



- (51) International Patent Classification:
G01C 21/16 (2006.01)
- (21) International Application Number:
PCT/US2014/030174
- (22) International Filing Date:
17 March 2014 (17.03.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/832,456 15 March 2013 (15.03.2013) US
- (71) Applicant: CAMBRIDGE MOBILE TELEMATICS [US/US]; 1 Broadway, 14th Floor, Cambridge, Massachusetts 02142 (US).
- (72) Inventors: GIROD, Lewis David; 67 Scituate St, Arlington, Massachusetts 02476 (US). BALAKRISHNAN, Hari; 9 Preble Gardens Road, Belmont, Massachusetts 02478 (US). MADDEN, Samuel Ross; 46 Newbury Street, Newton, Massachusetts 02459 (US).
- (74) Agent: ROHLICEK, J. Robin; Floor 2, 321 Summer Street, Boston, Massachusetts 02210 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]

(54) Title: INFERENCE OF VEHICULAR TRAJECTORY CHARACTERISTICS WITH PERSONAL MOBILE DEVICES

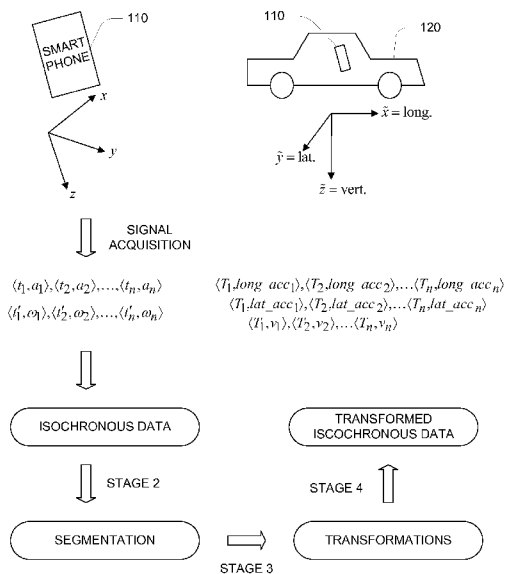



FIG. 3

(57) Abstract: Accurate longitudinal acceleration, lateral acceleration (perpendicular to the principal direction of motion, and velocity, is inferred by processing raw data from a commodity three-axis accelerometer that may be oriented arbitrarily in a moving vehicle (or carried by a moving user), and whose orientation and position may change arbitrarily during the motion. The approach is applicable to a range of applications, including insurance telematics, driver behavior and risk assessment, and road surface quality assessment.



Published:

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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

INFERENCE OF VEHICULAR TRAJECTORY CHARACTERISTICS
WITH PERSONAL MOBILE DEVICES

Cross-Reference to Related Applications

[001] This application claims the priority of U.S. Application No. 13/83,456, filed March 15, 2013. The contents of the application are hereby incorporated by reference in its entirety.

[002] This application is related to U.S. Patent 8,457,880, titled “Telematics Using Personal Mobile Devices,” issued June 4, 2013, which is incorporated herein by reference.

Background

[003] This invention relates to inference of vehicular trajectory characteristics with mobile computing devices, and in particular to inference of vehicular acceleration and velocity using commodity devices such as smartphones. The acceleration computed includes a longitudinal component in the principal direction of movement of the vehicle, and may be positive, zero, or negative (indicating deceleration), as well as a lateral component perpendicular to the longitudinal in the plane of movement. A noteworthy aspect of the invention is that it uses raw sensor data from a mobile device that is not in a known or fixed orientation in the vehicle, and allows the orientation and placement to change from time to time. The invention is battery-efficient and well-suited for battery-efficient operation on personal mobile devices such as smartphones.

[004] Inference of characteristics of a vehicle’s trajectory, including velocity and acceleration, has previously been addressed using inertial measurement and positioning systems. With inertial systems, accelerometers (e.g., one or multiple axis or degree of freedom linear and/or rotational accelerometers) that are rigidly placed or fastened to the vehicle in a known position and orientation have been used to determine velocity by numerically integrating acceleration signals. Positioning systems, such as the Global Positioning System (GPS), have been used to infer velocity. Integration of accelerometer signals may have technical limitations, for example, related to “drift” whereby a velocity or position estimate may have progressively increasing error. Inference of velocity characteristics using a positioning system may have technical limitations such that inference or detection of acceleration or rapid changes in velocity may be impossible because the reported velocity

estimates may be smoothed, preventing sudden acceleration and deceleration events from being inferred.

[005] Personal mobile devices, for example, smartphones, personal digital assistants, and the like may have integrated accelerometers and positioning system receivers. Use of such devices for inference of a vehicle's trajectory may have limitations, for example, related to an uncontrolled, unpredictable, and changing placement and orientation of the device relative to the vehicle, and related to high power consumption required for frequent acquisition of position estimates from a positioning system, particularly GPS.

Summary

[006] In a general aspect, this document discloses a method to produce accurate longitudinal acceleration (in the principal direction of motion), lateral acceleration (perpendicular to the principal direction of motion), and velocity, by processing raw data from a commodity three-axis accelerometer that may be oriented arbitrarily in a moving vehicle (or carried by a moving user), and whose orientation and position may change arbitrarily during the motion. This method is applicable to a range of applications including insurance telematics, driver behavior and risk assessment, and road surface quality assessment.

[007] In another aspect, in general, a method for obtaining accurate longitudinal and lateral acceleration information, as well as velocity information, in a battery-efficient way, from a commodity mobile device such as a smartphone or mobile phone, without requiring a known or fixed placement and orientation. Additionally, the method includes automatically calibrating or adapting to variations between phones and the inevitable errors in the raw acceleration data caused by internal imperfections and external factors such as temperature. In some examples, this calibration and adaptation is achieved by processing three-axis accelerometer data augmented with occasional GPS-based velocity samples.

[008] In another aspect, in general, a method for producing accurate longitudinal and lateral acceleration, as well as velocity estimates, from a three-axis accelerometer whose orientation and placement are unknown such that the orientation and placement may change during a drive because the device may be moved relative to the user or vehicle. In some examples, the method includes automatically inferring rest periods to estimate the orientation of the device relative to the direction of gravity, as well as the longitudinal axis of the vehicle, and use that to produce the desired results. In some

examples, to estimate velocity, the method employs integration of corrected longitudinal acceleration values.

