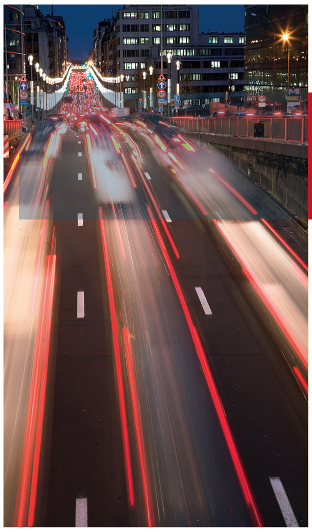

Editor: **Rosaldo J.F. Rossetti,** University of Porto, rossetti@fe.up.pt

Dynamic Coordination in Fleet Management Systems: Toward Smart Cyber Fleets

Holger Billhardt, *Universidad Rey Juan Carlos*
Alberto Fernández, *Universidad Rey Juan Carlos*
Lissette Lemus, *Artificial Intelligence Research Institute, Spanish National Research Council*
Marin Lujak, *Universidad Rey Juan Carlos*
Nardine Osman, *Artificial Intelligence Research Institute, Spanish National Research Council*
Sascha Ossowski, *Universidad Rey Juan Carlos*
Carles Sierra, *Artificial Intelligence Research Institute, Spanish National Research Council*

"Fleet management system" (FMS) is a term used for a wide range of solutions for different fleet-related applications in the fields of transportation, distribution, and logistics. It comprises target-based planning, as well as supervision and control of fleet operations based on available transportation resources and application constraints. FMSs have as an objective to reduce risk, increase quality of service, and improve a fleet's operational efficiency while minimizing its costs.[1]

A key problem in FMS operations is *fleet-route planning*, where different transport orders need to be aggregated into tours of fleet vehicles so that the resulting schedule is both efficient and robust while meeting the constraints put forward in customer requests. The main challenge at the tactical level is to support decision making based on seasonality, trends, changing customer mix, and demand. At operational and real-time levels, the challenge is to respond to daily dynamics, such as traffic, weather, employee absence, equipment breakdown, new orders, and order adjustments.

Approaches to fleet planning typically focus on the development of near-optimal plans using various types of effective vehicle-routing algorithms, which can be either static or dynamic.[2–4] Fleet schedules designed a priori with *static route planning* assume the following: all relevant data is known before the planning starts, short- and long-term decisions have the same importance, and the time available for creation, verification, and implementation of route plans is of minor importance. The use of an initial fleet schedule, although necessary, is by no means sufficient because it might not cope adequately with unexpected events during execution, such as traffic delays, vehicle breakdowns, road works, and new customer requests or the cancellation of preexisting ones, which causes fleet delays, unexpected costs, and poor customer service. *Real-time dynamic FMSs* are needed to handle unexpected events—that is, to detect deviations from the initial dispatch plan and adjust the schedule accordingly by suggesting effective rerouting immediately. In this context, timely decisions are very important because the time available for verification, correction, and implementation of changed route plans is often very short.[5,6]

Real-time FMSs have been applied to a broad variety of domains, including emergency vehicles (fire trucks, ambulances, and so on), police cars, taxis, commercial delivery vehicles, courier fleets, public transport fleets, and freight railcars.[7–11] However, state-of-the-art FMS solutions are *centralized* and require vehicle fleet operators to send low-level commands remotely to the fleet's drivers and their vehicles. Even though some dynamics of the environment are accounted for, certain decisions can't

1541-1672/14/$31.00 © 2014 IEEE
Published by the IEEE Computer Society

be reconsidered because doing so could complicate the assignment procedure and potentially compromise the fleet's response time. In addition, changes in the environment aren't always communicated fast enough to FMS operators to help them make decisions in a timely manner. The FMS's dependence on adequate central operator decisions, therefore, compromises its robustness and hinders effective scalability.[12]

The technological advances in sensors, communication and networking technologies, and geographic information systems enables fleet operators to be informed about unexpected changes in fleet operation almost at the time that they occur, and thus allows for increased levels of dynamicity in operational decisions. Moreover, the increased performance of small-scale, energy-efficient computing devices allows for delegating part of fleet decision making to the fleet's vehicles, enabling a more decentralized FMS architecture that gives more autonomy to the vehicles and their drivers and, thus fosters the system's reactiveness. Here, we sketch our work in the field of real-time FMS and point to developments that we believe are going to take place in the near future. We put forward our vision of conceiving FMS as smart cyber-physical systems, and illustrate the idea in the field of electro-mobility, where drivers of smart e-motorbikes (*cyber vehicles*),[13] equipped with an intelligent communication device (*cyber helmet*), are coordinated by means of a next-generation FMS.

