Long Tail Hardware: Turning Device Concepts Into Viable Low Volume Products

Steve Hodges and Nicholas Chen
Microsoft Research

At a time when the technology industry is embracing powerful new algorithms and cloud computing, continued innovation in hardware is essential. In addition to the growing storage and computation requirements of data centers and edge computers, hardware devices provide the critical gateway through which our systems receive input and provide output. Whether it is intentional user interaction, continuous context sensing, situated information display, environmental monitoring, or industrial control, we are more dependent on interactive and embedded hardware products than ever. Indeed, as the Internet of Things grows, many predict a dramatic growth in both the number and type of such devices. A key factor in this growth is the ability of those working in hardware to develop innovative devices with new forms and functions.

Hardware development can be split into two phases: first, a period of ideation, prototyping and design iteration leads to new device concepts; and then fruitful concepts transition beyond the basic prototype. The latter phase typically involves creating hundreds or thousands of copies of a prototype—either preproduction evaluation samples or a fully fledged low-volume product. The hardware device research community and the industry it serves have developed many tools and techniques to aid in the ideation, prototyping, and design iteration phase mentioned above. However, based on our first-hand experience coupled with the observation of others, we frequently see a bottleneck in the subsequent phase—the transition from a working prototype to a viable product. The challenge of this bottleneck, its root causes, and the potential benefits of overcoming it are the subject of this column.

Digital Object Identifier 10.1109/MPRV.2019.2947966
Date of current version 21 January 2020.

INTRODUCING THE LONG TAIL

Before examining the challenges of creating a viable interactive or embedded hardware product, it is insightful to understand how to characterize the success of different products within a

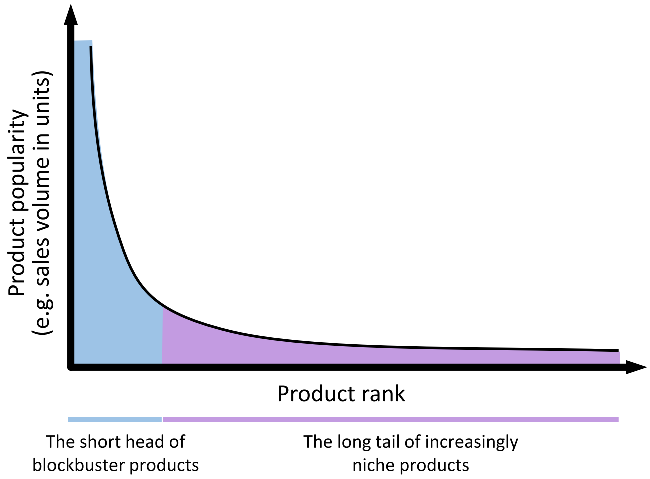

Figure 1. Plotting a rank-frequency distribution often demonstrate an inverse power law relationship with its characteristic short head and long tail.[1]

market. A common approach is to plot a graph of the sales volume of each different product in that market (on the *y*-axis) against product rank (on the *x*-axis). This is called a rank frequency distribution. Zipf's law[1] tells us that this distribution often follows an inverse power law, as illustrated in Figure 1. Such a distribution can be split into two parts: the "long tail" to the right represents a large number of niche (sometimes called boutique or custom) items, each of which is sold in relatively small quantities; whereas the "head" at the left contains just a few very popular items (the blockbusters), each of which is sold in large volumes.

Market dynamics—factors like customer demand, availability of supply, pricing, and competition—affect not only the total quantity of products sold but also the shape of the distribution. In some markets, the tail is particularly heavy, meaning that in aggregate it amounts to a lot of units. In other markets, the tail is truncated. A good example of a truncated market is the traditional movie industry. Here, there is a virtuous cycle of positive feedback where established film studios with prearranged global theatre distribution produce blockbusters almost as a matter of course. This generates large revenues to invest in future productions that are also likely to be successful—a process that economists call preferential attachment. However, at the other end of the spectrum would-be niche film producers struggle to find any outlet for their films, creating a sharp cut-off in the rank-frequency distribution.