[009] In another aspect, in general, a method for producing accurate longitudinal and lateral acceleration, as well as velocity estimates, from a three-axis accelerometer, guarding against spurious results caused by the movement of the device relative to the vehicle. The method includes inference of such movements and automatically detecting them as not corresponding to legitimate vehicular or user acceleration. In some examples, this method uses the three-axis acceleration data, optionally three-axis gyroscope data, and optionally information from the mobile phone about the current foreground application and whether a phone call is in progress. Additionally, this method can be used to detect when the user is engaged in potentially risky driving behavior, such as texting or dialing while driving.

[010] In another aspect, in general, a method for classifying periods of time into three categories: (1) periods of time when the mobile device is not moving in a reference frame of the vehicle and the vehicle is in motion, and (3) periods of time when the mobile device is moving in a reference frame of the vehicle, is being carried by the user, or is not in the vehicle at all, and (3) periods of time when the mobile device is not moving in the reference frame of the vehicle and the vehicle is not moving. In some examples, the acceleration data is classified or segregated hierarchically into these three categories to produce accurate longitudinal and lateral acceleration and velocity estimates for periods of time corresponding to a particular category or categories.

[011] In another aspect, in general, a method for producing accurate longitudinal and lateral acceleration, as well as velocity estimates, from a three-axis accelerometer even when the device traverses hills and slopes of unknown grades, across a range of speeds. In some examples, the method further handles variations in the values reported by the accelerometer (there are many reasons for such drift, including changes in the device's temperature across its typical operating range). In some examples, the method periodically or continuously estimates the gravity vector in the reference frame of the three-axis accelerometer, providing a method to correct for the slope on which the vehicle or user is traveling.

[012] In another aspect, in general, a method for estimating rapid lateral and longitudinal acceleration for forensic purposes, such as to detect the direction and magnitude of collisions with the vehicle. In some examples, a collision report that represents rapid changes in acceleration or deceleration is output, and that report may

include an estimate of the angle and speed prior to and/or at which the collision occurred.

[013] In another aspect, in general, a method for combining velocity generated according to aspects outlined above with estimates of velocity produced by other means, such as from a positioning system (e.g., GPS) in devices on the phone or car. Periodic and/or on-demand estimates of longitudinal velocity from such means are combined with accelerometer-based estimates to increase accuracy of velocity estimates. In some examples, the timing of these estimates is determined in part by an estimate of the uncertainty in the velocity estimate derived only from acceleration.

[014] In another aspect, the processing of the data to infer vehicular characteristics occurs after the entire drive has been completed, or on several minutes of data at a time.

[015] In another aspect, in general, high-frequency acceleration data is pre-processed on a mobile device to produce summaries, which are transmitted to a server for further processing.

[016] Other aspects include a device (e.g., a smartphone) configured to perform a method outlined above, or software (e.g., a smartphone app) stored on a non-transitory machine-readable medium for causing a data processor (or a client/server architected processing system) to perform the steps of a method outline above.

[017] Advantages of one or more aspects include improved battery usage in a mobile device, for example, by limiting use of a energy-intensive positioning system such as GPS, limited or requiring no prior calibration or training for a device, and not requiring human input or intervention to produce accurate estimates of longitudinal and lateral acceleration and velocity. Other advantages include the ability of the accurate estimates to improve the quality of trajectory inference described in U.S. Patent 8,457,880, titled "Telematics Using Personal Mobile Devices."

[018] Other features and advantages of the invention are apparent from the following description, and from the claims.

Description of Drawings

[019] FIG. 1 is system diagram.

[020] FIG. 2 is a block diagram of a smartphone.

[021] FIG. 3 is a signal flow diagram.

Description

Overview

[022] Referring to FIG. 1, an embodiment of a telematics system is described below in the context of use of a personal mobile device 110, such as a smartphone, that travels in a vehicle 120, to determine characteristics of a vehicle's trip through a road network, and in particular, to track velocity and acceleration profile through the road network. In this embodiment, the smartphone is not necessarily rigidly attached to the vehicle frame of reference; for example, the smartphone may be in the pocket of a driver or passenger in the vehicle, and may be moved from time to time, changing its placement and orientation. However, it should be understood that in other embodiments, other devices, which may be fixed or embedded in a vehicle, could be used in similar ways. In addition, the techniques may be applied to other tasks, for example, tracking other forms of transportation such as bicycling, bus travel, train travel, pedestrian walking, running, etc.

[023] Referring to FIG. 2, the mobile device has inertial sensors 260, for example, a three-axis accelerometer or a 6-degree-of-freedom (DOF) inertial measurement system (three-axis accelerometer and a gyroscope) fixed to the smartphone's frame of reference. Even though such inertial sensors are not fixed to the vehicle, they can nevertheless provide useful information about the vehicle's motion. The mobile device optionally has and makes use of yet other sensors 262, such as a compass, as well as a microphone, which may provide some information about the vehicle, for example, based on engine noise. The device 110 also includes a processor and storage 270, which in some embodiments performs various computations described below. The mobile device 110 optionally has a position measurement capability, for instance, using radios 240 that receive and/or send wireless signals with a fixed communication system 140 (see FIG. 1) such as cellular position estimation using a cellular radio, and/or a WiFi receiver for WiFi position estimation, and/or a positioning sensor 230 which receives position information with a positioning system reference 130 (see FIG. 1) such as via a Global Positioning System (GPS) receiver.

[024] Generally, the smartphone executes an application that performs at least the task of acquiring the sensor data and possibly part (or all) of the task of inferring the vehicle's acceleration and/or velocity as it travels through a road network. This inference is based on the outputs of the three-axis accelerometers 260, optionally aided by outputs of the gyroscope, positioning system 230, and other sensors 262. In some examples, the techniques described in U.S. Application Serial No. 13/687,479, titled "Telematics Using Personal Mobile Devices," are used to augment the inferred

velocity and/or acceleration information that is based on the accelerometers. In some examples, the results of the method described herein are used to augment the trajectory and mileage methods described in U.S. Application Serial No. 13/687,479.

