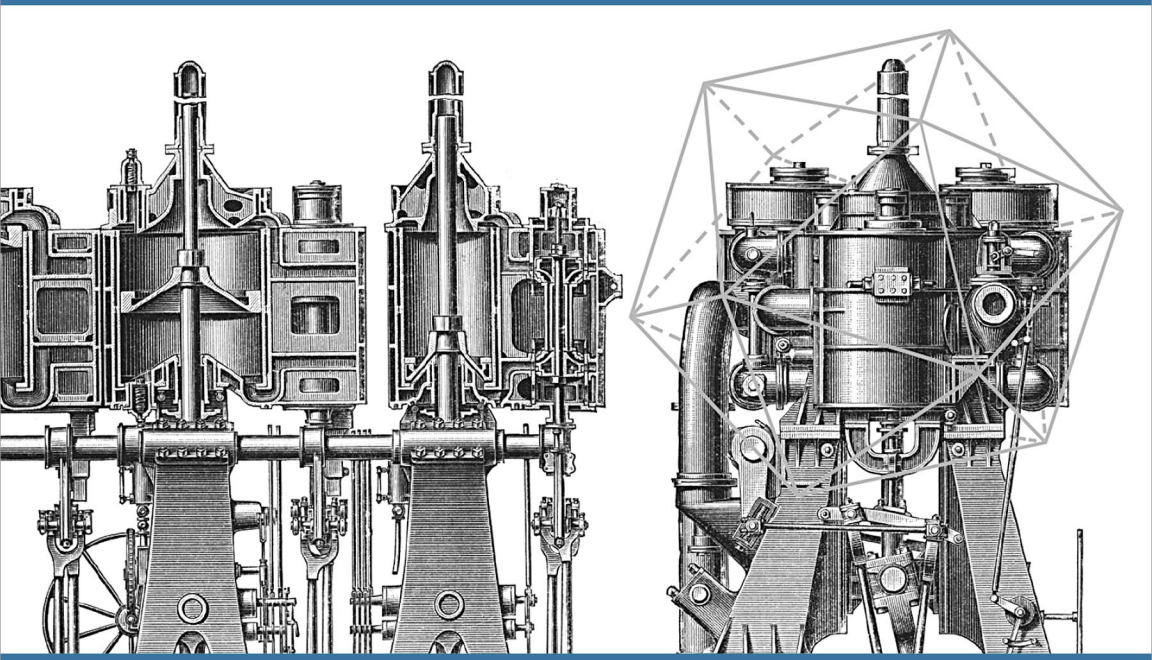
O'REILLY®

Software & Hardware Collide

The Best of Solid



Jon Bruner, Glen Martin, Matthew Gast,
Tim O'Reilly, Kipp Bradford,
Jim Stogdill & Andy Fitzgerald

www.dbooks.org



Stay up-to-date with the O'Reilly Solid Newsletter

Learn about the latest in the rapidly emerging area of the connected world, Internet of Things, and disruptive innovation.

See a sample:

oreilly.com/go/**solidnews**

O'REILLY®

Solid

SOFTWARE/HARDWARE EVERYWHERE

©2014 O'Reilly Media, Inc. The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. 14651

Software & Hardware Collide

The Best of Solid

Jon Bruner, Glen Martin, Matthew Gast, Tim O'Reilly, Kipp Bradford, Jim Stogdill, and Andy Fitzgerald

Beijing • Cambridge • Farnham • Köln • Sebastopol • Tokyo O'REILLY®

www.dbooks.org

Software & Hardware Collide

by Jon Bruner, Glen Martin, Matthew Gast, Tim O'Reilly, Kipp Bradford, Jim Stogdill, and Andy Fitzgerald

Copyright © 2014 O'Reilly Media . All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*http://my.safaribooksonline.com*). For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

Editor: Mike Loukides

April 2014: First Edition

Revision History for the First Edition:

2014-04-24: First release

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps.

While every precaution has been taken in the preparation of this book, the publisher and authors assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

ISBN: 978-1-491-90463-3

[LSI]

Table of Contents

Introduction 1

Software, Hardware, Everywhere 3

More 1876 than 1995 7

The New Bot on the Block 11

If This/Then That (IFTTT) and the Belkin WeMo 15

I, Cyborg 23

#IoTH: The Internet of Things and Humans 27

The Industrial Internet of Things 33

Trope or Fact? Technology Creates More Jobs than It Destroys 39

Architecture, Design, and the Connected Environment 49
Architecture and Design 50

iBeacon Basics 55

The Lingering Seduction of the Page 61
First Stop: The Library 61
Reading Brains 63
The Trouble with Systems (and Why They're Worth It) 65

www.dbooks.org

Fumbling toward System Literacy 66

Let There Be (Intelligent) Light. 69

Introduction

Jon Bruner, Co-chair, O'Reilly Solid

Over the last few years, members of the O'Reilly community have increasingly turned their attention to hardware. It's an engaging area of exploration—one with the tactile satisfaction of building something, whether it's a tool with commercial applications or a Saturday afternoon project.

For many of our developer friends, hardware has emerged as a way to extend their software into the physical world. It's a way to take the instrumentation, learning, and optimization that they've built in the era of "big data" into areas that they haven't touched before—like thermostats, cars, and medical devices.

We've also heard from companies that are experienced in building these sorts of devices. They're looking for software thinkers to help them add intelligence to their machines—to bring the mindset of the Web onto the factory floor.

It's getting much easier to design, prototype, build, and market hardware, and it's getting easier to integrate software and hardware into fluid packages of intelligence, beauty, and intuitive design. Hardware and software are coming together, and it's becoming essential for each side to understand the other.

Solid comes from that realization, and from the stories and theories of our alpha-geek friends, some of which are compiled here. As you read through them, consider what becomes possible in the collision between real and virtual.

www.dbooks.org

Software, Hardware, Everywhere

Jon Bruner

Real and virtual are crashing together. On one side is hardware that acts like software: IP-addressable, controllable with JavaScript APIs, able to be stitched into loosely-coupled systems—the mashups of a new era. On the other is software that's newly capable of dealing with the complex subtleties of the physical world—ingesting huge amounts of data, learning from it, and making decisions in real time.

The result is an entirely new medium that's just beginning to emerge. We can see it in Ars Electronica Futurelab's Spaxels, which use drones to render a three-dimensional pixel field; in Baxter, which layers emotive software onto an industrial robot so that anyone can operate it safely and efficiently; in OpenXC, which gives even hobbyist-level programmers access to the software in their cars; in SmartThings, which ties web services to light switches.

The new medium is something broader than terms like "Internet of Things," "Industrial Internet," or "connected devices" suggest. It's an entirely new discipline that's being built by software developers, roboticists, manufacturers, hardware engineers, artists, and designers.

Ten years ago, building something as simple as a networked thermometer required some understanding of electrical engineering. Now it's a Saturday-afternoon project for a beginner. It's a shift we've already seen in programming, where procedural languages have become more powerful and communities have arisen to offer free help with programming problems. As the blending of hardware and software continues, the physical world will become democratized: the ranks of people who can address physical challenges from lots of different backgrounds will swell.

www.dbooks.org

The outcome of all of this combining and broadening, I hope, will be a world that's safer, cleaner, more efficient, and more accessible. It may also be a world that's more intrusive, less private, and more vulnerable to ill-intentioned interference. That's why it's crucial that we develop a strong community from the new discipline.

Solid, which Joi Ito and I will present on May 21 and 22, will bring members of the new discipline together to discuss this new medium at the blurred line between real and virtual. We'll talk about design beyond the computer screen; software that understands and controls the physical world; new hardware tools that will become the building blocks of the connected world; frameworks for prototyping and manufacturing that make it possible for anyone to create physical devices; and anything else that touches both the concrete and abstract worlds.

The business implications of the new discipline are just beginning to play out. Software companies are eyeing hardware as a way to extend their offerings into the physical world—think, for instance, of Google's acquisition of Motorola and its work on a driverless car—and companies that build physical machines see software as a crucial component of their products. The physical world as a service, a business model that's something like software as a service, promises to upend the way we buy and use machines, with huge implications for accessibility and efficiency. These types of service frameworks, along with new prototyping tools and open-source models, are making hardware design and manufacturing vastly easier.

A few interrelated concepts that I've been thinking about as we've sketched out the idea for Solid:

- **APIs for the physical world.** Abstraction, modularity, and loosely-coupled services—the characteristics that make the Web accessible and robust—are coming to the physical world. Open-source libraries for sensors and microcontrollers are bringing easy-to-use and easy-to-integrate software interfaces to everything from weather stations to cars. Networked machines are defining a new physical graph, much like the Web's information graph. These models are starting to completely reorder our physical environment. It's becoming easier to trade off functionalities between hardware and software; expect the proportion of intelligence residing in software to increase over time.
- **Manufacturing made frictionless.** Amazon's EC2 made it possible to start writing and selling software with practically no capital

investment. New manufacturing-as-a-service frameworks bring the same approach to building things, making factory work fast and capital-light. Development costs are plunging, and it's becoming easier to serve niches with specialized hardware that's designed for a single purpose. The pace of innovation in hardware is increasing as the field becomes easier for entrepreneurs to work in and financing becomes available through new platforms like Kickstarter. Companies are emerging now that will become the Amazon Web Services of manufacturing.

- **Software intelligence in the physical world.** Machine learning and data-driven optimization have revolutionized the way that companies work with the Web, but the kind of sophisticated knowledge that Amazon and Netflix have accumulated has been elusive in the offline world. Hardware lets software reach beyond the computer screen to bring those kinds of intelligence to the concrete world, gathering data through networked sensors and exerting real-time control in order to optimize complicated systems. Many of the machines around us could become more efficient simply through intelligent control: a furnace can save oil when software, knowing that homeowners are away, turns down the thermostat; a car can save gas when Google Maps, polling its users' smartphones, discovers a traffic jam and suggests an alternative route—the promise of software intelligence that works above the level of a single machine. The Internet stack now reaches all the way down to the phone in your pocket, the watch on your wrist, and the thermostat on your wall.
- **Every company is a software company.** Software is becoming an essential component of big machines for both the builders and the users of those machines. Any company that owns big capital machines needs to get as much out of them as possible by optimizing their operation with software, and any company that builds machines must improve and extend them with layers of software in order to be competitive. As a result, a software startup with promising technology might just as easily be bought by a big industrial company as by a Silicon Valley software firm. This has important organizational, cultural, and competency impact.
- **Complex systems democratized.** The physical world is becoming accessible to innovators at every level of expertise. Just as it's possible to build a web page with only a few hours' learning, it's becoming easier for anyone to build things, whether electronic or

www.dbooks.org

not. The result: realms like the urban environment that used to be under centralized control by governments and big companies are now open to innovation from anyone. New economic models and communities will emerge in the physical world just as they've emerged online in the last twenty years.

- **The physical world as a service.** Anything from an Uber car to a railroad locomotive can be sold as a service, provided that it's adequately instrumented and dispatched by intelligent software. Good data from the physical world brings about efficient markets, makes cheating difficult, and improves quality of service. And it will revolutionize business models in every industry as service guarantees replace straightforward equipment sales. Instead of just selling electricity, a utility could sell heating and cooling—promising to keep a homeowner's house at 70 degrees year round. That sales model could improve efficiency and quality of life, bringing about incentive for the utility to invest in more efficient equipment and letting it take advantage of economies of scale.
- **Design after the screen.** Our interaction with software no longer needs to be mediated through a keyboard and screen. In the connected world, computers gather data through multiple inputs outside of human awareness and intuit our preferences. The software interface is now a dispersed collection of conventional computers, mobile phones, and embedded sensors, and it acts back onto the world through networked microcontrollers. Computing happens everywhere, and it's aware of physical-world context.
- **Software replaces physical complexity.** A home security system is no longer a closed network of motion sensors and door alarms; it's software connected to generic sensors that decides when something is amiss. In 2009, Alon Halevy, Peter Norvig, and Fernando Pereira wrote that having lots and lots of data can be more valuable than having the most elegant model. In the connected world, having lots and lots of sensors attached to some clever software will start to win out over single-purpose systems.

These are some rough thoughts about an area that we'll all spend the next few years trying to understand. This is an open discussion, and we welcome thoughts on it from anyone.

More 1876 than 1995

Glen Martin

Philadelphia's Centennial Exposition of 1876 was America's first World's Fair, and was ostensibly held to mark the nation's 100th birthday. But it heralded the future as much as it celebrated the past, showcasing the country's strongest suit: technology.

The centerpiece of the Expo was a gigantic Corliss engine, the apotheosis of 40 years of steam technology. Thirty percent more efficient than standard steam engines of the day, it powered virtually every industrial exhibit at the exposition via a maze of belts, pulleys, and shafts. Visitors were stunned that the gigantic apparatus was supervised by a single attendant, who spent much of his time reading newspapers.

"This exposition was attended by 10 million people at a time when travel was slow and difficult, and it changed the world," observes Jim Stogdill, general manager of O'Reilly's upcoming Internet-of-Things-related conference, Solid.

"Think of a farm boy from Kansas looking at that Corliss engine, seeing what it could do, thinking of what was possible," Stogdill continues. "When he left the exposition, he was a different person. He understood what the technology he saw meant to his own work and life."