TECHNOLOGY SHAPES THE TAIL

In microeconomic theory, it is recognized that technology affects the dynamics in a market.[2] For example, the growth of internet search and online shopping has given rise to new ways for consumers to discover and access products. In the case of downloadable digital products, such as e-books, music, and movies, this has significantly changed the shape of the rank frequency distribution. In these markets, more of the niche products in the long tail are now selling because consumers can more easily discover and obtain them. For example, Anderson describes how the now-defunct movie rental company Blockbuster historically reported that about 90% of in-store movie rentals were new releases, whereas only 30% of Netflix rentals were new releases because Netflix customers predominantly rent a large variety of back catalog movies.[3] Similarly, there are more sales of long-tail books online than in real stores. For example, Barnes & Noble found that long-tail titles represented just 1.7 percent of its in-store sales, but the same titles took 10 percent of its online sales.[3]

Just as technological development has introduced new product marketing and delivery mechanisms, it has also eased product creation. In the case of movies, the development of affordable digital cinematography cameras and PC-based video editing have dramatically lowered the barrier to entry for aspiring filmmakers.

Another example is the software market: since the turn of the century, technological progress in software development has had a dramatic effect on the dynamics in this market. Development tools have become easier to use and cheaper to obtain, while agile, open-source development processes, platforms, and communities that enable collaboration have become firmly established. As a result, an app can now be written by a relative newcomer to the industry—perhaps someone who primarily wants to address a personal need.

Another technological development that has helped to enable many software products is cloud computing. The ready availability of cloud services removes the requirement for a company to make an up-front investment in server infrastructure, again lowering the barrier of entry to newcomers and allowing new ideas to be trialed quickly and easily. Many cloud computing providers readily support dynamic scaling and geographical redistribution should a software product or service prove successful.

In the case of mobile software, all the technologies on the previous page combine for even greater effect. New products can be developed more quickly and easily than ever before, the services that power them can be deployed with little up-front expense, and the "app store" marketing and distribution model provides a channel to millions if not billions of potential users. Even if a product only attracts hundreds of users a year, it can still deliver useful revenue for the producer, effectively adding weight to the long tail and increasing customer choice along the way.

While the above examples of emerging long-tail markets are all based on products that can be delivered digitally, it is worth noting that technological developments have enabled a tail for physical goods too.[3] Amazon and Alibaba's e-commerce platforms are the most visible examples of how marketplaces that combine online search, a large customer base, digital payment processing, and a fulfillment infrastructure have changed the face of retail for physical goods. Companies like Etsy take a similar approach to enable sales of low volume boutique goods.

HARDWARE DESIGN CONTINUES TO GET EASIER

The last two decades of innovation in software design tools, services, and processes have been mirrored in the hardware industry. Powerful and well-documented hardware design tools have been developed. These are easy to use and install—and in some cases even run in a web browser, e.g., Upverter. Entry cost has fallen and fully featured design tools such as Autodesk's Fusion 360 are now available on a monthly subscription, while others are free. At the same time, online communities relating to design such as Electrical Engineering Stack Exchange have become established, providing a way to resolve design questions and troubleshoot issues.

The physical prototyping process for hardware has also become more accessible. On the electronics side, off-the-shelf boards from micro:bit to Raspberry Pi and plug-together hardware platforms such as .NET Gadgeteer,[4] Seeed's Grove system, and MikroElektronika's Click boards allow an initial prototype to be pulled together in an afternoon. When more refinement is needed, readily available ICs and electronic modules increasingly abstract away underlying complexity. Custom printed circuit boards (PCBs) are inexpensive and quick to source, and companies such as CircuitHub and Seeed Studio provide an online service to support the necessary PCB assembly. Sophisticated hardware simulation and modeling tools are available, and some—such as the LTSpice electronic circuit simulator—are free. In terms of mechanical prototyping, improvements in rapid prototyping equipment and services mean that laser cutting, three-dimensional (3-D) printing, and computer-controlled machining are less expensive and more readily available than ever.[5]

In addition, new technologies are continually being introduced to further accelerate electronics development. For example, the recently published AutoFritz electronics computer-aided design system[6] provides a kind of "auto complete" for electronic circuit design. It presents the designer with a list of suggested components and circuit configurations based on the software's analysis of circuits from thousands of electronic component datasheets and existing open-source electronic designs.

Collectively, these on-going innovations in the hardware development space mean that nonexpert individuals and organizations can prototype all manner of interactive and embedded devices.[5] Even high-school students can design and build useful and innovative hardware by combining an off-the-shelf microcontroller board, some external components, and a 3-D-printed enclosure. If current trends hold, newcomers will be able to get started even more easily and will be able to build more sophisticated devices, while experienced developers will be able to work more productively.

