(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
12 September 2019 (12.09.2019)

WIPO | PCT

(10) International Publication Number
WO 2019/170616 A1

(51) International Patent Classification:
H04L 29/08 (2006.01) *G06F 21/55* (2013.01)
H04L 29/06 (2006.01) *G06N 99/00* (2019.01)
H04L 12/40 (2006.01)

(21) International Application Number:
PCT/EP2019/055341

(22) International Filing Date:
04 March 2019 (04.03.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
18159885.5 05 March 2018 (05.03.2018) EP

(71) Applicant: **BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY** [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventors: **EL-MOUSSA, Fadi**; Ground Floor, Faraday Building, 1 Knightrider Street, London EC4V 5BT (GB). **SMITH, Karl**; Ground Floor, Faraday Building, 1 Knightrider Street, London EC4V 5BT (GB).

(74) Agent: **BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY, INTELLECTUAL PROPERTY DEPARTMENT**; Ground Floor, Faraday Building, 1 Knightrider Street, London EC4V 5BT (GB).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,

(54) Title: VEHICULAR NETWORK SECURITY

FIGURE 2

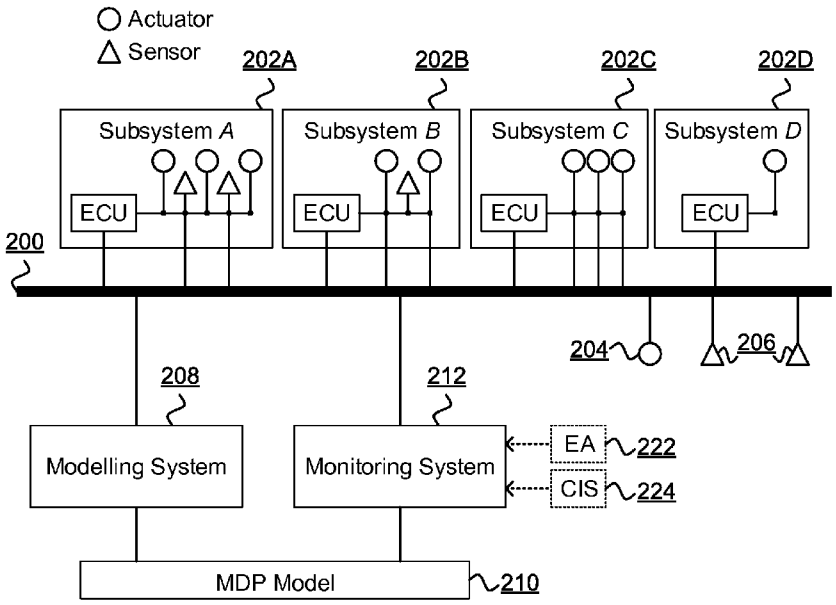

(57) **Abstract:** A computer implemented security method operable with a communications network in a vehicle, the network communicatively connecting devices including sensors and actuators in the vehicle such that information provided by sensors and states of actuators are determinable by data communicated via the network, the method comprising the steps of: defining a Markov decision process model for the vehicle, the model specifying states of the vehicle and actions constituting transitions between states, wherein a state of the vehicle is indicated by information provided by one or more sensors and a state of one or more actuators, and an action corresponds to a change in the information provided by one or more sensors and/or a change to a state of one or more actuators, each action having associated a probability of occurrence; determining, by accessing data communicated via the network, a current state of the vehicle in the model; accessing data communicated via the network; responsive to the accessed data indicating an action to change the

WO 2019/170616 A1

[Continued on next page]

HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

vehicle state to a new state, determining, from the model, a probability of the action; responsive to a determination that the determined probability falls below a predetermined threshold probability, generating an indication that the vehicle state transition is anomalous.

Vehicular Network Security

The present invention relates to a security method for a vehicle. In particular, it relates to security of a vehicle communications network.

Modern automobiles are increasingly provided with data networks providing communication between vehicle subsystems. Such networks can be used to provide services such as engine and transmission management, actuator controls and advanced features such as driving assistance including parking assistance, collision detection, collision avoidance, automated and/or assistive braking, adaptive cruise control and the like.

Data networks are susceptible to attack and vehicles are susceptible to negligent control. In vehicular applications, the consequences of unauthorised, malicious and/or negligent access to a vehicle data network can be serious including a risk to the safety and security of a vehicle and ultimately a risk to life.

Consequently, there is a need to mitigate these risks.

The present invention accordingly provides, in a first aspect, a computer implemented security method operable with a communications network in a vehicle, the network communicatively connecting devices including sensors and actuators in the vehicle such that information provided by sensors and states of actuators are determinable by data communicated via the network, the method comprising the steps of: defining a Markov decision process model for the vehicle, the model specifying states of the vehicle and actions constituting transitions between states, wherein a state of the vehicle is indicated by information provided by one or more sensors and a state of one or more actuators, and an action corresponds to a change in the information provided by one or more sensors and/or a change to a state of one or more actuators, each action having associated a probability of occurrence; determining, by accessing data communicated via the network, a current state of the vehicle in the model; accessing data communicated via the network; responsive to the accessed data indicating an action to change the vehicle state to a new state, determining, from the model, a probability of the action; responsive to a determination that the determined probability falls below a predetermined threshold probability, generating an indication that the vehicle state transition is anomalous.

Preferably, the predetermined threshold probability is based on a set of all possible actions in the current vehicle state defined in the model and probabilities associated with each action in the set such that the predetermined threshold probability serves to identify relatively improbable actions.

Preferably, the predetermined threshold probability is a proportion of a mean probability for all actions in the set.

Preferably, the Markov decision process model is defined empirically based on observation of vehicle states and actions in acceptable use.

Preferably, the Markov decision process model is defined by a reinforcement machine learning method based on observation of vehicle states and actions in acceptable use.

Preferably, the reinforcement machine learning method is a Q-learning reinforcement method based on a predetermined definition of vehicle states and actions and using a reward function wherein actions leading to desirable and/or acceptable vehicle states are attributed relatively greater reward so as to inform the Q-learning method to assign relatively greater probabilities to actions with relatively greater rewards.

Preferably, the method further comprises, responsive to the anomalous indication, precluding the transition of the vehicle to the new state by interrupting the action.

Preferably, the action involves a change in state of at least one actuator, and interrupting the action includes preventing the action by preventing a change of state of the at least one actuator.

Preferably, the action arises from a change in the information provided by one or more sensors, and interrupting the action includes forcing a transition of the vehicle state to a third state as a reactive state.

