

Minimal Risk Manoeuvre Strategies for Cooperative and Collaborative Automated Vehicles

Victoria Vu
Semcon Sweden AB
Göteborg, Sweden
victoria.vu@semcon.com

Stig Ursing
Semcon Sweden AB
Göteborg, Sweden
stig.ursing@semcon.com

Carl Bergenhem
Qamcom Research and Technology AB
Göteborg, Sweden
0000-0002-6903-0756

Fredrik Warg
RISE Research Institutes of Sweden
Borås, Sweden
0000-0003-4069-6252

Fredrik Sunnerstam
Agreat AB
Göteborg, Sweden
fredrik.sunnerstam@agreat.com

Irina Cosmin
Agreat AB
Göteborg, Sweden
irina.cosmin@agreat.com

Anders Thorsén
RISE Research Institutes of Sweden
Borås, Sweden
0000-0001-7933-3729

Jimmy Holler
Epiroc Rock Drills AB
Örebro, Sweden
jimmy.holler@epiroc.com

Abstract—During the last decade, there has been significant increase in research focused on automated vehicles (AVs) and ensuring safe operation of these vehicles. However, challenges still remain, some involving the cooperation and collaboration of multiple AVs, including when and how to perform a minimal risk manoeuvre (MRM), leading to a minimal risk condition (MRC) when an AV within one of these systems is unable to complete its original goal. As most literature is focused on individual AVs, there is a need to adapt and extend the knowledge and techniques to these new contexts. Based on existing knowledge of individual AVs, this paper explores MRM strategies involving cooperative and collaborative AV systems with different capabilities. Specifically, collaborative systems have the potential to enact local MRCs, allowing continued productivity despite having one (or several) of its constituents encounter a fault. Definitions are given for local and global MRCs and what it means for MRMs, and illustrative examples are presented for each type of system.

Index Terms—Automated vehicles, automated driving systems, cooperative vehicles, collaborative vehicles, minimal risk manoeuvre, minimal risk condition, degraded operation, safety.

I. INTRODUCTION

It is widely believed that automated vehicles (AVs), or vehicles equipped with an automated driving system (ADS), have the potential to positively transform the transportation sector by enhancing both safety and efficiency. An ADS is a combination of hardware and software used to perform a vehicle's operational and tactical functions [1]. While much of the work done so far has been focused on AVs acting individually, the possibility of systems, or system of systems (SoS), involving multiple cooperative or collaborative vehicles, have

the potential to improve, e.g., efficiency, safety and energy consumption [2]. Smoother traffic flow and reduced energy requirements can be achieved by coordinating movements and exchanging relevant traffic updates, while sharing information about planned trajectories or detected hazards can improve safety. When it comes to use cases such as construction sites, harbors or mines, multiple vehicles of different types can complement each other to solve shared tasks.

Despite many forward strides, however, the deployment of AVs continues to present challenges, specifically with regards to the safe handling of complex operations within real-world scenarios. To better understand AV concepts, taxonomies and definitions have been developed [1], [3]. These help organize the different concepts and provide a common framework for analysis and discussion, including how to address safety assurance. However, much of the existing work focuses on individual road vehicles. To fully leverage the potential of a multi-AV system, it is imperative to adapt these concepts to the cooperative and collaborative framework.

To ensure safe operation, an AV must: (i) perform safely within the operating conditions it is designed to handle, i.e., its operational design domain (ODD) [1], (ii) prevent operation outside of its defined ODD, and (iii) safely handle performance critical failures, or situations where the vehicle risks exiting the ODD while the driving automation feature is activated. The term minimal risk condition (MRC) has been defined for the individual-AV context as a stable, stopped state to which the vehicle has been brought in order to reduce the risk of a crash [1], [3] in response to condition (iii). Nominal functionality involves the dynamic driving task (DDT), defined as all-real time operational and tactical functions needed to

This research has been supported by the Strategic vehicle research and innovation (FFI) programme in Sweden, via the project SALIENCE4CAV (ref. 2020-02946).

operate a vehicle, [1] performing as expected. The capability to transition between nominal functionality and the MRC is termed a minimal risk manoeuvre (MRM) [3] or DDT fallback [1]. In this paper we use the term MRM.

There has previously been a lack of consideration for how the concepts of MRM and MRC can be adapted to vehicles participating in cooperative or collaborative activities (in this paper, we refer to such a vehicle as a constituent). For instance, if a digger and truck are jointly working to move resources from location A to location B and the digger were to malfunction, what happens to the truck?

To reduce productivity losses, cooperative and collaborative systems would ideally continue to function to the extent possible, while remaining within the constraint of safety, even if one or a few of the constituents are incapable of continuing normal operations. Thus, there is a strong interest in adopting strategies that involve halting one or more constituents to establish a safe state while allowing for continued operation by other constituents, albeit possibly with reduced performance.