Dynamic Fleet Management

We propose to employ an event-based architecture for dynamic fleet management. We applied this architecture to the coordination of a fleet of ambulances in a medical emergency scenario and show experimentally that our proposal outperforms a non-dynamic approach.

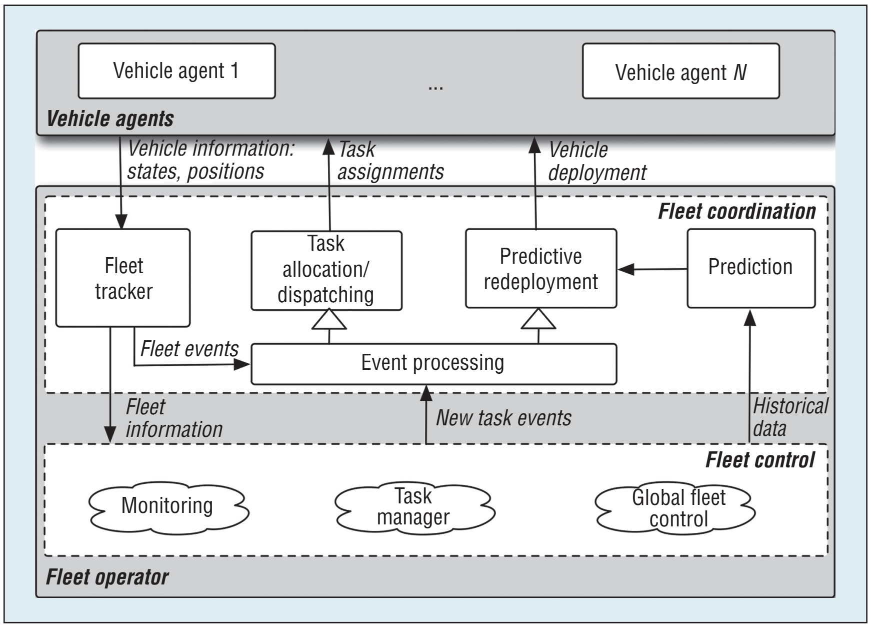

Figure 1. Event-based architecture for dynamic fleet management. The top layer contains the vehicles, modeled as agents; the second layer represents the fleet coordination modules; and the third layer includes other components necessary for normal fleet operation.

Event-Based FMS Architecture

Fleet operators face two main problems: task allocation and redeployment. The *allocation* problem consists of determining which vehicle should be sent to serve a given task. *Redeployment* consists of relocating vehicles in the region of influence in a way that new tasks can be reached quickly at a low cost. Both issues are particularly challenging in dynamic environments, as continuously arriving new tasks might require attendance, and the fleet's current situation might change due to external influences. To maximize vehicle utilization and improve service quality in such environments, task allocation and vehicle redeployment should also be accomplished in a dynamic manner, adapting the fleet's coordination seamlessly to upcoming events and changing demands. To adequately capture real-time requirements in such a scenario, we set out from an event-driven approach.[14]

Figure 1 depicts our architecture for dynamic fleet management. It contains three basic layers: the top layer contains the vehicles, modeled as agents; the second layer represents the fleet coordination modules; and the third layer includes other components necessary for normal fleet operation (components for monitoring, task management, global fleet control, and so on).

In the fleet coordination layer, a fleet tracker follows the vehicles' operational states and positions. (We assume that vehicles have the ability to send their current positions on a regular basis and to inform about changes in their operational states.) The fleet tracker informs the event-processing module about any changes in the fleet that would require an adaptation of task allocations or the deployment of idle vehicles. If necessary, it triggers task allocation and predictive redeployment modules. The task allocation module, when executed, recalculates the optimal global assignment of all pending tasks (in the current moment) to vehicles based on a set of assignment criteria (depending on application domain). The predictive redeployment module