The present invention accordingly provides, in a second aspect, a computer system including a processor and memory storing computer program code for performing the steps of the method set out above.

The present invention accordingly provides, in a third aspect, a computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of the method set out above.

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram a computer system suitable for the operation of embodiments of the present invention;

Figure 2 is a component diagram of a vehicle communications network for a vehicle in accordance with embodiments of the present invention;

Figure 3 is an illustrative depiction of an exemplary Markov decision process model suitable for embodiments of the present invention;

Figure 4 is a flowchart of a method operable with a communications network in a vehicle to generate an indication that a vehicle state transition is anomalous in accordance with embodiments of the present invention;

Figure 5 is a component diagram illustrating the operation of the modelling system in accordance with embodiments of the present invention; and

Figure 6 illustrates a exemplary partial Markov decision process model in accordance with an embodiment of the present invention.

The modern automobile may have many electronic control units (ECU) for various subsystems. Commonly, an ECU is provided for engine control. Additional ECUs can be provided for other functions and features of a vehicle such as transmission, airbags, antilock braking/ABS, cruise control, electric power steering, audio systems, power windows, doors, mirror adjustment, battery and recharging systems for hybrid/electric cars, etc. Some of these form independent subsystems though communications among others can be essential.

A vehicle subsystem may need to control actuators and/or receive feedback from sensors. For example, an auto start/stop facility can require sensor input to inform vehicle speed, steering angle and engine temperature. An electronic parking brake providing a “hill hold” function can require sensor input from tilt sensors and road speed sensors. Such sensors can have additional function, such as tilt sensors employed for a burglar alarm system, or road speed sensors employed by antilock braking, engine control or traction control functions. Parking assist can respond to engagement of reverse gear with information provided by a transmission control unit. There may also be an associated adjustment of door-mirrors to assist with reversing. Auto lane assist and collision avoidance systems can require information from external sensors (such as may be used for parking) and/or cameras capable of providing proximity data and road feature data for a system such as a lane departure warning system. Such information can also be used to inform an automatic braking or brake assist facility, vehicle steering assistance for collision avoidance, or adaptive cruise control including either or both of speed and proximity management and lane assistance through assisted steering.

Sensors and actuators in vehicles increasingly interoperate by communicating via a vehicle communications channel such as a common communications bus, wired or wireless communication means and/or regulated notification and/or communication modulators. Such communication channels are hereinafter referred to as a communications network in the vehicle. For example, the communications network can be a wired, wireless or combination computer network for communicating information from at least a subset of sensors and communicating signals and information to (and potentially from) at least a subset of actuators. Some sensors and/or actuators may be grouped into vehicle subsystems such as an engine management system, a transmission control system, a ventilation system or the like, and such systems can optionally associate with or include a control unit such as an ECU. The function, purpose and application of any particular ECU is beyond the scope of this description though it can conceivably include one or more of, inter alia: a controlling function for controlling sensors/actuators; a communication function for communicating with sensors/actuators such as providing an interface between sensors/actuators and the communications network, such interface may provide, for example, data format conversion, data standardisation or normalisation, data collection and/or aggregation; a monitoring function for sensors/actuators; an alerting function; intelligence such as one or more of software, hardware, firmware or other logic operable to manage, trigger, operate or instruct actuators, such management may be responsive to sensors within the subsystem or information received via the communications network from sensors, actuators and/or other subsystems; and other functions as will be apparent to those skilled in the art.

Communications networks are vulnerable to attack, unauthorised access, malicious access, malicious or negligent communication or interception. For example, vehicle actuators may be controllable via a vehicle communications network so putting a vehicle and its occupants at risk of harm. Vehicle security can be compromised putting a vehicle at risk of unauthorised physical access and theft. Vehicle subsystems can be modified, manipulated or maliciously or negligently controlled which can affect vehicle occupants and other road users risking injury and possible death.

Embodiments of the present invention include a modelling system for building a Markov decision process (MDP) as a model of the operation of a vehicle. MDPs specify states and actions constituting transitions between states, each action having a probability of occurrence, as is known in the art. In accordance with embodiments of the present invention, an MDP model is applied such that a state of a vehicle as indicated by information provided by one or more sensors and a state of one or more actuators serves to identify a state in the MDP for the vehicle. Actions for transitioning between states in the MDP correspond to a

change in the information provided by one or more sensors and/or a change to a state of one or more actuators.

In embodiments of the present invention, the MDP model is generated to include a probability associated with each action. The probabilities can be determined empirically and/or based on a reinforcement learning approach as is described below. Thus, the model can be generated during a training, learning or modelling phase of operation of a vehicle. Subsequently, for the vehicle in normal operation, the model is employed to identify deviations from predetermined acceptable state transitions as anomalous state transitions, such anomalies serving, for example, to trigger responsive, remedial, protective or alerting actions. Thus, embodiments of the present invention further include a monitoring system operable with the MDP model to monitor a state of a vehicle and actions triggering state transitions and to determine deviations from predetermined acceptable state transitions based on action probabilities and a predetermined threshold probability. For example, actions arising in the vehicle that have a probability in the MDP model below a threshold probability can be identified as anomalous. Most preferably, the threshold probability is based on probabilities of all actions in a set of all possible actions in a current vehicle state such that the predetermined probability threshold serves to identify relatively improbable actions. For example, a threshold probability can be predetermined as a predetermined proportion of a mean probability of all actions from a current vehicle state, such as 20%, 25%, 30%, 35%, 40%, 45%, 50%, 55% or 60% of the mean probability. In some embodiments, the predetermined portion of a mean probability can itself be determined based on one or more attributes, statistics or features of the set of actions for a vehicle state, such as one or more of: a number of actions in the set; a standard deviation of the probabilities of actions in the set; a range of probabilities of action in the set; and other attributes of probabilities of actions in the set.

Thus, in this way embodiments of the present invention identify anomalous vehicle state transitions, responsive to which reactive measures can be taken. Such reactive measures can include protecting the vehicle, its occupants and/or other road users by one or more of: preventing, precluding or inhibiting a state transition of the vehicle state; forcing a state transition to a different vehicle state such as a reactive state as a state of the vehicle for alerting vehicle operators, protecting the vehicle, its occupants or other road users, evading or mitigating a certain or possible harmful outcome such as a collision, or other reactive states as will be apparent to those skilled in the art.

