The Serval Project: Practical Wireless Ad-Hoc Mobile Telecommunications

Dr. Paul Gardner-Stephen,
Rural, Remote & Humanitarian Telecommunications Fellow,
Flinders University and Founder, Serval Project, Inc.

July 22, 2011

1 Introduction

The purpose of this work is to give an introduction into the genesis and motivation behind the Serval Project, which the author believes to be the first practical mesh mobile telephony platform. As such, it focusses more on the why and the what, rather than the fine detail of the how. Fuller explainations of the various protocols, features and technologies developed under the Serval Project will be the subject of subsequent papers.

The remainder of this section provides the several motivations and the rationale behind the creation of the Serval mesh mobile telephony platform, code named Serval BatPhone, and in the process explores several representative use-cases. Section 2 provides brief introductions to each of the key technologies and features incorporated in the Serval BatPhone software, or in some cases, that are in the process of being incorporated. It is anticipated that the features described in this paper will be essentially finished and available in the relevant source code repositories by the time that this paper is published[1]. Section 3 explores how these features address the various use-cases introduced in this initial section. Initial results from trials of the BatPhone technology and associated discussion is presented in Section 4, with future directions and conclusions constituting the final two sections.

1.1 Mobile Communication Is Valuable

Mobile communications have become an intrinsic component of modern economies and societies because of the many benefits they afford. Indeed, more than 5 billion people have a cell phone as of 2011, and in the countries with the highest GDP-per-capita levels the penetration rate of cellular telephony services

[1] `http://github.com/servalproject`; On-line; Accessed 14 June 2011.

exceeds 100%[2]. For those who are able to enjoy the regular use of a cellular telephone the benefits extend into the economic and social spheres.

Economically, the ability to communicate over distance has the potential to save time, transport costs and the expenditure of human energy. Communicating over distance also allows for improved economic efficiency by allowing for better coordination of actions. For example, primary producers, e.g., farmers and fishermen may contact markets to determine what species and what markets will result in the greatest income for them, and correspondingly, the greatest potential value to the markets through provision of the most desired goods. Moreover, the ability to contact markets potentially enables primary producers to negotiate more effectively with intermediaries who then on-sell their goods at markets. This is because if the primary producer can gain accurate knowledge of market conditions they are able to possess information that will empower them to identify and react to unfair pricing on the part of the intermediaries, resulting in a market that is nearer to ideal. These benefits are in addition to the ability to access and share information that stimulates and acts as an enabling technology for all manner of enterprise which are beyond the scope of this paper. Recent studies confirm that mobile telephone, computer and Internet access all have a positive effect on GDP [6].

Socially, the ability to communicate over distance is continuously demonstrated to enable individuals, communities and nations to sustain familial and cultural connections. Indeed in the richer nations a substantial percentage of younger people are now quite dependent on mobile telecommunications to manage and facilitate their social lives. This extends beyond simple telephony to an increasing range of data-oriented services. This trend began with the Short Message Service (SMS) of the Global System for Mobile-communications (GSM) which carriers initially believed was not useful due to its short message limit (144 eight-bit bytes or 160 seven-bit characters). In addition to SMS, Internet services have become ubiquitous, such as email and Facebook, which are oriented around longer messages, and Twitter, which is oriented around SMS and has even tighter bounds on message length. Together with general Internet access itself, these services have become popular tools continuing the trend towards digital mobile communications services. Indeed, the United Nations Human Rights Council has effectively declared that Internet access is a human right, derived from the basic human right of freedom of expression[9].

Carrier's belief that SMS was not socially useful and is has negligible provision costs is evidenced that in the early years of GSM networks it was possible to access SMS services for free in most markets. Were either of these statements false, carriers would have charged for the service. Indeed, it was not perhaps until the free availability of this service led to the broad adoption of "texting" that carrier's saw that they had a product they could charge for, and indeed they now typically levy charges that are quite out of proportion to the cost of provision. They are able to do this because the services are now in demand, and the

[2] `http://www.itu.int/newsroom/press_releases/2010/06.html`; On-line; accessed 3 June 2011.

oligopoly or monopoly of carriers introduces a limited supply that the carriers are presumably comfortable to accept. Consequentially SMS messages typically cost around US$0.01 - US$0.25 depending on the market, which corresponds to a data cost of between about US$70,000 and US$1,820,000 per giga-byte – that is about 10 to 400 times more expensive than using a premium satellite internet service, and has been quoted as being several times more expensive than the cost for NASA to receive space imagery from the Hubble Space Satellite[3]. Indeed, it would be cheaper to fly first-class from Sydney to London return, carrying 32MB of data on a USB memory stick, capable of holding perhaps an hour of music, or a dozen or so photographs, to be hand delivered to its recipient, than to send the same data as SMS messages[4].

It should nonetheless be noted that mobile carriers do not generally make particularly obscene profits compared to other business sectors, and thus while there is almost certainly room for improvement in fair pricing, there remains that the carriers have to raise a certain level of revenue to meet their operating costs, and thus if carriers were to reduce the cost of SMS messages to something reasonably derived from the cost of provision that one of two things must occur. Either, the volume of SMS traffic will increase, thus replacing the lost revenue, or they will need to increase costs elsewhere, or equivalently, remove subsidies on some other service they provide.

Evidence points strongly in the direction of SMS volumes increasing to replace the lost revenue, as appears to be the case in the Philipenes where SMS costs were reduced to approximately US$0.005 without harming carrier profitability[8]. Indeed, in Australia the prevelance of "mobile caps" which include a very large SMS allowance suggest that carriers realise this, but are unwilling to actually reduce the list price for some reason, perhaps fear that they will be unable to lift it again in a competitive market should the experiment turn out poorly for them.

This theme of market distortion due to the natural oligopoly or monopoly of mobile network operators in any given market, and the artificial economy of scarcity that frequently results, is a theme which shall be revisited later in this paper. Meanwhile, we can infer that mobile telecommunications is clearly valued, as proven by the ability that carriers are able to leverage substantial charges, otherwise no one would pay their often disproportionate fees.

1.2 Current Cellular Service Provision Costs Have a Floor

We have suggested that SMS messages are routinely charged a rate which bears little relationship to its cost of provision, the retail charges for monthly cellular telephony subscriptions seem to bear a more direct relationship to the costs of operation. Minimum monthly subscription charges are surprisingly similar

[3] `http://www.physorg.com/news129793047.html`; On-line; Accessed 15 June 2011.

[4] Assuming a cost of A$0.25 per international SMS, which is typical of Australian mobile carriers for domestic in 2011, 32MB of data would cost A$58,254 to send by SMS. A return first-class ticket from Sydney Kingsford Smith (SYD) to London Heathrow (LHR) cost A$19,382 as quoted by `http://qantas.com.au` on 29 June 2011.

Table 1: Monthly Cellular Subscription And Basic Usage (30 Minutes Air-Time) Costs For Major Australian Cellular Carriers. Plans with contracts longer than one month are not considered.

Carrier	Plan	Cost	Air-time*	% GDP per capita per annum
Telstra	Casual Plan $30[6]	$30	34 min	1.05%
Optus	Optus Online Cap BYO $19[7]	$19	111 min	0.67%
Vodafone/ Three	Three $20 Sim Only Cap [8]	$20	180 min	0.70%

* The air-time is calulated from the included call value divided by the per-minute cost to call another cellular phone in Australia. Note that in Australia receiving a telephone call on a cellular phone is generally free. We have not considered the effect of flag-fall charges, which can be considerable for the cheaper plans.

globally, especially when related to GDP per capita. In Australia, one of the richest countries (2009 GDP per capita US$34,327[5]), the cheapest monthly subscriptions that include a reasonable number of calls (arbitrarily defined as 30 minutes per month) is approximately A$19 (US$18) per month (See Table 1), or approximately 0.67% of GDP per capita. In much of Africa the minimum monthly spend varies between prices that are similar to those in Australia either in terms of dollars or GDP per capita (Table 2). It is pleasing to see that in some markets carriers now offer considerably cheaper plans, around the AUD$1/month range, however this is still rather piece-meal, and some of the poorest countries, such as Burkina Faso, have some of the highest cellular services costs compared with GDP per capita.

It could be that the similarity in pricing reflects a global approach to pricing of cellular telephony subscriptions, and there is likely an element of that, given the suggestive uniformity around 1% GDP per capita per annum for 30 minutes of air-time per month. However, this does not explain the plateau effect that occurs at lower GDP per capita levels where the fundamental cost seems to be somewhere above AUD$1 per month. There is also evidence to suggest that many carriers charging at such low levels are in fact making losses[12], a situation they are perhaps tolerating to build market-share, or to stimulate expansion of their markets to the lower-income populations in their markets. Nonetheless, we are greatly encouraged to see the price being driven down to such low levels on the basis that it may help many millions.

