Málaga, Diciembre 2009

Informe Ejecutivo

TÍTULO: ACP-1.0: El problema ACP y su resolución mediante tres MOEAs del estado del arte

RESUMEN: Este entregable presenta el problema real de la planificación automática de celdas (*Automatic Cell Planning* o ACP) con el que se trabajará en el proyecto DIRICOM. La adaptación de tres MOEAs del estado del arte, NSGA-II, SPEA2 y PAES, para su resolución también se considera. En concreto, el entregable aborda tres instancias reales correspondientes a un tramo de autovía y dos escenarios urbanos.

OBJETIVOS:

1. Presentar el problema de la planificación automática de celdas.
2. Adaptar tres MOEAs del estado del arte para la resolución del mismo (codificación y operadores genéticos específicos).

CONCLUSIONES:

1. Los algoritmos son capaces de encontrar soluciones de gran calidad para el problema ACP que, además, son de gran interés para los diseñadores de redes de telefonía.
2. Los resultados muestran que los algoritmos que utilizan recombinación de soluciones (NSGA-II y SPEA2) permiten encontrar mejores soluciones.
3. NSGA-II ha mostrado ser superior estadísticamente a SPEA2 para algunos escenarios concretos.

RELACIÓN CON ENTREGABLES:

PRE: —

CO: —

Málaga, December 2008

Executive Summary

Title: ACP-1.0: A real-world ACP problem and its resolution with three state-of-the-art MOEAs

Abstract: This deliverable presents the real-world automatic cell planning problem (ACP) that is being considered within the DIRICOM project. Three state-of-the-art MOEAs, namely NSGA-II, SPEA2, and PAES, have been adapted tor the resolution of this problem. In particular, the deliverable addresses three real world instances that correspond to a highway scenario and two urban zones.

Goals:

1. Presentar el problema de la planificación automática de celdas.
2. Adaptar tres MOEAs del estado del arte para la resolución del mismo (codificación y operadores genéticos específicos).

Conclusions:

1. The algorithm are able to reach high quality solutions for the ACP problem which are also of great interest for the network designers.
2. The results show that the algorithms that use a the recombination operator (NSGA-II and SPEA2) can reach better solutions.
3. NSGA-II has shown to perform statistically better than SPEA2 for some given scenarios.

Relation with deliverables:

PRE: —

CO: —

A real-world ACP problem and its resolution with three state-of-the-art MOEAs

DIRICOM

December 2009

1. Introduction

Planning and managing a cellular phone system make engineers face many challenging optimization problems ([13]). Assuming that the business planning activities are already completed, i.e., choosing customer segments, network technology to be used, etc., one of the most significant technical optimization problem is the radio network planning ([8]), also known as the *Automatic Cell Planning* (ACP) problem, the *network dimensioning* problem, or the *capacity planning* problem. Indeed, the foundation of a well-performing cellular network is the basic radio platform since it is the part of the network which is closer to mobile users. Also, it has a clear benefits for the operators since they reduce the infrastructure costs and, at the same time, increase revenue and user satisfaction.

In the initial deployment of a cellular network, the ACP problem lies in selecting the locations of base stations (BTSs) from a set of candidate sites, as well as their parameter settings, in such a way that a number of network requirements are satisfied. These requirements include the maximization of the covered area and the traffic capacity, while the infrastructure cost is to be minimized. Configuring BTSs is not a simple task, since it implies setting up many configuration parameters, such as the antenna type, the emission power, and/or the tilt and azimuth angles. However, cellular networks need to be adapted to the strongly competitive telecommunication industry: new services, new equipment technologies, increasing system capacity, etc. Even in relatively mature cellular markets, these issues force the deployment of additional sites not only to enhance the system capacity but also to provide increased levels of in-building coverage as mobile users expect to be offered service in all geographical area. The ACP problem therefore holds both for the second generation of cellular phone systems, GSM (*Global System for Mobile* communication, [9]), and its enhanced releases GPRS ([5]) and EDGE ([3]), as well as for the current third generation networks, UMTS (*Universal Mobile Telecommunication System*, [12]).