Figure 1 is a block diagram of a computer system suitable for the operation of embodiments of the present invention. A central processor unit (CPU) 102 is

communicatively connected to a storage 104 and an input/output (I/O) interface 106 via a data bus 108. The storage 104 can be any read/write storage device such as a random access memory (RAM) or a non-volatile storage device. An example of a non-volatile storage device includes a disk or tape storage device. The I/O interface 106 is an interface to devices for the input or output of data, or for both input and output of data. Examples of I/O devices connectable to I/O interface 106 include a keyboard, a mouse, a display (such as a monitor) and a network connection.

Figure 2 is a component diagram of a vehicle communications network 200 for a vehicle in accordance with embodiments of the present invention. The communications network 200 is illustrated as a singular network though it will be appreciated that a combination of two or more suitable interoperating network mechanisms may be employed including wired and or wireless networks. The vehicle communications network 200 is provided for a vehicle such as an automobile, a transport vehicle or the like, and provides communications facilities between sensors and actuators. Notably, the communications network 200 can additionally provide communication facilities between other components though these are beyond the scope of this description. Sensors and actuators can operate in direct communication via the network 200 such as sensors 206 and actuator 204. Alternatively, sensors and actuators can operate as part of one or more subsystems 202A-202D of the vehicle. A subsystem is a collection of one or more sensors and/or actuators and can be optionally associated with an electronic control unit (ECU). Sensors and actuators within subsystems may communicate via the network 200 directly, via an ECU or conceivably both.

Examples of sensors can include, inter alia: speed sensors; position, location and/or orientation sensors; incline sensors; environmental sensors such as humidity, temperature, rain, snow, fog, wet, light, dark or other such similar sensors; traction sensors; emission sensors; occupant sensors such as presence sensors (e.g. seat sensors), safety device sensors (e.g. seatbelt sensors), occupant characteristic sensors (e.g. seatbelt extension measurement, weight sensors, height sensors and the like); proximity and object detection sensors such as reversing sensors, vehicle proximity sensors, range detectors and the like; mass air flow sensors; engine speed sensors; tyre pressure sensors; fuel level, fuel quality, fuel temperature, fuel pressure, fuel-air-mix and/or other fuel sensors; voltage and/or current sensors; telematics sensors; crankshaft angle sensors; camshaft angle sensors; exhaust oxygen sensors; exhaust gas recirculation valve position sensors; spark knock sensors; fuel vapour pressure sensor; engine coolant temperature sensor; engine compartment temperature sensor; manifold air temperature sensor; transmission oil temperature sensor; anti-lock brake sensor; windshield crack sensor; door or tailgate ajar sensors; air conditioning pressure sensor; steering angle and/or position sensor; crash and/or impact sensors; pedal,

transmission or switch change sensors; and other sensors as will be apparent to those skilled in the art.

Examples of actuators can include: engine actuators such as actuators to adjust engine speed, performance, activation state and the like; steering actuators such as actuators to adjust steering direction, angle, responsiveness and the like; transmission actuators such as actuators to adjust a manual or automatic gearbox including a gear or range selection, clutch operation, vehicle direction and the like; illumination actuators for changing the state of internal and or external indicators such as lights and the like; environmental facilities such as actuators for operating wipers, heaters, heated vehicle areas or parts such as passenger or engine compartments, vehicle glazing or panels; ventilation actuators to adjust ventilation facilities such as fans, air-intake settings, air conditioning, air recirculation, air temperature and the like; braking actuators such as all-wheel, individual wheel or subset combination braking, antilock braking facilities and the like; driver and/or occupant information facilities such as displays, indicators, screens, lights, sounds, buzzers, audible messages and the like; feedback facilities such as haptic feedback for drivers and/or occupants and the like; traction control facilities; and other actuators as will be apparent to those skilled in the art. Notably, an actuator has associated characteristics that may be adjustable, adaptable, modifiable, configurable or otherwise changeable. For example, wipers can be inactive, active at a first speed, active at a second speed, intermittently active and so forth. Information regarding, about or a status of such characteristics for an actuator can be communicated via the network to other components communicating thereby so as to represent a “state” of the actuator”. E.g. wipers can be in an “intermittently active” state. Furthermore, there may be multiple characteristics of an actuator, and each characteristic can have a particular status (e.g. wiper activity characteristic can be “true” to indicate that a wiper is active; and wiper speed characteristic can be “2” to indicate a second speed of wiper action). The state of any particular actuator is a set of one or more statuses of characteristics of the actuator. Notably, the state of an actuator is different to, and distinct from, the status of a vehicle in a MDP model since the state of an actuator relates only to the particular actuator. Further, a subsystem may have a state that may be derived from, compiled of, consist of or aggregated from characteristics of one or more actuators within the subsystem and/or information from one or more sensors in the subsystem.

Figure 2 also includes a modelling system 208 as a hardware, software, firmware or combination component for modelling the operation of the vehicle in a training mode of operation in which the vehicle is operated in a predetermined acceptable manner of operation. For example, in the training mode of operation, the vehicle may be operated in a manner such that the network 200 and actuators and sensors of the vehicle are isolated from

potential sources of malicious or negligent access, intrusion or the like, such as within a controlled, managed and/or contained environment. The modelling system 208 is operable to generate an MDP model 210 modelling a plurality of vehicle states for the vehicle and, for each state, one or more actions for transitioning the vehicle to another state, the actions having associated probabilities.

In one embodiment, the set of vehicle states and actions between states for the vehicle are predefined, such as from a manufacture of the vehicle or from a prior analysis of the vehicle subsystems, actuators and sensors. In such embodiments, the modelling system 208 is configured to generate the MDP model 210 by determining probabilities for association with the actions within set of actions for each state.

Additionally, or alternatively, the set of vehicle states and actions between states can be at least partly determined by an observation and/or learning process such as an empirical process based on observation of the vehicle operating in the training mode of operation according to the predetermined acceptable manner of operation of the vehicle.

The predetermined acceptable manner of operation of the vehicle can be defined and/or determined by one or more of, inter alia: a manufacturer of the vehicle; vehicle standards organisations and agencies including government, safety, academic and other organisations; driver, owner or operator preferences for the vehicle; a database of acceptable, preferred, or otherwise identified vehicle states and/or actions; or other methods as will be apparent to those skilled in the art. Generally, in embodiments of the invention, the predetermined acceptable manner of operation corresponds to operations of the vehicle resulting in the vehicle being placed in to vehicle states that are preferred in a production mode of operation of the vehicle (i.e. in normal use after the training mode of operation).

