Connected Places Catapult

Safety Plan

Testing in a Car Park -
Autonomous Valet Parking

November 2019

CATAPULT
Connected Places

Executive Summary

This document describes the plan for activities to be completed to ensure the safety of the Autonomous Valet Parking (AVP) project. The aims are:

- Ensure an acceptable level of safety is achieved for those involved in the project;
- Ensure an acceptable level of safety for members of the public in the vicinity of the AVP vehicle;
- Provide an example of a safety case for a vehicle operating at higher levels of automation such that learning from this project can be used to inform the safety cases for future projects.

The report examines what evidence will be required to show that acceptable safety has been achieved for any given testing activity, and how that evidence will be collected, such that when all the evidence is taken together, there is a convincing argument that the project as a whole is safe.

The safety of the project is divided into two separate areas:

- **Operational Safety** considers how safe procedures will be used to deploy the system in its surrounding environment. It includes:

  - Method Statements for each type of test to describe how they should be undertaken to ensure safety, this includes a Risk Assessment;
  - Ability of the Safety Driver to intervene, bearing in mind driver training/ competence, driver awareness of the system limits, and the safety of the Human-Machine Interface (HMI);

- **System Safety** considers the Functional Safety and 'Safety of the Intended Function" (SOTIF) of the vehicle itself, particularly focussed upon the automated driving system but also considering safety of the base vehicle. It includes:

  - Verification that the vehicle meets a comprehensive set of safety requirements when subjected to specific Test Cases;
  - Validation to ascertain whether the vehicle performs safely during extended testing mileages;
  - Configuration Control procedures to ensure that updates are introduced safely.
  - The 'verification' aspect includes a significant amount of work to generate the requirements prior to carrying out the Test Cases to provide evidence that the vehicle is safe on the basis that it meets those requirements. This includes a functional safety analysis of the system and includes a review of relevant legislation and standards.

The 'validation' aspect will include a reporting spreadsheet to log any incidents that occur during testing. Testing will start in highly controlled environments, building up progressively to more challenging scenarios as assurance is gained.

Table of Contents

List of Figures

Notice

By using this safety report ("the Report") produced by the Connected Places Catapult ("CPC") you accept this disclaimer in full. The Report has been prepared in good faith on the basis of information, findings and analysis of our specific research activity entitled "Autonomous Valet Parking". All information contained in the Report is provided "as is" and CPC does not guarantee or warrant the accuracy, reliability or completeness of the information in the Report or its usefulness in achieving any particular outcome or purpose. CPC does not owe a duty of care to any third-party readers.

You are responsible for assessing the relevance and accuracy of the content of this publication. You must not rely on the Report as an alternative to seeking appropriate advice. and nothing in the Report shall to any extent substitute for consultation with an appropriately qualified advisor. You must obtain professional or specialist advice before taking, or refraining from, any action on the basis of the content of the Report.

To the fullest extent permitted by law, CPC excludes all conditions, warranties, representations or other terms which may apply to the Report or any content in it, whether express or implied. CPC will not be liable to any user for any loss or damage, whether in contract, tort (including negligence), breach of statutory duty, or otherwise, including without limitation loss of or damage to profits, sale business, revenue, use, production, anticipated savings, business opportunity, goodwill, reputation or any indirect or consequential loss or damage. Nothing in the Report excludes or limits CPC's for any liability that cannot be excluded or limited by English law.

Any entity seeking to conduct autonomous vehicle trials will need to develop and publish a safety case specific to their own trials (as specified by the government's Centre for Connected & Autonomous Vehicles (CCAV) Code of Practice for Automated Vehicle Trialling) and gain permission to do so.

FUNDING:

The Autonomous Valet Parking project is part-funded by the Centre for Connected and Autonomous Vehicles (CCAV), delivered in partnership with Innovate UK. It is part of the government's £100 million Intelligent Mobility Fund, supporting the Future of Mobility Grand Challenge.

As a key part of the UK government's modern Industrial Strategy, the Future of Mobility Grand Challenge was announced in 2017 to encourage and support extraordinary innovation in UK engineering and technology, making the UK a world leader within the transport industries.

This includes facilitating profound changes in transport technologies and business models, to make the movement of people, goods and services across the nation greener, safer, easier and more reliable.

AUTHORISATION:

ACTION	SIGNATURE BLOCK	NAME AND POSITION WITHIN CPC
Written by:	Maysun Hassanaly	Maysun Hassanaly Systems Engineer
Reviewed by:	Lovedeep Brar	Lovedeep Brar Senior Systems Engineer
Authorised by:	Stuart Rowell	Stuart Rowell Principal Systems Engineer

RECORD OF CHANGES:

REALEASED TO	VERSION	REASON FOR CHANGE	DATE
Parkopedia	0.1	First Draft	01/03/19
Parkopedia	0.2	Second Draft	April 2019
Parkopedia	1.0	First Release	07/06/19
Parkopedia	1.1	Second Release	09/09/2019
Parkopedia	1.2	Review for testing in car parks	27/10/19
Parkopedia	2.0	Third Release	18/11/19

1. Introduction

Project Description

The project is led by Parkopedia, with the University of Surrey and the Connected Places Catapult (CPC) as project partners. Its objective is to deliver a proof of concept involving an autonomous vehicle (AV) that will fulfil the valet parking function. The AV will navigate to a free parking space utilising indoor parking maps, autonomously execute the parking manoeuvre and respond to a summons request by navigating the vehicle back to the driver.

