COPYRIGHT ON THE INTERNET

or: What a Difference a Command Makes

by Campbell Vertesi

Copyright on the Internet is a contentious issue, to say the least. Everyone seems to have an opinion; everyone shouts that they are fighting for progress, and for the good of society as a whole. They shout, without ever clarifying just what the issue is in the first place. Music companies and their advocates would have you believe that it is a battle between piracy and ownership rights, and place the future of the free market and quality entertainment as the stakes. File sharers and their ilk insist that it is a fight between freedom and the greedy music industry, with the future of the Internet and free expression at stake. Why, complains the business world, can't the online community understand the reasoning behind strong copyright law? Why, protests the file sharing world, don't the companies just “get it?” Both sides have their rhetoric, but so far no one has managed to really define the base issue at stake here. How can we ever hope to reach a consensus, if we cannot even agree on the question we're arguing about in the first place? The only thing we can agree on, it seems, is that digital copyright is somehow a critical issue in the new medium of the Internet.

The trouble is, it is not really the question of digital rights that these people are fighting over. Rather, this is a classic clash of paradigms; a conflict between two competing views of the world. A big tip off should be that neither side is arguing rationally. Both are so entrenched in their own view of the world that any alternative poses an enormous threat, and provokes a strong defensive reaction. This essay will help to clarify exactly

where these two paradigms conflict, and in doing so will offer new perspective on the battle at hand, and on some possible solutions.

"CP": THE REVOLUTIONARY COMMAND

There is at least one fundamental difference between the computer world, and the "real world". This difference completely remodels our perspective on the universe, and yet it is embodied in a command so simple that even the most novice computer users take it for granted. This command is "copy".

The Copy command makes your computer a radically different place from the real world. Suddenly, everything is by nature abundant. A computer could care less if one person has a picture, or if a dozen people each have their own copy. The cost of duplication is so low as to be practically nonexistent, and each copy is a perfect duplicate of the first.

Imagine for a moment what life would be like in the real world if we all had the "copy" command – it would be something like on "Star Trek," where every household has its own "replicator" to satisfy the needs of the inhabitants. We all remember watching Captain Picard approach a computer panel, commanding "Tea. Earl Grey. Hot." On his command, the computer would recall a pattern for this sort of tea, and produce a copy – a "replicate" of this perfect cup of Earl Grey tea. Such a world would mean an end to famine and poverty. Everyone could have anything they ever wanted, without any consequence. This is the world inside your computer.

The Internet was designed to take advantage of this "copy-ability". The

ability of computers to make infinite copies is central to the way the Internet works. For example, every time you look at a web page, your computer is actually showing you a **copy** of the original. Every time you read an email, you are actually viewing just a **copy** of what was sent. Whatever the content, the only way anything is ever transmitted on the internet is through copying. It is almost impossible to keep a true "original" single copy of anything on the internet.

This fact can be hard for many people to understand. Most of us can think of some kind of Internet content that certainly seems to keep only a single copy. There is an entire software industry out there that specializes in this exact ability. Real Networks, of RealPlayer fame, is an excellent example. When you watch streaming video on RealPlayer, it seems like you're viewing the single original copy. After all, you can't save the movie to your computer, and it makes a big show of transmitting the video as you watch it. What is actually going on however, is more complex. As you watch your video, a copy is actually being made on your computer. The RealPlayer program is designed to delete each second of that copy after you view it, so only a few seconds of the video are actually present on your computer at any given instant. This creates the illusion of a unique copy. There are many programs in existence (called "stream rippers"), which are simple modifications of RealPlayer, to remove that little "delete" instruction. This leaves a complete copy of the video on the user's computer. In the end, even companies that specialize in protecting content can't keep anything as a unique copy online.

At least, an individual program can't preserve a single, unadulterated

copy online. Perhaps we must put our copy-protection scheme deeper into the computer, into the Operating System itself, perhaps. What if Windows itself were designed to get rid of this copy-ability? Microsoft, the biggest of all software giants, has been tackling just this problem, and even they have given up. This Goliath of the computer industry now works under the assumption that it is **impossible** to preserve a single copy of anything n line. They are instead focusing on so-called Rights Management technologies – software to ensure that although there may be millions of copies of a file, only an authorized user can use those copies. If you can't make sure that the information won't be copied, at least make sure that it can only be read by the appropriate authorities. Apparently, this “Copy-ability” is a powerful enough force on the Internet that even the largest software giants surrender in the fight against it.

