Proceedings of the Federated Conference on
Computer Science and Information Systems pp. 1193–1198

ISBN 978-83-60810-51-4

# Energy optimisation of the wireless access network through aggregation of mobile terminals

Hanno Hildmann, Sébastien Nicolas, Fabrice Saffre
Etisalat BT Innovation Centre
Khalifa University, PO Box 127788, Abu Dhabi, United Arab Emirates
Email: {firstname.lastname}@kustar.ac.ae

***Abstract*—We suggest the assignment of mobile terminals (MT) to base stations (BS) such that the number of expendable BSs is maximised (US Patent 13/339,763 (pending) [1], [2]). The investigation uses an existing implementation as base case and investigates a variety of specific scenarios; (simulated) computational results are provided for the individual test scenarios.**

## I. Introduction

CELL phone coverage is achieved through interlocking areas of coverage from a population of base stations (BS). These areas of the individual BS overlap and mobile terminals traveling through these fields are "handed over" from BS to BS. Depending on the degree of overlap, terminals can be shifted from one BS to another even when they are not moving.

This research is investigating the assignment of terminals to BS such that the number of active BS is minimised. The assumption is that by freeing as many BS as possible (without compromising on QoS, the quality-of-service) provided to the individual terminals, some financial and operational benefit can be gained and that the overall power consumption of the network is loosely proportional to the number of active BS. In the literature, [3] suggests performance regulations and site solutions as two of the three approaches to reduce energy consumption (improving the hardware being the third) while [4] and [5] support the claim that BSs constitute a major part of the overall power consumption of the network.

In this article, we investigate the potential benefits (in terms of active BS, coverage and number of handovers) which can be achieved under varying scenarios, i.e. for different ranges, user density, BS topology, user speed and terminal flexibility.

### *Overview*

The paper opens with some background on base station networks and mentions the opportunities for improvement (Sec. II). Section III provides the model, while Section IV briefly discusses the implementation thereof and describes the simplifying assumptions that were made. The majority of the article is dedicated to discussing the extent of optimisation possible and the potential ramifications of subjecting the system to it (Sec. V). The paper closes with a conclusion and future work (Sec. VI) and provides the relevant references.

## II. Background

Over the last decades cellular phones have become a pervasive technology, their supporting infrastructure has grown with them and (BS) are now a common feature of urban as well as rural landscapes. Due to the nature of the process, there are many properties of BS networks that can be considered sub-optimal, for example the location of the individual BS or the coverage they provide. With the rapid increase in MT usage as well as the ever changing landscape of cities and the providers, and thus the corresponding networks, the optimal placement—in itself an undefined term—of BS within such a network is difficult or even impossible.

### *A. Redundant capacity*

The operators of these networks aim to provide full coverage and comprehensive quality of service even (or especially) during peak operation times and therefore operate networks which are dimensioned to accommodate the peak volumes of terminals and traffic. Due to this, we conjecture that there is room for optimisation in times when the capacity of the network is not fully required. While network operators are notoriously uneasy about actually shutting parts of their network down, it can be argued that the optimisation of the number of customers associated with the individual BS, as well as ability to increase the number of redundant nodes (BS), has various benefits. The industry [6] is continuously increasing its engagement towards environmental friendly solutions [7]; furthermore the presented approach is making steps towards compliance with future greenICT regulations, due to increasing political and public pressure on the industry.

### *B. Making use of existing operational functionality*

Existing networks already have the following properties:

- the BS in a network are not always used to the maximum of their potential capacity,
- the fields of coverage of BS do overlap and
- handover mechanisms are already part of the existing implemented hardware and functionality.

Our approach is thus a straight forward use of these properties: We aim to reassign certain active MT associated with some BS `A` to another BS `B` (which can provide coverage for the terminal), such that `A` will eventually not be required to serve any terminal at all. Once this is achieved, we will consider `A` *inactive*. Inactive BSs will continue to *listen* and can become active again once their service is required to protect the quality of service constraints (i.e. all mobile terminals within the area of coverage of the network must be serviced at any time).