The AVP system will be developed and demonstrated on a drive-by-wire enabled vehicle provided by StreetDrone.

The AVP project aims to achieve this goal by:

1. Developing automotive-grade maps required for autonomous vehicles to localise and navigate within a multi-storey car park.
2. Developing localisation algorithms - targeting a minimal sensor set of cameras, ultrasonic sensors and inertial measurement units - that make best use of these maps.
3. Develop the safety case
4. Prepare for AVP car park trials
5. Engage with stakeholders to evaluate perceptions around AVP technology
6. Demonstrating this technology in a variety of car parks in the UK.

Work Packages

Within the AVP project, the CPC is responsible for developing a safety case and Systems Engineering deliverables to support the project until the demonstration phase.

Safety Deliverables

The Systems Engineering and Safety deliverables are listed below:

- **Concept of Operations (ConOps):** Provides the intended operation of the system, and a high-level system description. This is a live document and will be continuously updated until the end of the system development and testing.
- **Requirements:** Lists the system requirements. All the requirements will be tested, validated and verified. This is a live document and will be continuously updated until the end of the system development and testing.
- **Risk Assessment and Method Statement (RAMS)**: Covers Operational Safety and ensures that the trials are carried out safely. Each project milestone will follow a separate RAMS.

- **Failure Mode and Effect Analysis (FMEA):** Highest level analysis of the entire system and subsystems based on the Functional Architecture**.** The FMEA considers single point failures and focuses on system-related deficiencies. The FMEA will produce safety goals, which will form the basis of Safety Requirements for the system. This is a live document and will be continuously updated until the end of the system development and testing.
- **Hazard Risk Assessment (HARA):** Hazard Analysis and Risk Assessment from the perspective of System Safety**.** The possible vehicle level faults identified in the FMEA are used to generate hazards here. Three separate risk assessments are provided for the same hazards, depending on whether the vehicle is deployed on a private track, in a secure area of a car park, or as part of a demonstration involving travelling through non-secure areas of the car park.
- **Safety Plan** (this document): considers both 'System Safety' (i.e. is the system designed to behave in a safe way and be robust against hazardous failures) and 'Operational Safety' (i.e. the external controls applied to ensure safety during trials)
- **Safety Case Summary:** summarises the evidence collected prior to commencing Autonomous Valet Parking testing in accordance with the Safety Plan. There will be a Safety Case for testing in a controlled environment, testing in car parks and the demonstration.

Schedule for Safety Deliverables

Figure 1 below illustrates the safety schedule of the AVP project. The project has been divided in four major milestones:

1. Collecting data in car parks
2. Testing in a controlled environment
3. Testing in car parks
4. Demonstration

The AVP project commenced in May 2018 and will finish in October 2020.

These deliverables are distributed as below:

- Square boxes indicate a one-off workshop or document
- Long arrows indicate live documents