The ACP problem, as an extension of the classical minimum cost set covering problem, has NP-hard complexity ([4]). Also, the number of configurations for each BTS is very large, thus leading to a huge search space. Even when the BTS parameters such as emission power, tilt and azimuth, which are inherently continuous, are discretized by only considering a subset of possible values. An additional issue emerges in this optimization problem: changing the configuration setting of any BTS may affect other BTSs. For instance, if the maximum emission power of a BTS b is reduced to decrease the signal interference in a given area of the network, other BTSs should hold the traffic capacity that has been left unsupported by b. If these other BTSs are already operating at their full capacity, the network would simply start dropping calls of mobile users. This means that making small local changes would require to recompute most of the network predictions.

The aim of this deliverable is to include an experimental study of EAs to solve real-world ACP instances. Three out of the most well known multiobjective EAs of the literature, namely NSGA-II ([1]), SPEA2 ([17]), and PAES ([6]) have been used. The goal here lies in facing the search engines of these algorithms to both a non-standard solution encoding and specialized genetic operators. A thorough statistical analysis of the results has shown that NSGA-II improves upon SPEA2 in several scenarios whereas the current literature has recently shown that the two algorithms perform almost the same.

2. Experimental Study

It has been found that the most widely known multiobjective EAs, namely NSGA-II, SPEA2, and PAES, have never been applied by using neither the specialiced ACP-targeted encoding nor the specialized genetic operators. It is therefore of relevant interest for researchers on ACP how these algorithm perform in these practical context. Three real-world instances and the most accurate ACP model, the Cell and Test Point model (see [10]). Special attention has been paid to the methodology and the statistical analysis of the results.

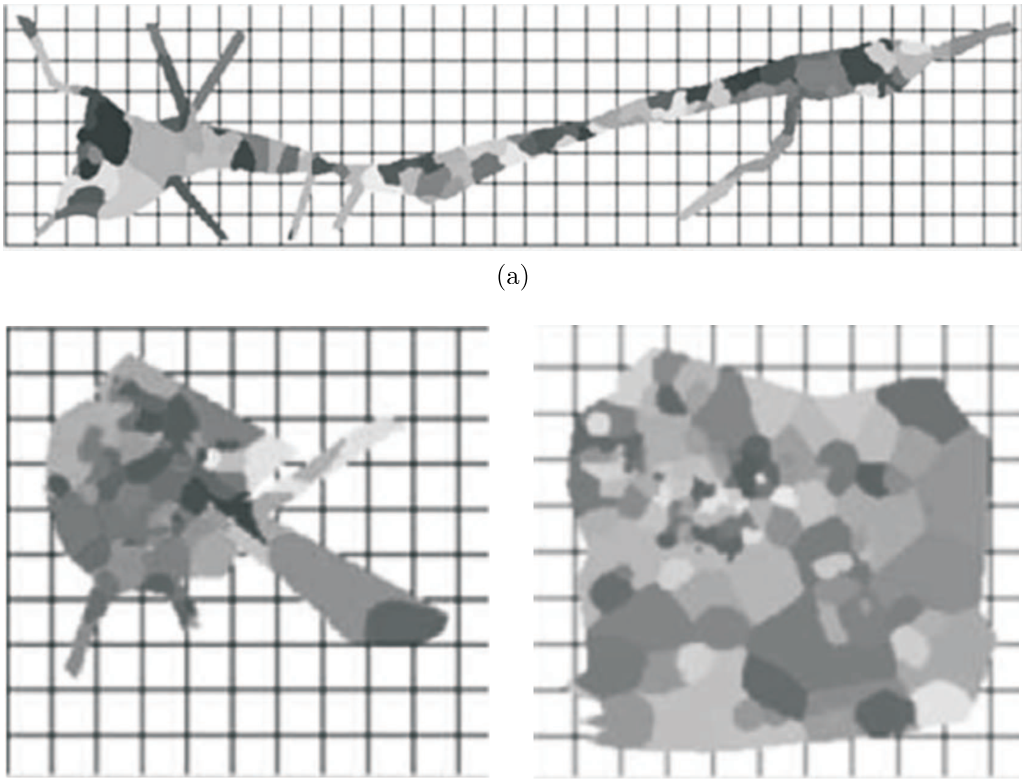

(a) (b) (c)

Figura 1: Topology of the instances (a) Arno 1.0, (b) Arno 3.0, and (c) Arno 3.1

2.1. ACP Model and Real-World Instances

The particular ACP optimization problem addressed here is stated as follows. Given a set of candidate sites for placing the BTSs of the cellular network, the problem lies in choosing the best localization for the BTSs and, for each BTS, selecting the type of antenna and their number. Finally, for each antenna, its maximum power of emission, tilt, and azimuth have to be set up. Three objectives are considered:

1. minimizing the network cost measured in terms of the number of sites used,
2. maximizing the amount of traffic held by the network, and
3. minimizing the interference provoked by overlapping cells,

subject to two constraints:

1. cover (a given percentage of the network area has to be provided with radio service), and
2. handover (ensuring the communication continuity from the starting cell to the target cell, when a mobile is moving toward a new cell).

