DESIGNING SMART THINGS

USER EXPERIENCE DESIGN FOR NETWORKED DEVICES

Mike Kuniavsky
UX Week
San Francisco
August 24, 2011

First, let me tell you a bit about myself. I'm a user experience designer and entrepreneur. I was one of the first professional Web designers in 1993. Since then I've worked on the user experience design of hundreds of web sites. I also consult on the design of digital consumer products, and I've helped a number of consumer electronics and appliance manufacturers create better user experiences and more user centered design cultures.

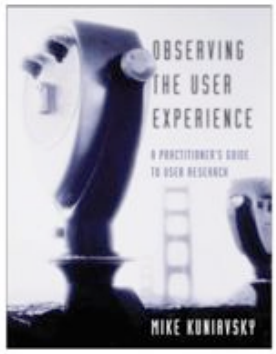

I sat out the first dotcom crash writing a book based on the work I had been doing. It's a cookbook of user research methods.

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

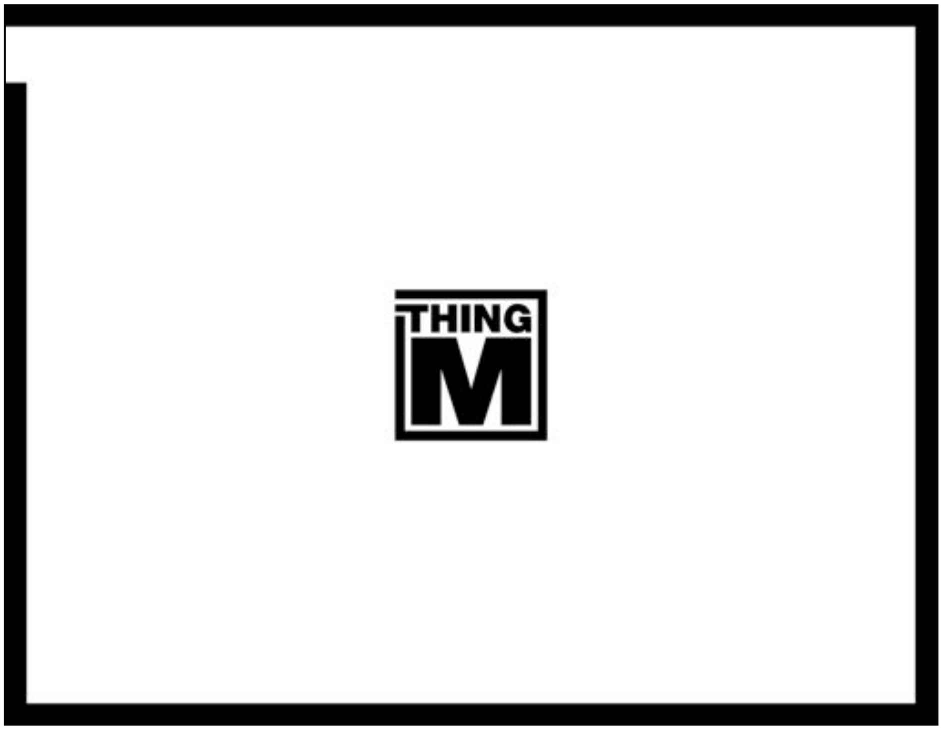

…and three years later I left it, and the Web behind, and founded a company with Tod E. Kurt called ThingM in 2006.

We're a micro-OEM. We design and manufactures a range of smart LEDs for architects, industrial designers and hackers.

I have a new startup called Crowdlight I’m trying to get off the ground and I’m currently consulting for the R&D lab of a major consumer electronics manufacturer.

This workshop 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.

This is a workshop on user experience design for networked devices, and I mean that in the broadest sense. My focus is not just on designing multi-touch apps for tablets or 10-foot UIs for connected TVs. That kind of screen design is part of it, but from my perspective it's a subset of a larger set of design challenges and possibilities around the design of digital devices that are connected to the internet. My goal is to think about the broader experience design possibilities created when any device becomes connected. This means rethinking the possibilities of many things from the ground up. That, of course, can't be covered in a single day workshop.

This picture, by the way, is of a visual designer, an interaction designer and an architect mocking up a new kind of clock in a workshop I ran a couple of years ago.

What we're going to try to do today is to give you a feel for how design of connected objects is different than the design of things that you may be familiar with, and to give you some concepts and tools that may help you with that. We will focus less on specific techniques than on thinking about how to deploy concepts and critically to ask questions about these sorts of projects so that you can be a better judge of your own designs and the designs of others.

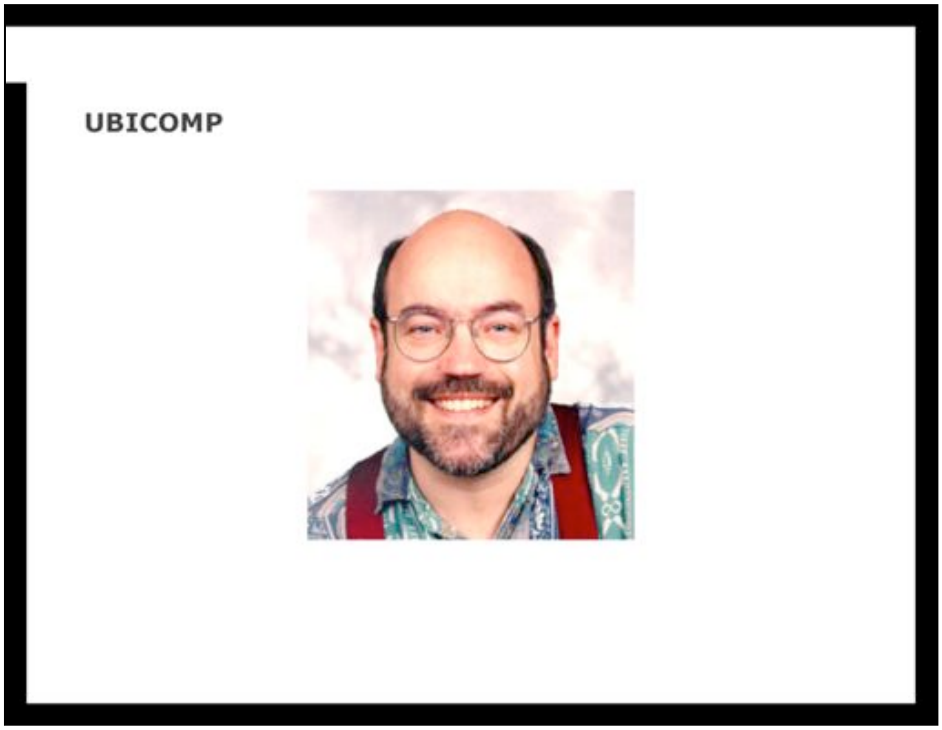

First, I'd like to set some foundational definitions, and that starts with what it is that we're talking about here. I define all digital connected device design as part of the same larger trend that was identified and named by the late Mark Weiser, then the CTO of Xerox PARC.

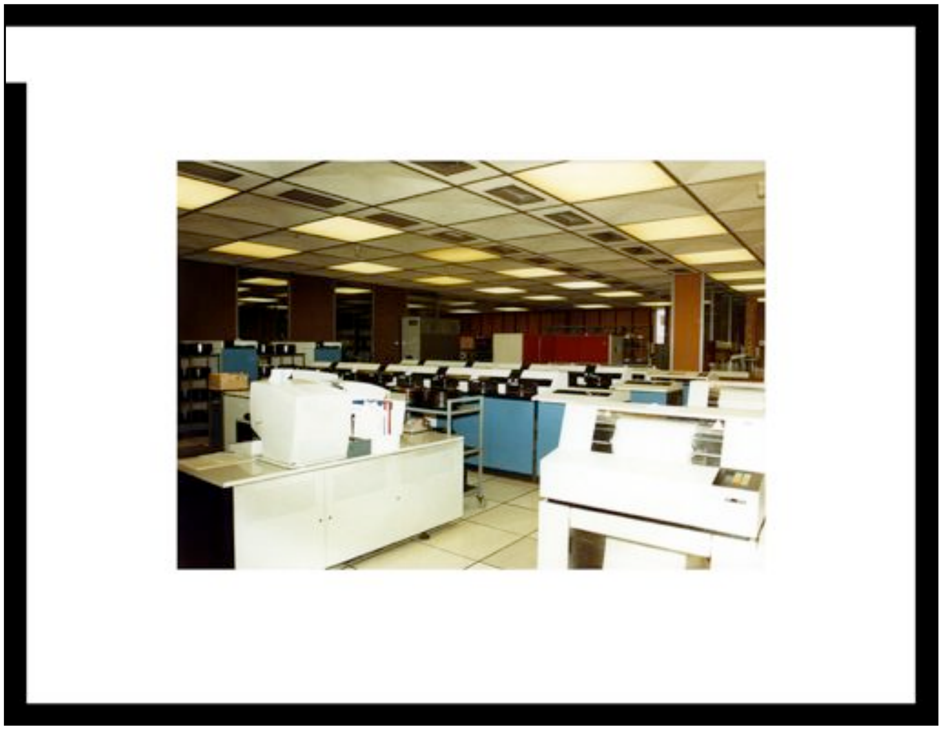

More than twenty years ago he envisioned a world that didn't have one big general purpose computer per household, but many computers distributed throughout the environment. He called this trend "ubiquitous computing," or "ubicomp."

Photo CC by Richard Pluck, found on Flickr.

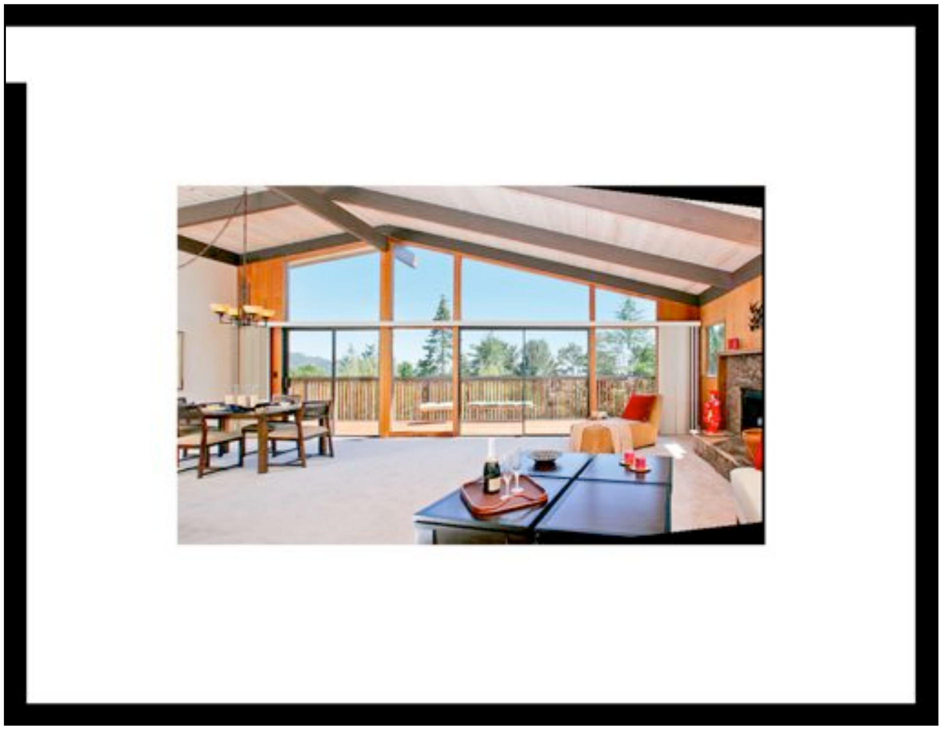

but could be, in his words, woven "into the fabric of everyday life until they are indistinguishable from it."

Photo: CC Mark Lacoste, original: http://www.flickr.com/photos/croco/914189024/

As electronics consumers we’re most clearly experiencing this as a proliferation of device form factors. Our general purpose computers now come in many shapes and sizes.

But I’m talking about a deeper change. Our relationship to our environment is fundamentally changing through the embedding of technology throughout our everyday environment. This looks like a typical San Francisco parking meter, but it’s actually part of an extensive network of overlapping services.

Streetline Networks is one such system that connects parking meters to sensors in the pavement that look like speed bumps but actually identify which parking spaces have cars in them communicating through base stations installed in street lights. This allows the city to be able to know which spaces are open when so that they can write tickets more efficiently and change the price of parking based on demand. It also allows parkers to get a real time map of where there are open parking spaces. This isn't some research project, it's an actual system that's installed in San Francisco. These are from the pamphlet Streetline published FIVE years ago. A similar system went on line in big parts of San Francisco earlier this year.

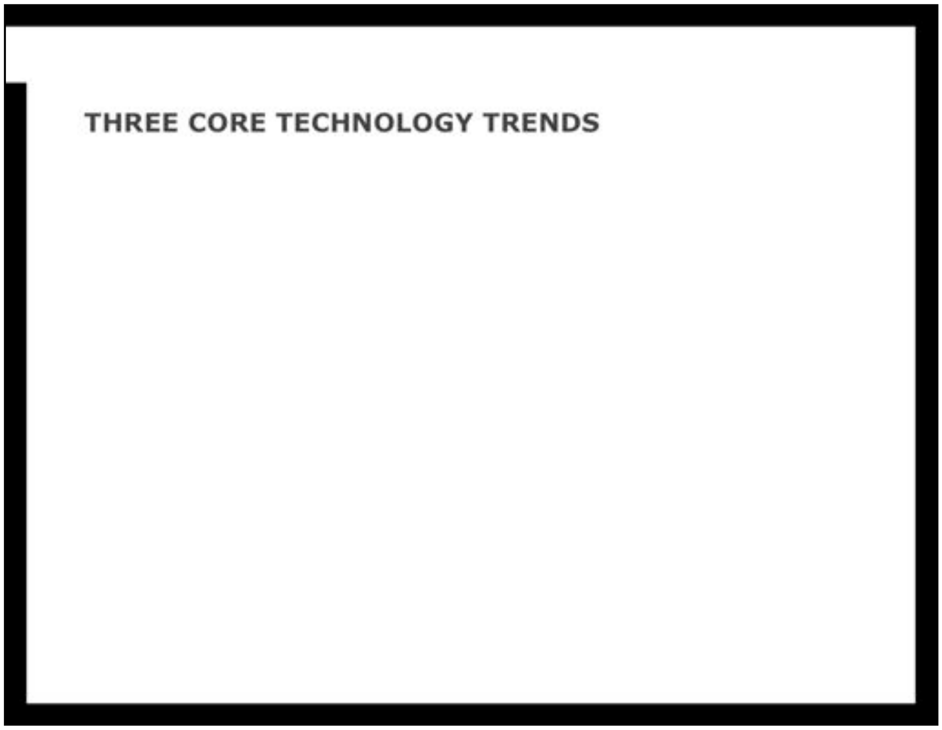

Now how did we get here? I believe that this is happening because of an intersection of three trends.

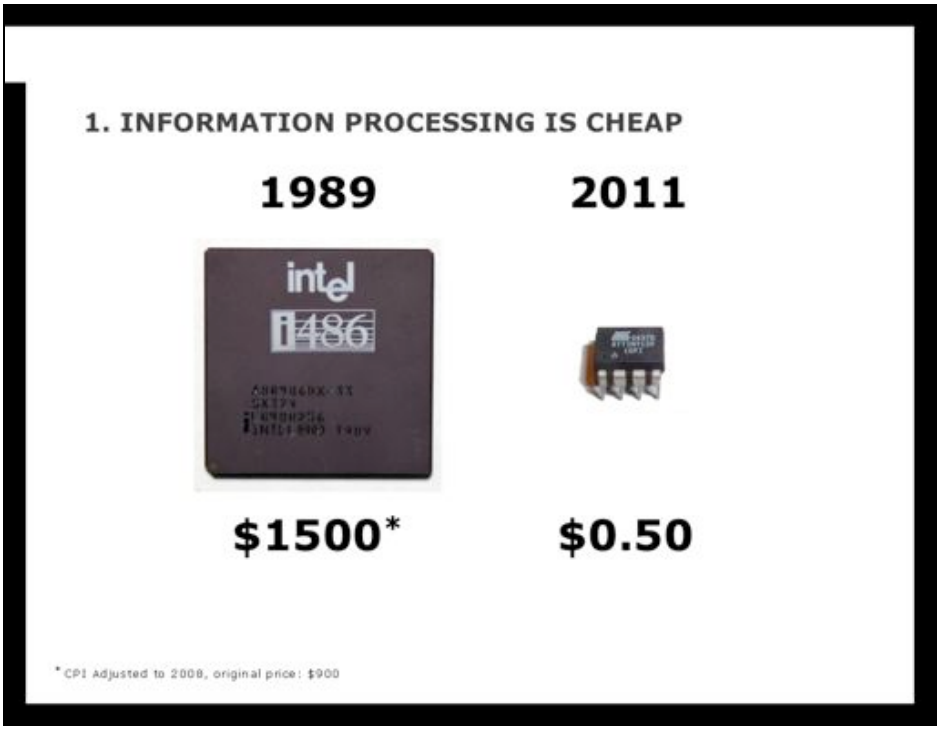

I want to start by talking about Moore's Law, since that's where all conversations about the implication of digital technology start. When people talk about Moore's Law, it's often in the context of maximum processing power. But it's actually something different. It's actually a description of the cost 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.

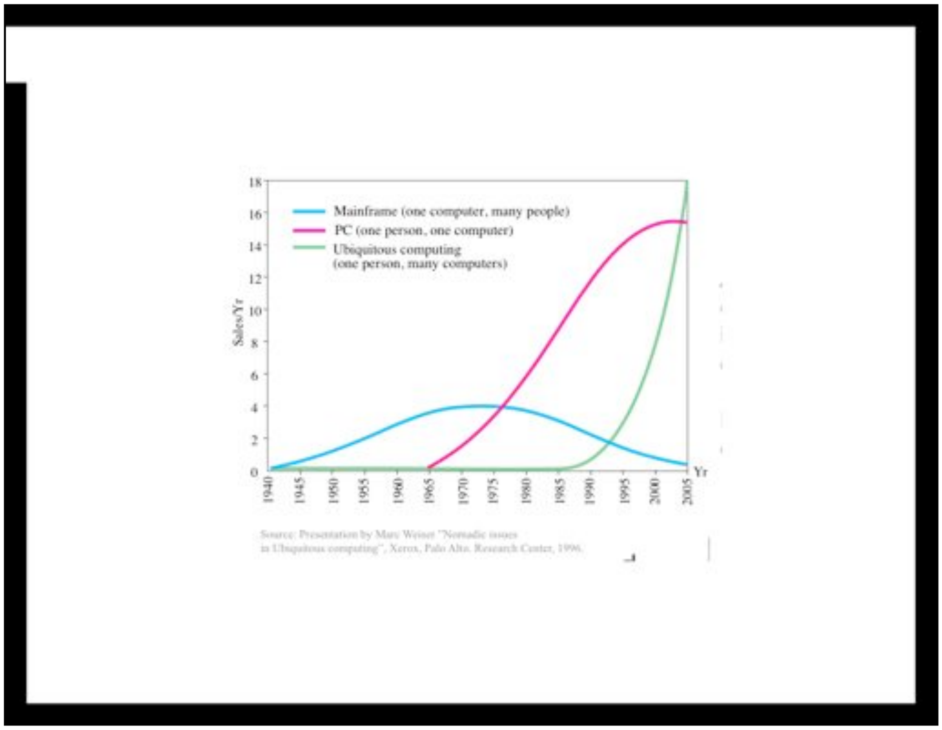

Here's Mark Weiser's diagram showing the shift from mainframes to ubiquitous computing from 15 years ago. He missed cloud services, so this isn't technically true, but it's generally a good model to think about how our world is changing because of all of the inexpensive applications for processing.

Basically what this is saying is that information processing used to be expensive and had to be limited to special devices, but now it is cheap and can be used in all kinds of novel situations. This means that you can now include powerful processing and networking in almost anything, and start rethinking the design of everything in terms of embedded digital technology I'll explore the implications of this again later in this talk, but first I want to talk about the other major technological changes that are driving ubicomp.

The other dominant trend right now is of course pervasive data communication. This is an image from Timo Arnall that's envisioning 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.