However, while these costs may be affordable as measured against the national average GDP-per-capita, we would argue that 30 minutes per month (about 1 minute per day, or potentially much less if flag-fall charges are levied) is insufficient to create an economy of abundance that enables communities to

[5] `http://gapminder.com`; On-line; Accessed 14 June 2011.
[12] `http://www.ictworks.org/tags/safaricom`; On-line, Accessed 29 Jun 2011.

Table 2: Monthly Cellular Subscription And Basic Usage (30 Minutes Air-Time) Costs For Selected African Cellular Carriers. Plans with contracts longer than one month are not considered.

Country	Carrier & Plan	Cost (AUD$)	Air-time	GDP per capita US$ (2009)	% GDP per capita per annum
South Africa	MTN Call-per-second[9]	$8.28*	30 min*	$9,141	1.11%
Kenya	Safaricom Uwezo Tarrif[10]	$1.26*	30 min*	$1,494	1.01%
Burkina Faso	AirTel Pre-pay tarrif[11]	$5.60*	30 min*	$1,234	5.45%

* Price is calculated by multiplying per-minute rate.

realise the full benefits of mobile telecommunications.

What these data suggest is that dominant cost of operating a cellular network is the recurrent operational costs, which are linked to local prices, rather than the cost of cellular plant, which is priced on the international markets. This seems a reasonable conclusion, and one that has been suggested by others previously [13], who have noted that to operate a cellular base station in such settings requires substantial effort to power, conect, and defend the communications assets.

Diesel generation is typically the only practical means of powering remote cell towers, which requires regular refueling and defence of the generator and fuel tanks which constitute attractive targets for theft. Solar panels are too valuable and difficult to defend at the scale required to provide sufficient power to power a cell tower. Moreover, solar power introduces the need for a large battery bank, or indeed a diesel generator to ensure continuity of supply at night or during poor weather. Thus fencing, lighting, cameras and other defensive structures are required. Security patrols and response teams may also be required to provide further protection of the assets.

These measures add not only to capital and operating costs but increase the energy needs of the installation, increasing the fuel load and generator size and thus attractiveness of the site for attack. The need for live security cameras also increases the back-haul link capacity to the site, further escalating energy consumption, spectrum usage and overall operating costs. These costs are further compounded in many instances by low population density, low economic density in terms of GDP per square kilo-metre, or both, mean that there are relatively few people from whom to recoup the cost of the cell tower in which

[13] `http://openbts.blogspot.com/2009/01/what-stuff-costs-part-2-opex.html`; Online; Accessed 14 June 2011.

they live in the shadow.

Thus, while labour costs may be lower in such nations, the economic inequalities, and low economic and population densities mean that the actual minimum cost per potential subscriber end up being roughly equivalent to urban areas in richer nations where the major costs probably come from spectrum licensing rather than physical protection and energy consumption. These factors push the real cost of provision in rural areas much higher, likely above the US$6/month mark[14], and well out of reach of the rural poor in many countries.

There is also another significant issue, which is that GDP-per-capita is misleading figure, as there is often very great variation between the richest and poorest in a country. For example, in 2006 in India, the GDP-per-capita varies from $732 in Bihar to $4,596 in Dehli[15], with the national GDP-per-capita for India in 2006 being $2,300[16]. Thus the poorest in most countries are unlikely to be able to afford cellular services.

These observations provide a potential explanation as to why the poorest billion or so do not own mobile telephones. It also suggests that there may be hundreds of millions more who, while they own a mobile telephone, are not able to afford to make full or regular use of them. However, it is not merely economic situation that affects the ability of people to enjoy the benefits of mobile telecommunications.

1.3 Unrest, Disaster & Emergency

Another significant and substantial population who lack adequate access to mobile telecommunications are those who encounter extra-ordinary events that can disable, disrupt or overwhelm mobile telecommunications infrastructure. This covers a variety of situations, including: (a) war, where infrastructure may be purposefully targeted; (b) terror attack, which is a specific form of war where combatants are not sanctioned by any state, and thus is more properly described as the land and air based equivalents to piracy; (c) adverse weather that may disrupt the logistical supply chain to mobile communications infrastructure, e.g., energy supply, or indeed affect the infrastructure itself, or both; (d)

[14] `http://openbts.blogspot.com/2009/01/what-stuff-costs-part-2-opex.html`; On-line; Accessed 14 June 2011.

[15] `http://www.gapminder.org/labs/gapminder-china-india-eu-usa/#$majorMode=chart$is;shi=t;ly=2003;lb=f;il=t;fs=11;al=30;stl=t;st=t;nsl=t;se=t$wst;tts=C$ts;sp=2.9316129032258;ti=2006$zpv;v=0$inc_x;mmid=XCOORDS;iid=pp59adS3CHWfKPVb7dEexFA;by=ind$inc_y;mmid=YCOORDS;iid=pp59adS3CHWeR0Ufcou95MQ;by=ind$inc_s;uniValue=20;iid=pp59adS3CHWcajNS5Y44uLw;by=ind$inc_c;uniValue=255;gid=CATID1;by=grp$map_x;scale=log;dataMin=240;dataMax=152481$map_y;scale=log;dataMin=1.569;dataMax=211$map_s;sma=50;smi=2$cd;bd=0$inds=;example=86`; On-line; Accessed 15 June 2011.

[16] `http://www.gapminder.org/world/#$majorMode=chart$is;shi=t;ly=2003;lb=f;il=t;fs=11;al=30;stl=t;st=t;nsl=t;se=t$wst;tts=C$ts;sp=5.59290322580644;ti=2007$zpv;v=0$inc_x;mmid=XCOORDS;iid=phAwcNAVuyj1jiMAkmq1iMg;by=ind$inc_y;mmid=YCOORDS;iid=phAwcNAVuyj2tPLxKvvnNPA;by=ind$inc_s;uniValue=8.21;iid=phAwcNAVuyj0XOoBL_n5tAQ;by=ind$inc_c;uniValue=255;gid=CATID0;by=grp$map_x;scale=log;dataMin=295;dataMax=79210$map_y;scale=lin;dataMin=19;dataMax=86$map_s;sma=49;smi=2.65$cd;bd=0$inds=i101_t002007,,,,;modified=75`; On-line; Accessed 15 June 2011.

disaster, such as earthquake, flood, bush-fire or fire-storm which may damage, inundate or isolate mobile communications assets via a variety of means, including disabling the back-haul that connects the assets to the rest of the network, reducing the effective range of the cellular signal, e.g., due to heat and smoke from fire-storms; (e) civil or other emergency where a surge in demand results which overwhelms the capacity of the infrastructure to provide service.

It is difficult to estimate the total number of people who are affected by one or more of these scenarios, but it would seem safe to estimate that it would be in the millions, and that the acute needs which result from these situations constitute a compelling need for a mobile telecommunications solution that at least partially addresses these use-cases.

1.3.1 When Goverments Deny Their Own Citizen's Mobile Telecommunications

There is another form of disruption to mobile telecommunications which has been seen during 2011, being the purposeful act of a government, de facto or otherwise, to disable telecommunications assets in order to prevent the general populace from communicating, perhaps from fear of commotion, riot, revolt or revolution, or of being held accountable for the conduct of government agents, such as riot police, through the recording and transmission of sound, images and/or video footage to other parts of the nation or to the international community. We mention this scenario, not because it is an intentionally selected use-case of our technology, but because any technology which can sustain or establish communications in the preceding use-cases will almost certainly be possible to use by citizens or subjects of a nation which so chooses, legally or illegally, justifiably or unjustifiably, to seek to prevent its populace from enjoying the benefits of mobile telecommunications at any time or place through disabling or preempting telecommunications infrastructure. We make no judgment on this, but do feel compelled to make readers aware of this, and refer the interested reader to [9].

1.4 Nomadic & Remote Populations

One population sector which has been systematically denied, not through purpose but through a disinterested neglect is that of nomadic and remote populations. These two populations are equivalent in so far as they both imply a geographical area which, on average, has too low a population and/or economic density, as expressed in GDP per area, such that infrastructure oriented mobile telecommunications is simply not economically feasibly to deploy. This makes it difficult to provide these populations with health care, education and a variety of other services that are vital to ensure equal opportunity and respectable life and health expectancy for these populations, whilst still allowing these populations to retain their cultural identity and land-bond, which are basic human rights.

