

Article

GASP: Genetic Algorithms for Service Placement in fog computing systems

Claudia Canali * , Riccardo Lancellotti 

Department of Engineering "Enzo Ferrari", University of Modena and Reggio Emilia, Modena, Italy;
{claudia.canali, riccardo.lancellotti}@unimore.it

* Correspondence: claudia.canali@unimore.it

Version September 11, 2019 submitted to Algorithms

Abstract: Fog computing is becoming popular as a solution to support applications based on geographically distributed sensors that produce huge volumes of data to be processed and filtered with response time constraints. In this scenario, typical of a smart city environment, the traditional cloud paradigm with few powerful data centers located far away from the sources of data becomes inadequate. The fog computing paradigm, which provides a distributed infrastructure of nodes placed close to the data sources, represents a better solution to perform filtering, aggregation and pre-processing of incoming data streams reducing the experienced latency and increasing the overall scalability. However, many issues are still open about the efficient management of a fog computing architecture, such as the distribution of data streams coming from sensors over the fog nodes to minimize the experienced latency. The contribution of this paper is twofold. First, we present an optimization model for the problem of mapping data streams over fog nodes, considering not only the current load of the fog nodes, but also the communication latency between sensors and fog nodes. Second, to address the complexity of the problem we present a scalable heuristic based on genetic algorithms. We carried out a set of experiments based on a realistic smart-city scenario: the results show how the performance of the proposed heuristic is comparable with the one achieved through the solution of the optimization problem. Then, we carried out a comparison among different genetic evolution strategies and operators that identify the uniform crossover as the best option. Finally, we perform a wide sensitivity analysis to show the stability of the heuristic performance with respect to its main parameters.

Keywords: Fog computing; Optimization model; Genetic algorithms; Sensitivity analysis

1. Introduction

In the last few years we witnessed an ever increasing popularity of sensing applications, characterized by the fact that geographically distributed sensors produce huge amounts of data that are then pushed towards the Internet core where cloud computing data centers are located to be processed. However, this traditional approach may cause excessive delays for those applications and services that require data to be processed with very low and predictable latency, such as those related to systems for smart traffic monitoring, support for autonomous driving, smart grid, or fast mobility applications (i.e., smart connected vehicle or connected rails). Another important observation is that not all the data needs to go to the cloud data centers, but in many cases data could be pre-processed, aggregated and filtered to store in network core only a reduced meaningful set of data, thus avoiding to unnecessarily stress the network infrastructure.

The emerging paradigm of fog computing represents a solution that can improve scalability and reduce application latency by extending cloud computing towards the edge of network [1,2]. While in the traditional cloud architecture (Fig. 1a), the data flows are sent directly from the *sensor layer* to the data center(s) in the *cloud layer*, in the fog infrastructure (Fig. 1b) tasks and services may be moved close to the sources of data to be processed thanks to the presence of an intermediate layer of *fog nodes*, located at the *network edge* interposed between the cloud data center(s) and the sources of data. The innovative paradigm of fog computing is

35 promising in addressing the still unsolved issues of cloud computing related to unreliable latency, lack of mobility
 36 support and location-awareness. However, realizing the fog computing full potential introduces several new
 37 challenges [3,4]. Many existing papers focus on balancing load distribution between fog and cloud resources:
 38 among them, we should mention the studies in [5,6] that tackle the issue of optimizing the allocation of the
 39 processing tasks coming from the fog nodes over the cloud infrastructure. To this aim different solutions are
 40 proposed such as the possibility to rely on horizontal communication among fog nodes to reduce the service
 41 delay through load sharing mechanism.

42 On the other hand, less attention was received by the lower level connecting the data sources to the fog nodes.
 43 In literature, indeed, many studies rely on the assumption that fog nodes communicate directly with sensors or
 44 mobile users through single-hop wireless connections [5] or that a domain of sensor nodes communicate with
 45 a specific and application-defined domain of fog nodes [6]. However, the choice of the fog node that receives
 46 and processes the data flow originated by a specific sensor may significantly affect the perceived latency and the
 47 overall performance of the application. Hence, we claim that mapping the data flows from the sensors over the
 48 available fog nodes for processing and filtering tasks represents a critical issue to guarantee a high QoS in terms
 49 of latency and/or response time.

50 The contribution of this paper is twofold:

- 51 1. We present an optimization model for mapping the incoming data flows (sensors workload) over the nodes
 52 of the fog layer: the proposed model considers not only the processing time on the fog nodes depending on
 53 the local load, but also the latency between sensors and fog nodes due to the communication delay of the
 54 geographically distributed infrastructure.
- 55 2. We propose an heuristic to provide a scalable solution to the problem of mapping sensors over fog nodes
 56 that could be applied to large-scale instances. The proposed heuristic is based on Genetic Algorithms
 57 (GAs), a method for solving both constrained and unconstrained optimization problems that relies on the
 58 natural selection process that drives biological evolution: those kind algorithms have been previously and
 59 successfully exploited in the context of cloud computing and Software-as-a-Service placement [7] but, to
 60 the best of our knowledge, it has never applied to fog computing infrastructures.

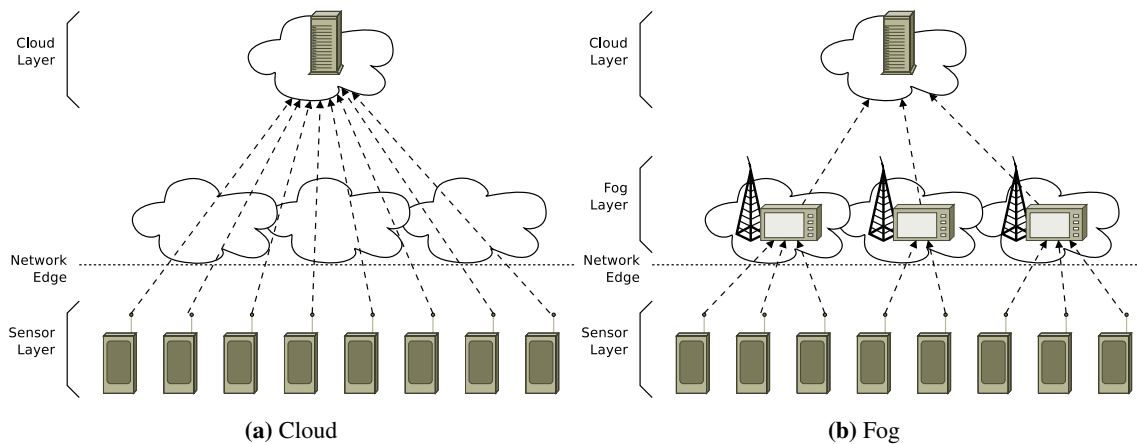


Figure 1. Cloud and Fog Infrastructures

61 This paper extends a previous study by the same authors [8], representing a clear step ahead for the following
 62 reasons:

- 63 • the proposal of a more streamlined model for the optimization problem;
- 64 • a in-depth discussion of the genetic algorithm features and of their adaptation to the specific context;
- 65 • a new and more significant experimental setup;
- 66 • new experimental results including a broad sensitivity analysis on the main heuristic parameters.

67 To evaluate the performance of the proposed solution, we consider a realistic smart city scenario as an
 68 example of a typical sensing environment where geographically distributed sensors produce data flows require

69 efficient processing for a wide range of possible applications, such as traffic monitoring and control, support
70 for autonomous driving and environmental sensing. Experiments were carried out on a geographic testbed
71 representing a fog architecture whose nodes are realistically placed in the streets of a small-sized city (roughly
72 180.000 inhabitants) in Emilia Romagna, Italy. The experiments show that the proposed heuristic can achieve
73 performance similar to the one of a commercial solver applied to an optimization problem for mapping the data
74 flows over the fog nodes. Then, we compare different genetic evolution strategies and operators to identify the
75 best options (for example, we show how the uniform crossover outperforms other crossover operator or we
76 demonstrate that the tournament selection is a better choice than the roulette selection). Finally, we evaluate
77 the stability of the heuristic performance with respect to parameters, such as the number of generations, the
78 probability of mutation and crossover, and the population size.

79 The remainder of this paper is organized as follows. Section 2 describes the problem formally defines
80 the considered optimization model, while Section 3 presents the heuristic algorithms proposed for solving the
81 problem. Section 4 describes the experimental testbed and results used to prove the viability of our approach.
82 Finally, Section 5 discusses the related work and Section 6 concludes the paper with some final remarks and
83 outlines open research problems.

84 2. Problem definition

85 This section describes the fog infrastructure and the problem of efficiently mapping the data flows coming
86 from the sensors over the available fog nodes. In the second part of the section, we present the optimization
87 model that formalized the mapping problem.

