Cell Switch-Off Algorithms for Spatially Irregular Base Station Deployments

Quoc-Nam Le-The, Tamer Beitelmal, *Member, IEEE*, Faraj Lagum, *Student Member, IEEE*, Sebastian S. Szyszkowicz, *Member, IEEE*, and Halim Yanikomeroglu, *Fellow, IEEE*

***Abstract*—In cell switch-off (CSO) research, the base station (BS) deployment is usually assumed to follow a regular grid. Some recent works considered instead a fully random (Poisson point process) BS deployment. Given that the best signal-to-interference ratio distribution can be achieved when BSs are located on a hexagonal layout, we study applying CSO to irregular network layouts with the objective of making the active BS locations as regular as possible, regardless of the irregularity of the original network layout. In this letter, we test the suitability of several CSO algorithms for this new objective; we also introduce a novel algorithm which performs very well when the number of BSs to switch off is high.**

***Index Terms*—Cell switch-off, green cellular networks, *p*-dispersion problem, stochastic geometry.**

I. Introduction

Cell switch-off (CSO) is an emerging approach for more energy-efficient cellular networks. CSO aims at reducing the number of active base stations (BSs) during periods of low demand while maintaining good coverage and user satisfaction. In most of the CSO literature, BS deployment is assumed to be either perfectly regular (i.e., BSs are placed on a triangular lattice (TL)) [1], [2] or completely random (i.e., BSs are placed according to a Poisson point process (PPP)) [1], [3], [4]. Due to geographic restrictions of site placement and the network planning, the actual BS placement falls somewhere in between [5], [6]. In other words, BS locations are most likely to have a repulsive structure. While a TL deployment produces an upper bound on the system performance, a PPP deployment results in a lower bound [7]. For BSs deployed with a certain amount of regularity, the performance of the network improves according to the increase in the spatial regularity of BS locations.

Keeping the active BSs as regular as possible is advantageous in terms of downlink SIR performance, reducing coverage holes, as well as optimizing the uplink [8], [9]. The fact that the best signal-to-interference ratio (SIR) distribution is achieved when BSs are located on a TL [6], [10], [11] presents a great opportunity for applying CSO to spatially irregular network layouts with the objective of making the active BS locations as regular as possible. In this letter, we consider interference-limited cellular networks, with uniform user distribution.

Contribution: This letter is one of the very few papers on CSO for irregularly-placed BSs, which brings this energy-saving approach (CSO) much closer to being practical and applied by cellular operators. In order to find algorithms for this problem, we relied on the p-dispersion problem (PDP), a well-known mathematical problem with the objective of activating p points such that the minimum distance between any two active points is maximized [12]. In this letter, we compare several practical PDP algorithms for CSO: most of these algorithms are presented for the first time in the CSO context, three of them were used in our earlier work [11]. We propose an entirely new CSO algorithm, designed for high switch-off percentages, with a demonstrated good performance against the upper bound. We recommend a rule-set for selecting the best algorithm in various operating regions. We demonstrate that the choice of the best algorithm is primarily influenced by the percentage of active BSs, but also, to a lesser extent, by the spatial regularity of the network BSs.

II. Problem Formulation

Consider a network with N BSs. In low traffic periods, only a fraction ρ of these BSs is required to support the users' demand, i.e., $(1-\rho)\times\ 100\%$ being the *CSO percentage*. Thus, a different set of ρN active BSs may be defined for different traffic load levels. The problem we are solving is to switch off a given percentage of BSs with the objective of maximizing the regularity of the remaining active BSs, so that the downlink SIR is maximized.

In order to compare the performance of the algorithms, we need metrics to quantify the regularity both of the initial BSs and the active BSs. We also need a point process model for placing BSs with a variable amount of spatial regularity.

