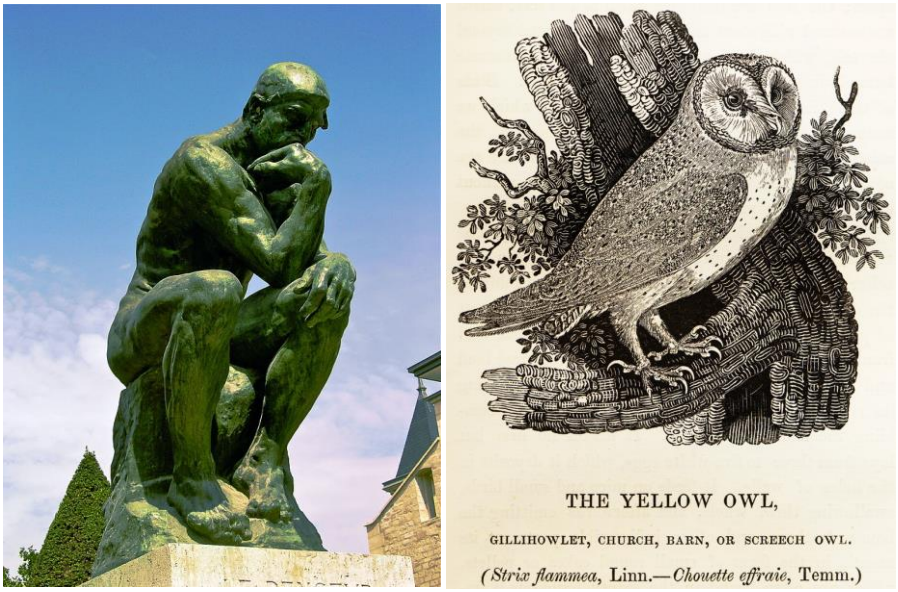

Figure 3: Left: A bronze statue "print" of The Thinker [14]. Right: A wood engraving "print" of Barn Owl [15].

In the world of digital art, the "physicality" workarounds include:

1. Embedding the idea into an object. E.g. a TV + VCR. Literally gluing the VCR lid shut. Or a stack of TVs. Or a phone (bitforms).
2. Making the idea thing physical: a print, 3D printed sculpture, etc.
3. Putting it into a USB stick or CD. Maybe creating an ad-hoc Certificate of Authenticity (COA). Put the USB stick in a fancy case. Signing the CD. Encouraging the owner not to share.

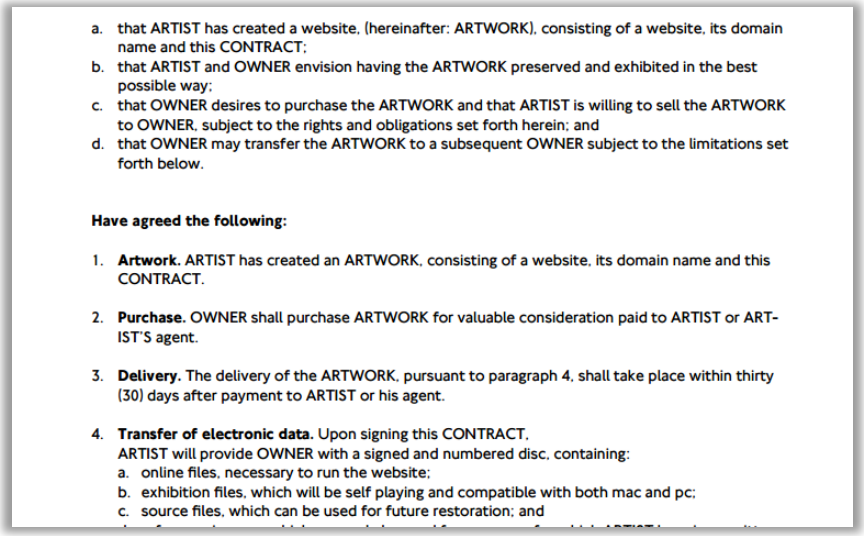

a. that ARTIST has created a website, (hereinafter: ARTWORK), consisting of a website, its domain name and this CONTRACT;
b. that ARTIST and OWNER envision having the ARTWORK preserved and exhibited in the best possible way;
c. that OWNER desires to purchase the ARTWORK and that ARTIST is willing to sell the ARTWORK to OWNER, subject to the rights and obligations set forth herein; and
d. that OWNER may transfer the ARTWORK to a subsequent OWNER subject to the limitations set forth below.

Have agreed the following:

1. **Artwork.** ARTIST has created an ARTWORK, consisting of a website, its domain name and this CONTRACT.
2. **Purchase.** OWNER shall purchase ARTWORK for valuable consideration paid to ARTIST or ARTIST'S agent.
3. **Delivery.** The delivery of the ARTWORK, pursuant to paragraph 4, shall take place within thirty (30) days after payment to ARTIST or his agent.
4. **Transfer of electronic data.** Upon signing this CONTRACT, ARTIST will provide OWNER with a signed and numbered disc, containing:
   a. online files, necessary to run the website;
   b. exhibition files, which will be self playing and compatible with both mac and pc;
   c. source files, which can be used for future restoration; and

Figure 4: A snippet of Rafael Rozendaal's art website sales contract [19]

All of these have been tried, none of these are wholly satisfactory, precisely because of the physicality aspect. For something that is naturally digital, we should be able to consume and share it digitally. It's 2015, we live our lives digitally. Sharing digitally is not just a fad, it's a way of life. Witness the social networks – 1.9 billion images are shared daily [10], 300 hours of video are uploaded every minute on YouTube [11], and so on. Most of us customarily show off pictures on our smartphones and tablets, never printing them. We watch movies streamed by boxes directly to our television sets. Art however, is still bound to the physical representation, because that's the only way that the creator currently has to be sure that her rights are not violated.

Besides the "physicality" workarounds, there's actually one partial solution to owning digital art—websites. If a person has a work of art on web domain x, then via a contract the owner of domain x can be the owner of the art. Figure 4 gives an example. It's a simple idea but a powerful concept. It's possible because there is *scarcity* on domain names – only one entity can own a domain name, e.g. amazon.com. The technology enabling the scarcity is the Domain Name System (DNS) [34], which at its core is simply a mapping from a domain name (e.g. amazon.com) to a number (e.g. 54.93.255.255). People trust the DNS because no one really controls it; it's administered by an organization called ICANN. Web domains have digital scarcity via a public ledger that is requires little extra trust by the user (trust minimizing system). It is a consensual arrangement and in fact, anyone could create an alternative registry tomorrow that would work in a similar manner. It would be near useless however, because there is too much of a critical mass of users that have already voted with their checkbooks and purchased and used domain names that work within the existing DNS system.

Websites for art are not a complete solution, however. Many digital art artifacts are not naturally websites, they're intrinsically images or videos or screensavers or some other code executing on some platform. To force such media into a website is contrived, like jamming a square peg in a round hole. Also, it can get expensive and annoying, as the owner is forced to pay continued registration fees every year. If the domain name contract is not renewed, the art may get removed. Websites also do not naturally support editions; there can be only one collector. Finally, it's not obvious what the best place to store the metadata is; in particular the ownership history (provenance) of the owners which is crucial to future values of the piece. It's difficult for collectors to resell a piece unless there is a good provenance record.

Because the question of owning digital art has remained unsolved or only partially solved (depending who you ask), the digital art market remains tiny. The size of the whole art market is $64B [18], which is comparable to video games ($75B) [16], film, or online advertising. The online art market was $1.57B, growing at 25%, expected to be $3.7B by 2018 [17]. In contrast, the 2014 digital art market is only currently estimated at 11% ($173M) of the online art market [17]. Photography is 46% ($720M) of the online art market [17].

What's interesting is that in the art world (and much of the world in general), there is confusion about the idea that there is even such a thing called "intellectual property" (IP). Questions like the following are commonly asked: "How can you own something intangible?" and "Is the file an object?" But, this is actually resolved by the *very nature* of "intellectual property". The "intellectual" part means that it's exactly about the ideas, and the "property" part is exactly about the ownership. IP is *exactly* about what the art world would call 'ephemeral.' IP was invented specifically to protect the creators – the artists, the writers, the musicians, and the like. Patents were protected to protect the inventors, to give them a

chance to monetize their invention, before the world at large could use it an unrestricted[3] and unlicensable manner. Copyright refers to content itself and describes the licensing rights (subject to the host country's copyright laws) and usually expires after a period of time. The origin of these laws date back to a battle in Ireland circa 1600, where 3,000 people died over what was essentially copyright rights to a book! [20]

2.5 How Do You Share 3D-Printing Designs?

3D printing is taking off, with a ton of creative uses. Many futurists believe that it will be the explosive technology of the future, with every consumer having a printer in their home to make purchased artifacts on site. Perhaps somewhat similar to a rudimentary Star Trek replicator. But if you look at the consumer 3D-printing design marketplaces like Shapeways (www.shapeways.com) and Thingiverse (www.thingiverse.com), very few professional designs are being shared. Currently, it's only hobbyist stuff.

> *"my conclusion is that **whatever you put on the internet you lose it**. Maybe keep the rights, but lose the power over it."*
>
> *-user on Shapeways blog [35]*

What's going on? It comes down to "3D printers will be to patent and trademark what Napster was to copyright" [21]. When Napster came along, every dorm in America was engaged in copying files, and routinely infringing copyrights. Content companies started demanding settlements for thousands of dollars, with the threat of even bigger judgments and legal fees if the complaint wound up in court. Lawmakers struggled to keep up. The law in all countries lags considerably behind the technology because of both the nature of the process and the ability of the lawmakers to understand and sufficiently generalize the new capabilities. This latency usually ranges from between 5 – 10 years. The Napster brouhaha made it clear that: (1) people who had never considered copyright law before would have to start paying attention, and (2) existing copyright law was poorly suited to enforcing non-commercial infringement by a large number of individuals.

