

Femtocell Coverage Optimisation Using Statistical Verification

Tiejun Ma and Peter Pietzuch

Department of Computing,
Imperial College London,
London SW7 2AZ, United Kingdom
{tma, prp}@doc.ic.ac.uk

Abstract. Femtocells are small base stations that provide radio coverage for mobile devices in homes or office areas. In this paper, we consider the optimisation of a number of femtocells that provide joint coverage in enterprise environments. In such an environment, femtocells should minimise coverage overlap and coverage holes and ensure a balanced traffic workload among them. We use statistical verification techniques to monitor the probabilistic correctness of a given femtocell configuration at runtime. If there is any violation of the desired level of service, a self-optimisation procedure is triggered to improve the current configuration. Our evaluation results show that, compared with fixed time, interval-based optimisation, our approach achieves better coverage and can detect goal violations quickly with a given level of confidence when they occur frequently. It can also avoid unnecessary self-optimisation cycles, reducing the cost of self-optimisation.

Keywords: Femtocell, Self-Optimisation, Statistical Verification

1 Introduction

Femtocells [5] are cellular base stations that cover small areas of tens of meters. They are low cost and low power devices that are normally installed by consumers in homes for better indoor mobile voice and data reception [3, 5]. A femtocell can be considered as a wireless cellular access point, which transfers data traffic through the home broadband connection to the operator's core network. By 2014, the deployment of femtocells is expected to reach around fifty million, providing service to more than a hundred million people [1].

An important challenge for femtocells is to optimise their radio coverage area dynamically. The goal is to achieve a desired level of performance for mobile transmission, avoid undesired interference and reduce power consumption. Providing optimal femtocell signal coverage is important to improve users' mobile usage experience as well as reduce service cost expenditure. Since femtocells are deployed by users themselves, they must also be able to self-configure all required parameters during operation with minimal user intervention [6].

In enterprise environments, a number of femtocells may be deployed together to achieve joint coverage. This is also done to cover a large area while balancing user load, minimising coverage gaps without signal and wasteful coverage overlaps between multiple femtocells. Self-optimisation adapts the configuration of radio parameters of femtocells at runtime, for example, by setting the power level of their radio signals. Achieving good configuration is challenging due to the diversity and dynamic nature of the deployment environment. Since femtocells are deployed in a decentralised fashion, self-optimising femtocells might interfere with each other. In addition, users of femtocells might move around at random, changing the workload and requiring dynamic hand-over between femtocells.

Current femtocell optimisation methods derive the average value of overlap, gap and workload within a fixed time interval and then compare this value to a predefined threshold, known as the optimisation goal [7, 9]. This approach achieves reasonable performance and adaptability using an evolving algorithm as shown in [9]. However, its convergence rate is slow because self-optimisation is only triggered at fixed time intervals. When violations of optimisation goals are rare, the algorithm cannot prevent unnecessary re-computation of parameters, which wastes energy.

In contrast to observing the average value of a desired property, *statistical verification* [11, 12, 15] is a technique that checks the probabilistic correctness and satisfaction of a system against its desired behaviour. It has two major benefits: first, it adopts a formal mathematical specification to describe the desired system behaviour without ambiguity; second, it measures a property more accurately than the average value because it considers the probabilistic likelihood that a property is violated. The possibility of false positives and negatives can be bounded below a required level [15]. When property violations occur frequently, a fast response to trigger re-optimisation is needed rather than waiting for a given timeout (i.e., false negatives). When violations are rare, statistical verification can also avoid unnecessary self-optimisation (i.e., false positives). Thus, statistical verification reduces energy consumption as well as durations of instability caused by self-optimisation.

In this paper, we show how to use a statistical verification technique to achieve self-optimisation goals for coverage optimisation in a joint femtocell scenario. In particular, the contributions of the paper are:

- We model a femtocell network as a stochastic discrete event-based system and formalise its desired behaviour using stochastic temporal logic. This model describes properties rigorously and can be verified using statistical verification.
- We propose a statistical verification technique based on hypothesis testing for verifying properties in femtocell coverage optimisation. When a property violation is detected, self-optimisation is triggered.
- We evaluate our approach in simulation and show that, compared with a fixed time, interval-based approach, it is effective and can provide a 25%–38% improvement in terms of reduced coverage overlaps and gaps.