The ACP model used is the Test Point model ([10]). Many reasons support this decision. This model is the most precise out of those presented in the literature, and therefore it is widely used in the telecommunication industry because it allows the different network objectives to be computed accurately. The interested reader for the all details on the model is referred to [15] and the references therein.

2.1.1. Problem Instances

Three real-world instances provided by France Telecom R&D have been tackled in this experimental study. They correspond with two well distinguished scenarios: the instance Arno 1.0 is a highway environment, whereas the instances Arno 3.0 and Arno 3.1[1] model two urban areas of different size. Their topology is respectively displayed in Figure 1a, Figure 1b, and Figure 1c.

Table 1 includes the most relevant values that characterize the three instances. It can be seen that they have an increasing size, ranging from 250 sites in Arno 1.0 to 743 sites in the Arno 3.1 instance (almost three times larger). It is important to note that not only the number of sites is increased but also the number of test points used to measure the quality of service in the working area.

[1]It is used as a greenfield design instead of an expansion of an existing infrastructure.

Instance	Arno 1.0	Arno 3.0	Arno 3.1
Total traffic (Erlangs)	3 210.91	2 988.12	8 089.78
Number of sites	250	568	747
Number of test points ($\|\mathcal{ST}\|$)	29 955	17 394	42 975
Number of traffic test points ($\|\mathcal{T}\|$)	4 967	6 656	21 475

Table 1: Several values that characterize the three instances tackled.

2.2. Multiobjective EAS Used

In this section, the three EAs that have considered in this study are briefly described. The implementation of these algorithms provided by jMetal ([2]) has been used. jMetal is a Java-based framework aimed at developing metaheuristics for solving multi-objective optimization problems[2].

The NSGA-II algorithm was proposed by [1]. It is a genetic algorithm based on obtaining a new population from the original one by applying the typical genetic operators (selection, crossover, and mutation); then, the individuals in the two populations are sorted according to their rank, and the best solutions are chosen to create a new population. In case of having to select some individuals with the same rank, a density estimation based on measuring the crowding distance to the surrounding individuals belonging to the same rank is used to get the most promising solutions.

SPEA2 was proposed by [17]. In this algorithm, each individual has a fitness value that is the sum of its strength raw fitness plus a density estimation. The algorithm applies the selection, crossover, and mutation operators to fill an archive of individuals; then, the non-dominated individuals of both the original population and the archive are copied into a new population. If the number of non-dominated individuals is greater than the population size, a truncation operator based on calculating the distances to the k-th nearest neighbor is used. This way, the individuals having the minimum distance to any other individual are chosen.

PAES ([6]) is based on a simple (1+1) evolution strategy that uses only a mutation operator. To find diverse solutions in the Pareto optimal set, PAES uses an external archive of nondominated solutions, which is also used to decide about the new candidate solutions. An adaptive grid is used as a density estimator in the archive (the maximun number of subdivisions used in this adaptive grid here is 5).

In order for the algorithms to be fairly compared, all of them share the same solution representation and genetic operators. Concretely, they use the ACP-targeted network encoding, the geographical crossover, and the multilevel mutation, they all are described below in the next section. The crossover and mutation rates respectively are$p_c = 0{,}9$ and $p_m = 1/L$, where L is the number of sites. The stopping condition has been set up to 25,000 function evaluations. Both NSGA-II and SPEA2, use an internal population of size equal to 100; the size of the external archive is also 100 in SPEA2 and PAES. This means that the algorithms will present, at most, 100 different network configurations to the decision maker (the network designer).

2.2.1. Individual Representation and Operators

Tentative solutions of the radio network design problem manipulated by the algorithms introduced previously encode the entire celular network configuration [15]. A multilevel encoding has been used in which level 1 handles the site activation, level 2 sets up the number and type of antennae, and level 3 configures the parameters of the BTS. Then, when a given site is enabled, one or more antennas are activated, always keeping either one single omnidirectional antenna or from one to three directive antennas. Figure 2 displays the hierarchical encoding used. As it can be seen, it is not a classical encoding so the genetic operators used by the EAs have to be properly designed.

