Modelling Small Cell Deployments within a Macrocell

Professor William Webb MBA, PhD, DSc, DTech, FREng, FIET, FIEEE[1]

Abstract

Small cells, or microcells, are often seen as a way to substantially enhance the capacity of cellular networks. Previous assumptions have been that by deploying a dense layer of small cells within a macrocell, capacity can be improved by an order of magnitude or more. However, there are complexities such as the need to share frequencies between macrocell and small cells, varying patterns of users, the balance between indoor and outdoor subscribers and the different options available within 4G for balancing loading. This paper describes a model that simulates the impact of small cell deployments in macrocells in a typical 4G network and shows that in some cases small cells can actually reduce capacity, while in the best case, maximum capacity gains are less than 100%.

Introduction

Cellular networks frequently need to grow their capacity as the data demand of users increases, often by 40-50% per year. There are many approaches to growing capacity including the addition of new frequencies, greater sectorisation of cells, use of multiple "MIMO" antenna systems and increasing the density of the macrocells. However, for some operators, especially those with relatively limited spectrum portfolios, these approaches have now reached their limit. The next step is often seen as the introduction of small cells (sometimes termed "microcells") within the coverage area of the macrocell.

It is often thought that adding small cells increases the capacity by approximately the number of small cells added, so adding 10 small cells within a macrocell would result in a ten-fold capacity increase. However, the experience of many mobile network operators (MNOs) has been far smaller gains with the result that small cells are not as extensively deployed as some predicted previously. This paper models small cell deployments in LTE networks in order to show why capacity gains are less than might be expected, and where optimal cell numbers might lie.

Frequency allocation for small cells

Each small cell will need to be assigned radio frequencies. Typically, when an MNO reaches the point that small cell deployments are being considered, they have already deployed all their available frequencies within the macrocell layer. They have the following options for assigning small cell frequencies:

1. Macrocell and small cells share the same carrier. (Termed "shared carrier".)
2. Macrocells and small cells each have a dedicated carrier created by splitting the original carrier in two. (Termed "dedicated carrier".)
3. Resource blocks (RBs) are set aside for users on the edge of small cells using enhanced inter-cell interference cancellation (eICIC) [1] where the macrocell does not transmit on these RBs. (Termed "RB".)

With a shared carrier both macrocell and small cell transmit at the same time on the same frequencies. For users close to one base station (eg the macrocell) but far from another (eg the small cell) this can work since the wanted signal level will be high and the interfering level low. But for a

[1] Professor Webb is a director of Webb Search Consulting. The work reported here was commissioned by Telefonica UK and performed in collaboration with NERA Economic Consulting.

user on the edge of the small cell interference can occur since they will experience a relatively weak signal from the small cell and a relatively strong one from the macrocell. The inclination of the network is then to hand them over to the macrocell, effectively reducing the coverage radius of the small cell and hence the percentage of users it can serve.

With a dedicated carrier, part of the frequency allocation is removed from the macrocell to be made available on the small cells. This resolves all interference issues between macrocells and small cells, but reduces the capacity of the macrocell and hence the combined macrocell/small cell combination for low numbers of small cells.

The eICIC approach sits somewhere in between these. It effectively dedicates sub-parts of a frequency band, as needed, to particular users who are towards the edge of the small cell. As a result, it might be thought it would deliver the highest performance. However, as the sections below discuss this is not always the case.

Simulation Environment

The simulation area is based on a sector of a cell. For simplicity the sector is assumed to be 90^o and square (rather than 120^o and pie-shaped). This makes placement of the small cells much simpler. This simplification does not materially change the results. Hence the macrocell is at the origin (0,0) of the square simulation area.

Small cells are then placed throughout the area. There are two approaches used here:

1. Random with minimal overlap. The location of the small cell is selected randomly such that the cell lies entirely within the macrocell area. The small cell is then tested for overlap with any other small cells already sited. If there is overlap a new random location is selected. If, after 100 attempts to find an overlap-free location, none can be found, then 10% overlap with other small cells is allowed and the process repeated. The allowed overlap percentage then increases to 20% and so on. The results of a typical deployment using this approach with 20 small cells in the macrocell area is shown in Figure 1. This approach is intended to mimic real-life where MNOs will seek to avoid overlap between small cells as far as possible but will be limited by the sites available for them to mount their base stations.
2. Hot spots. Here a number (*n*) of hotspots are assumed within the simulation area. The first *n* small cells are placed to cover these hotspots (with the ability to be offset by a chosen amount to reflect the reality of siting constraints). Any remaining small cells are placed randomly as above. Optionally, overlap with small cells covering hotspots can be given a higher penalty value such that overlap is less likely.