This paper explores the application of MRMs and MRCs in the context of cooperative and collaborative AVs. Although the terms cooperative and collaborative are often used interchangeably, this work considers the terms distinct, using the same classification as [4]; Cooperative AVs involve multiple AVs interacting for mutual benefit, with each vehicle still focused on achieving its individual strategic goal. On the other hand, collaborative vehicles, often centrally controlled by a traffic management system (TMS), work in harmony to complete a joint task termed the common strategic goal. We also assume that constituents of any SAE J3016 automation level, defined by the specific capabilities and limitations of the vehicles and its automation system [1], may belong to a cooperative or collaborative system as long as it has the capabilities to perform the role that it has been assigned.

State of the art literature on the subject of MRM and MRC, though mostly focused in the individual AV context, is reviewed; challenges and opportunities for MRM and MRC strategies for cooperative and collaborative AVs are identified; and the implications for each of the classes of cooperative and collaborative vehicles with different interaction characteristics (see Table I) are analysed. The analysis is facilitated by considering a number of use cases, also used throughout the paper to exemplify the principles.

The concepts of *global and local MRCs* are introduced to enhance the coordination and overall productivity of the system while maintaining safety. While global MRCs shut down an entire system when dependencies between constituents make productivity no longer possible or when safety is severely compromised, local MRCs affect merely one or a group of constituents, enabling systems to maintain a level of productivity not previously possible. We also define *concerted MRMs* as MRMs jointly performed by several AVs to reduce risk during these transitional manoeuvres.

This paper is organized as follows: Sec. II provides information on the Background and Related Work, Sec. III defines the concepts of MRMs and MRCs for cooperative

and collaborative AVs, and Sec. IV provides examples of cooperative and collaborative MRC classes. Finally, Sec. V summarizes the conclusions and outlines future work.

II. BACKGROUND & RELATED WORK

To facilitate the proposed taxonomy transfer from individual to cooperative and collaborative AVs, we first discuss the existing literature on individual vehicles in Section II-A, and then relevant existing work on cooperative and collaborative vehicles in Section II-B.

A. Individual Automated Vehicle Domain


With individual AVs, we refer to AVs that operate independently with individual strategic goals, using their own sensors and possibly connected infrastructure or cloud services, but without direct exchange with other AVs to guide operation.

1) *Definition of MRM and MRC*: Two standards - SAE J3016 [1] and ISO TR 4804 [3] - both provide definitions for the concepts of MRM and MRC. However, there are slight variations in these definitions, with the former requiring a vehicle to reach a full stop before considered to have achieved MRC, whereas the latter states that MRCs can include time-limited degraded modes and even allow for recovery to nominal operation. Transition between different MRCs can occur during this degraded operation. To denote a stable, stopped state with deactivated ADS, ISO 4804 introduces the term final MRC. SAE J3016 also uses the term DDT fallback to denote a response (either by a user or the ADS) to achieve MRC, whereas ISO TR 4804 refers to this action as an MRM.

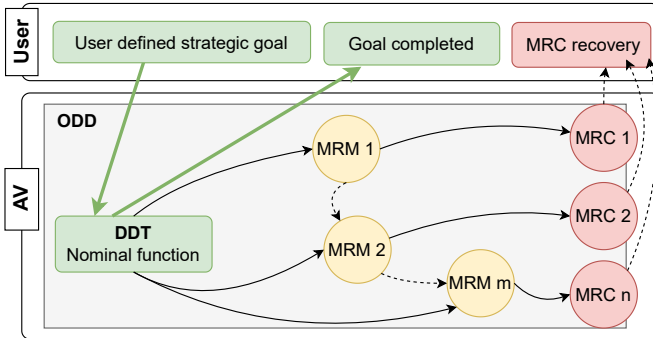
Gyllenhammar et al. [5] expand upon these definitions. First, an MRC is clearly described as a change of strategic goal (referred to as strategic mission by Gyllenhammar et al.) when the original goal cannot be fulfilled. Strategic goals, typically defined by the user, involves “trip planning, such as deciding whether, when and where to go, how to travel, best routes to take, etc” [1]. As illustrated in Fig. 1a, tactical (i.e., manoeuvring a vehicle in traffic during a trip such as changing lanes) and operational (i.e., instinctive split-second decisions such as braking) decisions are constrained by the strategic goal. In a situation with an unrecoverable ADS, performance-critical system failure or nearing ODD exit, the ADS must shift the strategic goal to achieving an MRC, and recovery from an MRC requires some form of user intervention. As long as the user defined strategic goal is pursued, even with limited capabilities, e.g., due to a faulty sensor, this is not an MRC. Thus, this definition also makes a clear distinction between MRM/MRC and degraded operation. In addition, in [5], the goal of an MRC is not just defined in terms of reducing the risk of a crash, but in line with safety standards in terms of acceptable risk: MRC is ‘a stable stopped condition at a position with an acceptable risk given the situation when the decision to enter MRC is taken. If an acceptable risk is not attainable, the position with the lowest risk should be selected.’ In order to argue that a certain MRC has an acceptable risk, one needs to account for (i) the risk of the selected position, (ii) the frequency with which this MRC might occur, and (iii)

the rate of resolving the MRC (e.g., if the vehicle is parked on the shoulder of the road, there is still some remaining risk which increases as the situation remains unresolved for longer periods of time). Finally, it is described that depending on the situation when a failure occurs, different MRCs will be required, e.g., for a road vehicle stopping at the nearest rest stop or immediately on the shoulder of the road, there could be two distinct MRCs. The rest stop gives a lower risk stopped position, but requires a more challenging MRM to achieve, which might not be possible depending on the type of failure. It may also be necessary to switch to a more easily achievable MRM and MRC during execution of an MRM, leading to a hierarchy of MRMs and MRCs as illustrated in Fig. 1b, e.g., if another failure occurs while driving towards the rest stop, it might be necessary to instead stop on the shoulder.