CROWDFUNDING AND INCUBATORS MITIGATE UP-FRONT COSTS

Of course, as mentioned in the introduction, the design and prototyping phase is not the only challenge in creating a hardware product—a working prototype is just the start. The next obstacle in hardware development is the need for up-front capital investment in materials and customized manufacturing equipment (tooling).

Another 21st-century phenomenon—crowdfunding—helps to address this. Websites such as Kickstarter and Indiegogo combine a prepayment mechanism with a storefront akin to that of online retailers and app stores. In this way, a hardware developer can receive monies a year or more ahead of product delivery, easing the management of development costs and mitigating the up-front costs associated with capital equipment and raw materials.

Crowd Supply is a crowd-funding organization that exclusively targets niche hardware products for both amateur and professional electronics developers. As such, the profile of most of their customers is similar to the profile of the product creators. Crowd Supply maintains close involvement with each product team: in addition to facilitating up-front payment via its crowd funding portal, the company provides know-how and experience, plus access to a network of trusted partners to maximize the chances of success for the projects it agrees to support.

Start-up incubators also provide hands-on support and mentoring, often coupled with venture capital (VC) investment—all in return for a stake in the start-up. Shenzhen-based Hax and Silicon Valley's Lemnos Labs and Highway 1 focus on hardware product incubation, drawing on their hardware expertise and contacts to minimize the chance of missteps. Hardware Club is a community-based venture firm whose members share experience, knowledge, and contacts in order to move more quickly while simultaneously reducing risk.

It is worth noting that incubators typically pursue products with high growth potential—things that will quickly move out of the long tail. Companies who want to avoid the pressure that often follows VC investment and grow a new business more organically can employ the services of an external design service or consultancy in order to access product development expertise. Companies such as Dragon Innovation and PCH International specialize in enabling the prototype to production transition for electronic products.

BUT HARDWARE IS STILL HARD

The collective benefits of online storefronts, accessible and powerful design tools, crowdfunding, incubators, consultancies, and next-day delivery networks all reduce cost and mitigate risk for those seeking to bring a new hardware product to market. But evidence suggests that they are often not enough. The flip-side of preferential attachment is that nascent devices often face a vicious circle. To be successful they must combine significant utility with compelling user experience, while simultaneously being robust and reliable in operation. Meeting these goals requires a large design and engineering investment, both for the product itself and for its manufacturing process—and that is hard to justify before a product is successful. As a result, there are many instances where new consumer hardware products fail.

One example comes from Central Standard Timing, a two-person start-up that designed the world's thinnest watch. Based on a sophisticated working prototype, Central Standard Timing had tremendous success on Kickstarter—they raised over a million dollars from backers who paid either $99 or $129 for a watch. In moving from prototype to product, they chose to partner with one of the world's most experienced electronics manufacturing partners. But ultimately, they were unable to create a reliable production process—less than ten thousand units was a small volume for their manufacturer, and ultimately the per-unit cost of the watch was rumored to be $260.[9] As a result, Central Standard Timing were not able to deliver the product to their backers and ultimately ceased operations.

TriggerTrap, a company that manufactured devices to allow SLR cameras to be triggered externally, provides a second example of the difficulty of moving from prototype to product. TriggerTrap's first hardware offering was based on an Arduino prototype and resulted in a successful Kickstarter campaign. However, they experienced difficulties sourcing enough of the displays used in their prototype to meet a production volume of nearly one thousand units. The part was no longer available and TriggerTrap had to convince the display manufacturer to instigate a special low-volume production run before they could fulfil their orders.[10] TriggerTrap's second hardware product—the Ada—also resulted in success on Kickstarter, raising £290,000. However, the company was unable to transition from their prototype and failed to deliver any Ada products. Although they brought in expertise from external consultancies, they still made costly mistakes and the consultancy fees were an additional drain on their resources.[11]

Central Standard Timing and TriggerTrap both failed due to difficulties navigating the transition from prototype to product. They leveraged many of the mechanisms previously described—such as powerful design tools, online storefronts, crowdfunding and experienced professional design, and manufacturing partners—but nonetheless they still underestimated the cost and complexity of transitioning from a refined prototype to a product that could be manufactured in volumes of hundreds to thousands.