The 1876 exposition didn't mark the beginning of the Industrial Revolution, says Stogdill. Rather, it signaled its fruition, its point of critical mass. It was the nexus where everything—advanced steam technology, mass production, railroads, telegraphy—merged.

"It foreshadowed the near future, when the Industrial Revolution led to the rapid transformation of society, culturally as well as economi-

www.dbooks.org

cally. More than 10,000 patents followed the exposition, and it accelerated the global adoption of the 'American System of Manufacture.' The world was never the same after that."

In terms of the Internet of Things, we have reached that same point of critical mass. In fact, the present moment is more similar to 1876 than to more recent digital disruptions, Stogdill argues. "It's not just the sheer physicality of this stuff," he says. "It is also the breadth and speed of the change bearing down on us."

While the Internet changed everything, says Stogdill, "its changes came in waves, with scientists and alpha geeks affected first, followed by the early adopters who clamored to try it. It wash't until the Internet was ubiquitous that every Kansas farm boy went online. That 1876 Kansas farm boy may not have foreseen every innovation the Industrial Revolution would bring, but he knew—whether he liked it or not—that his world was changing."

As the Internet subsumes physical objects, the rate of change is accelerating, observes Stogdill. "Today, stable wireless platforms, standardized software interface components and cheap, widely available sensors have made the connection of virtually every device—from coffee pots to cars—not only possible; they have made it certain."

"Internet of Things" is now widely used to describe this latest permutation of digital technology; indeed, "overused" may be a more apt description. It teeters on the knife-edge of cliché. "The term is clunky," Stogdill acknowledges, "but the buzz for the underlying concepts is deserved."

Stogdill is quick to point out that this "Internet of Everything" goes far beyond the development of new consumer products. Open source hardware and software already are allowing the easy integration of programatic interfaces with everything from weather stations to locomotives. Large, complicated systems—water delivery infrastructure, power plants, sewage treatment plants, office buildings—will be made intelligent by these software and sensor packages, allowing real-time control and exquisitely efficient operation. Manufacturing has been made frictionless, development costs are plunging, and new manufacturing-as-a-service frameworks will create new business models and drive factory production costs down and production up.

"When the digital age began accelerating," Stogdill explains, "Nicholas Negroponte observed that the world was moving from atoms to bits

—that is, the high-value economic sectors were transforming from industrial production to aggregating information."

"I see the Internet of Everything as the next step," he says. "We won't be moving back to atoms, but we'll be combining atoms and bits, merging software and hardware. Data will literally grow physical appendages, and inform industrial production and public services in extremely powerful and efficient ways. Power plants will adjust production according to real-time demand, traffic will smooth out as driverless cars become commonplace. We'll be able to track air and water pollution to an unprecedented degree. Buildings will not only monitor their environmental conditions for maximum comfort and energy efficiency, they'll be able to adjust energy consumption so it corresponds to electricity availability from sustainable sources."

Stogdill believes these converging phenomena have put us on the cusp of a transformation as dramatic as the Industrial Revolution.

"Everyone will be affected by this collision of hardware and software, by the merging of the virtual and real," he says. "It's really a watershed moment in technology and culture. We're at one of those tipping points of history again, where everything shifts to a different reality. That's what the 1876 exposition was all about. It's one thing to read about the exposition's Corliss engine, but it would've been a wholly different experience to stand in that exhibit hall and see, feel, and hear its 1,400 horsepower at work, driving thousands of machines. It is that sensory experience that we intend to capture with Solid. When people look back in 150 years, we think they could well say, 'This is when they got it. This is when they understood.'"

www.dbooks.org

The New Bot on the Block

Glen Martin

Fukushima changed robotics. More precisely, it changed the way the Japanese view robotics. And given the historic preeminence of the Japanese in robotic technology, that shift is resonating through the entire sector.

Before the catastrophic earthquake and tsunami of 2011, the Japanese were focused on "companion" robots, says Rodney Brooks, a former Panasonic Professor of Robotics at MIT, the founder and former technical officer of IRobot, and the founder, chairman and CTO of Rethink Robotics. The goal, says Brooks, was making robots that were analogues of human beings—constructs that could engage with people on a meaningful, emotional level. Cuteness was emphasized: a cybernetic, if much smarter, equivalent of Hello Kitty, seemed the paradigm.

But the multiple core meltdown at the Fukushima Daiichi nuclear complex following the 2011 tsunami changed that focus abruptly.

"Fukushima was a wake-up call for them," says Brooks. "They needed robots that could do real work in highly radioactive environments, and instead they had robots that were focused on singing and dancing. I was with IRobot then, and they asked us for some help. They realized they needed to make a shift, and now they're focusing on more pragmatic designs."

Pragmatism was always the guiding principle for Brooks and his companies, and is currently manifest in Baxter, Rethink's flagship product. Baxter is a breakthrough production robot for a number of reasons. Equipped with two articulated arms, it can perform a multitude of tasks. It requires no application code to start up, and no expensive software to function. No specialists are required to program it; workers

www.dbooks.org

with minimal technical background can "teach" the robot right on the production line through a graphical user interface and arm manipulation. Also, Baxter requires no cage—human laborers can work safely alongside it on the assembly line.

Moreover, it is cheap: about $25,000 per unit. It is thus the robotic equivalent of the Model T, and like the Model T, Baxter and its subsequent iterations will impose sweeping changes in the way people live and work.

"We're at the point with production robots where we were with mobile robots in the late 1980s and early 1990s," says Brooks. "The advances are accelerating dramatically."

What's the biggest selling point for this new breed of robot? Brooks sums it up in a single word: dignity.

"The era of cheap labor for factory line work is coming to a close, and that's a good thing," he says. "It's grueling, and it can be dangerous. It strips people of their sense of worth. China is moving beyond the human factory line—as people there become more prosperous and educated, they aspire to more meaningful work. Robots like Baxter will take up the slack out of necessity."

And not just for the assemblage of widgets and gizmos. Baxter-like robots will become essential in the health sector, opines Brooks—particularly in elder care. As the Baby Boom piglet continues its course through the demographic python, the need for attendants is outstripping supply. No wonder: the work is low-paid and demanding. Robots can fill this breach, says Brooks, doing everything from preparing and delivering meals to shuttling laundry, changing bedpans and mopping floors.

"Again, the basic issue is dignity," Brooks said. "Robots can free people from the more menial and onerous aspects of elder care, and they can deliver an extremely high level of service, providing better quality of life for seniors."

Ultimately, robots could be more app than hardware: the sexy operating system on Joaquin Phoenix's mobile device in the recent film "Her" may not be far off the mark. Basically, you'll carry a "robot app" on your smartphone. The phone can be docked to a compatible mechanism — say, a lawn mower, or car, or humanoid mannequin—resulting in an autonomous device ready to trim your greensward, chauffeur you to the opera, or mix your Mojitos.

YDreams Robotics, a company co-founded by Brooks protégé Artur Arsenio, is actively pursuing this line of research.

"It's just a very efficient way of marketing robots to mass consumers," says Arsenio. "Smartphones basically have everything you need, including cameras and sensors, to turn mere things into robots."

YDream has its first product coming out in April: a lamp. It's a very fine if utterly unaware desk lamp on its own, says Artur, but when you connect it to a smartphone loaded with the requisite app, it can do everything from intelligently adjusting lighting to gauging your emotional state.

"It uses its sensors to interface socially," Artur says. "It can determine how you feel by your facial expressions and voice. In a video conference, it can tell you how other participants are feeling. Or if it senses you're sad, it may Facebook your girlfriend that you need cheering up."

Yikes. That may be a bit more interaction than you want from a desk lamp, but get used to it. Robots could intrude in ways that may seem a little off-putting at first—but that's a marker of any new technology. Moreover, says Paul Saffo, a consulting professor at Stanford's School of Engineering and a technology forecaster of repute, the highest use of robots won't be doing old things better. It will be doing new things, things that haven't been done before, things that weren't possible before the development of key technology.

"Whenever we have new tech, we invariably try to use it to do old things in a new way—like paving cow paths," says Saffo. "But the sooner we get over that—the sooner we look beyond the cow paths—the better off we'll be. Right now, a lot of the thinking is, 'Let's have robots drive our cars, and look like people, and be physical objects.'

But the most important robots working today don't have physical embodiments, says Saffo—think of them as ether-bots, if you will. Your credit application? It's a disembodied robot that gets first crack at that. And the same goes for your resume when you apply for a job.

In short, robots already are embedded in our lives in ways we don't think of as "robotic." This trend will only accelerate. At a certain point, things may start feeling a little—well Singularity-ish. Not to worry—it's highly unlikely Skynet will rain nuclear missiles down on us anytime soon. But the melding of robotic technology with dumb things nevertheless presents some profound challenges—mainly because robots and humans react on disparate time scales.

www.dbooks.org

"The real questions now are authority and accountability," says Saffo. "In other words, we have to figure out how to balance the autonomy systems need to function with the control we need to ensure safety."

Saffo cites modern passenger planes like the Airbus 330 as an example.

"Essentially they're flying robots," he says. "And they fly beautifully, conserving fuel to the optimal degree and so forth. But the design limits are so tight—if they go too fast, they can fall apart; if they go too slow, they stall. And when something goes wrong, the pilot has perhaps 50 kilometers to respond. At typical speeds, that doesn't add up to much reaction time."

Saffo noted the crash of Air France Flight 447 in the mid-Atlantic in 2009 involved an Airbus 330. Investigations revealed the likely cause was turbulence complicated by the icing up of the plane's speed sensors. This caused the autopilot to disengage, and the plane began to roll. The pilots had insufficient time to compensate, and the aircraft slammed into the water at 107 knots.

"The pilot figured out what was wrong—but it was 20 seconds too late," says Saffo. "To me, it shows we need to devote real effort to defining boundary parameters on autonomous systems. We have to communicate with our robots better. Ideally, we want a human being constantly monitoring the system, so he or she can intervene when necessary. And we need to establish parameters that make intervention even possible."

If This/Then That (IFTTT) and the Belkin WeMo

Matthew Gast

Like most good technologists, I am lazy. In practice, this sometimes means that I will work quite hard with a computer to automate a task that, for all intents and purposes, just isn't that hard. In fits and starts for the past 10 years, I have been automating my house in various ways. It makes my life easier when I am at home, though it does mean that friends who watch my house when I'm gone need to be briefed on how to use it. If you are expecting to come into my house and use light switches and the TV as you do every place else, well, that's why you need a personalized orientation to the house.

In this post, I'll talk briefly about one of the most basic automation tasks I've carried out, which is about how the lights in my house are controlled.

The humble light switch was invented in the late 19th century, with the "modern" toggle switch following in the early 20th century. The toggle switch has not changed in about 100 years because it does exactly what is needed and is well understood. The only disadvantage to the toggle switch is that you have to touch it to operate it, and that means getting off the couch.

My current home is an older home that was previously a storefront. All the plumbing is on one wall of the house, which I sometimes jokingly call "the wall of modernity." The telephone demarc is on the opposite side, though I've made some significant modifications to the telephone wiring from the days when I was using Asterisk more ex-

www.dbooks.org

tensively. As a large open space, there is not much built-in lighting, so most of the lights are freestanding floor lamps.

The initial "itch" that I was scratching with my first home automation project was a consistent forgetfulness to turn off lights when I went to bed—I would settle into bed up in the loft and then realize that I needed to get up and turn off a light. At that point, remote controlled lights seemed pretty darn attractive.

My first lighting automation was done with Insteon, a superset of the venerable 1970s-era X10 protocol. With Insteon controls, I could use an RF remote and toggle lights on and off from a few locations within the house.

With the hardware in place, I turned to finding software. An RF remote is a good start, but I wanted to eliminate even that. What I found is that the home automation community doesn't really have off-the-shelf software that "just works" and lets you start doing things out of the box. I used Mister House for a while, but at the time I was using it, the Insteon support was pretty new. I worked with a friend writing some automation code—well to be honest, being a tester for him, but we were solving problems unique to our homes, not writing something that was ultimately going to be the answer for our parents.

Earlier this year, I was introduced to IFTTT, which stands for "If This, Then That." It is the closest thing I have seen to a generic software stack that can readily be used for home automation purposes. IFTTT can be used for much more than just home automation, but I'm going to stick to that restricted use for the purpose of this post.

There is much to like about IFTTT, so I'll focus on just a few attributes:

- Programming without knowing programming. I've taken enough computer science courses to know how to write a program, but I'm not good enough to be a professional programmer. IFTTT lets you write rules—essentially, programs that control objects in the physical world—without the high bar of learning specialized syntax or all the system administration work that goes along with running a program.
- Extending control beyond my home. Traditional home automation has been based on sensors inside the home. In order to respond to something, there needs to be a sensor to gather that data. To turn the lights on after dark, you need a sensor that tells your house it's dark. To do something when it's raining, you need a

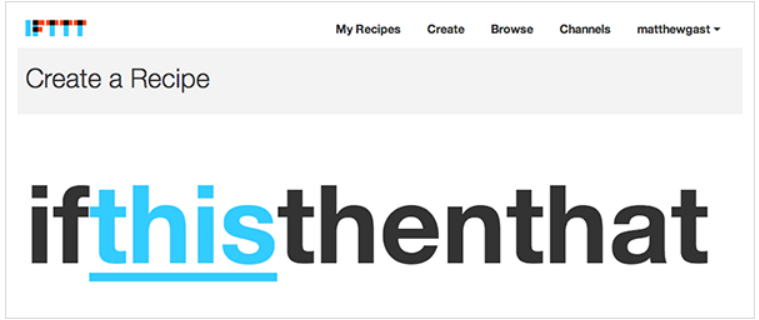
sensor that gets wet, and so forth. As you'll see, IFTTT lets you choose a location and use information from Internet services to assemble the context.

