Paul Gardner-Stephen, Romana Challans, Jeremy Lakeman, Andrew Bettison, Patrick Lieser, Ralf Steinmetz, Flor Alvarez, Matthew Lloyd. *Productizing Humanitarian Telecommunications Research: A Case Study of the Serval Mesh Extenders*. In Proc. 7th Global Humanitarian Technology Conference (GHTC), IEEE, 2017. ISBN:978-1-5090-6046-7

Productizing Humanitarian Telecommunications Research: A Case Study of the Serval Mesh Extender

Paul Gardner-Stephen and Romana Challans and Jeremy Lakeman and Andrew Bettison
College of Science & Engineering
Flinders University, Australia
Email: paul.gardner-stephen@flinders.edu.au

Patrick Lieser and Ralf Steinmetz
Multimedia Communications Lab, TU Darmstadt, Germany
Email: {patrick.lieser, ralf.steinmetz}@kom.tu-darmstadt.de

Flor Alvarez
Secure Mobile Networking Lab, TU Darmstadt, Germany
Email: falvarez@seemoo.tu-darmstadt.de

Matthew Lloyd
Serval Project Inc.

***Abstract*—The Serval Mesh Extender is a low-cost open-source infrastructure-independent telecommunications relay device developed to support telecommunications during and following disasters, as well as in remote and isolated locations. The Mesh Extender has been under development for five years, and is just now transitioning from a primarily University research system, into a mass-producible and deployable humanitarian telecommunications product. This has forced the research team to consider numerous challenges and trade-offs that are substantially common to this type of activity, including industrial design, supply-chain formation, inventory management, electronics design and manufacture, tooling for injection-molding and planning around the variability of research funding. In this paper, we describe these challenges, together with insights and mistakes made and lessons learned during the process, in a format intended to benefit other researchers seeking to productize their research for the common-good of humanity, without relying on traditional profit-oriented commercialization pathways. That is, we provide practical advice for those seeking to make the fruit of their research as widely available as possible, and as affordably as possible, for when charging what the market can bear is not appropriate or conscience.**

I. Introduction

This paper has come about as part of a pilot of the Serval Mesh in Vanuatu, funded through the Australian Department of Foreign Affairs and Trade (DFAT) Pacific Humanitarian Challenge [1]. For that pilot, there was a need to productize the Serval Mesh Extender. The Serval Mesh Extender itself is a low-cost communications relay device, designed to support the operation of the Serval Mesh [2]–[6], and thus, to support humanitarian and remote telecommunications in a variety of use-cases.

Ordinarily, the path to productizing research is to enter into a partnership with a new or established commercial entity, and then follow the path well worn to the release of a hardware product, even if the path is not necessarily smoothly paved, or well signposted. However, in our case, the humanitarian imperative makes this unattractive for a number of reasons.

Principally, we realized that to maximize our humanitarian impact, we need to minimize the price of the resulting product. That is, the commercial model of charging what the market could bear was not a sensible option. However, by abandoning the ordinary path to commercialization, we were faced with the reality of having to drive the entire process ourselves from our small research group. In this regard, it was fortunate that the department where the Serval Mesh Extender was being developed also hosts industrial design, and sufficient experience was available to enable us to commence on this journey.

The remainder of this paper documents key points and observations from that process, beginning with the genesis of the Serval Mesh Extender, and continuing to the present. The Serval Mesh Extender is not yet a finished product, and thus this story is necessarily incomplete. However, it is our hope that it may, perhaps, provide some rudimentary lighting and signage for that path less traveled of non-commercial productization, which is of particular relevance to the humanitarian sector. The tone and approach is therefore purposely intended to be candid and accessible, so as to maximize its usefulness, and, hopefully, be enjoyable to read as you follow the story of our journey so far.

It should also be mentioned that this paper is written from the perspective of the first author. The contributions of the other authors are principally in enabling the project to proceed to this point, and in providing various forms of support in the path to productization. It is only just that they be recognized in this way. However, it is the misadventures of the first author that are described in this paper.

A. Relationship to the Pacific Emergency and Disaster Telecommunications Coalition

This paper is one of a related series of papers [1], [7]–[14] concerning the Serval Project and related initiatives to be presented at GHTC 2017. These papers collectively represent a subset of the initiatives of a regional coalition to transform the affordability and reach of disaster and remote telecommunications technologies in the Pacific and beyond, based out of the Resilient Telecommunications Laboratory at the Flinders University, Adelaide, Australia. This coalition, temporarily named the Pacific Emergency and Disaster Telecommunications Coalition, is in the process of being constituted involving stakeholders from the Pacific and beyond, including operational NGOs, governmental institutions, international organizations and research groups. It is hoped that more will be able to be reported on this coalition at the GHTC gathering.

II. Involve an Operational Organization

The development of the Serval Mesh and Mesh Extenders has been based on a close working relationship with New Zealand Red Cross. This has been invaluable, because it has allowed us to jointly determine the functional requirements, and to ensure that we are making a device that will actually be of use in humanitarian situations. It also allowed for the development of additional ideas and technologies that would not have arisen if we had been working independently, for example, the Succinct Data data compression and satellite uplink system [6], that allows structured data collection using the Magpi online data collection system that ordinarily requires internet access, in remote locations lacking internet coverage.

