Efficient Mobile Mesh Networking:

Attractions, Myths and Techno-Economic Roadmap to Successful Commercial Innovation

S G Methley, M Crisp, J Newman
Plextek Ltd
Cambridge, UK
{sgm, mc, jn}@plextek.com

P A Ramsdale
STA
Cambridge, UK
peter.ramsdale@spectrade.co.uk

M Rio, S Bhatti
University College
London, UK
{m.rio@ee,s.bhatti@cs}.ucl.ac.uk

A Atefi
Ofcom
London, UK
ahmad.atefi@ofcom.org.uk

***Abstract*— This paper examines four scalability hypotheses of interest for mobile meshes via the following questions:**

'Do meshes self-generate capacity as new nodes join?'

'Are meshes more spectrally efficient?'

'Do directional antennas confer *significant* benefits for *hand-helds* below 3.5GHz?'

'No' is the answer because these hypotheses, whilst having a theoretical basis, can be shown to rely on inappropriate real world assumptions. However the following hypothesis is found to be true:

'Meshes may improve spectrum utilisation'.

Importantly however, there remain properties of meshes which make them uniquely attractive, such as coverage extension. However this raises a further question over the ability of a mobile mesh architecture to provide a guaranteed quality of service.

Finally, the wider aspects of commercial innovation are considered.

Keywords- mobile mesh, cellular multi-hopping, efficiency, quality of service, spectrum management

I. INTRODUCTION

The UK Office of Communications (Ofcom) recently commissioned a consortium of industry and academia to investigate the reality of mobile meshes in the bands below 3.5GHz [1]. Such an activity is termed 'sensemaking' by strategists, where the aim is to establish an initial position despite confusing evidence: Ofcom wished to examine the validity of the many competing mesh performance claims in the literature, since subsequent strategic and economic analysis could develop important policy conclusions from such technical claims.

Perhaps the largest initial attraction of mobile meshes is that they can be entirely unplanned in pure form. This is useful to the military and to disaster recovery teams who neither need infrastructure access for content nor want to rely on its presence for operation. It is far less clear what these benefits could lend to the roll-out of a mass market mesh network. On the other hand, to a service provider or regulator, the lure of a network which promises no planning phase must be high and thus must merit investigation.

It has often been said, as if it were a truism, that meshes increase capacity. The reasoning is usually along the lines of *each new user brings additional capacity to the mesh*, or *each new user effectively becomes a base station.* This paper critically examines such statements and aims to separate the reality from a 'something for nothing' type of mythology.

Nonetheless we find that, ultimately, meshes do retain some strongly attractive features, notably in the area of coverage, where they offer complementary performance to that of cellular systems. It is for this reason that meshes should find application in some scenarios, as part of a larger picture of mobile access technology.

In this paper we attempt to summarise several of the main points of a larger investigation [1]. The core approach begins via an examination of assumptions made by key papers in the literature – and establishing their relevance to mobile meshes under 3.5GHz. This focus is key to the paper's findings. Whilst we do not wish to overstate the case, the results are not all as might be expected from mesh 'folklore'.

Next, the real-world potential benefits of multiple hopping and antenna directionality are evaluated from practical first principles. Following these analyses we support meshing as a technique whose prime advantage is coverage extension. However we note that there are important caveats for service quality due to a performance which is dependent on the uncertainty in user movements. Finally we introduce five scenarios where we believe meshes can be innovative.

Remaining enablers and barriers to success are evaluated in the contexts of technology, human factors and regulatory framework.

II. HYPOTHESES - CAPACITY AND SCALABILITY

Ofcom wished to test the following widely proclaimed benefits of multiple hop mesh networks:

- capacity self-generation
- spectral efficiency
- omni-vs.-directional antenna benefits
- spectrum utilisation
- coverage extension

The first four are dealt with in this section, whilst coverage is dealt with in the following section.

A. Hypothesis Testing ... "that customers self-generate capacity"

There would be huge attractions to having 'self-generation of capacity' in a radio network. Notably, that the network is self-sustaining and that it could avoid the so-called 'tragedy of the commons' (the exhaustion of network resources due to over-use).

We believe misinterpretation of some published work may have led to several unfortunate myths concerning 'self-generation of capacity'. Four published approaches are reviewed below and, whilst each presents a coherent argument based on its stated assumptions, it will be shown that those assumptions do not always translate well to practical applications. The four approaches examined are*:*

Approach	Assumption Challenged
Grossglauser and Tse [2]	Unbounded delay
Gupta and Kumar [3]	Strict localisation of traffic
Shepard [4]	Unbounded spectrum
Negi and Rajeswaren [5]	Unbounded spectrum

1) Grossglauser and Tse [2]

This paper was taken as the basis for an economics and regulatory policy paper [6] which postulates many benefits if a 'tragedy of the commons' could be avoided by using mobility itself to increase the capacity of a network.

The model [2] specifically uses the mobility of nodes to act as intermediate 'couriers' of data between source and destination. Datagrams are passed from source nodes to near neighbours and delivery occurs when the courier nodes encounter the target recipients. Under this idealised model the per-node throughput remains constant, i.e. such a network is fully scalable in terms of capacity.

However, a clear consequence of this model is that the end-to-end packet delivery delay is related to the transit time of nodes moving throughout the area covered by the mesh. Statistically the mean delivery time is of order of 2d/v where d is the diameter of the mesh network and v the mean velocity of nodes within it. In a practical situation, the courier nodes may *never* encounter the recipient, in which case traffic is never delivered. The authors accept that this is clearly not acceptable for voice, or other real-time communications, and so direct the concept to non-critical store-and-forward messaging applications. It seems that this caveat may often be overlooked.