The remainder of the paper is organised as follows. In the next section, we describe related work on femtocell coverage optimisation and statistical verification. Section 3 introduces the goals of femtocell coverage optimisation in more detail. Section 4 provides the formal model and definitions for femtocell coverage optimisation. Section 5 presents evaluation results. We give conclusions and possible future work in Section 6.

2 State of the Art

2.1 Femtocell Coverage Optimisation

Due to the predicted adoption of femtocells, researchers have begun to consider the problem of coverage optimisation [4, 6–8]. However, all of this work focuses on single femtocell coverage optimisation for small-area residential users rather than on multiple femtocells that achieve joint coverage in large enterprise environments. The goals are to provide good indoor coverage, preventing signals from leaking outdoors [6], and to increase the flexibility in deployment locations [7, 8].

For multiple femtocells, the main optimisation goal is to reduce coverage overlaps and gaps, as well as to balance the workload among femtocells. Ho et al. [9] propose to use of genetic programming for coverage optimisation of a group of femtocells. Their approach finds a suitable signal power to control a femtocell’s coverage area. The authors show that genetic programming with adjustment of signal power at a fixed time intervals can balance the average workload of femtocells, with only small coverage gaps and overlaps. Their work is closest to ours in terms of the scenario considered for joint femtocells coverage optimisation. However, our focus is on selecting an optimal trigger time for the signal power adjustment algorithm using statistical verification. We want to carry out fast yet accurate violation detection with probabilistic guarantees of optimisation goals. This avoids the drawbacks of approaches that use fixed time intervals and average values, which may take a long time to achieve optimisation goals and cannot avoid unnecessary self-optimisations.

2.2 Statistical Verification

Sen et al. [12] first considered verification of systems based on statistical hypothesis testing. The authors assume that the system has already been deployed so that execution traces cannot be controlled (black-box system). Verification is based on the observation of system behaviour, and they show the effectiveness of verification models based on continuous time Markov Chains. The idea of verifying a black-box system would be suitable for femtocell coverage optimisation when the self-optimisation algorithm is unknown, such as the evolving algorithm proposed in [9].

Younes and Simmons [15] extend Sen et al.’s approach in order to verify probabilistic and time-based properties of black-box systems, as generalised by

semi-Markov decision processes. The authors adopt the *sequential probability ratio test* (SPRT) and verify system traces as stochastic discrete event based models. Such a model is suitable for modelling violations of femtocell optimisation goals. Thus we use a stochastic discrete event based model in this paper. Lee et al. [11, 13] propose a monitoring and checking framework based on SPRT, which verifies the probabilistic correctness of a system based on its implementation at runtime. The authors show that such an approach can be adopted to verify probabilistic properties of wireless sensor networks. The idea of monitoring and checking implementations at runtime is suitable for the coverage optimisation of femtocells. Therefore, we adopt the same SPRT-based verification approach.

Next we give a short overview on the SPRT statistical verification approach. SPRT was first developed by Wald [14] as a likelihood-based test for observed data. It is designed for continuous monitoring. In practice, it is often evaluated at frequent but discrete time intervals. SPRT defines a *Null* (H_0) hypothesis that expresses the accepted value for the parameter under test and the *Alternative* (H_1) hypothesis as the unacceptable value for such parameter. Let \mathbb{N}^+ be the set of positive integers. SPRT takes a sample of pre-determined size $m \in \mathbb{N}^+$ and, based on the result obtained from drawing this fixed-size sample, it makes a decision regarding these two hypotheses as follows.