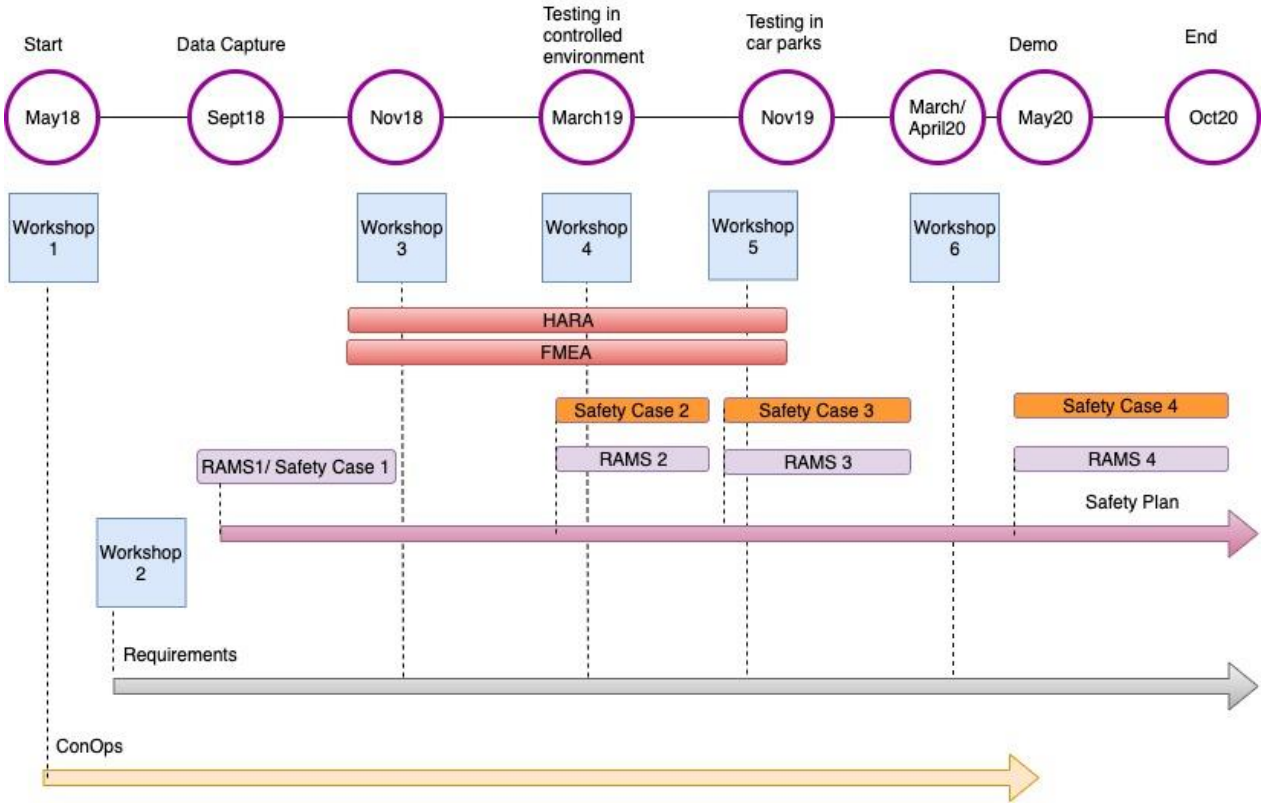

Figure 1 AVP Safety Deliverables Schedule

Scope of Safety Work

This section lays out the scope for what the safety work package will deliver, as opposed to the scope for what the Autonomous Driving (AD) System will be capable of; the latter is described by the 'Concept of Operations Document'. An 'Is/ Is Not' summary of the safety work is set out below:

Is

- Operational Safety (i.e. Safe methods of testing the AV);
- System Safety of the Intended Function (SOTIF) of the Autonomous System (i.e. ensuring that when performing as intended without faults, the design is suitably safe);
- Collecting appropriate evidence to demonstrate safety;
- Presenting a safety argument to show how the individual evidence combines to prove that the project, as a whole, is acceptably safe.

Is Not

- Workshop safety during modification, installation or maintenance work upon the base vehicle(s);
- Office safety during desk-based engineering or administrative activities;
- Analysis or mitigation of risks not associated with safety (e.g. financial risks, reputational risks);

Responsibilities

The overall responsibility for the safety deliverables lies with the CPC.

Parkopedia is the lead partner for the AVP project as a whole, and also lead several work packages, which includes the simulation, mapping development and physical testing of the vehicle, i.e. the activities that carry the risks being considered by the safety documents. As such, Parkopedia assume responsibility for ensuring that testing of the AV technology is conducted safely, and for ensuring that testing undertaken is in accordance with all safety requirements and procedures laid out under the safety documents.

2. Safety Evidence for each Trial Activity

Table 1 shows the level of safety analysis required for each test location and activity, with a particular focus on testing in a controlled environment. The required safety analysis is based on previous CAV projects and common Systems Engineering practices. As the level of detail involved in the safety analysis should be proportional to the expected risk, lower risk activities are marked as minimal safety analysis expected, since the risk is similar to typical manual driving and therefore the safety documentation should be broadly equivalent to the detail in a typical corporate 'driving for work' policy.

	Activity (required evidence must be in place BEFORE this activity commences)					
Evidence	**Manual driving: Data gathering**	**Automated: Testing in a controlled environment**	**Automated: AVP testing in car parks**	**Automated: Demonstration**	**Description**	**Responsible**
Risk Assessment and Method Statement (RAMS) 1 for driving around in car parks using manual mode	X				**Minimal Safety Analysis** needed. Risk comparable to normal driving	CPC
RAMS 2		X			**Moderate Safety Analysis** needed. Involves operation of relatively immature AV, but in controlled environment	CPC
Trial Plan – Testing in a controlled environment		X			Plan showing detailed test cases and test aims	Parkopedia
Test Report - (Testing on Private Track)			X		Report showing that test cases were carried out and no safety issues were identified	Parkopedia
RAMS 3			X		**Detailed Safety Analysis** needed.	CPC

	Activity (required evidence must be in place BEFORE this activity commences)					
Evidence	**Manual driving: Data gathering**	**Automated: Testing in a controlled environment**	**Automated: AVP testing in car parks**	**Automated: Demonstration**	**Description**	**Responsible**
					Relatively mature AV operated in an environment which is only partially controlled.	
Trials Plan – Testing in a Car Park			X		Plan showing detailed test cases and test aims	Parkopedia
Briefing/checklist		x	X	x	Checklist for actions to be performed before trials	Parkopedia/CPC
Test Report – (Testing in Secure Area of Car Park)				x	Report showing that test cases were carried out and no safety issues were identified	Parkopedia
RAMS 4				X	**Detailed Safety Analysis** needed. Relatively mature AV operated in an environment which is only partially controlled.	CPC
Trial Plan – Demo				x	Plan showing detailed test cases and test aims	Parkopedia
Requirements Review	X	X	X	X	**Confirm all requirements applicable to the forthcoming stage of testing have been signed off in the requirement spreadsheet**	CPC, Parkopedia
Failure Modes and Effect Analysis (FMEA)			X	X	To generate safety goals	CPC
System Safety Argument (HARA)			X	X	To generate safety goals, identify hazards and rate the risk	CPC