88 2.1. Mapping problem

89 Our problem concerns the management of data flows in a fog infrastructure such as the one shown in Fig. 1b.
90 The infrastructure, that we assume to be deployed in a smart-city scenario, is composed of three layers: a *sensor*
91 *layer* that produces data (represented as a set of wireless sensors at the bottom of the figure), a *fog layer* that is
92 responsible for a preliminary processing of data from the sensors (second layer in the figure), while a *cloud layer*
93 that is the final destination of the data (at the top of the figure). The underlying application logic involves the
94 typical services of a smart city scenario. Sensors collect information about the city status, such as traffic intensity
95 or air quality [9]. Such data should be collected at the level of a Cloud infrastructure to provide value-added
96 services such as traffic or pollution forecast. The proposed fog layer intermediates the communication between
97 the sensors and the cloud to provide scalability and reliability in the smart city services.

98 In our model, we assume a stationary scenario where a set of similar sensors \mathcal{S} are distributed over an area
99 (we consider the sensors to be not moving, although a different scenario, where mobility is taken into account
100 can be easily introduced in our model). Furthermore, we assume that sensors are producing data at a steady rate,
101 with a frequency that we denote as λ_i for the generic sensor i (for a summary of the symbols used in the model,
102 the reader may refer to Tab. 1). The fog layer consists of a set of nodes \mathcal{F} that receive the data from the sensors
103 and perform operations on them. These operations typically include pre-processing of the data, such as filtering
104 and/or aggregation, or may include some form of analysis to identify anomalies or problems as fast as possible.
105 The refined data samples from the fog nodes are then sent to a cloud platform where additional analysis is carried
106 out and where all the information is stored. These additional analysis tasks are typically highly expensive from a
107 computational point of view.

108 As the problem concerning the management of large cloud data centers has been widely addressed in
109 literature [10,11], we do not consider the inner details of the cloud layer in our problem modeling, such as the
110 computation time at the level of the cloud data center. Similarly, we do not consider the network-based latency
111 due to the communication between the fog nodes and the cloud data center, as we assume this latency to be the
112 same for all fog nodes and hence having no impact on the optimum mapping solution. Instead, we focus our
113 attention to the problem of coordinating the communication of the elements in the sensor layer with the nodes in
114 the fog layer. Specifically, we want to guarantee a high QoS, in terms of fast response. To this aim, we model the
115 response time according to a queuing theory-based formulation, considering that the response time has two major
116 contributions that should be taken into account:

- 117 • Network-based latency due to the communication between the sensor and the fog nodes. We denote this
- 118 value as $\delta_{i,j}$ where i is a sensor and j is a fog node.
- 119 • **Computation time on the fog node.** According to queuing theory, this time depends on the computation
- 120 cost of the request (we denote as $1/\mu_j$ the time to process a packet of data from a sensor on fog node j ,
- 121 relying on the typical queuing theory notation) and on the rate λ_j of the jobs incoming at fog node j , where
- 122 the value of λ_j is the sum of all the outgoing data rates λ_i of all the sensors i that are communicating with
- 123 the fog node j .

124 For the sake of clarity we summarize the symbols used throughout the paper in Tab. 1.

125 2.2. Optimization model

126 The main problem in the considered fog scenario is how to map on the fog nodes the data flows coming
 127 from the sensors. To this aim, we define an optimization problem where we use as the main decision variable a
 128 matrix of boolean flags $x_{i,j}$. In our model $x_{i,j} = 1$ if and only if sensor i is sending data to fog node j , otherwise
 129 $x_{i,j} = 0$. As the function of fog nodes is to pre-process the incoming data performing filtering and aggregation,
 130 we consider that all the data of a sensor must be sent to the same fog node and cannot be distributed across the
 131 fog layer.

132 Again, the reader may refer to Tab. 1 for a summary of the parameters used in our model.

Table 1. Notation.

Symbol	Meaning/Role
Decision variables	
$x_{i,j}$	Sending data flow from sensor i to fog node j
Model parameters	
\mathcal{S}	Set of sensors
\mathcal{F}	Set of Fog nodes
λ_i	Outgoing data rate from sensor i
λ_j	Incoming data rate at fog node j
$1/\mu_j$	Processing time at fog node j
$\delta_{i,j}$	Communication latency between sensor i to fog node j
Model variables	
i	Index of a sensor
j	Index of a fog node

133 The optimization model to address the previously-described problem can be formalized as follows, with
 134 an approach similar to the problem of allocating requests over a distributed infrastructure, such as VMs on a
 135 cloud [10,12,13]. In particular, we introduce a matrix of boolean decision variables $X = \{x_{i,j}\}$ that is used to
 136 define the objective function and the constraints as follows:

$$\min obj(X) = \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{F}} x_{i,j} \cdot \left(\frac{1}{\mu_j - \lambda_j} + \delta_{i,j} \right) \quad (1)$$

subject to:

$$\lambda_j = \sum_{i \in \mathcal{S}} x_{i,j} \cdot \lambda_i \quad \forall j \in \mathcal{F}, \quad (2)$$

$$\sum_{j \in \mathcal{F}} x_{i,j} = 1 \quad \forall i \in \mathcal{S}, \quad (3)$$

$$\lambda_j < \mu_j \quad \forall j \in \mathcal{F}, \quad (4)$$

$$x_{i,j} = \{0, 1\}, \quad \forall i \in \mathcal{S}, j \in \mathcal{F}, \quad (5)$$

137 In the problem formalization, the objective function 1 aims at reducing the total (and hence the average)
 138 latency and processing time from every sensor to the fog computing nodes. The expression of response time
 139 used for our objective function is consistent with other studies in literature focusing on distributed cloud
 140 infrastructures [14]. Specifically, the average processing time is derived from Little's result applied to a M/G/1
 141 model and considers just the average arrival frequency λ_j and the processing rate μ_j of each fog node j . This
 142 definition of the response time has been widely adopted in literature, for example in [14]. The second part
 143 of the objective function, that is the latency contribution, captures effectively the communication delay of a
 144 geographically distributed infrastructure using the latency $\delta_{i,j}$.

145 Together with the objective function, we have a set of constraints. Equation 2 defines the incoming load λ_j
 146 on each fog node j . Constraint 3 means that for every sensor i , we direct its output to one and only one fog node.
 147 Constraint 4 guarantees that, for every fog node j , we avoid a congestion situation, where the incoming load λ_j
 148 exceeds the processing capability μ_j of that node. Finally, constraint 5 defines the boolean nature of the decision
 149 variables $x_{i,j}$.

150 3. Heuristic algorithm

151 The optimization problem defined in the previous section aims to map sensors over fog nodes. This model
 152 can be processed using commercial solvers, like CPLEX or KNITRO [15], that have been successfully used in
 153 similar problems [16]. However, in this paper we explore the opportunity to develop a specific heuristic to tackle
 154 the problem.

155 When considering heuristic algorithms, multiple options are available. Greedy heuristics tend to be quite
 156 fast, but their performance may depend heavily on the inherent nature of the problem due to the risk, common to
 157 several gradient descent methods, of being stuck in local minima. With respect to this problem, the combination
 158 of a non-linear objective function and a feasibility domain that is not guarantee to be convex may hinder the
 159 application of greedy solutions. On the other hand, the dimensionality of the problem, with potentially many fog
 160 nodes and many sensors, may reduce the performance of branch and bound approaches, that have a large decision
 161 variable space to explore. Since our goal is to provide a general and flexible approach to tackle this problem, we
 162 focus on meta-heuristics that should be more able to adapt to a wide and heterogeneous set of problems [17].
 163 We focus on evolutionary programming (and in particular on genetic algorithms – GAs) because this class of
 164 heuristics has already provided good results in problems with similar characteristics, such as the allocation of
 165 VMs on a cloud infrastructure [7].

166 3.1. Genetic Algorithms overview

167 We now provide an overview of the use of genetic algorithms (GAs) to solve an optimization problem.
 168 When using GAs, we consider a *population* of *individuals*. Each individual encodes a solution of a problem
 169 as a *chromosome*: for example a generic individual k will be represented as a its chromosome C^k . In turn, the
 170 chromosome C^k is a sequence of a fixed number *genes*, that is $C^k = \{c_l^k\}$, and each gene represents a parameter
 171 that characterize the solution for that individual.

