



HAL
open science

Enhancing Automotive User Experience with Dynamic Service Orchestration for Software Defined Vehicles

Pierre Laclau, Stéphane Bonnet, Bertrand Ducourthial, Xiaoting Li, Trista Lin

► **To cite this version:**

Pierre Laclau, Stéphane Bonnet, Bertrand Ducourthial, Xiaoting Li, Trista Lin. Enhancing Automotive User Experience with Dynamic Service Orchestration for Software Defined Vehicles. 2024. hal-04505345v1

HAL Id: hal-04505345

<https://hal.science/hal-04505345v1>

Preprint submitted on 14 Mar 2024 (v1), last revised 26 Sep 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Enhancing Automotive User Experience with Dynamic Service Orchestration for Software Defined Vehicles

Pierre Laclau^{1,2}, Stéphane Bonnet¹, Bertrand Ducourthial¹, Xiaoting Li² and Trista Lin²

¹Heudiasyc, CNRS, Université de Technologie de Compiègne, France {firstname.lastname@utc.fr}

²Stellantis, Vélizy-Villacoublay, France {firstname.lastname@stellantis.com}

Abstract—With the increasing demand for dynamic behaviors in automotive use cases, Software Defined Vehicles (SDVs) have emerged as a promising solution by bringing dynamic onboard service management capabilities. While users may request a wide range of services during vehicle operation, background tasks such as cooperative Vehicle-to-Everything (V2X) services can activate on-the-fly in response to real-time road conditions. In this dynamic environment, the efficient allocation of onboard resources becomes a complex challenge, in order to meet mixed-criticality onboard Quality-of-Service (QoS) network requirements while ensuring an optimal user experience. Additionally, the ever-evolving real-time network connectivity and computational availability conditions further complicate the process. In this context, we present a dynamic resource-based onboard service orchestration algorithm that considers real-time in-vehicle and V2X network health, along with onboard resource constraints, to select degraded modes for onboard applications and maximize user experience. To enable dynamic orchestration, we introduce the concept of Automotive eXperience Integrity Level (AXIL) which expresses a runtime priority for non-safety-critical applications. This algorithm produces near-optimal solutions while significantly reducing execution time compared to straightforward methods as demonstrated by simulation results. With this approach, we aim to enable efficient onboard execution for a user experience-focused service orchestration.

Index Terms—Software Defined Vehicle (SDV), User Experience (UX), Service Oriented Architecture (SOA), Optimization, Network Aware Orchestration, In Vehicle Network (IVN).

I. INTRODUCTION

In the last few years, automotive manufacturers have started to develop a range of complex features such as dedicated infotainment app stores, automated driving, and Vehicle-to-Everything (V2X) services. This trend invites automotive manufacturers to continuously deliver these features through seamless Over-the-Air (OTA) updates, which contributes to longer lasting vehicles in the hope to bring the industry closer to sustainable transportation systems. As a result, future vehicles may resemble ‘*Smartphones on Wheels*’ where users can dynamically request services throughout the vehicle lifetime while automakers continuously integrate background services such as Advanced Driver Assistance Systems (ADAS), cooperative multi-vehicle services, smart grid, and more [1, 2].

To support this paradigm shift, the automotive industry is undergoing a rapid transformation toward Software Defined Vehicles (SDV) [3]. Previously, Electric and Electronic (E/E) architectures were hardware-defined by integrating many

single-function Electronic Control Units (ECU) into domain-centric networks. However, the increasing number of ECUs as well as more demanding network and OTA requirements are reaching the limits of current architectures [4]. These changes are motivating a global shift towards Zonal Oriented Architectures (ZOA), where fewer High Performance Computing (HPC) ECUs are expected to group and manage multiple heterogeneous functions. ZOA allows for a centralized and therefore cost-effective reservation of computational resources for future updates and features. In addition, Ethernet facilitates service-oriented communications and increases bandwidth [5]. Together, these efforts toward centralization using fewer and virtualized ECUs can help reduce resource usage such as embedded hardware needs and runtime energy consumption.

This new hardware approach must be controlled by a software stack capable of applying OTA updates and dynamically switching the onboard software context to provide relevant applications. A common solution is to deploy a Service Oriented Architecture (SOA) which handles high-level communications through publish-subscribe protocols and the orchestration of services based on user requests, vehicle context, and available updates, allowing for dynamic reconfigurations of software and network resources [6]. However, SOA does not address the guarantee of Quality of Service (QoS) constraints such as hard real-time latency and jitter for safety-critical flows. Hence, the onboard architecture also relies on an embedded infrastructure based on technologies such as Software Defined Networking (SDN) and Time Sensitive Networking (TSN) that can be used to handle dynamic routing and hard real-time scheduling of network resources on a per-flow basis [7].

While reconfigurations are feasible, the simultaneous activation of the rising number of applications from app stores and updates could face limitations due to onboard constraints, potentially hindering on the user experience (UX). The onboard constrained resources, coupled with strict automotive requirements [8], may impose limitations on the sets of possible combinations of concurrently active applications. Hence, the vehicle may need to prioritize and schedule specific app combinations within its physical capacities. In certain scenarios, it could become necessary to temporarily degrade the quality of service allocated to certain applications or even shutdown some of them, ideally depending on UX impacts [9].

However, there is currently no method for selecting which applications to degrade or shutdown when the physical capabilities are not sufficient to allocate all desired services. Besides, the automotive industry must ensure that a user will have an enjoyable experience as defined by contract-based

This work was supported by Stellantis under the collaborative CIFRE framework UTC/CNRS/PCA (ANRT contract n°2021/0865) with Heudiasyc.




Fig. 1: Block diagram of the onboard orchestration platform. Our goal is to orchestrate the selection of runtime modes (including shutdown) for each application to remain within the available onboard, V2X, and edge offloading resources while maximizing the user experience (defined by our AXIL rating).

Service Level Agreements (SLAs) established at the time of vehicle purchase [10]. Depending on the customer’s vehicle model, profile, and subscriptions, automakers may define different runtime priorities to adapt to the user expectations (i.e. higher-level SLAs should guarantee higher QoS). Currently, the automotive industry does not have a methodology to dynamically adapt the vehicle’s QoS depending on the context.

Without any orchestration methodology, one solution would be to statically allocate the maximum network requirements for each application under the worst network conditions. However, this introduces unnecessary overhead and constrains automakers to include fewer services running in the vehicle.

We propose a different approach by defining a concept called “Automotive eXperience Integrity Level” (AXIL) that takes inspiration from the well-known Automotive Safety Integrity Level (ASIL) standard [11]. While ASIL helps engineers optimize the “safety to cost ratio” of onboard features, AXIL aims to help shape the automotive digital offering by introducing a measurable “user experience to onboard resources ratio” for each application. We propose to let each application developer define several runtime modes that require different hardware and network resources with varying levels of functionality. Intermediate modes can be seen as *degraded* modes with reduced functionality. By associating an AXIL rating to each runtime mode, the onboard software infrastructure will then be capable of activating the relevant runtime modes to meet all onboard constraints while maximizing the overall onboard experience, thereby optimizing resource usage. With this approach, the onboard orchestrator may find that slightly degrading some applications to make room for more applications improves the overall perceived experience for passengers.

