*The follow is a script written in preparation for a talk given at an internal Philips Design Probes conference on Livable Cities. I was introduced as an artist and digital research, and asked to speak to the following questions:*

- *Can non-government networks step up when governments fail to provide basic services?*
- *How much will the digital layer and the physical layer melt into each other?*
- *What will be the influence of self-organized digital communities on cities?*

Hi!

So, I think "digital researcher" is the most politely euphemistic title someone has ever given to me. I'm a hacker. I look at systems and take them apart, see how they work, and put them back together so they work differently. A lot of the time, what I'm looking at is how they break. I've got a bit of a knack for breaking things, really.

In my day job, I'm a senior security associate with Stach & Liu, a small US computer security consultancy, and most of the systems I break are computer systems — large software applications and web sites run by banks, insurance agencies, and software companies. These days, I mostly work on a pretty high level, dealing with the intent and rules of the system rather than the actual code.

My father was an architect, however, and I've always found buildings, cities, and infrastructure fascinating. Over the past few years, I've been looking at them more and more like the systems I see at work. Cities have intents and rules and power structures too, and just like computer systems, they break.

---

Before I dig further into infrastructure and how it fails, I'd like to offer one caveat, and frame the rest of the talk a little:

Infrastructure, urbanism, and livability are very clearly political issues as much as technical ones. I'm speaking from a very specific political viewpoint throughout this talk, one which I recognize may be complicated for an entity like Philips at some levels, but bear with me.

Right now, 60 million people every year die of poverty, and every single one of these deaths is entirely preventable. We kill them, with our lifestyles, and with the concentrations of capital that we've helped to create; everyone in this room has the blood of 60 million people a year on their hands.

We are running out of oil, gas, and coal. How long we manage to go before we run into serious societal collapse due to energy availability is largely irrelevant. It may be ten years, it may be twenty, it may be forty, but it is extremely unlikely that the lifetime of the people in this room will not include the start of a catabolic collapse; if through a miracle, it doesn't happen to us, it will happen to our children. No society has ever survived the depletion of its fundamental resource stock intact, and it is gross hubris to believe that we will magically be different.

Assuming we do somehow shift entirely to a permanently renewable energy base, we still have to come to terms with the fact that a single-planet lifestyle, the lifestyle that if everyone shared it, would allow us to live in a sustainable fashion using the resources of only the single Earth we have, looks like Cuba or parts of Mexico.

I bring this up now, because it sets the stage for the rest of this discussion. Our task, and we have no choice but to accept it, is to learn how to be poor gracefully — to wash the blood of those 60 million off of our hands and to ensure that the cities we all live in are livable for everyone, rich and poor, present and future.

---

A city is a power structure. The structure may be flat or heavily hierarchical, it may have a single root or be multi-modal, it may be latent or active, but it is present. The same is true for any group of humans, from a family to a corporation to a nation state. Infrastructure is the body of a city, its corporeality.

We build infrastructure to support cities, and not the other way around. We build cities for a thousand human reasons, but for any given individual, access to that power structure is often high on the list. Choices about the infrastructure of a city are made by and mediated through its power structure, and the form of that power structure is reflected in the infrastructure the city creates. Organizations with tightly-coupled power structures create tightly-coupled infrastructure, and loosely-coupled organizations create loosely-coupled infrastructure.

Better communications infrastructure and more flexible abstractions of value allow the creation of larger power structures. Can you imagine trying to run a large multinational with a trading wing if you had to fly someone to a banking center every night with a half-ton of gold to cover nightly operating loans?

The story of the 19th and 20th centuries is one of the creation of larger and larger power structures, and more and more efficient ways to distribute that power — the will of the power structure and the resources, people, and products that those power structures interact with. The drive for those power structures to grow in size is a fundamental feature of the logic of capitalism. It's embedded in the money those power structures use to grow.

After two centuries of amazing growth, we are starting to see worrying cracks. Previously reliable systems are suddenly becoming unreliable in unexpected ways. Not everywhere, or all at once, of course, but nowhere is immune, either. In the US, transport systems are in horrible shape. In Europe, the Euro is suddenly showing severe structural problems. Everywhere, our hydrocarbon dependency is suddenly looming large. Why are things starting to fall apart?