- One of the cleanest and easiest designs I've seen. I know everybody brags about design, but IFTTT is so simple that a colleague of mine in sales uses IFTTT extensively. (Insert your own joke about the technical competence of sales representatives here.)

To act in the physical realm of my house, though, IFTTT requires devices that can take action. The Belkin WeMo is a family of products, including a remote-controlled electrical outlet, light switch, and motion sensor. The WeMo uses Wi-Fi to connect to my home network, so it can be controlled by any device on my network, or any service on the Internet.

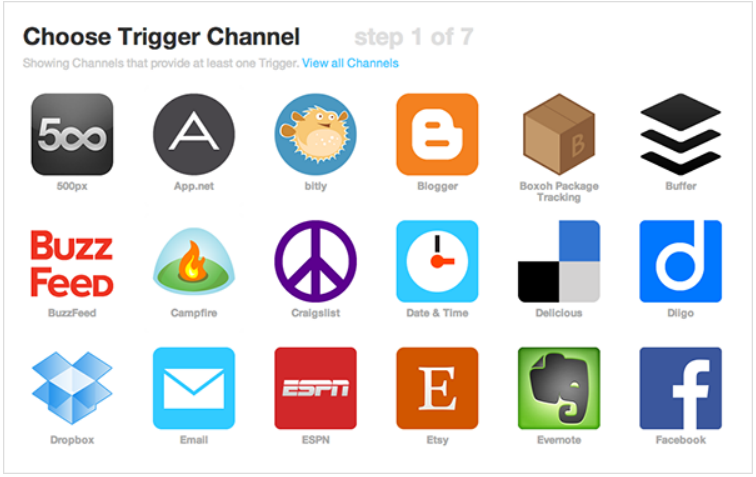
So, let's start out with a simple idea: I have a light, and I'd like to turn the light on when it's dark. In traditional home automation, I can easily do that by picking a time to turn the light on and using the same time every day. If I pick a time that is appropriate for the winter, though, the light will come on too early in June. If I pick a time that is right for June, I'll have to get up and turn the lights on in December when it gets dark. Or, using traditional home automation, I could find a source of sunset times, pull them in manually, and hope the data feed I'm using never changes.

That is, until IFTTT. In this example, I'll show you how to set up a WeMo to have a light that turns on at sunset. IFTTT mini-programs are called "Recipes" and consist of a trigger and an action. Building a recipe is a straightforward guided process that begins by picking the trigger:



www.dbooks.org

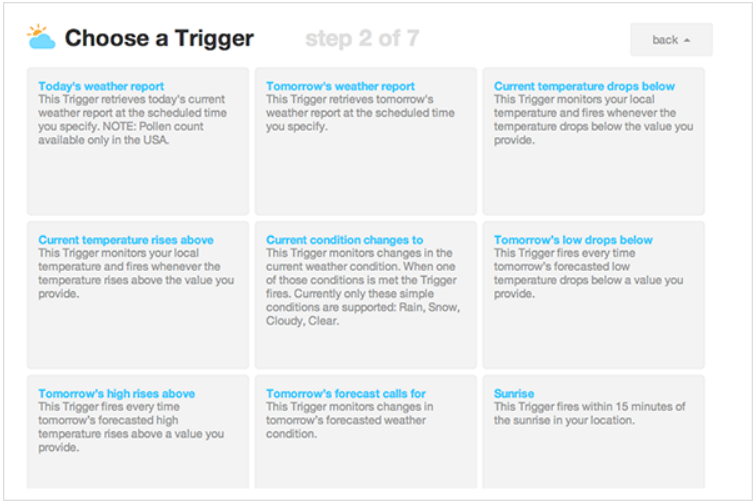
IFTTT's interaction with the world is organized into "channels," and more channels are being added on a routine basis. The first step in creating a recipe is to choose the trigger:



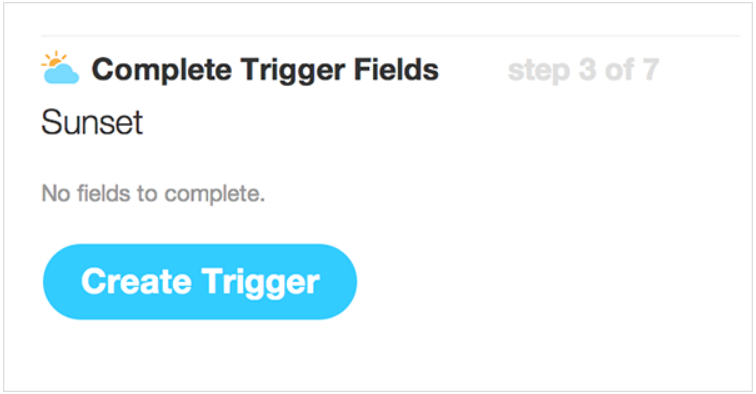
Triggers are organized alphabetically. I cut the screen shot off after F because I wanted to get down to W for Weather and WeMo. The first time you select the weather channel, you'll be prompted to set a location; I chose my home in San Francisco:



The weather channel has a number of components that you can choose to use as a trigger. Some are expected, such as the weather forecast. However, the weather channel also allows you to take actions when conditions change. The weather data used by IFTTT is also rich enough to have pollen count and UV index, which are not always readily available from many forecasts. For the purpose of our example, we'll be using the "sunset" option, though it's easy to imagine creating triggers to control lights when the sky becomes cloudy or turning on air filtration when the pollen count rises:



Some triggers require additional configuration. Sunset is straightforward because it's known for the location that you chose:

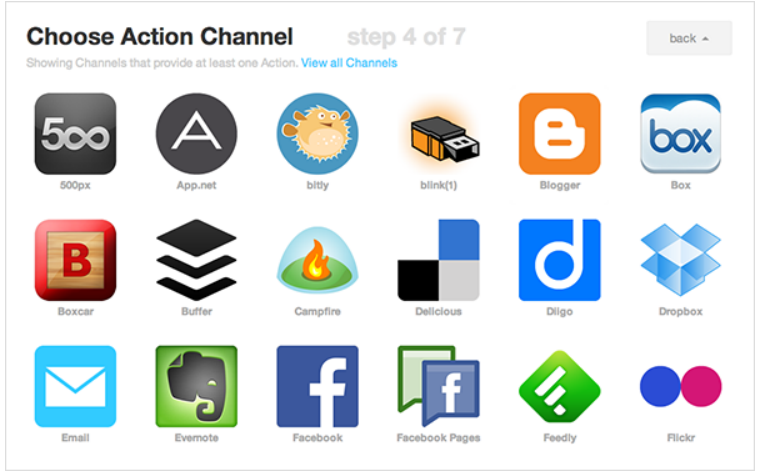


Here at step three, the first part of the rule is done. We have set up a rule that will fire every day at sunset. IFTTT helpfully fills in the "this" part of our rule in a way that makes it obvious what we're doing:

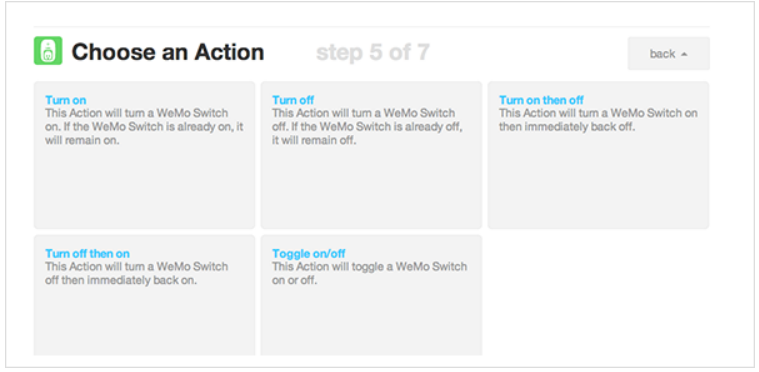


www.dbooks.org

The second part of writing a rule is to lay out the action (the "that") to take when the trigger fires. Once again, we choose a channel to take the action, choosing from the large number of channels that are available:

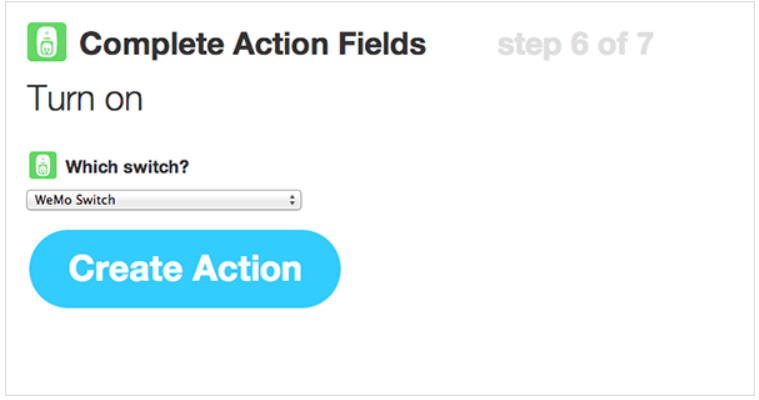


We're trying to control lights plugged into a Belkin WeMo, so we'll scroll down to "W" and pick the WeMo. It has the actions that you might expect from what is essentially a power switch: turn off the outlet, turn on the outlet, blink the outlet, or toggle the outlet. In our case, we want to turn on the outlet to turn on the light:

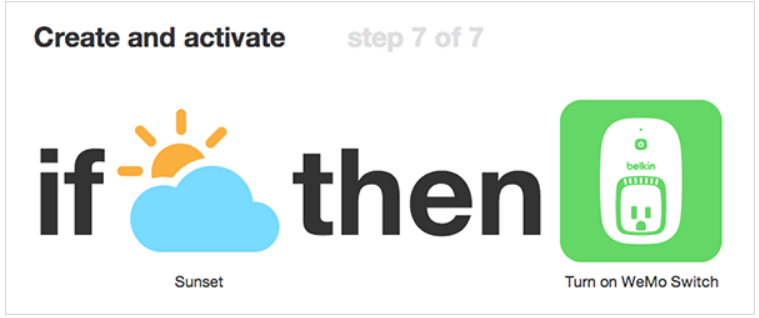


From the menu of options, choose "turn on." IFTTT supports having multiple WeMo switches, each of which can be named. A drop-down allows you to choose the switch being controlled by the rule so that many different devices can be controlled. For example, a coffeepot

might be turned on when an alarm goes off, different lights might be controlled by different rules, and so forth:



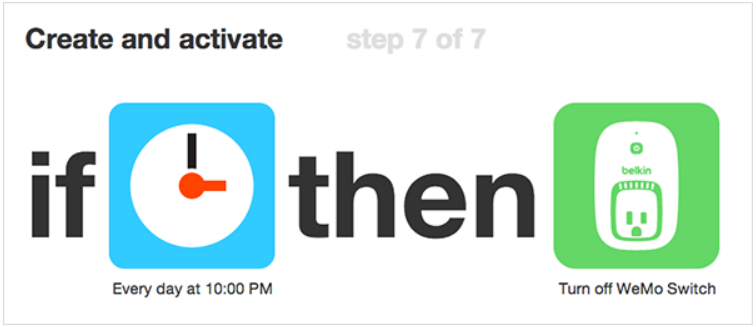
The last step in creating a rule is to finalize the action. IFTTT helpfully displays the rule in full in a nice simple form. Is there any doubt what the rule we've just created will do:



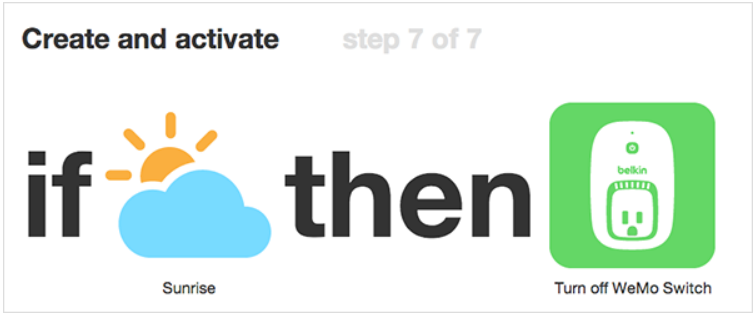
In my personal setup, I use a rule that turns off the light at 10 p.m. as a reminder to go to bed. There is not yet an easy way to say "turn off the light two hours after it turned on" because IFTTT doesn't hold much state (yet).

