Open Channels Project

Harold S. Henry
425 / 208 6230
harold@openchannelsproject.org

White Paper January 1, 2007

Low-Friction Channels:
Realizing the Internet's Commercial Potential

Abstract

An opportunity exists today to establish a new global marketplace that could dominate the Internet — a practicable set of infrastructure improvements could create open sales channels far superior to anything in existence today.

Channel utility is bounded by the highest point of friction in the channel. Anything that impedes product flow is a source of friction, including excessive cost, barriers to entry, lack of discoverability, uncertainty, inconvenience, danger of theft, and so on. To create a open channel, *all* significant high-friction points must be eliminated or the channel friction will remain high.

This paper describes the infrastructure improvements needed to remove all significant points of Internet channel friction end to end. Most of the technologies needed have already been implemented in other contexts, and the few substantive technical innovations needed are relatively simple, but the resulting channels would be so dramatically better than existing ones as to make most current Web commerce models obsolete over time.

To be successful, these channels must be created and maintained as a communal infrastructure like the Internet itself, by a kind of international public utility that charges only enough to be self-sustaining. The explicit objective of the project from the start must be to create a level commercial playing field at the lowest practical friction and the highest utility to buyers and sellers alike.

Not only will this approach encourage support from communities and governments around the world, it will also discourage fragmentation of the marketplace by commercial ventures, since there will be little or no money to be made by controlling the channels themselves.

It is not too much to claim that the resulting marketplace can have a dramatic, even revolutionary positive effect on the entire world economy.

Open Channels Project (www.openchannelsproject.org)
7511 329th Avenue SE
Fall City, WA 98024-6700
U.S.A.

As of: Friday, July 19, 2013 at 11:38 PM

Contents

Introduction

As long as Internet commerce continues to mimic earlier business paradigms, it fails to realize its economic potential.

The most successful business models so far (Google, eBay, Amazon, etc.) have been converging on a new buying and selling paradigm that will continue to shift markets and resources, but to realize this new paradigm fully, changes are required to provide:

- More trust — more reliable identity, anonymity, and security.
- Lower cost — make it less expensive to create and use channels.
- Better discovery — continue the current shift away from push advertising to less expensive and more effective pull discoverability.
- Better decision support — a more effective and reliable evaluation infrastructure to facilitate consumer selections.
- Better product delivery.
- Ease and fun —the entire process of online buying and selling should be effortless and entertaining.

Current Internet giants are providing primitive precursors of the kind of infrastructure being proposed here, often at the cost of significant investment. The new infrastructure would make all precursors obsolete, causing a major realignment of resources. Even more significant would be the effect on businesses whose current profitability depends on controlling bricks-and-mortar channels that could be replaced under the new online model.

Such a shift would eliminate a lot of jobs and create more new ones. It would realign entire industries. By eliminating middlemen, it would increase per-capita productivity. By creating new information and entertainment markets, it would vastly expand both the prosperity of scientists, inventors, and artists, and the breadth and quality of information and entertainment available to everyone.

It would shift a significant amount of world trade out of the hands of companies and into the hands of individuals.

It would make the generation of information more lucrative to the people who do it, while greatly lowering the cost of information to the rest of the world.

It would take us a huge step closer to making quality education universally accessible.

It might even help draw us together as a planet.

Required Set of Infrastructure Improvements

In 2006, online sales accounted for about 6% of all retail sales, and Forrester Research projected in May of this year that total online sales in 2007 will exceed $250 billion. Online markets in Brazil, Russia, India and China are all growing very rapidly — in January of this year, the Internet Society of China reported that Chinese online spending had grown 47% in 2006, to $35.5 billion (276.8 billion yuan). According to a recent Jupiter Research estimate, 1.1 billion people now have Internet access, and that number will grow to 1.5 billion people by 2011 — 22 % of the world population.

IDC analyst and vice-president Carol Glasheen projected in April of this year that by 2010, global e-commerce spending will exceed $10 trillion.