In this paper, we present our contributions in several steps. Section II introduces our AXIL metric to assess the perceived user experience of each available feature, facilitating effective orchestration. Then, Section III introduces an easy-to-represent model of the onboard resource requirements. Section IV presents an efficient algorithm to select modes that maximize the overall user experience while staying within resource limits. We describe a specialized heuristic that strikes a balance between solution quality and execution speed, enabling real-time onboard calculations. Finally, Sections V and VI present our implementation in a simulated environment to compare it with other approaches. Figure 1 summarizes our approach by identifying the components of our orchestration platform.

II. BACKGROUND

Traditionally, vehicles are manufactured with a pre-defined static allocation of resources for onboard applications. To ensure a safe and consistent state across the entire vehicle software stack, engineers also define rule-based local state management strategies for each ECU and software layer to account for unexpected events such as link failures, invalid software states, or unexpected packet receptions. Based on the surrounding local state and events, these strategies apply degraded modes on each application with reduced functionality [12]. However, with the arrival of dynamic use cases and updates with thousands of applications, it may become too difficult to control the exact combinations of active applications requested by the user (e.g. infotainment and ADAS features), vehicle (e.g. health monitoring and data collection), or road context (e.g. cooperative V2X services and edge offloading) [13]. In this context, the traditional rule-based approach lacks the adaptability required to handle unexpected or novel situations, making it inadequate for addressing the intricate and rapidly changing network traffic requirements.

The industry is currently shifting toward dynamic onboard service orchestration using the new capabilities of SOA to enable more efficient resource usage, reduced architecture complexity, and less engineering efforts for state management [14]. This new paradigm enables vehicle-wide monitoring [15] for centralized and context-aware global orchestration. With the help of reconfigurable technologies such as TSN and SDN for the in-vehicle network, it is now possible to dynamically change the behavior and resource allocations of onboard software components through reconfigurations [16, 17]. Furthermore, several research initiatives propose to allocate network and computing resources using atomic vehicle-wide reconfiguration strategies to guarantee safe transitions even while driving [1, 18, 19]. Hence, while it may be feasible to automate the generation of local state management rules, these methods rely on complex and intensive engineering processes. A new challenge then arises: how can we build a centralized and autonomous state management stack to enable these future dynamic orchestration mechanisms?

Vehicles can execute applications of two categories. First, *safety-critical* (SC) services require Quality-of-Service (QoS) guarantees such as real-time computing and time-sensitive network flows with pre-defined latency, jitter, and isolation requirements. They also require deterministic behavior across the entire vehicle to ensure a safe and reliable execution for any external context. Second, *best-effort* (BE) applications are usually user-focused and offer onboard features to enhance user experience. While some of these applications can require resources with QoS allocations (e.g. video network flow for a video streaming app), others can seamlessly run with a minimum bandwidth requirement without being affected by higher latencies and jitter caused by other apps [20, 21].

Hence, the vehicle could dynamically allocate resources for best-effort applications depending on the available resources.

Our previous work [22] presented a methodology to pre-compute configurations for safety-critical applications. It consists in generating resource allocations to guarantee all QoS

network requirements based on a fixed onboard E/E topology description. As configuration scheduling takes time and safety certifications require deterministic and predictable vehicle states, our work pre-computes these configurations offboard by predicting the combinations of active applications onboard, then pre-downloads them to the vehicle. With this method, the vehicle can safely switch between configurations for different execution contexts while always satisfying QoS requirements.

However, with the rising number of non-safety-critical and increasingly dynamic applications, future use cases may include thousands of apps which could reduce the effectiveness of vehicle state prediction. The method of [22] cannot guarantee that a valid configuration has been pre-calculated for every combination of active applications. Additionally, resource availability can change at runtime due to (1) the dynamic activation of safety-critical apps with reserved network and computing resources thereby restraining best-effort applications, and (2) outside conditions such as V2X network congestion or 4G/5G connectivity losses which can severely impact the functionality of applications [21]. Hence, these best-effort applications must adapt to the ever-changing remaining unused resources (or pre-allocated resources with QoS, e.g. video network flows). This work addresses these limitations by extending [22] with dynamic suboptimal onboard decisions.

Research has already widely explored dynamic resource-based service orchestration strategies, notably in the cloud [23-25] and more recently automotive [26, 27] industries with bridges between the two for edge and fog computing [28, 29]. These initiatives can allocate services to computing nodes (i.e. server or ECU) depending on node-specific peripherals or resource requirements defined by each application developer such as CPU, RAM, GPUs, actuators, and network usage.

However, cloud-native solutions are loosely constrained and therefore only consider simple metrics such as CPU and memory, as opposed to the highly embedded automotive environment with real-time traffic shaping, V2X network constraints, and more. To the best of our knowledge, while some automotive contributions have enhanced the resource models [27], no research considers the selection of degraded or runtime modes as an additional degree of freedom for more orchestration granularity. For instance, the vehicle could temporarily lower video streaming quality on network congestion events or even activate more applications with newly freed resources.

To the best of our knowledge, there is no existing approach that connects user experience to dynamic vehicle orchestration. Therefore, we propose to take inspiration from the well-adopted ASIL standard and extend it for non-safety-critical applications. While ASIL guides engineering design choices to dimension the hardware based on each feature's safety requirements, our proposed AXIL metric can facilitate automated orchestration algorithms for optional user-oriented features.

When allocating applications to specific software contexts or ECUs, automakers can either pre-define them or implement one of many existing scheduling strategies [27]. Here, we propose to add an onboard runtime selector alongside this orchestration step to dynamically apply degraded modes onboard thereby maximizing UX. In this work, we aim to define a light-weight algorithm to select runtime modes, which can then be

triggered onboard on context or resource availability changes.

We believe that the unique challenges posed by an environment with thousands of applications are not addressed by existing research. By defining an automated centralized orchestrator with these objectives, we add a global optimization layer of runtime mode selection instead of local rule-based state management, thereby optimizing the onboard user experience and resource usage. Our work aims to bridge this research gap by optimizing the runtime conditions of services in real time to support (1) a large number of applications, (2) maximized user experience, and (3) minimal onboard overhead.

III. PROBLEM FORMULATION

We propose a new resource-based orchestration methodology that allows for granular and real-time runtime mode selection for each application requested by the user or vehicle. In this section, we describe the problem by (A) setting our initial assumptions, (B) defining an app store with runtime modes and dependencies, (C) proposing a data structure to represent the vehicle state, (D) defining the optimization objective, and (E) identifying the NP-Hard nature of the presented problem.

A. Assumptions and goals

Our approach is based on the assumption that each application has already been assigned to a host software context or ECU, either manually or with the help of an automated onboard orchestrator. We believe this to be a reasonable assumption for the coming architectures and use cases, as services are usually highly dependent on hardware capabilities such as infotainment apps in the High Performance Computer (HPC) or external networking services in the Connectivity Control Unit (CCU). Our goal is to define an efficient algorithm capable of real-time onboard decisions to automatically adjust the onboard functionalities when the onboard or external network resources cannot allocate the currently requested resources. The algorithm must maximize the user experience at all times.