Even though it is not an actual search operator, generating the initial population strongly depends on both the given optimization problem at hand and the representation used. In this work, a fully random strategy has been considered. It works as follows: each site is randomly activated. If so, a random configuration for this site is then generated. First, either a omnidirectional antenna or several directive antennas are installed. In the former setting, a random value for the antenna power has to be chosen. In the latter one, the number of BTSs has to be set up (it ranges from one to three). Then, for each directive BTS, its loss diagram (either short or large directive) as well as the power, tilt, and azimuth values are randomly selected.

Let us now describe the genetic operators, which are called geographic crossover and multilevel mutation. It should be noted that they are not well known stardard operators but specially designed ones as long as they have to deal with the solution encoding proposed.

The geographic crossover operates at level 1 in the encoding hierarchy. It is based on exchanging the sites located within a given radius around a randomly chosen site. Let L_i and r be the site and the radius randomly selected. Then, all the sites whose distance to L_i is lesser than r are exchanged between the two parent solutions that are the input of this operator. The entire configuration at level 2 and 3 of the encoding is directly inherited. The geographic crossover is always applied with a probability of 0.9 (in both NSGA-II and SPEA2).

The multilevel mutation has the capability of updating any level of the encoding hierarchy. Each time a solution has to be mutated, only one single site is modified at one unique level of the hierarchy. The operator works as follows.

[2]jMetal is freely available to download at the following URL: `http://jmetal.sourceforge.net/`.

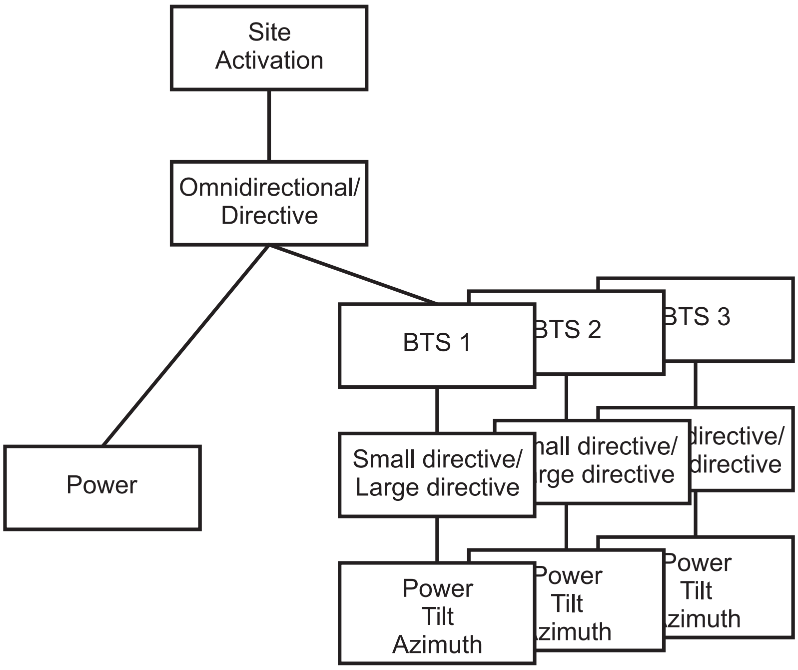

Figura 2: Hierarchical encoding of the individuals.

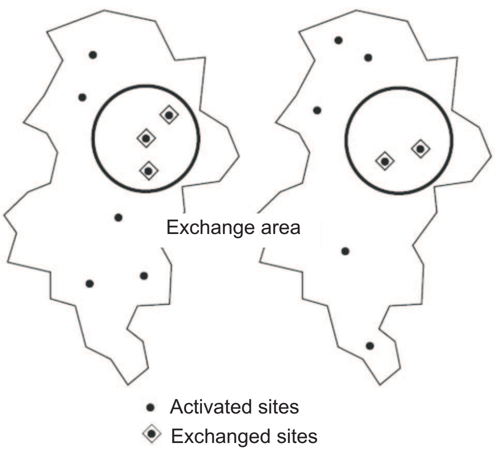

Figura 3: Geographic recombination: sites located within the radius are exchanged.