At the same time, there are serious problems impeding Internet commerce:

- Security problems continue to grow, as documented below, with no solutions in sight.
- Studies in early 2004 identified shipping costs and the inconvenience of entering necessary shipping and billing information as one of the biggest deterrents to online shoppers.
- Ongoing battles over copyright, piracy and DRM issues, combined with high prices of one kind or another, has disillusioned many people with online music purchasing.
- A lack of accountability and customer service has made big vendors such as eBay and Amazon less attractive to many people in recent years.
- As retail sites proliferate, Google has become less convenient as a way to find things. Searches now return so many hits that the one you're looking for can be many pages deep in the results. Many people also find it hard to evaluate new vendors, and know whether they are trustworthy.

The central premise of this white paper is that a carefully crafted package of infrastructure elements can remove points of friction from one end of an Internet sales channel to the other and thereby revolutionize Internet commerce. A corollary is that no subset of these elements would achieve that effect because significant points of friction would still remain.

The infrastructure elements in question fall into the following categories:

- A set of processes and protocols that makes transactions both secure and convenient.
- A set of interconnected distributed systems and peer-to-peer technology that achieves high performance at low cost.
- A distributed, rules-based data model that makes the system logic easy to modify, extend and maintain.
- Built-in transaction handling with support for microbilling and persisted history.
- Pull discoverability supported by a commercial, market-driven system of buyer decision support, along with smart categorization and search capabilities.
- An integrated and streamlined shipping process.
- A rich visual shopping environment that runs on local machines, with highly customizable and extensible shop creation, complex consumer preference and profile support, and an intuitive, entertaining user interface that can be extended or replaced by third parties.

These infrastructure improvements are described in more detail below.

Security and Trust

The purpose of security is to establish trust, particularly in financial transactions. One of the largest current points of friction in the Internet is a lack of security, and consequently of trust. As *eMarketer* analyst Ben Macklin observed in a report this June:

> Today, Internet security issues are less about viruses or denial-of-service attacks, and more about spyware, malware, phishing and social engineering — all designed to extract personal information from unwary users to provide financial gain for increasingly sophisticated e-predators. Cyber criminals are no longer pimply-faced teenagers who want to show off to their friends, but sophisticated gangs connected to organized crime whose primary motive is money.

Available data show that U.S. Internet users are adjusting their behavior to avoid identity theft and fraud by visiting fewer sites, avoiding online banking and restricting their online purchases. Based on Harris polling, Gartner and Consumer Reports data, Macklin estimated the costs of online security concerns in 2006 as follows:

30%	of U.S. adult Internet users limit online purchases to prevent identity theft.
24%	of U.S. adult Internet users limit online banking to prevent identity theft.
32.7 million	U.S. adult Internet users will not use online banking due to security concerns
$1.91 billion	U.S. e-commerce revenues were lost in 2006 due to online security concerns
$8.43 billion	of economic damage resulted nationwide in the U.S. from malware in 2006.

Add to this widespread privacy concerns about the kind of personal information that is constantly being gathered, harvested and sold, and it is little wonder that trust is so low.

Most existing types of online fraud and privacy invasion could be prevented, however, if there was a system that could reliably do the following:

- Identify the sender and recipient of every message.
- Maintain anonymity between users.
- Apply up-to-date profile information to restrict a user's ability to act.
- Apply simple rules.

It's not enough that the software and hardware implementing the system support these security objectives, though — in addition, the system's administrators must have the resources and commitment to resolve disputes, put a stop to abuses, and where necessary, track down and prosecute fraud, theft, and illegal or destructive use of the system. This is not to say that administrators need be involved in censoring content that may be popular in some communities but offensive in others — avoiding such content at an individual level can easily and in some cases automatically be achieved using rules acting on content metadata and user profile preferences.

The resulting level of security would be high enough to support many types of commercial transaction not currently entrusted to the Internet, and would allow significantly more automation of transactions than is now possible, not only between consumers and suppliers, but also within the supply chain.

Digital Rights Management

Digital rights management (DRM) has become an increasing concern to channel owners who sell digital content, because security techniques are generally ineffective at preventing piracy. Basically, digital content must be rendered into a clear-text form to be consumed, and in clear-text form can easily be pirated. Even if decryption occurred within computer monitors, clear-text signals would be relatively easy to capture electronically.

