(12) UK Patent Application (19) GB (11) 2535718 (13) A

(43) Date of A Publication 31.08.2016

(21) Application No: 1503082.8

(22) Date of Filing: 24.02.2015

(51) INT CL:
G06Q 10/04 (2012.01) ***G06Q 50/30*** (2012.01)

(56) Documents Cited:
None

(58) Field of Search:
Other: **No search performed: Section 17(5)(b)**

(71) Applicant(s):
Addison Lee Limited
35-37 William Road, LONDON, NW1 3ER,
United Kingdom

(72) Inventor(s):
Larry Scicluna
Andrey Glaschenko

(74) Agent and/or Address for Service:
Venner Shipley LLP
200 Aldersgate, LONDON, EC1A 4HD,
United Kingdom

(54) Title of the Invention: **Resource management**
Abstract Title: **Computer-implemented method and apparatus for allocating a vehicle resource to a future vehicle requirement**

(57) A method and apparatus for allocating a vehicle resource to a customer's future vehicle requirement are disclosed. The method comprises storing vehicle requirements, which include at least a start time and a pickup location. The required vehicle size and type may also be stored. If, in respect of a booking, the current time has a predetermined relationship with respect to the start time of the booking, a pool of candidate vehicles, which comprises fewer vehicles than the number of vehicles that are available, is formed for possible fulfillment of the booking. For each vehicle in the pool of candidate vehicles, a score is calculated that is related to the suitability of the vehicle to the booking. For one or more of the vehicles in the pool of candidate vehicles, a journey time from a first location to the pickup location included in the booking is calculated. The vehicle score and the calculated journey time are used to determine whether there is a need to allocate a booking. The steps of calculating the vehicle score and the journey time and determining whether there is a need to allocate the booking are repeated until there is a need to allocate the booking, at which point the booking is allocated to the vehicle with the best suitability score. The apparatus is configured for implementing the method.

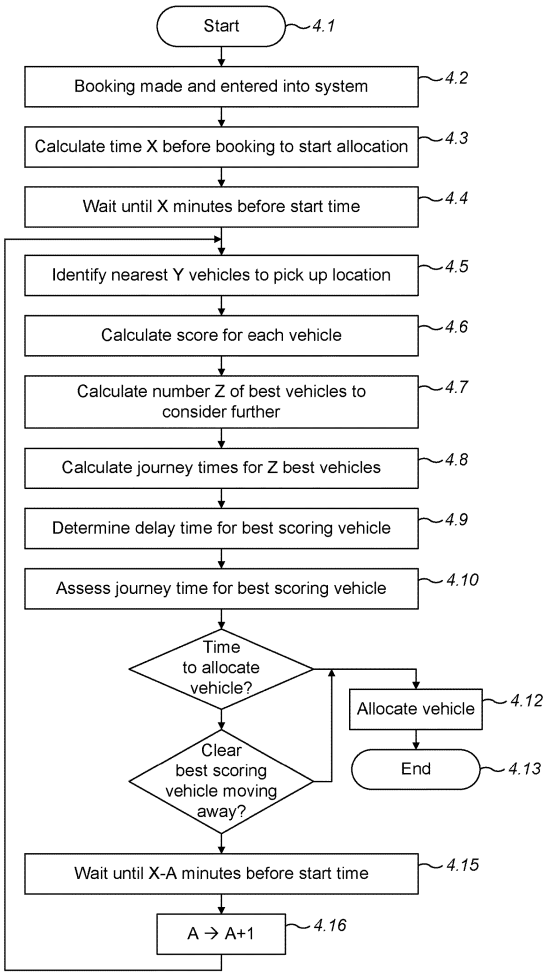

FIG. 4

GB 2535718 A

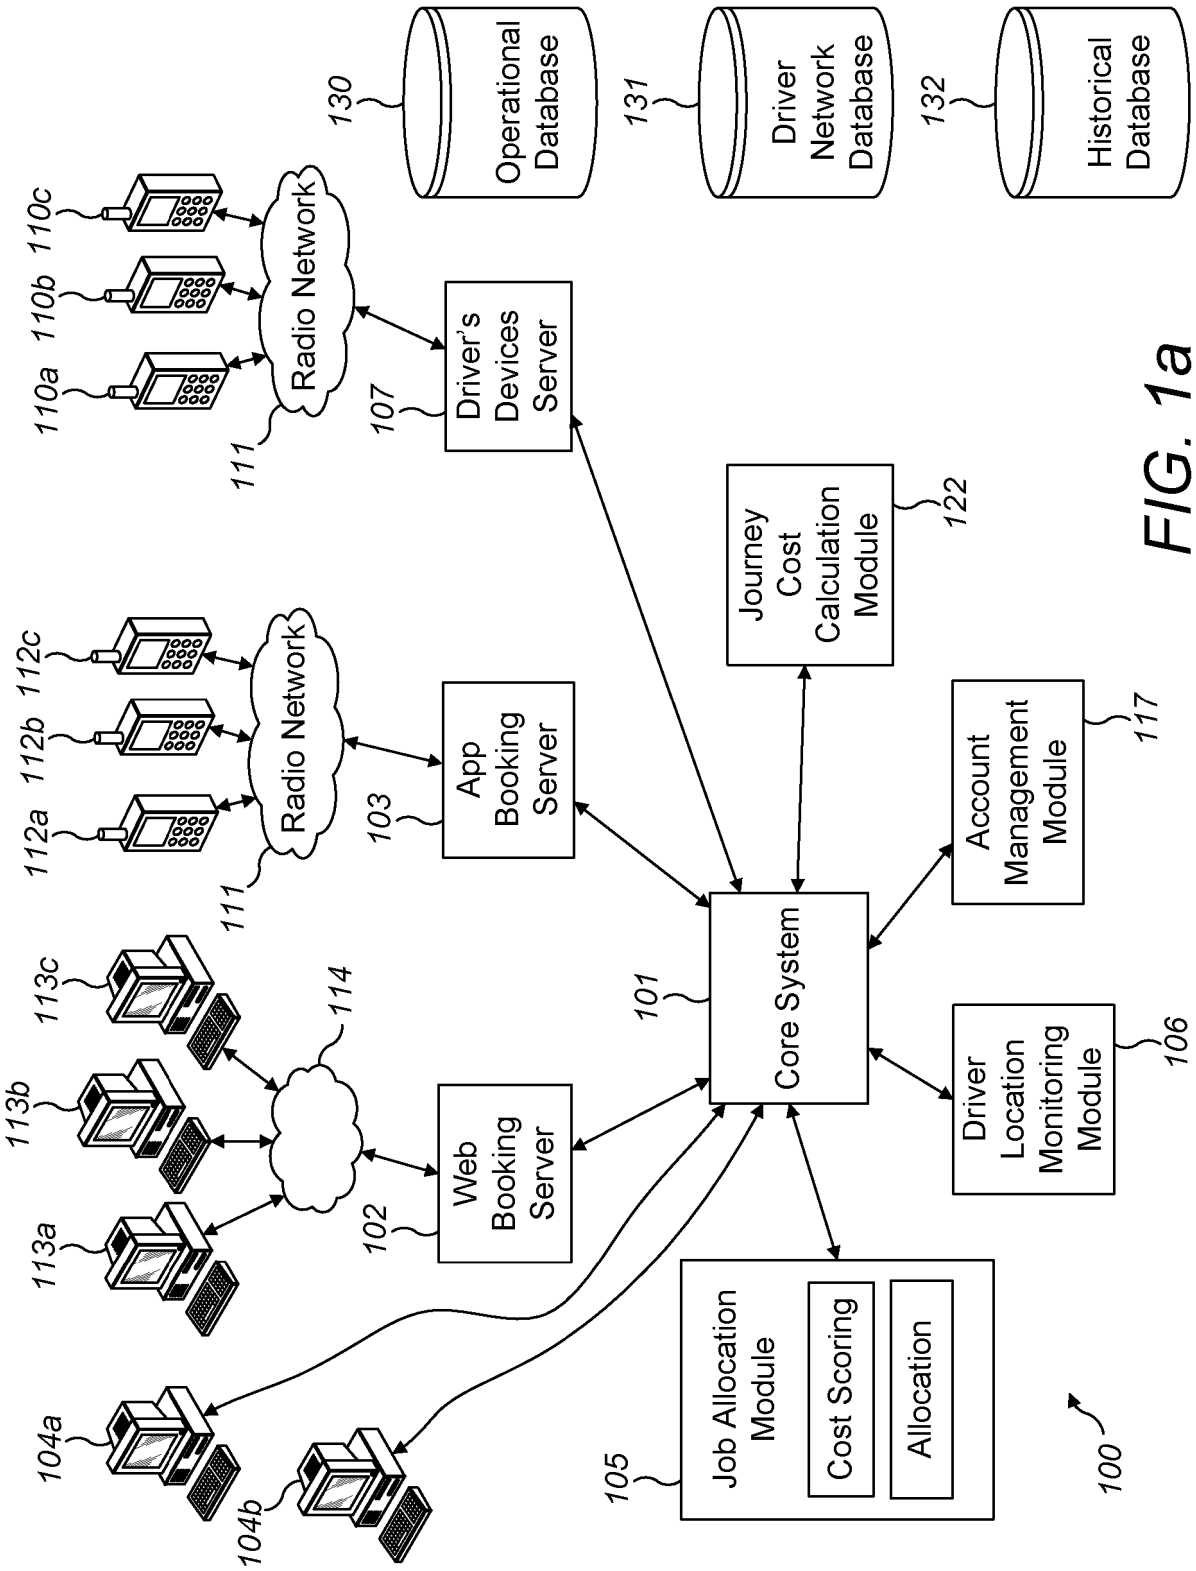
Operational Database
130
Driver Network Database
131
Historical Database
132
Radio Network
111
Driver's Devices Server
107
Journey Cost Calculation Module
122
App Booking Server
103
Account Management Module
117
Core System
101
Web Booking Server
102
Driver Location Monitoring Module
106
Job Allocation Module
Cost Scoring
Allocation
105
100
110a
110b
110c
112a
112b
112c
113a
113b
113c
104a
104b
114