We argue that this approach is possible from a hardware point of view, reducing the required effort to implement it to the matter of upgrading the software for both the supporting infrastructure as well as the BSs in the field. Specifically, there is no need to make any changes to the mobile terminals so that their active participation or consent is not required.

## III. The model

In what follows we investigate the approach and its potential benefits on the basis of a model, i.e. for a simplified network in which we omit a variety of real world aspects, none of which being relevant for the investigation. In line with this we abstract from all technical details as well as hardware issues and consider only 5 parameters for our investigation:

1) *density* ($d$): the number of MT in the simulation. The mobile terminals are placed randomly (see Figure 2), their initial placement is in no means dependant on the location of the available BS. As the area over which the BS network is distributed is the same for all simulations the number of $d$ MT determines the density of users.
2) *range* ($r$): the range if a BS. BS are assumed to have a perfect circular area of service, defined by its radius $r$.
3) *speed* ($s$): distance $s$ a MT is traveling in one time step. MTs are directed randomly at locations in the simulation area and travel towards them until they reach them and are assigned a new destination. We refer to [8] and [9] for the claim that users exhibit a significant probability for having a few highly frequented locations and that users follow simple and reproducible patterns.
4) *topology* ($t$): the distribution of the BS over the simulation area. Two basic networks are deployed and tested: the perfect hexagonal network (used as the benchmark and base case) and a random distribution of BS (Figure 3) where the system tries to enforce a minimum exclusion distance $t$ imposed on the (otherwise random) distribution of BS. It does so by trying 10 random locations and checking if one of them satisfies an exclusion zone x of the demanded value. If it fails, it does 10 further tries with a smaller exclusion zone x/1.5. This process of reducing the expected exclusion zone repeats 10 times after which a random location is chosen.
5) *usage* ($u$): represents the probability that an unengaged terminal becomes engaged. This gets evaluated each time step for each terminal and if this happens, it stays engaged for 0 to 100 time steps (flat probability).

The degree of optimisation (i.e. the potential benefits of the approach) is measured and expressed by 3 values:

1) *active BS* ($a$): the number of BS which are active when the convergence of the system has been reached. This is the main measure of the benefits of the system. [5] states that only about 5% of the power consumption of an active BS are used to transmit data (provide service).
2) *coverage* ($c$): the number of mobile terminals (out of a total of 1000—as compared to e.g. 100.000 investigated in real scenarios [8], [9]) which are assigned to a BS. This is the main measure of the QoS experienced by the MTs and should be kept maximal and controllable.
3) *hand-overs* ($h$): the number of hand-overs of MTs between the BS. In the literature, handovers can be used as criteria for comparative studies between networks (e.g. [10]). Since handovers consume energy and can sometimes be noticed by the user (e.g. during a call), they affect QoS and should be minimized and controlled.

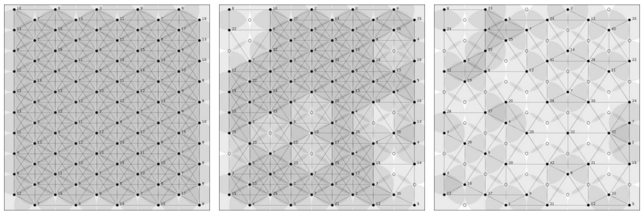

Fig. 1. Evolution over time of the optimisation of a hexagonal base station network, from all BS active (left) to an optimised state with a percentage of the BS passive (right). The mobile MT are removed to enhance readability.

With respect to $c$ it should be noted that due to the manner in which the networks are (automatically and randomly) created by the model, there may be small areas where there simply is no coverage. This is in line with the real world; in such cases no service can be provided by the network, neither before nor after the optimisation. See Figure 7 where the worst case overall coverage is given as 99.4%.

## IV. Implementation

The model was implemented in Java as a Monte Carlo simulation, taking its input parameters from an XML configuration file, and writing out the evolution of the system to a comma-separated CSV file. More specifically, the simulation outputs at each time step the 3 values above ($a$, $c$, $h$), in addition to computing the convergence time, which is the time it took the system to stabilize (notwithstanding random fluctuations). The latter is used as a debugging and fine-tuning value, to sanity-check the consistence amongst various simulations. The tool also offers real-time parameter modification and visualisation / graphing capabilities, to provide quick exploration of the parameter space, and visual debugging / tuning.