One aspect of IFTTT that I recently appreciated was how easy it is to change the rule. When going out of town, I decided to have the lights on all night so that pet sitters wouldn't have to figure out all the ways in which lights could be turned on. I simply deactivated the "turn off at 10 p.m." rule that was like this:

www.dbooks.org



I then replaced it with a rule that turned off the lights at sunrise. (The lights in question are a low-power strand of LED lights.)



Could I have accomplished the same tasks in other home automation systems? Absolutely, but it would have taken me much longer to get to my end goal, and I would have had to do significantly more testing to believe that my automation would behave as expected. The combination of IFTTT and the WeMo makes setup much easier and more accessible.

I, Cyborg

Glen Martin

There is an existential unease lying at the root of the Internet of Things—a sense that we may emerge not less than human, certainly, but other than human.

Well, not to worry. As Kelsey Breseman, engineer at Technical Machine, points out, we don't need to fret about becoming cyborgs. We're already cyborgs: biological matrices augmented by wirelessly connected silicon arrays of various configurations. The problem is that we're pretty clunky as cyborgs go. We rely on screens and mobile devices to extend our powers beyond the biological. That leads to everything from atrophying social skills as face-to-face interactions decline to fatal encounters with garbage trucks as we wander, texting and oblivious, into traffic.

So, if we're going to be cyborgs, argues Breseman, let's be competent, sophisticated cyborgs. For one thing, it's now in our ability to upgrade beyond the screen. For another, being better cyborgs may make us—paradoxically—more human.

"I'm really concerned about how we integrate human beings into the growing web of technology," says Breseman, who will speak at O'Reilly's upcoming Solid conference in San Francisco in May. "It's easy to get caught up in the 'cool new thing' mentality, but you can end up with a situation where the point for the technology is the technology, not the human being using it. It becomes closed rather than inclusive—an 'app developers developing apps for app developers to develop apps' kind of thing."

Those concerns have led Breseman and her colleagues at Technical Machine to the development of the Tessel: an open-source Arduino-

www.dbooks.org

style microcontroller that runs JavaScript and allows hardware project prototyping. And not, Breseman emphasizes, the mere prototyping of 'cool new things'—rather, the prototyping of things that will connect people to the emerging Internet of Things in ways that have nothing to do with screens or smart phones.

"I'm not talking about smart watches or smart clothing," explains Breseman. "In a way, they're already passé. The product line hasn't caught up with the technology. Think about epidermal circuits—you apply them to your skin in the same way you apply a temporary tattoo. They've been around for a couple of years. Something like that has so many potential applications—take the Quantified Self movement, for example. Smart micro devices attached right to the skin would make everything now in use for Quantified Self seem antiquated, trivial."

Breseman looks to a visionary of the past to extrapolate the future: "In the late 1980s, Mark Weiser coined the term 'ubiquitous computing' to describe a society where computers were so common, so omnipresent, that people would ultimately stop interfacing with them," Breseman says. "In other words, computers would be everywhere, embedded in the environment. You wouldn't rely on a specific device for information. The data would be available to you on an ongoing basis, through a variety of non-intrusive—even invisible—sources."

Weiser described such an era as "…the age of calm technology, when technology recedes into the background of our lives…" That trope—calm technology—is extremely appealing, says Breseman.

"We could stop interacting with our devices, stop staring at screens, and start looking at each other, start talking to each other again," she says. "I'd find that tremendously exciting."

Breseman is concerned that the Internet of Things is seen only as a new and shiny buzz phrase. "We should be looking at it as a way to address our needs as human beings," she says, "to connect people to the Internet more elegantly, not just as a source for more toys. Yes, we are now dependent on information technology. It has expanded our lives, and we don't want to give it up. But we're not applying it very well. We could do it so much better."

Part of the problem has been the bifurcation of engineering into software and hardware camps, she says. Software engineers type into screens, and hardware engineers design physical things, and there have been few—if any—places that the twain have met. The two disciplines

are poised to merge in the Internet of Things—but it won't be an easy melding, Breseman allows. Each field carves different neural pathways, inculcates different values.

"Because of that, it has been really hard to figure out things that let people engage with the Internet in a physical sense," Breseman says. "When we were designing Tessel, we discovered how hugely difficult it is to make an interactive Internet device."

Still, Tessel and devices like it ultimately will become the machine tools of the Internet of Everything: the forges and lathes where the new infrastructure is built. That's what Breseman hopes, anyway.

"What we would like," she muses, "is for people to figure out their needs first and then order Tessels rather than the other way around. By that I mean you should first determine why and how connecting to the Internet physically would augment your life, make it better. Then get a Tessel to help you with your prototypes. We'll see more and better products that way, and it keeps the emphasis where it belongs—on human beings, not the devices."

www.dbooks.org

#IoTH: The Internet of Things and Humans

Tim O'Reilly

Rod Smith of IBM and I had a call the other day to prepare for our onstage conversation at O'Reilly's upcoming Solid Conference, and I was surprised to find how much we were in agreement about one idea: so many of the most interesting applications of the Internet of Things involve new ways of thinking about how humans and things cooperate differently when the things get smarter. It really ought to be called the Internet of Things and Humans—#IoTH, not just #IoT!

Let's start by understanding the Internet of Things as the combination of sensors, a network, and actuators. The "wow" factor—the magic that makes us call it an Internet of Things application—can come from creatively amping up the power of any of the three elements.

For example, a traditional "dumb" thermostat consists of only a sensor and an actuator—when the temperature goes out of the desired range, the heat or air conditioning goes on. The addition of a network, the ability to control your thermostat from your smartphone, say, turns it into a simple #IoT device. But that's the bare-base case. Consider the Nest thermostat: where it stands out from the crowd of connected thermostats is that it uses a complex of sensors (temperature, moisture, light, and motion) as well as both onboard and cloud software to provide a rich and beautiful UI with a great deal more intelligence.

- While you can schedule your heating manually, you can also let the Nest "learn" when you get up in the morning and turn on the heat, and when you go to bed and turn it off. After a week or so, it will "understand" and repeat the pattern.

www.dbooks.org

- When you are away, the Nest will notice the absence of movement and automatically turn off the heat.
- When the moisture level is high, the thermostat will adjust the temperature to a lower level than you told it you wanted in order to achieve the right perceived temperature. (When humidity is high, it seems warmer than the thermostat alone would notice.)
- If you have a forced hot-air system, the Nest will remind you when to change the filters based on the number of hours the system has been running.
- The Nest uses external weather data to help explain when your energy usage is abnormally low or high and compares your usage with that of other customers.

So, let's generalize the #IoT paradigm as sensors + network + actuators + local and cloud intelligence + creative UI for gathering both explicit and implicit instructions from humans.

Note that any part of this pattern can vary. For example, some #IoT applications have strong, constant connectivity, while others will increasingly have intermittent connectivity and a lot of autonomy. As the resin.io blog put it, they will be "Strong Devices, Weakly Connected." A fully autonomous robot is our model for this kind of #IoT device.

But let me play devil's advocate with the question: is Uber an #IoT application? Most people would say it is not; it's just a pair of smartphone apps connecting a passenger and driver. But imagine for a moment the consumer end of the Uber app as it is today, and on the other end, a self-driving car. You would immediately see that as #IoT. Using this thought experiment, one way to think of the present Uber is as an example of what Eric Ries calls "concierge minimum viable product"—that is, a product where you emulate some of the functions with humans before you build them in software.

This is a powerful way to think about the Internet of Things because it focuses the mind on the human experience of it, not just the things themselves. I'm very fond of the Aaron Levie Tweet about Uber: "Uber is a $3.5 billion lesson in building for how the world **should** work instead of optimizing for how the world **does** work." That is precisely the lesson that Internet of Things designers need to learn: how does a smart thing make it possible to change the entire experience and workflow of a job we do in the real world?

How do quantified self sensors allow us to change how we think about our health care system? How will self-driving cars change transportation and logistics? How might as familiar a sensor as a camera change how we store our stuff? How might power tools like saws, drills, and routers work if they were "smart"?

Long before we get to fully autonomous devices, there are many "halfway house" applications that are really Internet of Things applications in waiting, which use humans for one or more parts of the entire system. When you understand that the general pattern of #IoTH applications is not just sensor + network + actuator but various combinations of human + network + actuator or sensor + network, you will broaden the possibilities for interfaces and business models.

For example, while we can envision a future of fully automated sensor-driven insulin pumps and other autonomous therapeutic devices, we are not there today, as Scott Hanselman explains. But that doesn't mean #IoT-related technology isn't a powerful tool for rethinking many of the ways we deliver health care. For example, sensors make it possible to do patient "observation" on an outpatient basis, while the health care team monitors those sensors with a tablet or smartphone. Or it might not be the patient who is instrumented, but his or her home, allowing seniors to age in place. IBM calls this the Patient Centered Medical Home (pdf).

Or consider Cargosense, a system for keeping track of environmental conditions for shipment of sensitive medical cargoes. A sensor package tracks the conditions during shipment, but there is, as yet, no opportunity for real-time adjustment. For now, the data is simply consumed via a tablet app that provides regulatory approval of the necessary conditions. This is still incredibly valuable.

How about Makespace? This startup is hardly #IoT at all, but it has that wonderful quality of understanding how a simple sensor, creatively applied, can make possible a complete rethinking of the interaction paradigm. Typical storage units are packed with jumbled boxes of forgotten stuff. What's in there? I can't remember. Makespace has the customer photograph what's in the boxes (we've forgotten that the camera is one of the most powerful sensors we carry about with us!), and then Uber-like, takes them away, to be retrieved on demand. The notion that it's possible to track what's in the box means that people themselves never need to visit their storage center, meaning that it can be located far away, with the contents returned at will.

www.dbooks.org

It is worth noting that the smartphone is the perfect halfway house #IoTH platform. It has a rich package of sensors (actually, far more sensors than the Nest), network connectivity, local data and intelligence, and easy access to cloud backends. And it has access on the other end to devices with all the same characteristics. The actuator—the "robot" that can act on the sensor data—can simply be a human on the other end, but it can be a human who is also augmented with that same sensor package. (For example, Uber depends on real-time sensing of the location of both passenger and driver via their smartphones.)

There is another axis to consider: sometimes the human provides input to the system explicitly (as I do when I turn down my Nest thermostat) and sometimes implicitly (as I do when I leave my house and the Nest notices I'm away.) The sensor package in the phone allows for a wide range of both implicit and explicit interactions: it notifies the Uber driver of my location without me having to do anything, but it lets me call or send a message for explicit communication.

Armed with this design pattern, let's look at a number of intriguing #IoTH applications and devices, each of which inserts the human into different parts of the process using the smartphone or tablet as a key link between human and the rest of the system.

- On the Moto X, Google uses sensor data from the phone ("you seem to be driving") to offer to read me my text messages, but then asks me explicitly whether I want it to do so. I can also wake up the phone without touching it, simply by talking to it and giving it simple commands. My "smartphone" has become a "smart thing," with its sensors used to help choose the best modality for the human to provide input.
- Applications like Waze, which collect real-time traffic data and predict your best route by considering the speed and location of the smartphones of its active users, will be directly connected to the car. Even short of self-driving cars, you can expect the car's mapping system to suggest a different route based on information sent automatically from other vehicles.

This pattern of human as part of the #IoTH system, of course, is not limited to applications that include a phone. For example:

- Taktia's smart hand tools let a human provide the gross motor function, but the robot provides the fine motor control to follow

exact patterns, which are input via machine vision. (This, incidentally, is how robotic surgery also works.)

- Google's self-driving car depends in part on the uploaded memory of humans who have previously driven those same roads in Google Street View cars—essentially, human drivers as cyborgs augmented with detailed location sensing and cameras.

My point is that when you think about the Internet of Things, you should be thinking about the complex system of interaction between humans and things, and asking yourself how sensors, cloud intelligence, and actuators (which may be other humans for now) make it possible to do things differently. It is that creativity in finding the difference that will lead to the breakthrough applications for the Internet of Things and Humans.

www.dbooks.org

The Industrial Internet of Things

Kipp Bradford

A few days ago, a company called Echelon caused a stir when it released a new product called IzoT. You may never have heard of Echelon; for most of us, they are merely a part of the invisible glue that connects modern life. But more than 100 million products—from street lights to gas pumps to HVAC systems—use Echelon technology for connectivity. So, for many electrical engineers, Echelon's products are a big deal. Thus, when Echelon began touting IzoT as the future of the Industrial Internet of Things (IIoT), it was bound to get some attention.



Figure 8-1. Photo: Kipp Bradford. This industrial burner from Weishaupt churns through 40 million BTUs per hour of fuel.

Admittedly, the Internet of Things (IoT) is all the buzz right now. Echelon, like everyone else, is trying to capture some of that mindshare for their products. In this case, the product is a proprietary system of

www.dbooks.org

chips, protocols, and interfaces for enabling the IoT on industrial devices. But what struck me and my colleagues was how really outdated this approach seems, and how far it misses the point of the emerging IoT.

Although there are many different ways to describe the IoT, it is essentially a network of devices where all the devices:

1. Have local intelligence
2. Have a shared API so they can speak with each other in a useful way, even if they speak multiple protocols
3. Push and pull status and command information from the networked world

In the industrial context, rolling out a better networking chip is just a minor improvement on an unchanged 1980s practice that requires complex installations, including running miles of wire inside walls. This IzoT would have been a breakthrough product back when I got my first Sony Walkman.