FIG. 1a

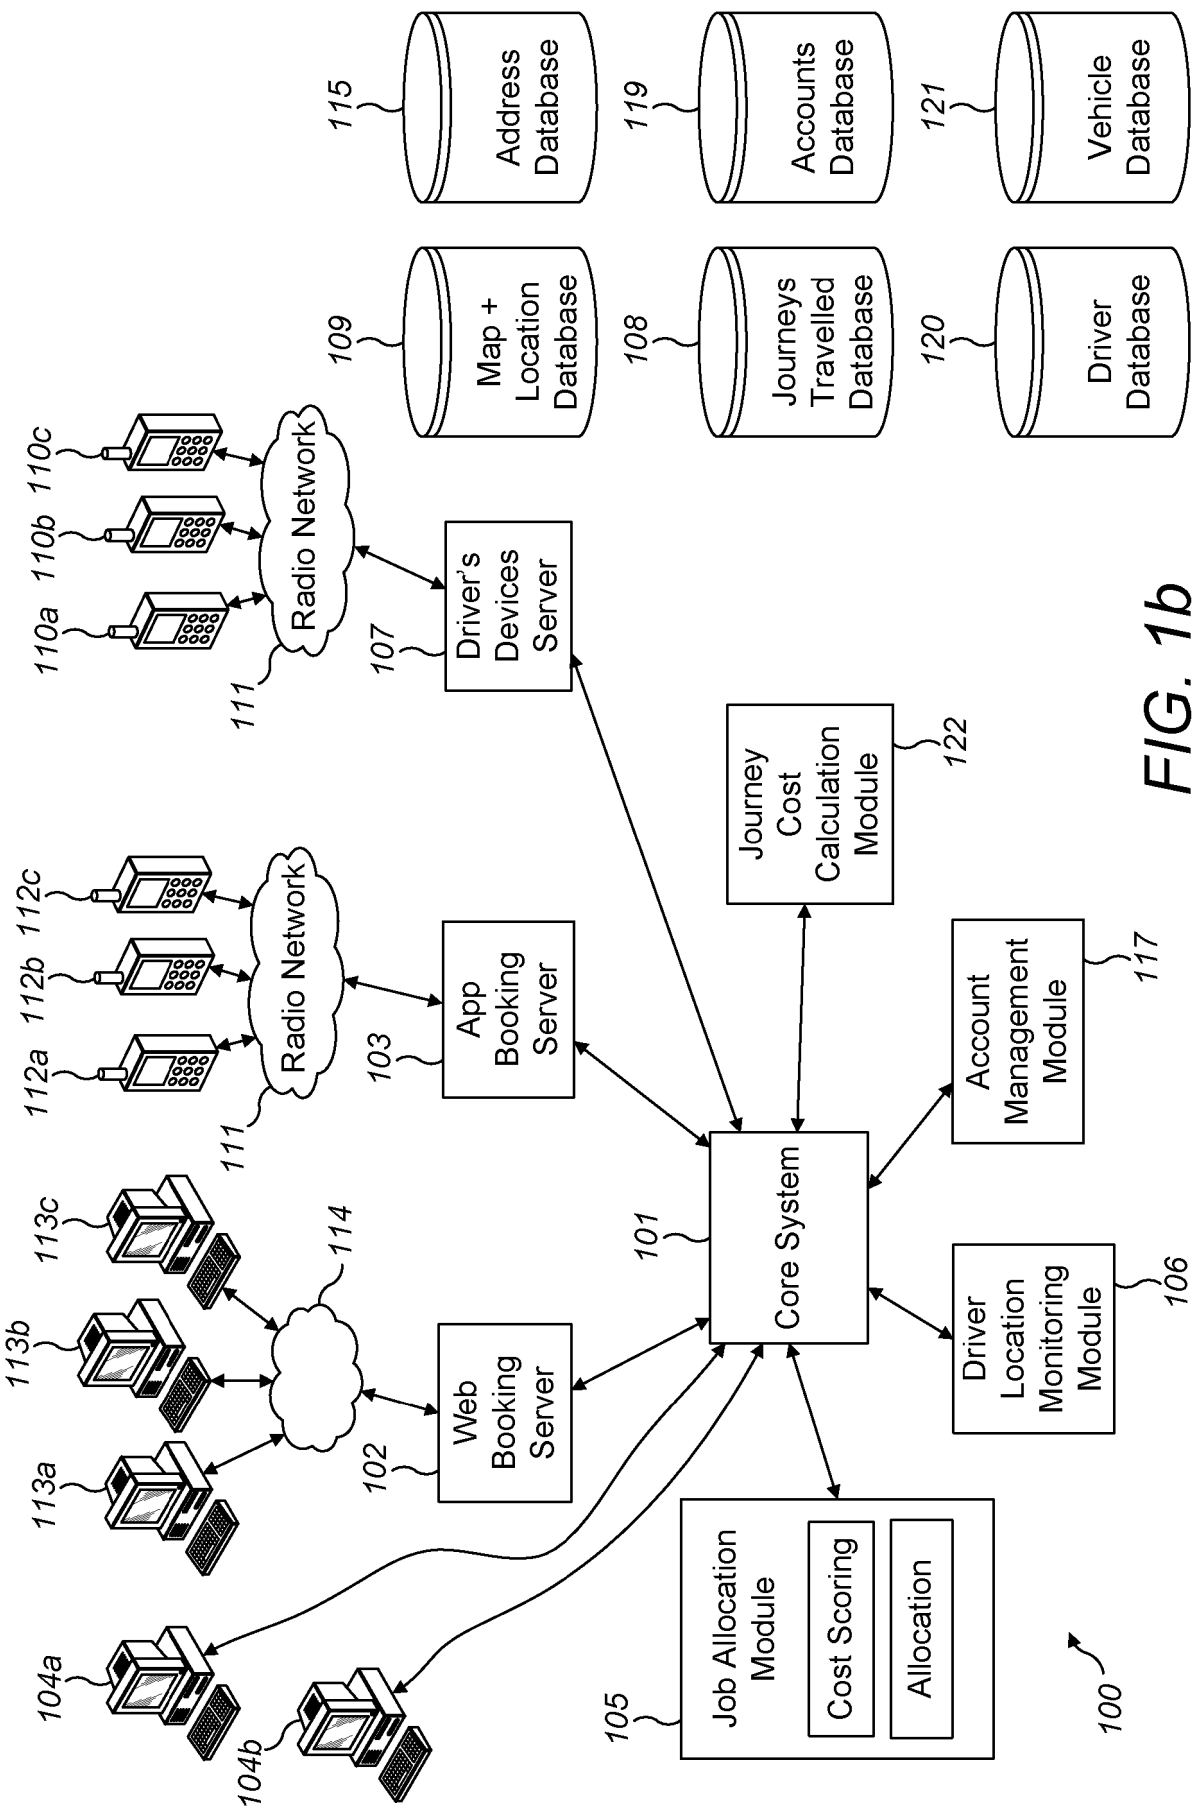
Address Database
115
Accounts Database
119
Vehicle Database
121
Map + Location Database
109
Journeys Travelled Database
108
Driver Database
120
Radio Network
111
110a
110b
110c
Driver's Devices Server
107
Journey Cost Calculation Module
122
Radio Network
111
112a
112b
112c
App Booking Server
103
Account Management Module
117
113c
113b
113a
114
Core System
101
Web Booking Server
102
Driver Location Monitoring Module
106
104a
104b
Job Allocation Module
Cost Scoring
Allocation
105
100

FIG. 1b

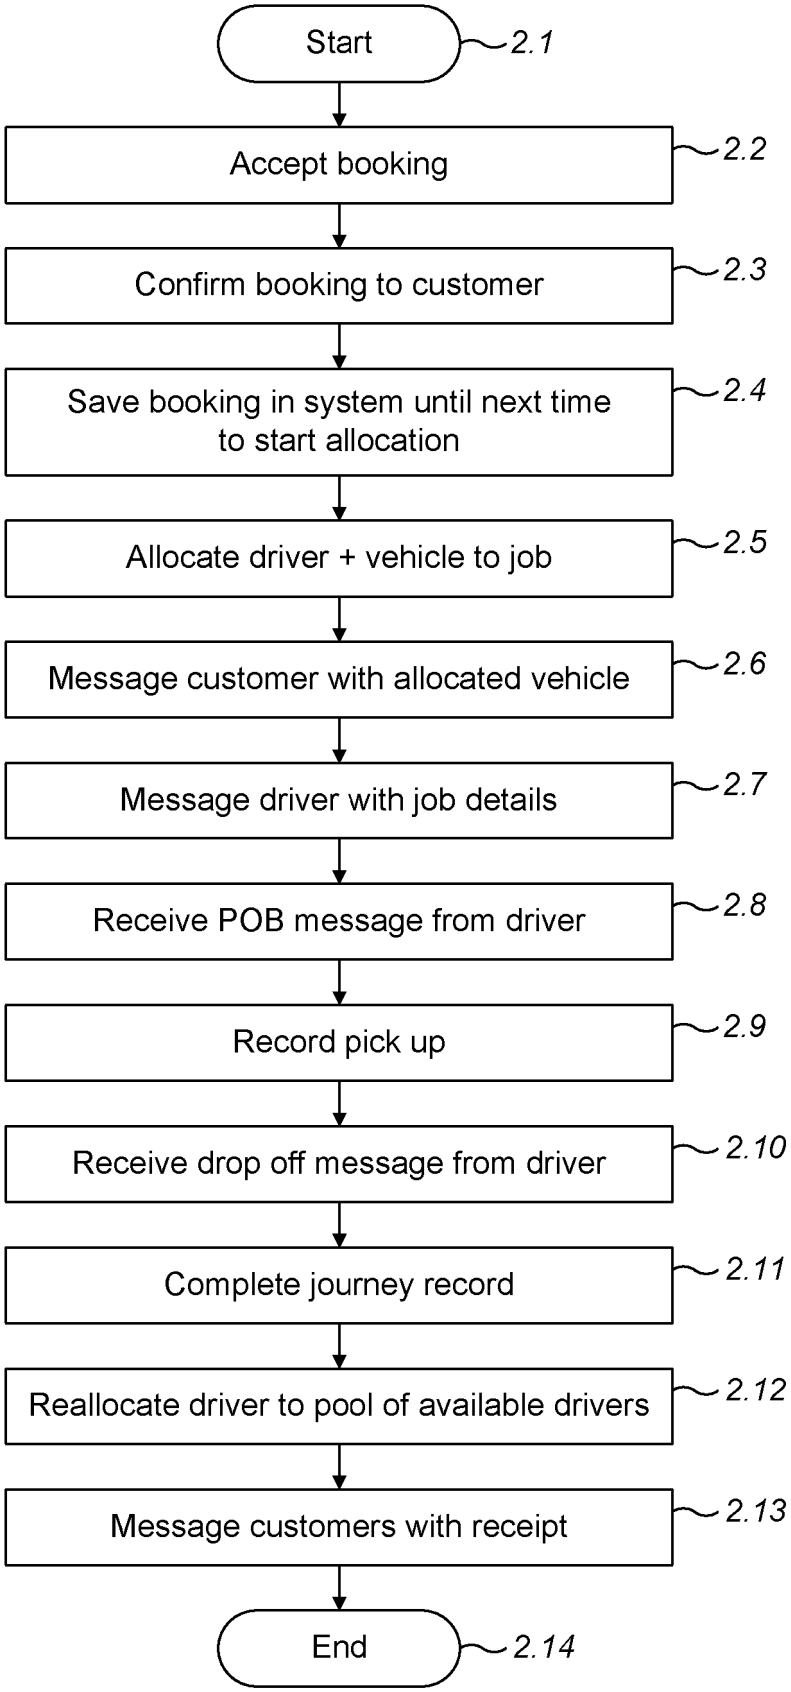
Start
2.1
Accept booking
2.2
Confirm booking to customer
2.3
Save booking in system until next time to start allocation
2.4
Allocate driver + vehicle to job
2.5
Message customer with allocated vehicle
2.6
Message driver with job details
2.7
Receive POB message from driver
2.8
Record pick up
2.9
Receive drop off message from driver
2.10
Complete journey record
2.11
Reallocate driver to pool of available drivers
2.12
Message customers with receipt
2.13
End
2.14

FIG. 2

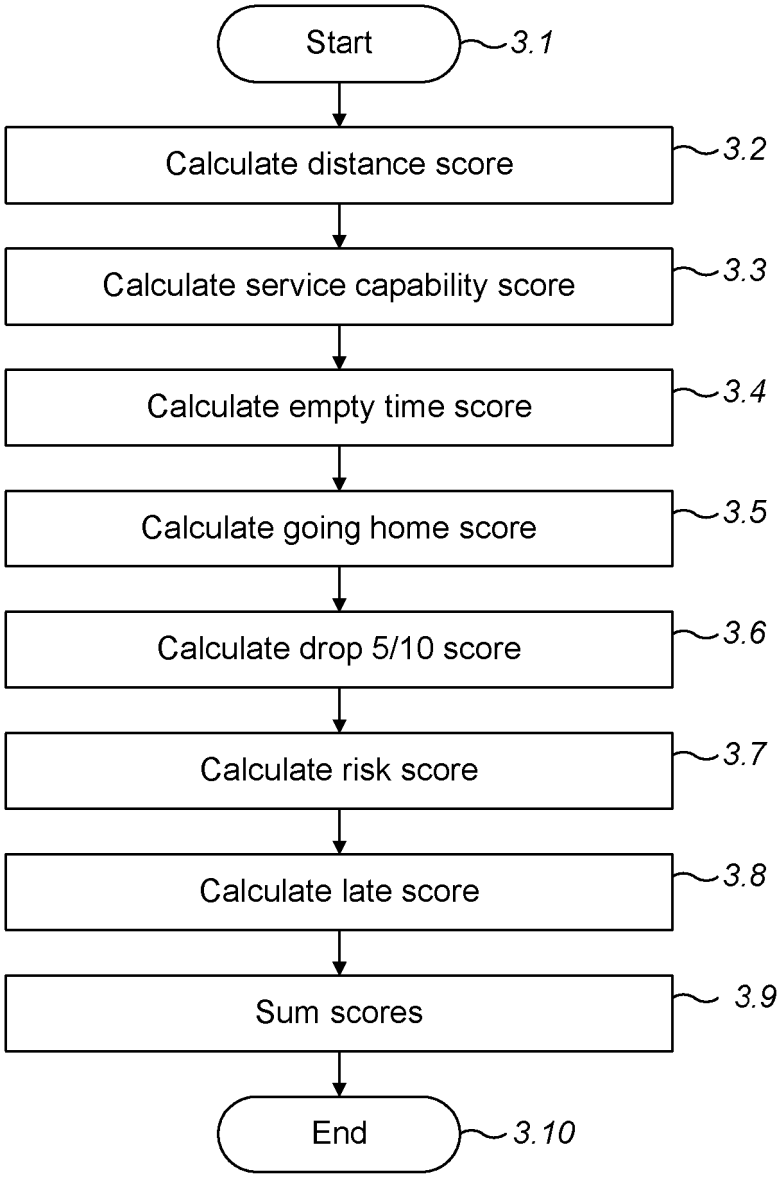
Start
3.1
Calculate distance score
3.2
Calculate service capability score
3.3
Calculate empty time score
3.4
Calculate going home score
3.5
Calculate drop 5/10 score
3.6
Calculate risk score
3.7
Calculate late score
3.8
Sum scores
3.9
End
3.10