A. Metrics

1) Geometry-Based Metric of Spatial Regularity: This metric was proposed in [6] and [13], and is used to measure the spatial regularity of the BS locations. It is the normalized coefficient of variation of the Delaunay triangulation edge lengths of a set of points. It is called C_{D}, and can be defined as

$$C_{\mathrm{D}} = \sigma_D/(k_D\mu_D), \tag{1}$$

where μ_D and σ_D are the mean and the standard deviation of the Delaunay edge lengths, respectively; and $k_D \cong 0.492$

Manuscript received February 13, 2017; accepted March 26, 2017. Date of publication April 4, 2017; date of current version June 19, 2017. This work was supported in part by the TELUS Corporation, in part by the NSERC USRA Scholarship, and in part by the Ministry of Higher Education and Scientific Research, Libya, through the Libyan-North American Scholarship Program. The associate editor coordinating the review of this paper and approving it for publication was I. Krikidis. *(Corresponding author: Tamer Beitelmal.)*

The authors are with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada (e-mail: quocnamlethe@cmail.carleton.ca; tamer@sce.carleton.ca; faraj.lagum@sce.carleton.ca; sz@sce.carleton.ca; halim@sce.carleton.ca).

Digital Object Identifier 10.1109/LWC.2017.2690677

TABLE I
RELATIONSHIP BETWEEN PERTURBATION RADIUS $\tilde{R}$ AND REGULARITY METRIC C_D

C_D	0.1	0.2	0.3	0.4	0.5
$\tilde{R}$	0.059	0.129	0.202	0.278	0.360
C_D	0.6	0.7	0.8	0.9	≈ 1
$\tilde{R}$	0.450	0.562	0.729	1.063	≥ 2

is the normalization factor such that for a PPP $C_D = 1$ on average [6]. C_D is thus a measure of spatial irregularity of points, where $C_D = 0$ is obtained for a TL and is said to be perfectly regular, and $C_D = 1$ is said to be completely spatial random [14], [15].

2) SIR-Based Performance Metric: The performance metric used in this letter for comparing CSO algorithms is the downlink SIR gain $G_{SIR}(X\%)$ in dB, which is the difference between the SIR for the best $X\%$ of users when the active BSs are selected according to a given algorithm, and the SIR in a network with the same number of BSs but deployed randomly (according to a PPP) [10]. In this letter, we measure the SIR gain (G_{SIR}) for 50% and 95% of users.

B. Perturbed Triangular Lattice for BS Locations Modelling

In cellular networks, the actual layout of BS locations lies somewhere in between the perfect TL and the PPP [6], [7], [10], [16], [17]. In [14], the perturbed triangular lattice (PTL) point process is used to model BS locations, which can sweep the whole range of regularity. We begin by generating a TL, where BSs are located on a hexagonal layout with inter-site distance η. This layout is maximally regular ($C_D = 0$). The PTL is obtained by displacing each point by an independent random vector. In the uniform variant of the PTL, which we use here, the displacement vector has a uniform distribution over a disk of radius R. The most interesting parameter is the normalized radius $\tilde{R} = R/\eta$, which controls the amount of regularity. The C_D metric is then a bijective function of $\tilde{R}$ with values shown in Table I, which are retrieved from [14].

III. ALGORITHMS

Being NP-hard [12], the PDP requires heuristic algorithms when the number of points N is large. Ten heuristic algorithms are described and compared in [12] to solve the PDP. In this letter, we investigate all of these algorithms (except the projection algorithm, which performs poorly [12]) to solve the CSO problem with the aim of maximizing the SIR by switching off some BSs to improve the regularity of the remaining active BSs. We also introduce a novel algorithm that performs best when the CSO percentage is high.

A. Proposed Algorithm: Triangular Lattice Fit (TLF)

Current research [6], [7], [10] suggests that the best SIR can be achieved with BSs placed in a TL; therefore we propose the TLF as a novel algorithm for the CSO problem at hand. For a given CSO percentage, the TLF algorithm starts by creating a regular TL of grid points with the required number of active BSs. This grid is rotated and shifted to best fit the real BS locations, by minimizing the sum of the distances from each grid point to its nearest BS. Once an optimal grid position is found, the nearest BS to each grid point is chosen to be active, and all other BSs are switched off.