For the implementation of the model, a number of simplifying assumptions were made, namely that:

- The range of a BS is perfectly circular and is not in any way obstructed by the terrain. Furthermore, it can not be altered with respect to capacity, direction or orientation, as it does happen in real networks.
- BS share the same frequency (unlike real networks), the overlap is strictly dictated by the circle's overlap area.
- All BS have the same range.
- A BS which is going inactive will remain so for a duration of 100 time steps. In real networks indeed, it takes a

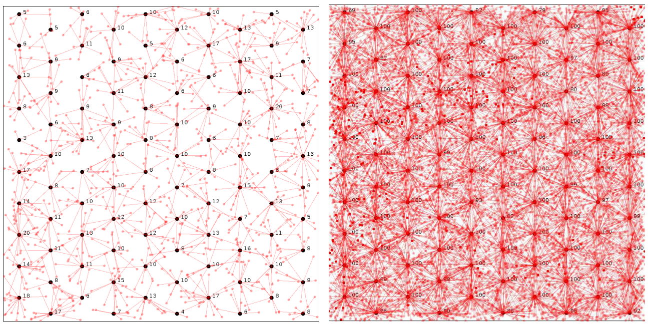

Fig. 2. Randomly placed terminals for $d=$ 100 (left) and $d=$ 1000 (right). The dots are the BS and the lines are connecting the MT to the BS.

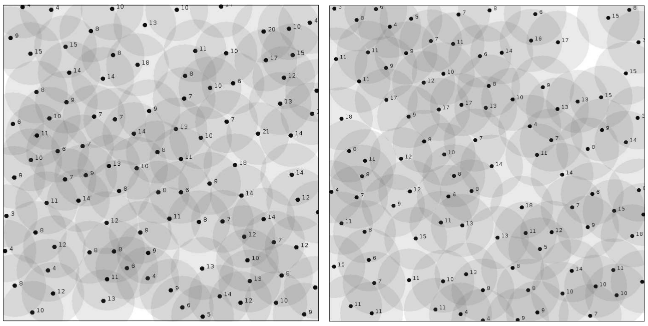

Fig. 3. Two examples of randomly generated non-hexagonal BS networks. The overlap of fields of network coverage created by these BS varies between single coverage and triple redundancy

significant amount of time to power-up the transceivers (of the order of the minute).

There are 100 BS in all simulations. Apart from graphs where $d$ is varying, we have simulated 1000 MT, initially located randomly in space. Every data point in each graph is obtained by averaging the results from 1000 simulations, bringing the variance down to 1% – meaning in practice a difference of 1 active or inactive base station. An approximate number of 500,000 simulations were run on 2 dedicated identical PCs over two weeks. A total of 60 MB of data was generated and aggregated into graphs; the most interesting of which are discussed in the next section.

## V. Results

As can be seen in Figure 4, the network quickly converges towards a state where only roughly 60% of the BS are active. This is also consistent with the drop in handover numbers since the system quickly converges to a stable situation. The minimal decrease in coverage (at most 0.6%) is effectively a tradeoff between the number of BS that can be deactivated, and the perceived QoS by the terminals. This parameter is fully controllable (see Figure 9). In some simulations (terminal movement and low BS range), an additional contributor to the QoS degradation is the fact that users will find themselves in an area where no coverage is possible anyway. In some

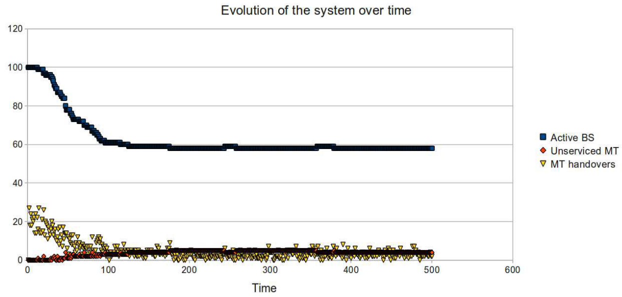


Fig. 4. Results: Decrease of number of active BS ($a$) and hand-overs ($h$) as well as variance in coverage ($c$) over the course of one (exemplary) run of the simulation. These results are representative for the evolution of the increasingly efficient optimisation of the network.

situations service can not be provided because the maximum number of 100 mobile terminals per BS has been reached.