FIG. 3

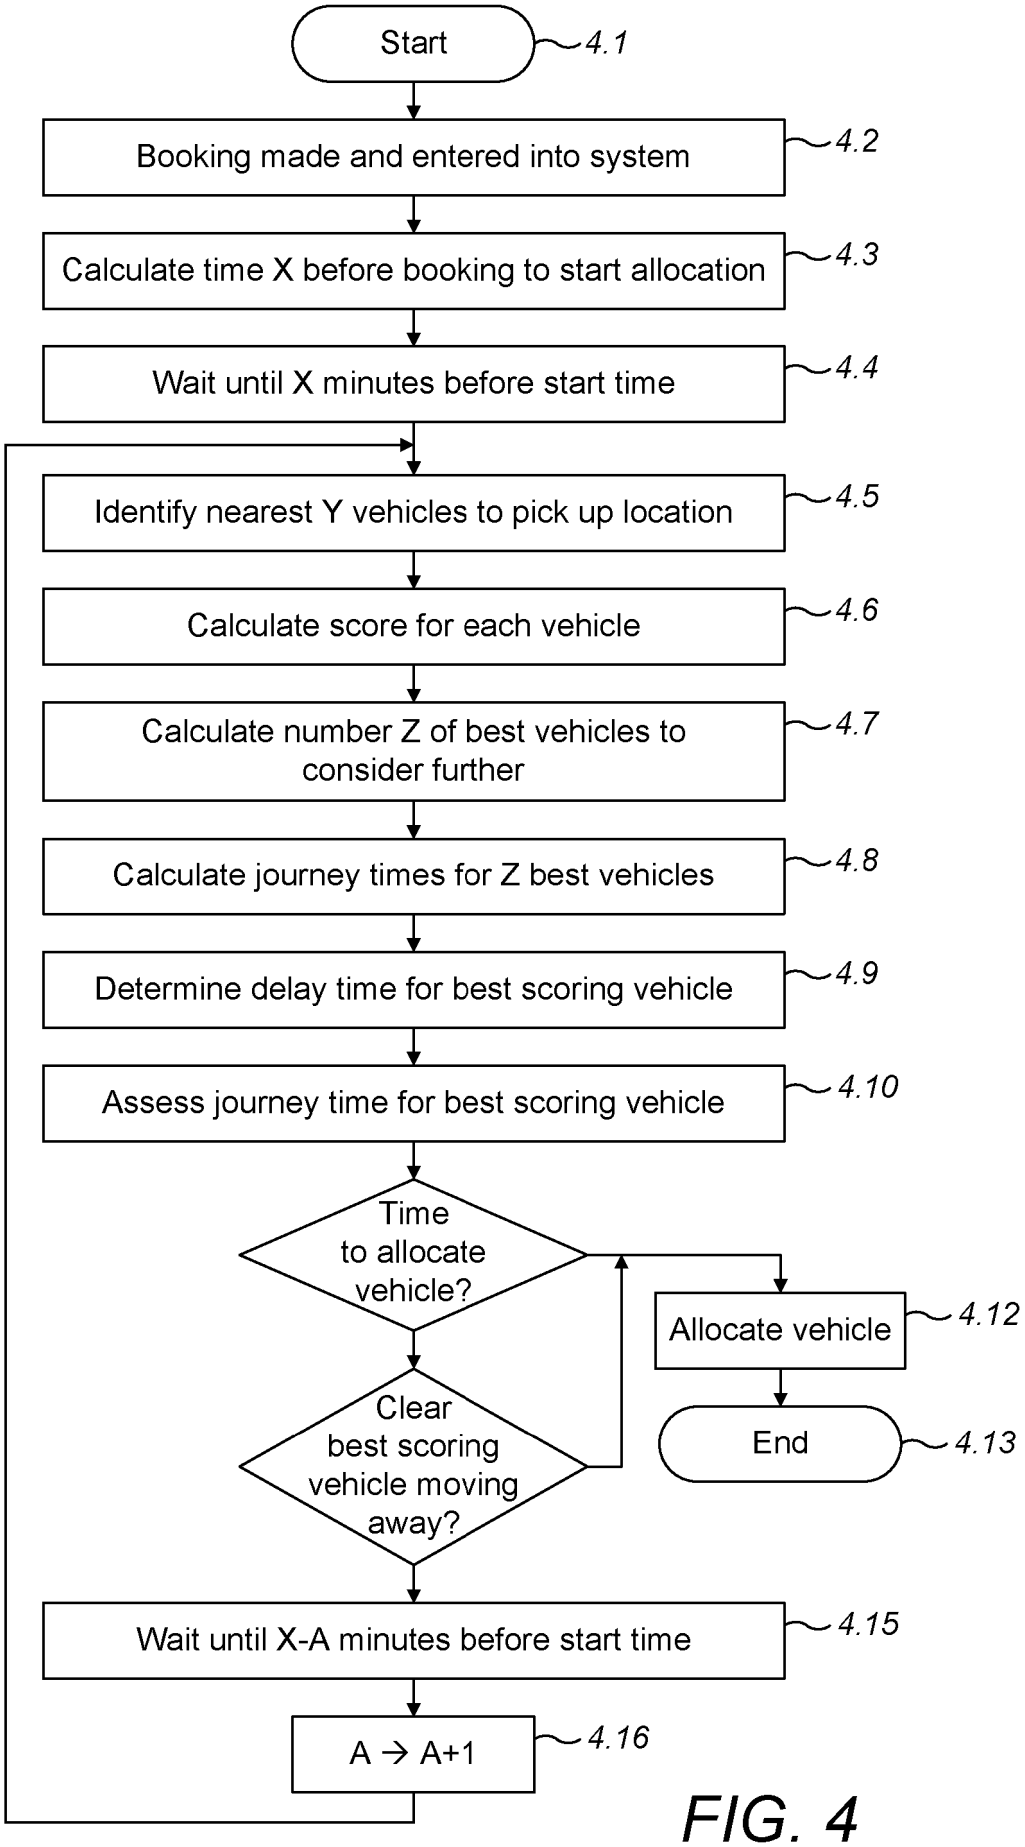
Start
4.1
Booking made and entered into system
4.2
Calculate time X before booking to start allocation
4.3
Wait until X minutes before start time
4.4
Identify nearest Y vehicles to pick up location
4.5
Calculate score for each vehicle
4.6
Calculate number Z of best vehicles to consider further
4.7
Calculate journey times for Z best vehicles
4.8
Determine delay time for best scoring vehicle
4.9
Assess journey time for best scoring vehicle
4.10
Time to allocate vehicle?
Allocate vehicle
4.12
Clear best scoring vehicle moving away?
End
4.13
Wait until X-A minutes before start time
4.15
A → A+1
4.16

FIG. 4

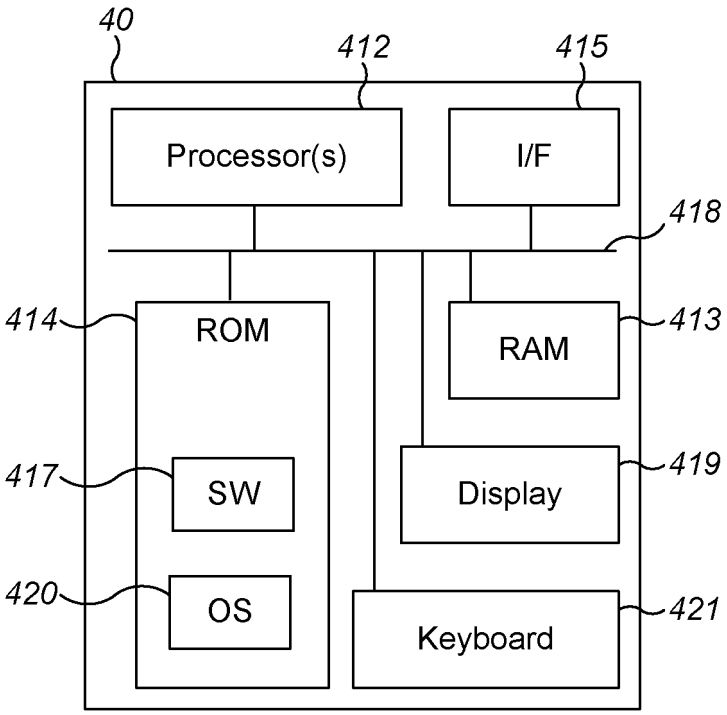
40
412
415
Processor(s)
I/F
418
414
ROM
RAM
413
417
SW
Display
419
420
OS
Keyboard
421

FIG. 5

Resource Management

Field of the Invention

The present invention relates to the management of vehicle resources in relation to vehicle requirements.

Background to the Invention

In the field of private hire vehicles and taxis, historically vehicles and drivers were allocated to bookings by a human operator, called a controller. The controller would be in contact with the drivers of vehicles of a fleet by voice channels (typically one shared voice channel) of a radio system, through which the controller and the drivers could speak with one another. The controller could ascertain the locations of drivers by requesting the drivers to provide their locations verbally. The controller could also enquire of drivers when they expected to drop passengers and thus become free again for fulfilling a booking. Drivers would be provided with details of bookings allocated to them by the controller verbally through the radio system. Such manual systems are still widely in use.

Addison Lee has for some time provided automated allocation of advance bookings in London, UK to drivers, although the way in which this is performed has not been visible to customers. Some third parties provide automatic allocation of instant bookings but do not provide automatic allocation of advanced bookings. The automatic allocation of advance bookings to drivers is a very different process to the manual process that was used previously and it introduces many challenges. Improved allocation can provide positive effects in a number of areas including the processing, memory and other computing and communication resources that are required. Other effects can include improved vehicle fleet utilisation and lower distances travelled in fulfilling customer bookings. Further advantages can be felt in terms of customer satisfaction and driver wellbeing and happiness. Techniques used in automatic allocation of instant bookings are generally not applicable to automatic allocation of advance bookings.

Summary of the Invention

A first aspect of the invention provides a computer-implemented method of allocating a vehicle resource to a future vehicle requirement, the method comprising:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles may comprise fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).

Step h) may comprise allocating the vehicle determined to have had the best score on the last performance of step d) to the first booking.

The method may comprise performing step e) for some but not all of the vehicles in the pool of vehicles formed at step c).

Calculating a journey time from the first location may comprise calculating a journey time from a current location of the vehicle to the pickup location included in the first booking, or it may comprise calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking.

The method may comprise choosing between:

1) calculating a journey time from the first location by calculating a journey time from a current location of the vehicle to the pickup location included in the first booking, and

2) calculating a journey time from the first location by calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking, based on a status of the vehicle.

The determining that a current time has a predetermined relationship with respect to the start time of the first booking may comprise determining that the current time is at least approximately a predetermined time before the start time of the first booking, for instance is between 15 and 40 minutes before the start time of the first booking. The predetermined time may be dependent on the pickup location for the first booking.

Forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise rejecting vehicles of the fleet that have a status that is inconsistent with the vehicle potentially being able to fulfil the booking.

Forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking.

Forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking.

Forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking in terms of the direct distance between the candidate vehicles of the fleet and the pickup location included in the first booking.

The predetermined number may be between 10 and 100.

Step g) may comprise performing steps c), d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking

Calculating a score that is related to the suitability of the vehicle to fulfil the booking may comprise calculating sub scores for each of plural factors and performing a mathematical operation on the sub scores.

Forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking and wherein calculating a journey time from the first location to the pickup location included in the first booking may comprise using a map database.

Calculating a journey time from the first location to the pickup location included in the first booking may comprise using a map database.

Calculating a journey time from the first location to the pickup location included in the first booking may comprise using historical or current data about traffic delays on one or more routes between the current location of the vehicle and the pickup location included in the first booking

Calculating a journey time from the first location to the pickup location included in the first booking may comprise using historical or current data about traffic delays or average speeds on one or more routes between the first location and the pickup location included in the first booking to calculate a fastest route between the first location and the pickup location included in the first booking

Determining whether there is a need to allocate the first booking may comprise determining whether a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

Determining whether there is a need to allocate the first booking may comprise determining whether the sum of an allocation buffer value and a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of

candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

The method may comprise:

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating the vehicle determined to have had the best score on last performance of step d) to the first booking.

The invention also provides apparatus for allocating a vehicle resource to a future vehicle requirement, the apparatus comprising means configured for:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles may comprise fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).

The means configured for performing step h) may comprise means configured for allocating the vehicle determined to have had the best score on the last performance of step d) to the first booking.

The apparatus may comprise means configured for performing step e) for some but not all of the vehicles in the pool of vehicles formed at step c).

The means configured for calculating a journey time from the first location may comprise means configured for calculating a journey time from a current location of the vehicle to the pickup location included in the first booking.

The means configured for calculating a journey time from the first location may comprise means configured for calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking. The apparatus may comprise means configured for choosing between:

1) calculating a journey time from the first location by calculating a journey time from a current location of the vehicle to the pickup location included in the first booking, and

2) calculating a journey time from the first location by calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking,

based on a status of the vehicle.

The means configured for determining that a current time has a predetermined relationship with respect to the start time of the first booking may comprise means configured for determining that the current time is at least approximately a predetermined time before the start time of the first booking, for instance is between 15 and 40 minutes before the start time of the first booking. The predetermined time may be dependent on the pickup location for the first booking.

The means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise means configured for rejecting vehicles of the fleet that have a status that is inconsistent with the vehicle potentially being able to fulfil the booking.

The means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise means configured for selecting a predetermined number of candidate vehicles that are geographically closest to the

pickup location included in the first booking. The means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise means configured for calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking. The means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise means configured for selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking in terms of the direct distance between the candidate vehicles of the fleet and the pickup location included in the first booking. The predetermined number may be between 10 and 100.

The means configured for performing step g) may comprise means configured for performing steps c), d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking

The means configured for calculating a score that is related to the suitability of the vehicle to fulfil the booking may comprise means configured for calculating sub scores for each of plural factors and performing a mathematical operation on the sub scores.

The means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking may comprise means configured for calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking and wherein the means configured for calculating a journey time from the first location to the pickup location included in the first booking may comprise means configured for using a map database.

The means configured for calculating a journey time from the first location to the pickup location included in the first booking may comprise means configured for using a map database.

The means configured for calculating a journey time from the first location to the pickup location included in the first booking may comprise means configured for using historical or current data about traffic delays on one or more routes between the current location of the vehicle and the pickup location included in the first booking

The means configured for calculating a journey time from the first location to the pickup location included in the first booking may comprise means configured for using

historical or current data about traffic delays or average speeds on one or more routes between the first location and the pickup location included in the first booking to calculate a fastest route between the first location and the pickup location included in the first booking

The means configured for determining whether there is a need to allocate the first booking may comprise means configured for determining whether a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

The means configured for determining whether there is a need to allocate the first booking may comprise means configured for determining whether the sum of an allocation buffer value and a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

The apparatus may comprise means configured for:

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating the vehicle determined to have had the best score on last performance of step d) to the first booking.

