

Chapter One

The Stigmergic Revolution

Several parallel developments are driving a trend toward the growing obsolescence of large, highly capitalized, hierarchical organizations, and the ability of networked individuals with comparatively cheap capital equipment to perform the functions formerly performed by such organizations. They include the drastically reduced cost of capital goods required for informational and material production, as well as drastically reduced transaction costs of coordinating efforts between individuals.

I. Reduced Capital Outlays

For most of the past two hundred years, the trend has been toward increasing capital outlays for most forms of production. The cost of the basic capital equipment required for production—the mass-production factory, the large printing press, the radio or TV station—was the primary justification for the large organization. The economy was dominated by large, hierarchical organizations administering enormous masses of capital. And the astronomical cost of production machinery was also the main justification for the wage system: production machinery was so expensive that only the rich could afford it, and hire others to work it.

In recent decades we've seen a reversal of this trend: a shift back from expensive, specialized machinery to inexpensive, general-purpose tools. Although this is true of both material and immaterial production—as attested by the recent revolution in garage-scale CNC machine tools¹—it was true first and most dramatically in the immaterial sphere.

The desktop computer is the primary item of capital equipment required for entering a growing number of industries, like music, desktop publishing and software design. The desktop computer, supplemented by assorted packages of increasingly cheap printing or sound editing equipment, is capable of doing what previously required a minimum investment of hundreds of thousands of dollars. In the words of Yochai Benkler: “The declining price of computation, communication, and storage have, as a practical matter, placed the material means of information and cultural production in the hands of a significant fraction of the world's population—on the order of a billion people around the globe.”² (Of course since that passage was written the proliferation of cheapening smartphones has probably more than doubled the latter figure.)

The growing importance of human capital, and the implosion of capital outlay costs required to enter the market, have had revolutionary implications for production in the immaterial sphere. In the old days, the immense outlay for physical assets was the primary basis for the corporate hierarchy's power, and in particular for its control over human capital and other intangible assets.

In many information and culture industries, according to Benkler, the initial outlay for entering the market in the days of “broadcast culture” was in the hundreds of thousands of dollars or more.

Since the introduction of the mechanical press and the telegraph, followed by the phonograph, film, the high-powered radio transmitter, and through to the cable plant or satellite, the capital costs of fixing in-

1 See Kevin Carson, *The Homebrew Industrial Revolution: A Low-Overhead Manifesto* (CreateSpace, 2010).

2 Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (New Haven and London: Yale University Press, 2006), p. 3.

formation and cultural goods in a transmission medium—a high-circulation newspaper, a record or movie, a radio or television program—have been high and increasing.³

The media of the broadcast era, for instance, were "typified by high-cost hubs and cheap, ubiquitous, reception-only systems at the end. This led to a limited range of organizational models for production: production in the information and entertainment industries was restricted to those who could collect sufficient funds to set up a hub."⁴ In the case of print periodicals, the increasing cost of printing equipment from the mid-nineteenth century on served as the main entry barrier for organizing the hubs. Between 1835 and 1850, the typical startup cost of a newspaper increased from \$500 to \$100,000—or from roughly \$10,000 to \$2.38 million in 2005 dollars.⁵ In other words, as the saying went, freedom of the press was great so long as you could afford to own a press.

The networked information economy, in contrast, is distinguished by "network architecture and the [low] cost of becoming a speaker."

The first element is the shift from a hub-and-spoke architecture with unidirectional links to the end points in the mass media, to distributed architecture with multidirectional connections among all nodes in the networked information environment. The second is the practical elimination of communications costs as a barrier to speaking across associational boundaries. Together, these characteristics have fundamentally altered the capacity of individuals, acting alone or with others, to be active participants in the public sphere as opposed to its passive readers, listeners, or viewers.⁶

Today most people, at least in the developed world, can afford to own a press. And thanks to smart phones, the same is rapidly becoming true for a large portion of the developing world's population as well.

In the old days, the owners of the hubs—CBS News, the Associated Press, etc.—decided what you could hear. Today you can set up a blog, or record a podcast, and anybody in the world who cares enough to go to your URL can look at it free of charge (and anyone who agrees with it—or wants to tear it apart—can provide a hyperlink to her readers).

The cultural authoritarianism that resulted from this state of affairs, as Clay Shirky points out, is unimaginable to someone who grew up with access to the Internet.

Despite half a century of hand-wringing about media concentration, my students have never known a media landscape of anything less than increasing abundance. They have never known a world with only three television channels, a world where the only choice a viewer had in the early evening was which white man was going to read them the news in English. They can understand the shift from scarcity to abundance, since the process is still going on today. A much harder thing to explain to them is this: if you were a citizen of that world, and you had something you needed to say in public, you couldn't. Period. Media content wasn't produced by consumers; if you had the wherewithal to say something in public, you weren't a consumer anymore, by definition. Movie reviews came from movie reviewers. Public opinions came from opinion columnists. Reporting came from reporters. The conversational space available to mere mortals consisted of the kitchen table, the water cooler, and occasionally letter writing....⁷

The central change that makes these things possible, according to Benkler, is that "the basic physical capital necessary to express and communicate human meaning is the connected personal computer."

The core functionalities of processing, storage, and communications are widely owned throughout the population of users.... The high capital costs that were a prerequisite to gathering, working, and com-

³ *Ibid.*, p. 51.

⁴ *Ibid.*, p. 179.

⁵ *Ibid.*, p. 188.

⁶ *Ibid.*, pp. 212-13.

⁷ Clay Shirky, *Cognitive Surplus* (New York: Penguin Press, 2010), pp. 60-61.

municating information, knowledge, and culture, have now been widely distributed in the society. The entry barrier they posed no longer offers a condensation point for the large organizations that once dominated the information environment.⁸

The desktop revolution and the Internet mean that the minimum capital outlay for entering most entertainment and information industries has fallen to a few hundred or a few thousand dollars, and the marginal cost of reproduction is zero.

The networked environment, combined with endless varieties of cheap software for creating and editing content, makes it possible for the amateur to produce output of a quality once associated with giant publishing houses and recording companies.⁹ That is true of the software industry, desktop publishing, and to a large extent even indie film (as witnessed by affordable editing technology and the success of projects like *Sky Captain*).

In the case of the music industry, thanks to cheap equipment and software for high quality recording and sound editing, the costs of independently producing and distributing a high-quality album have fallen through the floor. Bassist Steve Lawson writes:

...[T]he recording process—studio time and expertise used to be hugely expensive. But the cost of recording equipment has plummeted, just as the quality of the same has soared. Sure, expertise is still chargeable, but it's no longer a non-negotiable part of the deal. A smart band with a fast computer can now realistically make a release quality album-length body of songs for less than a grand....

What does this actually mean? Well, it means that for me—and the hundreds of thousands of others like me—the process of making and releasing music has never been easier. The task of finding an audience, of seeding the discovery process, has never cost less or been more fun. It's now possible for me to update my audience and friends (the cross-over between the two is happening on a daily basis thanks to social media tools) about what I'm doing—musically or otherwise—and to hear from them, to get involved in their lives, and for my music to be inspired by them....

So, if things are so great for the indies, does that mean loads of people are making loads of money? Not at all. But the false notion there is that any musicians were before! We haven't moved from an age of riches in music to an age of poverty in music. We've moved from an age of massive debt and no creative control in music to an age of solvency and creative autonomy. It really is win/win.¹⁰

As the last statement suggests, it may well be that most of the revenue loss to the music industry has fallen, not on actual performers, but on the rentiers and middlemen in the record companies themselves.

Networked distribution models have already gone a long way toward challenging and supplanting older models. For example the alternative rock group Radiohead marketed an album (*Rainbows*) directly over the Web, making it available for free and accepting whatever contributions downloaders saw fit to give. This would seem to be an ideal approach for independent artists, compared to the difficulty of making it through the record company gatekeepers and then settling for the royalties paid out after all the middlemen take their cut. It only requires, for all intents and purposes, a cheap website with a PayPal button. I have personal experience with a similar approach to publishing books, making them available for free online and selling hard copies through an on-demand publisher. And outside the blockbuster market, most writers and musical artists probably know more than the in-house marketing experts at the big content companies about their own niche markets. So they can do a better job marketing their own material virally to their target audiences through blogs, email lists and social networks than they would relying on the by-the-numbers efforts of the publishers' in-house promoters.

⁸ *Ibid.*, pp. 32-33.

⁹ *Ibid.*, p. 54.

¹⁰ Steve Lawson, "The Future of Music is... Indie!" *Agit8*, September 10, 2009 <<http://agit8.org.uk/?p=336>>.

This approach undermines the business model of the old record and publishing companies, and probably does cut into the revenues of their old stables of blockbuster artists. It's probably becoming a lot harder for another Stephen King or Mick Jagger to make megabucks because of competition from the networked distribution model, and surely a lot harder for the old gatekeeper corporations to make the giant piles of money they used to.

But if it's harder for the big boys to make gigantic piles of money, it's easier for a lot more little ones to make modest piles. Endless possibilities result from all the things they can now do for themselves, at virtually zero cost, that formerly only a highly capitalized record or publishing company could do for them.