When referring to individual AV MRMs and MRCs in this paper, we assume the original definition from J3016 with the added detail from [5]. We do not use the taxonomy of MRC and final MRC from ISO TR 4804.



(a) Decision hierarchy for individual AV. Higher levels restrict what decisions are valid at a lower level (adapted from [5]).



(b) A set of MRCs will be available depending on the circumstances when nominal function can not be maintained (adapted from [5]).

Fig. 1: MRM and MRC for an individual AV.

2) *Fallback Strategies*: Several papers provide suggestions of MRMs for specific scenarios. Emzivat et al. [6] propose an MRM for AVs with SAE level 4+ automation experiencing perception failures driving on roads where a stop would be considered dangerous (e.g., tunnels, roads lacking shoulder). This is done by enacting speed profiles which favour low speed rear end collisions with vehicles in front of the ego vehicle. The potential benefits of cooperative vehicles are acknowledged, but communication is believed prone to security threats as well as time delays, hence focus is on individual fallback strategies.

Wie Xue et al. [7] investigate MRM strategies involving an individual AV experiencing front-facing sensor failures while traveling on a highway. The approach suggested brings the vehicle to a safe stop on the shoulder of the road or at a designated parking area.

Aside from papers examining specific scenarios, the concept of varied MRMs is explored by Yu and Luo, who introduce fallback strategies involving three degraded levels with seven different fallback strategies. The authors consider varying types of failures that a human user might encounter in conjunction with the behaviours one might take to negate these events. This information is used as a base to outline different functional failures that an ADS may encounter (e.g., actuator, decision, perception, ODD, power supply) and the fallback scenarios (e.g., reducing speed, parking on the side of the road or continuing to a safe location) it could use as response.

Svensson et al. [8] discuss the difficulties of ensuring that an AV can stop safely when encountering a failure. The authors refer to a "safe stop" as a type of MRM that is less hazardous, preferably leading to an MRC outside of an active lane. In contrast, the term "emergency stop" is used to denote stopping the vehicle "as fast as possible, at severe hazards." Both a "safe stop" and "emergency stop" are types of manoeuvres which lead to an MRC. The paper specifically considers MRM trajectory planning for AVs. An optimal control problem formulation for trajectory planning is proposed, which is then used to evaluate various ways of achieving a safe stop (as defined by the authors), with a second algorithm using pre-planned routes to help the car halt safely in real time.

Tong, Solmaz, and Horn [9] propose a run-time monitoring device and motion planning algorithm (termed MonDev) that is capable of supervising a vehicle's ADS and triggering an MRM when appropriate. Unlike previous solutions, the proposed framework does not need a redundant planner but instead incorporates safe stop planning into the existing framework.

The surveyed works outline various technical ways that MRMs could be executed. However, they focus solely on individual AVs and do not consider how these concepts might work when multiple vehicles are operating in tandem.

B. Cooperative and Collaborative Vehicle Domain

Using the classification in [4], which in turn extends the taxonomy for cooperative systems from SAE J3216 [10], in cooperative systems, multiple AVs work together for mutually beneficial reasons, with each vehicle still focused on achieving its own individual goal. For example, cooperation among several AVs at an intersection could enhance traffic flow, reducing congestion to help each vehicle reach its intended destination with greater ease. In contrast, collaborative AVs work in tandem to achieve a common strategic goal. An example is a digger and truck working collaboratively to load and transport materials from point A to point B. Table I provides details and explanations of the various cooperative and collaborative classes. A common solution for collaborative AVs, especially

for non-road use cases such as mines, harbors, or construction sites, is to use a traffic management system (TMS) to control overall operations (i.e., the orchestrated class in Table I). The TMS acts as a mediator between vehicles by providing instructions when needed, while also controlling the system borders as outlined in [11]. Using a TMS in the digger/truck example, the vehicles could be reliant on the TMS to inform each vehicle of its respective status and responsibilities. If the digger were to experience a failure, the TMS would then communicate with the truck, providing new instructions, e.g., commanding it to assist another digger or park at a designated location. In other contexts, the terms traffic management center (TMC) or fleet management system are used rather than TMS, but they provide similar functionality.

While this paper explores the implementation of MRMs and MRCs for vehicles within both the on road and machinery domain, the concepts themselves do not currently exist within the latter. Instead, machines are regulated by the Machinery Directive (MD) [12], with the concept most parallel to MRMs and MRCs being that of the emergency stop, which requires an immediate halt of the vehicle experiencing a failure. The examples outlined in this paper may not strictly be in line with the current MD. Proposals to change the MD are out of scope for this paper, but could be investigated in future studies.