Another aspect of the invention provides a computer-implemented method of allocating a vehicle resource to a future vehicle requirement, the method comprising:

responding to detecting that a current time is a predetermined time ahead of a start time for a customer booking that constitutes the future vehicle requirement by identifying Y candidate vehicles that have potential to fulfil the booking,

scoring the Y candidate vehicles as to their suitability to fulfil the booking,

choosing Z of the Y candidate vehicles based on their scores, wherein the value of Z is lower than the value of Y,

calculating road journey times for the Z vehicles but not for all Y vehicles,

using the calculated road journey time for at least one vehicle to determine whether to allocate the vehicle, and

repeating the scoring and calculating steps until it is determined to allocate the vehicle.

Yet another aspect of the invention provides a computer-implemented method comprising:

calculating road journey times for one or more of Z vehicles each associated with a location that meets a predetermined criterion with respect to a pickup location for a private hire vehicle booking,

using the calculated road journey time for at least one vehicle to determine whether to allocate the vehicle, and

repeating the calculating and using steps until it is determined to allocate a vehicle to the booking.

Another aspect of the invention provides a computer-implemented method of allocating a vehicle resource to a future vehicle requirement, the method comprising:

identifying Y candidate vehicles that have potential to fulfil the future vehicle requirement,

scoring the Y candidate vehicles as to their suitability to fulfil the booking,

choosing Z of the Y candidate vehicles based on their scores, wherein the value of Z is lower than the value of Y,

calculating road journey times for the Z vehicles but not for all Y vehicles,

repeating the scoring and calculating steps at least once and then allocating a vehicle to fulfil the vehicle requirement.

Another aspect of the invention provides apparatus for allocating a vehicle resource to a future vehicle requirement, the apparatus comprising means for:

responding to detecting that a current time is a predetermined time ahead of a start time for a customer booking that constitutes the future vehicle requirement by identifying Y candidate vehicles that have potential to fulfil the booking,

scoring the Y candidate vehicles as to their suitability to fulfil the booking,

choosing Z of the Y candidate vehicles based on their scores, wherein the value of Z is lower than the value of Y,

calculating road journey times for the Z vehicles but not for all Y vehicles,

using the calculated road journey time for at least one vehicle to determine whether to allocate the vehicle, and

repeating the scoring and calculating steps until it is determined to allocate the vehicle.

Another aspect of the invention provides apparatus comprising means for:

calculating road journey times for one or more of Z vehicles each associated with a location that meets a predetermined criterion with respect to a pickup location for a private hire vehicle booking,

using the calculated road journey time for at least one vehicle to determine whether to allocate the vehicle, and

repeating the calculating and using steps until it is determined to allocate a vehicle to the booking.

Another aspect of the invention provides apparatus for allocating a vehicle resource to a future vehicle requirement, the apparatus comprising means for:

identifying Y candidate vehicles that have potential to fulfil the future vehicle requirement,

scoring the Y candidate vehicles as to their suitability to fulfil the booking,

choosing Z of the Y candidate vehicles based on their scores, wherein the value of Z is lower than the value of Y,

calculating road journey times for the Z vehicles but not for all Y vehicles,

repeating the scoring and calculating steps at least once and then allocating a vehicle to fulfil the vehicle requirement.

Brief Description of the Figures

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

Figure 1a is a schematic diagram of a system for management of a private hire vehicle service according to various aspects of the present invention;

Figure 1b is a schematic diagram of a different configuration of the Figure 1 system for management of a private hire vehicle service according to various aspects of the present invention;

Figure 2 is a flow chart illustrating overall operation of the system in fulfilling a booking through providing a private hire vehicle, and is performed by the system of Figure 1a or the system of Figure 1b;

Figure 3 is a flow chart illustrating calculating a score for a combination of a vehicle/driver pair in relation to a booking, and is performed by the system of Figure 1a or the system of Figure 1b;

Figure 4 is a flow chart illustrating allocation of a vehicle/driver pairing in relation to a booking, and is performed by the system of Figure 1a or the system of Figure 1b; and

Figure 5 is a schematic diagram illustrating components of a server forming part of the Figure 1a or Figure 1b system.

Detailed Description of Some Embodiments of the Invention

In brief, vehicles are allocated to bookings in a way that provides effective management of the fleet vehicles whilst minimising distance travelled by vehicles in fulfilling the bookings and whilst providing a relatively low burden on computer system resources that perform the allocation. In more detail, a pool of candidate vehicles is identified from the total fleet in a low computational cost manner. The candidate vehicles are assessed for their suitability and the most suitable vehicle is allocated, however it is allocated only when it is determined that it needs to be allocated so as to arrive at the pickup location in time to fulfil the booking. Because advanced processing is performed only for the most suitable candidate vehicle, the relatively computationally expensive process of determining road journey times is performed a relatively small number of times. Before allocation, the candidate vehicles are reassessed such that a different vehicle may be found to be the most suitable vehicle. Reassessment and allocation of the most suitable vehicle at a relatively late time allows better use of the vehicle and other resources of the fleet.

Figure 1a is a schematic diagram of a system for management of a private hire vehicle service according to various aspects of the present invention.

The system 100 includes a number of interconnected components, as will now be described. The system 100 includes at its centre a core system 101. This comprises one or more server computers running system software that ensures smooth operation of the system 100.

Key functions of the system 100 are bookings, allocation of a private hire vehicle to a booking, vehicle and driver management, account management and record keeping.

The booking function is provided primarily by a web booking server 102, an application booking server 103 and call centre terminals 104A and 104B, all of which are coupled to the core system 101.

The allocation function is provided primarily by a job allocation module 105, with information from other parts of the system 100.

The system includes database functions. In particular, an operational database 130 stores records that relate to general operation of the system 100. A driver network database 131 stores records that relate to drivers and vehicles that are managed by the system 100. Lastly, a historical database 132 stores records that have been archived from the operational database 130. Archiving of records from the operational database 130 to the historical database 132 occurs periodically and only records that are no longer needed for general operational use are archived.

The vehicle and driver management function is provided primarily by a driver location monitoring module 106 and a driver's devices server 107, with reference to the driver network database 131 as well as other components of the system 100.

The account management function is provided primarily by an account management module 117, utilising accounts information stored in the operational database 130 along with other components of the system 100.

The operational database 130 stores details of every account held with the operator of the system 100. Each account is identified by an account number stored in the

operational database 130. The accounts information stored in the operational database 130 may also include an account name, such as a company name and contact details for the company. The accounts information stored in the operational database 130 stores credit card details and/or other payment details so that payment can be taken from the account holder if permitted. A password and/or PIN (personal identification number) is associated with each account and stored with the accounts information in the operational database 130. Furthermore, a list of persons authorised to make bookings on the account may be stored, and optionally profiles for the individual authorised persons within the accounts.

The accounts information stored in the operational database 130 may also include a contact name and telephone number of a person who should be contacted in case of problems with the account. The accounts information stored in the operational database 130 includes information regarding invoicing preferences, for example the frequency of invoicing, date on which the invoice should be sent, the monthly/weekly credit limit and what information from each booking should be included on the invoice. The accounts information stored in the operational database 130 may indicate whether each account is active, or if it has been placed on hold. An account may be placed on hold by a credit control department and this may prevent further bookings being made on the account. Historical data of spending on the account may also be stored in the operational database 130, or this may be stored in the historical database 132.

The record keeping function is provided primarily by the historical database 132, although the operational database 130 and the driver network database 131 also provide some record keeping function.

In brief, a fleet of private hire vehicles is managed by the system 100. Each vehicle has a respective record in the driver network database 131, as will now be described.

The driver network database 131 stores information about every vehicle in the fleet. The registration number (license plate number) of each vehicle is stored in the driver network database 131. This may be used to identify each vehicle. Alternatively or in addition, a unique identifier separate from the registration number may be assigned to each vehicle as the primary means of identification within the driver network database 131.

Each Service is defined according to its vehicle type, capacity and other characteristics. In some embodiments, these types are "Up to 4 passengers", "Up to 7 passengers", "Electric vehicle", "VIP" and "Up to 4 passengers with luggage". The driver network database 131 stores the type of each vehicle and may also store a corresponding number or short string of characters which represents each type. Any special equipment such as a baby seat or the ability to accommodate a wheelchair is also identified in the driver network database 131. The driver network database 131 indicates the current driver to whom the car is assigned, although the driver/vehicle allocation changes from time to time.

The driver network database 131 stores the make and model and optionally the colour of each vehicle. The driver network database 131 also indicates the current status of the vehicle. In some embodiments, the status is chosen from "Driver Pool", meaning that the car is in use by a driver, "Free Pool", meaning that the car is not currently being used and is free to be allocated to a driver and "Workshop", meaning that the car is undergoing maintenance or repair. The driver network database 131 also stores the private hire license number (PCO) for each vehicle and the date on which this license expires as well as the road tax, vehicle insurance and MOT (vehicle roadworthiness certificate) expiry dates if appropriate. Examples of other data which may be stored are the date of purchase of the vehicle, the price paid for the vehicle, the date of manufacture, the supplier of the vehicle, warranty information and the date of the last inspection/maintenance.

Historic data about each vehicle may also be stored in the driver network database 131, such as a record of the previous registration numbers and a record of the previous drivers who were assigned to the vehicle. The service history and details of any accidents and repairs may also be stored.

The fleet of private hire vehicles is driven by a pool of drivers, each of which has a record in the driver network database 131.

The driver network database 131 stores information about each of the drivers registered with the operator of the system 100. The information relating to drivers includes personal data such as name, contact details (including phone number, home address), date of birth, next of kin and driver account data. Driver status information may be stored to indicate whether a driver is active or inactive, whether the driver has been

allocated a vehicle etc. Each driver may also be assigned an individual and unique identifier as a means of identifying the driver. Callsigns may also be used to denote drivers and are stored in the driver network database 131, although callsigns can be changed and reallocated between drivers as long as the same callsign is not in use by two drivers at the same time.

Driver account data includes an account number. Other information may include a driver's insurance details, a driver's length of service in the fleet, details of parking fines, historical wage information, holiday leave, driver diary information, information regarding payment collections from drivers, driver's licence number, national insurance (social security) number, details relating to a driver's taxicab licence (such as Public Carriage Office (PCO) licence), driver banking details (account number, sort code etc.). Miscellaneous information such as details of any allergies, smoker/non-smoker etc. may also be stored in the driver network database 131. Information regarding driver equipment such as a serial number of the driver's device 110 allocated to the driver, and mobile phone number of their driver's device 110 and their private mobile phone may also be recorded. Statistical information such as date of last job or historical earnings data may be recorded in the driver network database 131, or this may be recorded in the historical database 132.

Information relating to payments to and from drivers may be stored in the driver network database 131. Payments to the driver include a driver's wages. Driver outgoings may include, for example, car wash charges, insurance premiums, PCO renewal fee, accident costs, vehicle rental. To assist in maintaining this information, a purchase ledger number and contract number relating to each driver may be stored.

Each driver has an associated driver's device 110, three of which are shown at 110A, 110B and 110C in the Figure. The driver's devices 110 are portable electronic devices that are provided with wireless communication facilities. The driver's devices 110 may take any suitable form, but typically are smart phones or personal digital assistants or such like. The driver's devices 110 include a display and one or more input devices such as a keyboard, a voice control module or a touch screen or any combination thereof.

The driver's devices 110 are connected to the driver's devices server 107 via radio network 111, which may for instance be a mobile phone (cellular phone) network. In this case, the driver's devices 110 are provided with subscriptions to the mobile phone

network such that they can send digital data to and from the driver's devices server 107. Additionally, messages are able to be passed between the driver's devices 110 and the driver's devices server 107 through other media, and in particular SMS (short message service) messages and optionally also MMS (multimedia message service) messages.

The radio network 111 may alternatively be a dedicated radio network, for instance a private mobile phone network or a private radio network of some other type.

Data may be communicated between the driver's devices 110 and driver's devices server 107 over any suitable communications link, for instance using a data channel of a cellular telephone network such as a GSM, EDGE, GPRS, UMTS, HSxPA or LTE network.

The driver's devices 110 are configured to report their locations to the driver network database 131 at regular intervals, for instance 30 second intervals. The driver's devices 110 include GPS (global positioning system) receivers, which calculate the locations of the driver's devices 110 under control of the software applications executing on the driver's devices 110. Alternatively, they may include some other positioning module or device that is operable to calculate the positions of the driver's devices 110 with a suitable level of accuracy and resolution.

A private hire vehicle may be booked by a customer in one of three ways. Firstly, a private hire vehicle may be booked in a telephone conversation with a call centre operator. In this case, the customer initiates a telephone call with a call centre, an agent of which operates one of the call centre computer terminals 104A and 104B. The call centre agent then operates the terminal 104A, 104B so as to make the booking of the private hire vehicle according to the customer's requirements. The customer's requirements are obtained verbally during the telephone conversation between the customer and the agent.

In the second option, the customer may make the private hire vehicle booking using a browser application on a computing device 113, three examples of which are shown at 113A, 113B and 113C in the Figure. Each of the computing devices 113 is connected to the web booking server 102 by a network 114, which may for instance be the Internet or another public or private network. The web booking server 102 includes web server functionality that causes display of suitable web pages by the browser of the terminal

113. The customer's requirements with respect to the private hire vehicle booking are obtained by the web booking server 102 through the provision of suitable pages to the computer terminal 113 requesting the provision of the required information by the customer. The information may be provided by the customer through free text entry through the use of drop down lists, radio buttons etc. Some information may be pre-filled into the web pages provided by the web booking server 102.