This isn't a problem confined to Echelon. It's a problem shared by many industries. I just spent several days wandering the floor of the International Air Conditioning, Heating, and Refrigeration (AHR) trade show, a show about as far from CES as you can get: truck-sized boilers and cooling towers that look unchanged from their 19th-century origins dotted the floor. For most of the developed world, HVAC is where the lion's share of our energy goes. (That alone is reason enough to care about it.) It's also a treasure trove of data (think Nest), and a great place for all the obvious, sensible IoT applications like monitoring, control, and network intelligence. But after looking around at IoT products at AHR, it was clear that they came from Bizzaro World.[1]

I spoke with an engineer showing off one of the HVAC industry-leading low-power wireless networking technologies for data centers. He told me his company's new wireless sensor network system runs at 2.6 GHz, not 2.4 GHz (though the data sheet doesn't confirm or deny that). Looking at the products that won industry innovation awards, I was especially depressed. Take, for example, this variable-speed compressor technology. When you put variable-speed motors into an air-conditioning compressor, you get a 25–35% efficiency boost. That's a

1. I'm not sorry for the Super Friends reference.

lot of energy saved! And it looks just like variable-speed technology that has been around for decades—just not in the HVAC world.



Figure 8-2. Photo: Kipp Bradford. Copeland alone has sold more than 100 million compressors. That's a lot of "things" to get on the Internet.

Now here's where things get interesting: variable-speed control demands a smart processor controlling the compressor motor, plus the intelligent sensors to tell the motor when to vary its speed. With all that intelligence, I should be able to connect it to my network. So, when will my air conditioner talk to my Nest so I can optimize my energy consumption? Never. Or at least not until industry standards and government regulations overlap just right and force them to talk to each other, just as recent regulatory conditions forced 40-year-old variable-motor control technology into 100-year-old compressor technology.

Of course, like any inquisitive engineer, I had to ask the manufacturer of my high-efficiency boiler if it was possible to hook that up to my Nest to analyze performance. After he said, "Sure, try hooking up the thermostat wire," and made fun of me for a few minutes, the engineer said, "Yeah, if you can build your own modbus-to-wifi bridge, you can access our modbus interface and get your boiler online. But do you really want some Russian hacker controlling your house heat?" It would be so easy for my Nest thermostat to have a useful conversation with my boiler beyond "on/off," but it doesn't.

Don't get me wrong. I really am sympathetic to engineering within the constraints of industrial (versus consumer) environments, having

www.dbooks.org

spent a period of my life designing toys and another period making medical and industrial products. I've had to use my engineering skills to make things as dissimilar as Elmo and dental implants. But constraints are no excuse for trying to patch up, or hide behind, outdated technology.

The electrical engineers designing IoT devices for consumers have created exciting and transformative technologies like z-wave, Zigbee, BLE, wifi, and more, giving consumer devices robust and nearly transparent connectivity with increasingly easy installation. Engineers in the industrial world seem to be stuck making small technological tweaks that might enhance safety, reliability, robustness, and NSA-proof security of device networks. This represents an unfortunate increase in the bifurcation of the Internet of Things. It also represents a huge opportunity for those who refuse to submit to the notion that the IoT is either consumer or industrial, but never both.

For HVAC, innovation in 2014 means solving problems from 1983 with 1984 technology because the government told them so. The general attitude can be summed up as: "We have all the market share and high barriers to entry, so we don't care about your problems." Left alone, this industry (like so many others) will keep building walls that prevent us from having real control over our devices and our data. That directly contradicts a key goal of the IoT: connecting our smart devices together.

And that's where the opportunity lies. There is significant value in HVAC and similar industries for any company that can break down the cultural and technical barriers between the Internet of Things and the Industrial Internet of Things. Companies that recognize the new business models created by well-designed, smart, interconnected devices will be handsomely rewarded. When my Nest can talk to my boiler, air conditioner, and Hue lights, all while analyzing performance versus weather data, third parties could sell comfort contracts, efficiency contracts, or grid stabilization contracts.

As a matter of fact, we are already seeing a $16 billion market for grid stabilization services opened up by smart, connected heating devices. It's easy to envision a future where the electric company pays me during peak load times because my variable-speed air conditioner slows down, my Hue lights dim imperceptibly, and my clothes dryer pauses, all reducing my grid load—or my heater charges up a bank of thermal storage bricks in advance of a cold front before I return home from

work. Perhaps I can finally quantify the energy savings of the efficiency improvements that I make.

My Sony Walkman already went the way of the dodo, but there's still a chance to blow open closed industries like HVAC and bridge the IoT and the IIoT before my Beats go out of style.

www.dbooks.org

Trope or Fact? Technology Creates More Jobs than It Destroys

Jim Stogdill

NOTE

Editor's Note

We're trying something new here. I read this back-and-forth exchange between Malcolm Gladwell and Bill Simmons, and decided we should give it a try. Or, more accurately, since we're already having plenty of back-and-forth email exchanges like that, we just need to start publishing them. My friend Doug Hill, author of Not So Fast: Thinking Twice About Technology, agreed to be a guinea pig and chat with me about a subject that's on both of our minds (and a lot of other people's): technology and the jobless recovery. We'll be diving into this topic again next week in a debate hosted at Strata. This post was lightly edited on 2/6/14 for clarity.

STOGDILL: I saw this tweet over the holidays while I was reading your book. I mean, I literally got distracted by this tweet while I was reading your book:

www.dbooks.org



It felt like a natural moment of irony that I had to share with you. In the article Ari Gesher references in his tweet, Vivek Whadwa obviously has an optimistic point of view, and Gesher was right to call out the inconsistency of his claims with our jobless recovery. I also recently read George Packer's *The Unwinding*, his enlightening and disturbing look at the human stories behind our current malaise, and frankly it seems to better reflect the truth on the ground, at least if you get outside of the big five metro areas. But I suspect not a lot of techno optimists are spending time in places that won't get 4G LTE for another year or two.

I'm not going to ask you what you think of the article because I think I already know the answer. I do have a few things on my mind, though. Is our jobless recovery a new structural reality brought about by more and more pervasive automation? Are Norbert Wiener's cybernetic predictions from the late 1940s finally coming true? He spent the 1950s telling politicians, union leaders, and anyone else who would listen that robots were the slave labor of the future, and that free men can't compete with slaves for jobs. Or is creative destruction still working, but just taking some time to adjust this time around? And, if one is skeptical of technology, is it like being skeptical of tectonics? You can't change it, so bolt your house down?

HILL: Your timing is good. The day your email arrived the lead story in the news was the latest federal jobs report, which told us that the jobless "recovery" continues apace. Jobless, that is.

The national conversation about the impact of automation on employment continues apace, too. Thomas Friedman devoted his *New York Times* column a couple of days ago to *The Second Machine Age*, the new book by Erik Brynjolfsson and Andrew McAfee. They're the professors from MIT whose previous book, *Race Against the Machine*, helped start that national conversation, in part because it demonstrated both an appreciation of automation's advantages and an awareness that lots of workers could be left behind.

I'll try to briefly answer your questions, Jim, and then note a couple points that strike me as somewhat incongruous in the current discussion.

Is our jobless economy a new structural reality brought about by more and more pervasive automation?

Yes. Traditional economic theory holds that advances in technology create jobs rather than eliminate them. Even if that maxim were true in the past (and it's not universally accepted), many economists believe the pace of innovation in automation today is overturning it. This puts automation's most fervent boosters in the odd position of arguing that technological advance will disrupt pretty much everything except traditional economic theory.

Are Norbert Wiener's predictions from the late 1940s finally coming true? Or is creative destruction still working, but just taking some time to adjust this time around?

Yes to both. Wiener's predictions that automation would be used to undermine labor are coming true, and creative destruction is still at work. The problem is that we won't necessarily like what the destruction creates.

Now, about those incongruous points that bug me:

1. First, a quibble over semantics. It's convenient in our discussions about automation to use the word "robots," but also misleading. Much, if not most, of the jobs displacement we're seeing now is coming from systems and techniques that are facilitated by computers but less mechanical than the robots we typically envision assembling parts in factories. I don't doubt that actual robots will be an ever-more-important force in the future, but they'll be adding momentum to methods that corporations have been using for quite awhile now to increase productivity, even as they're reducing payrolls.

www.dbooks.org

2. It's commonly said that the answer to joblessness is education. Our employment problems will be solved by training people to do the sorts of jobs that the economy of the future will require. But wait a minute. If it's true that the economy of the future will increasingly depend on automation, won't we simply be educating people to do the sorts of jobs that eliminate more jobs?
3. Techno optimists argue that our current employment problems are merely manifestations of a transition period on the way to a glorious future. "Let the robots take the jobs," says Kevin Kelly, "and let them help us dream up new work that matters." Even on his own terms, the future Kelly envisions seems more nightmarish than dreamlike. Everyone agrees automation is going to grow consistently more capable. As it does, Kelly says, robots will take over every job, including whatever new jobs we dream up to replace the previous jobs we lost to robots. If he's right, we won't be dreaming up new work that matters because we want to, but because we'll have no choice. It will indeed be a race against the machines, and machines don't get tired.

One more thing. You asked if being skeptical of technology is like being skeptical of tectonics. My first thought was to wonder whether anybody is really skeptical of tectonics, but given the polls on global warming, I guess anything is possible — more possible, I think, than a reversal of the robot revolution. So yeah, go ahead and bolt the house down.

STOGDILL: Let me think where to start. My problem in this conversation is that I find myself arguing both sides of the question in my head, which makes it hard to present a coherent argument to you.

First, let me just say that I enter this discussion with some natural inclination toward a Schumpeterian point of view. In 1996, I visited a Ford electronics plant in Pennsylvania that was going through its own automation transformation. They had recently equipped the plant with then-new surface mount soldiering robots and redesigned the electronic modules that they produced there to take advantage of the tech. The remaining workers each tended two robots instead of placing parts on the boards themselves.

Except for this one guy. For some reason I've long forgotten, one of the boards they manufactured still required a single through-board capacitor, and a worker at that station placed capacitors in holes all day. Every 10 seconds for eight hours, a board would arrive in front of

him, he would drop a capacitor's legs through two little holes, push it over a bit to make sure it was all the way through, and then it was off to the next station to be soldiered. It was like watching Lucy in the Chocolate Factory.

I was horrified—but when I talked to him, he was bound and determined to keep that job from being automated. I simply couldn't understand it. Why wouldn't he want to upgrade his skills and run one of the more complex machines? Even now, when I'm more sympathetic to his plight, I'm still mystified that he could stand to continue doing that job when something else might have been available.

Yet, these days I find myself losing patience with the reflexive trope "of course technology creates more jobs than it destroys; it always has. What are you, a Luddite?" The Earth will keep rotating around the sun, too; it always has, right up until the sun supernovas, and then it won't.

Which isn't to say that the robots are about to supernova, but that arguments that depend on the past perpetuating into the future are not arguments—they're wishes. (See also, China's economy will always keep growing at a torrid rate despite over-reliance on investment at the expense of consumption because it always has). And I just can't really take anyone seriously who makes an argument like that if they can't explain the mechanisms that will continue to make it true.

So, to my thinking, this boils down to a few key questions. Was that argument even really true in the past, at least the recent past? If it was, is the present enough like the past that we can assume that, with retraining, we'll find work for the people being displaced by this round of automation? Or, is it possible that something structurally different is happening now? And, even if we still believe in the creative part of creative destruction, what destructive pace can our society absorb and are there policies that we should be enacting to manage and ease the transition?

This article does a nice job of explaining what I think might be different this time with its description of the "cognitive elite." As automation takes the next layer of jobs at the current bottom, we humans are asked to do more and more complex stuff, higher up the value hierarchy. But what if we can't? Or, if not enough of us can? What if it's not a matter of just retraining—what if we're just not talented enough? The result would surely be a supply/demand mismatch at the high end of the cognitive scale, and we'd expect a dumbbell shape to develop in our

www.dbooks.org

income distribution curve. Or, in other words, we'd expect new Stanford grads going to Google to make $100K and everyone else to work at Walmart. And more and more, that seems like it's happening.

Anyway, right now I'm all question, no answer. Others are suggesting that this jobless recovery has nothing to do with automation. It's the (lack of) unions, stupid. I really don't know, but I think we—meaning we technologists and engineers—need to be willing to ask the question "is something different this go-round?" and not just roll out the old history-is-future tropes.

We're trying to create that conversation at least a bit by holding an Oxford-style debate at our next Strata conference. The statement we'll be debating is: "Technology creates more jobs than it destroys," and I'll be doing my best to moderate in an even-handed way.

By the way, your point that it's "not just robots" is well taken. I was talking to someone recently who works in the business process automation space, and they've begun to refer to those processes as "robots," too—even though they have no physical manifestation. I was using the term in that broad sense, too.

HILL: In your last email you made two points in passing that I'd like to agree with right off the bat.

One is your comment that, when it comes to predicting what impact automation will have on employment, you find yourself "arguing both sides of the question." Technology always has and always will cut both ways, so we can be reasonably certain that, whatever happens, both sides of the question are going to come into play. That's about the only certainty we have, really, which is why I also liked it when you said, "I'm all question, no answer." That's true of all of us, whether we admit it or not.