Just as individuals who had never thought about copyright were blindsided by industry and legals, 3D printing will bring trademark and patent law into homes. Trademarks are symbols or word(s) representing a company or product. Patents cover inventions. Like copyright, trademark and patent laws are poorly suited to dealing with mass infringement. This will lead to litigation, followed by new precedents, then lobbying and changes to laws. Along the way, there will be plenty of confusion and misinformation and as in the previous examples the losers will be the honest consumers. The less honest ones will do what they have before, ignore the copyright laws and just pirate what they want.

While it's impossible to predict exactly how it will play out, the fight between Napster and copyright law may prove a bellwether of the 3D printing legal wrangling to come.

[3] Patents were offered as a compromise - if you tell the world what you've built, we give you a monopoly for 20 years. If you don't tell the world, if it gets out, you have no protection under the law.

In short, creators are reluctant to share on 3D printing marketplaces because the framework to protect their creations is immature. A better solution needs to be presented to this burgeoning and potentially lucrative marketplace.

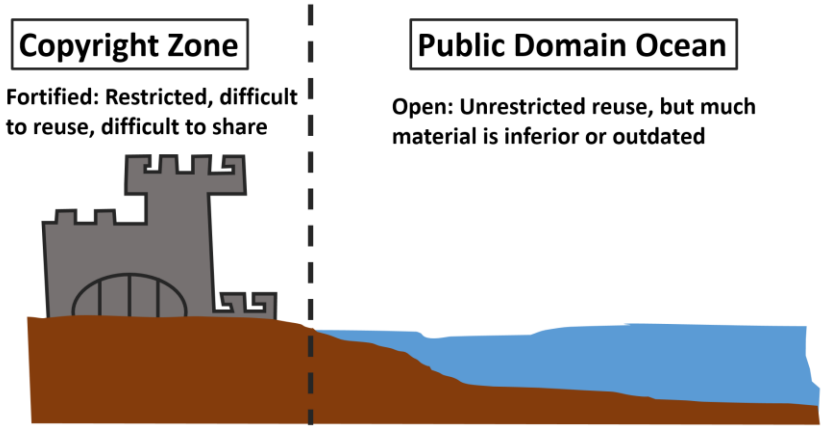

Figure 5: The internet's two ownership zones: copyright and public domain, each which has issues. Adapted from [36]

2.6 Two Zones on the Internet

All the issues we discussed so far are symptoms of a larger, broader issue: ownership of digital content is broken, especially on the internet. As Figure 5 illustrates, two ownership zones on the Internet have emerged. One is the "Copyright Zone", where content is restricted, hard to reuse and difficult to share (unless you want a letter from the RIAA). The other zone is the "Public Domain Ocean" which has open, unrestricted use, but much of the material is inferior or outdated (maybe Ricky Martin "Life" is open? Nope – copyright in the US is life of the creator plus 70 years, so "Life" will be protected for 70 years after his death.) Not only that, but the effective lifetime of copyright and trademark protection is a moving target in many countries. In the United States, the so-called "Sonny Bono trademark law" extended the protection, primarily to avoid having Disney's Mickey Mouse trademark move into the public domain. The mouse again becomes free in 2023, so expect to see similar shenanigans in the United States a few years before the next deadline.

With the current framework, all stakeholders involved have challenges:

- **Creators** – it's hard to get fairly compensated. If you have 1M views for your YouTube videos, making only 20 bucks may not seem fair. Furthermore, sharing means losing control. It's like having the title and keys to your car, but not knowing where it's parked: in theory you own it, in practice you can't use it in the intended way.
- **Collectors** and audiences – people who want to use content legitimately and pay a fair amount, whether it be a photo or a movie, should be able to negotiate for various cost-calibrated licenses, rather than being forced to choose between nothing and piracy. It should be possible to re-sell a piece of digital property with confidence, whether it be a digital art work (thus establishing provenance) or a concert ticket. It should also be noted that existing license

strategies are antiquated. The proliferation of digital devices leads to unanticipated usages of content. For example, screensavers from photographs and ringtones from songs were never anticipated as revenue sources.

- **Connectors** – marketplaces should be able to focus their energies on connecting the musicians with the music lovers, artists with collectors, authors with readers, and so on; rather than spending ridiculous amounts of money with lawyers simply to lock down their rights to market copyrighted material. It should be easy for a new marketplace to access great content, so the winners are the ones that do the best connecting and user experiences, rather than the ones with the best lawyers and licensing agreements. Let a thousand marketplace flowers bloom!

Where's my stuff?

This simple statement has two parts. "**Where**" means there's **no visibility**. Users don't know where and how much their digital property is being used on the internet. "**My stuff**" refers to the notion that the legals around copyright protection are painful to negotiate: it's difficult for the lay person to securely establish and transfer rights (e.g. ownership of an edition or rights to certain uses) to others. On the Internet, everyone may be able to be an artist, but only those with lawyers can make money from it. It needs to be as simple as selling a painting to a buyer in a gallery.

The Internet does a poor job of answering the question, "where's my stuff?" Ownership of digital property is a mess. The next section explores how we arrived at such a situation in the first place.

3 The Internet and Its Discontents

3.1 WWW History

This section explores why ownership of digital property on the internet is in such a sorry state. If we look closely, we see that much of it is due to the 1989 design of the World Wide Web (WWW). For all the WWW's strengths, it does have its flaws. Let's dive deeper.

In the 1960s, the US Department of Defense (DoD) funded work on packet network systems. One of these was the ARPANET, which would evolve and eventually use the Internet Protocol [22][23]. Usage grew exponentially, with developments like email and ftp making communication and content accessible to the desktop workstation.

In March 1989, Tim Berners-Lee wrote a proposal for "a large hypertext database with typed links" [24]. He considered several names, including "Information Mesh", "The Information Mine" or "Mine of Information", but settled on "World Wide Web" [25]. (Imagine an alternate universe where "Minecraft" is a Web editor…)

By Christmas 1990, Berners-Lee had built all the tools necessary for a working Web [26]:

- the HyperText Transfer Protocol (HTTP)
- the HyperText Markup Language(HTML)
- the first Web browser (named WorldWideWeb, which was also a Web editor)
- the first HTTP server software (later known as CERN httpd)
- and the first Web pages that described the project itself.

This was an amazing one and a half years of work. Like his fellow countryman Isaac Newton, Berners-Lee was knighted for his work. Figure 7 shows example output from an early non-graphical web browser (Lynx). Figure 6 shows the very first page of the WWW that was written by Berners-Lee in 1990 and is still maintained by CERN (page was rendered on Chrome).

Since its invention, the WWW has grown profoundly. The WWW slice of the internet has revolutionized so many aspects of our lives, it's hard to keep track. Figure 8 shows a snapshot of connections on the WWW. Today there are 48 billion Google-indexed web pages [31], 280 million domain names and just short of a billion websites.

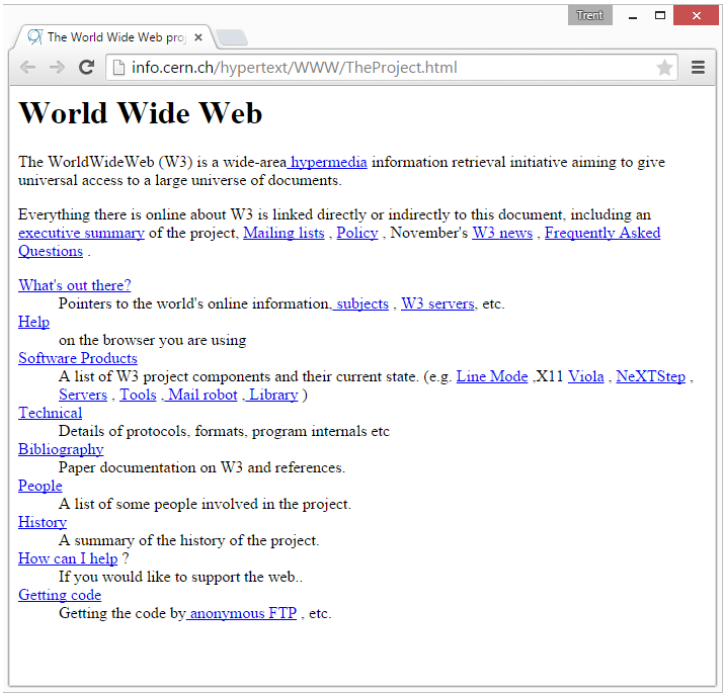

Figure 6: The very first page of the World Wide Web [28], rendered on a modern browser (Chrome) [29]

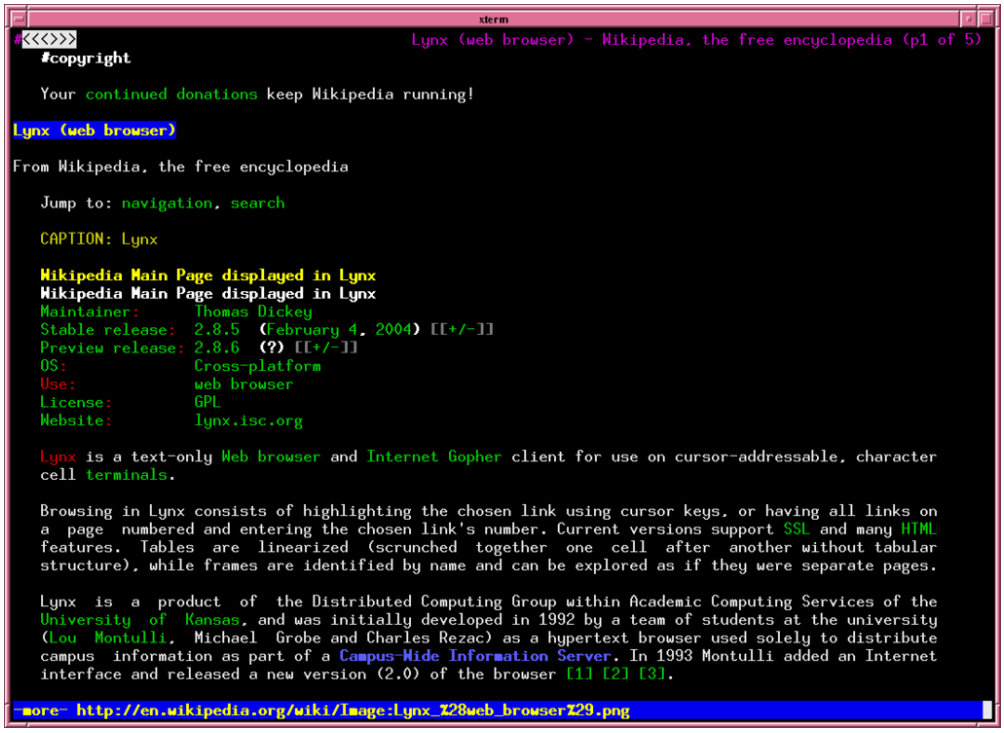

Figure 7: Rendering of wikipedia.org with an early (line mode) web browser called Lynx [27].