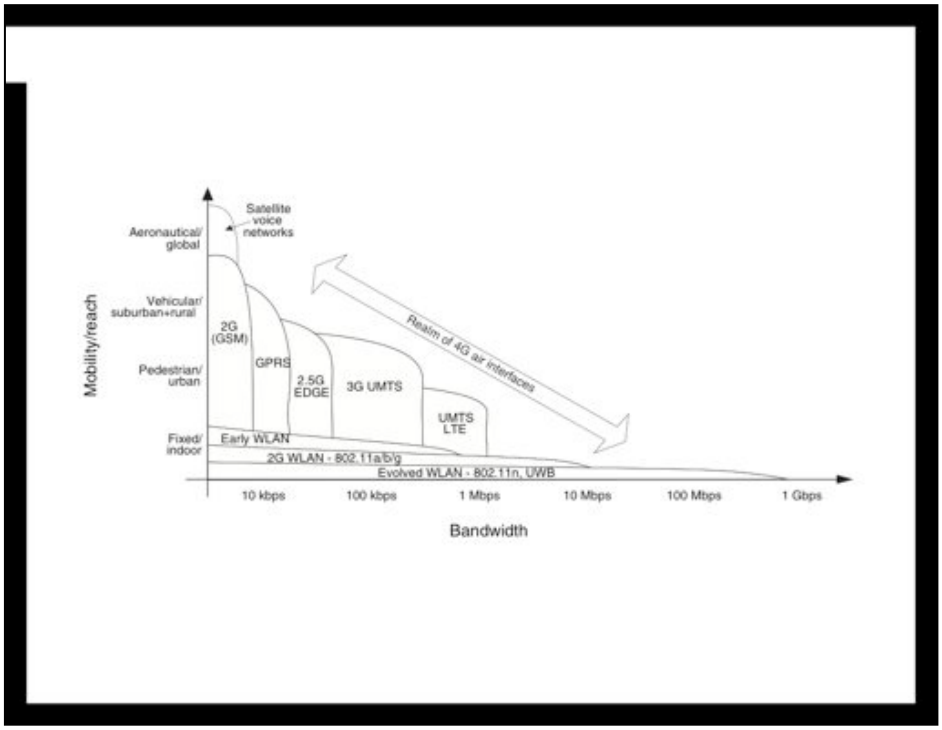

There are a vast number of networking technologies. They trend to trade off three things, distance, bandwidth and power. If you maximize any one of them, you minimize the other two. So if you want something that's fast, and has long range, it'll require a lot of power. If you want something that uses very little power, it'll have to be either slow or close, or, more likely, both.

Here’s what Cisco estimates the trend of wireless data traffic is going to look like. The baseline here is last year. For comparison, to represent the amount of data for 2008 you need a line that’s 1/8th as thick as that small line on the left and 2006 was 1/24th as thick. You get the idea. Wireless data has gotten pervasive and, judging by this level of adoption, very cheap.

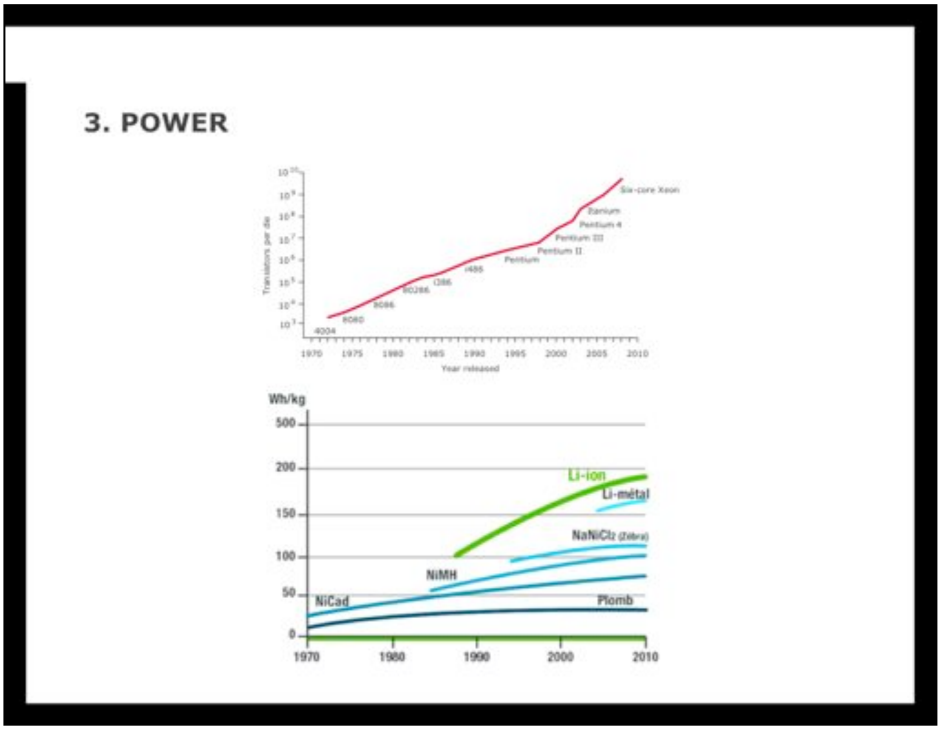

Which brings me to power, which is not so much a trend as an anti trend. I’m sure you’re familiar with the fact that battery technology has not advanced as fast as processing power, but let me show you exactly how much. In the time that processing power increased by a factor of ten million, battery efficiency increased by a factor of, let’s be generous, ten. Probably closer to four. This means that in practice many of the things we can theoretically do with processors, we can’t do in practice because of batteries. If there’s a brake on the advance of ubiquitous computing, it’s power.

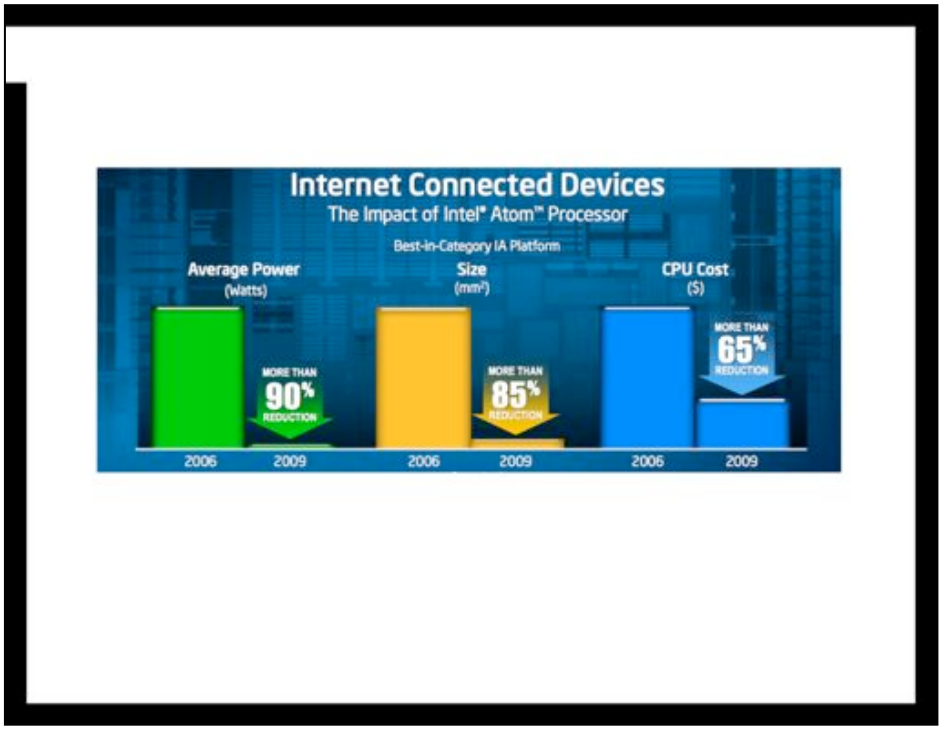

The combination of these factors has created a shift away from raw processing power to the application of processing, which has led CPU manufacturers to emphasize different things.

Here's a slide from a 2009 talk from Paul Otellini, the CEO of Intel. Notice that instead of talking about numbers going up, processor manufacturing has become all about pushing numbers down. Instead of competing on doing more with more, they are now competing on doing the same with less. Less power, smaller size, and lower cost.

What these manufacturers are doing is that they're emphasizing the context-specific use of information processing, rather than raw throughput and they're aggressively creating new classes of processors that enable information processing to happen with smaller amounts of energy.

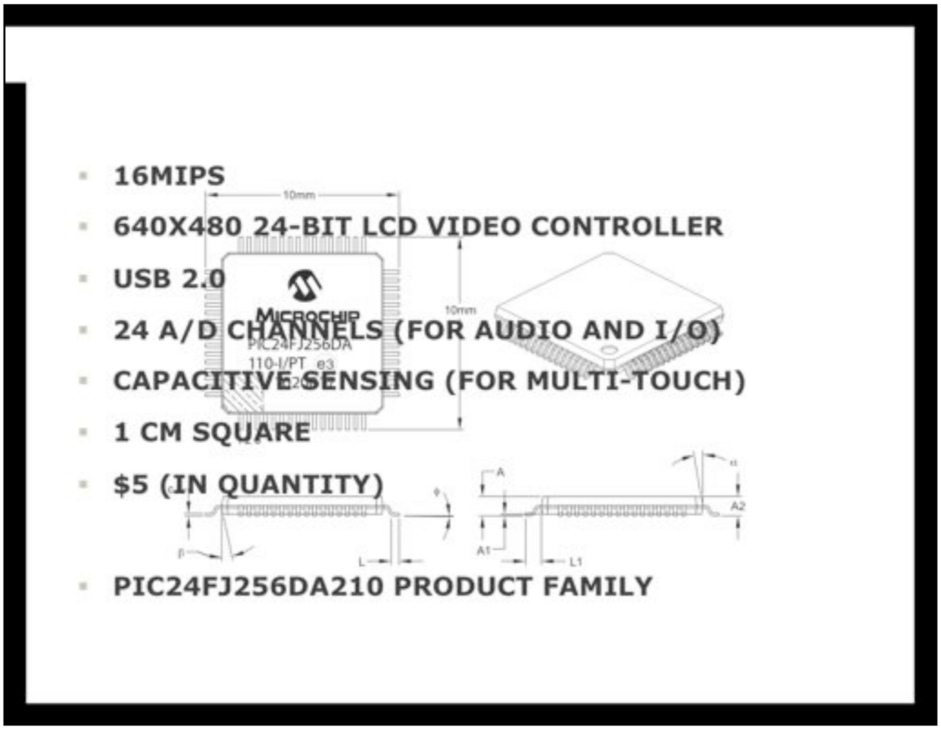

This new System on a Chip from Microchip has about as much processing power as that initial 486, but is also has an onboard video controller that can drive a VGA-class screen, a USB controller for peripherals, a 24-channel analog to digital converter for sensor, and a capacitive sensing driver that can drive a touch screen. It costs about $5, uses less power than a keyring LED flashlight, and fits on a chip the size of your fingernail. It's also not unusual. Almost every semiconductor maker makes similar products.

Ok, so that may have seemed like an obvious beginning: sure, processing is cheap, networking is pervasive, and we have specialized chips, but we knew that. True, but revolutions rarely come completely unexpectedly. The pieces are all around for us to see, but it's a set of circumstances that puts them together. I think that we hit the tipping point to ubiquitous computing in 2005. That's the year Apple put out the iPod Shuffle, Adidas launched the adidas_1 shoe and iRobot launched the Roomba Discovery, their second generation model. That was the year that it began to make sense to create devices that compete through information processor-enabled behavior. The Tickle Me Elmo Extreme, which came out in 2006, is a prime example of this. It's a toy that creates its competitive advantage, that justifies its $80 introductory price in a world of $20 plush toys, by using information processing.

It is 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.

In other words, information is quickly becoming a material to design with.

If you look at what happened when the price of extracting aluminum dropped by two orders of magnitude 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. It becomes a design material.

This capability of everyday objects to make autonomous decisions and act using arbitrary information 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

SUMMARY

- Our relationship to our environment is fundamentally changing through embedded information technology and networking.
- Ubiquitous computing is the most common name for this change (though it has others: pervasive computing, ambient intelligence, the Internet of Things, etc.)
- It's a result of:
  - Cheap processing
  - Cheap, pervasive networking
  - Low-power, but still powerful, embedded processors
- Which came together to create a tipping point (roughly speaking in 2005).
- People are using these technologies to compete by embedding information processing into a wide variety of everyday objects.
- In effect treating information as a novel material to design with.

Q&A

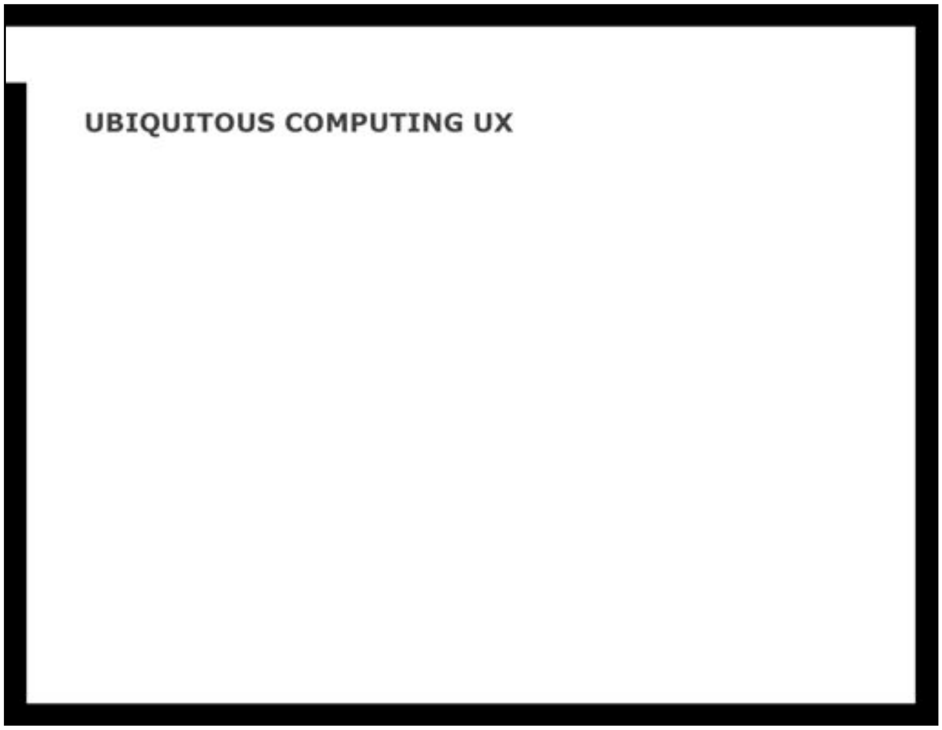

Now what does this mean for the user experience design of such devices.

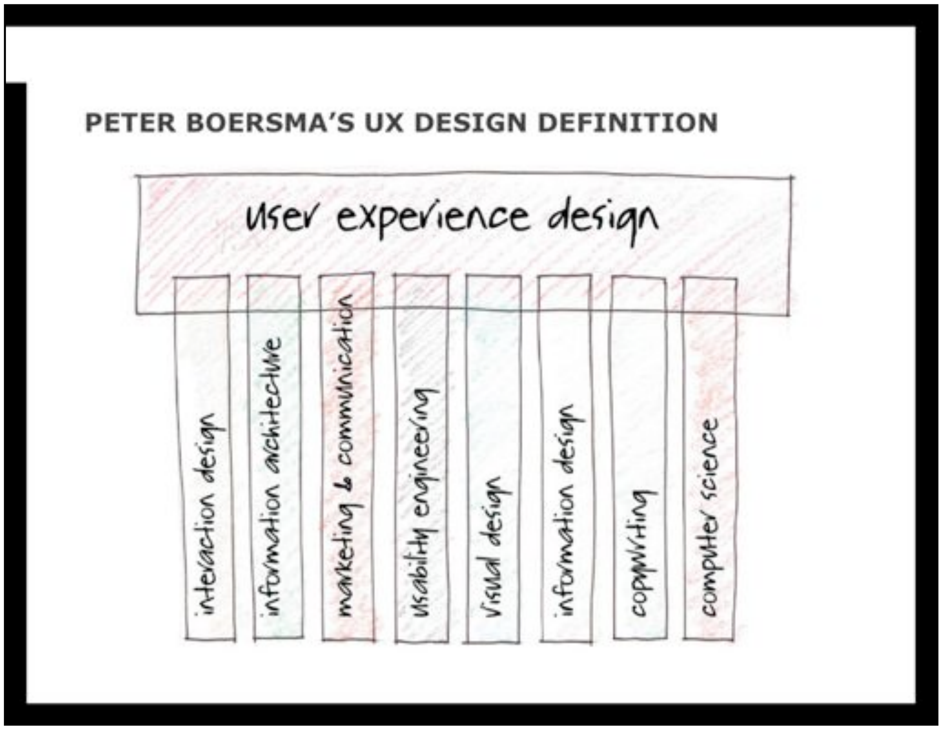

First, let me define user experience as it relates to ubicomp. In 2004 Peter Boersma, now of AP's Amsterdam office, defined UX design as a combination of eight disciplines. Interaction design, Information Architecture, marcomm, usability engineering, visual design, information design, copywriting and CS. This is a very accurate description of primarily screen-based experience design.

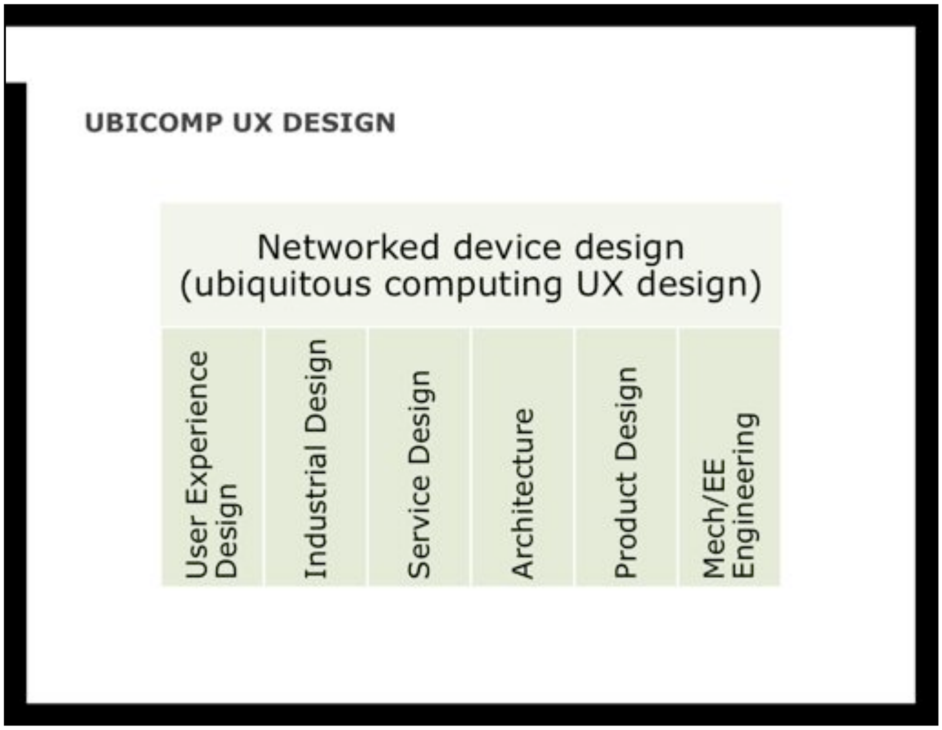

When working with ubiquitous computing devices, the landscape adds to this. In addition to all of the things Peter mentioned in 2004, there are now considerations of the physical design of devices, what services they connect to, how they operate in space, how they're created as marketable products, and what the technical capabilities of the specific technologies involved are, since these devices can no longer assume a generic computer platform.

In other words, we're moving from a world where the basic building blocks are pixels controlled by a single processor, where the design challenge is how to arrange the pixels appropriately and how to change that arrangement through time to a world where the basic building blocks are made of atoms and controlled by many processors. In this world, the main challenges are what shape to make out of the atoms and how to get the blocks to talk to each other.

This is picture of literally digital building blocks made by Modrobotics, but this same principle applies when you have a phone, a smart TV, a tablet and a self-checkout kiosk at the supermarket, and you'd like them all to work together in some way.

There are some big UX trends that I feel are important to understand with our current environment of networked objects.