---

There are a lot of factors contributing to it, and they're all pretty fundamental.

The pressure to grow is expressed in two directions: First, in the expansion of existing operations or the creation of new ones — more mills, more factories, more consultants, more product lines; and second, in an increase in efficiency in existing operations — higher margins, less waste, lower taxes. This efficiency is coercive; if an organization does not or cannot grow further and it does not or cannot become more efficient, it will be outcompeted and be unable to maintain its position in the power hierarchy, its share of resources. For cities, this is a core feature of globalization — the drive to become more attractive to business than other similar cities. This means lower taxes, less burdensome restrictions, and thus less capital for the city to embody its infrastructure with, fewer resources available for infrastructure that doesn't directly contribute to urban competitiveness.

Corporations perceive potential efficiency gains or new opportunities for capital expansion as surplus value in the world, which they have an impetus and even an onus to extract as much of as possible. If, for example, a maintenance budget can be cut, it will be, until a system fails and forces that budget's expansion.

A key feature of neoliberalism has been pushing this concept of surplus value into government. If a government provides a service and a private corporation could extract surplus value from the provision of that service, neoliberal logic dictates that the service should be divested because the exploitation of that additional surplus value is a moral good and a requirement of urban competitiveness. Of course, whenever possible, the government should continue to assume any risk if the service fails.

This has lead to a situation where we have systematically underinvested in all of our existing infrastructures, whenever and wherever possible. Money is much more likely to go into building new infrastructure which brings with it, in theory, new revenue streams. We can picture our collective infrastructure as suffering from a cumulative economic deficit.

In the long run, of course, this is ridiculous. The costs of infrastructural failures greatly outweigh the gains from shaving maintenance budgets. Furthermore, maintenance budgets are predictable, and catastrophic failures aren't. Much infrastructure has somewhat predictable life spans after which it needs replacement or renewal; a lot of the physical infrastructure we're currently coasting on was built after the war, and is seriously aging. Understanding this, however, requires a very long-term perspective, one that thinks in terms of centuries.

Businesses think in terms of quarters, or at best years. Supply chains operate on the factor of hours or days. Elected officials think at most a few years at a time, in electoral cycles. Even civil servants, once they've got the experience to understand long-term thinking, are unlikely to think beyond the end of their careers, say twenty years — still a fifth of where they'd need to be. As the pace of business and career mobility increases, all of this is speeding up further. There is a time-scale mismatch in the way we think about large systems that has led us to seriously underestimate the severity and likelihood of their failures.

Much of our infrastructure is seriously over capacity, and we cannot fix this. Our road systems, for instance, are routinely handling far more traffic than they were designed for. In the West, we've found that keeping up is financially impossible; traffic will increase to consume available capacity for roads, and building new heavy transit lines is prohibitively expensive.

There are other, longer term forces at play here as well, however.

The infrastructural wealth of the first world was built on a colonial legacy. We have been able to do as much as we have because over the past five hundred years, we have taken everything worth having from everyone else in the entire world and killed everyone who got in the way. We have used that massive wealth to build for ourselves a standard of living unimaginable by those we have stolen from. We pretended that the gain from this theft was regular income, instead of a one-time boon. It's not, and that money is mostly spent and not coming back. I don't know if you've noticed, but the West is increasingly going broke.

We are also facing an environmental deficit, where we have been operating in a completely unsustainable manner for a very long time. This has been both part and parcel of colonialism, but it's been going on for much longer. We can see in the historical record instance after instance of civilizations which fell when they exhausted the critical resources they depended on through over-use. Our infrastructure and our society have been designed around a set of assumptions which we are seeing will no longer hold true. Fixing this means that we cannot only repair and patch our existing infrastructure and attempt to keep up with capacity, but instead we need to look at radically reinventing it, at the same time that we are completely out of money with which to do this, cannot conceive of the scope of the problem or act on the time scale required to fix it, and are operating in the context of a system which favors individual gain over social good.

We have a bit of a problem here.

---

You can't solve a problem you don't understand, so let's take a look for a minute at how things fail.