	Activity (required evidence must be in place BEFORE this activity commences)					
Evidence	**Manual driving: Data gathering**	**Automated: Testing in a controlled environment**	**Automated: AVP testing in car parks**	**Automated: Demonstration**	**Description**	**Responsible**
Verification and Validation of Requirements			X	X	To verify and validate all requirements have been met and tested	CPC, Parkopedia
StreetDrone Safety Document	X	X	X	X	Reference document	StreetDrone
Incident Reporting Spreadsheet		X	X	X	To be kept up to date on a rolling basis, if and when incidents occur	Parkopedia

Table 1 Testing and Trials Activities and Evidence

The following table lists the activities of the AVP project and testing locations:

	Activity			
Test location	**Manual driving: Data gathering**	**Automated: Testing in an open environment**	**Automated: AVP testing in car parks**	**Automated: Demonstration**
Guildford Farnham Road car park	X			X
Turweston Aerodrome		X		
NCP London Bridge	X		X	X
Other (to be completed at a later stage)				

Table 2 Testing Locations - to be completed

3. Argument

Top Level

The safety argument consists of two aspects, the system safety and the operational safety, which together will ensure that the overall goal is met, i.e. that the AVP project is acceptably safe as a whole. Many items

of evidence are required to demonstrate this, but ultimately, they can be viewed as building blocks of four main 'pillars' which the safety case is built upon:

- Verification – safety requirements are created and the vehicle is tested to confirm it meets them
- Validation – the safety performance is monitored as mileage is accumulated
- Safe Working Practices – Method Statements are created to describe safe practices for each test
- Safety Driver Intervention – it is ensured that the driver is able to override when appropriate

It is expected that an R&D prototype vehicle will make occasional errors because of unexpected or unintended behaviour when faced with unprogrammed scenarios, meaning that the Safety Driver acting as a back-up system, plus other procedures such as marshals to ensure the safety of the test area, are key to the AVP safety case. The emphasis for this project is therefore on Operational Safety, although System Safety is considered in recognition of the need for suitable approaches to be developed, informed by R&D trials such as AVP, in order for production-level robustness to be attained when such systems are commercialised.

Figure 2 below illustrates the overall safety argument using Goal Structuring Notation (GSN), a graphical method of displaying how evidence is combined to make up an overall argument, which is widely used within safety engineering. The top goal is for the overall project to be safe, which is met if both System Safety and Operational Safety are acceptable. System Safety is acceptable if Verification and Validation are both completed, and Operational Safety is acceptable if Safe Working Practices are used and the Safety Driver is competent. The circles at the bottom represent 'solutions' required to meet the goals above. Note the 'context' shapes on the left; these indicate that Verification must be seen in the context of the FMEA (Failure Modes and Effects Analysis), the HARA (Hazard Analysis and Risk Assessment) and the System Safety Requirements, all of which are part of the process by which suitable test cases are arrived at so that verification testing can take place.

The 'solutions' and the 'context' elements shown in the diagram therefore form the safety evidence required to show that the overall project is acceptably safe, although note the scheduling in Table 1 (columns 2 to 5); not all evidence is required to be in place for every trial, e.g. no System Safety testing will have taken place prior to the Data Collection or Controlled Environment (i.e. private track) testing commencing. Both the 'Safe Working Practices' and the 'Safety Driver Competent' goals are satisfied by the RAMS (Risk Assessment & Method Statement) for each trial (note that table 1 includes four different RAMS documents to cover the different stages), meaning that the RAMS covers the entire goal of 'Operational Safety is Acceptable'.