We are obligated, nonetheless, to take what Norbert Wiener called "the imaginative forward glance." For what it's worth then, my answer to your question, "Is something different this go-round?"—by which you meant, even if it was once true that technological advancement created more rather than fewer jobs, that may no longer be true, given the pace, scale, and scope of the advances in automation we're witnessing today—is yes and no.

That is, yes, I do think the scope and scale of technological change we're seeing today presents us with challenges of a different order of magnitude than what we've faced previously. At the same time, I think

it's also true that we can look to the past to gain some sense of where automation might be taking us in the future.

In the articles on this issue you and I have traded back and forth over the past several weeks, I notice that two of the most optimistic, as far as our automated future is concerned, ran in the Washington Post. I want to go on record as denying any suspicion that Jeff Bezos had anything to do with that.

Still, the most recent of those articles, James Bessen's piece on the lessons to be learned from the experience of America's first industrial-scale textile factories ("Will robots steal our jobs? The humble loom suggests not") was so confidently upbeat that I'm sure Bezos would have approved. It may be useful, for that reason, to take a closer look at some of Bessen's claims.

To hear him tell it, the early mills in Waltham and Lowell, Massachusetts, were 19th-century precursors of the cushy working conditions enjoyed in Silicon Valley today. The mill owners recruited educated, middle-class young women from surrounding farm communities and supplied them with places to live, houses of worship, a lecture hall, a library, a savings bank, and a hospital.

"Lowell marked a bold social experiment," Bessen says, "for a society where, not so long before, the activity of young, unmarried women had been circumscribed by the Puritan establishment."

The suggestion that the Lowell mills were somehow responsible for liberating young women from the clutches of Puritanism is questionable — the power of the Puritan church had been dissipating for all sorts of reasons for more than a century before factories appeared on the banks of the Merrimack—but let that go.

It is true that, in the beginning, the mills offered young women from middle-class families an unprecedented opportunity for a taste of freedom before they married and settled down. Because their parents were relatively secure financially, they could afford to leave them temporarily behind without leaving them destitute. That's a long way from saying that the mills represented some beneficent "social experiment" in which management took a special interest in cultivating the well-being of the women they employed.

Thomas Dublin's *Women at Work: The Transformation of Work and Community in Lowell, Massachusetts, 1826-1860* tells a different story. Women were recruited to staff the mills, Dublin says, because they

www.dbooks.org

were an available source of labor (men were working the farms or employed in smaller-scale factories in the cities) and because they could be paid less than men. All supervisory positions in the mills were held by men. Also contrary to Bessen's contention, the women weren't hired because they were smart enough to learn specialized skills. Women tended the machines; they didn't run them. "To the extent that jobs did not require special training, strength or endurance, or expose operatives to the risk of injury," Dublin says, "women were employed."

How much time they had to enjoy the amenities supposedly provided by management is another question. According to Dublin, mill workers put in 12 hours a day, six days a week, with only three regular holidays a year. As the number of mills increased, so did the pressure to make laborers more productive. Speedups and stretch-outs were imposed. A speedup meant the pace of the machinery was increased, a stretch-out meant that each employee was required to tend additional pieces of machinery. Periodic cuts in piece wages were another fact of mill life.

Because of their middle-class backgrounds, and because they were accustomed to pre-industrial standards of propriety, the first generation of women felt empowered enough to protest these conditions, to little avail. Management offered few concessions, and many women left. The generation of women who replaced them were less likely to protest. Most had fled the Irish famine and had no middle-class homes to return to.

I go into this in some detail, Jim, because it's important to acknowledge what automation's fundamental purpose has always been: to increase management profits. Bold social experiments to benefit workers haven't figured prominently in the equation.

It's true that factory jobs have, in the long run, raised the standard of living for millions of workers (the guy you met in the Ford electronics plant comes to mind), but we shouldn't kid ourselves that they've necessarily been pleasant, fulfilling ways to make a living. Nor should we kid ourselves that management won't use automation to eliminate jobs in the future, if automation offers opportunities to increase profits.

We also need to consider whether basing our economy on the production and sale of ever-higher piles of consumables, however they're manufactured, is a model the planet can sustain any longer. That's the essential dilemma we face, I think. We must have jobs, but they have to be directed toward some other purpose.

I realize I haven't addressed, at least directly, any of the questions posed in your email. Sorry about that—the Bessen article got under my skin.

STOGDILL: That Bessen article did get under your skin, didn't it? Well, anger in the face of ill-considered certainty is reasonable as far as I'm concerned. Unearned certainty strikes me as the disease of our age.

Reading your response, I had a whole swirl of things running through my head. Starting with, "Does human dignity require meaningful employment?" I mean, separate from the economic considerations, what if we're just wired to be happier when we grow or hunt for our own food?—and will the abstractions necessary to thrive in an automation economy satisfy those needs?

Also, with regard to your comments about how much stuff do we need—does that question even really matter? Is perpetual sustainability on a planet where the Second Law of Thermodynamics holds sway even possible? Anyway, that diversion can wait for another day.

Let me just close this exchange by focusing for just one moment on your point that productivity gains have always been about increasing management profits. Of course they have. I don't think that has ever been in question. Productivity gains are where every increase in wealth ever has come from (except for that first moment when someone stumbled on something useful bubbling out of the ground), and profit is how is how we incent investment in productivity. The question is how widely gains will be shared.

Historically, that argument has been about the mechanisms (and politics) to appropriately distribute productivity gains between capital and labor. That was the fundamental argument of the 20th century, and we fought and died over it—and for maybe 30 years, reached maybe a reasonable answer.

But what if we are automating to a point where there will be no meaningful link between labor and capital? There will still be labor, of course, but it will be doing these abstract "high value" things that have nothing whatsoever to do with the bottom three layers of Maslov's hierarchy. In a world without labor directly tied to capital and its productivity gains, can we expect the mechanisms of the 20th century to have any impact at all? Can we even imagine a mechanism that flows the value produced by robots to the humans they sidelined? Can unemployed people join a union?

www.dbooks.org

We didn't answer these questions, but thanks for exploring them with me.

Architecture, Design, and the Connected Environment

Andy Fitzgerald

Just when it seems we're starting to get our heads around the mobile revolution, another design challenge has risen up fiercer and larger right behind it: the Internet of Things. The rise in popularity of "wearables" and the growing activity around NFC and Bluetooth LE technologies are pushing the Internet of Things increasingly closer to the mainstream consumer market. Just as some challenges of mobile computing were pointedly addressed by responsive web design and adaptive content, we must carefully evaluate our approach to integration, implementation, and interface in this emerging context if we hope to see it become an enriching part people's daily lives (and not just another source of anger and frustration).

It is with this goal in mind that I would like to offer a series of posts as one starting point for a conversation about user interface design, user experience design, and information architecture for connected environments. I'll begin by discussing the functional relationship between user interface design and information architecture, and by drawing out some implications of this relationship for user experience as a whole.

In follow-up posts, I'll discuss the library sciences origins of information architecture as it has been traditionally practiced on the Web, and situate this practice in the emerging climate of connected environments. Finally, I'll wrap up the series by discussing the cognitive challenges that connected systems present and propose some specific measures we can take as designers to make these systems more pleasant, more intuitive, and more enriching to use.

www.dbooks.org

Architecture and Design

Technology pioneer Kevin Ashton is widely credited with coining the term "The Internet of Things." Ashton characterizes the core of the Internet of Things as the "RFID and sensor technology [that] enables computers to observe, identify, and understand the world—without the limitations of human-entered data."

About the same time that Ashton gave a name to this emerging confluence of technologies, scholar N. Katherine Hayles noted in *How We Became Posthuman* that "in the future, the scarce commodity will be the human attention span." In effect, collecting data is a technology problem that can be solved with efficiency and scale; making that mass of data meaningful to human beings (who evolve on a much different timeline) is an entirely different task.

The twist in this story? Both Ashton and Hayles were formulating these ideas circa 1999. Now, 14 years later, the future they identified is at hand. Bandwidth, processor speed, and memory will soon be up to the task of ensuring the technical end of what has already been imagined, and much more. The challenge before us now as designers is in making sure that this future-turned-present world is not only technically possible, but also practically feasible—in a word, we still need to solve the usability problem.

Fortunately, members of the forward guard in emerging technology have already sprung into action and have begun to outline the specific challenges presented by the connected environment. Designer and strategist Scott Jenson has written and spoken at length about the need for open APIs, flexible cloud solutions, and the need for careful attention to the "pain/value" ratio. Designer and researcher Stephanie Rieger likewise has recently drawn our collective attention to advances in NFC, Android intent sharing, and behavior chaining that all work to tie disparate pieces of technology together.

These challenges, however, lie primarily on the "computer" side of the Human Computer Interaction (HCI) spectrum. As such, they give us only limited insight into how to best accommodate Hayles' "scarce commodity"–the human attention span. By shifting our point of view from how machines interact with and create information to the way that humans interact with and consume information, we will be better equipped to make the connections necessary to create value for indi-

viduals. Understanding the relationship between architecture and design is an important first step in making this shift.



Figure 10-1. Image by Dan Klyn, used with permission.

Information Architect Dan Klyn explains the difference between architecture and design with a metaphor of tailoring: the architect determines where the cuts should go in the fabric, the designer then brings those pieces together to make the finished product the best it can be, "solving the problems defined in the act of cutting."

Along the way, the designer may find that some cuts have been misplaced–and should be stitched back together or cut differently from a new piece. Likewise, the architect remains active and engaged in the design phase, making sure each piece fits together in a way that supports the intent of the whole.

The end result–be it a well-fitted pair of skinny jeans or a user interface–is a combination of each of these efforts. As Klyn puts it, the architect specializes in determining *what* must be built and in determining the overall structure of the finished product; the designer focuses on how to put that product together in a way that is compelling and effective within the constraints of a given context.

Once we make this distinction clear, it becomes equally clear that user interface design is a context-specific *articulation* of an underlying information architecture. It is this IA foundation that provides the direct connection to how human end users find value in content and functionality. The articulatory relationship between architecture and design creates consistency of experience across diverse platforms and

www.dbooks.org

works to communicate the underlying information model we've asked users to adopt.

Let's look at an example. The early Evernote app had a very different look and feel on iOS and Android. On Android, it was a distinctly "Evernote-branded" experience. On iOS, on the other hand, it was designed to look more like a piece of the device operating system.

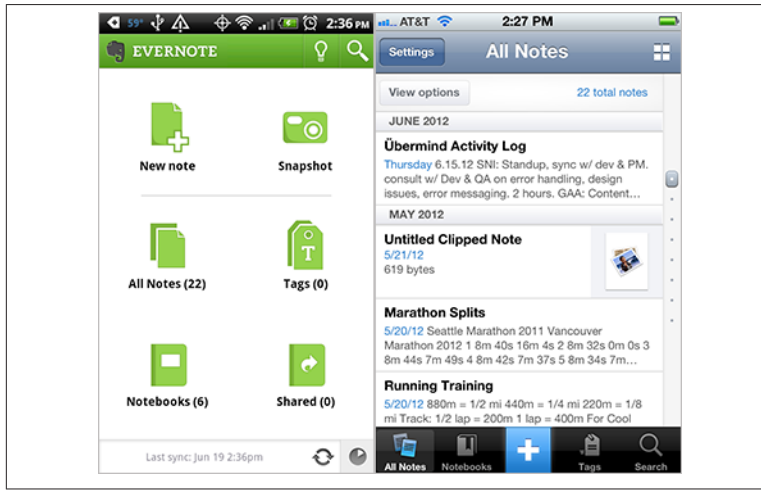


Figure 10-2. Evernote screenshots, Android (left) versus iOS.

Despite the fact that these apps are aesthetically different, their architectures are consistent across platforms. As a result, even though the controls are presented in different ways, in different places, and at different levels of granularity, moving between the apps is a cognitively seamless experience for users.

In fact, apps that "look" the same across different platforms sometimes end up creating architectural inconsistencies that may ultimately confuse users. This is most easily seen in the case of "ported applications," where iPhone user interfaces are "ported over" whole cloth to Android devices. The result is usually a jumble of misplaced back buttons and errant tab bars that send mixed messages about the effects of native controls and patterns. This, in turn, sends a mixed message about the information model we have proposed. The link between concept-rooted architecture and context-rooted design has been lost.

In the case of such ports, the full implication of the articulatory relationship between information architecture and user interface becomes

clear. In these examples, we can see that information architecture always happens: either it happens by design or it happens by default. As designers, we sometimes fool ourselves into thinking that a particular app or website "doesn't need IA," but the reality is that information architecture is always present — it's just that we might have specified it in a page layout instead of a taxonomy tool (and we might not have been paying attention when that happened).

Once we step back from the now familiar user interface design patterns of the last few years and examine the information architecture structures that inform them, we can begin to develop a greater awareness (and control) of how those structures are articulated across devices and contexts. We can also begin to cultivate the conditions necessary for that articulation to happen in terms that make sense to users in the context of new devices and systems, which subsequently increases our ability to capitalize on those devices' and systems' unique capabilities.

