

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
19 February 2004 (19.02.2004)


PCT

(10) International Publication Number
WO 2004/015369 A2

- (51) International Patent Classification⁷: **G01C**
- (21) International Application Number:
PCT/US2003/025194
- (22) International Filing Date: 11 August 2003 (11.08.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/402,178 9 August 2002 (09.08.2002) US
- (71) Applicant (for all designated States except US): **INTER-SENSE, INC.** [US/US]; 1 North Avenue, Burlington, MA 01803 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **FOXLIN, Eric** [US/US]; 285 Highland Avenue, Arlington, MA 02476 (US).
- (74) Agents: **FEIGENBAUM, David, L.** et al.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: TRACKING, AUTO-CALIBRATION, AND MAP-BUILDING SYSTEM



(57) Abstract: A navigation or motion tracking system includes components associated with particular sensors, which are decoupled from a tracking component that takes advantage of information in the sensor measurements. The architecture of this system enables development of sensor-specific components independently of the tracking component, and enables sensors and their associated components to be added or removed without having to re-implement the tracking component. In a software implementation of the system, sensor-specific software components may be dynamically incorporated into the system and the tracking component is then automatically configured to take advantage of measurements from the corresponding sensors without having to modify the tracking component.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

TRACKING, AUTO-CALIBRATION, AND MAP-BUILDING SYSTEM

Cross-Reference to Related Applications

[01] This application claims the benefit of U.S. Provisional Application No. 60/402,178, filed August 9, 2002, titled "Localization, Auto-Calibration, and Map-Building," the contents of which are incorporated herein by reference.

Background

[02] This invention relates to tracking, navigation, pose estimation, localization, auto-calibration, scene modeling, structure-from-motion and/or map-building based on sensor inputs.

[03] Tracking or navigation systems often make use of measurements from sensors to aid in determining a location ("localization") or an orientation (attitude and heading) or a pose (position and orientation) of an object such as a person, a vehicle or a robot as it navigates in an environment, such as within the bounds of a building. A variety of types of sensors are available for such systems, including sensors that measure a relative location between a sensor and a target. An example of such a sensor/target combination is an acoustic emitter (target) and a microphone array (sensor) that can determine a direction of arrival of an acoustic signal broadcast from the emitter. Different types of sensors measure different aspects of the relative pose of a sensor and a target, such as a range, direction, or relative orientation. Different sensors may have different measurement characteristics that affect the mapping between the relative pose of a sensor and a target and the measurement values provided by the sensor. These characteristics can include uncertainty or noise characteristics of the measurement values.

[04] Systems have been developed that use Kalman Filtering techniques to incorporate information in sensor measurements to track the position or orientation of an object, typically also using information about the dynamic characteristics of the object. The implementation of such Kalman Filtering techniques is often complex, and typically requires detailed knowledge of the measurement characteristics of the specific sensors used in tracking the object.

[05] Some navigation systems perform simultaneous localization and mapping (SLAM), also known in the field of computer vision as structure-from-motion (SfM). The mapping aspect

relates to determining the locations of fixed landmarks or beacons in the environment while at the same time using sensor measurements from those fixed landmarks to assist in localization of the object. As an example, when a robot navigates an uncharted territory, such as in a Mars rover mission, or in an underground mining or undersea operation, the robot may determine its location relative to the surrounding environment. If a complete map of the terrain is not available in advance, the robot may observe landmarks, build a map based on the landmark observations, and determine its location on the map that it has constructed so far. The landmarks may be man-made markers or natural features of the terrain.

[06] As another example, an automated factory may use robots to move materials and products among different locations. Beacons, such as ultrasound emitters, or graphic markers having special patterns, may be placed at various locations in the factory. The robots may have sensors, such as ultrasound receivers, laser range finders, cameras, or pattern recognition devices, for determining their positions relative to reference points in the factory environment. The locations of the reference points may not be known in advance, so the robots may update their maps of the factory based on inputs from the sensors, and navigate through the factory based on their updated maps.

[07] It may be desirable to also perform automatic calibration of sensors during the ongoing process of localization of the object. For example, various types of sensors may have different types of calibration parameters, such as measurement biases and scale factors. Examples of calibration parameters are focal lengths or distortion parameters of a camera lens, or alignment of a camera relative to the vehicle carrying it. The Kalman Filter implementation may estimate the sensor calibration parameters using a common infrastructure that is used to determine the location of the vehicle. As with the localization and mapping approaches, the characteristics of the calibration parameters are typically reflected in the implementation of the Kalman Filter techniques. Some systems combine localization, mapping, and auto-calibration.

Summary

[08] In a general aspect, the invention features a navigation or motion tracking system in which components associated with particular sensors are decoupled from a tracking component that takes advantage of information in the sensor measurements. The architecture of this system

enables development of sensor-specific components independently of the tracking component, and enables sensors and their associated components to be added or removed without having to re-implement the tracking component. In a software implementation of the system, sensor-specific software components may be dynamically incorporated into the system and the tracking component is then automatically configured to take advantage of measurements from the corresponding sensors without having to modify the tracking component.

[009] In general, in one aspect, the invention features a method for tracking an object that includes coupling a sensor subsystem to an estimation subsystem. The sensor subsystem enables measurement related to relative positions or orientations of sensing elements. Configuration data is accepted from the sensor subsystem, and the estimation system is configured according to the accepted configuration data. The method includes repeatedly updating a state estimate, including accepting measurement information from the sensor subsystem, and updating the state estimate according to the accepted configuration data and the accepted measurement data.

[010] This and other aspects of the invention may include one or more of the following features.

[011] Coupling the sensor subsystem to the estimation subsystem includes coupling software modules each associated with one or more of the sensing elements.

[012] Each of the software modules provides a software interface for receiving information related to an expected sensor measurement and providing measurement information that depends on the received information.

[013] Each of the software modules implements calculations that are independent of a representation of the state in the estimation subsystem.

[014] The state estimate characterizes an estimate of a location of the object.

[015] The state estimate characterizes configuration information for one or more sensing elements fixed to the object.

[016] The configuration information for the one or more sensing elements fixed to the object includes information related to position or orientation of the sensing elements relative to the object.

[017] The configuration information for the one or more sensing elements fixed to the object includes operational parameters for the one or more sensing elements.

[018] The state estimate characterizes configuration information for one or more sensing elements fixed in an environment of the object.

[019] The configuration information for one or more sensing elements fixed in the environment of the object includes a map of the locations of the sensing elements.

[020] Repeatedly updating the state further includes providing to the sensor subsystems information related to an expected sensor measurement, and wherein accepting the measurement information from the sensor subsystem includes accepting information related to an actual sensor measurement.

[021] Providing the information related to an expected sensor measurement includes providing information related to a relative geometric configuration of two of the sensing elements.

[022] Providing information related to a relative geometric configuration of the two of the sensing elements includes providing information characterizing a relative location of the sensing elements.

[023] Accepting the information related to an actual sensor measurement includes accepting information enabling the estimation subsystem to calculate a difference between the actual measurement and the expected measurement.

[024] Accepting the information related to an actual sensor measurement includes accepting information for correlating measurements and geometric relationships between sensing elements.

[025] The information for correlating measurements and geometric relationships between sensing elements includes a mapping between a relative pose of the sensing elements and a sensor measurement.

[026] The mapping between the relative pose of the sensing elements and the sensor measurement characterizes a linear mapping.

[027] Accepting the information related to an actual sensor measurement includes accepting information characterizing an uncertainty in the actual measurement.

[028] The information characterizing the uncertainty in the actual measurement includes parameters of a statistical distribution of an error of the actual measurement.

[029] Repeatedly updating the state further includes selecting a pair of sensing elements for measurement, and providing an identification of the selected pair to the sensing subsystem.

[030] Selecting the pair of sensing elements includes selecting the elements according to an expected utility of a measurement associated with the elements to the updating of the state.

[031] Repeatedly updating the state further includes updating the state according to the accepted information related to an actual sensor measurement.

[032] Repeatedly updating the state further includes updating the state according to accepted measurements from inertial sensors.

[033] Updating the state estimate includes applying a Kalman Filter approach.

[034] Each of the sensing elements includes at least one of a sensor and a target.

[035] The target includes an active device that interacts with the sensor.

[036] The target includes at least one of a man-made signal reflector and a natural feature of an environment.

[037] The object is selected from a group consisting of a vehicle, a robot, a person, a part of a person, a flying object, a floating object, an underwater moving object, an animal, a camera, a sensing apparatus, a helmet, a tool, a piece of sports equipment, a shoe, a boot, an article of clothing, a personal protective equipment, and a rigid object having a dimension between 1 nanometer to 109 meters.

[038] The state estimate includes information related to a position or an orientation of the object relative to a reference coordinate frame.

[039] In general, in another aspect, the invention features a tracking system includes an estimation subsystem, and a sensor subsystem coupled to the estimation subsystem. The sensor subsystem is configured to provide configuration data to the estimation subsystem and to provide measurement information to the estimation subsystem for localizing an object. The estimation

subsystem is configured to update a location estimate for the object based on configuration data and measurement information accepted from the sensor subsystem.

[040] This and other aspects of the invention may include one or more of the following features.

[041] The sensor subsystem includes one or more sensor modules, each providing an interface for interacting with a corresponding set of one or more sensing elements.

[042] The interface enables the sensor module to perform computations independently of an implementation of the estimation subsystem.

[043] The interface enables the estimation subsystem to perform computations independently of an implementation of the sensor modules.

[044] The tracking system also includes a navigation subsystem to navigate the object in an environment based on the location estimate for the object.

[045] In general, in another aspect, the invention features a sensor module that includes a sensor interface for communicating with a measurement sensor, and a communication interface for communication with an estimation system. The sensor module is configured to receive information related to an expected sensor measurement over the communication interface, receive a measurement signal over the sensor interface, and provide measurement information based on the measurement signal over the communication interface.

[046] This and other aspects of the invention may include one or more of the following features.

[047] The sensor module is configured to provide information over the communication interface related to an uncertainty in the measurement information.

[048] The received information related to an expected sensor measurement includes a predicted pose of a sensing element relative to the measurement sensor.

[049] In general, in another aspect, the invention features a method that includes enumerating a set of sensing elements available to a tracking system that includes an estimation subsystem that estimates a position or orientation of an object, and providing parameters specific to the set of

sensing elements to the tracking system to enable the estimation subsystem to be configured based on the parameters specific to the subset of sensing elements.

[050] This and other aspects of the invention may include one or more of the following features.

[051] The method includes generating a sequence of candidates of pairs of sensing elements selected from the set of sensing elements, the sequence based on an expected utility of a measurement associated with the elements to the estimation subsystem.

[052] The method includes selecting a pair of sensing elements from the sequence of candidates, the selected pair of sensing elements being ready to make a measurement at the time of selection of the pair or at a predefined time after the time of selection of the pair, the selected pair having a highest expected utility of a measurement among the sequence of candidates.

[053] The set of sensing elements includes at least one sensor and at least one target, the sensor making a measurement with respect to the target.

[054] The target includes a natural feature in an environment.

[055] In general, in another aspect, the invention features a method that includes computing an estimate of a pose of a target element relative to a sensor element based on an estimate of a pose of a tracked object relative to an environment having affixed thereto either the sensor element or the target element. The computing of the estimate of the pose of the target element relative to the sensor element is also based on an estimate of a pose of the affixed element relative to the tracked object and the other element relative to the environment. The method also includes computing an estimate of a measurement of the target made by the sensor based on the estimate of the pose of the target relative to the sensor, making an actual measurement of the target by using the sensor, computing a deviation between the actual measurement and the estimated measurement, and generating a new estimate of the pose of the tracked object based on the deviation.

[056] This and other aspects of the invention may include one or more of the following features.

[057] The method also includes computing a first observation matrix that characterizes a linearized model of a function relating the measurement made by the sensor to the pose of the target relative to the sensor.

[058] The method also includes computing a second observation matrix that characterizes a linearized model of a function relating the pose of the target relative to the sensor to the estimate of the pose of the tracked object relative to the environment.

[059] The also includes computing an observation matrix that characterizes a linearized model of a function relating the measurement made by the sensor to the pose of the tracked object relative to the environment by combining the first observation matrix and the second observation matrix.

[060] In general, in another aspect, the invention features a method that includes estimating a first value associated with a pose of a first sensing element relative to a second sensing element. The first sensing element is fixed to an environment and the second sensing element is fixed to an object being tracked. One of the first and second sensing elements is a sensor and the other is a target. The method includes estimating a second value associated with a pose of the second sensing element relative to the first sensing element, determining which of the first and second sensing elements is the sensor, and generating an innovation of a measurement of the target made by the sensor based on the first value when the second sensing element is the sensor.

[061] This and other aspects of the invention may include one or more of the following features.

[062] The method also includes generating the innovation based on the second value when the first sensing element is the sensor.

[063] Estimating the first value and estimating the second value are performed by a process ignorant of which of the first and second sensing elements is a sensor.

[064] In general, in another aspect, the invention features a method that includes estimating a calibration parameter of a sensing element that is either a sensor or a target, the sensing element being fixed either to an environment or to an object being tracked. The method includes determining whether the sensing element is the sensor or the target, assigning the calibration

parameter as a sensor calibration parameter when the sensing element is a sensor, and generating an innovation of a measurement of a target made by the sensing element based in part on the sensor calibration parameter.

[065] This and other aspects of the invention may include one or more of the following features.

[066] The method also includes assigning the calibration parameter as a target calibration parameter when the sensing element is a target, and generating an innovation of a measurement of the sensing element made by a sensor based in part on the target calibration parameter.

[067] Estimating the calibration parameter is performed by a process ignorant of whether the sensing element is a sensor or a target.

[068] In general, in another aspect, the invention features a method of using multiple sensors in a tracking system. The method includes providing an estimation subsystem, coupling one or more sensor modules to the estimation subsystem, each associated with a different set of one or more sensors. The method includes configuring the tracking system, which includes providing configuration information from each of the sensor modules to the estimation subsystem regarding the characteristics of the sensors associated with the sensor module, and configuring the estimation subsystem using the provided configuration information. The method includes maintaining estimates of tracking parameters in the estimation subsystem, including repeatedly passing data based on the estimates of the tracking parameters from the estimation subsystem to one or more of the sensor modules, receiving from the one or more sensor modules at the estimation subsystem data based on measurements obtained from the associated sensors, and the data passed to the sensor modules, and combining the data received from the one or more sensor modules and the estimates of the tracking parameters in the estimation subsystem to update the tracking parameters.

[069] This and other aspects of the invention may include one or more of the following features.

[070] The data passed from the estimation subsystem to one or more of the sensor modules includes an estimate of the pose of a target relative to a sensor that was calculated by the

estimation subsystem using an estimate of the pose of a tracked object relative to a frame of reference fixed to an environment.

[071] The data passed from the estimation subsystem to one or more of the sensor modules does not include the estimate of the pose of the tracked object relative to the frame of reference fixed to the environment.

[072] Providing the estimation subsystem includes providing a module that is configurable to use different sets of sensor modules coupled to it.

[073] Maintaining estimates of the tracking parameters in the estimation subsystem includes using a stochastic model in the estimation subsystem.

[074] Using a stochastic model includes implementing some or all of a Kalman filter in the estimation subsystem.

[075] Implementing some or all of the Kalman filter includes updating error estimates using linearized models of the sensor system.

[076] Implementing some or all of the Kalman filter includes implementing a distributed Kalman filter, wherein each of a plurality of components of the distributed Kalman filter is associated with a different subset of the sensor modules.

[077] One of the components of the distributed Kalman filter is associated with a subset of sensor modules consisting of sensor modules that are affixed to a tracked object.

[078] One of the components of the distributed Kalman filter is associated with a subset of sensor modules consisting of sensor modules which are affixed to an environment.

[079] One of the components of the distributed Kalman filter is not associated with any sensor modules.

[080] Implementing the distributed Kalman filter includes implementing a Federated Kalman Filter.

[081] Providing configuration information from the sensor modules includes providing information characterizing a type of a sensor associated with a sensor module.

[082] Providing configuration information from the sensor modules includes providing information characterizing a position or an orientation of a sensor associated with a sensor module.

[083] Providing configuration information from the sensor modules includes providing information characterizing one or more calibration parameters of a sensor associated with a sensor module.

[084] In general, in another aspect, the invention features a machine-accessible medium, which when accessed results in a tracking or navigation system that tracks or navigates, respectively, an object, performing operations that includes enumerating a set of sensing elements available to the tracking or navigation system, where the sensing elements available to the tracking or navigation system includes at least one of an inside-out sensor and an outside-in sensor. The inside-out sensor is fixed to the object and makes measurements with respect to a target fixed to an environment. The outside-in sensor is fixed to the environment and makes measurements with respect to a target fixed to the object. The machine-accessible medium, when accessed, results in the tracking or navigation system configuring an estimation module of the tracking or navigation system based on an enumeration of the set of sensing elements available to the tracking or navigation system so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, or a combination of inside-out and outside-in sensors depending on the sensors available. The machine-accessible medium, when accessed, results in the tracking or navigation system repeatedly updating an estimated pose of an object based on measurements from the set of sensing elements available to the tracking or navigation system.

[085] This and other aspects of the invention may include one or more of the following features.

[086] The sensing elements available to the tracking or navigation system include range sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, range sensors, or any combination of the above sensors.

[087] The sensing elements available to the tracking or navigation system include inertial sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, inertial sensors, or any combination of the above sensors.

[088] The sensing elements available to the tracking or navigation system include dead reckoning sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, dead reckoning sensors, or any combination of the above sensors.

[089] In general, in another aspect, the invention features a tracking or navigation method that includes receiving sensor configuration information indicating a set of sensing elements available to a tracking or navigation system, and configuring a data processing module of the tracking or navigation system based on the sensor configuration information to selectively perform one of (a) receiving data from at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the inside-out bearing sensor, (b) receiving data from at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the outside-in bearing sensor, and (c) receiving data from at least one inside-out bearing sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the outside-in bearing sensor and the inside-out bearing sensor.

[090] This and other aspects of the invention may include one or more of the following features.

[091] The tracking or navigation method also includes configuring the data processing module to selectively perform one of (d) receiving data from at least one range sensor, and updating an estimated pose of an object based on data received from the range sensor, (e) receiving data from at least one range sensor and at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the range sensor and the inside-out bearing sensor, (f) receiving data from at least one range sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the range sensor and the outside-in bearing sensor, and (g) receiving data from at least one range sensor, at least one

outside-in bearing sensor, and at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the range sensor, the inside-out bearing sensor, and the outside-in bearing sensor.

[092] In general, in another aspect, the invention features an apparatus that includes an estimation module to estimate a pose of an object based on measurement data from sensing elements, the estimation module configured to enable selective performance of (a) receiving data from at least one inside-out bearing sensor, and updating an estimated pose of an object based on the data received from the inside-out bearing sensor, (b) receiving data from at least one outside-in bearing sensor, and updating an estimated pose of an object based on the data received from the outside-in bearing sensor, and (c) receiving data from at least one inside-out bearing sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on the data received from the outside-in bearing sensor and the inside-out bearing sensor.

[093] In general, in another aspect, the invention features an apparatus that includes an estimation module to estimate a pose of an object based on measurement data from sensing elements, the estimation module configured to enable selective performance of one of (a) updating an estimate of the position or orientation of the object relative to an environment, (b) updating an estimate of the position or orientation, relative to the object, of at least one sensing element fixed to the object, and (c) updating an estimate of the position or orientation, relative to the environment, of at least one sensing element fixed in the environment.

[094] In general, in another aspect, the invention features a computer program that is capable of performing any combination (i.e., one, two, three, four, five, ..., etc.) of a set of functions, and selects which function or functions to perform at runtime based on configuration information. In this aspect of the invention, the set of functions can feature any subset of the following: (i) receive data from one or more inside-out bearing sensors; (ii) receive data from one or more outside-in bearing sensors; (iii) receive data from one or more range sensors; (iv) receive data from one or more inertial or dead reckoning sensors; (v) receive data from one or more inside-out bearing sensors and one or more outside-in bearing sensors; (vi) receive data from one or more bearing sensors and one or more range sensors; (vii) update an estimate of the position or orientation of a tracked object relative to a reference coordinate frame; (viii) update estimates of the position or orientation relative to the tracked object of one or more sensors or targets affixed

to the tracked object; (ix) update estimates of the position or orientation relative to the reference coordinate frame of one or more sensors or targets fixed in the reference coordinate frame; (x) update estimates of calibration parameters of one or more sensors; and (xi) update estimates of size or color or amplitude or frequency or pattern or category of one or more targets.

[095] Implementations of the invention may include one or more of the following features.

[096] The configuration information is obtained from configuration files stored on a computer readable medium accessible to the computer program.

[097] The configuration information describing a particular sensor is obtained from a nonvolatile electronic memory device physically attached to said sensing device.

[098] The computer program is capable of performing more than one of the functions in groups (i) to (vi).

[099] The computer program is capable of performing more than one of the functions in groups (vii) to (xi).

[0100] The computer program is capable of simultaneously performing more than one of the functions in groups (i) to (vi).

[0101] The computer program is capable of simultaneously performing more than one of the functions in groups (vii) to (xi).

[0102] The computer program is capable of simultaneously performing one or more of the functions in groups (i) to (vi), and one or more of the functions in groups (vii) to (xi).

[0103] Aspects of the invention may include one or more of the following features.

[0104] The system is implemented in whole or in part in software.

[0105] The system provides an application programming interface to a software application that makes use of tracking or mapping data.

[0106] The application programming interface to the application does not necessarily expose the details of the sensor configuration.

[0107] Sensor modules are implemented in software.

[0108] The sensor modules each provides an application programming interface that is used by the system to communicate with the sensor modules.

[0109] The sensor module application programming interface does not depend on the specific characteristics of the set of sensors associated with that sensor module.

[0110] The sensor modules are implemented as dynamically loaded software modules.

[0111] The sensor modules are automatically detected by the system.

[0112] The sensor modules provide sensor parameters to a central estimation subsystem.

[0113] The sensor parameters include parameters that identify a basic type of a sensor, such as 2-D bearing, 1-D bearing, range, etc.

[0114] The sensor parameters include parameters that identify a specific type of a sensor, such as make and model, etc.

[0115] The sensor parameters are provided according to the IEEE 1451.2 Smart Transducer Interface standard.

[0116] The sensor module provides measurement related parameters for the sensors.

[0117] The measurement-related parameters include parameters of an observation model.

[0118] The parameters of an observation model include an observation matrix suitable for a Kalman Filter.

[0119] The estimation subsystem automatically configures a Kalman Filter using the provided sensor parameters.

[0120] The estimation subsystem repeatedly provides state-related information to the sensor modules.

[0121] The state-related information includes location information.

[0122] The location information includes location information related to a location of a target relative to a sensor.

[0123] The location information includes location information related to a fixed location in an environment.

[0124] The state-related information includes bias parameter information.

[0125] The sensor modules repeatedly accept the state-related information and combine that information with sensor measurement information, and pass the combined information to the estimation subsystem.

[0126] The combined information includes innovation data for the Kalman Filter.

[0127] The estimation subsystem implements an Extended Kalman Filter.

[0128] The estimation subsystem implements a distributed Kalman Filter.

[0129] The distributed Kalman Filter includes separate components associated with different sets of one or more sensing elements.

[0130] The distributed Kalman Filter includes a Federated Kalman Filter.

[0131] The estimation subsystem includes a sensor fusing module for combining information from different sets of sensors.

[0132] The estimation subsystem includes a localization filter.

[0133] The estimation subsystem includes a mapping filter.

[0134] The estimation subsystem includes a simultaneous localization and mapping filter.

[0135] The estimation subsystem includes an auto-calibration filter.

[0136] The estimation subsystem includes a simultaneous localization and auto-calibration filter.

[0137] The estimation subsystem includes a simultaneous localization, auto-calibration, and mapping filter.

[0138] The estimation subsystem includes a selector for selecting one of the localization, and auto-calibration, and mapping related filters.

[0139] The system independently switches between simultaneous tracking and auto-calibration, simultaneous tracking and mapping, tracking only, etc.

[0140] The estimation subsystem performs auto-calibration during an initial interval.

[0141] The estimation subsystem performs auto-calibration while in a pre-surveyed area for which map information is available.

[0142] The estimation subsystem stops performing auto-calibration after the initial interval, or after it leaves the pre-surveyed area.

[0143] Aspects of the invention enable a framework in which a user can rapidly implement a wide variety of user-configurable multi-sensor fusion systems for tracking, auto-calibration, auto-mapping, or any combination of the above. Advantages of one or more aspects of the invention may include one or more of the following:

[0144] Sensor versatility: The architecture allows use of various combinations of types and qualities of sensors -- inside-out, outside-in, mixed, with or without inertial or dead-reckoning sensors.

[0145] Plug and track functionality: The architecture allows users to configure systems by plugging together self-identifying and self-describing smart sensor modules.

[0146] Scalability: The architecture enables deployment of systems with large numbers of tracked vehicles and/or large maps.

[0147] Flexibility: The architecture permits implementation of systems with one on-board processor per vehicle, or tracking systems in which a regional processor tracks all vehicles in the area.

