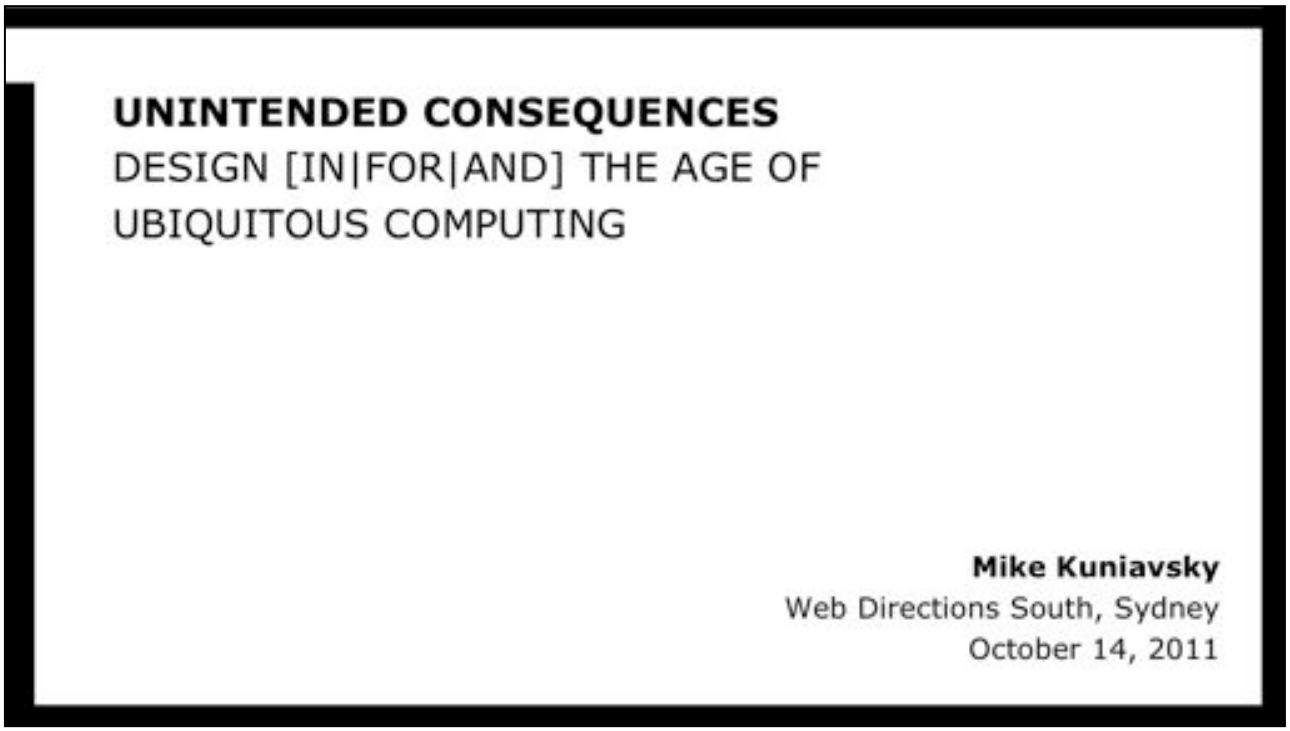

Good morning! Thank you very much for inviting me. I've heard great things about this event for years and it's an honor to be here. Today I'll be talking about ubiquitous computing and, very broadly speaking, design.

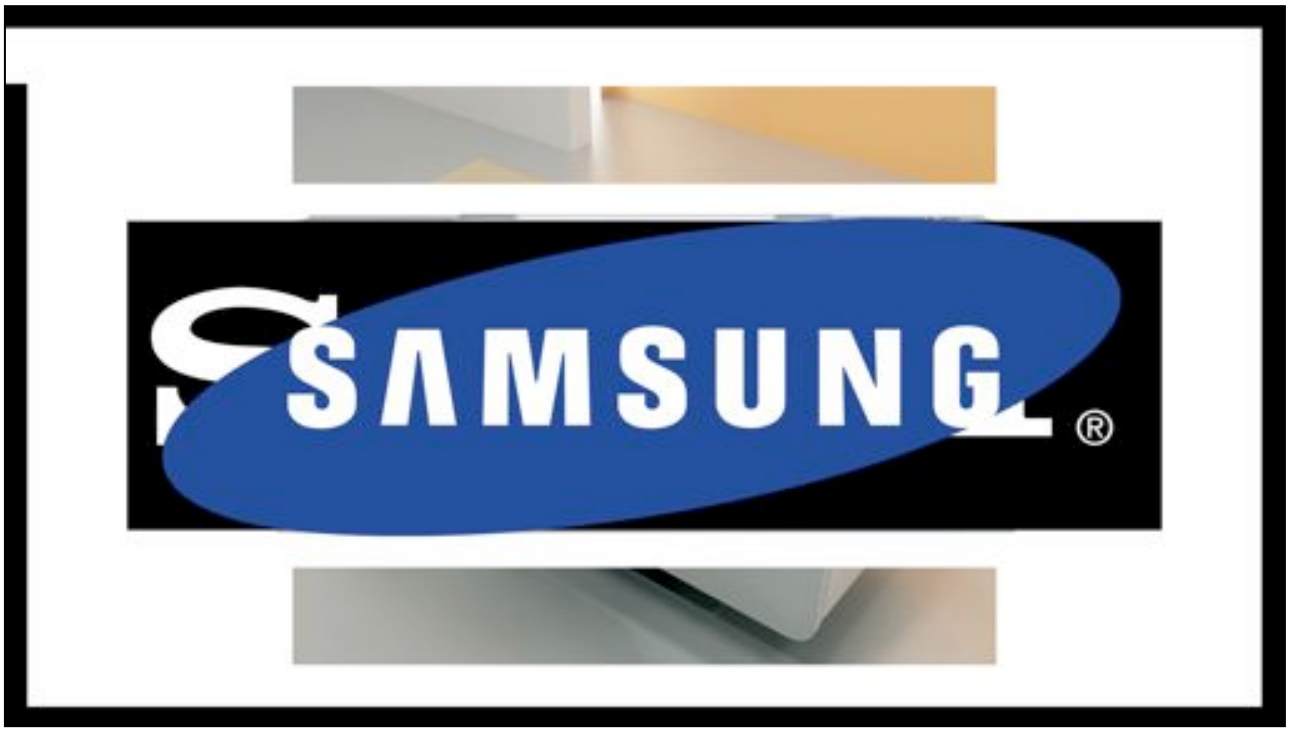

First, let me tell you a bit about myself. I'm a user experience designer. I was one of the first professional Web designers in 1993. I've worked on the design of hundreds of web sites and many digital consumer products. I also regularly work with companies to help them create more user centered design cultures so they can make better products themselves.