Digital rights can only be protected effectively at an individual level if consumers perceive the product to be worth the price you're charging for it — owning it must seem more convenient and appealing than pirating it.

Consider: How many people would borrow a paperback novel from a friend and copy it on a copier? It isn't copyright laws that deter them, it's a cost/benefit analysis. In library copy rooms, people copy $150 scholarly monographs day in and day out.

The problem of protecting digital rights is not a security problem, it's a pricing and product definition problem. Security can (and should) make stealing a product more inconvenient, but much more important is to make ownership a benefit. None of this is news, of course.

The system proposed here will support existing DRM security, including imperceptible watermarking of visual content and other techniques for detecting and preventing illegal distribution, but its main contribution to protecting rights of authorship will be to create direct channels between content creators and consumers in which creators are well-compensated, consumers pay prices they regard as fair, and ownership is more convenient than theft.

Low-Cost Collaborative Networks

The network infrastructure needed to support transaction volumes that are orders of magnitude higher than any today, at a level of security currently used only by financial institutions, must obviously be extremely performant and scalable. At the same time, in order to be able to evolve smoothly over time to take advantage of new technologies and meet new needs, it must also be flexible, modular and easy to extend.

The system will be composed of an innovative peer-to-peer cloud that handles much of the transaction negotiation workload. The advantage of using a peer-to-peer network is that the need for back-end resources at the edges of the cloud can be greatly reduced. High performance and scalability can be achieved without huge server farms.

Existing Internet giants such as Amazon, eBay, and Google all have had to invest heavily in their back ends, spending billions of dollars on server farms and server software to handle the volume of traffic that their sites generate. This investment represents a barrier to entry to competitors but is also a liability as technology advances — Amazon's recent woes illustrate the difficulty of responding to new challenges when you have a huge existing code base.

Google, by contrast, has built a relatively simple, general-purpose clustering system that allows it to grow and change much more easily. The open channel system proposed here is somewhat similar to Google's approach, but is even more performant, flexible, and scalable because it is more distributed and grows organically with the number of clients in the system.

Although the open channel system would still require server capacity at the edges of the peer-to-peer cloud so as to authenticate users' identities and current capabilities, publish content, and guarantee persistence of transaction and ownership data, the processing power and resources of the cloud would greatly reduce the size and bandwidth required.

Smart Data

The peer-to-peer network under discussion is unusual in allowing rules-based logic to be attached to data. As a result, the entire cloud can serve as a distributed, high-performance rules engine, providing a level of secure, granular programmability that has many advantages:

- Many facets of the system can be safely enhanced, extended and changed simply by adding or changing rules, without having to rewrite any infrastructure code. This will allow the system to evolve far more smoothly and easily than is normally the case with large-scale systems of this kind.
- Transaction support can be extended and customized without changing the underlying infrastructure code.
- Many of the kinds of functions envisioned loosely in the "semantic Web" concept can be realized straightforwardly, without special programming.
- Security rules relating to data can be set flexibly at a per-record or per-field level.
- Messaging services such as email and instant messaging can easily and securely be implemented within the system (no more spam!).
- The network can tune itself to improve performance.

Built-in Transaction Handling

One of the key enabling technologies for open channels is built-in end-to-end transaction support. In a crude way, this is what eBay offers in combination with its PayPal subsidiary, enabling one person to offer something for sale relatively easily and another person to buy it relatively easily.

The transaction support envisioned here for open channels must:

- Keep the cost per transaction very low.
- Be based on a standard pricing structure.
- Remain accessible over the life of the product.

Microbilling

Keeping the cost per transaction very low will allow transaction billings down to a scale of $0.01 USD (microbilling). More than any other single factor, effortless underlying support for microbilling will make it possible to sell digital content effectively online because it will allow pricing that both suppliers and consumers consider fair.

Without microbilling, the choice is generally between paying bricks-and-mortar prices for digital content ($25 per book, $15 per CD, $1 per song, etc.), and getting the content for free. One of the problems with "free" is that free subsidizes no investment by the suppliers in what they supply. They must support their contributions in other ways, and this significantly limits the quantity and quality of what is available.