[0148] Algorithm versatility: The architecture defines interfaces between modules to allow individual modules to be updated to take advantage of the newest large-scale simultaneous localization and mapping algorithms. The architecture allows the use of sub-map partitioning.

[0149] Function versatility: The navigation system may switch between simultaneous tracking and auto-calibration, simultaneous tracking and map-building, tracking only, or simultaneous tracking and map-building and auto-calibration.

[0150] Object-oriented design: The architecture facilitates reuse of the modules (e.g., PSE drivers or update filter modules) when the top-level navigation system design is changed.

[0151] Other features and advantages of the invention are apparent from the following description, and from the claims.

Description of Drawings

[0152] FIG. 1 shows a vehicle navigating an environment.

[0153] FIGs. 2 and 3 are block diagrams of a navigation system.

[0154] FIG. 4 shows different frames of reference.

[0155] FIG. 5 shows an overall architecture of a navigation system.

[0156] FIG. 6 is a block diagram of a measurement management unit.

[0157] FIG. 7 is a block diagram of software modules of the navigation system.

[0158] FIG. 8 shows a file directory structure used by the navigation system.

[0159] FIG. 9 shows a flow diagram of logic performed by a configuration manager program.

[0160] FIG. 10 shows a block diagram of an architecture for a multiple vehicle navigation system.

[0161] FIG. 11 illustrates applications of the navigation system.

[0162] FIG. 12 shows a distributed Kalman filter.

Description

1 System overview

[0163] Referring to FIGs. 1 and 2, a navigation system 90 tracks or navigates a vehicle 100 in an environment (or “galaxy”) 106, such as within a factory building. Navigation system 90 (see FIG. 2) includes sensors 103 that provide measurement data and a data processing unit 190 that processes data provided by the sensors. The navigation system 90 builds a map of the

environment, calibrates the sensors, and determines the location of the vehicle in the galaxy frame of reference. The navigation system 90 tracks the position and the orientation (together referred to as the 6-dimensional “pose”) of vehicle 100 based on both inertial measurements as well as sensor measurements between sensing devices or targets in the vehicle 100 and sensing devices or targets that are fixed in the environment 106.

[0164] The terms “environment” and “galaxy” is used interchangeably in the description below. The environment can be one, two, or three dimensional. For example, environment 106 may be a track that winds through a factory. Environment 106 may be a land mass, an ocean floor, a factory floor, a room, a building, a town, an airspace, an ocean body, or an underground tunnel. Environment 106 may be moving relative to earth. For example, the environment may be the interior of an aircraft carrier or a space surrounding a space shuttle orbiting the earth. Vehicle 100 may be, for example, a land vehicle, a water vehicle, an aircraft, a spacecraft, a robot, a person, a part of a person, a flying object, a floating object, an underwater moving object, an animal, a camera, a weapon, a handheld object, a sensing apparatus, a helmet, a tool, a medical instrument, a display, a piece of sports equipment, a shoe, a boot, an article of clothing, a personal protective equipment, or some other object. The terms vehicle and navigation are used for simplicity, but should not be construed to limit the scope of the invention, which is equally intended for tracking any object for any purpose. Therefore, the terms “navigation system” and “tracking system” will be used interchangeably in the description below, as will the terms “vehicle” and “tracked object”.

[0165] The pose of the vehicle is estimated in a fixed frame of reference in environment 106, which is referred to as the galaxy frame of reference or the navigation frame of reference (“N-frame”). In one version of the system, the galaxy frame of reference is an inertial frame of reference. Inertial sensors 104, which are part of an inertial measurement unit (IMU), provide measurements for inertial tracking of the vehicle by an inertial navigation subsystem.

[0166] Navigation system 90 makes use of two general classes of sensor measurements. One class of sensor measurements, called interoceptive measurements, are used for dead reckoning. Interoceptive measurements are made within the vehicle, without reference to any elements external to the vehicle. This class includes inertial measurements (e.g., from linear or angular accelerometers or from gyroscopes) and relative displacement measurements (e.g., from wheel

encoders or optical flow sensors). In the discussion below, one or more inertial measurement units (IMUs) provide the dead-reckoning measurements to the navigation system. However, the system architecture has been designed to support the substitution of other interoceptive sensors, so the abbreviation IMU may also be interpreted as “interoceptive measurement unit” where the context permits.

[0167] A second class of sensor measurements, called exteroceptive measurements, provide measurements between elements on the vehicle and elements external to it. Within this class, some measurements are relative to the reference frame (“earth referenced” measurements), such as GPS, altimeter, compass, and gravitometer measurements. Other sensor measurements are related to pairs of sensors and targets that are fixed to the vehicle and targets or sensors fixed in the environment (“map-referenced” measurements).

[0168] The sensors and targets fixed to the vehicle or fixed in the environment are collectively referred to as “pose sensing elements” (PSEs). The PSEs fixed to the vehicle 100 are referred to as “mobile” PSEs (MPSEs) 102 and the PSEs fixed in the environment 106 are referred to as “fixed” PSEs (FPSEs) 108. Sensor measurements generally relate to measurements between pairs of PSEs (indicated by the dashed lines), one MPSE and one FPSE, and depend on aspects of the relative (6-dimensional) pose of the two PSEs.

[0169] As introduced above, PSEs generally fall into two groups: sensors and targets. Examples of sensors include cameras and microphones. Examples of targets include target images (e.g., a bulls-eye) and acoustic beacons (sound emitters). It is not necessary that the sensor PSEs be fixed to the vehicle and targets fixed in the environment. For example, sensors may also be fixed in the environment such as tracking cameras fixed in a factory environment that track movements of vehicles moving on the factory floor.

[0170] Targets may be active or passive. An example of an active target is an ultrasound speaker that emits ultrasound signals. An example of a passive target is a circular marking pasted on the ceiling of a room, or a natural landmark such as a corner of a building, or rocks in a terrain.

[0171] For earth-referenced sensors (such as a GPS receiver), the sensor is an MPSE, with no specific target FPSE being associated with measurements by that sensor.

[0172] In various vehicle configurations, a wide variety of sensors and targets may be used. Navigation system 90 is modularized and configurable such that different sensors and targets can be selected for a particular vehicle without necessarily having to re-implement and re-test the implementation of the navigation system. The architecture provides a “plug and track” capability in which sensors and targets and their associated software drivers can be “plugged” into the navigation system 90, which then makes use of the sensors and targets in tracking the vehicle. For example, an acoustic sensor such as a directional microphone may be “plugged” into the navigation system, and range and directionality characteristics of the microphones are automatically used by the navigation system 90 in tracking the vehicle 100.

[0173] Navigation system 90 also has the capability of automatically determining values of PSEs’ intrinsic calibration parameters as well as estimating the extrinsic calibration parameters (poses) of PSEs (and optionally the IMUs) in a parent frame of reference (B-frame for MPSEs and N-frame for FPSEs). Examples of sensor intrinsic calibration parameters are camera lens distortion parameters or additive biases or scale factor error parameters in a range measuring system. The term “auto-calibration” is generally used in the discussion below to include estimating the intrinsic and/or extrinsic parameters of mobile sensors or targets relative to the vehicle, while the term “mapping” is generally used in the discussion below to include estimating the intrinsic or extrinsic parameters of fixed sensors or targets relative to the environment. This is slightly different from common usage, in which estimating the calibration parameters of sensors, even those fixed in the environment, would usually be called calibration rather than mapping.

[0174] When an MPSE is installed on vehicle 100, the information on the pose of the MPSE relative to the vehicle is not necessarily provided to the system or may be manually entered to limited accuracy by a user. While in auto-calibration mode, as the navigation system processes sensor inputs using that MPSE, the pose of the MPSE is updated. Such updating can also compensate for changes in pose of an MPSE, for example, because of a shifting of the MPSE during use of the vehicle.

2 Pose Sensing Elements (PSEs)

[0175] The navigation system 90 supports a wide variety of types of sensors to be used in tracking and navigation applications, each sensor having a different performance, operating environment, range, size, cost, update rate, and penetration capability.

[0176] As introduced above, PSEs provide two types of measurements. One type of measurements can be described as “map-point-referenced discrete update observables.” Such measurements are made with respect to point references in the environment, and contribute directly to the environment mapping process. For example, ultrasonic sensors, Lidar (Light Detection And Ranging), radar, and raw GPS “pseudorange” data can be used to measure ranges. Laser scanners, linear CCD (charge coupled device) sensors, and phased array sensors can be used to measure one-dimensional bearing angles. Imaging sensors, PSD (position sensitive device) sensors, quadcells, and pan/tilt servos can be used to measure two-dimensional bearing angles. Doppler radar and phase-coherent acoustic sensors can be used to detect range rates. TDOA (time difference of arrival) RF or acoustic sensors and radar can be used to measure range differences. Active source magnetic trackers and electric field trackers can be used to measure dipole field components.

[0177] A second type of measurements provided by PSEs can be described as “earth-referenced discrete update observables.” Such measurements are made with respect to physical properties of the environment. For example, processed GPS signals and altimeters can be used to measure Cartesian positions or aspects thereof. Magnetic compasses and gyrocompasses can be used to measure headings. Magnetometers and gravimeters can be used to measure fields that vary by location in the environment.

[0178] In the description below, when performing a bearing angle measurement, the “sensor” is defined as the device at the fulcrum of the bearing angle, even if it is actually a source rather than a detector of a signal. Thus, in the case where a laser scanner (or other light projector) and a photodetector are used to measure a bearing angle, the “sensor” is the laser, and the “target” is the photodetector. For an “inside-out” measurement, the sensor is attached to the vehicle and the target is fixed in the environment. For an “outside-in” measurement, a sensor fixed in the environment observes a target fixed to the moving vehicle.

[0179] For map-point referenced PSEs, a measurement z_{ij} between a particular MPSE j and FPSE i depends on the relative pose λ_{ij} of the sensor and the target. In general, this relationship takes the form of a non-linear function $z_{ij} = h_{ij}(\lambda_{ij})$. In general, this function depends on the characteristics of the sensor and may depend on the characteristics and/or the identity of the target. The function $h_{ij}(\lambda_{ij})$ is referred to as the observation function of the sensor. The value λ_{ij} is measured in the sensor's frame of reference. This sensor-target pose, λ_{ij} , depends on the poses of the MPSE relative to the vehicle, the FPSE relative to the galaxy, and the vehicle relative to the galaxy.

[0180] The pose λ_{ij} can be expressed using a number of different representations of the 6-degrees of freedom in the pose. In the present version of the system, the Cartesian coordinates x , y and z of the target in the sensor's frame of reference represent translation, and a 3x3 rotation matrix which transforms vectors from the target's frame of reference to the sensor's frame of reference represents orientation. Another potentially useful representation of the pose λ_{ij} between the i^{th} FPSE and the j^{th} MPSE is expressed as $\lambda_{ij} = (\theta_i, \phi_i, \theta_j, \phi_j, l_{ij}, t_{ij})$. Here θ_i is the angle between the FPSE principal axis and a line connecting the FPSE and the MPSE, and l_{ij} is the length of the line. The azimuthal angle ϕ_i of this line is measured around the i^{th} FPSE's principal axis (forward or x-axis) using a right hand rule and the FPSE's y-axis as the starting mark. Angles θ_j, ϕ_j are a symmetrically defined off-axis angle and azimuthal angle of the departing ray from the j^{th} MPSE. The final term, t_{ij} , is the twist of the imaginary rigid pole connecting the centers of the two PSEs about a swivel joint in the pole center, where the twist of each pole half relative to its PSE is fixed by imagining that the pole flattens out like a canoe paddle at the end, and the intersection of this blade with the PSE's y-z plane lies along the y-axis. The advantage of this unusual pose representation is that it is symmetrical – to reverse the direction of reference and find the inverse pose, just switch indices i and j .

[0181] For earth-referenced sensors, such as GPS sensors, there is no specific FPSE involved in the measurement. In this case, the subscript i is ignored and the pose of the MPSE is relative to a fixed point in the galaxy frame of reference.

[0182] To provide a simple example of a sensor measurement, suppose a sensor measures a separation of the MPSE and FPSE according to the equation $\mathbf{z}_{ij} = l_{ij}^{-3}$. In this case, the observation function takes the form $h_{ij} \left(\begin{pmatrix} \theta_i & \phi_i & \theta_j & \phi_j & l_{ij} & t_{ij} \end{pmatrix} \right) = l_{ij}^{-3}$. More complex observation functions arise when the sensor measurement depends on more than simply the absolute separation of the sensor and target. Also, note that in general, the observation $\mathbf{z}_{ij}$ is a vector-valued quantity.

[0183] The system makes use of an observation model in which the measured observation is assumed to be corrupted by an additive observation noise as $\mathbf{z}_{ij} = h_{ij}(\lambda_{ij}) + \mathbf{v}_{ij}$. When the observation is vector-valued, the components of the observation noise $\mathbf{v}_{ij}$ may be correlated, and the covariance of this noise is a matrix $\mathbf{R}_{ij}(\lambda_{ij})$ which, as is noted, may depend on the relative pose of the sensor and the target.

[0184] In operation, it is useful to linearize the observation model about an “operating point” (e.g., an estimated pose $\hat{\lambda}_{ij}$) in the form $\mathbf{z}_{ij} = \mathbf{H}_{ij}(\hat{\lambda}_{ij}) \lambda_{ij} + \mathbf{v}_{ij}$. (With a 6-dimensional representation of the pose, and an m -dimensional observation, the observation matrix has 6 columns and m rows.) It is also useful to compute the difference between the expected measurement based on the estimated pose, $\hat{\mathbf{z}}_{ij} = \mathbf{h}_{ij}(\hat{\lambda}_{ij})$, and the actual sensor measurement. This difference, $\delta \mathbf{z} = \mathbf{z}_{ij} - \hat{\mathbf{z}}_{ij}$, is referred to as the “innovation” in the measurement.

[0185] As discussed more fully below, a sensor and its associated components are responsible for accepting an estimated pose $\hat{\lambda}_{ij}$, computing the matrices $\mathbf{H}_{ij}(\hat{\lambda}_{ij})$ and $\mathbf{R}_{ij}(\hat{\lambda}_{ij})$ that depend on this estimated pose, making an actual measurement $\mathbf{z}_{ij}$, and computing the innovation $\mathbf{z}_{ij} - h_{ij}(\hat{\lambda}_{ij})$. Note that these operations and computations may be very specific to the sensor,

and the computations may involve considerations of complex non-linear characteristics of the sensors.

[0186] As introduced above, the navigation system may also estimate intrinsic calibration (bias) parameters for the PSEs. In such a case, the observation function is extended to include these arguments as $\mathbf{z}_{ij} = h_{ij}(\lambda_{ij}, \mathbf{k}_i, \mathbf{k}_j)$, where $\mathbf{k}_i$ and $\mathbf{k}_j$ are vectors of calibration parameters for the FPSE i and MPSE j respectively. The observation model is extended to include the calibration parameters:

$$\mathbf{z}_{ij} = \begin{bmatrix} \mathbf{H}_{\hat{\lambda}_{ij}} & \mathbf{H}_{\hat{\mathbf{k}}_i} & \mathbf{H}_{\hat{\mathbf{k}}_j} \end{bmatrix} \begin{bmatrix} \lambda_{ij} \\ \mathbf{k}_i \\ \mathbf{k}_j \end{bmatrix} + \mathbf{v}_{ij}.$$

3 Tracking system architecture

[0187] Referring to FIG. 2, sensors 103 include PSE devices and active targets 105 and IMU 104. Measurements by PSEs are made between pairs of PSEs, one PSE functioning as a sensor and the other PSE functioning as a target. In one example, the PSE sensor is fixed on vehicle 100 and the PSE target is fixed in the environment 106. In another example, the PSE sensor is fixed in the environment 106 and the PSE target is fixed on vehicle 100. In another example, where multiple vehicles are used, the PSE sensor is fixed in one vehicle and the PSE target is fixed in another vehicle. Measurements by PSEs allow the navigation system 90 to determine the vehicle's position relative to the environment or relative to other vehicles.

[0188] The interoceptive sensors (inertial or other dead-reckoning sensors) in IMU 104 make measurements without reference to anything outside the vehicle. For instance, integration of inertial measurements allows navigation system 90 to determine the vehicle's position relative to the vehicle's position at a prior time. If the vehicle's position in the galaxy frame of reference is known at the prior time, then the inertial measurements allow the navigation system 90 to make an estimate of the current location of the vehicle 100 in the galaxy frame of reference. Thus, when measurements between pairs of PSEs are not possible, such as when vehicle 100 navigates through a portion of the environment where PSEs in the environment are all blocked by obstacles, navigation system 90 can still keep track of the location (including orientation, and

possibly other state variables of the vehicle such as velocity or rotation rate) of the vehicle based on the inertial measurements.

[0189] In the data processing unit 190, sensor-specific modeling is separated from the generic sensor fusion algorithms used to update system states. Specifically, sensor specific computations are isolated in PSE drivers 120. A meta-driver 122 provides an interface between the PSE drivers 120 and a sensor fusion core 200, which does not need to include sensor-specific aspects.

[0190] The PSE drivers 120 provide interfaces to PSE devices 105. PSE drivers 120 are software modules, which may be written by manufacturers of PSE devices 105 independently of the specific implementation of the sensor fusion core, and are implemented as shared object library files, such as “.dll” (dynamic link library) files under Windows or “.so” files under Unix. As will be described in more detail below, PSE drivers 120 include information and interfaces that are specific to the PSE devices 105, and data and code needed for computation of the linearized observation matrices, observation noise covariance matrices, and expected sensor measurements and/or innovations as described above.

[0191] Data processing unit 190 includes a sensor fusion core (SFC) 200 that processes information from sensors 103. SFC 200 includes a state update module 214 that updates a vehicle state 202, a vehicle map 204, and an environment map 206. The vehicle state includes estimates of the location and/or orientation of vehicle 100, as well as other relevant state variables for the vehicle, which is denoted as a vector $\mathbf{x}$. The vehicle map 204 includes a map (i.e., the poses) of MPSEs in the vehicle's frame of reference (B-frame), denoted as a vector $\mathbf{M}$ of the stacked pose vectors $\mathbf{m}_j$ of the individual MPSEs. The environment map 206 includes a map containing poses of FPSEs in the galaxy frame of reference (N-frame), denoted as a vector $\mathbf{F}$ of the stacked poses $\mathbf{f}_i$ of the individual FPSEs. The overall state is denoted as the stacked vector $\mathbf{X} = (\mathbf{x}^T, \mathbf{M}^T, \mathbf{F}^T)^T$.

[0192] In versions of the system that estimate calibration parameters k_i for PSE_{*i*}, each individual FPSE “vector” $\mathbf{f}_i$ is composed of a position, orientation, and intrinsic parameter (bias) fields:

$$\mathbf{f}_i = \begin{bmatrix} \mathbf{f}_i^p \\ \mathbf{f}_i^o \\ \mathbf{f}_i^k \end{bmatrix} = \begin{bmatrix} \mathbf{f}_i^{\chi} \\ \mathbf{f}_i^k \end{bmatrix},$$

where $\mathbf{f}_i^{\chi}$ stands for the combined pose states and $\mathbf{f}_i^k$ stands for the bias state. Similar notation is used for MPSE_j.

[0193] The vehicle state 202, vehicle map 204, and environment map 206 are stored in a memory. The term “vehicle state”, depending on context, will refer either to the state of the vehicle or a portion of the memory storing the vehicle state. Likewise, the term “vehicle map”, depending on context, will refer either to a map of the MPSEs on the vehicle or a portion of the memory storing the vehicle map. The term “environment map”, depending on context, will refer either to a map of the FPSEs in the environment or a portion of the memory storing the environment map.

[0194] Data processing unit 190 includes a meta-driver 122 that is used as a bridge between PSE drivers 120 and SFC 200. By dividing data processing unit 190 into portions specific to PSE devices 105 and a portion specific to updating the states of the navigation system 90, the navigation system can be easily reconfigured depending on the latest versions of device drivers and/or update algorithms. The meta-driver 122 allows PSE devices 105 and PSE drivers 120 to be designed without knowledge of the inner workings of the SFC 200, and SFC 200 can be designed without knowledge of the specific details of the PSE devices 105 and PSE drivers 120. The details of the meta-driver 122 will be described below.

[0195] Referring to FIG. 3, the state update module 214 of the SFC 200 includes a dynamic reference unit (DRU) 220 for processing interoceptive measurement data and a sensor processing module 222 for incorporating exteroceptive measurement data. DRU 220 processes inertial and other dead-reckoning measurements from IMU 104, while sensor processing module 222 processes measurements from PSE devices 105. The DRU 220 and sensor processing module 222 in combination updates the vehicle state 202, vehicle map 204, and environment map 206 using an iterative process. In operation, the DRU 220 may make use of multiple measurements from the IMU 104 between uses of measurements from the PSE devices 105. As described more

fully below, SFC 200 periodically (e.g., once each time an exteroceptive measurement datum is available) uses a PSE sensor measurement to update the state of the system.

[0196] DRU 220 processes inertial measurements from IMU 104 and maintains the vehicle state 202. Sensor processing module 222 controls and processes PSE sensor measurements based on an estimate of the vehicle's state from the DRU 220, and determines adjustments to be made to the vehicle state 202, based on measurements from PSE sensors 105, which it passes in a feedback arrangement back to the DRU 220. The sensor processing module 222 maintains the vehicle map 204 and the environment map 206 which it uses in determining the amount of adjustment.

[0197] In navigation system 90, measurements from IMU 104 are made at constant intervals, e.g., 180 times per second. Measurements from PSE devices 105 are made whenever possible (i.e., when at least one measurement can be performed between a pair of PSEs). At each iteration, zero or one measurement requests between sensors and targets are sent to the meta-driver, and during the same iteration, zero or one PSE sensor measurements resulting from a previous request may be received from the meta-driver and used to update the state of the system. However, the limitation of one request and one result received per cycle is not intrinsic to the architecture, and may be increased to a larger number in a future implementation.

[0198] In addition to maintaining an estimate of the vehicle state, the DRU 220 calculates and provides the state transition matrix Φ_x , and the process noise covariance matrix Q_x , to the sensor processing module which uses them to update an estimate of a covariance P_{xx} of the estimation error of this vehicle state estimate. The term P_{xx} will be referred to as the vehicle error covariance. The sensor processing module 222 uses the covariance and provides time updates to the covariance based on the matrices Φ_x and Q_x on each iteration. It also provides measurement updates to the covariance on each iteration in which a measurement is received and processed.

[0199] The sensor processing module 222 includes a measurement management unit (MMU) 304 and an update filter 306. MMU 304 selects pairs of PSEs for measurements by taking into account the location of the vehicle, as well as a utility of different measurements that might be made. For instance, when the vehicle is at different locations, different pairs of PSEs may be in

range of one another and of the pairs of PSEs that could be used, different pairs may be able to provide more useful information based on the characteristics of the sensors and/or the geometry of their relative locations. MMU 304 communicates with PSE driver 120 through meta-driver 122 to obtain the measurements from PSE devices 105 and corresponding sensor parameters, and provides processed sensor measurements and sensor parameters to the update filter 306.

[0200] When the navigation system 90 is powered up, MMU 304 or another component called sfShell, to be described later, calls the meta-driver 122 and requests that the meta-driver enumerate the sensing hardware. Meta-driver 122 loads available PSE drivers, such as 308 and 309, from a driver database 310. The PSE drivers are currently implemented as “.dll” files and assigned names having a format that allows the meta-driver 122 to detect them. The term “PSE drivers 120” refers to the PSE drivers collectively.

[0201] Meta-driver 122 then requests that each PSE driver 120 enumerate the PSE devices 105 that are available to navigation system 90 through that PSE driver and collect configuration information from those sensors. One PSE driver may control one or more than one PSE devices. Each PSE driver may control one or more types of PSE devices. For example, PSE driver 308 may control ultrasonic sensors and emitters, and PSE driver 309 may control laser range finders. Each PSE driver contains interfaces to, and possibly some information about parameters related to the PSE devices. The PSE devices may also store information about themselves. Meta-driver 122 receives the configuration information from PSE drivers 120, compiles the information and outputs a hardware configuration file “HW.cfg.” The hardware configuration file lists all of the PSE devices available and their configuration information.

[0202] In one example, the PSE driver is written by the manufacturer of the PSE sensor associated with the PSE driver. A PSE driver may also control PSE devices made by different manufacturers.

[0203] By using meta-driver 122 to enumerate the PSEs available upon power-up of navigation system 90, the navigation system is able to automatically reconfigure itself and continue to perform accurate map building and navigation when PSE devices or IMUs are added or removed from the system. This capability is referred to as “plug-and-track”.

[0204] During initialization, navigation system 90 reserves states in the state vectors processed by the update filter 306 to accommodate the states of the PSE devices that are found during the enumeration process. During active tracking, navigation system 90 can also add new states (e.g., when new PSE devices are found) or delete states (e.g., when PSE devices are determined to be unavailable).