The modelling system 208 thus determines probabilities for association with actions for each vehicle state in a set of vehicle states, with actions having a greater probability indicating a more typical, acceptable and/or safe action at a vehicle state than actions having a lower probability. In one embodiment, a reinforcement machine learning algorithm is employed to determine the probabilities associated with actions in the set of vehicle states and actions to arrive at the MDP model 210. For example, the Q-Learning reinforcement method (as first instroduced in “Learning from Delayed Rewards” (Watkins, King’s College, 1989) and described in “Reinforcement Learning: An Introduction” (Sutton and Barto, 2017), in particular section 6.5) can be employed based on a definition of vehicle states and possible actions using a reward function in which actions leading to predetermined desirable and/or acceptable vehicle states are attributed relatively greater reward. In this way, the Q-Learning method can, for example, assign relatively greater probabilities to actions with

relatively greater reward. The reinforcement learning method can use back-propagation of rewards such that states leading to states having high reward also enjoy some reduced proportion of the reward, and therefore improvements to the probabilities of actions linking those states, to emphasise the desirable output state and paths leading to it. Thus, in this way, the MDP model 210 is generated including vehicle states, actions at vehicle states and probabilities of each action.

Figure 3 is an illustrative depiction of an exemplary Markov decision process model 210 suitable for embodiments of the present invention. The model 210 of Figure 3 includes five states S_1 to S_5 each having actions (outgoing arrows), with each action having associated a priority p. For example, a vehicle in state S_1 will transition to state S_2 with a priority of 1 (i.e. certainty). Subsequently, the vehicle will transition to state S_6 with a priority of 0.8 or to state S_1 with a priority of 0.2. This, in this way the vehicle states, relationships between the vehicle states, and the probability of transitions between vehicle states is represented by the MDP model 210.

Returning to Figure 2, the vehicle of Figure 2 also includes a monitoring system 212 as a software, hardware, firmware or combination component for accessing data communicated via the network 200 to determine the vehicle state in the MDP model 210. Further, the monitoring system 212 is operable to access data communicated via the network 200 to identify an indication of an action occurring in, with or by the vehicle to change the vehicle state to a new state.

Data accessed or obtained by the monitoring system 212 via the network 200 can include, for example, indications of actuator or subsystem state received by, from or obtained by the monitoring system 212 from subsystems 202A-202D or actuators. Additionally, or alternatively, such data can include sensor information obtained from, or communicated by, subsystems 202A-202D or sensors.

In some embodiments, at least some data communicated via the network 200 may be secured by, for example, encryption or the like. In such embodiments the monitoring system 212 operates as a trusted component operable to access such data by, for example, decrypting the data, such as by the monitoring system 212 having access to one or more cryptographic keys (or a separate security component having or accessing such keys) for decrypting such data.

Using accessed data communicated via the network 200, the monitoring system 212 determines a vehicle state (or a most likely vehicle state in the event of only partial or insufficient information) of the vehicle in the MDP model 210. Furthermore, using such data,

the monitoring system 212 identifies indications of an action that would transition the vehicle from its current vehicle state to a new vehicle state. In one embodiment, the occurrence of such an action is determined with reference to the MDP model 210 by identifying a set of all actions available for the vehicle in the current vehicle state and, for each action, identifying a subsequent state that would be achieved should the action take place. Each state in the set of identified subsequent states would constitute a potential new state of vehicle, and one of the potential new states is selected from the set of identified subsequent states based on the data received via the network 200.

By way of example, consider the exemplary partial Markov decision process model 210 of Figure 6. This model includes a current state S_1 in which vehicle speed is greater than 30kph, the accelerator is depressed and a second gear is selected. This information is determined by the monitoring system 212 from data communicated via the network 200. For example: the vehicle speed can be determined from data originating from a speed sensor; the state of the accelerator can be determined from an accelerator pedal sensor; and one or more states of actuators in the transmission system will indicate the current gear. In this vehicle state S_1, if the monitoring system 212 accesses data communicated via the network 200 indicating that the conditions of vehicle state S_5 were being met then the monitoring system 212 can conclude that action A_5 to transition to state S_5 is being performed. For example, such data could be indicated by: transmission actuator state changes to reverse gear; a rear-screen wiper actuator having a state of “active”; and a reverse parking sensor system actuator having a state of “active” and/or proximity sensor data being received from reverse parking sensors. Thus, in this way, the monitoring system 212 is operable to determine a current vehicle state and a new vehicle state with reference to the MDP model 210.

In use, the monitoring system 212 is further operable to determine a probability of a detected action with reference to the MDP model 210. Thus, in the previous example where the modelling system 212 concluded action A_5 occurs to transition the vehicle to vehicle state S_5, the monitoring system 212 also determines the probability of action A_5 with reference to the MDP model 210 which, according to Figure 6, is p=0.01. Thus, the monitoring system 212 determines that a probability of the observed action A_5 occurring for the vehicle in current state S_1 is 0.01. Subsequently, the monitoring system 212 determines if the probability of the determined actions falls below a predetermined threshold probability.

Preferably, the predetermined threshold probability serves to identify relatively improbable actions. In one embodiment, the predetermined threshold probability can be predefined as a constant or adjustable value for all actions or for all actions of a particular vehicle state. In

other embodiments, more sophisticated approaches are adopted in which the predetermined probability is evaluated or calculated based on a set of all possible actions in the current vehicle state. For example, the predetermined probability can be a proportion of a mean probability for all actions in the set, such as 50% of the mean probability. In this way, the relative improbability of a particular action depending on a number and range of probabilities for actions in the set is better assessed.

For example, with reference to Figure 6, all actions in state S_1 include actions A_0 to action A_7 having probabilities 0.1, 0.1, 0.2, 0.2, 0.05, 0.01, 0.14 and 0.1 respectively. Thus, the mean probability for all actions in state S1 is 1/7 or approximately 0.14. Taking an example threshold probability as a proportion of 50% of the mean probability, the threshold probability is thus 50% of 1/7 which is 1/14 (approximately 0.07). Thus, it can be seen that the probabilities of actions A_4 and A_5 fall below the threshold probability and are therefore deemed by the monitoring system 212 as sufficiently improbable that they constitute anomalous actions.

An anomalous action is an improbable action (determined based on a predetermined threshold probability as described above) indicating that a vehicle state transition is likely anomalous because it is not determined to be sufficiently probable based on the MDP model 210. Thus, in this way, the monitoring system 212 can identify anomalous vehicle state transitions, such identification being suitable for triggering responsive and/or reactive measures.