Let X_t be a random variable representing the number of observations at time t . With the classical SPRT, tests are performed continuously at every time t as additional data are collected. The test statistic is the likelihood or log-likelihood ratio defined as

$$\frac{p_{1m}}{p_{0m}} = \prod_{t=1}^m \frac{P(X_t|H_1)}{P(X_t|H_0)} \quad \text{or} \quad \log\left(\frac{p_{1m}}{p_{0m}}\right) = \sum_{t=1}^m \frac{P(X_t|H_1)}{P(X_t|H_0)}. \quad (1)$$

This test statistic is sequentially monitored for all values of $t > 0$, until either H_0 is rejected by $\frac{p_{1m}}{p_{0m}} \geq (1 - \beta)/\alpha$ or until H_1 is rejected by $\frac{p_{1m}}{p_{0m}} \leq \beta/(1 - \alpha)$, where α, β are type I and type II errors. With this stopping rule, the Null hypothesis is falsely rejected with probability α when it is true, while the Alternative hypothesis is falsely rejected with probability β when it is true.

SPRT only needs a minimum number of external event observations to reach decisions with bounded probability of false positives and negatives [14]. This allows SPRT to be applied to black-box systems and make fast yet accurate decisions. Moreover, SPRT is computationally inexpensive. All of these features make it particularly suitable for the femtocell coverage optimisation problem.

3 Coverage Optimisation in Femtocells

Fig. 1 shows a simple femtocell network with four femtocells jointly providing mobile signal coverage in an office area. Each femtocell uses home broadband access as a backhaul to the mobile operator's core network [5]. For a femtocell, the coverage radius \mathcal{R} indicates its radio coverage area and the change of signal power δ is used to adjust the radius. When a user equipment (UE), such as a mobile device, enters the femtocell network, it performs a handover to a femtocell,




Fig. 1. Illustration of the Femtocell Coverage Optimisation Problem

which is then used to transmit data. This results in improved mobile service and traffic offloading from the core network. However, there may be coverage gaps and coverage overlap between femtocells. When a UE is not covered by any femtocell (coverage gap), it has to handover to the core network again for its data transmission. When a UE is covered by more than one femtocell (coverage overlap), interference increases and this results in energy being wasted. Thus both coverage overlaps and gaps are undesired behaviour and should be avoided.

Each femtocell should serve a similar number of UEs to achieve load balancing. Therefore, the aim of femtocell coverage optimisation is to minimise coverage overlap and coverage gap, and also to balance the workload among femtocells. Solving this problem involves adjusting the signal power in order to control the coverage area of femtocells and deciding when such a signal power adjustment should be triggered. In this paper, we mainly address the latter problem by adopting statistical verification. The following is a list of desired optimisation goals that we consider:

- Minimising the femtocell coverage gaps D . Coverage gaps can be inferred from the number of handovers between femtocells and the core network, which represents the number of times a UE enters a coverage gap;
- Balancing the load L among femtocells in a group to prevent overload or under-utilisation;
- Minimising the coverage overlaps O by reducing the signal power.

Here we give a simple mathematical formulation of this problem. Let D_{thr} , O_{thr} , L_{thr} be the required upper bound for femtocell coverage gap, overlap and load, respectively. Let $\mathbb{F}(D; O; L; D_{thr}; O_{thr}; L_{thr}) \xrightarrow{\delta} \mathcal{R}$ be a function (in short as $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$), of which the input parameters are the current coverage performance observations (D , O , L) and the expected coverage perfor-

mance $(D_{thr}, O_{thr}, L_{thr})$. Then the output parameter is a suitable power adjustment δ , which enables a femtocell to achieve the expected \mathcal{R} that satisfies $(D_{thr}, O_{thr}, L_{thr})$. Such a signal power re-computation decision is made independently by each individual femtocell.

4 Statistical Verification for Femtocell Networks

To achieve the optimisation objectives given in Section 3, verification goals are used to check the probabilistic satisfaction of coverage optimisation. These verification goals can be represented as *eventually* $\forall n \in \mathbf{Femtocells} \models (D < D_{thr}) \wedge (L < L_{thr}) \wedge (O < O_{thr})$ where $n \in \mathbb{N}^+$ is the number of femtocells in the network. These properties are typically expressed in temporal logics. We adopt the continuous stochastic logic (CSL) [2] as a formalism for expressing quantitative properties. This is because CSL is semantically capable of specifying property such as “with probability of at most 0.1” ($\mathbb{P}_{\leq 0.1}(\Phi)$). For the femtocell coverage optimisation problem, the desired properties are formulated as Equations (2)–(4):