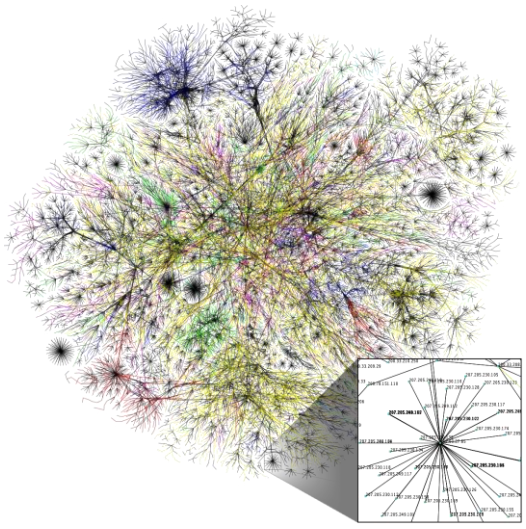

Figure 8: Partial map of the Internet based on the January 15, 2005 data found on opte.org. Each line is drawn between two nodes, representing two IP addresses.[30]

3.2 Coding Horrors and Uni-Directional Links

But… (and this is a big but and I cannot lie) it turns out that a design choice of the WWW is at the heart of ownership issues on the web: uni-directional links. Let's explore with an example.

Let's say that we create a work of art, like in Figure 9. We like it and want to share it with my friends. We want to put it online. We have these options on how we'd like others to use or not use it:

- First, we might not state anything, which would leave potential users confused about the source of the image, and how they might actually re-use it in a legal way.
- Second, we could give the image to the public domain, ideally with legals to make it perfectly clear, such as a Creative Commons license [32].
- Third, we could declare that we are the copyright holders ("© Trent McConaghy and David Holtzman 2015"), which implies that for others to use it in some contexts (e.g. a commercial website), they must ask our permission. From that, we could grant them certain rights under copyright, for example the right to publish it on their own website, maybe for a fee.[4]

Let's say we choose option 3, and put the work online. There's a few things that could happen next (Figure 10).

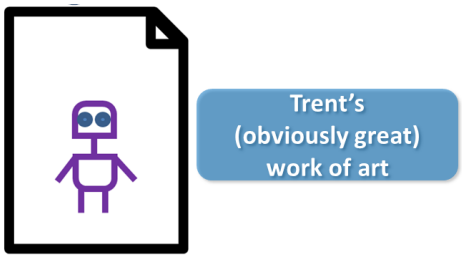

Figure 9: An example work of art. A robot, of course. Let's say we put it on a web page.

First, someone could copy the image and put on their own web page, perhaps even post it for sale, and not give Trent any attribution (Figure 10 top). We wouldn't know about this, unless perhaps we did regular runs of Google Image Search -- a huge pain. The ease-of-use balance is tilted towards the person who's using the work without consent. This would be ok if we've made our work public domain, but certainly not ok if we're trying to make a living (or at least beer money) with our art work.

This scenario could continue: someone else could copy the image and attribute (and even pay) the first copier (Figure 10 right.) We've heard many stories of this happening, usually by frustrated artists and graphic designers. The ultimate injustice is someone else getting paid for your work with no compensation or attribution.

In another scenario, let's say someone copies the image and does give proper attribution, e.g. give our name etc. and ideally also a url to the original work (Figure 10 bottom center). Great! However, because the WWW only has uni-directional links (pointing just one direction), we *still* wouldn't know about the usage.

The WWW on its own does nothing to address "where's my stuff": it doesn't address visibility of usage of content, and it is completely ignorant of legals.

The uni-directional links of the WWW are both its greatest strength, and its greatest drawback.

[4] For purposes of this example, let's pretend that we all live in the same country with the same copyright laws.

"The current world wide web does basically one thing: simple, stupid, mindless hyperlinks. But even that alone was enough to build a functional and useful internet for the world."

-Jeff Atwood, Coding Horror Blog [37]

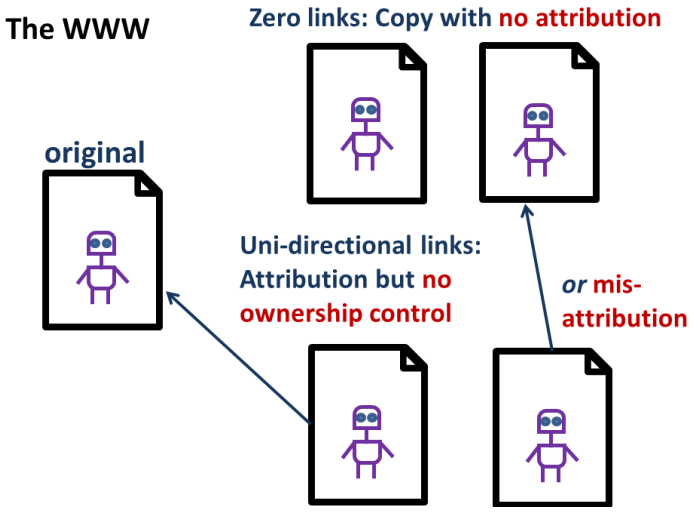

Figure 10: When art is shared on the WWW, various attribution issues arise.

3.3 Pre-WWW History

Does the WWW idea of hypertext have to be the only way? This section examines history of the Internet going back much further than the 1989 design of the WWW. In a vision from 1965, Ted Nelson described a system:

"[Consider] a unified ... service that would provide storage and publication services, and manage ... royalty payment on a ... fair basis that would facilitate unrestricted virtual republishing"

-Ted Nelson, on a vision from 1965 [36]

It's amazing that this specification from a system dates from 1965. This is no coincidence. Ted Nelson was in fact one of the leading thinkers on hypermedia in the 1960s, and remains that way today[5]. He

[5] As of April 2015, Ted Nelson continues to work on Xanadu, with as emphatic a vision as ever [40].

coined “hypertext”. In this vision, Nelson was looking out for the creators, and the audiences.

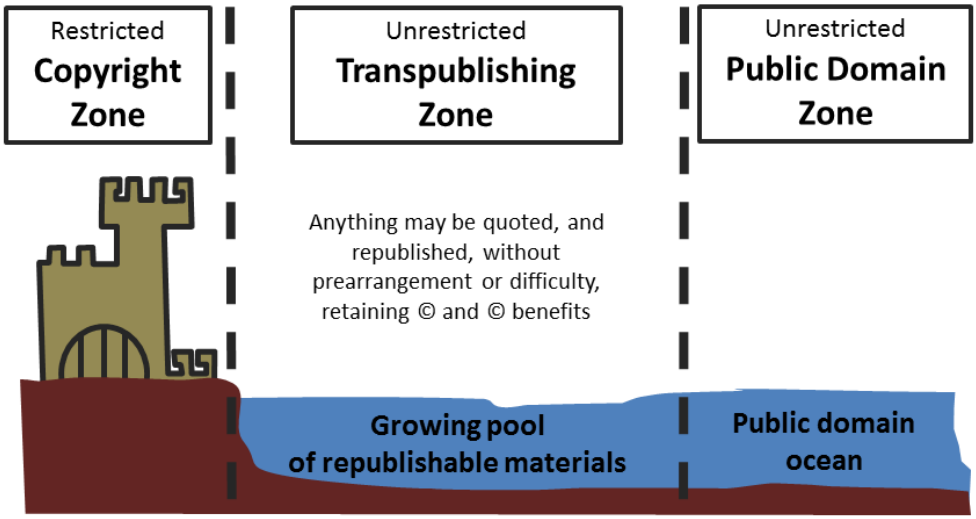

Figure 11: The Xanadu vision was (and is) to create a middle ground between the restricted copyright zone and the unrestricted public domain zone

Nelson imagined “a **new middle realm**, one which renders copyright benign and flexible ... a **win-win system, as it is beneficial both to rights holders and to users**, in a way that other copyright systems are not beneficial to users.” In this so-called “transpublishing zone”, anything could be easily reused, while retaining copyright and copyright benefits. This new zone would get the best of both worlds: copyright rights that are actually easy for the lay person to use. Figure 11 illustrates. Nelson published this vision in 1965 [38].

To realize this vision, Nelson designed a hypertext system with two key elements: bi-directional links and baked-in copyright (Figure 12 left). On the “where’s my stuff” question, the bi-directional links ensured that the creators had visibility, and the baked-in copyright handled the “my stuff” part (Figure 12 right).

From the initial designs, Nelson and colleagues worked to create prototypes of this hypertext. Figure 13 left shows a mockup of Xanadu, focusing on the bidirectional links.