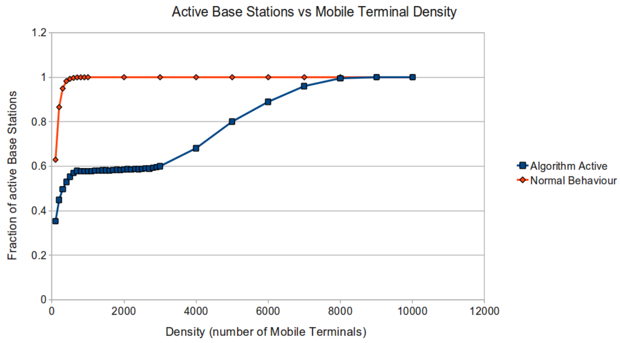


Fig. 5. Density ($d$) vs. active BS ($a$), average final solution of 1000 simulations in a hexagonal network with and without optimisation.

We now investigate the impact of the individual parameters introduced in Section III and report on the outcome of the simulations. Where appropriate we have investigated the outcome of the simulation for identical settings but for the case when all BS are active (if servicing at least one MT) and where no coordinated effort is made to hand mobile terminals over. In order to showcase the potential benefit, the results are compared to this benchmark case (see Figures 5 and 12).

### A. Increasing number of users (density)

With respect to the *density* ($d$) of the users (Figure 5), we notice a plateau for the number of active BS around 60% which is maintained until about the time when the number of mobile terminals in the simulation exceeds the maximum capacity of the active 60% of the BSs. These results are for the hexagonal network and under an even distribution of random placement of the users. The benchmark case when no optimisation is attempted is provided for comparison. A real world network and a more biased movement pattern (cf. [5], [11]) will not yield these results during peak times, however, for off-peak times even larger benefits can be expected.

The 3D diagram above shows a drop of the ridge with increasing range due to the degree of overlap between the

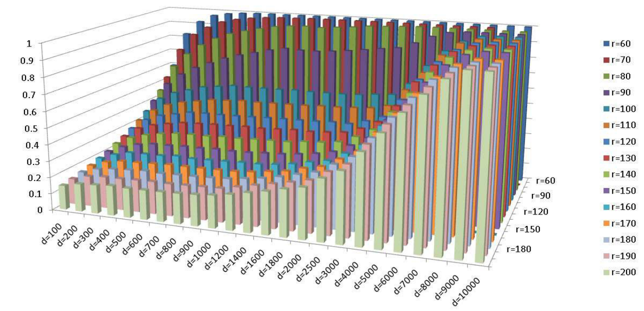


Fig. 6. The impact of the range of the BS on the potential benefit (i.e. the plateau shown in Fig. 5). As expected, for smaller ranges the ability to hand over terminals to BS which are still below maximum capacity is diminished, resulting in the disappearance of the plateau.

individual BS's fields of service (see Figure 8); with increasing redundancy, larger percentages of optimisation are achieved.

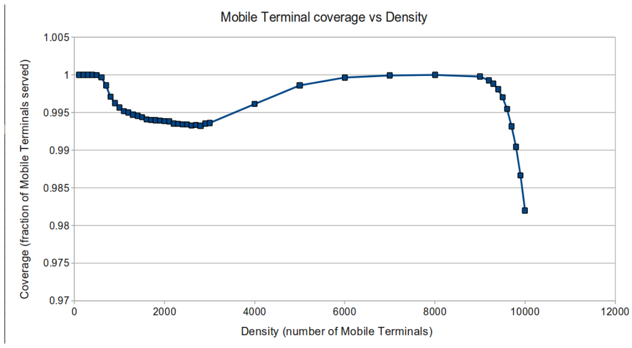


Fig. 7. Results: Density ($d$) vs. Coverage ($c$), average final solution of 1000 simulations in a hexagonal network. These results are obtained for the QoS constraint set to high (i.e. 99%)

The graph in Figure 7 shows the relation between $d$ and $c$. The sharp drop towards the end is due to the saturation of the individual BS: When approaching the maximum network capacity, an increasing number of terminals will be in areas where the coverage is already saturated. Note that if we expect this latter region, the decrease in QoS is limited to 0.6% only.