This basic distinction between architecture and design is not a new idea, but in the context of the Internet of Things, it does present architects and designers with a new set of challenges. In order to get a better sense of what has changed in this new context, it's worth taking a closer look at how the traditional model of IA for the web works. This is the topic to which I'll turn in my next post.

www.dbooks.org

iBeacon Basics

Matthew Gast

As any programmer knows, writing the "hello, world" program is the canonical elementary exercise in any new programming language. Getting devices to interact with the world is the foundation of the Internet of Things, and enabling devices to learn about their surroundings is the "hello world" of mobility.

On a recent trip to Washington D.C., I attended the first DC iBeacon Meetup. iBeacons are exciting. Retailers are revolutionizing shopping by applying new indoor proximity technologies and developing the physical world analog of the data that a web-based retailer like Amazon can routinely collect. A few days ago, I tweeted about an analysis of the beacon market, which noted that "[beacons] are poised to transform how retailers, event organizers, transit systems, enterprises, and educational institutions communicate with people indoors"—and could even be used in home automation systems.

I got to see the ground floor of the disruption in action at the meetup in DC, which featured presentations by a few notable local companies, including Radius Networks, the developer of the CES scavenger hunt app for iOS. When I first heard of the app, I almost bought a ticket to Las Vegas to experience the app for myself, so it was something of a cool moment to hear about the technology from the developer of an application that I'd admired from afar.

After the presentations, I had a chance to talk with David Helms of Radius. Helms was drawn to work at Radius for the same reason I was compelled to attend the iBeacon meetup. As he put it, "The first step in extending the mobile computing experience beyond the confines of that slab of glass in your pocket is when it can recognize the world

www.dbooks.org

around it and interact with it, and proximity is the 'Hello' of the Internet of Things revolution."

For such a source of excitement in the industry, proximity is a simple protocol. Periodic "beacon" transmissions contain three numbers: (1) a UUID, which is typically used for an organizationally unique identifier, such as a company, (2) a major number, and (3) a minor number. Major and minor numbers are assigned in whatever way the organization desires to use them. A typical usage would be in a company with many stores: the major number would be used for a store number and the minor number might be used for a department or display.

With such a limited protocol, most of the exciting work happens inside an application running on your mobile device. Beacons themselves have no ability to learn about geographic coordinates like latitude/longitude or street addresses; to translate the three numbers in a beacon's broadcast into a location, you need to pull from something else, such as a web service. They are not able to find out what else is in the neighborhood because they are transmit-only devices. It is possible to use a beacon's broadcasts to perform "ranging," or estimating the distance between a device and a beacon, but that is a function performed by the underlying operating system and cannot be done in a power-efficient manner.

The CES scavenger hunt application is a simple listener that implements a "virtual punch card." Nine locations on the CES show floor were equipped with beacons, and conference attendees who visited all nine locations completed a virtual card in the application and were entitled to receive a gift from the show. Each of the nine beacons had a unique signature, decoded by Alasdair Allan and Sandeep Mistry in Make magazine.

With such a simple protocol, you don't need much to get started. Bluetooth 4.0 is the first version that supports low-energy beacons, and hardware is widely available. You can use a dedicated USB device plugged into a computer or a small computer like the Raspberry Pi, a self-contained USB device like Radius' RadBeacon, or software running on a computing device that has the appropriate Bluetooth hardware. Radius produces a free tool for iOS called Locate for iBeacon as well as an inexpensive Mac app called MacBeacon.



Figure 11-1. Radius' RadBeacon, attached to an Apple USB charger.

At the DC meetup, Radius was giving away RadBeacons, and I was lucky enough to win one. It's a tiny device. In this photo, the RadBeacon is plugged into an Apple USB charger, and you can see it dwarfed by the electronics to supply power through the USB interface.

Configuring the RadBeacon is a snap. Apple's iBeacon specifications don't prescribe a method of configuration, so some may be configured by their USB hosts. The RadBeacon is configured over Bluetooth. Although the iBeacon spec is transmit-only, the management can take place using two-way communications over Bluetooth.

RadBeacon's configuration screen allows you to set the three numbers in the beacon tuple (UUID, major, and minor). Illustrated in the screen shot below, I used the article from *Make* to put in the beacon details from one of the beacons the CES app looks for. As a result, if you open the CES app, you'll be virtually on the show floor in Las Vegas, and you can even earn badges. When I set up the RadBeacon in a hotel room recently, it took me about five minutes to unpack the device, get it running, and pretend to be on the floor in Las Vegas.

www.dbooks.org

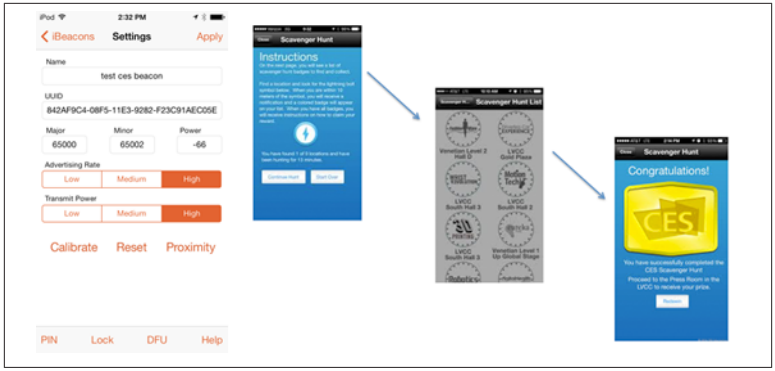


Figure 11-2. Left: CES beacon settings; Right: CES Scavenger Hunt.

With everything in a beacon being totally open to anybody with a receiver, security is not a strong suit of the protocol. (If you want to investigate your Bluetooth neighborhood, download and install the free LightBlue app from the iOS app store.) As David Helms noted, the security model is "essentially worthless," and he advocated that developers "design solutions on the assumption that the beacon information is accessible to third parties." Defensive tactics depend on the value of the information in the app; for the CES scavenger hunt, Radius didn't implement extensive security measures, Helms explained, but instead they used "an audit trail that validated the game progress of the players based on time and location to ensure reasonably fair game play, which was a cost-effective strategy to meet the goal." With more at stake, Helms said, developers might need to include more sophisticated security mechanisms to address the potential for interception, spoofing, and the limited ability to ensure that a client device can be a trusted computing platform.

Although the technology of iBeacon is not complex, incorporating proximity into existing or new applications may be a challenge. In the past decade, I've seen many organizations try to use Wi-Fi for proximity, which is at best a rough guess. Helms pointed out that more accurate location information based on beacons allows application developers to "design amazing experiences for their customers and users by drawing on deep knowledge of their particular market."

Of the many things I've learned from O'Reilly, one of the most important is to follow the developers. I started learning about BLE and beacons because of the excitement and exchange of ideas. In our conversation, David confessed to being "blown away with the amazing

new ideas and projects using proximity and iBeacons that thousands of developers are building every day." I am totally with him on that point, and I'll be returning to the use of proximity information in application development as I continue to learn more.

www.dbooks.org

The Lingering Seduction of the Page

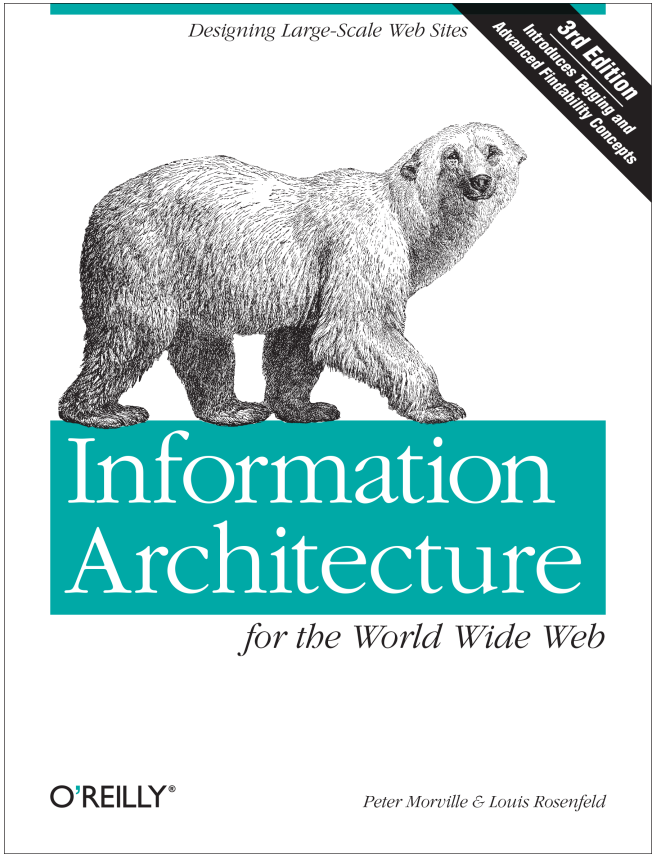
Andy Fitzgerald

In an earlier post in this series, I examined the articulatory relationship between information architecture and user interface design, and argued that the tools that have emerged for constructing information architectures on the Web will only get us so far when it comes to expressing information systems across diverse digital touchpoints. Here, I want to look more closely at these traditional web IA tools in order to tease out two things: (1) ways we might rely on these tools moving forward, and (2) ways we'll need to expand our approach to IA as we design for the Internet of Things.

First Stop: The Library

The seminal text for Information Architecture as it is practiced in the design of online information environments is Peter Morville's and Louis Rosenfeld's *Information Architecture for the World Wide Web*, affectionately known as "The Polar Bear Book."

www.dbooks.org



First published in 1998, The Polar Bear Book gave a name and a clear, effective methodology to a set of practices many designers and developers working on the Web had already begun to encounter. Morville and Rosenfeld are both trained as professional librarians and were able to draw on this time-tested field in order to sort through many of the new information challenges coming out of the rapidly expanding Web.

If we look at IA as two faces of the same coin, The Polar Bear Book focuses on the largely top-down "Internet Librarian" side of information design. The other side of the coin approaches the problems posed by data from the bottom up. In *Everything is Miscellaneous: The Power of the New Digital Disorder*, David Weinberger argues that the fundamental problem of the "second order" (think "card catalogue") organization typical of library sciences-informed approaches is that they fail to recognize the key differentiator of digital information: that it can exist in multiple locations at once, without any single location

being the "home" position. Weinberger argues that in the "third order" of digital information practices, "understanding is metaknowledge." For Weinberger, "we understand something when we see how the pieces fit together."

Successful approaches to organizing electronic data generally make liberal use of both top-down and bottom-up design tactics. Primary navigation (driven by top-down thinking) gives us a birds-eye view of the major categories on a website, allowing us to quickly focus on content related to politics, business, entertainment, technology, etc. The "You May Also Like" and "Related Stories" links come from work in the metadata-driven bottom-up space.

On the Web, this textually mediated blend of top-down and bottom-up is usually pretty successful. This is no surprise: the Web is, after all, primarily a textual medium. At its core, HTML is a language for marking up text-based documents. It makes them interpretable by machines (browsers) so they can be served up for human consumption. We've accommodated images and sounds in this information ecology by marking them up with tags (either by professional indexers or "folksonomically," by whomever cares to pitch in).

There's an important point here that often goes without saying: the IA we've inherited from the Web is textual—it is based on the perception of the world mediated through the technology of writing; herin lies the limitation of the IA we know from the web as we begin to design for the Internet of Things.

Reading Brains

We don't often think of writing as "technology," but inasmuch as technology constitutes the explicit modification of techniques and practices in order to solve a problem, writing definitely fits the bill. Language centers can be pinpointed in the brain—these are areas programmed into our genes that allow us to generate spoken language—but in order to read and write, our brains must create new connections not accounted for in our genetic makeup.

In *Proust and the Squid*, cognitive neuroscientist Maryanne Wolf describes the physical, neurological difference between a reading and writing brain and a pre-literate linguistic brain. Wolf writes that, with the invention of reading "we rearranged the very organization of our brain." Whereas we learn to speak by being immersed in language,

www.dbooks.org

learning to read and write is a painstaking process of instruction, practice, and feedback. Though the two acts are linked by a common practice of language, writing involves a different cognitive process than speaking. It is one that relies on the technology of the written word. This technology is not transmitted through our genes; it is transmitted through our culture.

It is important to understand that writing is not simply a translation of speech. This distinction matters because it has profound consequences. Wolf writes that "the new circuits and pathways that the brain fashions in order to read become the foundation for being able to think in different, innovative ways." As the ability to read becomes widespread, this new capacity for thinking differently, too, becomes widespread.

Though writing constitutes a major leap past speech in terms of cognitive process, it shares one very important common trait with spoken language: linearity. Writing, like speech, follows a syntagmatic structure in which meaning is constituted by the flow of elements in order —and in which alternate orders often convey alternate meanings.

When it comes to the design of information environments, this linearity is generally a foregone conclusion, a feature of the cognitive landscape which "goes without saying" and is therefore never called into question. Indeed, when we're dealing primarily with text or text-based media, there is no need to call it into question.