1) *Definition of MRM and MRC*: SAE J3216 [10] details different cooperation classes related to AVs (see cooperation classes in Table I), though there are no expectations outlined for MRMs or MRCs. To the best of our knowledge, there is still no standard or otherwise well established definition of MRM and MRC for cooperative or collaborative vehicles.

2) *Improved Safety and Efficiency*: A considerable amount of literature can be found on the potential to improve road safety and efficiency with cooperative or collaborative AVs. When it comes to safety, the most cited advantages focus on prevention (collision avoidance and traffic accident prevention through control and coordination of DDTs) in the context of a few well studied scenarios like platooning, intersection crossing, lane changing and merging.

Coll-Perales et al. [13] address the problem of multiple vehicles engaging in MRM at the same time within the same area due to a failure of transfer of control (ToC). The authors investigate the impact on traffic safety and efficiency of MRM execution when triggered by decentralized environmental notification message (DENM) based or manoeuvre coordination message (MCM) based ToC management schemes. There is a focus on vehicle-to-everything (V2X) communication to assist Connected and Automated Vehicles in executing a ToC and MRM, with the infrastructure initiating a spatial distribution of ToCs and informing vehicles of locations where they could execute an MRM into an MRC if the user fails to take over.

Malik, Khan and El-Sayed [2] provide a survey related to cooperative/collaborative driving. The authors propose a taxonomy and review present approaches, such as vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication, while also examining the use case of platooning, specifically with regards to electing a leader for the group

of vehicles during various situations. It should be noted that the authors do not follow the distinct definition of cooperative and collaborative interaction used in this paper.

Wang et al. [14] discuss using cooperative vehicle fleet technology to improve safety and efficiency near freeway bottlenecks. The focus is on reducing rear-end collision risks near recurrent bottlenecks using various resolution strategies based on vehicle types, time requirements and intentions. The paper proposes control strategies to establish rules and the appropriate vehicle disbandment when a vehicle formation must be changed.

Chintakunta and Akbaş [15] present an approach for cooperative connected and autonomous vehicles to support emergency vehicles in traffic. The authors formulate the trajectory planning problem for an emergency vehicle and map it onto a graph, proposing an optimization problem to accomplish cooperative planning of the surrounding vehicles.

Härfner et al. [16] evaluate the performance of two specific cooperative communication protocols for manoeuvre communication. The authors assess failure risks and mitigation strategies in conjunction with communication. Proposals such as sending updates of planned trajectory or extending an update to indicate an emergency are suggested.

III. DEFINING MRMs AND MRCs FOR COOPERATIVE AND COLLABORATIVE AVS

The definition by Gyllenhammar et al. [5] discussed in Section II-A1 cleanly defines MRC as distinct from nominal operation by defining it as a change in strategic goal. As the definition applies to individual AVs, it follows that if an AV is unable to continue pursuing its strategic goal, the goal will not be achieved. However, with multiple interacting AVs, goal achievement becomes less straightforward. Depending on the goals and current situation, if one AV must go to MRC, the others may or may not be able to continue. For reasons of productivity, it is also typically desirable to keep as much as possible of the operation running, even if there is some productivity loss. In this context, how do we define MRM and MRC, and what is the difference between MRC and degraded operation? In Section III-A we propose definitions for MRM and MRC in the cooperative and collaborative context, in Section III-B we discuss how to clearly distinguish performance degradation from MRC, and in Section III-C we revisit the concept of MRC hierarchies.

A. MRMs, Global MRCs, and Local MRCs

Consider a system consisting of a digger and a truck working collaboratively to move material from location A to location B. If a failure were to occur, forcing one of the vehicles to stop, the goal could not be completed. However, if the system consisted of two pairs of diggers and trucks, and only one digger needed to go to MRC, its partnered truck could continue functioning alongside the remaining digger. The remaining digger would instead be paired with two trucks, allowing for continued, although somewhat reduced, productivity. To be able to describe both of these situations, we

define two types of MRCs for cooperative and collaborative use cases (still assuming the basic definition of MRC from J3016 [1]):

Definition 1. *If a failure or near ODD exit condition makes it necessary to bring all constituents to MRC, completely abandoning the shared strategic goal/all individual strategic goals, this is a **global MRC**. User intervention is needed to recover stopped AVs and set a new strategic goal.*

Definition 2. *If a failure or near ODD exit condition makes it necessary to bring one or more constituents to MRC, but either the full strategic goal, or part of the strategic goal, can be continued by the remaining constituents, this is a **local MRC**. User intervention is required to recover stopped AVs.*

An MRM is the transitional manoeuvre to get to an MRC. In the example above, if one of the trucks experiences a sensor failure, risk could be reduced during this transition by informing the other AVs of the problem and jointly planning manoeuvres that aid the faulty AV in achieving the best possible MRC, avoiding accidents or obstructions during the MRM:

Definition 3. *An MRM is the response by a specific AV to achieve an MRC. **Concerted MRMs**, jointly performed by several AVs to reduce the risk during the transitional manoeuvres, are possible for some cooperative and collaborative systems. Concerted MRMs must result in MRC for at least one constituent involved.*

A local MRC is defined as affecting one or a group of constituents. Fig. 2 illustrates that this might be thought of as a hierarchy. The levels not only have the potential to affect productivity, but also influence the safety case (a structured argument, supported by evidence, that provides a clear and traceable account of how an AV system meets its safety requirements [17]). On the lowest level, MRCs may only affect an individual AV, but there may be intermediate levels where groups of AVs must jointly achieve an MRC as a consequence of a failure or near ODD exit. One might expect that fewer (in the extreme case only the global) MRC alternatives will generally result in a simpler safety case, but lower average productivity. Allowing more fine-grained alternatives reduces productivity impact, but increases the number of MRC strategies that needs to be proven safe. There may be exceptions, however, where shutting down larger groups of vehicles involves a higher risk than stopping one or a few.

Determining which AVs should be affected by a local MRC may stem from practical concerns, e.g., stopping all activity in a certain area may allow for a simpler safety case than having operational AVs that must consider the presence of stopped vehicles sharing the same space. It might also follow from dependencies between AVs. Similarly to any dependent failure analysis (see e.g., ISO 26262 [18]) there might be either *cascading failures*, where a failure in one AV prevents another from continuing, such as the single digger/truck example




Fig. 2: Global and local MRCs.

above, or it might be *common cause failures*, where the same root cause forces a group of AVs into MRC, e.g., several forklifts moving containers in a harbor and heavy rain incapacitates their perception (heavy rain is outside the ODD of these AVs), causing all to simultaneously go to MRC.

B. MRC, performance degradation, and tactical adaptations

The introduction of global and local MRCs also brings into question the concept of degradation. Defining the term for individual AVs is relatively straightforward, where degradation refers to a failure that does not prevent the AV from continuing to operate safely and pursue its strategic goal, though perhaps operation is further restricted to a more confined ODD [5] and/or lower performance.

For degradation of a cooperative or collaborative system however, two perspectives may be considered, that of the constituent and the system as a whole. Consider four cases: (i) The single digger/truck example again, and posit that the long-range radar on the truck malfunctions. The truck can continue to work, but is restricted to a lower speed since it has to rely on other sensors with a shorter detection range. The performance is permanently (i.e., repair is needed, the system cannot recover by itself) degraded. (ii) The perception range on the truck is temporarily reduced due to heavy rain, causing the truck to slow down, but it can resume its normal speed once the rain ceases. (iii) The double digger/truck example, where one digger breaks down and stops, but the remaining vehicles adapts such that the remaining digger works in tandem with two trucks. (iv) A platoon of five vehicles collaborating to achieve a common strategic goal of transporting goods from one location to another via public roads. The platoon leader has an extended responsibility for perception in relation to the trailing trucks due to its location in the platoon and superior field of view. A faulty forward-looking sensor makes it unable to handle the platoon leader task, but it can be a follower; as it does not require the faulty sensor. The platoon adapts by selecting a new leader and can continue the mission with the same speed and capacity.

In (i)-(iii) there is a performance degradation in the system as a whole compared to ideal circumstances. In (i) and (ii) there is also a performance degradation from the affected constituent perspective, but in (iii) there is instead an MRC from the constituent viewpoint. In (iv) there is no degradation at all from the SoS perspective, but a permanent performance degradation from the viewpoint of the affected constituent.

With the same reasoning as in [5], both (i) and (ii) are situations where the AVs tactical decision-making should be able to cope with capability degradation. Whether it is due to changes in the environment or due to a malfunction is not the critical question, but rather whether or not it is a capability change that can be diagnosed and handled by the tactical decisions without abandoning the user defined strategic goal. If not, it is an ADS performance-critical failure and the only option is to go to an MRC. However, there is one difference in the fact that (i) involves a permanent failure while (ii) is a condition that does not require user intervention to resolve. To distinguish the need for intervention to recover to nominal performance we define:

Definition 4. *If the operational capability of one or more constituents is permanently reduced due to a failure, but the degradation can be handled by adapting tactical decisions such that all constituents are still operational and the strategic goal is not abandoned, the cooperative or collaborative system has a **permanent performance degradation**.*

(iii) and (iv) are only applicable for some of the classes of cooperative and collaborative vehicles. (iii) is what we in Section III-A defined as a *local MRC* rather than a performance degradation. (iv) is an interesting case, where the faulty vehicle can operate as a follower in a platoon, but can, due to the faulty sensor, no longer operate on its own or in the role of platoon leader. As a constituent in the platoon it has a permanent performance degradation which, if attempting to operate without a lead vehicle, may force it to an MRC.