$$\phi := \mathbb{P}_{\geq \theta_D}(true \mathbb{U}^{\leq t}(D < D_{thr})), \quad (2)$$

$$\varphi := \mathbb{P}_{\geq \theta_L}(true \mathbb{U}^{\leq t}(L < L_{thr})), \quad (3)$$

$$\psi := \mathbb{P}_{\geq \theta_O}(true \mathbb{U}^{\leq t}(O < O_{thr})), \quad (4)$$

where \mathbb{U}^t is the bounded until path formula and t is a specified time. We regard t as t_{now} to indicate that a property should be satisfied when it is evaluated. θ_D , θ_L and θ_O represent the desired level of confidence of each property. For example, Equation (2) represents the probability of D being less than D_{thr} with the confidence of no less than the required level θ_D until now. Similar meaning applies to Equations (3)–(4) for O and L . If any of these properties are violated, then $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$ should be triggered.

4.1 Modelling Overlap, Load and Gap Violations as Stochastic Discrete Event Processes

For checking probabilistic satisfaction of given properties, such as Equations (2)–(4), statistical verifiers need to be evaluated against the trace of occurred coverage overlap, coverage gap and overload events. Here we give a formal description of how the trace of these events is modelled.

Let Ω be a sample space. \mathcal{F} is a collection of subsets of Ω . Let $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ be a *probability measure* defined on \mathcal{F} . Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A stochastic process is a collection of random variables $X = \{X_t : t \in \mathbf{T}\}$ with index set \mathbf{T} .

For the property of a femtocell, there are two types of states, in which a property is either violated (*true*) or not (*false*). Thus it can be represented as a Bernoulli process, where each random variable X is a Bernoulli trial, mapping $\Omega \rightarrow \{true, false\}$, such that $X = true$ with probability p and $X = false$ with

probability $1 - p$ ($p \in \mathbb{P}$). The probability distribution of a sum of n Bernoulli trials with parameter p can be represented as a binomial distribution with parameter n and p as follows:

$$B(n, p) = \sum_{j=0}^n \binom{n}{j} p^j (1 - p)^{n-j}. \quad (5)$$

For femtocell coverage overlap O , load L and gap D , let X_O , X_L and X_D be three random variables. X_O represents the occurrence of an overlap event, where $X_O = false$ when no overlap event occurred, $X_O = true$ when an overlap event occurred. The maximum allowed probability $\mathbb{P}(X_O = true)$ of such a Bernoulli process is O_{thr} . Similarly, X_L and X_D represent the occurrence of an overload or gap event, respectively.

In this way, we model system states of a femtocell as a stochastic discrete event system (denoted as \mathcal{M}), which is similar to the work introduced in [15]. The state space of such a system can be represented as sequential observations at times t_1, t_2, \dots , for example, a sequence of overlap events. The state change is discrete rather than continuous and is triggered when these events are detected. The verification procedure is to compute whether the occurrence probability of these events is above the predefined threshold. Thus the verification of each desired property shown in Equations (2)–(4) can be achieved by using the verification method shown in Equation (1).

It is worth mentioning that, in the real world, X_O , X_L and X_D are not completely independent. When reducing a femtocell's coverage area, overlap and load are reduced, but the possibility of gaps happening is increased and vice-versa. However, when triggering self-optimisation using statistical verification, the decision is made independently according to the observation of each property. Thus we consider X_O , X_L and X_D as independent observations.

4.2 Femtocell Coverage Verification Procedure

Based on the stochastic discrete event models from the previous section, we describe in detail how to detect violations using statistical verifiers.

For a femtocell, let trigger $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$ at an arbitrary time t be a binary decision $\mathcal{D}_t \in \{true, false\}$, $t \in \{1, 2, \dots\}$, where *true* represents triggering this function and *false* represents the opposite decision. Let V_1, V_2, \dots, V_N ($N \in \mathbb{N}^+$) be verifiers. Let P_1, P_2, \dots, P_N be properties, which are used to describe desired behaviours of a femtocell network. Each verifier V_i monitors a given property P_i ($i \in N$). At time t , a verifier V_i observes a random variable X_t according to P_i . Then, let H_0^i be the Null hypothesis of the observed property P_i that is violated and H_1^i be the Alternative hypothesis.