Particularly with content produced by individuals, though — music, books, scientific and technical articles, etc. — there's no need to charge a lot of money. If a writer gets fifty cents per download of a novel, that's as much as a publisher would generally pay in royalties. If a song costs five cents instead of a dollar, most kids could afford to download material by the artists they love, and the artists would make out as well as or better than they do now.

Commercializing the free Internet's culture would not be so hard if prices could be lower. Up until now, there's been no way (and little will) to make that easy.

Standard Pricing Structure

Transactions would be based on a standard pricing structure having these elements:

- The amount the supplier will receive at the time of any sale (set by the supplier).
- The amount paid to whoever actually makes the sale (also set by the supplier). This applies whether the item is sold by itself or is aggregated as part of a bundle.
- A transaction charge (set by the system to cover money transfer and infrastructure costs). Possibly this can be eliminated by subsidizing the system in other ways.
- Shipping cost, where applicable (determined based on product size and weight, standard shipping tables published by the system, and a handling cost set by the supplier).
- Sales tax, where applicable, calculated at time of sale.

For example, if you were selling a short story you'd written, you might set a pricing structure like this:

Amount you get from every sale of any kind: $0.06
Amount a mediator (shop or aggregator) gets from every sale:.....$0.03
Channel transaction cost:.....$0.01
Invariant total retail price of the short story:.....$0.10

Such a price structure would make selling digital content profitable for content creators. If a million copies of your short story sold, you'd make $60,000. It would also reward mediators — the shops that carried your story would make $30,000 from those sales.

To take another example, you might price a photograph at:

Amount you get from every sale:.....$0.02
Amount a mediator (shop or aggregator) gets from every sale:.....$0.01
Channel transaction cost:.....$0.01
Invariant total retail price of the image:.....$0.04

If 200,000 people downloaded your photograph, you'd make $4,000. It's quite likely that 200,000 people would pay that price for a beautiful picture, especially if the process were effortlessly easy and safe.

In the case of a scientific or technical article, the value of mediation might be lower, and you might set a price structure like this:

Amount you get from each sale: $0.10
Amount a mediator (shop or aggregator) gets from each sale: $0.01
Channel transaction cost:.....$0.01
Invariant total retail price of the article:.....$0.12

If only 1,000 people read your article, you'd make $100, but if it were widely read throughout the world, you could easily make tens of thousands of dollars from it. Income at this level would be enough to buy you a measure of independence from academic bureaucracy, and encourage you to pursue research of interest to the world community.

Having the underlying system automatically and securely handle this kind of transaction structure not only makes buying and selling much easier, it also makes possible an effective marketplace for filtering and discovery of content, as we discuss below.

Pull Discovery, Not Push

One of the key differences between Internet shopping and bricks-and-mortar shopping is the way products are discovered. In the bricks-and-mortar world, push advertising rules, and its uses are well-understood. On the Internet, by contrast, most shopping is mediated by Google, which is why the company is so successful. If you are looking for something, you search for it — the shopper is much less passive in this process. Google serves essentially a market unifier for the entire Internet.

Although the smart structure of the network proposed here will support high-performance browsing down a wide variety of categorization hierarchies, as well as high-performance search, these are not enough to provide the kind of discoverability that will make the new system successful. What will be needed is commercial decision support provided by experts.

The Need for Better Product Filtering

The existence of global open channels will quickly make available an overwhelming number and variety of product options.

Consider the case of digital books: If you could publish a novel simply by paying $25, filling in a form, and uploading your manuscript, thousands of new books would become available every day, many of them mediocre or dreadful. The same goes for music. How could a consumer find the few gems in such a vast slush pile of unfiltered products?

Customer reviews are what Amazon uses to help buyers evaluate books and other products, but this doesn't work very well. It's all too easy for friends of an author to insert glowing tributes at the head of the list, or highly opinionated people to write misleading reviews. Better are the vendor reviews that Amazon, eBay, and others implement, to let you evaluate whether you wish to buy from a given third-party seller, but they're mainly useful for weeding out the worst vendors.

The best way to guarantee high-quality decision support for customers is to let experts make good money doing it well — and this is exactly what the standard pricing structure described above is designed to do. The mechanism, as discussed below, is the shop in the online marketplace.

