Efficient models configuration for an electric vehicle energy management software

Borjan Tchakaloff, Sébastien Saudrais, Jean-Philippe Babau

Efficient models configuration for an electric vehicle energy management software

Borjan Tchakaloff *†, Sébastien Saudrais * and Jean-Philippe Babau †
* CERIE, ESTACA, F-53000 Laval, France
firstname.lastname@estaca.fr
† Univ. Bretagne Occidentale, UMR 6285, Lab-STICC, F-29200 Brest, France
firstname.lastname@univ-brest.fr

***Abstract*—Energy management of electric vehicles has been the focus of recent research to allow optimal engine and battery usage. Many efforts have been realised to use the trip knowledge – or a prediction of it – to provide the best vehicle efficiency. Yet few works take into account the embedded devices and the vehicle global Quality of Service. The ORQA framework has a vehicle systemic approach, its purpose is to generate an architecture to counter the range anxiety and offer the best quality effort based on the driver preferences. The work described in this paper is about the off-line configuration of the ORQA framework to match a target vehicle characteristics and abilities. With a more precise configuration, the on-line execution of ORQA is optimised. Two leads are presented to reduce the computation time needed to explore the solution space on-line. The final result is an energy management software tuned for a specific targeted vehicle which offers a driving strategy and a control of the embedded devices matching the driver destination and preferences.**

I. Introduction

Though they still represent a small part on the global market, the electric vehicles have reached the worldwide market. The main issue raised against the technology is the short autonomy this kind of vehicle provide, hence a certain reluctance about their usage (commonly known as range anxiety). Current vehicles can handle a few hundred kilometres. One way to deal with the embedded energy issue of the electric vehicles is through software. An Energy Management System (EMS) is a high level software monitoring and managing an environment through specific-purposed components. EMS are commonly used in (full- and hybrid-) electrical vehicles, though they mainly manage only the engine and ignore the end-user provided Quality of Service (QoS). In order to offer an efficient energy management and to take into account the user-related QoS, an electric vehicle EMS has to consider every embedded devices and the user expectations.

The ORQA framework [1] offers to tackle the global energy management while providing a QoS as good as possible. In ORQA, each embedded device is characterised at design phase by its energy consumption(s) and its quality(ies), if applicable. The framework offers to realise a components architecture which will elaborate on-line a solution to achieve the driver request (to reach a destination) while providing the best possible vehicle QoS. The main idea is to limit by software the devices and engine usage. The framework architecture and process are presented in [1]. But the solution space of an optimal solution in which operates ORQA exponentially increases with the routes amount and the number of devices. To reduce the solution space, the challenge is to propose efficient models and an associate configuration for the framework.

This paper presents the decision models of ORQA, how they can reduce the solution space when correctly configured, and how it impacts the on-line computation. An extensive off-line search is realised on generated data to identify possible solution space reductions. Two approaches are presented, one which reduces the search input dimension, and the other which approximates the solution space. The designer is invited to browse through the results and to select the final configuration. In the end, the solution domain is tuned according to the target vehicle characteristics and abilities, while still providing various viable solutions.

The rest of the paper is structured as follows. Section II is the ORQA background, what the framework is and how it operates. Section III details the decision models of the framework. Section IV presents the two reduction approaches. Section V evaluates the approaches with a use-case and compares the obtained results. The related works are compared in Section VI, and we finally conclude the paper in Section VII.

II. ORQA structure and usage

ORQA is a framework to set up an Energy Management System (EMS) for electrical vehicles. ORQA proposes specific models and a methodology which leads to an energy manager connected to the vehicle embedded software at run-time. The idea behind ORQA is to offer a global service to the driver for a specific journey while maximising the device services. Also, the destination point should be reached quickly while saving the battery as much as possible.

First, ORQA provides a library of predefined models to characterise the energy consumption of different physical devices embedded in an electric vehicle. The library contains energy models for the electric engine, the climate control, the lighting system, and the entertainment system. For each journey, the possible routes are modelled: depending on the route characteristics (like speed constraints and duration), one can limit or not engine usage to reach the destination point. The driver preferences are introduced to integrate user priorities between devices usage. They belong in the system memory and can be updated at need. The on-line usage of

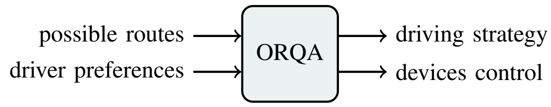

Fig. 1. Overview of the ORQA Energy Management System.

ORQA results in a driving strategy to reach the destination under constraints over the vehicle velocity and the devices usage. Another result is the corresponding device controller configuration (see Fig. 1 which illustrates the global overview of ORQA EMS), according to the driving strategy.