B. AXIL: User Experience Metric

As one cannot improve what cannot be measured, we propose to define a new metric to assess the user experience of each application. This metric can be defined by each application developer and extended to each degraded runtime mode to tweak the onboard functionality with enhanced granularity. We call this metric AXIL for **A**utomotive **eX**perience **I**ntegrity **L**evel. AXIL enables the vehicle to select which functionalities to activate for each application. We identify three parameters to assess the priority of a particular feature offered by an app:

- **Ease of substitution:** How easily an average user can obtain the same service provided by the onboard feature if the corresponding feature is not provided by the vehicle. For instance, optional navigation features could be easily offered by the user's smartphone instead, while cooperative services may require to be executed onboard.
- **Exposition:** Relative assessment of how often an average user may use the feature compared to the feature set. These values compare the frequency of use of each feature to the total number of features in the vehicle.

TABLE I: Definition of AXIL. Just like ASIL, AXIL combines three parameters to assess the runtime priority of a service and defines a priority order (AXIL D being the highest priority).

E_1	E_2	E_3			
		Minimal	Low	Medium	High
Easy	Rare	-	-	-	-
	Low	-	-	-	-
	Medium	-	-	-	A
	High	-	-	A	B
Medium	Rare	-	-	-	-
	Low	-	-	-	A
	Medium	-	-	A	B
	High	-	A	B	C
Difficult	Rare	-	-	-	A
	Low	-	-	A	B
	Medium	-	A	B	C
	High	A	B	C	D

Priority order: - < A < B < C < D

Legend →	E_1	E_2	E_3
ASIL	Controllability	Exposition	Severity
AXIL	Ease of Substitution	Exposition	Quality of Exp.

- **Quality of Experience (QoE):** Relative measurement of user convenience or enjoyment when using the feature. Automakers may conduct UX studies and define guidelines to calculate the corresponding QoE values, although the methodology is out of the scope of this paper.

Then, the original structure defined by ASIL in Table I is an excellent candidate to assign a priority rating based on these three parameters. In the table, AXIL uses the same structure as ASIL to combine the three previous parameters into one global runtime priority metric. Hence, an AXIL D feature has the highest priority in this definition. Naturally, this concept can be extended with more granular levels and even be dependent on the user profile and preferences. In the rest of this paper, we assume that the AXIL rating can be generalized to any value in \mathbb{R}^+ with higher values corresponding to higher priorities. See Figure 2 for an illustration of a particular instance.

More broadly, the AXIL rating can be defined as a function of the vehicle state and user profile. For instance, a navigation app may have a higher AXIL rating when the vehicle is in a city compared to a highway. This outcome could be achieved with existing methods used by web applications, which can ask for regular user ratings of their experience (e.g. call quality for online meeting platforms). Vehicles could integrate the same feedback system, where A/B testing methods would study the effectiveness of ratings to gradually converge toward maximized user satisfaction, although not in the scope of this work. Hence, we assume in this work that ratings can evolve dynamically throughout the vehicle lifetime.

C. Problem description

1) *Applications:* Let $\mathcal{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ be the set of applications available in the app store, including background services, each with their own independently defined QoS network and other runtime requirements. We partition \mathcal{A} in two disjoint subsets $\mathcal{A}_{SC} \subseteq \mathcal{A}$ for safety-critical (SC) apps and $\mathcal{A}_{BE} \subseteq \mathcal{A}$ for best-effort (BE) services such that $n_{SC} + n_{BE} = n$ with $n_{SC} = |\mathcal{A}_{SC}|$ and $n_{BE} = |\mathcal{A}_{BE}|$. We assume




Fig. 2: Random problem instance generated with 20 applications and at most 5 modes. Each cell (i, j) represents mode j of \mathcal{A}_i and is assigned a runtime priority X_i^j (along with a vector of resource requirements M_i^j defined in Section III-C3).

that requirements are correctly provisioned to orchestrate SC apps while also attempting to allocate QoS resources for BE apps without guarantees, with methods such as our previous work [22]. This work focuses on allocating BE apps while (1) the vehicle dynamically reserves resources for SC apps and (2) external conditions such as connectivity continuously evolve.

2) *Resources:* The vehicle can be modelled as a list of r resources $\mathcal{R}_1, \dots, \mathcal{R}_r$. The maximal hardware capacity of the vehicle is denoted as a vector of r positive reals R^{\max} such that $R^{\max}[i] \in \mathbb{R}^+$ denotes the maximal capacity of resource \mathcal{R}_i . Resources can describe the CPU and memory usages for each ECU, available best-effort bandwidth for each physical network link in each link direction, external system capacity such as V2X and edge computing availability, and more.

However, resource availability can dynamically change at runtime, such as when external conditions (e.g. V2X network saturation) or internal events (e.g. resource allocation change due to a change of SC apps thereby impacting BE bandwidth availability) occur. Hence, we define the current BE capacity of a vehicle as a vector R of r elements in \mathbb{R}^+ . We have:

$$0 \leq R[i] \leq R^{\max}[i] \quad i \in 1, \dots, r$$

3) *Runtime modes:* In this work, we consider that each application $\mathcal{A}_i \in \mathcal{A}_{BE}$ can independently define m_i runtime (or degraded) modes. Each mode offers a different level of functionalities and resource requirements. For instance, a video streaming application may define multiple video quality options with corresponding network bandwidth usages. For each application \mathcal{A}_i , we denote each mode j as \mathcal{A}_i^j and define M_i^j as its corresponding vector of resource requirements. This is a vector of r positive real elements satisfying:

$$0 \leq M_i^j[k] \leq R^{\max}[k] \quad k \in 1, \dots, r$$

We assume that \mathcal{A}_i^1 is the nominal mode with the maximum features, resource requirements, and AXIL rating of \mathcal{A}_i . Each following mode provides fewer features using less resources:

$$\begin{aligned} \forall i = 1, \dots, n_{BE} \quad \forall j, j' = 1, \dots, m_i \\ \forall k = 1, \dots, r \quad j < j' \implies M_i^j[k] \geq M_i^{j'}[k] \quad (1) \\ \text{and } \exists k \in \{1, \dots, r\}, \quad M_i^j[k] > M_i^{j'}[k] \end{aligned}$$

In Figure 2, AXIL ratings also follow the previous constraint as higher modes must have lower priorities. Each cell (i, j) is also associated with its vector of resource requirements M_i^j .

4) *Mode dependencies*: In a realistic automotive environment, applications will also have dependencies to other applications. In our context, each runtime mode may define a minimum mode for another application to be active. For instance, the video streaming app may require a minimum display and network activity level managed by their respective background services. Hence, we define a mode dependency graph $D = \{V, E\}$ with nodes $V = \{\mathcal{A}_i^j \mid \forall 1 \leq i \leq n_{\text{BE}}, 1 \leq j \leq m_i\}$ as a Directed Acyclic Graph (DAG) to avoid cyclic dependencies. There is an oriented edge $(\mathcal{A}_i^j, \mathcal{A}_{i'}^{j'})$ in E if and only if \mathcal{A}_i^j requires $\mathcal{A}_{i'}^{j'}$.

D. Optimization problem

The previous definitions describe the problem instance. However, at any time, the vehicle may request to launch any combination of applications $\mathcal{A}_{\text{req}} \subseteq \mathcal{A}_{\text{BE}}$ with $n_{\text{req}} = |\mathcal{A}_{\text{req}}|$.