B. Greedy Algorithms

These algorithms start with a set of BSs and then iteratively select one BS to add to or remove from the solution set. The two common varieties are *greedy construction* (GC) and *greedy deletion* (GD). GC initializes the solution set with the two furthest BSs and iteratively activates the BS that maximizes the minimum distance to other BSs already in the solution set. GD initializes with all BSs active and finds the nearest two active BSs then switches off the BS that has the shortest distance to its second nearest neighbour, doing so iteratively. The *semi-greedy deletion* (SG) is similar to the GD but the selection among the two nearest BSs being random. The three greedy algorithms were implemented in [11] as CSO algorithms. While GC and GD result in good performance, SG was found to perform poorly, and hence we exclude it from this discussion.

C. Neighbourhood Algorithms

The neighbourhood algorithms are similar in construction to the Matern Hard-Core Process of Type I (MHC-I) [6], where we can think of radius r as the hardcore radius. They start with one BS in the solution set and iteratively append another BS at each step [12]. However, they have a restriction that any added BS must be outside of the *neighbourhood* of all BSs in the solution set. A neighbourhood is defined as a circle with radius r centered at the BS. Thus, if a BS is added to the solution set, all BSs in its neighbourhood are eliminated. The objective of the neighbourhood algorithms is to generate a solution with a guaranteed BS-separation distance of radius r. The value of r needs to be chosen carefully to obtain the desired number of BSs. In our simulations, we set the initial value of r to equal the maximum distance between any two BSs, and then conduct a binary search to obtain a number of active BSs that is close to ρN. The algorithms terminate when there is no more BSs to add, and if the number of selected BSs is not exactly ρN, then we need to adjust the solution by running a GC or GD algorithm.

Each neighbourhood algorithm has a different criterion for selecting the new BS to add to the solution set. The three neighbourhood algorithms are *the first point outside the neighbourhood (FP)* (where the next BS is chosen arbitrarily), *the closest point outside the neighbourhood (CS)* (where the next BS is the one with the smallest sum distance to all BSs in the solution set), and *the furthest point outside the neighbourhood (FS)* (where the next BS is the one with the largest sum distance to all BSs in the solution set).

D. Interchange Algorithms

The interchange algorithms initialize with a random solution set that contains the desired number ρN of active BSs. Then BSs from the solution set might be repeatedly interchanged with BSs from outside the set (inactive BSs). The algorithm