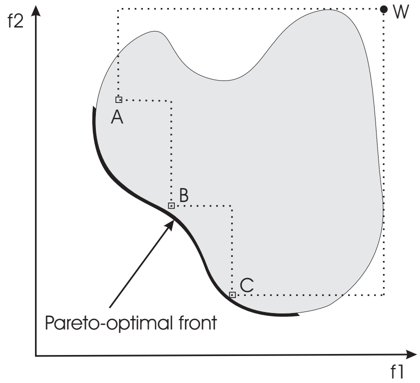

Figura 4: The hypervolume enclosed by the non-dominated solutions.

First, it randomly chooses a network site, L_i, that will undergo mutation. Once selected, five different mutations can be performed on L_i:

- L_i activation toggling. If L_i is activated, then it is just deactivated. On the other hand, if L_i is deactivated, then an entire random configuration for L_i is generated (the same as in the generation of the initial solutions).
- BTS power tuning. It requires L_i to be activated. It randomly chooses a BTS of L_i and then the power is randomly changed to one of its adyacent values (either lower or upper).
- BTS tilt tuning. The same as the power tuning, but changing the tilt angle.
- BTS azimuth tuning. The same as the power and tilt tuning, but modifying the azimuth angle.
- BTS diagram tuning. This mutation also requires the site to be activated. The goal of this operator is to change the BTS type, that is, from an omnidirectional BTS to several directive BTSs, or viceversa. The configuration for each newly generated BTS is randomly selected.

This multilevel mutation is used in the three algorithms with probability of $1/L$, where L is the number of sites.

2.3. Experimental Methodology

In order to measure the performance of the three multiobjective EAs used, the quality of their resulting nondominated set of solutions has to be considered. Since the exact Pareto fronts of the three ACP instances are unknown, a quality indicator which does not require this information beforehand is mandatory. This is why the *Hypervolume*, HV, indicator ([18]) has been chosen in this work. HV is one out of the most well-known and widely used indicator in the literature that allows two desirable features of Pareto fronts (convergence and diversity) to be measured with one single value. Convergence refers to the closeness of the approximated set towards Pareto optimal solutions whereas diversity considers the spread-out of this set of solutions along the nondominated front.

The hypervolume calculates the volume (in the objective space) covered by members of a non-dominated set of solutions Q (the region enclosed into the discontinuous line in Fig. 4, $Q = \{A, B, C\}$) for problems where all objectives are to be minimized. Mathematically, for each solution $i \in Q$, a hypercube v_i is constructed with a reference point W and the solution i as the diagonal corners of the hypercube. The reference point can simply be found by constructing a vector of worst objective function values. Thereafter, a union of all hypercubes is found and its hypervolume (HV) is calculated:

$$HV = volume\left(\bigcup_{i=1}^{|Q|} v_i\right). \tag{1}$$

Higher values of the hypervolume metrics are desirable. Since this indicator is not free from an arbitrary scaling of the objectives, the following normalization procedure has been performed in order to guarantee a fair comparison among the algorithms. First, all the nondominated solutions reached in all the executions of all the algorithms (for the same ACP instance) involved in the experimental study are collected into one single nondominated set. Dominated solutions are then removed. This resulting set can be considered the reference Pareto optimal set to be used as the basis for the normalization. By using the extreme values of this Pareto front, all the objectives of the approximated fronts reached by algorithms are mapped to [0.0,1.0] so that their scales have all the same influence in the HV indicator.

Since stochastic algorithms are considered and the results have to be provided with statistical significance, 30 independent runs for each algorithm and each problem instance have been done. The HV indicator is then computed for each of the approximated fronts (after the aforementioned normalization procedure), thus obtaining 30 HV values

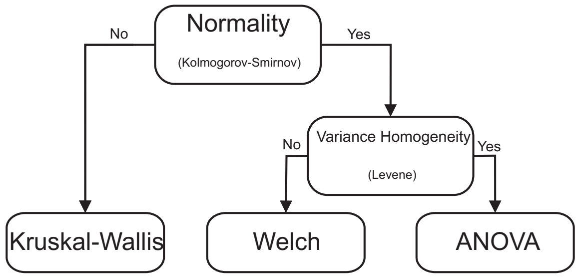

Figura 5: Statistical analysis performed in this work.

Table 2: HV values obtained by NSGA-II, SPEA2, and PAES.

