

An Efficient Monte Carlo-based Probabilistic Time-Dependent Routing Calculation Targeting a Server-Side Car Navigation System

Emanuele Vitali, Davide Gadioli, Gianluca Palermo, *Member, IEEE*,
 Martin Golasowski, João Bispo, Pedro Pinto, Jan Martinovič, Kateřina Slaninová,
 João M. P. Cardoso, *Senior Member, IEEE*, Cristina Silvano, *Fellow, IEEE*

Abstract—Incorporating speed probability distribution to the computation of the route planning in car navigation systems guarantees more accurate and precise responses. In this paper, we propose a novel approach for dynamically selecting the number of samples used for the Monte Carlo simulation to solve the Probabilistic Time-Dependent Routing (PTDR) problem, thus improving the computation efficiency. The proposed method is used to determine in a proactive manner the number of simulations to be done to extract the travel-time estimation for each specific request while respecting an error threshold as output quality level. The methodology requires a reduced effort on the application development side. We adopted an aspect-oriented programming language (LARA) together with a flexible dynamic autotuning library (mARGOT) respectively to instrument the code and to take tuning decisions on the number of samples improving the execution efficiency. Experimental results demonstrate that the proposed adaptive approach saves a large fraction of simulations (between 36% and 81%) with respect to a static approach while considering different traffic situations, paths and error requirements. Given the negligible runtime overhead of the proposed approach, it results in an execution-time speedup between 1.5x and 5.1x. This speedup is reflected at infrastructure-level in terms of a reduction of around 36% of the computing resources needed to support the whole navigation pipeline.

Index Terms—High Performance Computing, Approximate Computing, Adaptive Applications, Smart Cities, Vehicle Routing

1 INTRODUCTION

IN smart cities, the trend is to combine and automate several common tasks to ease the life of citizens. Among these tasks, traffic estimation and prediction plays a central role: it is used not only to avoid traffic congestion, which allows to have predictable travel times, but also to reduce car emissions. Considering the rising wave of self-driving cars, the amount of car navigation requests will increase rapidly together with the need of real-time updates and processing on large graphs representing the urban network. This trend imposes larger and more powerful computing infrastructures composed of HPC resources.

Concerning the algorithmic problem, car navigation is one of the main problems of applied theoretical research. The Dijkstra's shortest path algorithm is used for finding the optimal path between two vertices in a weighted graph representing a road network. Apart from single navigation between two points, navigation algorithms are used in various systems for solving larger optimization problems, like route planning for a fleet of package delivery vehicles, waste collection management or traffic optimization in a smart city [1]. Definition of the optimal path is based on the

type of used weights of the graph edges. The shortest path is based on the geographical distance between two adjacent vertices of a graph. The fastest path is based on time needed to cross a particular edge. There might be more complex criteria, however their description is out of the scope of this paper. Time needed to cross a particular stretch of road can be affected by various elements, such as accidents, traffic congestion, road work and so on. At the basic level, upper legal limit of speed is used, based on the assumption that each vehicle travels with the same speed. This can be vastly inaccurate due to the natural behavior of traffic.

With increasing availability of historical traffic monitoring data, there are several research efforts to determine average speed on road networks by using statistical analysis and various models. However, a single speed value is still not very useful as it does not reflect the stochastic behavior of the traffic. The probability distribution of the speed at a certain time enables to incorporate low probability real world events that can cause major delays and affect traffic over vast areas. By incorporating probability distribution to the computation, the system can compute the probability of arrival time within a certain time-frame which can be useful for more precise route planning. This problem is called *Probabilistic Time-Dependent Routing* (PTDR).

A scalable algorithm for solving the PTDR problem based on Monte Carlo simulations has been presented in [2] [3] and represents the base for our work. In particular, the algorithm uses probability distributions of travel time for the individual graph edges to estimate the distribution of the total travel time and it is integrated to an experimental

- E. Vitali, D. Gadioli, G. Palermo, and C. Silvano are with Dipartimento di Elettronica, Infomazione e Bioingegneria, Politecnico di Milano, Italy.
- J. Bispo, P. Pinto and J. Cardoso are with Faculty of Engineering (FEUP), University of Porto, Portugal.
- Martin Golasowski, Jan Martinovič and Kateřina Slaninová, are with IT4Innovations, VŠB - Technical University of Ostrava, Czech Republic.

Manuscript received July 31, 2018; revised Month XX, 201X.

server-side routing service. This service is deployed on an HPC infrastructure to offer optimal performance for a large number of requests as needed by the smart city context. The PTDR algorithm employed in this work simulates a large number of vehicles driving along a determined path in a graph at a particular time of departure. The speed of vehicles on individual roads is sampled from the speed probability distribution (also called speed profile) associated to the graph edge. The number of samples is a parameter that directly affects the informational value of the output as well as its computational requirements. Given the large amount of requests to be served, even small changes in the workload can affect the overall HPC system efficiency. While the original version was based on a worst-case tuning of the number of samples [2], and given that a reactive approach [4] is not a viable solution due to the overheads, in this paper we present a proactive method for dynamically adapting the number of samples for the Monte Carlo (MC) based PTDR algorithm.

In particular, the main contributions of this paper can be summarized as follows:

- We propose a methodology for self-adapting the PTDR algorithm presented in [2] [3] to the input data in a proactive manner, maximizing its performance while respecting the output quality level;
- We propose a probabilistic error model used to correlate the input data characteristics with the number of samples used by the Monte Carlo algorithm;
- We adopted an aspect-oriented programming language to keep separated the functional version of the application from the code needed to introduce the adaptivity layer.

The remainder of this paper is organized as follows. Section 2 provides an overview of the related works, while Section 3 provides an introduction to the Monte Carlo approach to solve the PTDR problem. Section 4 and Section 5 describe the proposed methodology respectively from the adaptivity and code integration point of view. Finally, Section 6 describes the experimental campaign to validate the proposed methodology and Section 7 concludes the paper.

2 RELATED WORK

Determining the optimal path in a stochastic time-dependent graph is a well-studied problem which has many formulations [5]. Our approach is closest to the *Shortest-path problem with on-time arrival reliability* (SPOTAR) formulation. It can be seen as a variant of the *Stochastic on-time arrival* (SOTA) problem, for which practical solution exists as shown in [6]. These algorithms have the objective of maximizing the probability of arriving within a time budget and are related to optimal routing in stochastic networks. However, there are not many solutions for the time-dependent variant of both of the problems. In [5] authors show practical results for the time-dependent variant of SOTA, simultaneously in [7] authors elaborate on the complexity of existing theoretical solutions of the SPOTAR problem and show how it can be extended with time dependency. There are many other papers which show various theoretical approaches for the SOTA problem, including some practical applications

[6] [8] [9] [10]. Solution to the SPOTAR problem based on *policy-based* SOTA as a heuristic is presented [9]. However, the authors make the assumption that the network is time-invariant, which is not true in real cases if considering long paths. The solution is also unusable in on-line systems as its scalability to graphs representing real-world routes is not sufficient.

Our approach follows the same philosophy presented in [2], [3] where the authors provides an approximate solution of the time-dependent variant of the SPOTAR problem based on Monte Carlo simulations. As shown in Section 3, our approach uses the k-shortest paths algorithm [11] [12] [13] to determine the paths for which the travel time distribution is estimated. This separation allows us to implement the approach in an online system which provides adaptive routing in real-time. Given the Monte Carlo nature of the algorithm, to improve the efficiency of the PTDR calculation we have two main alternatives [14]. The first is the sampling efficiency, while the second is the sampling convergence. In both cases, the algorithm optimization is reached by exploiting the iterative nature of the Monte Carlo simulation. To obtain the sampling efficiency, several techniques have been proposed to determine what is the next sample to be evaluated to maximize the gathered knowledge [15], [14]. However, in the implementation under analysis this has been discarded because our goal is to exploit the parallelism of the underlying HPC architecture [3] that excludes any iterative approach to the Monte Carlo. For the same reason also the approaches that require a statistical property evaluation after every iteration [4], checking if the error is acceptable, cannot be considered acceptable. Both approaches would be too time-consuming and, how it has been already analyzed in [2], for the specific problem the number of samples has to be chosen a priori in a proactive rather than in a reactive manner.