The vehicle devices model characterises the devices embedded in the target vehicle. ORQA considers the target vehicle as a *system* composed of a set of known devices. A *device* is a software function that controls a physical device. A device may function at different operating states to realise its service. An *operating state* has a certain power requirement to operate, that is a *power function* that can be constant or parameterised by environmental data. The energy consumption of an operating state basically depends on the duration the device operates in it. The whole energy consumption of a device is directly given by the sum of its operating states energy consumptions. ORQA differentiates devices that are mandatory to operate the vehicle (they are handled automatically by the vehicle) and those that are not. One can control the latter devices by selecting their operating states, thus changing the way they realise their service and hence the Quality of Service they offer. Each operating state of these controlled devices (the *controlled devices*) has a *quality* attribute denoting its QoS. A *devices combination* is a map such as each controlled device has an operating state. ORQA defines two extreme combinations: 1) the least consuming combination and 2) the nominal combination. The *least consuming combination* corresponds to every controlled devices operating at their least consuming operating state (in most cases, switched off). Whereas the *nominal combination* is the combination of the controlled devices operating at their best quality operating state. ORQA does not interfere with the mandatory devices except for the engine. It manages the qualified devices at run-time by constraining their operating states.

The routes model represents the different paths between the departure point and the arrival point, divided into steps. A *route* is a collection of steps and belongs to an environment. A *step* is defined by its initial and final velocities, its distance, and its slope. The *initial velocity* of a step is the final velocity of the preceding step, or 0 if there is none (that is, the step is the first of the route). An *environment* defines usual initial and final velocity ranges, and common step distances. Default environments are *urban*, *rural*, and *motorway* as defined by the Common Artemis Driving Cycles (CADC) [2] European project.

At run-time, the ORQA Energy Management System handles the driver request. The EMS operates as follow:

1) The driver defines a destination point;
2) The possible routes are retrieved or generated from the system (e.g. from the GPS unit);
3) Each possible driving strategy is evaluated (consumption, duration, quality): a driving strategy is composed of a route and a devices combination;
4) A driving strategy is chosen and the driver is informed;
5) The embedded energy manager controls the devices on-line following the proposed strategy. Dedicated broking components control the devices behaviour.

For the fourth step (the driving strategy selection), the ORQA framework has to define an efficient ranking of different driving strategies. We propose in this paper a simple ranking based on scoring functions. And because ORQA is executed on-line, the ranking computation has to be efficient. To limit the search space, we propose to reduce off-line the range of the different parameters to evaluate in the third step. We now present the different necessary models to evaluate the driving strategies.

III. ORQA MODELS FOR DECISION

The logic to select the driving strategy is embedded in the EMS. It is used on-line to determine which driving strategy to choose. The parameters of the choice are the velocity coefficients, the driver preferences and the scoring functions.

A. The velocity coefficients

The possible routes of a journey are retrieved from the system. A route models the nominal driving conditions of a journey: nominal velocities and no traffic congestion. A route is refined by applying a velocity coefficient to its nominal velocities. A *velocity coefficient* is a percentage and represents how much a vehicle is slowed down compared to a nominal driving. The new route is a *variation* of the nominal route and represents a different driving condition. This refinement allows the EMS to explore different driving conditions of the retrieved routes.

B. The driver preferences

The triggering event of the ORQA process is the driver request to reach a destination. Aside from the destination point itself, two sets of parameters can be set: the controlled devices ranking and the consumption policy. The *devices ranking* is a distribution of a finite amount over each controlled device reflecting the driver preferences. The default devices ranking is set at design phase but the driver must be able to update it for his own preferences. For instance, if there is not enough energy to operate the controlled devices in the nominal combination, then another less consuming combination is selected according to the importance of the devices for the driver. The *consumption policy* is the driver policy to match for the proposed solution. It is a named policy with optional constraints over the duration and the energy consumption. A consumption policy defines weights for the main results of a solution: duration, consumption, and quality. The *consumption policy weights* are used to rank the computed solutions as described later in the scoring functions. The duration constraint is expressed as a maximal *delay* to the nominal solution (e.g. $2\times$). The consumption constraint is expressed as a minimum *level of energy* left in the battery (e.g. 20%).

C. The scoring functions

The driving strategies are created by mapping each route to each devices combination. We introduce a ranking for the driving strategies to select the best one according to the consumption policy defined by the driver. The *score* of a driving strategy is defined by its main results: duration, consumption and quality. These results come from the route evaluation against devices combinations. Each one of the main results is passed to a corresponding scoring function. A *scoring function* normalises a main result over every driving strategies and vehicle capacities (e.g. the consumption of a route against the maximum level of energy stored in the battery). Depending on which main result is involved, the scoring function returns the best score with the minimum (duration and consumption) or the maximum (quality) value. The consumption policy constraints are taken into account by the scoring functions, an out-of-range result discards its driving strategy. For example, if the driver has set a duration constraint of $1.5\times$, any route that has a $\frac{route\ duration}{nominal\ route\ duration}$ ratio over 1.5 is discarded. The scoring functions results are combined by a weighted arithmetic mean. The meaning weights are defined by the consumption policy from the driver request. So the score S of a route is:

$$\begin{cases} S = S_T \cdot w_T + S_E \cdot w_E + S_Q \cdot w_Q & \text{feasible strategy} \\ S = 0 & \text{discarded strategy} \end{cases} \tag{1}$$

where S_i is a scoring function, w_i is a weight defined by the current consumption policy, and the indices $_{T}$,$_{E}$,$_{Q}$ respectively represent the duration, the consumption, and the quality results. A score has a null value if at least one of its results is discarded (i.e. it does not satisfy the constraints). The best scoring strategy is given back to the driver as the proposed driving strategy.