Booking through the web booking server 102 may require the customer to login to a web portal before they can make their booking. The logging in may require the entering of a username and a password or PIN number. Through the control of a web session by the web booking server 102, for instance using cookies provided to the computer terminals 113, the booking can be known to have been validly made by virtue of the customer having being logged in to the web booking server at the time the booking was made.

The final way in which a customer can make a booking of a private hire vehicle is using a dedicated software application that is installed on and running on a portable communications device 112, three of which are shown at 112A, 112B and 112C in Figure 1. The portable communications devices 112 may take any suitable form, but typically are smart phones, feature phones, tablet computers or personal digital assistants or such like. The communication devices 112 are coupled to the application booking server 103 by a radio network 111, which may be the same as the radio network 111 described above with relation to the driver's devices 110 and the driver's devices 107.

The application is configured to provide a user interface that allows the customer to provide the software application with the information required to make the private hire vehicle booking. For instance, the software application, when executed, may cause the display of interactive pages that allow the customer to select or enter the required information. The software application is configured also to communicate the information relating to the booking that has been provided by the customer to the application booking server 103. If based on information provided by the customer it is determined that the application booking server 103 requires additional information, the software application running the mobile device 112 is configured to provide an interactive display to the customer such that the customer can provide the information, following which the software application causes it to be provided to the application booking server 103.

The customer may be required to log in to the software application on the mobile device 112, prior to making a booking. Logging in to the software application may require a username and a password or PIN number. Alternatively, the username may be entered during set up of the application and may not need to be entered subsequently when this software application is executed. If the username is not required to be entered, the user may log in to the software application simply by entering the password or PIN number.

The information about the private hire vehicle booking that is obtained during the booking process is as follows.

- Customer details. The customer details may be the name of the customer or an identifier that uniquely identifies the customer within the operational database 130.
- Service type. This indicates the category of vehicle. For instance, the service type may indicate a vehicle of a standard type and having four seats, or a vehicle of a standard type and having seven seats. The service may alternatively indicate a VIP vehicle, or an environmentally-friendly (electric or hybrid) vehicle (also known as a green vehicle).
- Journey type. The journey type may be a single (one-way) trip, or it may be a wait and return trip. The journey type may alternatively be a journey including multiple pick-up locations or multiple drop off locations or both multiple pick-ups and multiple drop off locations. The journey type may alternatively indicate that it is a pick-up from an airport or a drop at an airport.
- Pick-up address. This indicates an address at which the customer is to be picked up at the beginning of the journey. The address is a natural language address. The address is selected from one of the plurality of addresses stored in a database. The addresses may be stored in the operational database 130 or the historical database 132, or they may be provided by an external address database service, for instance geo.me or qas.co.uk. The addresses each have associated therewith a verified coordinate location expressed in latitude and longitude. Multiple databases may be used (in a hierarchical fashion) for address lookup. The pick-up address may be selected by the customer in any suitable way, with the most appropriate way depending on whether the customer is using the software application on their mobile device, using the web

booking service or using an agent in a call centre. If the journey type is an airport pick-up type, the pick-up address indicates the airport and terminal and optionally flight number.

- Drop off address. The drop off address again is selected from one of multiple addresses stored in the database and is selected by the customer in any suitable way. If the journey type is an airport drop off type, the pick-up address indicates the airport and terminal and optionally flight number.
- Pick-up date and time. This indicates a time and date which the customer requires the journey to start. Alternatively, the date and time may indicate ASAP (as soon as possible), if the customer requires the private hire vehicle at the earliest opportunity.

Optional information regarding the booking includes the following.

- Customer's reference. This can be provided for instance as free text or selected from a drop-down menu. If a reference is provided, this information can be included in an account statement against a journey at a later date.
- Additional comments. This is free text that provides any potentially relevant information, and may be provided to the driver once the booking has been allocated.

The system 100 comprises a journey cost calculation module 122. The cost calculation module 122 executes software code which determines the price for a requested journey, during the booking process and prior to vehicle allocation. Journey cost calculation is performed at the time of a booking and the result returned to the customer requesting the booking. The resulting cost for the journey is provided before the customer confirms the booking.

The journey cost calculation module 122 uses a number of different ways of calculating the base cost of the journey. The module 122 may set a fixed price for some journeys. These are agreed in advance with a particular account customer for journeys between pre-determined points. The cost calculation module 122 checks whether the booked journey and customer meet the requirements for a fixed price tariff. If the conditions are not met, then another pricing method is used. The cost calculation module 122 may use zonal pricing if a fixed price is not used. Where every point on the journey is within a defined zone, zonal pricing can be used. If neither fixed pricing nor zonal pricing is

used, or if the conditions for their application are not met, then the cost calculation module 122 may use an A to B (A-B) pricing method. The A-B method may specify the number of units between points A and B. A unit price depending on the type of vehicle etc. is then used to calculate the price. If there is no A-B record for a particular journey, the crow fly (direct) distance (i.e. the length of a straight line between the pick-up and drop-off locations) is used to calculate the base cost for the journey. This method may use map grid references or alternatively may be based on GPS data, i.e. the latitude and longitude of the pick-up and drop-off points.

The cost calculation module 122 may retrieve all the map and location information needed to make these calculations from the historical database 132. The historical database 132 may store a detailed geospatial model of a particular region, such as a city. As an alternative, or in addition to the methods described above, the cost calculation module 122 may use the real road distance for the journey, which is calculated using the road map from the historical database 132 and a route planning algorithm. Different rates may be used for different parts of a single journey. For example a first per mile rate may be used for the first 10 miles of a journey and a second per mile rate may be used for the rest of the journey. The historical database 132 may also store information regarding speed limits and historical traffic data. This information may also be used by the cost calculation module 122 to calculate an estimated time for the journey. The estimated journey time may then form the basis of the cost calculation.

Other criteria used by the cost calculation module 122 when calculating the price are the type of vehicle (VIP, green, 7-seater etc.) including any special facilities the vehicle has, the method of payment and the date and time of the journey. The cost calculation module 122 may also apply a flat "pick-up fee" for every journey.

The cost calculation module 122 may also determine how much of the fare charged to the customer is passed to the driver. This may be a simple percentage of the total fare or a more complex calculation based on one or more of journey time, distance, waiting time and number of passengers.

The allocation function allocates a vehicle and driver to a booking. The allocation function is described in some detail below. In brief, a vehicle and driver are allocated to the booking, and the associated customer, having regard to a number of factors including the pick-up location specified in the booking, the drop off location specified

in the booking, the service type specified in the booking, the date and time specified in the booking, the geographical distribution of the vehicles that are managed by the system 100, the demand for vehicles that are managed by the system 100 and information relating to the drivers.

The allocation function is automatic insofar as it does not require any manual involvement once the booking has been made. Once a job has been allocated to a particular driver and a particular vehicle, this is recorded in the operational database 130 with an indication that the journey has not yet been travelled.

The vehicle and driver management function includes a number of features. These include the monitoring of vehicle in terms of distance travelled etc. and ensuring that they provided for mechanical servicing at appropriate times. Drivers are managed also to ensure that documentation relating to private hire vehicle licenses, insurance etc. is in place. Additionally, the function maintains a record of hours worked and jobs performed, along with any other relevant information.

The accounts management function acts to manage information relating to customer's accounts with the operator of the system 100. This includes the maintenance and management of information such as authorised users, credit limits, invoicing requirement etc.

The record keeping function acts to store various information that is created by or observed by the system 100. This information includes information about bookings yet to be fulfilled, which is included in the operational database 130.

The sequence of steps that are performed by the system during execution of a job will now be described with reference to Figure 2. This shows execution of a relatively simple job in which there is one pick-up location, one drop off location, no driver reallocation and no variation in the journey. Additionally, journey costing and invoicing are not covered by this Figure.

The operation starts at step 2.1. At step 2.2, a booking is accepted by the system 100. Alternative ways for accepting a booking are described above with reference to Figure 1.

At step 2.3, the booking is confirmed to the customer, for instance by email. The message sent to the customer by the system 100 includes a booking reference number and some or all of the information relating to the booking, including the pick-up and drop off locations, the date and time of pick-up etc.

At step 2.4, the system 100 saves the booking until it is time to start allocation. For an ASAP job, allocation may commence straight away. Otherwise, allocation may start a fixed time before the date and time specified for pick-up in the booking information, as is described in more detail below.

At step 2.5, the system allocates a driver and a vehicle to the job. This is discussed in relation to Figure 3 and Figure 4 below. On allocating the vehicle, the status of the vehicle and driver is changed from "Available" to "Allocated". This prevents the vehicle/driver being allocated to a different job until the status changes to a suitable status.

At step 2.6, the system 100 sends a message to the customer with details of the allocated vehicle. The message includes text such as 'Your vehicle is on its way'. The message also indicates the job number, which may be the same as the booking number. The message also indicates the identity of the vehicle, so that it can be readily identified by the customer. The identity of the vehicle may be indicated for instance by the registration or licence plate that is provided on the vehicle. It may also indicate the make and model of the vehicle, and/or the colour of the vehicle. Additionally, the message includes information by which the customer can contact the driver that has been allocated to the job. For instance, it may include the mobile telephone number of the driver. Providing the mobile telephone number of the driver allows the customer to call the driver with any comments or questions that they may have before the customer is collected by the vehicle. Additionally, the message includes a hyperlink to a webpage at which the location of the vehicle is shown on a map. This allows the customer to identify where the vehicle is at any stage between the vehicle being allocated to the job and the customer being collected by the vehicle.

At step 2.7, the system 100 sends a message to the driver with details of the job. The message includes various pieces of information including the name of the customer. This allows the driver to confirm the customer when the driver meets the customer at the pick-up location. The message also includes the pick-up location and the drop off

location. The pick-up location and drop off location may be provided in the message in such a way that they can be extracted by the driver's device 110 and automatically placed into a navigation application that is present on the driver's device 110. This allows the driver to commence the provision by the driver's device 110 of navigation guidance to the pick-up location in response to the driver selecting the pick-up location by way of an input on the driver's device 110. Similarly, after the customer has been collected at the pick-up location, the driver can cause the device 110 to commence providing route guidance to the drop off location by providing a suitable input on the driver's device 110.

At step 2.8, the system 100 receives a POB (passenger on board) message from the driver. This message is transmitted by the driver's device 110 in response to the driver indicating that they have collected the customer from the pick-up location. The option to indicate POB status is provided to the driver once the driver device 110 determines that the vehicle has arrived at the pick-up location, or is within a predetermined radius (e.g. 50m) of the pick-up location and has become stationary. However, the sending of the POB message from the driver's device 110 is not automatic. In this step, the status of the vehicle/driver is changed from "Allocated" to "POB".

Following receiving the POB message from the driver, the system 100 at step 2.9 records that the customer has been picked up. Next, the system 100 receives a drop off message from the driver at step 2.10. This is message is sent by the device 110 after the driver indicates to the driver's device 110 that the customer has been deposited at the drop off location. The option to indicate that the customer has been dropped off may be provided to the driver upon the driver's device 110 determining that the vehicle has reached the drop off location or is within a predetermined radius (e.g. 50m) of the drop off location and has become stationary. However, the sending of the drop off message from the driver's device 110 is not automatic.

After the drop off message has been received from the driver's device 110 at step 2.11, the system 100 completes a journey record for the journey in the operational database 130 (the record was created during the booking process). The record of the journey stored in the operational database 130 includes the following information. The record includes the pick-up address and the drop off address. The information also includes the pick-up time and date and, if different, the booking time and date. The record also includes the drop off time and date, as detected by the system 100 in response to

receiving the drop off message from the driver at step 2.10. The record also includes the cost of the journey, in terms of financial value.

The record also includes the travelled distance, which is not the crow fly (direct) distance between the pick-up and drop off locations but instead is the road distance travelled by the vehicle. The record also includes the journey time, in terms of minutes and seconds. The record also includes vehicle type information that indicates the type of vehicle that performs the journey.

The record also includes the booking information relating to the journey, which may include information about the identity of the customer that made the booking, the time of making the booking, the mode of making the booking (e.g. web, application or call centre) and any other relevant information relating to the booking.

Next, at step 2.12 the driver and vehicle are reallocated to the pool of available drivers. This is achieved by changing the status of the vehicle/driver to "Available" from "POB".

The customer is then messaged with a receipt for the journey travelled, if required, at step 2.13. Lastly, the operation ends at step 2.14.

A method of scoring a vehicle against a booking will now be described with reference to Figure 3. The scoring process of Figure 3 is performed by the job allocation module 105.

The operation starts at step 3.1. In brief, different scores are calculated at steps 3.2 to 3.7, and at step 3.8 the scores are summed together. Clearly, it will be appreciated that the scores may be calculated in any order, and may be calculated wholly or partly in parallel.

At step 3.2, a distance score is calculated. The distance score allows the distance between a vehicle and the pick-up location of the booking to be taken into account when scoring the vehicle against the booking. The distance score is calculated as the distance between the current position of the vehicle and the pick-up address. The distance has the unit of miles, but it may alternatively be kilometres. The distance is calculated as the distance that will need to be travelled by the vehicle to reach the pick-up address, taking into account road layout, one way streets etc. This is known as the

road distance. The shortest route from the vehicle to the pick-up address is used for the distance location, even if this is not the quickest route. The route and the road distance thereof are calculated by the system 100 using information from the historical database 132. It is the last recorded position of the vehicle that is used in the distance score calculation.