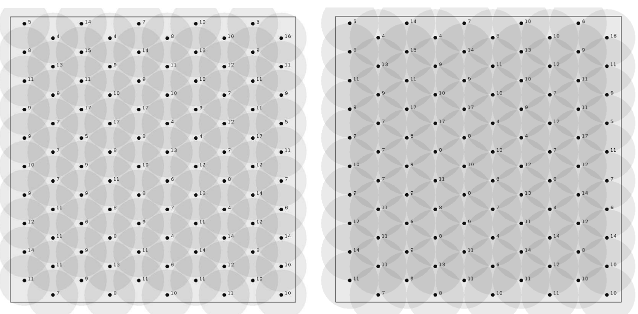

Fig. 8. A prefect hexagonal network with $r$ 90 (left) and 100 (right). Note the overlapping fields of coverage: Above $r = 100$ there is increasing quadruple redundancy, i.e. increasing areas for which there are 33% more BS to shift mobile terminals to than before; this explains the ridge in Figure 6.

## B. *Varying range of the BSs (range)*

With respect to the *range* ($r$), Figure 9 (Range vs. Coverage) and Figure 10 (Range vs. active BS) show the impact of increasing the range of the BS. These results are for a hexagonal network and the relationship between range and performance is directly related to the redundancy of service provided, i.e. the extent to which the fields of service of the individual BS overlap. Figure 8 shows that *for the hexagonal network* the redundancy becomes partially quadruple when the range exceeds 90 (i.e. 100 and above). This means that there are increasing areas where a mobile terminal can be assigned to one of 4 BS in order to free the other three.

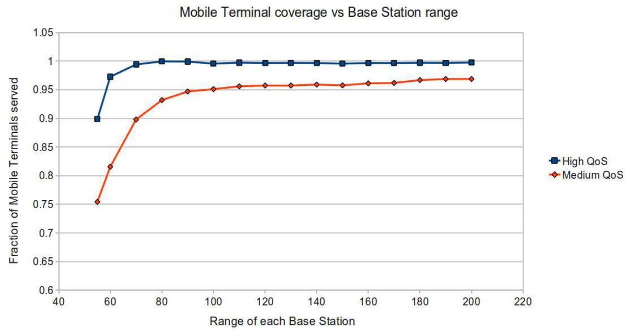


Fig. 9. Results: Range ($r$) vs. Coverage ($c$), average final solution of 1000 simulations in a hexagonal network. The two graphs show the results for varying quality of service commitments: 99% (higher) versus 95% (lower).

Regarding the service level provided by the network (i.e. QoS) investigations (shown in Figures 9 and 10) were using different settings to enforce the the degree to which the network *has* to provide absolute coverage to all terminals. The commitments considered were 99% (high) versus 95% (medium) (cf. Figure 7). For increasing range the results converge. As the system is dynamic this has an effect on the flexibility of the BS, and is investigated here to indicate that compromising on QoS for the customer (Figure 9) provides decreasing benefits for increased range (Figure 10).

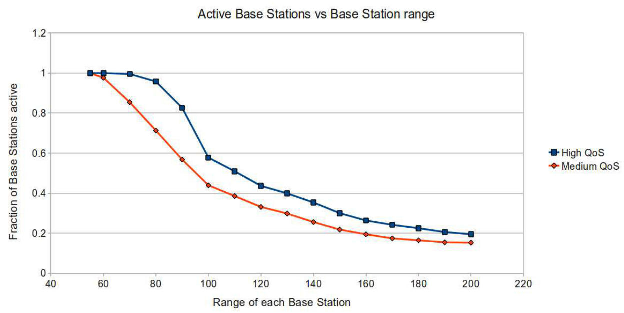


Fig. 10. Results: Range ($r$) vs. active BS ($a$), average final solution of 1000 simulations in a hexagonal network. The two graphs show the results for varying quality of service commitments: 99% (higher) versus 95% (lower).