[025] The output of a three-axis accelerometer 260 is a time-series sequence of the form

$$\langle t_1, a_1 \rangle, \langle t_2, a_2 \rangle, \dots, \langle t_n, a_n \rangle$$

where each t_i is a distinct timestamp and each a_i is a three-tuple $\langle a_{xi}, a_{yi}, a_{zi} \rangle$, giving the x, y, z components of the acceleration sample estimated by the mobile device at time t_i . Note that these components are in the frame of reference defined by the mobile device, and will in general not be the same as the frame of reference of the vehicle or be aligned on any of the axes with the earth's gravity vector. In addition, this acceleration time-series may optionally be augmented with a time-series of the form

$$\langle t'_1, \omega_1 \rangle, \langle t'_2, \omega_2 \rangle, \dots, \langle t'_n, \omega_n \rangle$$

where each ω_i gives the angular velocity components of the device as measured by a gyroscope at time t'_i . Further, this data may optionally be augmented with occasional velocity samples (or “fixes”) obtained from an accurate position sensor such as GPS, at certain points in time. In general, the number of such velocity fixes may be much smaller than the number of acceleration samples because obtaining accurate velocity samples may be impossible or consume high power. This data may optionally be augmented with acoustic data from a microphone, and/or with data from a magnetic compass on the mobile device.

[026] The method produces the following output streams:

$$\langle T_1, long_acc_1 \rangle, \langle T_2, long_acc_2 \rangle, \dots, \langle T_n, long_acc_n \rangle$$

$$\langle T_1, lat_acc_1 \rangle, \langle T_2, lat_acc_2 \rangle, \dots, \langle T_n, lat_acc_n \rangle$$

$$\langle T_1, v_1 \rangle, \langle T_2, v_2 \rangle, \dots, \langle T_n, v_n \rangle,$$

which are the time series of the longitudinal acceleration (in the principal direction of the vehicle's motion), lateral acceleration (in the direction perpendicular to the principal direction of the vehicle's motion and on the plane of motion), and velocity of the vehicle.

[027] Referring to FIG. 3, we first provide an overview and summary of the different stages of the method. Then, we describe the internal details of each stage. The method proceeds in four main stages:

- Stage 1: Pre-processing.
- Stage 2: Segmentation.
- Stage 3: Per-segment computation of the vehicle's acceleration in the principal direction of movement and in the perpendicular direction.
- Stage 4: Computation of the vehicle's velocity.

[028] The pre-processing is an optional step that makes the input time-series data isochronous (i.e., equally spaced in time). The reason this step is helpful is because the input data may not be sampled at equally-spaced times, or may have some outages. It is possible to apply the remaining stages without this stage, but pre-processing simplifies subsequent processing.

[029] Stage 2 segments the isochronous data stream, producing non-overlapping segments in which the mobile device has a fixed orientation with respect to the reference frame of the vehicle. Note, however, that this orientation can be arbitrary and has to be inferred by the method. The data in each of these segments must be rotated into the vehicle's reference frame in order to estimate the correct vehicle acceleration and speed in the direction of travel.

[030] Stage 3 runs on each segment determined in stage 2. This stage determines the rotation matrix to convert an acceleration vector in the mobile device's frame of reference to the vehicle's frame of reference. The three axes in the vehicle's frame of reference, when the vehicle is not on an incline, are (i) the direction of gravity, (ii) the longitudinal direction, and (iii) the lateral direction. Because the mobile device has a fixed orientation within each segment (from the previous steps), the desired rotation matrix for a given segment is unique.

[031] Stage 3 has four key steps:

- Step 3.1 Determine the direction of the earth's gravity.
- Step 3.2 Determine the acceleration in the longitudinal direction, i.e., in the "forward" direction aligned with the vehicle's primary direction of motion.
- Step 3.3 Determine the acceleration in the lateral direction.
- Step 3.4 Correct for gradually changing phone orientations that may occur within one segment and for drives on inclines.

[032] Stage 4 takes the longitudinal acceleration results from the previous stage and integrates it to compute vehicular velocity estimates.

- Step 4.1 Identify periods where the vehicle is at rest, using the variance of the acceleration data.
- Step 4.2 Identify other absolute velocity samples (e.g., from GPS or other sensors).
- Step 4.3 Using known points as boundary conditions, estimate velocity as the integral of longitudinal acceleration.
- Step 4.4 Between any two points of known time and location, the estimated velocity is corrected such that the mean velocity matches the mean velocity as determined by the most likely distance based on map data and the difference in time between points.

[033] We now describe the details of each of the above stages.

Stage 1: Pre-processing to make the data isochronous.

[034] There are several ways to make the data isochronous. On mobile device platforms where each sample is tagged with a timestamp, the input data is resampled at a defined rate and missing samples are interpolated. On platforms where sample times are not provided by the system, the method computes the median time difference between sample event arrival times, and then assume that the samples are sampled at that implied rate. It should be noted that several other methods of making the input isochronous may be used, e.g., using the mean time difference, or using the mean or median but ignoring gaps above a certain threshold value.

Stage 2: Segmentation

[035] In this stage the method determines periods of time where the mobile device was moved relative relative to the vehicle. To detect movements of the phone, the method relies on noticing a significant change in the direction of gravity. The method takes a certain number of seconds of acceleration data, and over that window, computes the median value of each of the reported acceleration axes. Note that these median samples may not be from the same sample vector, but nonetheless the median values tend to be from similar vectors. Then the method computes the maximum distance between the median vector and the other sample vectors in the window. If this distance exceeds a threshold, then the device is considered to have been re-positioned. One may determine the threshold empirically or using regression, which is the preferred approach. One could also use online learning with some minimal supervision. When gyroscope data is available from the device, spikes in the gyroscope's time-series data are used to identify the likely times when the device was

moved, and the resulting orientation changed. Similarly, a change in the orientation output of an inertial measurement unit (IMU) algorithm resident on the device also indicates a device movement (an IMU produces a three-axis orientation output -- roll, pitch, yaw -- generally derived from fusing acceleration, gyro, and compass data). If compass data is available, a sudden change in the reported direction (after filtering out noise) is indicative of device movement and orientation change of the device relative to the vehicle's frame of reference. The method uses these sensor inputs when they are available.

[036] Whenever a candidate phone movement event has been detected, whether by a change in gravity, a large spike in angular velocity, or the output of the orientation output of an IMU algorithm, the method corroborates the motion by checking for a significant change in average orientation resulting from the motion. This check is performed averaging the vectors over several seconds before and after the suspected motion. If the angle between these two average vectors is below a threshold, and the original motion detection was small, then the motion detection is tagged as a false positive. This technique allows us to lower the threshold on the detection algorithm and increase our accuracy.