1) *Resource model*: The previously defined problem can be arranged in a 3D matrix called P to store the resource requirements (up to r) for each runtime mode (up to $m_{\text{max}} = \max_{i \in \{1, \dots, n_{\text{req}}\}} m_i$ defined as the maximum number of modes among the apps \mathcal{A}_{req}) of each application (up to n_{req}):

$$P_{i,j,k} = \begin{cases} M_i^j[k] & \text{if } j \in \{1, \dots, m_i\} \\ 0 & \text{otherwise} \end{cases}$$

2) *User experience*: When the current available resources R cannot allocate all requested applications \mathcal{A}_{req} , the vehicle must select a subset of applications to run. Our approach is to select a mode for each application to maximize the user experience. For each application, each runtime mode can be assigned an AXIL rating to represent a launch priority focused on UX. This information can be stored for all applications in a single 2D matrix X of size $n_{\text{req}}, m_{\text{max}}$ with values in \mathbb{R}^+ .

3) *Objective function*: Our goal is to select one mode for each application to maximize the user experience while respecting the resource and dependency constraints. The solution can be represented as a vector C of selected runtime modes:

$$C[i] = \begin{cases} 0 & \text{if } \mathcal{A}_i \in \mathcal{A}_{\text{req}} \text{ is off} \\ j & \text{if mode } \mathcal{A}_i^j \text{ is active} \end{cases} \quad (2)$$

With this efficient encoding, the vehicle will then activate the set of final modes $\{M_i^{C[i]} \mid \forall C[i] \neq 0\}$. Finally, the objective is to maximize the sum of AXIL ratings of each activated mode, while remaining under the current resource budgets R and guaranteeing all dependency constraints in D :

$$\operatorname{argmax}_C \sum_{i=1}^{n_{\text{req}}} X[i, C[i]] \quad (\text{note } \forall i, X[i, 0] = 0) \quad (3)$$

$$\text{with } \forall k \in \{1, \dots, r\}, \quad \sum_{i=1}^{n_{\text{req}}} P_{i, C[i], k} < R[k]$$

$$\text{and } \forall (M_i^j, M_k^l) \in E \quad C[i] \leq j \implies C[k] \leq l$$

Once we have found the optimal runtime modes, the process is repeated whenever (1) the requested vehicle state \mathcal{A}_{req} changes, (2) at least one budget increases to allow for higher-end runtime modes, or (3) at least one budget updates to a value below its current resource usage.

E. Problem Complexity

This problem can be reduced to a variant of the *knapsack* problem [30] to demonstrate its NP-Hard nature. In the case where applications only define one runtime mode, the problem is equivalent to defining 1 bin of capacity R with r dimensions. Then, n_{req} items must be placed in the bin (i.e. runtime modes). Each item is a vector of r resource requirements and has a value defined by its AXIL rating. With the equivalent goal of placing the items with maximized sum of item values in the bin while remaining within the bin capacity on all dimensions, our problem is therefore NP-Hard.

In the case of multiple runtime modes, the complexity increases. Each app can be seen as a set of m_i items. The problem is then to select one item per app. While close to the *multiple-choice knapsack* problem [31], the problem is different as the items must also follow dependency constraints.

To the best of our knowledge, there is no existing work that solves this problem in an embedded context. In this paper, we propose to use a heuristic to find a solution in polynomial time.

IV. METHODOLOGY

Our approach enables a dynamic self-organizing orchestration methodology to automatically balance the onboard "experience-to-resource" ratio for best-effort applications.

The purpose of the *Runtime Orchestrator*, illustrated in Figure 1, is to manage the selection of runtime modes to achieve the previous optimization goal. It is called when the *State Orchestrator* (out of the scope of this work) changes the combination of requested applications. Once a runtime mode is assigned to each application, the vehicle then launches, stops, or reconfigures all applications to reach the desired state.

However, our proposed approach requires an onboard algorithm to efficiently select modes within the vehicle resources. It also requires a fast response with an optional timeout t if an important application needs to be launched quickly.

In this work, we designed an efficient algorithm in three steps. First, we use a SAT-based solver to obtain a performance baseline. Then, we use a genetic algorithm to improve the ratio between solution quality (i.e. sum of AXIL ratings of running modes) and execution speed. Finally, due to limitations identified in Section IV-B, we propose a specialized heuristic to improve the results further and enable onboard implementation. We describe the algorithms in the following paragraphs and compare their performances in Section VI.

A. SAT Solver

First, we implemented our problem in an SMT (SAT-based) solver to obtain a performance baseline [32]. We use the well-established Z3 solver [33] with the objective function and constraints defined in Section III. This approach is guaranteed to return a valid solution if one exists. However, it has an exponential complexity in the worst case. We set a one-hour computation timeout in all scenarios for practical reasons. To model the solution in SAT, we declare C as an unknown vector of variables constrained according to the runtime mode indexes and mode dependencies. Our objective function can be directly translated as a minimization problem, and Z3 will efficiently explore the tree of solutions until the optimal solution is found.

Setup	Parameter	Value
Fixed	Population size n_G	50
	Ratio of elitist candidates	10%
	Mutation rate	1.15%
Variable	G_a Maximum generations	100 to 2000
	G^b Auto-stop after stagnation	None to 500

TABLE II: Parameters for the genetic algorithm after calibration. See Figure 3 for our choice on the variable parameters. The solver S and heuristic H do not require any parameters.

B. Genetic Algorithm

With the aim to embed our methodology onboard, the NP-hard nature of the problem requires a heuristic approach. While several meta-heuristics exist with the potential to solve our problem, we chose to focus on a genetic algorithm G [34] for several reasons. First, it is a well-known and widely used approach in the literature. Second, it is capable of evolving in a highly constrained environment, which is required in our case as mode dependencies create a discontinuous search space. Third, the vehicle can stop the algorithm at any time and use the best solution found so far in case of a timeout, which can yield a reasonably performant solution as the algorithm improves quicker in the first iterations. This flexibility cannot be obtained with other meta-heuristics, e.g. simulated annealing.

Defined parameters are summarized in Table II. While the algorithm could be extended with many state-of-the-art variants (e.g. adaptive parameters, diversity preservation), we selected a simple implementation for clarity. The algorithm defines a fixed size population $\{C_1, \dots, C_{n_G}\}$, a mutation rate to avoid overfitting, and a ratio of elitist candidates to improve convergence speed. The notation G_a^b refers to a genetic algorithm with a the maximum number of generations and, if set, b the number of generations to wait before stopping when the score has not improved (or when reaching the timeout t).

For each candidate C_i , the fitness function assesses the score of each candidate and is simply defined as the sum of AXIL scores for each selected mode, i.e. Equation 3. The algorithm starts with a random population of n_G candidates, then iteratively generates new candidates by selecting two parents using their scores as selection probability and applying a crossover and mutation operator. The crossover operator defines how two parent genomes should create two children candidates. We chose to select a random pivot point and swap the modes of the two parents to produce two new candidates. The mutation operator randomly selects a mode and shifts its index by 1. The best candidate is then returned as the final solution.

However, this approach has several limitations. First, the algorithm requires to store several tens of candidates in memory, which is not practical for integration in embedded systems. Second, the algorithm is not guaranteed to return a valid solution, which could hinder on the user experience and bring the vehicle to an undefined state. Finally, as indicated by the a and b parameters, there is no indication on when to stop the algorithm. This is problematic as the algorithm may run for a long time without improving the solution. We discuss these limitations in greater detail with results in Section VI.