Although therefore limited in application in its basic form, we suggest the technique might be enhanced to reduce the transport delay and increase the probability of message delivery by nodes retaining a database of all other nodes they have had contact with and so selecting courier(s) on the basis of those that have had recent contact with the recipient.

2) Gupta and Kumar [3]

Their key conclusion is that capacity is shared amongst mesh nodes such that the upper bound for the average throughput $\lambda(n)$ obtainable by each node for a randomly chosen destination is of order of $W/\sqrt{(n.\log n)}$ bits/sec for the defined Random Network with the Physical Model. Thus the per-user throughput decreases with increasing node population.

Other authors, e.g. [8], have suggested other dependencies on the order of proportionality with n, but all models agree that average per-user throughput diminishes towards zero as the number of nodes increases, thus the mesh network does not scale indefinitely (and hence does not self-generate capacity).

It is interesting to consider what parameters, if any, might be changed to avoid this demise. Using a model [9] the dependencies on system parameters can be logically and simplistically stated as:
average throughput $\lambda(n)$ is proportional to functions of (γ, W, G/β, $1/L$, $1/r$, A, and $1/n$)

> where γ = propagation attenuation law, W = channel transmission rate, G = channel processing gain, β = required signal to noise ratio, L = mean end-to-end path length, r = mean per-hop link length, A = area covered by network, n = number of nodes

This implies that unless one or more of the parameters grows with n then per-user throughput will be asymptotic to zero:

- *W* cannot grow arbitrarily large because of thermal noise constraints and limits on transmission power.
- *G/β* depends on the properties of the communication system and increasing it generally makes it necessary to decrease *W*.
- Reducing hop length *r* (e.g. by constraining transmit power) increases spatial re-use but at the expense of increased hop-count and hence increased relay traffic. It transpires [3, 9] that the preference is to reduce *r* to increase spatial re-use. But there is a limit here in that if r is too small then the network can become

disconnected, i.e. minimum r is related to the inverse of node density (*A/n*).

- In random traffic flow models with uniform node density the mean end-to-end communication path length, *L*, is assumed to grow with coverage area *A* (*L* proportional to √*A*). This reduces capacity because of increased hop count. Thus, if one could conceive of services with more localised traffic (e.g. amongst localised communities) then *A/L* will increase more rapidly with increasing *A*. This will help to improve scalability.
- The remaining parameter that might scale with n is the area *A*. [9] suggests that three factors are required to achieve a non-zero throughput with increasing n: (i) the attenuation law γ needs to be greater than 3, (ii) the hop count *H* needs to be independent of n, (iii) area, *A*, needs to increase with n (i.e. the node *density* needs to be nearly constant or reducing with increasing *A*). However, (iii) requires that as the subscriber base increases those subscribers spread themselves out more thinly. It is not easy to see on what basis this might happen in any practical deployment.
- The propagation attenuation law γ strongly influences the above conclusions. A higher attenuation factor γ will permit higher throughput capacity [3, 9].

From the above list of options, one can see that there appears to be very little prospect of avoiding the asymptotic reduction in per-user throughput with increasing subscriber base. The analysis of [3] and others assumes a random association between source and destination nodes. Thus path lengths range from nearest neighbour (one-hop) to the full diameter of the area covered (many hops), and so, as the network size increases geographically and/or in terms of node-density, the number of hops per path must increase. This is one of the primary factors which cause the reduction in capacity with increasing number of nodes.

It is clear, then, that if traffic flows were more localised amongst neighbouring nodes, regardless of the geographic size of the network, then the number of hops per path would not increase *pro rata* with size and so the network would scale better, but we wonder how such a situation could be guaranteed in a real world deployment.

3) Shepard [4]

This paper has a relatively ‘out-of-the-box’ approach in suggesting a mesh in which collisions are not fatal for the MAC. It sees multiple concurrent transmissions as a signal-to-noise issue, rather than a requirement to back off and try again. It does this by using spread spectrum transmission, hence multiple transmissions simply raise the noise floor, as in any CDMA system. A complete theory is proposed to enable meshes to scale to millions of nodes. The problem is that it is extremely spectrally inefficient, due to the large processing gain required and in any case the predicted throughput of a large mesh is still only in the several kb/s range.

4) Negi and Rajeswaren [5]

A broadly similar approach to Shepard [4] with some similar problems is that of using “infinite” spectral bandwidth, for example in the ultra wide bandwidth (UWB) sense.

B. Hypothesis Testing ... “that mobile meshes are more spectrally efficient”

One of the traditionally used scenarios for suggesting that mesh operation into an Access Point might be more spectrally efficient than a regular PMP cell is the concept that increased throughput can be achieved over a series of short hops rather than one long hop. We shall demonstrate that this is only true for an idealised single-path scenario, and is diminished by the dissimilar antenna gains of Access Points and mobiles.

For the case of hopping between nodes of like type: If two hops of roughly equal length replace a single hop as shown in Figure 1 then:

- only half the time-bandwidth product of spectral resource is available for each hop, and this acts to reduce the delivered data rate by a factor of 2
- but as each hop is half the length of the original link, the link budget is improved. This improvement can be used to improve spectral efficiency either by increasing the transmission rate on each hop or reducing the transmit power. For example, in a third-law propagation environment the link budget is improved by x8 (~9dB); this would permit a four-fold increase in transmission rate by changing from QPSK to QAM64. Alternatively, with spread-spectrum the coding gain could be reduced to realise a similar increase in transmission rate.