In one embodiment, the identification of an anomalous vehicle state transition triggers a process of precluding, preventing the vehicle state transition or of reverting, reversing or returning the vehicle to its previous state. This response to the anomalous identification is practical in certain circumstances where the new vehicle state arising from the anomalous transition can be readily reversed, such as where an actuator is modified (e.g. activated, deactivated, adjusted) – the modification can be undone. However, preventing or precluding the anomalous vehicle state transition may not be possible in all cases, such as where the new vehicle state represents a state that cannot be returned to its immediately preceding state – such as the vehicle being steered into imminent danger.

Thus, in some embodiments, the identification of an anomalous vehicle state transition can trigger a further state transition to a reactive or responsive state such as a safety state in which the vehicle is operated to mitigate the effects of the anomalous state transition. For example, a state or series of states constituting an evasive manoeuvre may be employed, or the vehicle may be stopped by rapid deceleration or the like.

The nature of the response to the identification of an anomalous vehicle state transition can differ depending on a determined seriousness of the transition. In some cases, the response may merely involve generating an alert or indication to the vehicle operator or other road users, or registering the anomaly in a vehicle log for later investigation.

In one embodiment, the modelling system 208 is arranged to generate potentially multiple MDP models 210 such as models aligned to different subsystems or sets of subsystems of the vehicle. For example, MDP models can include: a model relating to the drive, engine and transmission system; a model relating to driver indications (signs, signals, displays, dials); a model relating to comfort and convenience (seating, heating, ventilation, air conditioning); and a model relating to vehicle safety (parking sensors, lane departure, adaptive cruise control, emergency braking). In such an embodiment, the vehicle can be in one vehicle state within each model at a particular point in time and the monitoring system 212 is configured to monitor the vehicle state in each MDP model to identify anomalous state transitions across all models.

Returning to Figure 2, the monitoring system 212 is also optionally operable to receive information from, and act responsive to, external components such as an environmental alerts (EA) component 222 and a central intelligence system (CIS) component 224. For example, the EA component 222 can communicate information to the vehicle regarding the environment surrounding the vehicle such as accident, incident, congestion or other information. Such information can be used to define states as part of the MDP model 210 such that states and probabilities of actions differ according to information received from the EA component 222. Similarly, the CIS component 224 can be a centralised component for vehicles providing information regarding global threats, known intrusion symptoms, driving behaviour that is known to be problematic etc. Such information received from the EA component 222 can additionally be used to interpret the MDP model 210 appropriately. In some embodiments, the EA 222 and CIS 224 data is additionally or alternatively used by the modelling system 208 to define a more sophisticated MDP model 210 for the vehicle.

Figure 4 is a flowchart of a method operable with a communications network 200 in a vehicle to generate an indication that a vehicle state transition is anomalous in accordance with embodiments of the present invention. At step 402 an MDP model 210 is defined for the vehicle, the model specifying states of the vehicle and actions constituting transitions between states. A state of the vehicle is indicated by information provided by one or more sensors and a state of one or more actuators, and an action corresponds to a change in the information provided by one or more sensors and/or a change to a state of one or more actuators. Each action has associated a probability of occurrence. At step 404 the method

determines a current state of the vehicle in the model by accessing data communicated via the network. At step 406 the method accessing data communicated via the network so that, at step 408, the method can determine if a state changing action has occurred. Such a state changing action can be identified by the indication of the vehicle entering a new state. At step 410, where a state changing action is determined to have occurred, the method determines a probability associated with the action from the MDP Model 210. At step 412 the method determines if the probability is below a threshold probability, and where it is below, the method generates an anomaly indication at step 414.

Figure 5 is a component diagram illustrating the operation of the modelling system 208 in accordance with embodiments of the present invention. Many of the features of figure 5 are identical to those described above with respect to Figure 2 and these will not be repeated here. Figure 5 illustrates one embodiment for the generation of the MDP model 210 of Figure 2 and, in particular, the definition of the set of states and actions of the MDP model 210. The modelling system 208 of Figure 5 receives vehicle state and action definitions 408 based on two data sources. A first data source is a manufacturer defined set of states and actions 404 that can be defined to represent the possible states and actions for a vehicle according to the manufacturer. Further, a second data source is information received from an empirically learned set of states and actions 410. The empirically learned information 410 can arise from use or operation of the vehicle in the training mode of operation and/or use of the vehicle in a production (in use) mode of operation. Both data sources constituting promising bases for the definition of vehicle states and actions 408 on which basis the modelling system 208 determines action probabilities as hereinbefore described.

Insofar as embodiments of the invention described are implementable, at least in part, using a software-controlled programmable processing device, such as a microprocessor, digital signal processor or other processing device, data processing apparatus or system, it will be appreciated that a computer program for configuring a programmable device, apparatus or system to implement the foregoing described methods is envisaged as an aspect of the present invention. The computer program may be embodied as source code or undergo compilation for implementation on a processing device, apparatus or system or may be embodied as object code, for example.

Suitably, the computer program is stored on a carrier medium in machine or device readable form, for example in solid-state memory, magnetic memory such as disk or tape, optically or magneto-optically readable memory such as compact disk or digital versatile disk etc., and the processing device utilises the program or a part thereof to configure it for operation. The computer program may be supplied from a remote source embodied in a

communications medium such as an electronic signal, radio frequency carrier wave or optical carrier wave. Such carrier media are also envisaged as aspects of the present invention. It will be understood by those skilled in the art that, although the present invention has been described in relation to the above described example embodiments, the invention is not limited thereto and that there are many possible variations and modifications which fall within the scope of the invention. The scope of the present invention includes any novel features or combination of features disclosed herein. The applicant hereby gives notice that new claims may be formulated to such features or combination of features during prosecution of this application or of any such further applications derived therefrom. In particular, with reference to the appended claims, features from dependent claims may be combined with those of the independent claims and features from respective independent claims may be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

CLAIMS

1. A computer implemented security method operable with a communications network in a vehicle, the network communicatively connecting devices including sensors and actuators in the vehicle such that information provided by sensors and states of actuators are determinable by data communicated via the network, the method comprising the steps of:

defining a Markov decision process model for the vehicle, the model specifying states of the vehicle and actions constituting transitions between states, wherein a state of the vehicle is indicated by information provided by one or more sensors and a state of one or more actuators, and an action corresponds to a change in the information provided by one or more sensors and/or a change to a state of one or more actuators, each action having associated a probability of occurrence;

determining, by accessing data communicated via the network, a current state of the vehicle in the model;

accessing data communicated via the network;

responsive to the accessed data indicating an action to change the vehicle state to a new state, determining, from the model, a probability of the action;

responsive to a determination that the determined probability falls below a predetermined threshold probability, generating an indication that the vehicle state transition is anomalous.