The ORQA framework has to take into account different parameters to choose the best strategy on-line: the velocity coefficients, the driver preferences and the scoring functions. In order to optimise the search, the framework has to be configured according to the vehicle capacities. We now propose two approaches to reduce the solution space during design time which lead to an on-line reduced complexity.

IV. Off-line configuration

The configuration of ORQA is based on the exploration of the velocity coefficients domain applied to each route. As the solution space is composed of the routes and their variations, the amount of velocity coefficients directly impacts the on-line search process. We define two approaches to accelerate the search process: a reduction of the search input dimension and an approximation of the route variations. We present both approaches in the remainder of this section. The former approach lies on the reduction of the velocity coefficients amount. It helps the designers determine an adequate group of velocity coefficients to represent the whole coefficients domain. The latter approach explores approximations of the route variations so as to avoid the whole variations computation.

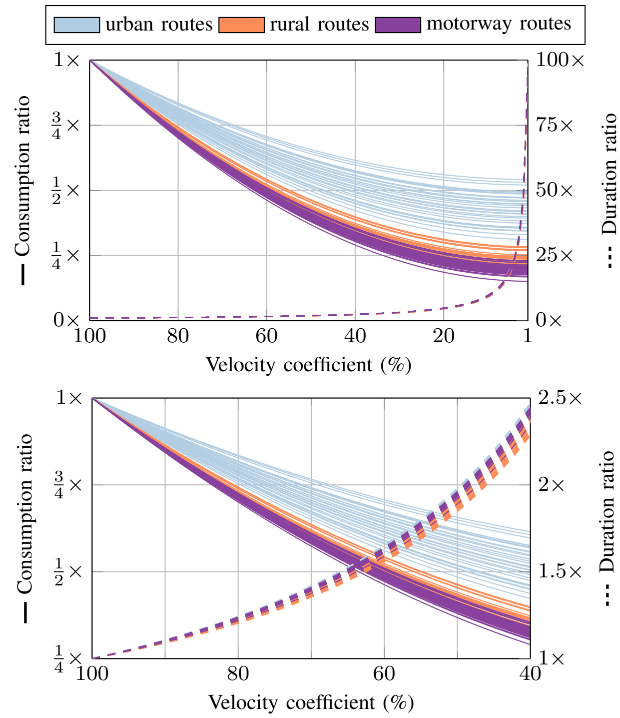

Fig. 2. Evolution of the routes ratios on (a) the complete domain of velocity coefficients and on (b) a domain restricted to $[40\%; 100\%]$.

These approaches can be used separately or combined together. It is possible to apply the approaches on each environment independently and embed on-line different reductions.

For each route, one can compute the duration and consumption which characterise the route with the ORQA devices models. We use the vehicle models defined in [1] (presented in the Evaluation section) which represent an urban electric car. To understand the approaches and their interests, we first present how duration and consumption evolve along the velocity coefficients. One hundred routes are generated per environment. Each route is refined a hundred times by velocity coefficients from 1% to 100%. This assures a good coverage of the velocity coefficients domain which is the complete input domain of the solution space. The duration and consumption are represented by relative values, called ratios, compared to the nominal routes (velocity coefficient of 100%). The duration ratio is in fact the delay compared to the nominal duration. Fig. 2.a illustrates the ratios variation for the generated routes with only one device in the vehicle, the engine. In the context of transportation, a delay of twice the nominal duration seems to be a reasonable limit for a journey. So, in the remainder of the paper, we limit the velocity coefficients domain under study to $[40\%; 100\%]$ as the example variations under 40% exceed a duration ratio of $2\times$. Fig. 2.b is a focus of Fig. 2.a on this restricted velocity coefficients domain.

A. Velocity coefficients of grouped results

The first approach is based on data clustering. It groups the variations to limit the on-line route exploration. The idea is to define a limited number of representative groups of velocity coefficients instead of using the whole domain. The partitioning is based on the *k-medoids* partitioning algorithm [3], [4] initialised with the *k-means++* algorithm [1] [5]. The input of the algorithm is a set of vectors. Each vector represents a velocity coefficient, it is characterised by the ratios of every routes for that coefficient. It is important to notice that duration and consumption ratios evolve independently (see Fig. 2), so they both have to be considered independently for partitioning. The partitioning algorithm yields a set of k (from 1 to the number of coefficients) groups of vectors. A set of routes are in the same group if they minimise the distortion metric. The *error rate* measures the relative error (the standard deviation) between a vector and one particular vector of the group, called the representative vector. As a vector represents a velocity coefficient, the k groups of vectors returned by the partitioning algorithm lead to k groups of velocity coefficients. Also, the representative vector of a group leads to the representative velocity coefficient of the group. If k is not fixed, a full range of clustering has to be realised to find the "best" groups possible.

B. Ratios approximation

In the second approach, we consider the fact that ratios, for different routes, evolve in a same way (see Fig. 2). We propose to approximate a set of ratios evolutions with a representative evolution, called the *approximation function*. So at each velocity coefficient (the evolution step), the ratios are represented by one ratio called the *approximated ratio*. The approximation functions resulting of this approach are given to the designers as hints to optimise the variation computation. Indeed, the complete evaluation of the variations is replaced by approximating the nominal route results, so the on-line complexity is greatly reduced. The duration and the consumption results evolve differently, their ratios are approximated independently. So it is possible that there is one approximation for the duration ratios and three for the consumption ratios (one per environment).