	Arno 1.0	Arno 3.0	Arno 3.1
Algorithm	$\bar{x}_{\pm\sigma_n}$	$\bar{x}_{\pm\sigma_n}$	$\bar{x}_{\pm\sigma_n}$
NSGA-II	$0{,}4424_{\pm 0{,}0218}$	$0{,}5115_{\pm 0{,}0134}$	$0{,}4515_{\pm 0{,}0122}$
SPEA2	$0{,}4431_{\pm 0{,}0180}$	$0{,}4950_{\pm 0{,}0110}$	$0{,}4409_{\pm 0{,}0141}$
PAES	$0{,}2450_{\pm 0{,}0682}$	$0{,}1805_{\pm 0{,}0277}$	$0{,}1102_{\pm 0{,}0241}$

for each pair algorithm/instance. Next, the following statistical analysis has been carried out. First, a Kolmogorov-Smirnov test is performed in order to check whether the values of the results follow a normal (Gaussian) distribution or not. If so, the Levene test checks for the homogeneity of the variances. If samples have equal variance (positive Levene test), an ANOVA test is done; otherwise a Welch test is performed. For non-gaussian distributions, the non-parametric Kruskal-Wallis test is used to compare the medians of the algorithms. Figure 5 summarizes the statistical analysis. A confidence level of 95 % (i.e., significance level of 5 % or p-value under 0,05) is considered in the statistical tests. This means that the differences are unlikely to have occurred by chance with a probability of 95 %.

As long as more than two algorithms are involved in the study, a post-hoc testing phase which allows for a multiple comparison of samples has been performed. The `multcompare` function provided by Matlab^{TM} has been used. It is worth mentioning here that providing the results with statistical significance has involved a remarkable computational effort. Indeed, the time required to carry out 25,000 function evaluations of any of the three ACP instances is longer than one day (so 30 independent runs $\times$ 3 algorithms $\times$ 3 instances $\times$ 1 day means 270 days of computation, at least).

2.4. Results

Table 2 includes the mean, $\bar{x}$, and the standard deviation, σ_n, over the 30 independent runs of the HV indicator reached by the approximated fronts of NSGA-II, SPEA2, and PAES on the three instances Arno 1.0, Arno 3.0, and Arno3.1. The cells with a gray background indicate the best (highest) values. The results of the pair-wise comparisons are included in Table 3. The "+" symbols in the tables indicate that statistical difference exist. The "−" symbol is used otherwise. A comparison with other multiobjective EAs on the same instances using the same representation and operators (e.g., [7, 15, 14]) has not been included because the experimental conditions are fairly different: parallel algorithms, hybridization with local search, and non-standard quality indicators to measure the quality of the Pareto fronts.

The main conclusion that can be drawn from Table 2 is that the GAs, i.e., NSGA-II and SPEA2, always approximate the Pareto fronts with higher (best) HV value. The point that PAES performs the worst leads to assume that recombining solutions becomes a critical issue in this context. That is, PAES is based only on the multilevel mutation operator without crossing over solutions, so therefore it is not able to properly explore the entire search space of the instances. Figure 6 displays this fact. It can be seen that solutions given by PAES are them all concentrated on small regions of front of the three instances, while NSGA-II and SPEA2 reach a wider set of network configurations. This provides the decision maker (e.g., the network designer) with a richer set of BTS settings equally optimal (in the Pareto sense). The availability of solutions is one of the most relevant contributions of using the multiobjective formulation. On the one hand, the radio network planner can choose the network configuration that best fits to his/her expectations. On the other hand, if a decision is already made and new information has to be considered (e.g., the budget for deploying the network is increased), the new network configuration is already computed (a set of nondominated solutions is available). Hence, the radio network designer can react very quickly by just choosing another nondominated solution that involves the installation of more BTSs. Diving a bit more into the PAES results, it can be observed that the larger the size of the instance, the greater the difference between the HV values with respect to NSGA-II and SPEA2. This means that it does not scales well with the current genetic operator configuration. Figure 6 also shown a graphical example of this fact. In the case of the Arno 1.0 instance (Figure 6a), it can be seen that PAES has converged towards the optimal Pareto front, but with a poor spread-out of solutions. For the Arno 3.0 and Arno 3.1 instances (Figure 6b and Figure 6c, respectively), PAES solutions are clearly dominated by

Table 3: Results of the post-hoc testing phase.