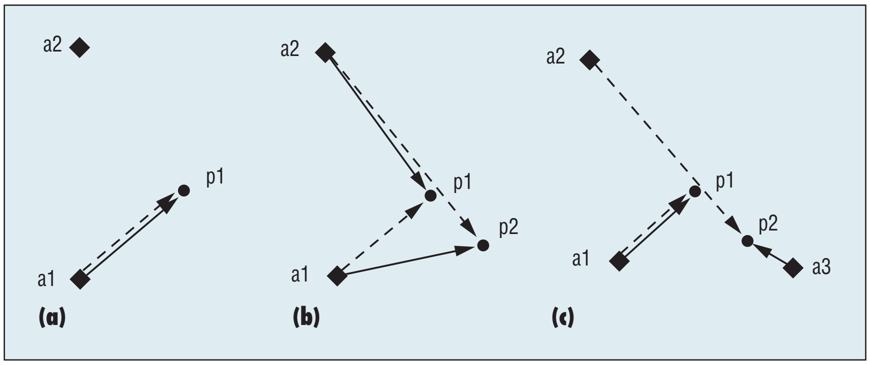

Figure 2. Ambulance assignment strategies. Dotted lines show the current first-come, first-served (FCFS) approach, while solid lines represent our assignment policy: (a) initial assignment, (b) a new patient appears, and (c) a previously busy ambulance becomes available.

calculates adequate positions for all idle vehicles at the current moment, taking into account predictions concerning the appearance of new tasks (based on historical data) and the fleet's current state.

Case Study: Ambulance Coordination in Madrid

To test our approach, we applied the architecture in a system for coordinating a fleet of ambulances for the Emergency Medical Assistance Service, SUMMA112 (www.madrid.org/cs/Satellite?pagename=SUMMA112/Page/S112_home), in Madrid, Spain—a region of about 8,000 km^2 with approximately 6 million inhabitants. Among other services, SUMMA112 maintains a fleet of ambulances that provides out-of-hospital assistance to patients in cases of emergencies. One of the main goals is to reduce response time (the time between when a patient calls and the moment the ambulance arrives) in life-threatening emergencies: shorter response times are directly correlated with lower mortality rates.

SUMMA112 currently employs a static approach to patient allocation and ambulance redeployment: calling patients are classified by a triage system into different severity levels, and patients at the highest level are assigned using the first-come/first-served (FCFS) principle—that is, the first patient in the system is assigned first, then the next patient, and so on. In each case, a patient is assigned to the closest available ambulance at that particular moment. After an ambulance has finished a mission, it returns to its base station and waits for a new assignment. The locations of ambulance base stations are fixed and have been chosen based on criteria such as population density and infrastructure.

Based on the architecture presented in Figure 1, we developed a prototype of a dynamic ambulance FMS in which ambulances and calling patients are the vehicle agents and new tasks, respectively. We concentrated only on the most severe emergency cases, those that are assisted with advanced life support units. Regarding the assignment of patients to ambulances (task allocation), we substitute the current static FCFS strategy with a reactive method in which existing assignments can be reconsidered on the fly. As Figure 2 illustrates, any patient who was already assigned to an ambulance might be reassigned to another one if this improves the average response time. In particular, a given assignment is recalculated if a new patient has appeared (see Figure 2b) or an ambulance has finished a previous mission (see Figure 2c). For this purpose, we used Dimitri Bertsekas's auction algorithm,[15] which assures an assignment that minimizes the key performance indicator (average ambulance response time) in a sufficiently fast manner.

With respect to ambulance redeployment, we use historical data to estimate the probability distribution of emergency cases in the region for different days and times of day (one-hour intervals). Based on this estimation, we calculate adequate waiting positions for all ambulances that are idle at a given moment. The waiting positions are dynamically recalculated if one of the following events occurs: an ambulance previously assigned to a patient becomes idle again (the mission is finished or the ambulance is de-assigned from a patient), an idle ambulance is assigned to a patient, or a different estimation of the probability distribution needs to be applied (every hour). We implemented the redeployment module based on the calculation of centroidal Voronoi tessellations, a geometric optimization technique that allows to estimate sub-optimal positions of a set of *generator points* in an Euclidean space and such that the weighted distance of all points in the space to the closest generator is minimized.[16]

To evaluate our dynamic approach's effectiveness, we tested it in a set of experiments analyzing the response times to emergency patients. For this purpose, we developed a simulation tool for emergency medical assistance (EMA) services, covering the whole assistance process—the emergence of patients, the schedule of an ambulance, the "in situ" attendance, and finally, the transfer of patients to hospitals—based on the information obtained from a well-calibrated, external route service. In our experiments, we considered a rectangle of 125 × 133 kilometres that covers the whole area of Madrid. We used 29 hospitals (all located at their real positions) and 29 ambulances with advanced life support (as currently used by SUMMA112). We