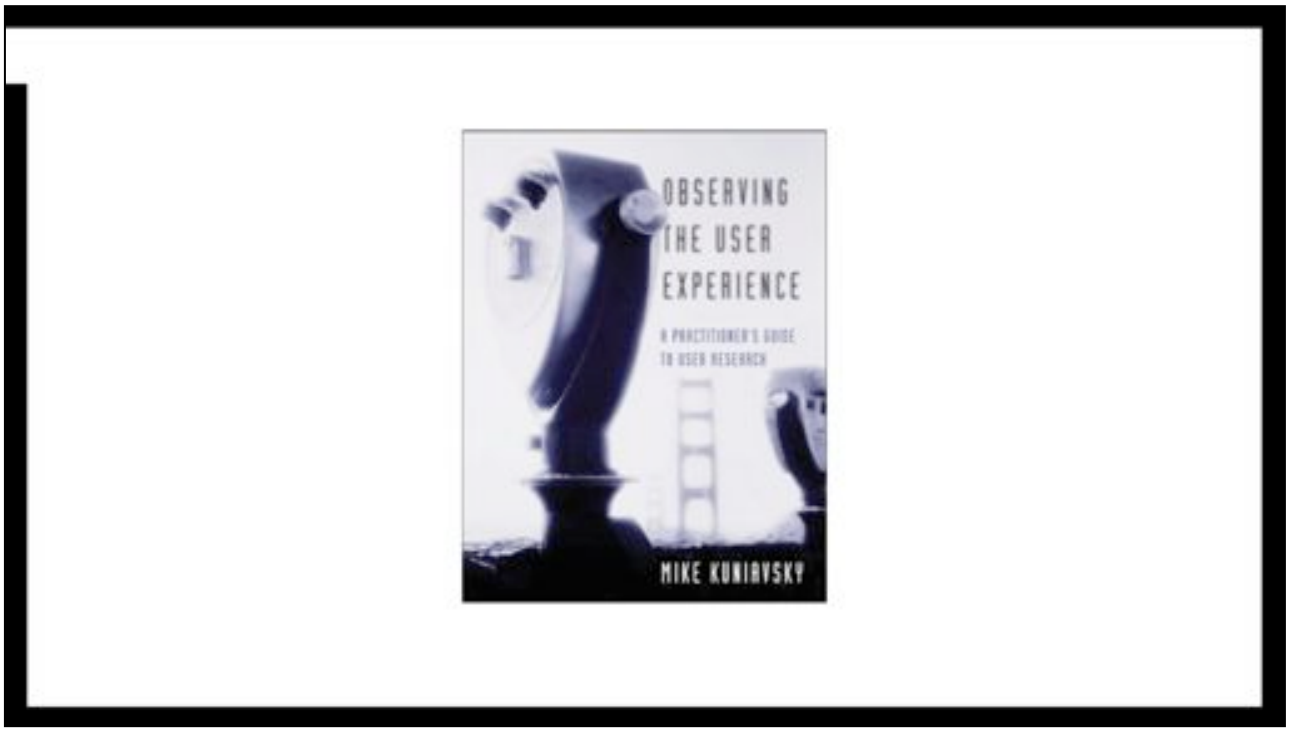

I sat out the first dotcom crash writing a book based on the work I had been doing. It's a cookbook of user research methods.

And 2001 I co-founded a design and consulting company called Adaptive Path.

…and three years later I left it, and I left the Web altogether, to found a company with Tod E. Kurt called ThingM in 2006.

We weren't sure what we were going to be but it's turned out that we're a micro-OEM. We design and manufactures a range of smart LEDs for architects, industrial designers and hackers.

This talk is based on my book on ubiquitous computing user experience design. It came out last September and it's called "Smart Things" and it's published by Morgan Kaufmann.

I want to start with a little history. I love the history of technology. This example comes from Harold Innes, who was a political economist and Marshall Mcluhan's mentor, wrote about technologies and empires. He has an interesting take on papyrus. According to him, it nearly brought down the Ancient Egyptian empire, and ended up changing it forewver. Before papyrus, writing in ancient Egypt was the process of slowly inscribing information permanently on immobile things like obelisks and tomb walls. Information moved slowly, formally. It was easily controlled and constrained.

When papyrus was invented, it seemed like a great idea for those in power. The pharaoh could administer his empire from a central location and wouldn't have to rely on messengers. Now he could send lots of precise instructions and scribes could write down complex ideas, such as those about geometry. But papyrus is not stone. It's easier to write on, orders of magnitude easier. So, people wrote more. A lot more. They were writing so much that they needed a less formal florid writing system, and more people learned to read and write. Suddenly, and by suddenly I mean over the course of hundreds of years, this meant that knowledge, and the control that comes with it, was no longer be centrally controlled. People started to get strange ideas. They started to ask why it was only the Pharoah who got to go to heaven. Scribes, the nerds of their era, were suddenly quite powerful. Surprisingly powerful. Dangerously powerful.

The Pharaoh—and I can't remember which dynasty this was, maybe the 19th?—decided that this was really endangering the stability of the Empire, which was under a lot of stress anyway. He needed to do something drastic. He made the all the Scribes report directly to him. They were elevated to the same level as priests and the position became hereditary and bureaucratic. No one else was allowed to write. Amazingly, this worked, and the literacy that was

The interesting thing is that the people who invented papyrus did not create it to threaten Egypt. Quite the contrary. And the scribes, they were just producing content. Moving symbols around. They were not intending to undermine their government.

No one involved intended to nearly topple Egypt with papyrus. There was nothing inherent in the technology that could have predicted this. No, it's that technology always, always has unintended consequences.

We who make technology have a strange perspective in its role in the world. We feel that because we make it, we understand it. We like to think we can predict where it will go and what it will do.

The problem is that our perspective is tiny and incremental. We usually miss the real deeply transformative change that happens outside our frame of reference. Often it's the people who create a technology that are the most surprised by its effects.

These are two small piece of Scott Weaver's toothpick sculpture of the Bay Area.

The whole thing looks roughly like this. It took him 30 years and a bazillion toothpicks.

As technologists, as human beings, really, we are great at seeing the details, but in many ways we're cognitively equipped to to see the whole. We're terrible at seeing emergent phenomena that come from the confluence of thousands of small things. Big social waves brought on by technology have to be nearly on top of us before we see them.

We're currently in the upslope to such a shift brought on by something familiar, something that we may think we have a handle on, but which is creating deep social shifts we couldn't have predicted.

I'm of course talking about Moore's Law, since that's where all conversations about the implication of digital technology start. When people talk about Moore's Law, it's often in the context of maximum processing power. But it's actually something different. It's actually a description of the cost of processing power. It's a model of how much more processing power we can fit into a single chip that's priced at a predictable pricing point this year than we could last year. This means that it's not just that processors are getting more powerful, it's that PROCESSING is getting cheaper.

For example, at the beginning of the Internet era we had the 486 as the state of the art and it cost $1500 in today's dollars. It's the processor that the Web was built for and with. Today, you can buy that same amount of processing power for 50 cents, and it uses only a fraction of the energy. That decrease in price is the same orders of magnitude drop as the increase in speed. This is not a coincidence, because both are the product of the same underlying technological changes.

What this means in practice is that embedding powerful information processing technology into anything is quickly approaching becoming free.

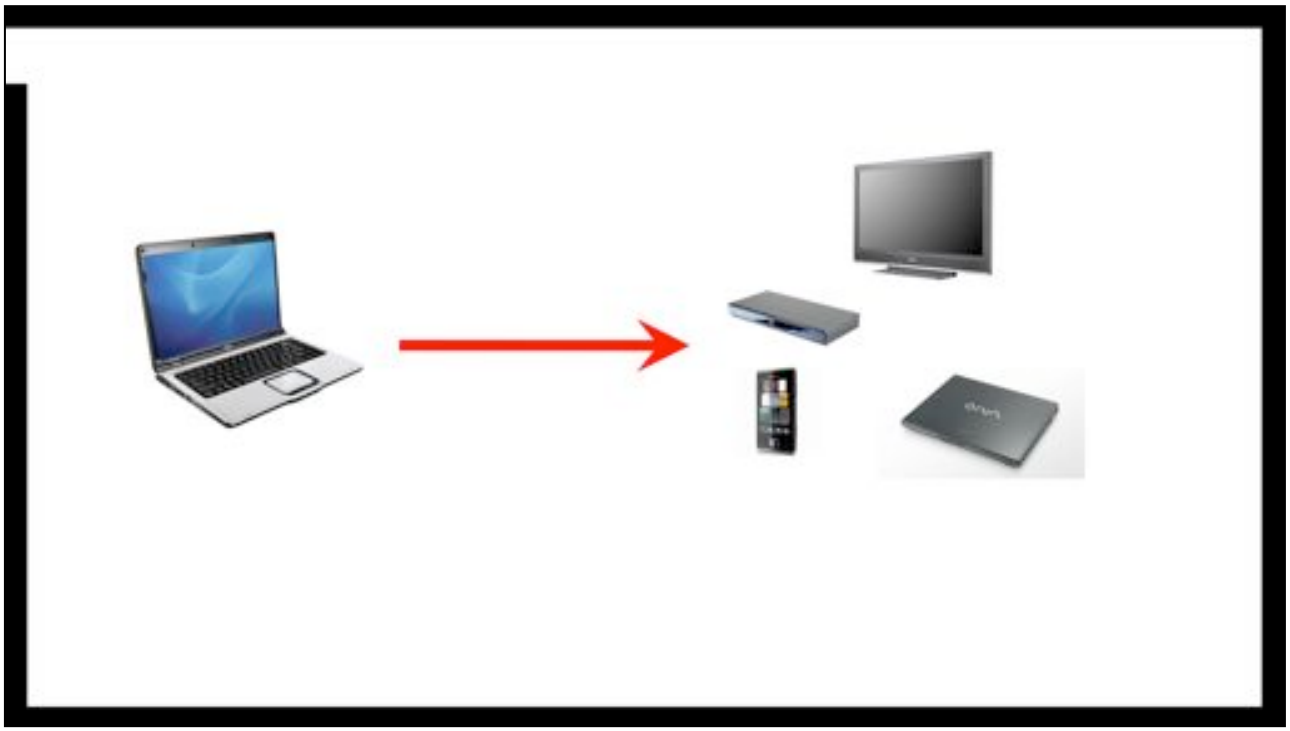

We see this most readily as a proliferation and a homogenization of digital devices because virtually any device can now do what every other device does. This is why we're seeing all of this churn in form factors, since the consumer electronics industry is trying to figure out how they can sell yet one more screen of a different size. Four years ago it was smart phones, three years ago it was all netbooks, two years ago it was tablets, now it's 7-inch tablets and connected TVs. They're all essentially the same device in different form factors.