172 The algorithm is initialized with a randomly generated population of individuals. To each individual is
 173 applied a *fitness function*, that is the objective function of the optimization problem. Hence, to each individual,
 174 we assign a fitness score that is the value of the objective function for that solution of the optimization problem.
 175 The population of individuals evolves through a defined number of *generations*, where the overall fitness of the
 176 population is increased using the following operators:

177 **Selection** decides if an individual is passed from the K^{th} generation to the $(K + 1)^{th}$. Selection typically uses
 178 the fitness score of each individual with the goal to discard unfit individuals. The typical approach in
 179 this case is to apply the fitness function to every individual (including new individuals generated through
 180 mutation and crossover) and to consider a probability of being selected for the next generation that depends
 181 on the fitness value.

182 **Mutation** is a (random) change in single gene or in a group of genes. In GAs, mutation plays the role of adding
 183 new genetic material with the main goal to explore new areas of the solution space.

184 **Crossover** is a merge of two individuals by exchanging part of their chromosomes. The main role of crossover
185 in GAs is to allow successful genes in the parent solutions to spread throughout the population.

186 These operators are combined to create an evolutionary strategy. In our analysis we consider the following
187 two evolutionary strategies:

188 **Simple strategy** is a strategy, described in [18], that, at every generation, samples the previous generation
189 population using the selection operator. In doing so, the fittest individuals (that are more likely to be picked
190 by the selection operator) are replicated, while the least fit individuals are left out of the new generation.
191 Next, the mutation and crossover operators are applied to the new population, where each individual can be
192 mutated or can undergo crossover with a given probability (defined as P_{mut} and P_{cx} , respectively). Mutated
193 individuals replace the originals, while the offspring replaces the parent individuals.

194 **Mu Plus Lambda strategy** ($\mu + \lambda$) is another popular approach in evolutionary programming [18]. In this
195 case, we apply the mutation and crossover operation to the previous generation population in order to
196 create a set of offspring with a size of λ individuals. Next, the selection operator is invoked on the original
197 population *plus* the offspring, with the aim to select μ individual for the next generation population (μ is
198 selected to maintain the population stable over the generations). It is worth to note that the symbols μ and
199 λ for the selection strategy are unrelated to the parameters with the same name introduced in the model
200 of Section 2. We preserve this notation, to match the definition of evolutionary strategies as described
201 in [18,19].

202 3.2. Selection operator

203 The selection operator is used together with the evolutionary strategy to define how individuals are passed
204 from a generation to the next. The selection operation is called multiple times on a pool of individuals (typically
205 from the older generation) and returns one individual for each invocation that will be passed in the current
206 generation. In our analysis, we focus on two selection operators:

207 **Tournament selection** is a selection operator where, for each individual to return, the operator picks randomly
208 N elements in the population and returns the fittest one among them.

209 **Roulette selection** is an operator that selects, for each individual to return, a chromosome from the population
210 with a probability that is proportional to its fitness score.

211 3.3. Mutation operator

212 The mutation operator is called on each individual within the population with a probability that we call P_{mut} .
213 Mutation alters one or more genes of the selected chromosome. Multiple mutation operators can be implemented.
214 In our study, we consider two types of mutation operators that are:

215 **Shuffle mutation** is a mutation where the mutated individual takes the genes of original chromosome applying
216 a permutation to them. An example of shuffle mutation is shown in the upper part of Fig. 2, where the
217 genes that are involved in the mutation are marked with a darker color and arrows are used to show the
218 genes permutation.

219 **Uniform Integer mutation** is a mutation operator in which one or more genes are modified. The value of the
220 affected genes is replaced using a uniform random distribution. In Fig. 2, this mutation operator is shown
221 in the lower part of the image and the affected gene is marked with a darker color.

222 3.4. Crossover operator

223 The crossover operation takes two individuals and mate them to create two new individuals (offspring) that
224 inherit their genes from the two parents. In the sensitivity analyses in Section 4, we will refer to the probability
225 of selecting an individual for a crossover operation as P_{cx} . There are several possible crossover operators. In our
226 analysis, we consider the following options:

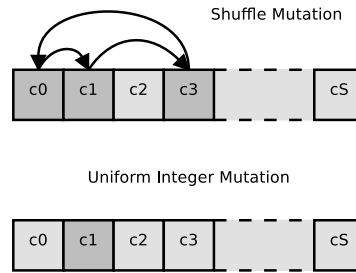


Figure 2. Examples of mutation operators

227 **Uniform crossover** is crossover operator where the offspring will inherit each gene from a randomly selected
 228 parent. In Fig. 3, the Uniform crossover operation is shown in the leftmost of the image. Each of the
 229 parents is characterized with a color. The offspring inherits each gene from one of the parents, and the
 230 parents from which the gene is inherited is marked with the same color of the parent.

231 **One Point crossover** is characterized by a splitting point in the chromosome is randomly selected. The two
 232 resulting sections of the chromosome (from beginning to the spitting point and from the spitting point to
 233 the end) are then inherited from the parents as shown in the center part of Fig. 3.

234 **Two Points crossover** is similar to the One Point version, but we have two splitting points, as shown in the
 235 rightmost part of Fig. 3.

236 **Uniform Partially Matched crossover (UPMX)** is a variant of the Uniform crossover operator proposed in [20]
 237 that takes into account the possible presence of identical genetic material between the parents.

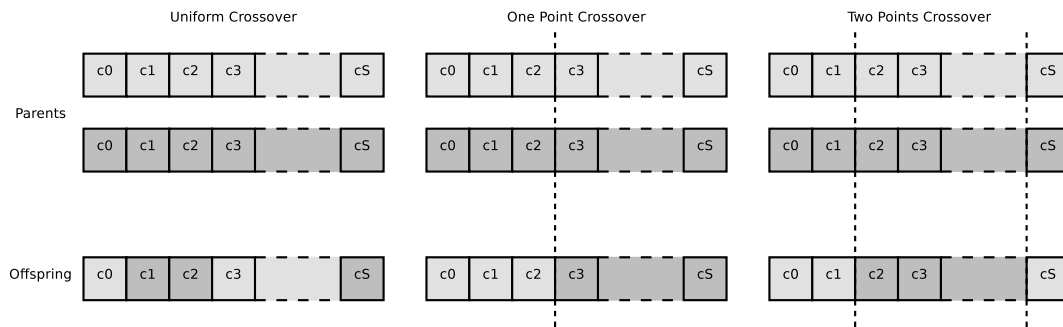


Figure 3. Examples of crossover operators

238 3.5. GA-based problem model

239 We now discuss how the a GA-based model of the problem can be derived from the optimization model
 240 described in Section 2.

241 The first critical choice is how to encode a solution in a chromosome. To this aim, we describe a chromosome
 242 as a sequence of S genes, with $S = |\mathcal{S}|$ being number of sensors. The generic i^{th} gene c_i is represented as an
 243 integer number from 1 to F (with $F = |\mathcal{F}|$ being the number of fog nodes in our infrastructure) and captures the
 244 information on which fog node will receive the output of that sensor. Chromosomes contains the information
 245 of the decision variable $x_{i,j}$ in the problem defined in 2. Specifically, we can define the generic i^{th} gene as
 246 $c_i = \{j : x_{i,j} = 1\}$. As only one fog node will receive data from sensor i , due to constraint 3 in the optimization
 247 model, we can map each possible solution of the optimization problem in Section 2 into a chromosome, with no
 248 conflicts or ambiguities.

249 The second critical problem is the definition of the fitness function for the evaluation of the chromosomes.
 250 This design choice is straightforward because, due to the mapping between chromosomes and optimization
 251 problem solutions, we can simply adapt the objective function 1 to the GA-based model and use this function for
 252 the evaluation of individuals.

253 Finally, we must take into account the constraints of the optimization problem. Constraints 3 and 5 are
 254 automatically satisfied by our encoding of the chromosomes. We need to implement also constraint 4 concerning

255 the fog node overload. We chose not to embed directly the notion of unacceptable solution in a genetic algorithm
 256 as it may hinder the ability of the heuristic to converge towards a solution. Instead, we added this constraint
 257 into the fitness function, so that an individual that contains a solution with one or more overloaded fog nodes
 258 is characterized by a high penalty and is, therefore, likely to exit the genetic pool in few generations. To
 259 define the value of the penalty, we refer to the model used in the problem definition. In particular, for most
 260 solvers the constraint 4 must be reformulated using a lesser-or-equal relationship, rather than the lesser-to. Hence,
 261 constraint 4 becomes $\lambda_j \leq \mu_j - \epsilon$, with $\epsilon = 10^{-5}$. We leverage this alternative formulation also for the GA-based
 262 model forcing the penalty in such a way that, if there is overload, the processing time contribution to the objective
 263 function becomes equal to $1/\epsilon$.