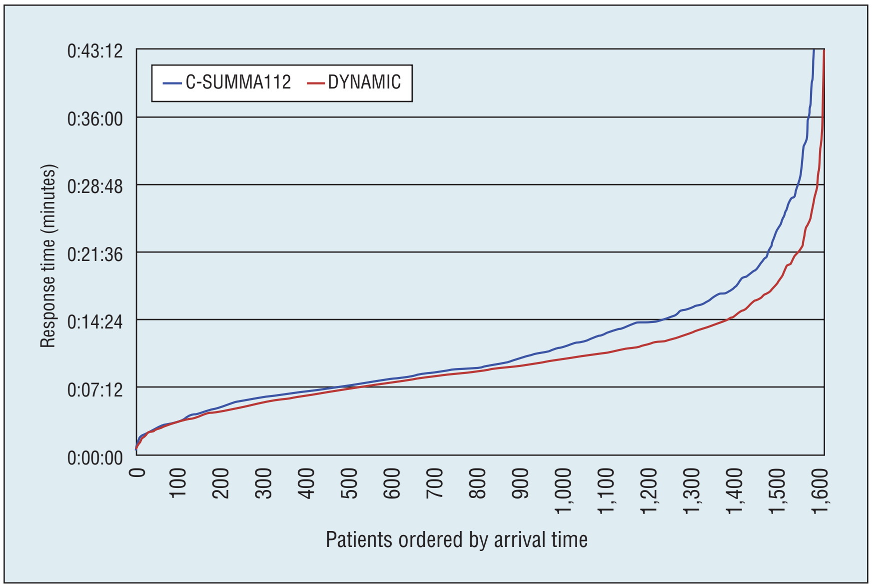

Figure 3. Dynamic versus static fleet management for ambulance coordination in Madrid. The response times in minutes over all patients (1,609 in total) for both the current FCFS coordination model (C-SUMMA112) and our dynamic coordination model (DYNAMIC) clearly show the benefits of our dynamic approach.

simulated the operation of the service for 10 different days (24-hour periods) with real patient data from 2009 provided by SUMMA112. We chose the days that gave us a good representation of high, medium, and low workloads.

Figure 3 compares the distribution of the response times in minutes over all patients (1,609 in total) for both the current FCFS coordination model (C-SUMMA112) and our dynamic coordination model (DYNAMIC). The results clearly show the benefits of our dynamic approach, which performs better for practically all response time ranges. Furthermore, the most important improvements can be observed in the ranges of higher response times. This is an important advantage, because it assures that more patients can be attended within given response time objectives. On average, the response times are 15.8 percent better in the DYNAMIC approach (9:54 versus 11:45 minutes). Especially in severe cases, a reduction of almost 2 minutes can be potentially lifesaving.

Toward Cyber Fleets

Prototypes of autonomous vehicles have been designed and tested, with the main challenges related to this new technology currently being studied in countless realistic and complex scenarios (www.cybercars.org).[17] At some point, this poses new challenges to FMSs as they attempt to manage fully autonomous vehicles in a decentralized manner. Based on currently available technologies, we're studying the impact of different types of sensors and driver assistance technologies on fleet management. In particular, with the goal of improving the efficiency, safety, and autonomy of vehicle fleets and their drivers, we propose an FMS as a smart, cyber-physical system (*cyber fleet*) made of cyber vehicles and drivers with cyber interfaces. In such a scenario, FMS decision making takes place both at the vehicle level (the drivers interact with their own and other cyber vehicles through cyber interfaces), as well as at the system level, where fleet operators can focus on more coarse grained management decisions for fleets that are potentially heterogeneous and large scale.

Figure 4 outlines our proposed FMS-based cyber fleet. The coordination cloud is similar to the dynamic FMS outlined in Figure 1, but many low-level events can be coped with locally in the cyber vehicles. That is, management and monitoring tasks are shared between the fleet operator and the cyber vehicles with their drivers.

We can illustrate the notion of cyber fleets through an example in the field of electro-mobility. The company GoingGreen (www.goinggreen.es), for instance, is deploying fleets of e-motorbikes in the city of Barcelona for vehicle sharing and home delivery purposes. The cyber fleet of e-motorbikes that we propose comprises three main components: cyber helmet (CH), cyber e-motorbike (CeM), and smart e-motorbike FMS (SeM-FMS).

Smart helmets are currently finding their way into the market (see Figure 5). However, the CH for a cyber fleet of e-motorbikes needs to go beyond the state of the art, insofar as it serves as a smart communication bridge between the driver and the vehicle, and between the driver and the SeM-FMS. For this purpose, it's equipped with additional communication outlets, a stereo camera, and a microphone, and it's connected to the CeM to take advantage of its computing capacity.