In this paper, we adopted autotuning techniques to face with the Monte Carlo efficiency problem. Classical autotuning approaches are based on code refactoring and loop parameterization to optimize extra-functional properties (e.g. power and performance) on a given target architecture (e.g. A-tune [16], SPIRAL [17], ATLAS [18], PowerDial [19]). More related to a Monte Carlo application, in [20] the authors developed a framework capable to automatically select for a Monte Carlo application the optimal parameters for GPU accelerated kernels, e.g. block size, thread numbers and other variables of the CUDA runtime. All those approaches are orthogonal to the proposed work that is more focused on reducing the number of samples, rather than tailoring the computation for a specific platform.

Two interesting works facing the autotuning problem considering the output accuracy, also on Monte Carlo simulations, are those presented in [21] and [22]. In both cases, they employ an ad-hoc DSL language as front-end for the developer, letting the compilation framework to customize the code and related execution. In particular, in [21], the authors propose a framework for driving the choice of extra-functional properties like the number of required processors, frequency and computation accuracy targeting the energy minimization. More related to tailoring the computation to the specific platform is the work in [22] where the authors enhanced the framework to include

compile time performance-accuracy trade-off. The proposed approach tries to face the problem from a different perspective. First, we intend to reduce as much as possible the intrusiveness on the target code, thus reducing also the load on the application developer, and second, our approach relies more on run-time and data-aware decisions, also avoiding compile-time decisions on performance-accuracy trade-offs.

A two-step approach for solving the Monte Carlo problem has been envisioned in [23]. Similar to our work, the authors suggest to have a first shot of a reduced number of samples to provide an initial approximate solution as fast as possible, and then to refine the output to the required accuracy in successive iterations. In the proposed context, this idea suffers from two main problems. First, it is suitable for scientific work-flows where an intermediate solution is used to trigger next computations, and it is not our case. Second, in the iterative phase, it suggests a reactive approach rather than a proactive one, that we already discussed to be necessary for the specific PTDR problem. Another two step approach is the one presented in [24]. In this case, the authors suggest building from each input a *small canary* – a statistically representative subset of the actual input – that is used to perform a parameter exploration at runtime. In this specific PTDR case, the approach suffers from a large overhead due to the canary extraction and approximation selection. In the proposed approach, we moved the two operations at design time by generalizing the approximation selection according to the *unpredictability function* computed at runtime. Moreover, the proposed approach has been designed to statistically guarantee the respect of an output error level.

Finally in [25], the authors present a proactive control of the application tuning as a constrained optimization problem that can be solved at run-time according to profiling information. The authors suggest using a Bayesian model to learn the accuracy of the computation according to data features that are mainly related to the data size. Despite the similarity of this previous work with one of the components we used for the integration (i.e. the mARGOt autotuner), the proposed approach is mainly related to the data-aware error prediction model for PTDR, where SLA constraints are considered, and to the seamless integration of the adaptivity concepts.

3 MONTE CARLO APPROACH FOR PROBABILISTIC TIME-DEPENDENT ROUTING

So far, many theoretical formulations and several algorithms have been developed for solving the problem of computing the travel time distribution [5]. In this paper we consider a *path-based* approach (SPOTAR) where the paths are known *a-priori* and travel-time distributions are determined subsequently for each one of the paths [26].

In the context of the complete traffic navigator application illustrated in Figure 1, our focus is on the efficient estimation of the arrival time distribution (PTDR - *Probabilistic Time-Dependent Routing* phase). More in detail, the three main steps of the application can be described as follows: (i) the first step consists of determining K *alternative paths* to be passed to the next steps. In the navigation scenario,




Fig. 1: The complete navigation infrastructure for serving a single request.

the identification of the shortest path is not enough to determine a good solution, if no traffic information have been considered. Thus alternative routes derived by algorithms for determining k -Short Paths with limited overlap have to be adopted in this step [12] [13] [11]. This first phase is out of the scope of this paper; (ii) For every path selected by the previous step (K -alternative paths), the computation of the travel time is done using the Probabilistic Time-Dependent Routing module. While the exact solution to the travel-time estimation (PTDR) has an exponential complexity, in this work we efficiently approximate the solution of the SPOTAR problem by adopting a Monte Carlo sampling approach [2]; (iii) The final step gathers the timing information provided by the k instances of the PTDR module for every single request and selects the best path to be given back to the user. Actually, this phase does not provide a single route but reorders the list of k paths determined by the first step according to the timing distributions determined in the second phase and user preference [27].

This three-step approach of the whole navigation application allows us to implement an approximate solution to the SPOTAR problem, which can be used online in a system to serve a large volume of routing requests.

Our definition of a probabilistic road network is similar to the definition of the stochastic time-dependent network as described by Miller-Hooks [26], with the exception of the segment travel times, which has been substituted by the speed probability distribution (*speed profile*) for a given time of departure within a week. Formally, it can be defined as follows. Let $G = (V, E)$ be a well connected, directed and weighted graph, where V is the set of vertices and E is the set of edges. Each vertex represents a junction or some important point corresponding to geospatial properties of the road, while edges represent the individual road segments between the junctions. Each path selected by the first phase of the application (i.e. K -Alternative paths) can be formally represented as a vector of graph edges $S = (s_1, s_2, \dots, s_n)$, while $S_p \subseteq E$ and n is the number of road segments in the path.

Using a travel time estimation function, we are interested in estimating the travel time θ as $\hat{\theta}_{S,t,P_S}$ where S is the given path, t is the departure times and P_S are the probabilistic speed profiles for the segments in S . More in detail, $t \in T$ is a departure time within a set of possible departure times which divide a certain timeframe to a set of intervals $T = \{t : t = n \cdot \phi, n \in \mathbb{N}\}$ [28], where the length of the interval ϕ is determined by input data. P is the set of probabilistic speed profiles for the entire graph edges E , where $P_S \subseteq P$. Each speed profile $p \in P$ is represented by a set of discrete speed values and assigned probabilities. The number of speed values depends on the method used for deriving the profiles from historical traffic monitoring data, while the minimum and maximum values represent

respectively the congestion speed and the free flow speed.

In our work, the time frame is set to *one week* and $\phi = 900s$ (15 minutes). This approach reflects traffic variations during the various hours of the day and for all the days of a week. By extending the time frame, other factors can be included, such as seasons or holidays. The number of speed values has been set to 4 levels according to the characteristics of the input data used for the creation of the speed profiles.

Focusing on the SPOTAR problem, we are not interested in a single travel time value θ but we require to calculate the probability distribution of the arrival time. Given the previous formalization of the problem, the travel time distribution can be estimated by traversing the path segments together while considering the speed profile distribution. In particular, we can define a tree where each layer represents a segment in the selected path [2]. The tree root is the starting segment, while the end segment is on the leaves. Each node in all layers of the tree has l children, where l is a number of the discrete speed values for each segment, and the tree depth corresponds to the number of selected path segments $|S|$. Each edge in the tree is annotated by the discrete speed value, its probability, and by the length of the considered segment. A travel time can be computed by a depth-first search (DFS) while selecting an arbitrary child node at each level of the tree. The travel time value is then the sum of the time spent in each segment (length/speed), while the probability of that value is the product of the probability on each edge of the traversal. Each traversal corresponds to a single car traveling along the entire path. The exact solution is obtained by an exhaustive search over all the possible paths between the root node and all the leaves. This approach is clearly not efficient since it scales exponentially with the number of segments in the path.

A Monte Carlo-based approach can be successfully employed in this case. By generating a large number of random tree traversals, enough samples can be obtained to estimate the final distribution. We define this final distribution, which is a collection of θ values ($\theta_1 \dots \theta_x$) obtained through the Monte Carlo simulation $MCS(x, i)$, where x is the number of random tree traversals, and i is the input set of the $\hat{\theta}$ function (i.e. S, t, P_S).