Both of the two approaches rely on approximation, so we introduce a metric (the error rate) to assess an approximation reliability. The error rate relies on the standard deviation of the results and is defined as follows:

$$error = \sqrt{\sum_{i}^{n} \frac{((x_i - x'_i)/x_i)^2}{n}} \tag{2}$$

where x_i is a computed result and x'_i is its corresponding estimation.

We present two different approaches to reduce the on-line search process complexity based on an adequate off-line configuration. Both approaches rely on a set of data gathered from generated routes and both consider the similarities between the routes variations. On the one hand, a specific amount of velocity coefficients is computed off-line to be later embedded in the EMS. On the other hand, ratios are approximated off-line to provide an on-line look-up table of ratios.

[1]The *k-medoids* algorithm output depends on the initial input. The *k-means++* is a classical initialisation method that maximise the initial input coverage.

V. Evaluation

We apply the two approaches on a use-case to evaluate their benefits over the complete solution space. The use-case contains three different routes, one per environment: urban (3.4km), rural (28.7km), and motorway (47.8km). The use-case vehicle is basically a city car with two optional devices: an air-conditioning unit and an entertainment system. The vehicle models are thoroughly described in [1]. The vehicle main characteristics are the following: a mass of 1200kg, a frontal area of 2.75m^2, an engine nominal power of 35kW, and a battery rated energy of 20kWh. We assume an outer temperature of 30°C and an inner objective temperature of 20°C. The driver preferences are equally distributed among the two optional devices. The consumption policy chosen by the driver for this use-case is conservative, it is a trade-off between high energy savings and good comfort. The weights of the duration, the consumption and the quality used to compute the route scores are respectively 0.2, 0.5, and 0.3. As the same vehicle is used to illustrate the previous section, we use the same generated routes to serve as the base data to the two reduction approaches. The scoring functions are defined as follows:

$$S_T = 100 + \frac{99 \cdot (T^{min} - T)}{T^{max} - T^{min}} \tag{3}$$

$$S_E = 100 + \frac{99 \cdot (E^{min} - E)}{E^{max} - E^{min}} \tag{4}$$

$$S_Q = 101 - \left(100 + \frac{99 \cdot (Q^{min} - Q)}{Q^{max} - Q^{min}}\right) \tag{5}$$

with the ^{min} and ^{max} exponents representing the minimal and maximal valid values (i.e. not discarded) found for the results. T, E, Q respectively represent the route duration, consumption, and quality. The scoring functions allow each result score to be in $[1; 100]$, or < 1 if discarded. The driving strategy scores follow (1).

A. Initial configuration

The initial configuration explores the complete solution space, one case per velocity coefficient and per devices combination. The two optional devices have respectively three and two operating states, so there are six possible devices combinations. And as there are sixty-one velocity coefficients in the studied domain ($[40\%; 100\%]$), the complete solution space is composed of 366 cases. For each case, the results are evaluated and the corresponding scores are computed. Fig. 5 shows the scores evolution for the three different route environments. For the urban route, the fifteen velocity coefficients in $[79\%; 93\%]$ yield the best scores. For the rural route, it is the thirteen ones in $[63\%; 75\%]$, while for the motorway route it is the fifteen ones in $[59\%; 74\%]$. The optimal driving strategies results are listed in Table I along their relative deviations.

TABLE I

OPTIMAL DRIVING STRATEGIES FOUND FOR EACH ROUTE IN THE COMPLETE SOLUTION SPACE. MEANS OF THE OPTIMAL RESULTS ARE DISPLAYED ALONG THEIR RELATIVE DEVIATIONS.

Route	Optimal velocity coefficient	Route results duration	consumption	quality
urban	$\in [79\%; 93\%]$	5"35' ±16'	243Wh ±14Wh	68% ±2%
rural	$\in [63\%; 75\%]$	40"28' ±2"09'	2113Wh ±132Wh	59.5% ±2%
motorway	$\in [59\%; 74\%]$	52"24' ±3"40'	3720Wh ±312Wh	58% ±2%

B. Application of the reduction approaches

We now apply the two approaches to the data coming from the generated routes. The explored approaches are applied to five configurations: 1) the urban routes, 2) the rural routes, 3) the motorway routes, 4) both the rural and the motorway routes, and 5) every routes (all of the environments). The other environments combinations (i.e. urban–rural and urban–motorway) are not shown here because of their low value for the study.