C. Hierarchy of MRCs

The concept of hierarchies of MRCs from [5], illustrated in Fig. 1b, can also be applied to cooperative and collaborative vehicles. For individual AVs, the main concern is selecting the appropriate MRC depending on the remaining capabilities. For the multi-AV case, additional concerns include which of the constituents need to reach MRC, and whether some sufficiently safe option is available where the impact on productivity can be limited. Consider a system of collaborating machines engaged in heavy lifting in a harbour. An automated crane removes cargo containers from a ship and several automated forklifts move the containers from the unloading zone, stacking them in a specific storage area. However, the sudden appearance of rain in combination with decreasing temperatures increases the risk of accidents due to slipping. Thus, the system aborts the common strategic goal of unloading, moving and stacking cargo, to instead initiate MRM1 into MRC1. This action results in a local MRC, where the crane halts its task, while the forklifts continue stacking containers already unloaded, before parking at a designated area. During MRM1 however, one of the forklifts indicates slipping and there is a new system decision to perform MRM2 into MRC2, resulting in a global MRC, with all forklifts stopping immediately and cargo being placed down on the ground. A local MRC is preferred for productivity reasons, as less human intervention is involved and the system still performs a portion of its task,

albeit at a reduced rate of productivity. On the other hand, the global MRC completely halts all productivity but provides the highest level of safety. Note that in this particular example there is only one local MRC (MRC1), with MRC2 being the global MRC. However, in larger systems consisting of more constituents, additional levels of MRC could exist (MRC1, MRC2, MRC3, etc.) leading to the shutdown of various parts of the system before enacting the global MRC.

IV. MRCs FOR COOPERATIVE AND COLLABORATIVE CLASSES

Cooperative type AVs consist of four distinct classes: status-sharing, intent-sharing, agreement-seeking and prescriptive, whereas collaborative types consist of three: coordinated, choreographed and orchestrated. Table I provides an explanation of these classes, alongside the defining characteristics of their MRMs/MRCs considering our definitions from Section III. In the subsections below, the MRMs/MRCs of each class are further explored through example scenarios.

A. Cooperative Class Examples

Status-sharing AVs share status information that could be used by other vehicles to make informed decisions, benefiting their own individual objectives. For instance, while working in a narrow mine, a truck reaches MRC, transmitting its stopped position to other machines in the vicinity. The receiving machines are then able to adjust their routes, avoiding the tunnel while still accomplishing their designated tasks. This scenario illustrates how a local MRC for status-sharing vehicles could function. Global MRCs however, do not exist in this context. Since each vehicle is only able to relay its current status, all vehicles must be responsible for making their own decisions.

Intent-sharing AVs are similar to the aforementioned status-sharing AVs, with the addition of informing others of its planned actions [10]. This extra information on what the ego vehicle intends to do allows other vehicles to act appropriately. Consider a car experiencing a perception failure, it broadcasts its intention of reaching MRC 500 metres ahead on the shoulder of the freeway. The surrounding vehicles are now aware of its plans and have the opportunity to adjust their trajectories accordingly, allowing the ego vehicle to smoothly enter MRC at its target location. Once again however, global MRCs are not possible as intent-sharing vehicles cannot force others into action or inaction, merely providing information and allowing others to behave as they see fit.

Agreement-seeking AVs enable vehicles to consent to manoeuvres together. Consider the example above, of a vehicle that wishes to reach MRC on the road shoulder, 500 metres ahead. With agreement-seeking, rather than only broadcasting its plans as intent-sharing does, the ego vehicle directly requests help, waiting for a response before enacting its manoeuvre [10]. If a positive response is received, a concerted MRM leads to an MRC. Otherwise, alternative plans must be considered, and confirmation sought once more. Local MRCs within this class involve an individual AV requesting help to achieve MRC, but global MRCs are also possible. Consider a

TABLE I: Taxonomy of cooperative and collaborative AVs (from [4]) with corresponding MRM/MRC characteristics.

Type	Class	Explanation	MRM/MRC
Several AVs with <i>individual</i> strategic goals.	<i>Status-sharing</i>	AVs share information, e.g. position, sensor data, or world model, to help other actors make their decisions (J3216 Class A).	Information of e.g., MRC status and position can be shared to improve safety (compare to red warning triangle used for road vehicles) (only individual MRC).
	<i>Intent-sharing</i>	AVs share intent for future actions (operational, tactical, strategic) to help other actors make their decisions (J3216 Class B).	Intent to go to MRC including e.g., location for stop and planned MRM trajectory can be shared to allow others to better adapt (only individual MRC).
	<i>Agreement-seeking</i>	AVs communicate to reach (voluntary) agreements with other actors in order to optimize some parameter(s) for mutual benefit (J3216 Class C).	Vehicle(s) going to MRC can seek agreement with adjacent agents to aid in safety reaching an MRC (local MRC). Multiple vehicles can coordinate actions to make sure all safely reach an MRC (global MRC). MRMs can be concerted.
	<i>Prescriptive</i>	AVs generally act individually, but accepts certain temporary prescriptive actions to achieve a tactical goal defined by another actor, e.g. the road operator (J3216 Class D).	Directing entity can instruct one or a few vehicles to go to MRC (local MRC), or instruct all vehicles in a specified area to go to MRC (global MRC). Directing entity may also instruct other actions that individual vehicles are unable to comply with, forcing them to instead go to MRC. MRMs can be concerted.
Several AVs with <i>common</i> strategic goal.	<i>Coordinated</i>	Constituent AVs communicate to reach agreements for how to jointly act in order to achieve a common strategic goal.	Vehicle(s) going to MRC may seek agreement with adjacent agents to aid in safely reaching an MRC (local MRC); remaining vehicles may coordinate to continue without missing constituent(s). Multiple vehicles can coordinate actions to make sure all safely reach an MRC (global MRC). MRMs can be concerted.
	<i>Choreographed</i>	Constituent AVs act individually but are designed to follow a common global scenario/goal, i.e., unlike coordinated vehicles, they do not rely on communication to perform the collaborative task.	Constituents designed to behave in a certain way for their MRM and MRC, or respond in a certain way to the MRM and MRC of others in line with the common strategic goal (local and global MRC).
	<i>Orchestrated</i>	Constituent AVs are directed by a single entity acting to achieve the strategic goal. The directing entity can be one of the constituent vehicles, or a separate system referred to as e.g., traffic management system (TMS).	Individual AV might unilaterally go to MRC, e.g., for system failures or lost communication with directing entity. Control of other constituents by directing entity to enact either a local or global MRC. Directing entity can direct remaining constituents to be able to continue without missing constituent(s) (local MRC), or safety stop all constituents (global MRC). MRMs can be concerted.