That's fine, but it's the most primitive of the transitions that's happening.

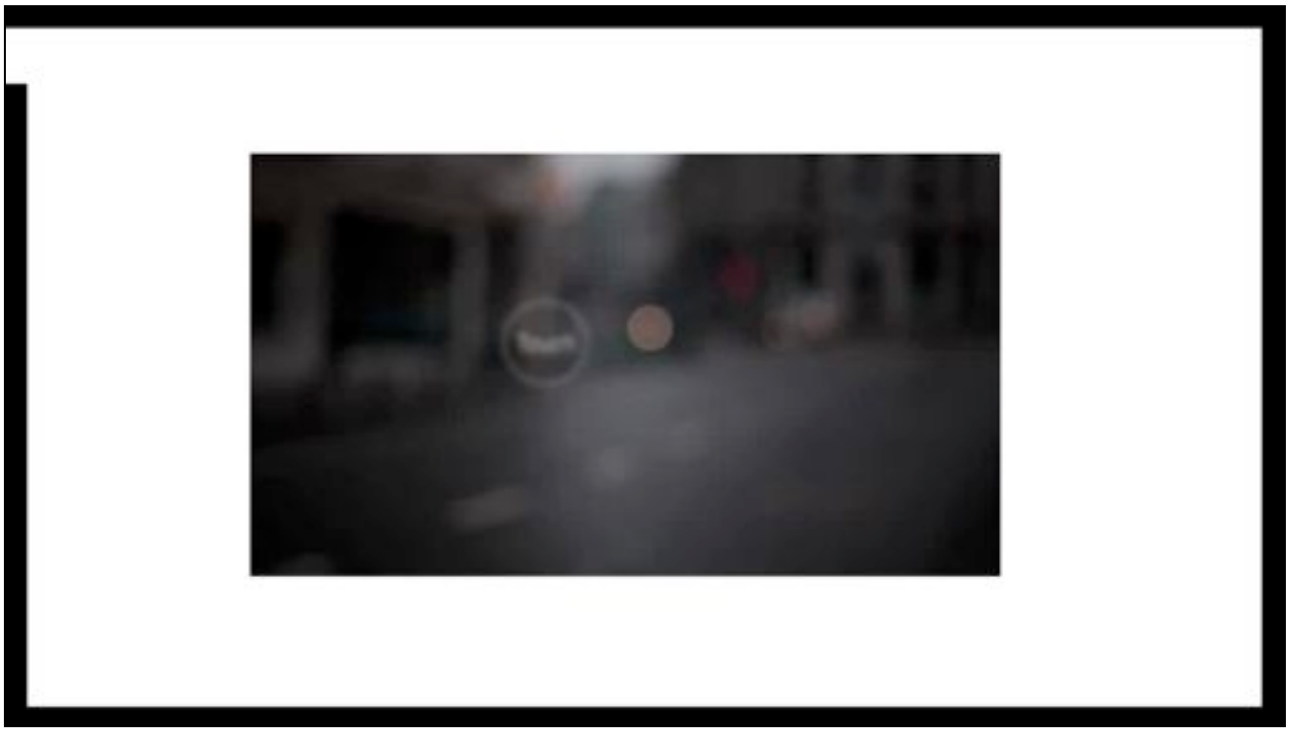

Simultaneously, the number of wireless networks in the world grew by several orders of magnitude.

This is a video by Timo Arnall that envisions how saturated our environment is with networks, and it's not even counting the mobile phone network, which covers just about everything. This means that virtually any device, anywhere can share data with the cloud at any time. People right now are excited about moving processing and data storage to the cloud and treating devices as terminals. That's certainly interesting, but it's also just the tip of the iceberg. That's like saying the steam engine is really great for pumping water out of mines. Yes, it's good at that, and also creating the industrial revolution.

It is thus no longer unthinkable to have an everyday object use an embedded processor to take a small piece of information—say the temperature, or the orientation of a device, or your meeting schedule—and autonomously act on it to help the device do its job better. Information processing is now part of the set of options we can practically consider when designing just about any object.

If you look at what happened when the price of writing fell, or when extracting aluminum became two orders of magnitude cheaper in the late 19the century, or when electric motors became significantly cheaper and smaller in the 1920s you see dramatic material and societal change. When something becomes cheap enough, when cost passes a certain tipping point, it quickly joins the toolkit of things we create our world with.

In other words, information has become a material to design with.

And with that, we have entered the world of ubiquitous computing, the world Mark Weiser roughly .

Because we have information as a design material, we no longer think it's crazy to have a processor that creates behavior in a toy, or for a bathroom scale to connect to a cloud service, or for shoes to collect telemetry.

This capability of everyday objects to make sophisticated autonomous decisions and acting using arbitrary information is new to the world and it is as deep an infrastructural change in our world as electrification, steam power, and mechanical printing. Maybe it's as big of a deal as bricks. Seriously, it's a huge change in how the world works, and we're just at the beginning of it.

Tickle me Elmo Extreme

Withings

adizero f50 miCoach

Today it's relatively simple to make a device sense the world with a great deal of precision.

There are thousands sensors that convert states of the world into electrical signals that can be manipulated as information. This also includes sensors that sense human intention. We call these "buttons", "levers", "knobs" and so on.

Our things can make physical changes in the world based on input. Devices made from the perspective of treating information as a design material can autonomously affect the world in a way that no previous material was capable of.

Information can be used to store knowledge about the state of the world and act on it later. This could be just a single piece of data.

Or it can encode very sophisticated knowledge about the world. This is a Blendtech programmable kitchen blender. With it you can program a specific sequence of blender power, speed and duration and associate that sequence with a button on the blender. it allows you to embed experience and knowledge about food processing into the tool, which can then produce that as a behavior, rather than requiring the operator to have that knowledge and develop the experience.

Why do this? Well, if you're Jamba Juice, which is a large US smoothie chain, your business depends on such programmable blenders so their staff don't have to be trained in the fine points of blending and their product is always consistent. Their profit margins depend on knowledge that's encoded into their blenders, knowledge that's accessed with a single button.

This is the control panel of Blend Tech's home blender. Blenders used to have buttons for different speeds. They described WHAT you were doing. Now, with embedded knowledge, it's about the desired end result. It's about WHY. The software handles the what.