a) Velocity coefficients of grouped results: The first approach is now applied to the data from the generated routes. Fig. 3 displays the evolution of the global error rate for the five configurations. The global error rate is the combination of the duration error rate and of the consumption error rate. The *elbow criterion* is a visual method to determine an adequate configuration based on the evolution of the error metric. In this method, the searched solution is visible when the results plot forms an important angle (the elbow). The external edge of the elbow points out the limit configuration from which the results look alike. We identify in each configuration the adequate number k of clusters with the elbow method, which is four for every configuration. The representative values of the configurations are a) $\{46\%, 60\%, 75\%, 92\%\}$, b) $\{46\%, 60\%, 76\%, 93\%\}$ and $\{46\%, 60\%, 77\%, 93\%\}$ for the c), d), and e) configuration. We notice that, in this example, the representative values of velocity coefficients are very similar for the five configurations. Also, though the ratios evolution looks linear in Fig. 2.b, both duration and consumption evolutions are taken into account to cluster the velocity coefficients. Moreover, compared to a regular discretisation of the velocity coefficients domain, the proposed approach provides the sufficient number of groups and more representative values. A naive by-4 discretisation would be $\{40\%, 60\%, 80\%, 100\%\}$ instead.

b) Ratios approximation: The second approach is now applied to the data from the generated routes. Fig. 4 displays the error rates of the approximation ratios along the velocity coefficients for each configuration. On the one hand, the duration approximations have a low error rate of at most 2.15% for the all-routes configuration. It means that the duration of a route variation can be approximated with less than 2.15% with a single approximation function on the duration. The "all routes" configuration is thus chosen to estimate on-line every duration ratios. On the other hand, the consumption approximations have more disparate error rates as can be seen on Fig. 4. The configurations respectively reach an error rate of 11.3%, 8.6%, 4.3%, 7.4%, and 20.4% at the last velocity coefficient (40%). It is up to the designers to act on these results as they are highly subjective. For instance, if we were to consider that a configuration should have a maximal consumption error rate of 10% at the last velocity coefficient, then the "rural–motorway routes" configuration would be acceptable for rural and motorway routes. But there would be no configuration for the urban routes. Hence no approximation function would be available on-line to approximate urban routes and their variations should be computed instead.

TABLE II

RESULTS OF THE BEST DRIVING STRATEGIES FOUND FOR EACH ROUTE BASED ON THE REDUCED INPUT DOMAIN (FIRST APPROACH).

Route	Optimal velocity coefficient	Route results duration	consumption	quality
urban	77%	6"11'	218Wh	63.5%
rural	60%	46"04'	1826Wh	55%
motorway	60%	57"40'	3305Wh	55%

C. Configuration I: velocity coefficients of grouped results

The different groups found by the first approach are close enough to allow one representative groups for every environment: $\{46\%, 60\%, 77\%, 93\%\}$. The six devices combinations are still explored, so there are 24 cases to evaluate per route. Fig. 6 shows the three different environments for the twenty-four cases. For the urban route, the 77% coefficient corresponds to the best score. For both of the rural and motorway routes, the 60% coefficient yields the best scores. We can see that for both the urban and the rural routes, sub-optimal driving strategies have been selected. The motorway route has an optimal score that is also one of those found in the complete evaluation. Their results are summed up in Table II. In summary, the urban route driving strategy differs from the mean optimal of 36 seconds (10%), 25Wh (11%), and 4.5 percent of quality. The chosen rural route differs from the mean optimal of 6 minutes and 24 seconds (14%), 287Wh (16%), and 4.5 percent of quality. Finally, the motorway route differs from the mean optimal of 6 minutes and 28 seconds (11%), 511Wh (15%), and 4 percent of quality.

D. Configuration II: ratios approximation

We choose to use the ratios approximations of each environment. Each nominal route is evaluated by ORQA and its variations results are approximated. So there are only 6 evaluations per route (one per device combination), and 360 approximations. Fig. 7 shows the scores resulting from these

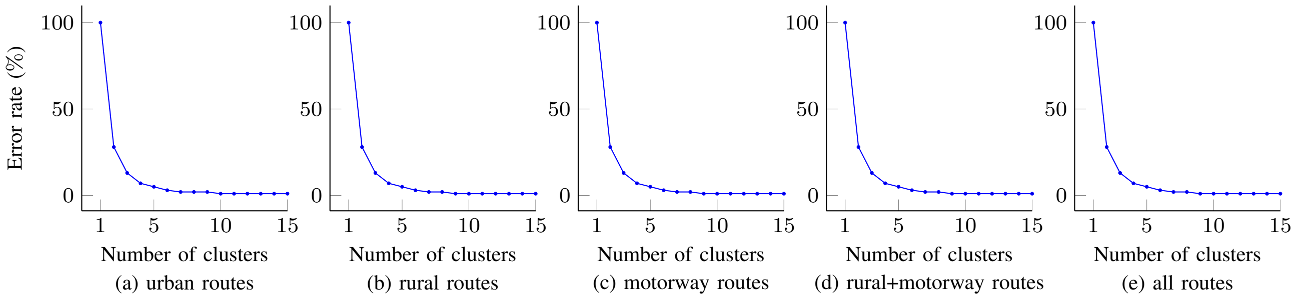

Fig. 3. Evolution of the global error rate per number of clusters (amount of velocity coefficients) for each configuration in the first approach.