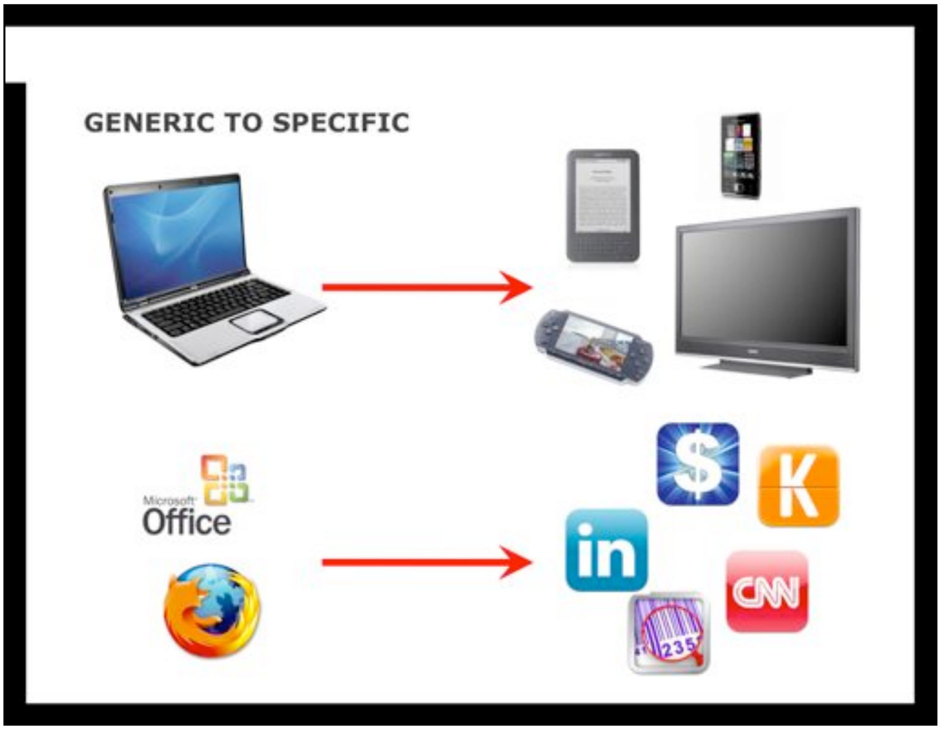

The result of cheap processing is a shift from generic devices and software to specialized devices and software. When computing was expensive, you had one or two general purpose devices that had deal with almost every situation. This necessitated design compromises that resulted in devices and software that could do almost everything, but did none of it well, and UX design was always a set of compromises about creating functionality within the constraints of an OS, an application environment or a browser.

Now that processing is so cheap, this is no longer true. You can now have a high degree of specialization. Your tool is now a tool BOX, a combination of 10, 20, or 30 computing devices and apps that you get for the price of that one expensive device ten years ago. You acquire new functionality as needed and every device and unit of software has a narrower purpose.

This fragmentation then creates a new set of challenges for users, which in turn become challenges for designers. Users no longer have to maintain two sets of UI standards—one for the device and operating system and the other for the application—in mind. Use is much more direct. You pick up a PSP, you know what it's for. You launch the CNN app and you know what content to expect there. However, it now creates a burden of deciding WHAT to put in your toolbox. We only have so much room in our backpacks and app docks. We're encountering new problems, such as finding which app or device does what we want, not which menu. It's a findability problem on the macro scale—should I

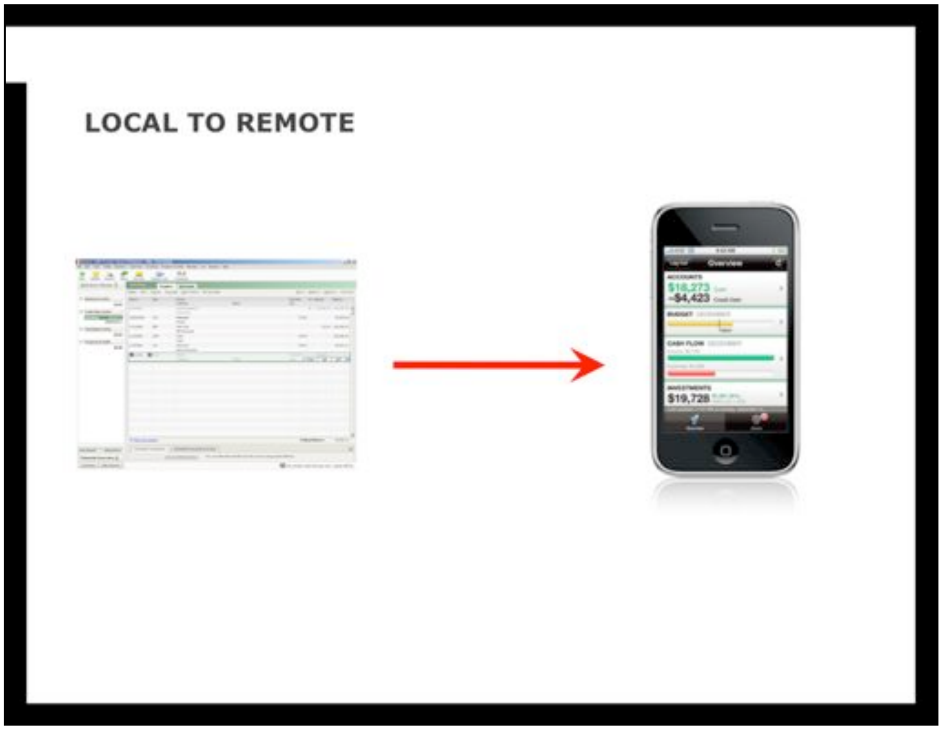

To this we add the effect of widely networked devices, which is to move value away from the local environment to a remote one.

The lasting legacy of the Web is that we have a shift in people's perception of the value digital technology from being primarily local to being primarily remote. The Web demonstrated that moving functionality online enables access to more compute power, continuous updates, real-time usage analytics, and (of course) social connections. It also created a shift in people's expectations. Today, most people understand that the experience you see on one device is often a part of something that's distributed throughout the world. There's no longer a need to pack everything into a single piece of software, and there's no expectation that everything will be there.

Again, this is great, but it too has created new problems. Perhaps it's temporary, but it's disconcerting for me when an application I use that's "on my Android phone" changes functionality after an overnight update. I'm not used to the applications that I "own" to shift functionality. I don't use it as a net app—it's a note-taking app or something and it's ON my phone—and suddenly it changes behavior. What if I don't want all of the changes? Our supposedly stable device, the thing we own, becomes a kind of slippery eel of changing

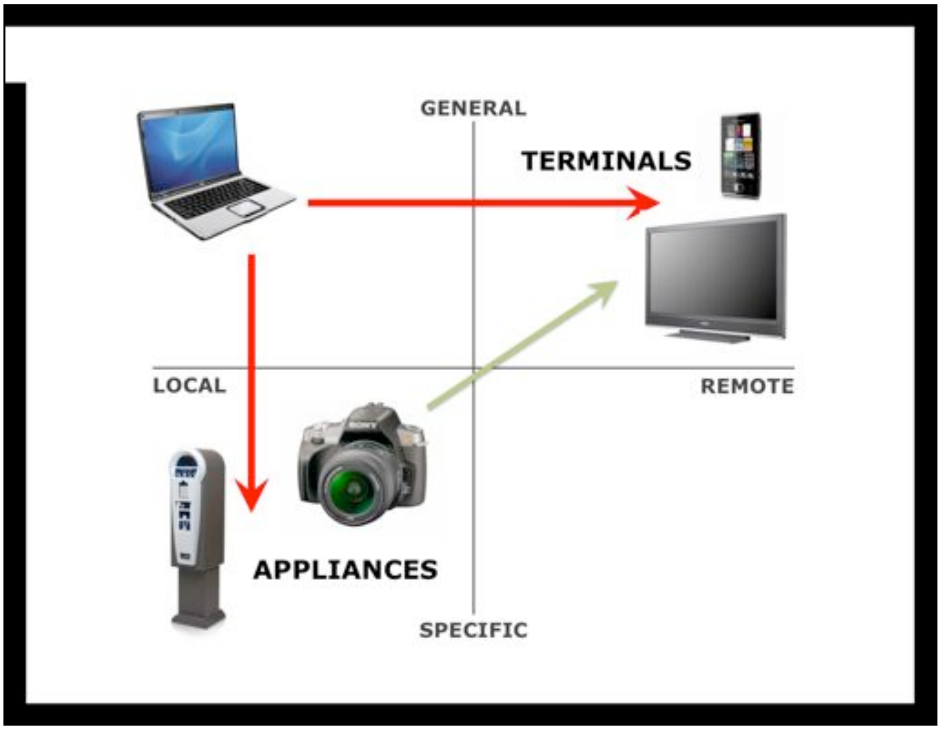

If we chart these two tends, two broad classes of digital products emerge.

If we follow the general to specific axis, we see a shift is to more narrow-function devices that are designed to do a small, specific set of things really well. They primarily differ in what those specific things are. I call these devices appliances.

If we follow the local to remote axis, we find general-purpose devices that do roughly the same set of things, and differ primarily in size. They exist to provide access to online services, in a form factor that's appropriate to the context in which they're used.

I call these devices terminals.

Now, while the digital world as a whole has seen an increase in the kinds of digital devices that exist, the consumer electronics industry has, on the whole, moved in one direction.

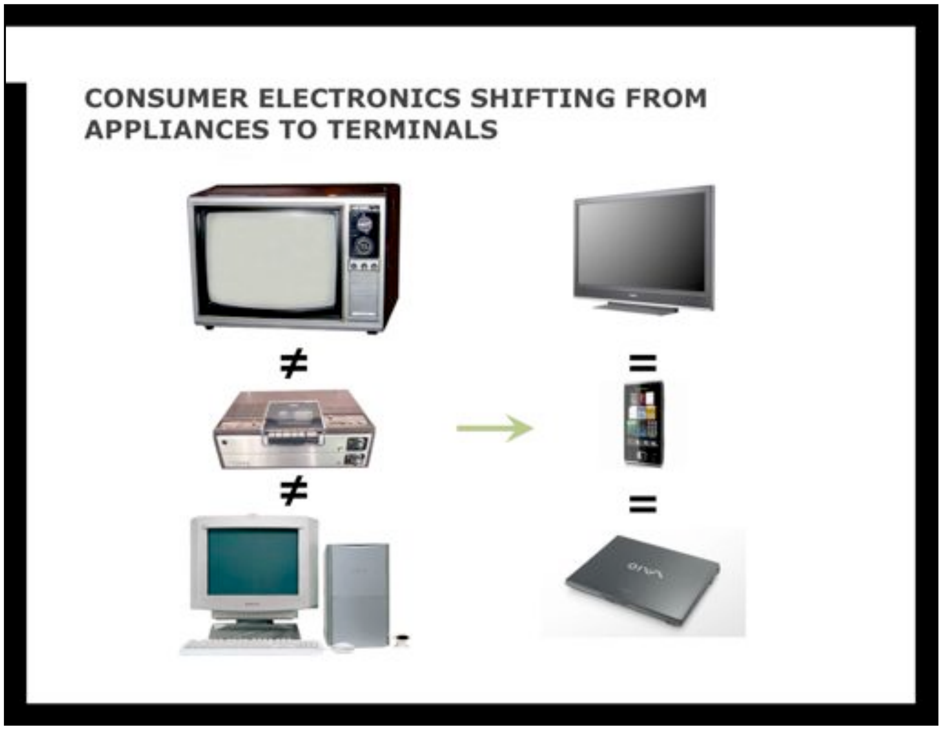

Let me look at consumer electronics a little closer, since they're the primary source of digital devices in our lives, and they're undergoing an enormous change that's becoming a huge challenge.

Consumer electronics used to make appliances, that were known as brown goods (in contrast to the white goods in the kitchen and laundry room).

These goods were of a specific narrow function. A TV was never going to be anything, but a TV, a VCR never anything but a VCR, a computer nothing but a general purpose computing device. Now consumer electronics is largely the business of building variations on the same thing in different form factors, turned to different use contexts. The devices themselves are nearly interchangeable in terms of what they can do, but they're tuned to specific ways that people use them, whether it's traveling, sitting in a living room or in focused interaction. This is great, but I think there's a limit to the effectiveness of this. Eventually, we're going to run out of rectangle form factors. We already have four—TVs, tablets, laptops, and phones—and we're maybe get one or two more. Perhaps watches on one end of the scale and movie screens on the other.

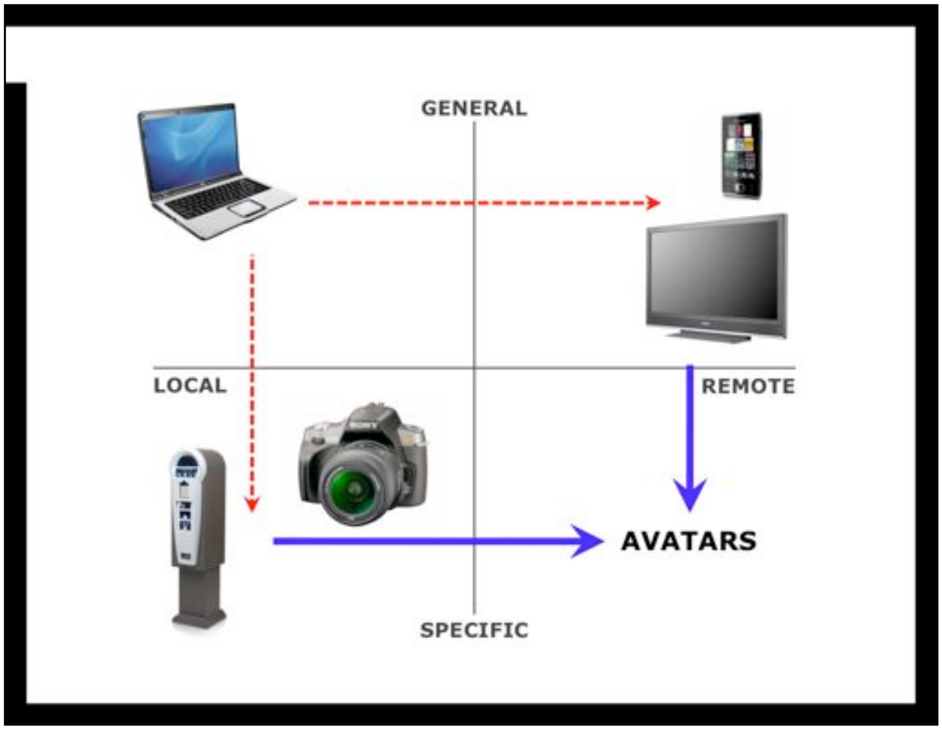

Back to this diagram.

I think that there's an even larger shift going on where devices are simultaneously specific AND deeply tied to online services. In this model, the service provides the majority of the value, and can be represented either as an inexpensive dedicated hardware device, an app running on a terminal, or anything in between.

It's an approach that combines the precision of appliances with the flexibility of terminals to create a fundamentally new class of products that can fill every possible niche where a service may be appropriate.

I call these devices service avatars.

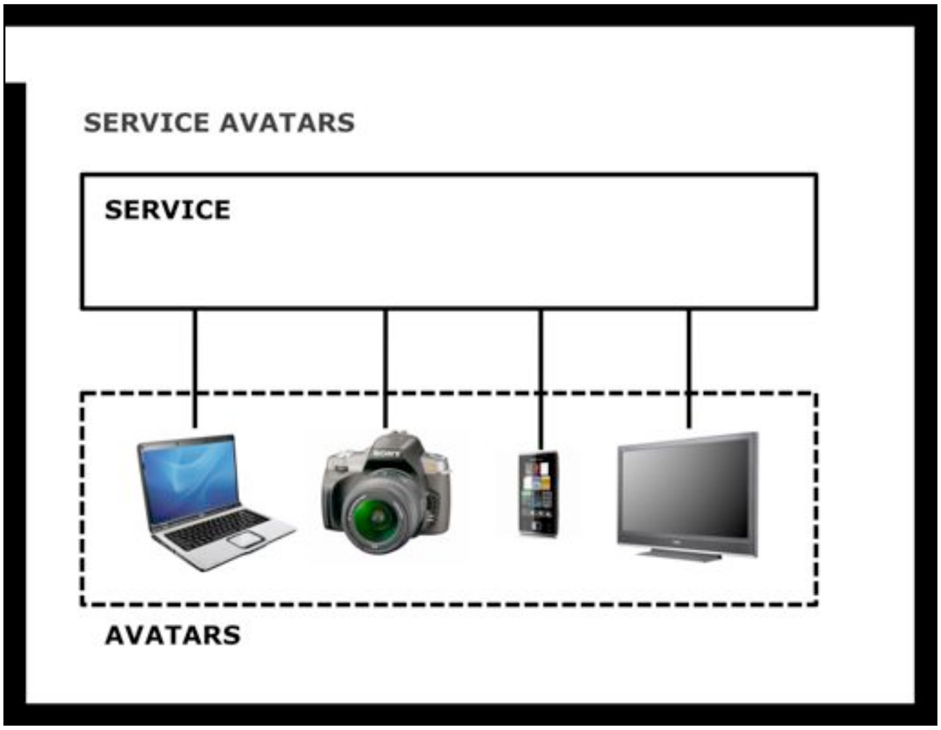

As value shifts to services, the devices, software applications and websites used to access it—its avatars—become secondary. A camera becomes a really good appliance for taking photos for Flickr, while a TV becomes a nice Flickr display that you don't have to log into every time, and a phone becomes a convenient way to take your Flickr pictures on the road.

Hardware becomes simultaneously more specialized and devalued as users see "through" each device to the service it represents.

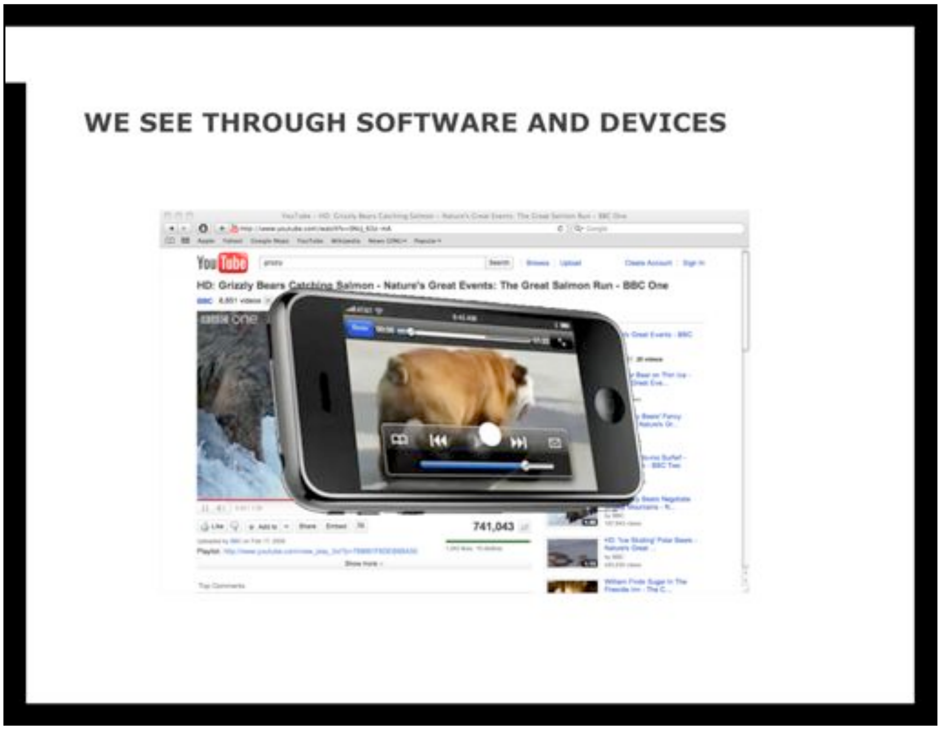

In effect now see through networked, service-dependent devices and software to the cloud-based services they represent. We no longer think of these services as being “online,” but services that we can access in a number of different ways, unified by brand identity and continuity of experience. This is a fundamental change in our relationship to both devices and software, since the 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.

This is how the local-to-remote axis links up with service design to create a new kind of user experience challenge, one that’s simultaneously about creating effective local experiences and integrated services across channels.