By partnering with an operational organization in particular, i.e., one that actually deploys into disaster zones, our feet were kept on the ground, and we were also able to make use of their periodic field exercises that they conduct to maintain their readiness, to test and receive feedback on new technologies.

We therefore encourage all humanitarian technology developers to seek similar partnerships when and where possible, to maximize the value and impact of their work. The Global Humanitarian Technology Conference itself is an excellent place to make those linkages.

III. Malnourished Production Principles

Toyota is famous for its lean production methodologies [15]. Chances are, however, that as a research project seeking nonprofit productization, you will have to not merely achieve lean production, but in fact malnourished or even starvation production methodologies.

Indeed, to push the food analogy further, academic research funding cycles is rather prone to feast and famine cycles. Whereas lean manufacturing could be compared to maximizing the calorific efficiency of the eater, malnourished production seeks to maximize survival when faced with an unreliable food supply. Overall calorific efficiency may, however, be worse.

What we have coined Malnourished Production Principles (MPP) have probably been practiced since time memorial and in every field. As a non-discipline, MPP is difficult to define. However, it is possible to draw out several guiding principle from our experience.

A. Expect Feasts and Famines

This principle is almost a description of what MPP is: There will be times when you have more resources than others. When you have more resources, see what advances you can make, that will endure the next famine cycle. This might be developing a next-generation prototype, generating quality promotional materials, traveling to meet with potential partners, or submitting grants, for example. Essentially the idea is to recognize when you have the capacity to move forward, and then based on the circumstances, make the most effective advances that you can, knowing that the time of feast is limited. Well executed, this can allow you to not merely protect the future of your endeavor, but to allow it to prosper during the next famine cycle (Genesis chapter 41). That is, use one feast cycle to setup the conditions for the next. Think hard about additional capability you can demonstrate, or which Achilles' Heel you can mitigate. Above all, be creative.

For the Serval Project, examples of feast-time activities have included: (1) implementing the security model and tailored transport protocols that have provided a solid foundation for future developments and funding applications; (2) developing successive generations of prototypes; (3) attending appropriate international meetings; and, most recently, (4) productizing the Mesh Extender for modest scale production, so that we can scale up our future piloting and deployment activities.

During the famine times, we have found a combination of work on software improvements, seeking to maintain community visibility through blogging and other communications, and seeking out the opportunities to deliver the next feast cycle to be productive. Again, the key is to think laterally, identify opportunities as they arise, and think about what tactical victories you can achieve to achieve the kinds of milestones that will increase your ability to bring the next feast closer. A challenge is to ensure that you maintain your long-term direction, balanced against seizing the opportunities that arise to avoid total starvation of the project.

Employment as an academic can be a productive defensive strategy, as when research income wanes, your own employment need not be in mortal danger, allowing a core of activity to continue. Tenure may be necessary to maximize the benefit of this, although even then the trade-off is that a lack of research income typically translates into a disproportionate increase in teaching load, which may make mortal danger feel like an attractive option.

A second strategy that we have accidentally adopted as a direct result of the famine-feast cycle, is to employ personnel on a fractional basis, so that they have multiple income sources. This has two beneficial effects. First, your minimal survival level of funding is reduced. Second, during feast times, you have an obvious path to scaling up, without having to recruit: simply increase the employment fraction of the appropriate personnel. However, not all people are willing to work under such circumstances, and there may be considerable lead times before employment fractions can be increased, especially for good employees who wish to be honorable to all their employers.

Finally, when possible, try to build up a buffer of discretionary funds that can be used to smooth out the famines, or be used to re-prime the pump at an opportune time. Maintaining funds to allow rapid grasping of emergent opportunities is a good policy, as opportunities tend to come along, provided you can wait long enough.

B. Benefits of Open-Sourcing and Advocacy

A further element that we discovered, is the value of open-sourcing technology. There are a number of benefits for this. First, it allows you to find partners, before you complete creation. Second, it can help you find contributors, who can help you build, test and refine your creations. Third, we have found that for some significant partners, such as well known international organizations, that they have been willing to collaborate with us precisely because our technology is open, and therefore avoids risk of vendor lock-in for them if they were to adopt our technologies. Fourth, by releasing your source code as it is developed (as compared to releasing it only on completion) allows for advocacy, community and partner building while you are building. Given the likelihood that you will not have sufficient funds in one serve to achieve your goals, this provides a helpful mechanism to maximize the potential to find the next meal of funding.

This issue of advocacy is significant, as it can make tremendous difference to the ability to obtain funding, find partners and generally amplify your impact. We have found it helpful to have someone who can serve as the visible figure head, i.e., ambassador for our projects. These benefits stem from several sources:

First, having a single person or a small group of people who are the primary voices talking publicly about a particular project, this tends to lead naturally to having relatively consistent messaging. Consistent messaging is simply more effective. Second, once the figure head or ambassador is known it is easier for media and other inquiries to find their way to your project. Also, it means that when the media do find someone to talk to, it is hopefully someone who can give accurate answers, and answers that help to advance the prospects of your project. Third, once media get to know that a particular source or "talent" is worth consulting, they will tend to consult again, including on matters that are only indirectly related. This gives further opportunity to raise the profile of the project and team. Finally, visibility, whether direct or indirect, helps industrial, governmental, humanitarian partners to find you.