The interaction between the CH and the driver has to be grounded in situational awareness: the former should refrain from communicating with the driver during difficult maneuvering operations or traffic situations, which require the driver's full attention. In particular, a traffic evaluation module ought to take into account traffic images received through the camera, the driver's current maneuvering complexity based on the CeM's GPS coordinates and the actual traffic state, weather conditions, the road infrastructure complexity, and CH sensor readings about the CeM's current state (acceleration,

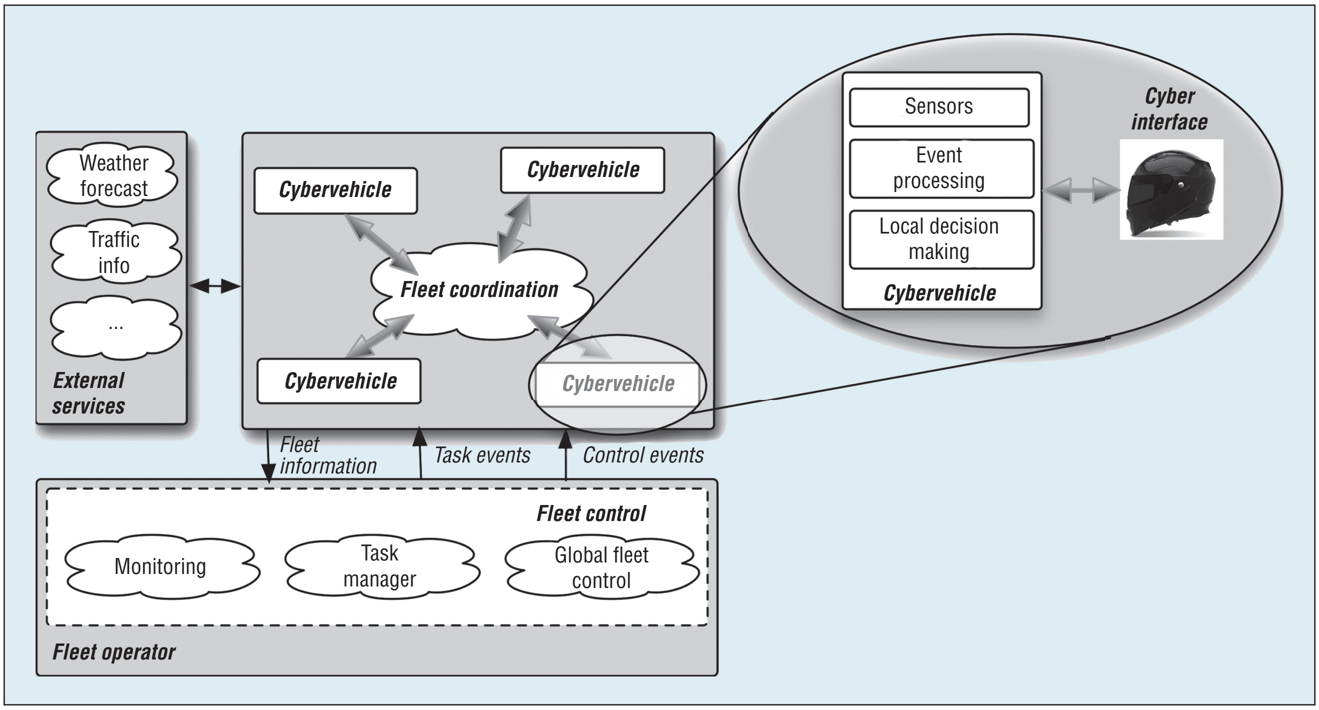

Figure 4. Structure of our fleet management system (FMS) for cyber fleets. The coordination cloud is similar to the dynamic FMS outlined in Figure 1, but many low-level events can be coped with locally in the cyber vehicles.

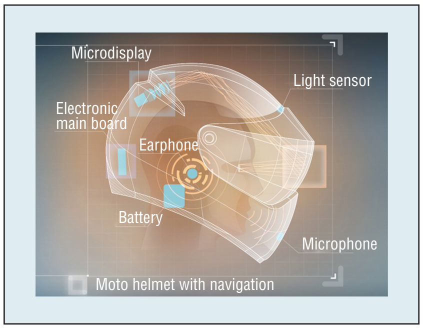

Figure 5. Smart motorbike helmet designed by the Russian company LiveMap (https://livemap.info).