Stage 3: Compute the vehicle's acceleration in the longitudinal and lateral directions.

Step 3.1 Determine the direction of the earth's gravity

[037] There are multiple possible ways to estimate the direction of the earth's gravity. One approach is to compute the average value of all the input vectors in the segment. (The x , y , and z components reported by the mobile device's accelerometer corresponding to the three-dimensional vector.) The average direction thus computed is an estimate of the "down" vector corresponding to the earth's gravity. This method will produce a correct answer if on average the vehicle does not have a net acceleration within the segment, and if on average the vehicle drives on level ground. Another method to estimate the direction of gravity is to compute the average vector during a rest period. This approach is more robust, but can be used only on a subset of all the input data. From time to time, the level assumption may not hold. In these cases, the output of trajectory mapping is used to determine the appropriate way to average the vectors to determine the "down" vector.

Step 3.2 Determine the acceleration in the longitudinal direction

[038] To compute the longitudinal direction, the method rotates the data such that the negative z axis points toward gravity as estimated in step 3.1, and then projects the data onto the x - y plane, and converts the resulting 2-D vectors into polar form $\langle r, \theta \rangle$.

It then analyzes the distribution of θ values for these 2-D acceleration vectors. However, these initial θ values suffer from two sources of error. First, there is invariably some error in the estimate of gravity from step 3.1. This error can cause a bias in the resulting data that appears in the 2-D plane as a changing bias in the origin. In addition, most accelerometers have drift in the origin from various factors including temperature. These errors in the origin value would generally cause the values of θ to be incorrect. To handle this problem, the method computes the derivative of the raw acceleration vectors, and then computes the probability distribution of these derivatives. This step eliminates the (potentially drifted) origin as a factor, capturing the trend in direction more precisely.

[039] The method then filters out “degenerate” values: when more or less at rest, accelerometers often report points with a change only in a single axis (e.g., x and y values are the same as the last point, but the z-value changes). These cases do not produce useful information; they are artifacts of the accelerometer hardware. The method removes these spurious points.

[040] In the resulting distribution, there are typically four peaks about 90 degrees apart, corresponding to the axes of the vehicle. However, depending on the type of driving, the dominant peak may be either forward, reverse, left, or right. The method must therefore determine which of these 4 peaks corresponds to the longitudinal (forward) direction. To resolve this question, consider all points with known velocity (e.g., any available GPS samples with velocity or rest periods). Consider all pairs of points separated by 1 point. Then, for each of these intervals, integrate the acceleration along each of the 4 candidate directions, to estimate a total change in velocity over the interval, and compare it to the actual change in velocity (as determined by the available velocity samples at either end of the interval). For each candidate direction, accumulate the absolute error, weighted by the absolute value of the slope of the actual velocity change (dv/dt). This weighting puts more weight on cases where the change in velocity is more significant. Choose the direction with the lowest total weighted error as the longitudinal direction.

[041] An alternate method to compute the longitudinal direction is to again compute the derivative of the acceleration time series (after eliminating the gravity component). Then, compute the distribution of the angles of these derivatives. The result will show a clear principal component in the forward direction of travel. The method selects the maximum value of this distribution and uses this value as the forward direction.

Step 3.3 Determine the acceleration in the lateral direction

[042] The lateral direction is the direction perpendicular to both gravity and the longitudinal directions. The acceleration component in this direction may be calculated easily because the total magnitude is known, and the values of two of the three components of acceleration have already been estimated.

Step 3.4 Correct for gradually changing phone orientations

[043] In some cases the phone may gradually change orientation, e.g., if it is located on a slippery surface that allows it to rotate. In such cases, Step 3.2 will result in a slowly or suddenly varying orientation value. The method detects this situation by computing the principal rotation over smaller windows.

[044] The method also includes corrections for inclinations encountered during a drive. As the vehicle drives, its inclination will change as it drives over hills and other slopes. As a result, the gravity vector changes over time in the reference frame of the vehicle. This change in the gravity vector will (erroneously) contribute a portion of the acceleration due to gravity to the lateral and longitudinal acceleration estimates, which in turn will cause drift and errors in the estimates. Estimating the true direction of gravity within the vehicle's reference frame is non-trivial because there is no clear distinction between legitimate acceleration of the vehicle and the acceleration component due to inclination. To address this challenge, the method assumes that changes in inclination occur over longer time scales than changes in vehicle acceleration, and that the bulk of the effect of gravity feeds into the longitudinal acceleration because most segments of road are not banked. The method therefore uses a smoothing filter to estimate the effect due to gravity with a large window. However, the method also takes into account rest periods, where the average vector is purely due to gravity. The smoothing function is applied to the longitudinal acceleration between rest intervals, and the portions near a rest interval are adjusted such that the final value of the smoothed function is equal to the unsmoothed value. The smoothed value is subtracted from the longitudinal acceleration to remove bias due to gravity.

Stage 4: Velocity estimation

[045] To estimate velocity, the method takes the longitudinal acceleration results from the previous stage and integrates them to compute vehicular velocity estimates. In this stage, the method first determines all potential absolute velocity samples, performs the integration, and then post-corrects based on map data.

Step 4.1 Determining rest periods

[046] A rest period is an interval of time when the vehicle was at rest, and therefore has an absolute velocity of 0. To determine a rest period, the method takes the acceleration data in windows of a few seconds each. It computes the variance of the 3-dimensional acceleration vectors over each such window. It then computes the distribution of the logarithms of resulting variance values. The log distribution tends to show multiple Gaussian components, which correspond to different vibration regimes experienced by the device. Empirically, vehicles in motion tend to experience variances in a measurable and known small range, so we fit a Gaussian model to the mode centered near that point and determine the fit parameters μ and σ . We then select the “rest-cutoff” value as a function of the μ and σ values and (for example, $\mu - 2\sigma$).

Step 4.2 Identifying other absolute samples

[047] Absolute velocity samples are acquired from other sources, such as periodic GPS samples, or using an audio analysis of engine noise (disclosed below). These are combined with the rest periods to form a complete set of velocity samples.