Clearly, there is a strong bias towards copying (which I call Copy-ability) built into computers and the Internet. This bias is so strong that even the most powerful software companies cannot design a way to combat it – they must treat it as a force of nature instead. But what difference does this make to businesses, and the real world? More importantly, what does all this have to do with the copyright debate?

CORPORATE COPIES – CP SUCCESS STORIES

Copy-ability is so fundamental to the way computers and the Internet work, that it is apparently impossible to fight against it. We can think of it as a sort of “ocean current” in the Internet. It seems that no matter what sort of boat you have, you cannot fight it. Rather, companies have had to deal with the consequences of Copy-ability as if it were a fact of life. But what if, instead

of fighting against Copy-ability or finding away to grudgingly live with it as Microsoft has, you could use it to your advantage? What if you could harness the power of the Copy-ability current to drive your business? A business that could sail with the current would be very successful indeed.

There have been at least three major examples of attempts to use Copy-ability to drive business, and every one of them has been the roaring success we would expect it to be. Not all of these people were conscious of what they were doing at the time, however.

The first example is everybody's favorite company, Microsoft. How did a small garage-based company in the 80's become the world's biggest software vendor, in less than a decade? Through Copy-ability. Bill Gates' key move was to allow his operating system, MS-DOS, to be easily copied. Though most software vendors were heartily investing in copy-protection schemes for their applications to make sure there were no illicit copies made, there was no protection on MS-DOS. People were practically ENCOURAGED to give it away illegally to their friends, and that's exactly what they did. Soon enough, everyone was using MS-DOS to run their computers. In this way, Microsoft used even the limited copy-ability of pre-Internet computers to gain market share. And it worked, possibly better than they had hoped. Their rival at the time, Apple, was selling more powerful computers with an easier interface, and were were first-to-market with the Personal Computer concept. In a traditional market, such an advantage is almost insurmountable – but this was not a traditional market. Apple was using the old methods of distribution, putting their best protection on their software. Because of this fact, Apple was all but

obliterated by the Copy-abled Microsoft in the 1980's. As we all know, Microsoft went on to gain a monopolistic control of the PC market, simply by allowing people to pirate their software. This was a Copy-ability driven business model.

Microsoft had incredible success using Copy-ability to distribute their Operating System, but that was before the age of the Internet. Where regular, disconnected computers simply have a bias towards copying, the Internet has an all out mandate. As the second example, Netscape was the first company to leverage the increased Copy-ability of the Internet to sweep market share, in an imitation of Microsoft's earlier success. By making their Netscape web browser a free download from the Internet, the company found themselves in complete control of the browser market almost overnight. Because anyone could download a copy for free, everyone did – and it didn't cost Netscape a penny. Just like Microsoft before them, Netscape found that by allowing their customers to freely take copies of their software, they could dominate the market. Copy-ability granted Netscape such total control that even when Microsoft resorted to illegal, anti-competitive tactics, it still took years to reduce Netscape's market share to 50%. The Copy-ability of the Internet helped Netscape take on one of the world's most powerful monopolies, and wage a war of attrition that lasted for several years. It seems safe to say now, that Copy-ability is a major boost to any company’s software distribution.

Eventually though, Netscape was bought by AOL, and stepped down from their fierce competitive stance against Microsoft. With total control of the Operating System and web browser markets, the latter company did an

about face. Microsoft has no more need for Copy-ability in their distribution; they have all the market share they need, thank you very much. Instead, Microsoft is now trying to make good on every copy of their software in use, through extensive copy-protection measures. Bill Gates should have learned his own lesson; this was exactly what his competitors were doing in the 1980's. Now Gates is faced with the significant challenge of my third example of taking advantage of Copy-ability: the Open Source community.

Linux and the Open Source world have taken the use of Copy-ability to a new level. Where previously, software vendors have used Copy-ability to power the distribution of their programs, the Open Source model uses Copy-ability to *create* their software as well. This concept takes some explanation for the uninitiated.