264 4. Experimental results

265 We now present the main findings of our research. We start by introducing the reference scenario of our
 266 experiments. Next, we compare the ability to achieve an optimal solution of the GAs-based approach with the
 267 AMPL-based [21] model solved using KNITRO [15]. In the remaining of the paper we discuss the sensibility of
 268 the GA-based heuristic with respect its main parameters: the number of generations, the evolutionary strategies
 269 and operators, the probability of mutation and crossover, and the population size.

270 4.1. Experimental testbed

271 To evaluate the viability of our proposal, we consider a fog scenario characterized by (1) a significant
 272 number of sensors; (2) a set of fog nodes, with limited computational power, that aggregate and filter the data
 273 from the sensors; (3) a cloud data center that collects the information processed by the fog nodes. Our testbed
 274 scenario is based on a smart city whose topology is based on a real Italian city (Modena) with a population in the
 275 order of 180,000 inhabitants.

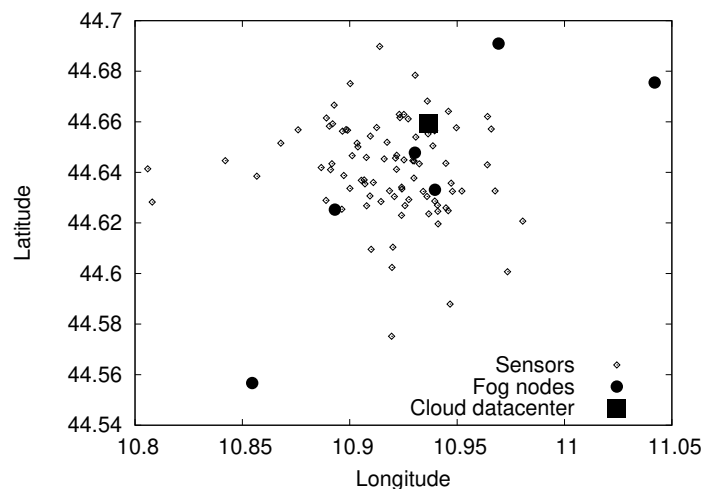


Figure 4. Smart city scenario

276 Our reference use case is a traffic monitoring application. Sensors, fog nodes and the cloud data center are
 277 shown in Fig. 4 to represent the smart city scenario. To gather data concerning traffic-related measures, such as
 278 the number of cars passing in each street (with their speed), we place wireless sensors on the main streets of
 279 the city. To model the application we referred to the Trafair Project [9], currently involving the city of Modena
 280 considered for the topology model.

281 The sensors map in Fig. 4 is created starting from the geo-referenced list of the streets of the cities
 282 characterized by higher volumes of traffic: we assume to have one sensor for each one of these streets.
 283 Furthermore, we selected a group of 6 buildings that host offices of the municipality and that are interconnected
 284 with a high speed metropolitan area network: each one of these buildings is assumed to host a fog node. The
 285 final scenario is composed of a 89 sensors and 6 fog nodes. The interconnection between fog nodes and sensors
 286 is characterized by a delay that we model using the euclidean distance between the nodes. The average delay is

287 in the order of 10 ms, that is consistent with a geographic network. Finally, we assume to have only one cloud
 288 data center, that is located where the Modena municipality data center actually is and is connected to the fog
 289 nodes through low latency links not considered in the optimization model.

Concerning the traffic model, we describe each scenario based on two metrics. The first metric is the average load of the system ρ , while the other parameter $\delta\mu$ represents the ratio between the average network delay (δ) and the average service time ($1/\mu$). Specifically, we define the two parameters as:

$$\rho = \frac{\sum_{i \in \mathcal{S}} \lambda_i}{\sum_{j \in \mathcal{F}} \mu_j} \quad (6)$$

$$\delta\mu = \frac{\sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{F}} \delta_{i,j}}{|\mathcal{F}| |\mathcal{S}|} \cdot \frac{\sum_{j \in \mathcal{F}} \mu_j}{|\mathcal{F}|} \quad (7)$$

290 In our experiments, we consider several scenarios, corresponding to different combinations of these
 291 parameters, in order to analyze the performance of the GAs-based solution for the sensor mapping problem. For
 292 example, a scenario where $\rho = 0.8$ and $\delta\mu = 0.01$ (corresponding to the bottom right corner of Fig. 5) represents
 293 a case where network delay is much lower than the average job service time, while the processing demand on the
 294 system is high. This means that the scenario is CPU-bound because managing the computational requests is likely
 295 to be the main driver to optimize the objective function. On the other hand, a scenario where $\rho = 0.2$ and $\delta\mu = 1$
 296 (top left corner of Fig. 5) is a scenario characterized by a low workload intensity and a network delay comparable
 297 with service time of a job, where it becomes important to optimize also the network contribution to the objective
 298 function. Throughout our experiments, we first consider the implementation of the model discussed in Section 2,
 299 using the AMPL language [21] and KNITRO [15] as the solver. Due to the nature of the problem, we were not
 300 able to let the solver run until the convergence. Instead, we placed a walltime limit of 120 minutes, with a 16 core
 301 CPU and 16 concurrent threads. We compare the results of this solver with the GAs-based implementation, that
 302 uses the Distributed Evolutionary Algorithms in Python (DEAP) framework [19]. Using the same framework, we
 303 also evaluate and compare several evolutionary strategies and operators, as described in Section 3.

Table 2. Default GAs parameters

Parameter	Value
Strategy	<i>Simple strategy</i>
Selection	<i>Tournament selection</i>
Mutation	<i>Uniform Integer mutation</i>
Crossover	<i>Uniform crossover</i>
Population size	200
Generations	300
P_{mut}	0.8%
P_{cx}	0.8%

304 When evaluating the GAs-based approach, we run the experiments 100 times and we report the main
 305 performance metrics in the form of average value and confidence interval. In particular, for the confidence
 306 interval we consider a span of $\pm 2\sigma$, where σ is the standard deviation, that accounts for $\approx 95\%$ of the samples.
 307 The genetic algorithm considers the default parameters shown in Tab. 2, that have been selected after some
 308 preliminary experiments. Furthermore, when considering the convergence speed, we consider as the convergence
 309 criteria the case of a fitness value within 1% of the optimum value obtained using the AMPL solver.

310 4.2. Comparison of Solver and GA-based approaches

The first analysis of our research compares the difference between the solution found by the genetic algorithm and the one obtained by the KNITRO solver. To this aim, we introduce as the main performance metric the discrepancy ϵ defined as follows:

$$\epsilon = \frac{Opt_S - Opt_{GA}}{Opt_S} \quad (8)$$

311 Where Opt_S is the value of the objective function for the best solution found by the solver, and Opt_{GA} is the
 312 value based on the best solution found using the genetic algorithm. It is worth to note that the run with KNITRO
 313 are not guaranteed to reach optimality because we stop the solver after a given amount of time. Hence, in some
 314 cases the genetic algorithm may outperform the solver, resulting in $\epsilon < 0$, as we will see in the results.

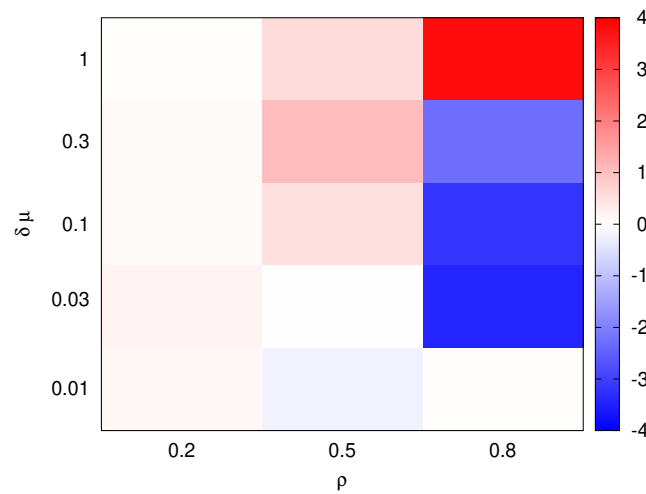


Figure 5. Performance comparison: ϵ [%]

315 Fig. 5 shows as an heatmap the value of ϵ for ρ ranging from 0.2 to 0.8 and $\delta\mu$ ranging from 0.01 to 1. In
 316 the color coded representation of ϵ , blue hues refer to a better performance of the genetic algorithm, while red
 317 hues correspond to better performance of the solver. From this comparison, we observe that the performance
 318 of the two approaches are similar, with the genetic algorithm providing slightly better performance in some
 319 cases (e.g., for $\rho = 0.8$, $\delta\mu = 0.03$ we have $\epsilon = -3.3\%$) and the solver prevailing in other cases (for $\rho = 0.8$,
 320 $\delta\mu = 0.01$ we have $\epsilon = 3.8\%$). It is worth to note that most differences occur for $\rho = 0.8$, that is when the risk
 321 of overloading the fog nodes is higher and the value of the objective function is highly variant with respect to the
 322 considered solution.