	ARNO	NSGA-II 1.0	NSGA-II 3.0	NSGA-II 3.1	SPEA2 1.0	SPEA2 3.0	SPEA2 3.1
SPEA2	1.0	−					
	3.0		+				
	3.1			−			
PAES	1.0	+			+		
	3.0		+			+	
	3.1			+			+

those of NSGA-II and SPEA2. All the previous claims are supported with statistical significance, as the "+" symbol shows in Table 3.

The second relevant finding of this experimental study emerges in the comparison of NSGA-II and SPEA2. The literature has shown that these two algorithms have a similar performance when solving the ACP problem using both the Demand node model ([16]) and the Test Points model ([11]). The thorough statistical tests performed here have revealed that, in some real-world scenarios (i.e., the Arno 3.0 instance), the *HV* values reached by approximated fronts of NSGA-II outperforms those of SPEA2 (with a 95 % of confidence level). In the case of Arno 1.0 and Arno 3.1, nothing can be concluded since the tests have not succeed at this level (see Table 3). In general, few works include statistical analysis that provide their conclusions with statistical confidence. So, in the context of these three real world ACP instances and the experimental settings, it can be concluded that NSGA-II performs better than SPEA2. Again, Figure 6 tries to graphically display this fact. It can be seen that, in the case of the Arno 1.0 and Arno 3.1 instances, the nondominated solutions computed by NSGA-II and SPEA2 almost get overlapped. However, the fronts of the Arno 3.0 instances (Figure 6b) show that NSGA-II reaches solutions in a region that the network configuration can hold a large amount of traffic (at the top of the front), where SPEA2 does not reach. This behavior appears in many of the fronts and it may be the reason for the difference between the *HV* values of these two algorithms.

Referencias

[1] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182 – 197, 2002.

[2] J. J. Durillo, A. J. Nebro, F. Luna, B. Dorronsoro, and E. Alba. jMetal: A java framework for developing multi-objective optimization metaheuristics. Technical Report ITI-2006-10, Departamento de Lenguajes y Ciencias de la Computación, University of Málaga, E.T.S.I. Informática, Campus de Teatinos, 2006.

[3] A. Furuskar, J.Ñaslund, and H. Olofsson. EDGE – enhanced data rates for GSM and TDMA/136 evolution. *Ericsson Review*, (1), 1999.

[4] C. Glasser, S. Reith, and H. Vollmer. The complexity of base station positioning in cellular networks. *Discrete Applied Mathematics*, 148(1):1 – 12, 2005.

[5] H. Granbohm and J. Wiklund. GPRS – general packet radio service. *Ericsson Review*, (1), 1999.

[6] J. Knowles and D. Corne. The pareto archived evolution strategy: A new baseline algorithm for multiobjective optimization. In *Proceedings of the 1999 Congress on Evolutionary Computation*, pages 9 – 105. IEEE Press, 1999.

[7] H. Meunier, E-G. Talbi, and P. Reininger. A multiobjective genetic algorithm for radio network optimization. In *Proceedings of the 2000 Congress on Evolutionary Computation*, pages 317 – 324, 2000.

[8] A. R. Mishra. *Fundamentals of Cellular Network Planning and Optimisation: 2G/2.5G/3G... Evolution to 4G.* Wiley, 2004.

[9] M. Mouly and M. B. Paulet. *The GSM System for Mobile Communications.* Mouly et Paulet, Palaiseau, 1992.

[10] L. Raisanen. *Multi-objective site selection and analysis for GSM cellular network planning.* PhD thesis, Cardiff University, 2006.

[11] L. Raisanen and R. M. Whitaker. Comparison and evaluation of multiple objective genetic algorithms for the antenna placement problem. *MONET*, 10(1-2):79 – 88, 2005.

[12] J. Rapeli. UMTS: Targets, system concept, and standardization in a global framework. *IEEE Personal Communications*, 2(1):30 – 37, 1995.