An administrator or other operator of the system 100 may apply a setting such that the distance score is always zero, in which case the distance between the vehicle and the pick-up location is not taken into account in the score calculation.

At step 3.3, a service compatibility score is calculated. The service compatibility score results in the taking into account of the car type preference that was specified in the booking against the type of the vehicle that is being scored. If the type of vehicle that is being costed is the same type as that is specified in the booking, or is consistent with that type, then the service compatibility score is zero. The service compatibility score takes a positive value if there is incompatibility between the service type of the booking and the type of vehicle that is being costed. In the case of the booking specifying a VIP and the vehicle being costed being a standard vehicle, a penalty of 500 may be provided as the service compatibility score. This penalty helps to ensure that a VIP vehicle will be provided to fulfil the booking if one is available, but if not then a standard car can be provided.

In the case of the booking specifying a standard four passenger vehicle, a penalty score of 50 points is provided for a vehicle that is a seven-seater vehicle. This helps in ensuring that the booking is serviced with a suitable car, but also contributes to avoiding the removal of a large capacity vehicle from the pool of available vehicles unnecessarily.

In the case of the booking being for a standard car and the vehicle type being a VIP car, a penalty score of 100 is provided. Similarly to the situation described in relation to the larger capacity vehicle, this helps to ensure that the booking is satisfied whilst not removing VIP vehicles from the available fleet unnecessarily.

At step 3.4, an empty time score is calculated. The empty time score allows the utilisation of the vehicle (and corresponding driver) to be taken into account in the scoring of the vehicle in relation to the booking.

The empty time score is calculated as the product of -1 and the time (in minutes) since the last job allocated to the car/driver combination was completed and a cost per empty minute. The cost per empty minute is in effect a weighting factor. The weighting factor may be set by an administrator of the system 100. For a vehicle that is in the state POB, the empty time score is zero.

The inclusion of an empty time score in the operation of Figure 3 helps to provide load balancing of the vehicles, and load balancing of the drivers. Vehicle load balancing helps to even out wear and tear on different vehicles in the fleet on a unit time basis. Load balancing of drivers is useful because it helps to prevent the likelihood of drivers performing too many consecutive jobs with insufficient breaks in between the jobs, and it also helps to reduce the likelihood that drivers will wait for long periods between jobs. Load balancing of drivers, through use of the empty time score in the costing operation, helps to prevent driver fatigue and thus improves safety.

At step 3.5, a going home score is calculated. If the status of the driver is 'going home', then a score is calculated. If the driver has some other state, then the going home score is zero.

If the driver's status is 'going home', the going home score is calculated as the product of -1 and the number of saved miles and a distance criteria. The saved miles component of the score provides a measure of how much closer to their home the driver would be if they fulfilled this booking. The saved miles component is calculated as the current distance to home (which is the road distance from the current location of the vehicle to the driver's home address) minus the distance between the drop off address and home (which is the road distance from the drop off location of the booking to the driver's home address). The distance criteria provides a weighting, and may be set by an administrator of the system 100.

The effect of the inclusion of the going home score is to increase the likelihood that a job will be allocated to a driver who is on the way to their home (for instance for a lunch break or having finished their shift) if the job would take the driver to a location that is nearer to their home. The magnitude of the score depends on the distance that would be saved, so a score is obtained if the drop off location is relatively closer to the driver's home address.

At step 3.6, a drop 5/10 score is calculated. For drivers that have a 'drop in 5' or a 'drop in 10' status, the drop 5/10 score has a positive value. For drivers that do not have a 'drop in 5' or a 'drop in 10' status, that is for drivers that are vacant and not allocated to a booking, the drop 5/10 score is zero. The status of the vehicle is set by the driver through their driver's device 110. In particular, when the driver's device 110 calculates that there are fewer than 10 minutes remaining in the journey to the drop off address, the driver's device 110 provides an option to the user to adopt the 'drop in 10' status. If the driver selects this option on the driver's device 110 (when the vehicle is stationary), the 'drop in 10' status is entered. Similarly, when the driver's device 110 detects that there are fewer than five minutes remaining in the journey to the drop off location, the driver's device 110 provides an option to allow the driver to select entering the 'drop in 5' status.

If the driver of the vehicle has a 'drop in 5' status, a score of 20 points is calculated. If the driver has the 'drop in 10' status, a score of 30 points is calculated.

The calculation of a drop 5/10 score allows vehicles that have a POB status (that is, they have a job in progress) to be considered for allocation to a booking. However, a penalty is applied to them with the result that they are less favoured than vehicles that are currently empty. This provides protection against the driver arriving late for the booking if there are unexpected delays in the previous journey.

At step 3.7, if the status of the vehicle is 'POB', 'drop in 5' or 'drop in 10' (i.e. the vehicle has a passenger on board and is currently fulfilling a booking), a risk score is calculated. In brief, the inclusion of the risk score allows vehicles that are not currently available to be allocated to a booking to be considered for the booking whilst taking into account the risk that they will not be able to fulfil the booking.

In some embodiments, the risk score is calculated according to:

$$\text{Risk Score} = \text{POB penalty} + \text{POB risk penalty} * \text{risk}$$

where

$$\text{risk} = \text{home risk} + (1 - \text{home risk}) * \text{shift end risk}$$

where

- home risk = base home risk / 100;
- base home risk =
    - 0, if distance to home > 6;
    - 50, if 3 < distance to home < 6;
    - 1000, if distance to home < 3;
- shift end risk = base shift end risk / 100;

and where:

Base shift end risk = current time - driver shift start time

Thus, the value of base shift end risk is a measure of the period of time elapsed since the driver started the current shift. The value increases, giving rise to a higher value of the risk score, as the shift length increases.

The value of the base home risk parameter varies according to the distance to the driver's home from the drop off location of the booking currently being fulfilled. It is banded such that it may take one of only a predetermined number of values, in this case 3. If the distance to the driver's home is low, the value of the base home risk is relatively high, and vice versa.

The value of the shift end risk parameter may be derived from a look up table in which bands are stored. The look up table may for instance take the following form:

Duration	Base shift end risk
0-2	20
2-3	25
3-4	30
4-5	40
5-6	50
6-7	60
7-10	80

>10	90

According to this look up table, a shift length of 2.6 hours provides a base shift end risk value of 25, and a shift length of 4.6 hours provides a base shift end risk value of 40.

In general terms, the risk score is a function of the distance between the drop off location of the booking currently being fulfilled and the home address of the driver of the vehicle. It is a function also of the time since the driver of the vehicle started their shift. The risk score is a measure of the likelihood that the driver will end their shift and go home after they have completed fulfilling their current booking.

At step 3.8, a late score is calculated. In brief, the inclusion of the late score introduces a penalty for vehicles that it is calculated would be likely to arrive late at the pickup location.

The late time score is calculated according to the following:

Late Time Score = late time * late time weighting factor

Where the value of the late time parameter is the number of minutes after the start time for the booking that the vehicle is predicted to arrive at the pickup location and the late time weighting factor is a configurable parameter that weights the contribution of the late time to the overall score for the vehicle/booking combination.

In general terms, the late time score is a function of how late it is expected that a vehicle would be if it were allocated to a booking at the current time.

As discussed in more detail below, the late score is not calculated for every vehicle. Instead, the late score is only calculated for a relatively small number of vehicles.

At step 3.9, the scores calculated in steps 3.2 to 3.8 are summed, to provide a total score for the driver/vehicle/booking combination. This score is then used in an allocation process, as will now be described in reference to Figure 4. The allocation process of Figure 4 is performed by the job allocation module 105.

Referring to Figure 4, the operation starts at step 4.1.

At step 4.2, a booking is made and entered onto the system. This corresponds to step 2.2 of Figure 2. The booking constitutes a vehicle requirement. Vehicles in the fleet of private hire vehicles constitute vehicle resources that can be allocated to the vehicle requirement constituted by a booking.

At step 4.3, the job allocation module 105 calculates a value of X for the booking. The value of X indicates a time before the booking start time at which to start the allocation process. The calculation of a value of X will now be described.

Advantageously, the value of X depends on the pickup location for the booking. The value of X may also depend on the booking priority, for instance based on whether the booking is an account booking (relatively high priority) or a cash booking (relatively low priority). The value of X may depend also on the identity of the account under which the booking was made, with relatively high value accounts having a higher priority than relatively low value accounts. For instance, the value of X may depend on a zone in which the pickup location is present. A relatively low value of X may be used for zones in which a high concentration of fleet vehicles usually is present, an intermediate value of X may be used for zones in which an intermediate concentration of fleet vehicles usually is present, and a high value of X may be used for zones in which a low concentration of fleet vehicles usually is present.

In London, a value of X=15 is considered to be suitable for pickup locations in central zones, for instance zones having postcodes starting with SE1, SE11, SE17, SW1, SW4, SW8, WC1, WC2, EC1, EC2, EC4, N1, W1, E1, E2 etc. A value of X=20 is considered to be suitable for intermediate zones, for instance zones having postcodes starting with SW6, SE8, N6 etc. A value of X=30 is considered to be suitable for pickup locations in outer zones, for instance zones having postcodes starting with SE6, SW20, TW10, HA9 etc. A value of X=60 is considered to be suitable for pickup locations in zones outside but surrounding London, for instance zones having postcodes starting with RH3, SG13, KT15. The value of X may be as high as 300 for pickup locations that are remote from London.

The value of X may vary for different zones depending on the priority of the booking. For instance, in London a value of X=25 is considered to be suitable for low priority (cash) bookings with pickup locations in zones having postcodes starting with SE21,

SE26, Sw16 and SW13 but a value of X=30 is considered to be suitable for high priority (account) bookings in the same zones.

Alternatively the value of X could be fixed or could be dynamically set by an operator or controller of the system 100. A fixed or set value may be independent of the pickup location or priority. A fixed or set value of X may be between 10 and 30 minutes and for instance the value of X may be 20. In this case, the allocation process starts 20 minutes before the scheduled pickup time.

A relatively low value of X means that the job allocation module 105 does not start to allocate the job until relatively soon before the booking start time. This equates to a lower amount of processing that is required to allocate a job if there are multiple vehicles that could fulfil the booking whilst allowing the best vehicle to fulfil the bookings. Because a low value of X is applied in zones that typically have a high concentration of vehicles, this provides advantages of commencing job allocation at a time when a number of vehicles may be able to satisfy the booking and so allows more options for allocating the best vehicle for the booking.

A relatively high value of X means that the job allocation module 105 starts to allocate the job a relatively long time before the booking start time. This provides a greater possibility to identify a vehicle to fulfil the booking without arriving late for the booking, and is particularly useful where the vehicle may have a significant distance to travel to the pickup location as a result of the relatively low vehicle density. Because a relatively high value of X is applied in zones that typically have a low density/concentration of vehicles, this provides advantages of commencing job allocation at a time when a number of vehicles may be able to satisfy the booking and so allows more options for allocating the best vehicle for the booking. This effect is advantageous also for high value customers, for which a higher priority is accorded.

At step 4.4, the job allocation module 105 waits until X minutes before the start time for the booking. This results in the allocation process being commenced an appropriate time before the start time (on the correct date).

At step 4.5, the job allocation module 105 identifies the Y vehicles that are nearest to the pickup location of the booking. The value of Y advantageously is between 10 and 100, and may for instance be 40. The vehicles are determined to be nearest if they have

shortest crow fly distance between their current location (which is their last reported location, as described above vehicle locations usually are reported every 30 seconds) and the location of the pickup address. The distance is calculated as the straight line distance between the latitude and longitude coordinates of the location of the vehicle and the location corresponding to the pickup address.

In step 4.5 all vehicles with an active status, i.e. 'available', 'going home', 'POB, 'drop in 5' or 'drop in 10' can be selected. Vehicles that have a status of 'off shift', 'allocated' (to a booking), or 'on break' (the driver is on a break) are not considered for selection. The result is a pool of candidate vehicles for the booking. The number of candidate vehicles typically is a much smaller number than the number of vehicles in the fleet that have one of the relevant statuses.

The use of crow fly distances in step 4.5 results in an appropriate number of vehicles being selected for possible allocation to the job but without requiring the processing needed to calculate road routes and road distances for each of the vehicles. The calculation of crow fly distances is computationally non-intensive compared to road distance formulation.

At step 4.6, a score is calculated for the vehicle/booking combination for each of the pool of candidate vehicles, i.e. the vehicles that were identified in step 4.5. The score is calculated as described above with reference to Figure 3. The result is a numerical value that is an indication of the suitability of the vehicle for the booking. Put another way, the result is a metric of the suitability of the vehicle for the booking.

At step 4.7, the job allocation module 105 calculates a number Z of vehicles for which to calculate journey times. This may be performed in any suitable way.

In some embodiments, the job allocation module calculates the value of Z to be the number of vehicles whose score calculated at step 4.6 falls below a threshold value (if a low score is better) or exceeds a threshold value (if a higher score is better). The threshold may be an absolute amount, for instance 100. Alternatively, the threshold may be related to the score of the best scoring vehicle. Where a low score is better, the threshold may be calculated as a multiple of the score of the best scoring vehicle and where a high score is better the threshold may be calculated as a fraction of the score of the best scoring vehicle. For instance, if the multiple is 3 and the best scoring vehicle

has a score of 20, the threshold is set at 60. Alternatively, the threshold may be set at a predetermined distance from the best scoring vehicle. For instance, if the predetermined distance is 40 and the best scoring vehicle has a score of 20, the threshold is set at 60. The value of Z is capped at a maximum proportion (for instance 25%) of the value of Y.