Algorithm 1: Specialized Heuristic H

Input: \mathcal{A}_{req} set of applications, D graph of dependencies, P matrix of resource requirements, R resource budgets

Initialization: Vector C with n_{req} zeros (all apps off);

```

while True do
   $C^{search} \leftarrow \emptyset$ ; Local in-loop candidate
   $S \leftarrow 0$ ; Local real to hold  $C^{search}$  score
  for  $\{\mathcal{A}_i \in \mathcal{A}_{req} \mid \forall C[i] \neq 1\}$  do
     $C^{new} \leftarrow C$ ; Local copy of  $C$ 
    if  $C^{new}[i] = 0$  then
       $C^{new}[i] \leftarrow m_i$ ; Enable  $\mathcal{A}_i$ 
    else
       $C^{new}[i] \leftarrow C[i] - 1$ ; Upgrade  $\mathcal{A}_i^{C[i]}$ 
    end
     $C^{new} \leftarrow \text{bumpDependencies}(C^{new}, D)$ ; Section IV-C
     $s \leftarrow \text{fitness}(\mathcal{A}_i, C^{new}, C, P)$ ; Equation 4
    if  $s > 0$  and  $s > S$  then
       $C^{search} \leftarrow C^{new}$ ;  $S \leftarrow s$ ; Save new best
    end
  end
  if  $C^{search} = \emptyset$  or  $S = 0$  then
    return  $C$  End when no improvement left
  end
   $C \leftarrow C^{search}$ ; New best candidate to beat in next loop
end

```

Output: Final candidate C

C. Specialized Heuristic

To overcome the limitations of the genetic algorithm, we propose a specialized heuristic H defined by Algorithm 1 that is guaranteed to return a valid solution although not necessarily optimal. The algorithm starts by selecting a unique candidate C^0 where all applications are inactive. Then, the algorithm iteratively 'climbs up' the runtime modes one-by-one, starting from the most degraded modes, using a custom fitness function until no higher runtime mode can be selected. It stops when at least one resource budget $R[i]$ has been reached (saturated).

Each iteration follows the following visual metaphor. Using the current candidate C , we evaluate the *value-to-resource-usage* ratio improvement of upgrading each application by one mode (if not already at their maximum mode). This can generate up to n_{req} candidates per iteration. Then, the one with the best fitness score is set as the next iteration reference.

To respect the dependency constraints, this algorithm leverages the property defined in Equation 1. When a runtime mode \mathcal{A}_i^j is selected for an application \mathcal{A}_i , every mode from other applications that depend on \mathcal{A}_i must also be selected (or any higher mode). This property is enforced by the *bumpDependencies* function which recursively raises the missing mode dependencies until the current candidate respects them all.

Finally, the fitness function compares two candidates and yields an improvement score. This score divides the AXIL difference by a cost function. The cost is defined as the sum of additional resource usages requested by the upgraded mode. If selecting the new mode implies that at least one resource becomes saturated, then the candidate is invalidated in Algorithm 1. Additionally, the cost is designed to be lower




Fig. 3: Simulation results for (A) one particular parameter combination to study different algorithm options. H is our specialized heuristic, G_a^b is a genetic algorithm with a the maximum number of generations and b the number of generations to wait before stopping when the best score hasn't improved (if set), and S is the exact solver based on Z3. We compare the (B) relative score, success rate, and (C) solve time of each algorithm for 1000 random instance repetitions. This figure serves as an illustration as the overall results greatly depend on the problem size. See Figures 4 and 5 for a general performance assessment.

when a resource is filled in its lower capacity to encourage the progressive allocation of resources. Hence, we use the integral of the resource usage difference. The cost function for one resource k with budget b used from u to u' is defined as:

$$f(u, u', b) = \begin{cases} 0 & \text{if } u' > b \\ u'^2 - u^2 & \text{otherwise} \end{cases} \quad (= \frac{1}{2} \int_u^{u'} x dx)$$

Then, the fitness function sums the cost of each resource to obtain the final score or returns 0 if one resource is saturated:

$$\text{fitness}(\mathcal{A}_i, C^{\text{new}}, C, P) = \begin{cases} 0 & \text{if } 0 \in K \\ \frac{X[i, C^{\text{new}}[i]] - X[i, C[i]]}{\sum_K} & \text{otherwise} \end{cases} \quad (4)$$

$$\text{with } K = \{f(M_i^{C^{\text{new}}[i]}[k]), M_i^{C[i]}[k], R[k] \quad \forall k=1, \dots, r\}$$

The algorithm stops when no higher runtime modes can be selected once resources are used (i.e. when all scores are 0).

D. Algorithmic Complexity

Contrary to the genetic algorithm G which stores and maintains many candidates, our heuristic H only stores three candidate solutions during the search, two being temporary.

Our proposed heuristic is of polynomial complexity. Each iteration must evaluate n_{req} candidates, each dependent on a fitness function that sums r costs. In the worst case, the algorithm may need to upgrade every available mode one by one. Hence, the overall complexity is $\mathcal{O}(m \times n_{\text{req}} \times r)$. However, due to its efficient search strategy, the algorithm is fast in practice as demonstrated by our results in Section VI.

Finally, the guarantee that H will always return a valid solution is due to its search direction enabled by Equation 1. The first iteration activates one of the lowest runtime modes with the best *AXIL-to-resource-usage* ratio. As long as there is at least one mode that does not immediately saturate a resource, this heuristic will certainly activate at least one application. Then, by iteratively upgrading modes, the algorithm slowly fills the resources. We rely on the extensive simulation results presented in Section VI to discuss the overall performance.

V. SIMULATION SETUP

To evaluate the performance of each algorithm (S , G , and H) and assess their embedded feasibility, we designed a simulation setup with randomly generated applications, resource requirements, mode dependencies, and network flows. The list of parameters required to generate a random instance along with selected values for our first study is shown in Figure 3A.

The problem instance is generated as follows. First, we model a fixed physical E/E topology as a star network of 5 ECUs connected by Ethernet links of 100Mbps. The central ECU could be seen as the main onboard HPC, and the others as zonal (ZCU) or communications (CCU) ECUs. Note that our approach is independent of the topology type. Then, we generate a random number of n applications to represent an app store. Each application is randomly assigned to one of the ECUs, which we assume is provided by an external onboard service orchestrator out of the scope of this work. Additionally, we generate a random dependency graph with a target density by iteratively adding edges randomly to a graph up to a given target density, and removing one edge per cycle if any appears.

Each app \mathcal{A}_i has a random number of m_i runtime modes below a maximum value m_{max} . For each dependency edge in the app-level graph, we generate a new mode-level dependency graph D . We generate a random number of edges between modes of the two apps as long as they do not cross. Then, for each edge in D , we generate a random number of network flows with random bandwidth requirements. Finally, each mode \mathcal{A}_i^j is attributed a random resource requirement $M_i^j[k]$ for each resource k , i.e. CPU, RAM, and energy consumption in this work. Note that values are generated within the bounds set in Figure 3A. They also respect all constraints defined in Section III. See Figure 2 for an example of AXIL ratings generated for a particular instance with values in Figure 3A.