As an independent scholar and author, I share Steve Lawson's view of things. From my perspective, the proper basis for comparison is the money I can make that I never could have made at all in the "good old days." In the good old days, I'd have—and have done—painstakingly put together a manuscript of hundreds of pages, and then put it away to gather cobwebs when I couldn't persuade the gatekeepers at a conventional publisher that it was worth marketing. Never mind whether the facsimile pdf's of my books available at torrent sites are costing me money (I don't think they are—I believe the free e-books are more like viral advertising). More importantly, if it weren't for digital publishing technologies and free publishing venues on the Internet, I would probably have lived and died doing menial labor with nobody anywhere ever hearing of my ideas. Thanks to digital culture, I'm able to make my work directly available to anyone in the world who has an Internet connection. If only a tiny fraction of the people who can read it for free decide to buy it, giving me a few thousand dollars a year in royalties, I'm richer by exactly that amount than I would have been in the "good old days" when my manuscripts would have yellowed in an attic.

That extra money may not be enough to support me by itself, but it's enabled me to pay off my debts and put away an F.U. fund equivalent to several months' wages. That probably puts me in a much better bargaining position vis-a-vis my employer than most people enjoy.

For every small full-time musician who has a harder time scraping by, and may have to supplement her performing revenues with a day job, I suspect there are ten people like me who would have spent their entire lives as (if you'll pardon the expression) mute inglorious Miltons, without ever making a goddamned cent from their music or writing, but who can now be heard. And for every blockbuster writer or musician who has a few million shaved off her multi-million dollar revenues as a result of online "piracy," I suspect there are probably a hundred people like me. No doubt people like Andrew Keen think the world is going to hell in a handbasket because so many riff-raff like me made it past the gatekeepers—but as far as I'm concerned that's another item to add to their T.S. lists.

As for the old broadcast media, podcasting makes it possible to distribute "radio" and "television" programming, at virtually no cost, to anyone with a broadband connection. As radio historian Jesse Walker notes, satellite radio's lackadaisical economic performance doesn't mean people prefer to stick with AM and FM radio; it means, rather, that the iPod has replaced the transistor radio as the primary portable listening medium, and that downloaded files have replaced the live broadcast as the primary form of content.¹¹

A network of amateur contributors has peer-produced an encyclopedia, Wikipedia, which Britannica sees as a rival.

There are enormous online libraries like Google Books and Project Gutenberg, as well as more specialized efforts like Marxists.org (which archives the collected works of Marx, Engels and Lenin, and of writers ranging from Kautsky to Luxemburg to Trotsky to C.L.R. James), the Anarchy Archives (extensive archives of most of the works of classical anarchism), and Constitution.org (including, among many other things, St. George Tucker's edition of Blackstone). In effect they give any kid with a smart phone, whether in the Third World or in an American ghetto, access to the equivalent of a university library. If one is willing and able to pay an annual subscription fee, there are enormous online collections of scholarly journals like JSTOR. And

11 Jesse Walker, "The Satellite Radio Blues: Why is XM Sirius on the verge of bankruptcy?," *Reason*, February 27, 2009 <<http://reason.com/news/show/131905.html>>.

rebellious scholars are in process of tearing down the paywalls and the textbook racket; scholars with JS-TOR memberships are providing articles for free to their peer networks. There are also services which strip DRM from college textbook pdfs which publishers make available for rental, so that they can be used indefinitely and distributed through torrent download sites.

The network revolution has drastically lowered the transaction costs of organizing education outside the conventional institutional framework. In most cases, the industrial model of education, based on transporting human raw material to a centrally located “learning factory” for processing, is obsolete. Forty years ago Ivan Illich, in *Deschooling Society*, proposed decentralized community learning nets that would put people in contact with the teachers they wanted to learn from, and provide an indexed repository of learning materials. The Internet has made this a reality beyond Illich's wildest dreams. MIT's OpenCourseware project was one early step in this direction. But most universities, even if they don't have a full database of lectures, at least have some sort of online course catalog with bare-bones syllabi and assigned readings for many individual courses.

Niall Cook, in *Enterprise 2.0*, describes the comparative efficiencies of software available outside the enterprise to the “enterprise software” in common use by employers. Self-managed peer networks, and individuals meeting their own needs in the outside economy, organize their efforts through social software and platforms chosen by the users themselves based on their superior usability for their purposes. And they are free to do so without corporate bureaucracies and their officially defined procedural rules acting as a ball and chain.

Enterprise software, in contrast, is chosen by non-users for use by other people of whose needs they know little (at best). Hence enterprise software is frequently a gold-plated turd. The IT departments I've encountered at all the places I've worked seemed to be run by people who'd read about “The Feds” in *Snow Crash* (or the Ministry of Information in *Brazil*), and decided to base a real-life human society on it, like that planet on *Star Trek* that stumbled across a book on 1930s gangland Chicago. And the corporate intranets looked like something designed by Rube Goldberg.

Blogs and wikis, and the free, browser-based platforms offered by Google and Mozilla, are a quantum improvement on the proprietary enterprise software that management typically forces on its employees. My OpenOffice CD cost me all of ten bucks, as opposed to \$200 for Microsoft Office. The kinds of productivity software and social software freely available to individuals in their private lives is far better than the enterprise software that corporate bureaucrats buy for a captive clientele of wage slaves—consumer software capabilities amount to “a fully functioning, alternative IT department.”¹² Corporate IT departments, in contrast, “prefer to invest in a suite of tools 'offered by a major incumbent vendor like Microsoft or IBM'.” System specs are driven by management's top-down requirements rather than by user needs.

...a small group of people at the top of the organization identify a problem, spend 12 months identifying and implementing a solution, and a huge amount of resources launching it, only then to find that employees don't or won't use it because they don't buy in to the original problem.¹³

Management is inclined “to conduct a detailed requirements analysis with the gestation period of an elephant simply in order to choose a \$1,000 social software application.”¹⁴ Employees often wind up using their company credit cards to purchase needed tools online rather than “wait for [the] IT department to build a business case and secure funding.”¹⁵ This is the direct opposite of agility.

It's just one particular example of the gold-plated turd phenomenon, in which stovepiped corporate design bureaucracies develop products for sale to other stovepiped corporate procurement bureaucracies, without the intervention of user feedback at any point in the process.

12 Niall Cook, *Enterprise 2.0: How Social Software Will Change the Future of Work* (Burlington, Vt.: Gower, 2008), p. 91.

13 *Ibid.*, p. 93.

14 *Ibid.*, p. 95.

15 *Ibid.*, p. 96.

As a result of all this, people are more productive away from work than they are at work. And management wonders why people would rather work at home using their own software tools than go through Checkpoint Charlie to use a bunch of klunky proprietary “productivity software” from the Whore of Redmond.

As Tom Coates put it, all these developments in the field of immaterial production mean that “the gap between what can be accomplished at home and what can be accomplished in a work environment has narrowed dramatically over the last ten to fifteen years.”¹⁶

II. Distributed Infrastructure and Ephemeralization

The larger and more hierarchical institutions become, and the more centralized the economic system, the larger the total share of production that will go to overhead, administration, waste, and the cost of doing business. The reasons are structural and geometrical.

At its most basic, it’s an application of the old cube-square rule. When you double the dimensions of a solid object, you increase its surface area fourfold (two squared), but its volume eightfold (two cubed). Similarly, the number of internal relationships in an organization increases as the square of the number of individuals making it up.

Leopold Kohr gave the example, in *The Overdeveloped Nations*, of a skyscraper. The more stories you add, the larger the share of floor space on each story is taken up by ventilation ducts, wiring and pipes, elevator shafts, stairwells, etc. Eventually you reach a point at which the increased space produced by adding stories is entirely eaten up by the increased support infrastructure.

The larger the scale of production, the more it must be divorced from demand, which means that the ostensible “economies” of large batch production are offset, and then more than offset, by the increasing costs of finding new ways of making people buy stuff that was produced without regard to preexisting orders.

The society becomes more and more like the Ministry of Central Services in *Brazil*, or The Feds in Neal Stephenson’s *Snow Crash*, and the distribution of occupations increasingly resembles the demographic profile of the promoters and middlemen in that crashed spaceship in *A Hitchhiker’s Guide to the Galaxy*, who founded the human race on Earth.

The only way out is a new standard of progress that doesn’t equate “growth” with larger institutional size and more centralization: scalable, distributed infrastructure, stigmergic organization, module-and-platform design configurations, and production capacity sited close to the point of consumption and scaled to demand.

Paul Hawken and the Lovinses, in *Natural Capitalism*, stated the general principle that when load-bearing infrastructures are built to handle the load at peak demand, most of the unit cost comes from the added infrastructure that comes from the increased usage during the tiny fraction of time when infrastructure experiences peak load. They gave the specific example of home heating, where enormous savings could be achieved by scaling capacity to handle only average usage, with additional demand handled through spot heating.