---

We can split failures into two categories; simple and chained. Simple failures are pretty easy to understand; a key structural member in a bridge breaks and the whole thing falls into a river. A few hundred people may die, but in the grand scheme of things, it's not a big deal unless you're on the bridge. Except, of course, when that bridge is also the link between your island and the mainland, and now your shipping networks are cut off and you're in immediate danger of starving to death. This is a chained failure, and the hyper-connected nature of modern infrastructure means we see more and more of these. Chained failures occur inside single infrastructural systems, as well as between them — large scale electrical blackouts are an obvious chained failure.

Connectivity in an infrastructural network is a form of optimization. A system which is highly efficient under normal conditions will likely be very tightly coupled. In the electrical grid, connectivity allows utilities to sell power between generation markets, centralize generation, and compete on price. Spare capacity or resources sitting around a system waiting to be deployed represent inefficiency. Unfortunately, they also represent lead-time to fix the system when a link breaks and the ability to recover from single-point failures. A network which has a lot of lead-time built into it isn't necessarily more robust against a given failure than one without the lead-time — the system will still fail in exactly the same way in the end, but that reduction in efficiency does give a practical protection against the effects of that kind of fragility.

As the complexity of our society has increased, so have the complexity of our infrastructure and the length of our supply chains. The US Department of Defense is the single largest infrastructural organization in the world, and it provides a lot of great examples here. The DoD pays $.50US per gallon for diesel fuel. However, it costs $400 to get that gallon of fuel to a battle tank in Afghanistan. The cost per-unit of the infrastructure exceeds the cost of the basic commodity by a factor of 800! Half of all DoD personnel and a third of their budget are spent on logistics, and 70% of the tonnage they move during deployment is fuel.

Modern infrastructural systems have become very heavily optimized, and can thus said to be very fragile. When failures happen on a much longer timeline than efficiency gains, centralized, efficient, cheap, and brittle systems look like a good deal. In fact, in our banking systems, we've gone one step worse — we've magnified the negative effects of failure, via leverage, to provide marginally more short term gain. In exchange for a decade or two of year on year 12% gains, we've racked up enough exposure to failure such that when the crash finally hit, it wiped out all profits for the entire history of the banking industry.

Chained failures tend to have an accelerating effect. A single trigger will hit, possibly entirely randomly, and the system will attempt to adapt, burning through what store resources have. When accumulated stress triggers the next failure, much of the system's resources will already have been expended trying to fix the first failure, making the second one move that much more quickly. Resources gone, the system won't be able to pad the effects of the second failure, bringing pressure on the next point to failure levels much more quickly.

When large organizations design systems, one of the key things they attempt to maintain is a unified conceptual map. Even if a system is built and operated in a distributed fashion, some group, somewhere, will understand how the entire system fits together and use that understanding to optimize both system design and day to day functioning; this is the core of cybernetic management theory. The more centralized the organization, the stronger the conceptual map will be. Strong central maps promote efficiency under normal order, and orderly response to anticipated failures, but they make the system blinder to unanticipated failures. Someone approaching the map with the right mindset can sometimes use it to find those failures, but for a design or operations team, the fact that the world has sharply defined edges discourages them from looking at those edges too closely.

Diversity is another frequent victim of large organizations. Standardizing on a single design reduces costs when a sufficient degree of coordination exists to drive the process, either inside a single organization or across a market. We've seen this come up again and again in agriculture, from the Irish potato famine to the current problems with Tropical Race 4 and bananas, but it's a problem in all types of systems, too. In a diverse environment, elements of a system may be failing all the time, but they rarely fail all at once. Furthermore, because the system is used to failures, it has a well-rehearsed response when one occurs. Frequent small failures can actually improve the health of the system overall, as it can learn and adapt from them.

---

Understanding how a system will fail is hard, because we're not always aware where the dependencies are.

Let's look at this from a very immediate perspective.

---

You are sitting here, right now, in this room. In the event of a serious structural collapse — say, the collapse of the Euro, there are about six ways that you can die. If you get too hot, too cold, too hungry, too thirsty, too sick, or too injured, you're dead. You can't fix any of these conditions on your own. Under the status quo, the systems that can fix any of these things fail or succeed at a much larger scale than any individual. We can divide the world into about four levels of scale

— the individual, the group (anything between a family and a village), the organization (a large company or a region), and the nation-state.