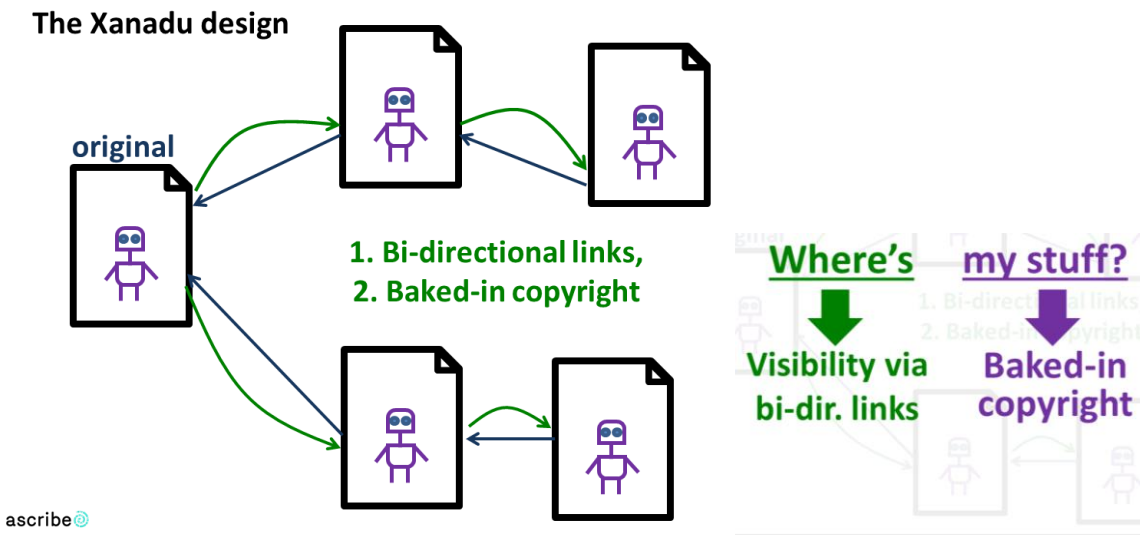

Figure 12: Left: the Xanadu design links hypertext pages in both directions. Right: the Xanadu design addresses the "where's my stuff?" question

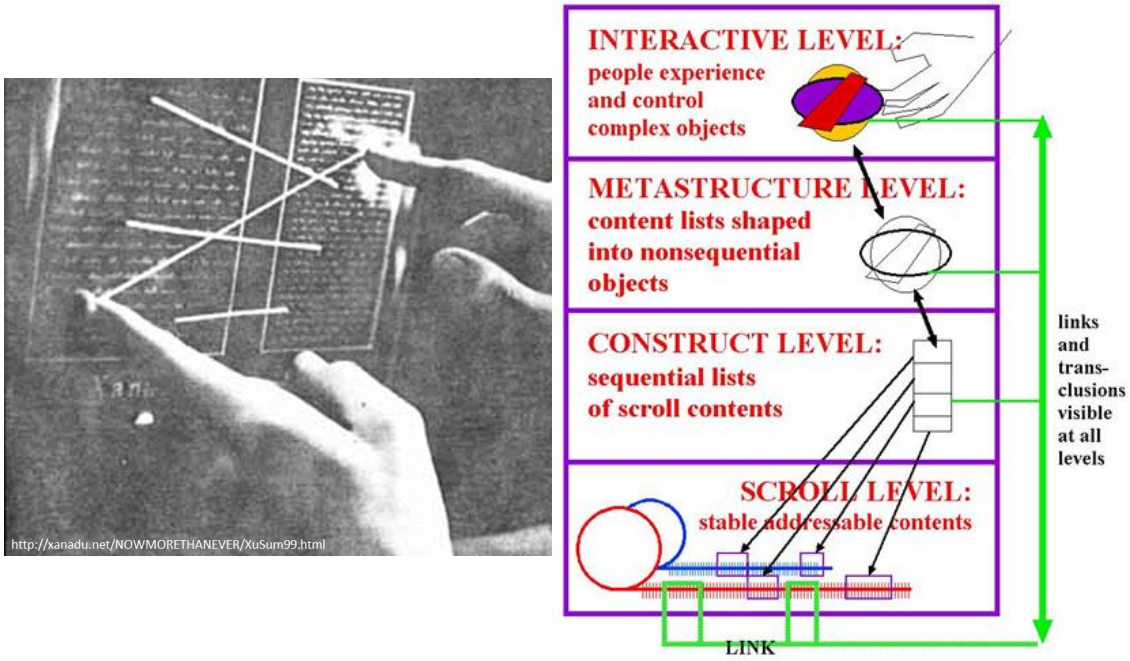

Figure 13: Left: 1972 mockup of Xanadu. Right: relation among levels of the Xanadu design

While Xanadu had grand vision, the implemented design turned out to be quite complicated. Figure 13 right illustrates the four levels of the design: interactive, metastructure, construct, and scroll. All these levels are connected by links and transclusions (snippets of text). All this complexity meant that Xanadu was very hard to build in its full-featured vision. Nelson kept at it: in the 1980s he actually started a company to work on building it. That company got bought by Autodesk a few years later. But building it proved so elusive that it acquired the dubious reputation of being the longest-running case of vaporware ever, even more than Duke Nukem Forever [33].

When the WWW was released in Christmas 1990, Nelson and his team were still at Autodesk. The WWW took off rapidly, however, and left Xanadu in the dust. A few years later, Autodesk disbanded the Xanadu project [39]. There's much to be said for something simple and deployable. The Xanadu project continues in fits and starts to this day, in an open source form [39].

> ***HTML is precisely what we were trying to prevent -- ever-breaking links … no rights management"***
>
> ***-Ted Nelson [36]***

So now we have the WWW, warts and all. It's shipping, it works, though it has obvious flaws.

So far, this paper has reviewed a host of symptoms, where with a broad brushstroke we can say that ownership of digital property (especially on the internet) is a mess. This is *despite* being anticipated since the 1960s, *and* designed for. But, the simplicity of the WWW won out. However, this means that the "where's my stuff" question remains unsolved on today's internet. Or is there a way? ….

4 Reviving the Vision of Ownership on the Internet

4.1 Introduction

Can something be done about ownership on the internet? The Internet and the WWW are widely deployed. There are billions of web pages. So, redesign of the WWW is not a practical option.

But: **can we *retrofit* the Internet for ownership?** (And realize the Xanadu aims in the process). This is a more practical undertaking. To get going, we don't need to rewrite Internet protocols. Rather, we can work with existing infrastructure and overlay new ownership meta data. We need to answer the question "where's my stuff", and unpacking that, we need to answer "where" (visibility) and "my stuff" (intellectual property).

Figure 14 illustrates our proposed design. It answers each half of the "where's my stuff" question by leveraging modern technologies and exploiting the existing legal system. For the "where" aspect, it automatically discovers the bi-directional links. For the "my stuff" aspect, it makes legals not only easy, but secure. Let us elaborate.

4.2 Addressing "Where"

To address the "where" half of the "where's my stuff" question**,** the ascribe approach is the following:

- Crawl the entire internet (220 Tb text)
- Similarity match against creator's content (10G+ images, 3D designs, ...). This is a machine learning problem, at Internet scale.

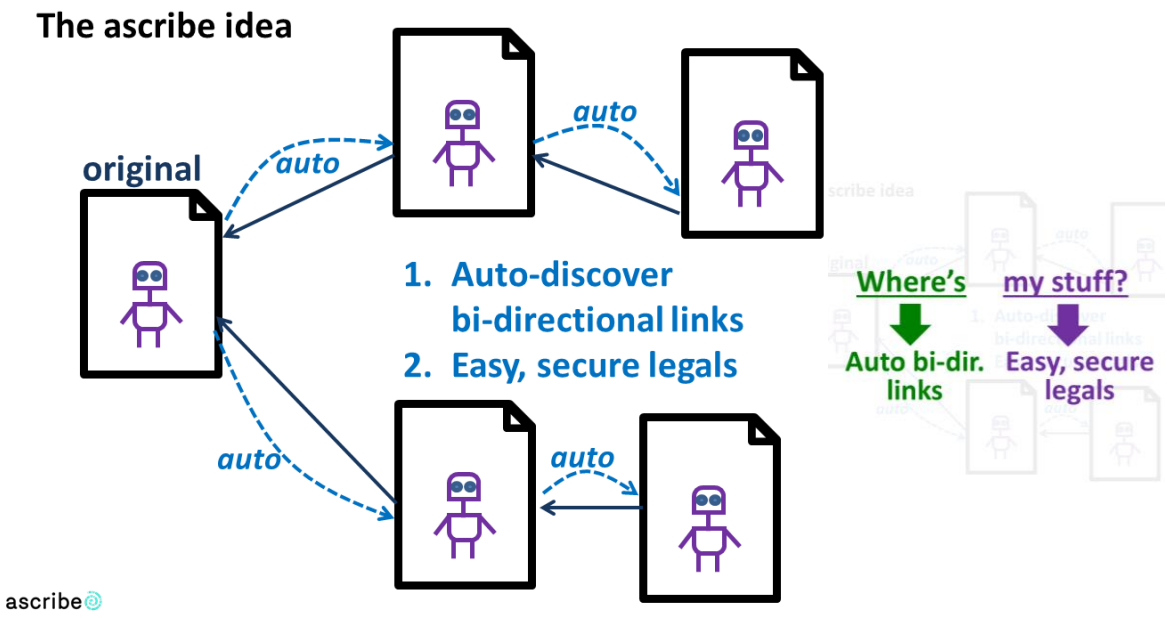

Figure 14: Left: like Xanadu, the ascribe design links hypertext pages in both directions. Right: like Xanadu, the design addresses the "where's my stuff?" question. However, the ascribe design employs very different technologies than Xanadu to achieve these goals.

The result is bi-directional links (just as Xanadu targeted), yet without the complexity of web page creators or maintainers having to maintain their own forward links to anyone that might reference them, and without the complexity of changing existing web protocols. Any solution that would result in bi-directional links will have to be done without requiring existing website owners to change anything. Our solution requires no cooperation from any third party, only the second party signatories to our service.

This is a machine learning similarity search problem. Figure 15 shows example results from ascribe's image similarity search engine. The top image is the reference image. Ascribe searched against this image and found that the most similar image (left, middle) was an exact match, i.e. distance of zero to the reference image. The image to its right is the next-most similar, also with a skyline, and brightness fading from left to right. Each subsequent image to the right (with wrapping) is the next-most-similar image.

Other companies do image similarity search, such as Google image search [41] or TinEye [42]. But the value proposition becomes truly novel when tied with *ownership* goals. We do similarity matching as reinforcement of owner rights, not as a curiosity feature for a search engine. Because this technology is intrinsic to our business model, ascribe will always devote considerable amounts of its programming resources to enhance this functionality across a broad swath of media types on behalf of the license owners. There are several use cases that emerge specifically for ownership, all which address the "where" question:

- **Tracking:** know who's using your work. Then it's up to the rights holder to make the call on what to do. Options include: letting go, asking for licensing fees or a tip, or a takedown request.
- **Screening:** when a person tries to register a work, cross-reference against other works that already exist to avoid duplicative registration.