It’s a basic principle of lean production: most costs come from five percent of point consumption needs, and from scaling the capacity of the load-bearing infrastructure to cover that extra five percent instead of just handling the first ninety-five percent. It ties in, as well, with another lean principle: getting production out of sync with demand (including the downstream demand for the output of one step in a process), either spatially or temporally, creates inefficiencies. Optimizing one stage without regard to production flow and downstream demand usually involves expensive infrastructure to get an

16 Tom Coates, “(Weblogs and) The Mass Amateurisation of (Nearly) Everything...” Plasticbag.org, September 3, 2003 <http://www.plasticbag.org/archives/2003/09/weblogs_and_the_mass_amateurisation_of_nearly_everything>.

in-process input from one stage to another, often with intermediate storage while it is awaiting a need. The total resulting infrastructure cost greatly exceeds the saving at individual steps. Inefficient synchronization of sequential steps in any process results in bloated overhead costs from additional storage and handling infrastructure.

More generally, centralized infrastructures must be scaled to handle peak loads even when such loads only occur a small fraction of the time. And then they must amortize the extra cost, by breaking user behavior to the needs of the infrastructure.

At the opposite pole is distributed infrastructure that is mostly distributed among the endpoints, with relations directly between endpoints rather than passing through a central hub, and volume driven entirely by user demand at the endpoints. Since the capital goods possessed by the endpoints are a miniscule fraction of the cost of a centralized infrastructure, there is no incentive to subordinate end-users to the needs of the infrastructure.

The classic example is Bucky Fuller's: the replacement of the untold millions of tons of metal in transoceanic cables with a few dozen one-ton satellites. The entire infrastructure consists of satellite dishes at the endpoints communicating — via free, immaterial ether! — to the satellites.

Likewise systems that replace the fiber optic backbone with satellite connections and last-mile mesh-works.

Also the enormous infrastructure tied up in the civil aviation system's central hubs and batch-and-queue processing, as opposed to small jets flying directly between endpoints.

Another example is mass-production industry, which minimizes unit costs by running its enormously costly capital-intensive machinery at full capacity 24/7, and then requires organizing a society to guarantee consumption of the full output whether consumers want the shit or not — what's called "supply-push distribution." If consumers won't take it all, you soak up surplus output by destroying it through a permanent war economy, sinking it into an Interstate Highway System, etc. — or maybe just making stuff to fall apart.

The opposite of mass-production is distributed production on the Emilia-Romagna model described by Charles Sabel and Michel Piore in *The Second Industrial Divide*, with the capital infrastructure distributed to the point of consumption and output geared to local demand. The transnational corporate model of outsourcing is an attempt to put this new wine in old bottles. It distributes the production facilities, but does so on the basis of local labor cost rather than the location of market demand. So it still relies on the centralized wholesale infrastructure of warehouses on wheels/containerships, scaled to peak load, to transfer goods from the distributed production sites to the point of final consumption. The pure and unadulterated distributed manufacturing model, on the other hand, does away with this infrastructure by siting production at the last-mile network of consumption.

As we will see in later in this chapter, the model of stigmergic organization in Wikipedia and open-source design—the central theme of this book—is an example of distributed infrastructure. Individual contributions are managed entirely by endpoint users, coordinating their efforts with the finished body of work, without the intermediary of a centralized institutional framework as in old-line activist organizations.

III. Distributed Infrastructure and Scalability

Another advantage of distributed infrastructure is that it is scalable; that is, each separate part is capable of functioning on its own, regardless of whether the rest of the system is functioning. When a centralized infrastructure fails at any point, on the other hand, the whole system is incapacitated.

A large dam project must be completed to give service, and if something in the environment changes half way through the project, there is little hope of adapting the project to the new circumstances. The entire risk is assumed at the start of the project, based on long term projections about the future in many

different domains, from energy demand through to geopolitical stability. On the other hand, an array of micropower projects could provide equivalent electrical services, and as the projects are each built, continuous assessment of the “right next move” can be made to suit learning from previous projects, response to changing demand, adoption of improved technologies or shifting priorities. Fundamentally, half a dam is no dam at all, but 500 of 1000 small projects is half way to the goal. A modular approach to infrastructure in an uncertain world just makes sense.¹⁷

Take, for example, a wireless meshwork in which the end users themselves act as routers. It doesn't matter how many individual nodes are taken out of commission; so long as any nodes remain within communication distance of each other, the whole system survives.

IV. Network Organization

As Johan Soderburg argues, “[t]he universally applicable computer run on free software and connected to an open network... has in some respects leveled the playing field. Through the global communication network, hackers are matching the coordinating and logistic capabilities of state and capital.”¹⁸

Until the early 1990s, there were many possible Internets. What makes the Internet the “Internet” we know is really the World Wide Web: all the billions of web pages linked together by hyperlinks. And depending on the institutional context in which hyperlinks had been introduced, the Web as we know it might never have existed. Tim Berners-Lee in 1990,

published a more formal proposal... to build a "Hypertext project" called "WorldWideWeb"... as a "web" of "hypertext documents" to be viewed by "browsers" using a client-server architecture. This proposal estimated that a read-only web would be developed within three months and that it would take six months to achieve "the creation of new links and new material by readers, [so that] authorship becomes universal" as well as "the automatic notification of a reader when new material of interest to him/her has become available." While the read-only goal was met, accessible authorship of web content took longer to mature, with the wiki concept, blogs, Web 2.0 and RSS/Atom.¹⁹

The Web as we know it is something that could never have been built as the unified, conscious vision of any institution.

It's interesting that most visions of the “Information Superhighway,” back in the 1970s and 1980s, imagined it as populated largely by large institutional actors of one kind or another, and its communications as mainly one-way. It would be built on the backbone of the Internet's packet-switching infrastructure, vastly expanded in capacity by a fusion of the telephone and cable TV industries into a single high-bandwidth fiber-optic network.

I recall seeing a speculative article in *TV Guide* in the late '70s, when I was just a junior high school kid, speculating on the science fictiony wonders that would soon be possible. Everyone would have a combination digital telephone-computer-radio-cable TV terminal as the main entertainment center in their home, cable of accessing streaming content—television programs, movies, music, digitized books and periodicals, etc.—presumably on a paid basis. The key actors providing this whiz-bang content would be libraries, media conglomerates, and government agencies.

The Internet envisioned by figures like Al Gore and Bill Gates was, despite the decentralized nature of the physical packet-switching process, very centralized in terms of the actors providing content. Their vision of the Internet was simply as a foundation for the Information Superhighway. The legal infrastructure for the Superhighway consisted of the Telecommunications Act of 1996, which eliminated barriers to tele-

17 Vinay Gupta, *The Global Village Development Bank* (2009/03/12 draft 2).

18 Johan Soderberg, *Hacking Capitalism: The Free and Open Source Software Movement* (New York and London: Routledge, 2008), p. 2.

19 “World Wide Web,” Wikipedia <http://en.wikipedia.org/wiki/The_web> Accessed September 27, 2011.

phone/cable mergers, and the Digital Millennium Copyright Act of 1998, which created the draconian system of copyright law needed for digital content providers to turn the Superhighway into a turnpike. Here's what Bill Gates had to say, as late as early 2000:

This new generation of set-top boxes that connects up to the Internet is very much part of that. The potential impact is pretty phenomenal in the terms of being able to watch a TV show whenever you want to. There will be so many choices out there. You've got to imagine that a software agent will help you find things that you might be interested in.

...The "TV guide" will almost be like a search portal where you'll customize and say, "I'm never interested in this, but I am particularly interested in that." It's already getting a little unwieldy. When you turn on DirectTV and you step through every channel—well, there's three minutes of your life.

When you walk into your living room six years from now, you'll be able to just say what you're interested in, and have the screen help you pick out a video that you care about. It's not going to be "Let's look at channels 4, 5, and 7." It's going to be something that has pretty incredible graphics and it's got an Internet connection to it.²⁰

But the Information Superhighway—in the sense of a fusion of telephone, cable, radio, and on-demand music and movies, accessed through a single digital home entertainment center, simply fizzled out. Instead, the World Wide Web took over the Internet.

Mike Masnick speculates on what the World Wide Web—if it could even be called that—would have looked like, had Tim Berners-Lee obtained a patent on the hyperlinked architecture of the Web. And his hypothetical description reads very close to the vision of *TV Guide*, Gore and Gates.

Where do you think the world would be today if the World Wide Web had been patented? Here are a few guesses:

- Rather than an open World Wide Web, most people would have remained on proprietary, walled gardens, like AOL, Compuserve, Prodigy and Delphi. While those might have eventually run afoul of the patents, since they were large companies or backed by large companies, those would have been the few willing to pay the licensing fee.
- The innovation level in terms of the web would have been drastically limited. Concepts like AJAX, real time info, etc. would not be present or would be in their infancy. The only companies "innovating" on these issues would be those few large players, and they wouldn't even think of the value of such things.
- No Google. Search would be dismal, and limited to only the proprietary system you were on.
- Most people's use of online services would be more about "consumption" than "communication." There would still be chat rooms and such, but there wouldn't be massive public communication developments like blogs and Twitter. There might be some social networking elements, but they would be very rudimentary within the walled garden.
- No iPhone. While some might see this as separate from the web, I disagree. I don't think we'd see quite the same interest or rise in smartphones without the web. Would we see limited proprietary "AOL phones?" Possibly, but with a fragmented market and not as much value, I doubt there's the necessary ecosystem to go as far as the iPhone.
- Open internet limited by lawsuit. There would still be an open internet, and things like gopher and Usenet would have grown and been able to do a little innovation. However, if gopher tried to expand to be more web like, we would have seen a legal fight that not only delayed innovation, but limited the arenas in which we innovated.²¹

20 "The Emperor Strikes Back—Bill Gates Interview," *Entertainment Weekly*, January 7, 2000

<<http://www.angelfire.com/nt/vapor/bginterview.html>>. Note—I corrected numerous errors, presumably the result of transcription or scanning error, from the online version.