2. The method of claim 1 wherein the predetermined threshold probability is based on a set of all possible actions in the current vehicle state defined in the model and probabilities associated with each action in the set such that the predetermined threshold probability serves to identify relatively improbable actions.

3. The method of claim 2 wherein the predetermined threshold probability is a proportion of a mean probability for all actions in the set.

4. The method of any preceding claim wherein the Markov decision process model is defined empirically based on observation of vehicle states and actions in acceptable use.

5. The method of any or claims 1 to 3 wherein the Markov decision process model is defined by a reinforcement machine learning method based on observation of vehicle states and actions in acceptable use.

6. The method of claim 5 wherein the reinforcement machine learning method is a Q-learning reinforcement method based on a predetermined definition of vehicle states and actions and using a reward function wherein actions leading to desirable and/or acceptable vehicle states are attributed relatively greater reward so as to inform the Q-learning method to assign relatively greater probabilities to actions with relatively greater rewards.

7. The method of any preceding claim further comprising, responsive to the anomalous indication, precluding the transition of the vehicle to the new state by interrupting the action.

8. The method of claim 7 wherein the action involves a change in state of at least one actuator, and interrupting the action includes preventing the action by preventing a change of state of the at least one actuator.

9. The method of claim 7 or 8 wherein the action arises from a change in the information provided by one or more sensors, and interrupting the action includes forcing a transition of the vehicle state to a third state as a reactive state.

10. A computer system including a processor and memory storing computer program code for performing the steps of any preceding claim.