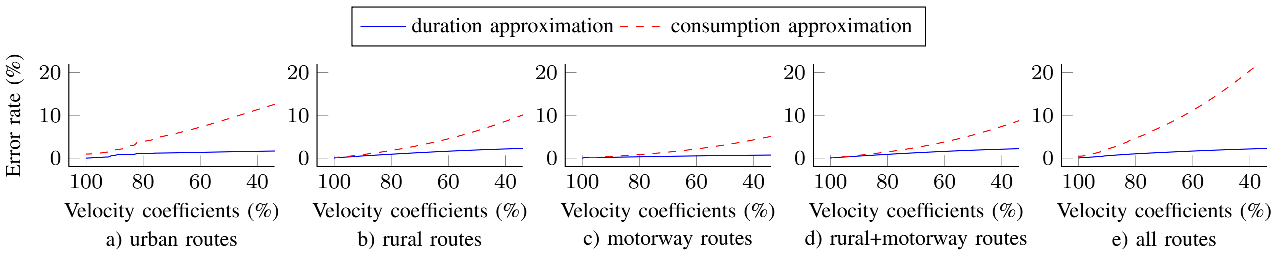

Fig. 4. Evolution of the duration and consumption error rates for each configuration in the second approach.

TABLE III
MEAN RESULTS OF THE BEST DRIVING STRATEGIES FOUND FOR EACH ROUTE BASED ON THE APPROXIMATED SOLUTION SPACE.

Route	Mean optimal velocity coefficient	Mean route results duration	consumption	quality
urban	89%	5"27'	254Wh	66.75%
rural	72%	38"41'	2142Wh	54%
motorway	69%	49"55'	4384Wh	69.5%

TABLE IV
RESULTS OF THE BEST DRIVING STRATEGIES FOUND FOR EACH ROUTE BASED ON THE COMPOUND OF REDUCED INPUT DOMAIN AND APPROXIMATED SOLUTION SPACE.

Route	Optimal velocity coefficient	Route results duration	consumption	quality
urban	77%	6"14'	221Wh	57.75%
rural	60%	45"49'	1719Wh	45%
motorway	60%	57"24'	3643Wh	60%

approximations. For the urban route, the thirteen velocity coefficients in $[79\%; 99\%]$ yield the best scores. For the rural route, it is the nine coefficients in $[68\%; 76\%]$, while for the motorway route it is the twelve ones in $[64\%; 75\%]$. The optimal scores of the approximated results share common coefficients with the complete evaluation scores. The mean results of the driving strategies are summed up in Table III. In summary, the urban route mean driving strategy differs from the mean optimal of 8 seconds (2%), 11Wh (4%), and 1.25 percent of quality. The chosen mean rural route differs from the mean optimal of 1 minute and 47 seconds (14%), 287Wh (5%), and 5.5 percent of quality. Finally, the mean motorway route differs from the mean optimal of 1 minute and 17 seconds (3%), 568Wh (13%), and 10.5 percent of quality. Both the urban and rural routes are within the relative deviations of the optimal results, aside from the quality level of the rural route. The motorway mean driving strategy has a duration in the optimal range but its high quality devices combination leads to a higher consumption (18% more).

E. Configuration III: compound approach

We now combine the approaches by approximating the ratios of the chosen coefficients. Each route has 6 evaluations and 24 approximations. Fig. 8 shows the scores resulting from these approximations. For the urban route, the 77% coefficient corresponds to the best score. For both of the rural and motorway routes, the 60% coefficient yields the best scores. We can see that for both the urban and the rural routes, sub-optimal driving strategies have been selected. The motorway route has an optimal score that is also one of those found in the complete evaluation. This is coherent with the first approach results, the four chosen coefficients give close results but not in the range of the optimal results. The results of the three routes are summed up in Table IV.

We evaluate the two reduction approaches against three different routes. We see that the new solution spaces are effectively reduced (24 evaluations against 366 in the first approach) or less complex to compute (6 evaluations and 360 approximations in the second approach). On the three examples, the obtained driving strategies are quite close to the optimal ones for the two proposed approaches. For the urban and the rural routes, we see that the error is within the relative deviation range when relying on the ratios approximation approach but not for the motorway route. The first approach, grouping the velocity coefficients, produces less accurate solutions which results variate more than 10% from the optimal ones. On the other hand, the compound approach (6 evaluations and 24 approximations) gives mixed results. They are coherent with the two approaches but too much loss is noted. This composition is not adequate in this use-case. The next step of this evaluation is a run-time monitoring of a real executing platform to tune the two approaches.

VI. Related Works

Energy consumption optimisation for and by software is a main-stream research domain, from data-centres to embedded systems or vehicles. In the domain of computing platforms, Snowdon et al. [6] offers a platform to trade performance and energy consumption of applications at an operating system (OS) level. They rely on an off-line characterised model to predict software performance and energy consumption. The consumption modelling is specific to the hardware-platform (CPU/bus/memory clock frequencies, CPU voltage, temperature sensors). They collect each process behaviour at run-time based on statistics provided by the OS kernel. A subjective policy is then dynamically applied to perform the trade-off for each process. The authors approach is similar to a systemic dynamic power management (DPM). In [7], several system-level DPM techniques are detailed. They advocate it is best to implement DPM at the higher software system-level (the OS). Such DPM benefits from the global view of the OS over on-going operations, and software-based management is more flexible than hardware-based. Druilhe et al. [8] presents an energy optimisation in digital homes by using consolidation over IT devices. They take into account the devices appearance, heterogeneity and the services QoS. Devices and services are characterised off-line. The system optimisation is triggered by the execution of new services and the devices appearance updates. ORQA shares the same philosophy of these approaches. As them, ORQA takes into account the vehicle controllers consumption and also considers every other devices consumption, including the engine which is the main consumer of vehicles. Thus, it enhances the vehicle manager to perform an overall optimisation at run-time.