21 Mike Masnick, "What If Tim Berners-Lee Had Patented The Web?" *Techdirt*, August 11, 2011

<<http://www.techdirt.com/articles/20110811/10245715476/what-if-tim-berners-lee-had-patented-web.shtml>>.

The Internet would have been a wasteland of walled-garden ISPs like AOL, with Usenet and BBSs grafted on. What Web there was would have been accessed, not by browsers or open search engines, but through portals like AOL or Yahoo!.

It's not necessary to speculate that something like that would surely had happened had Berners-Lee not been first to the draw. It was happening, in fact. As recounted by David Weinberger, the software company for which he was vice president of strategic marketing at the time was in process of developing a proprietary document format with embedded links, when it was caught off-guard by the Mosaic browser. As the developers attempted to reassure themselves, their software was far more polished and professional-looking, and had better capabilities, than Mosaic. But deep down, they knew that Mosaic's lack of "bells and whistles" was more than compensated for by its openness.

With our software, a publisher could embed a link from one document to another, but the publisher had to own both documents. That's fine if you're putting together a set of aircraft maintenance manuals and you want to make all the cross-references active, so that clicking on one brings up the page to which it's referring. But those links had to be compiled into the system. Once the document was published, no more links could be added except by recompiling the document. And, most important, the only people who could add new links were those working for the publisher. If you were an aircraft mechanic who had discovered some better ways to clean a fuel line, you had no way to publish your page with our system and no way to link it to the appropriate page in the official manual.

The Web, on the other hand, breaks the traditional publishing model. The old model is about control: a team works on a document, is responsible for its content and format, and releases it to the public when it's been certified as done. Once it's published, no one can change it except the original publisher. The Web ditches that model, with all its advantages as well as its drawbacks, and says instead, "You have something to say? Say it. You want to respond to something that's been said? Say it and link to it. You think something is interesting? Link to it from your home page. And you never have to ask anyone's permission." ...By removing the central control points, the Web enabled a self-organizing, self-stimulated growth of contents and links on a scale the world has literally never before experienced.²²

Although some idiots like Rupert Murdoch still attempt to gainsay it, the basic organizing principle of the Web as envisioned by Berners-Lee is that you can link to another person's website without having to ask permission or secure her cooperation.²³

It was actually the collapse of Web 1.0 in the dot-com bubble, and with it most of the hopes of the "visionaries" of the 1990s for enclosing the Web as a source of revenues, that created the space in which the decentralized vision of Web 2.0 could be fully realized. As Foundation for P2P Alternatives founder Michel Bauwens described it:

All the pundits were predicting, then as now, that without capital, innovation would stop, and that the era of high internet growth was over for a foreseeable time. In actual fact, the reality was the very opposite, and something apparently very strange happened. In fact, almost everything we know, the Web 2.0, the emergence of social and participatory media, was born in the crucible of that downturn. In other words, innovation did not slow down, but actually increased during the downturn in investment. This showed the following new tendency at work: capitalism is increasingly being divorced from entrepreneurship, and entrepreneurship becomes a networked activity taking place through open platforms of collaboration.

The reason is that internet technology fundamentally changes the relationship between innovation and capital. Before the internet, in the Schumpeterian world, innovators need capital for their research, that research is then protected through copyright and patents, and further funds create the necessary factories. In the post-schumpeterian world, creative souls congregate through the internet, create new software, or any kind of knowledge, create collaboration platforms on the cheap, and paradoxically, only

²² David Weinberger, *Small Pieces Loosely Joined: A Unified Theory of the Web* (Cambridge, Mass.: Perseus Publishing, 2002), vii-ix.

²³ *Ibid.*, pp. 52-53.

need capital when they are successful, and the servers risk crashing from overload. As an example, think about Bittorrent, the most important software for exchanging multimedia content over the internet, which was created by a single programmer, surviving through a creative use of some credit cards, with zero funding. But the internet is not just for creative individual souls, but enables large communities to cooperate over platforms. Very importantly, it is not limited to knowledge and software, but to everything that knowledge and software enables, which includes manufacturing. Anything that needs to be physically produced, needs to be ‘virtually designed’ in the first place.²⁴

The ghost of Web 1.0, still clanking the chains of Al Gore's vision, had to be laid to rest before the World Wide Web could reach its full potential. (Of course no less a figure than Jaron Lanier is still weeping for Tammuz, and we may need more incantations and holy water against Apple's attempts to call up its ghost like the witch of Endor.)

The Web's many-to-many communications capabilities have enabled networks to coordinate the actions of self-directed individuals without the transaction costs of traditional hierarchies. Benkler explained the implications of networked communications, combined with the near-universal distribution of capital goods for information and cultural production:

...the technical architectures, organizational models, and social dynamics of information production and exchange on the Internet have developed so that they allow us to structure the solution to problems—in particular to information production problems—in ways that are highly modular. This allows many diversely motivated people to act for a wide variety of reasons that, in combination, cohere into new useful information, knowledge, and cultural goods. These architectures and organizational models allow both independent creation that coexists and coheres into usable patterns, and interdependent cooperative enterprises in the form of peer-production processes.²⁵

In other words, it's stigmergic organization (about which more below)—what Weinberger calls “small pieces loosely joined.”

Networked crowdsourcing venues like Kickstarter have radically lowered the costs of aggregating capital even when total outlays are still beyond the means of the average individual. That means that, in addition to circumstances like information and cultural production where physical capital outlays are minimal because the desktop computer is the main item of capital equipment—even when the costs of the physical capital required for production are non-trivial—the transaction costs of aggregating the required investment capital from a number of small contributors, or of putting a user community in touch with the owners of spare capacity of capital goods, are much lower.

But whether capital outlay requirements are large or small, network technology has had a revolutionary effect on the transaction costs of traditional organization.

That was true even back in the 1990s, when the Internet was dominated by static institutional websites. Email, both individual and in discussion lists, was a powerful tool for networked organization. The forms of culture jamming described by Naomi Klein in *No Logo*, themselves unprecedented and revolutionary in her day, were an outgrowth of the possibilities of the Web 1.0 of the 1990s. But the rise of Web 2.0, and the free platforms it made available, increased the possibilities exponentially. To quote Benkler again:

What we are seeing now is the emergence of more effective collective action practices that are decentralized but do not rely on either the price system or a managerial structure for coordination.... [The networked environment] provides a platform for new mechanisms for widely dispersed agents to adopt radically decentralized cooperation strategies other than by using proprietary and contractual claims to

24 Michel Bauwens, “Asia Needs a Social Innovation Stimulus Plan,” *P2P Foundation Blog*, March 23, 2009 <<http://blog.p2pfoundation.net/asia-needs-a-social-innovation-stimulus-plan/2009/03/23>>.

25 Benkler, *The Wealth of Networks*, pp. 105-106.

elicit prices or impose managerial commands.... What we see in the networked information economy is a dramatic increase in the importance and the centrality of information produced in this way.²⁶

Consider the drastically lowered costs of aggregating people—in large or small numbers, with common interests or concerns of any kind—into affinity groups or movements for the sharing of information and taking concerted action. Clay Shirky cites the example of Voice of the Faithful, a Catholic lay organization formed to fight priestly sexual abuse:

Had VOTF been founded in 1992, the gap between hearing about it and deciding to join would have presented a series of small hurdles: How would you locate the organization? How would you contact it? If you requested literature, how long would it take to arrive, and by the time it got there, would you still be in the mood? None of these barriers to action is insurmountable, but together they subject the desire to act to the death of a thousand cuts.²⁷

Because of the delays and costs involved, going from a couple dozen people in a basement to a large and global organization in six months is inconceivable without social tools like websites for membership and e-mail for communication.

I can remember, as a grad student in the 1980s, experiencing that “series of small hurdles” in dealing with a completely different—but analogous—situation. If I heard of some periodical in my area of interest that the university library didn't carry, the only way to find out more about it was to dig through the latest installment of *Ulrich's Periodicals Directory*, send a query letter soliciting information about the price of sample issues, wait several weeks for a response, send in the money, and wait several more weeks for my sample. Or I could go through one of the scholarly reference sources like *Social Sciences Index* to look for the kinds of articles that appeared in that journal, and order some likely-sounding ones through Interlibrary Loan to see if a subscription would be worth it.