### C. Mobility of users (speed)

Mobile terminals are rarely fixed in the real world and exhibit complex movement patterns. In the present paper, we have taken a look at the influence of speed and the results are shown in Figures 11 and 12. For reasonably low values of the speed, the efficiency (number of active BS) is slightly better that the benchmark, because we effectively introduce noise in the system, thereby removing potential locks due to the current random topology of BSs and terminals. This comes at a cost in terms of QoS however, since 0.2% less customers get serviced. This is expected since it takes 100 time steps for BS to be reactivated.

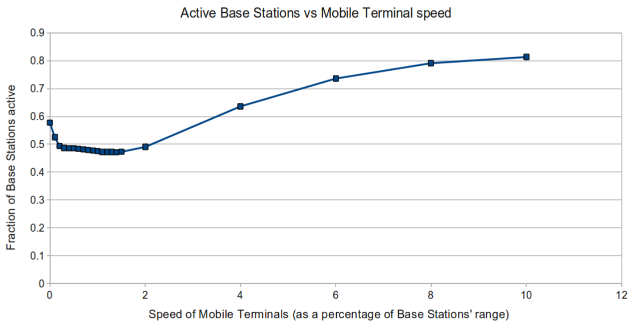


Fig. 11. Results: Speed ($s$) vs. active BS ($a$), average final solution of 1000 simulations in a hexagonal network.

On the other end of the spectrum, it is clear from the graphs that the BSs cannot cope with too much mobility so that this decreases both the efficiency and the quality of service. Effectively, BSs can never be passivated since there will always be a number of new terminals appearing in their vicinity. Note that this situation is unrealistic in most mobile telephony networks though, since users don't normally move by more than 1% of a BS range every second. Investigations into special cases of BS networks build to accomodate high speed travel connections (motorways, high speed trains) are outwith the scope of this work, (see Sect. VI on Future Work).

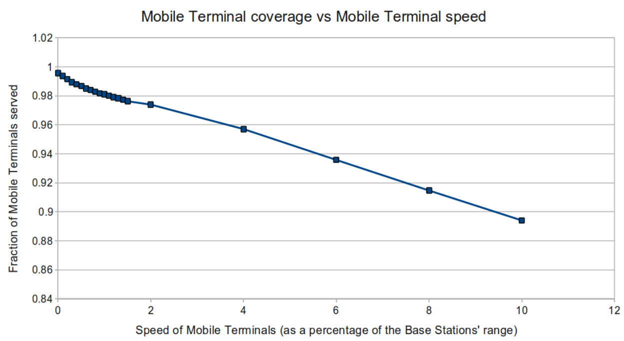


Fig. 12. Results: Speed ($s$) vs. Coverage ($c$), average final solution of 1000 simulations in a hexagonal network.

### D. Layout of the network (topology)

When considering the influence of the placement of BS over performance, the most noticeable effect is that for very low values of the exclusion range, a larger fraction of BS is able to go into hibernation (see Figure 14). However, this is the result of their more random placement, which entails very high concentration in some places (facilitating aggregation) and complete lack of coverage (blind spots) in others. This in turn leads to comparatively poor average QoS seen in Fig. 13.

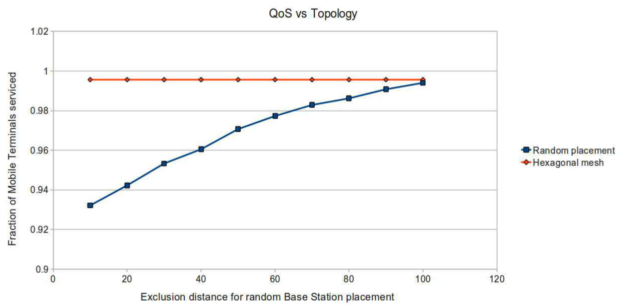


Fig. 13. Results: Topology ($t$) vs. Coverage ($c$), average final solution of 1000 simulations in a hexagonal network.