323 Using this heatmap, we select two relevant cases that will be used throughout the remaining of the paper to
 324 analyze the stability of the genetic algorithm behavior: we consider an intermediate value for the load ($\rho = 0.5$),
 325 while for the delay impact we consider **three scenarios in the range from $\delta\mu = 0.01$ to $\delta\mu = 1$ to fully explore**
 326 **the variations in problem properties from the solver point of view.**

327 4.3. Convergence analysis of GAs

328 The second critical analysis concerns the impact of the number of generations on the performance of the
 329 genetic algorithm. We carry out this analysis for the two previously presented two scenarios, that is $\rho = 0.5$,
 330 $\delta\mu = 0.01$ and $\rho = 0.5$, $\delta\mu = 1$.

331 In this (and in the subsequent) evaluations we consider the previously introduced discrepancy between the
 332 GAs and the solver (ϵ). The value of ϵ is measured at every generation for the GA (compared with the final output
 333 of the solver). This allows us to evaluate if the population is converging over the generations to an optimum.

334 Fig. 6 presents the results of this analysis. Specifically, we consider the evolution of ϵ for **the considered**
 335 **scenarios** through 300 generations of the genetic algorithm. The graph shows also an horizontal line at the value
 336 of 1%: we consider that the genetic algorithm has reached convergence when $\epsilon \leq 1\%$ and we consider the

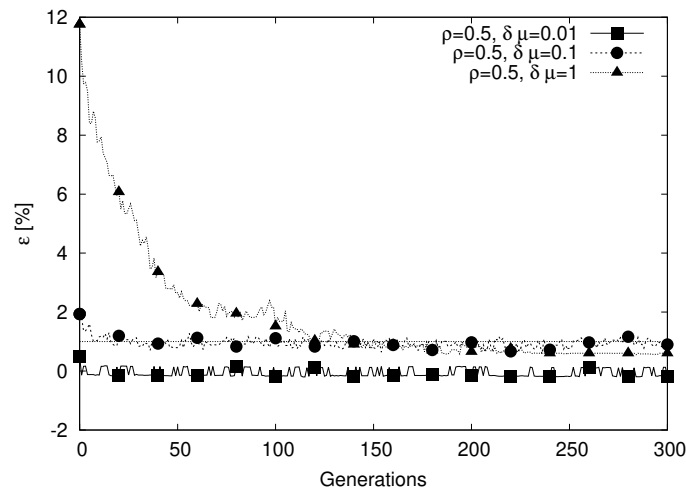


Figure 6. Convergence analysis

337 generation when this condition is verified as a metric to measure of how fast the algorithm can find a suitable
 338 solution. It is worth to note that, given the generally low value of ϵ , this convergence criteria returns values quite
 339 similar to the more common formulation where convergence is defined on the basis of being close (i.e., within
 340 1%) to the best result achieved with the genetic algorithm itself.

341 Comparing the three curves we observe two highly different behaviors. On one hand, for the curve
 342 characterized by $\delta\mu = 1$ (curve marked with filled triangles in Fig 6) we observe a clear descending trend.
 343 At the opposite end of the behavior spectrum, the curve for $\delta\mu = 0.01$ (curve marked with filled squares),
 344 is characterized by a value of ϵ that is very low right from the first generation and remains stable over the
 345 generations. The third curve ($\delta\mu = 0.1$, marked with filled circles) is similar to the case with $\delta\mu = 0.01$, even if
 346 a small descending trend can be observed in the first generations. The reason for this behavior can be explained
 347 considering the nature of the problem, where the objective function depends on two main contributions: the
 348 processing time (that depends mainly on the ability of the algorithm to distribute fairly the sensors among the fog
 349 nodes) and the network latency, that depends on the ability of the algorithm to map the sensors on the closest fog
 350 node. When $\delta\mu = 0.1$ and $\delta\mu = 0.01$, the impact of the second contribution is reduced by one or two orders of
 351 magnitude compared to the case $\delta\mu = 1$. Hence, any solution that provides a good level of load sharing will
 352 be close to the optimum. As the genetic algorithm initializes the chromosomes with a random solution (with
 353 a uniform probability distribution), it is likely to have one or more individual right from the first generation
 354 that provide very good performance. Further evolution of the population provides a limited benefit (due to the
 355 reduced weight of the network latency component) explaining the stable values of ϵ over the generations. The
 356 case when $\delta\mu = 1$ is remarkably different as the two contributions to the objective function have a comparable
 357 impact. Hence, the genetic algorithm must evolve through the generations to explore a large space of solutions to
 358 find good individuals letting the population evolve.

359 4.4. Comparison of evolutionary strategies and operators

360 We now discuss how the considered evolutionary strategies and operators affect the performance of the
 361 genetic algorithm. In these analyses, we will consider two main metrics, that are the best value of ϵ over the
 362 generations introduced in Section 4.2 (we recall that we run the algorithm for 300 generations), and the number
 363 of generations required to reach convergence, introduced in the Section 4.3. For each of the two metrics we
 364 record the average value over the 100 repetitions of the experiments and a confidence interval that accounts for
 365 roughly 95% of the results.

366 Tab. 3 shows the value of the considered metrics as a function of different evolutionary strategies, selection
 367 operators, mutation operators, and crossover operators. Each of these options has been introduced and described
 368 in Section 3.

Table 3. Evolutionary strategies

Strategy or Operator	$\rho = 0.5, \delta\mu = 0.01$		$\rho = 0.5, \delta\mu = 0.1$		$\rho = 0.5, \delta\mu = 1.0$	
	ϵ [%]	# Gen.	ϵ [%]	# Gen.	ϵ [%]	# Gen.
Evolutionary strategies						
<i>Simple</i>	-0.20 ± 0.01	0.00 ± 0.00	0.53 ± 0.07	0.00 ± 0.00	0.54 ± 0.15	23.43 ± 3.82
$\mu + \lambda$	-0.27 ± 0.01	0.00 ± 0.00	0.02 ± 0.05	0.00 ± 0.00	0.65 ± 0.17	21.36 ± 3.68
Selection Operators						
<i>Tournament</i>	-0.21 ± 0.01	0.00 ± 0.00	0.52 ± 0.08	0.00 ± 0.00	0.54 ± 0.13	23.85 ± 3.69
<i>Roulette</i>	0.23 ± 0.27	0.00 ± 0.00	1.54 ± 0.26	0.00 ± 0.00	11.25 ± 0.56	300.00 ± 0.00
Mutation Operators						
<i>Shuffle</i>	-0.23 ± 0.01	0.00 ± 0.00	0.48 ± 0.04	0.00 ± 0.00	1.92 ± 0.29	30.29 ± 5.79
<i>Uniform Int</i>	-0.21 ± 0.01	0.00 ± 0.00	0.53 ± 0.08	0.00 ± 0.00	0.54 ± 0.15	22.82 ± 3.66
Crossover Operators						
<i>Uniform</i>	-0.20 ± 0.01	0.00 ± 0.00	0.52 ± 0.07	0.00 ± 0.00	0.53 ± 0.13	24.21 ± 3.48
<i>One Point</i>	-0.24 ± 0.01	0.00 ± 0.00	0.28 ± 0.08	0.00 ± 0.00	0.81 ± 0.19	27.50 ± 4.26
<i>Two Points</i>	-0.23 ± 0.01	0.00 ± 0.00	0.34 ± 0.07	0.00 ± 0.00	0.78 ± 0.17	26.27 ± 4.10
<i>UPMX</i>	-0.27 ± 0.01	0.00 ± 0.00	0.11 ± 0.07	0.00 ± 0.00	1.07 ± 0.22	27.38 ± 4.28

A first result from the values in the table is the confirmation of the main difference between the considered scenarios. When $\delta\mu = 0.01$ and $\delta\mu = 0.1$, the convergence occurs after just a few generations (typically right in the initial population) because the balancing part of the solution is reached easily, while the network part, that receives most benefit from population evolution, plays a minor role in the objective function. As a consequence, the impact of most considered options (that drive the evolution of the population) is reduced. On the other hand, the scenario with $\delta\mu = 1$ provides a more significant comparison of the alternatives. It also worth to note that the best solution may depend on the scenario, as the inner nature of the problem may change as the balance in the contributions to the objective function shifts.