- **Measuring value:** the number of shares (and where they're shared) can be a proxy for value. This can be incredibly useful to the rights holder, or the consignee who's helping to sell the work. This is only true some of the time: an artist selling unique editions can charge more if there are 5 million views rather than 50; on the flip side it may be harder to charge for a movie rental if there are copies everywhere.

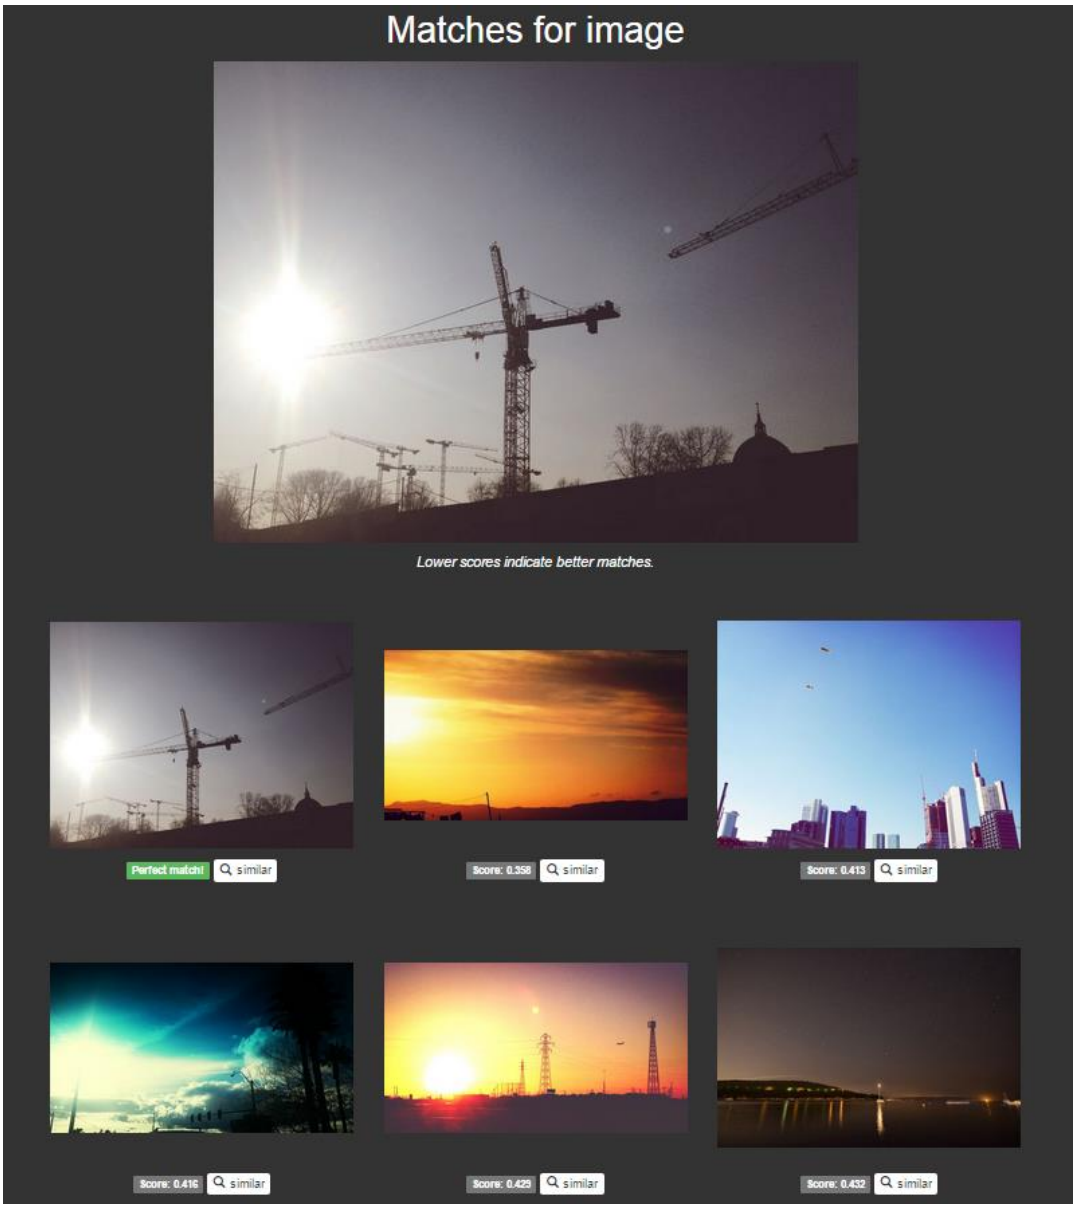

Figure 15: Example results of image similarity search

4.3 ADDRESSING "MY STUFF"

To address the "my stuff" part of "where's my stuff", the ascribe approach is **a registry for easy** & **secure legals**.

4.3.1 Easy Legals via TOS

The ascribe Terms of Service (TOS) [45] makes the legals easy to use by creators and consumers. It incorporates actions like “I claim copyright rights” (creator, when registering a work into ascribe) and “I transfer copyright rights, or a license” (old owner to new owner). One can think of this as copyright in a box, easy to use by the layperson (creator, consumer). Artists are not usually programmers or lawyers. Ascribe handles the complexities of the licensing process, leaving them free to do what they do best.

Ascribe worked closely with lawyers in several countries as well as its own in-house general counsel to develop a TOS that addresses these issues.

PayPal enabled electronic transfer of money between individuals, allowing money to flow across the internet from and to non-corporate end-users in the ‘long tail’ [46]. Similarly, ascribe’s legals enable electronic transfer of digital property between individuals, allowing digital property to flow across the internet and into the long tail.

4.3.2 Evidence of copyright via blockchain

We secure ownership transactions by time-stamping evidence of ownership actions onto a *trusted registry* using existing and proven technology—the bitcoin blockchain [1].

The blockchain can be viewed as a database where anyone can add information, but no one can delete it and no own “owns it” [49]. The bitcoin protocol itself is geared towards payment processing [1]. Of course, our interest is not payments, but ownership transactions.

There are other protocols that are designed to pass around non-financial tokens, such as OpenAssets [43] and Counterparty [44]. However, the existing protocols do not support our needs for ownership processing for two reasons: the feature set is different, and the legals are different:

- **Feature set is different.** We need to distinguish unique editions, such as edition 1/10 of a photograph or print being different than edition 2/10. Other protocols consider those simply as quantities, e.g. 10 e-gold. We also need to support these ownership actions: consigning, un-consigning, loaning, renting, and more. The other protocols do not support those needs. In the future, we will be adding a richer taxonomy of limited use rights and licenses.
- **Legals are different.** Most non-financial protocols assume have the implicit assumption that the value being passed around is the coins in the blockchain itself, as if it were bitcoin. Hence the label “colored coins” that these protocols are often called. For our case, this is not appropriate. Code is not law; the blockchain is not law; law is law. We are not passing around coins that are somehow magically transformed into tokens of value. The ascribe terms of service already handles the legals of the ownership transfer. If the blockchain is going to be used, it must provide extra value. As we discuss shortly, this value is in the time-stamping.

Existing blockchain protocols were not sufficient to meet our needs. So, we developed a new protocol called SPOOL. SPOOL is a protocol for documenting transactions relating to ownership of digital property. SPOOL stands for Secure Public Online Ownership Ledger. SPOOL is used for *time stamps on ownership transactions*. These recordings can and will be used as evidence in court in case of an ownership dispute. Artists sometimes take a USB stick of their work, mail it to themselves, and keep the envelope sealed, all to prove that they had access to that file at that point in time. The blockchain

accomplishes the same thing, but in a much more convenient fashion. Interestingly, the blockchain has already been accepted as evidence in a court of Law, as part of the infamous Silk Road trials [2].

SPOOL is not the law. It must be used with a contract concerning transfer of copyright that participants have agreed to; in our case that is the ascribe Terms of Service (TOS)[6]. Since our protocol runs on the Internet, it creates a marketplace that flows across borders, but its power comes from its combination with local copyright laws, which provides for truly global distribution with sovereign intellectual property protection.

Because ascribe uses the blockchain, it means that the rights holder doesn't have to rely on ascribe as the keeper of the ownership records. Even if ascribe completely disappeared, all users would retain their copyright and licensing privileges because the untamperable records are all on the blockchain for the world to see. If this were not true, then people would not trust their currency to Bitcoin, which has a worldwide value of over a billion US dollars.

4.4 SPOOL Protocol Details

This subsection provides detail about SPOOL to give the reader a flavor of the protocol. The two main ownership transactions are registering a piece of work, and transferring copyright rights. Figure 16 illustrates both. A bitcoin transaction ≥1 inputs, and ≥1 outputs. The transactions have special choices for inputs and outputs.

Figure 16 left shows a *registration* transaction, when the work has three editions (1/3, 2/3, 3/3). The input address is a public ascribe address, in the role akin to a certificate authority. There are six outputs. The first output is a hash of information defining the work: the file containing the work, the artist name, the title, and the year. The second, third, and fourth outputs are the addresses "owning" edition 1/3, 2/3, and 3/3 respectively. By definition of the TOS, whoever has the private key to one of those addresses is the owner of that work. The fifth output is the verb of the transaction; we use the bitcoin OP_RETURN to embed the word "register". The final output is change. (It's a quirk of bitcoin: you have to send *all* your money from the ascribe address and then get most of it back. So be it.)

[6] We considered a Ricardian Contract [3] approach, but postponed implementation. This a contract that is readable by both computers and humans, and legally binding. This contract could be part of a protocol, such that the contract is embedded with every transaction. The potential benefit is that it allows for a decentralized store of the contract itself. But as we've discussed, going fully decentralized can hurt usability and adds complexity. Once the contract is released, it's fully baked and cannot be changed easily. This being said, there are actually ways to update the contract while maintaining decentralization [4]; we will consider these for the future.