11. A computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of a method as claimed in any of claims 1 to 9.

FIGURE 1

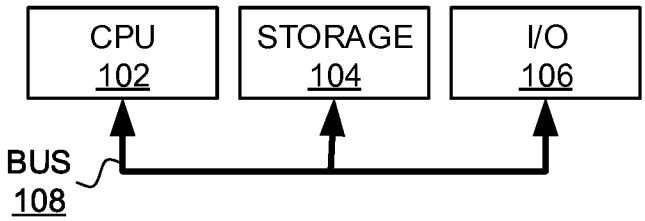

FIGURE 2

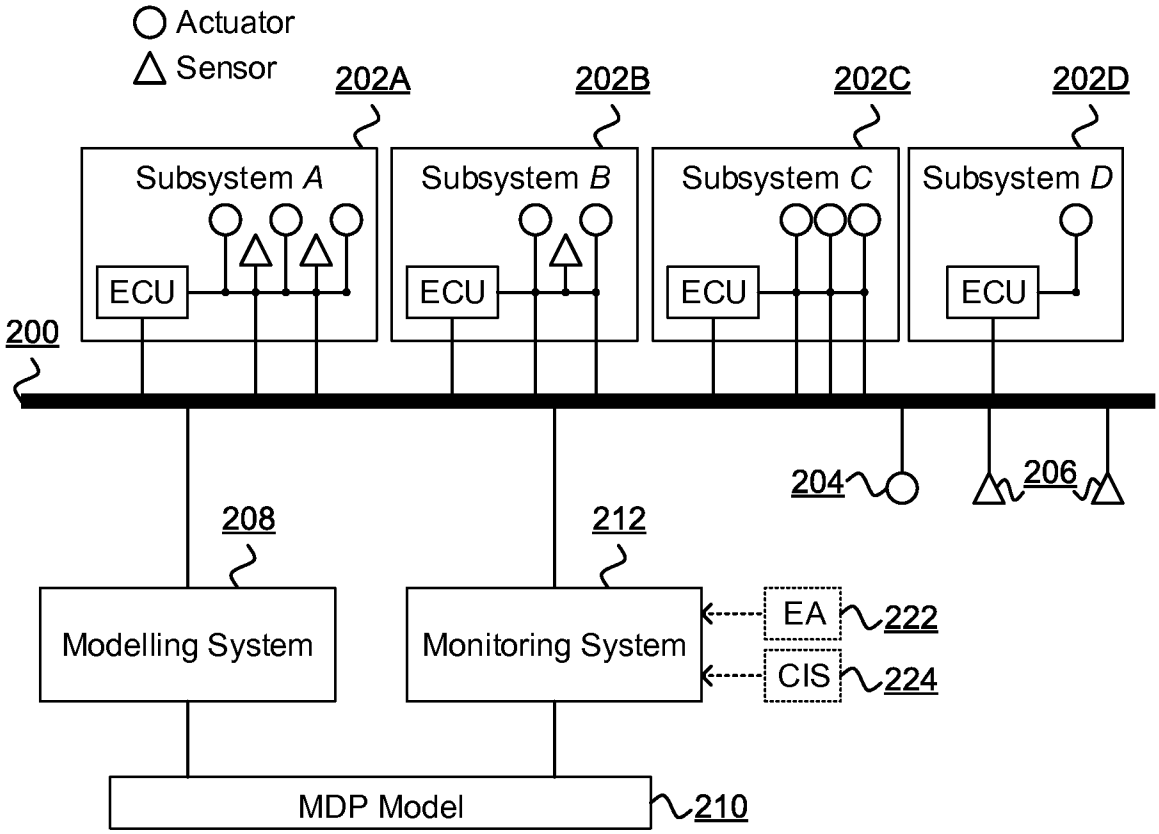

FIGURE 3

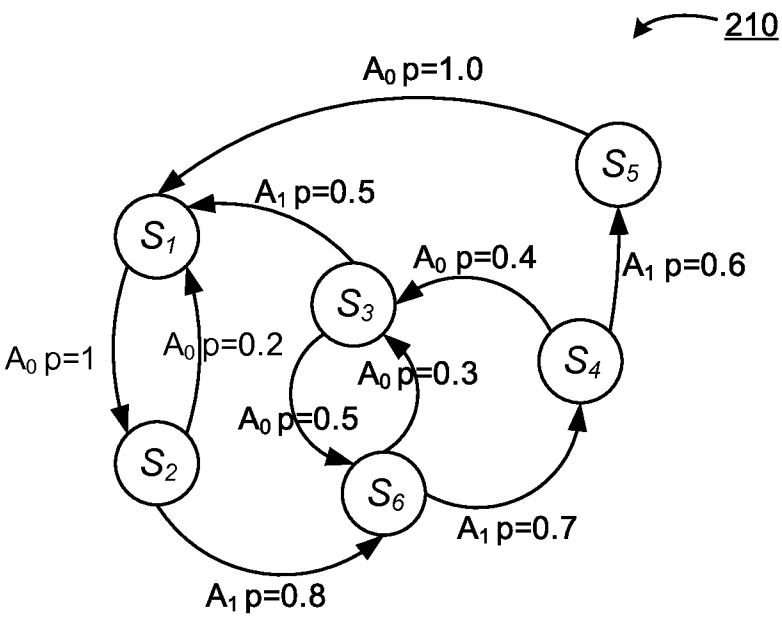

FIGURE 5

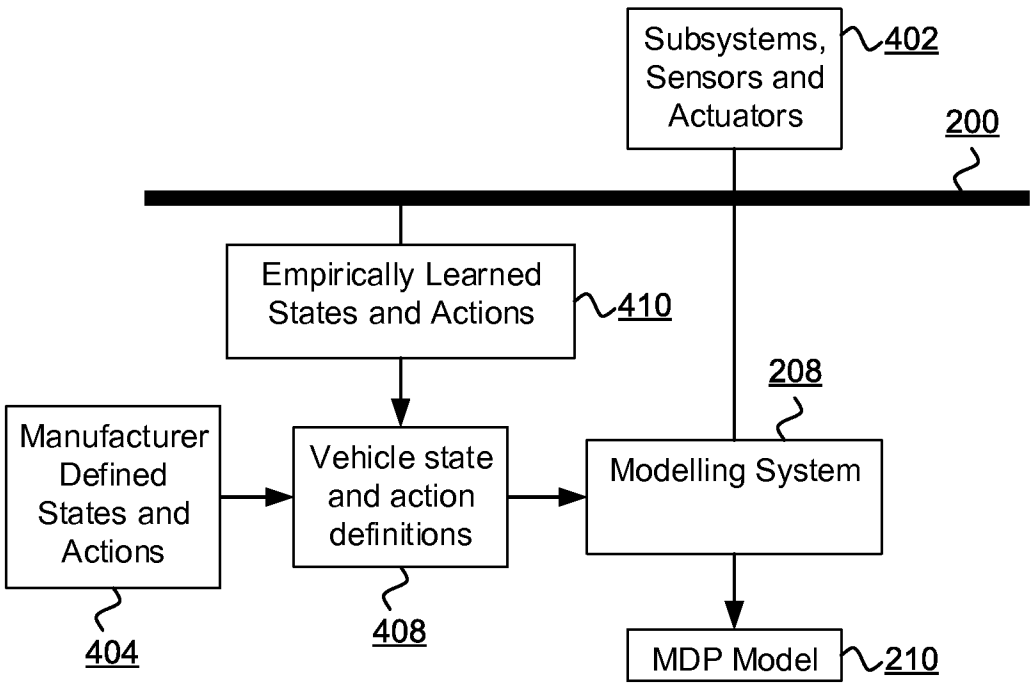

FIGURE 4

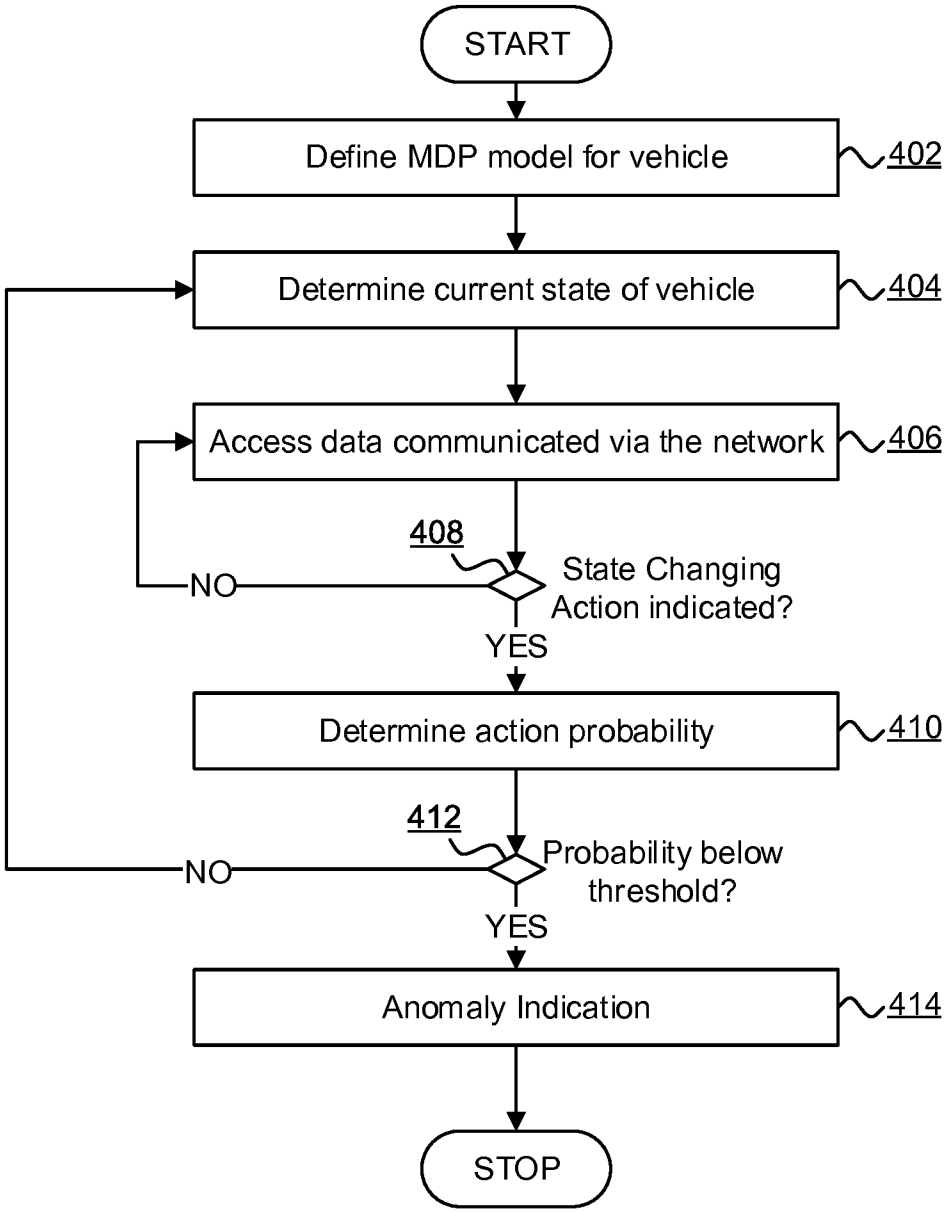
START
Define MDP model for vehicle
402
Determine current state of vehicle
404
Access data communicated via the network
406
408
State Changing Action indicated?
NO
YES
Determine action probability
410
412
Probability below threshold?
NO
YES
Anomaly Indication
414
STOP