Given that travel times usually have a long tailed distribution due to inherent properties of the traffic (e.g. rare events such as accidents) a large number of samples is needed to estimate the travel time distribution with a sufficient precision. Regarding the definition of the number of samples for the Monte Carlo simulation, the particular implementation of the PTDR kernel cannot rely on a runtime stability analysis of the output. Each tree traversal (a sample of the Monte Carlo simulation) is totally independent to the others, thus this problem is perfectly suitable for parallel computing architectures, such as modern CPUs or accelerators. To efficiently exploit this parallelism, it is necessary to know a-priori the number of travel time estimations required to build the final distribution.

To summarize, the PTDR algorithm can be seen as in Figure 2, where all the information regarding the request are provided to the Monte Carlo simulation (MCS) capable to return the predicted travel time distribution for the given route.




Fig. 2: The original approach for PTDR routing based on Monte Carlo simulations to derive the travel time distribution.

4 THE PROPOSED APPROACH

The Monte Carlo simulation is designed to use a given number of samples x for every run. Based on a conventional approach, this number is selected according to the worst-case analysis and it is the lowest number of samples always able to reach a target precision [3]. In this section, we present the proposed technique adopted to select at runtime the number of samples for the Monte Carlo simulation according to the input data characteristics.

Before moving on the methodological part, let us better define the specific context of the problem. In particular, even if we are interested in the travel time distribution, our goal is to know a value τ_i to guarantee with a certain probability that the travel time will be within that value: $P(\theta < \tau_i) \geq y$ where i has been defined as the input set of the travel-time function. The value τ_i is the output of the PTDR phase. In the following, we characterize that value with an additional property $\tau_{i,y}$, where y is the probability that the travel time will be lower than τ .

Using the Monte Carlo simulation, we can estimate the value of $\tau_{i,y}$ using x samples as follows $\hat{\tau}_{i,y}^x = MCS(x, i, y)$. In particular, we estimate the value $\hat{\tau}_{i,y}^x$ by selecting the y -th percentile of the finite-sample distribution obtained from the Monte Carlo simulation (i.e. if $y = 95\%$ then $\hat{\tau}_{i,y}^x$ is the 95th percentile of the distribution).

In the context of this work, we are interested in minimizing the execution time of the function MCS , while limiting the prediction error defined as $error_{i,y}^x = \frac{|\tau_{i,y} - \hat{\tau}_{i,y}^x|}{\tau_{i,y}}$. In particular, the target problem can be expressed as follows:

$$\begin{aligned} & \text{minimize} \quad cpu_time_x^x \\ & \text{subject to} \quad error_{i,y}^x \leq \epsilon \end{aligned} \quad (1)$$

where ϵ represents the upper bound on the computation error. We want this error to be relative to the output of the MCS, that is the desired percentile of the predicted travel time. In this way, we can abstract from the actual path. Given the tight correlation between the execution time and the used number of samples x , the previous problem can also be simplified by considering the minimization of x instead of the cpu_time . According to Monte Carlo properties, we can derive that $\tau_{i,y} \equiv \hat{\tau}_{i,y}^\infty$, where $\hat{\tau}_{i,y}^\infty$ is the output of the MCS function computed using an infinite number of samples. Thus, we can rewrite the error as

$$error_{i,y}^x = \frac{|\hat{\tau}_{i,y}^\infty - \hat{\tau}_{i,y}^x|}{\hat{\tau}_{i,y}^\infty} \quad (2)$$

Due to the Monte Carlo properties [29], the value $\hat{\tau}_{i,y}^x$ is a random variable, asymptotically normally distributed with mean $\mu_{\hat{\tau}_{i,y}^x}$ and standard deviation $\sigma_{\hat{\tau}_{i,y}^x}$. In particular, according to the central limit theorem [30], while considering enough values of samples the mean value does not depend on the number of Monte Carlo simulations, and the standard deviation decreases by increasing the number of Monte Carlo simulations. Given that, we can define the error as characterized by a normal distribution with mean 0 and a standard deviation $\sigma_{\hat{\tau}_{i,y}^x} / \mu_{\hat{\tau}_{i,y}^x}$. In the following, we refer to the standard deviation of the error as $\nu_{\hat{\tau}_{i,y}^x} = \frac{\sigma_{\hat{\tau}_{i,y}^x}}{\mu_{\hat{\tau}_{i,y}^x}}$. This expression is the same as the coefficient of variation (relative standard deviation) of the result of the Monte Carlo simulation.

According to the probabilistic nature of the problem, we cannot guarantee that the error will be always below ϵ . However, this can be done relaxing the error constraint by introducing a confidence interval (CI) level. In particular, given the normal distribution of the error, the selected confidence interval can be correlated with the expected error:

$$P(\text{error}_{i,y}^x \leq \epsilon) \geq CI \implies \text{error}_{i,y}^x \leq n(CI) \times \nu_{\hat{\tau}_{i,y}^x} \leq \epsilon \quad (3)$$

where $n(CI)$ is a value that express the confidence level (e.g. $n(68\%)=1$, $n(95\%)=2$ and $n(99.7\%)=3$ derived from the 1-3 σ -intervals of the normal distribution). Thus, if we decrease the number of Monte Carlo simulations used to derive $\hat{\tau}_{i,y}^x$, on one hand we decrease the execution time of the application, but on the other hand we are also reducing the accuracy of the results, having a larger value for the coefficient of variation $\nu_{\hat{\tau}_{i,y}^x}$.

An additional problem is derived from the fact that $\hat{\tau}_{i,y}^x$ is input dependent. This means that it is not possible to predict the possible Monte Carlo error for unknown paths, according to the number of samples. To deal with this, we found a feature u_i of the inputs i that can be used to quickly estimate the number of samples necessary to keep the error below the threshold ϵ . The idea is to evaluate the error by using u_i instead of the actual i so that we can transform the original problem as

$$\text{error}_{i,y}^x \leq n(CI) \times \nu_{\hat{\tau}_{u_i,y}^x}. \quad (4)$$

The feature u_i has been called *unpredictability*, since it represents a set of characteristics of the inputs i (road, starting time,...) that provides information about how complex is the prediction of $\tau_{i,y}$, therefore it is also related on how many samples are required to satisfy a certain error and confidence level. More details on the unpredictability feature are presented in Section 4.1.

Given that the error is not anymore related to the specific input set i but only to the feature u_i , the number of samples needed to satisfy the constraint can be easily extracted by $\nu_{\hat{\tau}_{u_i,y}^x} \leq \frac{\epsilon}{n(CI)}$. A profiling phase on a set of representative inputs can be used to then extract the values of $\nu_{\hat{\tau}_{u_i,y}^x}$, that will be used to determine the correlation between the unpredictability function and the error. More details on the profiling phase including the prediction function are presented in Section 4.2.

To summarize, the proposed methodology adds an adaptivity layer on top of the Monte Carlo simulation (see




Fig. 3: The proposed adaptive approach for PTDR routing based on Monte Carlo simulations and dynamic choice of the number of simulations.

Figure 3) to quickly determine at runtime the right number of samples for each request that satisfies the required accuracy. In particular, a feature-extraction procedure estimates the unpredictability value from the input data of the request (path, starting time and segment speed-profiles). The dynamic autotuner combines this data feature with the profiled knowledge and the extra-functional requirements to configure the Monte Carlo simulation.

4.1 Unpredictability Feature

Given that the extraction of the data feature from inputs should be done at runtime, its computation should not be a costly operation. Otherwise, the benefit of speeding up the computation phase by reducing the number of Monte Carlo samples would be reduced by the data feature extraction overhead, eventually making the whole approach meaningless.

From the experimental results, we have found that a measure of the unpredictability of the path can be extracted by a simple statistical property of the set of travel times θ extracted by a quick Monte Carlo simulation: the coefficient of variation. Intuitively the more the results are spread out, the more the route is hard to predict, thus to have a precise estimation of the distribution, and in particular the percentiles, we need a higher number of samples.