Today, I just Google the title of the journal, and most likely it's got a website with an index of past issues. I can instantly get a pdf of any article of interest through an online indexing service like SSRN (albeit at a hefty price). Better yet, I can shop for a free torrent download of the same article from a file-sharing site, or ask a friend with a JSTOR or SSRN membership to email me a pdf as an attachment. Soon, dedicated sharing sites with indexed academic articles available free for scholars will probably be as common as mp3-sharing sites—much to the chagrin of the Copyright Nazis in the academic publishing industry.

The cumulative effect is that a rapidly increasing share of the functions previously carried out by corporations and by the state can now be effectively carried out by what Marx and Engels, in *The Communist Manifesto*, called the “associated producers”—without any bureaucratic intermediation. As Matthew Yglesias put it:

In its 1875 Gotha Program, the German Social Democratic Party “demand[ed] the establishment of socialistic productive associations with the support of the state and under the democratic control of the working people.”

It turns out that finding a feasible way to do that for industrial age enterprises was fairly problematic. And yet their arguments that such associations would be beneficial remain compelling. Meanwhile, the Internet makes it much easier for individuals to form socialistic productive associations without a ton of explicit support of the state.²⁸

Another result of the reduced threshold for communications in networks is the drastic increase in speed of propagation.

26 Benkler, *The Wealth of Networks*, p. 63.

27 Shirky, *Here Comes Everybody* (Penguin, 2008), p. 151.

28 Matthew Yglesias, “Actually Existing Internet Communism,” *Yglesias*, November 9, 2010 <<http://yglesias.thinkprogress.org/2010/11/actually-existing-internet-communism/>>.

Not only have connective technologies increased the breadth of influence of agents in the network, they have also severely increased the speed of propagation. Smart mobs are essentially a rapid cascade of coordinated action. "Whenever a new communications technology lowers the threshold for groups to act collectively, new kinds of institutions emerge.... We are seeing the combination of network communications and social networks."²⁹

V. Stigmergy

Networked organization is based on a principle known as stigmergy. "Stigmergy" is a term coined by biologist Pierre-Paul Grasse in the 1950s to describe the process by which termites coordinate their activity. Social insects like termites and ants coordinate their efforts through the independent responses of individuals to environmental triggers like chemical markers, without any need for a central coordinating authority.³⁰ It was subsequently applied to the analysis of human society.

Up to about 1990, the notion of stigmergy appears to have remained limited to a small circle of researchers studying the behavior of social insects. However, one of these insect specialists, Jean-Louis Deneubourg, was also a member of the "Brussels School" of complex systems, headed by the late Nobel Prize in chemistry, Ilya Prigogine. In this interdisciplinary environment, it became clear that stigmergy was a prime example of spontaneous ordering or self-organization, and as such potentially applicable to complex systems other than insect societies.³¹

Applied by way of analogy to human society, stigmergy refers primarily to the kinds of networked organization associated with wikis, group blogs, and "leaderless" organizations configured along the lines of networked cells.

The termites do not communicate about who is to do what how or when. Their only communication is indirect: the partially executed work of the ones provides information to the others about where to make their own contribution. In this way, there is no need for a centrally controlled plan, workflow, or division of labor.

While people are of course much more intelligent than social insects and do communicate, open access development uses essentially the same stigmergic mechanism...: any new or revised document or software component uploaded to the site of a community is immediately scrutinized by the members of the community that are interested to use it. When one of them discovers a shortcoming, such as a bug, error or lacking functionality, that member will be inclined to either solve the problem him/herself, or at least point it out to the rest of the community, where it may again entice someone else to take up the problem.³²

Mark Elliott, whose doctoral dissertation is probably the most thorough and comprehensive treatment of stigmergy to date, contrasts stigmergic coordination with social negotiation. Social negotiation is the traditional method of organizing collaborative group efforts, through agreements and compromise mediated by

29 Paul Herzog, "21st Century Governance as a Complex Adaptive System" (Political Science, University of Utah), p. 4. *Proceedings PISTA*. Informatics and Society Series (Orlando: International Institute of Informatics and Systemics, 2004). The quote is from Howard Rheingold, *Smart Mobs and Antiwar Protests* (2003).

30 Mark Elliott, "Stigmergic Collaboration: The Evolution of Group Work," *M/C Journal*, May 2006 <<http://journal.media-culture.org.au/0605/03-elliott.php>>.

31 Francis Heylighen, "Stigmergy as a Universal Coordination Mechanism: components, varieties and applications." To appear in: Lewis, Ted & Marsh, Leslie (eds), *Human Stigmergy: Theoretical Developments and New Applications* (Studies in Applied Philosophy, Epistemology and Rational Ethics, Springer, 2015) <<http://pespmc1.vub.ac.be/papers/stigmergy-varieties.pdf>>, p. 2.

32 Heylighen, "Why is Open Access Development so Successful? Stigmergic organization and the economics of information," draft contribution to B. Lutterbeck, M. Bärwolff & R. A. Gehring (eds.), *Open Source Jahrbuch 2007* (Lehmanns Media, 2007) <<http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>>, p. 7.

discussions between individuals. The exponential growth in the number of communications with the size of the group, obviously, imposes constraints on the feasible size of a collaborative group, before coordination must be achieved by hierarchy and top-down authority. Stigmergy, on the other hand, permits collaboration on an unlimited scale by individuals acting independently. This distinction between social negotiation and stigmergy is illustrated, in particular, by the contrast between traditional models of co-authoring and collaboration in a wiki.³³ Individuals communicate indirectly, “via the stigmergic medium.”³⁴

The distinction between social negotiation and stigmergic coordination parallels Elliott's distinction, elsewhere, between “discursive collaboration” and “stigmergic collaboration.” The “discursive elaboration of shared representations (ideas)” is replaced by “the annotation of material and digital artefacts as embodiments of these representations.” “Additionally, when stigmergic collaboration is extended by computing and digital networks, a considerable augmentation of processing capacity takes place which allows for the bridging of the spatial and temporal limitations of discursive collaboration, while subtly shifting points of negotiation and interaction away from the social and towards the cultural.”³⁵

Modular, building-block structures are ubiquitous. Why? Because such a structure “transforms a system's ability to learn, evolve and adapt... Once a set of building blocks... has been tweaked and refined and thoroughly debugged through experience... then it can generally be adapted and recombined to build a great many new concepts... Certainly that's a much more efficient way to create something new than starting all over from scratch. And that fact, in turn, suggests a whole new mechanism for adaptation in general. Instead of moving through that immense space of possibilities step by step, so to speak, an adaptive system can reshuffle its building blocks and take giant leaps.” A small number of building blocks can be shuffled and recombined to make a huge number of complex systems.³⁶

If you start with a large number of modular individuals, each capable of interacting with a few other individuals, and acting on other individuals according to a simple grammar of a few rules, under the right circumstances the modular individuals can undergo a rapid phase transition, according to systems theorist Stuart Kauffman: “The growth of complexity really does have something to do with far-from-equilibrium systems building themselves up, cascading to higher and higher levels of organization. Atoms, molecules, autocatalytic sets, et cetera.” And the new higher level entities, in turn, can interact among themselves, perhaps creating another autocatalytic phase transition to a higher level.³⁷

Gus diZerega's discussion of spontaneous orders closely parallels the concept of stigmergy:

Hayek and Polanyi identified the basic processes that generate spontaneous orders. They arise from networks of independent equals whose actions generate positive and negative feedback that help guide future actors in pursuing their own independently conceived plans, thereby continuing the feedback process. Each person is a node within a network and is linked by feedback, with each node free to act on its own. The feedback they generate minimizes the knowledge anyone needs about the system as a whole in order to succeed within it.

All spontaneous orders possess certain abstract features in common. Participants are equal in status and all are equally subject to whatever rules must be followed to participate within the order. All are free to apply these rules to any project of their choosing. Anything that can be pursued without violating a rule is permitted, including pursuing mutually contradictory goals. Finally, these rules facilitate cooperation among strangers based on certain broadly shared values that are simpler than the values ac-

33 Elliott, “Stigmergic Collaboration.”

34 Mark Elliott, “Some General Off-the-Cuff Reflections on Stigmergy,” *Stigmergic Collaboration*, May 21, 2006 <<http://stigmergiccollaboration.blogspot.com/2006/05/some-general-off-cuff-reflections-on.html>>.

35 Mark Elliott, *Stigmergic Collaboration: A Theoretical Framework for Mass Collaboration*. Doctoral Dissertation, Centre for Ideas, Victorian College of the Arts, University of Melbourne (October 2007), pp. 9-10

36 M. Mitchell Waldrop, *Complexity: The Emerging Science at the Edge of Order and Chaos* (New York, London, Toronto, Sydney, Tokyo, Singapore: Simon & Schuster, 1992), pp. 169-170.