[Image of Six Ways to Die map]

---

Let's take a look at just one condition, hunger, as an example.  In order to eat, you need food.  To get food, unless you're a self-sufficient farmer, you need a working system of exchange, which in the West means a functioning currency market, which operates at the state level.  If you've got that, you need to get food from a farmer to you.  This requires a road or rail network, which requires fuel, which requires fuel markets and refineries.  If you are a self-sufficient farmer, you require rule of law sufficient to guarantee you land security, which puts you back up at the organization level.

[Image of Six Ways to Die Hunger map]

---

This set of concepts is called a Simple Critical Infrastructure Map.  It tries to figure out what an actor at each level needs to survive in the short term.  Groups, organizations, and nation states all have their own immediate needs to stay alive.  None of these are sufficient in the long-run, but they are an effective triage to buy time.  Most of the work we do in our day to day lives isn't aimed at basic survival, but if we try to analyze the entire package at once, we get mired in all the complexity of the modern world.  With SCIM, we get a simple roadmap for taking care of immediate needs.  There's a lot more to SCIM, once you start looking at the interconnections between the levels, and I'm happy to pass on more in-depth resources later.

Every single one of the requirements identified by SCIM needs a strategy that ensures resilience.  The closer you get to the individual, the more immediate the needs are.  The smaller the scale that an individual's needs can be satisfied at, the more that individual can do to work to get that need satisfied.  I'm not going to go through the entire structure right now, because what matters is that we understand that there are specific critical paths where resilience must be maintained.

[Image of larger SCIM map]

---

Now that we understand where our existential vulnerabilities are, we need a tool that will let us identify the failure points in the systems which keep us alive.  Ideally, we'd like a tool which gives us actionable information about mediations for those failure points, or that at least lets us compare different solutions for the same need.

Conveniently, there's an analysis tool that lets us do exactly that. Most of my personal research over the past seven years has been into threat modeling. At the most general level, it's a methodology that lets us understand how a system can fail. It lets us model the different intentions of multiple actors, what constitutes a failure of the system for each of them, and how they'd like the system to respond to failure. Having understood how the system is intended to work, we can explore what happens at each step if something goes wrong.

This lets us see cases that are likely to cause cascading failures, even across apparently unrelated systems (within whatever scope we define for our model). It gives us some insight into places where we can buffer against possible failures for a system that doesn't meet our objectives, and gives us a way to capture the knowledge of what prevents the system from failing.

This model is adaptable to a wide variety of situations and systems — in addition to the more traditional software systems that I work on professionally, I've done work toward threat models for things as diverse as the Icelandic national constitution (adjunct to the now-derailed constitutional assembly), and the combined software and human process system for emergency information sharing, Ushahidi.

Threat modeling can't stand alone for infrastructural analysis, but it doesn't need to — there's a plethora of existing tools for analyzing the efficiency of supply chains and systems, computing actual resource needs, etc., but it gives us a unique view of failure and of the power structures embedded in a system.

---

What threat modeling is best at is getting at this ill-defined thing I'm calling the "shape" of a failure (it's hard to talk about a lot of this, because we don't have a vernacular for large systems failure). A system like the electrical grid has very particular feel to its large-scale failures, providing the classic example of progressive overloading of different links. A supply chain, on the other hand, has a different feel as a small disruption creates larger and larger delays moving down the line, causing other branches of the chain to pack up. The financial crisis mixed these two, while adding in a whole new layer of misinformation, secrecy, and distrust causing the perceived stability of different parts of the system to fluctuate wildly.

Used carefully, threat modeling may let us design systems so the shape of their failures will be readily managed at a human scale. This question of scale is absolutely critical to building resilient systems, especially in the face of large shifts in available resources. Large systems require steady resource streams to maintain functionality. When systems can only be repaired at the organization or nation-state level, a non-functional organization spells disaster for all the individuals who rely on it. When a system fails at the individual level and can be repaired at that level, not only does the failure affect far fewer people, but those people can fall back on social resources which are weak or non-existent above the small group level. Our humanity is the ultimate resource for resilience.

However, just as infrastructure mirrors the socioeconomic power structures of the societies which build it, infrastructure is a tool of social enforcement. A decentralized system that mirrors a weak power structure or a system which is resilient at the human level is a poor enforcement tool, and this fact means that if we want to fix systems and enable individuals and small groups, we have to fight not only the institutional tendencies of corporations and governments, but their direct interests as well.