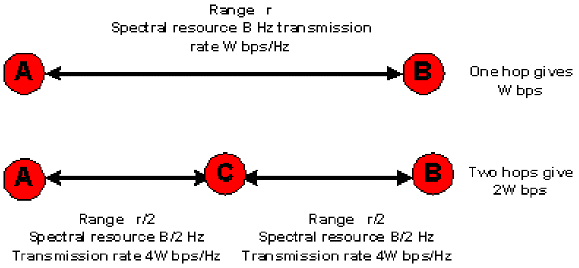

Figure 1 Two-hop vs. one-hop rate improvement between mesh nodes

This example implies that twice as much data can be transferred using two shorter hops, i.e. spectral efficiency is doubled. But this only prevails when the path length is exactly halved. If instead there is asymmetry in the two-hop path lengths then the link-budget gain in the longer hop will diminish and so the higher rate becomes unsupportable. This “sweet spot” in the path length split is illustrated in the graph of link budgets in Figure 2.

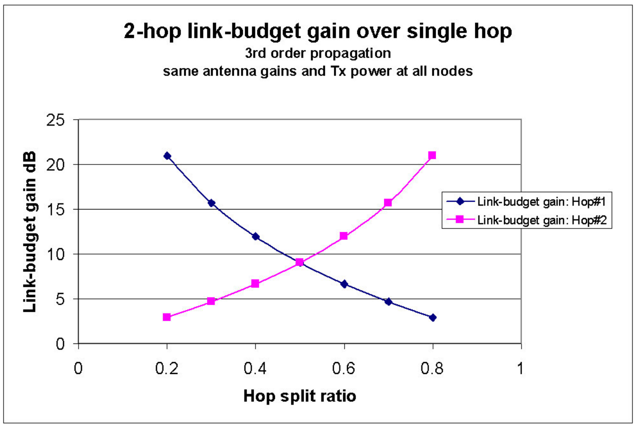

Figure 2: Two-hop link budget gain over single hop

But the comparative performance is further eroded for the case of multi-hopping into a mesh Access Point or cellular base station as represented in Figure 3.

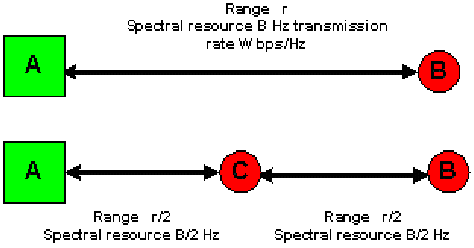

Figure 3: Two-hop vs. one-hop into high gain Access Point

The hop(s) between mobiles lack the higher antenna gain and height of the link into the Access Point (item **A** in Figure 3). Due to this imbalance the "sweet spot" no longer occurs at the 50:50 path-length split. The graph of Figure 4 illustrates this for the case when the Access Point antenna gain is just 13dB above the mobile nodes' gain – the "sweet spot" has moved to approximately 75:25 path length ratio and the optimal link budgets on the two hops are only about 4dB above the single-hop case. With this small link-budget gain the transmission rate might be little more than doubled. Thus the best case throughput rate of this two-hop route is roughly the same as the single-hop route.

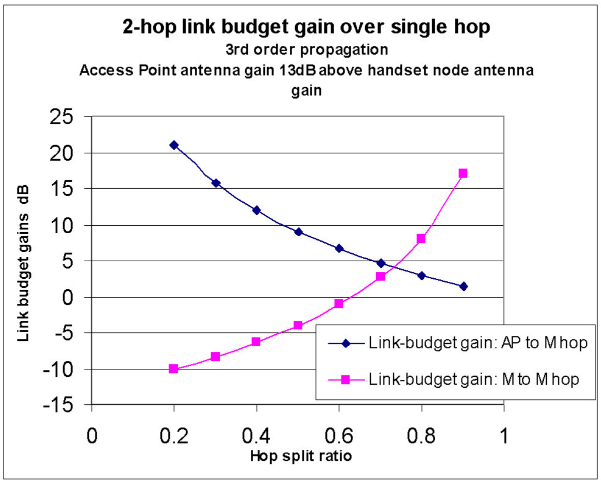

Figure 4: Two-hop vs. one-hop link budgets with high antenna gain

A further implicit assumption in the above simplified analysis is that the multi-hop path length is the same as the single hop length. In practice this may not be the case; nodes will be unevenly distributed and routes may circumvent building and terrain clutter. The detrimental effect of increased route length is illustrated in the graph of Figure 5 which illustrates the reduction in link budget gain at the "sweet spot" of Figure 4 as the route length is increased

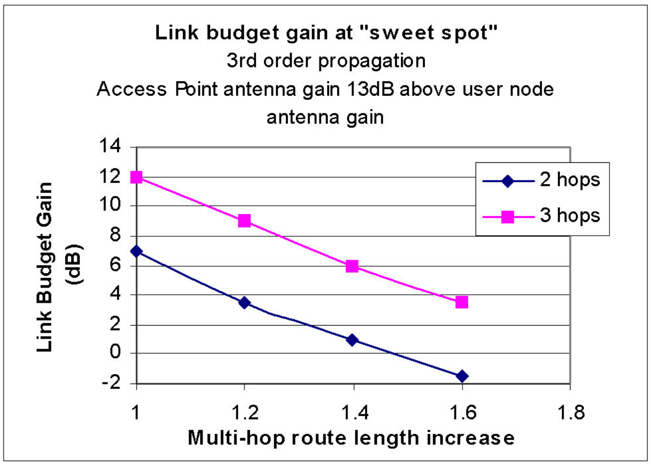

Figure 5 "Sweet Spot" link budget gain vs. extension in total route length

C. Hypothesis Testing ... "that directional antennas confer significant benefits for mobile mesh networks below 3.5GHz"

A starting point in the analysis is to consider an idealised antenna having negligible side lobe responses. This can be represented by the "flat top" model – where the antenna beam in the azimuth (horizontal) plane is represented as an arc of a circle subtending an angle equal to the 3dB beam width of a polar response. This leads to a simplistic interfering / non-interfering alignment of beams as illustrated in Figure 6:

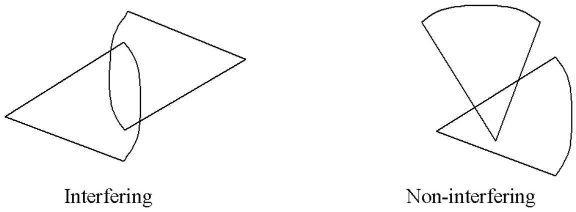

Figure 6: Interference model for directional antennas

For a network of randomly deployed nodes equipped with such antennas, the theoretical upper limit on the improvement of throughput capacity is as large as $4\pi2/\alpha\beta$ [10] (where α and β are the beam widths of the transmit and receive antennas respectively). However, for any practical antenna, and more so for mobile/hand-held products in the bands of interest here (0.5-3.5 GHz), there will be a finite side lobe response which will seriously erode the gains anticipated.

The key manifestation of this finite side lobe response in the network is to extend the interference boundary around nodes

[10]. The physical extent of this boundary is governed also by the attenuation factor of the propagation environment. If an antenna has a mean side lobe level which is κdB below the main beam then, in a propagation environment with attenuation rate γ (i.e. path loss proportional to $(\text{range})^{\gamma}$), the differential coverage range, Δ_R, between main beam and side lobe is given by [10]:

$$\kappa = 10.\,\gamma.\log(1/\Delta_R) \qquad (1)$$

It is postulated, from practical work at Plextek Ltd and data from the antenna-supply industry, that for mobile/hand-held products operating below approximately 6GHz the side lobe response is unlikely to be more than about 10dB-15dB below the main beam. So, taking a likely figure for side lobe level of κ=13dB, in a fourth-law propagation environment Δ_R is only 0.5. Thus, the interference boundary for the side lobes is only half that within the main beam.

Considering the case of 90° beam widths with -13dB side lobes this implies a capacity gain in the region of x3.3, compared to a theoretical maximum gain of x16 for the zero-side lobes case. This illustrates the detrimental effect of finite side lobe levels.

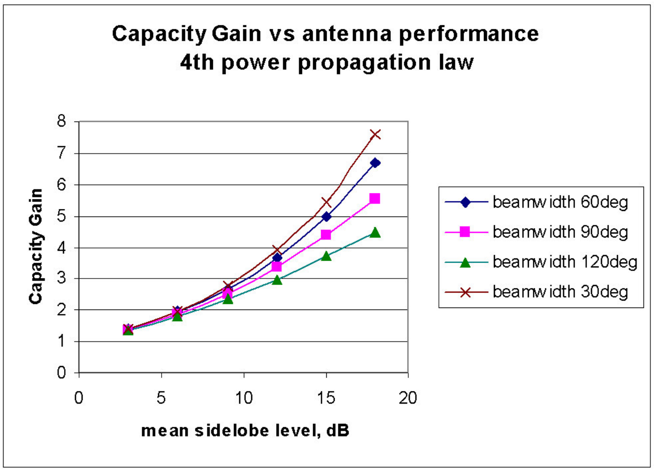

Figure 7: Capacity gain vs. antenna performance

Figure 7 illustrates that the capacity gain factor is a more sensitive function of side lobe level than it is of beam width. Furthermore, as beam width is reduced the side lobe level dominates performance, thus indicating that there is little benefit in decreasing beam width without equal attention to reducing side lobe levels, which returns us to the practical barriers first stated.

This capacity gain performance is also a function of the propagation environment. The range difference, Δ_R, between main beam and side lobes decreases with increasing propagation attenuation law and so the benefit of side lobe attenuation diminishes. This follows from the premise that, for a given density of nodes, the ratio of the number of nodes residing inside the main beam coverage area to the number residing in the side lobe coverage area diminishes with increased propagation law. This effect is illustrated in Figure 8 for a beam width of 90°. From this it can be seen that the benefit from antenna directionality decreases with increasing propagation attenuation factor.

Note that the capacity gain curves in Figure 8 are normalised to the omni-directional antenna case for each respective propagation law. Thus the curves do not imply that the network has the same capacity for an omni-directional antenna, independent of propagation law; there are scale-factors to be applied to the vertical axis for each curve. In fact the corollary to this is that the high attenuation environments enable greater spatial reuse and hence higher spectral efficiency than a low attenuation environment [3, 9]. However, the low attenuation environment will reap more benefit from the use of directional antennas, because of reduced interference over the longer propagation ranges.