The unpredictability function is defined as $u_i = \sigma_{\theta_i}^x / \mu_{\theta_i}^x$ where σ_{θ_i} and μ_{θ_i} are evaluated on a MCS done with the minimum number x of samples allowed at runtime. It is important to note that σ_{θ_i} is the variance of the travel times extracted by a single Monte Carlo simulation on the minimum number of samples. We calculate the unpredictability function together with the first set of Monte Carlo samples to further reduce the overhead introduced by the data feature extraction. In particular, we will use this first short Monte Carlo run to determine if there is a need of further samples (and how many) to satisfy the error constraint.

To validate the usage of u instead of i , we performed the Spearman correlation test [31] between the unpredictability value and the value of $\nu_{\hat{\tau}_{i,y}^x}$ used in the calculation of the expected error for different values of x and y over a wide range of inputs sets i . In all cases, the correlation values were larger than 0.918 showing a p-value equal to 0. These correlations confirm our hypothesis and the p-values proves that the results are statistically significant.




Fig. 4: Integration flow outlining the two main LARA aspects and related actions: original code enrichment and autotuner configuration.

4.2 Error Prediction Function

To predict the expected error for a specific configuration according to the data feature u , we need to extract $\hat{\nu}_{\tau_{u_i, y}}^x$ from profiling data. We run the Monte Carlo simulation several times for each configuration in terms of number of samples. In particular, we decided to use values ranging from 100 samples up to 3000. The two numbers have been derived from the observation that 100 samples is the minimum to have the estimation of the percentile for the distribution, while 3000 is the number of samples that has already been found good enough to satisfy the worst case conditions on the previous work [32]. Between the two values, we selected 2 more sampling levels corresponding to 300 and 1000, that have been derived by considering that the Monte Carlo error decreases as $1/\sqrt{n}$ [33]. Thus in our case at each sampling level, we have that the error is almost halved.

We run each set of Monte Carlo simulations with the same configuration in terms of the number of samples on a large set of inputs i (i.e. roads, starting time ...), and we extract $\nu_{\tau_{u_i, y}}^x$ and u_i from each single configuration. Then we create a predictor $\hat{\nu}_{\tau_{u_i, y}}^x$ as the *quantile regression* [34] over the extracted data. The use of quantile regression enhances the robustness of the model in the context of its use. Indeed, we are not interested in predicting an average value as final result, but we want to use it for the inequality formula $\hat{\nu}_{\tau_{u_i, y}}^x \leq \frac{\epsilon}{n(CT)}$. In this case, a higher value of the quantile with respect to 50th (the purely linear regression), guarantees a higher robustness in satisfying the previous inequality. The quantile used for the regression is an additional parameter that can be explored to trade-off robustness and performance.

5 INTEGRATION FLOW

While previous section introduces the proposed methodology from the end-user perspective, thus considering execution time and elaboration error, this section focuses on the application developer perspective by presenting the integration flow proposed to enhance the target application with limited effort. The proposed integration flow enforces a separation between the functional and extra-functional concerns using an Aspect-Oriented Programming Language to inject the code needed to introduce the adaptivity layer in the target source code.

On one hand we use the mARGOt [35] framework to dynamically tune the application, thus implementing the

```

1 // Load data
2 Routing::MCSimulation mc(edgesPath, profilePath);
3 auto run_result = mc.RunMonteCarloSimulation(samples, startTime);
4 ResultsStats stats(run_result);
5 Routing::Data::WriteResultSingle(run_result, outputFile);
6 return 0;

```

Listing 1: Original source code before integrating the adaptivity layer

adaptivity concepts presented in Section 4 and thus transforming the target application as highlighted in Figure 3. mARGOt is an open source dynamic autotuning library¹ that enhances an application with an adaptation layer. In particular, it provides to the application the most suitable configuration of software knobs, according to application requirements and to the execution environment evolution. The configuration selection is done by relying on application knowledge gathered at profile time. In this context, we use mARGOt to select the number of samples that minimizes the execution time, provided that they are enough to lead to an error below a certain threshold. In particular, the selection is done by considering the unpredictability value of the current path, and using as application knowledge the design-time model described in Section 6.1.

On the other hand, we hide all the complexity for code manipulation to the application developer by using LARA [36] as a language to describe user-defined strategies, and its Clava compiler² for source code analysis and transformation. LARA is a Domain Specific Language inspired by Aspect-Oriented Programming concepts. It allows a user to capture specific points in the code based on structural and semantic information, and then analyze and act on those points. This produces a new version of the application, leaving the original unchanged and separating the main functional concerns from those specified in LARA. Clava is a C/C++ source-to-source compiler based on the LARA framework. The compilation analyses and code transformations are described in scripts written in the LARA language.

In this work, we use Clava to perform two main tasks: first, to enrich the original source code with the required autotuner glue code, and second, to configure the autotuner library according to application requirements. Figure 4 depicts the transformation process, from the original source code, to the final application and highlights the two main LARA aspects used. To further clarify the evolution of the application code and related aspects, Listing 1–4, present respectively the original source code, the two LARA aspects used to enhance the target application, and the final enriched code.

In particular, the code in Listing 2 shows the aspect needed to configure mARGOt, producing an autotuning library tailored according to the application requirements. In lines 9–16, we define the *num_samples* tunable software knob, the *unpredictability* feature that we want to observe, the *error* metrics and the goal (i.e. the Service Level Agreement, *error* < 3%) that in mARGOt is a condition that can be used later to define the optimization problem. Once the knobs, metrics, and data features have been defined,

1. Project repository: https://gitlab.com/margot/_project/core
2. Project repository: <https://github.com/specs-feup/clava>

```

1 aspectdef McConfig
2   /* Generated Code Structure*/
3   output codegen end
4
5   /* mARGOt configuration */
6   var config = new MargotConfig();
7   var travel = config.newBlock('ptdrMonteCarlo');
8
9   /* knobs */
10  ptdrMonteCarlo.addKnob('num_samples', 'samples', 'int');
11  /* data features */
12  ptdrMonteCarlo.addDataFeature('unpredictability', 'float',
13    MargotValidity.GE);
14  /* metrics */
15  ptdrMonteCarlo.addMetric('error', 'float');
16  /* goals */
17  ptdrMonteCarlo.addMetricGoal('my_error_goal', MargotCFun.LE,
18    0.03, 'error');
19
20  /* optimization problem */
21  var problem = ptdrMonteCarlo.newState('problem');
22  problem.setStarting(true);
23  problem.setMinimizeCombination(MargotCombination.LINEAR);
24  problem.minimizeKnob('num_samples', 1.0);
25  problem.subjectTo('my_error_goal', 1);
26
27  /* creation of the mARGOT code generator for the following code
28    enhancement (McCodegen aspect) */
29  margoCodeGen_ptdrMonteCarlo = MargotCodeGen.fromConfig(config, '
30    ptdrMonteCarlo');
31 end

```

Listing 2: LARA aspect for configuring the mARGOT autotuner

we can proceed with the creation of the multi-objective constrained optimization problem that the autotuner has to manage (lines 18–23). In mARGOt optimization problems are called states (line 19). It is because mARGOt gives the possibility to define multiple optimization problems (only one can be the *default* one, line 22) and also to switch among them according to dynamic conditions. The constraints can be generated as in line 23, where the number represents the priority of the constraint. In case of more than one constraint, if the runtime is unable to satisfy both of them, it will relax the low priority one. Lines 21–22 define the objective function. Given that in this case the objective is the minimization of the number of samples, the aspect describes it as a linear combination (line 21) of the *num_samples* knob only by using a linear coefficient equal to 1 (line 22). mARGOt and LARA integration aspects permit to build different types of combined objective functions (e.g. linear or geometric combinations). Finally, line 26 builds the LARA internal structure *margoCodeGen_ptdrMonteCarlo* that is then used to create the mARGOt configuration file and code generator.