Finally, each value is stored in a 3D resource matrix P of r layers, i.e. made of 2D matrices for each resource. In total, there are $r = 21$ layers, namely the global energy usage requirement, CPU and memory requirements for each ECU, and network bandwidth usage for each direction of the

5 Ethernet cables. Bandwidth requirements for each flow are added in each link layer that is part of the packet route.

Each algorithm receives a given randomly generated problem instance as input and returns both their best candidate solution C and total solve time. The latter includes initialization steps such as creating assertion rules in the solver or population generation in G . H only sets a vector of n_{req} zeros.

VI. RESULTS

The simulation results greatly depend on the problem instance parameters. We first present the results for a single set of parameters to understand the main trends in Figure 3. Then, Figure 4 shows the evolution of performance depending on the number of requested applications, which is the main parameter that can change dynamically. Finally, we present a general performance assessment through a heatmap in Figure 5 for all parameter combinations. A clear trend will be observed.

With 1000 instances generated using parameter values in Figure 3A, Figure 3B shows the distribution of the total AXIL score for each algorithm relative to the value found by the solver (set as 100%). Figure 3C shows the solving time of each algorithm. Since algorithm G does not provide any stop condition, we studied 6 variants with different fixed stop points or once the score does not evolve after a number of iterations.

We observe that both the performance and solve times are similar between our heuristic H and G . However, while G_{100} runs faster than H , it only finds a valid solution in 13% of all instances. All other G variants take slightly more time to return a solution than H and have a success rate of 60 to 67%, while H always returns a valid solution with relative good performance. While there is a clear advantage of using H in this scenario (parameters set in Figure 3A), we then extended the comparison with other parameter combinations.

Since the problem is defined by many parameters, we chose to only vary the number of applications, modes, and density of applications. We believe this to be a reasonable simplification as all other parameters correspond to usually fixed app-level resource requirements. Figure 4 focuses on studying the sensitivity of results on the number of applications requested to launch. We believe that this parameter may vary the most at runtime. With 100 instances per problem size, we observe that the AXIL scores found by all algorithms are very close, which indicates good performance for both G and H . However, all variants of G fail after a *frontier* of complexity while H always succeeds. Note that the 1h timeout for S is reached at the last point for all instances. After this point, the performance of H cannot be compared, although it can still be used in practice.

Finally, Figure 5 shows a heatmap of the performance of each algorithm for all parameter combinations. Each cell represents the ratio of the median value from 30 instance repetitions on performance, solve time, and success rate. When a cell is empty, either S or G failed to find a solution for all repetitions while H always finds a solution for every instance. Hence, the success rate corresponds to S or G . Both algorithms fail after reaching their own frontier of problem complexity.

We observe that H is always the best performing algorithm in terms of performance and solve time. G is faster than H for




Fig. 4: Performance evaluation by moving the number of applications requested to launch from 5 to 100, with other parameters kept identical as Figure 3A. Each color follows the legend from Figure 3B. Lines represent the median value of 100 repetitions and corresponding areas the Q1 to Q3 quartiles.

small problem sizes, mainly with fewer mode dependencies, although the absolute computing time could be negligible for both cases. This is due to the fact that G is a genetic algorithm that relies on a population of candidates to evolve towards a better solution using random crossover operations. When dealing with many mode dependencies, the search space becomes discontinuous (like holes in cheese) and the population struggles to generate valid children. This is also why G is faster than H for small problem sizes, as crossovers and mutations are very effective under low constraints.

While the problem could potentially be solved by more efficient approaches than our proposed heuristic, the proposed algorithm in this work is both fast and always returns a valid solution. This is a key requirement for onboard runtime decisions, as the algorithm must always return a valid solution within a reasonable time frame. In case where H fails to find a solution within the timeout t , we can use its latest candidate solution as a fallback and optionally continue the search later.

Further implementation options in an industrial setting can be considered. For instance, the algorithm can be initiated with a few pre-selected applications to guarantee their activation. Background dependencies can also have AXIL ratings set to zero to only enable them when required by other features.

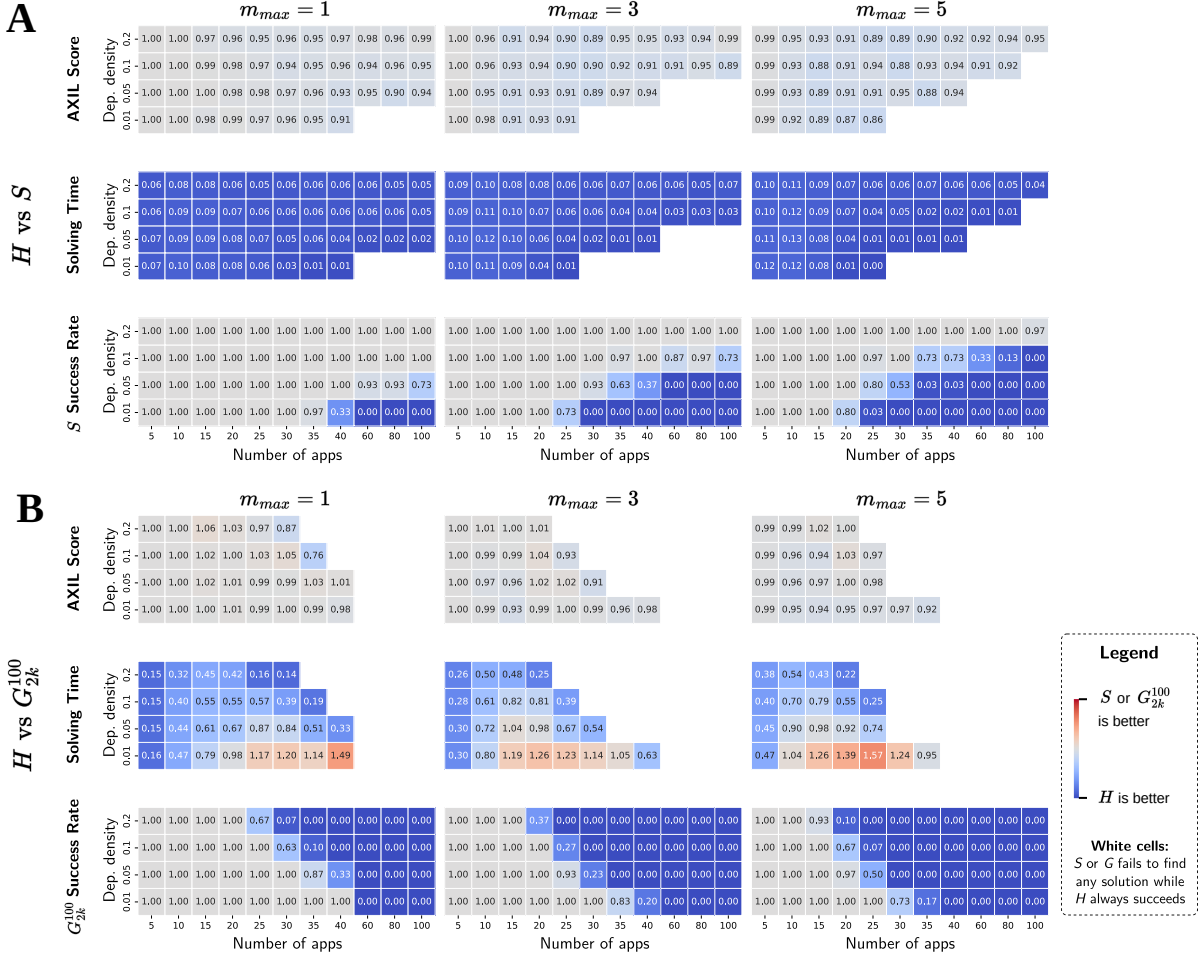


Fig. 5: Multi-parameter performance comparison of our proposed heuristic algorithm H with the (A) ILP solver S and (B) genetic algorithm G_{2k}^{100} chosen manually as the best competitor variant. Each cell represents the ratio of the median value from 30 instance repetitions on performance, solve time, or success rate. When a cell is empty, either S or G failed to find a solution for all repetitions while H always finds all solutions. Hence, the success rate corresponds to S (1h timeout) or G .