[0205] By using the meta-driver 122, even though the existence of available PSE device 105 are relied upon in the calculation of inputs to update filter 306, the implementation of the sensor fusion core 200 is separated from the implementation of the devices 105, and thus it is possible to replace one type of sensor 105 with another without the need to redesign any components of the sensor fusion core 200. When the PSEs are added or removed from the system, the meta-driver 122 enumerates the PSEs connected to the system, and the SFC 200 accommodates them automatically, even if the new PSEs have structurally different measurement models or triggering policies compared to previous ones.

[0206] When navigation system 90 starts the process of tracking the location of vehicle 100, it first acquires an initial vehicle pose estimate in a process called acquisition, to be described later, and then enters tracking mode. In tracking mode, MMU 304 communicates with meta-driver 120 to determine which PSE pairs (including target and sensor) are available to make measurements during the next measurement cycle. MMU 304 selects a pair of PSEs from among the pairs of PSEs that are available to make measurements. MMU 304 makes the selection based on an "information gain" that represents the utility (or usefulness) of a measurement by the pair of PSEs to navigation system 90.

[0207] After MMU 304 selects the pair of PSEs that can make a measurement having the highest information gain, the MMU sends a message to meta-driver 122 to cause a PSE driver to trigger a measurement by the pair of PSEs. The actual measurement is generally available after a number of cycles due to the speed at which measurements can be made.

[0208] In choosing the pair of PSEs for making a measurement, MMU 304 first generates a "wishlist" of pairs of PSEs. The wishlist lists all of the pairs of PSEs consisting of a sensor which is available to the MMU and ready and a target which the MMU calculates to be within the sensing range of the sensor.

[0209] The MMU 304 sends the wishlist to an arbiter that is part of meta-driver 122. The arbiter identifies which pairs of PSEs on the wishlist are available to initiate a measurement in the next cycle. A particular pair may be marked by the arbiter as unavailable either because the pair of PSEs are still performing a measurement from a previous cycle or because the sensor or target is shared with one or more other vehicles. For shared devices, the arbiter tries to select a sensor-target pair which provides the highest information gain to as many of the requesting SFCs as possible. The arbiter sends a message to each requesting MMU, indicating that the MMU is allowed to request a measurement using that chosen pair in the next cycle. In the next cycle, the MMU sends a message to the meta-driver to cause a PSE driver corresponding to the selected pair of PSEs to trigger a measurement. The actual measurement may take many cycles to complete. Therefore, in each cycle, the MMU typically receives a measurement that it requested a number of cycles earlier. The sensor processing module 222 includes an update filter 306 to compute the amount of adjustment that needs to be made to the vehicle state 202, vehicle map 204, and environment map 206 after each measurement is made by a PSE sensor.

[0210] In the current version of the navigation system, the update filter 306 implements a complementary extended Kalman filter (CKF). The complementary extended Kalman filter estimates errors in the underlying states. The underlying states themselves are maintained outside of the CKF. Underlying states of the vehicle are maintained in the DRU. Underlying states used for intrinsic and extrinsic parameters of sensors and targets attached to the vehicle or to the galaxy are maintained in a vehicle map management module and galaxy map management module respectively. The CKF is decomposed into several separable modules by making use of a distributed Kalman filtering strategy called the Federated Kalman filter, which is modified to fit the complementary extended Kalman filter architecture.

[0211] The communication between sensor processing module 222 and meta-driver 122 is in the sensor's frame of reference. The sensor processing module 222 has available to it the current vehicle state (including pose), the vehicle map and environment map. The vehicle pose together with the vehicle map determine the poses of the MPSEs in the global frame of reference. Therefore, the sensor processing module has sufficient information to compute the relative pose of any pair of an MPSE and an FPSE.

[0212] The complementary extended Kalman filter computes incremental updates $\delta \mathbf{X}$ of underlying system states $\mathbf{X} = (\mathbf{x}^T, \mathbf{M}^T, \mathbf{F}^T)^T$ based on the innovations in the sensor measurements. To perform the computation, update filter 306 uses knowledge of the observation matrices that are linearized at the “operating point” (i.e., estimated system state $\hat{\mathbf{X}}$) and the observation noise covariance matrix $\mathbf{R}$.

[0213] Navigation system 90 is designed so that linearization of the nonlinear measurement model equations are performed by the PSE drivers 120, and sensor processing module 222 performs the linear computations associated with the CKF. In this way, sensor processing module 222 can be implemented without knowledge of the characteristics of the PSE devices 105. The PSE devices can be designed without knowledge of the specific details of the sensor processing module 222. PSE devices 105 and drivers 120 also do not have to know their positions and orientations relative to vehicle 100 or the galaxy 106.

[0214] Navigation system 90 is designed so that update filter 222 performs computations with coordinates expressed in the galaxy frame of reference, whereas PSE drivers 120 perform computations with coordinates expressed in the PSE sensor frame of reference. Update filter 306 operates in the galaxy frame of reference allowing navigation system 90 to determine the location of vehicle 100 in the galaxy frame of reference. In order to allow PSE devices 105 to be designed independently, each PSE sensor 105 operates in its own frame of reference. The MMU 304 provides the coordinate transformations that are necessary to allow measurements processed by PSE drivers 120 to be used by update filter 306, and to allow pose predictions made by update filter 306 to be used by PSE drivers 120. The following provides a more detailed description of the coordinate transformations.

[0215] FIG. 4 shows relative poses among galaxy 106, vehicle 100, and PSEs 338. The figure shows a measurement between a j -th MPSE (MPSE _{j}) 109 and an i -th FPSE (FPSE _{i}) 107. A galaxy frame of reference (N-frame) 330 is fixed relative to galaxy 106. A vehicle frame of reference (B-frame) 332 is fixed relative to vehicle 100. An MPSE frame of reference (j -frame) 334 is fixed relative to the MPSE _{j} 109 and an FPSE frame of reference (i -frame) 336 is fixed relative to the FPSE _{i} . The pose of vehicle 100 relative to the N-frame 330 is denoted as $\mathbf{x}$. The

pose of FPSE_i 107 relative to the N-frame 330 is denoted $\mathbf{f}_i$. The pose of MPSE_j 109 relative to the B-frame 332 is denoted $\mathbf{m}_j$. The pose of MPSE_j 109 relative to FPSE_i 107 is denoted as λ_{ij} . The measurement made between MPSE_j 109 and FPSE_i 107 is denoted $\mathbf{z}_{ij}$.

[0216] So far we have discussed a measurement between the j-th MPSE and the i-th FPSE without specifying which is the sensor and which is the target. In the case of an inside-out measurement, the MPSE is a sensor and the FPSE is a target. In this case, the sensor frame or s-frame is the j-frame, and the target frame or t-frame is the i-frame. In that case we denote the pose of the target relative to the sensor, by $\lambda_{st} = \lambda_{ji}$. This pose may be represented, for example, by a translation vector and rotation matrix $\{r_{j \rightarrow i}^j, C_i^j\}$. For an outside-in measurement, the s-frame is the i-frame and the t-frame is the j-frame, so $\lambda_{st} = \lambda_{ij} = \lambda_{ji}^{-1}$ where the inverse pose may be calculated by inverting the rotation matrix to get $C_j^i = (C_i^j)^T$ and reversing the translation vector and expressing it in the i-frame: $r_{i \rightarrow j}^i = C_j^i(-r_{j \rightarrow i}^j)$. For the symmetric 6-parameter representation of relative pose described earlier, reversing the pose referencing direction is achieved by switching the indices i and j.

[0217] PSE drivers 120 are designed to make measurements of targets relative to the S-frame of sensors, so that a PSE does not have to know whether it is attached to the vehicle or the galaxy, or know its pose relative to either one in order to make a measurement. This simplifies the design of PSEs and their drivers.

[0218] During tracking, MMU 304 computes a prediction $\hat{\lambda}_{ji}$ of the pose of FPSE_i 107 relative to MPSE_j frame of reference 334. This involves combining the estimates of $\hat{\mathbf{x}}$, $\hat{\mathbf{m}}_j$, and $\hat{\mathbf{f}}_i$. The MMU then inverts this pose if necessary in order to obtain the sensor-target relative pose prediction $\hat{\lambda}_{st}$ as described in the previous paragraph. This is the sensor-target pose prediction sent by the MMU to the PSE driver 120. The PSE driver performs its computations in the sensor frame of reference 334, ignorant of actual values or estimates of $\mathbf{x}$, $\mathbf{m}_j$, or $\mathbf{f}_i$. The PSE driver

120 estimates the predicted measurement $\hat{z}_{st}$ (which is a function h of the relative pose prediction $\hat{\lambda}_{st}$, the bias $\mathbf{k}_s$ of the PSE sensor, and the bias $\mathbf{k}_t$ of the PSE target), and produces a measurement error covariance matrix $\mathbf{R}$ which characterizes the expected degree of measurement noise in the planned measurement. PSE driver 120 sends to MMU 304 the predicted measurement $\hat{z}_{ij}$, $\mathbf{R}_{ij}(\lambda_{ij})$, and the sensor-specific component of the observation matrix

$$\mathbf{H}_\lambda = \frac{\partial \mathbf{h}(\lambda)}{\partial \lambda}$$

[0219] MMU 304 then uses the chain rule to convert the partial observation Jacobian $\mathbf{H}_\lambda$ into the final observation matrices $\mathbf{H}_x$, $\mathbf{H}_m$, and $\mathbf{H}_f$ needed by the Kalman filter in the update filter 306 to update $\mathbf{x}$, $\mathbf{m}_j$, and $\mathbf{f}_i$, respectively. Because of the cross-correlation matrices $\mathbf{P}_{MM}$ and $\mathbf{P}_{FF}$ maintained in the update filter, this also allows the navigation system to update estimates for the poses of other sensors $\mathbf{m}_{j'}$, $j' \neq j$, and targets $\mathbf{f}_{i'}$, $i' \neq i$.

[0220] MMU 304 then obtains an actual measurement from the meta-driver 122, and passes the innovation, the $\mathbf{R}$ matrix, and the final $\mathbf{H}$ matrices to the update filter 306 for processing.

[0221] Some sensors function in a way that a prior knowledge of the relative position of the target can make the measurements more accurate. For example, when a camera is used as a sensor, having knowledge of an estimate of the target location allows the camera to be aimed at the target so that the target is near the center of the camera's field of view where there is less distortion, thus providing more accurate information on the location of the target. The estimated pose of the target relative to the camera sensor that is provided by the MMU 304 to the PSE driver 120 can be used for such adjustment of the field of view.

3.1 Meta-driver with loadable sensor drivers

[0222] The meta-driver 122 provides uniform access for the SFC 200 to obtain information related to the PSEs. The meta-driver 122 is implemented as an application programming interface (API) that supports a number of function calls.

3.1.1 Discovery of device drivers and devices

[0223] One function call (e.g., sfMetaEnumerate()) invokes the enumeration process and returns a list of PSEs available to navigation system 90. The meta-driver 122 searches driver database 310 and loads PSE drivers (e.g., 308 and 309). Meta-driver 122 invokes each PSE driver to enumerate its PSEs. Each PSE driver is responsible for managing a particular type of PSEs (which can include different sub-types of PSEs, and can be made by different manufactures). The PSE drivers check the vehicle hardware to determine which PSE devices of its particular type(s) are connected to the vehicle. Each PSE driver reports back to the meta-driver 122, and the meta-driver 122 returns with a hardware configuration list that lists all of the PSEs available to the navigation system 90.

[0224] During enumeration, a handle which consists of a driver ID number and device ID number) is given to each PSE. For example, if there are 1000 PSEs on the ceiling of a room and 6 PSEs on the object being tracked, there would be 1006 handles. Meta-driver 122 would enumerate 1006 PSEs and furnish as much information about them as it was able to obtain from their respective PSE drivers 120.

3.1.2 Plug and track feature

[0225] A key feature of the navigation system 90 is the separation of modules specific to PSEs and modules specific to updating the states and maps. A separation between the PSEs and the update filters is desirable because there are different kinds of PSEs, each having different measurement characteristics. The measurement characteristics affect how the measurements are used in the update process. Due to the separation, PSEs can be designed without knowledge of the updating process. The modules specific for updating can be designed without knowledge of the PSE characteristics. A new PSE can be "plugged" into the navigation system and the navigation system will be able to recognize and use the new PSE.

3.2 Overall operation

[0226] At any given time, the tracking system 90 (or navigation system) is in one of three primary modes of operation: initialization, acquisition, or tracking. When system 90 powers up, it starts in initialization mode, which includes the enumeration and self-configuration procedures

described above. Once a valid system configuration has been achieved and stored in the requisite configuration files as described in the software architecture section later, the system moves into acquisition mode. In this mode, the goal is to find out the approximate location or pose of the vehicle in the galaxy. This initial pose estimate allows the system 90 to enter the tracking update loop, which involves a recursive estimation filter that uses an approximate initial guess of states.

[0227] When the acquisition procedures have found an initial estimate of the vehicle pose, the estimate $\mathbf{x}(0)$ is transferred into the DRU and the system enters tracking mode, in which it iteratively executes the tracking update loop. The system 90 remains in this tracking mode until some conditions occur, such as when the vehicle loses track of its location or pose information. This may happen when all of the exteroceptive measurements fail (for example, due to occlusion of line-of-sight sensors, or because the vehicle is out of range of all the FPSEs) for a prolonged period of time. If there is a prolonged period of time without external measurements, the vehicle pose estimate maintained in DRU 220 by dead-reckoning may drift to the extent that it becomes unusable. In that case, the system 90 returns to the acquisition mode.

[0228] The decision to re-acquire may be made automatically by the MMU 304 based on the pose estimation error covariance $\mathbf{P}_{xx}$ or the recent history of measurement returns and acceptance/rejection of measurement data or both. Alternatively, a command may be sent by a user to the system 90 to force it to re-acquire. The system 90 can transmit the diagonal elements of $\mathbf{P}_{xx}$ to a host application, along with its pose estimates $\hat{\mathbf{x}}$. This allows the user to take into account the system's own degree of uncertainty about its current pose, along with other factors outside of the system 90, in deciding to force the system to re-acquire.

3.2.1 Initialization

[0229] Initialization includes the plug-and-track procedures that have already been discussed, by which the meta-driver identifies the IMUs and PSE devices that are present and stores configuration files that will be used by the SFC 200 to govern its operation, including acquisition and tracking. The details of how the configuration files are created, and how the system obtains user input to fill in missing information that could not be obtained by the meta-driver during enumeration, are discussed more fully below in the section on software architecture.

Initialization occurs on power-up of the system 90. A command can also be sent to force the

system 90 to re-initialize, such as in the event that PSE devices are added or removed while the system 90 is running.

3.2.2 Acquisition

[0230] During acquisition, a scheduler in the MMU 304 sends requests for measurements to the meta-driver 122, and but instead of instructing a sensor to make a measurement to a specific target, MMU 304 asks the sensor to scan its operating volume and report all targets that it can identify. In one example, the scheduler continues to send scan requests to each sensor in turn until enough measurements are returned to perform an acquisition. While this scanning sequence is in process, each returned measurement is saved to a memory stack inside of an acquisition module within the MMU 304. On each cycle, an algorithm within the acquisition module analyzes the set of measurements that have been returned recently, and determines whether the set of measurements is sufficient to solve for the vehicle pose..

[0231] For example, four 2D bearing measurements from inside-out bearing sensors (such as cameras) on the vehicle to known targets in the environment is sufficient to solve a 4-point pose-recovery algorithm (PRA) known in the computer vision literature as the perspective n-point problem. Alternatively, measurements from two outside-in cameras at known poses in the environment to a single target on the vehicle are sufficient to triangulate the position of that target. If the position of three targets on the vehicle can be measured (e.g., by using 3 measurements from each of two outside-in cameras), then the orientation of the vehicle can be deduced from the position of the three known MPSE targets. In a similar example, three range measurements from an acoustic emitter on the vehicle to three non-collinear acoustic receivers on the ceiling can be used to tri-laterate the position of the emitter MPSE. Also, combinations may be used, such as one outside-in camera measurement combined with one range measurement, or a combination of inside-out and outside-in measurements.

[0232] Depending on the set of measurements returned, the acquisition module invokes an appropriate PRA routine, such as triangulation, trilateration, solution of the perspective-n-point problem, or a combination algorithm. The PRA routine returns a failure condition (e.g., because the measurements given to the PRA routine were not sufficiently consistent to produce a high-confidence pose estimate), or returns a presumably valid initial vehicle pose estimate $x(0)$.

[0233] If the PRA routine fails to determine the pose of the target, the MMU 304 remains in acquisition mode and continues to scan for more measurements. If the PRA routine succeeds, the MMU 304 sends the estimate $\mathbf{x}(0)$ to the DRU and switches the system into tracking mode. In order to initialize an complementary filter in update filter 306, the DRU 312 calculates an initial uncertainty for the vehicle state estimate, $\mathbf{P}_{\mathbf{xx}}(0)$, based on the confidence reported by the PRA routine, and sends $\mathbf{P}_{\mathbf{xx}}(0)$ to update filter 306. One example of causing the update filter 306 to initialize its $\mathbf{P}_{\mathbf{xx}}$ to $\mathbf{P}_{\mathbf{xx}}(0)$ is to send a $(\Phi_{\mathbf{x}}, \mathbf{Q}_{\mathbf{x}})$ message to the update filter with $\Phi_{\mathbf{x}}$ set to a matrix of all zeros, and $\mathbf{Q}_{\mathbf{x}}$ set to $\mathbf{P}_{\mathbf{xx}}(0)$.

3.2.3 Tracking update loop

[0234] In the tracking mode, the navigation system 90 repeatedly makes measurements and updates an estimate of a location of vehicle 100 in an iterative process. Each iteration may have a duration of, e.g., 1/180 second.

[0235] Within the primary system mode called tracking, the SFC can operate in a number of tracking modes. In a first mode of operation, navigation system 90 performs simultaneous localization and map-building (SLAM), and updates vehicle state 202 and environment map 206 in the iterative process. In a second mode of operation, navigation system 90 performs simultaneous localization and calibration (SLAC), and updates vehicle state 202 and vehicle map 204 in the iterative process. In a third mode of operation, navigation system 90 performs simultaneous localization, mapping, and calibration, and updates the vehicle state 202, vehicle map 204, and the environment map 206 in the iterative process. In a fourth mode of operation, navigation system 90 performs localization only (LO) by using existing vehicle map 204 and environment map 206, and does not update the maps; only estimates of the vehicle state 202 are updated.

[0236] The update filter 306, DRU 220, and MMU 304 form a feedback loop for updating the system states, as summarized in the following steps.

- a) The DRU 220 updates and stores its estimate of the vehicle state ($\mathbf{x}$) using inputs from IMU 104. DRU 220 generates $\Phi_{\mathbf{x}}$ and $\mathbf{Q}_{\mathbf{x}}$, which represent the error propagation during a measurement cycle and sends them to update filter 306.

- b) The update filter 306 updates $\mathbf{P}_{xx}$ based on Φ_x and $\mathbf{Q}_x$, and sends $\mathbf{P}_{xx}$ to MMU 304.
- c) MMU 304 interacts with meta-driver 122 and receives a measurement z from which it forms innovation δz , partial Jacobians $\mathbf{H}$ in the sensor frame of reference, and $\mathbf{R}$. MMU 304 converts the partial Jacobians into observation matrices $\mathbf{H}$ in the galaxy frame of reference and sends them along with $\mathbf{R}$ and the innovation to update filter 306.
- d) Update filter 306 generates an incremental update, $\delta \mathbf{X}$, of the system state $\mathbf{X}$ based on δz , $\mathbf{H}$'s, and $\mathbf{R}$. Update filter sends $\delta \mathbf{x}$ to DRU 200.
- e) The DRU 220 computes an updated estimate of the vehicle state $\mathbf{x}$.
- f) etc.

[0237] FIG. 5 summarizes the sensor fusion core architecture described above. SFC 200 includes MMU 304, DRU 220, an environment map manager 310 that maintains environment map 206, a vehicle map manager 314 that manages the vehicle map 204, and update filter 306. The update filter 306 may be a centralized complementary Kalman filter that updates $\delta \mathbf{x}$, $\delta \mathbf{M}$, $\delta \mathbf{F}$, $\mathbf{P}_{xx}$, $\mathbf{P}_{xM}$, $\mathbf{P}_{xF}$, $\mathbf{P}_{Mx}$, $\mathbf{P}_{Fx}$, $\mathbf{P}_{FF}$, and $\mathbf{P}_{MM}$.

[0238] The underlying states of the navigation system 90 are maintained by the DRU 220 (which maintains the vehicle state $\mathbf{x}$), the vehicle map manager 314 (which maintains the vehicle map, $\mathbf{M}$), and the environment map manager 310 (which maintains the environment map, $\mathbf{F}$). The update filter 306 estimates errors in the underlying MPSE poses and biases, and sends the corrections vector $\delta \mathbf{M}$ back to the vehicle map manager 314, which applies the corrections to the appropriate MPSE fields with appropriate translation of orientation format.

[0239] The update filter 306 sends error correction vector $\delta \mathbf{F}$ back to the environment map manager 310, which applies the corrections to the appropriate FPSE fields. Update filter 306 sends error correction vector $\delta \mathbf{x}$ to the DRU, which applies the corrections to update the vehicle state $\mathbf{x}$.

[0240] The following describes an example of the data that is passed between the modules in a measurement cycle during active tracking.

[0241] At the start of the measurement cycle, the IMU 104 sends (660) a measurement to the DRU 220. DRU 220 may access IMU data through the meta-driver 122, or through a separate IMU meta-driver, or access directly. DRU 220 integrates the measurement and updates the vehicle state $\mathbf{x}$. The DRU 220 also calculates a state transition matrix $\Phi_{\mathbf{x}}$ and process noise matrix $\mathbf{Q}_{\mathbf{x}}$, which are sent (665) to the update filter 306. Based on $\Phi_{\mathbf{x}}$ and $\mathbf{Q}_{\mathbf{x}}$, the update filter 306 updates the error covariance $\mathbf{P}_{\mathbf{xx}}$ using $\mathbf{P}_{\mathbf{xx}}(k+1) = \Phi_{\mathbf{x}}(k)\mathbf{P}_{\mathbf{xx}}(k)\Phi_{\mathbf{x}}(k)^T + \mathbf{Q}_{\mathbf{x}}(k)$, where k represents the time step. Update filter 306 sends (670) $\mathbf{P}_{\mathbf{xx}}$ to MMU 304, which uses $\mathbf{P}_{\mathbf{xx}}$ to calculate the innovation covariance matrix $\mathbf{S}$.

[0242] As will be described in detail below, measurement matrices $\mathbf{H}_{\mathbf{x}}$, $\mathbf{H}_{\mathbf{m}_j}$, and $\mathbf{H}_{\mathbf{f}_i}$ are partitioned into generic and model-specific parts, keeping the generic operations inside the MMU 304 and pushing the sensor/target details out into the PSE drivers 120. As a result, the architecture is extensible by outside parties. A sensor manufacturer who understands its own product and can model its output, but knows nothing of Kalman filtering or SLAM theory, can follow some instructions and write a sensor driver that would enable the sensor to work with the sensor fusion core and be optimally fused with the inertial sensors and any other sensors that are attached.

[0243] The MMU 304 exchanges a sequence of messages with the meta-driver 122 (shown as 600, 605, 610, 615, 625, and 635, which will be described in more detail below in FIG. 6), and eventually receives a measurement $\mathbf{z}$ from the meta-driver 122. MMU 304 computes an innovation $\delta\mathbf{z}$, which is the difference between the actual measurement $\mathbf{z}$ and the measurement that the PSE driver 120 would have predicted $\mathbf{h}(\hat{\lambda})$ based on the estimate of the relative pose $\hat{\lambda}$ of the PSE sensor and PSE target. The innovation may be the result of a measurement request that was made several cycles earlier, because a measurement may take several cycles to complete.

[0244] The MMU 304 also receives linearized measurement Jacobians $\mathbf{H}_{\lambda}$, $\mathbf{H}_{\mathbf{k}_s}$, and $\mathbf{H}_{\mathbf{k}_t}$ (which are linearized about the operating points $\hat{\lambda}_{st}$, $\hat{\mathbf{k}}_s$, and $\hat{\mathbf{k}}_t$.) and the noise covariance $\mathbf{R}$ of

the observation noise. MMU 304 converts $\mathbf{H}_\lambda$, $\mathbf{H}_{k_s}$, and $\mathbf{H}_{k_t}$ (which are in the sensor frame of reference) into $\mathbf{H}_x$, $\mathbf{H}_{m_j}$, $\mathbf{H}_{f_i}$, $\mathbf{H}_{k_i}$, and $\mathbf{H}_{k_j}$ (which are in the galaxy frame of reference). MMU 304 sends (675) δz , $\mathbf{H}_x$, $\mathbf{H}_{m_j}$, $\mathbf{H}_{f_i}$, $\mathbf{H}_{k_i}$, $\mathbf{H}_{k_j}$ and $\mathbf{R}$ to the update filter 306.