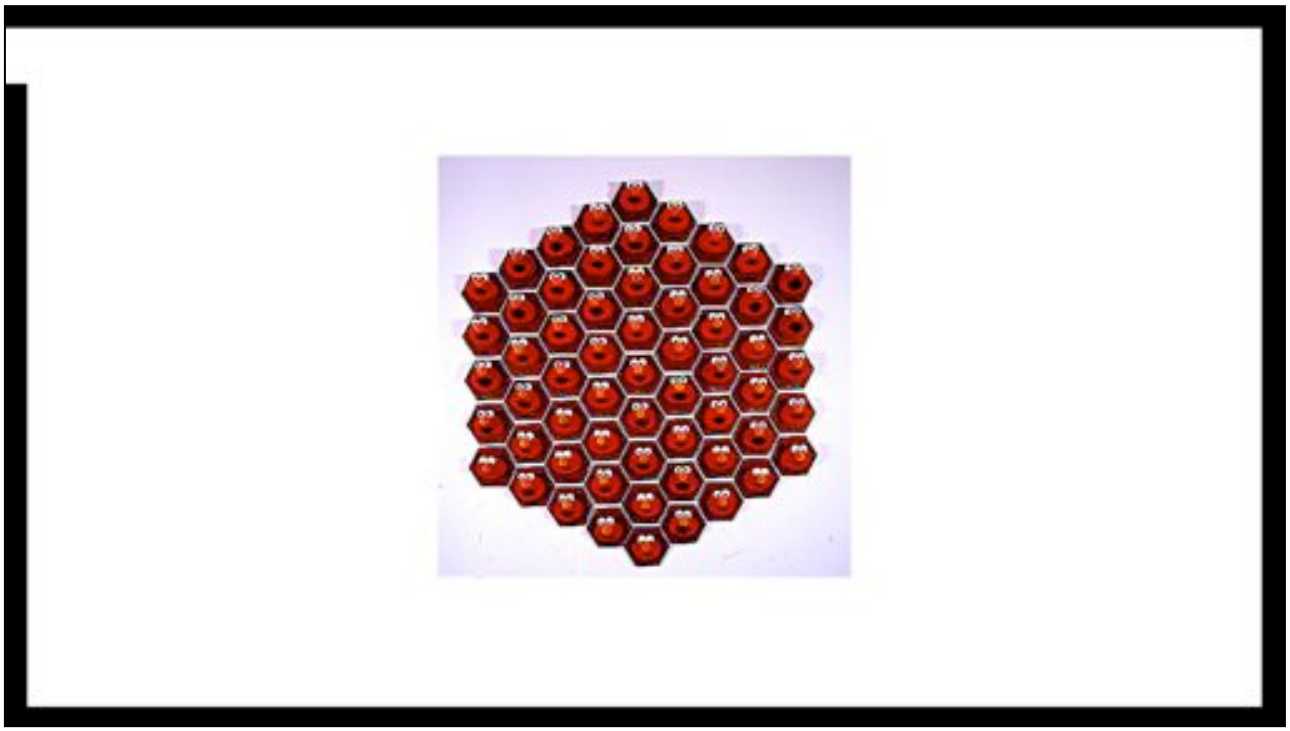

One of the most transformative qualities of information is that it can be duplicated exactly and transmitted flawlessly. This has already changed the music and video industry forever, as we know.

But it also means that device behavior can be replicated exactly. We've become acclimated to it, but--stepping back--the idea of near-exact replication in a world full of randomness and uncertainty is a pretty amazing thing, and is a core part of what makes working with information as a material so powerful.

Image: N-Trophy, 2000-2003, Kelly Heaton, Feldman Gallery: http://www.feldmangallery.com/pages/exhsolo/exhhea03.html

Finally, and most profoundly, things made with information do more than just react, they can have behavior.

Information enables behavior that's orders of magnitude more complex than possible with just mechanics, at a fraction of the cost. This is a modern small airplane avionics system. It consists of a bunch of small fairly standard computers running special software. It's a bit like a flight simulator that actually flies.

Found on: http://www.vansairforce.com/community/showthread.php?t=51435

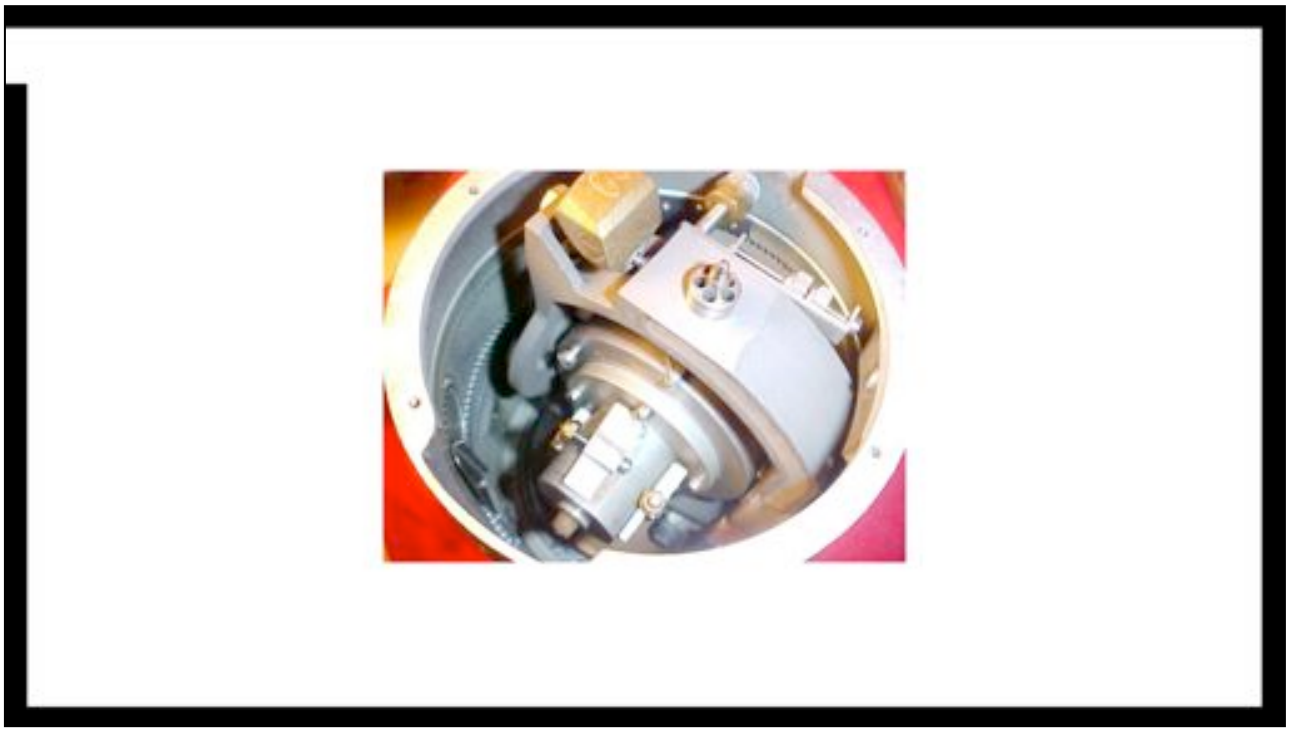

Compare that to a traditional gyroscopic autopilot, what it replaced. Every component is unique, it does very little, and to change its behavior you have to completely reengineer it.

When you make something with information, you enable that thing to exhibit behaviors that are vastly more sophisticated than what was possible with any previous material.

That is the wave that's basically on top of us.

So what can we as designers do in this situation?

Well, we're possibly the luckiest ones.

For the last 20 years we've been building a digital representation of the world on the Internet. We call the the Web, and if you look at it as a unit, it's a rough and unorganized, but fairly complete model of most things in the world and how they interact.

Until now, however, it was disjoint from the thing that it was modeling. We left it up to people to make the connection between this map of the world with the world itself. We had to resort to things like stickers to tell people in the real world that a given object, or location, had an information shadow in the cloud.

But that's quickly changing.