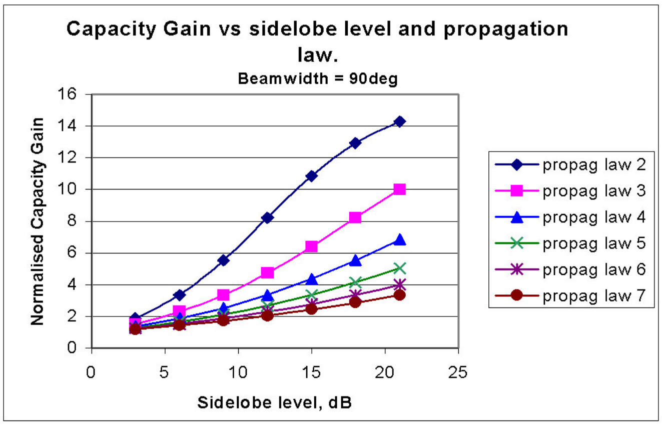

Figure 8 Capacity gain vs. propagation attenuation factor

As a final aside, another way of visualising the potential benefits of directional antennas is to consider that a mesh may be made physically or logically. Physical meshes are those made by physical level constraints, for example by directional antennas or perhaps by constraint of the signal path by terrain or medium. The wired internet is clearly a perfect physical mesh in that transmitting on one link does not interfere with any others. On the other hand a logical mesh is configured above the physical layer. There is not necessarily any physical constraint to a station's neighbours imposed by the system. Omni-directional antennas in an open field could be connected as a logical mesh, although their interference footprint would clearly be quite different from that of a physical mesh and the full benefits of physical meshing should not be expected. (The unrealised benefit would be that the omni-directional antennas could equally be re-configured as e.g. point to multi-point, in the logical sense, if that were ever required.) The relevance of this discussion to the preceding argument is that it is how closely a true physical mesh may be approached which is important, be that via the degree of antenna directionality available, or otherwise.

D. Hypothesis Testing ... "that meshes could improve spectrum utilisation"

This hypothesis relates to the wider issue of spectrum utilisation, rather than simple specific spectrum efficiency. It suggests that the spectrum may be better utilised by having short line-of-sight (LoS) mesh links use 'less precious' spectrum e.g. up to 6 GHz.

In [1] three key factors point to mobile mesh networks offering opportunities for use of higher frequency bands:

i. They are not necessarily more spectrally efficient than current cellular systems operating in the 2GHz region (cf. I.B). Thus they might usefully be allocated less commercially precious spectrum.
ii. To achieve useful per-user throughput, the relaying capacity of mesh nodes needs to be high (a corollary of II.A). Thus meshes need access to large allocations of bandwidth.
iii. The potential of increased end-to-end throughput by using multi-hop vs. single-hop is best realised when there is a high propagation path loss at the chosen frequency of use (cf. I.B, II).

III. MESH COVERAGE

One clear claim for mesh networking is that it helps coverage by allowing multiple hops around obstacles. This is of most use in a cluttered environment, where a powerful base station would otherwise have to be used to combat the large variance in path loss. Figure 9 [11] compares meshing around obstacles via a 'forwarder' to gain an advantage in throughput.

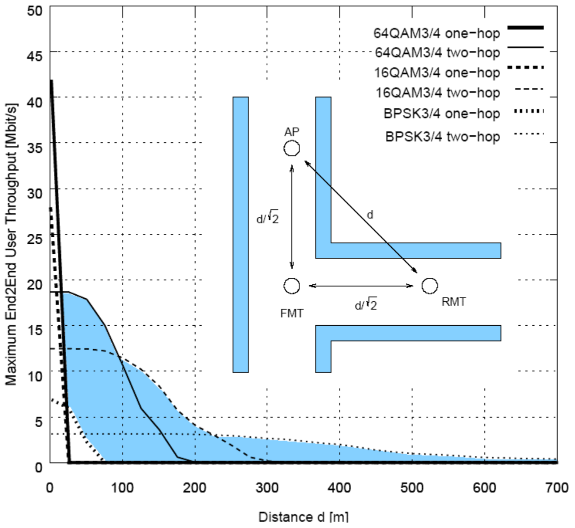

Figure 9 Use of a 'forwarder' to 'skip' around obstacles as in a mesh network

Figure 9 presents a specific argument for hopping around obstacles, rather than multi-hopping *per se*, and may be understood as follows: the situation being modelled is as drawn in the figure inset. The objective is to communicate from the AP (access point) to the RMT (remote mobile terminal), which is sited around a corner. This may be done in one of two ways:

- as a single hop, distance d, 'directly'
- as a two-hop, twice distance d/√2, 'hopping' around the corner via the FMT (forwarding mobile terminal)

Additionally, the effect of adapting the modulation scheme is also shown.

The potential gains to be had in all the two-hop cases are shown within the total blue shaded area of the graph which is made up of the gains from each two-hop case listed in the graph legend. The white area, enclosed by the axes and the shaded area, is the single hop performance, which is very limited as might be expected. Although the results of Figure 9 relate to extending the infrastructure mode of HiperLAN/2 in the 5GHz band, the implications of this simulation are of a general nature and thus may be applied to other hopping situations, such as a mesh network.

Clearly, in the real world, propagation scenarios may be more or less extreme than Figure 9. The best system solution thus depends on the target environment. In general however, meshes should show benefits where the problem faced is one of coverage due to a cluttered or shadowed environment. Most often this is associated with a dense urban area. Potentially, therefore, meshes could be also deployed as 'hot zones' in a sea of cellular, due to their complementary coverage characteristic. We would define this as an 'access mesh', but this might also be included under the collective of 'cellular with multi-hopping'.

Hopping around corners or obstacles is perhaps the greatest benefit offered by a mesh.

IV. COULD A MESH GUARANTEE A QUALITY OF SERVICE?

Accepting the limitations of II.A to I.D, we have stated that we still believe that meshes have much to offer from a coverage point of view. Hence the next question to ask is whether quality of service is possible within a mesh, despite user movement. Poor quality of service would be at odds with users' expectations, as they will require guaranteed service for inelastic applications like video or other applications intolerant of delay-variation or throughput-variation.

Royer et al [12] noted the result of a well known publication which stated that a fixed node with the 'magic number' of six fixed neighbours could be shown to possess the optimum trade-off between transmission power and self-interference for best throughput [13]. The question was whether a similar optimum carried across to mobile networks.