[0245] Based on δz , $\mathbf{H}_x$, $\mathbf{H}_{m_j}$, $\mathbf{H}_{f_i}$, $\mathbf{H}_{k_i}$, $\mathbf{H}_{k_j}$ and $\mathbf{R}$, the update filter 306 updates δx , $\delta \mathbf{M}$, $\delta \mathbf{F}$, $\mathbf{P}_{xx}$, $\mathbf{P}_{xm}$, $\mathbf{P}_{xf}$, $\mathbf{P}_{mx}$, $\mathbf{P}_{mm}$, $\mathbf{P}_{mf}$, $\mathbf{P}_{fx}$, $\mathbf{P}_{fm}$, and $\mathbf{P}_{ff}$. The update filter 306 sends δx to the DRU 220, sends $\delta \mathbf{M}$ and $\{\mathbf{P}_{mm}\}$ to the vehicle map manager 314, and sends $\delta \mathbf{F}$ and $\{\mathbf{P}_{ff}\}$ to the environment map manager 310. The notation $\{\mathbf{P}_{ff}\}$ refers to the set of $\{\mathbf{P}_{f1f1}$, $\mathbf{P}_{f2f2}$, ...} each of which is an individual 3x3 (or 6x6 for full PSE pose auto-calibration/auto-mapping) pose uncertainty of one PSE. Based on $\delta \mathbf{M}$, the vehicle map manager 314 updates $\mathbf{M}$. Based on $\delta \mathbf{F}$, the environment map manager updates $\mathbf{F}$. Based on δx , the DRU 220 updates the vehicle state x .

[0246] The sequence in which the matrices or variables are updated is not critical. Each component of the SFC 200 uses the most current information it can get to update the variables that it maintains.

[0247] FIG. 6 shows a sequence of messages that are passed among the update filter 306, MMU 304, meta-driver 122, and PSE driver 120.

[0248] Step 600: MMU 304 makes a wishlist of PSE pairs based on what PSEs are available and within range of each other. MMU 304 calculates the information gain B for each PSE pair based on an algorithm that determines the marginal benefit of the measurement to the vehicle. For example, a measurement having a lower measurement noise covariance may have a higher information gain than a measurement having a higher measurement noise covariance. In one instance, a measurement from a laser range finder may have a higher information gain than a measurement from an ultrasonic range finder, because using a laser is intrinsically more accurate than using ultrasound. In another instance, where the ultrasonic range finder is much closer to its target than the laser range finder to its corresponding target, the measurement from the ultrasonic

range finder may have a higher information gain. In another instance, where navigation system 90 needs information about a range that is in the forward direction of vehicle 100, a PSE sensor having a lower accuracy and pointed in the forward direction may produce a measurement having a higher information gain than a measurement made by a more accurate sensor pointing in the sideways direction. Thus, the information gain may depend on both the type of the sensor and the circumstances of the measurement.

[0249] The optimal selection of which sensor-target pair to process on the current cycle is difficult to make independently without considering all possible sequences of measurements and comparing the results which would be achieved at the end of the sequence. This results in a complex scheduling problem. An example of a suboptimal but efficient algorithm is to optimize the marginal benefit of an individual selection assuming the past selections have already been made, and the current selection will be the last. The marginal benefit will be the decrease in some risk function associated with the tracking quality, for example the trace of the vehicle state estimation error covariance $\mathbf{P}_{xx}$. This would require obtaining the measurement matrices $\mathbf{H}_{ij}$ and $\mathbf{R}_{ij}$ for each candidate sensor-target pair from the meta-driver and calculating the benefit of that measurement pair using:

$$B_{ij}(\mathbf{H}_{ij}, \mathbf{R}_{ij}) = \text{trace} \left\{ \mathbf{P}_{xx} \mathbf{H}_{ij}^T \left(\mathbf{H}_{ij} \mathbf{P}_{xx} \mathbf{H}_{ij}^T + \mathbf{R}_{ij} \right)^{-1} \mathbf{H}_{ij} \mathbf{P}_{xx} \right\}.$$

[0250] Different choices of the risk function would lead to different expressions for marginal benefit. To reduce the amount of computation, some ad hoc or heuristic algorithm may be used for the benefit function, such as giving higher benefit to sensor-target pairs that differ more from recently used sensor-target pairs in terms of location, direction or measurement model.

Additionally the benefit score of each sensor-target pair may be modified to take account of detection probability, for example by down-weighting targets which recently failed to be detected or returned unusable measurements and which may therefore still be occluded (or poorly lit, or inoperative).

MMU 304 sends the wishlist, including sensor and target device handles and benefit score for each pair, to the arbiter 352 of the meta-driver 122.

[0251] Step 605: If there is only one vehicle, the meta-driver 122 ignores gain B, and returns indicators showing which PSEs are available.

If there are multiple vehicles, the arbiter 352 of the meta-driver 122 looks at the requests for shared hardware from all of the MMUs, and select a pair of PSEs that will result in the highest information gain for all of the vehicles. The meta-driver 122 returns an indicator showing which shared PSE pair is available. The meta-driver 122 also marks all the privately held PSE pairs available unless they are busy.

[0252] Step 610: Based on the indicators sent back from the arbiter 352, the MMU 304 selects an MPSE j and an FPSE i to make the measurement (e.g., by selecting the pair on the wishlist having the highest information gain among those the arbiter indicated are available). MMU 304 makes an estimate of the pose $\hat{\lambda}_{ij}$ based on estimates $(\hat{\mathbf{x}}, \hat{\mathbf{m}}_j^{\mathcal{Z}}, \hat{\mathbf{f}}_j^{\mathcal{Z}})$ of the poses of the vehicle, MPSE $_j$, and FPSE $_i$.

[0253] The notations $\mathbf{m}_j^{\mathcal{Z}}$ and $\mathbf{f}_j^{\mathcal{Z}}$ refer to the poses of the MPSE $_j$ and FPSE $_i$, respectively. The notations $\mathbf{m}_j^k$ and $\mathbf{f}_j^k$ refer to the biases (or calibration parameters) of the MPSE $_j$ and FPSE $_i$, respectively.

[0254] Before MMU 304 sends the estimate of the pose (and optionally, calibration vectors) to the library server 354, MMU 304 uses the (i, j) tags which identify the measurement pair in the SFC to look up the (“drvId,” “devIdSensor,” and “devIdTarget”) tags which identify the pair in the meta-driver by determining which PSE driver (identified by the “drvId” tag) will be controlling the measurement, which PSE is the sensor (identified by the “devIdSensor” tag), and which PSE is the target (identified by the “devIdTarget” tag).

[0255] MMU 304 then performs a switch process to convert $\hat{\lambda}_{ij}$ (relating to the pose of MPSE $_j$ with respect to FPSE $_i$) into $\hat{\lambda}_{st}$ (relating to the pose of the target relative to the sensor). By using the MMU 304 to perform this switch process, the PSE driver 120 can be designed without knowledge of whether the measurement is inside-out or outside-in. The PSE driver 120 is designed to trigger a sensor to make a measurement of a target. The PSE driver 120 can be ignorant of whether the sensor is fixed to vehicle 100 and the target fixed to environment 106, or

whether the sensor is fixed to environment 106 and the target fixed to vehicle 100. This allows greater flexibility in the design and deployment of the PSE devices (and their associated hardware drivers).

[0256] MMU performs the switch process as follows. If FPSE_i is a sensor (such as in an outside-in measurement), then $\hat{\lambda}_{st} = \hat{\lambda}_{ij}$, $\hat{\mathbf{k}}_s = \hat{\mathbf{f}}_i^k$; $\hat{\mathbf{k}}_t = \hat{\mathbf{m}}_j^k$. If FPSE_i is a target (such as in an inside-out measurement), then $\hat{\lambda}_{st} = \hat{\lambda}_{ji}$, $\hat{\mathbf{k}}_s = \hat{\mathbf{m}}_j^k$; $\hat{\mathbf{k}}_t = \hat{\mathbf{f}}_i^k$.

MMU 304 also computes the partial Jacobian components of the pose $\lambda_{st}(\mathbf{x}, \mathbf{m}_j^z, \mathbf{f}_i^z)$ with

respect to $\mathbf{x}$, $\mathbf{m}_j^z$ and $\mathbf{f}_i^z$ (i.e., $\Lambda_{\mathbf{x}} = \frac{\partial \lambda_{st}}{\partial \mathbf{x}} \bigg|_{\hat{\mathbf{x}}, \hat{\mathbf{m}}_j^z, \hat{\mathbf{f}}_i^z}$, $\Lambda_{\mathbf{m}_j^z} = \frac{\partial \lambda_{st}}{\partial \mathbf{m}_j^z} \bigg|_{\hat{\mathbf{x}}, \hat{\mathbf{m}}_j^z, \hat{\mathbf{f}}_i^z}$, and

$\Lambda_{\mathbf{f}_i^z} = \frac{\partial \lambda_{st}}{\partial \mathbf{f}_i^z} \bigg|_{\hat{\mathbf{x}}, \hat{\mathbf{m}}_j^z, \hat{\mathbf{f}}_i^z}$), which the MMU 304 saves in a stack in the memory to be used later

(in step 620 below). MMU 304 sends $\hat{\lambda}_{st}$ and the (drvId, devIdSensor, devIdTarget) tags to the model library server 354 of the meta-driver 122. If auto-calibration or auto-mapping is performed, the MMU 304 also sends $\hat{\mathbf{k}}_s$ and $\hat{\mathbf{k}}_t$ to the meta-driver 122.

[0257] Step 615: Based on the drvId number, the model library server 354 of the meta-server 122 calls the model lookup function from the appropriate PSE driver 120 and passes in $\hat{\lambda}_{st}$.

Based on $\hat{\lambda}_{st}$, the model lookup function of the PSE driver 120 computes $\hat{\mathbf{z}}_{st}$, $\mathbf{H}_{\lambda}$, $\mathbf{H}_{\mathbf{k}_s}$, and $\mathbf{H}_{\mathbf{k}_t}$ using:

$$\hat{\mathbf{z}}_{st} = \mathbf{h}(\hat{\lambda}_{st});$$

$$\mathbf{H}_{\lambda} = \frac{\partial \mathbf{h}(\lambda)}{\partial \lambda} \bigg|_{\hat{\lambda}};$$

$$\mathbf{H}_{\mathbf{k}_s} = \frac{\partial \mathbf{h}(\lambda)}{\partial \mathbf{k}_s} \bigg|_{\hat{\mathbf{k}}_s}; \text{ and}$$

$$\mathbf{H}_{\mathbf{k}_t} = \left. \frac{\partial \mathbf{h}(\lambda)}{\partial \mathbf{k}_t} \right|_{\hat{\mathbf{k}}_t}.$$

The model lookup function of PSE driver 120 also computes the measurement noise matrix $\mathbf{R}(\lambda_{st})$. These calculations involve the specific nonlinear observation model h and noise characteristics of the sensor device, and these calculations are handled in the specific PSE driver written to deal with that type of sensor device.

[0258] Model library server 354 receives $\hat{\mathbf{z}}_{st}, \mathbf{H}_\lambda, \mathbf{H}_{\mathbf{k}_s}, \mathbf{H}_{\mathbf{k}_t}$ and $\mathbf{R}$ from the PSE driver 120 and returns them to the MMU 304. The terms $\mathbf{H}_\lambda, \mathbf{H}_{\mathbf{k}_s}, \mathbf{H}_{\mathbf{k}_t}$ each represents portions of the observation matrix in the sensor's frame of reference, and each includes information on how the measurement will depend on sensor-target pose, sensor biases, and target biases, respectively (linearized in the vicinity of the current estimates of those parameters).

[0259] Step 620: The values $\mathbf{H}_\lambda, \mathbf{H}_{\mathbf{k}_s}, \mathbf{H}_{\mathbf{k}_t}, \hat{\mathbf{z}}_{st}$, and $\mathbf{R}$ provided by model library server 354 are associated with the (drvId, devIdSensor, devIdTarget) tags. Also, the partial Jacobians $\mathbf{H}_\lambda, \mathbf{H}_{\mathbf{k}_s}, \mathbf{H}_{\mathbf{k}_t}$ are in the sensor frame of reference. The MMU 304 converts $\mathbf{H}_\lambda, \mathbf{H}_{\mathbf{k}_s}, \mathbf{H}_{\mathbf{k}_t}$ into the galaxy frame of reference by using the chain rule to combine the partial Jacobians from the model lookup with the saved Jacobians ($\Lambda_{\mathbf{x}}, \Lambda_{\mathbf{m}^\chi}$, and $\Lambda_{\mathbf{f}^\chi}$), and using a second switch process:

$$\mathbf{H}_{\mathbf{x}} = \mathbf{H}_{\lambda_{st}} \Lambda_{\mathbf{x}};$$

if FPSE_i is a sensor (outside-in measurement), then

$$\mathbf{H}_{\mathbf{m}} = \begin{bmatrix} \mathbf{H}_{\mathbf{m}^\chi} & \mathbf{H}_{\mathbf{m}^k} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{\lambda_{st}} \Lambda_{\mathbf{m}^\chi} & \mathbf{H}_{\mathbf{k}_t} \end{bmatrix},$$

$$\mathbf{H}_{\mathbf{f}} = \begin{bmatrix} \mathbf{H}_{\mathbf{f}^\chi} & \mathbf{H}_{\mathbf{f}^k} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{\lambda_{st}} \Lambda_{\mathbf{f}^\chi} & \mathbf{H}_{\mathbf{k}_s} \end{bmatrix};$$

if FPSE_i is a target (inside-out measurement), then

$$\mathbf{H}_m = \begin{bmatrix} \mathbf{H}_{m\chi} & \mathbf{H}_{mk} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{\lambda_{st}} \Lambda_{m\chi} & \mathbf{H}_{k_s} \end{bmatrix},$$

$$\mathbf{H}_f = \begin{bmatrix} \mathbf{H}_{f\chi} & \mathbf{H}_{fk} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{\lambda_{st}} \Lambda_{f\chi} & \mathbf{H}_{k_t} \end{bmatrix}.$$

(The i, j subscripts were omitted in the above equations.)

[0260] The MMU 304 then computes innovation covariance matrix S using:

$\mathbf{S} = \mathbf{H}_x \mathbf{P}_{xx} \mathbf{H}_x^T + \mathbf{H}_m \mathbf{P}_{mm} \mathbf{H}_m^T + \mathbf{H}_f \mathbf{P}_{ff} \mathbf{H}_f^T + \mathbf{R}$. The innovation covariance represents the degree of uncertainty in how closely the actual measurement will match the prediction, due to a combination of measurement noise (with uncertainty $\mathbf{R}$) and uncertainty in the prediction (resulting from uncertainties $\mathbf{P}_{xx}$, $\mathbf{P}_{mm}$ and $\mathbf{P}_{ff}$ in the vehicle pose, MPSE pose and FPSE pose). It will be used in step 630 to calculate a search window or a range about the estimated measurement that the sensor should search for the target.

[0261] The MMU 304 places $\mathbf{H}_x$, $\mathbf{H}_{m_j}$, $\mathbf{H}_{f_i}$, and $\mathbf{R}$ and $\mathbf{S}$ (which will be needed again for χ^2 acceptance testing when the measurement result comes back, as described later) in the stack along with (i, j) tags to keep track of which PSE pairs are associated with the measurement matrices.

[0262] Step 622: A previous measurement $\mathbf{z}'_{st}$ of (sensor s' and target t') is sent from a PSE driver 120 to the meta-driver 122, and queued by the meta-driver 122.

[0263] Step 625: The MMU 304 sends (drvId, devIdSensor, devIdTarget) tags, innovation covariance matrix S, and estimated measurement $\hat{\mathbf{z}}_{st}$ to the meta-driver 122 and requests a measurement.

[0264] Step 630: The meta-driver 122 forwards the request for a measurement by the PSE pair identified by the (drvId, devIdSensor, devIdTarget) tags to the appropriate PSE driver 120 as

designated by the drvId value. It also sends the predicted measurement $\hat{z}_{st}$ and the innovation covariance S . PSE driver 120 uses S to compute a measurement search region about $\hat{z}_{st}$, and then performs the measurement, returning only targets within this search region.

[0265] Step 635: The meta-driver 122 sends the measurement currently available, $z'_{s't'}$, along with its (drvId', devIdSensor', devIdTarget') tags to the MMU 304. The measurement z can sometimes represent a set of measurement results, because a sensor may find several targets within the search area (a.k.a. range gate), makes measurements with respect to every one of them, and returns all of the measurement results.

[0266] Step 640: The MMU 304 performs data association to select the most probable measurement result from the set, converts (s', t') to (i', j'), and matches the measurement $z'_{i'j'}$ with the H_x' , $H_{m_j'}$, $H_{f_i'}$, and R' and S' stored in the stack. The MMU 304 also computes the innovation $\delta z'_{i'j'} = z'_{i'j'} - \hat{z}'_{i'j'}$. And it applies a χ^2 acceptance test to the innovation using the innovation covariance S' to make sure the measurement is not an outlier before feeding it to the update filter 306.

Step 645: The MMU sends (i', j') tags, innovation $\delta z'_{i'j'}$, measurement matrices H_x' , $H_{m_j'}$, $H_{f_i'}$, and noise covariance R' to the update filter 306. Optionally, MMU 304 sends a timestamp that represents when the measurement actually occurred. The update filter 306 can use the timestamp to effectively backtrack and apply the measurement in the update filter 306 at the right moment in the recent past.

[0267] Step 650: A number of cycles after the meta-driver requested a measurement from a PSE driver 120, the PSE driver 120 sends back the measurement z_{st} , which is subsequently sent to the MMU 304.

3.2.4 Dynamic Reference Unit

[0268] At initialization of navigation system 90, DRU 220 or meta-driver 122 searches for interoceptive sensor drivers, and invokes the interoceptive sensor drivers to determine whether there are interoceptive sensors, such as inertial sensors, wheel encoders and dead-reckoning sensors.

[0269] Based on the type of interoceptive sensors available, DRU 220 determines whether tracking is desired in two-dimensional or three-dimensional space, and chooses an appropriate integration algorithm and error model. DRU 220 then asks each inertial sensor driver for a number of standard parameters appropriate to the integration model. Information about an IMU may include, for example, white noise and random walk amplitudes, root-mean-square initial uncertainty estimates for gyro and acceleration biases, ramps, misalignments, scale factors, nonlinearities, etc. Based on the sizes of these initial uncertainties, the DRU 220 decides whether to create estimable states in the state vector for each parameter. If DRU 220 does create states, it will cause the states to be estimated and corrected by update filter 306 simply by modeling their effects in the state transition matrix Φ_x and process noise matrix Q_x , and setting appropriate initial uncertainties for the states in $P_{xx}(0)$.

[0270] In one example, IMU 104 may use smart sensor modules that are compatible with an IEEE 1451.2 interface, which specifies a transducer electronic data sheet (TEDS) resident in an EEPROM of the sensor that identifies to the sensor driver the types of sensors present, and provides calibration data. In this situation, DRU 220 queries the sensor driver to obtain factory calibration data about the inertial sensor stored in the TEDS. Based on this data, DRU 220 compensates for sensor biases, scale factors, and misalignments.

[0271] When non-IEEE 1451.2 compliant inertial sensors are used, the sensor driver emulates the TEDS with user-supplied files.

4 Software architecture and multi-vehicle tracking support

[0272] The PSEs are represented by a class of PSE objects, which may be sensors or targets. Each PSE object has the following attributes:

(1) Pose (location and orientation) relative to the map of the environment or vehicle to which the PSE is attached.

(2) Pose uncertainty.

(3) Bias parameters vector **k**. This is used in the auto-calibration process.

(4) **k**-vector uncertainty.

(5) Basic type, which may be one of the following:

RS (range sensor)

RT (range sensor target)

2DBS (2-D bearing sensor)

2DBT (2-D bearing target)

1DBS (1-D bearing sensor)

1DBT (1-D bearing target)

RRS (range rate sensor)

RRT (range rate sensor target)

DS (dipole field component sensor).

DT (dipole target or source)

(6) Specific type, which includes the manufacturer name and model number. (This is used by the scheduler to look up information about what other device types it can measure to, which ones it cannot be used simultaneously with, max update rate, response time, field of view (FOV) and range limitations, approximate resolution, etc.).

(7) Unique identifier (such as a HW device serial number, or barcode of a vision fiducial target)

(8) Color.

(9) Size. (Identifier, color, and size attributes may be used differently by different specific types of devices. For example, a fiducial may use the size field to represent the diameter of the target, while a smart camera may use it to represent the FOV, and an architectural corner feature may ignore it.)

(10) Driver number (of the PSE driver that enumerated it and controls or uses it).

(11) Device handle (assigned by the driver on enumeration and used to access this PSE through its PSE driver. This is relevant for actively controlled hardware devices.)

(12) Status (ready, busy, etc.).

(13) Membership (attached to galaxy [0] or to a particular vehicle [n])

[0273] In the present version of the system, SFC 200, meta-driver 122, and PSE drivers 120 shown in FIG. 2 are implemented as software dynamic link library (“.dll”) files.

[0274] When a particular PSE is used in a measurement, function calls are made to the “.dll” library corresponding to the appropriate PSE driver 120.

[0275] FIG. 7 shows a block diagram of the software drivers of the tracking/navigation system 90.

Sensor Fusion Shell (“sfShell.dll”)

[0276] Tracking system 90 includes a sensor fusion shell (“sfShell.dll”) 412 that provides a unified access API for an application program 424 to obtain tracking data from one or more sensor fusion cores (“sfCore.dll”) 200. As an example, application program 424 may use the outputs of navigation system 90 to control the movements of vehicle 100. In this example, the navigation system residing on the vehicle requires data about the pose of a single vehicle (itself) and needs to instantiate one sensor fusion core 200. In another example, a tracking system 90 may be used to track the poses of 15 limb segments on a human dancer for character animation. Since this application requires tracking poses of 15 separate “vehicles”, all on the same computer and all reporting data to the same animation program, the sensor fusion shell would invoke 15 separate instances of the sensor fusion core .dll 200. The sensor fusion shell 412 also provides a unified access for a configuration manager (“sfCfgMgr.exe”) 416 (described in more detail below) to control the configuration of the SFCs operating under a given shell.

[0277] In describing the software modules, the terms “sensor fusion core” and “sfCore.dll” will be used interchangeably. The same applies to other software modules.

[0278] At start-up, the application program 424 opens the sensor fusion shell 412. The sensor fusion shell 412 opens the meta-driver (“sfMeta.dll”) 122, and causes it to enumerate all the sensor and target devices currently available. If necessary, the sensor fusion shell 412 invokes

the configuration manager 416 to clarify the configuration of each SFC, such as which sensor/target data will be controlled by and available to which SFCs.

[0279] For hardware devices private to the SFCs, the sensor fusion shell 412 passes pointers to the sfMeta.dll functions (122) into each SFC to allow the SFCs to directly access the hardware drivers over which the SFCs have exclusive control. For shared hardware devices, the sensor fusion shell 412 passes pointers to shell functions that mediate between SFCs and invoke the sfMeta.dll functions in turn.

[0280] For a simple single-vehicle self-tracking configuration, an SFC 200 may be operated stand-alone without any shell. In this case, it invokes the meta-driver 122 directly instead of through the shell 412. The meta-driver 122 provides hardware arbitration functions to mimic those in the shell, so that downward interfaces of the SFC 200 are identical whether accessing the meta-driver 122 through shell functions or directly. This option is reserved to allow future embedded operation of the SFC 200 in mobile robots with limited memory resources or which may not be running an operating system which supports shared libraries like ".dlls".

Meta-driver

[0281] A meta-driver 122 provides a unified access API for accessing device drivers 120. Device drivers 120 include drivers that control PSEs or interoceptive measurement units. As illustrated, a single device driver (e.g., 402) may control both PSEs (e.g., 406) and IMUs (e.g., 408).

[0282] The meta-driver 122 communicates with two types of drivers. One type of drivers control local hardware devices that are used by vehicles (or equivalently SFCs) running on the same computer on which the local device drivers reside. This type of drivers are provided as dynamic link library (".dll") files 402. Each .dll file 402 may control a number of PSEs 406 or a number of IMUS 408 or a combination. Another type of drivers control hardware that is potentially shared by SFCs running on remote computers. This type of drivers are called regional hardware drivers 404, and are provided as separate executable files, which may be running on the local machine or on a remote machine. For regional drivers, the meta-driver 122 communicates with the drivers through a socket interface.

Local drivers

[0283] Local drivers implemented as “.dll” files (402) provide the mechanism to allow third-party developers to extend the functionality of the tracking system by adding new devices (sensors or active targets) to be fused with the IMU 104 and/or with other devices already supported. To support a new device, the developer writes a “.dll” file conforming to the application programming interface (API) described later in this document, names it with a name in the format “sfDrvNNN*.dll”, and places the driver in the appropriate directory.

[0284] In this naming convention, “NNN” represents a three-digit integer from 001 to 999, and is a number not already in use. The author is responsible for checking the directory for existing .dll files and picking a driver number not already taken. If the first 8 characters of a name are not “sfDrvNNN,” then the file is ignored and not used as a driver, even if it is a valid .dll file. Extra characters after the first 8 are ignored and may be used by the developer to add a descriptive title to the driver, for example “sfDrv002_isenseSerialSmartCameras.dll.”