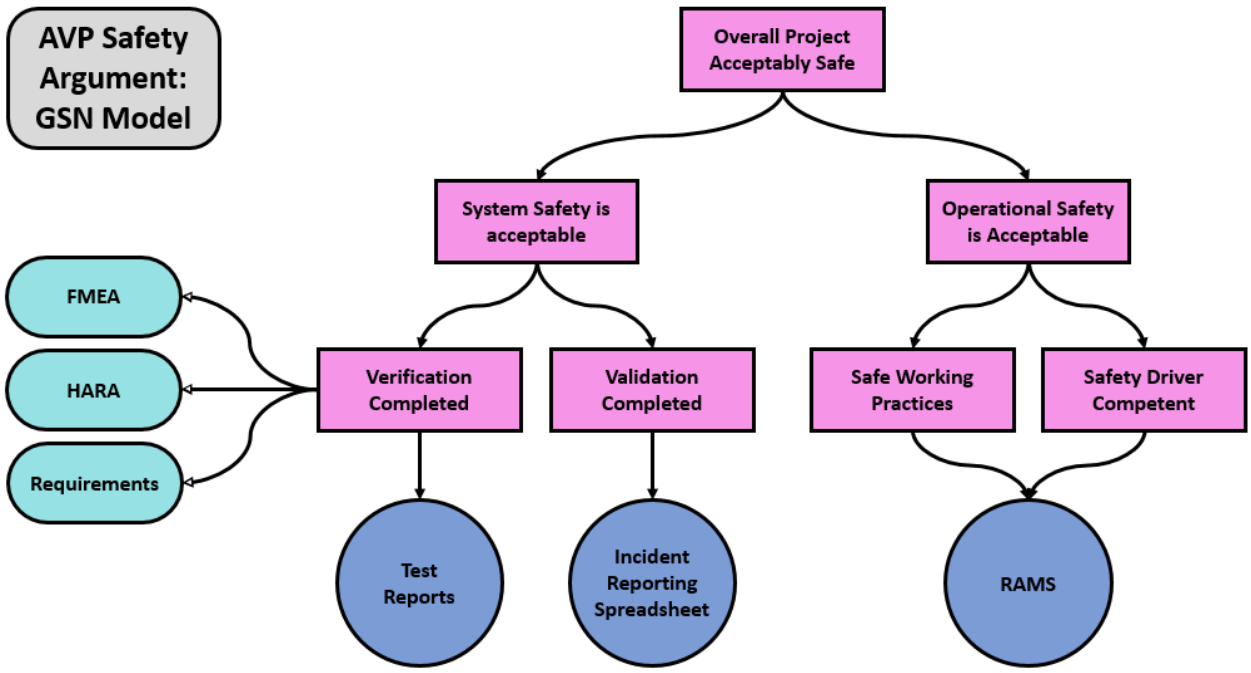

Figure 2 GSN Diagram showing System Safety and Operational Safety

System Safety

Verification

The ultimate output of the verification is a Test Report showing that test cases have been carried out to verify that all the safety requirements have been met. In order for this to give reasonable assurance of safety, a large volume of work is needed prior to this testing commencing, to ensure that the set of requirements generated is suitably comprehensive. This takes two streams: an FMEA (Failure Modes and Effects Analysis) carried out with assistance from Parkopedia, StreetDrone and University of Surrey via brainstorming at a Safety Round-Table meeting, and a HARA (Hazard Analysis and Risk Assessment) of the system, also completed with assistance from Parkopedia and StreetDrone. Both of these documents also form safety evidence themselves. A review of all relevant legislation and standards that was undertaken by CPC for another project was also used as a source of requirements for AVP; this provided a useful prompt for requirements, even though the regulations themselves are not technically applicable to AVP as the project exclusively uses private land for testing.

In addition to this, the safety of the base vehicle is also considered; the StreetDrone vehicle is based upon a production Renault Twizzy, which has been subjected to European Community Whole Vehicle Type Approval Tests for quadricycles. It is therefore only necessary to review the safety of modifications that have been made to the base vehicle, as to repeat any of the safety tests from the type approval process would be prohibitively expensive and time consuming whilst providing little or no safety benefit. The review of the modifications should cover whether safety has been compromised (e.g. sensors located where they could contact an occupant or a pedestrian in an accident). Functional safety of the modifications to the base vehicle undertaken by StreetDrone has been discussed, and it was agreed that they have suitably thorough processes, based upon ISO26262, for it to be reasonable and proportionate for the AVP project to accept the safety of these changes without requiring further evidence, meaning

that only the changes made by the AVP project itself shall be considered further within the scope of the safety deliverables.

Validation

Validation forms a vital component of the test plan as, unlike verification (which is targeted at specific test cases already identified as being of interest), validation allows 'unknown unknowns' to be captured, i.e. it provides a filter to capture potential issues that were not foreseen when the requirements were being generated. This is achieved by accumulating a large volume of testing (in terms of hours or mileage) and monitoring to see if acceptable performance was achieved, logging any issues and taking remedial action as appropriate, regardless of whether the issue relates to a requirement that was previously captured. For production autonomous systems, extremely large volumes of testing will be required to gain reasonable assurance that the system is acceptably safe, but this will not be necessary for AVP, due to the R&D nature of the project and the presence of a Safety Driver at all times.

The initial approach to validation for AVP is for simulation testing to take place before the autonomous system is ever tested in a physical vehicle, as this allows a level of assurance without exposing anybody to physical risks. Success of this shall be monitored via verbal communication, but a formal written report shall not be required, bearing in mind the need to ensure evidence collection is proportionate to the risk, and the need to prioritise resources. Once appropriate control behaviour is achieved, physical testing can commence in a controlled environment on private ground. When the system has accumulated a reasonable number of hours of testing whilst maintaining acceptable safety performance (note that target thresholds have not been set at this stage as it was not felt that there is sufficient evidence to base numerical targets upon, so a subjective judgement on acceptability must be made once each test phase is complete), testing within a secure area of a car park can begin, and once this has been satisfactorily completed, the full demonstration, potentially including autonomous driving through areas of the car park that are open to the public, can commence.