In many cases it is the first-nation populations of many other wealthier countries are, the dominant population. Examples may be found in Australia, Canada, the United States of America, Russia and in Arabian countries among others. At the risk of gross generalisation, these first-nation populations often retain a strong cultural attachment to specific regions of their enclosing countries, and either continue to live in areas remote from the rest of the population, or travel the land nomadically to some degree, or often, both. Those populations who have similar land bonds, but to land where populations and sedantry civilisation has built up, whilst they may suffer other difficulties, typically have access to the infrastructure serving those population centres, and thus are not included in this population.

1.5 WMNs as a Complement To Infrastructure-Oriented Cellular Services

What is clear from the preceeding discussion is that there are various populations who are not well served by the current infrastructure oriented approach to mobile telecommunications. However, it is also obvious that the infrastructure oriented approach does work very well for areas of high economic and density, albeit if the natural oligopoly of mobile carriers results in pricing distortion. Thus appropriate corrective action is not to replace the infrastructure-oriented model, but to supplement it with something that does not depend on static infrastructure, so that the under-served populations can have their needs met.

If the supplementary approach must not depend on infrastructure, on carriers unable or uninterested in solving the problem, on peace emerging in troubled places, then there would seem to be only one viable approach: the cell phone must be the only infrastructure, physical or administrative, that is required to form and operate such a supplementary network. This implies a system that can operate on some sort of cell phone hardware, does not require the use of licensed spectrum, does not require a carrier to relay calls between caller and callee, does not require a carrier or other authority to allocate telephone numbers, and is completely self-organising. Mobile Ad-hoc Networks (MANETs), also known as Wireless Mesh Networks (WMNs) has the potential to address these requirements. For such an effort to succeed, it is vital that the barriers to entry to be removed so far as possible.

1.5.1 Finding Spectrum

Perhaps the greatest challenge to implementing a WMN mobile telephony platform is finding a suitable air-interface that does not require acts of spectrum allocation on the part of any government. Indeed, for it to be truly successful, it depends on the availability of a common spectrum allocation in at least a substantial number of nations, and that it is possible to use to send and receive broadcast packetised traffic. This is complicated by the stated requirement that the spectrum must be usable on existing cell phone hardware. Fortunately there are solutions to this problem.

Table 3: Industrial, Scientific and Medical (ISM) Radio Bands, but note that these data are very generalised. Region 1 consists of most of Africa, Europe, the Middle-East and Russia. Region 2 consists of the Americas.

Band	Centre	Notes
6.765–6.795 MHz	6.780 MHz	Subject to local acceptance .
13.553–13.567 MHz	13.560 MHz	
26.957–27.283 MHz	27.120 MHz	
40.66–40.70 MHz	40.68 MHz	
433.05–434.79 MHz	433.92 MHz	Region 1 only and subject to local acceptance.
902–928 MHz	915 MHz	Region 2 and some others only.
2.400–2.500 GHz	2.450 GHz	Most commonly used for microwave ovens and 802.11 WiFi.
5.725–5.875 GHz	5.800 GHz	Commonly used for 802.11 WiFi.
24–24.25 GHz	24.125 GHz	
61–61.5 GHz	61.25 GHz	Subject to local acceptance.
122–123 GHz	122.5 GHz	Subject to local acceptance.
244–246 GHz	245 GHz	Subject to local acceptance.

The Industrial, Scientific & Medical (ISM) bands represent the most widely standardised permissive spectrum allocations globally. While there are a number of ISM bands (see Table 3[17]), there are two that have particular potential for a WMN applications in cell phones, for reasons which will be explained. These are the ISM 915MHz band, covering the UHF band from 915MHz - 928.5MHz, and the ISM 2400MHz band at 2400 MHz - 2500 MHz in the microwave band.

The ISM915 band is usually licensed up to 1 watt transmit power, and is thus commonly used for low-speed long-range serial communications, achieving ranges of up to 1 km in urban areas and up to 40 km at lower speed in more ideal conditions, such as in remote areas and over water. Typical applications include some wireless telephone handsets and baby monitors. The ISM915 band also sits conveniently between the GSM 850 uplink and downlink bands, which means that it should be possible for most cell phone chipsets to transmit and receive on that band using their existing GSM base band radios. However, the ISM915 band has the limitation that it is not available in all jurisdictions, and in the jurisdictions where it is available, the precise bounds of the band vary. For example, in countries that use GSM900, such as Australia, only the upper half of the band is typically available. A further difficulty that applies to the development stage of realising the utilisation of the ISM915 band is that the documentation required to program the base-band processors in cell phones is not readily available.

[17] `http://en.wikipedia.org/wiki/ISM_band`; On-line; Accesed 14 June 2011.

The long distances that it is possible to transmit using the ISM915 band also present a two-edged sword. On the positive side, it is extremely helpful in areas where the handset density is low, communication using the ISM915 band has the potential to allow the formation of WMNs that would not be possible with shorter-range communications. However, as handset densities increase, effects such as the hidden transmitter problem become significant. Also, the relatively low data-rates possible in the ISM915 band mean that bandwidth starvation will set in much earlier, and with much smaller mesh sizes. Thus, while the ISM915 band has potential in the long term, especially for in rural and remote areas, the initial focus on creating a functional system has been on another of the ISM bands: the ISM2400 band.

The ISM2400 band is perhaps the most widely standardised permissive spectrum allocation ever, and this achievement is a credit to those who strove to make this a reality. On the positive side, the ISM2400 band is extremely wide, and being in the microwave band can readily support high data rates, as evidenced by the proliferation of WiFi devices capable of up to 300mbits per second over a 40MHz wide channel. However, the ISM2400 band is subject to strong absorbtion by water, is notorious for multi-path interference due to its short wave-length, and is strongly poluted by microwave ovens. Indeed, the difficulties with the ISM2400 band are perhaps the reason why the band was made available to the general public, as it thus had little value to mobile carriers and other interests around the world. The fact that the band is nonetheless useful in the hands of the general public has been demonstrated by the incredibly strong adoption of WiFi over ISM2400, dispite its various handicaps, including those introduced by the 802.11 specifications themselves, particularly around secutity.

Indeed, dispite the many documented issues with the WiFi standards, it represents the only relatively long range ISM2400 air interface that is practical for WMN telephony, due to its combinations of relatively high bit rates, near ubiquitous deployment, and provision of a mesh-compatible mode, commonly called ad hoc mode. This is inspite of the poor standardisation and spotty implementation quality of ad hoc mode in WiFi drivers, and failings with the ad hoc mode specification itself, particularly around cell management that can result in interrupted communications, e.g., due to cell splitting.

Thus, while ad hoc mode WiFi over ISM2400 is not ideal, it does come sufficiently close to a simple broadcast transmission and reception system that can support a reasonably capable WMN mobile telephony solution, and was the only credible choice for doing so at modest cost and in reasonable time. For this reason the Serval Distributed Telephony Technology has been created using this platform as the first supported air-interface.

An issue that is common to any spectrum is the effect that over-use has on reducing effective range by increasing the noise floor. Emergency and disaster situations frequently result in the deployment of significant WiFi assets, sometimes at illegal power levels, which will limit, but not destroy, the usefulness of our technology. However, this only occurs once the relief agencies arrive. Until then, most disasters that are likely to knock out cell towers are also likely to knock out power to microwave ovens and WiFi access points, and thus the noise floor may be surprisingly low, and consequentially unexpectedly good range may be obtained with WiFi. For instance, a test carried out by David Rowe on behalf of VillageTelco.org demonstrated that with a little ingenuity and planning it is possible to obtain a good WiFi link over distances exceeding 2km, without the use of illegal power levels or directional antenna [18]. Rowe's test relied on raising a WiFi based mesh devices onto 8m poles on piers separated only by open waters, thus ensuring reasonably clear Fresnel zones and near-optimal conditions.

This may seem outrageous to achieve in a disaster zone or any other realistic situation, however this is not the case. Mesh networks actively discover the best route, without requiring any action on the user's part. Thus, having even one device located at a suitable vantage point can make considerable difference. This may be as simple as attaching a meshing cell phone to a long stick or pole.

1.5.2 Cell Phone as Infrastructure

If a complementary WMN based mobile telephony platform is to be realised, it must be able to operate, at least to some degree, using cell phones as the only components of the network. To do otherwise is to reintroduce the dependence on infrastructure. It has already been described how the decision was made to use ad hoc WiFi in the ISM2400 band as it is the only feasible spectrum allocation and air interface that is readily available. Thus it becomes apparent that the solution must rely on WiFi-enabled cell phones in the first instance, and that these cell phones must be use ad hoc mode, rather than rely on access points for comunication. Rather frustratingly, several of the major cell phone operating systems do not provide APIs for making use of ad hoc mode in the flexible manner that is required to create a WMN based mobile telecommunications solution. This is perhaps because there has been limited demand, and other features always take priority. Thus, all WiFi enabled cell phones could participate in a WMN mobile telecommunications network, but may require support from the software or hardware vendors to achieve this. We now summarise our understanding of the current situation for the major platforms, which we confess varies in detail from platform to platform based on where our efforts have been concentrated, in particular we have not made significant explorations into the feasibility of implementing our system on the Blackberry/RIM or WinCE platforms. Other platforms are considered below.