Systems will fail, however, and when they do, resilience doesn't just give us technical options and a better chance at survival, it gives us social options other than paying more and more obeisance to the powers that be in hopes that they'll turn the lights back on and give us food. Resilience can make us more equal partners in what happens, good and bad; fragility makes us starving peasants.

---

Now that we've spent some time looking at what happens when classical systems fail, let's look at how modern digital or "soft" infrastructure can be different.

---

Just as there's a huge variety of hard infrastructure, soft infrastructure takes many forms, which cover a spectrum of territory from the edges of hard infrastructure on one side to the social sphere on the other. Obvious examples of things closely connected to hard infrastructure are things like smart electrical grid systems, where electric meters can communicate with control systems to give precise information on real-time power demand, and can communicate with appliances to trigger automated load-shedding during periods of peak power use. On the other end, we see things like Yelp or Foursquare that change how people interact with others in the context of the city, and change how they find available resources. Somewhere in the middle, we see a huge range of monitoring systems, from the apparatus of the security state through things like air quality monitors.

All infrastructure exhibits patterns of behavior, processes, and procedures that are part and parcel of its operation. With soft infrastructure, that behavior is moved from the community that operates the infrastructure into the infrastructure itself. Soft infrastructure embodies behavior. In some cases, not only rote procedure is internalized, but even adaptive behavior, such that the system may change its own structure or operation over time without outside intervention. In order for this to happen, behavior and memory is formalized in soft infrastructure. In some cases, this can be very good — a software system can be unbiased and radically egalitarian, but it can just as easily be bad, with rules applied without mercy or appeal. Like any hard infrastructure, a soft system reflects the power structure which defined it, but here it acts to enforce that power structure actively and not passively. On the other hand, as it is still infrastructure, it has the advantage over other forms of active power enforcement of omnipresence, of being in the background of all things at all times.

---

In order to understand how these systems do and don't work, we need to understand networks.

There are a few properties of networks which are especially interesting from our perspective. First and foremost, networks are decentralized structures. It's not always obvious from the outside how different this is, and how core it is to networks as conceptual structures. On a properly designed network, no node is any more important than any other node. This obviously runs very counter to traditional organizational structures, and represents the most disruptive feature of network-oriented systems. One key property of a properly decentralized infrastructure is separability, wherein any subset of the system can continue to function divorced from the rest of the system.

Networks, and the Internet in particular, are a culture as much as a technical system. The culture mirrors this same decentralized, anti-hierarchical nature. One of the things that I hear frequently from institutions trying to grapple with the changes that the Internet is bringing to their world is "How can we use the Internet to enable our institution." The Internet does not enable institutions. The Internet, by and large, has absolutely no use for institutions not of its making, which do not reflect its nature. It gleefully vivisects institutions, takes the few things it needs from their guts, and recreates them in its image. The old boundaries that those institutions defended are generally casualties in this process. I don't mean to claim a value judgment in this; it simply is.

If you need a model to understand what an organization that's native to the Internet looks like, don't go to CIA-astroturfed Facebook groups, even if they're having a real impact on the world, or even Twitter, with its culture of modern celebrity. Look at Anonymous, the leaderless, aimless, amorphous vigilante hacker group. Not all Internet organizations will be equally abrasive of course, but there is a core nature there, the lack of fixed organization or indeed any fixed structure, membership determined only through participation, as a group largely unknowable except through its actions, transmitted as much by infection with a cultural meme as by any more coherent growth, which very accurately reflects what the Internet looks like.

Size in networks is an interesting thing. Networks scale extremely well. For a piece of physical infrastructure, going from a thousand users to a million users involves a massive investment in expansion and redesign. For a well-designed network, it involves a few tweaks, and can literally be done overnight in many cases. On the other hand, because of decentralization, the removal of layers of hierarchy, and the ability to make decisions in an immediate, transparent, and often automated fashion allows network services to operate with much less overhead than non-networked versions.