[048] Dynamic inclusion of velocity data: Drift in velocity estimates grows quickly and “true” velocity reduces additive errors. There is, however, a cost to collecting an absolute reference (in particular an energy cost to use GPS). The method uses an online algorithm to determine when to acquire an absolute reference point. Define a counter at T seconds, and deduct 1 for each passing second. Reset the counter to T seconds whenever an absolute reference is used (e.g. GPS, audio, or rest period). In addition, every time the acceleration transitions from $-a$ to $+a$ or from $+a$ to $-a$ (where a is 1.5 m/s^2 , for example) deduct a penalty from the counter (e.g., 20 seconds). This serves to increase the rate of reference points when there is a substantial swing in acceleration.

Step 4.3 Integrate to estimate velocity

[049] Next consider the absolute velocity samples to be boundary conditions for the velocity estimate. The estimate defines the velocity to be known exactly at these points, and then computes an estimate over the points falling between adjacent velocity samples, based on the numeric integral of longitudinal acceleration. The values of the boundary conditions are used to determine a constant bias to the acceleration that forces the estimate to converge to the known velocity at the boundaries. For any pair of velocity samples spanning a period of unknown acceleration values (e.g. due to missing data), the velocity is not estimated. Once the

method has produced longitudinal acceleration estimates, it computes velocity estimates. The method numerically integrates the acceleration using rest periods and occasional absolute velocity estimates as boundary conditions. These absolute velocity estimates can come from sources such as GPS or using an audio analysis of engine noise (disclosed below). Given two boundary velocities, the method computes the integral of the acceleration between those two points. The total change in velocity computed from the integral is compared to the change of velocity between the boundaries. The difference between these changes is considered to be a uniform drift in acceleration and is deducted or added uniformly to meet the two boundaries. If the velocity estimate ever drops below 0, the minimum velocity is assumed to be an intermediate boundary condition at 0 and the estimates bounded by this new intermediate point are recomputed.

4.4 Correction of velocity estimates using map-matching output from trajectory inference method

[050] A map-matching approach, for example, as described in *U.S. Patent 8,457,880*, returns a most-likely map-matched trajectory corresponding to cellular or WiFi positioning data or intermittent GPS data. The method maps known time and location pairs collected from GPS onto the most-likely trajectory, and the distance along the trajectory between adjacent GPS points is computed. The timestamps associated with a GPS point provide the length of time over which such distance was traveled. Hence, one can compute the map-matched distance between successive points and divide by the time interval between the observations to estimate an average speed. Next, the mean of the estimated velocity is computed and compared to that based on GPS and the map-matched trajectory. If the correction is reasonably small, an adjustment function is applied to the velocity estimates to correct the mean value to match the value based on the map.

[051] This adjustment is made by constructing a curve formed by the function

$$\sqrt{\frac{n}{s-1}} \sqrt{\frac{s-1-n}{s-1}} \text{ for } n = 0, \dots, s-1$$

where s is the number of points in the region. Other functions with a similar shape and properties could also be used.

[052] Scale this function so that the integral of the function is equal to the difference between our estimate and the average velocity determined from the map, and add it into our estimate to correct it.

[053] Finally, differentiate the corrected velocity estimates to construct a matching, corrected acceleration time series.

Audio analysis for velocity determination.

[054] The approaches described above, which make use of accelerometer data, can be augmented using audio data captured by the smartphone. The theory of operation underlying the audio-based velocity determination method is that the frequency of noise associated with driving (engine noise, road noise, etc) is proportional to the velocity of the vehicle. This is because the velocity is proportional to the rotation of the tires and any consistent vibrations that are correlated to wheel rotation will have this property of proportional frequency. This is true of both consistent road textures as well as properties of the engine and drivetrain, although the properties of the drivetrain are generally going to be more consistent over time. These properties of the drivetrain may be correlated to the rotation of the wheels (e.g., noise in the axles, bearings, constant velocity joints (front wheel drive), or transfer case (4-wheel drive). They may also be correlated to the rotation of the engine, e.g., belts, bearings, rotating parts, combustion noise, exhaust noise, transmission noise. Note that depending on the state of the transmission, the rotation of the engine may relate to velocity at a different scale than the wheels. Note also that the particular noise spectrum and signature will vary for different cars. This means that some form of calibration may be required, but on the other hand it also allows for a possible means of vehicle identification, which is often of value in a variety of applications.

[055] Audio analysis methodology. The principle of analysis is to analyze a short-term-fourier-transform (STFT), represented as a spectrogram. First, sample the audio, then filter and downsample it; regardless of the audio sampling rate, apply a k th order butterworth filter at f_1 Hz, then downsample to f_2 Hz. The result is a maximum available frequency of F Hz ($F = 250$ usually), which for privacy concerns is well below speech frequencies, so there should be no privacy concerns. However, the frequency content of engines tends to be well within this range. Once downsampled, compute the STFT using a narrow Hann window that covers W samples, and a wider 1024-point (e.g.) FFT to increase the frequency resolution, and a step of 50 samples to produce 10 spectra per second.

[056] Then apply a frequency-tracking algorithm to locate and track the shifting frequency of a narrowband source in the data. This algorithm begins with the maximum magnitude frequency and tracks it by selecting successive maximum values from within a tight range around the previous estimate. Another approach is to use information from acceleration to guide the tracking; track multiple hypotheses

concurrently using hints from acceleration to guide. This results in a time series of frequency over time, which we assume is proportional to velocity. Another approach is to apply frequency tracking to a power cepstrum of the FFT columns in the STFT. This approach is similar to the “cepstral coefficients” methods commonly used in speech recognition. A power cepstrum computes the FFT power spectrum of the log of the FFT power spectrum of the original signal. This method is used to coalesce harmonic energy and to more readily identify pitch changes.

[057] Take the longitudinal acceleration computed previously and correlate it against short clips of the derivative of the frequency time series data. Select out clips where the correlation is the highest; use a least squares fit to compute the estimated scale factor from the acceleration data (units in m/s^2) to frequency, to determine the ratio to convert frequency to velocity.

[058] Calibration process for audio: collect these “least noisy” frequency traces (i.e. those with highest correlation to acceleration measurements) and capture GPS data points concurrently with the audio. For each “good” trace segment, compute the ratio required to convert frequency to velocity. These ratios should form a distribution with several modes, where each mode corresponds to a gearing ratio. Record these gearing ratios for further use. In addition, collect the source audio data corresponding to each ratio and compute a summation of all of the spectra data for that ratio, first scaling each spectrum according to actual speed. This summation will give you a signature of the engine at that gearing ratio. This signature can be used to enhance the frequency-tracking algorithm as well as to identify the vehicle.