37 *Ibid.*, pp. 316-317.

tually motivating many people when they participate. Compared to human beings, spontaneous orders are “value-thin.”³⁸

David de Ugarte quotes the Rand theorists John Arquilla and David Ronfeldt, in “Swarming and the Future of Conflict.” “[N]etwar,” they say,

is a privateers’ war in which many small units “already know what they must do”, and are aware that “they must communicate with each other not in order to prepare for action, but only as a consequence of action, and, above all, through action.”³⁹

Critics of “digital communism” like Jaron Lanier and Mark Helprin, who condemn network culture for submerging “individual authorial voice” in the “collective,” couldn’t be more clueless if they tried. Stigmergy synthesizes the highest realizations of both individualism and collectivism, and represents each of them in its most completely actualized form, without qualifying or impairing either in any way. Michel Bauwens uses the term “cooperative individualism”:

this turn to the collective that the emergence of peer to peer represents does not in any way present a loss of individuality, even of individualism. Rather it ‘transcends and includes’ individualism and collectivism in a new unity, which I would like to call ‘cooperative individualism’. The cooperativity is not necessarily intentional (i.e. the result of conscious altruism), but constitutive of our being, and the best applications of P2P, are based on this idea.⁴⁰

Stigmergy is not “collectivist” in the traditional sense, as it was understood in the days when a common effort on any significant scale required a large organization to represent the collective, and the administrative coordination of individual efforts through a hierarchy. But it is the ultimate realization of collectivism, in that it removes the transaction cost of concerted action by many individuals.

It is the ultimate in individualism because all actions are the free actions of individuals, and the “collective” is simply the sum total of individual actions. Every individual is free to formulate any innovation she sees fit, without any need for permission from the collective. Every individual or voluntary association of individuals is free to adopt the innovation, or not, as they see fit. The extent to which any innovation is adopted results entirely from the unanimous consent of every voluntary grouping that adopts it. Each innovation is modular (meaning the project “can be broken down into smaller components... that can be independently produced before they are assembled into a whole”⁴¹), and may be adopted into any number of larger projects where it is found useful. Any grouping where there is disagreement over adoption may fork and replicate their project with or without the innovation.

In this regard it attains the radical democratic ideal of *unanimous* consent of the governed, which is never completely possible under any representative or majoritarian system. Consent—the individual’s participation in the decisions that affected her—was the central value of Jeffersonian democracy. The smaller the unit of governance, and the closer it was to the individual, the closer it approached the ideal of unanimous consent to all acts of government. Hence Jefferson’s ward republics, whose chief virtue was the increased role of each individual in influencing the outcome of policy. But this ideal can only be fully attained when the unit of governance is the individual. So majority rule was the lesser evil, a way to approximate as closely as possible to the spirit of unanimous consent when an entire group of people had to be bound by a single decision. Stigmergy removes the need for any individual to be bound by the group will. When all

38 Gus DiZerega, “Outlining a New Paradigm,” *Cosmos and Taxis* 1:1 (2013), p. 9.

39 David de Ugarte, *The Power of Networks : An Illustrated Manual for People, Collectives, and Companies Driven to Cyberactivism*. Translated by Asunción Álvarez (n.d.), p. 62 <<http://deugarte.com/gomi/the-power-of-networks.pdf>>.

40 Michel Bauwens, “Individuality, Relationality, and the Collective in the P2P era,” *P2P Foundation Blog*, May 15, 2010 <<http://blog.p2pfoundation.net/individuality-relationality-and-the-collective-in-the-p2p-era/2010/05/15>>.

41 Benkler, *The Wealth of Networks*, p. 100.

group actions reflect the unanimous will of the participants, as permitted by stigmergic organization, the ideal of unanimous consent is finally achieved in its fullness.

Group action is facilitated with greater ease and lower transaction costs than ever before, but all “group actions” are the unanimous actions of the participating individuals. As described by Heather Marsh:

With stigmergy, an initial idea is freely given, and the project is driven by the idea, not by a personality or group of personalities. No individual needs permission (competitive) or consensus (cooperative) to propose an idea or initiate a project. There is no need to discuss or vote on the idea, if an idea is exciting or necessary it will attract interest. The interest attracted will be from people actively involved in the system and willing to put effort into carrying the project further, not empty votes from people with little interest or involvement. Since the project is supported or rejected based on contributed effort, not empty votes, input from people with more commitment to the idea will have greater weight. Stigmergy also puts individuals in control over their own work, they do not need group permission to tell them what system to work on or what part to contribute.

The person with the initial idea may or may not carry the task further. Evangelizing the idea is voluntary, by a group that is excited by the idea; they may or may not be the ones to carry it out. It is unnecessary to seek start up funding and supporters; if an idea is good it will receive the support required.... Secrecy and competition is unnecessary because once an idea is given, it and all new development belongs to anyone who chooses to work on it. Anyone can submit work for approval, the idea cannot die or be put on hold by personalities; acceptance or rejection is for the work contributed, not the person contributing it. All ideas are accepted or rejected based on the needs of the system.

Responsibility and rights for the system rest with the entire user group, not just the creators. There is no need for people to leave the system based on personality conflicts as there is no need for communication outside of task completion and there are usually plenty of jobs with complete autonomy. As no one owns the system, there is no need for a competing group to be started to change ownership to a different group....

Communication between nodes of a system is on an as needed basis. Transparency allows information to travel freely between the various nodes, but a formal relationship or communication method is neither necessary nor desirable. Information sharing is driven by the information, not personal relationships. If data is relevant to several nodes it will be immediately transmitted to all, no formal meetings between official personalities are necessary.

Any node can disappear without affecting the network, and the remaining necessary functionality of that node can be taken up by others. Nodes which find they are performing the same tasks will likely join, or one will be rendered obsolete by lack of use. New nodes are only created to fulfill a new need or provide greater functionality; it is inefficient to have the same task performed twice, and that only occurs if a second group discovers an alternative method that the first group is unwilling to adopt. In that case, the best system will win the most support from the user group, the other will die or remain as a valued alternative. Any user can contribute to the node which best matches their interests and abilities, or contribute to multiple nodes.

A new system of governance or collaboration that does not follow a competitive hierarchical model will need to employ stigmergy in most of its action based systems. It is neither reasonable nor desirable for individual thought and action to be subjugated to group consensus in matters which do not affect the group, and it is frankly impossible to accomplish complex tasks if every decision must be presented for approval; that is the biggest weakness of the hierarchical model. The incredible success of so many internet projects are the result of stigmergy, not cooperation, and it is stigmergy that will help us build quickly, efficiently and produce results far better than any of us can foresee at the outset.⁴²

Another remarkable thing about stigmergic coordination is that free riders are not a problem; all actions are voluntarily undertaken out of self-interest, and their service to the individuals undertaking them and to the group is not lessened by the fact that others free ride without contributing.

42 Heather Marsh, “Stigmergy,” *GeorgieBC's Blog*, December 24, 2012 <<http://georgiebc.wordpress.com/2012/12/24/stigmergy-2/>>.

In the stigmergic paradigm, the common good (e.g. Wikipedia, or a network of trails and roads connecting common destinations) is gradually built up via the cooperation implicit in stigmergically coordinated actions. Free riders may profit from this common good without putting in any effort in return. However, the benefit derived from a stigmergic trace does not in general reduce the value of that trace. For example, an ant that follows a pheromone trace laid by others without adding pheromone of its own does not by that action make the pheromone trace less useful to the other ants. Similarly, a person who downloads a piece of open source software without contributing to the development of that software does not impose any burden on the software developers. Thus, in a situation of stigmergy, a free rider or “defector” does not weaken the cooperators, in contrast to situations like the Prisoners’ dilemma or Tragedy of the Commons.⁴³

In short, as Michel Bauwens describes it, “Peer production is based on the elimination of permission-asking and a shift to the self-selection of tasks...”⁴⁴

A good example is Raymond’s “Bazaar” model of open-source development, as illustrated in a hypothetical case by Benkler:

Imagine that one person, or a small group of friends, wants a utility. It could be a text editor, photo-retouching software, or an operating system. The person or small group starts by developing a part of this project, up to a point where the whole utility—if it is simple enough—or some important part of it, is functional, though it might have much room for improvement. At this point, the person makes the program freely available to others, with its source code—instructions in a human-readable language that explains how the software does whatever it does when compiled into a machine-readable language. When others begin to use it, they may find bugs, or related utilities that they want to add (e.g., the photo-retouching software only increases size and sharpness, and one of its users wants it to allow changing colors as well). The person who has found the bug or is interested in how to add functions to the software may or may not be the best person in the world to actually write the software fix. Nevertheless, he reports the bug or the new need in an Internet forum of users of the software. That person, or someone else, then thinks that they have a way of tweaking the software to fix the bug or add the new utility. They then do so, just as the first person did, and release a new version of the software with the fix or the added utility. The result is a collaboration between three people—the first author, who wrote the initial software; the second person, who identified a problem or shortcoming; and the third person, who fixed it. This collaboration is not managed by anyone who organizes the three, but is instead the outcome of them all reading the same Internet-based forum and using the same software, which is released under an open, rather than proprietary, license. This enables some of its users to identify problems without asking anyone’s permission and without engaging in any transactions.⁴⁵

Nevertheless, the creation of value itself is inherent in the network as an entity—a form of network effect that is more than the sum of the individual parts. Michael Hardt’s and Antonio Negri’s discussion of value production on the commons is relevant here:

Once one adopts the standpoint of the common, many of the central concepts of political economy have to be rethought. In this context, for instance, valorization and accumulation necessarily take on a social rather than an individual character. The common exists in and is put to work by broad, open social networks. The creation of value and the accumulation of the common, then, both refer to an expansion of social productive powers....