VII. DISCUSSION

By introducing AXIL, we propose a way of expressing a launch priority based on the user’s perceived quality of experience for each runtime mode. Our approach aims to solve the problem of orchestrating a high number of applications in a highly constrained embedded system. This concept led us to build a resource-based model of the onboard vehicle state, which in turn enables optimization algorithms to maximize the total onboard AXIL score given the resource constraints.

This approach can be particularly useful in two situations. First, the user and background services can be requested to launch without limit, and our algorithm selects the best experience-to-resource ratio to accommodate for the current global onboard state. This reduces the need for state management which is not mandatory for non safety-critical applications thereby reducing engineering effort. It also increases the number of active applications by potentially executing more of them in slightly degraded modes. Second, AXIL ratings can change depending on the user preferences and subscriptions, which impacts the autonomous runtime decisions and tailors

the experience to the user. Hence, SLAs can be respected by tweaking AXIL ratings based on the user profile or contract.

Our heuristic starts with all applications disabled and gradually improves runtime modes based on an AXIL and resource-based cost function. This approach offers great potential and flexibility in an industrial setting, allowing automakers to initiate the algorithm with a few pre-selected modes thereby guaranteeing their activation. They are guided by marketing strategies, user preferences, state management constraints, safety-critical prerequisites, or V2X service priorities.

Additionally, while this work focuses on onboard resources, our model allows for a flexible incorporation of various other resource types to enhance the overall system. For instance, edge computing resources near the current vehicle location or internal peripherals on each ECU can be added as additional resources. Finally, automakers can add a “digital eco-mode” onboard by considering energy consumption as an auxiliary resource, empowering users to select their preferred trade-off between user experience and environmental impact. This extended metric could include network infrastructure and data center usage to evaluate each runtime mode’s all-round impact.

VIII. CONCLUSION

We have presented a methodology for orchestrating non-safety-critical onboard applications based on the available vehicle resources to cover the dynamic future automotive use cases. Once developers define several degraded modes and dependencies for each application, our approach then selects the highest possible modes without exceeding the current onboard resources while maximizing the overall user experience. Our proposed heuristic achieves a performance-to-speed ratio that enables onboard adjustments at runtime while also guaranteeing to always return a valid solution, contrary to traditional approaches like using a SAT solver or a genetic algorithm. Finally, the vehicle calls our algorithm on context or resource availability changes and dynamically applies the calculated runtime modes for a continuous user experience.

Future work will aim to (1) validate the industrial feasibility of our proposal using a physical test bench with realistic use cases and resources, (2) evaluate our algorithm's potential to influence on the onboard energy consumption, (3) dynamically optimize mode selection based on a future timeline of predicted events or resource usage changes, and (4) simulate a multi-vehicle environment to account for cooperative V2X use cases and edge computing resource distribution.