In a copy-enabled world, a group project is defined differently from in the real world. In the real world, if ten people work together to build a car, they will all have to share one car when they're finished. There is therefore a limit on the number of people who will be willing to work together in that way, since the reward seems to decrease the more workers are involved. After all, who wants to share a car with 100 other people? But what if we had the Copy command at our disposal? We could all contribute, and each get our own entire car in exchange for our small amount of labor. Heck, we could give cars to everyone we know and ask for their feedback. Any feature that anyone adds to their car could be likewise distributed to the rest of us, at no cost to anyone. In the end, you could build a glovebox and get an entire car in exchange – simply for giving away copies of your glovebox. Group labor is incredibly efficient in a Copy-enabled world.

This is the Open Source method of software development. You can take the result of thousands of programmers' efforts, all for free. The only price is the expectation that you will contribute what you can back to the community – even if all you can contribute is your feedback. If you have any problems with it, you can fix it yourself and share your results with the world, or let a programmer know and they will help you do it, or perhaps take care of the problem themselves. Once one person has solved a problem in this way, no one else will ever have that problem again, because it has been repaired in the original program itself. Imagine if your car's carbeurator bit the dust, and by replacing it yourself, you prevented the same problem for every other car in the world.

Open Source is driven by copy-ability. Although most people only contribute a small amount to the community of developers, they each get their own copy of the sum total of the group's labor in return. Because computers and the Internet are designed for this kind of copying, it's easy for me to contribute even something so small as an email with my comments on a program, and get exponential multiples of my work back. I can't program my way out of a paper bag, but the Copy-ability of the Internet allows me to reap the rewards of millions of hours of talented programmers' work.

Open Source projects use Copy-ability to power not only their distribution, but also their development processes. In an analysis of Open Source projects then, we should expect products with superior distribution to those under “traditional” licensing. We should also expect to see the effects of a Copy-enabled development process in superior, or at least highly

competitive products.

These expectations are borne out by many popular Open Source projects. Superior distribution is evidenced by projects such as Apache web server, which runs 67% of the world's web sites – more than every other brand of web server combined (According to Netcraft.com's 2003 survey). Their leading competitor is Microsoft's web server, under a development model that is not copy-enabled. Microsoft runs only 23% of the world's web sites. Open Source email agents such as Sendmail run a similarly large proportion of email servers. Though email servers are more difficult to profile, a common estimate puts Open Source email servers at more than 70% of email infrastructure for the entire Internet. Both Apache and Sendmail compete against many other professionally developed products. That projects with zero advertising budget have taken such powerful command of their markets should be strong evidence that the Copy-abled distribution system is working well for the Open Source model.

Proving that Copy-enabled development creates superior products is more difficult. After all, how do you rank quality objectively? One can glean a reasonably objective view of the quality of a product by analyzing its stability; in other words, how often does this program crash? The Netcraft survey quoted above found that the 50 most stable web sites (as judged by average uptime) were **all** running operating systems with Open Source development models. Furthermore, 94% of these most stable servers were using the above-mentioned Apache web server software.

These numbers go well beyond statistical significance. They overwhelmingly favor products developed through Open Source methods as

vastly superior to their traditional counterparts. Even Microsoft admits openly that they cannot match the sheer numbers of developers – and implicitly, the quality of development – of Open Source projects such as Apache. Apparently, Copy-ability is a very powerful force in software development as well as distribution.

These numbers are astonishing to a traditional market mindset, and probably appear doctored. Consider though, what we have seen with Copy-enabled distribution. When the small upstart Microsoft Copy-enabled their products, they wiped out everyone else in the market. Apple had every possible traditional business advantage on their side, and the simple fact of Copy-ability tore them to pieces practically overnight. Netscape Copy-enabled it's distribution, and without ever charging a penny to their users, held their ground against flagrant violations of anti-trust law, and all of the might that one of the biggest companies in the world could throw at them. Seeing the effect that Copy-ability has on distribution, perhaps the success of a Copy-enabled development model is not so shocking after all.

The existence of a copying “current” on the Internet should now be proven to the reader beyond a shadow of a doubt. By harnessing this power, businesses have leapt from obscurity to stardom. This power has allowed a group of volunteers to produce a product that appears to outdo many professionally developed products. But what does Copy-ability online have to do with copyrights?