Next users are placed randomly across the simulation area. Where there are hotspots, then the selected percentage of users are placed within the coverage area of these hotspots. Users are also assigned to be indoors or outdoors using a percentage selected in the simulation. Indoor users are then assigned randomly to a floor within the building of between level 0 and 5. However, within hotspots all users are assigned to level 0 on the assumption that the small cell would have been sited to be able to capture all the traffic in the hotspots (eg in a stadium or shopping mall).

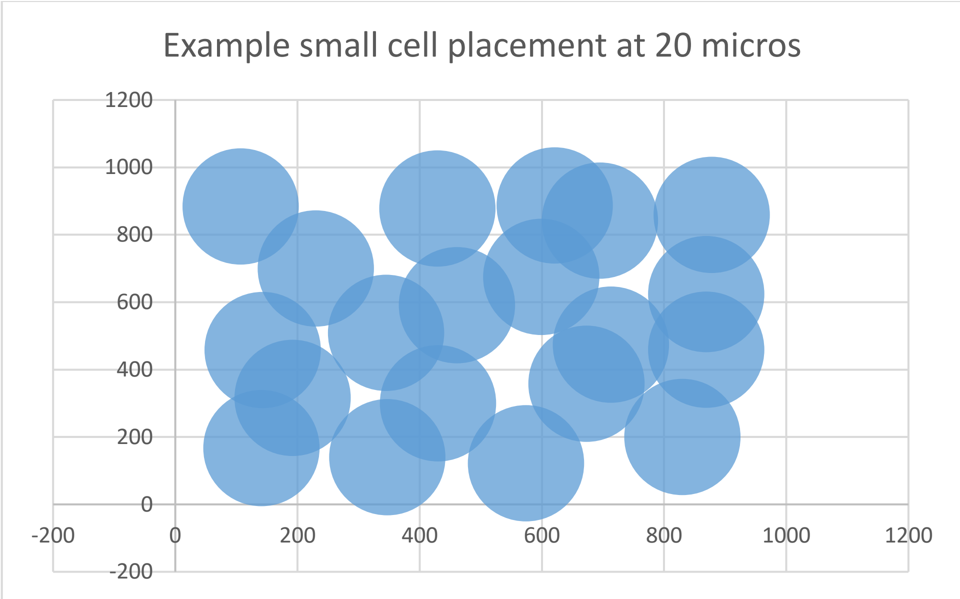

Figure 1 – A typical distribution of small cells in a dense deployment scenario

Note that only outdoor small cells are considered. Indoor small cells can be effective but MNOs typically find it difficult to gain access to the buildings to install them, and uneconomic in that one small cell per floor of each building might be needed. Instead, users tend to self-provide coverage with Wi-Fi which meets their data needs although there may be some issues with voice calls.

User data rates

For each user their maximum downlink data rate is calculated. This is determined according to the signal-to-interference ratio (SINR) using a best-fit curve to the performance of a typical LTE system – essentially a look-up function that takes the SINR and returns the data rate in bits/s/Hz.

The process of determining the SINR is as follows, with most steps being further explained below.

- Determine whether the user is using the macrocell or a small cell. This becomes the "serving cell" and all others are "interfering cells".
- Calculate the signal level from the serving cell.
- Calculate the interference level from all interfering cells – potentially including the macrocell and other small cells.
- Calculate the noise floor according to standard equations.
- The SINR is signal level minus interference and noise.

The determination as to whether the user is camped onto a macrocell or small cell depends on the deployment strategy as follows:

- Shared. The model selects the cells with the strongest signal level.
- Dedicated. As above, the model selects the cells with the strongest signal level.
- RB. The model determines the difference between the small cell and macrocell and if this difference exceeds a user-set threshold, it selects the small cell. This allows small cells to optionally be preferred even when their signal level is lower than the macrocell, effectively extending their range.