The answer was 'no'; there is no single optimum power-efficiency trade off for a range of node mobilities, although there is an optimum for each value of node mobility. Essentially, as node mobility increases it was found that the node density, in terms of connectivity, needed to rise. Otherwise, the mesh could eventually break apart or 'partition' [14]. More mobile nodes need more guaranteed neighbours. The clear issue for the mesh network planner is that network connectivity, hence quality of service, depends on parameters outside his control – the users.

More recently Nilsson [15] showed that connectivity also depends on the traffic level within a multi-hop mesh. His conclusion was that both predicted traffic levels and node mobility/densities would need to be known at the planning stage, for viable network design.

Lungaro and Wallin [16] expanded on Lungaro [17], where notably they added a caveat to his earlier work on mesh-like cellular multi-hopping. They noted that the principal drawback of the scheme was that, due to the 'partially uncontrolled infrastructure' (i.e. the user nodes), multi-hopping was 'not able to guarantee a quality of service'. They proposed that a simple solution was the addition of nodes, as relays, by the network operator. In fact they proposed a system with three node types: access node, relay node and user node. Their conclusion goes beyond Nilsson's. Knowing traffic and user mobility is still not enough to guarantee quality of service within a mesh of user nodes - infrastructure must be added.

We note that the addition of some infrastructure is necessary in any case where access beyond the mesh is required, e.g. to the Internet. We expect such access meshes will dominate completely over isolated pure meshes for public deployments. The fixed part of an access network could be arranged to provide (i) seeding of a new deployment (to ensure connectivity whilst user number ramps up), (ii) reduction of network dependence on user mobility and traffic levels, as well as (iii) access to the wider Internet.

V. INNOVATION

The previous sections of this paper have addressed and commented on technology aspects of mesh networks and identified the properties which make them unsuitable for certain applications yet suitable for others. Achievement of the best overall result will also require consideration of regulation, creation of standards, development of policies concerning harmonisation effects, etc., each from the point of view of how they might encourage or discourage the complex innovation process, Figure 10.

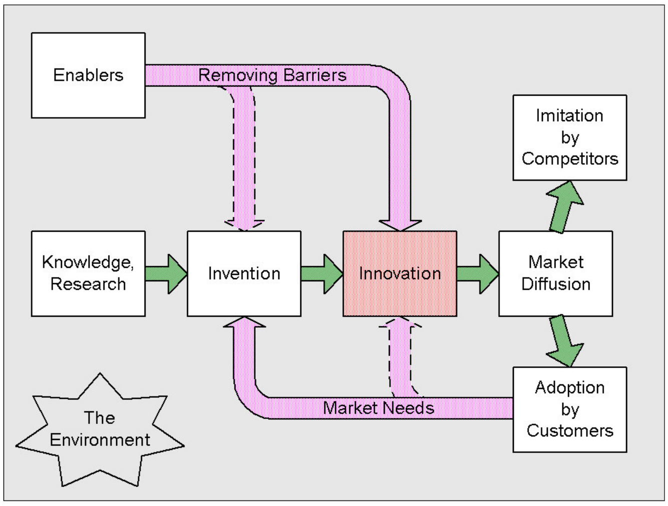

Figure 10 Innovation as part of a larger process

Innovation is not simply the introduction of new techniques and technologies; it is also about doing things in a new way. To encourage innovation in any field it is necessary to understand all the enablers and barriers in order to see what becomes possible by forcing change. Figure 10 shows that necessity is not only the 'mother of invention', but also of innovation. In other words the market drives the need.

The discussion is structured into three categories:

- a definition of the desired end result of innovation
- consideration of enablers for innovation
- consideration of barriers to innovation.

A. *Innovation - What, Where and How?*

It is important to be clear on the goal and the general strategy to achieve it. This paper takes the goal as achieving the deployment of access mesh networks in particular, i.e. those with access points to external content. Moreover the likely scenarios are taken to be those identified in [1]. These are:

- cellular with multi-hopping
- WLAN hotspot extension
- community networking, e.g. [22]
- home and office indoor networking
- zero-infrastructure environments

The latter two scenarios may not be familiar: for home and office networking we observe that a closed user group, peer to peer network within a cluttered environment is the requirement and this plays to the strengths of a mesh. By zero-infrastructure environments, we mean those places where infrastructure is yet to be installed or where a permanent infrastructure is not needed, e.g. a new technology park or a temporary conference venue.

The strategic 'how' factor is addressed via an evaluation of the enablers and barriers, within the specific context of our stated goal. We believe that the very creation of the likely scenario list above already shows how mesh networks can be strongly innovative in removing barriers. What follows identifies those barriers which may remain and the further enablers which need consideration.

B. *Enablers for Innovation*

Enablers are discussed under the following groupings:

- technology
- human factor aspects
- technology understanding.

1) *Technology*

In Figure 10, these include the remaining inventions necessary to feed into the innovation process. In some cases the inventions are very difficult technically and hence a solution is sought but not yet available. In other cases the need for the invention may not yet be clear. The physical layer is fairly well served by existing technology, although there is substantial scope for inclusion of the emerging cognitive radio technology. But substantial further work is required in the fields of MAC and routing protocols to improve the efficiency of carrying delay sensitive traffic. Alongside these are requirements for more advanced system modelling tools for mobility and traffic flows.

2) *Human factor aspects*

Some user behaviour and expectations will affect the successful uptake of mesh networking. A mesh specific problem is the power consumption of a battery powered device. Because mesh nodes are required to relay traffic which may not be their own, especially when close to an access point, more demands are made on battery performance. A situation may occur where a user who is not generating his own traffic may nonetheless experience high battery drain due to relaying the traffic of others. Further problems could arise depending on the user's response to such a situation. If the user turns off his node when not using the device, then the mesh availability may be compromised for many other users. Clearly some provision needs to be made to encourage good group behaviour from individual users. Reward schemes have been suggested [18], but may carry a large traffic overhead.