FIGURE 6

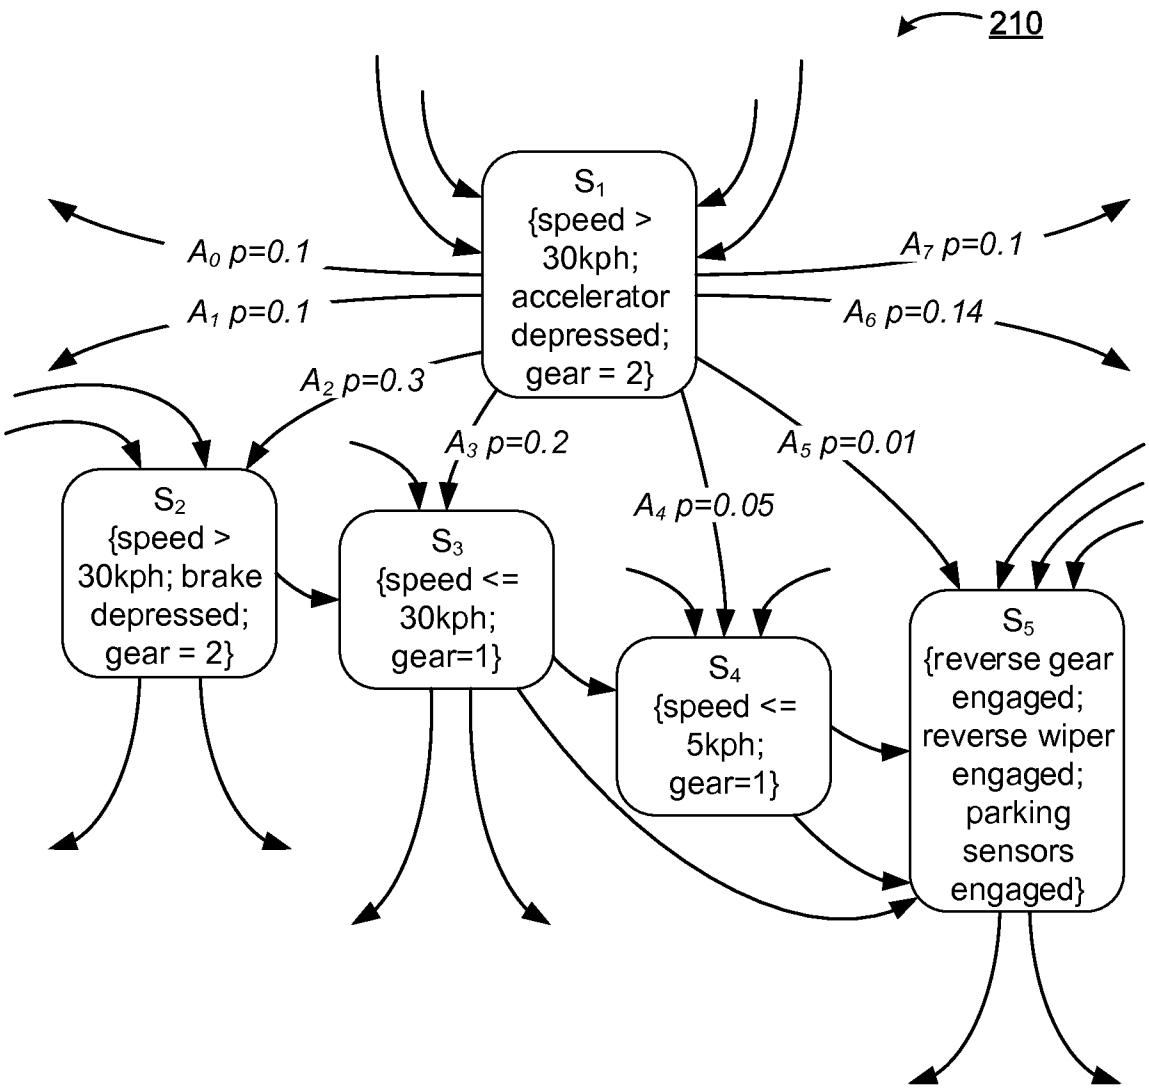
210
S1 {speed > 30kph; accelerator depressed; gear = 2}
A0 p=0.1
A1 p=0.1
A2 p=0.3
A3 p=0.2
A4 p=0.05
A5 p=0.01
A6 p=0.14
A7 p=0.1
S2 {speed > 30kph; brake depressed; gear = 2}
S3 {speed <= 30kph; gear=1}
S4 {speed <= 5kph; gear=1}
S5 {reverse gear engaged; reverse wiper engaged; parking sensors engaged}

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2019/055341

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/08 H04L29/06 H04L12/40
ADD. G06F21/55 G06N99/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L G06F G06N H04W

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	NARAYANAN SANDEEP NAIR ET AL: "OBD_SecureAlert: An Anomaly Detection System for Vehicles", 2016 IEEE INTERNATIONAL CONFERENCE ON SMART COMPUTING (SMARTCOMP), IEEE, 18 May 2016 (2016-05-18), pages 1-6, XP032917406, DOI: 10.1109/SMARTCOMP.2016.7501710 abstract page 1, left-hand column, lines 9-25 page 1, right-hand column, lines 14-33 page 3, left-hand column, lines 22-36 page 3, right-hand column, lines 27-42 page 4, left-hand column, lines 1-10,29-34 page 4, right-hand column, lines 1-41 ----- -/--	1-11

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
14 March 2019	22/03/2019

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Yamajako-Anzala, A

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2019/055341

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2017/279834 A1 (VASSEUR JEAN-PHILIPPE [US] ET AL) 28 September 2017 (2017-09-28) paragraphs [0002], [0031], [0063]	1-11
A	Anonymous: "Q-learning Machine learning and data mining Machine learning portal", , 22 February 2018 (2018-02-22), XP055498973, Retrieved from the Internet: URL:https://en.wikipedia.org/w/index.php?title=Q-learning&oldid=826975938 [retrieved on 2018-08-10] the whole document	1-11

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2019/055341

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2017279834 A1	28-09-2017	US 2017279834 A1	28-09-2017
		US 2017279835 A1	28-09-2017