Efficient Mobile Mesh Networking: Testing Scalability Hypotheses

S.G.Methley*, M. Crisp*, J. Newman*, P.Ramsdale§, M. Rio¶, S.Bhatti¶, A.Atefi †

*Plextek Ltd, UK, {sgm,mc,jn}@plextek.com, §STA, UK, peter.ramsdale@spectrade.co.uk,
¶UCL, UK, m.rio@ee.ucl.ac.uk, s.bhatti@cs.ucl.ac.uk, †Ofcom, UK, ahmad.atefi@ofcom.org.uk

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

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:

'May meshes improve spectrum utilisation?'

1 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.

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 first principles. Finally we note that 'less precious' spectrum e.g. up to 6GHz could usefully be utilised by mesh systems.

2 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

2.1 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 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

<u>Grossglauser and Tse [2]</u>

This paper was taken as the starting point for an economics paper [6] which postulates many benefits if a 'tragedy of the commons' could thereby be avoided.

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 missed.

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.

<u>Gupta and Kumar [3]</u>

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 $c2W/\sqrt{(n.\log n)}$ 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 $\sqrt{A}$). 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.

<u>Shepard [4]</u>

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.

Negi and Rajeswaren [5]
A broadly similar approach with some similar problems is that of using "infinite" spectral bandwidth, for example in the ultra wide bandwidth (UWB) sense.

2.2 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 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.

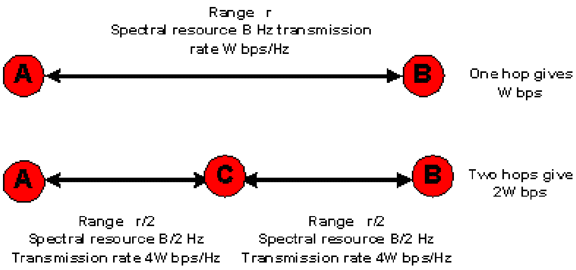

Figure 1 Two-hop vs. one-hop rate improvement between mesh nodes

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 budget in Figure 2.

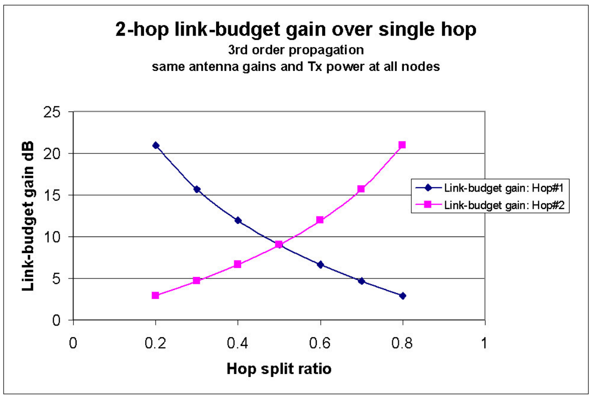

Figure 2: Two-hop link budget gain over single hop

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.

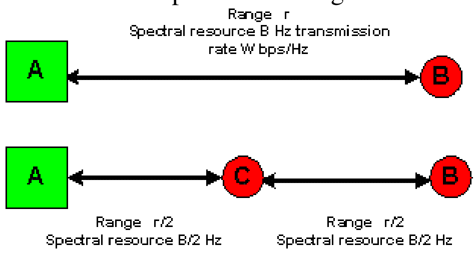

Figure 3: Two-hop vs. one-hop into high gain Access Point

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.