An incident reporting spreadsheet will be maintained to keep an up to date log of issues encountered during any phase of testing. In so doing, it will be ensured that the complexity of scenarios presented to the vehicle is progressively increased throughout the life of the project, but that increases don't occur until the vehicle has reached a reasonable level of performance in the previous stage of testing.

Operational Safety

Operational safety is concerned with how the system is deployed, and consists of two strands:

- Test Conduct, for which safe methods of working are set out in a RAMS (Risk Assessment and Method Statement) document, which describes the risks presented by the testing, and the procedures in place to mitigate against them. There is one RAMS document for each testing activity.
- Ability of the Safety Driver to intervene, which is also covered in the RAMS document applicable to each phase of testing, and is dependent upon:
    - Confirmation that all safety drivers hold an appropriate class of driving licence;
    - Knowledge of the AVP system's Operational Design Domain (i.e. the limits within which the system is intended to operate) – for example, if the vehicle can't detect or navigate around dynamic obstacles, the safety driver needs to know that they would need to intervene when the vehicle encounters this scenario;

- Knowledge of how to disable and/ or override the autonomous system;
- A Safety Driver Acclimatisation Policy to ensure that all safety drivers are familiar with the AVP Human-Machine Interface, including experience of overriding the system, prior to testing in car parks. Any new safety drivers joining midway through the project would therefore have to gain experience on a private track before taking responsibility for the vehicle in a car park. The purpose of this step is both to ensure that the driver is capable of making an intervention, but also that the vehicle is capable of accepting an intervention and responding in an appropriate way (e.g. not fighting against the driver).

Evidence Collection

Some pieces of evidence do not take the form of a one-off report or statement, but instead will take the form of a spreadsheet where data is entered on an ongoing basis throughout the project, as and when appropriate. For example, the Incident Reporting Spreadsheet will be available to enter data upon from the start to the end of physical testing of the vehicle, and includes a procedure to describe what level of incident requires an entry to the spreadsheet, and which incidents do not need to be reported beyond that of the testing organisation.

Any partner listed as responsible (as shown in Table 1) for a document that is completed on an ongoing basis has an ongoing responsibility to ensure that it is kept up to date throughout the duration of the project.

It is important to ensure that the time and cost expended in creating safety evidence is proportionate to the safety benefits, and therefore, given the limited scale and scope of the project, formal reports are not required for all evidence; for example, test reports could be in the form of a spreadsheet showing a test matrix that has successfully been completed. However, key deliverables such as the RAMS for each phase of testing, the requirements spreadsheet, the FMEA and the HARA (detailed in the next chapter) will be formally documented, and wherever possible will be published online to assist other projects.

The Safety Case will not be a single document, but instead a collection of documents that, when taken together, indicate suitable safety in line with the justifications described in this safety plan. This will require progressively more evidence to be put in place as the trial moves on to more challenging scenarios; the evidence required prior to the commencement of each new stage of testing is detailed in Table 1.

4. System Safety Analysis

Background

Although this project will feature a safety driver ready to override the system and will only operate in tightly-controlled environments (i.e. Operational Safety forms the key to the Safety Case), there is still the potential to reduce risks by increasing the robustness of the system such that the safety driver has to intervene less often; such an improvement in robustness would also significantly improve the quality of the demonstration. Furthermore, if such systems are to be industrialised, it is essential that system safety

is developed to the point that a safety driver is not required. The AVP Safety Case therefore includes significant consideration of System Safety.

There were two main approaches to analysing system safety, the FMEA and the HARA, which are described subsequently.

FMEA

FMEA, or Failure Modes and Effects Analysis, is a widely-used technique within safety engineering. The architecture of a system is evaluated, considering each possible fault that could occur within the system and identifying how that fault will affect the rest of the system, including any failures or other undesirable behaviours that the fault could cause the system as a whole to display. This is referred to as a 'bottom-up' or 'inductive' approach, as it starts with the fault (at a low level) and works up to the effect on the high-level behaviour of the system, as opposed to 'top-down' or 'deductive' analysis (e.g. Fault Tree Analysis), which starts at the high-level failure and works downward to identify the fault(s) that form the root cause.

In the case of AVP, a Functional Architecture Diagram was developed with the help of Parkopedia, identifying the subsystems that make up the overall system, and the sub-subsystems that make up each subsystem. This architecture can be seen in the ConOps (Concept of Operations) document. A workshop was then held with the project consortium, where participants were asked to identify what possible failures could occur in each sub-subsystem, what the local effect would be on the subsystem in which it resides, and whether this would result in any undesired vehicle-level behaviours. This activity was undertaken as a Post-It Note exercise, with subsequent follow-up work to close off open points then completed by Parkopedia and CPC.

It was decided not to apply severity and likelihood scores to each failure (this is commonly referred to as Failure Modes, Effects and Criticality Analysis, FMECA) due to the difficulty in estimating reasonable scores (each failure typically causes a low chance of a high severity accident, a higher chance of a lower severity accident, plus an infinite number of variations on the spectrum between these points), the limited resources available, and the fact that there would have been little that the creation of such prioritisation scores would achieve in the project anyway. However, Safety Goals (i.e. high-level requirements for what should be done to mitigate the risks) were logged for each fault, and these were incorporated into the Requirements Spreadsheet, with corresponding test cases and acceptance criteria being added.