Shops Provide Market-Driven Expert Human Mediation

The main function of shops in the online marketplace envisioned here is to select, review, and recommend products for a particular market segment. In the case of tangible goods, shops may also compete on price, although digital content will always sell at prices set by content providers.

A shop will be relatively easy to open and inexpensive to maintain, but building a customer base will require establishing and maintaining a good reputation. An individual with special expertise and/or taste will be able to make good money filtering products for a discerning market segment.

To take a music example, imagine a grandmother in South Dakota who has always been a heavy-metal head. She pays $50 to open a heavy-metal music shop in the online marketplace, and every day she spends a couple of hours listening to songs published in her category since the previous day. Most of them are terrible, but her ear is good once she turns up her hearing aid, and she picks out the few that rock. Before long, heavy-metal heads all over the world start realizing that a subscription to Grandma's weekly song list is well worth the dollar or so it costs, and before you know it, Grandma's shop is bringing her an income in six figures, while the little-known artists she's picking rise to stardom.

The reason this works is that Grandma doesn't have to ask anyone's permission to sell a song. She can't change the price, and she doesn't have to worry about the complexities of dividing up the resulting sales or calculating sales tax where applicable. All she has to do is pick songs, and maybe write pithy reviews full of four-letter words.

Paying people well to filter, select, validate, review, aggregate, and recommend products is absolutely essential to making a marketplace of open channels work smoothly. Sellers will find it wise not to stint on the percentage they allow for mediation, except in cases where demand is so high and their name so well known that they can keep all the profits for themselves.

Shops can be found by browsing through an extremely rich and complex kind of "Yellow Pages" provided by the marketplace interface (as described below), and can be searched for and filtered using various criteria. In addition, the system will help create another layer of mediation for larger market segments in the form of newsletters that review shops and keep people up to date about what's happening in their area of interest. In the longer term, publications of this sort will be profitable, self-sustaining businesses in themselves.

This kind of expert mediation would inevitably supplant generic search as the primary model for locating products — Google would almost certainly have to find a new business model!

Benefits of Pull Discovery

The dominance of pull discoverability over push advertising constitutes a crucial, fundamental difference between Internet commerce and its bricks-and-mortar predecessors. Right from the beginning, the Internet has allowed niche and specialized markets to be served in a way local retailers can't afford to do. Most bricks-and-mortar channels are explicitly managed so as to exclude products for which demand falls below a given threshold, even though the excluded markets may be quite sizable on a national or global level. The Internet's search engines have allowed retailers to sell to these neglected markets much more effectively than ever before, at little or no cost. Amazon's book-selling strategy is based on this principle, and eBay has taken garage sales international on this basis.

Continuing the shift from push advertising to pull discovery primarily benefits consumers and small vendors. Product quality, price, service, and reputation count more than advertising budget, allowing smaller vendors to enter the market more easily.

The significant amount of money that used to be needed for push advertising can now go to profits on one hand and lower prices on the other.

An Integrated, Streamlined Shipping Process

In collaboration with leading package carriers and governments around the world, and taking advantage of the channel infrastructure, the open channels project will promote the growth of a low-cost, reliable global package shipment network to allow tangible goods to flow easily through the system.

The security of the channels, the validated authentication of buyers and sellers, the built-in transaction support, and the ease of integrating with other systems will provide a convenient foundation for building an inexpensive global shipping network that can aggregate package shipment efficiently through hubs and gateways established by interested governments. This will allow people around the world to buy and sell tangible items through the channels conveniently, giving buyers full visibility of shipping costs and delivery dates before making a purchase.

A Rich Shopping Environment

Another important change required to lower channel friction to the point that the new system becomes irresistible is a really good shopping interface. Since we first described this concept in 2001, the game *Second Life* has validated some of these premises in practice, and now Sony is introducing a shopping experience mediated by the Playstation.

Pluggable User-Interface Modules

We envision a shopping user-interface that consists of three independent components:

- A “mall” navigation module.
- A “shop” navigation module.
- A “shop” setup and maintenance module.

Each of these components is intended to run locally on consumers’ machines so as to take full advantage of high-end graphics capabilities, although the system also ships a fallback Web interface for each.