mine experiencing an unfortunate fire. All machines within the system must be evacuated and parked outside in a safe zone. To do this, the machines must communicate and agree on who goes first, in which order, and where exactly each individual should achieve its safe stopped position.

The last cooperative AV class, **prescriptive**, involves individual vehicles that must do as directed [10]. Consider two machines, one large, the other small, working within a mine tunnel. As the small machine works, its larger counterpart requires passage through the narrow tunnel. Both are unable to fit, thus, the larger machine communicates with the smaller, directing it to engage in a local MRC within the nearest pocket. Global MRCs are also possible in the prescriptive class of AVs. When parts of a road have been washed away due to flooding, road authorities can ensure safety by forcing all vehicles entering the area to achieve MRC with the use of a TMS. Note that while agreement-seeking requires all parties to agree to the requested action, prescriptive allows an entity to force constituents into MRC.

B. Collaborative Class Examples

The **coordinated** class of collaborative AVs is similar to the agreement-seeking class of cooperative AVs, but coordinated AVs share a common strategic goal. [4]. A system with multiple trucks and a single digger operates as follows: the digger deposits material into truck A until it is full and signals the digger to stop. Truck A travels to a designated area, emptying the material there, while truck B takes its

place to be filled by the digger. During its travels, truck A experiences an MRC, forcing the remaining trucks to agree on a new route to avoid potentially colliding with truck A. Coordinated AVs are able to communicate with each other when a failure occurs within the system, consenting to an alternative safe solution. While local MRCs in the coordinated class enable continued productivity by shutting down certain parts of the system, global MRCs may also be necessary. Consider a system where humans work alongside AVs and the AVs must have continuous communication and access to the individual's location. If all constituents within the system lose sight of the person, every constituent must shut down to ensure maximum safety. Dependency is also a factor that should be considered. Continued productivity is influenced by how dependent the system is on the failing constituent. In the aforementioned example of a digger and trucks, the constituent entering MRC was merely one of many trucks, thus enabling the system to continue functioning with a local MRC. If instead, the lone digger had experienced a failure, the remaining constituents would be rendered useless while posing a potential safety hazard by continuing to function unnecessarily. Thus, the digger could communicate with the trucks, requesting they navigate to a designated parking area, where further communication would be required between the trucks to negotiate how and where each constituent should safely achieve the global MRC.

Whereas all other classes involve communication of some

kind – whether that be one-way or two-way, requiring a response or not requiring a response – **choreographed** AVs involve no communication to perform their collaborative task [4]. Each constituent within the system has an assigned task, with the system itself designed to handle specific situations without communicating. Returning to the previously mentioned digger and trucks, in a choreographed setting, these machines do not communicate to complete a shared task or to achieve MRC. Instead, the trucks may have a weight limit, once that weight limit is reached, they depart for their new destination. The digger, rather than waiting for a signal from the truck, could have a designated number of scoops it must complete before halting. In such a system, if one vehicle stops, the remaining may just observe. However, the system must also be designed to handle certain situations. Perhaps the constituents are designed to recognize that, if a truck does not check into the deposit destination within a certain time period, a failure is assumed and all trucks must automatically take an alternate, predetermined route. On the other hand, a global MRC could involve all constituents immediately halting when a truck fails to arrive at the deposit destination. In such a choreographed setting, both local and global MRCs need to have been previously designed into the system in order to be properly executed.

Lastly, the **orchestrated** class involves multiple constituents directed by a single entity (e.g., one of the other constituents or a separate system such as a TMS), to achieve a shared strategic goal [4]. In the previous scenarios, we have discussed how MRCs may be accomplished for multiple trucks working with one digger. To exemplify MRCs in this situation under the orchestrated class, a TMS could be in charge of communication between all machines. Now, if truck A were to reach MRC on its way towards the deposit destination, the TMS would be informed. With this information, the TMS could then send out a new route command to all functioning trucks, ensuring productivity continues while maintaining safety. The global MRC could be triggered in a similar way. If the sole digger experiences a failure and reaches MRC, the TMS would be informed, communicating with all trucks in the system and requesting they either halt immediately or engage in a concerted MRM, navigating to a designated safe area before reaching MRC.