General improvements in radio technology will continue to benefit mesh networks. In particular, techniques which increase battery life, such as lower voltage rails, higher efficiency PAs, and more efficient modulation methods (energy/bit), will particularly benefit mobile meshes.

3) *Technology understanding*

Both users and operators need to understand the technology in terms of what it could do for them. Currently, much of this thinking may be dominated by individual perception.

In the UK, the major cellular operators use GSM and UMTS for their 2nd and 3rd generation mobile technology. Multi-mode handsets which add WLAN access to the service offering are becoming available. In the UK, one operator already offers a separate, extensive WLAN hotspot service, and the business aspects of the differing services are beginning to be understood.

However there remain concerns about combining services which offer very different performance guarantees. Some operators believe that they have established a particular service expectation for their 'brand'. So, for a cellular operator, this means that there are levels of service in terms of coverage, availability, dropped calls (including calls with hand-overs), consistency, etc. The list of parameters continues to grow with data and IP based calls, such that delay and throughput, for example, also need to be at or above the 'brand' level on a regular basis.

On the other hand the WLAN hotspot service provider offers a service, some aspects of which attract lower expectations, but yet is still acceptable to their users. This arises from experience and tolerance of the Internet as a 'best efforts' service and the assumed coverage area of a WLAN radio hotspot being conservative.

With these mindsets, how likely is it that mesh networks will be considered as a suitable technology for delivering part of a service? Clearly, if mobile mesh techniques are to form part of a cellular network, the operator must ensure that customers' expectations are not disappointed by variability or lower than expected performance levels (section III). It is sometimes considered better not to offer coverage rather than to offer a poor service and, especially when rolling out a network, sites are only introduced when they can offer contiguous coverage.

An enabler for innovation would be to first evaluate these 'soft' aspects of rolling out a mesh based service.

C. *Removing barriers to innovation*

1) *Spectrum management model*

Ofcom's policy on spectrum management is based on the Cave Report [19] as developed further in the Spectrum Framework Review [20, 21]. Ofcom has changed its approach from command-and-control to the use of market mechanisms plus some licence exempt use. With the market mechanisms approach, Ofcom is still responsible for partitioning the spectrum and making it available - an auction being the preferred route. Ownership and use is then driven by market forces. Secondary trading is expected and spectrum should be acquired by those who value it most highly. Over time spectrum blocks may be sub-divided or aggregated.

This encompasses trading and liberalisation activities and raises issues of radio interference. Owners of spectrum need to know what their spectrum usage rights are, as this impinges on the quality of service which they can provide. The issues are:

- spectrum quality
- levels of real interference
- levels of perceived interference
- protection / adjacent channel ratios
- the effects on interference of new RF structure & usage levels.

The above leads to the need for carefully defined spectrum usage rights: The existing usage of an allocation is referred to as specific spectrum usage rights. Further rights, called restrictive spectrum usage rights, are permitted provided that neighbouring users do not suffer increased interference, although initially these rights may be too restrictive to allow efficient usage. Hence it is proposed that new specific spectrum usage rights are developed together by agreement with neighbours, potentially involving compensation elements.

Ofcom's new policy is designed to speed innovation.

2) *Legacy licence effects*

When used in a mesh network, nodes are effectively acting as both infrastructure (relays) and customer equipment (access).

In a TDD system there is no physical difference in the transmissions from the node, regardless of whether it is at the end or within part of a chain of radio transmissions. However, such systems may fall victim to older legislation concerning the physical elevation of infrastructure transmitters. In short it may transpire that different rules could apply to the same radio, depending on whether it is said to be functioning as a user or backhaul node.

For an FDD system there is a further issue in that the operator has a licence to transmit on the down-link frequencies and the customer equipment is approved for transmissions on the up-link frequency. When acting as a relay, customer equipment then requires some licence extension or some means by which it is considered to be being operated by someone as the agent and under the control of the operator, fully in line with the provisions of the appropriate legislation.

These examples are from the UK, but similar circumstances may prevail elsewhere.

D. Summary

To precis, we note specifically that Ofcom are implementing a move away from the old, slow command-and-control method of spectrum management towards market mechanisms. This should be expected to enable faster innovation.

We note also that the softer issues of user perception should not be ignored when planning a service and that radio deployment approval processes need to be updated to cope with the capabilities of mesh.

VI. CONCLUSIONS

The main contribution of this paper is the rationalisation and clarification of many competing mesh performance claims within the literature for the specific case of mobile meshes below 3.5GHz. This is important since such technical claims can form the basis of economic and policy planning. We both examined existing work and used analysis from practical first principles.

From the hypotheses tested we conclude that mesh subscribers cannot self-generate capacity at a rate sufficient to maintain a target level of per-user throughput, regardless of network size and population. We further conclude that there remain fundamental trade-offs between throughput, capacity and delay, which cannot be dismissed easily. Thirdly, whilst network capacity can be improved through the use of directional antennas, for handheld devices the extent of directionality is limited since the high side lobes levels associated with such small antennas severely limit the improvements in spatial re-use that would otherwise be possible. Finally we note that spectrum utilisation could be improved by operating meshes within higher, less precious spectrum, e.g. up to 6GHz.

Noting that coverage extension is a likely application for mesh, we note that the quality of service within a mesh is user dependent, in terms of both mobility and traffic level. Unless specific additional steps are taken to mitigate this e.g. via additional fixed nodes, an operator will be unable to provide service level agreements. Such mitigation will need to be tuned to the actual mobility and traffic circumstances in each case, if they are known. The provision of appropriate additional infrastructure could also be arranged to provide scalability, so forming a viable access mesh.