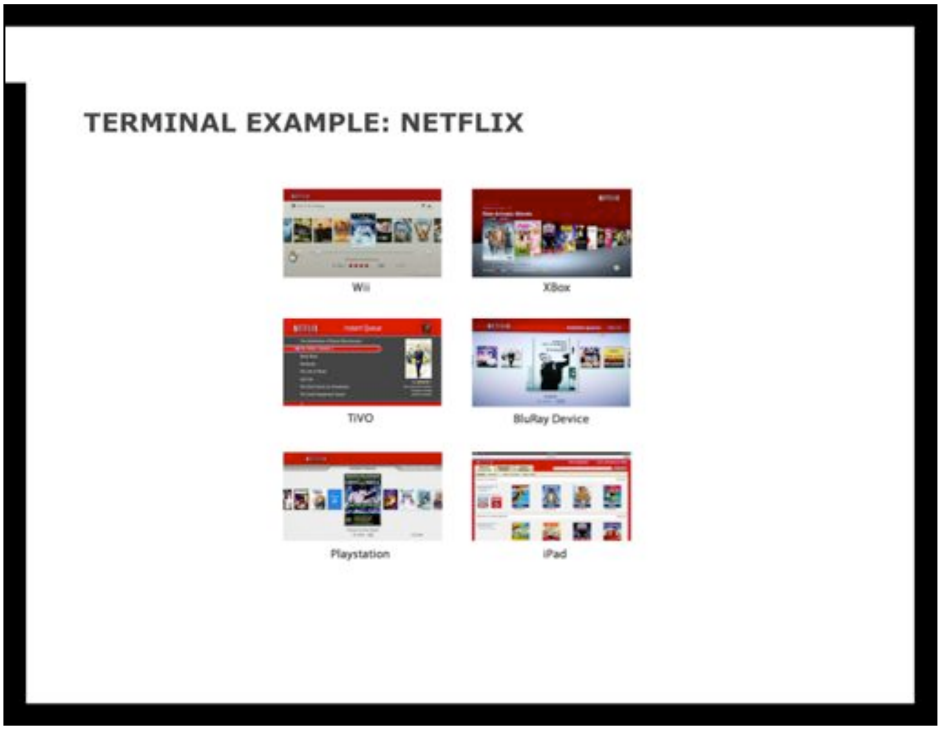

For example, you can now get Netflix on virtually any terminal that has a screen and a network connection. You can pause a Netflix movie on one terminal and then upause it on another. This may feel a bit novel, but it also seems natural. Why?

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.

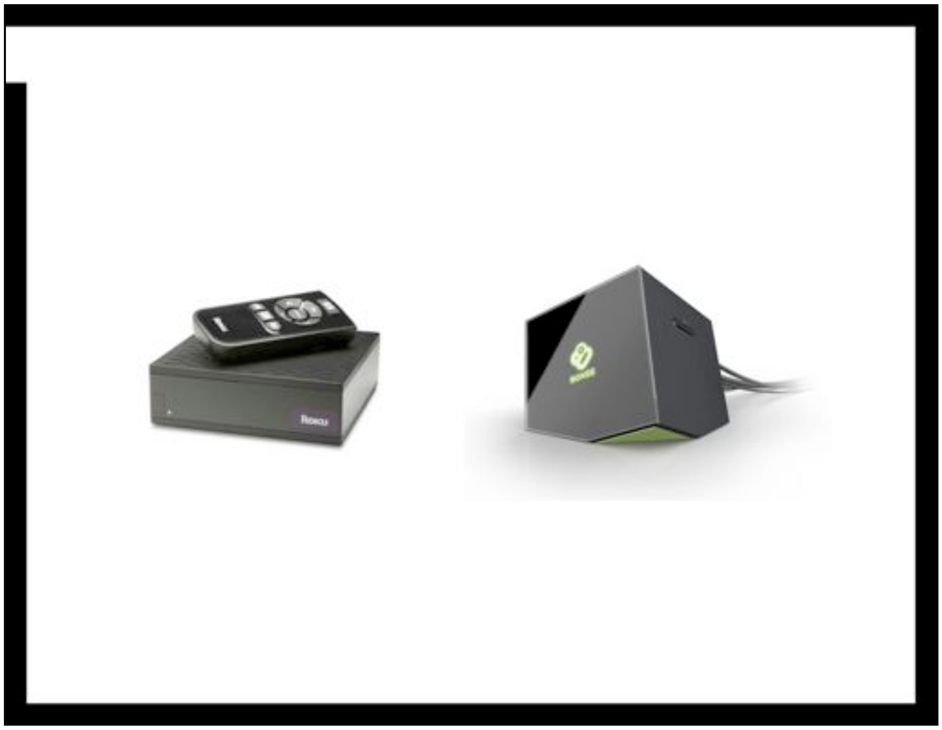

Netflix appliances are created for a single reason: to make it easier to access Netflix. That's what Roku does. It turns every terminal that's not already Netflix enabled into a Netflix terminal. The Boxee box does that for the Boxee service. The new Apple TV does it for iTunes.

Another example is the Kindle. Here's a telling ad from Amazon for the Kindle, another pure, and largely terminal-based 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." I really like the Kindle avatar experience, too. You can read on the phone on your way home, close the app, open it on your laptop and it picks up where you left off on the phone. You don't think of it as two separate things, but as one thing that exists in two places.

Let me give you another example. This is Vitality's Glowcap, which is a wireless network-connected pill bottle appliance that's an avatar to Vitality's service for increasing compliance to medicine prescriptions. When you close the cap, it sends a packet of information through a mobile phone-based base station to a central server and it starts counting down to when you next need to take your medicine. When it's time, it lights up the LED on the top of the bottle.

However, the real power is in the packet of data it sends. That packet opens a door to the full power of an Internet-based service. Now Vitality can create sophisticated experiences that transcend a single piece of software or a single device.

For example, another avatar of the Vitality service is an online progress report that can be used interactively or delivered by email. It's like Google Analytics for your medicine.

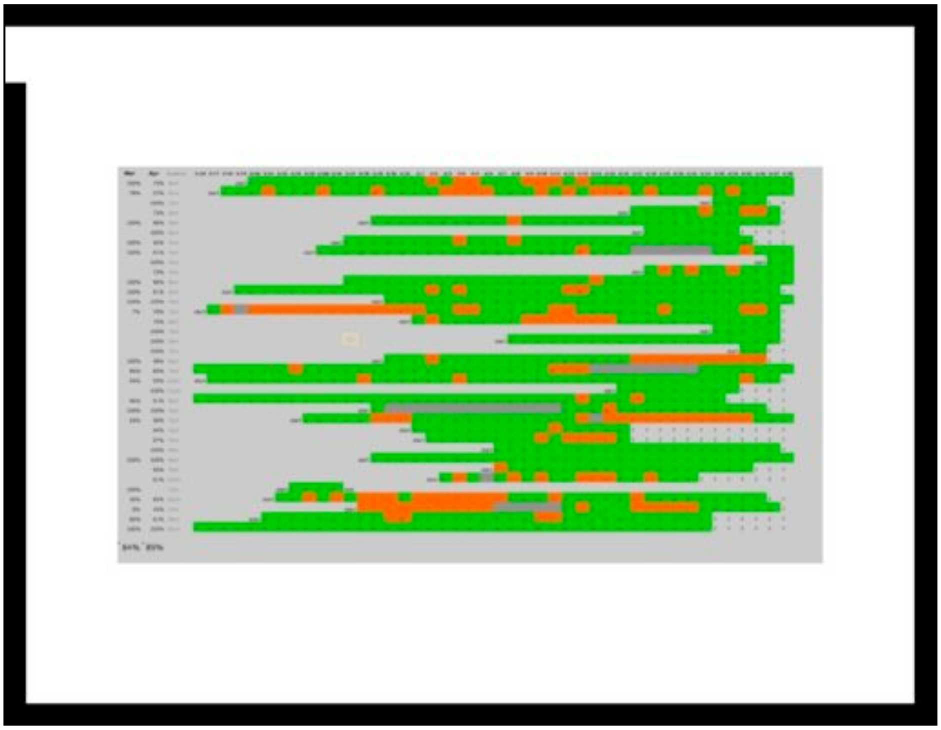

Health care practitioners get yet another avatar that gives them long-term and longitudinal analytics about compliance across medications and time.

To me, this kind of conversation between devices and net services is where the real power of The Internet of Things begins.

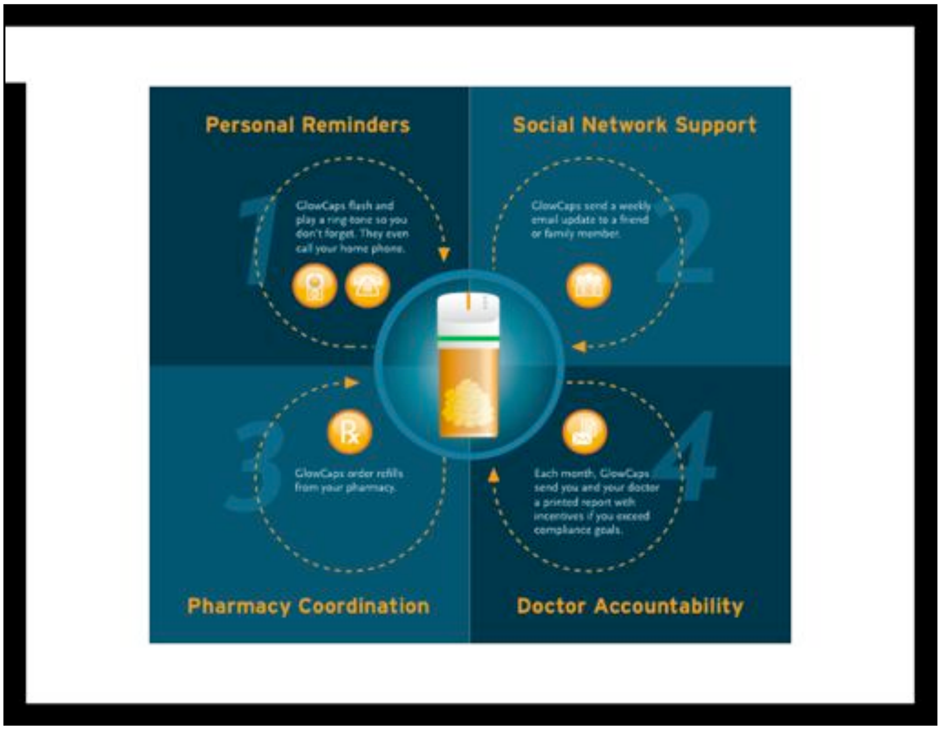

Vitality has developed a complete system around this service that includes a social component, and different avatars for patients, patients families, health care practitioners and pharmacies. Each avatar looks different and has different functionality, but they're perceived, and designed as a single system.

Another example. Nikeplus started as a service with a couple very simple avatars—the iPod, a shoe sensor and a Web site—now the service has morphed to encompass a wide variety of devices, use contexts and uses. They've even gamified the experience, so now you can play a game where you capture territory based on your exercise performance. Once the core value of the service was defined—in this case the automatic collection, analysis and sharing of physical fitness data—and a couple of core use cases were worked through, they could build and extend the platform in a relatively straightforward way into whatever they believed was an appropriate new use context.

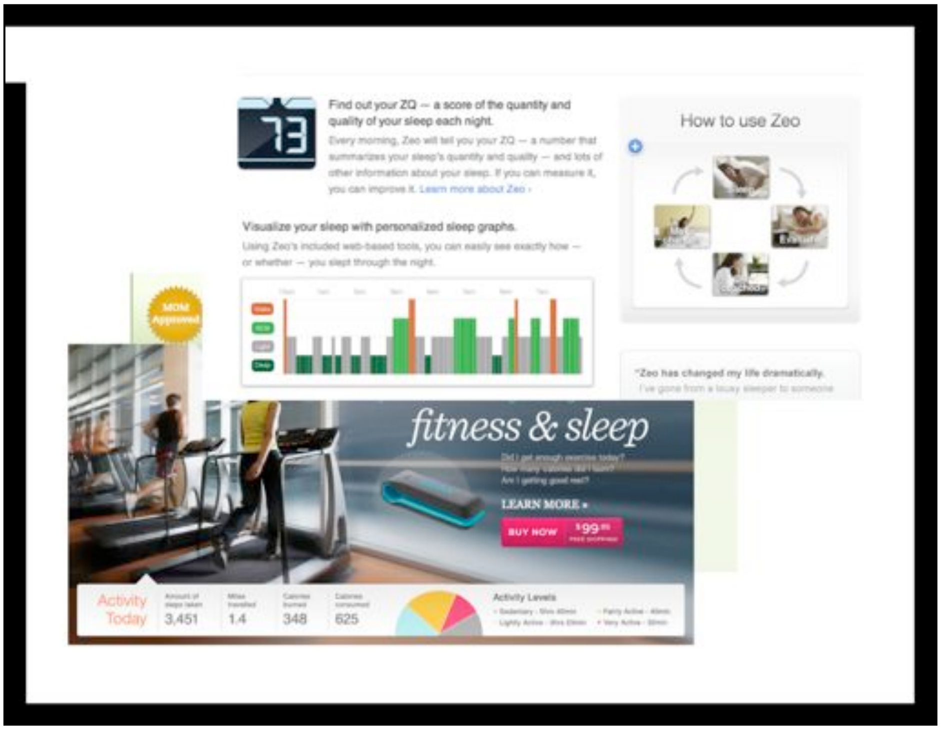

There are now many examples of services that have hardware avatars in the physical fitness and health space. There's the Withings connected scale, Green Goose's bike computer, Fitbit's pedometer and Zeo's sleep sensor. All of which depend on their online components to create their core value.

Let me change gears here and introduce a second major concept, machine readable digital identification and tracking, that I think is very important when thinking about how to design ubicomp networked objects.

Manufactured things have long had identifying marks, from silversmiths' hallmarks to barcodes.

These are the link between the object and information about the object and every object that has one exists simultaneously in the physical world and in the world of data.

Photo CC from http://www.flickr.com/photos/dumbledad/298650884/

I call this data the object’s information shadow. Until recently, accessing the information shadow was very difficult. The world of objects and the world of information shadows were separated by the difficulty of access. In a store, you didn’t know what the barcode meant, the store did, because only the store had the database and the hardware. And even they only knew a small part of what’s going on because a barcode only identifies the class of objects, not the individual object.

When Amazon extended ISBN to create their ASIN system they suddenly allowed anyone to reference any product Amazon sells or has ever sold. Tom Coates likened such codes as handles that we can use to grab information shadows and do interesting things with them, such as having conversations about them, getting more information about them. Amazon has built a large portion of their business around the fact that people point at their objects in a million ways, but at the core of that is always the ASIN.

The tipping point here is that we're about to enter a world where we can not only just point at objects, but have digital conversations with them by querying their information shadows.

For example, wine has a very rich information shadow. There’s a huge amount of structured information about where it was made, what it’s made of, how it was made, what critics think about it, etc. In addition, every bottle is a social object. There’s a community of collectors, aficionados, etc.

You can see some of the possibilities when you look at all the information about a single object in Amazon. That, however, is about a class of objects.

We now have the technology to uniquely see the information shadow of every object you’re looking at.

Each object is unified with its information shadow and you can query it. You can now know about where it is made, is it a real Gucci, what it is made of, what your friends think of it, how much it sells for on Ebay, how to cook it, how to fix it, how to recycle it, whether it will go with your mother’s drapes, whatever. Any information that’s available about an object can now be available immediately.

Source: Yottamark

Until the recent past, there was a fairly clear distinction between an object, a digital representation of that object and the metadata about that object. Now that distinction has sufficiently blurred so that there is a range of objects that exist to varying degrees as information shadows.

Some things have dematerialized almost completely.

When was the last time you thought of a plane ticket as a physical thing?

You obviously can't dematerialize a cantaloupe like that, but the blurring of physical and virtual objects caused by access to information shadows is transforming the world.

This is the beginning of mashups between the physical world and the data world.

The previous examples were of relatively static ways of looking at information shadows: a unique ID, whether it's a QR code or an RFID creates a relatively straightforward local experience: one device, generally a terminal, reads an ID and then allows you to view or manipulate the information shadow that ID is applied to.

But the point of the first of the initial trends is that processing is becoming cheaper. Information shadows don't have to be static things and the objects that they're the shadows of don't have to be static. Conceivably, every object that has an information shadow can update it itself.

For example, you can check on the status of your Amazon order because hundreds of devices, hundreds of appliances, are being used to track nearly every single atom Amazon is responsible for. Right now they' re using barcodes. The FedEx Sensaware smart tag has a bunch of sensors, a GPS and the equivalent of a phone in it for sending data about where a package is and what conditions it' s traveling in. It is an appliance for updating the information shadow of the package it's attached to with a wide range of telemetry.

When you put any digital appliance together with a network and a cloud-based

This is from Green Goose, a sensor platform based here in San Francisco. They sell these stickers that are actually tiny computers with a wireless transmitter and a sensor pack. They create information shadows for things that don't have them already. You create the meaning for the sensors. These are available right now.

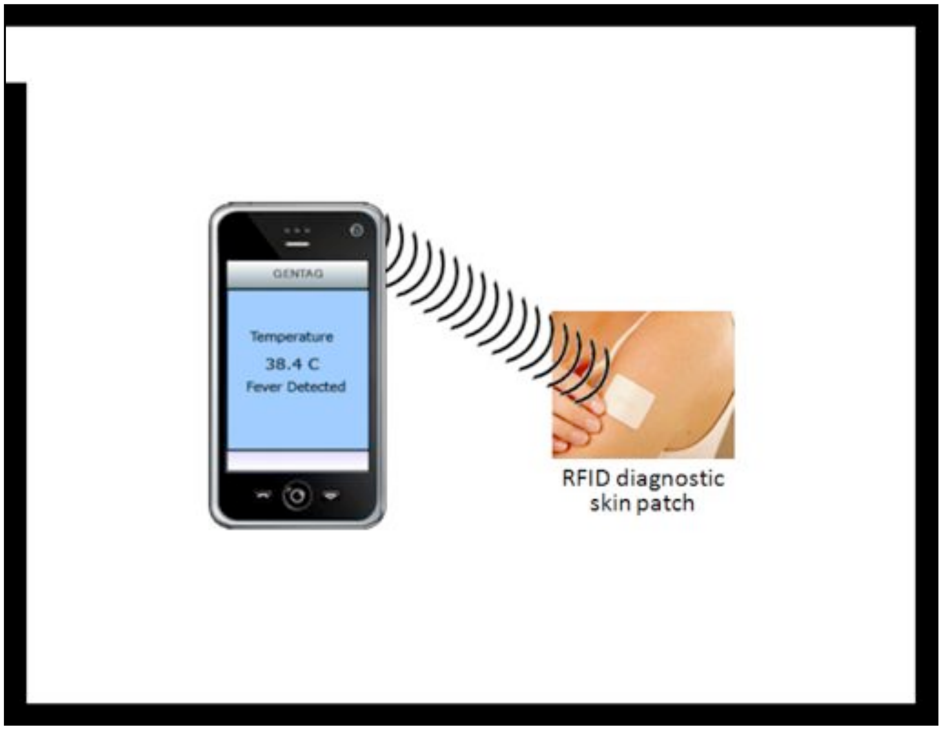

This is a stick on patch measures temperature and then transmits it to NFC phones. It came out earlier this year. The company, Gentag, claims they're developing patches that can test for pregnancy, the AIDS virus, drugs, allergens and certain types of cancers. In real time.

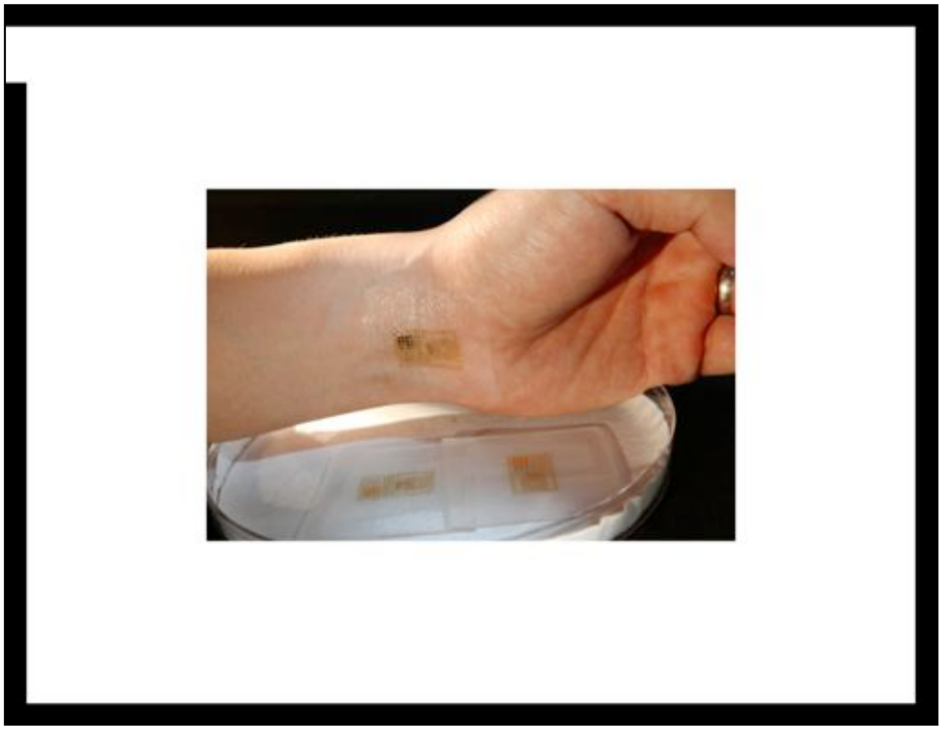

Here’s another one that was just announced by the University of Illinois that can monitor heart activity, brain waves, muscle activity, etc. Again, it transmits the telemetry wirelessly in real time to devices that then transmit it to the cloud.