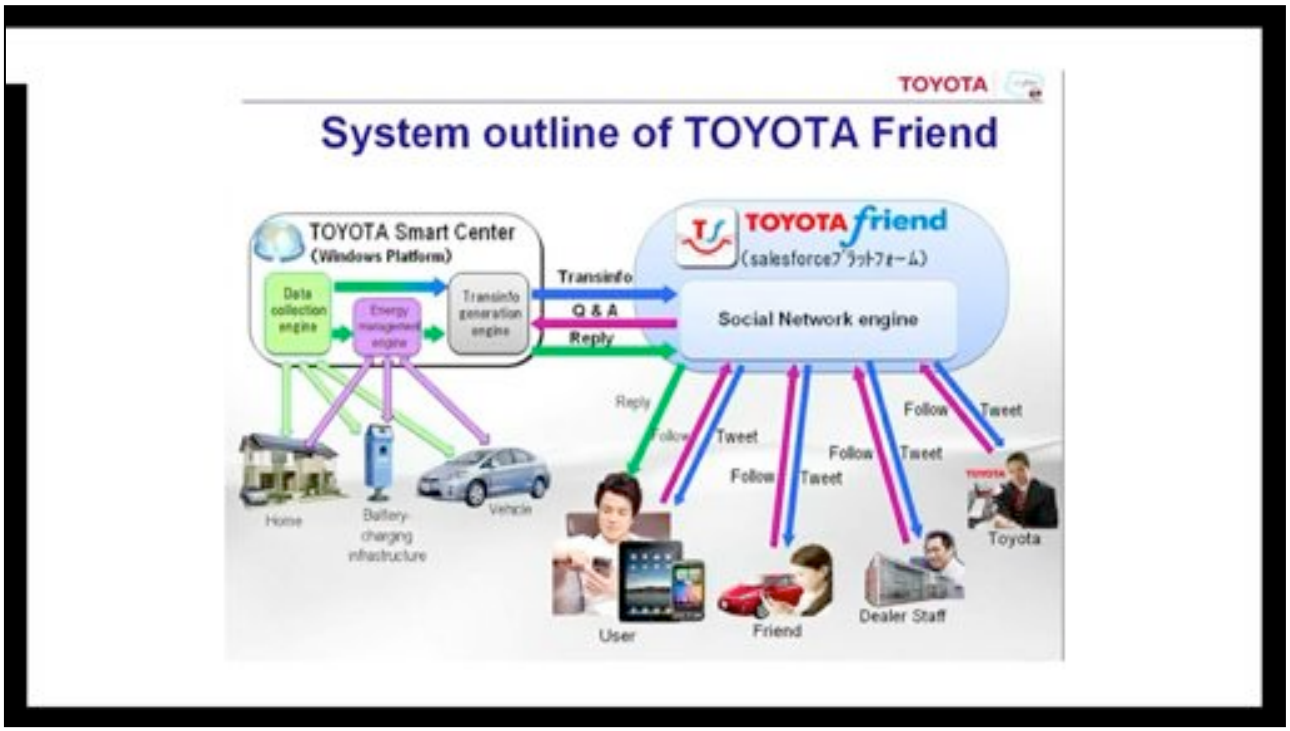

Here's Toyota and Salesforce's plan for having your car continuously embedded in both Toyota corporate's network and your social network. The factory can update the car firmware remotely and the car can text you when it's done charging. The information shadow of the object, it's representation in the cloud, and the object have been glued together.

For Web designers this is great news. As the model of the world and the world merge, as the map and the territory become increasingly intertwined, who knows the most about the map? It's us. We're been swimming in it longer than anyone else. And as things are increasingly made using information as a material thanks to the inclusion of cheap processing and networking, we're the ones who know how to design for it.

Colliding galaxies, NASA

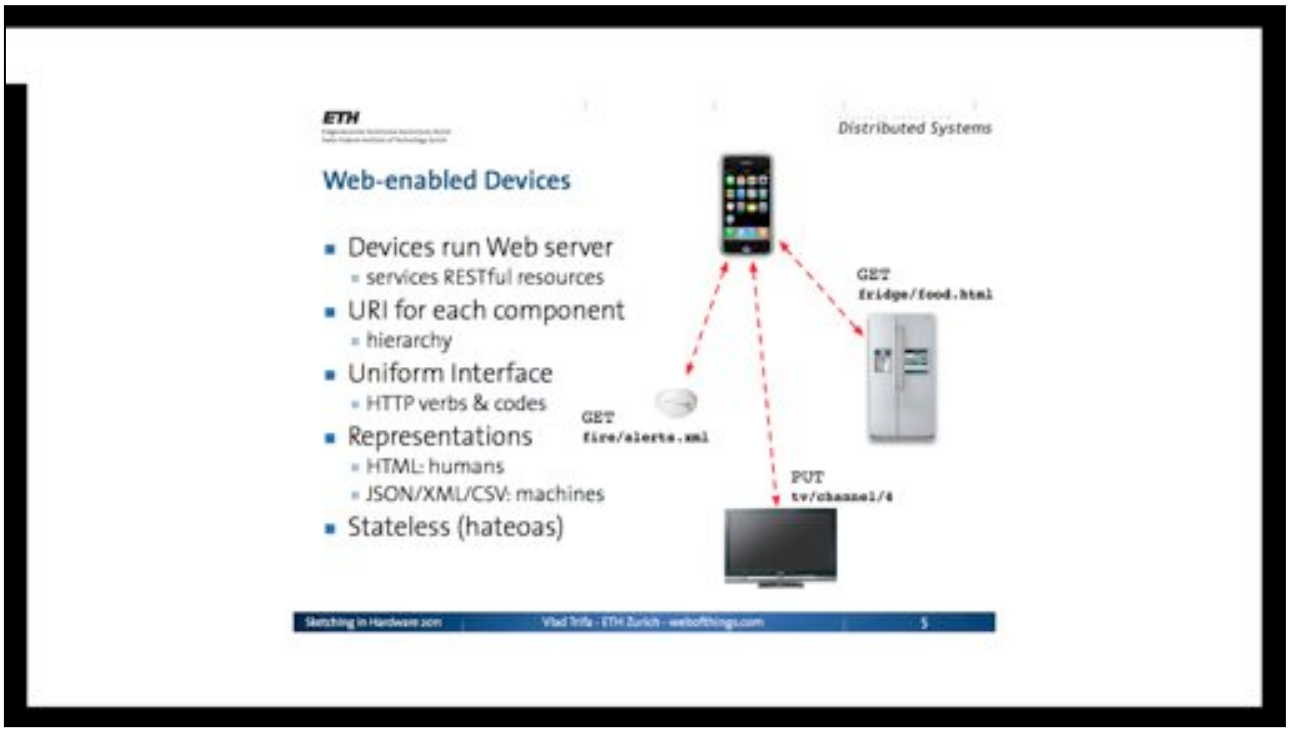

Because we're way ahead of the curve in terms of figuring out how digital things should talk to each other.

Everything that communicates needs to do in some standard way, and increasingly that way looks a lot like the Web. Here's a slide from a project by Vlad Trifa and Dominique Guinard of ETH Zurich. They've build a middleware layer that makes every physical objects look basically like a Web site. They call it, appropriately, The Web of Things. It doesn't sound particularly farfetched. They're just applying stable technical standards that were developed when Web servers were as powerful as today's smart TVs to things like, well, smart TVs.

This allows us to transfer our skills easily, since we can now mash up objects how we mash up web sites.

That's a way of treating devices from afar as you would Web sites, but people's use of devices close by is also becoming more Web-like.

When devices are used to access online services people begin to see through them to the online world they provide access to, rather than looking at them as tools in their own right. In many ways we no longer think of experiences we have on devices as being "online" or "offline," but as services that we can access in a number of different ways, unified by brand identity and continuity of experience. Our expectation is now that it's neither the device nor the software running on it that's the locus of value, but the service that device and software provide access to.