Of course, an experienced team with a track record of hardware product development is already well versed in the transition from prototype to product. But even then, it is still possible to underestimate the scale of the challenges involved. Our final example comes from SenseCam, which started as a wearable camera research project[12] before being productized by an established device manufacturer. Despite the company's experience with hardware manufacture, they ultimately found the vicious circle too difficult to overcome. The first and second versions of the SenseCam product (the Revue and Autographer, respectively) successfully sold in modest volumes to researchers, clinicians, and enthusiasts. But as the product design evolved on a trajectory toward a compelling and competitively priced consumer product, it became clear that an investment of tens of millions of dollars was necessary to reach sufficient economies of scale.[13] This level of investment was ultimately not justified given the risk inherent in a new product category such as a wearable stills camera, and production ceased.

THE REPLICATION CHALLENGE

We believe that many of the difficulties in creating a new hardware product, as encountered in the examples above, stem from a fundamental challenge. Unlike digital products which can be replicated in a simple way requiring almost no resource and resulting in perfect copies, the process of replicating a physical device is complex and incurs cost, and no-matter how much is spent on manufacturing, the copies are subtly different. We call this "the replication challenge."

Compared to a smartphone app that can be released to an audience of millions of people when development and feature testing are complete, scaling a hardware device from "I have one that works" to "anyone can buy one"—i.e., transitioning from a working prototype to a viable product—is much harder. Although this difference between software and hardware productization is easy to describe, the number and complexity of steps required to take a device to production is much harder to comprehend until experienced first-hand. Essential activities associated with the replication challenge[8] include:

- finding reliable suppliers for all components and materials;
- accommodating component tolerances;
- designing and building the necessary tooling;
- building an efficient and reliable manufacturing process;
- controlling and accommodating manufacturing variability;
- selecting and managing manufacturing partners;
- instigating and maintaining manufacturing quality control; and
- adapting to changes in pricing and availability of components and services.

These activities are by no means unique to electronic device production—they apply to the production of nearly all physical products. However, they are worth highlighting for two reasons: first, many of them involve more complexity and cost for electronic devices than for other physical products; and second, many of these activities are unfamiliar to those steeped in the software-side of technology.

The replication challenge does not only apply to high production volumes—in fact as volumes increase beyond ten thousand units a year many well-established manufacturing processes that leverage economies of scale become viable, so the cost of the above activities is more easily absorbed. Instead, replication is particularly challenging at lower volumes. For consumer products, the difficulties typically start at around one hundred pieces per year when the same ad-hoc craft production techniques used to make prototypes are no longer sufficient. A compounding factor with low volume manufacturing is its natural reliance on batch production, which introduces the further challenge of periodically reinstigating the entire manufacturing process and supply chain.

As a result, our sense is that it is increasingly rare for low volume consumer devices to be viable in the market. While Zipf's Law predicts a long tail of demand for niche hardware, the economics of production constrain the economic viability of such products, and the tail is truncated. As a result, consumer choice is limited to a relatively small number of high-volume devices. These are typically made by large companies who can justify the sufficient resource required to bring them to market and sustain them, and often only when they amplify or otherwise complement an existing product line.[7] In this environment, the recent successes in the interactive and embedded hardware space include the well-known smartwatches and voice assistants from major companies. These have created enough momentum in the market to transition from vicious to virtuous positive feedback and generate enough revenue to warrant on-going refinement.

CHARACTERISTICS OF LONG TAIL HARDWARE SUCCESSES

Of course, there are examples of commercially viable low volume consumer devices. Armed with a better understanding of what is necessary to transform a hardware prototype into a low volume product—and some of the pitfalls to look out for—it is interesting to examine some of the characteristics of these successful products.

Like SenseCam, Circuit Stickers[14] started life in research and quickly transitioned into a product through the efforts of Chibitronics, a small self-funded start-up. From the outset, Chibitronics understood the importance of reliable and cost-effective supply and manufacturing, recognizing that the design complexity of manufacturing tooling can out-weigh that of the product itself.[15] They had realistic expectations of production costs and were able to fulfill their crowd-funding obligations.

Subsequent to this, Chibitronics has successfully established a small range of niche products that sit in the long tail. In doing so, the company has started to turn positive feedback to its advantage: it is now able to amortize investments creating relationships with suppliers, manufacturing partners, and distributors, across multiple products. Although no individual product is manufactured in huge numbers, the aggregate volume across multiple products is large enough to negotiate better pricing and attract more commitment from component suppliers and manufacturing partners.