[0285] Each driver may access one or more devices (i.e. sensor or active target PSEs or IMUs). At startup, each driver is asked to enumerate all of the devices that the driver controls, and give each device a unique integer handle that will be persistent from boot-up to boot-up. Then, for each listed device, the meta-driver requests a hardware description record which specifies as many of the PSE attributes as it can. The PSE driver may try to obtain this information from a transducer electronic data sheet (TEDS) or non-volatile memory that is physically located in the hardware device, or from a driver-specific configuration file or using any other mechanism. It may be difficult to obtain some of the attributes values. For example, pose relative to the vehicle or the galaxy may only be available for factory assembled fixed device clusters.

[0286] The meta-driver 122 compiles all the lists of information about each device from each driver into one master list and outputs it to a hardware configuration file “HW.cfg” 410.

[0287] The drivers provide a hardware abstraction layer so that the SFC 200 does not need to know how the devices are interfaced to the computer, or where the signal processing is performed. For example, one driver may communicate with a network of IEEE 1451.2 smart sensor modules over CAN bus or Ethernet. Each of these modules contains its own processor with nonvolatile memory that contains all the self-descriptive information listed above. The driver implements the code to use the calibration data contained in the transducer electronic data

sheet (TEDS) to perform sensor compensation. Another driver may access a few smart cameras through the serial ports. This driver is responsible for checking on the available serial ports to see which if any has a smart camera attached. The image processing and sensor compensation may be performed in firmware on the smart cameras.

[0288] As another example, consider a driver that controls a network of acoustic emitter beacons through proprietary ISA bus hardware. As yet another example, there may be a driver that accesses a number of ordinary video cameras (e.g., cameras that do not process the images) over a firewire or USB bus. This driver then includes all the image processing and lens compensation code in its “.dll” file. For these drivers, the meta-driver does not need to know which port the hardware is physically attached to, or whether the information describing the device originated from the device or from a configuration file read by the driver. This is abstracted by the drivers and an integer device handle is returned, so that each device after initialization can be accessed using just a driver number and device number.

Sensor Fusion Core (“sfCore.dll”)

[0289] Each SFC 200 makes function calls to the standard function pointers passed in by the “sfShell.dll” 412 that will be used by “sfShell.dll” 412 and/or “sfMeta.dll” file 122 to initialize the hardware, lookup model information, initiate measurements, retrieve measurement data, and perform arbitration. During its measurement scheduling process, each SFC 200 is required to provide “sfShell.dll” file 412 with a "wish list" of PSEs for its next scheduling cycle. That list is used by the hardware mediator 418 to determine which of the PSEs may actually be scheduled for update.

Meta Driver (“sfMeta.dll”)

[0290] Meta driver 122 is the access point to all the hardware devices, local and regional. The SFC 200 receives function pointers from “sfShell.dll” 412 to enable the SFC 200 to obtain data from hardware devices. Hardware devices are marked exclusive or shared between SFCs. A regional hardware is considered shared. If a hardware device is exclusive, then SFC 200 will be able to obtain the data directly from the meta driver 122 using the function pointer. If a hardware device is shared, then is it the responsibility of “sfShell.dll” 412 to obtain the data from meta driver 122. Calls made from SFCs 200 simply retrieve the data from “sfShell.dll” 412.

[0291] FIG. 8 shows an example of a directory structure of the driver files and configuration files that will be described below.

Configuration Files

[0292] The sensor fusion core 200 uses seven files to control its operation:

- 1) Hardware configuration file ("HW.cfg").
- 2) Default vehicle PSE configuration file ("defaultVehiclePSEs.cfg").
- 3) Default galaxy PSE configuration file ("defaultGalaxyPSEs.cfg").
- 4) Auto-calibrated vehicle configuration file ("autoVehiclePSEs.cfg").
- 5) Auto-mapped galaxy configuration file ("autoGalaxyPSEs.cfg").
- 6) System parameter configuration file ("sysPar.cfg").
- 7) Device type rule configuration file ("typeRules.cfg").

HW.cfg

[0293] The hardware configuration file ("HW.cfg") 410 is a flat file that is output by the meta-driver 122 when it completes initialization. "HW.cfg" 410 contains a compilation list with one line for each device found. The line contains the driver number and device number (to be used thereafter by the meta-driver 122 to access the device) and other descriptive information about the device arranged in the following order:

Col 1	Col 2	Col 3	Col 4-6	Col 7-9	Col 10-18	Col 19-21	Col 22-31
Driver #	Device #	Status/ membership	X, y, z translation	X, y, z uncertainty	Rotation matrix	Rotation uncertainty	K parameters

32-41	Col 42	Col 43	Col 44	Col 45	Col 46
K param uncertainty	Basic Type	Specific type	Serial number	Color	Size

[0294] If certain fields of information can not be obtained from the device drivers, they are filled in with a question mark in the “HW.cfg” file, which causes the configuration manager program to prompt the user to fill in this missing data.

defaultVehiclePSEs.cfg & defaultGalaxyPSEs.cfg

[0295] Configuration files “defaultVehiclePSEs.cfg” 420 and “defaultGalaxyPSEs.cfg” 420 are flat files in the same format as “HW.cfg” 410. Files 420 show which PSEs belong to the galaxy and which PSEs belong to the vehicle. Each line from “HW.cfg” 410 is copied to one of files 420, and the missing fields (indicated by “?”) are filled in with valid data. The files 420 do not change unless the hardware configuration changes. Pose refinements or new PSEs discovered by the SFC 200 while the SFC is running are saved to “autoVehiclePSEs.cfg” and “autoGalaxyPSEs.cfg” files 422.

autoVehiclePSEs.cfg & autoGalaxyPSEs.cfg

[0296] Configuration files “autoVehiclePSEs.cfg” 422 and “autoGalaxyPSEs.cfg” 422 are flat files in the same format as “defaultVehiclePSEs.cfg” 420 and “defaultGalaxyPSEs.cfg” 420.” Non-zero lines present in these files override the corresponding lines in the default files. If they do not correspond to any lines in the default files, they are treated as additional PSEs. The new information obtained by the SFC 200 during auto-calibration or auto-mapping is saved into the files 422. Thus, if something goes wrong with the auto-mapping process, or the user wishes to start mapping again in a different room, he may delete either or both of the files 422, and the system will start tracking next time using the defaults files 420, which still retain the original hardware configuration information that was created by the “sfcfgMgr.exe” 416 program using data from the “HW.cfg” file 410 and user input if needed.

sysPar.cfg

[0297] A configuration file “sysPar.cfg” 422 is a text file containing <name = “value”> pairs, comments and section delimiters. This file is used to override default system parameters that govern how sfc.exe behaves. For example, a user can add a line such as StartupGalaxyMap = “defaultGalaxyPSEs.cfg”

to over-ride the normal behavior, which is to combine all the lines of “defaultGalaxyPSEs.cfg” 420 and “autoGalaxyPSEs.cfg” 422, using the values from the latter where there is a conflict. Other parameters that may be set in this file include prediction, perceptual post-filter modes, genlock parameters, simulation modes, and parameters that govern how the system decides when to switch between SLAC, SLAM, and LO.

typeRules.cfg

[0298] A configuration file “typeRules.cfg” uses <name = “value”> pairs and lists the rules associated with each specificType of device that are needed by the scheduler to schedule that type of device. Some of these rules are:

- (1) basicType membership.
- (2) Compatibility list (which specificTypes of targets are detectable by a given specificType of sensor or vice versa. Since the specificType identification numbers are non-overlapping with the basicType numbers, one may list a basicType, and that will tell the scheduler that all specificTypes belonging to that basicType can be scheduled.)
- (3) Conflict list (which specificTypes or basicTypes can not be scheduled at the same time as this specificType).
- (4) Max repetition rate.
- (5) FOV and or range limitations, or how to calculate them from the identifier, color and size attribute fields.
- (6) Approximate resolution (used by the scheduler to decide which sensor/target pair will yield maximum information gain). A more exact expression for resolution depending on the specific location of the target relative to the sensor is implemented by the model response function of the driver dll when it returns an R matrix for use in the Kalman filter.
- (7) Uncertainties of the 10 K-parameters.
- (8) Formulas for how to apply the K-parameters to compensate sensor biases.

[0299] SpecificTypes are intended to become officially registered (e.g., through a registration web page that can be set up to allow third party developers to submit a set of type rules for a new specific type, and then receive a unique specificType ID code that is not already used by anyone.) This allows a developer to put a new device on the market and have it become compatible with future users of the architecture (and current users, who has to re-load the master “typesRules.cfg” file from the web-page in order to use the new devices). This registration process allows a new device to declare itself incompatible or compatible with certain older devices that are already registered in the database. A small number of user specificType codes are reserved so that a developer can perform in-house testing of new device types before registering them.

Operation of sfCfgMgr.exe and interaction with sensor fusion system

[0300] When the navigation system boots up, a sensor fusion shell program “sfShell.dll” begins by asking the meta-driver 122 to query all the local drivers 402 and regional drivers 404 for status. The drivers go out and find all the connected hardware (e.g., PSEs 406 and IMUs 408), and return enumerated lists of devices that the meta-driver 122 concatenates into a master list and outputs as the file “HW.cfg” 410. The program sfShell.dll then invokes the configuration manager program “sfCfgMgr.exe,” which ensures that the “default*.cfg” and “auto*.cfg” files are updated, consistent with the actual hardware devices, and represent a viable tracking configuration. When these criteria are met, program “sfCfgMgr.exe” returns a success signal to the program “sfShell.dll,” which then proceeds to invoke the appropriate number of instances of sfCore.dll and cycle through them to acquire measurements and track locations of vehicles 100. If the criteria cannot be met, program “sfCfgMgr.exe” returns a failure signal to abort the program sfShell.dll, and issues an error message to the user to help with identification and correction of the problem.

[0301] Note that sfCfgMgr.exe is drawn as a standalone executable interfaced to sfShell.dll through a socket interface. This allows it to be invoked and run on a remote computer from the sensor fusion system itself. This capability could be very useful because the configuration manager 416 requires a GUI to obtain from the user any missing configuration information that could not be resolved by analyzing the information in HW.cfg. Since the tracking system 90 data processing unit 190 may be running on an embedded computer with minimal or no user

interface capabilities, the ability to configure it from a remote computer such as a PC would be very useful. Alternatively, part of the configuration management logic could be implemented within the sfShell.dll program (or a .dll loaded by it), and the user-interface components of it may be implemented in a separate program. This might allow the portion running in sfShell to make its best guess about configuration if or when the user cannot be consulted.

[0302] Referring to FIG. 9, a flow diagram illustrates a process 430 implemented by “sfCfgMgr.exe”:

[0303] Step 431: the program “sfCfgMgr.exe” is started.

[0304] Step 432: the program “sfCfgMgr.exe” determines whether “HW.cfg” has changed, as compared to the “lastGoodHW.cfg.” If not, skip to step 456.

[0305] Step 434: If “HW.cfg” has changed, then determine whether any new lines have been added. If not, skip to step 446.

[0306] Step 436: If new lines have been added, then determine whether the new lines match inactive PSE lines in “defaultVehiclePSEs.cfg” or “defaultGalaxyPSEs.cfg” or “autoVehiclePSEs.cfg” or “autoGalaxyPSEs.cfg”? If yes, then in step 438, re-activate the lines in the default configuration files and/or the auto configuration files where they appear. Then skip to step 446.

[0307] Step 440: If new lines do not match inactive PSE lines in the default files, determine whether there is enough information in the new lines to automatically assign the new lines to a vehicle or the galaxy. If yes, then skip to step 444.

[0308] Step 442: If there is not enough information in the new lines to automatically assign the new lines to a vehicle or the galaxy, open a graphical user interface (GUI) to ask the user to assign membership of new devices and fill in missing pose or other information.

[0309] Step 444: Copy new lines to “defaultVehiclePSEs.cfg” or “defaultGalaxyPSEs.cfg” depending on whether the new PSE has been identified as belonging to the galaxy or to a vehicle.

[0310] Step 446: Determine whether any lines have been subtracted from “HW.cfg.” If not, skip to step 450.

[0311] Step 448: If there are lines that have been subtracted from "HW.cfg," set status inactive in default and auto configuration files for any lines that were subtracted in "HW.cfg."

[0312] Step 450: Determine whether there are enough active PSEs available in the default or auto configuration files to track. If not, then in step 452, tell the user to add additional hardware and abort with failure condition.

[0313] Step 454: Copy "HW.cfg" to "lastGoodHW.cfg."

[0314] Step 456: Return control to calling program with success condition.

[0315] When the system is used for the first time, there are "sysPar.cfg" and "typeRules.cfg" configuration files. When the program "sfShell.dll" starts, it asks its meta-driver 122 to find all the existing "sfDrvNNN*.dll" drivers 402, and in turn ask each driver to enumerate all the hardware it can find that is physically connected and responding. The complete list is output to the file "HW.cfg" 410. The program "sfShell.dll" then invokes the program "sfCfgMgr.exe." Because "HW.cfg" 310 has changed (relative to the non-existent or empty "lastGoodHW.cfg"), the "sfCfgMgr.exe" program will attempt to automatically create the files "defaultVehiclePSEs.cfg" and "defaultGalaxyPSEs.cfg," using information obtained from the drivers about how each sensor or target device is disposed relative to the vehicle or the galaxy (in general, this is possible when the user has plugged in factory assembled stations that contain this information in their EEPROM descriptors).

[0316] If the information in the devices is not sufficient, the "sfCfgMgr.exe" program prompts the user to fill in the required information, and then creates the default files. If the resulting default configuration is deemed valid, the "sfCfgMgr.exe" program then passes control back to the calling program. Since there are no auto files to override the default files, the SFC 200 will begin tracking using the information in the default files.

[0317] After the SFC starts tracking, the default files contain the approved, fully-assembled hardware description for the configuration. Typically, this information does not need to change unless the user re-assembles the vehicle or the active FPSEs in the galaxy. Information about passive fiducial arrays (or natural features) is automatically generated by the auto-mapping function of the SFC and stored in the "autoGalaxyPSEs.cfg" file without harming the default files. If the user wishes to take the same vehicle to a new location and begin auto-mapping, he

can delete or rename the “autoGalaxyPSEs.cfg” file. Likewise, if he wishes to re-calibrate the vehicle, he can delete the “autoVehiclePSEs.cfg” file..

5 Using the system for simulation

[0318] An advantage of the described software system is that it can be used to perform Monte Carlo simulation and covariance analysis of a contemplated tracking configuration before developing or installing any actual hardware. The software can be used to decide what types of new devices to develop and what performance specifications they require for any given application and device layout.

[0319] A simulation is performed by using the following process:

- (1) Add a line `<mode = “simulated”>` in “sysPar.cfg” file.
- (2) Add a line of the form `<simIMUtype = “imuIC2”>` in “sysPar.cfg” file.
- (3) Manually construct the files “defaultVehiclePSEs.cfg” and “defaultGalaxyPSEs.cfg” to represent the “true” configuration of PSEs to be simulated. Unlike normal operation where “defaultGalaxyPSEs.cfg” might contain just a seed map, in simulation mode, it contains the full list of the PSEs that might be used during the course of the entire simulation. To inform the simulation that only the first 4 PSEs are pre-known for an auto-mapping experiment, give those 4 low covariance values, and give the others high covariance values. The covariances will be ignored by the sensor/target meta-driver, which will treat the pose values in the default files as absolute truth for the purpose of creating simulated measurements. However, the SFC will add random errors to the PSEs proportional to covariance values when it initializes its internal data store maps, and it will pay attention to the covariances in deciding which PSEs it can schedule. As it adds new PSEs to the “known” portion of the map, it will save the reduced covariances and refined pose estimates into the “autoGalaxyPSEs.cfg” file, which can be compared afterwards to the “defaultGalaxyPSEs.cfg” file to determine how closely the auto-mapping process converged to the true PSE positions.

- (4) Copy a mat-file containing the desired simulated trajectory data to the directory where the SFC resides and name it "trajectory.mat."

[0320] The currently supported IMU types include, e.g., "imu100", "imuIC2," "imuIT2," "imuCommercial," "imuTactical," "imuNavigation," "imuStrategic," "imuPerfect," "imuNone," and "imuCustom." If "imuCustom" is selected, then many additional parameters defining simulated IMU performance are added to the "sysPar.cfg" file. The non-custom types already contain pre-defined combinations of parameters (e.g. angle random walk, bias random walk, initial bias uncertainty, scale factor uncertainty, etc.) which are used to simulate the performance of different grades of IMU.

[0321] In the simulation mode, the "sfShell.dll" program skips the startup steps of asking the meta-driver to enumerate hardware and output the concatenated list of devices to the file "HW.cfg" and invoking "sfCfgMgr.exe" file. Thus, it proceeds immediately to acquisition and tracking using the user-generated "default*.cfg" files. In the simulation mode, if any hardware devices are connected, they are ignored. The drivers are passed an exact sensor-target pose (which is calculated in the MMU using vehicle pose data that comes from the "trajectory.mat" file) and exact true poses of the MPSE and FPSE (which come from the "defaultVehiclePSEs.cfg" and "defaultGalaxyPSEs.cfg" files respectively) and told to provide simulated data (by adding biases and noise).

[0322] In the normal mode, the initial vehicle and galaxy PSE maps are loaded into internal data store memory by first reading in all lines from the default files, and then reading in the auto files if they exist and overwriting the data from the default files. Therefore, both auto files are deleted before running a simulation, since they contain PSE pose values which no longer match the truth values defined in the default files. To simulate an ordinary tracking experiment with no auto-mapping or auto-calibration, the PSE covariances in the default files are set to zero.

6 Alternative system architectures

[0323] Tracking system 90 can be used to track a single vehicle or multiple vehicles. Multiple vehicles can be used to track different parts of an object, such as body parts of a person. Multiple vehicles may also be used to track a group of people or robots or objects in a common

environment. Multiple vehicles can also increase the speed and accuracy of map building since more data is available from the sensors of multiple vehicles to collaboratively construct the map.

Where multiple vehicles are tracked, the states of different vehicles can be updated centrally or the processing may be distributed between multiple SFCs running on the same data processing unit or separate data processing units. In a centralized multi-vehicle navigation system, one SFC receives sensor inputs from all vehicles, performs tracking and auto-calibration, and map building for all vehicles. States of multiple vehicles are appended to the state vector of the single SFC and propagated through the update filter to create an optimal algorithm for collaborative mapping.

[0324] Referring to FIG. 10, in a fully distributed multi-vehicle system, the navigation system utilizes multiple vehicles 470 in which each vehicle has its own data processing unit 190, with its own hardware drivers 476 (which could include a meta-driver 122, local drivers 402 and/or regional drivers 404) and SFC 200 to perform its tracking and possibly auto-calibration or auto-mapping. Each vehicle 470 builds its own private environment map 474 representing the portion of the environment 106 it sees. Each vehicle has hardware drivers 476 that collect data from PSEs private to the vehicle, and possibly also from PSEs that are shared among vehicles through an infrastructural hardware driver 480 which communicates with each vehicle 470 during the enumeration process to notify the vehicles what shared PSEs are available.

[0325] Each vehicle 470 has a measurement scheduler 478 that schedules which PSEs to use for each measurement cycle. A regional scheduler mediator 482 communicates with the measurement scheduler 478 of each vehicle to resolve conflicts in utilizing shared PSEs.

[0326] The private environment maps 474 developed by one vehicle 470 may or may not be shared or merged with the private environment maps 474 developed by other vehicles 470. A regional map reconciliation filter 484 reconciles the private environment maps based on a number of factors. For example, if certain PSE position estimates from one vehicle are known to be more accurate, either because its sensors are more accurate or because it is located at a more favorable geometric location with respect to certain targets, the private environment map from that vehicle or a portion of the private environment map relating to those certain targets may be given more weight in the map-merging process. The position estimates and associated

covariances for a given PSE provided by different vehicles may be fused using the information fusion form of Kalman filter update equations, such as given in section 2.4 of the appendix. If the fused map or portions of it are redistributed to the vehicles after a merge, algorithms should be designed to avoid potential filter divergence problems that could result from violation of the Kalman filter assumption of uncorrelated measurement noise.

[0327] In an alternative distributed system architecture, each vehicle has its own SFC to update its own states, but multiple SFCs are running on one data processing unit 190, as illustrated in FIG. 7. This data processing unit may be physically connected to one of the vehicles, or it may a separate computer receiving data by equivalent connections from all the vehicles.

[0328] In an alternative distributed system architecture, the responsibility of central command may transfer from one vehicle to another. Each vehicle independently stores the system states, and under certain circumstances, such as when the vehicle currently serving as central command malfunctions or when it loses communication with other vehicles, another vehicle takes over and assumes the responsibility of the central command.

[0329] In an alternative distributed system architecture, each vehicle has its own SFC and storage that stores the system states. Depending on circumstances, one or more vehicles may serve as central commands that actively receives sensor inputs from other vehicles and updates the states for itself and other vehicles. The vehicles not acting as central commands pass their sensor inputs to the central commands, receive the system states returned from the central commands, and store the system states in a local storage.

[0330] Initially, when all of the vehicles are within communication range, one vehicle may serve as the central command. As the vehicles navigate the environment and roam farther and farther away from each other, one or more vehicles may break off communication. When a vehicle breaks off communication with the central command, the vehicle activates its SFC and uses its own SFC to proceed with the task of localization, calibration, and map building. The central command, after noticing that a fellow vehicle has broken off communication, can either treat the sensors from the fellow vehicle as unavailable, or may reconfigure its update filter to reflect the changes.

[0331] When the fellow vehicle regains communication with the central command, a fusion process may be performed in which the system states of the fellow vehicle is fused with the system states of the central command, and then the fellow vehicle cedes the responsibility of updating the system states to the central command.

[0332] Because the navigation system implements an enumeration process to determine availability of sensors and a self-configuration process to set up the vehicle and environment maps according to sensor configurations, there is much flexibility in how the multiple vehicles cooperate to obtain the most accurate measurements. For example, the central command may dispatch “sentries” to navigate closer to a target in order to obtain more accurate measurements of the target. The sentry may send back data to be fused with data owned by the central command, or the sentry may simply become another target that can be used as a navigation reference point.

[0333] In this architecture, each vehicle has a sensor fusion core (SFC) associated with it to perform its tracking and auto-calibration, and to build its own private version of the part of the map it sees, which may or may not be shared or merged with map fragments developed by other vehicles. The SFC may either be a dedicated processor running on-board the vehicle, or just a process or thread running inside a server somewhere. Each SFC has a DRU containing the vehicle state, a vehicle map containing pointers to all the PSEs attached to the vehicle, and an environment map containing pointers to all the environment-fixed PSEs the vehicle is currently using.

[0334] The following describes an example of how a multi-vehicle navigation system determines which pair of PSEs to use for measurement in each measurement cycle. In a first scenario, only one measurement is made during each cycle. Referring to FIGs. 3, 6, and 8, the MMUs 304 in all the SFCs 200 send their wishlists to an arbiter 418, which can be part of sfShell 412. For any hardware which is not under the direct control of the local sfShell, the arbiter may in turn consult a regional mediator 482. Having received wish lists from all vehicles in the current cycle, the arbiter picks a pair of PSEs that are available to make a measurement in the next cycle so that the measurement will result in the highest overall information gain to the navigation system 90, taking all vehicles into consideration.

[0335] On the next cycle, the arbiter sends a message to the MMU 304 in each of its SFCs, indicating that the MMU 304 is allowed to make a measurement using the selected pair of PSEs. In turn, the MMU 304 sends a message to the meta-driver to setup a PSE driver corresponding to the selected pair of PSEs to perform a measurement. Once all the SFCs have made a request to setup the same measurement (as they have been directed to do by the arbiter) the sfShell sends a trigger command to the meta-driver to initiate the actual hardware measurement process. The actual measurement may take many cycles to complete. Therefore, in each cycle, the MMU 304 typically receives a measurement that it requested a number of cycles earlier. For shared measurements, the measurement is requested by multiple vehicles, completed once by the meta-driver, then the same result is returned to all the requesting vehicles to utilize in their respective update filters.

[0336] In the second scenario, different vehicles are allowed to make different measurements simultaneously as long as the measurement by one vehicle does not interfere with measurements by other vehicles. Each MMU 304, in addition to generating a wishlist of shared PSEs, also generates a wishlist of private PSEs. The term "private" means that the PSE is not shared by SFCs of other vehicles.

[0337] An example of a shared PSE is an ultrasound emitter whose chirps can be received by microphones on multiple vehicles. When there are multiple ultrasound emitters in close range, typically only one emitter is allowed to chirp at a time so as to prevent interference from the multiple emitters. Each microphone has its own timer and can determine the distance to the emitter by measuring a time difference between the trigger of the chirp and the reception of the chirp.