As the exclusion range reaches values similar to the distance separating neighbouring BS in the hexagonal mesh, apparent performance gradually decreases while QoS increases. Although counter-intuitive, the fact that the fraction of active base stations eventually becomes higher than in the benchmark (hexagonal topology) can be understood when considering the BS placement algorithm. Indeed, for higher exclusion range, the pattern becomes something of a hybrid between an approximation of the regular mesh (ensuring few blind spots) and a random distribution. This leads to heterogeneous overlap in different parts of the network, meaning that some BS have very few or no opportunities for hibernation, as they are the only available access point in their area.

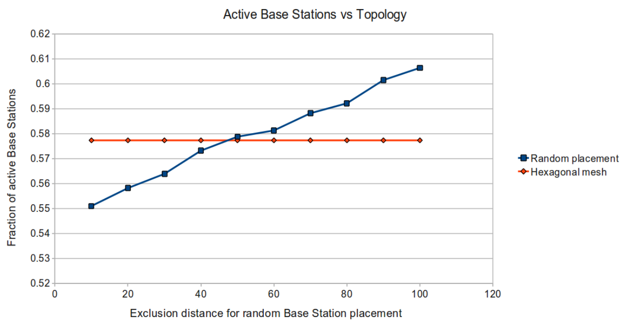


Fig. 14. Results: Topology ($t$) vs. active BS ($a$), average final solution of 1000 simulations in a hexagonal network. The intersection of the graphs occurs (as expected) at a range where the random placement is more or less forced to approach a hexagonal network (due to distance constraints).

### *E. Impact of non-transferable terminals (usage)*

Finally, when it comes to the amount of (non-transferrable) terminals in *use* ($u$), the results presented are not surprising. In our simulation, a terminal which is *in call* can not be transferred from one BS to another. Therefore, as expected, the number of active BS increases with the probability of the terminals in use (Figure 15); consequently, as the optimisation of active BS is hindered, the number of handovers decreases with increasing probability for $u$ (Figure 16).

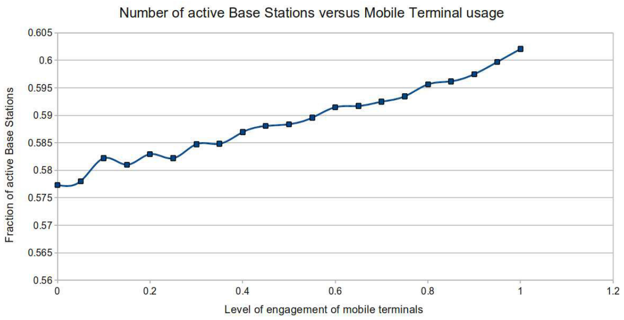


Fig. 15. Results: Usage ($u$) vs. active BS ($a$), average final solution of 1000 simulations in a hexagonal network. As expected, with increasing difficulty to hand over terminals from BS to Bs the number of active BS increases.

We refer to future work (Sect. VI) aiming to use realistic network topologies and traffic patterns [12] (as provided in detail in e.g. [11] for overall demand - voice call only - and distribution of demand over a 24h period on exemplary infrastructure in the USA); future work will see close collaboration with ETISALAT in Abu Dhabi and for now such investigations are outside the scope of this article.

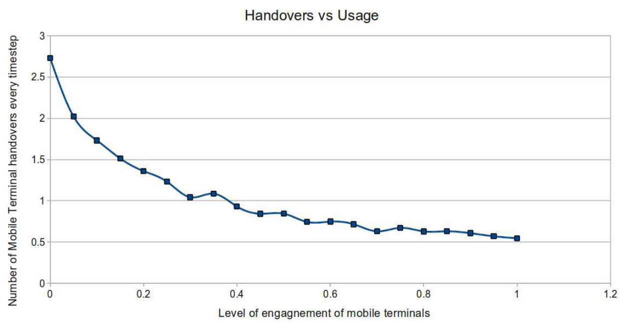


Fig. 16. Results: Usage ($u$) vs. Handovers ($h$), average final solution of 1000 simulations in a hexagonal network. With increasing number of terminals unable to switch BS the number of handovers decreases.