Symbian and Windows Mobile 6 are similar from our perspective, in that we have strong reason to believe that it is quite possible to port our software to these platforms, and that the relatively flexible or permissive security models of these platforms would enable us to manipulate the WiFi interface and manage the network routing table. Indeed there are applications that implement these capabilities for quite different purposes, often providing WiFi tethering, i.e., acting as a cellular to WiFi adapter to allow laptop computers to use a phone's

[18] `http://www.villagetelco.org/2010/08/v1-3-antenna-testing/`; On-line; Accessed 14 June 2011.

cellular data services. Examples include JoikuSpot[19] on Symbian and Windows Mobile WiFi Router[20]. The drawback with these platforms is that they are both officially being dropped by their vendors for use on new devices in the near future. Symbian does retain some attraction, however, as it is likely to be present on a large number of Nokia devices for some time, especially in developing countries and on lower-end devices.

Conversation with a representative from Microsoft in January 2011 indicated that it would likely be some time before Windows Mobile 7 would have the necessary APIs to accomodate our software.

The Apple iOS platform has the ability to join an existing ad hoc wireless network, but not to create one, thus Apple iOS devices could participate in such a WMN based telephony solution, provided that they come into contact with at least one other mesh device or Apple chose to enhance their ad hoc network support. It would be possible to prototype the necessary capabilities by jail-breaking an iPhone, however this is not an ideal solution for distribution, especially as Apple have been suspected to take action to disable phones that have been jail-broken. Also, the iPhone is a relatively expensive cell phone, thus limiting the potential market to wealthier populations, thus missing the very populations we most wish to help.

The Android platform does not have support for ad hoc WiFi via official APIs. Similar to the iOS platform, it is possible to wait for Google to add support to some future version of Android, or to obtain root, i.e., super-user or administrator, access on the phone to perform these operations at a level below the public APIs. While some Android handset vendors have been known to push updates that revoke root access, this is relatively rare, and does not result in permanent disabling of the device. Also, there exists relatively mature software for easily obtaining root access on a wide range of Android phones. Finally, Android phones are produced my many vendors and are available at relatively low cost, including some devices below AUD$100 (for the Huawei IDEOS U8150, network locked), and indeed Android represents the platform with the highest sales as of 2011[21].

Our research group already has extensive UNIX development experience, enabling us to leverage the opportunity to manipulate the WiFi hardware on Android phones at a low level. Combined with the low cost of the handsets has made it the logical choice for our initial implementation. It remains our intention to port our software to all major mobile platforms, including the deprecated platforms, to maximise the opportunity for people to benefit from our technology.

---

[19] `http://www.joiku.com/products/joikuspot_light`; On-line; accessed 3 June 2011.

[20] `http://www.wmwifirouter.com/`; On-line; accessed 3 June 2011.

[21] Nielsen Mobile Insights, National (USA), March 2011 and Millennial Media April 2011 report.

1.5.3 Automatic Network Configuration and Carrier-Independent Self-Allocation of Telephone Numbers

Irrespective of the mobile platform used, it is necessary for each phone to be allocated an IP address for use on the WiFi interface. Similarly, it is necessary to have a means of mapping telephone numbers to handsets. What confounds this process is that to create a robust infrastructure-free mobile telecommunications solution is that it is necessary to be able to perform these functions without reference to any authority, server or infrastructure. Security and convenience are rather contradictory, and thus sensible trade-offs must be made. In this case, the convenience of being able to immediate claim your telephone number versus the security for a caller knowing that they will be connected to the correct person.

The approach taken by the Serval Project is to not inhibit claiming of telephone numbers, as that would prevent communication in isolated settings. Rather, such claiming of a telephone number without any attesting authority creates and untrusted identity. Any party contacting that untrusted identity will receive a warning message that the identity of the untrusted identity cannot be confirmed. This moves the otherwise complex software problem into a simple and frequently performed human cognition problem of voice recognition of the untrusted identity. At the conclusion of a call, the calling party can be prompted to indicate if the untrusted identity was indeed who they were expected to be, and if so, to record that decision. It is then possible for the calling party to sign a digital certificate expressing confirmation that the subscriber should indeed have the telephone number concerned. This can then be presented by the hithero untrusted identity to any other calling party. If they have separately elected to trust the judgment of the previous calling party, then the presence of the certificate can be used to subpress the warning.

1.6 Summary

It seemed to us that there is compelling need and utility for a mobile telecommunications system that can continue to operate without the support of any infrastructure, phyiscal or organisational, and thus allow communications to continue, or be possible for the first time, in a variety of scenarios. Moreover, is appeared that the necessary preconditions existed for us to create a system that has the potential to achieve this goal. The remainder of this paper reports on the progress of this endeavour, and the technologies we have created to progress towards the goal.

2 The Serval Distributed Telephony Model

The following four sub-sections describe the key technologies that underpin the Serval Distributed Telephony Model, which is our attempt to create an architecture that addresses the issues with the incumbant exclusively infrastructure oriented approach to mobile telephony in particular. The manner in which these interact to create communications solutions that address the use-cases decsribed in the previous section is discussed in Section 3.

2.1 VillageTelco Foundation

The Serval Project can be considered to be a derivative of The Village Telco project[22], specifically of their Mesh Potato[23] device. The Mesh Potato is an unusually robust meshing WiFi router with integrated FXS analog telephone port designed to provide local fixed-line style communications. As the initial target market for the Mesh Potato is developing regions in Africa and Asia, it is designed to tolerate up to 240v mains on any port, including telephone and ethernet ports, and to operate on a minimalist energy budget met by feeding from 9v - 35v DC into either a dedicated port or onto power-over-ethernet or power-over-phone-line ports. The Mesh Potato software consists of an embedded Linux distribution which includes an instance of Asterisk[24], an open-source PBX software suite that drives the telephone line, and relays calls between Mesh Potatoes.

A pivotal innovation by the Village Telco team was to configure the Asterisk instance so that it can operate in a fully-distributed and zero-idle-traffic manner. This is achieved by having the dialing plan based on IP address. Thus, to dial one Mesh Potato from another one dials the IP address, or more usually the last octet of the IP address. This causes the Asterisk instance on the caller's Mesh Potato to attempt a SIP connection to a well-known extension on the callee's Mesh Potato, which then causes the callees telephone to ring. This architecture allows Mesh Potatoes to communicate without any prior knowledge of the network, and without any central authority, although IP address collision is a hazard in truly decentralised deployments.

As will be seen later, the Serval Project extends this foundation by adding a telephone number to SIP address mapping layer (described in Sub-Section 2.2 below), thus converting dialled numbers into SIP addresses that are then contacted using the same method as that employeed on the Mesh Potato. Indeed, it is possible for a Mesh Potato to dial a mobile telephone running the Serval Project software, and vice-versa. The open-source nature of both projects also means that the Serval telephone number mapping layer is able to be integrated into the Village Telco software distributions, and this is likely to occur before the end of 2011, subject to availability of resources.

2.2 Serval Distributed Numbering Architecture: ARP For Phone Numbers

IPv4 space has been effectively exhausted as of early 2011[25], and also suffers from a small overall address space which represents a particular problem for

[22] `http://villagetelco.org`; on-line; accessed 5th June 2011.
[23] `http://www.villagetelco.org/mesh-potato/`; on-line; accessed 5th June 2011.
[24] `http://asterisk.org`; On-line; Accessed 14 June 2011.
[25] `http://www.nro.net/news/ipv4-free-pool-depleted`; On-line; Accessed 14 June 2011.

automatic network configuration in a fully distributed mesh network where it is not possible to have a central authority allocating addresses from a pool. The only practical approach is to have each device randomly select an IP address from some pool of addresses. Even if all 2^{32} IPv4 addresses were available, the birthday paradox[3] will result in an unacceptable probability of address clashes once meshes of more than a few tens of thousands of users emerge. Even hash-based approaches using ethernet addresses or some similar unique identifier from the handset itself will only serve to delay the process, even without accepting that duplicate ethernet addresses do exist. Fundamentally, the problem is that 2^{32} is less than the number of cell phone subscribers[26], and even less than the total human population as of 2011. IPv6 has certain attractions, however, it remains that there are many devices that do not support IPv6, and so for the time being we must resign ourselves to using IPv4, and take appropriate steps to mitigate the risk of address collission. This is accomplished by separating subscriber identities and telephone numbers from the IP addressing layer.