Figure 6. E-motorbike belonging to GoingGreen's fleet (www.goinggreen.es).

velocity dynamics, wheel orientation, and so on). The identification of the traffic situation is possible through image recognition, fusion of data received from different helmet and vehicle sensors, and sensor knowledge extraction. In addition, the CH can determine the mode of communication (audio communication through microphone, video presentation of data on the helmet's augmented reality display, or a combination of both) and inform the driver about his or her tasks, as well as CeM and traffic conditions (malfunctions, battery, driving performance, security alerts, weather, traffic accidents, traffic jams, alternative routes, and so on).

The e-motorbikes that GoingGreen currently deploys (see Figure 6) are already equipped with simple sensors and basic data-processing capacity. We can enhance them with additional data sources such as accelerometer, proximity (laser) sensors, stereo cameras, and so forth, and will turn them into a CeM by endowing them with additional computing power. With this configuration, the CeM will perform real-time sensor data extraction, fusion, and reasoning, and communicate with the driver through the CH connected to the vehicle's battery and to the SeM-FMS through standard wireless communication. Some of the exemplary vehicle processes are forecasting the residual battery autonomy with a specific driver profile, maintaining a driver profile based on driving habits, and networking with other vehicles and SeM-FMSs in the system for task and work break distribution, contingency coverage, and so on. The CeM assists decision making about the mission's execution. It might receive and directly execute commands from the SeM-FMS about maximum speed limit, maximum acceleration, and engine blocking, so as to enhance energy efficiency and vehicle security. But it can also suggest directly to the driver information such as the most adequate charging spots.

The SeM-FMS is the computational platform that ultimately satisfies fleet objectives. Its level of decentralization in decision making is customizable to the fleet owner's preferences and