In the case of embodied experience in physical space, however, we natively bring to bear a perceptual apparatus which goes well beyond the linear confines of written and spoken language. When we evaluate an event in our physical environment—a room, a person, a meaningful glance—we do so with a system of perception orders of magnitude more sophisticated than linear narrative. JJ Gibson describes this as the perceptual awareness resulting from a "flowing array of stimulation." When we layer on top of that the non-linear nature of dynamic systems, it quickly becomes apparent that despite the immense gains in cognition brought about by the technology of writing, these advances still only partially equip us to adequately navigate immersive, physical connected environments.

The Trouble with Systems (and Why They're Worth It)

I have written elsewhere in more detail about challenges posed to linguistic thinkers by systems. To put all of that in a nutshell, complex systems baffle us because we have a limited capacity to track system-influencing inputs and outputs and system-changing flows. As systems thinking pioneer Donella Meadows characterizes them in her book *Thinking in Systems: A Primer*, self-organizing, nonlinear, feedback systems are "inherently unpredictable" and "understandable only in the most general way."

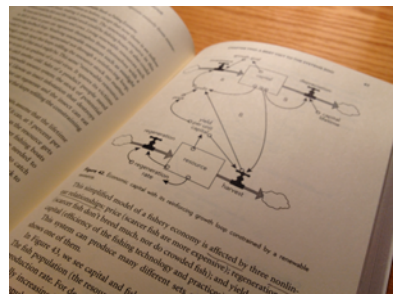


Figure 12-1. Photo: Andy Fitzgerald, of content from Thinking in Systems: A Primer, by Donella Meadows.

According to Meadows, we learn to navigate systems by constructing models that approximate a simplified representation of the system's operation and allow us to navigate it with more or less success. As more and more of our world—our information, our social networks, our devices, and our interactions with all of these—becomes connected, our systems become increasingly difficult (and undesirable) to compartmentalize. They also become less intrinsically reliant on linear textual mediation: our "smart" devices don't need to translate their messages to each other into English (or French or Japanese) in order to interact.

This is both the great challenge and the great potential of the Internet of Things. We're beginning to interact with our built information environments not only in a classically signified, textual way, but also in a physical-being-operating-in-the-world kind of way. The text remains—and the need to interact with that textual facet with the tools

www.dbooks.org

we've honed on the Web (i.e., traditional IA) remains. But as the information environments we're tasked with representing become less textual and more embodied, the tools we use to represent them must likewise evolve beyond our current text-based solutions.

Fumbling toward System Literacy

In order to rise to meet this new context, we're going to need as many semiotic paths as we can find—or create. And in order to do that, we will have to pay close attention to the cognitive support structures that normally "go without saying" in our conceptual models.

This will be hard work. The payoff, however, is potentially revolutionary. The threshold at which we find ourselves is not merely another incremental step in technological advancement. The immersion in dynamic systems that the connected environment foreshadows holds the potential to re-shape the way we think—the way our brains are "wired"—much as reading and writing did. Though mediated by human-made, culturally transmitted technology (e.g., moveable type, or, in this case, Internet protocols), these changes hold the power to affect our core cognitive process, our very capacity to think.

What this kind of "system literacy" might look like is as puzzling to me now as reading and writing must have appeared to pre-literate societies. The potential of being able to grasp how our world is connected in its entirety—people, social systems, economies, trade, climate, mobility, marginalization—is both mesmerizing and terrifying. Mesmerizing because it seems almost magical; terrifying because it hints at how unsophisticated and parochial our current perspective must look from such a vantage point.

As information architects and interface designers, all of this means that we're going to have to be nimble and creative in the way we approach design for these environments. We're going to have to cast out beyond the tools and techniques we're most comfortable with to find workable solutions to new problems of complexity. We aren't the only ones working on this, but our role is an important one: engineers and users alike look to us to frame the rhetoric and usability of immersive digital spaces. We're at a major transition in the way we conceive of putting together information environments. Much like Morville and Rosenfeld in 1998, we're "to some degree all still making it up as we go along." I don't pretend to know what a fully developed information architecture for the Internet of Things might look like, but in the spirit

of exploration, I'd like to offer a few pointers that might help nudge us in the right direction—a topic I'll tackle in my next post.

www.dbooks.org

Let There Be (Intelligent) Light

Glen Martin

New technologies often manifest their most dramatic effects through things that are commonplace, even prosaic. Consider the electric light: it's ubiquitous and, well, boring. But meld it with some modern technology and you get intelligent lighting—wirelessly networked LED lights augmented by software and sensors.

Early adopters have included creators of Las Vegas shows and productions, but in the big picture, entertainment is a mere sideshow. Intelligent lighting's greatest impacts will be in the commercial and industrial sectors: warehouses, office buildings, factories, cold storage plants, hospitals—any place that encompasses large spaces and employs a lot of lights.

That's because smart lighting is highly efficient lighting. A PG&E study conducted at a 44,800 square-foot Ace Hardware distribution center in Rocklin, California, confirmed that an intelligent LED system used up to 93% less energy than the "dumb" metal halide lights that formerly lit the building. And an Escondido, California, brewery outfitted a new building addition with an intelligent lighting system that uses 86% less energy than the T8 fluorescent fixtures specified in the original design. According to the case study, the LED system will secure project payback in less than two years; avoids the re-lamping, re-ballasting and mercury disposal costs that are an inevitable corollary to high-intensity fluorescent lights; and contributes dramatically toward the addition's LEED Silver rating under the U.S. Green Building Council's standards.

"This is really about creating an 'Internet of light,'" says Allison Parker of Digital Lumens, a smart lighting firm. "It's about harvesting the huge

www.dbooks.org

amount of data generated by people working in a building where light is required, and using that information to both support their needs and maximize efficiencies at many levels."

Digital Lumens, which has already installed smart lights in more than 100 million square feet of space, employs a proprietary technology called LightRules, which integrates power data from all other non-lighting systems and circuits for a comprehensive portrait of a building's energy use. This information can be used to extrapolate the ways people interact with each other, the tools of their trades, and the buildings that surround them. Often, the data reveal errors in earlier suppositions—and provide unsuspected opportunities.

"One of our clients, Atlas Box & Crating, was contemplating buying a new baling machine," says Parker. "That's an expensive piece of capital equipment. While commercial spaces often have building energy management systems, most industrial spaces don't. So having that capability in LightRules, which can serve as an energy dashboard, was invaluable to Atlas. A manager merely looked at the energy data from the existing system usage patterns and identified a shift when their baling machine was largely unused. Instead of buying more equipment, they were able to rearrange shifts to meet their goals."

According to Digital Lumens case studies, a number of their clients have realized drastically reduced energy bills via data gleaned from the LightRules system. For instance, Creed Monarch, a Connecticut-based producer of precision-machined alloy components, reported an immediate 75% savings in lighting energy costs at its manufacturing facility, with ultimate savings projected at 90%. Likewise, Associated Grocers of New England and Ben E. Keith Foods in Fort Worth each saved 90%, and Vector Aerospace Helicopter Services of Richmond, British Columbia, reported 72% savings at its aircraft maintenance facility.

Energy cost savings are achieved through analysis of data gathered by the LightRules system, which measures key metrics such as localized energy use in an aisle, zone, or room; occupancy patterns; and the temperature across the facility. Managers also have access to an interactive map of their facility that offers up insights such as high lighting-use areas, where daylight light is harvested most, peak activity times, and projected monthly savings that could be realized through various modifications to the lighting program.

As with many technology changes in the workplace, there are behavioral science considerations as well. Workers accustomed to "dumb" lights may need a period of adjustment to cope with an intelligent system, Parker observes.

"It's a kind of red-carpet effect," Parker says. "If the lights go on as you enter a room and go off as soon as you pass, it can be disconcerting. So we advise managers to program a two or three minute delay before the lights go off. That can be reduced as people get used to the system."

Intelligent lights can also be tweaked to coordinate with existing ambient light from outdoor sources, and to dovetail with the circadian rhythms of company staffers.

"With our systems, you can adjust the light intensity and color spectrum to mimic the progression of natural light," Parker explains. "For example, at the beginning of the day, the lights can be bright and clear white-yellow, like the morning sun. That's stimulating, energizing light. As the day progresses, the lighting can gradually shift to the warmer, dimmer, more golden part of the spectrum—the part that corresponds to the afternoon sun."

Further, intelligent lights can preserve more than your sense of well-being.

"Vaccines, beer, produce, and meat can degrade under high UV lights," says Parker. "So, you might adjust your lights to express more UV in the workplace, but minimize it where susceptible products are stored."

Intelligent lighting has just begun making inroads into the economy. That presents great opportunities, in that there's a lot of low-hanging fruit waiting to be picked. But it also poses challenges. For one thing, most executives simply don't understand the full significance of lighting to their bottom lines.

"If you're a technophile, it can be easy to forget just how uninstrumented large parts of the economy are," says Parker. "What we're really talking about is creating a new infrastructure. At this point, the biggest challenge is getting through the financial gauntlet, educating the top people in the commercial and industrial sectors and government so they're willing to make the necessary investments. That means we're going to have to frame the message in different ways to tailor it specifically for each segment—Class A office, deep industrial, health care. The data, after all, is unique to each sector, so the ways the message is presented will be different."

www.dbooks.org

Still, the core of the message—vastly enhanced efficiencies—translates across all sectors, adds Parker. As an example, she cites heavy industry.

"How do you get industrial production to come back to the United States? Practically speaking, by shaving pennies until it makes economic sense. And one of the best ways to do that is through lighting. To a very large degree, once you establish really dramatic energy savings, many other benefits follow."

But refining messaging isn't enough, Parker adds: the lighting industry itself must be willing to change.

"Lighting has never had a standards organization," she says. "We're not used to changing bits; we're not really good at playing well with others across broad applications. Obviously, that's going to have to change. I see the connected world as a wheel, and the lighting industry as a spoke in that wheel—one of many spokes. For that wheel to function properly, all the spokes have to be connected to it, and they all have to work in concert with one another."

You can see the LightRules system in action at our Solid Conference demo pavilion.

About the Authors

Jon Bruner is a data journalist who approaches questions that interest him by writing and coding. Before coming to O'Reilly, where he is editor-at-large and co-chair of the Solid program, he was data editor at *Forbes Magazine*. He lives in New York, where he can occasionally be found at the console of a pipe organ.

Glen Martin covered science and environment for the *San Francisco Chronicle* for 17 years and has freelanced to more than 50 magazines, including *Wired*, *Audubon*, *Discover*, *The Utne Reader*, *Men's Journal*, *Science Digest*, *National Wildlife*, *BBC Wildlife*, *Outside*, *Sierra*, *The Financialist*, *Reader's Digest*, and publications for the University of California at Berkeley, Stanford University, Notre Dame University, the University of Colorado, and San Francisco State University. His latest book, *Game Changer: Animal Rights and the Fate of Africa's Wildlife*, was published in 2012 by the University of California Press.

Matthew Gast is the director of software product management at Aerohive Networks. He has been active within the Wi-Fi community and has served as a leader on several industry standards committees, including as chair of the current revision of the 802.11 standard. He has written extensively about Wi-Fi, including three books for O'Reilly. In his spare time, Matthew is typically near an airport flying gliders.

Tim O'Reilly is the founder and CEO of O'Reilly Media Inc. Considered by many to be the best computer book publisher in the world. O'Reilly Media also hosts conferences on technology topics, including the O'Reilly Open Source Convention, Strata: The Business of Data, the Velocity Conference on Web Performance and Operations, and many others. Tim's blog, the O'Reilly Radar "watches the alpha geeks" to determine emerging technology trends, and serves as a platform for advocacy about issues of importance to the technical community. Tim is also a partner at O'Reilly AlphaTech Ventures, O'Reilly's early stage venture firm, and is on the board of Safari Books Online, PeerJ, Code for America, and Maker Media, which was recently spun out from O'Reilly Media. Maker Media's Maker Faire has been compared to the West Coast Computer Faire, which launched the personal computer revolution.

Kipp Bradford is an educator, technology consultant, and entrepreneur with a passion for creating new products as well as finding new applications for existing technologies. He was the founder or cofound-

www.dbooks.org

er of start-ups in the fields of transportation, consumer products, HVAC, and medical devices, and holds numerous patents for his inventions. Kipp cofounded Revolution By Design, Inc, a non-profit education and research organization dedicated to empowerment through technology and co-organizes Rhode Island's mini Maker Faire. As the Senior Design Engineer and Lecturer at the Brown University School of Engineering, Kipp teaches several engineering design and entrepreneurship courses. He is the chair of the Rhode Island Entrepreneurship Faculty group and serves on the boards of The Steel Yard and AS220. He is also on the technical advisory board of *MAKE Magazine* and is a Fellow at the College of Design, Engineering and Commerce at Philadelphia University.

A lifelong technology practitioner **Jim Stogdill** is finding this media thing ridiculously fun. In a previous life he traveled the world with the U.S. Navy. Unfortunately from his vantage point it all looked like the inside of a submarine. He spends his free time hacking silver halides with decidedly low-tech gear.

Andy Fitzgerald is an Associate Design Director and the User Experience Competency Lead at Deloitte Digital in Seattle. Andy has spent the better part of a decade massaging truculent bits of information into difficult digital spaces. His more recent work focuses on mobile and on designing for effective experiences across diverse digital touchpoints.