[0338] When one or more PSE are shared among SFCs, it is necessary to arbitrate among the needs of the SFCs to determine which PSE will result in a measurement having the highest overall benefit to the navigation system. For example, when there are two ultrasound emitters, the first emitter may be closer to a first vehicle but far away from a second vehicle. The second emitter may be farther to the first vehicle than the first emitter, but the average distance to both vehicles may be smaller than that of the first emitter. Thus, triggering the second emitter may result in a measurement having a higher overall information gain to the navigation system. Of course, the benefit of a measurement also depends on the level of need in the vehicles, so if one

vehicle as a large pose estimation error covariance in a particular direction, the arbiter may schedule a beacon in that direction, even if there are other beacons which could provide slightly higher information gain to a larger number of vehicles.

[0339] The MMUs 304 send the shared and private wishlists to the arbiter on each cycle. For each PSE pair on the private wishlist, the arbiter immediately marks it available if it is not busy, thus allowing it to be scheduled on the current cycle. For the shared wishlist, however, it waits until it has received the shared wishlists from all vehicles at the end of the cycle, then makes a decision, and returns that decision to the vehicles in the next cycle. Thus, on each cycle, the MMU will be authorized to use any of its private PSEs that are not busy, or to use the selected pair involving shared hardware from what it requested last cycle. The MMU can choose either one of its private PSE pairs or the shared pair, whichever yields highest information gain for its own vehicle. However, if the shared pair yields almost as much gain as the best private pair, it may choose the shared pair because it knows that opportunity will only be available for a very limited time. The shared device will now be activated as long as at least one vehicle decides to request a pair involving it. On a given cycle, some vehicles may choose to ignore the authorized shared measurement opportunity and schedule private measurements using different non-interfering sensors. Different vehicles can be making different measurements using different sensors at the same time, as long as they do not interfere with each other. The arbitration process is designed to prevent vehicles from scheduling measurements that may interfere.

[0340] As an example, a navigation system may be used to track body parts of a person in a room where a number of ultrasound speakers are installed on the ceiling. During the tracking process, only one ultrasound speaker may be triggered to chirp at a time. A first vehicle to be tracked may be placed on the person's head, a second vehicle to be tracked may be placed on the person's hand, and a third vehicle to be tracked may be placed on the person's leg. The vehicles contain microphones to measure the ultrasound chirps. The MMU of each vehicle generates a wishlist of pairs of PSEs (each pair including a speaker and a microphone). In generating the wish list, the MMU assigns a selfish information gain value (i.e. the gain for its own vehicle) to every combination of speaker-microphone pair.

[0341] At the k^{th} cycle, the MMUs of the vehicles send their wishlists to the arbiter 418. The arbiter responds by specifying a pair of PSEs which includes a shared speaker that was selected

based on the wishlists submitted at the $(k-1)^{\text{th}}$ cycle. During the k^{th} cycle, the arbiter gathers the wishlists from the MMUs and uses an algorithm to decide which speaker will maximize the information gain of as many vehicles as possible on the $(k+1)^{\text{th}}$ cycle. In this example, only one speaker chirps, and one microphone in each vehicle makes a measurement. At the $(k+1)^{\text{th}}$ cycle, the arbiter sends messages to tell the MMUs which speaker will chirp and which microphones have been granted permission to measure the chirp. After all of the MMUs send setup messages to ultrasonic hardware driver preparing it to chirp the same speaker and measure the time of flight to their different respective microphones, the sfShell will send a trigger message to the ultrasonic hardware driver. The driver then simultaneously chirps the speaker and starts the counters in all the microphones that have been setup for the measurement. A number of cycles later, the ultrasonic hardware driver returns the measurement results to the MMUs.

7 IMU and DRU automatic configuration

[0342] The DRU is responsible for predicting the estimate of vehicle pose forward at each time step, and providing the time varying linearized Φ and Q matrices to the update filter 306 that it will use to propagate forward the covariance matrix P_{xx} representing the uncertainty in the vehicle pose estimate. The DRU may propagate the predicted pose based on inputs from interoceptive sensors or without any sensor inputs.

[0343] If there are no sensor inputs to the DRU, the pose prediction will be based on some assumed dynamical process model of the object being tracked. Two popular such process models are the constant velocity model driven by white-noise acceleration, or the constant acceleration model driven by white-noise jerk.

[0344] If the enumeration process has identified the existence of any interoceptive sensors, it will also obtain parameters from the TEDS in the sensors describing what type of sensors they are and what measurement characteristics they possess, and save all these parameters into an IMU.cfg file. This file is read by the DRU and allows it to configure itself in an appropriate mode to perform the pose propagation by integrating the measurements received from the sensors on each cycle. The parameters describing the stochastic model of the interoceptive sensors are used by the DRU to formulate the appropriate matrices Φ and Q . For example, if

the interoceptive sensors identify themselves as a standard IMU providing measurements of $\Delta\theta$ and Δv on each cycle, then the DRU will implement inertial integration equations for updating the pose estimate and implement for Φ a state transition matrix representing the propagation of pose and velocity errors as a result of inertial sensor errors, and it will calculate Q based on such IMU parameters as gyro angle random walk, bias random walk, initial bias uncertainty, scale factor uncertainty, misalignment uncertainty, and all the equivalent parameters for accels. On the other hand, for odometry sensors, there would be different integration equations implemented, a different state transition matrix generated, and different stochastic parameters required to generate the process noise matrix.

8 Applications of tracking, auto-calibration and map-building system

[0345] Referring to FIG. 11, navigation system 90 can be used in many applications, such as to navigate a forklift 500 in a factory environment, to track a helicopter 502 or a soldier 504 in a battle field, or to track a person 506 in an indoors environment. An inside-out measurement can be used in which a sensor 508 attached to person 506 makes a measurement with respect to a target 510. An outside-in measurement can be used in which a sensor 512 attached to the environment makes a measurement with respect to a target 514 attached to a person 516.

[0346] In some applications, the primary goal of the system is to provide a means of navigation or tracking. If auto-mapping is used at all in such an application, it is done to provide a reference map for the purpose of the tracking. The map may be discarded after the tracking or navigation activity is complete, or it may be saved for future tracking or navigation activities in that environment. These types of applications are frequently discussed in the robotics literature under names like simultaneous localization and map-building (SLAM) or concurrent mapping and localization (CML).

[0347] In other applications, the map itself may be the primary product of interest. For example, a camera may be moved through a building, finding existing natural features such as corners and adding them to the map as PSEs. The map representation could be augmented to include information about lines connecting the corners together. Eventually a 3D model of the building will be formed which can be converted into a CAD model for as-built documentation and use in architectural planning. If the camera captures images and stitches together fragments of images

to create texture maps which can be associated with the polygons in the CAD model, the result would be a virtual reality (VR) model of the environment, created rapidly and either automatically or semi-automatically with minimal user interaction. These types of applications, where the goal is more to capture the 3D structure of the environment rather than track a vehicle's motion, are often discussed in the computer vision literature under names like structure-from-motion or scene modeling.

9 Alternative embodiments

[0348] 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.

[0349] For example, referring to FIG. 12, the update filter 306 can be a distributed Kalman filter, such as a Federated Kalman filter, that includes a number of filter components. Update filter 306 may include a SLAC filter (which stores and updates $\delta\mathbf{x}$, $\delta\mathbf{M}$, $\mathbf{P}_{\mathbf{xx}}$, $\mathbf{P}_{\mathbf{xM}}$, and $\mathbf{P}_{\mathbf{MM}}$), a SLAM filter (which stores and updates $\delta\mathbf{x}$, $\delta\mathbf{F}$, $\mathbf{P}_{\mathbf{xx}}$, $\mathbf{P}_{\mathbf{xF}}$, and $\mathbf{P}_{\mathbf{FF}}$), and an LO filter (which stores and updates $\delta\mathbf{x}$ and $\mathbf{P}_{\mathbf{xx}}$).

[0350] The SLAC filter 364 estimates errors in the underlying MPSE poses and biases, and sends the corrections vector $\delta\mathbf{M}$ back to the vehicle map manager 314, which applies the corrections to the appropriate MPSE fields with appropriate translation of orientation format.

[0351] Similarly, the SLAM filter 362 sends error correction estimates back to the environment map manager 310, which applies the corrections to the appropriate FPSE fields with appropriate translation of orientation format.

[0352] The LO filter 360 is used when the user or the navigation system 90 decides that both auto-calibration and mapping are done, and turns off the SLAC filter 364 and SLAM filter 362. Any one, two or three of the filters may be active at a given time.

[0353] Because each of the three local filters in this federated filter bank is a reduced-order Kalman filter which leaves out certain states that were present in the full-order centralized filter,

techniques such as a Schmidt-Kalman filter algorithm or just “bumping up” the noise assumed in the measurement model should be used to assure stability and best performance. To facilitate this, during the incorporation of a measurement between $FPSE_i$ and $MPSE_j$, the SLAM filter 362 sends $P_{f_i f_i}$ (error covariance of $FPSE_i$) to LO filter 360 and SLAC filter 364 and the SLAC filter 364 sends $P_{m_j m_j}$ (error covariance of $MPSE_j$) to LO filter 360 and SLAM filter 362.

[0354] When more than one of the local filters 360, 362, and 364 are running in parallel, the fusion filter 366 is used after all the local filters have completed their measurement updates to fuse their individual estimates of vehicle pose error $(\delta x_{SLAM}, P_{xx_{SLAM}})$, $(\delta x_{SLAC}, P_{xx_{SLAC}})$, and $(\delta x_{LO}, P_{xx_{LO}})$ into a fused estimate $(\delta x_{Fused}, P_{xx_{Fused}})$. Then δx_{Fused} is sent to the DRU to correct the vehicle pose states, and $P_{xx_{Fused}}$ is sent back to each of the local filters to perform the fusion reset equations of the federated filter algorithm. The fusion reset equations in the SLAM filter result in another correction δF being sent to the environment map manager, and likewise for the SLAC filter. Further details and references concerning the federated filter and distributed Kalman filtering are in the appendix.

[0355] In order to accommodate large galaxy maps, there may be many SLAM filters within the distributed filter bank, each responsible for updating a submap or portion of the entire map. These submaps may be disjoint or overlapping. One, several, or all of the SLAM submap filters may be active at a given time. For example, the environment map could be partitioned into rooms, and when the vehicle is in a particular room, only the SLAM filter which updates PSEs in that room is active. When the robot transitions from one room to another, the federated filter fusion reset equations are used to transfer the vehicle state estimate covariance matrix into the new room's SLAM filter, and update the stale P_{xF} and P_{FF} matrices in the new filter based on the change in P_{xx} which has occurred compared to the stale P_{xx} matrix left in the new SLAM filter since the last time it was active. This map transition process is similar in effect to the submap transition algorithm of Leonard and Feder. For smoother transitions as the robot moves between submaps, there may be several submap SLAM filters around the robot active at any given time, and as the robot moves forward, submaps behind it de-activate and ones in front activate.

[0356] If there will be a maximum of N submap filters active at any given time, then the bank of filters in the federated filter implementation would only need to contain N SLAM filters. When the vehicle moves, the SLAM filter corresponding to the submap behind it de-activates first and transfers its latest information into the environment map manager. Then, a group of PSEs in the area the vehicle is approaching are transferred from the environment map manager into the now vacant SLAM filter memory store.

[0357] Data processing unit 190 can either reside in vehicle 100 or be located outside of vehicle 100. Measurements from sensors 102 and IMU 104 may be transmitted wirelessly to a server computer located at a base station, the server computer updates the vehicle states, vehicle maps, and environment maps based on the measurements, and sends the updated vehicle states, vehicle maps, and environment maps back to vehicle 100 for storage in a memory in vehicle 100. The vehicle states, vehicle maps, and environment maps may also be stored at the base station, so vehicle 100 has little on-board computation power, and only carries sensors 102, IMU 104, communication devices, and perhaps motor control mechanism for controlling the movement of vehicle 100.

[0358] A vehicle with higher processing power may be used with vehicles with less processing power. The vehicle with less processing power sends measurement data to the vehicle with higher processing power, which updates the vehicle states, vehicle maps, and environment maps, and sends the updated information back to the vehicle with less processing power. In this way, lower cost navigation robots can be used with a more expensive navigation robot to explore a terrain.

[0359] The term "vehicle" can mean any object, from as small as an atom to as large as a planet. In a multi-vehicle system used to track movements of a person, a vehicle can be a tracker mounted on the person's head, and another vehicle can be a second tracker mounted on the person's wrist.

[0360] The auto-calibration process may also refine calibration parameters provided by the manufacturers of the sensors, where the calibration parameters have limited accuracy.

[0361] The pose of a target relative to a sensor does not necessarily have to be expressed as $\lambda_{ij} = (\theta_i, \phi_i, \theta_j, \phi_j, l_{ij}, t_{ij})$. Other representations may be used as well. For example position of

target relative to sensor may be specified using Cartesian or polar or other coordinate systems. Orientation may be expressed with euler angles, quaternions, rotation matrices, or rotation vectors. For some PSEs, orientation parameters may not be important, in which case no corresponding filter states need be allocated. For example, for a circular fiducial mark, the orientation is only used to determine visibility, and perhaps to calculate a centroid correction for perspective distortion. It should suffice to enter an approximate orientation and not update the orientation during the iterative measurement cycles. By comparison, for a camera, because the aim of camera lens is important, the orientation are included in the state vector so that measurements from the camera can be interpreted, and if auto-calibration is desired the extended Kalman filter can refine the alignment.

[0362] The pose of the vehicle relative to the galaxy, the FPSEs relative to galaxy, or the MPSEs relative to the vehicle may also be expressed in different formats or with fewer than 6 degrees of freedom. For example, in 2D navigation problems it is common to express the robot pose using just x, y and heading.

[0363] Various methods of auto-calibration may be used. In one implementation, navigation system 90 gathers information from a number of measurements from different sensors, and determines whether measurement values from a particular MPSE have a high percentage that are consistently offset from an estimated measurement value computed based on measurements from other sensors. The navigation system 90 calculates an estimate of the pose of the particular MPSE relative to the vehicle so that, using the estimated pose, the measurement values from that particular sensor more closely match the estimated measurement values based on measurements from other sensors. The auto-calibration parameters may be initialized to very accurate values through precise mechanical alignment fixtures, and then no auto-calibration are performed, or they may be initialized to approximate values and then refined by the real-time iterative auto-mapping process described above. Alternatively, they may be estimated by a batch process after collecting a certain amount of data.

[0364] When an MPSE is first installed on vehicle 100, instead of having a user manually enter the MPSE pose information, the MPSE may communicate with other sensors upon initialization of the navigation system 90 to compute an initial estimate of its pose relative to the vehicle. The

initial pose estimate is subsequently calibrated by the navigation system 90 using measurements from other sensors.

[0365] The navigation system 90 may automatically calibrate measurement biases inherent in the MPSEs. For example, a specification sheet of an ultrasound range finder may specify that readings from the range finder are to be deducted by a preset value. Due to manufacturing tolerances, the readings may actually have to be deducted by a different value. Such biases may be calibrated using the method described above for calibrating poses of the MPSEs.

APPENDIX A

Towards a Generalized Sensor Fusion Architecture for Distributed Real-time Tracking, Auto-calibration and Environment Mapping

Eric Foxlin

InterSense Incorporated, Burlington, MA, USA

ericf@isense.com

Abstract

This paper discusses the design of a very general architectural framework for systems that involve tracking the motion of users by fusing dead-reckoning sensors such as an inertial navigation system with aiding sensors such as sonar or computer vision which correct the motion tracking system by making measurements to external landmarks while simultaneously constructing a map of the positions of those landmarks. The goals of the architecture are to permit easy "plug-and-track" user configuration of a wide variety of sensor combinations including IMUs, GPS, range sensors, inside-out bearing sensors, outside-in bearing sensors, etc., and to provide compatibility with multiple sensor networking standards, distributed sensor fusion algorithms, and implementation strategies. A decentralized Kalman filter based on Carlson's Federated filter algorithm is used to decouple the auto-mapping, auto-calibration and tracking filters to produce a more flexible and modular architecture.

1. Introduction and motivation

There is an enormous range of applications that require measuring the position and/or orientation of moving objects. This is usually called navigation or localization when a person, robot or vehicle determines it's own position, and called tracking when an outside system monitors the position or orientation of moving objects, as with air traffic control, missile detection radar, wildlife tracking, and human motion capture.

Other applications require creating a map of an environment or at least surveying the locations of certain landmark points in the environment. This may range from large scale aerial photogrammetry and GIS to more local surveying and architectural as-built documentation, all the way down to 3D modeling of small products.

In an increasing range of applications, including augmented reality (AR) and mobile robotics, there is interest in both tracking (or navigating) the moving entity and mapping it's surrounding environment. In both applications, the map-building is necessary in support of the tracking function because the tracking operates by making relative measurements between the vehicle and certain reference points in the environment. (From here on, we will use the term vehicle to refer to the moving object, even if it is an HMD or hand-held stylus.) In AR applications, the environment map is also needed for the purpose of registering the virtual objects to the real world.

When navigating in a previously uncharted territory with no GPS availability, such as Mars rover missions, underground mining or undersea operations, often the only means of navigating is to make use of landmark observations, and the map of the landmarks is not available in advance. This has led to extensive research on the topic of simultaneous localization and map-building (SLAM) by the robotics community. Given the extreme difficulty and expense of specialized surveying to obtain landmark locations in preparation for tracking or navigation activities, SLAM is likely to become the predominant approach to both tracking and environment modeling in AR, robotics, and other potentially large markets which rely heavily on 3D tracking and modeling and are currently held back by the lack of practical means for performing these functions with high accuracy on a large scale.

Although most theoretical research on SLAM algorithms is relatively unconcerned with the specific types of sensors used, every application requires different types of sensors adapted for the operating environment and performance requirements of the application. Each specific sensor will have it's own measurement model with unique distortions and biases that must be included in the tracking filter algorithms for maximum accuracy. Thus we are really interested not just in SLAM, but in simultaneous tracking, auto-calibration, and environment mapping (STACEM) for

hybrid multisensor/multitarget systems. Although SLAM is a much better acronym, the word "localization" seems to connote position on a 2D map, which is the typical example used in SLAM research. The term "tracking" will be used from here on to remind the reader that we are primarily targeting 6-DOF position and orientation tracking, although the proposed architecture must also support lower-dimensional problems.

1.1 Structure of centralized KF solution to SLAM

Smith, Self and Cheeseman (1987) introduced a solution to the SLAM problem based on an Extended Kalman Filter (EKF). For a map with N fiducial points, the state vector consists of the vehicle state $\mathbf{x}$ augmented with the position states of all the fiducials:

$$\mathbf{X} = [\mathbf{x}_v^T \quad \mathbf{f}_1^T \quad \mathbf{f}_2^T \quad \cdots \quad \mathbf{f}_N^T]^T,$$

where $\mathbf{f}_i$ is the position of the i^{th} beacon in the map. The estimation error covariance matrix associated with this augmented state vector has the structure

$$\mathbf{P}_{\text{XX}} = \begin{bmatrix} \mathbf{P}_{\text{xx}} & \mathbf{P}_{\text{xf}_1} & \mathbf{P}_{\text{xf}_2} & \cdots & \mathbf{P}_{\text{xf}_N} \\ \mathbf{P}_{\text{f}_1\text{x}} & \mathbf{P}_{\text{f}_1\text{f}_1} & \mathbf{P}_{\text{f}_1\text{f}_2} & \cdots & \mathbf{P}_{\text{f}_1\text{f}_N} \\ \mathbf{P}_{\text{f}_2\text{x}} & \mathbf{P}_{\text{f}_2\text{f}_1} & \mathbf{P}_{\text{f}_2\text{f}_2} & \cdots & \mathbf{P}_{\text{f}_2\text{f}_N} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{\text{f}_N\text{x}} & \mathbf{P}_{\text{f}_N\text{f}_1} & \mathbf{P}_{\text{f}_N\text{f}_2} & \cdots & \mathbf{P}_{\text{f}_N\text{f}_N} \end{bmatrix}$$

in which the diagonal blocks represent the covariance matrices of the vehicle position estimate and all the fiducial position estimates. The off-diagonal blocks represent the cross-covariances between the vehicle and each fiducial, and between each fiducial and all the other fiducials. These cross-covariances are essential to the stochastic mapping approach. If they are ignored (i.e. treated as zero), each successive measurement will be incorrectly interpreted as new independent information and the filter will diverge. On the other hand, provided that the full covariance matrix is maintained, and updated together with the state vector using the standard EKF procedures, the optimal stochastic mapping algorithm has been shown to perform very well in many applications.

Because the fiducials are assumed to have constant position, the state transition and process noise matrices for the EKF time update step will have diagonal structure:

$$\Phi(k) = \begin{bmatrix} \Phi_x(k) & & & \\ & \mathbf{I} & & \\ & & \ddots & \\ & & & \mathbf{I} \end{bmatrix} \quad \mathbf{Q}(k) = \begin{bmatrix} \mathbf{Q}_x(k) & & & \\ & \mathbf{0} & & \\ & & \mathbf{0} & \\ & & & \mathbf{0} \end{bmatrix} \quad (1.1)$$

which reduces the computational load of the time update from the usual $O(N^3)$ to $O(N)$. During each measurement update cycle, the vehicle makes a measurement to one of the fiducials, say $\mathbf{f}_i$, which is used in the standard EKF measurement update procedures to update the mean estimates $\hat{\mathbf{x}}$ and $\hat{\mathbf{f}}_i$, and the full covariance matrix $\mathbf{P}$. This updating of $\mathbf{P}$ requires $O(N^2)$ operations and becomes the rate limiting step for large N .

1.2 Limitations of centralized EKF approach

The centralized EKF algorithm summarized above is optimal, but it tends to constrain the architecture of systems that use it. The Kalman filter must be specifically designed for a certain set of sensors. The measurement function $\mathbf{h}$ depends on the type of sensor and how it is positioned on the vehicle. The measurement noise matrix $\mathbf{R}$ must be tuned as a function of $\hat{\mathbf{x}}$ and $\hat{\mathbf{f}}_i$ to provide a level of noise representative of the sensor performance, bumped up appropriately to account for unmodeled biases, of which there are always some. By sensor biases we mean any constant parameters that affect the measurement, even in complex nonlinear ways, such as camera lens distortion parameters. A given type of sensor also has its own requirements or rules governing the frequency with which it can acquire data or how to trigger its sources, so a fairly specific measurement scheduler and data screening algorithm are also coupled to the KF implementation. Because the details of every sensor are relied upon in the design of the Kalman filter, it is usually not a straightforward task to replace one type of sensor with another. Any time the hardware is changed, the filter must be retuned at a minimum, and reprogrammed if the sensor has a structurally different measurement model or triggering policy. The reprogramming and revalidation effort can be significant if new states need to be added to reflect new or different types of bias terms.

The stochastic mapping formulation is intended to address the problem of recursive map-building, not sensor calibration, and no explicit provisions are made for sensor bias states. Either all the sensors must be calibrated by other means in the factory and their biases compensated before the measurements enter the Kalman filter, or the bias states can be appended to the vehicle state vector. In the latter case, the designer must decide which bias states to include, and the filter will always be estimating these parameters, even after they become sufficiently well estimated that no further gains can be made. The user cannot reconfigure the system with more or fewer sensors, because that would change the dimension of the vehicle state vector and require reprogramming the function which generates the matrices $\Phi_x(k)$ and $\mathbf{Q}_x(k)$.

Finally, the algorithm as described is designed for tracking a single vehicle in a fairly limited area. Additional vehicle states could be appended to the state vector, and that would indeed lead to an optimal algorithm for collaborative mapping. Unfortunately, this would mean that the tracking code for all the vehicles would have to run on the same computer, which might be a problem if each vehicle requires its own on-board navigation processor. New fiducials can be added to the map as the vehicle enters a new area, but the computational burden is $O(N^2)$, and after a few hundred fiducials the system may slow down to an unacceptable update rate.

1.3 Goals of this work

This paper presents preliminary design notions for an architecture that tries to overcome some of these limitations and provide a framework in which to develop an extremely wide variety of user-configurable multisensor/multitarget sensor fusion systems for applications which require tracking, auto-calibration, auto-mapping or any combination of these.

The goals of the architecture are

- Sensor versatility: allow fusion any combination of types, quantities and qualities of sensors, including inside-out, outside-in or a mixture.
- "Plug & track": allow user to configure systems by plugging together self-identifying and self-describing smart sensor modules.
- Permit implementation of either navigation systems with on-vehicle processing per vehicle, or tracking systems in which a regional processor tracks all the vehicles in the region.
- Scalability: enable systems with large numbers of tracked vehicles and/or large maps.
- Algorithm versatility: provide a top-down architectural specification which defines the interfaces between modules, but allows the designers of the individual modules freedom to choose specific sensor fusion algorithms and to later change those algorithms without changing other modules. In particular, the architecture shall neither force nor prevent the use of algorithms based on sub-map partitioning.
- Function versatility: at any time (even while running) the system may switch its function between simple tracking only (STO), simultaneous tracking and auto-calibration (STAC), simultaneous tracking and environment mapping (STEM), or simultaneous tracking, auto-calibration, and environment mapping (STACEM)

- Object-oriented design: facilitate re-use of the modules even if the top-level system architecture needs to change.