The second aspect (shown in Listing 3) aims at integrating the proposed methodology in the target application. It takes as input (line 3) the target function call that we want to tune, the mARGOt code generator produced by the previous aspect (Listing 2), and the number of samples needed to evaluate the unpredictability feature. In line 6, we query the code to identify the statement (*stmt*) including Monte Carlo function *call* as target join point to be manipulated. Lines 7–17 contain the actual manipulation actions done on the selected join point *stmt* of the target code. It is composed of mainly two different types of operations. First, to integrate the mARGOt calls for initializing the library and updating the software knob (Lines 10 and 14). Second,

```

1 aspectdef McCodegen
2   /* Target function, mARGOt code generator from McConfig aspect, #
3     samples for feature extraction */
4   input targetName, margoCodeGen_ptdrMonteCarlo,
5     unpredictabilitySamples end
6
7   /* Target function call identification */
8   select stmt.call{targetName} end
9   apply
10  /* Target Code Manipulation */
11  /* Add mARGOt Init*/
12  margoCodeGen_ptdrMonteCarlo.init($stmt);
13  /* add unpredictability code */
14  $stmt.insert before UnpredictabilityCode(unpredictabilitySamples);
15  /* Add mARGOt Update */
16  margoCodeGen_ptdrMonteCarlo.update($stmt);
17  /* Add Optimized Call Code */
18  $stmt.insert replace OptimizedCall(unpredictabilitySamples);
19  end
20 end
21
22 /* Unpredictability extraction code */
23 codedef UnpredictabilityCode(unpredictabilitySamples) %{
24   auto travel_times_feat_new = mc.RunMonteCarloSimulation([[
25     unpredictabilitySamples]], startTime);
26   ResultStats feat_stats(travel_times_feat_new, {});
27   float unpredictability = feat_stats.variationCoeff;
28 }% end
29
30 /* Optimized MonteCarlo call */
31 codedef OptimizedCall(unpredictabilitySamples) %{
32   auto run_result = mc.RunMonteCarloSimulation(samples - [[
33     unpredictabilitySamples]], startTime);
34   run_result.insert(run_result.end(), travel_times_feat_new.begin(),
35     travel_times_feat_new.end());
36 }% end

```

Listing 3: LARA aspect for inserting the application-specific glue code (unpredictability extraction) and the required mARGOt calls

```

1 // Load data
2 Routing::MCSimulation mc(edgesPath, profilePath);
3 auto travel_times_feat_new = mc.RunMonteCarloSimulation(100, startTime)
4 ;
5 ResultStats feat_stats(travel_times_feat_new, {});
6 float unpredictability = feat_stats.variationCoeff;
7 if(margot::travel::update(samples, unpredictability)) {
8   margot::travel::manager.configuration_applied();
9 }
10 auto run_result = mc.RunMonteCarloSimulation(samples - 100, startTime);
11 run_result.insert(run_result.end(), travel_times_feat_new.begin(),
12   travel_times_feat_new.end());
13 ResultStats stats(run_result);
14 Routing::Data::WriteResultSingle(travel_times_new, outputFile);
15 return 0;

```

Listing 4: Target source code after the integration of the adaptivity layer

to *insert* the glue code (LARA *codedef*) for calculating the unpredictability (line 12 and lines 21–25), and to *replace* the original Monte Carlo call with the optimized one that does not repeat the unpredictability samples (line 16 and lines 28–31).

Overall, in this specific instance of integration, we used 53 lines of LARA to generate 221 lines of C++ code. However, the advantage cannot be only considered from a numerical point of view (>4x in terms of line of codes). In fact, three are the main reasons to justify this approach. First, the user does not need not to worry about the details of the mARGOt configuration files and low-level C++ API, but can instead focus on the high-level interface available in LARA that results to be more declarative on the target problem (as shown in Listing 2 and Listing 3). Second, this approach

reuses information between the integration's several steps. There is mARGOt-specific information that should be provided by the user in several places like the configuration files and when using the autotuning API (e.g. the name of the autotuner block and the knobs and data features). By using the high-level LARA aspects, users only define this information once, saving time and possibly resulting in fewer production errors. Third, this approach leverages on a separation of concerns between the original code (functional description) and the autotuning code (extra-functional optimization). All the extra-functional optimizations, including problem definition (optimization targets and constraints), are kept separated and users do not have to modify the original source. In this way, the original developer does not need to be involved with all the optimization process and tools, thus permitting the functional development and extra-functional optimization to run in parallel.

6 EXPERIMENTAL RESULTS

In this section, we show the results of applying the proposed methodology to the PTDR algorithm. The platform used for the experiments is composed by several nodes based on the Intel Xeon E5-2630 V3 CPUs (@2.8 GHz) with 128 GB of DDR4 memory (@1866 MHz) on a dual channel memory configuration. First, we show the results of the model training for estimating the expected error (see Section 6.1). Then, in Section 6.2 we validate the approach by verifying the respect of the error constraint ϵ . We compare the proposed approach with respect to the original version that takes a static decision on the number of samples (see Section 6.3). Finally in Section 6.4 we discuss the overhead introduced, while in Section 6.5 we evaluate the optimization impact when considering the entire navigation service at system-level.

6.1 Training the Model

The first actual phase of the methodology is done off-line and it consists of training the error model ($error_{i,y}^x$) presented in Section 4.2 by using a different number of samples. For training the quantile regression, we used profiling data extracted by running the PTDR algorithm on a training set. This training dataset has been built using random requests done on 300 different paths across the Czech Republic in different time-slots, thus considering different speed-profiles for each segment of the paths. All these requests have been done for all the 4 levels of sampling used in this paper (i.e. 100, 300, 1000 and 3000, as described in Section 4.2). The output of the model training is represented in Figure 5. The points in the three plots represent the results obtained from the profiling runs. The lines represent the quantile regression lines, thus the model that will be used at runtime. The three sub-figures are different in terms of the quantile value used for the regression. Figure 5a, Figure 5b and Figure 5c represents the regression results using respectively the 50th, 75th and 95th quantile. We can see that the three regressions are slightly different since we pass from a more permissive one in Figure 5a, where almost half of the points are below the corresponding regression lines, to the most conservative one in Figure 5c where only a few points are

above. Analyzing in depth the data, we can see that the coefficients of the lines of the quantile regression are almost doubled passing from 75th to 95th percentile (e.g. for 100 samples, the coefficients pass from 0.27 to 0.38, while for 3000 samples they pass from 0.049 to 0.071).

The extracted models are now ready to be used at runtime by the dynamic autotuner to select the minimum number of samples that satisfy the error constraint for the given unpredictability value. The results shown in Section 6.2 will demonstrate the effectiveness of the proposed method at runtime.

6.2 Validation Results

The set of validation results presented in this section are reported to demonstrate how the dynamic tuning of the number of samples satisfies the error constraints. For doing this, we randomly generated 1500 requests to the enhanced PTDR module for routes on the Czech Republic at different starting times. These requests are different from the ones used in the training phase of the model. We validate the approach by using three different quantile regressions (on 50th, 75th and 95th quantile), two different target errors ϵ (3% and 6%) and a confidence interval (CI) for the error constraint equal to 99% (i.e. $n(99\%) = 3$). The error has been derived by considering a run of the Monte Carlo simulation on the same input set by using 1 million of samples, thus enough to be considered a good estimation of the actual travel time distribution. Then, we selected as error the maximum between different key percentiles: 5th, 10th, 25th, 50th, 75th, 90th and 95th percentile.

The results are reported in Figure 6 and Figure 7 respectively for an error constraint ϵ equal to 3% and 6%. The two figures show the error results for each run with respect to the unpredictability feature extracted on the path. Each dot in the plots represents a PTDR request, while its shape depends on the number of samples used for the Monte Carlo simulation. In most of the cases, the actual error is below the target error. As it was expected considering the same value of the error constraint ϵ , the more conservative is the quantile regression, the less are the points that violate the constraint error. For the data we processed, in all cases the number of times the error constraint is not respected is within the selected CI (99%). At the same time, moving from a less conservative quantile regression (e.g. 50th percentile) towards a more conservative one (e.g. 95th percentile), it is possible to note how the threshold values for selecting the same number of samples shifts to the left. As an example, by considering an error constraint $\epsilon = 3\%$ (see Figure 6), the maximum unpredictability value for having 300 samples moves from 0.075 to less than 0.06 respectively when considering the quantile regression from the 50th percentile, up to the 95th quantile. Similar is the case when we consider an error constraint $\epsilon = 6\%$ (see Figure 7), where the same threshold moves from an unpredictability of 0.15 to 0.14 and 0.11 when using the 50th, the 75th and the 95th as quantile value for the regression. Finally, it is clearly visible the difference in terms of number of samples between the two cases with different ϵ values. Indeed, while for ϵ equal to 3% (Figure 6) only a very small fraction of the cases use 100 samples and there is a not negligible fraction of cases where




Fig. 5: Training of the error model by using different number of samples and quantile regressions




Fig. 6: Validation of the proposed approach by using 3% as target error and different percentiles for the quantile regression




Fig. 7: Validation of the proposed approach by using 6% as target error and different percentiles for the quantile regression

3000 samples are employed. For ϵ equal to 6% (Figure 7) in some cases only 100 samples are required.