In some other embodiments, step 4.7 involves looking up a value of a setting that is stored in the system. The value of the setting may be set manually or automatically. The value may be 1. Advantageously, the setting is between 3 and 10.

In still further embodiments, the value of Z is set as a predetermined proportion of the value of Y. For instance, the value of Z may be set at 0.1 times, 0.2 times or 0.3 times the value of Y.

However it is calculated, the value of Z is lower than the value of Y. As such, road distance calculation is not performed for all of the Y vehicles that were determined to be closest to the pickup location at step 4.5.

At step 4.8, a journey time for each of the Z best scoring vehicles is calculated. The Z best scoring vehicles are the vehicles whose scores are better than the other vehicles that were scored at step 4.6. The journey time for a vehicle is the journey time from its current location to the pickup location associated with the booking.

The journey time for each of the Z best scoring vehicles is calculated at step 4.8 by the job allocation module 105 or it may alternatively be calculated by the core system 101. The journey time is calculated by predicting journey times for a small number of the most direct routes between the vehicle location and the pickup address, and choosing the shortest journey time. If one route can be determined to be the best route, then the journey time for that route is taken as the journey time between the vehicle location and the pickup location. Alternatively, the journey with the smallest time cost may be identified and used to determine the journey time. For vehicles that have the status 'drop in 5' 'drop in 10' or 'POB', the journey time is calculated from the drop off location of the booking being fulfilled, rather than the vehicle's current location.

The prediction of journey times may use actual or estimated averages of transit times of segments of the journey (which are stored in the map and location database 109) added

together to produce a journey time for the required journey. In the simplest embodiments, the segment transit times are estimated from known speed limits (maximum permitted speeds) and junction counts (numbers of traffic lights, give way junctions etc.) in the segments. In other embodiments, historical averages are stored in the map and location database 109 or historical database 132 and are used in the journey time calculation. In the most sophisticated embodiments, historical averages for different times of day and different types of day (weekday versus Saturday versus Sunday) are stored in the map and location database 109 and are used in the journey time calculation. The more sophisticated embodiments typically require more data to be collected, processed and stored in the map and location database 109. However, they provide more accurate journey time calculations, in particular by taking into account issues such as weekday rush hour traffic conditions and generally taking into account local traffic volume and flow patterns. However it is performed, the calculation of a journey time is a relatively computing resource intensive step in part because it requires identification of one or more best routes between two locations.

Step 4.8 also involves calculating a late score for each off the Z best vehicles. Calculation of the late score is described above with reference to Figure 3. This involves determining how late the vehicle would be if it were allocated to the booking at the current time. This takes into account the calculated journey time and the time remaining before the start time of the booking. This may or may not take into account the allocation buffer, which is discussed below. If any of the vehicles scores a non-zero late score, this is taken into account in the overall score for the vehicles and may change which vehicle is the best scoring vehicle for the booking.

At step 4.9, an allocation buffer value is calculated for the best scoring vehicle. The allocation buffer provides a margin or cushion for the vehicle to arrive at the pickup location at the start time. The value of the allocation buffer may be calculated in any suitable way.

The allocation buffer value calculated in step 4.9 may for instance be calculated using a predetermined margin or a predetermined multiplier. A larger allocation buffer value may be calculated for higher priority bookings, for instance those bookings originating from high value customers.

The allocation buffer value may be calculated as the calculated journey time multiplied by a predetermined factor, which may have a value between 0.2. and 0.5 and most advantageously is between 0.3 and 0.4. Alternatively, the allocation buffer may be independent of the time remaining to the booking start time or the calculated journey time. It may be a fixed value, for example with a value between 1 and 3 minutes. The value may instead be calculated as the sum of a number of values which are set on a zone basis, e.g. each zone is allocated a value (which may vary according to the time of day and the day of the week) and the allocation buffer is the sum of the values for the zones through which the vehicle would need to travel to reach the pickup location.

The allocation buffer value may additionally include an element that derives from the status of the vehicle. For instance, an additional amount may be added for vehicles with the status 'drop in 5' (for instance an additional 5 minutes), 'drop in 10' (10 minutes) and 'POB' (estimated time remaining until the vehicle reaches the drop off location).

At step 4.10, the job allocation module assesses the journey time for the best scoring vehicle, in particular to assess whether the vehicle needs to start the journey to the pickup location in order to arrive there at the start time. The assessment involves the following calculation:

Optimum time to leaving = booking start time – (journey time + allocation buffer)

Where:
optimum time to leaving is the time of day that would be optimum for the vehicle to start the journey to the pickup location,
booking start time is the start time for the booking,
journey time is the journey time value calculated for the best scoring vehicle at step 4.8, and
allocation buffer is the allocation buffer value calculated for the best scoring vehicle at step 4.9.

The difference in time between the current time and the optimum time for leaving is then calculated and the result is compared to a threshold. The result of the comparison

of the value to the threshold is used in the next step to determine whether it is time to allocate the vehicle to the booking.

The use of an allocation buffer value in the calculation increases the chances that, on average, vehicles will arrive to fulfil bookings on time whilst minimising the chances that vehicles will arrive too early and thus be a non-utilized resource whilst waiting for the booking start time.

At step 4.11, it is determined whether the current time is the time to allocate a vehicle to the booking. The threshold mentioned above in relation to step 4.10 may for instance be 1 (1 minute) or zero. If it is calculated at step 4.11 that the difference in time between the current time and the optimum time for leaving is less than the threshold, then a positive determination is made. This indicates that the vehicle needs to start now (or within the next minute) their journey to the pickup location in order to arrive in time (having regard to the allocation buffer value).

On detecting at step 4.11 that it is time to allocate a vehicle to the booking, the best scoring vehicle is allocated to the booking at step 4.12 before the operation ends at step 4.13.

If it is detected that it is not time to allocate the vehicle, the operation continues to step 4.14. Here, it is determined whether there is a clear best scoring vehicle and if so whether that vehicle is moving away from the pickup location. A clear best scoring vehicle is one that has a score that is significantly better than (e.g. a threshold amount away from or multiple/fraction of) the score of the next best scoring vehicle. A determination as to whether it is moving away is made by comparing the road distance from the vehicle to the pickup location to a previously calculated road distance, or by comparing the calculated journey time to the last calculated journey time. On making a determination that the best scoring vehicle is moving away from the pickup location, the operation proceeds to allocate the best scoring vehicle to the booking at step 4.12. Step 4.10 may be dependent on other conditions being present, for instance relating to the time remaining to the start time of the booking and the journey time to the pickup location.

By making the assessment at step 4.10 in respect of the best scoring vehicle, it is the vehicle that is best suited to the booking that is determined to be required to be

allocated to the booking even if that vehicle is not the vehicle that is closest to the pickup location or has the shortest journey to the pickup location. This allows the optimisation of the allocation of fleet vehicles to bookings, and it allows it to be achieved in an intelligent way. Additionally, it allows modification of the way in which vehicles are selected for fulfilling bookings through adjustment only of the cost scoring process (of Figure 3), and does not require adjustment of the allocation process (of Figure 4).

After step 4.14, at step 4.15 the job allocation module 105 wait until X-A minutes before the pickup time. The value of A was initialised at 1 when the operation began at step 4.1 so on first performance of step 4.15 the operation waits until X-1 minutes before the start time of the booking. Once this time has been reached, the value of A is incremented at step 4.17 and the operation returns to step 4.5, where the nearest Y vehicles to the pickup location are identified. An effect of steps 4.16 and 4.17 and the return to step 4.5 is that nearest vehicles are identified and scored once every minute until a vehicle is allocated to the booking.

On subsequent execution of step 4.5, vehicles that no longer have one of the relevant statuses ('POB', 'available', 'drop in 5' or 'drop in 10') are not identified as candidate vehicles and are not scored at step 4.6.

On subsequent performance of step 4.6 in relation to a given booking, many of the same vehicles will be scored but usually they will score differently to the last occasion on which they were scored. Vehicles will score worse (with a higher score) in some instances, for instance if they have moved away from the pickup location by a significant distance. Vehicles will score better (with a lower score) in many instances, for instance if the status of the driver has changed from 'drop in 5' to 'available'. Also, wait times necessarily will be longer and so the wait time score (which is negative) will reduce the total score for many vehicles. As such, on subsequent performance of step 4.7 different vehicles may be in the Z best vehicles and a different vehicle may be the best scoring vehicle. In this case, it is a different vehicle for which the journey time to the pickup location needs to be calculated so that it can be determined at step 4.11 if it is time to allocate the vehicle to the booking.

Whether or not different vehicles form the pool of Z best vehicles are identified on subsequent execution of step 4.7, or Z has a different value, the calculation of the

journey times is made again for the vehicles at step 4.8. If the journey time for the best scoring vehicle has already been calculated for the same or a very similar location of any of the Z best scoring vehicles (which occurs when the vehicle is stationary or has not moved significantly), step 4.7 may use the previously calculated journey time instead of calculating the journey time again. A journey time for a vehicle may be taken from a journey time calculated for a different vehicle either currently or very recently if there is a substantial match in vehicle locations.

Whether or not a different vehicle is the best scoring vehicle on subsequent execution of step 4.6, at step 4.11 it is determined whether it is time to allocate the booking, based on the journey time to the pickup location and the time remaining to the start time for the booking etc. Since step 4.11 necessarily is performed at a later time than previous performance of the step (1 minute later in this example), it may on the latest performance of step 4.11 be determined that it is time to allocate the vehicle even if the best scoring vehicle has not moved, in particular when time has elapsed such that the vehicle needs to start its journey to arrive at the pickup location on time. Once it is determined that it is time to allocate the vehicle, the best scoring vehicle is allocated to the booking at step 4.12. This is the case regardless of whether it is predicted that the vehicle could arrive at the pickup location on or before the start time. The vehicle that is allocated to the booking at step 4.12 will, by virtue of the scoring operation of Figure 3, be the vehicle that is best suited to the booking. It may be that the circumstances are such that the vehicle will be allocated at a time when it could arrive late at the pickup location. This may occur for instance if the best vehicle changed status from 'on break' or 'off shift' to 'available' at a late stage in the allocation process or if there is a severe shortage of vehicles near the pickup location for the booking. However, the inclusion of the allocation buffer means that the vehicle may not arrive late.

The configuration of the job allocation module 105 to allocate a vehicle to the booking at the last minute, or 'just in time', increases the flexibility of allocation of vehicle resources of the fleet. It also contributes to reducing the overall mileage that is travelled by the vehicles of the fleet in order to satisfy the bookings that are received by the system 100. In the absence of the dynamic and flexible allocation described above, vehicles would be allocated the bookings less optimally, resulting in an increased overall distance travelled of the vehicles of the fleet. Reduced distance travelled equates to lower fuel consumption leading to lower carbon emissions, lower vehicle

wear and tear, lower chance of accidents, lower driver fatigue, and generally lower costs.

The use of crow fly (direct) distance calculation in selecting candidate vehicles whilst using calculated journey times in determining the best vehicle from the pool of candidate vehicles minimises system resource use whilst providing effective allocation because it requires relatively little journey time calculation.

It will be appreciated from the above that steps 4.5 to 4.11 are repeated until a vehicle is allocated to the booking. The number of times that the steps are repeated depends on the initial value of X, which dictates how long before the start time of the booking the allocation process of Figure 4 begins, and the number of minutes before the start time that the vehicle is allocated to the booking. For bookings in central city locations where there are relatively large numbers of vehicles and journey times to potential pickup locations may be short, bookings may be allocated only a small number of minutes, for instance 2, 3 or 4 minutes, before pickup times. For bookings in less central locations, where there may be relatively few vehicles and a low vehicle density, bookings may be allocated significantly longer before the start time for the booking, for instance 12, 15 or 18 minutes before the start time. For bookings with pickup locations where there are no vehicles, so a vehicle needs to travel a long distance, bookings may be allocated many minutes, potentially a number of hours, before the start time. By selecting values of X for different zones, the allocation process of Figure 4 will typically run for approximately the same period of time across the fleet at any given time. During busy periods or during periods when there is a relatively high number of active vehicles, bookings may be allocated closer to the start times of bookings and so the allocation process may run for longer.

For vehicle fleets with relatively low vehicle densities, having regard to the covered geographical area, a higher value of X may be appropriate. Advantageously, the value of X, which indicates the number of minutes prior to the start time that the allocation process begins, may be set by an administrator of the system. The value of X may be settable by an administrator on a per zone basis. A global setting may also be applied, by which the value of X is altered for all zones by a given amount or factor.

In some embodiments, the scores of vehicles in pre-allocation processes for other bookings are taken into account when allocating a vehicle to a booking. In particular,

before allocating the best scoring vehicle to a booking, the job allocation module 105 determines how the vehicle ranks in other bookings for which vehicles are being scored but which have not yet had a vehicle allocated to them. If a vehicle is the best scoring vehicle against a first booking but is the best scoring vehicle against a second booking by a larger margin, the job allocation module 105 acts to allocate the vehicle to the second booking even if at the current time the vehicle would be allocated to the first booking so as to arrive at the pickup location on time to fulfil the booking. This may be achieved by allocating the vehicle to the second booking or by applying a penalty score for the vehicle against the first booking such that it ceases to be the best scoring vehicle against that booking.

Whether the system 100 includes one or multiple servers, each server includes a number of features as will now be described with reference to Figure 5. Figure 5 shows one server 40. If the system 100 comprises plural servers, multiple versions of the Figure 5 server 40 are connected together.