While in the early stages of a project the workload of such an ambassador may not be high, as a project grows over time, this workload can increase significantly. For the Serval Project responding to inquiries and related activities often consumes more than one day per week, and often involves conference calls at all hours. Nonetheless, we consider that overall the benefits far outweigh the disadvantages.

IV. History of the Serval Mesh Extender, or, Don't Overestimate What You Can Do In A Year, Nor Underestimate What You Can Do In A Decade

Very early in the history of the Serval Project, the staff of the Shuttleworth Foundation offered me a number of very sage pieces of advice. As I set out on a rather ambitious 12 month program of activity, they gently told me not to over-estimate what can be accomplished in a year, but also not to under-estimate what can be achieved in a decade. This was of course their way of trying to tell me that I was trying to do too much in a year. Twelve months later I finally understood what they meant, as my plans lay only partially completed. However, in the following six years, we have made considerable progress, more than I would have expected.

A. Genesis of the Mesh Extender and First Prototype

The genesis of the Mesh Extender arose from our attempts to enable Android smart-phones to communicate directly, to form self-organizing ad-hoc mobile telecommunications networks. In early versions of Android, it was possible to use ad-hoc Wi-Fi to enable true peer-to-peer communications over distances of up to a couple of hundred meters per hop. However, this required rooting phones, and also dramatically increased the power consumption of the phones. The power consumption was so extreme, that battery life was typically reduced to as little as an hour. The final nail in the coffin, however, was that later versions of Android have made it effectively impossible to use ad-hoc Wi-Fi on newer phones.

Therefore we set about designing a helper device, that would act as both a Wi-Fi access point and ad-hoc Wi-Fi node, so that there would no longer be a need to root phones. A welcome side-effect was that the power consumption of the mobile phones was returned to normal levels, because they were able to use the energy-saving features of 802.11n infrastructure mode. The power consumption of the helper device would only be similar to that of a single phone running ad-hoc Wi-Fi, so as soon as more than one phone was involved, the overall energy consumption of the system would be reduced.

It was also observed that if we had a separate device, we could add other radio types, to increase the range of communications. The RFD900 [16] was identified as a small, low-cost, low-power and high-performance radio that could operate in the 915MHz ISM band. This was added to the system, and a simple Serval Mesh synchronization driver for the RFD900 written, combined with a battery and a plastic tub (Figure 1), and the Serval Mesh Extender was born as a concept.

Fig. 1. The original Serval Mesh Extender prototype, during an indoor test at the local supermarket, where the plastic tub was purchased. This prototype was built entirely from off-the-shelf products.

1) Avoid Even Superficial Likeness to Bombs: At this point, the Mesh Extender, known at the time as the Mesh Helper, was obviously only a very crude prototype, with an unfortunate superficial likeness to an Improvised Explosive Device (IED), due to the combination of the large black plastic slab, the proliferation of wires, antennae and blinking lights. More correctly, the prototype looked like what movies have taught people to think that bombs look like. Unfortunately in this modern media-saturated era, the subtle difference between substance and form may only become apparent during post-mortem.

The cosmetic likeness to an IED was not merely of theoretical concern: We had a few tense moments on different occasions where explanations had to be given to customs and border protection officials when demonstrating it internationally. Also, special care had to be taken when, for example, meeting with partners to test the Mesh Extenders on the National Mall in Washington, D.C., so that we would not attract unwanted attention from the Department of Homeland Security. It was perhaps a good thing, that it was only *after* taking the units out onto the National Mall, that they told the first author about the snipers based on all of the surrounding buildings.

We would therefore suggest that, to the extent possible, that future projects avoid building devices that look like bombs, or look like what movies have trained people to think bombs look like. As the problem is a purely psychological one, solution could be to simply use an opaque container, and affix appropriate, clear messaging to the exterior [17].

At this point in time, the prototypes demonstrated the concept of the Mesh Extender, but were also rather functionally limited. Specifically, the standard firmware of the RFD900 radio at the time supported only point-to-point links, not ad-hoc links. This meant that only a pair of Mesh Extenders could be used, and that it was not possible to establish multi-hop links using larger numbers of Mesh Extenders. However, even with this limitation, the promise of the Mesh Extender was sufficiently established to allow us to obtain further development funds from a variety of sources. Fitting with our low-cost production ideals, these funds came from humanitarian and philanthropic sources, rather than commercial sources.

2) Pitfalls of Commercial Funding: This is an important point, because if the funds were commercial, the commercial partner would naturally seek to obtain a return on their research and development investment. Given that the total research and development investment in the Serval Mesh Extender and underlying technologies exceeds two million Australian dollars, and the relatively small expected initial market size of the Mesh Extender, would necessitate a prohibitive pricing model, at which point the potential market would likely reduce to zero, effectively killing the project.

That is, if you wish to end with a low-cost humanitarian project, careful consideration should be given to the funding sources accessed, to ensure that it can be cost-effective enough to succeed. Fortunately, there are a number of humanitarian and philanthropic funding sources available.