Work has been done to optimise the energy consumption of vehicles, whether focused on the engine, the embedded devices, and even their controllers. Optimising the driving process to minimise the trip energy consumption can be related to optimal trajectory with constraints on the engine power [9], [10] and trip duration [11]. These approaches focus on the engine and its optimal operation. Several techniques are used such as backward search (Mensing et al.), optimal control problem (Petit et al.), and inversion-based approach (Dib et al.). In [12], Katoen et al. presents a model-based approach to determine an energy-optimal software/hardware deployment found off-line. The authors approach is to systematically explore the design space of possible mappings and rank them along their energy consumption. The result is a hardware deployment with a minimal energy consumption which is the embedded devices controllers have an energy-optimal consumption. ORQA offers to take into account both the engine operation and the embedded devices to perform a better trip. As the control is realised on-line, the search must be executed fast. So ORQA relies on a complete off-line overview of the solution space to fasten the results computation.

VII. Conclusion

The paper presents techniques to model and configure an electric vehicle Energy Management System. We outline the ORQA framework based on a systemic view of the vehicle. It generates a dedicated EMS featuring the global energy management while providing a QoS as good as possible. We introduce the decision models defining the evaluation part of ORQA. We present two different approaches to effectively configure the decision models off-line. They are based on extensive results matching the targeted vehicle capacities. One offers to reduce the search input dimension (6% of the complete solution space evaluated) while the other proposes to approximate the search space (2% evaluated with a whole coverage, but with less precision). The approaches can be used exclusively or in combination. The compound approach is even less complex than the second one but also cumulates the loss of both approaches. We see on the three route examples that the compound approach is not interesting for this use-case vehicle. The approaches application leads to a suitable configuration optimised for the vehicle EMS.

The current framework approach to reduce the vehicle velocity is to apply one coefficient throughout the trip. A variable coefficient approach offers a more realistic driving but introduces a new dimension to the search space. This enhancement therefore requires a new configuration setting to be validated. Also, we plan to realise a complete phase of measures on a real vehicle controller for the different approaches. Then the two proposed approaches can be evaluated against the initial solution space on the execution time.

References

[1] B. Tchakaloff, S. Saudrais, and J.-P. Babau, "ORQA: Modeling Energy and Quality of Service within AUTOSAR Models," in *Proceedings of the 9th international ACM Sigsoft conference on Quality of software architectures - QoSA'13*. Vancouver, British Columbia, Canada: ACM Press, 2013, p. 3.

[2] M. André, "The ARTEMIS European driving cycles for measuring car pollutant emissions." *The Science of the total environment*, vol. 334-335, pp. 73–84, Dec. 2004.

[3] L. Kaufman and P. J. Rousseeuw, "Clustering by means of medoids," *Statistical Data Analysis Based on the L 1-Norm and Related Methods. First International Conference*, pp. 405–416 416, 1987.

[4] S. Theodoridis and K. Koutroumbas, *Pattern Recognition, Fourth Edition*, 2008, vol. 11.

[5] D. Arthur and S. Vassilvitskii, "k-means++: The advantages of careful seeding," *SODA '07 Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035, 2007.

[6] D. C. Snowdon, E. Le Sueur, S. M. Petters, and G. Heiser, "Koala - A Platform for OS-Level Power Management," in *Proceedings of the fourth ACM european conference on Computer systems - EuroSys '09*. New York, New York, USA: ACM Press, 2009, p. 289.

[7] L. Benini, A. Bogliolo, and G. De Micheli, "A survey of design techniques for system-level dynamic power management," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 8, no. 3, pp. 299–316, Jun. 2000.

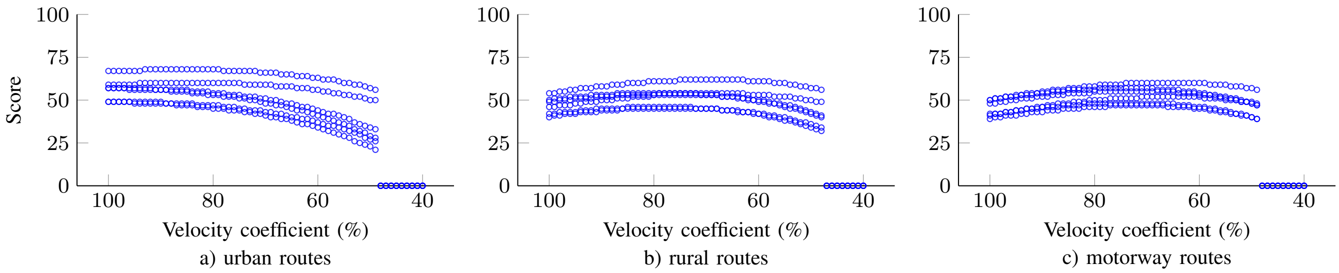

Fig. 5. Scores of the three routes based on the complete solution space.