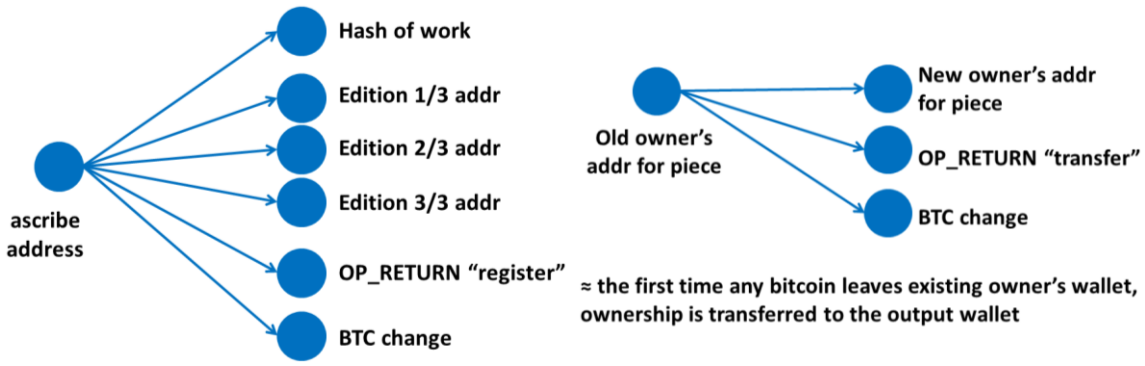

Figure 16: SPOOL transactions. Left: register transaction on 3 editions. Right: ownership transfer transaction

Figure 16 right shows an *ownership transfer* transaction; for example transferring ownership of edition 2/3 of a work from Alice to Bob. The basic idea is that the first time any bitcoin leaves the existing owner's wallet (e.g. Alice), then ownership is transferred to the output wallet (Bob). We also include an OP_RETURN to embed the verb "transfer". The final output is change.

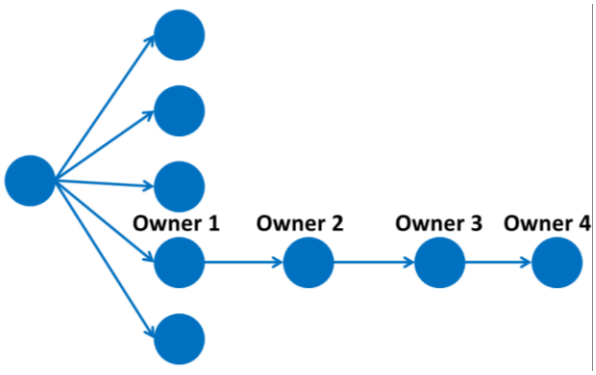

Figure 17: In SPOOL, provenance emerges naturally

In the art world, and other domains from automobiles to houses, ownership history (provenance) is extremely important. If there are gaps in the ownership history of an art work, car, or house then the value of the work is diminished. Traditionally, there was no way to track the ownership history of a digital edition. SPOOL enables this trivially, as Figure 17 illustrates. If you know the address of any of the owning addresses, then you can find (e.g. with a blockchain explorer) all the other addresses that have owned the address. SPOOL enables perfect digital provenance.

Each automobile, iPhone, and bicycle ever manufactured have their own unique serial numbers. This helps tracking the item over its lifetime. SPOOL makes it possible for each unique edition to have its own ID (serial number). Figure 18 illustrates how. As Figure 18 left shows, SPOOL defines piece ID as simply the bitcoin address of the original owner. So, even if the ownership changes, the piece ID stays the same. Figure 18 right emphasizes how a work's hash and an edition are bound: they are both defined in the same atomic bitcoin transaction (registration), and the protocol makes their relationship clear.

Figure 18: About piece IDs in SPOOL. Left: each edition of work gets a unique ID. Right: the register transaction binds the work hash to all piece IDs

SPOOL facilitates easy searching. One can take a work's hash and find its fingerprint on the blockchain, and from that find all the addresses that own each edition[7]. Or, one can start with any piece ID and find the whole ownership history. One can even start with any owning address of any edition, and from that easily see the whole history of each edition.

However, while the transactions are public, they are also pseudonymous. The owners' identities are private by default on the blockchain. The ascribe web app (and API) is also private by default. It's not browsable. However, if he wishes, the owner can make his ownership public, e.g. by tweeting the bitcoin address he owns, or a url pointing to the piece detail page on ascribe.io.

For the interested reader, the full SPOOL specification is open-sourced at www.github.com/ascribe/spool [47]. Furthermore, we have made an open-source reference implementation in python available at www.github.com/ascribe/pyspool [48].

4.5 ASCRIBE TECH STACK

Figure 19 shows the ascribe technology stack. The ascribe service is consumed by either marketplaces (top left) to bring benefits to their users; or directly by users on the ascribe web app. Marketplaces and the ascribe web app both consume the ascribe REST API. Ascribe ownership servers implement the functionality exposed by the API. The servers implement answers to "where's my stuff" via the three verticals below. On the bottom right, the "where" to auto-discover bidirectional links is by ascribe crawl data and similarity search against content. On the bottom middle is the ascribe Terms of Service to make

[7] We considered putting the hash inside the OP_RETURNs in order to facilitate blockchain maintenance. However, this would have prevented the ability to search based on a piece's hash. So we chose searchability.

the legals easy.

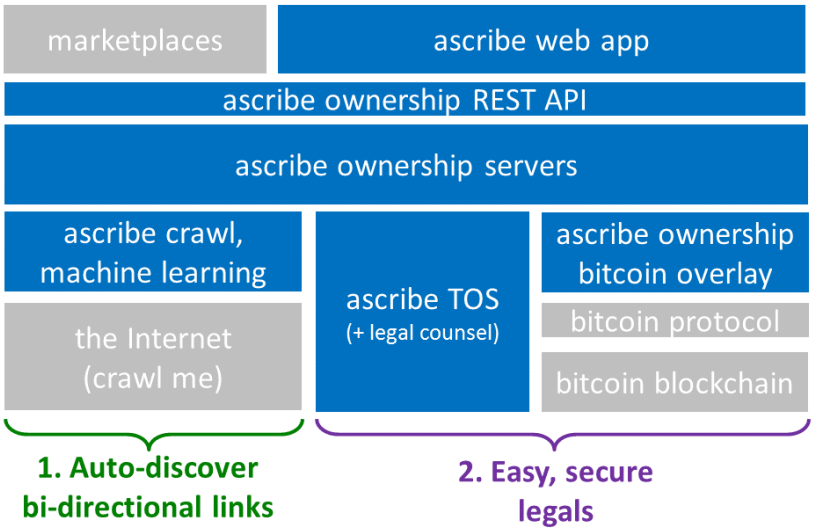

Figure 19: ascribe technology stack

The bottom right vertical of Figure 19 shows how the blockchain is used to make the ownership transactions securely time-stamped. Specifically, the ascribe servers construct bitcoin transactions according to the dialect of the ascribe ownership bitcoin overlay (SPOOL protocol). This is special subset of the bitcoin protocol. Those transactions are pushed to bitcoin servers, and ultimately stored onto the blockchain (by the winning bitcoin miner). Sometimes a bitcoin transaction doesn't go through the first time; for example different miners have different transaction-accepting rules. If the transaction doesn't go through, the ascribe system will periodically push the transaction onto the bitcoin network, until it does go through.

4.6 ASCRIBE INTERFACES

Figure 20 illustrates the three ways that users or developers can interact with the ascribe service. The web app (Figure 20 left) is for individual creators (artists, graphic designers, photographers, writers, ..) who want to register, consign, and archive their work directly. It's also for individual galleries, and for collectors. The REST API (Figure 20 center) is typically for marketplaces of digital goods (art, photography, 3D, …), archives, and other businesses to answer "where's my stuff" for their users, and themselves. The SPOOL protocol (Figure 20 right) is for adventurous Bitcoin hackers, or for other businesses looking to develop their own ownership services. At ascribe we are actively working with

other bitcoin organizations around SPOOL.

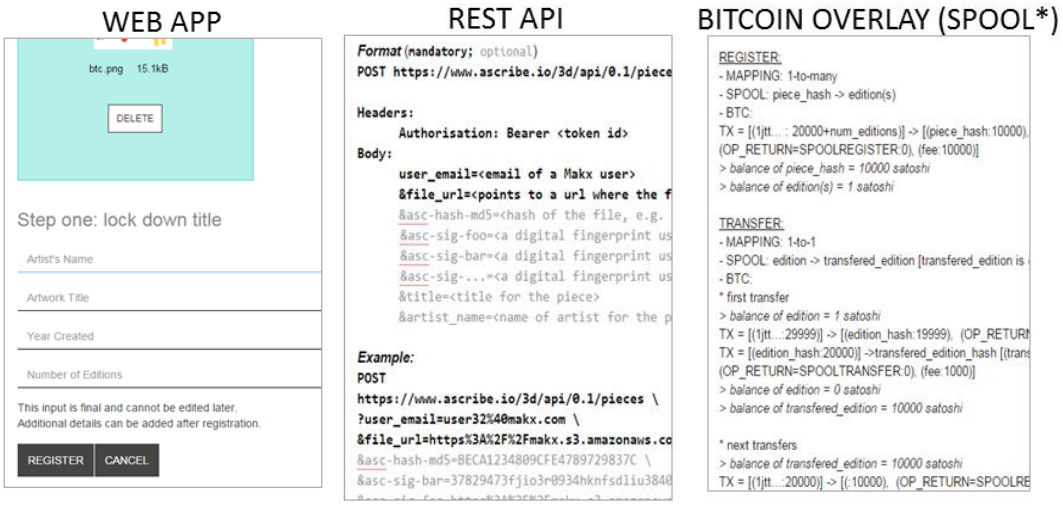

Figure 20: ascribe interfaces. Left: Web app. Center: REST API. Right: Bitcoin blockchain overlay (SPOOL)

4.7 Benefits

The benefit of this system is ownership control. When one has ownership control, there are several ways to monetize (beyond ads). These include:

- Selling limited editions (a la prints). The artist creates e.g. 5 or 500 editions of a single work, and sells each individually. The owner of each edition gets a limited set of beneficial-use licensing rights; but among them are the ability to enjoy the edition, and to re-sell the edition. In this way, the use case mirrors the traditional use case for selling limited-edition photos, prints, sculptures, and etchings.
- Streaming (like Spotify, Netflix). Users pay a fee (e.g. monthly) and in return can access content. The provider passes on the revenues to the creators. Typically, the content cannot be downloaded.
- Tipping or pay-what-you-want (a la Radiohead [50]). The artist makes the work available for all to use, without any real restrictions. But when a consumer likes the work, they can tip the creator. Blockchain technologies may facilitate micropayment-size tipping.