Starting from the evolutionary strategies, we observe that both the simple strategy and the $\mu + \lambda$ alternative provide similar results, both in terms of objective function value and in terms of convergence. However, due to the additional complexity of the $\mu + \lambda$ strategy (that adds additional parameters to the algorithm) we consider the more straightforward simple strategy as the best option.

Switching to the selection operators, we observe a major difference in the performance of the two alternatives (roulette and tournament selection). In the considered problem, the roulette selection tends to keep over the generations also unfit individuals, rather than purging them from the genetic pool. This hinders the convergence, as shown by the results of our experiments where the roulette selection is outperformed by the alternative in every considered scenario both in terms of ϵ and in terms of number of generations for convergence. Indeed, the roulette selection provides a good behavior when there is a difference that spans order of magnitude in the objective function, which is unlikely to occur in this type of problem. The tournament operator, instead, provides a good ability in filtering individuals from a generation to the next, guaranteeing lower values of ϵ (we recall that the best value of ϵ is computed over the span of 300 generations, so, even if convergence occurs in the initial population, we may experience a positive impact of the strategy on the final value of ϵ).

If we consider mutation operators, we observe another effect of the changes in the problem nature with the different scenarios. For the $\delta\mu = 0.01$ and $\delta\mu = 0.1$, the Shuffle operator is better, because it preserves the most important characteristics that is the load balancing. However, this behavior has a negative impact in the $\delta\mu = 1$ scenario as it hinders a more free exploration of the solution space. On the other hand a uniform mutation operator provides better performance for the $\delta\mu = 1$ scenario at the expenses of a less effective search for the optimum when $\delta\mu = 0.01$ or $\delta\mu = 0.1$.

397 Finally, we compare the four crossover operators. The basic performance of these operators are quite similar,
 398 with a limited difference in both ϵ and in the number of generations to reach convergence. We also observe the
 399 opposite behavior between the scenarios where the load balancing is more important than the search for the
 400 closest fog nodes ($\delta\mu = 0.01$ and $\delta\mu = 0.1$) compared to the opposite case where distance reduction through
 401 optimized topology plays an important role in performance ($\delta\mu = 1$). However, due to the reduced impact of this
 402 parameter, we consider as the best option the uniform crossover, due to its fast and simple implementation and
 403 due to its stable performance over the different scenarios.

404 4.5. Sensitivity to population size

405 As an additional analysis we evaluate how the population size affects the performance of the genetic
 406 algorithm. To this aim we change the population size from 50 to 500 individuals and we measure the difference
 407 in the objective function ϵ , the generations required to reach convergence and the execution time considered as
 408 the time elapsed until the algorithm reaches the convergence criteria (the main results are shown in Fig. 7). For
 409 each metric also show the confidence intervals, represented with error bars in the graphs of Fig. 7.

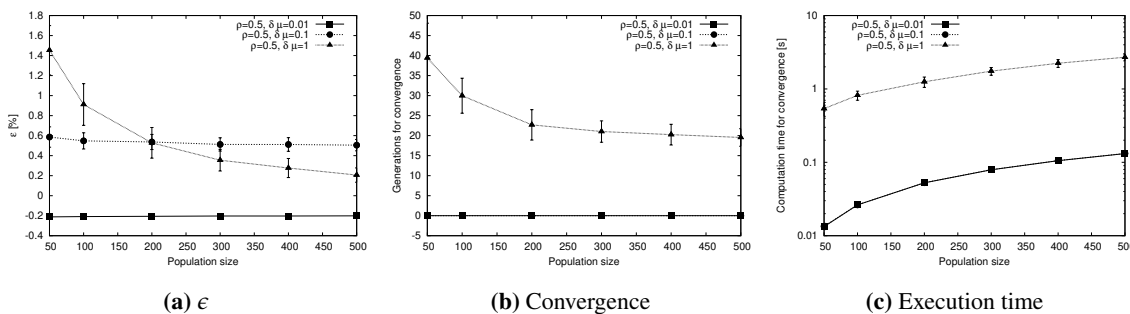


Figure 7. Sensitivity to population size

410 The first significant result, shown in Fig. 7a comes from the evolution of ϵ with the population size: in this
 411 case we show that, for the more challenging scenario $\delta\mu = 1$, increasing the population has a positive impact as
 412 it helps reducing the best value of ϵ reached. On the other hand, the other scenarios $\delta\mu = 0.01$ and $\delta\mu = 0.1$
 413 show that the impact of the population size is less evident, as the exploration of the solution space improving the
 414 existing solutions has a limited effect. In a similar way, the behavior of the other main metric, that is the number
 415 of generations to reach convergence, remains stable for the scenarios $\delta\mu = 0.01$ and $\delta\mu = 0.1$, while drops with
 416 the population when $\delta\mu = 1$, due to the more efficient exploration of the solution space when the population is
 417 higher.

418 Another interesting result, shown in Fig 7c, concerns the time to execute the genetic algorithm until
 419 convergence is reached (as the result of this experiment provides values spanning over multiple orders of
 420 magnitude, we use a logarithmic scale for the y axis). We observe that, for $\delta\mu = 0.01$ and $\delta\mu = 0.1$, the time
 421 to reach convergence grows linearly (the logarithmic scale results in a curve that is not a straight line) and
 422 corresponds with the setup time of the algorithm, as convergence is reached in the first generation, as shown in
 423 Fig. 7b. For the case where $\delta\mu = 1$ the number of generations to convergence decreases with the population
 424 size. However, this effect is not enough to compensate the higher computation time needed to handle a larger
 425 population, resulting in a curve that is monotonically increasing with the population size, even if the growing rate
 426 is less evident in the first part of the graph (i.e for populations of 200 individuals or less) compared with the other
 427 curves.

428 4.6. Sensitivity to mutation and crossover probability

429 As a final analysis, we consider the sensitivity of the genetic algorithm to the two probabilities that define
 430 the evolution of the population, that is the mutation probability P_{mut} and the crossover probability P_{cx} .

431 Concerning the mutation probability, Fig. 8 shows how this parameter affects the ability of the algorithm to
 432 reach a value close to the solver-based output through the ϵ metric (Fig. 8a) and how many generations it takes to
 433 converge (Fig. 8b). We recall that, as the scenarios $\delta\mu = 0.01$ and $\delta\mu = 0.1$ achieve a value of $\epsilon < 1\%$ in the

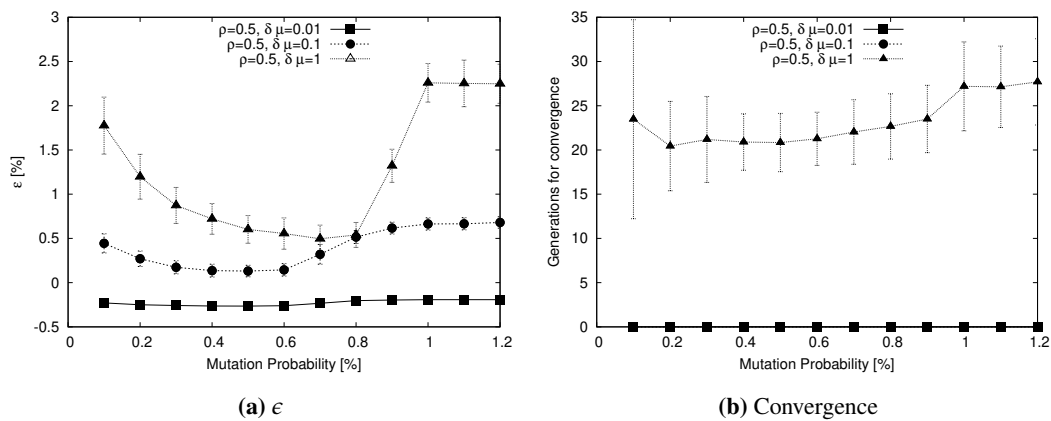


Figure 8. Sensitivity to Mutation Probability P_{mut}

434 first generation, the number of generations to reach convergence is not meaningful in these cases. Focusing on
 435 the ϵ metric, we observe for the scenarios a U-shaped curve that is more evident as $\delta\mu$ increases. In this curve
 436 both low values ($P_{mut} \leq 0.2\%$) and high values ($P_{mut} \geq 0.9\%$) result in poor performance while values in the
 437 range $0.4\% \leq P_{mut} \leq 0.8\%$ result in low values of ϵ .