THE CONFLICT

So Copy-ability can make monopolies out of nobodies, and make a thousand separate geeks produce some quality programs. Who cares? If you have ever used the Internet, then **you** care. Copy-ability is not limited to use by companies or programmers. Most everyone on the Internet interacts in some way with some sort of Copy-enabled community, not altogether unlike the Open Source community. Every time you view a website, you are participating in a small way in a Copy-enabled universe. Every time you post to a message board, send a virus warning to your friends, or leave a comment on a website, you are contributing in a much more significant manner. When you do any of these things, your thoughts, your comments, or information – what the real world calls your "intellectual property" – may be freely copied by anyone in the world in the form of email forwards and web searches. Anybody can see your contribution to the Internet. What's more, this probably doesn't bother you. In fact, most people contribute in this manner because they have already received so much in return, and they feel that it is morally *right* to contribute.

This is a little bit hard to believe. How can you be contributing to a Copy-enabled world that you don't even know about? A good example can be found in the rating systems employed by so many Internet vendors. My mother buys all sorts of things on Amazon.com, but not before she has read a product's user reviews through to the end! When she receives her product, she typically adds her own this thoughts to the site. In this way, she is contributing her intellectual property to a Copy-enabled community. My mother also refuses to trade on Ebay without seeing someone's buyer or seller rating first. This is another Copy-enabled community in which she is participating. Almost any community on the Internet is Copy-enabled,

simply because it must be to survive. Mailing lists, public forums and job-search sites are all Copy-enabled communities.

Have you ever interacted with the Internet in this way? Have you ever posted a product review, rated an Ebay seller, or participated in an online chat session? Have you ever forwarded an email joke to your friends, or subscribed to a mailing list? This sort of information sharing activity is ubiquitous to the Internet, and is eventually encountered by anyone who "surfs the web."

Social psychologists will tell you, that anyone who spends a significant amount of time interacting in a particular community gradually develops the mindset of that community. According to the mindset required for participating in the Internet, sharing information is good, and is to be encouraged. If you spend a significant amount of time interacting with the Internet, chances are that you feel good when you post to Ebay warning about a fraudulent seller. If you receive a virus warning, it is the *right thing to do* to pass it on to your friends. You might even get excited when an essay you wrote is published on a well known news site like Newsforge. The fact that sharing information – your intellectual property – is a good thing in your mind is a dead giveaway of the Internet paradigm.

The Internet paradigm is essential for the functioning of most every Internet community, from Ebay's seller rating system to more complex communities such as Kazaa. Unfortunately, in the real world, this mindset is not only counter-intuitive; it's completely non functional. In the real world, intellectual property is very highly valued, with everything from college tuitions to copyrights as obstacles to various types of information. Scientists

working for private companies rarely, if ever, publish their work in academic journals (a form of information-sharing community). In fact, in the United States, everything you record in any way – with typewriter, tape deck or scribbled napkin – is automatically copyrighted to you. Information sharing is anathema to many real world communities. So, if information sharing works so well on the Internet, then why don't companies publish their information publicly? Why is it a good thing to give your programming away a la Open Source, but a bad thing to publish your pharmaceutical research?

Remember, we're talking about the real world now, without the Copy command. In the real world, everything is by nature scarce. There are only a certain number of cars in the world, and so it is important to determine exactly who owns which cars. Or carrots, or CD players, or chairs. Property is an essential part of our interaction with the real world, because scarcity is a fact of life out here. We use property as a measurement of value and of wealth, and so we apply the same rules of physical property ownership to intellectual property. In a world without Copy, it makes no sense for a research corporation to publish their results, since the value of those results is directly related to how few people know them. If they told people what they discovered, research companies would be out of a job, the same way a farmer would be out of a job if he gave away all of his carrots for free.

In the real world, everything is by nature scarce. You deal with it or die. This fact is a cornerstone of every culture and economic perspective in the world. It must be, because scarcity is a fact of existence. Even under communism, where everything belongs to the community, in the end food and objects must be portioned out to individuals. After all, even if everyone

owns the last cookie, only one person gets to eat it.