3) Academic and Student Co-Development: Academic researchers have the further advantage that they can work with students to undertake the research and development as a natural part of their occupations. Care must be taken to ensure in such cases that intellectual property is properly managed, to ensure that no student nor University commercialization body can veto the release of the product. Early discussions are the key here.

In the case of the Serval Project, the University was engaged early in the process, to ensure that all parties understood and agreed that the final product would be open-source and publicly released. Similarly, when necessary, students may be required to sign intellectual property agreements, where the give permission for their contributions to be released under open-source license terms. Many Universities have intellectual property counseling services for students, although not all of those may be familiar with open-source licenses and methodologies.

4) Opportunity and Hazards of Crowd-Funding: Crowd-funding is also a potential option, although it requires considerable work, and there is a large element of luck. Expect to devote several person months to the effort. Our efforts in this space were only somewhat successful, raising approximately US$13,500 of a US$300,000 goal, plus the attendant increase in publicity which has its own value. As we had chosen a flexible funding goal, we received those funds, and commenced a very modest program of updating the Serval Mesh Extender. As recent crowd-funding debacles have shown, e.g., [18], it is important to make sure that you set achievable goals. It is very easy for fulfilling perks to require many times more time and effort than the funds raised, and to generate considerable public hostility.

Fortunately, for a humanitarian project, you are in a position to appeal to people's humanity, and offer perks that are either substantially intangible, e.g., personal thank-you emails, of negligible financial cost, such as sending a post-card, or otherwise unrelated to the technology you are developing, such as exclusive dinners with the development team. Stickers and t-shirts are also possible, and we did offer those. However, even for these sorts of seemingly simple perks, do not underestimate the costs and complications involved. To this day, we still have several supporters we have not been able to contact to deliver their perks, and the first author discovered just how long it takes to write a personal note on even a modest number ($<$100) of post-cards. Fortunately, for the more common perks, such as t-shirts and stickers, a whole industry has popped up to service those undertaking crowd-funding schemes, not unlike the hardware and service industries that have spontaneously developed around gold rushes in the past. Indeed, the ecosystem and economy surrounding the "crowd funding rush" bears many similarities to that of the various gold rushes, however, that is a topic beyond the scope of this paper.

We would specifically counsel against offering your product as a perk, unless you already have a working final version first, as this invokes considerable risks and challenges. There are, after all, very good reasons why commercial entities sell products for $2\times$ - $4\times$ their production cost, and they are experienced in this field. Academic researchers, if they are nonetheless possessed with the idea of offering yet-to-be-realized products as perks in a crowd-funding campaign, should allow at least $4\times$ production costs, or undertake very detailed modeling of costing their time, transport, consumables and many other costs before proceeding. If you are concerned about how potential backers will respond to such a large mark-up, simply be transparent and explain that you are applying such a mark-up, precisely because there are risks, and you want to make sure that everyone gets what they have paid for, without anyone being disappointed. In crowd-funding honesty and transparency are the key to trust and satisfaction, whether you are making computers, potato salad [19]–[21] or something entirely different [22]–[26].

B. Second Prototype

In response to the limitations of the first prototypes of the Mesh Extender, including a desire to avoid unnecessary post-mortems, we set about designing a revised design. The objective of this second design was to make it possible for third-party contributors to easily and cheaply build and use. However, our resources at the time were extremely constrained due to the limited success of our crowd-funding campaign.

What we achieved, however, was to engage a supporter from Hungary who designed us a 3D-printable case, so that first impressions of the revised Mesh Extenders would make people think of communications before detonations (Figure 2). Internally, it was, however, just the combination of an off-the-shelf wireless router running custom firmware, an RFD900 UHF packet radio, designed for use in model aircraft, and a USB memory stick (Figure 3). Essentially we were forced to continue to use off-the-shelf parts due to our very limited budget.

Fig. 2. A second generation Serval Mesh Extender prototype, looking more communicative and less explosive than the first. The 3D-printed exterior gives the appearance of a more advanced prototype. Internally, however, it still consists entirely of carefully selected off-the-shelf components.

Because of our very limited resources at this time, we elected to make the battery external. This was not an entirely negative development, as it allowed us to have the Mesh Extender powered from any USB power supply, which is very flexible.

Our resource limitations also meant that we continued to use an off-the-shelf wireless router as our main PCB. This caused us a number of problems. First, the PCB that was most appropriate to our purposes lacked on-board bulk storage, so we had to resort to a USB memory stick. This increased power consumption by about 20%, and introduced a fascinatingly wide variety of failure modes. Our favorite was revealing a bug in the firmware of the common micro-form-factor USB memory sticks in production at the time.