Improving efficiency across multiple products, as compared with the efficiency of manufacturing any one in isolation, is termed "economy of scope." While economies of scale are characterized by volume, economies of scope are efficiencies formed by variety. The hardware crowdfunding specialist Crowd Supply also leverages economies of scope by drawing on a specific set of partners to manufacture many of the products in its portfolio. In a similar vein, Crowd Supply takes advantage of its partnership with Mouser, a large electronic component distributor, to secure reliable and competitively-priced supply.

A characteristic of Crowd Supply's wide range of successful products is that many take the form of a "bare board"—a PCB that has components soldered to it but comes with no enclosure. This is perfect for Crowd Supply's target audience of hardware developers and hobbyists. It also avoids the high up-front cost often associated with designing and manufacturing tooling for an enclosure, which can be thousands, if not tens of thousands, of dollars for each plastic piece. In contrast, TriggerTrap, which needed an enclosure because the company was targeting a different audience, ran into difficulties transitioning from the rapid prototyping used to create its prototypes to manufacturing processes suitable for a product. Ultimately, the company could not reconcile the tooling cost of injection molding with its relatively low production volume.[11]

Central Standard Timing did not require injection molding for its product but faced a different challenge. In order to deliver one of the thinnest electronic devices ever produced, the company needed to use nonstandard electronics assembly techniques. Unlike PCB assembly, which is still cost effective even when scaled down to low volumes, the direct chip-to-flex process required costly tooling and introduced nonstandard steps in the manufacturing process, making it error prone and unreliable. Relying on leading-edge manufacturing technology in this way is particularly risky for start-ups.[13] In contrast, the established and well-understood manufacturing process for a standard PCB bare board product results in few defects and can be readily transferred between manufacturers, stimulating healthy competition.

UNLOCKING THE LONG TAIL: A CALL TO ACTION

Based on our analysis of some recent successes and failures in taking new hardware concepts to market, we do not believe that turning an idea for a new device into a working prototype is a limiting factor in new product introduction. Instead, we believe that the bottleneck is the transition from a working prototype to a viable product. Therefore, we encourage those working in the field of device hardware to join us in tackling this transition. Collectively, we would like to work toward three broad goals.

Improved Teaching Materials and Education Programs

We need richer learning materials regarding the productization process: new ways to share the knowledge that today largely resides in the design consultancies and large companies with the first-hand experience of device manufacture. Fortunately, there are already a few books that cover the topics raised in this paper; we encourage those interested to read "The Hardware Hacker" by Andrew 'bunnie' Huang,[15] "Prototype to Product: A Practical Guide for Getting to Market" by Andrew Cohen[8] and "IoT hardware from prototype to production" by Richard Marshall, Lawrence Archard, and Steve Hodges.[16]

We would like to see these books complemented by other learning materials. For example, curated virtual tours of device manufacturing facilities would provide valuable insights to those who cannot visit in person. We also encourage universities to provide more coverage of topics relating to the replication challenge in undergraduate- and graduate-level level courses and in professional development programs.

Stronger Communities and Tighter Integration Between Partners

We see an opportunity to extend today's established online communities and professional networks in a way that allows newcomers to the device hardware space to engage with each other and with established players so that they can more easily form the partnerships that are vital to the delivery of a successful product.

We imagine that these partnerships will increasingly span the globe. In the short term, we would like to streamline access to resources in today's leading electronics manufacturing regions. For example, the geographical proximity of capital equipment, skilled and unskilled labor, and raw materials in locations such as Shenzhen provides unrivaled economies of scope. Indeed, much of the Shenzhen ecosystem naturally lends itself to batch production: there are thousands of small factories specializing in different aspects of manufacturing from injection molding to PCB assembly.

At the same time, we would like to learn from existing electronics hubs so that we can share best practices. Exchanges like the annual trip the MIT Media Lab organizes to Shenzhen are one example of how this might be done. Improved communication and remote collaboration could also help bridge the cultural, geographical, and language barriers. Eventually, we imagine that additional electronics manufacturing hubs will emerge around the world.