The full analysis can be viewed in the AVP FMEA document.

Hazard Analysis and Risk Assessment (HARA)

The 'AVP System HARA' document is a Hazard Analysis and Risk Assessment from the perspective of System Safety, as opposed to Operational Safety (i.e. it looks at faults within the system and how the system would respond if presented with a situation that is outside its limitations, but does not look at operational procedures for how tests are carried out; this is covered in the Risk Assessment & Method Statement for that particular test). A 'Hazard' refers to the situation that causes a possibility of death or injury, and a risk is an instance of such a hazard, including an assessment of the likelihood that the risk

materialises and the severity of the consequences; in simple terms, the HARA therefore identifies what could go wrong and then assesses how acceptable this is.

The possible vehicle level faults identified in the Failure Modes and Effects Analysis (FMEA) were used to generate hazards; other hazards were identified via brainstorming, prior experience and knowledge of the system limitations. This covers both hazards due to direct failures (ISO26262) and unintended and hazardous but otherwise failure-free system deficiencies (SOTIF).

Three separate risk assessments are provided for the same hazards, depending on whether the vehicle is deployed on a private track, in a secure area of a car park, or as part of a demonstration involving travelling through non-secure areas of the car park.

Much experimentation was carried out with the scoring system; as per the FMEA, it was identified that a traditional Probability versus Severity approach to Risk Assessment isn't feasible for such projects as there is a continuum of possible values that could be used, from the less likely but more severe accident permutations to the more likely and less severe ones. The approach that was finally settled upon was thus:

The Scoring for Probability (P) is based on both the probability of the hazardous vehicle behaviour occurring AND the probability that other traffic and obstacles around the vehicle poses a threat at the moment that the hazard is triggered such that it is reasonably possible for an incident to result. This is multiplied by the score for severity (S) of the resulting incident, and the controllability (C) a low score indicating that the driver has a good chance of being able to prevent or mitigate the incident, to obtain an overall Risk Score (R). A schematic of the approach is shown in Figure 3, and the scoring system is detailed in Figure 4.

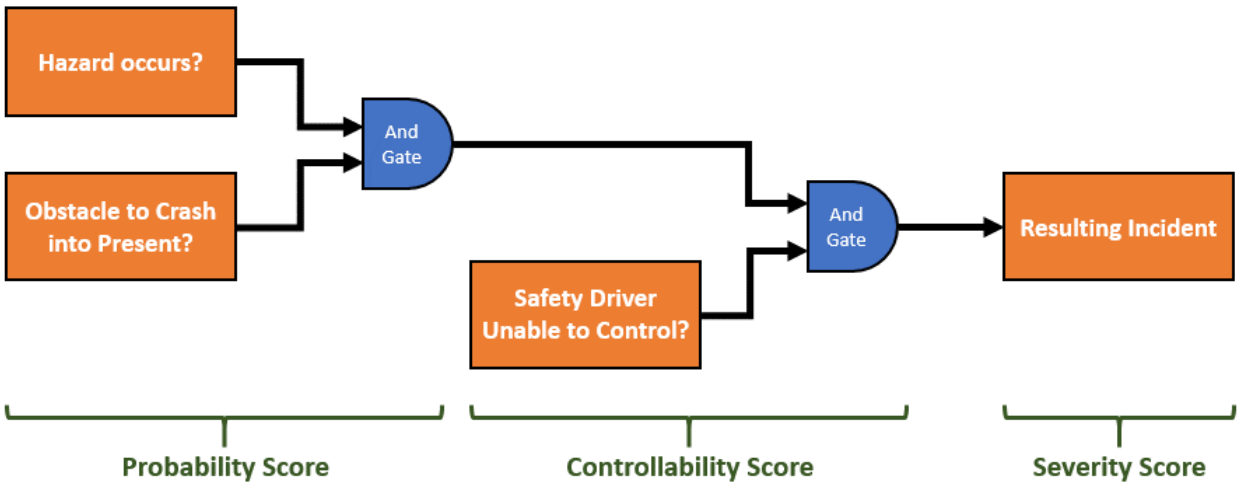

Figure 3 Schematic representation of the HARA Scoring System

Probability	
1	Very unlikely
2	Unlikely
3	Possible
4	Likely
5	Very Likely

Severity	
1	Minor injury – No first aid treatment required
2	Minor injury – Requires First Aid Treatment
3	Injury - requires GP treatment or Hospital attendance
4	Major Injury
5	Fatality

Probability considers the likelihood of 2 factors both required for an incident; how likely is the hazard to occur, and how likely is it that the surrounding environment will have the right combination of features to cause this hazard to result in an actual incident (e.g. oncoming traffic, pedestrian near vehicle etc.)