We were using the micro-form-factor USB memory sticks, so that the unit as a whole would remain compact and with pleasing proportions, and be easy to print on even a small 3D printer. However, the firmware on those memory sticks turned out to be intolerant to the particular IO access patterns that the Serval Mesh software made. Even without unexpected power loss, it was common for the memory stick's firmware to corrupt

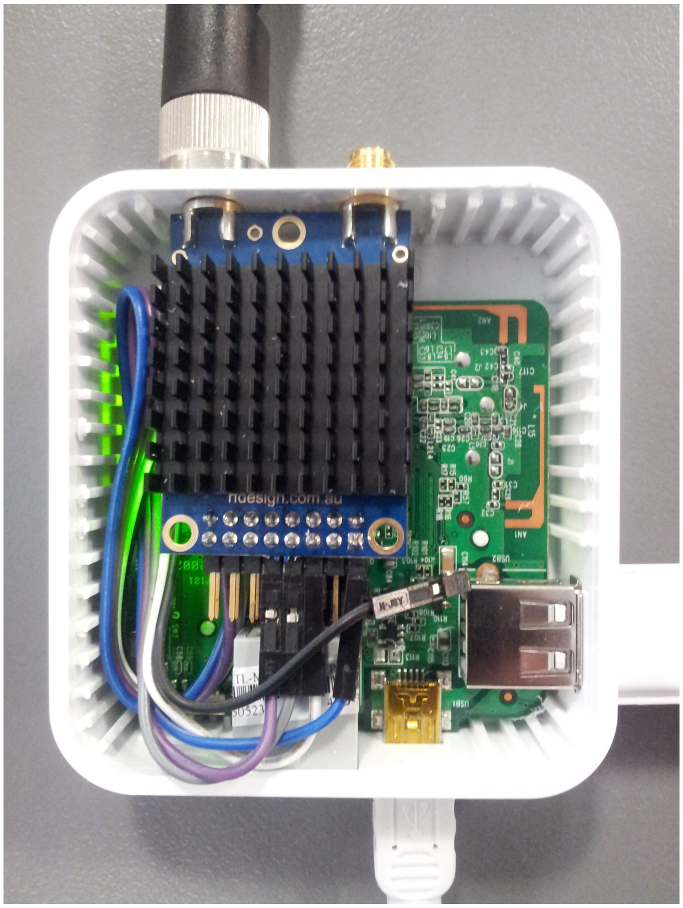

Fig. 3. A second generation Serval Mesh Extender prototype, lacking custom 3D-printed case, revealing how it was designed around a low-cost off-the-shelf wireless router and UHF packet radio module, designed for model aircraft.

some low-level data structure, which would trigger the USB memory stick to enter some kind of safe-mode. The problem is that the safe-mode manifested itself as the memory stick going permanently and irreparably read-only. The half-life of these memory sticks was only about 16 hours of operation in a Mesh Extender.

Problems like the short half-life of USB memory sticks in our prototypes are extremely difficult to predict ahead of time. Adequate resourcing allows them to be engineered around, however, that is rarely a luxury in the academic and non-profit research and development process. This low capacity for unexpected technical problems is one of the many reasons we recommend extreme caution if running a crowd-funding campaign.

Despite these challenges, we achieved our goal with the second generation prototype, which was to create a device that was recognizable, and that could be used to demonstrate the potential of the Serval Mesh and Mesh Extender concept in a variety of situations. It also allowed the developing community around the Serval Mesh to build their own units, and replicate our work, to considerable benefit, and continued development of the community and public interest in the project.

It was, however, still far from a product: It was not weather-proofed for outdoor installation, the half-life of the USB memory sticks was highly problematic, and the units were fiddly to build. Building even a dozen units required several days of effort, and was simply not scalable.

C. Third, and (Hopefully) Manufacturable Prototype

These limitations with the previous prototypes, combined with improved funding circumstances through the Pacific Humanitarian Challenge (PHC), led us to create a third generation prototype. The PHC award did not, itself, require us to create a new manufacturable generation of the Mesh Extender. However, following Malnourished Production Principles, we concluded that the best course of action was to make the maximum advances that we could during the feast part of the cycle.

Taking into account the labor cost of building the approximately 100 Mesh Extenders that we would require for our PHC pilot project, we concluded that it would only be moderately more expensive to design a productized version of the Mesh Extender, that could, hopefully, be realistically manufactured in reasonable quantities at an acceptable cost. This process was not without trade-offs. First, it required us committing what little discretionary funding we had from other sources. Second, while it felt like a feast to us, the entire PHC award was still smaller than a regular commercial product development budget – and we had considerable other costs to factor, in particular considerable travel and logistical costs due to the location of the pilot in rural and remote areas of Vanuatu, a country which itself is remote and relatively expensive to access, even from Australia.

We were fortunate enough to have access to product design experience in our department. Without this, I doubt that we could have proceeded down this path. Gaining contacts to local and competent industrial designers, tool makers and injection molders for the case made all the difference. We were able to quickly gain a reasonable upper-bound on the costs and time-lines involved in our particular case, which enabled us to make the decision to proceed with a reasonable degree of confidence.

Even so, there were several mis-adventures along the way. First, the first author was called to a visiting fellowship at another University for several months, which slowed the process down at a critical time. Second, everything took longer than we expected, and in some cases, than we were led to believe. In some cases, this was because suppliers were providing us with considerable discounts on their normal rates due to the humanitarian nature of our project, and therefore, quite reasonably, scheduled higher-margin work ahead of ours.