In the Serval system, individual subscribers are identified by a Subscriber ID (SID), which takes the form of a public key in an appropriate asymmetric cryptography system. The subscriber's phone retains the private key so that transactions can be authenticated and encrypted. This is the basic unit of addressing in the Serval Distributed Numbering Architecture (Serval DNA). Each SID and associated information is stored in a small database maintained by the Serval software on the mobile telephone, called the Home Location Register (HLR) by analogy to the structure which serves a similar purpose on GSM/3G cellular networks.

Each SID has one or more associated telephone numbers referred to as Direct Inward-Dialled (DID) numbers, and one or more locations (usually SIP addresses) on which the subscriber can be reached. By convention, the location field is usually "4000@IP", where 4000 is the well-known local telephone extension, as used on Village Telco Mesh Potatoes, and IP is substituted with the IP address of the callee.

In practice, the IP field is left blank in the location record, and is filled in with the appropriate IP address depending on which interface the request for resolution arrived. This allows for devices to have multiple network connections simultaneously, e.g., mesh and cellular data service, and present the correct location in response to requests from each. This approach also facilitates the changing of the IP address of the phone in the unlikely event that an IP address collission does occur, and also allows for the use of different underlying address families in future, e.g., IPv6 instead of IPv4.

The Village Telco uses BATMAN[2] as the underlying mesh routing protocol, a decision which the Serval Project has followed in the immediate term for interoperability, and because BATMAN has proven to have good performance in the Village Telco's initial Mesh Potato deployments. However, work is currently underway to enable the use of olsrd[27], and other mesh routing algorithms will

[26] `http://www.itu.int/ITU-D/ict/statistics/`; On-line; Accessed 14 June 2011.
[27] `http://olsr.org/?q=olsr_on_serval`; On-line; Accessed 14 June 2011.

be considered on a case-by-case basis to increase the potential for interoperation with existing mesh network deployments.

In the future, it is planned to adapt the mesh operation to a user-land overlay relying solely on RFC922[4] local broadcast of UDP packets, i.e., packets addresed to 255.255.255.255, for all communications. This would allow all devices on a Serval mesh to use the same IP address, or possibly addresses on incompatible networks, and still be able to communicate. This also makes better use of the naturally broadcast nature of wireless communications to allow multiple phones to receive a single message, a feature that has significant utility for messaging, software and data distribution over a mesh, such as the MeshMS and Rhizome services described later in this paper.

Whatever the mesh routing algorithm used, the creation of a mesh mobile telephony with acceptable scalability is that the mesh routing protocols must be able to handle routing in a small patch of a larger mesh without consuming all available bandwidth. For BATMAN, this modification should be relatively easy to achieve, as each node can simply announce the highest-scoring subset of those nodes it has received announcements from, and possibly to modify time-to-live (TTL) handling of announcements, such that announcments originating nearby take priority over those originating further afield. This could be done by employing packet filtering rules that drop BATMAN packets with TTL = 255-(maximum hops), where maximum hops is periodically recalculated based on the number of packets forwarded per second.

The process of establishing a mesh telephone call then becomes one of broadcasting the DID, i.e., called number, over the mesh and requesting any corresponding SIDs. It is possible that no SIDs will be obtained, in which case the call cannot be completed over the mesh, as the dialled party is not reachable. Otherwise, one or more SIDs will be obtained. If exactly one SID is received, the call can be progressed by requesting the location the SID can be reached at, usually a SIP address, and then having Asterisk request a connection to that address. For efficiency, the request protocol allows the DID->SID->location resolution to occur in a single transaction where a DID is supplied, and the reply will contain both the SID and location information.

This simple case where a single response is received is illustrated in Figure 1 where telephone A wishes to call telephone number 5552600. It begins by broadcasting a Serval DNA request for the location (SIP address) of the subscriber to that number (1). Telephone C responds with a SIP address that can be used to call that number (2). Finally, telephone A established the call using the supplied SIP address (3).

2.3 Security and Authentication

Note that the preceeding description of the Serval DNA protocol does not provide any assurance that the party being called is in fact the party that will respond and that will receive the call. Indeed, as described, the protocol makes it trivially simple for anyone to claim any phone number they may choose. The alternative is to have numbers allocated by an authority. However, this may

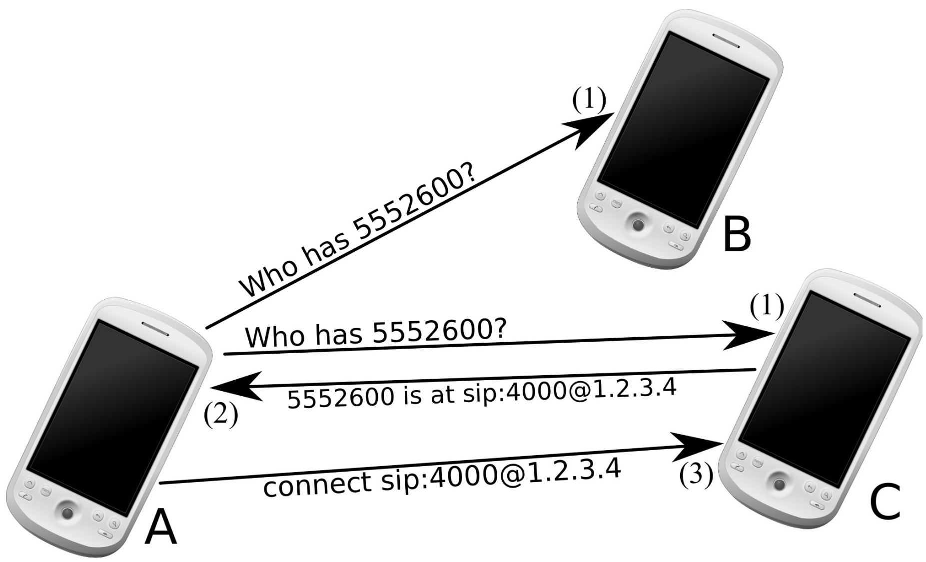

Figure 1: Simple Call Resolution Using Serval Distributed Numbering Architecture.

not be possible during a disaster, and it may be impractical for persons living in remote locations. Also, the creation of a number issuing authority creates a central point of failure, and would introduce an ongoing operating expense. Thus a solution is required that still allows immediate claiming of telephone numbers, but introduces appropriate assurances of authenticity. The Serval Project has two complementary solutions to this problem.

The first solution is to require users to record a voice signature, e.g., by having their name recorded by the telephone. A callers phone can request the transmission of the voice signature, and play it to the caller and prompt them if they wish to connect. This can be done for each SID, thus also allowing for discrimination between multiple claimants of a single telephone number.

This approach shifts the caller identification problem from software, where it is essentially impossible given the lack of a central authority, to the brain of the caller, where for persons known to the caller, it is likely to be trivial. This method is not without fault, as it is possible for a malicious party to request the voice signature of a 3rd party, and then use it to impersonate the 3rd party. For this reason the Serval software presents a brief message indicating that it cannot be assured that the call will connect to the correct party, and that the caller should take appropriate steps to verify the identity of the party at the other end of the call.

The very significant value of this approach is that it does allow people to claim their existing telephone numbers immediately, and without recourse to an authority, allowing the instant establishment of communications whether in a disaster, remote location, or using a replacement telephone. Nonetheless, improvement is possible through appropriate use of computer memory and cryptographic functions as will now be described.

The first improvement is that at the conclusion of a call we can ask the caller whether they are satisfied that the person was the rightful owner of the telephone number. If so, then the DID:SID pairing can be remembered, and calls to that DID can proceed without confirmation if the SID matches that remembered. At the other extreme the caller may elect to never be connected to that SID when requesting the DID. In between, the user can remain unconvinced one way or another, and no action needs to be taken, and subsequent calls to the called SID will be treated the same as though no previous call had been made.

The second improvement is to have each SID the public key of a public-private key pair. This offers many benefits. First, for authentication it prevents SID spoofing, as responses to DID mapping requests can be signed with the responding SID's private key. Second, once the caller is satisfied that the callee is who they claim to be, or more correctly, have a rightful claim on their DID, their phone can sign a digital certificate with their private key and present this automatically to the callee's phone. This allows the callee to present this certificate to other callers in future. This allows the creation of trust networks.

These measures can also be used to authenticate Caller ID (CID) information from an inbound call, with untrusted DID information either explicitly marked, or simply not presented.