Innovation could be encouraged by attending to technology enablers, which consist of many protocol and modelling aspects. The softer issues of user perception and behaviour become relevant for some unique mesh system properties which are not present under regular cellular systems. Barriers to innovation which can be removed are the tight, slow command-and-control model of spectrum management in favour of market mechanisms. Unforeseen barriers may also exist within radio legislation which predates mesh networking proposals.

In summary, we still believe that whilst mobile meshes do not live up to all the claims in the literature, they can be very beneficial in the area of coverage extension. In fact, hopping around corners or obstacles is perhaps the greatest benefit offered by a mesh. We suggest that meshes are equally applicable to extending the coverage of WLAN hotspots as they are to cellular multi-hopping, within certain limitations, and should be seen as integral to any 4G or ‘beyond 3G’ vision. They should also find application in home and office indoor networking and community networks.

ACKNOWLEDGMENTS

The authors would like to thank William Webb of Ofcom for his encouragement and colleagues at Plextek (Graham Maile, Bill Garland, Peter Massam, David Cox), Ofcom (Stephen Jones) and UCL (Lionel Sacks, John Mitchell) for their inputs to the project.

REFERENCES

[1] Methley S G, Crisp M, Newman J, Rio M, Bhatti S, Ramsdale P, *“Study of Efficient Mesh Networks”*, Ofcom Contract Ref. C31400/001, Ofcom, to be published.

[2] Grossglauser M, Tse D, “Mobility increases the capacity of ad hoc networks” *Proceedings of IEEE INFOCOM '01*, Anchorage, Alaska, April 2001.

[3] Gupta P and Kumar P R, "The capacity of wireless networks", *IEEE Trans. Information Theory*, vol 46, pp. 388-404, Mar. 2000.

[4] Shepard T, A Channel Access Scheme for Large Dense Packet Radio Networks”, *Proceedings of ACM SIGCOMM'96*, August 1996, Stanford University, California.

[5] Negi R, Rajeswaren A, "Capacity of power constrained ad-hoc networks," *Proc. IEEE INFOCOM*, Hong Kong, China, Mar 2004.

[6] Jones S, Levine P and Rickman N, *“The Economics of Radio Spectrum”*, Ofcom report, 2003. Also *Fifth Wireless World Conference*, 15-16 July 2004, University of Surrey, Guildford, UK.

[7] Sharma G, Mazumdar R R, “On Achievable Delay/Capacity Trade-offs in Mobile Ad Hoc Networks,” *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, Cambridge, UK, March 2004.

[8] "*High Performance Mesh for Dense Metro WLAN*", MeshDynamics white paper available at <http://www.meshdynamics.com/wPresentations.html>

[9] Arpacioglu O, Zygmunt J.H : "On the Scalability and Capacity of Wireless Networks with Omnidirectional Antennas". *Third International Symposium on Information Processing in Sensor Networks* (IPSN 2004), Berkeley, California (April 27–28, 2004)

[10] Yi S, Pei Y, Kalyanaraman S, "On the capacity improvement of ad hoc wireless networks using directional antennas", *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, Annapolis, Maryland, USA, June 01-03, 2003.

[11] Esseling N, Weiss E, Kramling A, Zirwas W, "A Multi Hop Concept for HiperLAN/2: Capacity and Interference", Proc. European Wireless 2002, Vol. 1, p.p. 1-7, Florence, Italy, 02/2002.

[12] Royer E M, Melliar-Smith P M, Moser L E, "An analysis of the optimum node density for ad hoc mobile networks", *Proc. IEEE Comms*, June 2001.

[13] Kleinrock L, Silvester J, "Optimum transmission radii for packet radio networks or why six is a magic number", *Proc. IEEE Nat. Telecoms Conf.*, pp4.3.1-5, Birmingham, Alabama, Dec. 1978.

[14] Hsieh H-Y and Sivakumar R, "Performance Comparison of Cellular and Multi-hop Wireless Networks: A Quantitative Study", *Proceedings of ACM SIGMETRICS*, June 2001.

[15] Nilsson A, "Performance analysis of traffic load and node density in ad hoc networks", *Proc. European Wireless (EW2004)*, Barcelona, Spain, February 24-27, 2004.

[16] Lungaro P, and Wallin E, "Coverage, Capacity and QoS Tradeoffs in Hybrid Multi-Hop Ad Hoc Cellular Access Systems", *Proceedings of the Affordable Wireless Services and Infrastructure, 1st Annual Workshop*, Tammsvik, Bro, Sweden, 3-4 June 2003.

[17] Lungaro P, "Coverage and capacity in hybrid multihop ad hoc cellular", Master of Science Thesis, TRITA-S3-RST-0315, Radio Dept. KTH, Stockholm, 2003.

[18] Salem N, Buttyan L, Hubaux J, Jakobsson, M. "A Charging and Rewarding Scheme for Packet Forwarding" ACM *MobiHoc*, Annapolis, Maryland, USA, 2003.

[19] Cave M, "Review of Radio Spectrum Management - An independent review for Department of Trade and Industry and HM Treasury", 6 March 2002. Available via http://www.ofcom.org.uk.

[20] Spectrum Framework Review, Ofcom, 23 November 2004. Available via http://www.ofcom.org.uk.

[21] Spectrum Framework Review Statement, Ofcom, 28 June 2005. Available via http://www.ofcom.org.uk.

[22] Lad M, Bhatti S, Kirstein P, *"Challenges: Opportunities For Coalition-Based Community Networking"*, submitted to the 11th Annual International Conference on Mobile Computing and Networking. Paper: 1568957932, 2005