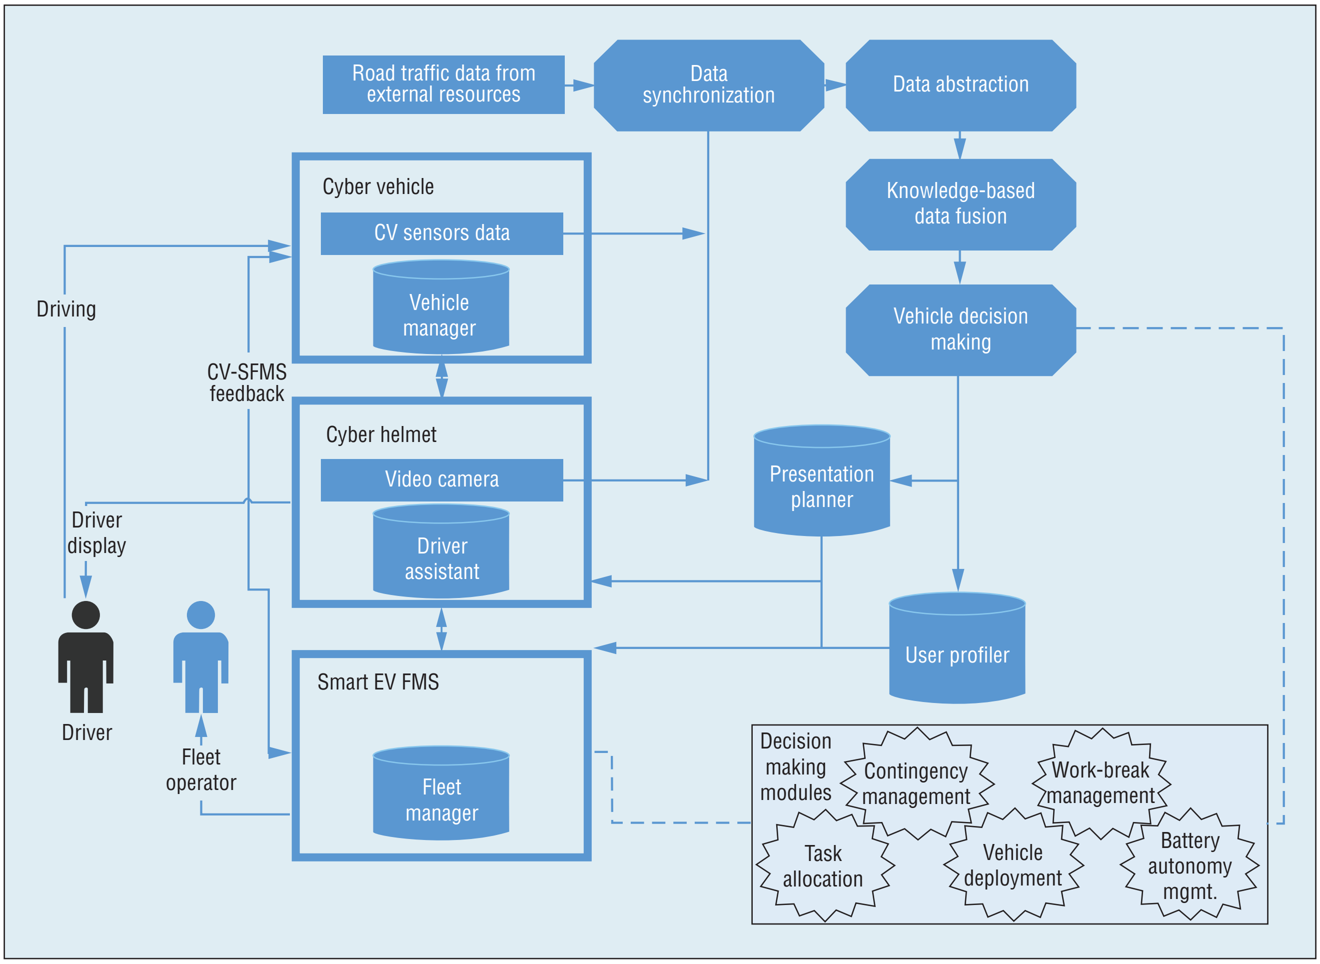

Figure 7. Cyber fleet's customizable decision-making structure. Depending on the assignment of decision-making modules to the fleet manager or vehicle fleet, different levels of decentralization can be achieved.

constraints; it can vary from a fully centralized to a subsidiary option (see Figure 7). Subsidiarity is an organizing principle of decentralisation, promoting the delegation of responsibilities to the smallest, lowest, or least centralized authorities capable of addressing an issue effectively (in our case, drivers and their CeMs).[18] In a centralized structure, all decisions regarding a fleet's strategic, tactical, and operational levels related to task allocation, vehicle deployment, battery autonomy, work break, and contingency management are controlled by the central fleet operator. In contrast, with fully decentralized control, the fleet operator only controls the overall fleet strategy (its mission and related constraints); CeMs manage key parts of mission execution and related operations at tactical and operative levels in real time through lateral interactions. Besides constraints emanating from the concrete organizational environment in which the SeM-FMS is embedded, the fleet's level of decentralization depends, for instance, on its size and dispersion over one or more regions of interest. To facilitate individual accounting for fleet performance, one of the tools for mission evaluation track's a personal driver profile record. If necessary, the SeM-FMS can undertake corrective actions on fleet vehicles and drivers to minimize performance degradation during sudden performance variations.

We plan to further explore cyber fleets of e-motorbikes in two case studies with GoingGreen. In home delivery, the task of electric motorcycles is to distribute products in an urban area by assigning CeMs to product pick-up and delivery tasks. The vehicle-sharing scenario refers to a type of vehicle rental for short periods of time, often a matter of hours. The principle of vehicle sharing is that individuals gain the benefits of private transportation without the costs and responsibilities of ownership. Instead, a private user accesses a fleet of vehicles on an as-needed basis. We also plan to integrate both business cases into one business solution, where a fleet of CeMs serves both purposes at the same time, and dynamically and seamlessly adapts to user demands to maximize vehicle utilization and increase profit gains. In future work, we also intend to extend our approach to mixed cyber fleets capable of managing heterogeneous fleets of traditional vehicles, cyber vehicles, and fully autonomous vehicles. ■

Acknowledgments

This work has been partially supported by the Spanish Ministry of Economy and Competitiveness through the projects "Agreement Technologies" (grant CSD2007-0022; CONSOLIDER-INGENIO 2010), "intelligent Human-Agent Societies" (grant TIN2012-36586-C03-02), and "Smart Delivery" (grant RTC-2014-1850-4).