6.3 Comparative Results with Static Approach

In this subsection, we demonstrate the advantages obtained by using the proposed approach with respect to the baseline version [2] where the number of samples are defined *a priori*. To provide a fair comparison, we extracted the number of samples to be used for the baseline version by using the training dataset.

For the 4 levels of sampling used in this paper (i.e. 100, 300, 1000 and 3000, as described in Section 4.2), we analyzed the cumulative distributions of the expected error (see Figure 8). We selected the minimum sampling level that

passes a certain threshold of the cumulative value before reaching the error constraint value ϵ . This threshold value has almost the same *robustness* meaning of the quantile regression value used in our approach. In the following, we are going to compare the proposed approach where the quantile regression model has been built over a certain percentile, with a static tuned version where the same percentile has been used as the threshold for the cumulative. If we use for the proposed approach the quantile regression at 95% we compare with the statically tuned version where the number of samples has been defined looking at the cumulative curve that reaches at least 95% before to the target error constraint. In particular looking at Figure 8, we can notice that for an error constraint $\epsilon = 6\%$ the




Fig. 8: Cumulative distribution of the error by using different numbers of samples over the training set

TABLE 1: Average number of samples for the validation set using different quantile regression values (columns) and different error constraints. The results are reported for the *baseline* and proposed *adaptive* versions.

ϵ		Average Number of Samples		
		50 th perc.	75 th perc.	95 th perc.
3%	baseline	1000	3000	3000
	adaptive	632 (-36%)	754 (-74%)	1131 (-62%)
6%	baseline	300	1000	1000
	adaptive	153 (-49%)	186 (-81%)	283 (-71%)

static tuning is set to 1000 samples for the entire percentile interval between 72th and 98th, while for values larger than 98th and smaller than 72th percentile (down to 7th) we have to consider respectively the configuration using 3000 and 300 samples. On the other side for $\epsilon = 3\%$ we select 3000 samples within the percentile interval 72th-97th, 1000 samples for percentile values smaller than 72th (down to 5th), while we need more than 3000 samples if the request is very tight on a percentile larger than 97th.

Table 1 shows the comparative results obtained by using the proposed adaptive technique with respect to the original version (*baseline*) with the statically defined number of samples obtained with the previously described analysis. In particular, Table 1 presents the average number of samples and gain with respect to the baseline for different values of error constraint ϵ and different percentiles used to build the predictive model and for the static tuning of the baseline. The results are obtained by running a large experimental campaign over randomly selected pairs of Czech Republic routes and starting times, different from those used for the training. While the routes have been randomly selected, we used a more realistic distribution of the starting time [37] [38].

In all the considered cases, the proposed approach reduces the number of samples by at least 36%, and up to 81%. As expected, the average number of samples for the proposed approach is lower when we relaxed either the error constraint (i.e. 6%) or the percentile used for building the model (e.g. 50th percentile). The lower gain for the configurations using the 50th percentile with respect to the cases using the 75th-95th percentile is due to the fact that

in the former case the baseline requires a lower number of samples with respect to the latter cases (i.e. 1000 vs 3000 for $\epsilon = 3\%$ and 300 vs 1000 for $\epsilon = 6\%$). Looking at the absolute numbers, it is possible to detect that even if the percentage gain seems higher with more conservative regressions (75th-95th), the actual average number of samples used is smaller with the more permissive quantile (50th).

The reduction in terms of the number of samples is directly reflected on the execution time reduction since there is a linear dependency, except for the overhead introduced by the dynamic autotuner. In particular, we observed an execution time speed-up between 1.5x and 5.1x. A more detailed analysis on the overhead is presented in Section 6.4.

To further show the benefits of the proposed methodology, Figure 9 shows the number of samples selected by the adaptive Monte Carlo simulation when the same request in terms of target path is performed every 15 minutes during the entire week. The used temporal interval is derived by the smaller time granularity (ϕ) we had for the database containing the speed profiles. The two plots (a) and (b) have been generated using respectively 3% and 6% as maximum target errors and for both experiments a quantile regression on the 75th percentile.

By looking at the number of samples requested by the adaptive version of the Monte Carlo simulation, we can easily recognize well-known traffic behaviors in both plots. The daily distribution on the weekdays is characterized by 2 main peaks determined by less predictable situations. The first around 7–8 am and the second around 4–5 pm. During the weekend the morning peak seems to be a bit postponed, while the afternoon one almost disappears. On the opposite, it is also visible how the evening hours result to be the most predictable ones.

This dynamic behavior that is captured by the enhanced version of the algorithm cannot be exploited by using the original (baseline) version. Following the same philosophy adopted in Table 1, the original version must be tuned by considering 3000 samples for the experiment in Figure 9(a) ($\epsilon = 3\%$) and 1000 samples for the experiment in Figure 9(b). In both cases, the static tuning results to be the larger number of samples selected from the proposed techniques, that instead is able to use it only when it is strictly required (e.g. during the traffic peaks). Also considering the static tuning to the *average* case (i.e. 1000 samples for the experiment in Figure 9(a) ($\epsilon = 3\%$) and 300 samples for the experiment in Figure 9(b)) is not a viable solution. This is because there are still a lot of sampling reduction possibilities in predictable moments that will not be captured, and more important, the prediction will not be able to satisfy the algorithm output quality during the most unpredictable periods. Finally, a fixed time-slot policy is also sub-optimal since the unpredictability strongly depends not only on the time of the request but also on the path characteristics (e.g. urban or countryside path, close or far from congested areas) and length (e.g. when it is expected the arrival in a congested area).

6.4 Overhead Analysis

While we widely describe in Section 5 how we reduced the integration overhead from the application developer

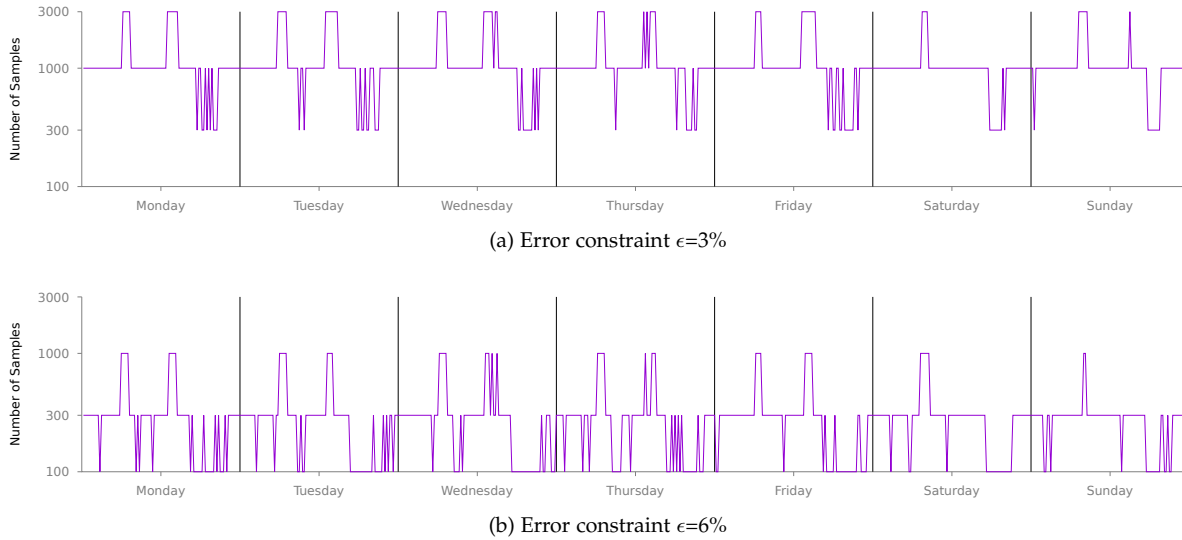


Fig. 9: Number of samples selected by the proposed adaptive method when the same request is performed every 15 minutes during the entire week.