One of the facts that make us rethink such concepts of political economy in social terms is that biopolitical production is not constrained by the logic of scarcity. It has the unique characteristic that it does not destroy or diminish the raw materials from which it produces wealth. Biopolitical production puts bios to work without consuming it. Furthermore its product is not exclusive. When I share an idea or image with you, my capacity to think with it is not lessened; on the contrary, our exchange of ideas

43 Heylighen, “Stigmergy,” pp. 29-30.

44 Michel Bauwens, “Interview on Peer to Peer Politics with Cosma Orsi,” *P2P Foundation Blog*, April 10, 2008 <<http://blog.p2pfoundation.net/interview-on-peer-to-peer-politics-with-cosma-orsi/2008/04/10>>.

45 Benkler, pp. 66-67.

and images increases my capacities. And the production of affects, circuits of communication, and modes of cooperation are immediately social and shared.⁴⁶

The synergy produced by the sharing of knowledge by the network is—in both senses of the word—a property of the network.

This has had revolutionary implications for the balance of power between networks and hierarchies, and almost unimaginably empowered individuals and small groups against large organizations.

In a hierarchy, all communications between members or between local nodes must pass through a limited number of central nodes. The only communications which are allowed to pass from one member or local node to another are those which meet the standards for distribution of those who control the central nodes. Only a few nodes within a hierarchy have the power to transmit; hence the use of the phrase “one-to-many” to describe its topology. The version of local news that appears in the local newspaper under the byline of a local journalist may be far superior in relevant detail and analysis, but it is the wire service version—even if far inferior in quality—which appears in local newspapers all around the world. It is only the communications approved by the Party Secretariat that are heard by all local cells of a party.⁴⁷

But in a distributed network, every node has the power to transmit, and any two nodes can communicate directly with each other without passing through a central node or obtaining the approval of whoever controls that node. A network is “plurarchical,” in de Ugarte's terminology, rather than democratic. Instead of the individual members simply selecting who controls the central nodes, “[s]omeone makes a proposal and everyone who wishes to join in can do so. The range of the action in question will depend on the degree to which the proposal is accepted. This system is called a pluriarchy....” Democracy is a “scarcity system” in which decision-making power is rivalrous: “the collective must face an either/or choice, between one filter and another, between one representative and another.” In a distributed network, on the other hand, decision-making power is non-rivalrous. Each individual's decision affects only herself, and does not impede the ability of others to do likewise. “Even if the majority not only disagreed with a proposal, but also acted against it, it wouldn't be able to prevent the proposal from being carried out.”⁴⁸ “[I]n the blogosphere,” de Ugarte writes elsewhere,

a space where the social cost of an extra post is zero, any blogger's publishing his or her information does not decrease anyone else's publication possibilities. The marginal cost is zero. The need to collectively decide what is published and what is not simply disappears. As opposed to scarcity logic, which generates the need for democratic decision, abundant logic opens the door to pluriarchy.

In such a universe, every collective or hierarchical decision on what to publish or not can only be conceived as an artificial generation of scarcity, a decrease in diversity, and an impoverishment for all.⁴⁹

Hardt and Negri describe the form of organization they call the “multitude”—as opposed to the monolithic “people,” the atomized, undifferentiated “masses,” and the homogenous “working class”—in terms that sound very much like stigmergy.

The people has traditionally been a unitary conception. The population, of course, is characterized by all kinds of differences, but the people reduces that diversity to a unity and makes of the population a single identity: “the people” is one. The multitude, in contrast, is many. The multitude is composed of innumerable internal differences that can never be reduced to a unity or a single identity—different cultures, races, ethnicities, genders, and sexual orientations; different forms of labor; different ways of living; different views of the world; and different desires. The multitude is a multiplicity of all these sin-

46 Michael Hardt and Antonio Negri, *Commonwealth*, pp. 283-284.

47 De Ugarte, *The Power of Networks*, p. 38.

48 *Ibid.*, pp. 39-40.

49 De Ugarte, *Phyles: Economic Democracy in the Network Century* (n.d.), pp. 18-19 <<http://deugarte.com/gomi/phyles.pdf>>.

gular differences. The masses are also contrasted with the people because they too cannot be reduced to a unity or an identity. The masses certainly are composed of all types and sorts, but really one should not say that different social subjects make up the masses. The essence of the masses is indifference: all differences are submerged and drowned in the masses. All the colors of the population fade to gray... In the multitude, social differences remain different. The multitude is many-colored, like Joseph's magical coat. Thus the challenge posed by the concept of multitude is for a social multiplicity to manage to communicate and act in common while remaining internally different.

Finally, we should also distinguish the multitude from the working class.... The multitude... is an open, inclusive concept. It tries to capture the importance of the recent shifts in the global economy: on the one hand, the industrial working class no longer plays a hegemonic role in the global economy...; and on the other hand, production today has to be conceived not merely in economic terms but more generally as social production—not only the production of material goods but also the production of communications, relationships, and forms of life. The multitude is thus composed potentially of all the diverse figures of social production.... [A] distributed network such as the Internet is a good initial image or model for the multitude because, first, the various nodes remain different but are all connected in the Web, and, second, the external boundaries of the network are open such that new nodes and new relationships can always be added.⁵⁰

The multitude, unlike the people, in traditional political philosophy cannot rule as a sovereign power because it “is composed of a set of *singularities*... whose differences cannot be reduced to sameness.” Yet “although it remains multiple, it is not fragmented, anarchical [word choice?], or incoherent.”⁵¹

Their description of the “common,” or background against which the multitude cooperates, is quite similar to Mark Elliot's conception of the stigmergic medium against which individuals coordinate their actions via markers.

Insofar as the multitude is neither an identity (like the people) nor uniform (like the masses), the internal differences of the multitude must discover *the common* that allows them to communicate and act together. The common we share, in fact, is not so much discovered as it is produced.... Our communication, collaboration and cooperation are not only based on the common, but they in turn produce the common in an expanding spiral relationship. This production of the common tends today to be central to every form of social production, no matter how locally circumscribed, and it is, in fact, the primary characteristic of the new dominant forms of labor today. Labor itself, on other words, tends through the transformations of the economy to create and be embedded in cooperative and communicative networks. Anyone who works with information or knowledge... relies on the common knowledge passed down from others and in turn creates new common knowledge.⁵²

Indeed, in their description of the swarming activity of the multitude, they appeal explicitly to the behavior of stigmergically organized termite colonies.⁵³

Hardt and Negri also attribute an internal tendency toward democracy to the multitude, in terms much like what David Graeber calls “horizontalism.” The modern history of resistance movements displays a shift from “centralized forms of revolutionary dictatorship and command” to “network organizations that displace authority in collaborative relationships” (this was written after the rise of the Zapatistas and the Seattle movement, but before the Arab Spring or the Occupy movement). Not only do resistance movements aim at the creation of a democratic society, but also tend “to create internally, within the organizational structure, democratic relationships.”⁵⁴

50 Michael Hardt and Antonio Negri, *Multitude: War and Democracy in the Age of Empire* (New York: Penguin, 2004), xiv-xv.

51 *Ibid.* p. 99.

52 *Ibid.* xv.

53 *Ibid.* p. 91.

54 *Ibid.* xvi.

The advantages of stigmergic organization go beyond resilience. Jean Russell coined the term “thrivability” to describe systems that are more than merely resilient.

Thrivability transcends survival modes, sustainability, and resilience. Thrivability embraces flow as the sources of life and joy and meaning, adds to the flow and rides the waves, instead of trying to nullify the effects. Each layer includes and also transcends the previous layer, expanding both interconnections as well as expanding system awareness as each layer hits limits and discovers that more forces are at work than can be explained within their purview.

She illustrates the distinction by contrasting descriptions of resilient and thrivable systems. Rather than simply withstanding or recovering quickly from difficulties, the thrivable organization is characterized by an “unfolding pattern of life giving rise to life”; it will “develop vigorously,” “prosper” and “flourish.” It is “anti-fragile”: that is, it gets better, generates and transformed when disturbed.⁵⁵

For a while I struggled a bit trying to picture examples of what her distinction between resilience and thrivability would mean in concrete terms. Then it hit me: stigmergic organizations are both resilient (because of distributed infrastructure and redundant pathways between nodes) *and* thrivable.

A stigmergic organization fits her description perfectly: “invites everyone to contribute their very best to making a world that not only works, it also produces joy, delight, and awe.” The reason is that it’s organized on a modular basis, and each discrete module of work is carried out by someone who volunteered to do it because it’s something they care about (often passionately) and they were empowered to do it without waiting for anyone else’s permission. So each task in a stigmergic organization is carried out by those most interested in it. Anyone who sees an opportunity for improvement, or has a eureka moment, can immediately jump in and get their hands dirty, and doesn’t have to work at it past the point where it ceases to be a joy for them.

To the extent that progress depends on the Shoulders of Giants Effect—people building on each other’s contributions—a stigmergic organization that facilitates collaboration, and does so without enforcing any barriers (like patents and copyrights) to making use of others’ ideas or creations, is the ideal embodiment of Russell’s idea of thrivability as promoting “growth on growth.”