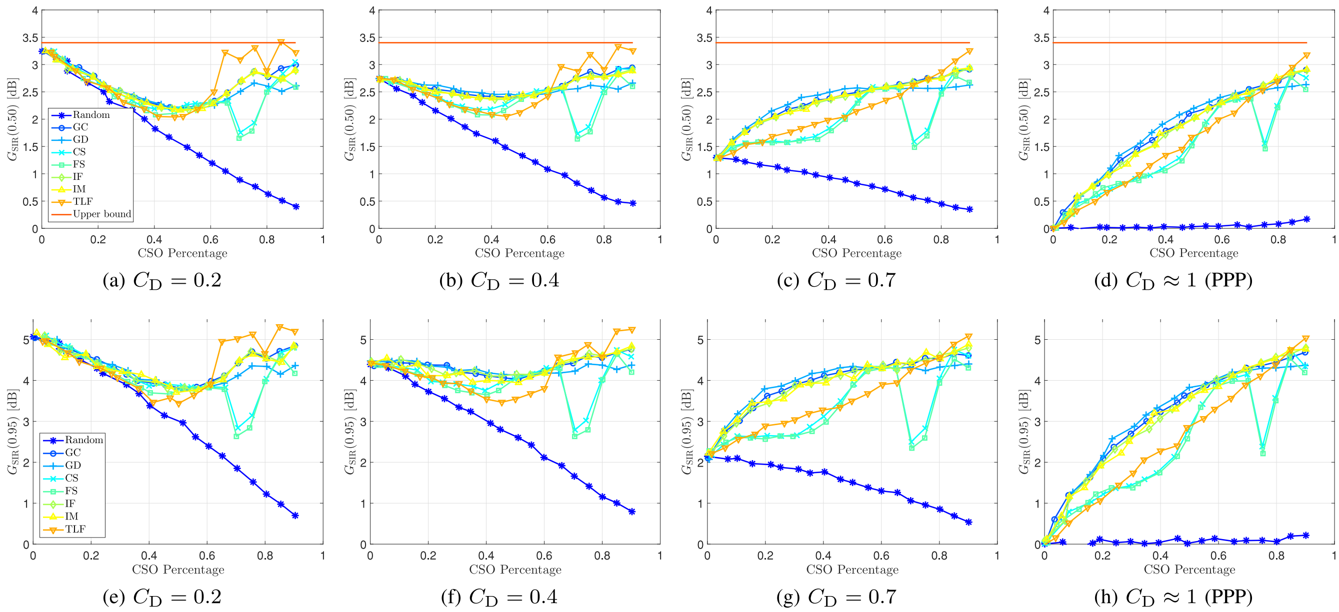

Fig. 1. The SIR gain G_{SIR} for 50% and 95% users as a function of the CSO percentage $(1 - \rho)$, for different deployment regularity (C_{D}) values.

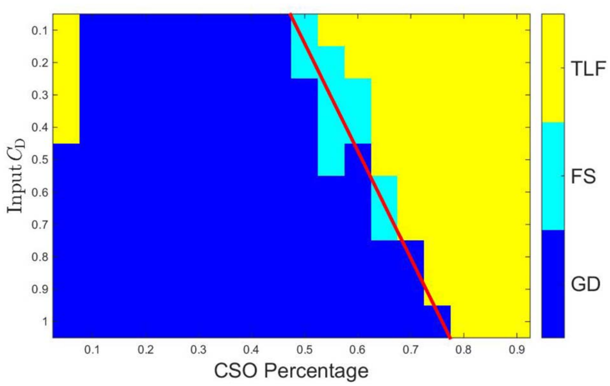

Fig. 2. Map of the best performing algorithms. The algorithm selection can be approximately based on the location with respect to the red line.

terminates when no more interchanges result in an increase in the regularity metric. The three interchange algorithms are *the first pairwise interchange* (IF), *the best pairwise interchange* (IM), and *simulated annealing* (SA). Whereas IF interchanges an active BS with the first inactive BS that would increase the minimum distance, IM interchanges with the inactive BS that would increase the minimum distance the most. A disadvantage of both IF and IM is that the solution could get stuck at a local maximum. SA was used in [12] to escape the local maximum by allowing occasional random (with a probability that decreases with time) interchanges that decrease the performance metric. However, we did not implement the SA here due to its prohibitive computational time, its several parameters that were complicated to tune correctly, and the high variability (instability) of its output solutions.

IV. Performance

A. Simulation Setup

First, we place 600 BSs in a square area according to a PTL and vary the spatial regularity from TL to PPP to obtain different amounts of spatial regularity C_{D}, by tuning $\tilde{R}$ according to Table I. All BSs are assumed to have the same transmit power, operating frequency, and are equipped with omnidirectional antennas. The performance is measured in terms of downlink SIR gain. Users are uniformly distributed inside the central area of the network. The received power is calculated assuming a pathloss exponent of $\alpha = 4$ with no shadowing and all of the links experience independent Rayleigh fading. Users are connected to the BS with the strongest received signal. After computing the SIR values of all users, the SIR gain G_{SIR} is then calculated for 50% and 95% of users. We compare the algorithms at different values of initial network regularity C_{D}.

We vary the CSO percentage from 5% to 90%. For each percentage, we found the number ρN of active BSs and place them according to a PPP. We consider the PPP as the baseline model, which is known to give the worst SIR performance for repulsive networks [11]. The algorithms' performances are compared in terms of the SIR gain G_{SIR} (with respect to a PPP deployment). For the same number of BSs, the best SIR is achieved with a TL placement, which is known to have a G_{SIR} of 3.4 dB (compared to the PPP) [18]; this value is considered as the upper bound on the algorithm performance [11]. As a lower bound, we switch off BSs randomly: any reasonable algorithm should at least perform better than this.