438 This behavior is explained considering the two-fold impact of mutations. On one hand, a low value of P_{mut}
 439 hinders the ability to explore the solutions space by creating variations in the genetic material. On the other hand,
 440 an higher mutation rate may simply reduce the ability of the algorithm to converge, because the population keeps
 441 changing too rapidly and good genes cannot be passed through the generations. The low values of P_{mut} has
 442 an interesting effect on the convergence curve for $\delta\mu = 1$ in Fig. 8b: low mutation rates result in a poor ability to
 443 explore the solution space, with a high variance in the number of generations to reach convergence as this value
 444 depends significantly on the initial population setup.

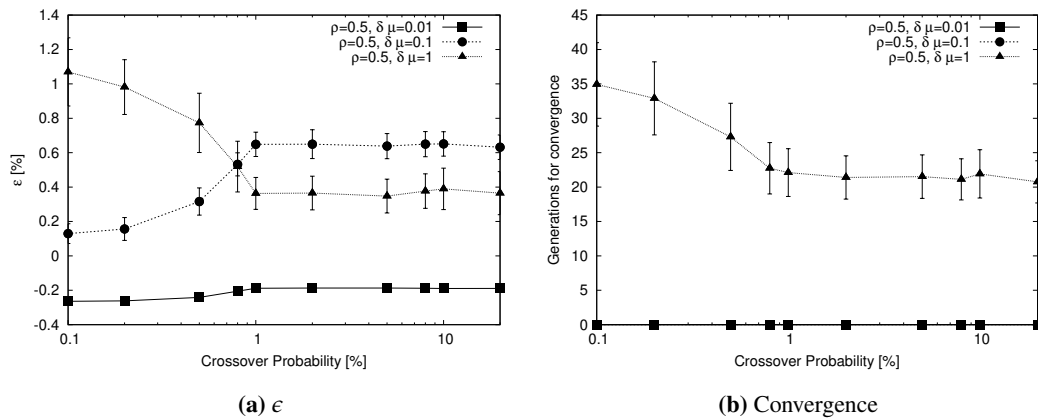


Figure 9. Sensitivity to crossover probability P_{cx}

445 The analysis shown in Fig. 9 evaluates the impact of the probability of selecting an individual for a crossover
 446 operation P_{cx} . Again we show both ϵ (Fig. 9a), and the generations to reach convergence (Fig. 9b) as a function
 447 of this parameter. Since we change the probability over a large range of values (from 0.1% to 20%), we use
 448 a logarithmic scale for the x -axis. From the analysis, we observe that also crossover probability has a major
 449 impact on the performance of the genetic algorithm, but this parameter has a different effect on the algorithm
 450 depending on the considered scenario. For $\delta\mu = 1$, low values of P_{cx} reduce the ability of the algorithm to
 451 converge rapidly, while for higher values (that is $P_{cx} \geq 0.8\%$) this behavior does not occur. The reason for this is
 452 that a low crossover probability hinders the distribution of good genes in the population, thus slowing down the
 453 improvement of the population. For $\delta\mu = 0.01$ and $\delta\mu = 0.1$, the effect is opposite because for the exploration
 454 of the solution space, mutation is more important than crossover and a high crossover rate interferes with the
 455 mutation operator (when crossover is applied mutation does not occur) hindering its action.

5. Related work

The explosive growth in the generation of data and the need for their processing to provide innovative services and applications has recently led researchers to focus on fog computing solutions to complement the cloud systems capabilities. To always exchange localized data from and to the remote cloud, indeed, tends to be inefficient under different points of view, thus motivating fog computing to partially process workload and data locally on fog nodes [3,5,22,23].

A survey discussing representative application scenarios and identifying various issues related to design and implementation of fog computing systems can be found in [3], while the study in [23] provides an overview of the core issues, challenges, and future research directions in fog-enabled orchestration for IoT services, focusing on smart cities as main motivating example of the research. Also our study considers the smart cities as a meaningful scenario where large amount of sensors and smart devices produce a huge volume of data on a geographically distributed area. Specifically, we focus on the specific issue of distributing the incoming workload over the fog nodes to minimize communication latency while avoiding overload.

Some existing studies focus on the issue of allocating the processing tasks coming from the fog nodes to the cloud nodes to optimize performance and reduce latency. Among these studies, Deng et al. [5] explore the tradeoff between power consumption and transmission delay in the fog-cloud computing system, formulating an optimization of the allocation problem among fog and cloud nodes. The study in [6] explicitly focuses on the issue of minimizing the service delay in IoT-fog-cloud application scenarios, proposing a delay-minimizing policy for fog nodes: in contrast to other proposals in literature, the proposed policy employs fog-to-fog communication to reduce the service delay by sharing load. Similarly, in [24] an offloading mechanism is proposed where a fog-to-fog collaboration based on the FRAMES load balancing scheme aims to reduce the overall latency experienced by the services. It is worth to note that these studies do not consider the issue of mapping data sources on the fog nodes: the fog nodes directly communicate with the mobile users through single-hop wireless connections using the off-the-shelf wireless interfaces (e.g., WiFi, Bluetooth, LR-WPANs, etc.), or the considered scenario envisions a domain of IoT nodes (in a factory, for instance) that communicate with a domain of fog nodes, associated with the specific domain application(s). On the other hand, our study focuses on the issue of optimizing the mapping of the workload coming from data sources over the fog nodes.

Few studies assume a flexible mapping of data sources over the fog nodes. The study in [25] proposes a framework to design services for smart buildings based on edge and fog computing paradigms, where the IoT sensors communicate data through the MQTT protocol: this study assumes a potential flexible mapping for design purposes but does not propose any specific (optimized) solution to address this issue. The studies in [26] and in [27] propose a coordination scheme between cloud and fog nodes applied to a healthcare-driven IoT application and to a real-time streaming application, respectively. In these studies, the data sources can choose to connect to the fog node or directly to the cloud data center according to specific conditions. However, in our solution the sensors always send data to an intermediate fog node to be selected within the fog layer to optimize the service performance.

Among the studies focusing on fog computing applied to the same context of our paper, in [22] a hierarchical 4-layer fog computing architecture is proposed for big data analysis in smart cities. The layered fog computing network exploits the natural characteristic of geo-distribution in big data generated by massive sensors, performing latency-sensitive applications and providing quick control loop to ensure the safety of critical infrastructure components. In this paper, the mapping between fog nodes and sensors is fixed: each fog node is connected to and responsible for a local group of sensors that cover a neighborhood or a small community.

The study in [28] considers Data Stream Processing (DSP) applications and, specifically, the so called operator placement problem, that is the allocation of DSP operator on fog nodes with the goal of optimizing the applications Quality of Service (QoS). The optimal DSP placement is modeled as an Integer Linear Programming (ILP) problem. In this case the authors made the assumption that it is possible to split the incoming data flow for parallel processing, while we consider generic applications where this assumption may not be true.

As regards the use of Genetic Algorithms (GAs), they have been successfully applied to the context of cloud computing in recent literature. The study in [7] exploits GAs to produce a suitable and scalable solution for the Software as a Service (SaaS) Placement Problem [7], while Karimi et al. [29] proposes a QoS-aware service

506 composition for cloud computing systems based on GAs. A previous version of the paper by the same authors
507 was presented in [8], proposing the use of a GA-based heuristic to map data flows over the fog nodes. However,
508 this paper represents a clear step ahead with respect to the previous work, with improvements regarding both the
509 theoretical contribution and the experimental setup, including new scenarios and sensitivity analysis.

510 6. Conclusions and future work

511 This paper presents a solution based on fog computing to address the issues of the typical scenario of a
512 smart city where sensors or smart devices disseminated over a geographic area produce a large amount of data.
513 We pointed out that a traditional cloud infrastructure, with all data flows converging on a single cloud data center
514 (or, at most, on few data centers) is at risk of network congestion. Moreover, as some applications in such
515 scenario are latency-sensitive (e.g. applications related to automated traffic management) or produce a bulk of
516 data that could create congestion at the network level (e.g., widespread sensors for environmental analysis) the
517 most suitable approach is to exploit a layer of intermediate fog nodes located as close as possible to the data
518 sources to perform pre-processing (e.g., filtering and aggregation) or latency-critical tasks.

519 The innovative paradigm of fog computing opens several new issues. In this paper we focus on the problem
520 of how to map the data streams produced by the sensors over the fog nodes. We provided a formal model for
521 the problem of minimizing the overall latency experienced in the system, considering both data transfer and
522 processing times. Furthermore, we proposed an heuristic algorithm, based on genetic programming to solve the
523 problem without the need to rely on an external solver.