Ok, so where do we get one of those cloud-based information shadow servers? Well, it just so happens that there are a number of them popping up right now.

This is Pachube, a free service that allows any arbitrary data stream from any net connected device to share that stream with any other device. It'll do the buffering, the protocol translation, the analytics, everything. You have it subscribe to an input stream like an RSS feed, or you subscribe to an output stream and off you go.

It's essentially a platform for creating mashups with physical objects, for connecting information shadows together.

It was used this spring to connect tiny personal digital radiation dosimeters all over Japan to measure radiation levels to a resolution inconceivable before. The service was put together within several days by Haiyan Zhang of IDEO and several other folks, essentially creating a mashup between Google Maps and a thousand different hardware devices. This points to the real power of the combination of device identification and pervasive networking.

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 can uniquely identify an object and attach it to an online service, you change the business model around the ownership of that object. It no longer has to be owned, but it can be an avatar of the service for as long as you're a subscriber to the service that it's an avatar of.

The old phone network is the classic example of this. People did not own their own phones in the US until 1984, when the old phone system was broken up. The phone was your avatar to the system. To set that kind of system up then was incredibly expensive, but now it's much more affordable. We are now seeing what's being called in some circles as the rise of the product service system, which is a system based on the delivery of value, rather than the sale of goods. Much of the product service system literature emphasizes sustainable manufacturing, but I think that's only a side effect to the notion of the dematerialization of everyday objects into service avatars.

Let me give you a couple of examples.

When you buy into a car sharing service such as City Carshare or Zip Car you subscribe to a service.

Each car is an avatar of its respected service, actively connected to the service at all times. You can only open the car and start the engine the service allows it, when the car has your permissions in its information shadow. The car logs whether it's been dropped off at the right location, and how far it's been driven. Your relationship with these cars becomes something different than with rentals and with ownership. It's like having your own car because you have access to it 24 hours a day, 7 days a week, with very little advance notice, but you can't leave your carseat in it, because it's not yours. It's different kind of relationship.

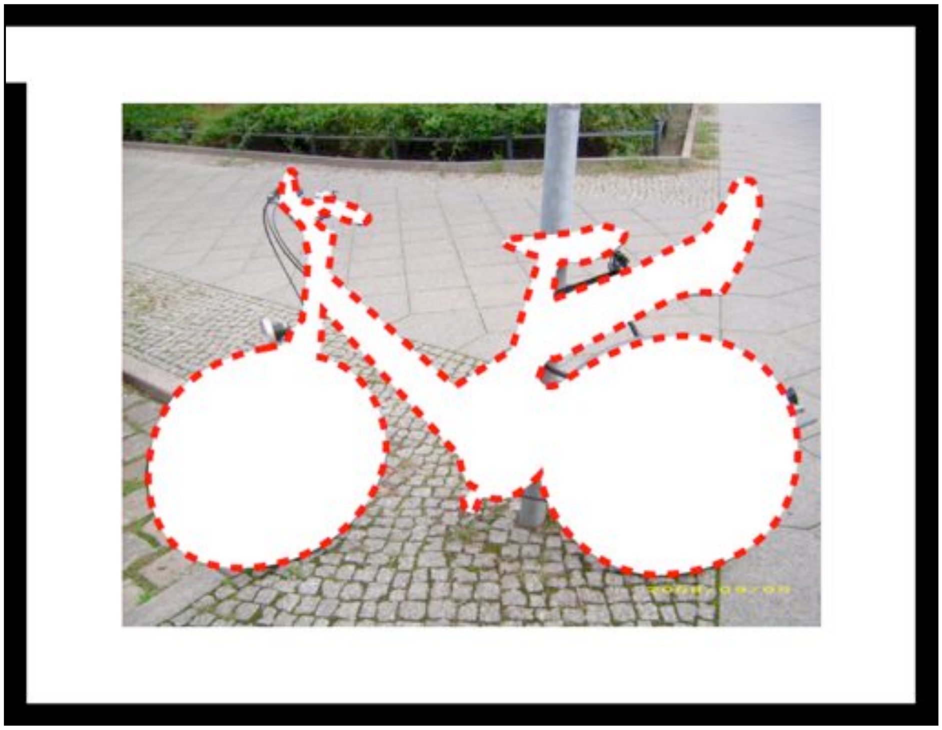

This is the German Call-a-Bike program, run by the rail service. You need a bike, you find one of these bikes, which are usually at major street corners. You use your mobile phone to call the number on the bike. It gives you a code that you punch in to unlock the bike lock. You ride the bike around and 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 bicycle service, its state maintained as part of its—and your phone's—information shadow.

See where I'm going?

Photo CC by probek, found on Flickr.

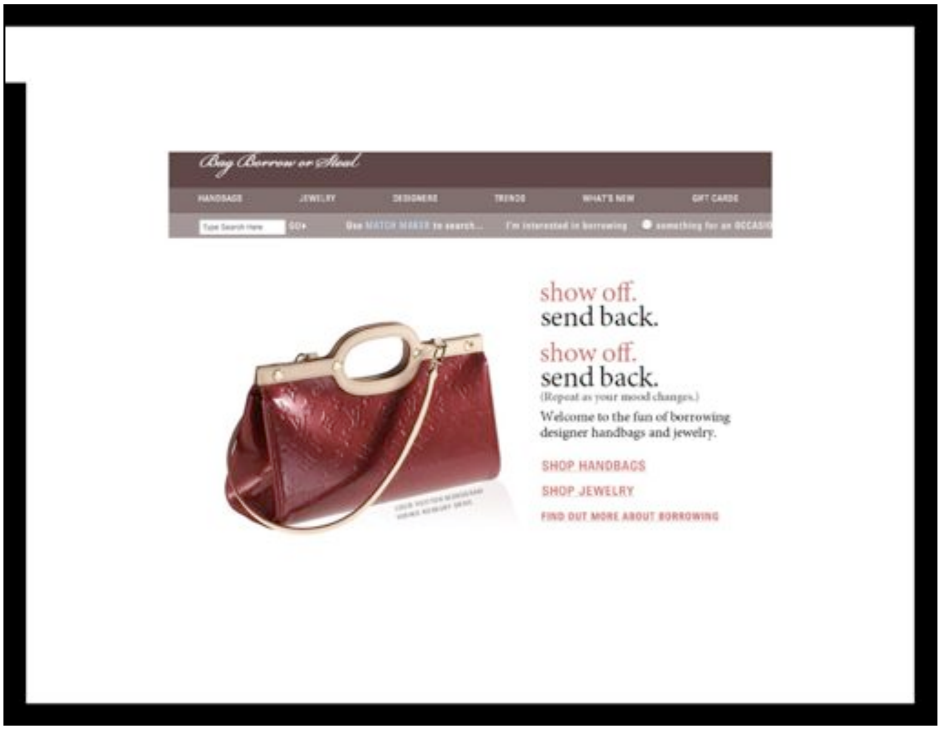

Here's another example that points to some exciting possibilities. Bag, Borrow or Steal is a designer purse subscription site. It works like Netflix, but for really expensive handbags.

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, btw, is also about five years old.

Photo CC by bs70, Flickr

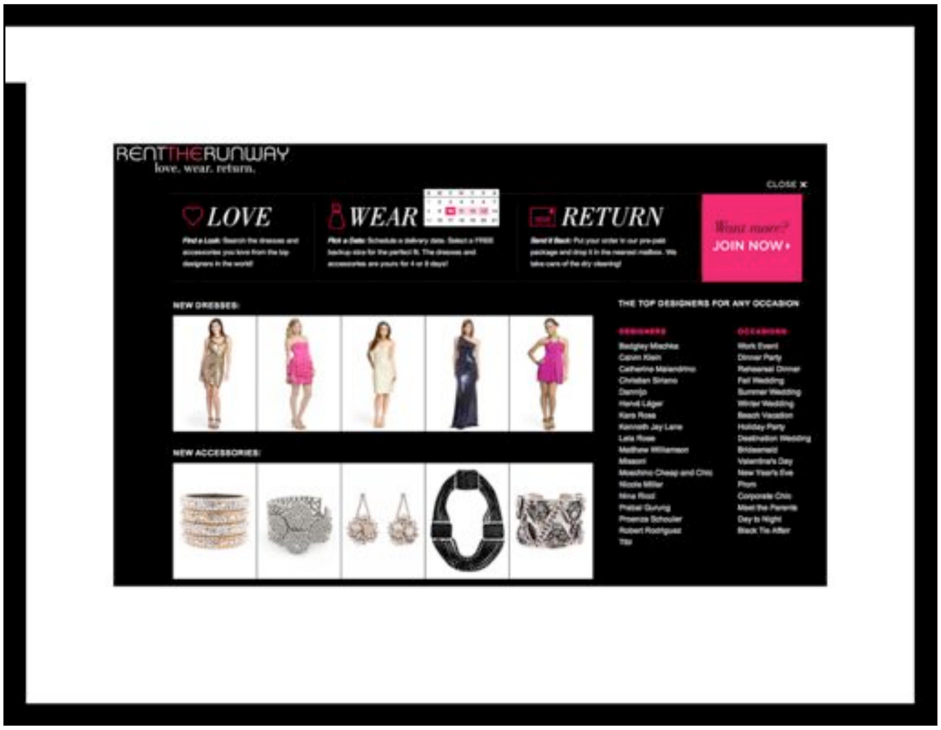

Here's another one called Rent the Runway that has expanded this idea to dresses and accessories.

How long until you get a subscription to the Gap 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 the Gap?

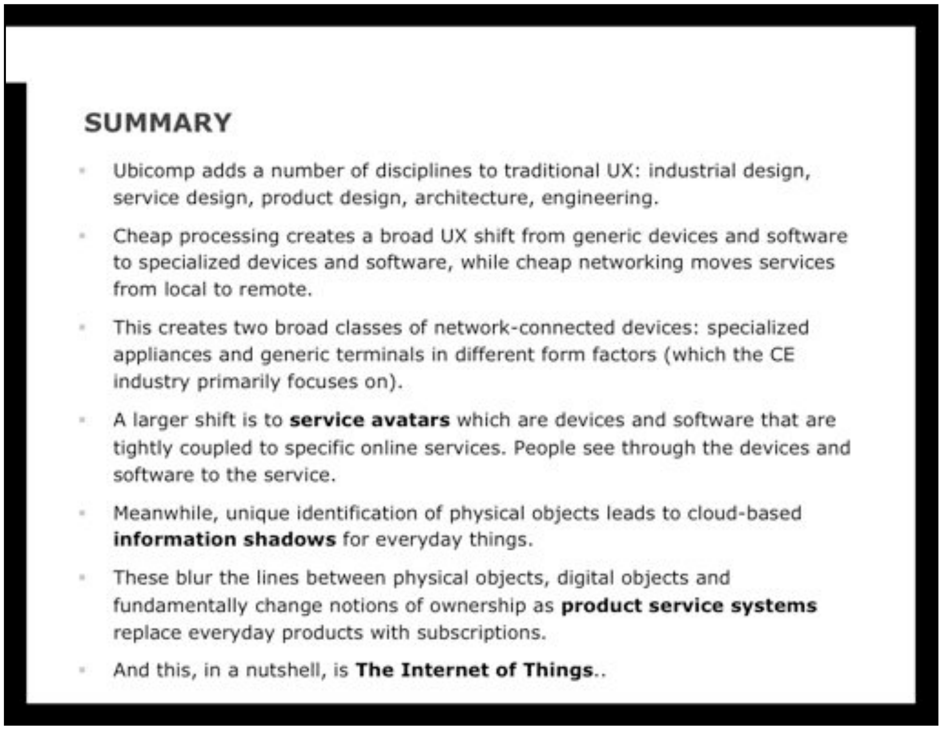

My goal in this review was to describe the general lay of the land in ubiquitous computing. We working in a complex environment with a number of interrelating factors, each of which represents both an opportunity for innovation and a challenge to the status quo in the design of consumer electronics. Things are moving fast and shifting our view of the world as they go. In effect, what we're seeing is the discovery of a new design material. Networked information processing is changing from being special thing that certain specialized devices do, to being a core building block, like plastic or aluminum, and a basic manufacturing process, like standardization in the creation of anything.

This is a huge and fundamental change, and we're just at the beginning of it. All of these other things are just symptoms of that one deep shift that we're going to see play out for the rest of our lives.

Q&A

15 MINUTE BREAK!

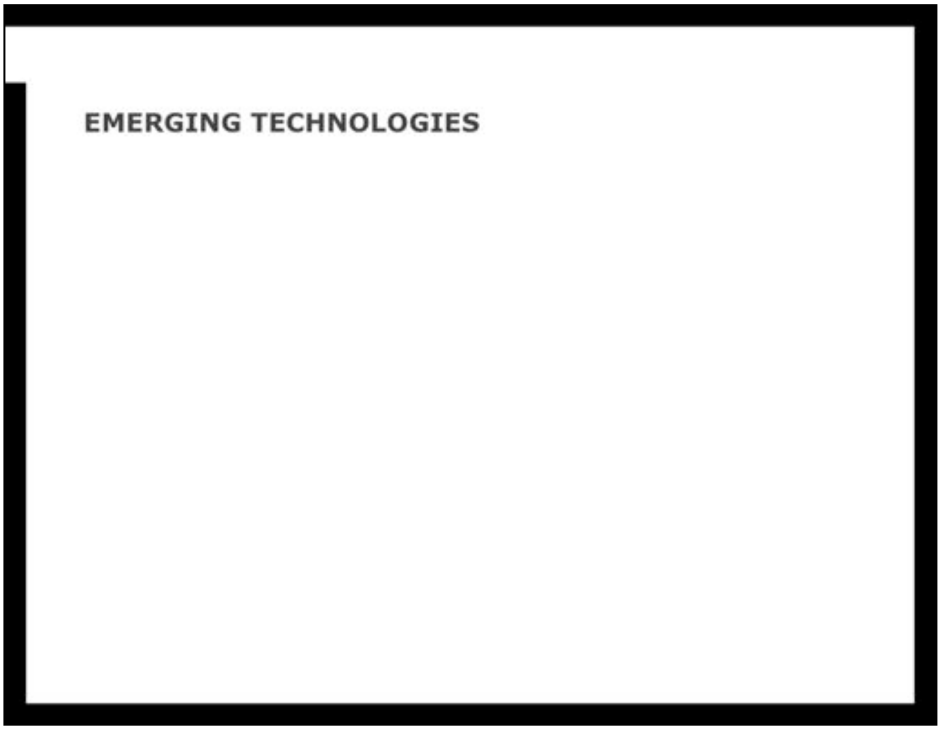

Now I'd like to do a tour of technologies that are maturing quickly into commercial-grade solutions and which have powerful capabilities. Now as we all know, there are a million new technologies appearing all the time and the vast majority of them go nowhere because the technology is not as easy to implement, or as powerful, as initially advertised. Many just have bad timing. There were a bunch of very clever things that CDROM makers did with their drives just before that medium became obsolete.

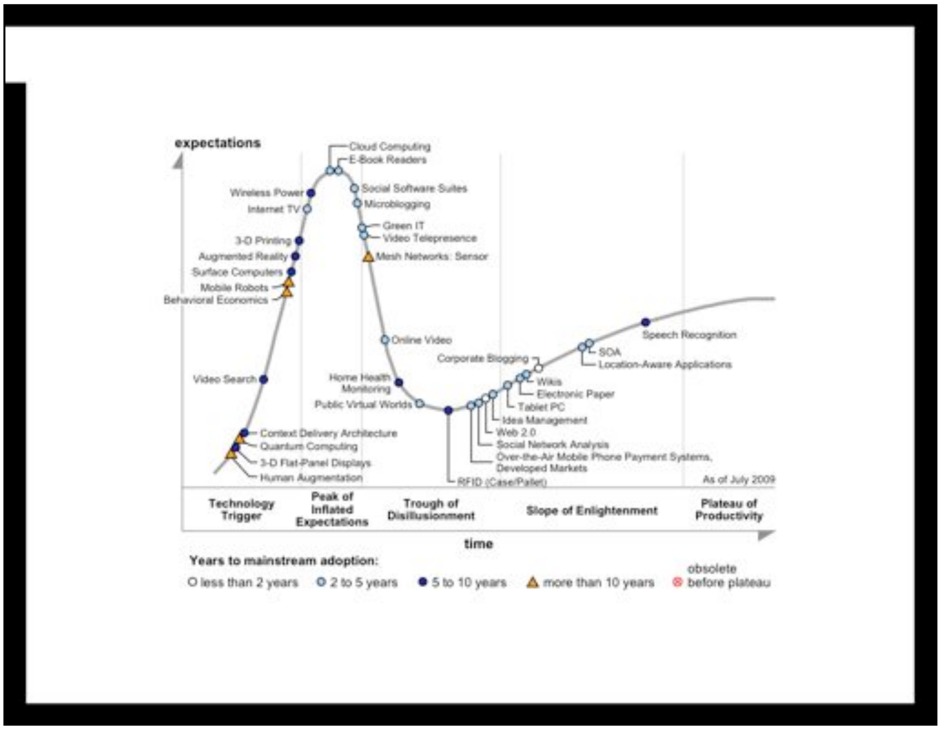

Gartner tracks these in the form of their hype cycle, and you can watch as various things go on it and fall off it. It describes the process technologies take in the public eye as a kind of hero's journey. Here's one from 2009.

The technologies I want to show you are in no particular order, but they're ones that either have made it past the Trough of Disillusionment or ones I believe will be able to make the jump to the Slope of Enlightment, but I wanted to give you a short tour of ubiquitous computing technologies that are either getting significant market penetration. You're probably familiar with some of these, but

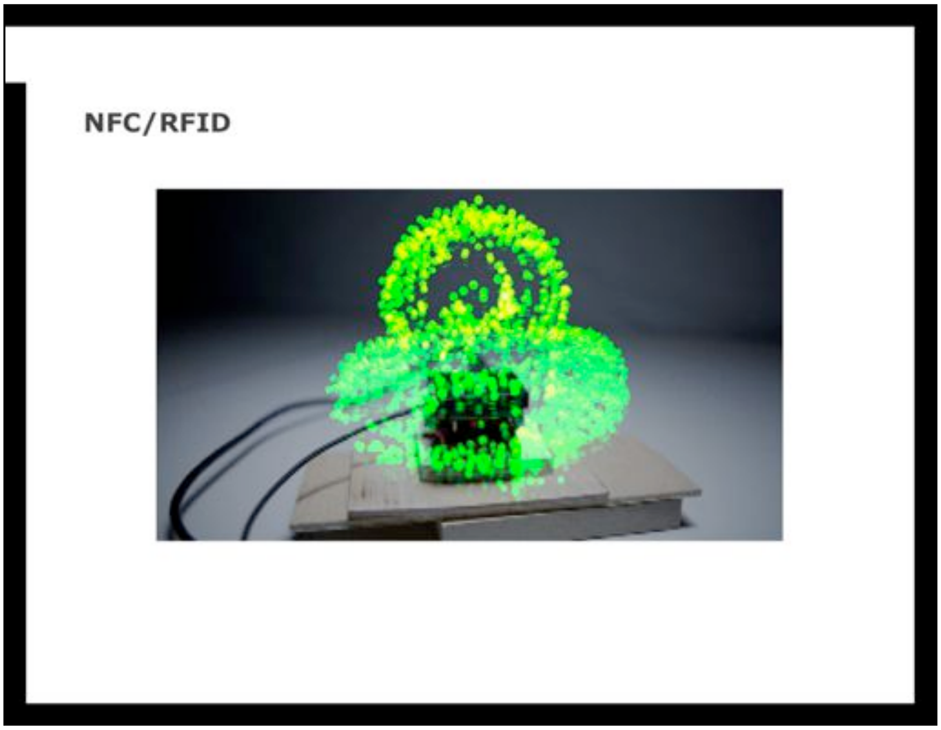

NFC, or near-field communication and RFID are related technologies. RFID is essentially a one-way version of NFC. It assumes a dumb device on one end and a smart one on the other. NFC assumes two smart devices.