[059] Once the gear ratios have been estimated for the vehicle, the ratios determined by fitting the measured acceleration to the derivative of the tracked frequency can be compared with the known ratios to find the nearest fit; that fit could be used to assume the known gear ratio and deduce the velocity from the frequency track.

[060] The signatures can also be used to deduce the correct gear ratio: given the frequency track, assume a gear ratio and then scale the signature data to match the observed frequency. Then, compute the degree of correlation between the scaled signature spectrum and the observed spectrum. The correct gear ratio would be the one that has the highest correlation with the signature. This can be used to disambiguate cases where the match based on acceleration data is not decisive.

[061] Signatures may also be used directly to improve detection. Rather than locating the maximum frequency magnitude in the tracking algorithm, instead the maximum correlation can be found, using scaled versions of the signature spectra.

Applications

[062] The approaches described above can be used in a variety of applications. In one application, the derived motion data in the frame of reference of the vehicle is used in the task of tracking a vehicle. Note that the vehicle can be an automobile, a bicycle, a boat, etc. For example, in order to limit power or computation cost of using a positioning system such as GPS, the smartphone-derived data is used in a map-based tracking task. In some such examples, both longitudinal velocity data, as well as rotation about the vertical axis (i.e., vehicle turning), are used to track the vehicle's trajectory through the map. In some examples, position estimates are acquired at times that are based on the tracking using the smartphone accelerometer data, for example, when the system determines that there is a high uncertainty of the vehicle's location. In this way, higher accuracy of tracking of a vehicle can be obtained for a given amount of energy consumption (battery life) by avoiding frequent use of relatively power intensive radio transmissions.

[063] In another application, the velocity estimates are recorded on the smartphone, and provided as a summary of the driving characteristics of the user (e.g., by periodic or on-demand uploading of summary data). In one example, the data may be used in various types of commercial relationships (e.g., contracts, usage-based or behavior-based insurance policies, etc.) established between one party (e.g., the driver and/or the owner of a vehicle) and another party (e.g., an insurance carrier, a vehicle owner, etc.), where the terms of the contract depend on the driving characteristics that are determined using the techniques described above.

Implementations

[064] Implementations of the approaches described above may use software that includes instructions stored on non-transitory computer-readable media. The instructions are used to control operation of a data processor (e.g., a computer, a virtual machine) to perform some or all of the functions described above. For examples, such instructions control operation of a processor in a smartphone and/or operation of a server computer in communication with the smartphone. In some implementations, the instructions a smartphone are provided (e.g., as an "app") over a wireless network for installation on the smartphone.

[065] It is to be understood that the foregoing description is intended to illustrate and not to limit the scope of the invention, which is defined by the scope of the appended claims. Other embodiments are within the scope of the following claims.

What is claimed is:

1. A method comprising inferring vehicular trajectory characteristics, said inferring comprising:
 - acquiring data representing a sequence of acceleration-related measurements of a device in an unknown and changeable orientation traveling in a vehicle;
 - forming a plurality of time segments of the acceleration-related data;
 - for each time segment of the plurality of time segments
 - determining data characterizing an orientation of the device relative to a frame of reference of the vehicle,
 - forming motion-related data in the frame of reference of the vehicle during the time segment, and
 - determining data characterizing trajectory characteristic of the vehicle in the time segment; and
 - storing data characterizing trajectory characteristic of the vehicle determined from the data characterizing the trajectory characteristics in multiple of the time segments.
2. The method of claim 1 wherein the device comprises a personal mobile device that is held or is on the person of a driver or passenger of the vehicle.
3. The method of claim 1 wherein forming the plurality of time segments comprises determining said segments such that the orientation of the device relative to the vehicle is relatively fixed as compared to changes of the orientation at transitions between segments.
4. The method of claim 1 wherein, for each time segment of the plurality of time segments, the method further comprises determining whether there is time interval in the time segment in which the vehicle is stationary.
5. The method of and of claims 1 through 4 wherein forming the one or more time segments comprises detecting times of substantial relative motion of the device relative to the vehicle, and forming the time segments between said detected times.

6. The method of any of claims 1 through 5 wherein determining data characterizing an orientation of the device comprises determining a fixed orientation of the device relative to the frame of reference of the vehicle.
7. The method of any of claims 1 through 5 wherein determining data characterizing an orientation of the device comprises determining a time varying orientation of the device relative to the frame of reference of the vehicle.
8. The method of any of claims 1 through 7 wherein determining data characterizing an orientation of the device in a time segment comprises first determining a vertical direction.
9. The method of claim 8 further comprising using a distribution of sensed motion in a plane orthogonal to the vertical direction to determine a direction of travel of the vehicle.
10. The method of claim 9 wherein determining a time interval in which the vehicle is stationary uses a threshold picked from statistics of the acceleration related measurements.
11. The method of any of claims 1 through 10 wherein the data characterizing trajectory characteristics of a vehicle comprise at least one of longitudinal acceleration, lateral acceleration, longitudinal velocity, and distance travelled of the vehicle.
12. The method of any of claims 1 through 11 wherein acquiring the data representing the sequence of acceleration-related measurements of the device comprises measuring acceleration values on multiple axes using accelerometers in the device.

13. The method of any of claims 1 through 12 wherein the inferring further comprises:
- accessing map data characterizing a road network;
 - processing the data characterizing the trajectory characteristics of the vehicle according to the map data to determine data characterizing traversal of the road network by the vehicle.
14. The method of any of claims 1 through 12 further comprising applying map traversal data in determining the data characterizing the trajectory characteristics of the vehicle.
15. The method of any of claims 1 through 14 wherein the inferring further comprises:
- acquiring data representing environmental characteristics of the device; and
 - wherein determining the trajectory characteristic of the vehicle includes using the data representing the environmental characteristics.
16. The method of claim 15 wherein the data representing environmental characteristics includes data representing an acoustic environment of the vehicle, where said data provides information related to velocity of the vehicle.
17. The method of any of claims 1 through 16 wherein the inferring further comprises:
- acquiring data from a gyroscope on the device; and
 - wherein determining the trajectory characteristic of the vehicle includes using the data from the gyroscope.
18. The method of any of claim 1 through 17 further comprising acquiring data from a gyroscope on the device and wherein forming of the time segments includes using the data from the gyroscope on the device to identify times at which an orientation of the device changed.