2.4 Serval Gateway: Connecting Isolated Meshes to Each Other and to the World.

The preceeding text describes how call resolution works on a mesh, which is able to be entirely stand-alone, self-administering and isolated from all else. Indeed, Serval mesh networks have been tested in remote areas of the Australian Outback[28] [1], as well as deep underground in caves[29], tunnels and underground hydro-electric power stations, without resort to any infrastructure. It is also possible to connect such an isolated mesh to the internet and/or Public Switched Telephone Network (PSTN).

The Serval DNA protocol allows for any device to provide a response to a solicitation to resolve a DID. To implement a gateway to allow out-bound dialing from a Serval mesh the gateway device need simply provide a response to every solitication, providing a SIP address that is valid from the caller's perspective. Gateway responses are differentiated from direct responses by having the extension of the SIP address begin with 41XX*, where X is any valid character, although sequences other than 4101* are reserved for future protocol extensions. Such responses are only considered by the caller if no other valid callee response is received.

In practice, this means that gateway devices will need to act as a way-point for the call, by offering a SIP address on the gateway's internal asterisk server

[28] `http://www.youtube.com/watch?v=K819ggBvkU0`; On-line, Accessed 29 June 2011.

[29] `http://www.youtube.com/watch?v=wYibygOdYMM`; On-line, Accessed 29 June 2011.

that when dialled triggers a call to the dialed number. This helps to traverse any network address translation (NAT) that may occur at the gateway, e.g., if the gateway's internet connection is via a cellular data service. This approach also allows nested layers of gateways forming a hierarchy linking very large numbers of small mesh networks, or overlapping patches of a larger mesh.

Gateways effectively broker in-bound calls when observed from the opposite side of the gateway. In some situations, this will be a DID obtained from a VoIP provider that will provide a hot-line number to dial into the mesh. In that situation, the hot-line number would greet callers with an interactive voice recording (IVR) that prompts them to enter the telephone number that they wish to call on the mesh, warn them that the identity of the called party cannot be assured, and then attempt to connect the call, or inform the caller if the desired number was not available on the mesh.

In other situations, the otherside of a gateway could be another mesh, in which case the operation of the gateway is effectively symmetric, relaying requests from one mesh to another, acting much as an ethernet bridge would. Care must be taken to ensure that the repeating of requests on other meshes does not result in duplicate requests, e.g., due to the presence of more than one gateway joining any two meshes. The Serval DNA protocol has a transaction field that can be remembered by a gateway, allowing apparently new requests to be compared with recent requests to prevent the gateway responding multiple times to any given request. It is also prudent to make use of a TTL on the request such that it will only spread through a limited number of gateways before expiring.

Call maintanence while roaming among meshes and mesh patches has been considered, and will be described in a later paper. Roaming from mesh to cellular and vice-versa, would require significant cooperation from carriers or the provision of some form of centralised infrastrcture, and is not a goal at this stage.

2.5 Serval MeshMS: General-Purpose Infrastructure-Free Short Messaging

Having covered basic telephony we now consider the provision of Short Message Services (SMS). It is in principle possible to make use of the broadcast nature of wireless networks to efficiently send multicast and broadcast messages to many recipients at the same time. We consider unicast SMS first, that is where one user of the system sends a message to exactly one other user of the system. Number resolution proceeds as for placing a telephone call, but instead of initiating a telephone call a Serval DNA packet is sent to the recipient containing the message, which is then acknowledged when received. This mechanism allows for messages of up to approximately 1KB, although shorter messages will have a greater chance of being delivered on a busy network due to the reduced probability of collision.

In the Android implementation, the Serval DNA service can directly trigger an intent which causes the Android system to acknowledge receipt of the system.

This functionality depends on SMSDroid[30] for reception. Sending SMS via non-cellular means on Android is complicated by the lack of a broadcast intent when the system is about to send a message. This contrasts with placing a telephone call where it is possible to intercept the call from the dialer when the user attempts to place a call. As a result of the lack of a suitable intent being broadcast, it is necessary to use an additional program, WebSMS[31] which allows the composition of SMS messages and their delivery via artibrary means via a plugin interface. The authors of WebSMS (private correspondance) have offered to include patches to their software distributions which will make the integration of these components more streamlined.

2.6 Serval Rhizome: Situation Awareness, Mapping, Disseminating Software, Updates & Messages Without Infrastructure

Previously it was observed that the naturally broadcast nature of wireless networks can allow efficient multicasting and broadcasting of messages. We generalise this by implementing a broadcast messaging platform that in some ways resembles Twitter[32], but without reliance on infrastructure. This platform allows the creation and dissemination of messages on a Serval mesh, and their appropriate processing in a manner similar to monitoring a Twitter feed. The content of the message itself can contain tags and keywords that can be used to filter received messages, with those of interest being presented to the user.

These messages may also contain geographic coordinates, in which case they can be pinned onto a map which can be queried in an integrated user-interface which will be present in the next release of the Serval software. This provides a Ushahidi[7]-like incident mapping and visualisation platform, but which can operate without access to the internet or any supporting infrastructure.

The platform has been designed to make it relatively easy to interconnect with Twitter and Ushahidi instances on the internet to allow information sharing, e.g., between people in a acute-phase disaster and authorities and relief agencies.

Message dissemination occurs through each node caching messages which have been seen recently in an on-phone database. At regular intervals the phone selects one or more records from the database and broadcasts them to the nodes immediate neighbours, who also perform the same process. Thus eventually all messages propagate throughout the entire mesh, regardless of size. Optimisations, such as not broadcasting messages that have been broadcast recently, and modulating the probability of each message being broadcast according to size, time stamp, sender and other criteria are under consideration.

To facilitate the exchange of larger files, a third message type exists which is that of File Manifest, which provides a (usually) signed certificate which

[30] `http://code.google.com/p/websmsdroid/`; On-line; Accessed 14 June 2011.
[31] `http://code.google.com/p/websmsdroid/`; On-line; Accessed 14 June 2011.
[32] `http://twitter.com`; On-line; Accessed 14 June 2011.

includes the name, type, version, checksum and other data regarding a file. These manifests are not broadcast by a node until it has the file itself. File reception is accomplished by sending a request to the peer that broadcast the manifest and requesting it to transfer a portion of the file in question. This is repeated until all portions of the file have been transferred. Other nodes may actively listen to the transmissions to collect the file data themselves, thus avoiding redundant transmission of data.

Nodes may implement schemes to select which files to request, cache, advertise and offer to peers. However, all nodes are expected to cache the most recent observed version of files that are signed by the Serval Project. Such files that represent newer versions of the Serval software installed on the phone will be offered for installation. This allows software updates to be distributed over the mesh, with no supporting infrastructure.

Additionally, the Serval software has a feature where it switches the phone to access-point mode, and offers the software required to join the Serval mesh to nearby phones. This allows for the Serval mesh telephony software to be both distributed and retro-installed in otherwise impossible situations where the is no supporting infrastructure, such as during the acute phase of a disaster or in remote or isolated locations, and enables many interesting and potentially life-saving use-cases.

2.7 Phone-To-Phone and Phone-To-Mesh Video Streaming

Almost all modern smart phones contain cameras capable of recording and compressing video in real-time. There are many situations where this capability would be useful for providing remote vision, for example in disaster relief, search and rescue, situation awareness, asset protection or recording of human rights violations and other illegal acts. Again, the ability to perform these functions without recourse to infrastructure or expensive equipment enables these capabilities to be used in situatons previously impossible. For example, a local community could record video of illegal police brutality, with the video streaming, live, to other phones on the mesh, thus rendering the confiscation of the phone of the recording party an ineffective ploy to prevent the release of the video. Indeed, un-manned phones could be placed strategically such that no person need be endangered at all. If the mesh were then connected via a gateway device the video could be streaming, live, to the internet thus thwarting almost any possibility of preventing the video reaching the global community, short of interferring with satellite internet services. It is hoped that this capability might serve to keep authorities accountable to the rule of law.

3 Addressing The Use-Cases

While there are many potential use-cases for the Serval BatPhone technology, such as allowing teenagers to text and call each other and exchange files without incurring cellular tarrifs, we concentrate now on four and describe briefly how the capabilities and features described in the preceeding section support and enable those use-cases.

3.1 Base-of-Pyramid Populations

For the purposes of this paper we consider base-of-pyramid populations to be those who cannot afford a traditional cellular service. For this population group, the Serval BatPhone software would allow them to create local mesh networks that could be used without cost, provided that they could afford compatible handsets. In mid-2011, it is possible to purchase compatible unlocked Android handsets for AUD$129[33], and compatible network locked handsets for under AUD$100 [34]. This is a significant reduction versus a year previously, when the cheapest compatible phones cost approximately AUD$250, and gives hope that the cost of handsets will not be a significant barrier. We also hold hope that one or more handset vendors will pursue the ISM915 path to allow the creation of very-low-cost handsets that are capable of participating in a Serval BatPhone mesh. The ability to create isolated mesh networks with no carrier or regulatory involvement makes it possible for poor communities to devise and implement their own telephony solutions, especially when the Serval BatPhone is combined with the Village Telco Mesh Potato to allow the creation of combined fixed and mobile telephony networks with no infrastructure cost.