You're all familiar with the basic idea of RFIDs, since they're the things in your building access cards and have been around for many years. The big plus of RFIDs is that they can inexpensively uniquely identify virtually any object so you can get at its information shadow. The downsides is that they can be replaced pretty easily with QR codes and other optically-readable and even cheaper identification schemes and, most importantly, that the range of action is very small. This is another image by Timo Arnall and the design firm BERG in London that shows the shape of the active reading area on an RFID reader. You can see it's on the order of inches, which is typical for the technology. NFC is about the same. The nice thing about NFC is that it allows touch-based interaction: you can exchange data by just touching two devices. That way you can use your phone, which is a trusted personal device, to introduce your TV to your camera or to replace your credit card. NFC is going to be built into a lot of phones in the next year with the intent of it becoming the basis of a new form of payment. We'll see, but it'll get the technology out there.

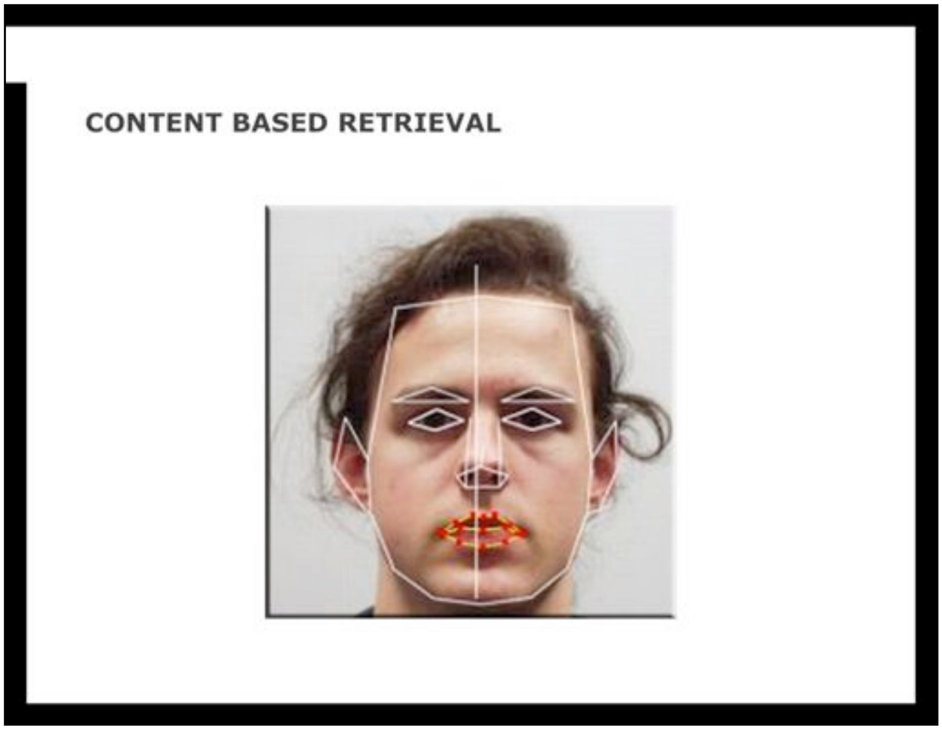

Fast processing has enabled the practical deployment of algorithms that can understand the content of non-textual digital data, so an image is not just seen as a bunch of pixels, but as a collection of meaningful objects. This is an incredibly hard problem that AI has been trying to solve for decades, but there's some real headway being made. Recognizing that there's a face in an image is now standard on most cameras, and there's a lot of progress being made in terms of recognizing whose face it is. Identifying brand logos is pretty standard, as is specific landmarks that appear in photos. The same kind of unique fingerprint extraction is happening in audio. Google has an audio API that is very good at deciphering what you're actually trying to say, or what song is playing, or what movie is playing, etc. We're still not yet at the point that the software can tell the difference between a fluffy white dog and a snowscape, which is trivial for people, but it's getting closer.

One particularly interesting application is automatic face beautification. If software can identify facial features and it has an idealized model of what people's faces look like it can adjusts photos and videos to make the face it recognizes match the idealized model without affecting the surrounding area. Thus, as far anyone but people who see you in person are concerned, you can always have beautiful skin and perfect features.

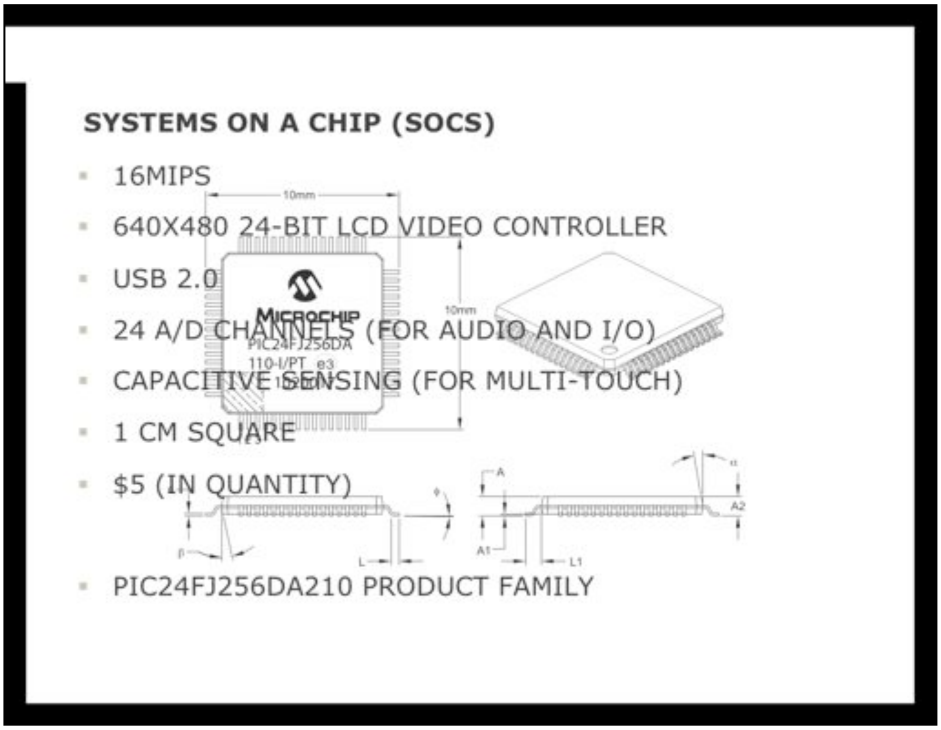

I talked a little about systems on a chip earlier. Putting the equivalent of a set of different kinds of processors on a single chip is relatively new, but robust and popular chipmaking philosophy. The microprocessor manufacturers have things called cores that are like object oriented chip descriptions. You want an ARM processor, Nvidia video, digital to analog conversion and a touch screen driver on a single piece of silicon, a chip fab will make you a single chip that has all of those components. You no longer have to go with what Intel or AMD will sell you. That's why Apple now has their own custom processors manufactured. The chip that's in your iPad or iPhone is not an Intel chip, it's made exclusively for Apple so that they can get exactly the functionality they want at low power…and incidentally control the user experience all the way down to the silicon, so iOS doesn't run on anyone else's processors.

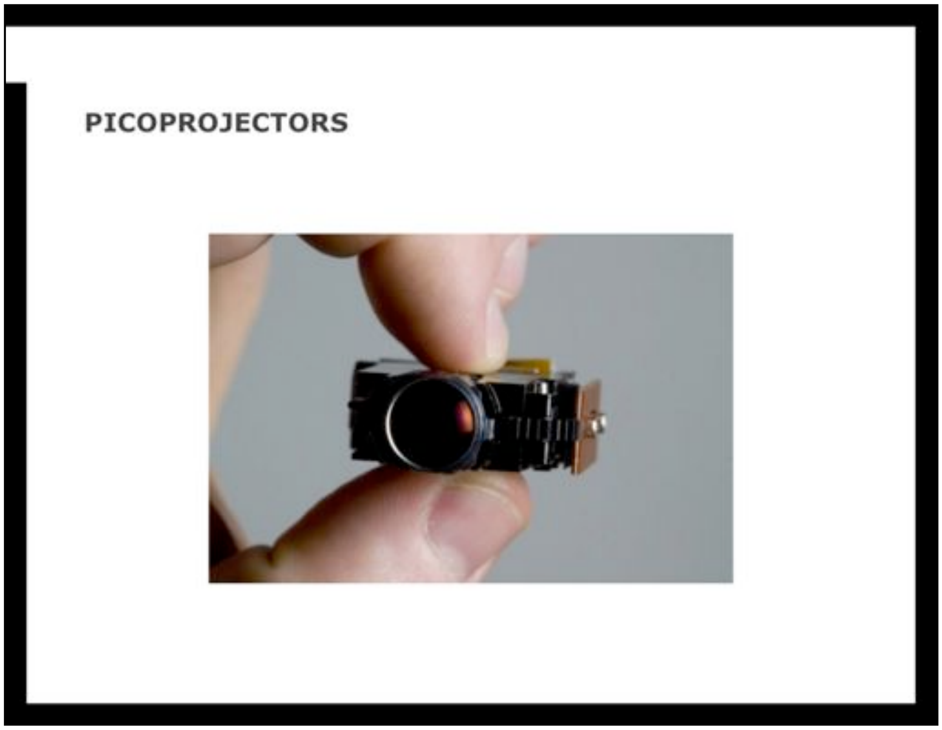

Tiny LED and MEMS-based projectors are coming out on a regular basis. You can now include a video projector in something the size of a sugar cube.

They still have heat and brightness issues, but that's improving every year. The great thing about these is that they mean that you can put small images everywhere. You can turn any surface into a display. Some even have motion tracking, so you can turn any surface into a multi-touch display. I've seen phones that have projectors built into them, but I think that the potential is much greater than that.

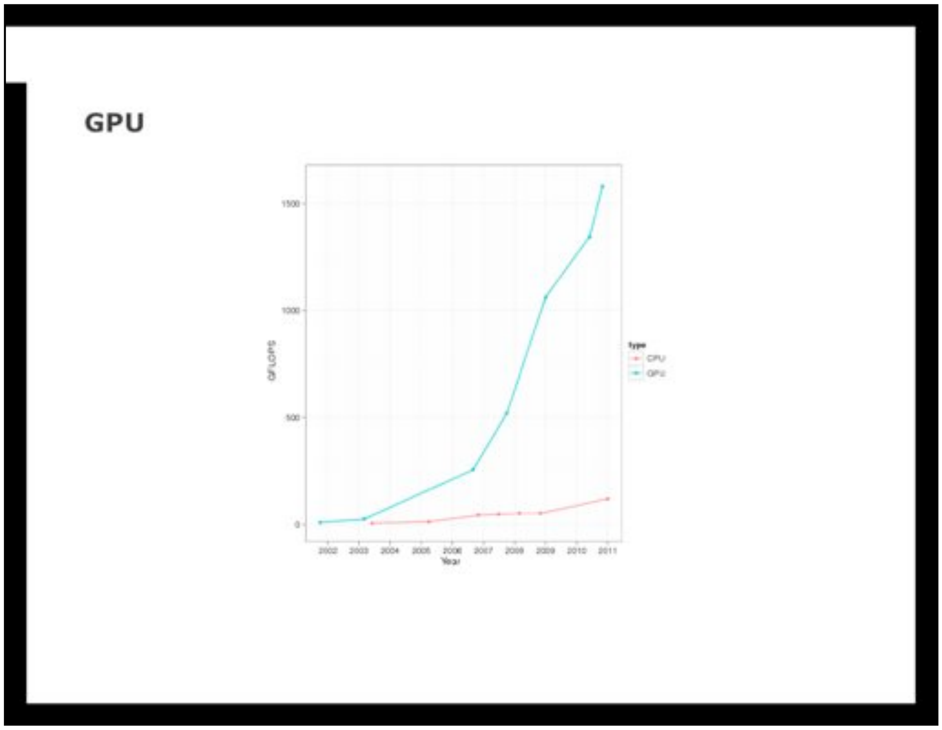

If you think that Moore's Law has made CPUs fast, you should check out GPUs, graphics processing units. That on the bottom is Moore's Law. In terms of raw processing power, in this case measured in gigaflops, they kill normal CPUs. The reason that we don't just use them for all processing is that they're designed as highly parallel processing machines. Writing parallel processing code is difficult and tasks that aren't easily parallelizable won't run faster on a GPU than on a CPU.

http://www.r-bloggers.com/cpu-and-gpu-trends-over-time/

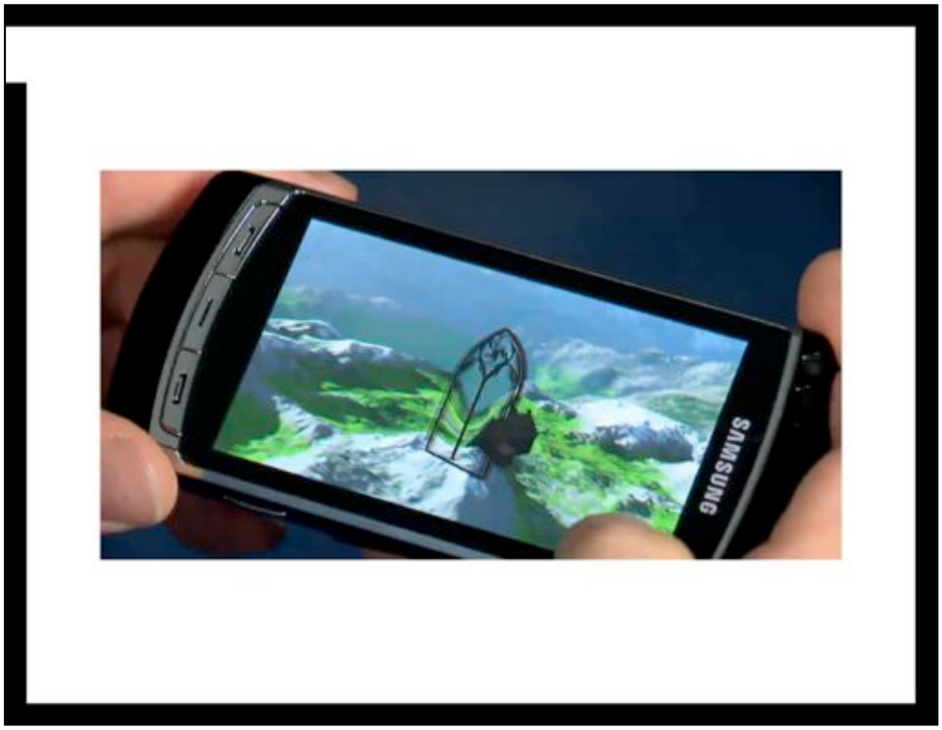

The upside is that it's possible to do all kinds of things with graphics, from layering generated scenes onto the existing world, which is what a lot of augmented reality applications do, to creating sophisticated visual effects for interfaces. Taking the Hidden Middle philosophy, this means that it's possible to take what were state of the art graphics five years ago and incorporate them for a small fraction of the cost and power consumption, and we can assume that this will continue in the future. This is also the technology that will allow a lot of the content retrieval techniques to be applied in a general sense, so that your phone—for example—will recognize who in all your photos appears multiple times and will offer to cluster such photos or videos together. Or it can recognize voice patterns and identify who is at a table talking. Or know the shape of the room you're in based on the echoes sound makes. That's what the Color application that got all the venture capital is aiming for, as I understand it. They're planning to collect a bunch of data about you, images, sound, location and use these algorithms to figure out who was where when so as to create an automatically generated social graph.

Microsoft’s Kinect is a big hit, and rightly so, the technology is clever and the implementation is pretty great. The Wii, its predecessor did a great job, too. I doubt that we’ll be running around like dancing chickens to control our devices in a general way soon, but I do believe that gestural input is here to stay. We do not yet have a stable vocabulary of meaningful gestures or the technology to recognize them under all circumstances, but I think there’s great potential there. Apple is very carefully and slowly introducing multitouch into its operating systems and I think that’s a good approach. I think that gestural input, or a kind of virtual direct onscreen manipulation, is a great way to interact with 10-foot interaction, while accelerometer-detected gestures with objects are a good way to give individual small objects input.

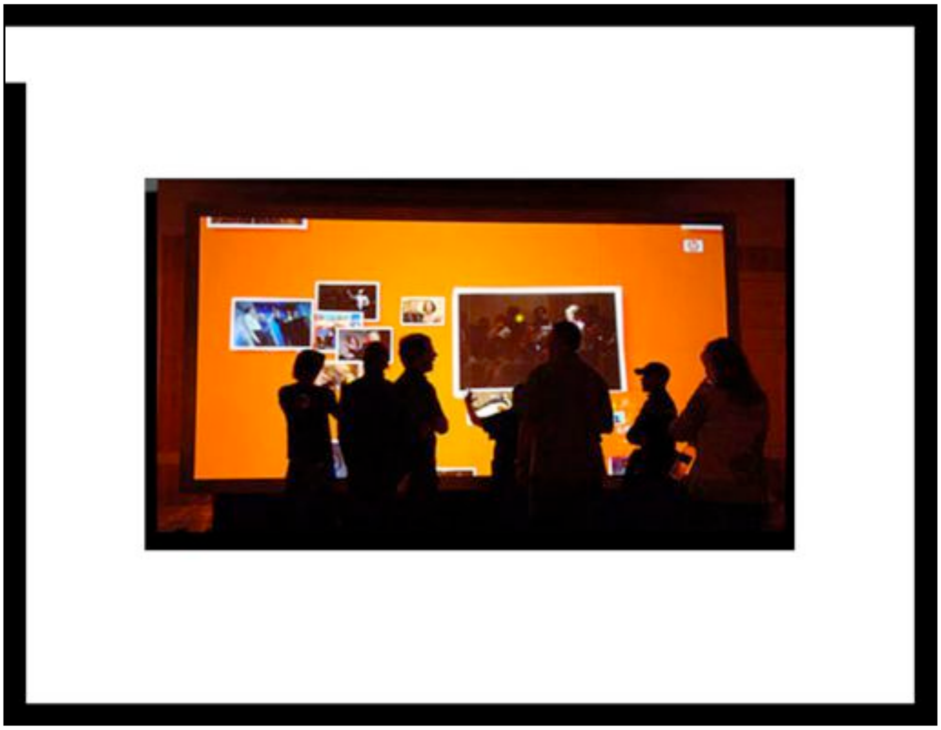

I also really like the multi-person experience of large multi-touch interactive surfaces, since they allow for the use of many small screen avatars at arbitrary sizes, and I think there are interesting possibilities when this technology is mixed with picoprojectors.

This is Stimulant Design’s wall-size multitouch display for HP.

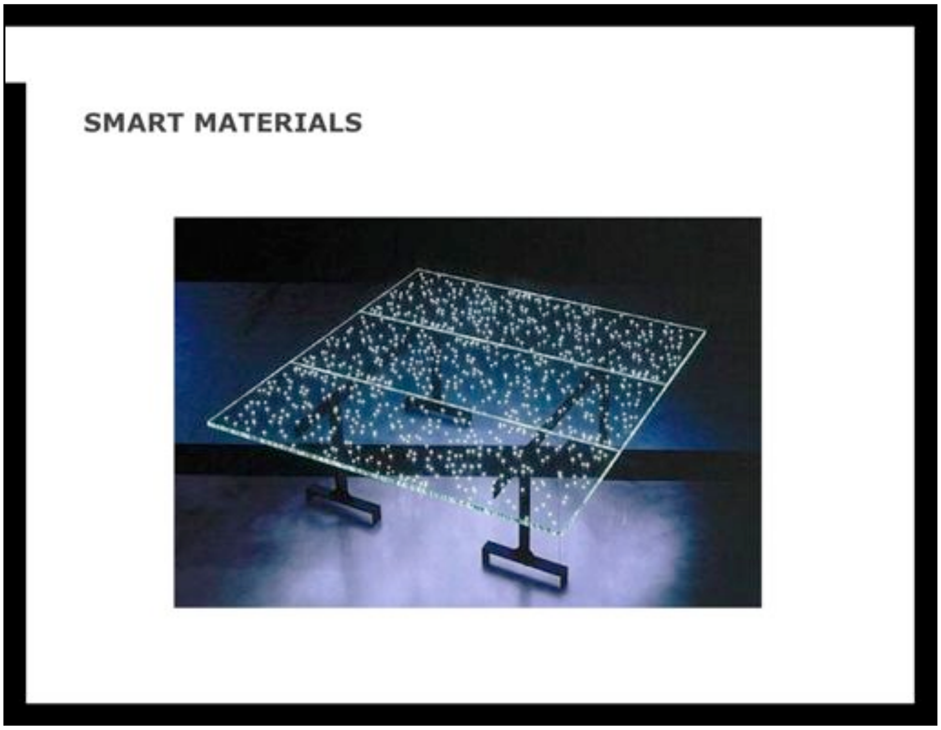

Increasingly, digital reactivity is being included in special materials. This is probably the most speculative technology here, since there are actually many technologies that fall under the umbrella of smart materials. These range from wall panels with embedded LEDs to paper that changes color based on how electricity is applied, to shape memory alloys that change shape based on how much heat is applied, to ceramic sandwich flooring that generates electricity as you walk on it. This is a table by lighting designer Ingo Maurer.