New Manufacturing Solutions for Small Batches

Finally, we believe there is an opportunity to develop new hardware manufacturing solutions optimized for low-volume production, for batch sizes of hundreds to thousands of units. Of course, manufacturing in low volumes will always incur a premium due to the need to amortize up-front development and tooling costs across fewer units. Similarly, batch manufacturing will never be as efficient as continuous production because of the fixed costs associated with changing a production line from one product to the next. However, we imagine new tools and processes that make better use of economies of scope by amortizing nonrecurring costs across multiple products.

Take the current difficulty of manufacturing high-quality enclosures in low volumes as a specific example. Perhaps this could be addressed by a modular approach that combines a standard library of injection-molded parts in a novel way to create a finished enclosure. Alternatively, it may be possible for designers of new products to reuse pre-existing injection-molding tooling, an approach that Seeed Studio calls "design from manufacture." Perhaps the same reuse philosophy could be used to reduce tooling costs in PCB manufacturing tests: the custom test fixture—or jig—typically required for PCB tests could be based on standard hardware and software components.

Of course, the ideas above are largely speculative. They are included simply to illustrate the potential for innovative manufacturing solutions to address the replication challenges.

OUTLOOK

In this paper, we have argued that the innovations that have enabled a long tail of products in several industries can be leveraged to do the same for hardware devices. While we see these established innovations as necessary, in the case of hardware, they are insufficient. Devices bring additional complexities in comparison with products that can be distributed digitally such as apps, books, and movies. One key difference is the replication challenge, especially at low volumes where economies of scale do not readily apply.

It is certainly possible to be successful with a niche hardware product, in the same way that niche apps, books, and movies have been possible since the inception of their respective markets. However, as things stand, low-volume manufacturing is hard and requires tough tradeoffs between complexity, refinement, and price. As a result, it is all too easy to fail. We believe that a sustained focus on the replication challenge can reduce this complexity and risk.

Our motivation is simple: we want to change the dynamics in the electronic device market to enable a diverse ecosystem of products that do not need to sell in tens or hundreds of thousands of units in order to be viable. Such cost-effective, low volume, batch manufacturing would benefit many parties across many domains. Researchers developing custom devices could deploy them more widely. Start-ups looking to grow a sizeable hardware business could get to market with less risk, giving them headroom for iteration in search of product-market fit. Small companies could manage a portfolio of niche, but viable, hardware products. Large companies could be more agile and less conservative, perhaps even trialing innovative hardware products to learn first-hand how they fare in the market.

We welcome feedback from the research community on the ideas presented here and encourage others to complement our ideas by considering how their work might support the growth of long-tail hardware. We will also continue to engage in dialogue with practitioners in the device industry. Ultimately, a focus on long tail hardware can overcome the tyranny of positive feedback that currently constrains innovation, thus creating a greater variety of viable hardware products that address markets that are currently underserved, and meeting the world's growing demand for interactive and embedded devices.

IN MEMORIAM

This article is dedicated to the memory of Gavin Zhao who lost his battle with cancer in August 2019. Gavin was a manufacturing engineer at electronics manufacturing company AQS, where he helped a great many people navigate the difficulties of low volume hardware manufacturing in China. In 2017, he was awarded an MIT Media Lab Director's Fellowship in recognition of his work opening the Shenzhen ecosystem up to others. Many conversations with Gavin helped shape our thinking around long tail hardware.

ACKNOWLEDGMENTS

The authors would like to thank the many people who have entertained our questions and ideas about electronics manufacturing over the years, and ultimately informed the ideas presented in this paper. Particular thanks go to: Gavin Zhao, AQS, Shenzhen; Haje Jan Kamps, Bolt, San Francisco; Dave Vondle, Central Standard Timing, Chicago; Bunnie Huang and Jie Qi, Chibitronics, Singapore; Andrew Seddon, Circuit Hub, London; Jewel Deng, Coolkit, Shenzhen; Josh Lifton and Darrell Rossman, Crowd Supply, Portland; Brooks Vigen, Digikey; Alex Gluhak and Ran Katzir, Digital Catapult, London; Gus Issa, GHI Electronics, Detroit; Ji Ke and Mike Reed, Hax, Shenzhen; Jerry Shi, Itead, Shenzhen; Eric Klein, Lemnos Labs, San Francisco; Liya Du, Microsoft, Shanghai; Anita Rao and Tarun Singh, Microsoft, Redmond; Phil Eade, Microsoft, Cambridge; Gibson Guo, Zoe He and Nicolas Schmitt, Microsoft, Shenzhen; MJ Shen, MJ Maker, Shenzhen; Zach Fredin, NeuroTinker, Minneapolis; Liam Casey and Alan Cuddihy, PCH International, Shenzhen; Simon Randall, Pimloc, London; Fraser Forbes and Jonathan Smith, Premier Farnell, Leeds; Joey Jiang, Ivy Li, Albert Miao, Eric Pan and Shuyang Zhou, Seeed Studio, Shenzhen; Nick Bolton, Vicon, Oxford. They would also like to thank Oliver Amft, James Devine, Rushil Khurana, Michal Moskal, and Nilay Patel for comments on earlier drafts of this column and Albrecht Schmidt for his enthusiasm to see these ideas come together.