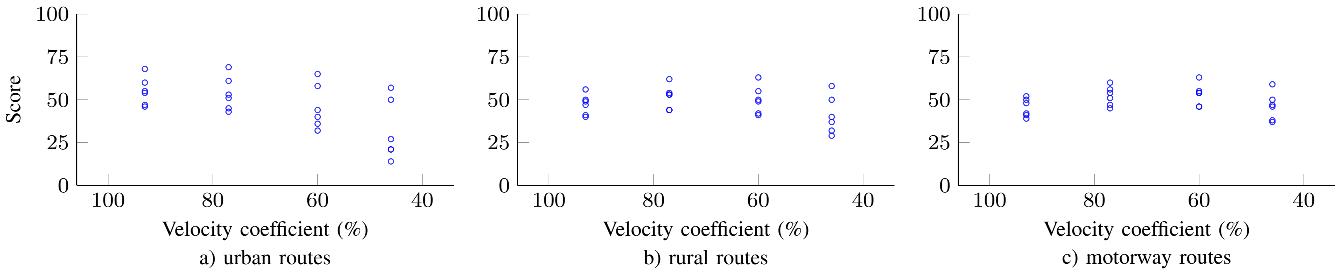

Fig. 6. Scores of the three routes based on the reduced input domain (first approach).

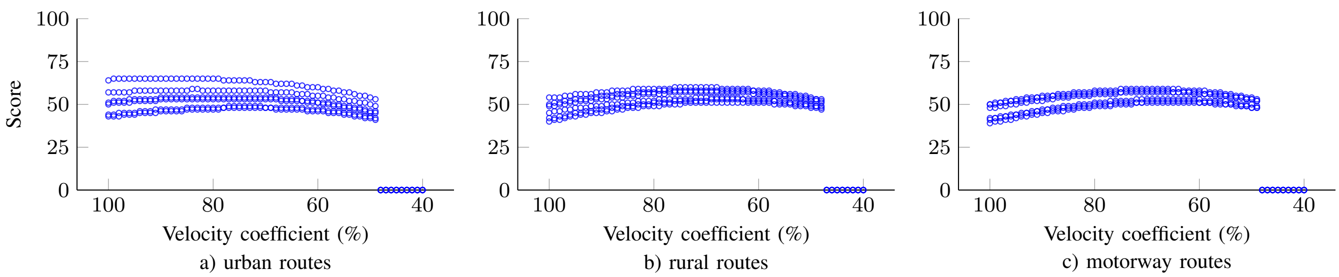

Fig. 7. Scores of the three routes based on the approximated solution space (second approach).

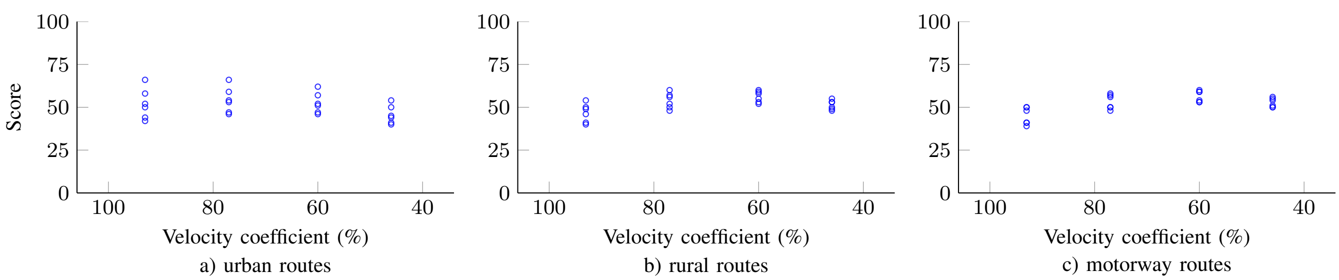

Fig. 8. Scores of the three routes based on both the reduced input domain and the approximated solution space (compound approach).

[8] R. Druilhe, M. Anne, J. Pulou, L. Duchien, and L. Seinturier, "Energy-driven consolidation in digital home," in *Proceedings of the 28th Annual ACM Symposium on Applied Computing - SAC '13*. New York, New York, USA: ACM Press, 2013, p. 1157.

[9] F. Mensing, R. Trigui, and E. Bideaux, "Vehicle trajectory optimization for application in ECO-driving," in *2011 IEEE Vehicle Power and Propulsion Conference*. IEEE, Sep. 2011, pp. 1–6.

[10] N. Petit and A. Sciarretta, "Optimal Drive of Electric Vehicles Using an Inversion-Based Trajectory Generation Approach," in *Proceedings of the 18th IFAC World Congress, 2011*, B. Sergio, Ed., Aug. 2011, pp. 14 519–14 526.

[11] W. Dib, A. Chasse, D. Di Domenico, P. Moulin, and A. Sciarretta, "Evaluation of the Energy Efficiency of a Fleet of Electric Vehicle for Eco-Driving Application," *Oil & Gas Science and Technology Rev. IFP Energies nouvelles*, vol. 67, no. 4, 2012.

[12] J.-P. Katoen, T. Noll, H. Wu, T. Santen, and D. Seifert, "Model-Based Energy Optimization of Automotive Control Systems," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013*. New Jersey: IEEE Conference Publications, 2013, pp. 761–766.