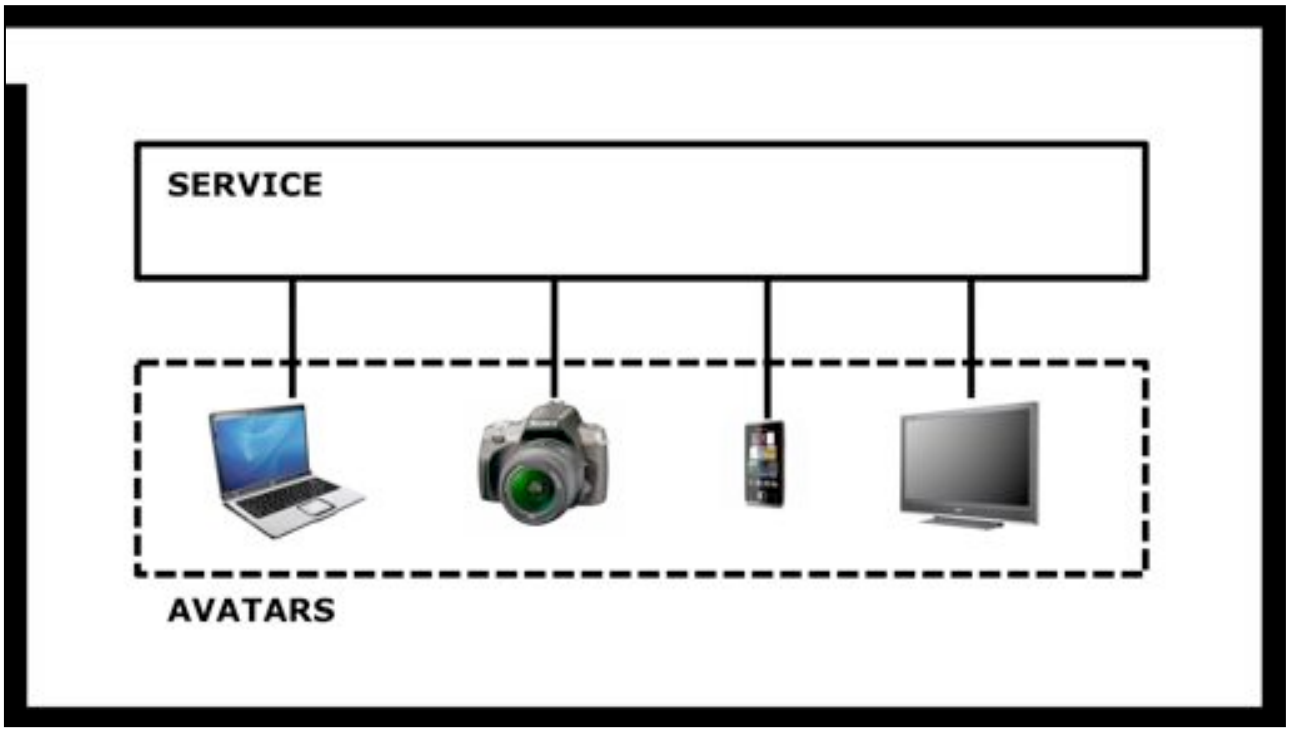

These devices become what I call “service avatars.” A camera becomes a really good appliance for taking photos for Flickr, while a TV becomes a nice Flickr display that you don’t have to log into every time, and a phone becomes a convenient way to take your Flickr pictures on the road.

Thus, the service and the device become increasingly inseparable and we who create the services effectively control the devices.

For example, you can now get Netflix on virtually any terminal that has a screen and a network connection. You can pause a Netflix movie on one terminal and then upause it on another.

Because to the Netflix customer, any device used to watch a movie on Netflix is just a hole in space to the Netflix service. It's a short-term manifestation of a single service. The value, the brand loyalty, and the focus is on the service, not the frame around it. The technology exists to enable the service, not as an end to itself.

This is one way that objects in the world and the digital online map are becoming the same thing, a thing that we as interaction designers, control.

Here's a telling ad from Amazon for the Kindle, which is one of the purest examples of a service avatar based user experience. This ad is saying "Look, use whatever avatar you want. We don't care, as long you stay loyal to our service. You can buy our specialized device, but you don't have to."

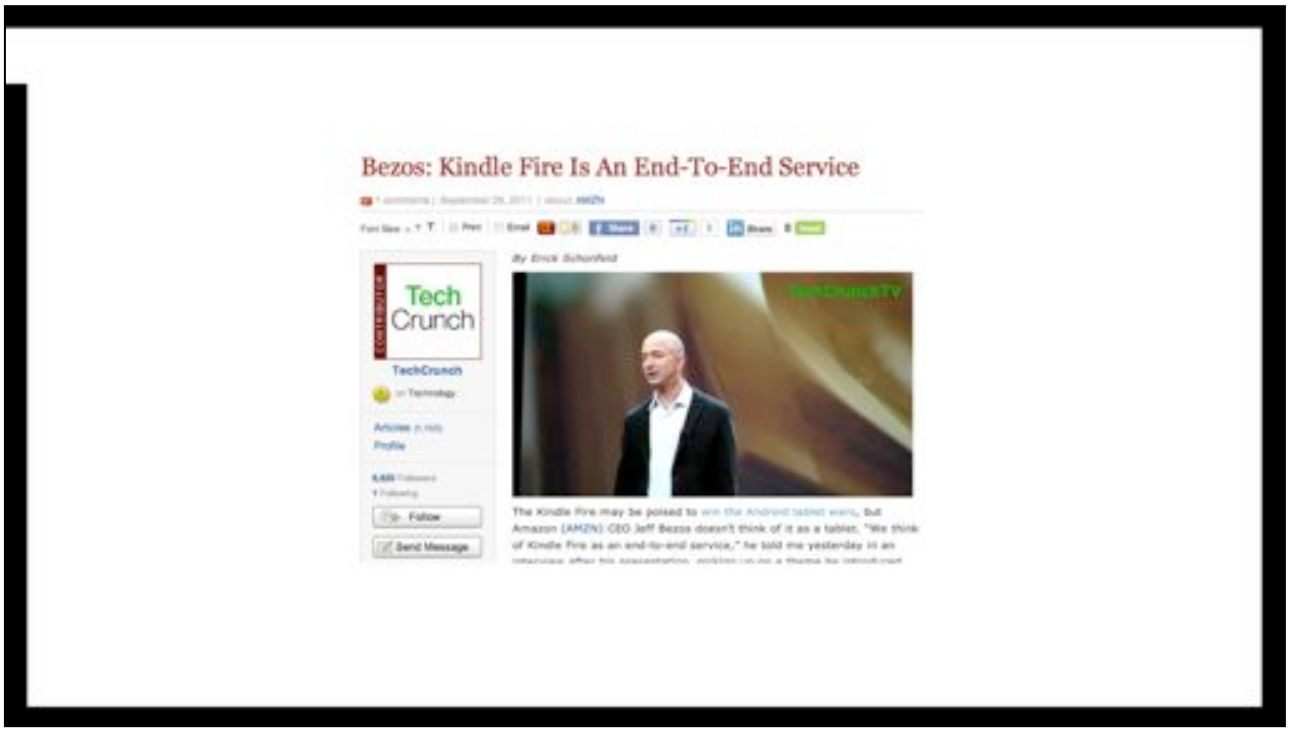

Jeff Bezos is now even referring to it in these terms.

This leads to another experience design conclusion. The core of the product is not the web site that you're designing, or the product you're designing—it's not any of the avatars of the service. The core is the service that lies underneath. The avatars reflect that service, they deliver the product in context-appropriate ways, and their design is very important since they are how people experience the service, but the most important part of the design is the itself.

Thus, when we are designing FOR the Web, we are increasingly designing for the world.

So what's the upshot of all of this? How do these pieces fit into place?

It's still pretty early, and—like I said, we're terrible at identifying emergent phenomena—so we don't really know what this ubicomp elephant looks like. We do, however, have some pointers to what kinds of changes we could see.

Source: Banksy's elephant.

For example, what happens when you mix information shadows and service avatars? You get a blurring between what's a product and what's a service.

When you sign up with a car sharing company like Flexicar or GoGet you become a subscriber to their service.

Each specific car is an avatar of its respected service, actively connected to the service at all times. You can use it any time you want, but you can only open the car and start the engine the service allows it. Your relationship with these cars becomes something different than either renting a car or owning one, sharing elements of both. It's a new kind of relationship that we don't yet have a good word for. And it's a relationship that's created by the capabilities of underlying technologies that didn't exist or were impractical 20 years ago.