V. CONCLUSIONS AND FUTURE WORK

This paper explores the concepts of MRMs and MRCs for systems of cooperative and collaborative AVs. The terms MRM, MRC and permanent performance degradation are defined for the cooperative and collaborative domain. The hierarchy of MRCs, previously explored by Gyllenhammar et al. [5] is adapted to cooperative and collaborative AVs. Local and global MRCs are introduced, with examples for the varying classes of cooperative and collaborative AVs provided.

For future work, we suggest running simulations on the concepts presented in this paper, adapting MRMs and MRCs to the machine domain and exploring whether a recovery from MRC can be safely handled without human intervention.

REFERENCES

- [1] SAE, *J3016 Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. SAE, Apr. 2021. [Online]. Available: https://doi.org/10.4271/J3016_202104
- [2] S. Malik, M. A. Khan, and H. El-Sayed, “Collaborative Autonomous Driving—A Survey of Solution Approaches and Future Challenges,” *Sensors*, vol. 21, no. 11, p. 3783, 2021.
- [3] ISO, *ISO*, 2020. [Online]. Available: <https://www.iso.org/standard/80363.html>
- [4] F. Warg, A. Thorsén, V. Vu, and C. Bergenheim, “A unified taxonomy for automated vehicles: Individual, cooperative, collaborative, on-road, and off-road,” 2023.
- [5] M. Gyllenhammar, M. Brännström, R. Johansson, F. Sandblom, S. Ursing, and F. Warg, “Minimal risk condition for safety assurance of automated driving systems,” in *CARS: 6th International Workshop on Critical Automotive Applications: Robustness & Safety*, 2021.
- [6] Y. Emzivat, J. Ibanez-Guzman, P. Martinet, and O. H. Roux, “Dynamic driving task fallback for an automated driving system whose ability to monitor the driving environment has been compromised,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, 2017, pp. 1841–1847.
- [7] W. Xue, B. Yang, T. Kaizuka, and K. Nakano, “A fallback approach for an automated vehicle encountering sensor failure in monitoring environment,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 1807–1812.
- [8] L. Svensson, L. Masson, N. Mohan, E. Ward, A. P. Brenden, L. Feng, and M. Törngren, “Safe stop trajectory planning for highly automated vehicles: An optimal control problem formulation,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 517–522.
- [9] K. Tong, S. Solmaz, and M. Horn, “A search-based motion planner utilizing a monitoring functionality for initiating minimal risk maneuvers,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 4048–4055.
- [10] SAE, *SAE J3216 Taxonomy and Definitions for Terms Related to Cooperative Driving Automation for On-Road Motor Vehicles*. SAE, Jul. 2021. [Online]. Available: https://doi.org/10.4271/J3216_202107
- [11] C. Bergenheim, M. Majdandzic, and S. Ursing, “Concepts and risk analysis for a cooperative and automated highway platooning system,” in *Dependable Computing - EDCC 2020 Workshops*. Cham: Springer International Publishing, 2020, pp. 200–213.
- [12] European Commission, “Directive 2006/42/EC of the European Parliament and of the Council of 17 May 2006 on machinery, and amending Directive 95/16/EC (recast) (Text with EEA relevance)” [Online]. Available: <http://data.europa.eu/eli/dir/2006/42/oj/eng>
- [13] B. Coll-Perales, J. Schulte-Tiggas, M. Rondinone, J. Gozalvez, M. Reke, D. Matheis, and T. Walter, “Prototyping and evaluation of infrastructure-assisted transition of control for cooperative automated vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6720–6736, 2022.
- [14] L. Wang, M. Yang, Y. Li, B. Wang, and J. Zhang, “Resolution strategies for cooperative vehicle fleets for reducing rear-end collision risks near recurrent freeway bottlenecks,” *Journal of Intelligent Transportation Systems*, 2022.
- [15] H. Chintakunta and M. Akbaş, “Spectrum analytic approach for cooperative navigation of connected and autonomous vehicles,” in *Proceedings of the 9th ACM Symposium on Design and Analysis of Intelligent Vehicular Networks and Applications*, 11 2019, pp. 97–104.
- [16] B. Häfner, J. Jiru, H. Schepker, G. Schmitt, and J. Ott, “Preventing failures of cooperative maneuvers among connected and automated vehicles,” *Computer Communications*, vol. 193, 07 2022.
- [17] UL, *ANSI/UL 4600:2019 Standard for Safety for the Evaluation of Autonomous Products (Draft)*. UL, 2019. [Online]. Available: <https://users.ece.cmu.edu/~koopman/ul4600/index.html>
- [18] ISO, *ISO 26262:2018 - Road vehicles – Functional safety*. International Organization for Standardization. [Online]. Available: <https://www.iso.org/standard/43464.html>