During the verification procedure, if H_0^i holds, then $D_t = true$, otherwise if H_1^i holds, $D_t = false$. Based on the verification rule (see Equation (1)), we have

Algorithm 1 Statistical verification based coverage optimisation algorithm

- 1: **List of verification goals in Equations (2)–(4):** $\mathcal{M} \models \phi, \varphi, \psi$;
 - 2: Input threshold θ for coverage gap, overlap and load $D_{thr}; O_{thr}; L_{thr}$;
 - 3: **Start statistical verifiers V_i for observing each given property**
 - 4: For each property $\in \{\phi, \varphi, \psi\}$, when an observation X_t is received
 - 5: **Verify each property against Equation (6)**
 - 6: **if** a given property is violated ($\mathcal{M} \not\models \{\phi, \varphi, \psi\}$) **then**
 - 7: Trigger the self-optimisation algorithm ($\mathbb{F} \xrightarrow{\delta} \mathcal{R}$).
 - 8: Output optimised parameters δ
 - 9: Clean observations X_t s for the verified property
 - 10: **end if**
 - 11: Continue the statistical property monitoring
-

the following forms:

$$\begin{aligned} &\text{accept } H_0^i \text{ if } p_t^i \leq \log\left(\frac{\beta}{1-\alpha}\right); \\ &\text{accept } H_1^i \text{ if } p_t^i \geq \log\left(\frac{1-\beta}{\alpha}\right); \\ &\text{continue if } \log\left(\frac{\beta}{1-\alpha}\right) \leq p_t^i \leq \log\left(\frac{1-\beta}{\alpha}\right), \end{aligned} \tag{6}$$

where p_t^i is the $\log\left(\frac{p_{im}}{p_{om}}\right)$ in Equation (1), α and β are the type I and type II errors respectively ($\alpha, \beta = 1 - \theta$), as introduced in Section 2.2.

With the verification procedure from Equation (6), each femtocell in the network has one or more verifiers and makes the decision of triggering $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$. For example, $H_0^D : \mathbb{P}(X_D = true) \geq D_{thr}$ represents the hypothesis that the occurrence probability of coverage gap exceeds the given threshold ($\mathcal{M} \not\models \phi$ or $\mathcal{M} \models \neg\phi$) and $H_1^D : \mathbb{P}(X_D = true) \leq D_{thr}$ represents the alternative hypothesis ($\mathcal{M} \models \phi$). When the H_0^D is accepted, a violation is detected, so the decision \mathcal{D}_t should be set to be *true* in order to trigger $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$.

Algorithm 1 gives the pseudo code for the statistical verification-based femtocell coverage optimisation procedure. Lines 1–2 specify properties as verification goals with a given confidence level θ . In lines 3–4, verifiers are initiated V_1, \dots, V_N for each given property. Lines 5–10 specify that if a violation is detected, $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$ is triggered. Here we keep $\mathbb{F} \xrightarrow{\delta} \mathcal{R}$ as a high-level abstraction, which can be substituted by any algorithm designed for this type of problem. After the optimisation procedure finishes, the previous observation history is cleared to refresh the verifiers. Eventually the femtocell network should become stable and achieve $\{\phi, \varphi, \psi\}$.

5 Evaluation

The goal of our evaluation is to examine the effectiveness of our statistical verification approach in terms of reducing the probability of coverage overlaps and

gaps, and achieving optimisation quickly as well as avoiding unnecessary re-optimisation. We investigate the impact of the proposed *statistical verification-based approach* (noted as *SVA*). The results obtained are compared with the performance of the *fixed-time interval approach* (noted as *FIA*). The underlying power adjustment algorithms are considered to be the same.

We implement a simulator for a four-femtocell network deployed within a plane office area as introduced shown in Fig. 1. The statistical verification module is implemented as part of Matlab 2009b. The occurred coverage overlap or gap events are collected by each femtocell and passed to its statistical verifier as input. For each statistical verifier, Equation (6) is implemented as the main function where its likelihood function follows Equation (1) and the probability distribution of the input data follows Equation (5). The violation of desired coverage overlap or gap threshold is used as the Null hypothesis in Equation (6).