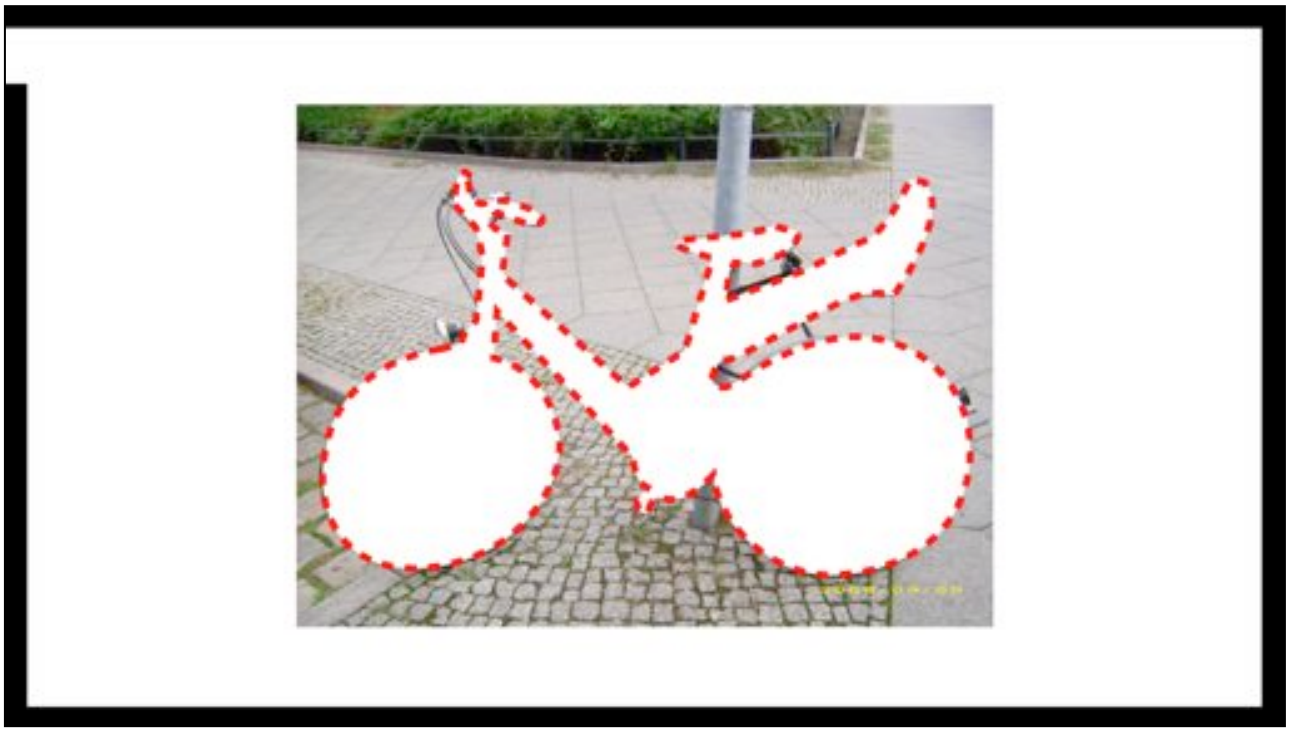

This is the German Call-a-Bike program, run by the rail service. You need a bike, you find one of these bikes, which are usually at major street corners. You use your mobile phone to call the number on the bike. It gives you a code that you punch in to unlock the bike lock. You ride the bike around and when you've arrived, you lock it. The amount of time you rode it automatically gets billed to your phone, by the minute. Each bike is an avatar of the Call-A-Bike service.

Photo CC by probek, found on Flickr.

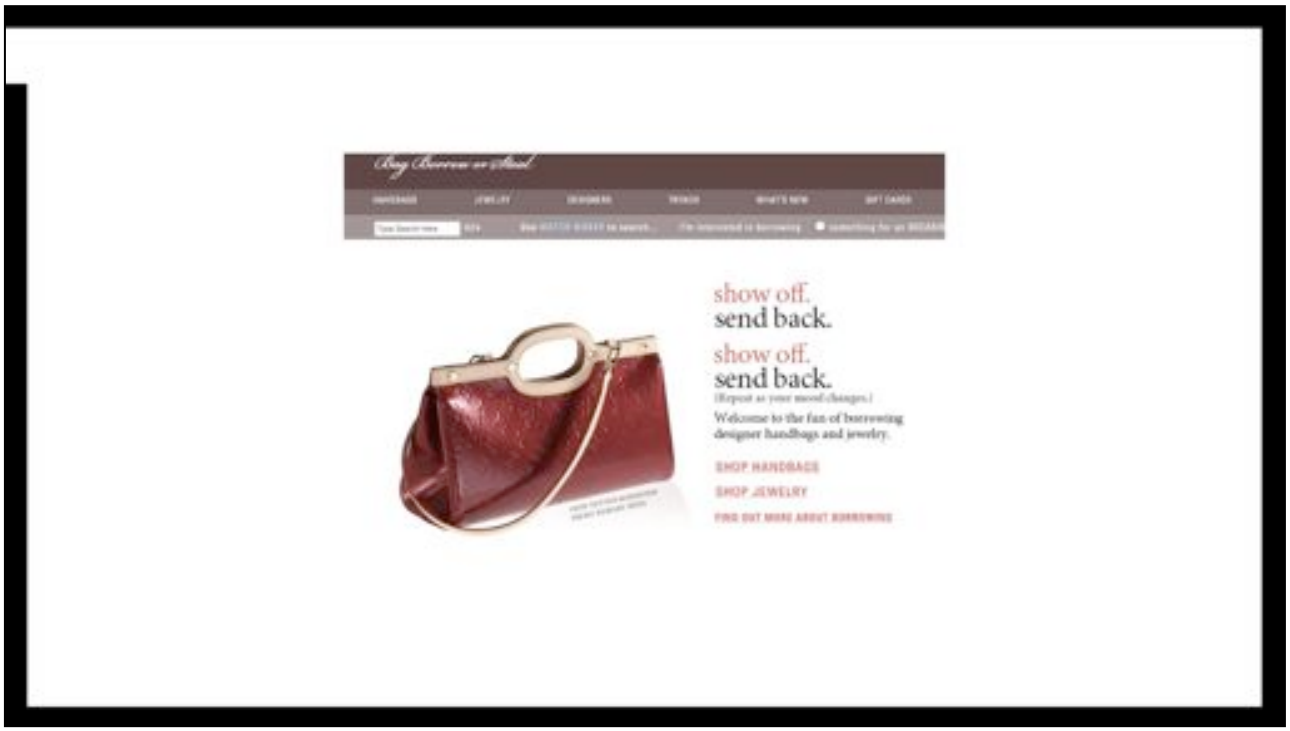

Here's another example that points to some exciting possibilities and that also straddles this model of not quite ownership and not quite rental. Bag, Borrow or Steal is a designer purse subscription site. It works like Netflix, but for really expensive handbags.

It's fashion by subscription. From a user-centered design perspective, it’s great. Here’s a class of infrequently-used, highly desired, expensive objects whose specific instantiation changes with the seasons. You don’t want a specific bag as much as you want whatever the current appropriate thing to fill the dotted line is, but actually keeping up with that fashion is expensive.

This service lets you own that bag possibility space without actually owning a single bag.

Photo CC by bs70, Flickr

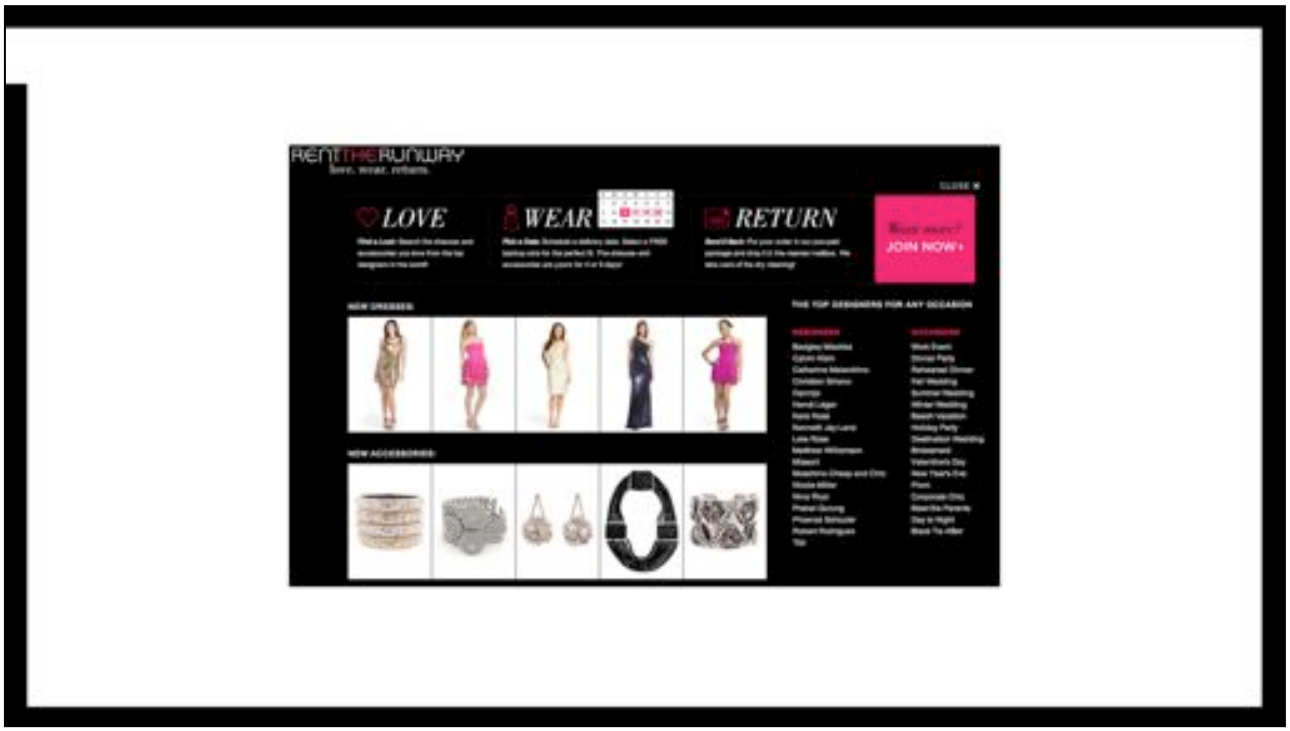

Here's another one called Rent the Runway that has expanded this idea to dresses and accessories.