## VI. Conclusion and Future Work

In this paper, we have explored the influence of several key parameters on the performance of an aggregation strategy exploiting coverage overlap in order to release (and "hibernate") a sizeable fraction of BS. Our results indicate that there is an extensive region of this multidimensional parameter space in which such a strategy would be feasible at an almost negligible cost in terms of QoS degradation.

It can reasonably be assumed that aggregating mobile terminals on a sub-set of BS would result in appreciable power savings in the wireless access network. Indeed, hardware vendors are increasingly including intelligent features into their products that allow them to detect situations in which they are under-used, and to respond by spontaneously lowering their energy consumption. Our findings show that, in identical conditions (in terms of terminal density and mobility, network topology, bandwidth availability etc.), the proposed aggregation strategy would statistically maximise the number and duration of such "low power" episodes across the network.

This may however have a hidden energy cost on the terminal side. Indeed, higher handover frequency and longer average distance to the BS will inevitably increase power usage by the devices. Measuring trade-offs between the two aspects, as well as quantifying the impact of transferring some of the combined energy consumption to battery-powered mobile devices, will be the subject of future work.

In addition, by collaborating closely with our industrial partners (Etisalat and BT), we hope to further increase realism by incorporating actual mobility and usage patterns into the model, and to be able to run experimental trials. As mentioned before, specific and extreme scenarios (like networks along high-speed train tracks) are left for future work as these investigations will need to be based on actual network and usage data in order to be meaningful.

## References

[1] Patent, "Method and system for determining allocation of clients to servers," US Patent No: 13/339,763 (pending), F. Saffre, H. Hildmann and S. Nicolas (inventors), Dec. 2011.

[2] H. Hildmann, S. Nicolas, and F. Saffre, "A bio-inspired resource-saving approach to dynamic client-server association," *IEEE Intelligent Systems*, vol. 27, no. 6, Nov. - Dec. 2012.

[3] J. T. Louhi, "Energy efficiency of modern cellular base stations," in *29th Int. Telecommunications Energy Conf. INTELEC*, 2007, pp. 475 – 476.

[4] J. Gong, S. Zhou, Z. Niu, and P. Yang, "Traffic-aware base station sleeping in dense cellular networks," in *$18^{th}$ International Workshop on Quality of Service (IWQoS).* Beijing: ISSN: 1548-615X, Print ISBN: 978-1-4244-5987-2, June 2010, pp. 1 – 2.

[5] E. Oh and B. Krishnamachari, "Energy savings through dynamic base station switching in cellular wireless access networks," in *IEEE Global Communications Conference. GLOBECOM*, Florida, Dec. 2010.

[6] GSM Association, "European Mobile Manifesto," Nov. 2009, http://gsmworld.com/documents/gsma_2015_mobile_manifesto.pdf.

[7] GSM Association, "Mobiles Green Manifesto," Nov. 2009, http://gsmworld.com/our-work/public-policy/mobile_manifesto.htm.

[8] M. C. González, C. A. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 5, pp. 779 – 782, June 2008.

[9] M. C. González, C. A. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns - supplementary information," *Nature*, vol. 453, no. 5, pp. 779 – 782, June 2008.

[10] A. Simonsson and T. Lundborg, "Large scale mobility measurement in two metropolitan networks," in *IEEE $19^{th}$ International Symposium on Personal, Indoor and Mobile Radio Communications. PIMRC.* Cannes: Print ISBN: 978-1-4244-2643-0, Sept. 2008, pp. 1 – 5.

[11] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *$3^{rd}$ IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks. DySPAN.* Chicago: ISBN: 978-1-4244-2016-2, Oct. 2008, pp. 1 – 11.

[12] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary user behavior in cellular networks and implications for dynamic spectrum access," *Communications Magazine, IEEE*, vol. 47, no. 3, pp. 88 – 95, 2009.