Openness is another core feature of network culture, and it goes hand in hand with decentralization. If any node on the network is a true peer of any other node, then it stands to reason that no node should be privileged over any other node in its ability to interact with the network. In the software world, it has repeatedly been the case that a software function becomes truly networked, in a decentralized and open fashion, billion dollar industries turn into hundred million dollar or smaller industries. The growth of the Internet and the new opportunities created by increased access have largely hidden this reduction, but it is very real.

Obviously, merely moving to a truly networked model will not change the math on something like the power grid in the same way; power generation has mostly-fixed energy costs, and the existing power grid is already relatively decentralized. However, moving to a fully decentralized system enables things like micro-generation at whole new levels, something we've already started to see, and it does so without any loss of efficiency beyond that imposed by physical constraints. This is critical, because with decentralization, we can gain separability and therefore a decrease in the scale of a functional system. This directly impacts resilience, driving the down the scale at which infrastructural problems can be fixed. Other soft systems, especially ones aimed around ownership, can provide large boosts to overall efficiency — car sharing systems are an obvious example here. Soft systems can allow us to use hard systems in new ways without needing to change the hard system at all, by solving coordination problems.

---

Networks are a new kind of power structure, one where the worth of the network, which is not owned by any individual group, exceeds the worth of any of its components many-fold, and in which the value of those components is related to their participation in the network, not to anything intrinsic in their position or status. We are still just starting to understand how powerful network-structured organizations can be. For instance, between Pirate Bay and Wikileaks, we now have a rough understanding of how powerful a file server can be.

To say that this new power structure has been resisted and co-opted by the neoliberal world is a massive understatement. Indeed, the entire recent history of the Internet (say, the last ten years) has been about its neoliberalization — the industrial revolution of the Internet, if you will.

Just as corporations have acted to capture more and more value from hard infrastructure over the past few decades, they have also attempted to reshape the Internet as a tool for value capture, and are already trying to do the same thing with soft infrastructure.

To do this, they have a variety of tools at their disposal. The most fundamental is the imposition of hierarchies of control and inequality on an otherwise flat network. The very idea that there are clients and there are servers on the Internet is a tool to subjugate an otherwise equal network. There is no reason, for instance, why my phone shouldn't serve data to other phones around it, why it shouldn't connect directly to another phone in reach of its radio when that makes sense topologically. However, if I could do that, it would be much more difficult for the phone company to justify capturing an often usurious amount of revenue from the system.

We see the same thing in the way the domain name system currently works (something which was not part of the original Internet, but which instead represents a successful capture of value by a corporate monopoly), in the way SSL certificates are issued for secure communications, in digital rights management as applied to digital media, and in services like Twitter and Facebook. These systems are not only perversions of the network intended to replicate hierarchy on an otherwise flat system, they actively degrade the value of the network both directly and indirectly.

Centralized systems represent single points of failure and are an unacceptable risk. If you entrust the entirety of your cultural heritage to DRM-enabled media and the company that manages the DRM goes bankrupt and falls offline, you lose access to that cultural heritage. If you trust a centralized system to manage your namespace and you say things that that system doesn't like, you risk becoming inaccessible as you are censored. If you trust a company to provide security for your communications, you risk a compromise of those communications when that company betrays your trust in order to save money.

Beyond purely functional issues, however, by enabling value capture at the core of network functions, the network becomes less efficient. For instance, it is much harder for an inclusive open scientific movement like the DIY biology community to exist when they have to contend with Internet paywalls in front of important journals put in place by publishing companies trying to capture value from the scientific community.

Another serious form of value capture is the differential privacy problem, which happens because of the massive scale difference between an entity like Facebook and an individual user. The data stream of one individual is only of value to them, but that same data across hundreds of millions of users becomes very valuable. Even if users had access to the same data about the corporation, very few of them (namely, only those users who happen to control other corporate-scale entities) would be able to make much use of it. This differential value exerts a corrosive effect on the network, making it less free as large entities exploit their information streams, and by extension, other users. This problem can happen even in a purely peer to peer network, if large scale and capital differentials exist. It can be solved, but only by engineering privacy into the network in fundamental ways.

The Internet as a cultural entity, after being enticed by the simplicity of the cloud and the shiny web 2.0 toys dangled in front of it, is starting to realize these dangers and to work at fixing the problems. We can only hope, for the future freedom, reliability, and security of the Internet, that it's not too late.