In contrast, on the Internet, everything is by nature abundant. Try as you might, you can't keep a single copy as just a single copy. Communities have arisen out of this copy-ability, and anyone who interacts with these communities is used to sharing their intellectual property freely. Of course these two perspectives are in conflict! Companies are used to functioning in a world where scarcity is a fact of life. Their business models are built around that fundamental assumption. They are therefore struggling as hard as they can to limit ownership of their on-line property, using fees and copy protection schemes. These are methods of artificially imposing scarcity on the on line world, where scarcity does exist the way it does in the real world. These methods work to a certain degree, permitting real world business models to operate with some success online, but they are not complete solutions. They rely on forcing the Internet world to mimic real world behavior, and therefore they are fighting upstream against Copy-ability. And if there is one thing we have learned, it is that you cannot fight Copy-ability on it's home turf of the Internet.

If you are used to the Internet perspective, where sharing your knowledge and experience is a good thing, measures like copy protection are frustrating, and don't make sense. Why don't the companies just "get it?" or "information is free!" you may cry. "We have a right to information!" It is very difficult to build up a concrete argument for exactly **why** information should be free, you just "know it." This should sound very familiar to many supporters of music download services such as Kazaa. If you participate in an advanced sharing community like a P2P network,

chances are you are operating quite firmly in the Internet's sharing mindset. This paradigm has Copy-ability as it's foundation, and is steadfastly rooted in the online world. This perspective is perfectly legitimate, as long as you are used to sharing information and are operating on the Internet. It makes no sense whatsoever though, if you are used to working in the outside world.

If you are used to working with the real world, your perspective is necessarily built on ownership. Especially in a capitalist society, possession is critical to your understanding of the world. If you are used to thinking in terms of ownership and property in this way, then copying on the scale for which the Internet is designed is an immensely destructive and threatening force. Reactions to it have been predictably strong. *"A machine in every house, all designed to copy everything that comes in contact? It must be stopped by any means possible!"* Though these responses may sound exaggerated, they are typical of many of the defensive reactions against the Copy-ability of the Internet. In their efforts to curb Copy-ability, companies and lawmakers are proposing laws that border on the ridiculous. One bill calls for the legalization of hacking into computers suspected of sharing copyrighted material. Vigilante justice at its best! Another proposal (on the Senate floor at this writing) advocates explosive devices in computers so that file sharers' computers can be **destroyed** remotely by content providers. These reactions are so over the top as to seem comical. They are real, however. Remember, Copy-ability threatens the foundations of possession-oriented thought. Extreme reactions should be expected.

Both sides feel personally threatened in this debate. The real world is

watching capitalism be torn apart by on line economics. How does supply and demand work if scarcity doesn't exist? On the Internet, supply is infinite, so cost should be zero. How do you make a profit on that? Furthermore, how do you have incentive if there is no property? These questions strike at the heart of our Western economic philosophy, and it is critical that they be answered.

On the other side, the Internet community is watching everything they value get torn down by these real world intruders. Record companies are suing 12 and 15 year olds for participating in an information-sharing culture. Websites are subpoenaed to be taken down for exposing information that companies don't want the public to see. This does not encourage information sharing! Everything that tells us that posting reviews to Amazon is *good*, tells us that this corporate behavior is *bad*. As more and more sites "go commercial," the free information-sharing culture that makes the Internet such a fantastic place is being eroded.

How can this issue be solved? How can the property-less, practically communist world of the Internet interact with the necessarily property-centered real world?

CONCLUSION - HOW DO YOU SELL TO A COMMUNIST?

We can now redefine the debate over copyright on the Internet. It is not a battle over the nature of ownership on the Internet, after all. Rather, it is a conflict of two opposing paradigms. One of these paradigms lives within the Copy-enabled world of the Internet, where freedom of information is expected. In this worldview, ownership is a moot point. This perspective

cannot understand why the companies don't “get it” that ownership doesn't make sense here.