B. Results

The performances of the algorithms are illustrated in Fig. 1 for selected C_{D} values, and are compared in Table II. We show the SIR gains $G_{\mathrm{SIR}}(0.50)$ and $G_{\mathrm{SIR}}(0.95)$, for 50% and 95% of users, respectively. For both SIR percentages, the ranking of the algorithms is similar.[1]

[1]The non-monotonic behaviour of the CS and FS algorithms is analogous to the non-monotonic behvaiour of the MHC-I as a function of r [6, Fig. 3].

TABLE II
SUMMARY OF ALGORITHMS

Initial	Name	Performance
Novel Algorithm [proposed in this paper]		
TLF	Triangular lattice fit	Performs the best in high CSO percentages; it is not so good for low CSO percentages.
Greedy Algorithms [11], [12]		
GC	Greedy construction	Performs well overall, but never the best.
GD	Greedy deletion	Performs well overall. It is the best for low CSO percentages ($< 55\%$).
SG	Semi-greedy deletion	Performs poorly compared to others.
Neighbourhood Algorithms [12]		
FP	First point outside	Not shown in Fig. 1 as it performs inferior to others.
CS	Closest point outside	Performs fairly well in high CSO percentages ($> 80\%$).
FS	Furthest point outside	Performs well overall, and it is the best neighbourhood algorithm.
Interchange Algorithms [12]		
IM	Best pairwise	Performs well in most regions.
IF	First pairwise	Performs well overall, but never the best. Also, it is faster than IM.
SA	Simulated annealing	Has prohibitive computational time, several parameters to tune, and high variance (instability) of performance.

Based on both CSO percentages and C_D values, we can choose the best performing algorithm as illustrated in Fig. 2. The three dominating algorithms are GD, FS and TLF. Although the main factor is the CSO percentage, C_D also has a moderate effect on the choice of algorithm.

For low CSO percentages ($< 50\%$), GD is the best algorithm for any input C_D value. For very high CSO percentages ($> 75\%$), the TLF is the best algorithm. For CSO percentages between 50% and 75%, the best performing algorithm is chosen based on the following formula:

$$C_D \gtrless 1.875 - 3.5\rho, \; 0.50 < (1-\rho) < 0.75, \qquad (2)$$

where ρ is the fraction of active BSs, i.e., $= 1-$(CSO Percentage), and C_D can be estimated for a given point set using (1). Based on (2) and Fig. 2, we have three regions, each has a preferred algorithm:

1) above the line ($C_D + 3.5\rho < 1.875$) : TLF,
2) below the line ($C_D + 3.5\rho > 1.875$) : GD,
3) on the line ($C_D + 3.5\rho \approx 1.875$) : FS.

Our proposed TLF algorithm outperforms all others for very high CSO percentages (within 1 dB SIR of the upper bound). Indeed, as ρN decreases, the probability of finding a subset of the BSs that almost matches a TL increases. Also, algorithms GC and IF are nowhere optimal, but overall they are good choices.

V. CONCLUSION

In this letter, we compared several algorithms for solving the recently proposed problem of CSO that considers spatially irregular BS deployments. While the previously proposed greedy algorithms GC and GD were found to be good choices overall, we found that all existing algorithms perform suboptimally for high CSO percentages. We thus proposed TLF as a novel algorithm which performs well (better than others) in very high CSO percentages. For low CSO percentages, we found that all algorithms - although very different in design-perform quite similarly.

The choice of algorithm is not only influenced by the desired CSO percentage, but also, to a lesser extend, by the spatial regularity of the entire network. If implementing only one algorithm is preferable, then the two algorithms GC and IF -although never the best- offer a good overall performance.

REFERENCES

[1] B. Rengarajan, G. Rizzo, and M. A. Marsan, "Energy-optimal base station density in cellular access networks with sleep modes," *Comput. Netw.*, vol. 78, pp. 152–163, Feb. 2015.