REFERENCES

- [1] M. Haerberle, F. Heimgaertner, H. Loehr, N. Nayak, D. Grewe, S. Schildt, and M. Menth, "Softwarization of Automotive E/E Architectures: A Software-Defined Networking Approach," in *2020 IEEE Vehicular Networking Conference (VNC)*, Dec. 2020, pp. 1–8.
- [2] X. Hu, K. Wang, X. Liu, Y. Sun, P. Li, and S. Guo, "Energy management for EV charging in software-defined green vehicle-to-grid network," *IEEE Communications Magazine*, vol. 56, no. 5, pp. 156–163, 2018.
- [3] F. Zhao, H. Song, and Z. Liu, "Identification and Analysis of Key Technical Elements and Prospects for Software-Defined Vehicles," SAE International, Warrendale, PA, SAE Technical Paper 2022-01-7002, Jan. 2022, ISSN: 0148-7191, 2688-3627.
- [4] V. Bandur, G. Selim, V. Pantelic, and M. Lawford, "Making the Case for Centralized Automotive E/E Architectures," *IEEE Transactions on Vehicular Technology*, vol. PP, pp. 1–1, Jan. 2021.
- [5] W. Zeng, M. A. S. Khalid, and S. Chowdhury, "In-Vehicle Networks Outlook: Achievements and Challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 1552–1571, 2016.
- [6] M. Rumez, D. Grimm, R. Kriesten, and E. Sax, "An Overview of Automotive Service-Oriented Architectures and Implications for Security Countermeasures," *IEEE access : practical innovations, open solutions*, vol. 8, pp. 221 852–221 870, 2020.
- [7] C. Lin, G. Han, J. Du, T. Xu, and Y. Peng, "Adaptive Traffic Engineering Based on Active Network Measurement Towards Software Defined Internet of Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 6, pp. 3697–3706, Jun. 2021, conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [8] H. Askaripoor, M. Hashemi Farzaneh, and A. Knoll, "E/E Architecture Synthesis: Challenges and Technologies," *Electronics*, Jan. 2022.
- [9] P. Dai, K. Liu, Q. Zhuge, E. H.-M. Sha, V. C. S. Lee, and S. H. Son, "Quality-of-Experience-Oriented Autonomous Intersection Control in Vehicular Networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 7, pp. 1956–1967, Jul. 2016, conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [10] X. Li, C. F. Chiasserini, J. Manges-Bafalluy, J. Baranda, G. Landi, B. Martini, X. Costa-Pérez, C. Puligheddu, and L. Valcarenghi, "Automated Service Provisioning and Hierarchical SLA Management in 5G Systems," *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4669–4684, Dec. 2021.
- [11] A. Frigerio, B. Vermeulen, and K. Goossens, "Component-Level ASIL Decomposition for Automotive Architectures," in *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, Jun. 2019, pp. 62–69.
- [12] S. Kugele, P. Obergfell, M. Broy, O. Creighton, M. Traub, and W. Hopfensitz, "On Service-Oriented Architecture for Automotive Software," in *2017 IEEE International Conference on Software Architecture (ICSA)*, Apr. 2017, pp. 193–202.
- [13] M. Schindewolf, H. Stoll, H. Guissouma, A. Puder, E. Sax, A. Vetter, M. Rumez, and J. Henle, "A Comparison of Architecture Paradigms for Dynamic Reconfigurable Automotive Networks," in *2022 International Conference on Connected Vehicle and Expo (ICCVE)*, Mar. 2022.
- [14] J. Firtunikj, M. Armbruster, and A. Knoll, "Run-Time Adaptive Error and State Management for Open Automotive Systems," in *2014 IEEE International Symposium on Software Reliability Engineering Workshops*, Nov. 2014, pp. 467–472.
- [15] X. Liao, Z. Wang, X. Zhao, K. Han, P. Tiwari, M. J. Barth, and G. Wu, "Cooperative Ramp Merging Design and Field Implementation: A Digital Twin Approach Based on Vehicle-to-Cloud Communication," *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [16] B. Shi, X. Tu, B. Wu, and Y. Peng, "Recent Advances in Time-Sensitive Network Configuration Management: A Literature Review," *Journal of Sensor and Actuator Networks*, vol. 12, no. 4, p. 52, Aug. 2023.
- [17] K. Halba, C. Mahmoudi, and E. Griffor, "Robust Safety for Autonomous Vehicles through Reconfigurable Networking," *Electronic Proceedings in Theoretical Computer Science*, vol. 269, pp. 48–58, Apr. 2018.
- [18] T. Hackel, P. Meyer, F. Korf, and T. C. Schmidt, "Software-Defined Networks Supporting Time-Sensitive In-Vehicular Communication," in *IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*, 2019.
- [19] T. Haugg, M. F. Soltani, T. Hackel, P. Meyer, F. Korf, and T. C. Schmidt, "Simulation-based Evaluation of a Synchronous Transaction Model for Time-Sensitive Software-Defined Networks," p. 9, 2021.
- [20] Y. Peng, B. Shi, T. Jiang, X. Tu, D. Xu, and K. Hua, "A Survey on In-Vehicle Time-Sensitive Networking," *IEEE Internet of Things Journal*, vol. 10, no. 16, pp. 14 375–14 396, Aug. 2023.
- [21] D. F. Külzer, S. Stańczak, and M. Botsov, "Novel QoS Control Framework for Automotive Safety-Related and Infotainment Services," in *2020 IEEE Wireless Communications and Networking Conference (WCNC)*, May 2020, pp. 1–7.
- [22] P. Laclau, S. Bonnet, B. Ducourthial, X. Li, and T. Lin, "Predictive Network Configuration with Hierarchical Spectral Clustering for Software Defined Vehicles," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, Jun. 2023, pp. 1–5.
- [23] R. Ranjan, B. Benatallah, S. Dustdar, and M. P. Papazoglou, "Cloud Resource Orchestration Programming: Overview, Issues, and Directions," *IEEE Internet Computing*, vol. 19, no. 5, pp. 46–56, Sep. 2015.
- [24] A. Tosatto, P. Ruiu, and A. Attanasio, "Container-Based Orchestration in Cloud: State of the Art and Challenges," in *2015 Ninth International Conference on Complex, Intelligent, and Software Intensive Systems*, Jul. 2015, pp. 70–75.
- [25] Y. Jiang, Z. Huang, and D. H. K. Tsang, "Challenges and Solutions in Fog Computing Orchestration," *IEEE Network*, May 2018.
- [26] N. Nayak, D. Grewe, and S. Schildt, "Automotive Container Orchestration: Requirements, Challenges and Open Directions," in *2023 IEEE Vehicular Networking Conference (VNC)*, Apr. 2023, pp. 61–64.
- [27] M. Schindewolf, D. Grimm, C. Lingor, and E. Sax, "Toward a Resilient Automotive Service-Oriented Architecture by using Dynamic Orchestration," in *2022 IEEE 1st International Conference on Cognitive Mobility (CogMob)*, Oct. 2022, pp. 000 147–000 154.
- [28] J. C. Nobre, A. M. de Souza, D. Rosário, C. Both, L. A. Villas, E. Cerqueira, T. Braun, and M. Gerla, "Vehicular software-defined networking and fog computing: Integration and design principles," *Ad Hoc Networks*, vol. 82, pp. 172–181, Jan. 2019.
- [29] P. Rani, N. Hussain, R. Khan, Y. Sharma, and P. Shukla, "Vehicular Intelligence System: Time-Based Vehicle Next Location Prediction in Software-Defined Internet of Vehicles (SDN-IoV) for the Smart Cities," Oct. 2021, pp. 35–54.
- [30] F. Gurski, C. Rehs, and J. Rethmann, "Knapsack problems: A parameterized point of view," *Theoretical Computer Science*, Jul. 2019.
- [31] V. Cacchiani, M. Iori, A. Locatelli, and S. Martello, "Knapsack problems — An overview of recent advances. Part II: Multiple, multidimensional, and quadratic knapsack problems," *Computers & Operations Research*, vol. 143, p. 105693, Jul. 2022.
- [32] W. Gong and X. Zhou, "A survey of SAT solver," *AIP Conference Proceedings*, vol. 1836, no. 1, p. 020059, Jun. 2017.
- [33] N. Bjørner, A.-D. Phan, and L. Fleckenstein, "vZ - An Optimizing SMT Solver," in *Tools and Algorithms for the Construction and Analysis of Systems*, ser. Lecture Notes in Computer Science, C. Baier and C. Tinelli, Eds. Berlin, Heidelberg: Springer, 2015, pp. 194–199.
- [34] P. Chu and J. Beasley, "A Genetic Algorithm for the Multidimensional Knapsack Problem," *Journal of Heuristics*, vol. 4, pp. 63–86, Jun. 1998.




Pierre Laclau is a PhD student at the University of Technology of Compiègne (UTC, France) and associated with Stellantis. He has a double degree of both Computer Science Engineering (with a specialization on real-time, embedded, and complex Systems) and a Master of Science (specialized on autonomous and intelligent systems). His research interests lie in multi-agent and complex systems architecture design applied to the development of next generation E/E and software architectures.



Stéphane Bonnet received the M.Sc. degree in computer science and the Ph.D. degree in control engineering from the Université de Technologie de Compiègne, Compiègne, France. He is currently a Research Engineer at this same institution and has contributed to several research projects, covering a wide array of related fields in embedded systems engineering, robotics, and communications. His current main interests are autonomous vehicles, related vehicle-to-vehicle and vehicle-to-infrastructure communication systems, and their impact on smart cities.



Bertrand Ducourthial Bertrand Ducourthial received the Ph.D. degree in computer science from the University of Paris-Sud, Orsay, France, in 1999. He joined the Heudiasyc Laboratory, Université de Technologie de Compiègne, Compiègne, France, where he became a Full Professor in 2010. His research deals with dynamic ad hoc networks (e.g., VANET and UAVs): modeling, networking, distributed algorithms, security, robustness, embedded software, etc. His work is funded by industrial, regional, national, and European projects.



Xiaoting Li is an In-Vehicle Network engineer in the Electrical and Electronics Systems department at Stellantis since 2018. She holds a Ph.D. degree in Computer Science from the National Polytechnic Institute of Toulouse, France in 2013. Prior to joining Stellantis, she worked in high education sector as an associate professor in real-time network and system domain. She is currently working on In-Vehicle Network technologies, with a particular focus on Ethernet and associated protocols such as AVB/TSN. She is also highly interested and involved

in research work on Software-Defined Vehicle for the next generation of E/E automotive architecture.



Trista Lin is an onboard IT architect and technical specialist in the E/E architecture design department at Stellantis since 2018. She is responsible for TCP/IP architecture design and protocol deployment. Prior to Stellantis, she has 8 years of industry experience in wireless network simulators and 3 years of research experience in open source software development. She holds a PhD degree in computer science from INSA Lyon (National Institute of Applied Sciences of Lyon), France, and a BS in mathematics and communications engineering

from National Tsing Hua University, Taiwan. Her research interests lie in IT solution adaptation for cars towards software-defined architectures and services.