Fig. 10: Evaluation of the execution time overhead due to the additional code for the proposed method with respect to the target Monte Carlo simulation by varying the number of samples.

point of view, this section clarifies the time-overhead introduced to obtain the proposed adaptivity. In particular, the additional computations that we add are related to the calculation of the $\nu_{\hat{x},y}$ and to the autotuner calls used to determine the right number of samples to be used. The initial 100 Monte Carlo samples, required to extract the data feature, are not part of the overhead given that they are reused (and thus discounted) to calculate the expected travel time (see Listing 4).

Figure 10 shows the overhead introduced by the proposed methodology compared to a set of Monte Carlo calculation by using a different number of samples (from 100 to 300, and 1M) over a set of paths among different locations in the three main cities in the Czech Republic. As expected, it is evident that the execution time is strictly correlated to the number of samples used for the travel time computation. The different paths we used are in a range between 300 and 800 segments long. When we fix the number of samples, the different number of segments is the main reason for the variability for the Monte Carlo simulation computing time.

Despite it is needed for every request, the overhead

introduced by our approach is almost negligible, i.e. more than 2 orders of magnitude less than the smaller Monte Carlo simulation with 100 samples. In particular, we found that the execution time of the data feature extraction and mARGO calls is comparable to the evaluation of a single sample of the Monte Carlo on a road composed of 200 segments.

6.5 System-Level Performance Evaluation

To quantify at system-level the effects of the proposed adaptive method, in this section we present an analysis done when considering that the efficient PTDR module is included in the full navigation pipeline shown in Figure 1. We built a performance model of the navigation pipeline by using the simulation environment Java Modeling Tools (JMT) [39]. JMT is an integrated environment for workload characterization and performance evaluation based on queuing models [40]. It can be used for capacity planning model simulation, workload characterization and automatic identification of bottlenecks. In particular, to build the simulation model of the queuing network, we considered one station for each of the modules composing the navigation pipeline and we added a fork-join unit to model the parallel PTDR evaluations of each alternative path found in the first stage.

The model, shown in Figure 11, has been annotated with values derived by the profiling of each module (K-Alternative path, PTDR, and reordering) and considering a value for K (the number of alternative paths to evaluate) equal to 10. Moreover, we made a resource allocation according to a load produced by up to 100K cars producing a request every 2 minutes. The latter number is in line with the consideration of having self-driving cars continuously connected with route planner, while the former has been derived by a simple estimation considering a Smart City such as the Milan urban area. Indeed, in this area the population is composed of around 4 Million people, every day it is estimated to have more than 5 Million trips, and




Fig. 11: Complete navigation pipeline modeled using JMT

only less than 50% are done by using public transportation [41], [42].

Under these conditions and considering the configuration with $\epsilon = 6\%$ and 95^{th} percentile, we found out that by adopting the proposed technique we have a 36% reduction in terms of number of resources needed to satisfy the target workload. In particular, we can differentiate 2 cases. The first one considering the number of resources needed to satisfy the steady-state conditions, and thus that the throughput in terms of input requests should be satisfied by all the stages. In this case, without the proposed optimization we would have needed at least 777 computing resources (cores). Among them, 400 cores (52% of the entire set) should be dedicated to PTDR. By applying the proposed technique only 497 cores are needed, reducing to 120 (24% of the entire set) those required for the PTDR stage. The second case considers a more dynamic environment where it is suggested to keep the average utilization rate of each station below 70%. While respecting this rule of thumb [43], the distribution of the system response time (the time passing from the navigation request to the response) results to be narrow, thus being able to better react to burst of requests. In this second case, without the proposed optimization we would have needed 1010 cores to allocate the entire pipeline. 572 of them (57% of the entire set) should be dedicated to the PTDR stage. By applying the proposed technique, 646 cores are enough to allocate the pipeline and out of them only 172 (26% of the entire set) are dedicated to the PTDR.

7 CONCLUSIONS

In this paper, we presented a novel approach for dynamically select the number of samples used in a Monte Carlo simulation to solve the Probabilistic Time-Dependent Routing problem. The proposed method quickly samples the input data to extract an unpredictability feature used to determine in a proactive manner the number of simulation to be executed while satisfying a certain error threshold. The runtime decision is based on a probabilistic error model – learned offline – correlating the unpredictability feature extracted from the data and the number of samples used by the Monte Carlo algorithm. Experimental results demonstrated that the proposed adaptive approach for the PTDR problem is able to save a large fraction of simulations (between 36%

and 81%) with respect to a static approach while considering different traffic situations, paths and error requirements. Considering the entire navigation pipeline, composed also of the k-alternative path and reordering stages, the adoption of the proposed technique guarantees a significant reduction in terms of computing resources. Finally, we adopted an aspect-oriented programming language (LARA) together with a flexible dynamic autotuning library (mARGOt) to reduce the effort necessary on the application developer for introducing the code needed to improve the execution efficiency.

ACKNOWLEDGMENTS

This work has been supported by European Commission under the grant 671623 FET-HPC-ANTAREX (AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems) and by The Czech Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project "IT4Innovations excellence in science – LQ1602" and by the IT4Innovations infrastructure which is supported from the Large Infrastructures for Research, Experimental Development and Innovations project "IT4Innovations National Supercomputing Center – LM2015070".

REFERENCES

- [1] P. Toth and D. Vigo, *Vehicle Routing: Problems, Methods, and Applications*. SIAM, 2014, vol. 18.
- [2] R. Tomis, L. Rapant, J. Martinovič, K. Slaninová, and I. Vondrák, "Probabilistic time-dependent travel time computation using monte carlo simulation," in *International Conference on High Performance Computing in Science and Engineering*. Springer, 2015, pp. 161–170.
- [3] M. Golasowski, R. Tomis, J. Martinovič, K. Slaninová, and L. Rapant, "Performance evaluation of probabilistic time-dependent travel time computation," in *IFIP International Conference on Computer Information Systems and Industrial Management*. Springer, 2016, pp. 377–388.
- [4] M. J. Gilman, "A brief survey of stopping rules in monte carlo simulations," in *Proceedings of the Second Conference on Applications of Simulations*. Winter Simulation Conference, 1968, pp. 16–20. [Online]. Available: <http://dl.acm.org/citation.cfm?id=800166.805256>
- [5] A. Agafonov and V. Myasnikov, "Reliable routing in stochastic time-dependent network with the use of actual and forecast information of the traffic flows," in *Intelligent Vehicles Symposium (IV), 2016 IEEE*. IEEE, 2016, pp. 1168–1172.
- [6] S. Samaranyake, S. Blandin, and A. Bayen, "A tractable class of algorithms for reliable routing in stochastic networks," *Procedia-Social and Behavioral Sciences*, vol. 17, pp. 341–363, 2011.
- [7] Y. M. Nie and X. Wu, "Shortest path problem considering on-time arrival probability," *Transportation Research Part B: Methodological*, vol. 43, no. 6, pp. 597–613, 2009.
- [8] M. Abeydeera and S. Samaranyake, "Gpu parallelization of the stochastic on-time arrival problem," in *High Performance Computing (HIPC), 2014 21st International Conference on*. IEEE, 2014, pp. 1–8.
- [9] M. Niknami and S. Samaranyake, "Tractable pathfinding for the stochastic on-time arrival problem," in *International Symposium on Experimental Algorithms*. Springer, 2016, pp. 231–245.
- [10] E. Nikolova, J. Kelner, M. Brand, and M. Mitzenmacher, "Stochastic shortest paths via quasi-convex maximization," *Algorithms-ESA 2006*, pp. 552–563, 2006.
- [11] A. Paraskevopoulos and C. Zaroliagis, "Improved Alternative Route Planning," in *ATMOS - 13th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems - 2013*, ser. OpenAccess Series in Informatics (OASISs), D. Frigioni and S. Stiller, Eds., vol. 33. Sophia Antipolis, France: Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Sep. 2013, pp. 108–122. [Online]. Available: <https://hal.inria.fr/hal-00871739>