[2] D. G. González, H. Yanikomeroglu, M. Garcia-Lozano, and S. R. Boque, "A novel multiobjective framework for cell switch-off in dense cellular networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Sydney, NSW, Australia, Jun. 2014, pp. 2641–2647.

[3] E. Altman *et al.*, "Stochastic geometric models for green networking," *IEEE Access*, vol. 3, pp. 2465–2474, Nov. 2015.

[4] J. Peng, P. Hong, and K. Xue, "Stochastic analysis of optimal base station energy saving in cellular networks with sleep mode," *IEEE Commun. Lett.*, vol. 18, no. 4, pp. 612–615, Apr. 2014.

[5] J. G. Andrews *et al.*, "What will 5G be?" *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.

[6] F. Lagum, S. S. Szyszkowicz, and H. Yanikomeroglu, "CoV-based metrics for quantifying the regularity of hard-core point processes for modeling base station locations," *IEEE Wireless Commun. Lett.*, vol. 5, no. 3, pp. 276–279, Jun. 2016.

[7] J. G. Andrews, F. Baccelli, and R. K. Ganti, "A tractable approach to coverage and rate in cellular networks," *IEEE Trans. Commun.*, vol. 59, no. 11, pp. 3122–3134, Nov. 2011.

[8] A. Kumar and C. Rosenberg, "Energy and throughput trade-offs in cellular networks using base station switching," *IEEE Trans. Mobile Comput.*, vol. 15, no. 2, pp. 364–376, Feb. 2016.

[9] T. Beitelmal, S. S. Szyszkowicz, and H. Yanikomeroglu, "Regular and static sector-based cell switch-off patterns," in *Proc. IEEE Veh. Technol. Conf. (VTC-Fall)*, Montreal, QC, Canada, Sep. 2016, pp. 1–5.

[10] A. Guo and M. Haenggi, "Spatial stochastic models and metrics for the structure of base stations in cellular networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 11, pp. 5800–5812, Nov. 2013.

[11] F. Lagum, Q.-N. Le-The, T. Beitelmal, S. S. Szyszkowicz, and H. Yanikomeroglu, "Cell switch-off for networks deployed with variable spatial regularity," *IEEE Wireless Commun. Lett.*, vol. 6, no. 2, pp. 234–237, Apr. 2017.

[12] E. Erkut, Y. Ülküsal, and O. Yenicerioğlu, "A comparison of p-dispersion heuristics," *Comput. Oper. Res.*, vol. 21, no. 10, pp. 1103–1113, Dec. 1994.

[13] M. Mirahsan, R. Schoenen, and H. Yanikomeroglu, "HetHetNets: Heterogeneous traffic distribution in heterogeneous wireless cellular networks," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 10, pp. 2252–2265, Oct. 2015.

[14] F. Lagum, S. S. Szyszkowicz, and H. Yanikomeroglu, "Quantifying the regularity of perturbed triangular lattices using CoV-based metrics for modeling the locations of base stations in HetNets," in *Proc. IEEE Veh. Technol. Conf. (VTC-Fall)*, Montreal, QC, Canada, Sep. 2016, pp. 1–5.

[15] B. Błaszczyszyn and D. Yogeshwaran, "Connectivity in sub-Poisson networks," in *Proc. IEEE Annu. Conf. Commun. Control Comput. (Allerton)*, Monticello, IL, USA, Sep. 2010, pp. 1466–1473.

[16] J. G. Andrews, "Seven ways that HetNets are a cellular paradigm shift," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 136–144, Mar. 2013.

[17] D. B. Taylor, H. S. Dhillon, T. D. Novlan, and J. G. Andrews, "Pairwise interaction processes for modeling cellular network topology," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Anaheim, CA, USA, Dec. 2012, pp. 4524–4529.

[18] M. Haenggi, "The mean interference-to-signal ratio and its key role in cellular and amorphous networks," *IEEE Wireless Commun. Lett.*, vol. 3, no. 6, pp. 597–600, Sep. 2014.