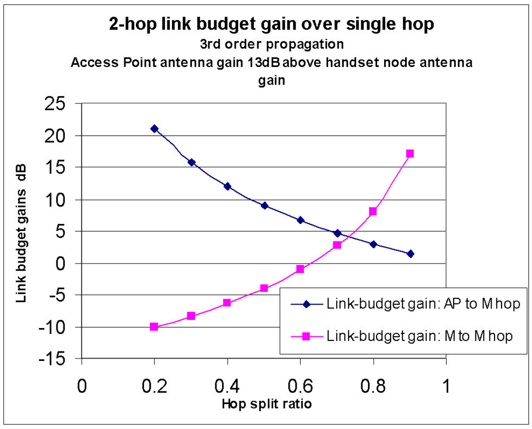

Figure 4: Two-hop vs. one-hop link budgets with high antenna gain

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

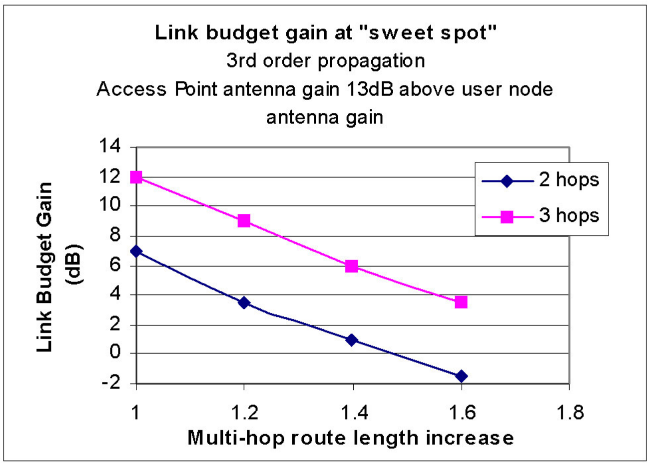

Figure 5 "Sweet Spot" link budget gain vs. extension in total route length

2.3 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:

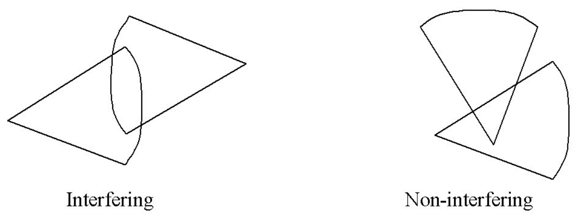

Figure 6: Interference Model for Directional Antennas

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 $4\pi 2/\alpha\beta$ [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:

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

It is postulated, from practical work at Plextek 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 gain of x16 for the zero-side lobes case. This illustrates the detrimental effect of finite side lobe levels.

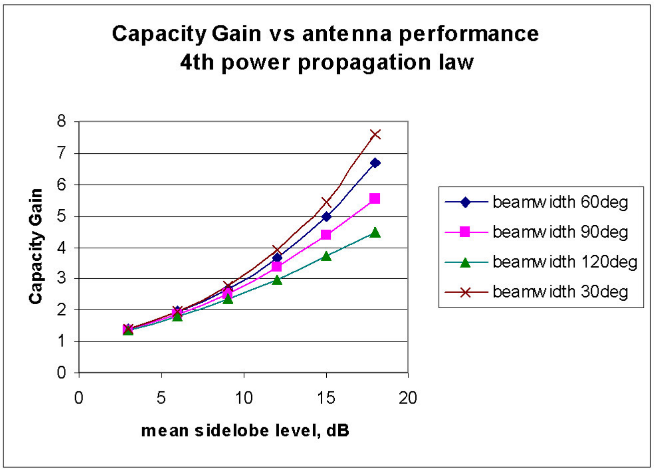

Figure 7: Theoretical Capacity Gain vs. Antenna Performance

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.

2.4 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. 2.2). 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 2.1). 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.

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 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. One way scalability could be achieved is by providing additional capacity in the form of an access network so forming an "Access Mesh". We further conclude that there remain fundamental tradeoffs 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 reuse that would otherwise be possible. Finally we note that spectrum utilisation could be improved by operating meshes within higher, less precious spectrum.

In summary, we still believe that whilst mobile meshes do not live up to all the claims in the literature, they can be very benficial in the area of coverage extension. 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.

Acknowledgements

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 WiFi*", 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.