In order to achieve some of these goals the stochastic mapping algorithm is recast from its standard form as presented in Section 1.1 into a complimentary EKF formulation with explicit states for intrinsic and extrinsic parameters of sensors and targets attached to the vehicle and those attached to the map. This is done in Section 2.3 based on a sensor/target object extraction made in Section 2.1 and 2.2. In section 2.4, the sensor fusion process for a given vehicle is decomposed into several separable modules by making use of a distributed Kalman filtering strategy called the Federated filter (Carlson, 1990) modified to fit our complimentary EKF architecture. All of the sensor and configuration specific modeling is separated out from the generic sensor fusion algorithms into two modules called the DRU and MMU. Section 3 discusses the hooks that are being planned to allow future expansion of the architecture to handle multi-vehicle and multi-submap systems. Section 4 highlights a few of the options for how the modules can be put together to form different types of systems.

This paper only provides a conceptual outline of a proposed sensor fusion architecture. The work is still ongoing and there are many details left to fill out. The goals are extremely ambitious, and much research and experimentation with distributed sensor fusion algorithms is required in order to design an architecture that has enough flexibility to handle most of the promising algorithms.

1.4 Related work

Klinker et al (2000) present a high-level architecture for AR systems that track their position by combining data from infrastructure-based sensors and wearable sensors. The block diagram shows some sensor fusion being done on each user's wearable computer, while other parts of the sensor fusion are accomplished on a room computer, with data passed across a wireless LAN. The paper is focused on AR system architecture and does not provide detail on the architecture of the sensor fusion subsystem. Conversely, Welch (1997) discusses Kalman filtering for tracking and auto-mapping, but does not strive for a generalized architecture. There is a huge body of literature on SLAM, a few examples of which will be cited below, mostly striving to find a scalable algorithm for large maps. The plug-and-play features based on self-describing sensor modules are closely related to the Smart Transducer Interface of the IEEE 1451.2 standard (IEEE, 1997).

2. Sensor Fusion Core Architecture

2.1 The sensor/target abstraction

There are a huge variety of different sensors used in tracking and navigation applications, based on almost all known electromagnetic, acoustic, and mechanical principles. Each sensor has different performance, operating environments, range, size, cost, update rate, penetration capabilities, and so on which make it more or less suitable for various applications. From the sensor fusion algorithm's point of view, however, most sensors belong to one of a few basic categories according to the type of measurement observable they produce:

Dead-reckoning continuous observables

- **Angular rate** (e.g. gyros)
- **Linear acceleration** (e.g. accelerometers)
- **Relative displacement** (wheel encoder, optical flow)
- **Relative direction** (steering encoder, optical flow)

Map-point-referenced discrete update observables

- **Range** (e.g. ultrasonic TOF, RF TOF, Lidar, radar, raw GPS, depth cameras, pull-string encoder)
- **1-D Bearing** (e.g. spinning laser scanner, linear CCD, 1-D phased array)
- **2-D Bearing** (e.g. imaging sensors, servoed pan/tilt beam tracker, scanning radar)
- **Range rate** (Doppler radar, phase-coherent acoustic)
- **Range difference** (e.g. TDOA acoustic or radar)
- **Dipole field component** (e.g. active source magnetic tracker, electric field tracker)

Earth-referenced discrete update observables

- **Cartesian position** (e.g. cooked GPS, altimeter)
- **Heading** (e.g. magnetic compass, dual-antenna GPS heading system, gyrocompass)
- **Speed** (optical flow, airspeed, waterspeed sensors)
- **Direction** (e.g. optical flow)
- **Homogeneous field** (magnetometers, gravimeters,)

We will assume that the sensor types listed under dead-reckoning are used for that purpose, and therefore enter into the Kalman filter as known control inputs in the time-update equations. Although this usually gives the best results, some authors have used these types of observables as Kalman filter measurements instead, and the architecture does not prevent them from being used as earth-referenced update observables instead. These will be called internal sensors, and their processing is handled separately in the Dynamics Reference Unit (DRU). The rest of the observables are treated as measurements and their data is fetched and interpreted by the Measurement Management Unit (MMU).

The map-point-referenced observables are the most common, and only these types of measurements actually contribute to the environment mapping process. All these observables make a measurement between two points; call them a sensor and a target. For convenience we will define

the sensor for a bearing angle measurement to be the device at the fulcrum of the bearing angle measurement, even though in the case of the laser scanner and related structured light projectors, this means the "sensor" is the light projector, and the "target" is the photodetector. Each measurement, z_{st} is a function of the pose λ_{st} between a particular sensor, s , and target, t , as well as the biases k_s associated with the sensor and the biases k_t associated with the target:

$$z_{st} = h_{st}(\lambda_{st}, k_s, k_t) \quad (1.2)$$

For an inside-out measurement, the sensor is attached to the vehicle and the target is fixed in the environment, while for an outside-in measurement a sensor fixed in the environment observes a target fixed to the moving vehicle.

2.2 The main object classes: PSEs, maps, filters

The environment map consists of a collection of sensors and targets, each having a pose with respect to the environment map coordinate reference frame, also called the navigation frame or n-frame. Likewise the vehicle has a number of sensors and targets rigidly attached to it, each with a pose relative to the vehicle frame, called the body-frame or b-frame. We define a class of objects called pose-sensing endpoints (PSEs) that may be sensors or targets of any type. Every PSE has the following attributes:

- Position of it's center of measurement, relative to the parent map to which it is attached.
- Orientation relative to parent map.
- Bias parameters vector k .
- Basic type. One of the following:
 - 2DBS (2-D bearing sensor)
 - 1DBS (1-D bearing sensor)
 - BT (bearing sensor target)
 - RS (range sensor)
 - RT (range sensor target)
 - RRS (range rate sensor)
 - RRT (range rate sensor target)
 - DT (dipole target or source)
 - DS (dipole field component sensor)
- Specific type. Mfg. and model number, which is used by the MMU to look up all the detailed information such as range, update rate, etc.
- Status (ready, busy, etc.)

The position of the PSE is always important, since it is impossible to make use of a measurement between two points without knowing where the points are. The orientation parameters may not all be important for certain PSEs. They are all important for bearing sensors, but the other types generally have circular symmetry about the principal axis, which means the roll is irrelevant to the measurement process. For ideal omnidirectional range

sensors, even the principal axis direction is unimportant. When these attributes are non-critical, their values can be set approximately, and no corresponding filter states need be allocated. For example, for a circular bull's-eye fiducial mark, the principal-axis vector is only used to determine whether the mark is visible or is too slanted relative to the view vector, and perhaps to calculate a correction for centroid offset due to perspective distortion. Since this is a second order correction, it should suffice to enter $[001]^T$ for all the fiducials mounted on the ceiling pointing down.

The PSEs are the atomic objects in this architecture. Each vehicle has a Sensor Fusion Core (SFC) associated with it to perform its own tracking and auto-calibration, and to build its own private version of the part of the map it sees, which may or may not be shared or merged with map fragments developed by other vehicles. The SFC may either be a dedicated processor running on-board the vehicle, or just a process or thread running inside a server somewhere. Each SFC has a DRU containing the vehicle state, a vehicle map containing pointers to all the PSEs attached to the vehicle, and an environment map containing pointers to all the environment-fixed PSEs of which the vehicle is so far aware.

2.3 Complementary EKF for estimating tracking, calibration and mapping errors

We now introduce some expanded notation to handle full STACEM with a complementary filter. The complementary filter (e.g. Brown & Huang, 1993) implements the extended Kalman filter (EKF) by using an ordinary linear Kalman filter to estimate errors in the underlying states of interest. A linear model for the propagation of error states is obtained by perturbation analysis from the nonlinear dynamics model of the underlying system. The underlying states are not part of the filter, and may be stored and propagated separately. In the proposed architecture, the underlying states consist of the vehicle states $\hat{\mathbf{X}}$ which are stored in the DRU, the states $\mathbf{M} = [\mathbf{m}_1^T \ \mathbf{m}_2^T \ \dots \ \mathbf{m}_M^T]^T$ of the M mobile PSEs (MPSEs) that are stored in the vehicle map manager, and $\mathbf{F} = [\mathbf{f}_1^T \ \mathbf{f}_2^T \ \dots \ \mathbf{f}_N^T]^T$ of the N fixed PSEs (FPSEs) that belong to the environment map manager. Although the vehicle map "state vector" $\mathbf{M}$ and environment map "state vector" $\mathbf{F}$ have been written using stacked vector notation, they are actually more likely implemented as arrays of pointers to PSE objects, each of which stores the actual state information (position, orientation, biases) in three separate fields of a structure. Which format is used internally to store the orientation depends on the PSE class implementation, but access methods must be provided to return the principal and secondary axis vectors, since these are used by the MMU

for measurement selection and prediction. Another method is provided which allows the PSE to update its underlying orientation representation with an incremental rotation vector returned from the complementary KF. Each individual PSE "vector" $\mathbf{f}_i$ is composed of position, orientation and intrinsic parameter (bias) fields:

$$\mathbf{f}_i = \begin{bmatrix} \mathbf{f}_i^p \\ \mathbf{f}_i^o \\ \mathbf{f}_i^k \end{bmatrix} = \begin{bmatrix} \mathbf{f}_i^x \end{bmatrix} \quad (1.3)$$

where $\mathbf{f}_i^x$ stands for the combined pose states.

Corresponding to the underlying system "state vector" $\mathbf{X} = [\mathbf{x}^T \ \mathbf{M}^T \ \mathbf{F}^T]^T$ there is an estimation error state vector $\delta\mathbf{X} = [\delta\mathbf{x}^T \ \delta\mathbf{M}^T \ \delta\mathbf{F}^T]^T$ which is a real vector and serves as the state vector of the complementary KF which maintains the covariance matrix

$$\mathbf{P}_{\mathbf{X}\mathbf{X}} = \begin{bmatrix} \mathbf{P}_{\mathbf{X}\mathbf{X}} & \mathbf{P}_{\mathbf{X}\mathbf{M}} & \mathbf{P}_{\mathbf{X}\mathbf{F}} \\ \mathbf{P}_{\mathbf{M}\mathbf{X}} & \mathbf{P}_{\mathbf{M}\mathbf{M}} & \mathbf{P}_{\mathbf{M}\mathbf{F}} \\ \mathbf{P}_{\mathbf{F}\mathbf{X}} & \mathbf{P}_{\mathbf{F}\mathbf{M}} & \mathbf{P}_{\mathbf{F}\mathbf{F}} \end{bmatrix}$$

In general, lowercase letters represent vectors and boldface uppercase letters represent matrices, but we have made an exception using the uppercase $\mathbf{X}$, $\mathbf{M}$, and $\mathbf{F}$ to represent stacked vectors. Further abusing notation, write $\hat{\mathbf{X}} = \mathbf{X} \oplus \delta\mathbf{X}$, where the circle around the plus sign reminds us that $\mathbf{X}$ (the "true" values of the underlying states) and $\hat{\mathbf{X}}$ (the estimated values of the underlying states) are not really vectors, or if they are, they are not of the same dimension as $\delta\mathbf{X}$ (due to the fact that underlying orientation requires at least 4 parameters to represent without a singularity, whereas $\delta\mathbf{X}$ contains just the three parameters of the small-angle error rotation vectors, in order to avoid redundant states in the Kalman filter). The position, orientation and bias components of each individual PSE error vector are stacked:

$$\delta\mathbf{F} = [\delta\mathbf{f}_1^T \ \delta\mathbf{f}_2^T \ \dots \ \delta\mathbf{f}_N^T]^T$$

$$\delta\mathbf{f}_i = \begin{bmatrix} \delta\mathbf{f}_i^p \\ \delta\mathbf{f}_i^o \\ \delta\mathbf{f}_i^k \end{bmatrix} = \begin{bmatrix} \delta\mathbf{f}_i^x \end{bmatrix}$$

and these are all real vectors this time. Corresponding notation applies to the MPSEs.

Now consider the time update of the complementary Kalman filter (CKF). As in Section 1.1, both the MPSE states and the FPSE states are assumed constant, so the

simplified Φ and Q matrices of Eq. (1.1) hold, and the computation time is linearly related to the number M of MPSEs and the number N of FPSEs:

$$\begin{aligned} \widehat{\delta \mathbf{x}}(k+1) &= \widehat{\delta \mathbf{x}}(k) + \Phi_{\mathbf{x}}(k) \widehat{\delta \mathbf{x}}(k) \\ \mathbf{P}_{\mathbf{xx}}(k+1) &= \Phi_{\mathbf{x}}(k) \mathbf{P}_{\mathbf{xx}}(k) \Phi_{\mathbf{x}}(k)^T + \mathbf{Q}_{\mathbf{x}}(k) \\ \text{for } j &= 1, \dots, M: \\ \mathbf{P}_{\mathbf{x}m_j}(k+1) &= \Phi_{\mathbf{x}}(k) \mathbf{P}_{\mathbf{x}m_j}(k) \\ \mathbf{P}_{m_j \mathbf{x}}(k+1) &= \mathbf{P}_{\mathbf{x}m_j}^T(k+1) \\ \text{end} \\ \text{for } i &= 1, \dots, N: \\ \mathbf{P}_{\mathbf{x}f_i}(k+1) &= \Phi_{\mathbf{x}}(k) \mathbf{P}_{\mathbf{x}f_i}(k) \\ \mathbf{P}_{f_i \mathbf{x}}(k+1) &= \mathbf{P}_{\mathbf{x}f_i}^T(k+1) \\ \text{end} \end{aligned}$$

Note that to perform these operations, the CKF does not need to know anything about the structure or physical meaning of the underlying states $\mathbf{X}$ or even the error state vector $\delta \mathbf{X}$ that it is estimating. It just needs to be handed the matrices $\Phi_{\mathbf{x}}(k)$ and $\mathbf{Q}_{\mathbf{x}}(k)$ which encapsulate everything that needs to be known about the performance and configuration of the inertial or dead-reckoning sensors, the integration algorithms applied to them and associated error propagation models, and the current vehicle pose and dynamics. This is why the calculation of $\Phi_{\mathbf{x}}(k)$ and $\mathbf{Q}_{\mathbf{x}}(k)$ is left to the DRU, which already has all this source information available because it has the job of reading the dead-reckoning sensor driver and integrating the rate and acceleration information to keep track of the current vehicle state.

For the measurement update step, it is assumed that on each measurement cycle the MMU obtains one map-point-referenced measurement $\mathbf{z}_{ij}$ between the i^{th} FPSE and the j^{th} MPSE or one earth-referenced measurement $\mathbf{z}_{0j}$ between the j^{th} MPSE and the "world". The MMU calculates and delivers the innovation $\mathbf{v}_{ij} = \mathbf{z}_{ij} - \hat{\mathbf{z}}_{ij}$ which is a vector of measurement dimension $m=1,2$ or 3. The MMU also calculates and delivers the three Jacobians which form the only nonzero entries in the linearized measurement matrix:

$$\mathbf{H}_{ij} = \begin{bmatrix} \mathbf{H}_{\mathbf{x}} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{H}_{m_j} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{H}_{f_i} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix}$$

$$\mathbf{H}_{\mathbf{x}} \equiv \frac{\partial(\mathbf{z}_{ij})}{\partial(\delta \mathbf{x}_v)} \quad \mathbf{H}_{m_j} \equiv \frac{\partial(\mathbf{z}_{ij})}{\partial(\delta \mathbf{m}_j)} \quad \mathbf{H}_{f_i} \equiv \frac{\partial(\mathbf{z}_{ij})}{\partial(\delta \mathbf{f}_i)}$$

with which the CKF computes its Kalman gains and updates its covariance matrix:

$$\begin{aligned} \mathbf{B}_{ij} &\equiv \mathbf{P}_{\mathbf{xx}} \mathbf{H}_{ij}^T = \\ &\begin{bmatrix} \mathbf{P}_{\mathbf{xx}} & \mathbf{P}_{\mathbf{x}m_1} & \dots & \mathbf{P}_{\mathbf{x}m_M} & \mathbf{P}_{\mathbf{x}f_1} & \dots & \mathbf{P}_{\mathbf{x}f_N} \\ \mathbf{P}_{m_1 \mathbf{x}} & \mathbf{P}_{m_1 m_1} & \dots & \mathbf{P}_{m_1 m_M} & \mathbf{P}_{m_1 f_1} & \dots & \mathbf{P}_{m_1 f_N} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{m_M \mathbf{x}} & \mathbf{P}_{m_M m_1} & \dots & \mathbf{P}_{m_M m_M} & \mathbf{P}_{m_M f_1} & \dots & \mathbf{P}_{m_M f_N} \\ \mathbf{P}_{f_1 \mathbf{x}} & \mathbf{P}_{f_1 m_1} & \dots & \mathbf{P}_{f_1 m_M} & \mathbf{P}_{f_1 f_1} & \dots & \mathbf{P}_{f_1 f_N} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{f_N \mathbf{x}} & \mathbf{P}_{f_N m_1} & \dots & \mathbf{P}_{f_N m_M} & \mathbf{P}_{f_N f_1} & \dots & \mathbf{P}_{f_N f_N} \end{bmatrix} \begin{bmatrix} \mathbf{H}_{\mathbf{x}}^T \\ \vdots \\ \mathbf{H}_{m_j}^T \\ \vdots \\ \mathbf{H}_{f_i}^T \\ \vdots \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{P}_{\mathbf{xx}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{P}_{\mathbf{x}m_j} \mathbf{H}_{m_j}^T + \mathbf{P}_{\mathbf{x}f_i} \mathbf{H}_{f_i}^T \\ \mathbf{P}_{m_1 \mathbf{x}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{P}_{m_1 m_j} \mathbf{H}_{m_j}^T + \mathbf{P}_{m_1 f_i} \mathbf{H}_{f_i}^T \\ \vdots \\ \vdots \\ \vdots \\ \mathbf{P}_{f_N \mathbf{x}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{P}_{f_N m_j} \mathbf{H}_{m_j}^T + \mathbf{P}_{f_N f_i} \mathbf{H}_{f_i}^T \end{bmatrix} \end{aligned}$$

$$\begin{aligned} \mathbf{S}_{ij} &= \mathbf{H}_{ij} \mathbf{P}_{\mathbf{xx}} \mathbf{H}_{ij}^T + \mathbf{R}_{ij} = \mathbf{H}_{ij} \mathbf{B}_{ij} + \mathbf{R}_{ij} \\ &= \mathbf{H}_{\mathbf{x}} \mathbf{P}_{\mathbf{xx}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{H}_{\mathbf{x}} \mathbf{P}_{\mathbf{x}m_j} \mathbf{H}_{m_j}^T + \mathbf{H}_{\mathbf{x}} \mathbf{P}_{\mathbf{x}f_i} \mathbf{H}_{f_i}^T + \\ &\quad \mathbf{H}_{m_j} \mathbf{P}_{m_j \mathbf{x}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{H}_{m_j} \mathbf{P}_{m_j m_j} \mathbf{H}_{m_j}^T + \mathbf{H}_{m_j} \mathbf{P}_{m_j f_i} \mathbf{H}_{f_i}^T + \\ &\quad \mathbf{H}_{f_i} \mathbf{P}_{f_i \mathbf{x}} \mathbf{H}_{\mathbf{x}}^T + \mathbf{H}_{f_i} \mathbf{P}_{f_i m_j} \mathbf{H}_{m_j}^T + \mathbf{H}_{f_i} \mathbf{P}_{f_i f_i} \mathbf{H}_{f_i}^T + \mathbf{R}_{ij} \end{aligned}$$

$$\mathbf{K}_{ij} = \mathbf{B}_{ij} \mathbf{S}_{ij}^{-1}$$

$$\begin{aligned} \mathbf{P}_{\mathbf{xx}}(+) &= (\mathbf{I} - \mathbf{K}_{ij} \mathbf{H}_{ij}) \mathbf{P}_{\mathbf{xx}}(-) (\mathbf{I} - \mathbf{K}_{ij} \mathbf{H}_{ij})^T + \mathbf{K}_{ij} \mathbf{R}_{ij} \mathbf{K}_{ij}^T \\ &= \mathbf{P}_{\mathbf{xx}}(-) - \mathbf{B}_{ij} \mathbf{K}_{ij}^T - (\mathbf{B}_{ij} \mathbf{K}_{ij}^T)^T + \mathbf{K}_{ij} \mathbf{S}_{ij} \mathbf{K}_{ij}^T \end{aligned}$$

Finally, the CKF can update its states:

$$\begin{aligned} \widehat{\delta \mathbf{x}}(+) &= \widehat{\delta \mathbf{x}}(-) + \mathbf{K}_{ij}^{\mathbf{x}} \mathbf{v}_{ij} \\ \widehat{\delta \mathbf{M}}(+) &= \widehat{\delta \mathbf{M}}(-) + \mathbf{K}_{ij}^{\mathbf{M}} \mathbf{v}_{ij} \\ \widehat{\delta \mathbf{F}}(+) &= \widehat{\delta \mathbf{F}}(-) + \mathbf{K}_{ij}^{\mathbf{F}} \mathbf{v}_{ij} \end{aligned}$$

where $\mathbf{K}_{ij}^{\mathbf{x}}$, $\mathbf{K}_{ij}^{\mathbf{M}}$, $\mathbf{K}_{ij}^{\mathbf{F}}$ represent the appropriate partitions of the Kalman gain matrix. At any time, normally after each measurement cycle is complete, the CKF can transfer $\widehat{\delta \mathbf{x}}$, $\widehat{\delta \mathbf{M}}$, and $\widehat{\delta \mathbf{F}}$ to the DRU, the vehicle map manager and the environment map manager respectively to correct the underlying states, and then reset its own error state estimates to zero.

2.4 Federated filtering algorithm for decoupling mapping and auto-calibration

The preceding section introduced the notation for a complementary Kalman filter that performs simultaneous tracking, auto-calibration and map-building. The centralized filter produces theoretically optimal results, but it limits the

flexibility of the architecture. The auto-calibration process and auto-mapping process are intertwined through the cross-correlations P_{MF} and cannot be used separately one at a time. In practice it may be desirable to first perform the auto-calibration of the vehicle's sensors using a small pre-surveyed map, then to freeze the vehicle map after it has converged below some small covariance level, and then set out to map the larger space using the pre-calibrated vehicle. Once the whole environment of interest is mapped it may even be beneficial to turn off the mapping and just use the system for real time tracking. These things are cumbersome with the centralized filter. A second drawback of the centralized filter is that the computation cost of the measurement update is proportional to the square of the total number of states. If the number of auto-calibration parameters is high (for example, the UNC Hi-Ball sensor has 26 views, each with 10 internal and 6 external parameters, totaling 416 states) then decoupling this from a similar sized map (about 138 LEDs) would almost halve the computation. Most importantly, the centralized filter does not scale well with large map sizes N . To overcome this, the architecture should facilitate the use of SLAM strategies based on breaking the map into smaller submaps (e.g. Leonard & Fedder, 2000). The first step is to remove the auto-mapping states from the SLAM filter so that they don't have to be replicated in each submap.

The proposed architecture exploits the Federated Kalman

filter algorithm (Carlson, 1990) to provide distributed and fault-tolerant filters for aircraft navigation systems. In the original application, there was no environment map, just a vehicle with an INS and a handful of earth-referenced sensors (GPS, baro-altimeter, SAR, radar altimeter, etc). Each sensor has a few bias states, and the centralized filter stacks them all up together with the vehicle INS states in the centralized state vector:

$$X_{centralized} = [x^T \quad k_1^T \quad k_2^T \quad k_3^T]^T.$$

In a federated system, separate local filters (LFs) are created for each of the three sensors, with state vectors:

$$X_{LF1} = [x_1^T \quad k_1^T]^T \quad X_{LF2} = [x_2^T \quad k_2^T]^T \quad X_{LF3} = [x_3^T \quad k_3^T]^T$$

Because there are redundant copies of the vehicle state vector x in each LF, the total information about x has to be split up amongst the LFs. This is done by multiplying $P_{xx}(0)$ by a factor γ (in this case $\gamma = 3$) and initializing all the LFs with this inflated covariance. Each LF then performs a time update, using $Q_{LFi} = \gamma Q$, and then processes the measurement from only its own sensor. Finally, a master filter (MF) fuses the common states of all the LFs together using the information fusion form of KF update:

$$P_{xxFused} = (P_{xxLF1}^{-1} + P_{xxLF2}^{-1} + P_{xxLF3}^{-1})^{-1}$$

$$\hat{x}_{MF} = P_{xxFused} (P_{xxLF1}^{-1} \hat{x}_{LF1} + P_{xxLF2}^{-1} \hat{x}_{LF2} + P_{xxLF3}^{-1} \hat{x}_{LF3})$$




Figure 1: Block diagram of the vehicle-level sensor fusion core (SFC) architecture showing data flow during tracking. Blocks which maintain state information show the names of their internal states in a data store icon.