The chance of success of this approach depends heavily on the ease of deployment, and considerable work is being undertaken by both the Serval Project and Village Telco to achieve this. The automatic network configuration measures employed in the Serval BatPhone software represents significant progress in that direction. The ability of the Serval BatPhone software to duplicate itself onto newly acquired handsets is also of significance here, as it allows the network to be grown by the local population in-situ, and without requiring the importation of equipment which is often costly and time-consuming.

Nonetheless, at this time it would realistically require someone with a degree of training to facilitate the creation and operation of such a network. This is not altogether bad, as it creates an opportunity for a micro-enterprise. Also, the relative simplicity of the system makes it reasonable to have local people training local people after a short period of time, a position that is supported by the achievement of this same goal by the Village Telco team working in Dili, Timor Leste, where there are now "second generation" local trainers who have been trained by locals and are training other locals.

[33] `http://www.ebay.com.au/itm/HUAWEI-U8150-ANDROID-V2-2-UNLOCKED-BRAND-NEW-SEALED-/150606681320?pt=AU_Mobile_Phones&hash=item2310db94e8` Huawei IDEOS U8150, unlocked, Buy-It-Now price AUD$129 plus postage. On-line; Accessed 14 June 2011.

[34] `http://www.crazyjohns.com.au/huawei-ideos-bl/`; Huawei IDEOS U8150, network locked, AUD$99 delivered. On-line; Accessed 14 June 2011.

3.2 Emergency, Disaster & Unrest Theatres

The emergency, disaster and unrest theatre use-case has much in common with the preceeding use case. This is because while victims of disaster may have financial wealth sufficient to procure new or additional equipment, it is effectively impossible to do so. Thus the ability of the Serval BatPhone software to be duplicated from phone to phone allows the build out of a network during the most acute phase of the disaster. Again, ease of use is key here, and immediate self-claiming of telephone numbers is vital, and is hoped to allow at least some families and friends in the theatre zone to reestablish contact, avoid the emotional paralysis of not knowing the condition of their loved ones, and so be better enabled to help others. In the longer term, it is hoped that the BatPhone software will come pre-installed on future models of cell phone, thus avoiding the need for a word-of-mouth spreading of the software.

The mapping, messaging and situation-awareness features of the BatPhone software are anticipated to be of great value in emergency and disaster theatres, however the challenge remains how to get users to realise that the software has these capabilities, and then to engage with them. One option is to have some mechanism where the software automatically activates and seeks the user's attention, perhaps when cellular coverage has been lost for some time, and/or on reception of some kind of emergency announcement message, e.g., as a result of one or more other BatPhone users activating some not-yet-existing feature in the software that causes it to express the need for the mesh to enter emergency mode. Some sort of voting or authentication algorithm would be required to reduce the likelihood of malicious parties raising a false alarm for whatever reason.

For relief agencies, the possibility of having their field staff pre-equipped with Serval BatPhone software on their cell phones offers many potential advantages. Appropriate maps, operational documentation and other material can be included in the installation, and custom applications can be built on the rhizome messaging platform to enable field workers to collaborate, coordinate and support their organisational processes, such as filling in electronic versions of forms instead of paper forms, and having the filled-in forms automatically upload via the nearest available gateway device, and receiving updated operational information via the gateways. The self-organisating nature of the mesh also means that multiple teams, perhaps from completely different organisations, working in a common geographic area will result in a stronger mesh, with potentially more gateway devices to provide capacity and redundancy.

The ability to add gateway devices is of value in emergency and disaster for tying the mesh, and the situation-awareness platforms into the internet to enable authorities and relief agencies to gain real-time situation-awareness, and to provide information and material into the theatre, for example up-to-date mapping tiles and resource availability information that will be of use to the people in the theatre zone. This ability to establish an effective two-way communications and situation-awareness platform using only existing cell-phones, if it proves to be as effective as we hope, has significant potential to save lives, reduce hardship and improve the overall effectiveness of disaster and emergency response.

3.3 Rural & Remote Communities

Rural and remote communities from our perspective can be characterised as areas or regions where traditional cellular coverage may be absent or marginal due to low population-density, geography, or both. In this situation the ability to maintain local communications is helpful, but it is interconnection to the PSTN and/or internet via gateway devices become the most attractive feature. For example, there are rural communities less than 5km from the edge of suburban Adelaide, the capital city of South Australia which have very poor to non-existant cellular coverage due to the heavily wooded and steeply undulating country.

Communities such as this are often capable and actively interested in the progress of their communities, and thus are willing to explore supplementary communications solutions. Thus it is entirely conceivable for these communities, for example, to find a location, possibly on a tall agricultural structure where there is cellular service, place a cellular BatPhone gateway at that point to make it possible for them to call out from the surrounding land, and to add additional mesh points as required to provide adequate coverage for as much of the community as possible.

The ability to self-claim telephone numbers, make outgoing calls via a gateway and to have an inbound hot-line are key features that enable this scenario. For these communities it is also possible for them to setup dedicated hot-lines per mesh user and have their cell phones divert to those numbers when not available on the regular cellular network, thus allowing a practically transparent bidirectional mesh telephony service using their existing telephone number.

While there is a cost in providing gateway devices and globally visible DID numbers, this may be offset in some situations by changing their cellular subscriptions from a premium operator with maximum network coverage to a lower-cost carrier that has a less extensive network, as they may no longer have need for the extended, but still inadequate coverage of the premium carrier. In Australia the cost differential for otherwise equivalent services from the most extensive cellular network versus the second most extensive can be 50% or more (see Table 1), providing evidence that such cost savings are possible.

The provision of mobile communications solutions does introduce ethical issues in terms of system reliability. The main factor being that the Serval BatPhone system offers "best effort" service rather than "carrier grade" service, and does not offer an emergency number service. However, the risks can be minimised through user education and by ensuring that people in such communities have cell phones which are capable of accessing the most extensive network, so that if they dial 112 or the local national cellular emergency number the phone will still use the cellular network. Tests have confirmed that dialling of these numbers, 110-119 internationally, and 000 in Australia, forces Android phones to dial via the cellular network, regardless of the presence of the BatPhone software, and indeed whether the cellular radio was previously turned off or not.

3.4 Nomadic Communities and Work-Forces

Nomadic communities and work-forces can be considered as the more extreme version of the rural and remote community use-case, where cellular back-haul is not possible. In such situations the same approach can be used as for rural and remote communities, allowing in-bound and out-bound dialling, however the internet back-haul will likely be satellite or some other dedicated solution, such as a fibre-optic link from underground to surface, rather than cellular. This flexibility in back-haul method allows for switching from satellite to cellular, and even to DSL or some other lower-cost means of connectivity when it is available, without requiring the end-users to manage complicated processes. All that is required is to have a gateway device of each type, and to turn on the cheapest gateway that is required to gain connectivity.

In a future release of the BatPhone software it is intended to allow the gateways to be price-tagged, so that all gateways can be operated simultaneously, but with only the cheapest path being used, thus simplifying operational deployment. The remaining user-configuration task is then to configure each gateway with the PSTN VoIP subscription required to make PSTN-terminated and receive PSTN-originated calls. While this task has been simplified as far as possible, it does still require the end-user to go to some effort to obtain such a VoIP subscription. It would make sense, purely from an ease-of-use perspective for the Serval Project to form partnerships with one or more VoIP providers and provide an integrated process to obtain and configure the subscription into a gateway device.

4 Preliminary Results & Discussion

To date the core telephony, outbound gateway and SMS functionality has been developed and tested in a variety of trials.

4.1 Remote Infrastructure-Free Surface Test

In July 2010 the core local mesh telephony function was tested without support from any infrastructure. This was performed by using several HTC Dream (also know as the Google G1) Android phones with an early version of the Serval BatPhone software installed. There was no cellular signal available where the trials were conducted at the Arkaroola Wilderness Sanctuary in the far North of South Australia. This trial demonstrated that the technology worked, and could deliver surprisingly good range in favourable conditions when there is no nearby source of 2.4GHz interference, good line-of-sight and dry air. Regretably, no precise range testing was performed on that occassion, however the engineer present at the test estimated that the effective range between phones was of the order of 500m, and likely >1km from ridge to ridge where the phones would have enjoyed clear Fresnel Zones.