References

1. V. Zeimpekis, G.M. Giaglis, and I. Minis, "Development and Evaluation of an Intelligent Fleet Management System for City Logistics," *Proc. 41st Ann. Hawaii Int'l Conf. System Sciences*, 2008; doi:10.1109/HICSS.2008.121.
2. W. Maden, R. Eglese, and D. Black, "Vehicle Routing and Scheduling with Time-Varying Data: A Case Study," *J. Operational Research Soc.*, vol. 61, no. 3, 2010, pp. 515–522.
3. D. Mardiasmo et al., "Asset Management and Governance: Analysing Vehicle Fleets in Asset-Intensive Organisations," *Proc. 1st Int'l Conf. Infrastructure Systems and Services: Building Networks for a Brighter Future* (INFRA), 2008; doi:10.1109/INFRA.2008.5439593.
4. V. Zeimpekis et al., "Dynamic Management of a Delayed Delivery Vehicle in a City Logistics Environment," *Dynamic Fleet Management*, Springer, 2007, pp. 197–217.
5. B. Eksioglu, A. Volkan Vural, and A. Reisman, "The Vehicle Routing Problem: A Taxonomic Review," *Computers & Industrial Eng.*, vol. 57, no. 4, pp. 2009, pp. 1472–1483.
6. V. Pillac et al., "A Review of Dynamic Vehicle Routing Problems," *European J. Operational Research*, vol. 225, no. 1, 2013, pp. 1–11.
7. S. Ossowski et al., "Decision Support for Traffic Management Based on Organisational and Communicative Multiagent Abstractions," *Transportation Research Part C*, vol. 13, no. 4, 2005, pp. 272–298.
8. M.-V. Belmonte et al., "Ontologies and Agents for a Bus Fleet Management System," *Expert Systems with Applications*, vol. 34, no. 2, 2008, pp. 1351–1365.
9. T.G. Crainic, M. Gendreau, and J.-Y. Potvin, "Intelligent Freight-Transportation Systems: Assessment and the Contribution of Operations Research," *Transportation Research Part C*, vol. 17, no. 6, 2009, pp. 541–557.
10. A. Goel, *Fleet Telematics: Real-Time Management and Planning Of Commercial Vehicle Operations*, Springer, 2007.
11. J. Andersen, T.G. Crainic, and M. Christiansen, "Service Network Design with Management and Coordination of Multiple Fleets," *European J. Operational Research*, vol. 193, no. 2, 2009, pp. 377–389.
12. J.F. Kurose and R. Simha, "A Microeconomic Approach to Optimal Resource Allocation in Distributed Computer Systems," *IEEE Trans. Computers*, vol. 38, no. 5, 1989, pp. 705–717.
13. J. Shi et al., "A Survey of Cyber-Physical Systems," *Proc. Int'l Conf. Wireless Communications and Signal Processing*, 2011; doi:10.1109/WCSP.2011.6096958.
14. D. Luckham, *Power of Events*, Addison-Wesley, 2002.
15. D. Bertsekas, "The Auction Algorithm: A Distributed Relaxation Method for the Assignment Problem," *Annals of Operations Research*, vol. 14, no. 1, 1988, pp. 105–123.
16. Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi Tessellations: Applications and Algorithms," *SIAM Rev.*, vol. 41, no. 4, 1999, pp. 637–676.
17. A. Awasthi et al., "Centralized Fleet Management System for Cybernetic Transportation," *Expert Systems with Applications*, vol. 38, no. 4, 2011, pp. 3710–3717.
18. S.L. Paterson and D.M. Brock, "The Development of Subsidiary-Management Research: Review and Theoretical Analysis," *Int'l Business Rev.*, vol. 11, no. 2, 2002, pp. 139–163.

Holger Billhardt is an associate professor in Computing Science and Artificial Intelligence at Universidad Rey Juan Carlos in Madrid. Contact him at holger.billhardt@urjc.es.

Alberto Fernández is an associate professor in Computing Science and Artificial Intelligence at Universidad Rey Juan Carlos in Madrid. Contact him at alberto.fernandez@urjc.es.

Lissette Lemus is a project manager at the Artificial Intelligence Research Institute, Spanish National Research Council (CSIC). Contact her at lissette@iiia.csci.es.

Marin Lujak is a post-doc researcher in the Artificial Intelligence Group at Universidad Rey Juan Carlos in Madrid. Contact him at marin.lujak@urjc.es.

Nardine Osman is a research fellow at the Artificial Intelligence Research Institute, Spanish National Research Council (CSIC). Contact her at nardine@iiia.csic.es.

Sascha Ossowski is a full professor of computer science and the Director ofe the Centre for Intelligent Information Technologies (CETINIA) at Universidad Rey Juan Carlos in Madrid. Contact him at sascha.ossowski@urjc.es.

Carles Sierra is a research professor at the Artificial Intelligence Research Institute, Spanish National Research Council (CSIC). Contact him at sierra@iiia.csic.es.