19. The method of claim 17 or claim 18 wherein the data from the gyroscope is used in the forming of the motion-related data in the frame of reference of the vehicle during the time segment.

20. The method of any of claims 1 through 19 wherein the inferring further comprises:

acquiring position or velocity data of the device using a radio or positioning sensor in the device; and

wherein determining the trajectory characteristic of the vehicle includes using the position data.

21. The method of claim 20 further comprising determining when to acquire the position or velocity data at least in part according to the determined trajectory characteristics determined from the acceleration-related measurements of a device.

22. The method of claim 20 further comprising determining when to acquire the position or velocity data at least in part according to the determined motion state of the device in the frame of reference of the vehicle.

23. The method of claim 20 further comprising:

accessing map data characterizing a road network;

processing the data characterizing the trajectory characteristics of the vehicle according to the map data to determine data characterizing traversal of the road network by the vehicle.

24. The method of any of claims 1 through 23 further comprising providing software to the device for execution on the device for causing a processor in the device to perform the inferring.

25. The method of any of claims 1 through 24 further comprising providing the data characterizing trajectory characteristic of the vehicle via a radio interface of the device.

26. The method of claim 25 wherein providing the data includes providing the data to a server computer and at least some of the steps of claim 1 are performed at said server.
27. The method of any of claims 1 through 26 further comprising providing the data characterizing the trajectory characteristics to a second party for use in usage-based insurance applications.
28. The method of any of claims 1 through 26 further comprising providing the data characterizing the trajectory characteristics to a second party with a contractual relationship with a first party comprising at least one of a user of the device and party in control of the vehicle.
29. The method of claim 28 wherein the contractual relationship is governed at least in part according trajectory characteristics of the vehicle.
30. The method of claim 28 further comprising forming the contractual relationship between the first party and the second party.
31. Software stored on a non-transitory machine-readable medium comprising instructions for causing a processor to perform all the steps of any of claims 1 through 30.
32. A personal mobile device configured to perform all the steps of any of claims 1 through 30.