For the initial deployment, we assume that the femtocells are already placed through automatic planning [10]. We ignore the required property of workload L_{thr} for simplicity. We implement UEs, which randomly appear in the femtocell network, to simulate user behaviour. We set the number of UEs to be 24, which has the same ratio of UEs and femtocells as in [9], where there are 50 femtocells and 300 UEs. We assume that the average time of a UE at a given location within the femtocell network is 90 seconds, chosen from an exponential distribution.

It is necessary to make assumptions about the spatial location of UEs. A normal distribution is a good choice due to the Central Limit Theorem—sufficiently many location samples from UEs can be approximately normal distributed. A uniform distribution is also a good model for the behaviour in an office environment because UEs are often equally distributed around the office area. We assume that UEs move within a square area around the centre of femtocells. Such a setting results in a higher overlap and gap possibilities than that of UEs moving across all possible locations in a femtocell network. Here we apply the same settings as in [9], setting $D_{thr} = 0.1$ and $O_{thr} = 0.3$. We adopt $\theta_D = \theta_O = 0.95$ as the verification confidence level.

5.1 Evaluation Results

The evaluation experiments were run on a MacOS 10.5.8 machine with an Intel 2.53 GHz core 2 CPU with 4 GB of RAM. We simulate 24 hours in the lifetime of a four-femtocell network.

Fig. 2 shows the impact of the occurrence probability of coverage overlap P_O and gap P_D under the SVA and FIA schemes. The results show that the SVA achieves a better satisfaction by lowering P_O and P_D . The SVA drastically reduces P_O by 25% and P_D by 38% under UEs following a normal distribution (Fig. 2(a)) compared to FIA. A similar 26% improvement in reducing the P_D for UEs with uniformly distributed movements is shown Fig. 2(b). However, SVA does not improve P_O and the overlap is more likely to be increased by about 24% compared to FIA. Although P_O tends to increase, the overlap is still lower than a threshold $O_{thr} \leq 0.3$. Furthermore P_D is lower than P_D of FIA.




Fig. 2. Overlap and Gap Probabilities for the Femtocell Network




Fig. 3. Average Number of Overlap and Gap Events per Femtocell

Analogous to the occurrence probability, the number of occurred overlap events X_O and gap events X_D can be also observed under two UE movement schemes. Fig. 3 shows that SVA successfully reduces the number of occurred overlap events and gap events in comparison to FIA. For UEs with normally distributed movements, Fig. 3(a) shows that SVA lowers the number of generated overlap and gap events. The number of gap events is also reduced under uniformly distributed movements (Fig. 3(b)). This is in agreement with the result from Fig. 2. Both results demonstrate that SVA achieves a better optimisation than FIA.

We also record the number of triggered power adjustments ($\mathbb{F} \xrightarrow{\delta} \mathcal{R}$). For FIA, within each fixed time interval, there is always one power adjustment if the desired level of overlap or gap probability has not been achieved. Fig. 4 shows the number of power adjustments per half hour using SVA. Note that for both normal and uniform distributions of UE movements, SVA detects violation initially rapidly and avoids unnecessary self-optimisations as whole network approaches a stable optimal state.




Fig. 4. Number of Self-Optimisations per 0.5 hour

5.2 Discussion

A limitation of our statistical verification approach is that, if a desired property is infeasible to achieve and violations occur very often, power adjustments are triggered frequently without being able to converge to the optimisation goal. To avoid such a problem, a limit on the number of self-optimisations during a given time interval can be added to SVA. This might slow down the convergence speed of the optimisation but it prevents the system from becoming unstable.

Moreover, different office environments might require different coverage overlap, gap and verification confidence values. A smaller overlap or gap threshold makes it harder for the femtocell network to achieve its self-optimisation goal. Higher verification confidence means that a verifier requires more evidence (i.e., overlap or gap events) to detect violations. It would be interesting to explore the impact of this parameter and we plan to do this as part of future work.