All the data involved in offering a product for sale, including names, categories, descriptions, reviews, images, video, specifications, pricing information, display information, and so forth, will be stored across the network in a standard international schema owned and published by the system.

In addition, all interfaces needed for setting up shops, offering products for sale and purchasing them will be owned and implemented by the system and fully documented. Not only will full specifications for all three components be published, their reference implementations will be open-source so as to encourage third-party developers to create commercial alternatives that are even more desirable.

As a result, any third-party developer can create a replacement module for any one of the three interface components independently of the other two.

Mall Navigation Scenario

There are any number of ways mall navigation could be designed — the one described below uses a metaphor that could be both easy to use and entertaining.

Upon entering the mall (perhaps by clicking on its icon on your desktop), you find yourself flying high in the air in a slow circle above a mountainous semitropical island in the middle of an azure sea.

The island is perhaps 20 miles long and 10 wide, with an unusual variety of terrain. A dramatic mountain range of snow-capped peaks runs down the middle, surrounded by forests, grasslands, lakes, fields, orchards, and white sand beaches. Human habitations include an historic port city, large towns and villages, and assorted palaces, tree-dwellings, hobbit-holes, a space-port, and other curiosities of tourist appeal. A lone lighthouse tower near the peak of the highest mountain is the library, visible from anywhere on the island.

At the bottom of the screen is your control panel, designed to let you navigate around the island easily using different views. You can zoom in and out using your mouse wheel. There are several readouts that let you travel instantly to your favorite spots, to places you've recently visited, and to new attractions, or that keep you up to date on information about your account.

The navigation paradigm here is mainly spatial. As you zoom down to smaller and smaller scale, more and more detail becomes visible, or alternatively when you use search to reach a location that interests you, zooming up to a larger scale shows you the surrounding territory of related subject matter.

Most shops will initially have their primary location in one of the cities on the island — later, some communities will seek to segregate themselves for cultural reasons of various kinds and will explicitly abandon the main cities.

A city of arts might itself be divided into boroughs with names (localized for different supported natural languages and cultures) such as *Musica*, *Literatura*, *Graphica*, *Studia*, and *Voluptaria*. Each borough might in turn be divided into neighborhoods, with names such as *Poetry*, *Fiction*, *Science Fiction*, *Mystery*, and *Romance* in the borough of *Literatura*, or *Classical*, *Jazz*, *Rock*, *Hip Hop*, *Country,* and *World* in the borough of *Musica*.

Each neighborhood in turn has a warren of streets, each with its own name and a character that is influenced by the streets it intersects and its proximity to other neighborhoods. Finally, there are shops along the streets, located according to their area of specialization.

Each storefront would have a standard size, regardless of how big or small the shop behind it might be. This would simplify storefront design and re-use of designs for shopkeepers.

In such a world, of course, the same shop can have storefronts in as many different locations as the owner is willing to rent — digital space folds easily to permit this. Thus a customer can potentially enter a shop from any of a number of directions and may encounter different displays depending on what door he or she uses.

As you "walk" down a street, the storefronts all look different because their owners have customized them heavily with signs, window displays and banners. The only available opportunity for push advertising outside of the storefront itself is be on "posters" that

appear in specified neighborhoods or more generally around the city. Revenue from such advertisements would offset infrastructure costs and reduce charges for everyone.

As desirable locations become crowded, more streets and alleys can be added, and shops can appear on the second and higher floors as well.

Each street corner would have a "kiosk" where you could find information about all shops on the street, including feedback from their customers, and read news about the market area served by that street.

Simple as this kind of geographical and spatial metaphor is, it would be intuitive to users because it's so similar to the way people shop in the bricks-and-mortar world, and could easily create a powerful, extensible navigation model.

Shop Navigation

Once you enter a shop, the shop navigation module takes over control. It should go without saying that the size and layout of a shop's interior need bear no relationship whatsoever to the size of its storefront — it would be perfectly possible to enter an urban shop and find oneself outdoors in a kind of orchard in the mountains with products hanging from the branches of ancient fruit trees.