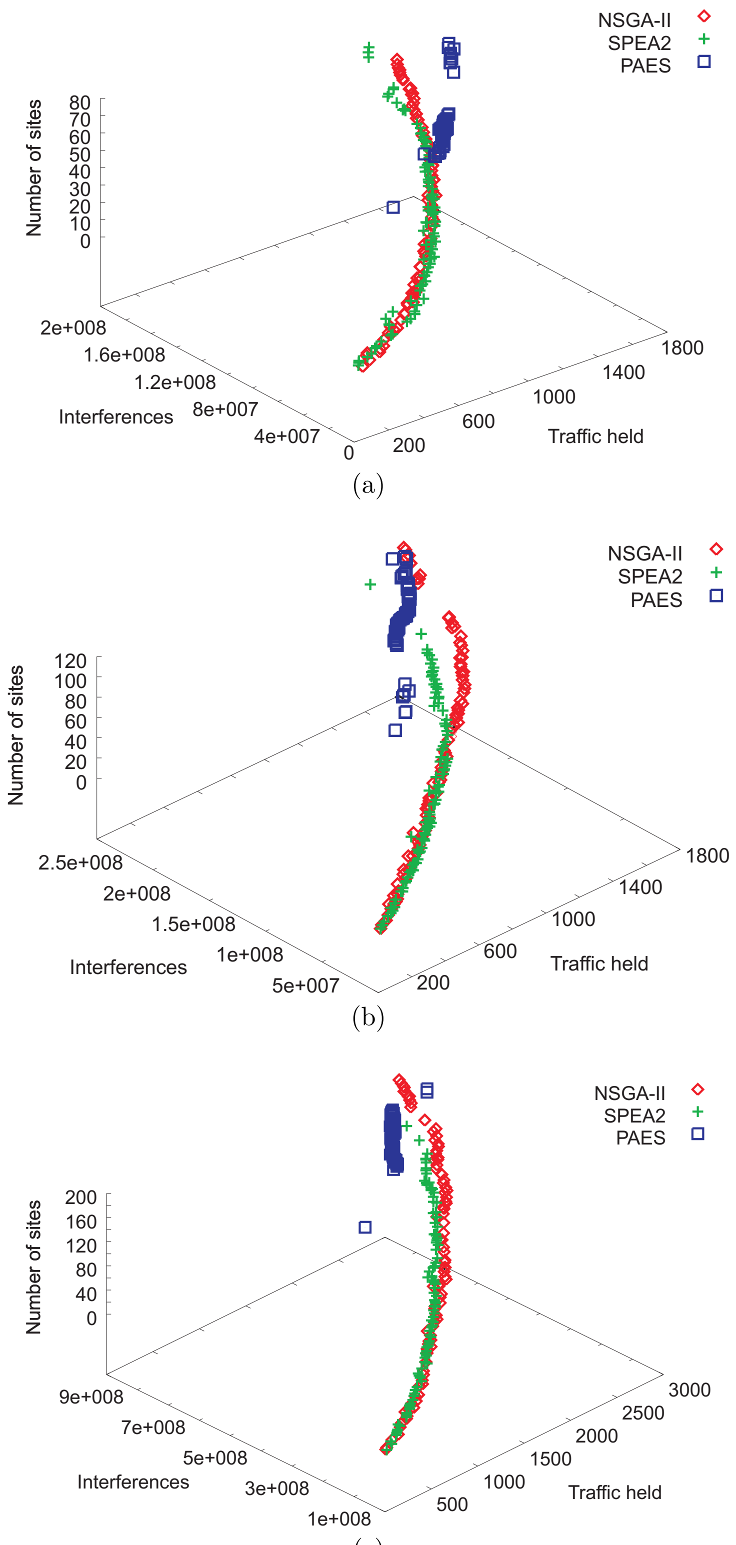

Figura 6: Approximated Pareto fronts of the instances (a) Arno 1.0, (b) Arno 3.0, and (c) Arno 3.1.

[13] M. G. C. Resende and P. M. Pardalos, editors. *Handbook of Optimization in Telecommunications*. Springer, 2006.

[14] E-G. Talbi, S. Cahon, and N. Melab. Designing cellular networks using a parallel hybrid metaheuristic on the computational grid. *Computer Communications*, 30:698 – 713, 2007.

[15] E-G. Talbi and H. Meunier. Hierarchical parallel approach for gsm mobile network design. *Journal of Parallel and Distributed Computing*, 66(2):274–290, 2006.

[16] N. Weicker, G. Szabo, K. Weicker, and P. Widmayer. Evolutionary multiobjective optimization for base station transmitter placement with frequency assignment. *IEEE Transactions on Evolutionary Computation*, 7(2):189 – 203, 2003.

[17] E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the strength pareto evolutionary algorithm. In K. Giannakoglou, D. Tsahalis, J. Periaux, P. Papailou, and T. Fogarty, editors, *EUROGEN 2001. Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems*, pages 95–100, Athens, Greece, 2002.

[18] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257 – 271, 1999.