6 Conclusions and Future Work

In this paper, we have shown how to use statistical verification to manage violations of optimisation goals in order to provide joint coverage in femtocell networks. Our proposed approach is based on a sequential likelihood ratio hypothesis test. It improves selection time to trigger the signal power adjustment procedure of femtocells, and in turn improves the satisfaction of desired coverage performance. Our evaluation results show that our approach performs better than a fixed time interval-based technique, resulting in a better satisfaction of the desired upper bound of overlap and gap occurrence probabilities. Our approach can help optimise a femtocell network faster and avoid unnecessary signal power adjustments when violations are rare. All these features are necessary for guaranteeing fast, accurate and autonomous femtocell coverage self-optimisation.

For future work, we will consider scenarios with larger enterprise femtocell networks and changes in femtocell locations in order to evaluate the scalability, convergence speed and stability of statistical verification. In addition, we want to exploit peak and off-peak usage patterns in femtocell networks, leading to more realistic workloads of daily and weekly patterns of mobile devices.

7 Acknowledgements

The work reported in this paper has formed part of the Flexible Networks area of the Core 5 Research Programme of the Virtual Centre of Excellence in Mobile & Personal Communications, Mobile VCE (www.mobilevce.com), and has been jointly funded by Mobile VCE's industrial member companies and the UK Government, via the Engineering and Physical Sciences Research Council.

References

1. Femtocell Market Status. <http://www.femtoforum.org/femto/Files/File/InformaFemtocellMarketStatusQ12010F.pdf> (2010)
2. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.: Verifying Continuous-time Markov Chains. *Lecture Notes in Computer Science* 1102, 269–276 (1996)
3. Chandrasekhar, V., Andrews, J., Gatherer, A.: Femtocell Networks: A Survey. *IEEE Communications Magazine* 46(9), 59–67 (2008)
4. Choi, S., Lee, T., Chung, M., Choo, H.: Adaptive Coverage Adjustment for Femtocell Management in a Residential Scenario. *Management Enabling the Future Internet for Changing Business and New Computing Services* pp. 221–230 (2009)
5. Claussen, H., Ho, L., Samuel, L.: An Overview of the Femtocell Concept. *Bell Labs Technical Journal* 13(1), 221 (2008)
6. Claussen, H., Ho, L., Samuel, L.: Self-Optimization of Coverage for Femtocell Deployments. In: *Proc. Wireless Telecom. Symp.* pp. 278–285
7. Claussen, H., Pivitt, F.: Femtocell Coverage Optimization Using Switched Multi-Element Antennas. In: *IEEE Int. Conf. on Comm.* pp. 1–6 (2009)
8. Duran, A., Carrasco, G.: UMTS Femtocell Performance in Massive Deployments: Capacity and GoS Implications. *Bell Labs Technical Journal* 14(2), 185–202 (2009)
9. Ho, L., Ashraf, I., Claussen, H.: Evolving Femtocell Coverage Optimization Algorithms Using Genetic Programming. In: *20th IEEE In. Sym. on Personal, Indoor and Mobile Radio Comm.* pp. 2132–2136. IEEE (2010)
10. Molina, A., Athanasiadou, G., Nix, A.: The Automatic Location of Base-stations for Optimised Cellular Coverage: A New Combinatorial Approach. In: *Vehicular Technology Conference, 1999 IEEE 49th. vol. 1*, pp. 606–610. IEEE (2002)
11. Sammapun, U., Lee, I., Sokolsky, O.: RT-MaC: Runtime Monitoring and Checking of Quantitative and Probabilistic Properties. In: *11th IEEE Int. Conf. on Embedded and Real-Time Comp. Sys. and Apps.* pp. 147–153 (2005)
12. Sen, K., Viswanathan, M., Agha, G.: On Statistical Model Checking of Stochastic Systems. In: *Computer Aided Verification.* pp. 266–280 (2005)
13. Sokolsky, O., Sammapun, U., Regehr, J., Lee, I.: Runtime Verification for Wireless Sensor Network Applications. In: *Runtime Verification* (2008)
14. Wald, A.: Sequential Tests of Statistical Hypotheses. *The Annals of Mathematical Statistics* 16(2), 117–186 (1945)
15. Younes, H., Simmons, R.: Statistical Probabilistic Model Checking with a Focus on Time-Bounded Properties. *Information and Computation* 204(9), 1368 (2006)