The other paradigm lives in the “real world,” and cannot fathom why the Internet community has such a hard time with perfectly natural property rules. The Internet paradigm makes no sense in real life, and similarly the real world paradigm is non-functional in Cyberspace. Both sides fail to understand that their problems are reflections of fundamental differences between the real world and the online one. The Internet is radically different from the “real world” in that it is specifically designed for freedom of information. The Internet is designed for Copy-ability. This difference is what makes a connection between the two worlds so difficult to create. With our new perspective in mind, the real problem that we discussed in the introduction becomes clear: how do you sell to a communist? How can someone who doesn’t understand the concept of property participate in a capitalist business model?

Our new perspective allows us to consider many new possibilities for how the business world can interact with the online world. The problem of incentive for example, can be approached with the currency of authorship. This means creating rules to give authors credit for their work. The Open Source movement relies heavily on this method, as does the popularity of weblogging (aka “blogging”). Authorship is also a proven incentive model in the real world, where universities and academia in general form highly successful authorship communities. In academia, just as in Cyberspace, normal market rules do not apply. No Physics professor is in it for the money, but every one would kill to be the next Einstein. Perhaps industry

can interact with the Internet using templates from existing business-academic relationships. At the very least, existing academic communities can provide some framework to help real-world rooted models cope with the Internet paradigm.

There are also some existing business models that already work on the internet. Advertising has found a comfortable niche in the trafficking of website "hits." The online adult content industry, for example, is completely self-sustaining, and is one of the most profitable areas of the Internet market today. Most of these websites never charge a cent to the end user, instead receiving payment from other websites for advertising placement and "click-throughs" on banner ads. Although it's true that many websites have had only limited success with this method, it is clear that innovative and targeted marketing have great possibilities as moneymakers on the Internet.

Service industries too, seem to function perfectly well in Cyberspace. The travel booking industry is an example that has migrated almost entirely to the Internet. Vendors of Open Source software such as Red Hat Linux post solid profits by selling technical support and updates to business and academic clients.

Unfortunately, industries that are based on information supply are still stuck. Publishers for instance, are acutely aware that there **must** be tremendous opportunity for them in the Internet, where information travels so freely. No one has come up with a satisfactory model, however. The record and movie industries are more widely publicized instances of the same problem. So far, the best they have done is to offer costly online

versions of existing record stores, offering lower than CD quality music for the same price you pay for a disk in a store. In the end, how can you sell your product to someone who can get it for free? How **do** you sell to a communist?

What many companies are learning, is that you cannot. It is impossible to sell a product to a communist. Companies must learn that online, value comes not from the product itself, but from the associated services. Itunes is an excellent example of a model that is (perhaps accidentally) enjoying some success online. Although most everyone knows how to download music for free from P2P networks such as Kazaa, Apple is selling songs at premium prices, mostly because they are not actually selling songs.

The reality is that with Itunes, no one is paying for music – they can get music for free from networks like Gnutella and Kazaa. Instead, customers are paying for the quality of service. For 99 cents per song, you get a guaranteed, high speed download of any song or album you want, with all of the correct labels and filenames. The store is integrated with your music player, which guarantees the ultimate in convenience shopping. As an added bonus, you get information about the artist you like, and recommendations of other music you might enjoy. Itunes is successful in competing with free, because although P2P services offer a free product, their downloads are often slow, and mislabeled. Furthermore, Kazaa certainly doesn't supply album and artist information the way Itunes does. As it turns out, many customers think that the extra services and convenience are worth the 99 cents.

It is clear that the debate over online copyright is not a black and white

issue, as the rhetoric would have us believe. Rather, it is a clash of paradigms, two incompatible views of the way information and property operate. In trying to apply the real world paradigm online, we have in some ways been trying to fit a square peg in a round hole. Though that peg has worked perfectly well in all of the square holes we encounter in the real world, it is completely nonsensical in the round pegged world of the Internet.

That being said, there is still hope that the two worlds can interoperate. There are definitely similarities between the square and round pegs that we should note. The currency of authorship and the value of services are tremendously powerful tools for forming business models. Likewise, there are enormous advantages in the round peg world, such as copy ability, and the unparalleled access to grass roots support. Clever and adaptive businesses are certainly capable of operating under business models that work on these principles, and many of them are well into that transition already. As for information supply industries such as the movie and recording industries, there seems no choice but to change drastically. Perhaps they should heed the example of Itunes, and sell their products as cleverly packaged services. Even the communists would buy that.