Controllability	
1	Very easily controllable
2	Easily controllable
3	Moderately controllable
4	Difficult to control
5	Not controllable

Overall Risk Scoring	
R >= 30	Unacceptable Risk - must be reduced before activity can begin
40 > R >= 30	Tolerable Risk (upper) - must be reduced if reasonably practicable within project resources
30 > R >= 18	Tolerable Risk (lower) - should be mitigated if possible to achieve with minimal additional burden
R < 18	Acceptable Risk - no further mitigation necessary
The overall Risk Level is scored by multiplying together the Probability, severity and (difficulty of) Controllability	

Controllability could be regarded as a measure of how difficult the scenario is for the safety driver to correct, i.e. the lower the score, the easier it is to control (the numbering was set out this was to allow multiplication to generate the Overall Risk Score

Figure 4 Numerical Scoring within the HARA

In line with Health and Safety Executive (HSE) Guidance, risks were rated as Unacceptable, Tolerable or Acceptable. However, the HSE approach advises that where risks are Tolerable, they should be demonstrated to be As Low As Reasonably Practicable (ALARP), assessed with techniques such as comparing the cost of fatalities to the cost per life saved of mitigation measures. For an R&D project, such an analysis isn't feasible, so it was deemed that a better approach to judging whether 'reasonably practicable' steps have been taken was to separate Tolerable risks into an Upper (amber) and Lower (yellow) band, with a higher burden of mitigation in the upper band.

The HARA included a column to log existing measures to mitigate against the risk, an assessment of the risk level using the scoring system above, and then where applicable, further mitigation measures added to reduce the score (particularly for risks in the red or amber categories) and a reassessment of the risk with these additional mitigation measures in place.

5. Updating the Safety Plan

This Safety Plan was developed in collaboration with Parkopedia as a draft version. The first version was released before testing in a controlled environment; this is the second version approved by CPC and Parkopedia before the physical testing of the complete system commenced. Subsequent updates will then be made for the duration of the particular testing phase, in response to lessons learned as the project progresses and in response to any changes in the system itself. Any such updates to the approved document will be listed as new versions within the version control section, and it shall be made clear what the changes are. These changes must be agreed by all Partners that are affected.

The validity of the safety assumptions made within the safety case and the associated evidence shall be reviewed after each phase of testing is complete, and the likelihood and severity of risks corrected if empirical evidence shows the initial estimates to be inaccurate.

Once testing is complete, the latest version of the document will become the final version and no further changes will be made. A separate Safety Case will be created for the following testing phases (testing in car parks and demonstration).

6. Conclusion

Any deviation from understood and predicted behaviour should be immediately recorded and investigated.

Safety is Everyone's Business and Historical Accidents

The history of the railways, aeroplanes and surgery provides many common elements in the causes of accidents.
Firstly, anyone present during the trial should be both able to stop it at any point, but feel confident to do so. There are surgical fatalities that could have been prevented if the junior nurses had not been ignored. Similarly, there have been preventable aeroplane crashes because the junior pilots were worried about speaking up. Both the surgical teams and airline crews now recognise this failing and practice the intervention of all members of the team.
To this end, everyone who will be present should attend a practice session where they say "STOP" out loud in response to a concern.

Safety Driver Awareness of failure, reaction times and training:

- Choice reactions take longer. Simple sports reactions are about 100 milliseconds, but adding choice and complexity slows the response time: an additional choice slows the reaction by another 50 milliseconds. The safety driver has been given a very complex job of deciding when to take an emergency action in an environment where the modes of failure may have different onsets. For example, a sudden swing of the path or a slow drift will both need an intervention, but the former needs a sudden one and the later will be a judgement of whether a limit has been exceeded.
- The various AI failure modes can be dealt with by the safety driver, if they make the correct timely decision. Any base vehicle failures are within the scope of normal, ordinary driving. In a non-autonomous car, if the hydraulic brakes fail, the driver should normally try the handbrake or do whatever they can to bring the vehicle to a safe stop. So the question becomes how to train the safety driver to make that correct timely decision for the autonomous vehicle.
- It is suggested that the safety driver for this demonstration practice recognising the most likely and serious failure modes by marking out the proposed site in an open area, then using the StreetDrone remote control to emulate the Autoware and its failure modes, i.e. sudden swings toward parked cars or slow drifts of course.There are of course an infinite number of ways for the demonstration to fail and only a few to succeed, but if the principal characteristics of the failures can be observed with practice, then the realistic threats can be ameliorated.

7. References

Document Name	Owner	Release date and version
CCAV Code of Practice for Automated Vehicle Trialling	CCAV	February 2019
Concept of Operations (ConOps)	CPC	20/12/2018 v2.0
Failure Mode and Analysis (FMEA)	CPC	18/11/2019 v1.0
Hazard Risk Assessment (HARA)	CPC	18/11/2019 v1.0
Safety Case 3	CPC	18/11/2019 v1.0
Risk Assessment and Method Statement 2 (RAMS 2)	CPC	07/06/19 v1.0
Risk Assessment and Method Statement 3 (RAMS 3)	CPC	18/11/2019 v1.0
Requirement Management Spreadsheet	CPC	18/11/19 v2.0

Table 3 References

CATAPULT
Connected Places