The signal level is calculated as follows:

- Macrocell. Two models are used. For distances above 1km (rarely encountered in practice since the model is generally user-set to have a maximum macrocell range of 1km) the Hata urban propagation model is used [2]. For distances below 1km, the Walfisch line of sight (LoS) urban model [2] is used.
- Small cell. A classic two-path microcell model with breakpoint at 100m is adopted. In addition, a further step-function increase in path loss is added at the assumed maximum range of the small cell.
- Indoor. A user-set percentage of subscribers are located indoors and a building penetration is added to the path loss. For macrocells, the penetration loss is constant at 15dB. For small cells the loss is assumed to be low near the cell where the angle of visibility into the building is high, rising to higher penetration levels as the distance increases and the angle of visibility down the street becomes increasingly oblique. The model assumes 10dB penetration up to 20m distance, rising at 0.2dB for each metre further from the transmitter (so at 70m the penetration loss would be 20dB). Penetration loss to users above floor 1 from small cells is assumed to be infinite since the small cell antenna is typically located below this level.
- Transmit power levels are assumed for macrocells and for small cells.

By the end of this process each user has an assigned data rate that they are able to receive at.

Network capacity

The users are assigned a desired data volume. In this simulation, this is set at a relatively high level equivalent to 5Gbytes/user/month and 2,000 subscribers in the 1km^2 sector to ensure the network is fully congested (which allows maximum network capacity to be determined). The time each user needs to receive their data is equal to the data volume divided by their determined data rate. This time is effectively their percentage of the cell capacity used (so if they need 60s to receive their typical hourly data requirements they use 1/60th of the capacity of their serving cell). This process continues until all the capacity of a cell is used at which point the data carried by that cell is totalled. The cell capacity is determined by the size of the carrier and any allocation set aside for resource blocks. The simulated capacity is then the sum of all the capacity across the small cells and the macrocell combined.

The associated network cost can be simply computed as the capital expenditure (capex) and operational expenditure (opex) for the macrocell, and the capex and opex for each of the small cells deployed. This allows metrics such as cost/busy hour Mbyte carried to be calculated.

Results

In generating results it is necessary to set percentages for the number of users indoors and the percentage of users located in hotspots within the macrocell. These percentages will vary from one macrocell to another and over time. For that reason a range of different scenarios are modelled below.

The model considers all possible spectrum assignment approaches described above and selects the optimal policy. As might be expected, for small numbers of microcells (typically less than two) a shared carrier is optimal. After that the model prefers a dedicated carrier until extremely large numbers of small cells (typically around 20) are reached when a RB approach with most RBs (70%) being used in the small cell is optimal. Further assessment of the results shows that the differences between the dedicated and RB approach is relatively small, suggesting that it is not critical which of these is chosen. The somewhat counter-intuitive nature of the results also indicates the complexity of the situation and the reason why detailed modelling is needed to understand the outcome.

The results for the scenario with three hotspots carrying 50% of traffic, 50% indoor users are plotted in the following chart:

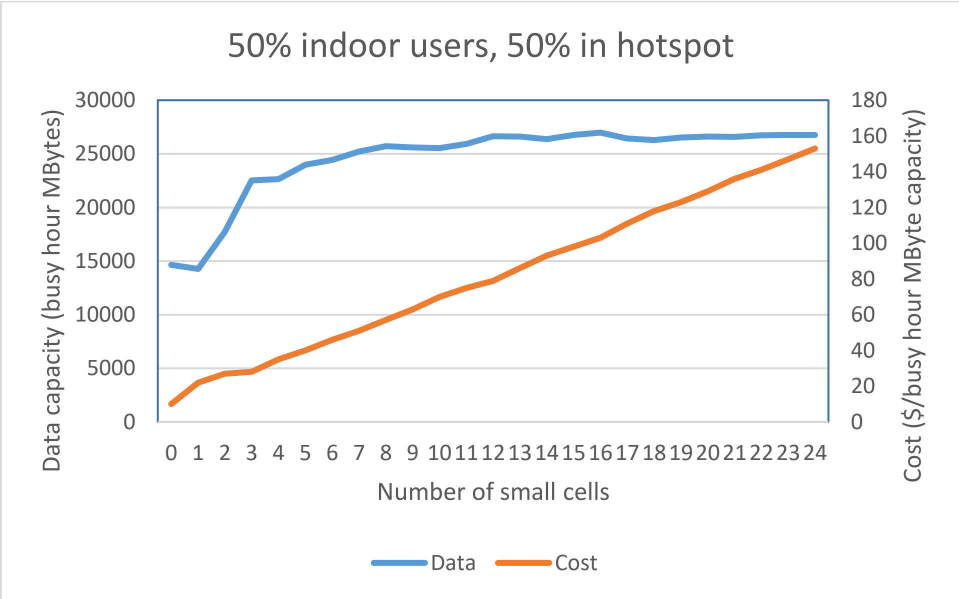

Figure 2 – Results for 50% indoor users and 50% spread across three hotspots

The results show:

- The first small cell actually reduces the capacity. This is because it reduces the macrocell capacity through generating interference by more than the capacity it adds.
- The second and third small cells which are targeted at hot spots, add substantial capacity as they can reuse the same frequencies used by the first small cell and so do not materially increase interference. These hotspots have been selected to be well-spaced around the macrocell and so do not have significant interference between themselves[2]. With three small cells the sector capacity has been increased by 50%.
- Going from four to around nine small cells provides some small gains. Gains are limited because the small cells cannot serve many of the indoor users and so do not attract large volumes of traffic. At this point the overall capacity increase is around 75%.
- Beyond this, capacity is essentially static as additional small cells increasingly overlap with existing small cells.
- Costs per unit of data carried rise throughout, being three times higher for three small cells, six times higher for nine small cells and over ten-fold beyond this. Hence, small cells are an expensive way of providing further capacity.

From these results we might conclude:

- Deployment of small cells in traffic hotspots can be effective.
- There is little point in deploying beyond the number of hotspots in a sector, and more generally beyond about three-four small cells per sector.

[2] If the hotspots were close together the results would be worse due to the interference between them.

- With a complete layer, capacity gains of around 75% on the case where there are no small cells are possible, but higher gains (e.g. 10x) cannot be achieved.
- Small cells significantly increase the cost per Mbyte of traffic carried, which would reduce profitability or require ARPU increase.

If there are no hotspots, the results are as shown in Figure 3:

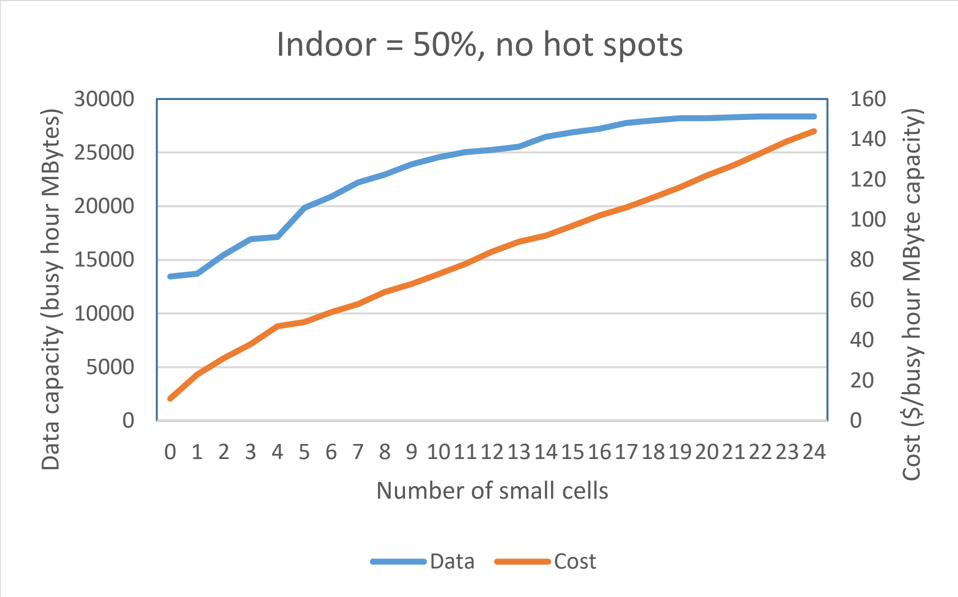

Figure 3 – Results for 50% indoor users and no hotspots

This shows a steadier climb in capacity to around 10 small cells, with a gradual plateauing out after that at a similar level of capacity increase as the previous scenarios. Costs rise somewhat steadily throughout and are higher at lower numbers of small cells as might be expected from the fact that the first few cells do not carry the same traffic volumes as when there are hotspots.