524 The proposed solution is evaluated in the context of a realistic smart-city scenario. The experiments show
525 the excellent performance of the proposed genetic algorithm to solve the mapping problem. Moreover, different
526 evolutionary strategies and genetic operators are compared to identify the best performing one in the considered
527 scenario. Finally, we demonstrate the stability of the proposed heuristic through a sensitivity analysis on its main
528 parameters, such as the number of generations, the probability of mutation and crossover, and the population size.

529 As future work, we plan to extend the current research taking into account more complex scenarios that
530 involve dynamic changes in the workload, for example to introduce mobility in the data sources or to consider
531 adaptive sampling techniques at the sensor level producing dynamic outgoing data rate. To support these news
532 scenarios, we plan to provide contributions not only in terms of the definition of the topology but also at the level
533 of adaptive algorithms proposal.

534

- 535 1. Liu, J.; Li, J.; Zhang, L.; Dai, F.; Zhang, Y.; Meng, X.; Shen, J. Secure intelligent traffic
536 light control using fog computing. *Future Generation Computer Systems* **2018**, *78*, 817 – 824.
537 doi:<https://doi.org/10.1016/j.future.2017.02.017>.
- 538 2. Sasaki, K.; Suzuki, N.; Makido, S.; Nakao, A. Vehicle control system coordinated between cloud and mobile edge
539 computing. 2016 55th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE), 2016,
540 pp. 1122–1127.
- 541 3. Yi, S.; Li, C.; Li, Q. A Survey of Fog Computing: Concepts, Applications and Issues. Proceedings
542 of the 2015 Workshop on Mobile Big Data; ACM: New York, NY, USA, 2015; Mobidata '15, pp. 37–42.
543 doi:10.1145/2757384.2757397.
- 544 4. Dastjerdi, A.V.; Buyya, R. Fog Computing: Helping the Internet of Things Realize Its Potential. *Computer* **2016**,
545 *49*, 112–116.
- 546 5. Deng, R.; Lu, R.; Lai, C.; Luan, T.H.; Liang, H. Optimal Workload Allocation in Fog-Cloud Computing Toward
547 Balanced Delay and Power Consumption. *IEEE Internet of Things Journal* **2016**, *3*, 1171–1181.
- 548 6. Yousefpour, A.; Ishigaki, G.; Jue, J.P. Fog Computing: Towards Minimizing Delay in the Internet of Things. 2017
549 IEEE International Conference on Edge Computing (EDGE), 2017, pp. 17–24. doi:10.1109/IEEE.EDGE.2017.12.
- 550 7. Yusoh, Z.I.M.; Tang, M. A penalty-based genetic algorithm for the composite SaaS placement problem in the Cloud.
551 IEEE Congress on Evolutionary Computation, 2010, pp. 1–8. doi:10.1109/CEC.2010.5586151.
- 552 8. Canali, C.; Lancellotti, R. A Fog Computing Service Placement for Smart Cities based on Genetic Algorithms. Proc.
553 of International Conference on Cloud Computing and Services Science (CLOSER 2019); , 2019.

- 554 9. Bigi, A.; Veratti, G.; Fabbi, S.; Ziven, O.; Po, L.; Ghermandi, G. Forecast of the impact by local emissions at an urban
555 micro scale by the combination of lagrangian modelling and low cost sensing technology: the TRAFair project.
556 Proc. of 19th International conference on Harmonisation within Atmospheric Dispersion Modelling for Regulatory
557 Purposes; , 2019.
- 558 10. Shojafar, M.; Canali, C.; Lancellotti, R. A Computation- and Network-Aware Energy Optimization Model for Virtual
559 Machines Allocation. Proc. of International Conference on Cloud Computing and Services Science (CLOSER 2017);
560 , 2017.
- 561 11. Shojafar, M.; Canali, C.; Lancellotti, R.; Abolfazli, S. An Energy-aware Scheduling Algorithm in DVFS-enabled
562 Networked Data Centers. Proceedings of the 6th International Conference on Cloud Computing and Services Science
563 - Volume 1 and 2, 2016, CLOSER 2016.
- 564 12. Noshay, M.; Ibrahim, A.; Ali, H. Optimization of live virtual machine migration in cloud computing: A survey and
565 future directions. *Journal of Network and Computer Applications* **2018**, *110*, 1–10. doi:10.1016/j.jnca.2018.03.002.
- 566 13. Duan, H.; Chen, C.; Min, G.; Wu, Y. Energy-aware scheduling of virtual machines in heterogeneous cloud computing
567 systems. *Future Generation Computer Systems* **2017**, *74*, 142 – 150. doi:https://doi.org/10.1016/j.future.2016.02.016.
- 568 14. Ardagna, D.; Ciavotta, M.; Lancellotti, R.; Guerriero, M. A Hierarchical Receding Horizon Algorithm for
569 QoS-driven control of Multi-IaaS Applications. *IEEE Transactions on Cloud Computing* **2018**, pp. 1–1.
570 doi:10.1109/TCC.2018.2875443.
- 571 15. *Knitro Website*. Available at <https://www.artelys.com/solvers/knitro/>, last accessed on 10th Jul 2019.
- 572 16. Canali, C.; Lancellotti, R. Scalable and automatic virtual machines placement based on behavioral similarities.
573 *Computing* **2017**, *99*, 575–595. doi:10.1007/s00607-016-0498-5.
- 574 17. Binitha, S.; Sathya, S.S.; others. A survey of bio inspired optimization algorithms. *International Journal of Soft
575 Computing and Engineering* **2012**, *2*, 137–151.
- 576 18. Back, T.; Fogel, D.; Michalewicz, Z. *Evolutionary Computation 1: Basic Algorithms and Operators*; CRC Press,
577 2002.
- 578 19. DEAP: Distributed Evolutionary Algorithms in Python, 2018. – <https://deap.readthedocs.io>.
- 579 20. Cicirello, V.A.; Smith, S.F. Modeling GA Performance for Control Parameter Optimization. Proceedings of the 2nd
580 Annual Conference on Genetic and Evolutionary Computation; Morgan Kaufmann Publishers Inc.: San Francisco,
581 CA, USA, 2000; GECCO'00, pp. 235–242.
- 582 21. AMPL: Streamlined Modeling for Real Optimization, 2018. – <https://ampl.com/>.
- 583 22. Tang, B.; Chen, Z.; Hefferman, G.; Wei, T.; He, H.; Yang, Q. A Hierarchical Distributed Fog Computing Architecture
584 for Big Data Analysis in Smart Cities. Proceedings of the ASE BigData & SocialInformatics 2015; ACM: New York,
585 NY, USA, 2015; ASE BD&SI '15, pp. 28:1–28:6. doi:10.1145/2818869.2818898.
- 586 23. Wen, Z.; Yang, R.; Garraghan, P.; Lin, T.; Xu, J.; Rovatsos, M. Fog Orchestration for Internet of Things Services.
587 *IEEE Internet Computing* **2017**, *21*, 16–24. doi:10.1109/MIC.2017.36.
- 588 24. Al-khafajiy, M.; Baker, T.; Al-Libawy, H.; Maamar, Z.; Aloqaily, M.; Jararweh, Y. Improving fog computing
589 performance via Fog-2-Fog collaboration. *Future Generation Computer Systems* **2019**, *100*, 266 – 280.
- 590 25. Ferrández-Pastor, F.J.; Mora, H.; Jimeno-Morenilla, A.; Volckaert, B. Deployment of IoT Edge and Fog Computing
591 Technologies to Develop Smart Building Services. *Sustainability* **2018**, *10*.
- 592 26. Maamar, Z.; Baker, T.; Faci, N.; Ugljanin, E.; Khafajiy, M.A.; Burégio, V. Towards a Seamless Coordination of
593 Cloud and Fog: Illustration through the Internet-of-Things. Proc. of the 34th ACM/SIGAPP Symposium on Applied
594 Computing, 2019, SAC '19.
- 595 27. Nair, B.; Saira Bhanu, S.M. Fog-Cloud Collaboration for Real-Time Streaming Applications: FCC for RTSAs. In
596 *Handbook of Research on the IoT, Cloud Computing, and Wireless Network Optimization*; IGI Global, 2019.
- 597 28. Cardellini, V.; Grassi, V.; Lo Presti, F.; Nardelli, M. Optimal Operator Placement for Distributed Stream Processing
598 Applications. Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems; ACM:
599 New York, NY, USA, 2016; DEBS '16, pp. 69–80. doi:10.1145/2933267.2933312.
- 600 29. Karimi, M.B.; Isazadeh, A.; Rahmani, A.M. QoS-aware Service Composition in Cloud Computing Using Data
601 Mining Techniques and Genetic Algorithm. *J. Supercomput.* **2017**, *73*, 1387–1415. doi:10.1007/s11227-016-1814-8.