- [12] T. Chondrogiannis, P. Bouros, J. Gamper, and U. Leser, "Alternative routing: k-shortest paths with limited overlap," in *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems, Bellevue, WA, USA, November 3-6, 2015*, 2015, pp. 68:1–68:4.
- [13] —, "Exact and approximate algorithms for finding k-shortest paths with limited overlap," pp. 414–425, 2017.
- [14] H. Janssen, "Monte-carlo based uncertainty analysis: Sampling efficiency and sampling convergence," *Reliability Engineering & System Safety*, vol. 109, pp. 123 – 132, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0951832012001536>
- [15] Q. Xu, M. Sbert, M. Feixas, and J. Sun, "A new adaptive sampling technique for monte carlo global illumination," in *2007 10th IEEE International Conference on Computer-Aided Design and Computer Graphics*, Oct 2007, pp. 191–196.
- [16] C. A. Schaefer, V. Pankratius, and W. F. Tichy, "Atune-il: An instrumentation language for auto-tuning parallel applications," in *Euro-Par 2009 Parallel Processing*, H. Sips, D. Epema, and H.-X. Lin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 9–20.
- [17] J. M. F. Moura, J. Johnson, R. W. Johnson, D. Padua, V. K. Prasanna, M. Püschel, and M. Veloso, "SPIRAL: Automatic implementation of signal processing algorithms," in *High Performance Extreme Computing (HPEC)*, 2000.
- [18] R. C. Whaley and J. J. Dongarra, "Automatically tuned linear algebra software," in *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, ser. SC '98. Washington, DC, USA: IEEE Computer Society, 1998, pp. 1–27. [Online]. Available: <http://dl.acm.org/citation.cfm?id=509058.509096>
- [19] H. Hoffmann, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard, "Dynamic knobs for responsive power-aware computing," in *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS XVI. New York, NY, USA: ACM, 2011, pp. 199–212. [Online]. Available: <http://doi.acm.org/10.1145/1950365.1950390>
- [20] A. Paukštė, "Monte carlo optimisation auto-tuning on a multi-gpu cluster," in *2012 2nd IEEE International Conference on Parallel, Distributed and Grid Computing*, Dec 2012, pp. 894–898.
- [21] V. Vassiliadis, C. Chaliou, K. Parasyris, C. D. Antonopoulos, S. Lalis, N. Bellas, H. Vandierendonck, and D. S. Nikolopoulos, "Exploiting significance of computations for energy-constrained approximate computing," *International Journal of Parallel Programming*, vol. 44, no. 5, pp. 1078–1098, Oct 2016. [Online]. Available: <https://doi.org/10.1007/s10766-016-0409-6>
- [22] J. Ansel, Y. L. Wong, C. Chan, M. Olszewski, A. Edelman, and S. Amarasinghe, "Language and compiler support for auto-tuning variable-accuracy algorithms," in *International Symposium on Code Generation and Optimization (CGO 2011)*, April 2011, pp. 85–96.
- [23] J. S. Miguel and N. E. Jerger, "The anytime automaton," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, June 2016, pp. 545–557.
- [24] M. A. Laurenzano, P. Hill, M. Samadi, S. Mahlke, J. Mars, and L. Tang, "Input responsiveness: Using canary inputs to dynamically steer approximation," in *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '16. New York, NY, USA: ACM, 2016, pp. 161–176. [Online]. Available: <http://doi.acm.org/10.1145/2908080.2908087>
- [25] X. Sui, A. Lenharth, D. S. Fussell, and K. Pingali, "Proactive control of approximate programs," in *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '16. New York, NY, USA: ACM, 2016, pp. 607–621. [Online]. Available: <http://doi.acm.org/10.1145/2872362.2872402>
- [26] E. Miller-Hooks and H. Mahmassani, "Path comparisons for a priori and time-adaptive decisions in stochastic, time-varying networks," *European Journal of Operational Research*, vol. 146, no. 1, pp. 67–82, 2003.
- [27] J. Martinovič, V. Snašel, J. Dvorský, and P. Dráždilová, "Search in documents based on topical development," in *Advances in Intelligent Web Mastering - 2*, V. Snašel, P. S. Szczepaniak, A. Abraham, and J. Kacprzyk, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 155–166.
- [28] M. Asghari, T. Emrich, U. Demiryurek, and C. Shahabi, "Probabilistic estimation of link travel times in dynamic road networks," in *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2015, p. 47.
- [29] J. M. Juritz, J. W. F. Juritz, and M. A. Stephens, ""on the accuracy of simulated percentage points."" 1983.
- [30] D. C. Montgomery and G. C. Runger, *Applied Statistics and Probability for Engineers*. John Wiley and Sons, 2003.
- [31] D. Zwillinger and S. Kokoska, *CRC Standard Probability and Statistics Tables and Formulae*. Chapman & Hall, 2000.
- [32] R. Tomis, J. Martinovič, K. Slaninová, L. Rapant, and I. Vondrák, "Time-dependent route planning for the highways in the czech republic," in *IFIP International Conference on Computer Information Systems and Industrial Management*. Springer, 2015, pp. 145–153.
- [33] P. P. Boyle, "Options: A monte carlo approach," *Journal of Financial Economics*, vol. 4, no. 3, pp. 323–338, 1977. [Online]. Available: <https://EconPapers.repec.org/RePEc:eee:jfinec:v4:y:1977:i:3:p:323-338>
- [34] R. Koenker, *Quantile Regression*, ser. Econometric Society Monographs. Cambridge University Press, 2005.
- [35] D. Gadioli, G. Palermo, and C. Silvano, "Application autotuning to support runtime adaptivity in multicore architectures," in *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*, 2015 International Conference on. IEEE, 2015, pp. 173–180.
- [36] J. M. Cardoso, J. G. Coutinho, T. Carvalho, P. C. Diniz, Z. Petrov, W. Luk, and F. Gonçalves, "Performance-driven instrumentation and mapping strategies using the lara aspect-oriented programming approach," *Software: Practice and Experience*, vol. 46, no. 2, pp. 251–287, 2016.
- [37] US Department of Transportation, Federal Highway Administration, *US Department of Transportation, Federal Highway Administration – Traffic Report*, 2014. [Online]. Available: <https://www.fhwa.dot.gov/policy/publications.cfm>
- [38] UK Department for Transport, gov.uk, *Average annual daily flow and temporal traffic distributions*, 2017. [Online]. Available: <https://www.gov.uk/government/statistical-data-sets/tra03-motor-vehicle-flow>
- [39] M. Bertoli, G. Casale, and G. Serazzi, "Jmt: Performance engineering tools for system modeling," *SIGMETRICS Perform. Eval. Rev.*, vol. 36, no. 4, pp. 10–15, Mar. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1530873.1530877>
- [40] E. D. Lazowska, J. Zahorjan, G. S. Graham, and K. C. Sevcik, *Quantitative System Performance: Computer System Analysis Using Queueing Network Models*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1984.
- [41] Milano Agenzia Mobilita' Ambiente e Territorio, *Annual Mobility Report*. [Online]. Available: <https://www.amat-mi.it/it/documenti/>
- [42] Marco Bedogni, Milano Agenzia Mobilita' Ambiente e Territorio, *Road Traffic Measures in The City of Milan*. [Online]. Available: <http://www3.gdos.gov.pl/Documents/Wizyty/W%C5%82ochy/Road%20Traffic%20Measures%20in%20the%20city%20of%20Milan.pdf>
- [43] M. Gribaudo, P. Piazzolla, and G. Serazzi, "Consolidation and replication of vms matching performance objectives," in *Analytical and Stochastic Modeling Techniques and Applications*, K. Al-Begain, D. Fiems, and J.-M. Vincent, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 106–120.