Each server 40 in the system 100 includes a processor 412. The processor 412 is connected to volatile memory such as RAM 413 by a bus 418. The bus 418 also connects the processor 112 and the RAM 413 to non-volatile memory, such as ROM 414. A communications interface 415 is coupled to the bus 418, and thus also to the processor 412 and the memories 413, 414. The interface 415 is connected to a radio network in any suitable way, for instance via the Internet or a local network. Within the ROM 414 is stored a software application 417, which includes program code that causes the server to perform the functions required of it. An operating system (OS) 420 also is stored in the ROM 414.

An output device such as a display 419 may be provided with the server 40. An input device such as a keyboard 421 may be provided with the server 40. This allows configuration, monitoring and updating by administrators and other users as required.

The server 40 may take any suitable form. Generally speaking, the server 40 comprises processing circuitry 412, including one or more processors, and a storage device 414, 413, comprising a single memory unit or a plurality of memory units. The storage device 414, 413 stores computer program instructions that, when loaded into the processing circuitry 412, control the operation of the server 40.

The term 'memory' when used in this specification is intended to relate primarily to memory comprising both non-volatile memory and volatile memory unless the context implies otherwise, although the term may also cover one or more volatile memories only, one or more non-volatile memories only, or one or more volatile memories and one or more non-volatile memories. Examples of volatile memory include RAM, DRAM, SDRAM etc. Examples of non-volatile memory include ROM, PROM, EEPROM, flash memory, optical storage, magnetic storage, etc.

Reference to "computer-readable storage medium", "computer program product", "tangibly embodied computer program" etc., or a "processor" or "processing circuit" etc. should be understood to encompass not only computers having differing architectures such as single/multi processor architectures and sequencers/parallel architectures, but also specialised circuits such as field programmable gate arrays FPGA, application specify circuits ASIC, signal processing devices and other devices. References to computer program, instructions, code etc. should be understood to express software for a programmable processor firmware such as the programmable content of a hardware device as instructions for a processor or configured or configuration settings for a fixed function device, gate array, programmable logic device, etc.

It should be realised that the foregoing embodiments are not to be construed as limiting and that other variations and modifications will be evident to those skilled in the art and are intended to be encompassed by the claims unless expressly excluded by the claim language when taking into account equivalents. Some such alternatives and modifications will now be described.

In the above, journey cost calculation is performed at the time of a booking and the result returned to the customer requesting the booking. Cost calculation may alternatively be performed ahead of the booking being made (for instance on the basis of an agreed tariff), at the end of fulfilment of the booking, or at a later time.

Additionally, the cost scoring of a vehicle against a booking may be performed in any suitable way. Also, cost scoring may be performed only once and the best vehicle allocated at that time, rather than cost scoring being performed until it is decided to allocate a vehicle.

Instead of selecting vehicles for the candidate pool of vehicles based solely on crow fly distance, vehicles may be selected based on zoning of their current location. For instance, in respect of a pickup location in a certain zone, all vehicles that are currently located in the same zone or a zone that borders that zone may be selected for inclusion in the pool of candidate vehicles. If zone information is stored in the map and location database 109 (or another database) as part of the current location information for vehicles, using zone information in the selection of candidate vehicles may be particularly computationally non-intensive. The use of zones also allows barriers (such as rivers) between geographically proximate areas to be taken into account when classifying whether or not zones are adjacent.

Instead of calculating the allocation buffer value only for the best scoring vehicle, the allocation buffer value may instead be calculated for the 2, 3 or 4 best scoring vehicles. Although this can increase the number of allocation buffer values that need to be performed, it may increase the likelihood that the allocated vehicle will arrive at the pickup location prior to the booking start time or at the booking start time.

Instead of waiting until a predetermined number of minutes before the booking start time before commencing allocation, the number of minutes may for instance be a function of the number of bookings with start times in a certain period, the number of 'available' vehicles in the fleet and/or the density of 'available' vehicles in the area of the pickup location.

Moreover, the disclosure of the present application should be understood to include any novel features or any novel combination of features either explicitly or implicitly disclosed herein or in any generalisation thereof and during prosecution of the present application or of any application derived therefrom, new claims may be formulated to cover any such features and/or combination of such features.

Claims

1. A computer-implemented method of allocating a vehicle resource to a future vehicle requirement, the method comprising:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles comprises fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).

2. A method as claimed in claim 1, wherein step h) comprises allocating the vehicle determined to have had the best score on the last performance of step d) to the first booking.

3. A method as claimed in claim 1 or claim 2, comprising performing step e) for some but not all of the vehicles in the pool of vehicles formed at step c).

4. A method as claimed in any preceding claim, wherein calculating a journey time from the first location comprises calculating a journey time from a current location of the vehicle to the pickup location included in the first booking.

5. A method as claimed in any of claims 1 to 3, wherein calculating a journey time from the first location comprises calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking.

6. A method as claimed in claim 1, comprising choosing between:

1) calculating a journey time from the first location by calculating a journey time from a current location of the vehicle to the pickup location included in the first booking, and

2) calculating a journey time from the first location by calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking,

based on a status of the vehicle.

7. A method as claimed in any preceding claim, wherein the determining that a current time has a predetermined relationship with respect to the start time of the first booking comprises determining that the current time is at least approximately a predetermined time before the start time of the first booking, for instance is between 15 and 40 minutes before the start time of the first booking.

8. A method as claimed in claim 7, wherein the predetermined time is dependent on the pickup location for the first booking.

9. A method as claimed in any preceding claim, wherein forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises rejecting vehicles of the fleet that have a status that is inconsistent with the vehicle potentially being able to fulfil the booking.

10. A method as claimed in any preceding claim, wherein forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking.

11. A method as claimed in claim 10, wherein forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises calculating a direct

distance between vehicles of the fleet and the pickup location included in the first booking.

12. A method as claimed in claim 10 or claim 11, wherein forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking in terms of the direct distance between the candidate vehicles of the fleet and the pickup location included in the first booking.

13. A method as claimed in claim 12, wherein the predetermined number is between 10 and 100.

14. A method as claimed in any preceding claim, wherein step g) comprises performing steps c), d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking

15. A method as claimed in any preceding claim, wherein calculating a score that is related to the suitability of the vehicle to fulfil the booking comprises calculating sub scores for each of plural factors and performing a mathematical operation on the sub scores.

16. A method as claimed in any preceding claim, wherein forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking and wherein calculating a journey time from the first location to the pickup location included in the first booking comprises using a map database.

17. A method as claimed in any preceding claim, wherein calculating a journey time from the first location to the pickup location included in the first booking comprises using a map database.

18. A method as claimed in any preceding claim, wherein calculating a journey time from the first location to the pickup location included in the first booking comprises using historical or current data about traffic delays on one or more routes between the current location of the vehicle and the pickup location included in the first booking

19. A method as claimed in any preceding claim, wherein calculating a journey time from the first location to the pickup location included in the first booking comprises using historical or current data about traffic delays or average speeds on one or more routes between the first location and the pickup location included in the first booking to calculate a fastest route between the first location and the pickup location included in the first booking

20. A method as claimed in any preceding claim, wherein determining whether there is a need to allocate the first booking comprises determining whether a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

21. A method as claimed in any preceding claim, wherein determining whether there is a need to allocate the first booking comprises determining whether the sum of an allocation buffer value and a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

22. A method as claimed in any preceding claim, comprising

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating the vehicle determined to have had the best score on last performance of step d) to the first booking.

23. Apparatus for allocating a vehicle resource to a future vehicle requirement, the apparatus comprising means configured for:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles comprises fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).

24. Apparatus as claimed in claim 23, wherein the means configured for performing step h) comprises means configured for allocating the vehicle determined to have had the best score on the last performance of step d) to the first booking.

25. Apparatus as claimed in claim 23 or claim 24, comprising means configured for performing step e) for some but not all of the vehicles in the pool of vehicles formed at step c).

26. Apparatus as claimed in any of claims 23 to 25, wherein the means configured for calculating a journey time from the first location comprises means configured for calculating a journey time from a current location of the vehicle to the pickup location included in the first booking.

27. Apparatus as claimed in any of claims 23 to 25, wherein the means configured for calculating a journey time from the first location comprises means configured for calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking.

28. Apparatus as claimed in claim 23, comprising means configured for choosing between:

1) calculating a journey time from the first location by calculating a journey time from a current location of the vehicle to the pickup location included in the first booking, and

2) calculating a journey time from the first location by calculating a journey time from a drop off location of a customer booking that the vehicle is currently fulfilling to the vehicle to the pickup location included in the first booking,

based on a status of the vehicle.

29. Apparatus as claimed in any of claims 23 to 28, wherein the means configured for determining that a current time has a predetermined relationship with respect to the start time of the first booking comprises means configured for determining that the current time is at least approximately a predetermined time before the start time of the first booking, for instance is between 15 and 40 minutes before the start time of the first booking.

30. Apparatus as claimed in claim 29, wherein the predetermined time is dependent on the pickup location for the first booking.

31. Apparatus as claimed in any of claims 23 to 30, wherein the means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises means configured for rejecting vehicles of the fleet that have a status that is inconsistent with the vehicle potentially being able to fulfil the booking.

32. Apparatus as claimed in any of claims 23 to 31, wherein the means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises means configured for selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking.

33. Apparatus as claimed in claim 32, wherein the means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises means configured for calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking.

34. Apparatus as claimed in claim 32 or claim 33, wherein the means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises means configured for selecting a predetermined number of candidate vehicles that are geographically closest to the pickup location included in the first booking in terms of the direct distance between the candidate vehicles of the fleet and the pickup location included in the first booking.

35. Apparatus as claimed in claim 34, wherein the predetermined number is between 10 and 100.

36. Apparatus as claimed in any of claims 23 to 35, wherein the means configured for performing step g) comprises means configured for performing steps c), d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking

37. Apparatus as claimed in any of claims 23 to 36, wherein the means configured for calculating a score that is related to the suitability of the vehicle to fulfil the booking comprises means configured for calculating sub scores for each of plural factors and performing a mathematical operation on the sub scores.

38. Apparatus as claimed in any of claims 23 to 37, wherein the means configured for forming a pool of plural candidate vehicles for possible fulfilment of the first booking comprises means configured for calculating a direct distance between vehicles of the fleet and the pickup location included in the first booking and wherein the means configured for calculating a journey time from the first location to the pickup location included in the first booking comprises means configured for using a map database.

39. Apparatus as claimed in any of claims 23 to 38, wherein the means configured for calculating a journey time from the first location to the pickup location included in the first booking comprises means configured for using a map database.

40. Apparatus as claimed in any of claims 23 to 39, wherein the means configured for calculating a journey time from the first location to the pickup location included in the first booking comprises means configured for using historical or current data about traffic delays on one or more routes between the current location of the vehicle and the pickup location included in the first booking

41. Apparatus as claimed in any of claims 23 to 40, wherein the means configured for calculating a journey time from the first location to the pickup location included in the first booking comprises means configured for using historical or current data about traffic delays or average speeds on one or more routes between the first location and the pickup location included in the first booking to calculate a fastest route between the first location and the pickup location included in the first booking

42. Apparatus as claimed in any of claims 23 to 41, wherein the means configured for determining whether there is a need to allocate the first booking comprises means configured for determining whether a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

43. Apparatus as claimed in any of claims 23 to 42, wherein the means configured for determining whether there is a need to allocate the first booking comprises means configured for determining whether the sum of an allocation buffer value and a calculated journey time from the first location to the pickup location included in the first booking for a vehicle determined in step d) to be better than scores for other vehicles in the pool of candidate vehicles meets a predetermined relationship with respect to a time remaining from the current time to the start time of the first booking.

44. Apparatus as claimed in any of claims 23 to 43, comprising means configured for:

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating the vehicle determined to have had the best score on last performance of step d) to the first booking.

45. A computer program comprising machine readable instructions that when executed by computing apparatus causes it to perform the method of claim of claims 1 to 22.

46. A non-transitory computer-readable storage medium having stored thereon computer-readable code, which, when executed by computing apparatus, causes the computing apparatus to perform a method of allocating a vehicle resource to a future vehicle requirement, the method comprising:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles may comprise fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).

47. Apparatus having at least one processor and at least one memory having computer-readable code stored thereon which when executed controls the at least one processor to allocate a vehicle resource to a future vehicle requirement by performing:

a) storing plural vehicle requirements in the form of plural customer bookings, each customer booking including at least a start time and a pickup location;

b) in respect of a first customer booking, automatically determining that a current time has a predetermined relationship with respect to the start time of the first booking;

c) in response to the determining, forming a pool of plural candidate vehicles for possible fulfilment of the customer booking, wherein the pool of candidate vehicles may comprise fewer than the number of vehicles that are available to fulfil the booking;

d) for each vehicle in the pool of candidate vehicles, calculating a score that is related to the suitability of the vehicle to fulfil the booking;

e) for one or more of the vehicles in the pool of candidate vehicles, calculating a journey time from a first location to the pickup location included in the first booking;

f) using both:

1) the scores calculated for the vehicles of the pool of candidate vehicles, and

2) the calculated one or more journey times,

to determine whether there is a need to allocate the first booking;

g) performing steps d), e) and f) at least two times for the first booking until it is determined at step f) that there is a need to allocate the first booking; and

h) following g), allocating to the first booking the vehicle determined to have had the best score or one of the best scores on last performance of step d).