400129

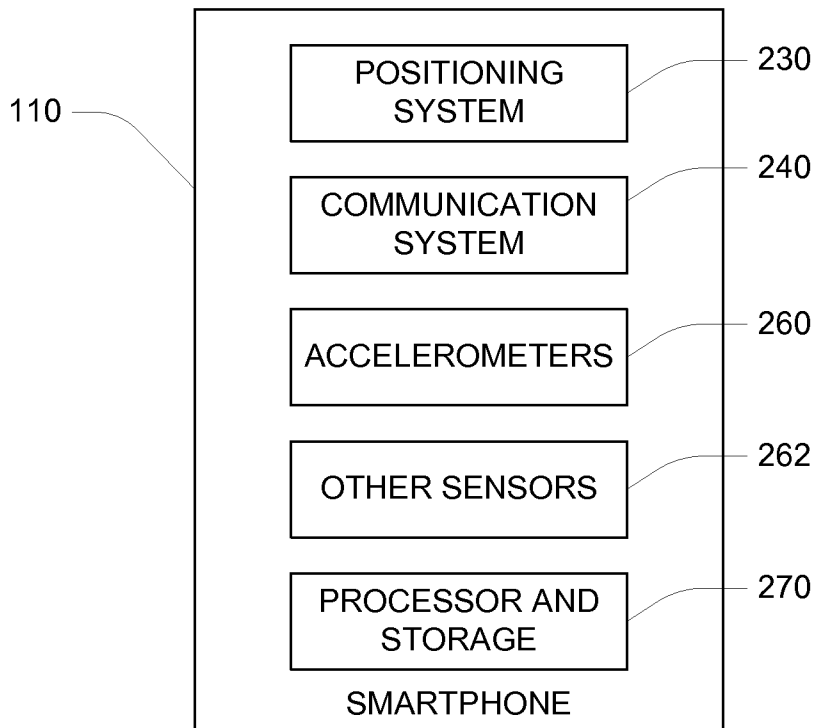



FIG. 2

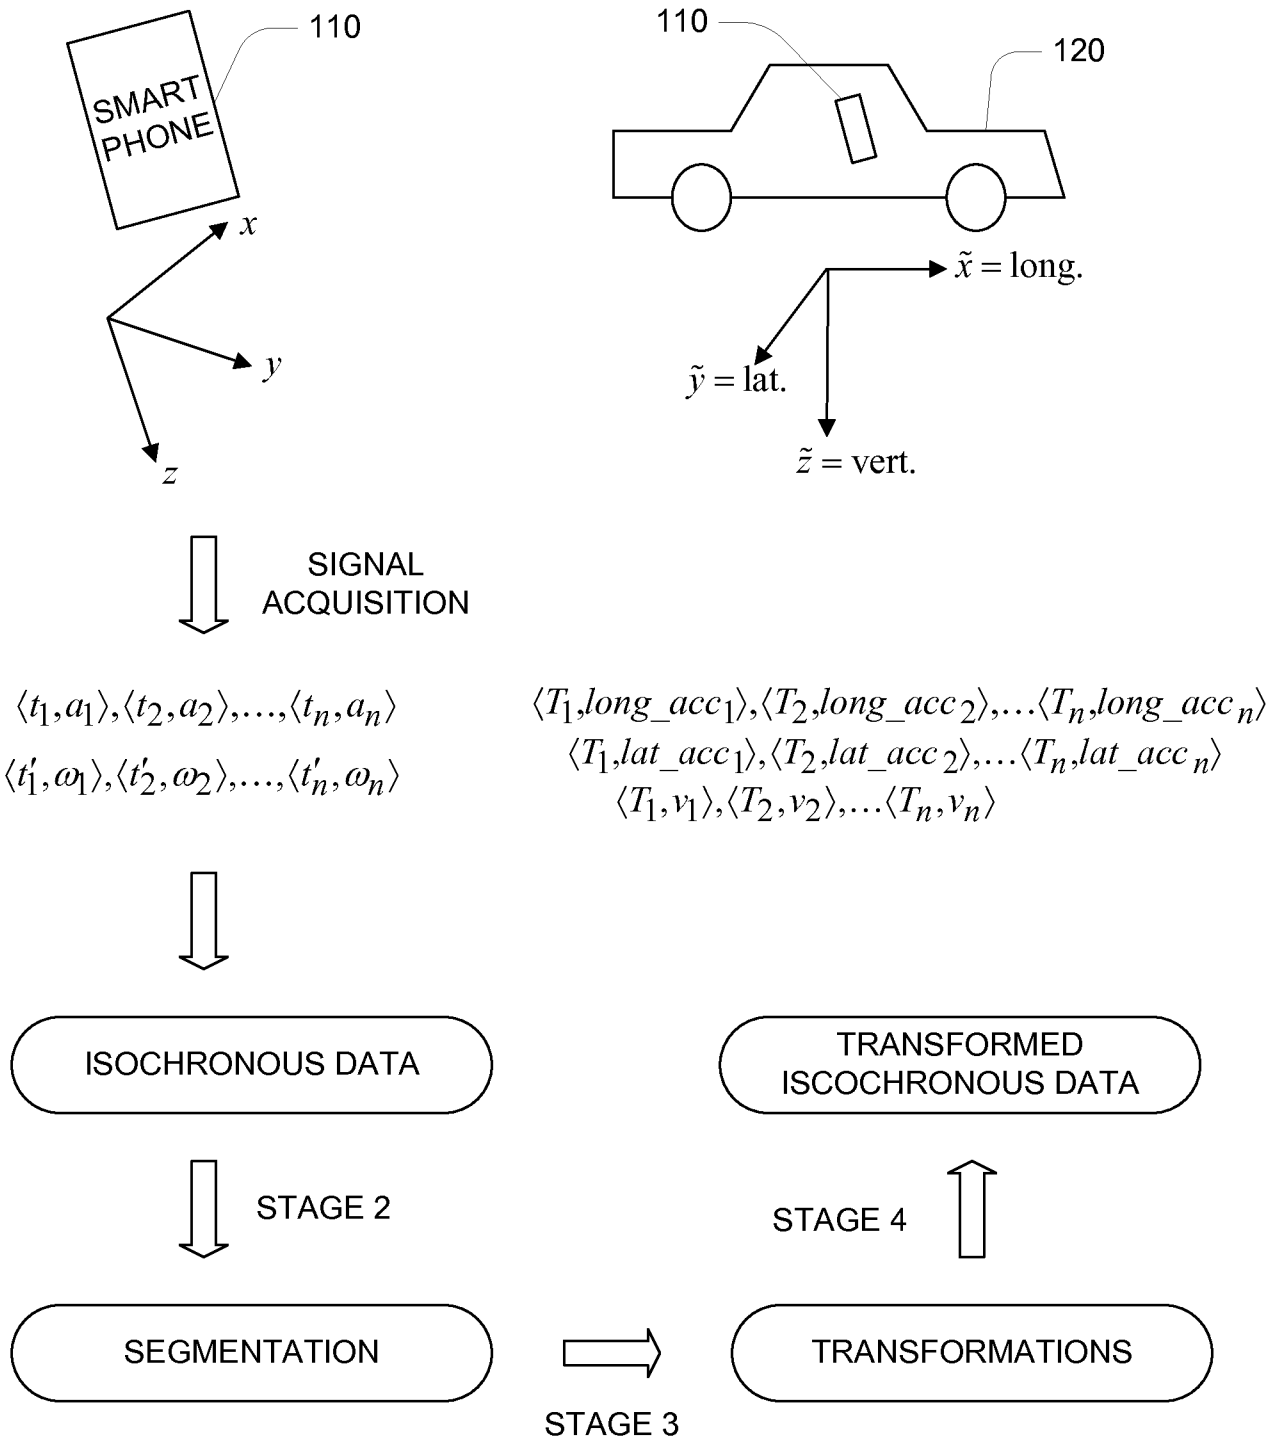


FIG. 3

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/030174

A. CLASSIFICATION OF SUBJECT MATTER
INV. G01C21/16
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G01C G01S

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 02/18873 A2 (MAGELLAN DIS INC [US]; BEGIN JOHN [US]; CHEOK KA C [US]) 7 March 2002 (2002-03-07)	1-4,6-9, 11-15, 17-26, 31,32
Y	paragraph [0001] - paragraph [0029] -----	5,10,16
X	EP 2 469 229 A1 (MAGNETI MARELLI SPA [IT]) 27 June 2012 (2012-06-27)	1,4, 8-12,14, 15,21, 22,25, 26,31,32
Y	paragraph [0005] paragraph [0029] - paragraph [0078] -----	5,10,16
Y	US 2012/029811 A1 (GARIN LIONEL J [US]) 2 February 2012 (2012-02-02) paragraph [0043] ----- -/--	5

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 19 August 2014	Date of mailing of the international search report 26/08/2014
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer van Doornum, A

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/030174

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2013/041585 A1 (CZOMPO JOSEPH [US] ET AL) 14 February 2013 (2013-02-14) paragraph [0033] -----	5
Y	WO 01/71372 A2 (YEOMAN GROUP PLC [GB]; AGNEW HUGH [GB]; GEAKE VINCENT [GB]; WOOLLEY RO) 27 September 2001 (2001-09-27) page 15, line 7 - page 25, line 16 -----	16
A	US 2012/245839 A1 (SYED ZAINAB [CA] ET AL) 27 September 2012 (2012-09-27) the whole document -----	1-26,31, 32

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/030174

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.: 27-30
because they relate to subject matter not required to be searched by this Authority, namely:
Article 17(2)(a)(i) jo. 17(2)(b) jo. Rule 39.1(iii) PCT : The claims relate to schemes, rules and/or methods for doing business.
2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2014/030174

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 0218873	A2	07-03-2002	AU 8707901 A
			EP 1315945 A2
			WO 0218873 A2

EP 2469229	A1	27-06-2012	EP 2469229 A1
			RU 2011152911 A
			US 2012330497 A1

US 2012029811	A1	02-02-2012	CN 103052864 A
			EP 2601479 A1
			JP 2013538345 A
			KR 20130059405 A
			US 2012029811 A1
			WO 2012018695 A1

US 2013041585	A1	14-02-2013	US 2013041585 A1
			WO 2013025348 A1

WO 0171372	A2	27-09-2001	AU 3942901 A
			WO 0171372 A2

US 2012245839	A1	27-09-2012	CA 2769788 A1
			EP 2503288 A2
			US 2012245839 A1