5 Examples of Usage

Here are some examples of how people are using ascribe.

5.1 Individual Creators

These people use ascribe because it allows them to securely share their digital creations, and sell unique digital editions. In many cases, ascribe works closely with these individuals to help us refine the vision, the product, and help spread the word.

- Jonathan Monaghan, a leading digital artist represented by bitforms gallery in New York has ascribed all of his work [51]. He says, “ascribe helps legitimize digital-based art work by providing a concrete system of authenticity. It is an important part of the ever-growing acceptance of computer-created work by the art world.”

- Artist Ella Frost not only ascribed her work, she also used ascribe to transfer ownership of some editions to ascribe, for use by ascribe on its landing page (ascribe is “eating its own dog food”) [52]. –
- A 15 year old designer has ascribed her designs for neck chains so that she can work with marketplaces, confident that her design is secure and protected.
- Other leading artists using ascribe include: Richie Culver, Valentin Ruhry, Harm van den Dorpel, Thomas Traum, and several others; typically along with galleries that represent them such as PAM Gallery and Neumeister Bar-Am.

Since launching the website in March 2015, as of early May 2015, ascribe has 600+ users and 2.500+ registrations of digital property. We’ve heard stories of people and companies registering digital art, trademarks, logos and icons.

5.2 Marketplaces, Institutions

ascribe has been working with marketplaces, galleries, libraries, archives, museums, and other institutions to allow them to achieve their objectives.

5.2.1 Art

- MAK Vienna added ascribed digital art to its collection [53] – a screensaver by artist Harm van den Dorpel. In fact, of the 100 editions of the piece, they bought 20 editions so that they could loan out many editions at once without worry.

- Video Art Channel (videoartchannel.com), a service for high-end locations such as hotels provides streaming of video art from the leading archives of Europe. Using ascribe, Video Art Channel and the Archives can easily share and track their video inventory.

- Berlin Art Prize (berlinartprize.com), an art prize competition to discover the latest work of Berlin-based artists. Using ascribe as their technology platform for submissions, hundreds of artists were able to protect their art for the future [54].

- Cointemporary (cointemporary.com), a “concept” digital art marketplace which accepts bitcoin as payment, is using ascribe as their backend for authenticating digital art [55]. When a sale is made, the ascribe registry records the legal transfer and collects the provenance history in anticipation of future value appreciation for collectors. Founder Valentin Ruhry says. “ascribe was something we’ve long been waiting for and eventually enabled us to start Digital Editions on Cointemporary."

- Direct2Artist (direct2artist.com) – a physical art marketplace based in Canada, is using ascribe to notarize and track the sale and history of physical art.

- New startups – people are creating new startups, still in stealth, for digital art marketplaces and auctions, where they’re using ascribe as the IP backbone. These startups are only possible now, because ascribe unlocks the ability to create authentic digital content with strong provenance.

ascribe is in discussions with dozens of other lead marketplaces, platforms, museums and foundations to ascribe, protect, and archive their collections of digital work.

5.2.2 Photography

- EyeEm (eyeem.com) – a photo sharing App with 20M users and 70M images is using ascribe to authenticate photos and track their usage across the internet. EyeEm can better secure their content and ensure that photos that are sold on their marketplace, or sent to Getty Images, is authentic.

ascribe has agreements with other photo marketplaces to ensure that exclusivity agreements are enforced, by giving transparency on whether images are elsewhere on the internet or in other marketplaces.

5.2.3 3D-Printed Designs and more

- Creative Commons France (creativecommons.fr) is working with ascribe to support copyleft, for the benefit of the Free Culture movement [56]. The new CC + ascribe service cc.ascribe.io enables users to share their CC-licensed work without worry of loss of attribution.

- Stilnest (stilnest.com) is a marketplace that sells high-end 3D printed jewellery is using ascribe register designs and enforce exclusivity agreements with their designers. Stilnest wants to give designers more security and reassurance that when they use Stilnest, their proprietary designs are protected and tracked.

ascribe has ongoing discussions with industries across a diverse range such as media, health, software and manufacturing to help them to register and secure their digital intellectual property on the ascribe registry and gain transparency on where their IP may be used.

There is significant, unrealized latent value in digital property and media companies, and until now, it has been difficult for buyers to obtain the rights to use or buy content. DRM doesn't reconcile the physics of bits. ascribe makes it possible for media providers to reach more customers and better monetize their existing inventory of content.

6 A VISION FOR THE FUTURE

Ascribe's vision goes beyond what we've built so far.

We envision nothing less than a comprehensive ownership layer for the internet as a registry for *all* media. We like to think of this as a "bigchain", i.e. combining the characteristics of big data and blockchain-based ownership for the best of both worlds [57][58]. Bigchain™ would have the following characteristics:

> ***Big* Registry with easy, secure legals.** The registry would include not only the time-stamped hashes of ownership transactions, but also the **metadata** (artist name, title, etc.) as well as the **data blobs** themselves (.mov, .png, etc.). Just as the current registry is owned by no one and accessible by all (because it's the bitcoin blockchain), we also envision this for the metadata and the accompanying descriptive information. With the appropriate incentives, this could evolve into an archival role that could serve as an unifying function for libraries, museums, and the Internet Archive (archive.org) of today.
>
> ***Even richer* visibility into usage of the digital property.** It would continue to support existing functionality of registration, transfer of ownership, loaning, renting, and sharing. It would support partial attribution among multiple parties. This includes attribution with the possibility of compensation for all creators, not just digital artists. Licensing could be commercial or Creative Commons. We are working on more complex licensing capabilities to support the evolving nature of international global media. There would be ways to increase visibility into how the work is reproduced, distributed, and what derivative works there are. Visibility will come from a combination of being built into the system "a priori" (e.g. Xanadu-style bi-

directional links or git-style version control), and visibility that is reverse-engineered “a posteriori” (e.g. web-scale search with linkage to registered artefacts for existing media).

The building blocks for bigchain are emerging as mashups of technologies for sharing large distributed files (distributed hash tables, BitTorrent protocol), modern code versioning systems (e.g. git [59]), and consensus [60][61] engines with an answer to Sybil attacks [62] (proof of work [63][64], proof of stake [65][66], or federations, with extra tricks like posting bonds [67][68]). Example mashups include InterPlanetary File System (IPFS) [69], Thelonius and Decerver [70], Project Alexandria [71], and Blockstore [72]. These are all examples of visionary work in the Bitcoin 2.0 ecosystem, work that is fully complementary to our vision of an ownership layer for the internet.

The applications of Bigchain extend beyond digital art and photography. A global ownership registry would include 3D designs, patents, e-books (including e-comics), music, graphic designs, movies, videos, stock certificates, and other digital content. It could also include physical content with similar ownership concerns such as land or vehicle titles [73].

7 CONCLUSION

Solving ownership of digital property, especially on the internet, is not an easy problem. But, it is one worthy of solving. Creators of digital content deserve to be fairly compensated for their work. This paper describes our implementation and proposes a fix that makes it easy to control one’s intellectual property. This solution has two pieces: a registry with easy, secure legals (with the help of the bitcoin blockchain) and visibility into usage / provenance of the content (with the help of internet-scale machine learning). Taken together, these pieces constitute “ownership processing” – a simple, turnkey solution to make ownership actions of digital property accessible to all.

REFERENCES

[1] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System”, 2009, https://bitcoin.org/bitcoin.pdf

[2] R. Brandom, “In the Silk Road trial, Bitcoin is a cop’s best friend,” Jan. 14, 2015, http://www.theverge.com/2015/1/14/7546669/silk-road-trial-bitcoin-tracking

[3] I. Grigg, “The Ricardian Contract”, Proc. First International Workshop on Electronic Contracting, IEEE, 2004, http://iang.org/papers/ricardian_contract.html

[4] I. Grigg and T. McConaghy, personal communication, April 2015

[5] N. Blomkamp, director, “Chappie”, introduces Chappie (character), Columbia Pictures, released March 6, 2015, http://www.imdb.com/title/tt1823672/

[6] C. Nolan, director, “Interstellar,” introduces TARS (character), Paramount Pictures, released November 7, 2014, http://www.imdb.com/title/tt0816692/

[7] N. Stephenson, “The Diamond Age”, introduces Young Lady’s Illustrated Primer (YLIP) (device), Bantam Spectra, 1995, http://en.wikipedia.org/wiki/The_Diamond_Age

[8] CEPIC Copyright Questionnaire, 2014, http://cepic.org/app/uploads/2014/11/cepic_consultation-document_FINAL26Fev.pdf

[9] B. Schneier, “Sony's DRM Rootkit: The Real Story”, Wired / Schneier on Security, November 2005, https://www.schneier.com/blog/archives/2005/11/sonys_drm_rootk.html

[10] N. Khedekar, “We now upload and share over 1.8 billion photos each day: Meeker Internet report”, May 29, 2014, http://tech.firstpost.com/news-analysis/now-upload-share-1-8-billion-photos-everyday-meeker-report-224688.html

[11] YouTube Press, Statistics, https://www.youtube.com/yt/press/statistics.html, last viewed April 27, 2015

[12] The Verge, 2013, http://www.theverge.com/2013/10/17/4844814/would-you-pay-2000-for-a-gif

[13] List of Thinker sculptures, http://en.wikipedia.org/wiki/List_of_Thinker_sculptures, last accessed Apr 21, 2015