Third, tool-making, that is, the process of carving out the 300KG block of tool-steel to make the injection molds, is at the end of the day a an art, involving very accurately milling away hard steel millimeter by millimeter. If the milling tool or the jig holding the block of steel shifts even slightly in the middle of the night, you end up with a misshapen hole in the steel. You are then faced with working out how to adjust the design (more cost and a bit more time) to accommodate this, or if your contract with the tool maker allows it, starting the milling process all over again with new block of steel (possibly no cost, but a lot more time).

We had such a situation, where the tool slipped by 0.2mm, causing a slight under-cut in one area. We were very grateful that we had access to our industrial designer and the tool makers to understand what impact this would have, and to have a solution recommended. In our case, we are using polycarbonate, which does not tolerate under-cuts in molds well, and so we ended up thickening one wall of the case by 0.2mm, effectively cutting the under-cut out, and saving a delay of 12 weeks to start the tooling process all over. In the end, the tool was complete (Figure 4), and we received our first samples of the injection-molded cases in time, although several months later than we had hoped. Even then, as is quite common, some post-machining was required to correct some minor problems with the parts.

Fig. 4. Injection-molding tools during the milling process (top), and on completion. The complete set of tools consisted of several such pieces.

We had the tools manufactured in China to contain costs, but through a local tool maker, so that we could avoid the complications of engaging directly with foreign suppliers for such a critical component of the process. However, in preparation for future times of famine, we have arranged for the tools to then be shipped to Australia, where they will be stored by a local plastic injection molding company. This will allow us to quickly and easily produce more cases as and when we need them, without fear of the tool going missing, being damaged or otherwise holding us to ransom.

D. Supply-Chain Formation

The electronics of the Mesh Extender were also overhauled at the same time as the case design. The goal was to produce an integrated circuit board that could be easily manufactured, without all the fiddly wiring to the UHF packet radio, and including a variety of features critical to our use-cases that demand unattended outdoor operation in hostile environments.

We chose to use the manufacturer of the RFD900 radio as our partner for the PCB design. This decision was based on two factors: First, we knew from the RFD900 that they were capable of producing excellent quality at a reasonable price. Second, we were able to involve one less supplier in the supply chain. Third, we were able to form a strategic partnership with them, as they know that every Mesh Extender will mean an extra sale for their flag-ship RFD900 product range. As a result, we were also able to obtain an in-principle agreement that they would act as retailer for the Mesh Extender when we reach that point, saving us the complication of having to deal with such matters ourselves. Thinking ahead to future famines, we have also made sure that we own the intellectual property in the PCB design, so that we retain full flexibility in future developments and opportunities as they arise.

As multiple iterations of circuit boards are an almost certainty, we found having a very structured approach to capturing and communicating errata to be extremely productive. The small time investment involved has easily repaid itself by allowing RFDesign to work more efficiently and with greater certainty as to what is required with each successive revision of the PCB.

Circuit design is also an inherently risky process, and we have had a number of challenges that we have had to solve, including one serious one that we are still in the process of resolving, again related to mass storage. These problems are not always to predict ahead of time, and sufficient time and budget contingency should be allowed for these. We are already at revision four, and expect at least a fifth revision of the PCB, before we can manufacture the prototypes for the PHC pilot. Revisions of the PCB will likely cost more than the initial design of the PCB. Certainly the time lost to this process is also significant. It is now six months since the first revision of the PCB. As the saying goes, hardware is hard.

However, despite these difficulties, the end result has been well worth it (Figures 6 and 7). The Mesh Extender now exists in a form that is deployable, and potentially ready for commercial availability. It has been clear to us as we have

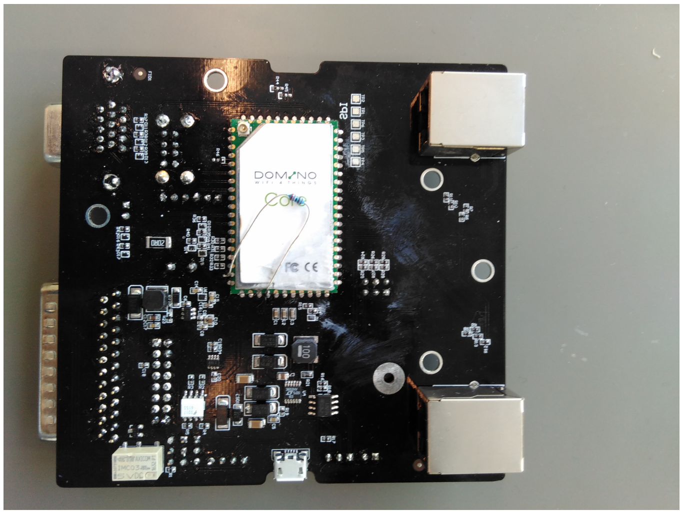

Fig. 5. An early revision of the PCB, showing hand-soldered modifications as we identified and corrected faults.

demonstrated this new third-generation Mesh Extender, that its clearly matured and robust form has a significant impact on first impressions, when parties are introduced to the Serval Mesh. We now find that a common response is people asking how much they cost, and where can they buy them, or inviting us to conduct a pilot in their country or organization

Fig. 6. A third-generation of the Serval Mesh Extender, in its natural environment, strapped to a coconut palm in Vanuatu.