This is a luminescent fabric from lumigram.

From http://www.lumigram.com/

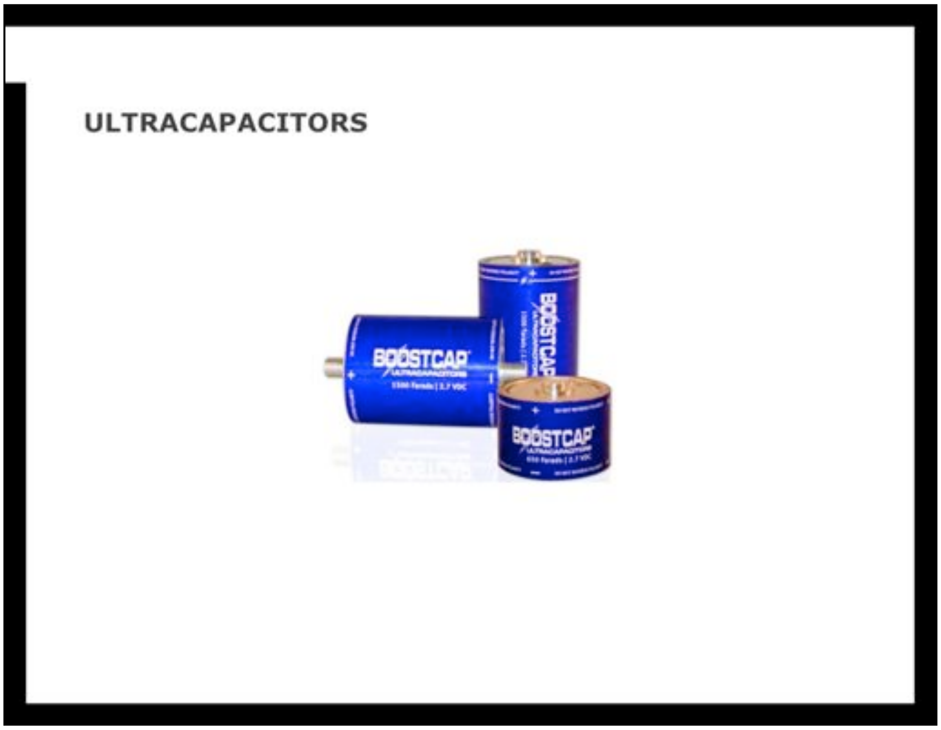

Since batteries are so inefficient, people have been looking at other ways to store energy. One is ultracapacitors, which use an entirely different method to store energy than do batteries. They recharge very quickly, on the order of seconds rather than hours, they can be coupled with an energy harvesting system that takes vibration energy or heat energy and trickles it into the battery, and they can store about half the amount of charge a conventional battery of the same size can store. The problem is that they're still really expensive and the price doesn't seem to be falling very quickly.

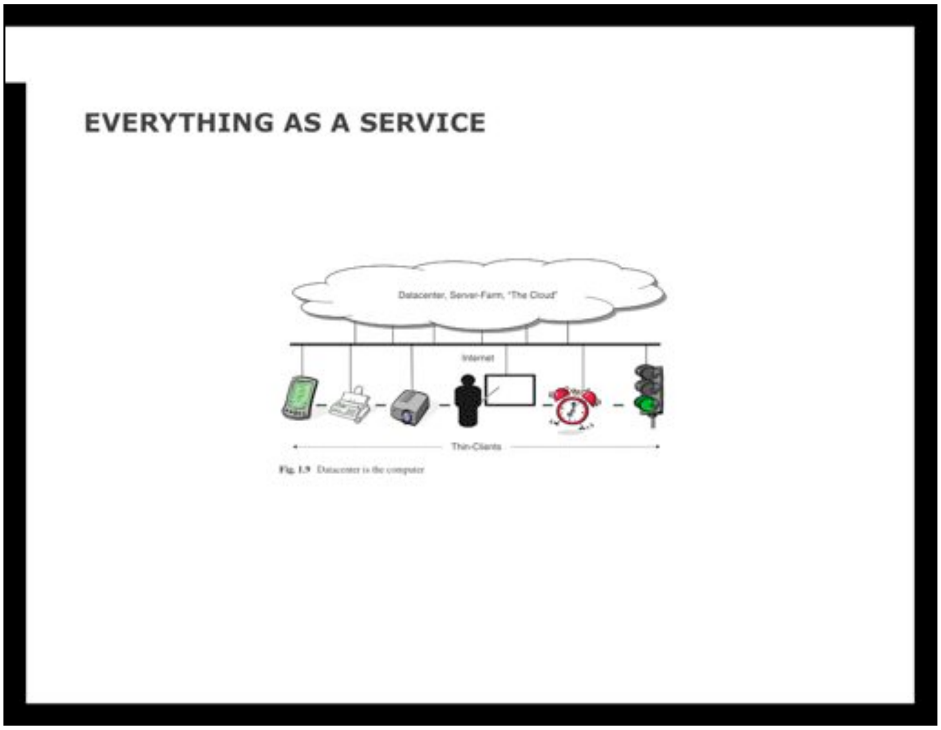

Finally we have cloud-based services, which make virtually any kind of processing you can do with a computer locally you available remotely, usually orders of magnitude faster and more efficiently than you could with a local device. This is what has enabled the explosion of so many startups in the last couple of years, because it means that many hard problems are solved by just paying someone per transaction to solve them. This, more than just connectivity between arbitrary devices, is what enables ubiquitous computing to have crossed the tipping point to viability. It certainly still has its problems—the interdependence of so many services means that you're entrusting a lot of the back end powering your user experience to people who you don't personally know—but it certainly enables rapid deployment and iteration of ideas.

Image Source: “Cloud Connectivity and Embedded Sensor Networks”

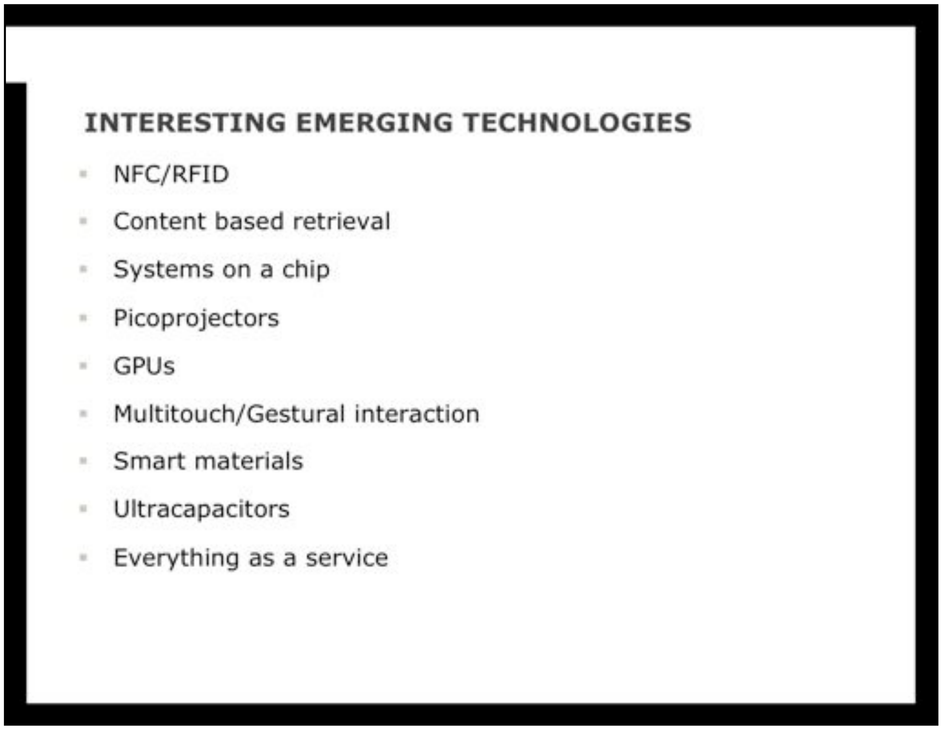

This is a bit of a grab bag of ideas, but I think it’s a good toolkit to start thinking about how these technologies can be used in designs to create profoundly new experiences.

Finally I want to list a number of application domains where there's a lot of interesting work going. Again, this list is idiosyncratic to my perspective, but these are the domains where I think some of the most interesting work is happening.

Wearable computing is the idea of using clothing or jewelry as a part of your computing ensemble. In sports people regularly wear specialized clothing or devices that collect data about their performance.

These are Ruth Kikin-Gil's buddy beads.

This is one of my favorite ubicomp products from the toy world. It’s called Clickables and it’s a product from a Hong Kong company called TechnoSource. It’s part of Disney’s Fairies initiative.

Source: Disney Clickables

Here's one of the ways that works: when two kids put their Clickables bracelets together, their avatars link up in Pixie Hollow, the online social network associated with the Fairies brand. This bridges the physical world of kids with their social network in a transparent and familiar way. All of the products in this line have such an online-offline existence. Another example: when you get one of the charm bracelets and you touch the charms to the USB-connected jewelry box, your fairy avatar gets a version of the same charm.

I think of appliances as a specialized kind of furniture, and the appliance business has been working hard, although not particularly successfully, at creating ubiquitous computing devices.

This is a series of Internet connected appliances by Salton, the people who brought you the George Foreman Grill. It was an experiment from about five years ago that points to some interesting ideas, but never quite got the UX right. For example, the microwave has a barcode scanner built in: when you scan some food, it goes out to their server, gets the cooking instructions and programs itself. That’s nice, but how hard is it to read the back of a box and type in one number?

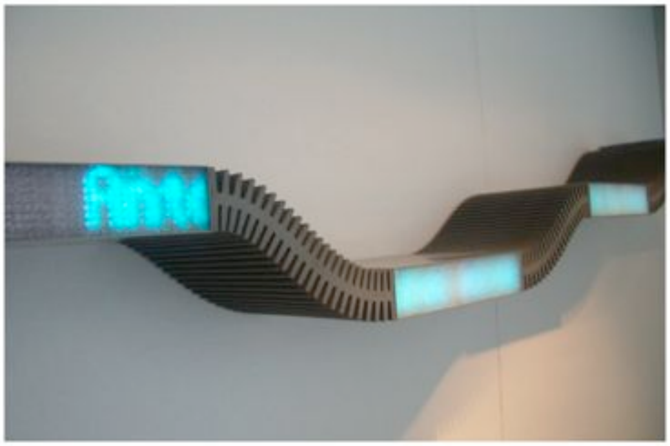

I think that what's more interesting is expanding the notion of furniture and our understanding of appliances by incorporating digital technology into it. This it Jean-Louis Frechin's bookshelf that's also a mirror of your text message feed.

Speaking of furniture, cars are a kind of room that moves around that's full of technologically augmented furniture. People expect them to have lots of technology in them, and I think that auto companies—who have been effectively making ubicomp devices for twenty years—are finally starting to figure that out.

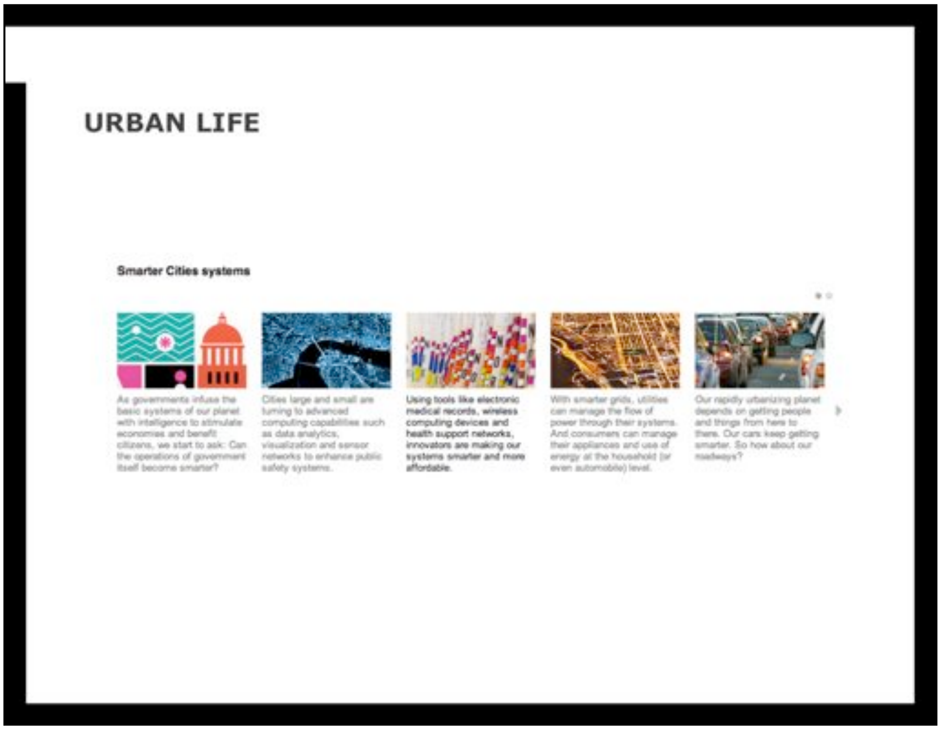

Cars and mobility and the relationship people have to cities is also a large area of development.

IBM's Smarter Cities initiative is a ubicomp initiative. They're treating cities as a mesh of different kinds of networks, social, infrastructural, financial and they're looking for ways that they can inject technology into those networks at a massive scale to create what are essentially new kinds of utility services based on ubiquious computing.

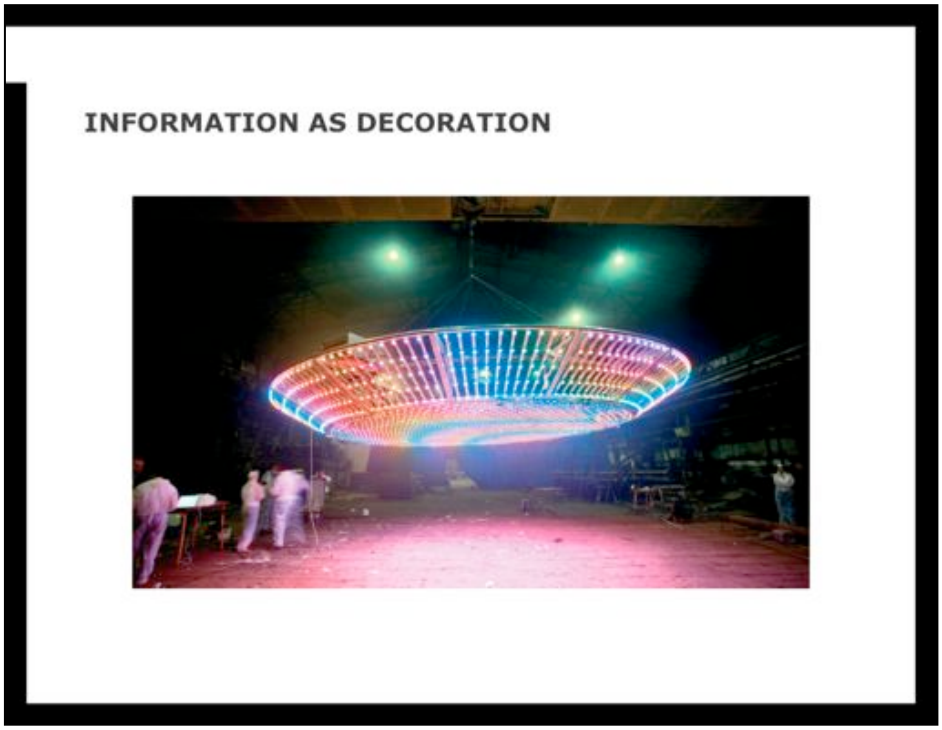

One byproduct of all of this information display embedded in the environment is the treatment of information display as an esthetic experience. I think this is going to go further and we're just going to use electronics as pure decoration, as has happened to pretty much every material before it.

You can call Vegasification, but I think it's pretty exciting that the surfaces of our world can begin to shimmer, move and react just to be beautiful.

Image: UFO by Cinimod Studio and Peter Coffin, 2009

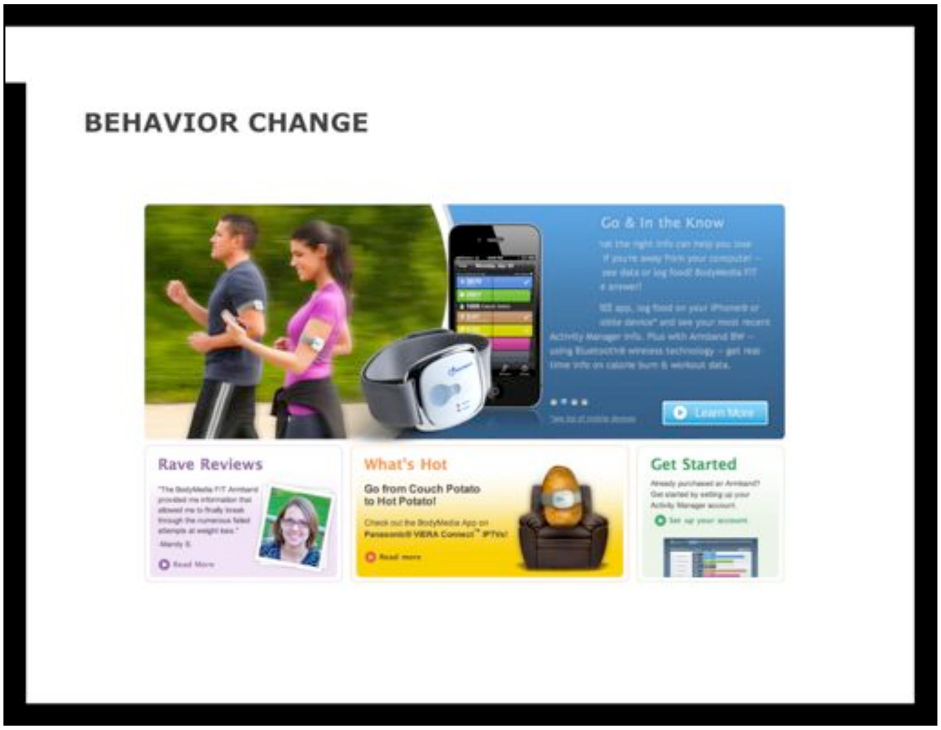

Moving from things to people, one of the major current uses of ubicomp is in technologies for behavior change. These come in many flavors, from Green Goose's original model, which converted bike miles you ride to dollars that you would have spent driving your car, to a lot of health-related products.

Here's Bodymedia. They tell you that the technology can help you go from couch potato to hot potato. How is it going to do that? Well, it's a combination of quantified self tracking and gamification, where they create a set of game-like extrinsic rewards based on automated sensing of body state.

. [Do I need to explain QS and gamification?] There are other more subtle interventions.

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 computation 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.

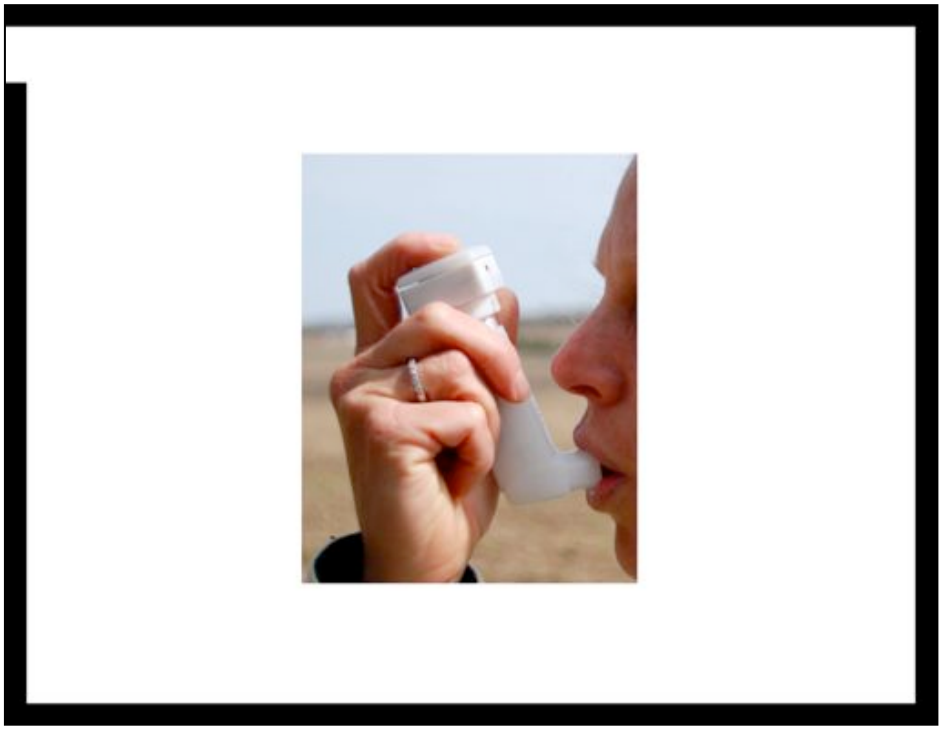

The Asthmopolis project uses GPS to track where people used their inhaler. This is designed to give the user some information about when and where they typically have asthma, so that they can change their behavior, while at the same time producing heat maps of high asthma areas for health care scientists. A lot of personal health technology is essentially behavior change technology, and the same kinds of ideas apply to it as to physical fitness, or eco consciousness.