[14] A. Horne, "The Thinker, Rodin” (talk). Original uploader was AndrewHorne at en.wikipedia - Transferred from en.wikipedia; transferred to Commons by User:Kafuffle using CommonsHelper.(Original text : I (AndrewHorne (talk)) created this work entirely by myself.). Licensed under Public Domain via Wikimedia Commons - http://commons.wikimedia.org/wiki/File:The_Thinker,_Rodin.jpg#/media/File:The_Thinker,_Rodin.jpg, last accessed Apr 21, 2015

[15] T. Bewick. Barn Owl (Tyto alba) in History of British Birds. 1797-1804, http://en.wikipedia.org/wiki/Wood_engraving, last accessed Apr 21, 2015

[16] "Global Games Market 2012-2016", Newzoo, July 15, 2013 http://www.newzoo.com/infographics/global-games-market-report-infographics/

[17] Hiscox Group, “The Hiscox Online Art Trade Report 2014”, April 28, 2014, http://www.hiscoxgroup.com/~/media/Files/H/Hiscox/reports/The-Hiscox-online-art-trade-report-new-version.pdf

[18] T. Maastricht, “Global art sales in 2014 break all known records, “Press Release, http://www.tefaf.com/DesktopDefault.aspx?tabid=15&tabindex=14&pressrelease=16959&presslanguage tef

[19] R. Rozendaal, “Art Website Sales Contract”. http://www.artwebsitesalescontract.com/, 2011-2014. Last accessed April 22, 2015

[20] G. McMullen, “The History of Copyright, Part 1: The Battle of the Book”, March 9, 2015, http://blog.ascribe.io/the-battle-of-the-book/

[21] G. McMullen, “Home Printing is Killing the Stuff Industry”, April 23, 2015, http://blog.ascribe.io/home-printing-is-killing-the-stuff-industry/

[22] “ARPANET”, Wikipedia, last accessed May 9, 2015, http://en.wikipedia.org/wiki/ARPANET

[23] “History of the World Wide Web”, Wikipedia, last accessed May 9, 2015, http://en.wikipedia.org/wiki/History_of_the_World_Wide_Web

[24] T. Berners-Lee, “Information Management: A Proposal”, March, 1989, World Wide Web Consortium, last accessed May 9, 2015, http://www.w3.org/History/1989/proposal.html

[25] T. Berners-Lee, “Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web”, San Francisco: Harper, Nov 7, 2000, p. 23. ISBN 9780062515872.

[26] T. Berners-Lee, “Frequently asked questions - Robert Cailliau's role”, World Wide Web Consortium, last accessed May 9, 2015, http://www.w3.org/People/Berners-Lee/FAQ.html#Cailliau

[27] "Lynx-wikipedia" by Original uploader was Tedickey at en.wikipedia. Later version(s) were uploaded by Kahlil88 at en.wikipedia. - Transferred from en.wikipedia; transferred to Commons by User:IngerAlHaosului using CommonsHelper.. Licensed under GPL via Wikimedia Commons - http://commons.wikimedia.org/wiki/File:Lynx-wikipedia.png#/media/File:Lynx-wikipedia.png

[28] T. Berners Lee, "World Wide Web", hosted by CERN, Switzerland, last accessed May 9, 2015, http://info.cern.ch/hypertext/WWW/TheProject.html

[29] Google Inc., "Google Chrome Browser", Version 42.0.2311.135, described at http://www.google.com/chrome/, last accessed May 9, 2015

[30] The Opte Project, "Visualization of Internet Routing Paths", Uploaded by Rezonanswoy, 2005, Last accessed May 9, 2015, http://commons.wikimedia.org/wiki/File:Internet_map_1024_-_transparent.png

[31] "The size of the World Wide Web (The Internet)", last accessed May 9, 2015, http://www.worldwidewebsize.com/

[32] Creative Commons, "Creative Commons Licenses", last accessed May 9, 2015, http://creativecommons.org/licenses/

[33] Wikipedia, "Duke Nukem Forever", last accessed May 9, 2015, http://en.wikipedia.org/wiki/Duke_Nukem_Forever

[34] Domain Name System (DNS), Wikipedia, http://en.wikipedia.org/wiki/Domain_Name_System

[35] Shapeways blog, http://www.shapeways.com/blog/archives/19789-copyright-3d-printing-and-you-what-does-it-all-mean.html

[36] T. Nelson, "Xanalogical Media: Needed Now More Than Ever", http://xanadu.net/NOWMORETHANEVER/XuSum99.html

[37] J. Atwood, "The Xanadu Dream", Oct 12, 2009, http://blog.codinghorror.com/the-xanadu-dream/

[38] T. Nelson, "A File Structure for the Complex, the Changing and the indeterminate." Proceedings of the ACM National Conference, 1965

[39] Wikipedia, "Project Xanadu", last accessed May 9, 2015, http://en.wikipedia.org/wiki/Project_Xanadu

[40] T. Nelson and T. McConaghy, personal communication, April 2015

[41] Google Inc., "Google Image Search", last accessed May 9, 2015, https://www.google.com/imghp

[42] TinEye Inc., "TinEye Reverse Image Search", last accessed May 9, 2015, https://www.tineye.com

[43] F. Theron, "Open Assets Protocol," version 1.0, Dec 12, 2013, last accessed May 9, 2015, https://github.com/OpenAssets/open-assets-protocol/blob/master/specification.mediawiki

[44] "Counterparty", last accessed May 9, 2015, http://counterparty.io/

[45] Ascribe Terms of Service, https://www.ascribe.io/art/terms/

[46] Chris Anderson, "The Long Tail: Why the Future of Business Is Selling Less of More," Hyperion, 2005, http://www.amazon.com/Long-Tail-Future-Business-Selling/dp/0935726942

[47] T. McConaghy and D. de Jonghe, SPOOL Protocol https://github.com/ascribe/spool

[48] R. Marques, "pyspool", https://github.com/ascribe/pyspool, last accessed June 9, 2015

[49] T. McConaghy, "Blockchain & Blue Ocean", March 14, 2015, http://blog.ascribe.io/blockchain-blue-ocean/

[50] E. Porter, "Radiohead's Warm Glow", New York Times, Oct. 14, 2007, http://www.nytimes.com/2007/10/14/opinion/14sun3.html?_r=0

[51] T. McConaghy, "Jonathan Monaghan's 'Escape Pod' at bitforms NYC, ownership secured with ascribe", March 25, 2015, http://blog.ascribe.io/johnathan-monaghan-at-bitforms-nyc/

[52] T. McConaghy, "In Getting Work by Ella Frost, We Just Ate Our Own Dog Food!", May 9, 2015, http://blog.ascribe.io/we-just-ate-our-own-dog-food/

[53] MAK Vienna, "MAK purchases digital art work by Harm van den Dorpel with Bitcoin", Press Release, April 23, 2015, http://www.mak.at/jart/prj3/mak/data/uploads/downloads/presse/2015/Harm_van_Dorpel_e.pdf

[54] T. McConaghy, "Berlin Art Prize – Wrapup on Submissions", April 12, 2015, http://blog.ascribe.io/berlin-art-prize-wrapup-on-submissions/

[55] T. McConaghy, "Born digital, authenticated digital, sold digital: Cointemporary + ascribe", April 12, 2015, http://blog.ascribe.io/cointemporary-ascribe/

[56] J. Coates, "Creative Commons France experiments with Ascribe to support copyleft through the Blockchain," Creative Commons Global blog, https://creativecommons.org/weblog/entry/45612

[57] T. McConaghy, "Blockchain Scalability Part I – The Problem", http://blog.ascribe.io/blockchain-scalability-part-i-the-problem/

[58] Trent McConaghy, “Blockchain, Throughput, and Big Data”, Bitcoin Startups Berlin, Oct 28, 2014 http://www.meetup.com/Bitcoin-Startups-Berlin/events/212170052/

[59] Wikipedia, “git (software)”, last accessed May 9, 2015, http://en.wikipedia.org/wiki/Git_%28software%29

[60] L. Lamport et al, Reaching Agreement in the Presence of Faults , J. ACM 27(2), April 1980

[61] L. Lamport et al, "The Byzantine Generals Problem“, ACM Transactions on Programming Languages and Systems 4 (3): 382–401, July 1982

[62] J. R. Douceur, “The Sybil Attack”, International workshop on Peer-To-Peer Systems, 2002. http://freehaven.net/anonbib/cache/sybil.pdf

[63] Wikipedia, “Proof-of-work system”, last accessed May 9, 2015, http://en.wikipedia.org/wiki/Proof-of-work_system

[64] A. Back, "Hashcash - Amortizable Publicly Auditable Cost-Functions", Tech Report, August, 2002, Adam Back, http://www.hashcash.org/papers/amortizable.pdf

[65] Wikipedia, “Proof of Stake”, last accessed May 9, 2015, http://en.wikipedia.org/wiki/Proof-of-stake

[66] Proof of Stake, bitcoin.it wiki, last retrieved Feb 4, 2015, https://en.bitcoin.it/wiki/Proof_of_Stake

[67] J. Kwon, “Tendermint: Consensus without Mining”, Draft v.0.6, fall 2014, http://tendermint.com/docs/tendermint.pdf

[68] V. Buterin, “Slasher: A Punitive Proof-of-Stake Algorithm”, Jan 15, 2014, https://blog.ethereum.org/2014/01/15/slasher-a-punitive-proof-of-stake-algorithm/

[69] J. Benet, “IPFS – Content Addressed, Versioned, P2P File System, 2014, http://static.benet.ai/t/ipfs.pdf

[70] Eris Industries, “The distributed application platform”, last accessed May 9, 2015, https://erisindustries.com/

[71] Blocktech, “Alexandria – The People’s Library, last accessed May 9, 2015, http://blocktech.com/

[72] Onename.org, “Blockstore: key-value store for name registration and data storage on the Bitcoin blockchain”, last accessed May 9, 2015, https://github.com/namesystem/blockstore

[73] M. Swan, “Blockchain: Blueprint for a new Economy”, O’Reilly, 2015