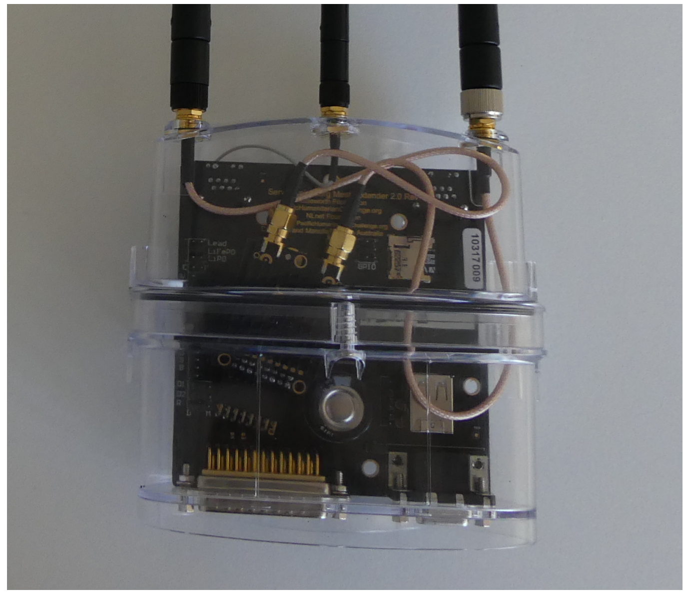

Fig. 7. A limited-edition transparent exemplar of the third-generation of the Serval Mesh Extender. The contrast of the black and gold interior with the function-exposing transparent case makes a particularly strong first impression.

E. Enabling Additional Use-Cases

Aware that scale drives economy, we also spent some time to think about other potential use-cases for the Mesh Extender, so that we might be able to increase production volumes, and thus reduce unit costs. We concluded that a kind of off-grid Internet of Things (IoT) capability would complement the existing off-grid communications capabilities of the Mesh Extender. This could be used in humanitarian use-cases, for example, to help monitor or operate equipment remotely. It could also be used in farms in remote parts of Australia and other countries, where cellular connectivity is not available, and where the Mesh Extender has the potential to be strongly competitive with satellite-based and RF-based IoT solutions.

F. Minimizing Product Range

Every product you sell means additional support, and for MPP operations, can mean additional inventory management challenges. The last thing you want as a financially constrained endeavor is to have a pile of the wrong model of device. This is particularly relevant for projects involving wireless communications where regulations differ from country to country. For the Mesh Extender we sought to address this problem in three complementary ways.

1) Socketed Radio Module: First, we simply socket the RFD900 radio on the PCB, so that an RFD868 can be substituted for Europe, or a future RFD434 or RFD170 could be used to access other radio frequencies as the need arises, without having to re-engineer the PCB, or having to worry about which frequency particular units can support.

2) Regulatory Information in Cable: Second, the power cable for the Mesh Extender also includes a serial EEPROM, that can be programmed with the desired frequency, transmit level, maximum duty cycle and other country and region-specific parameters. This means that for all countries that can use a single model of radio, for example, the RFD900 in Australia, New Zealand, USA, Canada, Vanuatu and others, the country specific frequency, if any, can be specified in the cable. For humanitarian deployments, this means that we could, for example, have cables pre-positioned in a number of countries in the Pacific region, and have a fewer number of the Mesh Extender units themselves positioned regionally.

3) Radio Serial Interface Loop-Back: Third, for situations where the internal radio is insufficient, for example when interfacing to a Codan or Barrett HF radio, we route the serial port that connects the processor to the radio module out and in the power cable port. This means that for HF radio use, all that is required is a custom power cable that routes the processor's serial port to the HF radio instead of to the internally fitted RFD900. The software in the Mesh Extender automatically detects the radio type at boot time. This further allows us to have a single model of Mesh Extender, that can be used for a diverse range of situations.

V. Conclusion

We summarize our advice to other practitioners as follows: (1) Involve a beneficiary of what you are trying to create, to make sure you create what they really need, i.e., follow the principles of human centered designed; (2) Realize that funding comes and funding goes: try to make discrete and identifiable advances when conditions permit, so that you are better placed to obtain more resources later; (3) Realize that product development is a long and time consuming process: don't get disheartened if it takes longer than you think; (4) Iterate towards your goal, rather than trying to do everything at once; (5) Try to find a partner to help bring the fruit of your efforts to market; (6) Engage with others trying to achieve similar goals and learn from them what you can; (7) By following an open-source/inclusive methodology, it is possible to build a collaborating community, that can help to implement, assess, refine and make use of your creations; (8) Advocate, socialize and publicize your work, including through informal outlets such as blogs, social media and broadcast media/press, to both increase impact and also opportunities for finding appropriate partners, funding opportunities and beneficiaries.

The productization of the Serval Mesh Extender is not yet complete, however, it has progressed considerably. In the process, we have encountered a number of challenges and situations that we hope will be of value to other practitioners who are seeking to create humanitarian innovations.