The tests covered three simulated use-cases: (a) remote search and rescue, placing several BatPhones and Mesh Potatoes along a ridge to provide several square km of blanket coverage to surrounding country to make it possible to contact an imaginary lost bush-walker; (b) an isolated remote work-force or recreational group keeping in contact with one another without relying on additional hardware or infrastructure beyond their cell phones; and (3) (re-)establishing fixed-line and mobile telephony in a remote community in response to a simulated disaster. All three use-cases were simulated by a fly-in-fly-out team in less than eight hours, the vast majority of which was taken up by a media film crew and travelling by road and air to the remote locations.

This final use-use case was remarkable in that in a matter of 20 minutes we had not only provided an alternate land-line service to the Arkaroola Wilderness Sanctuary administration building, but provided the entire Arkaroola Village with mobile telephone service for the first time, and we were able to telephone the administration building from the open space around the village.

This process did reveal several faults with the early software, in particular inefficient implementation of aspects of the voice signature transfer protocol and the use of shell scripts in the Asterisk dial plan resulted in call connection times of up to 90 seconds. This was remedied by replacing he shell scripts with an Asterisk plugin that could perform the number resolution in as little as 250ms, and never more than three seconds, and by improving the voice signature transfer protocol. Call connection times for mesh-local calls are now practically instant, and much faster than placing a similar cellular call.

The operators of the sanctuary commented that even using UHF radio and repeaters that there are numerous communications black-spots that it would be advantageous to remedy using the technology presented in this paper.

4.2 Underground Testing

Radio propagation is known to be poorer in underground spaces than in unbounded open spaces for a variety of reasons, including multi-path fade where echoes of the signal dramatically raise the interference floor, and thus reduce the effective link-budget[11]. Thus is was desirable to gain some understanding on how a WiFi mesh would actually behave underground. Three tests and an above-ground control were performed.

The above-ground control consisted of establishing a call between two BatPhones on a long straight suburban road devoid of motor traffic and then having the caller walk along until the call began to break up, and then walk back until audio resumed. This would be considered the maximum effective range, and was measured at 165m using two HTC Dream phones.

The first underground test involved HTC Dream phones, one at the entrance to a cave, and the other carried into the cave as far as possible. The maximum straight-line distance available from the cave entrance was 42m, and the call did not break up. This was further than the ~30m that had been suggested by the literature[35]. It is possible that the rough, pocked surface of the lime-stone cave reduced the multi-path fade potential by absorbing a large proportion of the incident signal, however, this remains unconfirmed.

The second underground test involved HTC Dream phones, this time in a 6m wide disused railway tunnel in the Adelaide Hills. In this test the maximum range was found to be about 90m, again significantly further than theory predicted. This tunnel has a flat concrete floor and curved brick lined walls/roof, and thus absorbtion due to surface roughness was not likely. Rather surprisingly it was discovered that cellular coverage penetrated at least part of the tunnel, around the 160m mark, where our range tests were being performed. This indicates that the tunnel wall was certainly not entirely reflective of UHF or Microwave radiation, and possibly that the tunnel roof was covered with relatively little overburden at that point. None the less we were left with the realisation that WiFi propagation underground could be surprisingly good.

The third underground test involved Huawei IDEOS U8150 phones in an underground hydro-electric power station which was shut down for maintenance. The power station consists of four levels, with large plant present on most levels preventing clear line of sight to all areas, and no appreciable transmission through the heavy steel and concrete floors. The main floors were approximately 20m $\times$ 40m, and using only two handsets as way-points it was possible to place calls from the main control room to most locations on the various levels. It was not possible to determine how the system would perform were the power station operating, which would be expected to produce considerable interference. The operators of the plant commented that simply being able to contact and find people quickly and easily, even if only during shut-down maintenance, was of considerable operational advantage, aside from the worker amenity advantages.

4.3 Software Availability

The positive obligation to provide freedom of communication under the Declaration of Human Rights[5], as reinforced by recent reports from the United Nations[9] creates an imperative upon us to ensure that the technology we create is unalienably available to those who have greatest need. For this reason the software components described in this paper has been released under GNU General Public Licenses[36][37], and the source code can be downloaded from
`http://github.com/servalproject`
It is anticipated that an Alpha or pre-Alpha release will be made later in 2011.

5 Future Directions

There are numerous paths of refinement that are being pursued or contemplated, several of which have been described in the preceeding text. Principally this will

[35] Private correspondance with Dr. Simon Williams.
[36] `http://www.gnu.org/licenses/gpl-2.0.html`; On-line, Accessed 29 June 2011.
[37] `http://www.gnu.org/licenses/gpl.html`; On-line, Accessed 29 June 2011.

take the form of consolidating the features which have been demonstrated in the existing software to produce a commercial-grade release that can be released to the public. This will require the alleviation of the requirement for root access. This will be tackled on two fronts simultaneously. First, the mesh routing system will be adapted to operate as a UDP overlay network, thus obviating one of the two requirements for root access. Second, the WiFi radio management system is being improved to allow it to fall-back to cycling between virtual-access-point and managed mode to allow one to two hop voice calls and multi-hop asynchronous communications. This differs from WiFi Direct in that it can work on pre-WiFi-direct devices as well as allow one-to-many relationships that make better use of the broadcast nature of the medium.

Concurrent with the on-going technical development, closed-trials will also be conducted with partners in the humanitarian and industrial sectors. Further opportunities will continue to be explored, for example the application of our system to seek to fulfil the aspiration to use mesh communications in mining safety [10] and other under-ground plant operations, while simultaneously improving worker amenity above and below ground, and improving operational efficiency by enabling convenient, familiar and low-cost communications in these otherwise difficult environments. It is hoped that the commercial applications of the technology can be used to derive revenue that will enable the technology to be refined and distributed to the benefit of the poorest and most vulnerable people of the world.

6 Summary & Conclusions

The current prototype software has repeatedly demonstrated that convenient mobile mesh telephony is possible using off-the-shelf mobile tephone handsets as the sole piece of infrastructure. This has been complemented with testimony to the utility of the system from those who have witnessed or participated in tests. Thus we conclude, that the technology has great promise, even while much work remains to be done to mature the software and technologies to the point where they will be ready for appropriate use by the general public. This work is being actively pursued by the Serval Project team based at Flinders University. It is our firm belief that this technology will make considerable impact to the lives and livelihoods of many millions of people the world over.

7 Acknowledgements

The production of this paper has been supported by grants and other support from The Awesome Foundation For Arts & Sciences (Boston Chapter), Flinders University, Serval Project, Inc., The Arkaroola Wilderness Sanctuary, The Australian National Science Weeks Grants and NLnet Foundation. The work described in this paper would not have been completed without the considerable contributions, energy and dedication of the students and volunteers working on

the Serval Project at Flinders University and on behalf of Serval Project, Inc., the not-for-profit body dedicated to making the fruits of this research available to all.

References

[1] Michael Adeyeye and Paul Gardner-Stephen. The village telco project: A reliable and practical wireless mesh telephony infrastructure. *EURASIP Journal on Wireless Communications and Networking*, in press.

[2] D Johnson, N Ntlatlapa, and C Aichele. A simple pragmatic approach to mesh routing using batman. In *2nd IFIP International Symposium on Wireless Communications and Information Technology in Developing Countries*, Pretoria, South Africa, Oct 2008.

[3] E. H. Mckinney. Generalized birthday problem. *The American Mathematical Monthly*, 73(4):pp. 385–387, 1966.

[4] J.C. Mogul. Broadcasting Internet datagrams in the presence of subnets. RFC 922 (Standard), October 1984.

[5] United Nations. *Universal declaration of human rights*. U.S. Govt. Print. Off., Washington :, 1949.

[6] United Nations. Measuring the impacts of information and communication technology for development. *United Nations Conference on Trade and Development (UNCTAD) Current Studies on Science, Technology and Innovation, No3*, 2011.

[7] Ory Okolloh. Ushahidi or 'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory Learning and Action*, 59:65–70, 2009.

[8] Maria Cecilia Paule. "'sms' capital of the world": A political economy of a wireless revolution in the philippines. Master's thesis, Simon Fraser University, 2004.

[9] Frank La Rue Special Rapporteur on the promotion & protection of the right to freedom of opinion & expression. Promotion and protection of all human rights, civil, political, economic, social and cultural rights, including the right to development. *United Nations Human Rights Council, Seventeenth Session, Agenda Item 3*, 2011.

[10] Song Wen, Wang Fei, and Dai Jianbo. A emergency communication system based on wmn in underground mine. In *Proceedings of the Internation Conference on Computer Application and System Modeling (ICCASM)*, 2010.

[11] Jim Zyren and Al Petrick. Tutorial on basic link budget analysis. Technical Report Application Note AN9804.1, Intersil, 1998.