Figure 4 shows the situation where there are no indoor users but 30% are in hotspots.

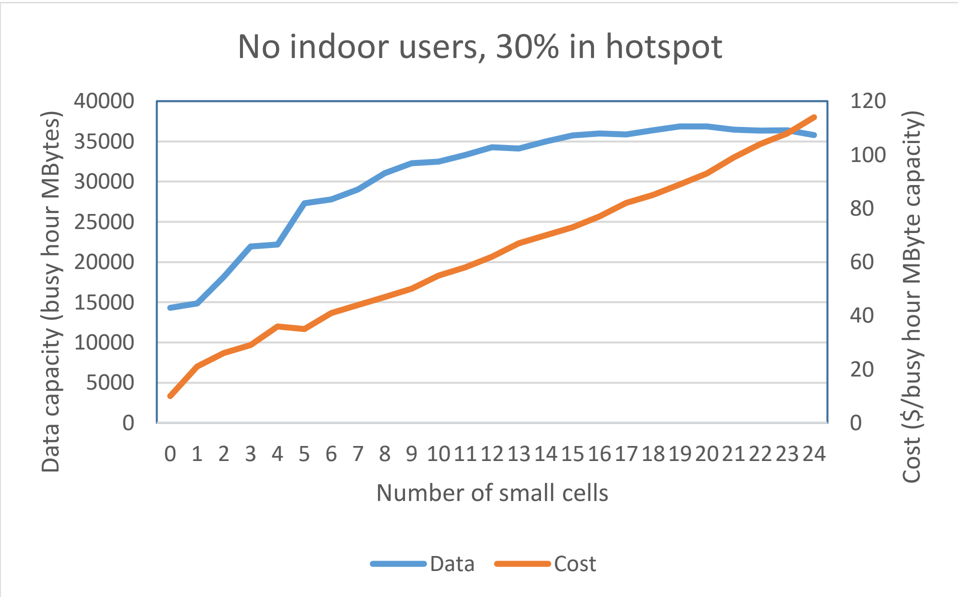

Figure 4 – Results for no indoor users and 30% in hotspots

Here we see a similar impact of the first three small cells as was the case for the other hotspot scenario then a steady climb to a plateau at about 12 small cells. This shows the greatest level of capacity increase, growing by just over 130%. The overall capacity increase is greatest as the small cells can access all subscribers in this case.

There is clearly a very large difference between the 1000% capacity improvement that might have been expected with 10 or more cells and the 75% or so that is seen in practice in the most likely scenarios. This is because:

- Frequencies have to be taken from the macrocell, reducing its capacity, but it is only the macrocell that can serve users in most buildings and in the gaps between small cells. If the macrocell becomes heavily congested, cell capacity falls.
- Small cells increasingly interfere with each other as they get closer together, reducing the effective capacity of each.
- Small cells, especially outside of hotspot areas, may not be able to attract many subscribers and hence may be under-utilised even while the macrocell is congested.

Conclusions

The key conclusions are:

1. Small cells are not a source of infinite capacity expansion. The best possible improvement is around 100% increase (2x) over a sectored 1km radius macrocell.
2. The optimal number and deployment strategy vary depending predominantly on the presence of hotspots in the sector and also the percentage of indoor subscribers. In most cases deploying more than around three small cells is not worthwhile.
3. A hot-spot strategy will nearly double the cost of carrying traffic in the sector on a $/bit basis compared to using a macrocell alone. A dense layer will result in a six-fold cost increase and a complete layer more than a ten-fold increase.
4. Capacity improvements beyond these levels will require indoor picocells. These have not been modelled but typically improve capacity owing to the shielding offered by the building which reduces interference.
5. The situation is complex, requiring a cell-by-cell evaluation of optimal strategy.

The implications for MNOs are significant. It is not possible to use outdoor small cells as a way to substantially add capacity in the manner previously thought. This could leave an MNO that has already deployed all of its spectrum and all other capacity enhancement approaches in a position where it is no longer able to grow capacity to meet growing demand absent being able to access additional spectrum.

References

[1] http://www.3gpp.org/technologies/keywords-acronyms/1576-hetnet

[2] Haslett C, “Essentials of radio wave propagation”, Cambridge University Press, 2008.