Acknowledgments

The authors would like to acknowledge the support of USAID (USA), the NLnet Foundation (Netherlands), Radio Free Asia (USA), the Shuttleworth Foundation (South Africa), the Arkaroola Wilderness Sanctuary (Australia), and the Australian Foreign Aid Program within the Australian Department of Foreign Affairs and Trade (Australia) and Associate Professor Sandy Walker.

References

[1] P. Gardner-Stephen, M. Lloyd, A. Bate, A. Cullen, and S. Farouque, "Piloting the serval mesh and serval mesh extender 2.0 in vanuatu: Preliminary results," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[2] P. Gardner-Stephen, "The serval project: Practical wireless ad-hoc mobile telecommunications," *Flinders University, Adelaide, South Australia, Tech. Rep*, 2011. [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.460.1778\&rep\-=-rep1\&type\-=-pdf

[3] P. Gardner-Stephen, J. Lakeman, R. Challans, C. Wallis, A. Stulman, and Y. Haddad, "Meshms: Ad hoc data transfer within mesh network," *International Journal of Communications, Network and System Sciences*, vol. 5, no. 08, p. 496, 2012.

[4] P. Gardner-Stephen, A. Bettison, R. Challans, and J. Lakeman, "The rational behind the serval network layer for resilient communications," *Journal of Computer Science*, vol. 9, no. 12, p. 1680, 2013.

[5] P. Gardner-Stephen, R. Challans, J. Lakeman, A. Bettison, D. Gardner-Stephen, and M. Lloyd, "The serval mesh: A platform for resilient communications in disaster & crisis," in *Global Humanitarian Technology Conference (GHTC), 2013 IEEE*. IEEE, 2013, pp. 162–166.

[6] P. Gardner-Stephen, A. Bettison, R. Challans, J. Hampton, J. Lakeman, and C. Wallis, "Improving compression of short messages," *Int'l J. of Communications, Network and System Sciences*, vol. 2013, 2013.

[7] P. Lieser, F. Alvarez, P. Gardner-Stephen, M. Hollick, and D. Boehnstedt, "Architecture for responsive emergency communications networks," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[8] W. Kongsiriwattana and P. Gardner-Stephen, "Historical distribution of duration of unplanned power outages in queensland: Insights for sustaining telecommunications during disasters," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[9] ——, "Eliminating the high stand-by energy consumption of ad-hoc wi-fi," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[10] Ghassan and P. Gardner-Stephen, "Scalable telecommunications over ultra-low-bandwidth radio backbones," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[11] M. Lloyd, A. Bate, A. Wallace, and P. Gardner-Stephen, "Zero recurrent-cost two-way satellite communications for humanitarian applications," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[12] R. Shuman, J. Kim, and P. Gardner-Stephen, "Optimising the question box for cost and local-manufacturability," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[13] J. Lakeman, R. Challans, M. Lloyd, A. Wallace, M. Schmittner, M. Hollick, and P. Gardner-Stephen, "A practical and secure social media facility for internet-deprived populations," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[14] K. Hawtin and P. Gardner-Stephen, "Self-configuring heterogeneous hf/uhf/wi-fi disaster communications networks," to appear Global Humanitarian Technology Conference (GHTC). IEEE, 2017.

[15] J. Liker, *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer, First edition*. McGraw-Hill, 2003.

[16] RFDesign, "Rfd900 radio modem data sheet," 2013. [Online]. Available: http://files.rfdesign.com.au/Files/documents/RFD900\%20DataSheet.pdf

[17] D. Adams, *The Hitchhiker's Guide to the Galaxy*. Pocket Books, 1979.

[18] "Zx spectrum vega plus backers complain of months-long refund delays," 2017. [Online]. Available: https://www.theregister.co.uk/2017/03/29/retro_computer_ltd_vega_plus_refunds/

[19] "Potato salad, kickstarter campaign," 2014. [Online]. Available: https://www.kickstarter.com/projects/zackdangerbrown/potato-salad

[20] B. Kuchera, "$55,000 "jok" potato salad kickstarter concludes with party, large donations to charity," 2014. [Online]. Available: https://www.polygon.com/2014/9/30/6872423/potato-salad-kickstarter-party-columbus

[21] B. Zack and T. Blackburn, *The Peace, Love and Potato Salad Cook Book*. Springhouse Press, 2016.

[22] J. Brown, "Combat kitchenware," 2013. [Online]. Available: https://www.kickstarter.com/projects/1901693996/combat-kitchenware

[23] T. Rushton, "5 o'clock shadow," 2012. [Online]. Available: https://www.kickstarter.com/projects/5oclockshadow/5-oclock-shadow?ref=category

[24] L. Canning, "Less car more go: the cargo bike documentary," 2014. [Online]. Available: https://www.kickstarter.com/projects/1887563980/less-car-more-go-the-cargo-bike-documentary
[25] StudioBananaThings, "Ostrich pillow," 2012. [Online]. Available: https://www.kickstarter.com/projects/ostrich-pillow/ostrich-pillow
[26] Matt, "Illumibowl 2.0: Toilet night light (shark tank upgrade)," 2016. [Online]. Available: https://www.kickstarter.com/projects/1585372093/illumibowl-2-toilet-night-light-sharktank-upgrade