There are all kinds of ways to navigate inventories of various kinds, from very simple interfaces of the kind current Internet stores use to much more complex and interesting ones emulating bricks-and-mortar stores or creating entirely new models.

Shop navigation models must let customers find any specific item easily, and also browse through all available items in a way that makes sense.

In addition, buying an item should always be trivially easy, especially since all the payment details are handled by the system's built-in transaction system.

Shop Management Scenario

Since shops perform an essential service in the system, without which consumers would be lost in a sea of products, it will be essential from the beginning to have a good shop management interface. Shop owners will need to be able easily to:

- Select and add inventory
- Add information about products in inventory
- Aggregate products into bundles
- Manage and remove inventory
- Design storefronts
- Design shop floor plan
- Design shop spaces and decoration
- Design product displays and arrangements

Shopping Support Services

The system would provide ways for third-party services to integrate with it. For example:

- You might be able to sign up with a multi-shopper online service that would allow you create an avatar and go shopping with friends, using speech or text messaging to chat and exchange opinions as you go.
- You might be able to sign up with a service that would aggregate product reviews and/or pricing for a tangible product of interest.
- You might be able to sign up with a service that would attempt to find and acquire for you a specific tangible product at or below a given price.

Economic Implications

Bricks-and-mortar distribution models have been refined over the past several centuries, and are generally taken for granted by now. A widely accepted rule of thumb states that the manufacturing cost of a product must not exceed 20% of its expected retail price when sold through bricks-and-mortar channels. Assuming a gross manufacturing margin of 50%, this implies that distribution accounts for 60% of the retail cost of goods.

The promise of the kind of Internet system proposed here is to cut channel costs to 25% or less of retail price, resulting in an overall productivity gain of 35% or more. This by itself is dramatic, but there's more:

- Real-time communication between buyer and manufacturer would make just-in-time manufacturing an option in many more situations, resulting in significant inventory reductions.
- As more shopping moves online over time, there would be a reduction in bricks-and-mortar retail space requirements, resulting in less need for parking lots, less traffic, reduced fuel use, and so on.
- Online shopping could eliminate shrinkage through shoplifting, which is significant in some retail sectors.
- The shift from push to pull advertising, and the increased role of reliable paid product mediation would result in refocusing competition from shelf-space and packaging to product quality and fully satisfying consumer needs.

A Global Public Service

It is important to understand that the open channel system proposed here must be created as a community service like a public utility, and not as a profit-oriented commercial venture, for reasons that are strategic as much as altruistic:

- Without world-wide community support of a sort that a commercial venture could never attract, create a global open channel system is virtually impossible.
- Even if a monetized subset of the channel system could be implemented, the profitability of the channel system would attract competitors, thereby fragmenting the marketplace and reducing both its usefulness to consumers and its profitability.

A public project will be able to attract support from individuals, institutions and governments that no profit-oriented venture could ever hope to receive, and without such community support, creating the channels on a global scale is probably not feasible.

Furthermore, the long-term success and usefulness of the channels depends on keeping them open rather than controlling and monetizing them. This may go against all the received wisdom about bricks-and-mortar business, but it's an implicit premise in the logic of Internet commerce, and will be crucial to preserving the unity of the marketplace that the channels create. Monetized channels will inevitably lead to a fragmented marketplace that could only partially be reintegrated by technologies such as Google search. A single, unified low-friction market would be far more useful and valuable to everyone.

The success of craigslist.com illustrates this — the site has so little commercial competition because there's so little money to be made competing with its free services. As a result, it remains virtually the only place a customer needs to go for the services it provides, which in turn vastly increases its utility to customers and its long-term value and success.

Those contributing to the project must understand that the open channels will offer enormous new opportunities for profit that derive from *using* the channels to sell products, not from charging for their use.

Conclusion

The infrastructure improvements described here are entirely within reach today. The secure network infrastructure is the most critical component, but many of the hard problems associated with it have already been solved. The shopping interfaces and other improvements are equivalent, when taken together, to creating a relatively simple new game and game engine. All of this could be achieved at a level of investment that is common in much smaller projects.

The greatest challenge will be to make this vision real to all those communities of suppliers and consumers worldwide who stand to benefit most from open commercial channels.