I wanted to finish with the ubicomp initiative that probably has had the most investment in it in the last five years and which you're probably most familiar with. This is, broadly speaking, the push to use ubicomp technologies to provide access to infrastructure services. This means paying for things, opening doors, tracking utility usage, etc. This is already a pretty embedded part of our world, and is getting increasingly so.

That's an early Nokia NFC payment image, the Bay Area clipper card and a zigbee smart meter house visualization device.

It also points to how these technologies insinuate themselves into life. They usually don't come with a big splashy new device that has a UX that much be mastered to get at the awesome new functionality, but in the form of an incremental digitization of everyday things that creeps along until, suddenly, mint.com can give you instant access to the information shadow of all of your financial transactions at a level that you didn't know was possible, but which was happening all along.

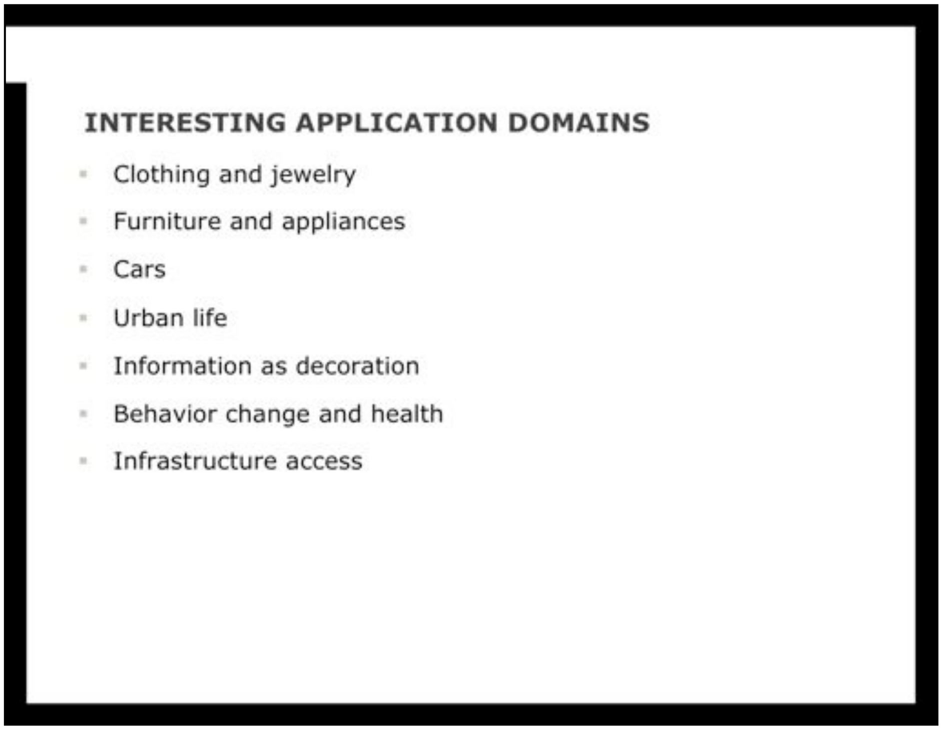

This list, and this presentation in general, is absolutely not exhaustive. I wanted to give you all an overview of both the possibilities, the uses and the challenges in the technology and the broaden the focus a bit beyond classic consumer electronics paths. As we go through the ideation and vision definition process, we may well end up in the expected places because of other constraints, but I want to encourage all of you to think about not just the products that we're making or going to be making, but the ecosystem of devices and services that all of these products exist in. The more we can design something that fits into a larger ecosystem, and perhaps defines a novel and valuable new part of the ecosystem, the more successful we'll be.

DISCUSSION!

The internet refrigerator has been reinvented approximately 50 times in the last 15 years, but has yet to get ANY traction at all in the market. Why?

LUNCH!

DISCUSSION!

What made the iPod successful? Unlike smart fridges, MP3 players were a known success, but none of them became the runaway hit like the iPod?

So how do we invent artifacts from the near future? How do we mitigate the risks so that we're more likely to make iPods than internet fridges? There's no obvious way to do that successfully, of course, and this whole field is so new that there are no best practices. This afternoon is an opportunity for all of us to explore some of these ideas together and see what happens.

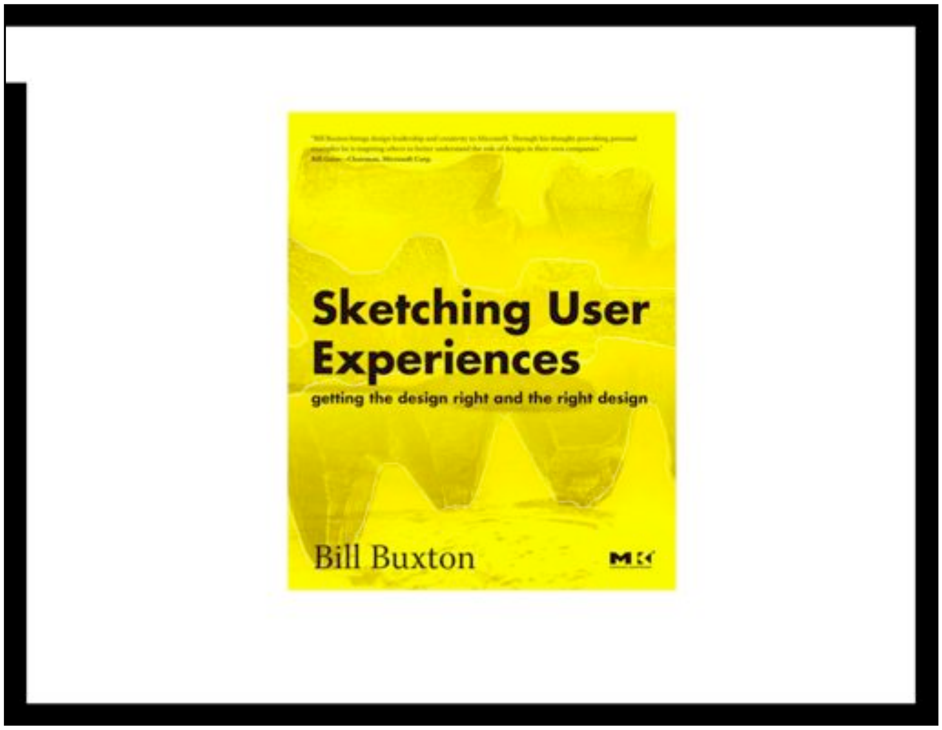

First, let me plug a book that you may have already read. It's Bill Buxton's "Sketching User Experiences." If you only read one UX book this year, this is the one. It's a great description of a way of thinking about how to create novel user experiences.

It starts with ideation. I'm not going to give you instructions about how to ideate, but the idea is to have lots of ideas. This is Martino Gamper's "100 chairs in 100 days" project, where he made 100 chairs from parts he found on the street, one per day.

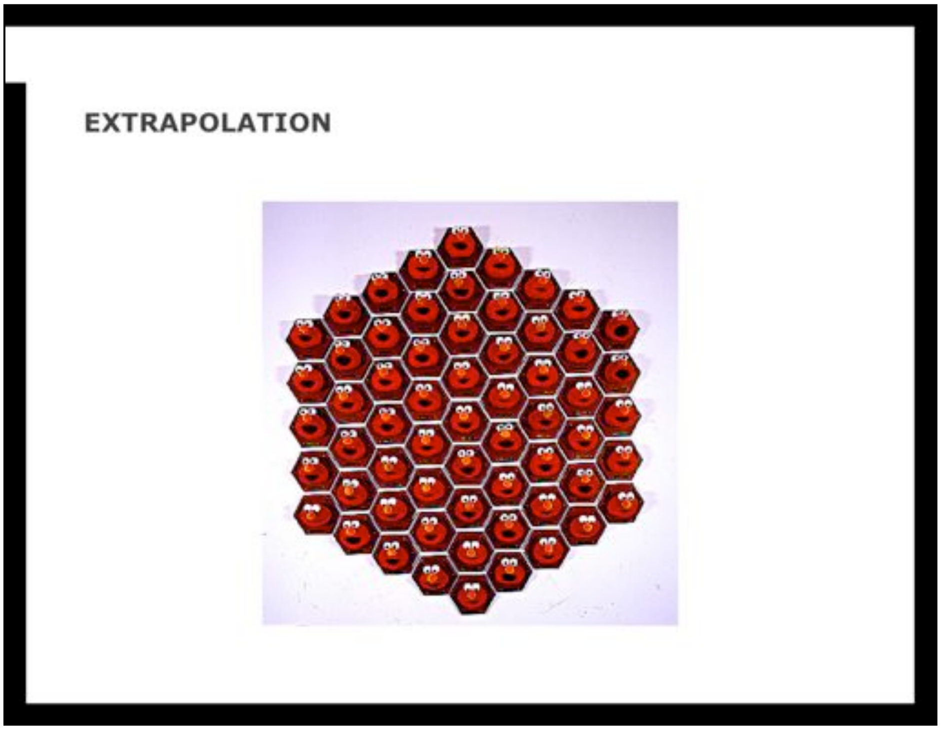

One technique is extrapolation. That's where you take a piece of data that you have and project it to get something new.

One effective form of extrapolation is to multiply something by 10 or 100. That's what Weiser was doing in 1988. What if we took this computer that's on our desk and said that we're going to have the equivalent of 10 of them. How will that change our lives? What if that camera becomes 1/100 as cheap? What if we have 1000 computers embedded in that wall?

Another extrapolation technique is a projection across demographics and time. For example, the folks whose 20s were spent on Facebook and in World of Warcraft are probably going to have a different expectations around personal information and narrative than people from earlier generations. Let's map the attitudes and behaviors of kids on 4chan to the office environment of 2020. What does that look like? We don't know that for sure, but when doing ideation, we can assume that some mapping will occur and use that as the basis for identifying problems people may be experiencing that can be solved with technology.

Image: N-Trophy, 2000-2003, Kelly Heaton, Feldman Gallery: http://

A second technique is thinking about different scales. Computers traditionally have interfaces that are person-scale, but there's no reason that has to be the case. At PARC under Weiser they defined the tab, pad and board as names for the scales of the devices they were developing. The iPad is an homage to that.

From Flickr: watch by funadium, box by ubermichael, phone booth by rastrus, room by bigpinkcookie

Scale	Label	Examples
1 cm	covert	RFID, nail polish, cochlear implant
10 cm	mobile	phone handset, portable media player, wallet
1 m	personal	chair, car, ATM, payphone, laptop
10 m	environmental	wall, door, chandelier
100 m	architectural	church clock, billboard, bus
1000 m	urban	street intersection, landmark, crowd

This is the scale I've been using. It's a set of definitions to talk about granularity and it helps us identify that works and doesn't work at various scales. Screens don't work when you approach the covert scale, which is why wrist TVs have never taken off. Buttons don't work well on the environmental scale and above, because they're too small relative to the object. You probably can't make anything that's designed to be immediately social at anything above the environmental level.

One of the biggest challenges in designing service avatars is moving an experience from from one avatar to another. We've probably all had that experience where a piece of data is in our phone and getting it from there to our laptop seems like an impossibility. You look at your phone and say, but it's just an inch away….can't my phone just borrow that big screen so I can continue doing this thing I was doing.

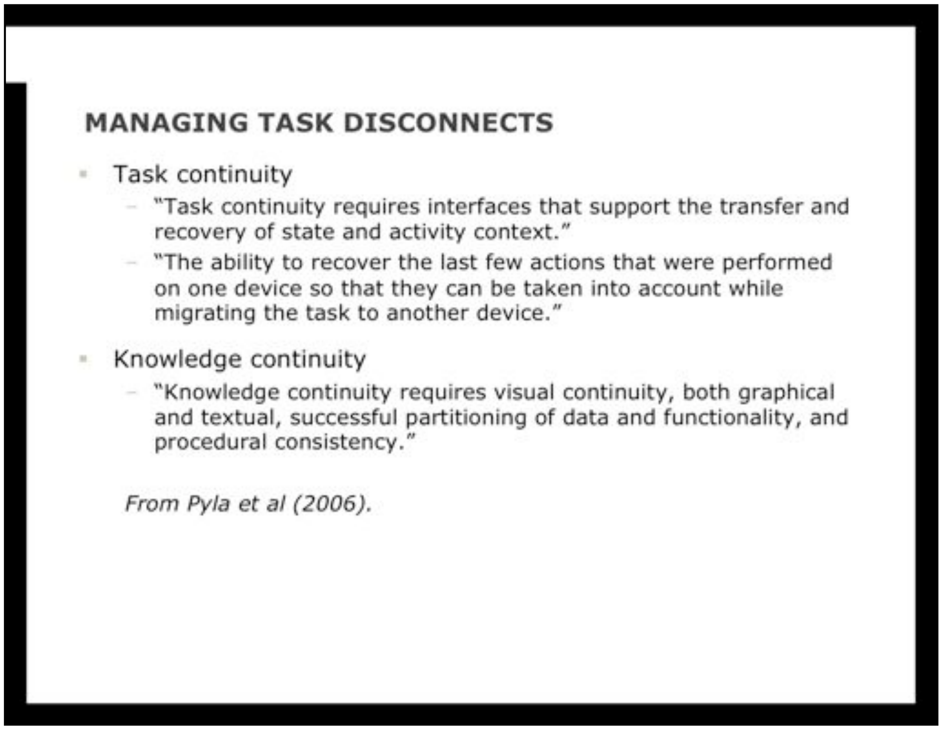

That's what Pardha Pyla and Manas Tungare called a "task disconnect." In your mind you have an idea of a single thing that you're doing and you want to continue doing that thing with whatever tool is available. Spending the extra 10 seconds or 20 seconds or 5 minutes recreating your mental state in a new environment completely brakes your thought flow. You've just experienced a task disconnect.

How do you manage that transition? Pyla and Tungare give some very general guidelines in the paper where they described this idea, but there's no clear way.

To bring up Amazon again, when it brings you to whatever page you were last looking at in a different avatar, that's managing a task disconnect.

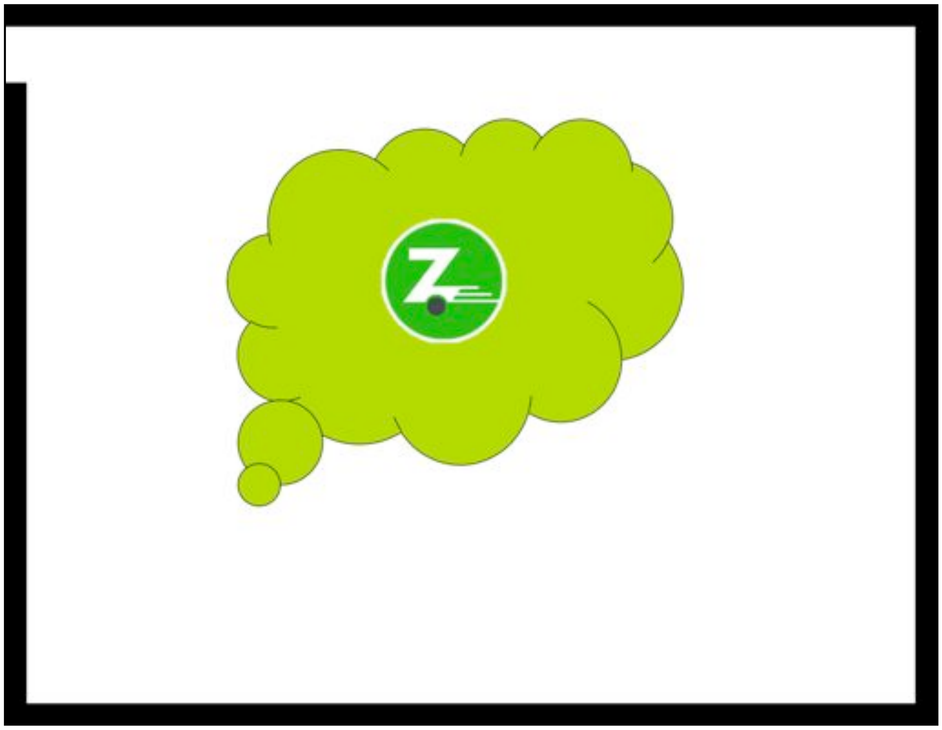

Another more varied example. When you want to reserve your car from Zipcar, you see through the browser to the Zipcar service on the other side. As you walk you can check on the status and location of your car with an iPhone app. When you get there, your key fob works and the car opens. At no time, theoretically, has the service interrupted your flow of thought around the process of getting your car. You never say, “Oh, wait, now I have to do this other thing with Zip car because I’m now using my phone.” It’s a very smooth experience precisely because they have managed the task disconnects well.

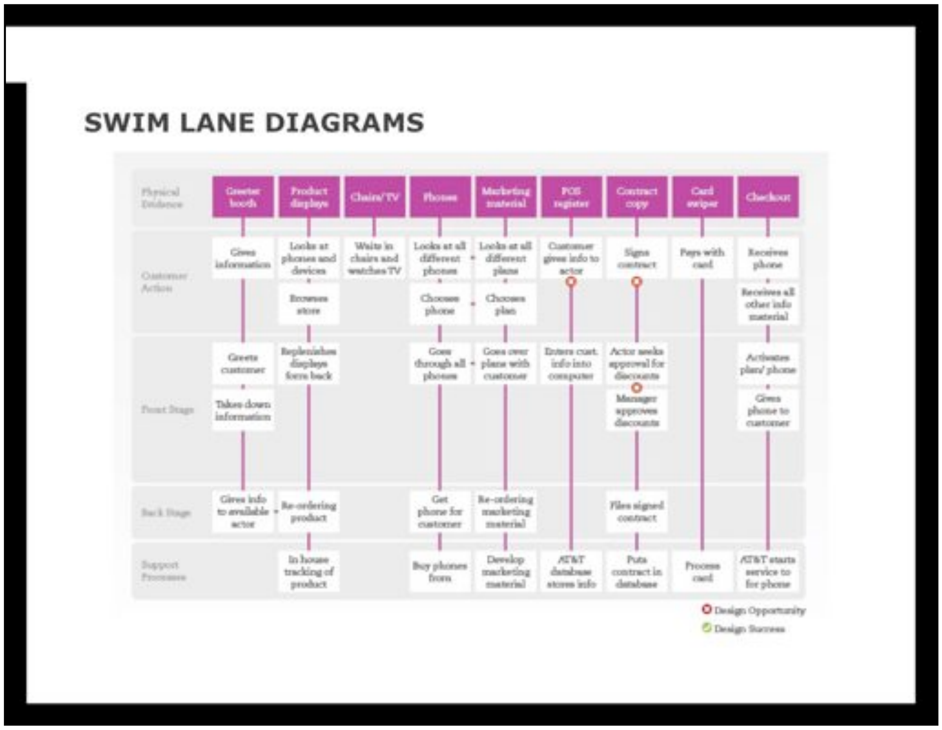

So how do you know where you need to manage those? Well one service design tool that's gaining popularity is the Swim Lane diagram. This one is by Iza Cross, who was a student until recently at the Savannah College of Art and Design's service design program. This type of diagram maps avatars of a service to what those avatars mean in terms of customer actions, service actions and back end technologies. I think it's probably the most useful of all of the service design tools right now in terms of understanding ubicomp UX design and it's a good way to organize what you know about the service you're designing. A task disconnect between two avatars happens when a task crosses between two of the vertical bars.

Source: http://www.izacross.com/portfolio/attmobile

FINAL Q&A AND DISCUSSION


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

Thank you.