Infrastructure that interacts with less purely-electronic systems moves on a slower time-scale than the Internet, even soft infrastructure, and it's much more subject to value-capture and the imposition of power hierarchies because of this. It is critical that the infrastructure of the future avoid this trap.

---

As the financial crisis is hitting governments, we're seeing a curious trend. In places, corporations are deciding that they can no longer profit from providing services, and so instead of pushing for that service to be privatized by the government, they're pushing for it to simply be gotten rid of altogether. The government can't be seen to simply be ceasing to offer services that the population is dependent on, so instead, we get the logic of devolution, of the "Big Society" movement in the UK. As it stands, of course, this is just an almost-transparent spin. Because government has been captured by the rich, we cannot have a meaningful tax structure, and so are left with austerity measures.

We know that network-organized groups can both function and create things of lasting value, so will they be able to take up this challenge? Only if resources and authority devolve with responsibility. Governments would love to devolve only responsibility, but still keep the door open to recapturing authority, on the off chance that a community actually does create a functional replacement for a service. However, in doing so they will attempt to impose the same kind of power hierarchy which proves problematic in so many existing places and subject the newly created services to the same logics of regulatory and value capture that they work under. If it is to have long-term impact, any infrastructure created in these cases must actively resist, in its fundamental design, co-option and subversion.

The goal with devolution is to devolve responsibility without resources. Obviously, this is not functional. Infrastructural projects evolved to replace government services will need to reclaim resources from the corporate entities that have captured them. For instance, a citizen-run pharmacy replacing government pharmacy benefits might consider working with the global DIY biology movement to break international drug patents and manufacture their own drugs in a distributed fashion, destroying the value captured by the large pharmaceutical companies. By crowdsourcing the manufacturing and ensuring that the knowledge required for production is available to everyone, it becomes much harder for the government to reassert authority.

While a preferable situation might be to simply regulate the pharmaceutical industry's prices and recapture sufficient revenue run the pharmacy benefit system as before, in our example that option is no longer open. The new solution may not become what was first intended, but once the change is in process, it is no longer under central control, by intention. True devolution is unavoidably a revolutionary, or at least insurrectionary, act.

Because of the distributed nature of the solution, this solution has potentially interesting side benefits — instead of requiring a single centralized and very long distribution chain, the possibility exists for truly distributed production. In the event of a system failure, a fix can be performed at a much lower level. Because there is little or no value capture in the system, affinity and mutual aid enter the picture in new ways. An open source distributed production network is not beholding to the dictates of profit, so, for instance, simply giving someone the medicine they need isn't a problem. Because the system is (potentially) operating on a very small, personal scale, this kind of commons can be maintained in the traditional manner, via social links.

If such a system is actually going to replace existing infrastructure, some coordination will be required, but only as much hierarchy will be needed as actually required by the flattest restating of the problem. Open infrastructure doesn't (necessarily) need to act as a revenue channel or a vessel for the enforcement of power structures, so it can be much flatter and lower overhead.

As soft infrastructure extends into the real world, cheap sensors let us build APIs for physical objects. We're seeing plenty of this with spimes, self-aware, Internet-enabled things that talk to other things, but where it gets interesting is when those APIs become mature enough that we can enforce API contracts on physical objects. Returning to the example of a bottom-up grid, this might be a system where every household-node in the electrical system both consumes and generates some electricity, and its neighbors share information about how much power it consumes and about the frequency, voltage, amperage, and noise levels of the power it puts out, allowing the system as a whole to make statements about each of its members.

The combination of network-as-power structure and API contracts for physical objects let us conceive of vernacular infrastructure. Much like vernacular architecture, it emerges, bottom-up; it is owned in bits and pieces by everyone and controlled by no one, but the logic of the network makes efficiency, appropriate redundancy, resilience, and quality control possible. The network distributes knowledge about how to be part of it, and trust emerges from the network, not from any central authority.

If we're very lucky, we'll be able to decrease the footprint of soft infrastructure to the point where it will continue to be viable even as the hard, value-capturing parts of our infrastructure undergo a profound contraction. There may not be any other way to retain what matters to us in the world.

Eleanor Saitta / @dymaxion

2011.4.30, Brooklyn/Eindhoven