Stigmergy is ideal for facilitating division of labor, with those best suited to a task selecting it for themselves. The Left—even those on the anarchist Left, who should know better—is plagued with the lionization of “activism” and guilt-tripping of anyone who lacks sufficient activist street cred. If your primary talent is writing or theory, according to this valuation, you’re a second-class Leftist. If you’re not “doing something”—which translates more or less into participating in demos—you’re a poser. But when viewed in light of the stigmergy paradigm, this view of things is just plain stupid. It makes far more sense for each person to do what she is best at, and let others make use of her contributions in whatever way is relevant to their own talents.

...if you have a large assembly of people who are forced to agree on every movement, including the mechanism for what constitutes such agreement, then you rarely achieve anything at all.

Therefore, as you build a swarm, it is imperative that everybody is empowered to act in the swarm just through what they believe will further its goals – but no one is allowed to empower themselves to restrict others, neither on their own nor through superior numbers.

This concept – that people are allowed, encouraged and expected to assume speaking and acting power for themselves in the swarm’s name, but never the kind of power that limits others’ right to do the same thing – is a hard thing to grasp for many....

55 Jean Russell, “Resilience Ain’t Enough,” *Thrivable.net*, February 6, 2013 <<http://thrivable.net/2013/02/resilience-aint-enough/>>.

As a result, somebody who believes the swarm should take a certain action to further its goals need only start doing it. If others agree that the action is beneficial, then they will join in on that course of action.

The key reasons the swarm should not be leaderless are two. You will notice that I refer to "its goals". Those come from you, the swarm's founder. If the swarm would be allowed to start discussing its purpose in life, then it would immediately lose its attracting power of new people – who, after all, feel attracted to the swarm in order to accomplish a specific goal, and not out of some general kind of sense of social cohesion.

The second reason is these very mechanisms, the swarm's culture of allowing people to act. These values will be key to the swarm's success, and those values are set and established by you as its founder. If the swarm starts discussing its methods of conflict resolution, meaning there is no longer any means to even agree when people will have come to an agreement, then the necessary activism for the end goal will screech to a halt.

Therefore, I believe that leaderless swarms are not capable of delivering a tangible change in the world at the end of the day. The scaffolding, the culture, and the goals of the swarm need to emanate from a founder. In a corporate setting, we would call this "mission and values".

But I also believe in competition between many overlapping swarms, so that activists can float in and out of organizations that best match the change they want to see in the world. One swarm fighting for a goal does not preclude more doing the same, but perhaps with a slightly different set of parameters....

You do the vision. The swarm does the talking.

Traditional marketing says that a message needs to stay constant to penetrate. My experience says that's total hogwash....

If somebody comes up to you and tells you a factual statement in a language that you identify as that of a group you dislike, you are very likely to discard that message as false, despite its actual truthness. In the same vein, if somebody that dresses, speaks, and acts in a manner consistent with your social standards tells you a factual statement, then you are likely to accept it as plausible and maybe examine it on its own merits later.

This sounds obvious. Yet, it has not been used in the marketing of ideas before swarms arrived on stage.

The recipe is ridiculously simple: communicate your vision to everybody, and let the thousands of activists translate your vision into words that fit their specific social context. Don't make a one-size-fits-all message that everybody has to learn. It will be a one-size-fits-none.

Let me give a tangible example. When I speak about the opportunities associated with the obsolescence of the copyright industry, I can do so in many different languages. If I were to speak about this before a liberal entrepreneur crowd, I would say something like this:

"There is tremendous opportunity in the cutting of this link from the value chain. The copyright industry intermediaries no longer add value to the end product or service, and so, on a functioning market, they are going to die by themselves. There is a problem here, as their monopoly prevents that. Therefore, we must assist in this cutoff, as removal of that overhead allows for growth of the overall market, future opportunities for the artist entrepreneurs, and for new jobs that take the place of the obsolete ones."

However, speaking to dark-red communist groups that celebrate the Red Army Faction as heroes, I would choose a different language:

"I think it is fantastic that the cultural workers have finally assumed control over their means of production, and that we finally have the ability to throw off the middlemen parasite capitalists who have been profiting for decades off of their hard labor. We should help our brothers and sisters to make this transition happen, and help them turn the captured middlemen profits into new jobs for our culture."

Factually, these two statements are completely identical. I am saying *the exact same thing*. But one wording would not work for the other group; you would get thrown out of the room and the curiosity in your swarm discarded for good.

Granted, these two settings are extreme contrasts to make a point. But even a subtle sign of not belonging can be enough to get your idea and vision discarded in a conversation.

This is why you need the activists – thousands of them – to translate your vision into many different social contexts as you have activists. Only then will you be able to electrify their friends with your vision, once that vision is clad in the language of that social context.

Don't think you can do this yourself for every setting. You don't master every nuance of language and social code. Nobody does. I may be able to switch languages rudimentarily from years of training in different settings, but I can't easily change appearance. If I arrive in a suit at a place for a presentation, and they turn out to be laid-back hippie types, then that's it. No word I say after that can change their perception of me.

It is also important, and imperative, that your activists not only are encouraged to translate your vision, but also to *interpret and apply it* to specific scenarios. In a political swarm, for example, that means they need to be able to translate general principles into specific policy on the fly, and express it in appropriate language for the context – always without asking permission. A three-activist rule can apply here, or you can empower everybody straight off the bat. When this starts to happen without any central planning and control, the swarm starts to really fly.

There will be people in the swarm who object to others' interpretations of the vision and general principles, of course. This brings us back to the distinction between empowerment of the activist self, versus the power to crack down on the work of others. The golden rule of the net springs to life: "If you see something you don't like, contribute with something you do like."⁵⁶

Vinay Gupta expressed the same idea in a couple of tweets:

Noble Saint Hexayurt does the heavy lifting, every hexayurt build makes four more likely.⁵⁷

I cannot save people, there are too many. I can give ideas and maybe some examples, but only an idea is big enough to help everyone.⁵⁸

Exactly. The primary bottleneck in today's world is not physical resources, but the transmission of knowledge. Why do something that I'm bad at, when the most cost-effective use of my time and talent is writing? Putting ideas together and propagating them is "doing something."

In sum, the transition to a society organized around stigmergic coordination through self-organized networks involves an exponential increase in agility, productivity and resilience.

The full power of such synergetic interaction supported by stigmergy is seen in complex, creative work environments, where different agents contribute different skills, experiences and perspectives. Here, the work done by one individual is enhanced by the work of others with complementary abilities in a way that the single individual never could have achieved. Wikipedia and communities developing open source software development are prime examples, having achieved results that could not even have been reached via hierarchical, command-and-control strategies of coordination. Smaller scale examples are people posting photos, ideas, artwork, or essays on their blog, Twitter feed, or Facebook page, and getting feedback from friends, followers, or strangers, which help them to further develop their insights, while inspiring these others to build further on their experiences. In such cases, the benefits that accrue to the "cooperators" are direct, concrete, and stimulating enough to motivate them to produce more of such "public traces" in their medium of choice (Wikipedia, Facebook, ...).

Thanks to the user-friendly electronic medium, the material and human cost of publishing such traces is nearly zero. This combination of strong motivation, minimal cost, and effective stigmergic coordination turns the medium into a powerful system for mobilizing joint action. The result is a rapidly expanding "collaborative commons," a virtual workspace for stigmergic (and more traditional) cooper-

56 Rick Falkvinge, "Selling Your Vision With a Swarm," *Falkvinge on Infopolicy*, February 18, 2012 <<http://falkvinge.net/2012/02/18/selling-your-vision-with-a-swarm/>>.

57 Vinay Gupta (@leashless), 6:26PM, August 5, 2012 <<https://twitter.com/leashless/statuses/232286550522724352>>.

58 Gupta, 6:31PM, August 5, 2012 <<https://twitter.com/leashless/statuses/232287814354624512>>.

ation that encompasses the planet. This world-wide stigmergic medium is presently developing into the equivalent of a global brain able to efficiently tackle the collective challenges of society.⁵⁹

* * *

Thus, stigmergy can be seen as a fundamental mechanism of self-organization: it allows global, coordinated activity to emerge out of local, independent actions. Like self-organization in general, stigmergy relies on feedback: action elicits action, via the intermediate of the trace. This feedback is typically positive, in that actions intensify and elaborate the trace, thus eliciting more intense and diverse further actions. The resulting virtuous cycle explains in part why stigmergic organization is so surprisingly effective, enabling the construction of complex structures—such as a termite hill, a network of trails, or a global encyclopedia—in a very short time, even when starting from scratch. When needed, feedback can also be negative: errors, disturbances or “overshoots” that make the trace deviate from its ideal shape will elicit actions that correct the deviation.⁶⁰

Our discussion, in the third section of Chapter Two, of the advantages of networks over hierarchies, can be read as a direct continuation of this section.

[Draft last modified October 5, 2014]

59 Heylighen, “Stigmergy” p. 31.

60 *Ibid.* p. 33.