How long until you get a subscription to Zara and instead of buying your clothes, you just pay a monthly fee to get whatever is seasonal for your type of work in your part of the world at your price point.

We already have Exactitudes and people seem quite comfortable with it. Why not turn it into a subscription business model for Zara?

Another effect, and one which may be the most profound of all, is how our increasing reliance on embedded algorithms shifts relationships of authority and responsibility. This isn't necessarily bad—I, for one, am happy to let Google Maps plot routes for me since it only gets it spectacularly wrong every once in a while—but the more we embed sensors in our world and use automatically processed information to make material changes in the world, the more power we implicitly give algorithms and the more authority we give their designers.

For example, San Francisco has instituted a dynamic parking pricing system called SFPark. Sensors that look like speed bumps are embedded in the pavement. They sense whether a car is in a given parking space or not. This information is uploaded to the cloud where three things happen to it: it serves as the data source for an app that shows drivers where there are empty spaces, it tells meter maids where there are cars with expired meters, and—most interestingly—it uses the parking frequency data to adjust parking prices dynamically. Their stated goal is that the algorithm will price the parking so that there's always two available spaces on every block. Theoretically, a spot in a busy part of town that costs 50 cents an hour at 5AM may cost $50 an hour by 1PM. The people that run this program in San Francisco understand the potential danger of letting such an algorithm run completely free and they've intentionally limited both the price range and how often it changes, but the fact that they felt they had to do that shows that a public negotiation with algorithms that control the world has already begun. You can see a similar negotiation happening with smart electrical meter pricing.

This kind of negotiation is happening all the way to the personal level, down to individuals and their relationship with themselves.

Right now the Quantified Self movement is quite popular in the San Francisco Bay Area. People are using a wide variety of sensors to measure things about themselves so that they can optimize their bodies and lives. Here's the cloud-connected pedometer from Fitbit, Bodymedia's multi-sensor cuff. The sleep sensor from Zeo. They're all designed to collect data about you, then process it, perhaps share it, and visualize it. They're great examples of service avatars made with information as a material. But there's something about them that unsettles me.

At their core, they're shifting intrinsic rewards, the positive internal drive for being healthier, getting better sleep, being more fit, to extrinsic rewards--making numbers go up. But those extrinsic rewards are controlled by algorithms, rather than their owners' judgment. What these products are saying, in effect, is that we can become the people we want to by giving up some of the control of our lives to these digital devices. Perhaps that's true—people depend on a lot of tools—but what results is a hybrid between a person with goals and a set of algorithms that purports to tell them whether those goals have been achieved. This is likely to have many unintended consequences. We trust algorithms and sensors because they look objective, but are they? How do we know?

This is the Water Pebble. It aims to reduce water usage by timing your shower and telling you then you hit your designated shower time. The way it works is that when you first use it, you push a button and take a shower. That sets the baseline. From then on it works like a shower timer. The algorithmic part of it comes in when, after a while, it starts slowly reducing the amount of time it gives you, so that you progressively build a habit of using less water.

My personal experience with it, however, is that its algorithm for behavior change doesn't match my ability to actually change. It reduced the amount of time it gave me to shower, and I was following along with it, until my change curve deviated from its. Instead of helping me change my behavior, it just sat there at in the shower drain blinking red and mocking me for not being good enough. I couldn't reason with it, I couldn't get it to change its algorithm to match my capabilities, so I stopped using it.

I'm not saying that we shouldn't enter into these relationships, but that they represent a deep shift in how we relate to the world. We shift our trust and the responsibility of making sense of the world to algorithms more than our own capabilities. We are likely going to spend the rest of our lives negotiating power relationships with embedded devices, in a way that no people have ever

And we can expect many unintended consequences. The designers of Facebook, Twitter, YouTube and text messaging did not, and could not have predicted, a new papyrus-level crisis in Egyptian government. And yet they provided the medium through which that revolution happened, largely confirming Ethan Zuckerman's assertion that any technology that can be used to share cute cat pictures can be used to overthrow a government.

We, those who grew up on the net and who design it, will be the ones who create ubiquitous computing, not the roboticists or network engineers, and ubicomp will fundamentally change the world and us along with it. We have tremendous power and enormous responsibility. And it's our responsibility to enjoy ourselves, make great stuff, take huge risks, and be thoughtful about the implications of what we're doing without ever forgetting that we have no idea what's going to happen next.

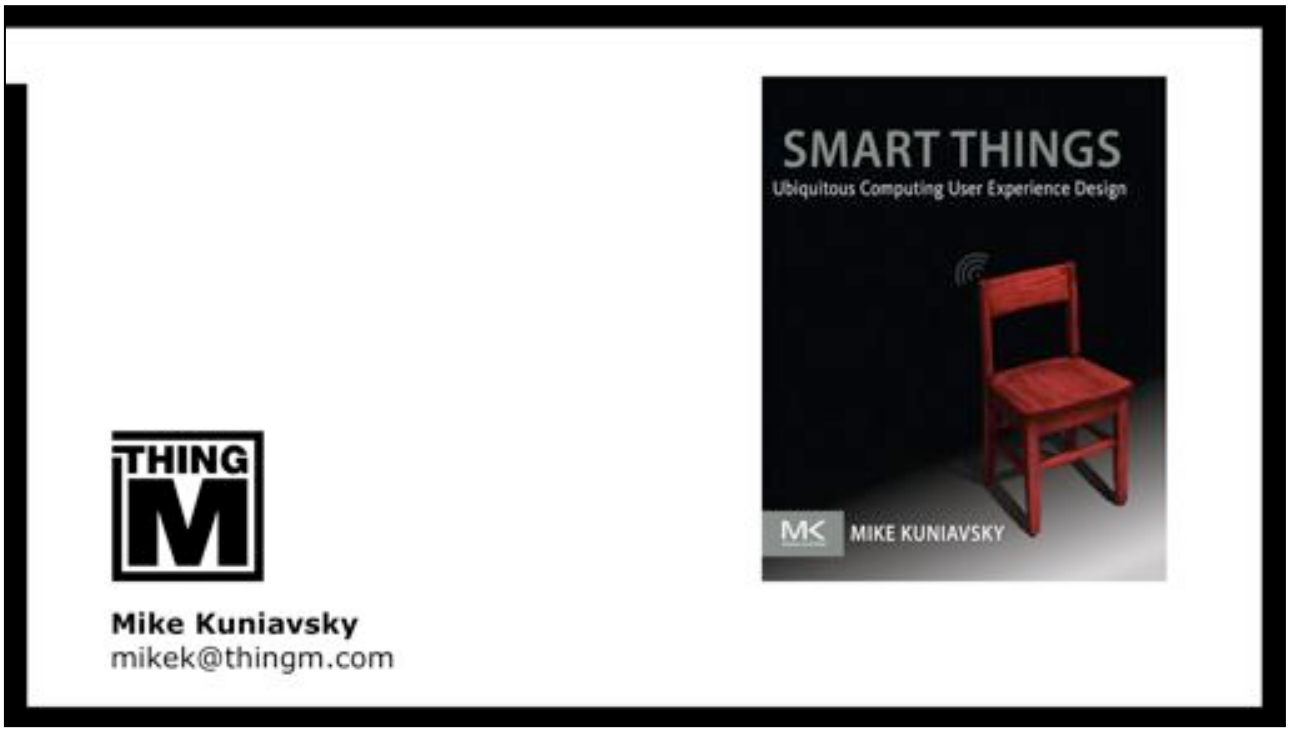
SMART THINGS
Ubiquitous Computing User Experience Design
MIKE KUNIAVSKY
THING M
Mike Kuniavsky
mikek@thingm.com

Thank you.