REFERENCES

1. Wikipedia. "Zipf's law." Accessed: Nov. 2019. [Online]. Available: https://en.wikipedia.org/wiki/Zipf%27s_law
2. J. M. Perloff, *Microeconomics*, 8th ed. London, U.K.: Pearson, 2018.
3. C. Anderson, *The Long Tail: How Endless Choice is Creating Unlimited Demand*. New York, USA: Hachette Books, 2008, ISBN 978-1401309664.
4. N. Villar, J. Scott, S. Hodges, K. Hammil, and C. Miller, "NET Gadgeteer: A platform for custom devices," in *Proc. Pervasive Comput.*, 2012, pp. 216–233.
5. S. Hodges, N. Villar, J. Scott, and A. Schmidt, "A new era for ubicomp development," *IEEE Pervasive Comput.*, vol. 11, no. 1, pp. 5–9, Jan.-Mar. 2012. doi: 10.1109/MPRV.2012.1.
6. J.-Y. Lo *et al.*, "AutoFritz: Autocomplete for prototyping virtual breadboard circuits," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2019, Paper 403. doi: 10.1145/3290605.3300633.
7. F. Manjoo, *The Gadget Apocalypse Is Upon Us*. New York Times, New York, NY, USA, Dec. 7 2016. [Online]. Available: https://www.nytimes.com/2016/12/07/technology/personaltech/the-gadget-apocalypse-is-upon-us.html
8. A. Cohen, *Prototype to Product: A Practical Guide for Getting to Market*. Sebastopol, CA, USA: O'Reilly Media, Aug. 2015.
9. S. McGlaun, "CST-01 watch project may be dead", SlashGear, Jun. 23 2015. [Online]. Available: https://www.slashgear.com/cst-01-watch-project-may-be-dead-23390137/
10. H. J. Kamps, "Hardware is Hard: Getting a Kickstarter project out the door", Medium Article, Jan. 22, 2015. [Online]. Available: https://medium.com/triggertrap-playbook/hardware-is-hard-getting-a-kickstarter-project-shipped-59c9596bdd7f
11. H. J. Kamps, "How Triggertrap's $500k Kickstarter campaign crashed and burned", Medium Article, Mar. 2, 2015. [Online]. Available: https://medium.com/@Haje/how-a-half-million-dollar-kickstarter-project-can-crash-and-burn-5482d7d33ee1
12. S. Hodges *et al.*, "SenseCam: A retrospective memory aid," UbiComp, 2006, pp. 177–193.
13. S. Randall, "Hacking the consumer electronics hardware start-up," Medium, Sep. 2018. [Online]. Available: https://medium.com/swlh/do-not-apply-if-you-want-an-easy-life-403e77a4a38f
14. S. Hodges *et al.*, "Circuit stickers: Peel-and-stick construction of interactive electronic prototypes," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 2014, pp. 1743–1746. doi: 10.1145/2556288.2557150.
15. Andrew "bunnie" Huang, "The Hardware Hacker: Adventures in Making and Breaking Hardware", 2017.
16. R. Marshall, L. Archard, and S. Hodges, *IoT hardware from prototype to production*, 2019. [Online]. Available: https://aka.ms/proto-to-product

Steve Hodges is a senior principal researcher at Microsoft, where he builds new embedded and interactive systems and devices. He received a PhD in computer vision and robotics from Cambridge University. Contact him at: shodges@microsoft.com.

Nicholas Chen is a senior program manager at Microsoft, where he is building the hardware ecosystem for Azure Sphere. He received his PhD from the University of Maryland at College Park. Contact him at: nchen@microsoft.com.