At this point all the information about $\mathbf{x}$ has been transferred into the MF, whose covariance $\mathbf{P}_{\mathbf{x}\mathbf{x}}^{Fused}$ represents the true system covariance, while all the information about each sensor's biases remains in its own LF. To start the next cycle, the vehicle information from the MF gets split up and distributed back to the LFs using Carlson's Fusion Reset equations. If the master filter has dimension less than the centralized filter, as in the version just described, there will be some loss of optimality, but in the examples shown in Carlson's papers the difference is not noticeable.

In the Federated filter, each measurement is processed by just one LF. However, the STACEM problem posed in Section 2.3 has an extra complication. If the STAC process is assigned to one LF and the STEM process to another LF, then each measurement will depend on states in both filters. We solve this by using the Principle of Measurement Reproduction (Tupsey, 1998), which states that a measurement $\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{v}$, $\mathbf{v} \sim N(\mathbf{0}, \mathbf{R})$ is exactly equivalent to two repetitions of the measurement, each with twice the original noise:

$$\begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{H} \\ \mathbf{H} \end{bmatrix} \mathbf{x} + \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix}, \quad \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} \sim N\left(\mathbf{0}, \begin{bmatrix} 2\mathbf{R} & \mathbf{0} \\ \mathbf{0} & 2\mathbf{R} \end{bmatrix}\right)$$

Thus, each measurement is sent to all the LFs, and each LF multiplies the measurement noise matrix by γ so that the information in the measurement is divided up between the LFs. Each LF then acts independently, as if the measurement it received were just for it, and bumps up its noise to account for states not included within itself, just as it would if it were an ordinary stand-alone filter tuned to account for unmodeled states by standard reduced-order KF design methods.

Figure 1 summarizes the architecture described above. The underlying states are maintained in the dynamics reference unit, the vehicle map manager, and the environment map manager. The STAC filter estimates errors in the underlying MPSE poses and biases, and sends the corrections vector $\delta\mathbf{M}$ back to the vehicle map manager, which bears responsibility for applying the corrections to the appropriate PSE fields with appropriate translation of orientation format. In like fashion, the STEM filter sends error correction estimates back to the environment map manager. A third LF for simple tracking only (STO filter) has been added, so that if the user or the system decides that both auto-calibration and mapping are substantially completed, both the STAC and STEM filters can be turned off. Any one, two or three of the LFs may be active at a given time, with γ set to the number active. Figure 1 only shows what data is passed between the

modules during active tracking. Much of the flexibility of the architecture stems from how it detects hardware and initializes itself, which will be touched on briefly in the following sections.

2.5 Dynamic Reference Unit (DRU) and Internal Sensor Driver

At initialization, the DRU queries the internal sensor driver to find out what kind of dead-reckoning sensors are available, if any. Most probably, the answer will be either 1) nothing, 2) an IMU, 3) a pair of differential IMUs, or 4) a set of wheel encoders. Based on the answer, the DRU automatically determines whether tracking is desired in 3D or 2D, and chooses an integration algorithm with the appropriate number of states. Next, the DRU asks the driver for initial uncertainty specifications of a number of standard parameters appropriate to the integration model. For example, if it is an IMU, these would include rms initial uncertainty estimates for gyro bias, gyro bias ramp, gyro misalignments, gyro scale factors, gyro quadratic nonlinearities, etc. and corresponding accel parameters. If the driver is unable to obtain this information from the IMU itself because it is an off-the-shelf device not specifically designed for compatibility with this architecture, then the driver will try to retrieve it from a file, and if that doesn't exist it will query the user and create the file. Based on the sizes of these initial uncertainties, the DRU decides whether to create KF states to estimate them in order to apply improved fine compensation in addition to the compensation already done in the IMU or the driver.

If smart sensor modules are specifically designed for use with this architecture, they may use an IEEE 1451.2 interface, which specifies a transducer electronic data sheet (TEDS) be resident in EEPROM on the transducer that can tell the application what type of sensors are present and provide calibration data. The driver then passes along to the DRU the part of the information it needs, and compensates for sensor biases, scale factors and misalignments for which factory calibration data is stored. For non IEEE 1451.2 sensor interfaces, the driver emulates this process, delivering to the DRU compensated run-time data and initialization uncertainty information.

Finally, during acquisition after initialization, the DRU obtains the initial pose estimate from the MMU over the dotted line connection in Figure 1. It also obtains a list of any shared bias states that affect all the external measurements of a given type, such as room temperature for ultrasonic systems, and adds these states to $\mathbf{x}$.

2.6 Measurement Management Unit (MMU) and external sensor/target meta-driver

Here is where we really exploit the sensor/target abstraction of Eq. (1.2) to partition the work of preparing

the measurement matrices H_x , H_{mj} , and H_{fi} into generic and model-specific parts, keeping the generic operations inside the MMU and pushing the sensor/target details out into the driver. The result is that the architecture is extendible by outside parties: a sensor manufacturer who understands their own product and can model its output, but knows nothing of Kalman filtering or SLAM theory, could follow some instructions and write a driver that would enable the sensor to work with the sensor fusion core and be optimally fused with the inertial sensors and any other aiding sensors that are attached. The goal is that this would be accomplished just by loading a driver and plugging in the sensor – truly Kalman filter design for dummies! The key to this partitioning was to invent a symmetrical and very sensor-friendly representation of the relative 6-DOF pose between the i^{th} FPSE and the j^{th} MPSE:

$$\lambda_{ij} = (\theta_i, \phi_i, \theta_j, \phi_j, l_{ij}, t_{ij}) \quad (1.4)$$

Here θ_i is the angle between the FPSE principal axis and the line of length l_{ij} connecting the FPSE and the MPSE.

ϕ_i is the azimuthal angle of this line around the FPSE's principal axis (forward or x-axis), using a right hand rule and the FPSE's y-axis as the starting mark. θ_j and ϕ_j are a symmetrically defined off-axis angle and azimuthal angle of the departing ray from the MPSE. t_{ij} is the twist of the imaginary rigid pole connecting the centers of the two PSEs about a swivel joint in the pole center, where the twist of each pole half relative to its PSE is fixed by imagining that the pole flattens out like a canoe paddle at the end, and the intersection of this blade with the PSE's y-z plane lies along the y-axis.




Figure 2: Measurement Management Unit

Figure 2 illustrates at a high level the functionality of the MMU. The first task of the MMU is to decide which measurements to request from the driver. The measurement scheduler consults the sensor/target status server in the meta-driver during initialization to find out what types of sensors exist with what capabilities, repetition rates, triggering protocols etc. It also obtains initial bias parameter vectors and uncertainties from any sensors that can produce a self-describing TEDS, and transmits the information to the vehicle and environment map managers for map initialization. During tracking, after receiving the vehicle map and the nearby portion of the environment map, it checks the server to see which sensors are busy or ready, which may in turn cause the local status server to consult with a regional status server if there are sensor or active target resources that are shared with other vehicles. It may also check with a local or regional geometric reasoning module provided by the application to predict which sensor-target pairs are likely to be occluded, suggest good places to look for natural features, etc. Having culled down the number of possible measurements that are available, it evaluates them using the current vehicle state and covariance information to see which will produce the greatest information gain (selfishly for this vehicle) and thus

selects the i^{th} FPSE and the j^{th} MPSE. It passes their underlying state information to the generic pose predictor which calculates the relative pose in the form of Eq. (1.4) and then takes advantage of the symmetry to swap theta and phi of the sensor into the first position. This predicted sensor-target pose is sent to the driver which selects the appropriate manufacturer-specific measurement model and calculates the measurement prediction and the noise matrix, taking into account all the fine details of the sensor and its uncompensated distortions known only by the manufacturer or the engineer who modeled the sensor and wrote its driver. The driver also returns Jacobians with respect to the pose parameters, evaluated at the predicted values. In the MMU, these get combined by the chain rule (i.e. matrix multiplication of Jacobians) with the generic parts of the Jacobians, $\Lambda_s, \Lambda_m, \Lambda_r$, to produce the H matrices needed by the complementary Kalman filters. Finally, the MMU calculates range gates (or image windows for a camera) for the predicted measurement, and sends them to the hardware driver to fetch an actual measurement. It performs nearest-neighbor data association (or a more sophisticated algorithm if needed) on the set of sensor returns, and sends the innovation to the KF for processing.

3. Multi-Vehicle Coordination Architecture

3.1 Distribution of computation to vehicle filters

Diagram of multiple vehicles communicating with a coordination center to coordinate use of shared infrastructural sensor/target resources, and to share map info.

Warning about map fusion introducing non-independent information. Possible workaround using Covariance Intersection (CI).

3.2 Implementation Options

Distributed versus Centralized Implementations
 Hardware Implementations
 Software Implementations
 Mixed Implementations
 Timing and Synchronization Issues

4. Conclusions and future work

A broad-brush outline has been presented for a new sensor fusion architecture that is being designed to facilitate rapid development and reconfiguration of systems that perform tracking, auto-calibration, and auto-mapping using a variety of different types of sensors fused together.

We are working on developing a vision-aided inertial tracker for see-through AR systems based on this architecture, which we plan to demonstrate at ISAR2001 as a practical complement to this paper if it is accepted.

Beyond that, there is much more theoretical work left to do to fill in the details of the architecture and achieve some of its loftier goals, including scalable mapping algorithms. The federated filter has been selected as a nice framework for experimentation with algorithms that involve a group of submaps of which only a subset are active at a time. However, more work is required to identify the most efficient yet consistent algorithms for swapping the active set of submaps, or to find algorithms for combining and repartitioning submaps dynamically. Also, future work must address the issues of how to merge the submaps generated by different vehicles at a regional coordination node, even though they may not be based exclusively on independent information sources, and how to re-integrate these merged maps back to the vehicles.

5. References

- Brown, R.G. & Hwang, P.Y.C. (1992). Introduction to Random Signals and Applied Kalman Filtering. New York: John Wiley & Sons, Inc.
- Carlson (1990). Federated square root filter for decentralized parallel processes. IEEE Transactions on Aerospace and Electronic Systems, AES 26:3
- IEEE Instrumentation and Measurement Society (1997). IEEE Std 1451.2-1997: Smart transducer interface for sensors and actuators- transducer to microprocessor communication protocols and transducer electronic data sheet (TEDS) formats.
- Klinker, G., Reicher, T. & Brugge, B. (2000). Distributed user tracking concepts for augmented reality applications. Proceedings of IEEE ISAR 2000, Munich
- Leonard, J.J., & Feder, H.J.S. (2000). A computationally efficient method for large-scale concurrent mapping and localization. To appear in Robotics Research: The Ninth International Symposium, J. Hollerbach & D. Koditschek (Ed.s) London: Springer-Verlag.
- Smith, R., Self, M., & Cheeseman, P. (1987). A stochastic map for uncertain spatial relationships. Fourth International Symposium on Robotics Research, MIT Press
- Tupyshev, V. A. (1998). A generalized approach to the problem of distributed Kalman filtering. AIAA Guidance, Navigation & Control Conference, AIAA 98-4309, Boston
- Welch, G. & Bishop, G. (1997). Single-constraint-at-a-time tracking. In SIGGRAPH 97 Conference Proceedings, ACM Annual Conference Series. ACM Press

What is claimed is:

1. A method for tracking an object comprising:
coupling a sensor subsystem to an estimation subsystem, said sensor subsystem enabling measurement related to relative locations or orientations of sensing elements;
accepting configuration data from the sensor subsystem;
configuring the estimation system according to the accepted configuration data;
repeatedly updating a state estimate, including
accepting measurement information from the sensor subsystem, and
updating the state estimate according to the accepted configuration data and the accepted measurement data.
2. The method of claim 1 wherein coupling the sensor subsystem to the estimation subsystem includes coupling software modules each associated with one or more of the sensing elements.
3. The method of claim 2 wherein each of the software modules provides a software interface for receiving information related to an expected sensor measurement and providing measurement information that depends on said received information.
4. The method of claim 3 wherein each of the software modules implements calculations that are independent of a representation of the state in the estimation subsystem.
5. The method of claim 1 wherein the state estimate characterizes an estimate of a location of the object.
6. The method of claim 1 wherein the state estimate characterizes configuration information for one or more sensing elements fixed to the object.

7. The method of claim 6 wherein the configuration information for the one or more sensing elements fixed to the object includes information related to position or orientation of said sensing elements relative to the object.

8. The method of claim 6 wherein the configuration information for the one or more sensing elements fixed to the object includes operational parameters for the one or more sensing elements.

9. The method of claim 1 wherein the state estimate characterizes configuration information for one or more sensing elements fixed in an environment of the object.

10. The method of claim 9 wherein the configuration information for one or more sensing elements fixed in the environment of the object includes a map of the locations of said sensing elements.

11. The method of any of claims 1 through 10 wherein repeatedly updating the state further includes:

providing to the sensor subsystems information related to an expected sensor measurement; and

wherein accepting the measurement information from the sensor subsystem includes accepting information related to an actual sensor measurement.

12. The method of claim 11 wherein providing the information related to an expected sensor measurement includes providing information related to a relative geometric configuration of two of the sensing elements.

13. The method of claim 12 wherein providing information related to a relative geometric configuration of the two of the sensing elements includes providing information characterizing a relative location of said sensing elements.

14. The method of claim 11 wherein accepting the information related to an actual sensor measurement includes accepting information enabling the estimation subsystem to calculate a difference between the actual measurement and the expected measurement.

15. The method of claim 11 wherein accepting the information related to an actual sensor measurement includes accepting information for correlating measurements and geometric relationships between sensing elements.

16. The method of claim 15 wherein the information for correlating measurements and geometric relationships between sensing elements includes a mapping between a relative pose of the sensing elements and a sensor measurement.

17. The method of claim 16 wherein the mapping between the relative pose of the sensing elements and the sensor measurement characterizes a linear mapping.

18. The method of claim 11 wherein accepting the information related to an actual sensor measurement includes accepting information characterizing an uncertainty in the actual measurement.

19. The method of claim 18 wherein the information characterizing the uncertainty in the actual measurement includes parameters of a statistical distribution of an error of the actual measurement.

20. The method of claim 1 wherein repeatedly updating the state further includes:
selecting a pair of sensing elements for measurement; and
providing an identification of the selected pair to the sensing subsystem.

21. The method of claim 20 wherein selecting the pair of sensing elements includes selecting said elements according to an expected utility of a measurement associated with said elements to the updating of the state.

22. The method of any of claims 11 through 19 wherein repeatedly updating the state further includes:

updating the state according to the accepted information related to an actual sensor measurement.

23. The method of claim 20 wherein repeatedly updating the state further includes:
updating the state according to accepted measurements from inertial sensors.

24. The method of any of claims 1 through 23 wherein updating the state estimate includes applying a Kalman Filter approach.

25. The method of claim 1 wherein each of said sensing elements comprises at least one of a sensor and a target.

26. The method of claim 25 wherein the target comprises an active device that interacts with the sensor.

27. The method of claim 26 wherein the target comprises at least one of a man-made signal reflector and a natural feature of an environment.

28. The method of claim 1 wherein the object is selected from a group consisting of a vehicle, a robot, a person, a part of a person, a flying object, a floating object, an underwater moving object, an animal, a camera, a sensing apparatus, a helmet, a tool, a piece of sports equipment, a shoe, a boot, an article of clothing, a personal protective equipment, a rigid object having a dimension between 1 nanometer to 10^9 meters.

29. The method of claim 1 wherein the state estimate comprises information related to a position or an orientation of the object relative to a reference coordinate frame.

30. A tracking system comprising:

an estimation subsystem; and

a sensor subsystem coupled to the estimation module and configured to provide configuration data to the estimation module and to provide measurement information to the estimation module for localizing an object;

wherein the estimation module is configured to update a location estimate for the object based on configuration data and measurement information accepted from the sensor subsystem.

31. The system of claim 30 wherein the sensor subsystem includes one or more sensor modules, each providing an interface for interacting with a corresponding set of one or more sensing elements.

32. The system of claim 31 wherein the interface enables the sensor module to perform computations independently of an implementation of the estimation subsystem.

33. The system of claim 31 wherein the interface enables the estimation subsystem to perform computations independently of an implementation of the sensor modules.

34. The system of claim 30, further comprising a navigation subsystem to navigate the object in an environment based on the location estimate for the object.

35. A sensor module comprising:

a sensor interface for communicating with a measurement sensor;

a communication interface for communication with an estimation system;

wherein the sensor module is configured to

receive information related to an expected sensor measurement over the communication interface,

receive a measurement signal over the sensor interface,

provide measurement information based on the measurement signal over the communication interface.

36. The sensor module of claim 35 wherein the sensor module is configured to provide information over the communication interface related to an uncertainty in the measurement information.

37. The sensor module of claim 35 wherein the received information related to an expected sensor measurement includes a predicted pose of a sensing element relative to the measurement sensor.

38. A method comprising:

enumerating a set of sensing elements available to a tracking system that includes an estimation subsystem that estimates a position or orientation of an object; and providing parameters specific to the set of sensing elements to the tracking system to enable the estimation subsystem to be configured based on the parameters specific to the subset of sensing elements.

39. The method of claim 38, further comprising generating a sequence of candidates of pairs of sensing elements selected from the set of sensing elements, the sequence based on an expected utility of a measurement associated with said elements to the estimation subsystem.

40. The method of claim 39, further comprising selecting a pair of sensing elements from the sequence of candidates, the selected pair of sensing elements being ready to make a measurement at the time of selection of the pair or at a predefined time after the time of selection of the pair, the selected pair having a highest expected utility of a measurement among the sequence of candidates.

41. The method of claim 38 wherein the set of sensing elements comprises at least one sensor and at least one target, the sensor making a measurement with respect to the target.

42. The method of claim 41 wherein the target comprises a natural feature in an environment.

43. A method comprising:

computing an estimate of a pose of a target element relative to a sensor element based on an estimate of a pose of a tracked object relative to an environment having affixed thereto either the sensor element or the target element,

the computing of the estimate of the pose of the target element relative to the sensor element further based on an estimate of a pose of the affixed element relative to the tracked object and the other element relative to the environment;

computing an estimate of a measurement of the target made by the sensor based on the estimate of the pose of the target relative to the sensor;
making an actual measurement of the target by using the sensor; and
computing a deviation between the actual measurement and the estimated measurement;
and
generating a new estimate of the pose of the tracked object based on the deviation.

44. The method of claim 43, further comprising computing a first observation matrix that characterizes a linearized model of a function relating the measurement made by the sensor to the pose of the target relative to the sensor.

45. The method of claim 44, further comprising computing a second observation matrix that characterizes a linearized model of a function relating the pose of the target relative to the sensor to the estimate of the pose of the tracked object relative to the environment.

46. The method of claim 45, further comprising computing an observation matrix that characterizes a linearized model of a function relating the measurement made by the sensor to the pose of the tracked object relative to the environment by combining the first observation matrix and the second observation matrix.

47. A method comprising:

estimating a first value associated with a pose of a first sensing element relative to a second sensing element, the first sensing element fixed to an environment and the second sensing element fixed to an object being tracked, one of the first and second sensing elements being a sensor and the other being a target;
estimating a second value associated with a pose of the second sensing element relative to the first sensing element;
determining which of the first and second sensing elements is the sensor; and
generating an innovation of a measurement of the target made by the sensor based on the first value when the second sensing element is the sensor.

48. The method of claim 47, further comprising generating the innovation based on the second value when the first sensing element is the sensor.

49. The method of claim 47 wherein estimating the first value and estimating the second value are performed by a process ignorant of which of the first and second sensing elements is a sensor.

50. A method comprising

estimating a calibration parameter of a sensing element that is either a sensor or a target, the sensing element being fixed either to an environment or to an object being tracked;

determining whether the sensing element is the sensor or the target; and

assigning the calibration parameter as a sensor calibration parameter when the sensing element is a sensor, and generating an innovation of a measurement of a target made by the sensing element based in part on the sensor calibration parameter.

51. The method of claim 50, further comprising assigning the calibration parameter as a target calibration parameter when the sensing element is a target, and generating an innovation of a measurement of the sensing element made by a sensor based in part on the target calibration parameter.

52. The method of claim 50 wherein estimating the calibration parameter is performed by a process ignorant of whether the sensing element is a sensor or a target.

53. A method of using multiple sensors in a tracking system comprising:

providing an estimation module;

coupling one or more sensor modules to the estimation module, each associated with a different set of one or more sensors;

configuring the tracking system, including

providing configuration information from each of the sensor modules to the estimation module regarding the characteristics of the sensors associated with the sensor module, and

configuring the estimation module using the provided configuration information;
maintaining estimates of tracking parameters in the estimation module, including
repeatedly
passing data based on the estimates of the tracking parameters from the estimation
module to one or more of the sensor modules,
receiving from said one or more sensor modules at the estimation module data
based on measurements obtained from the associated sensors, and the data
passed to the sensor modules, and
combining the data received from said one or more sensor modules and the
estimates of the tracking parameters in the estimation module to update
the tracking parameters.

54. The method of claim 53 wherein the data passed from the estimation module to one or more of the sensor modules includes an estimate of the pose of a target relative to a sensor that was calculated by the estimation module using an estimate of the pose of a tracked object relative to a frame of reference fixed to an environment.

55. The method of claim 54 wherein the data passed from the estimation module to one or more of the sensor modules does not include the estimate of the pose of the tracked object relative to the frame of reference fixed to the environment.

56. The method of claim 53 wherein providing the estimation module includes providing a module that is configurable to use different sets of sensor modules coupled to it.

57. The method of claim 53 wherein maintaining estimates of the tracking parameters in the estimation module includes using a stochastic model in the estimation module.

58. The method of claim 57 wherein using a stochastic model includes implementing some or all of a Kalman filter in the estimation module.

59. The method of claim 58 wherein implementing some or all of the Kalman filter includes updating error estimates using linearized models of the sensor system.

60. The method of claim 58 wherein implementing some or all of the Kalman filter includes implementing a distributed Kalman filter, wherein each of a plurality of components of the distributed Kalman filter is associated with a different subset of the sensor modules.

61. The method of claim 60 wherein one of the components of the distributed Kalman filter is associated with a subset of sensor modules consisting of sensor modules that are affixed to a tracked object.

62. The method of claim 60 wherein one of the components of the distributed Kalman filter is associated with a subset of sensor modules consisting of sensor modules which are affixed to an environment.

63. The method of claim 60 wherein one of the components of the distributed Kalman filter is not associated with any sensor modules.

64. The method of claim 60 wherein implementing the distributed Kalman filter includes implementing a Federated Kalman Filter.

65. The method of claim 53 wherein providing configuration information from the sensor modules includes providing information characterizing a type of a sensor associated with a sensor module.

66. The method of claim 53 wherein providing configuration information from the sensor modules includes providing information characterizing a position or an orientation of a sensor associated with a sensor module.

67. The method of claim 53 wherein providing configuration information from the sensor modules includes providing information characterizing one or more calibration parameters of a sensor associated with a sensor module.

68. A machine-accessible medium, which when accessed results in a tracking or navigation system that tracks or navigates, respectively, an object, performing operations comprising:

enumerating a set of sensing elements available to the tracking or navigation system, the sensing elements available to the tracking or navigation system including at least one of an inside-out sensor and an outside-in sensor, the inside-out sensor being fixed to the object and makes measurements with respect to a target fixed to an environment, the outside-in sensor being fixed to the environment and makes measurements with respect to a target fixed to the object;

configuring an estimation module of the tracking or navigation system based on an enumeration of the set of sensing elements available to the tracking or navigation system so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, or a combination of inside-out and outside-in sensors depending on the sensors available; and

repeatedly updating an estimated pose of an object based on measurements from the set of sensing elements available to the tracking or navigation system.

69. The machine-accessible medium of claim 68 in which the sensing elements available to the tracking or navigation system include range sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, range sensors, or any combination of the above sensors.

70. The machine-accessible medium of claim 68 in which the sensing elements available to the tracking or navigation system include inertial sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, inertial sensors, or any combination of the above sensors.

71. The machine-accessible medium of claim 68 in which the sensing elements available to the tracking or navigation system include dead reckoning sensors, and configuring the estimation module includes configuring the estimation module so that the estimation module can process measurement information from either inside-out sensors, outside-in sensors, dead reckoning sensors, or any combination of the above sensors.

72. A method comprising:
receiving sensor configuration information indicating a set of sensing elements available to a tracking or navigation system;

configuring a data processing module of the tracking or navigation system based on the sensor configuration information to selectively perform one of

- (a) receiving data from at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the inside-out bearing sensor,
- (b) receiving data from at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the outside-in bearing sensor, and
- (c) receiving data from at least one inside-out bearing sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the outside-in bearing sensor and the inside-out bearing sensor.

73. The method of claim 72 further comprising configuring the data processing module to selectively perform one of

- (d) receiving data from at least one range sensor, and updating an estimated pose of an object based on data received from the range sensor,
- (e) receiving data from at least one range sensor and at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the range sensor and the inside-out bearing sensor,
- (f) receiving data from at least one range sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on data received from the range sensor and the outside-in bearing sensor, and
- (g) receiving data from at least one range sensor, at least one outside-in bearing sensor, and at least one inside-out bearing sensor, and updating an estimated pose of an object based on data received from the range sensor, the inside-out bearing sensor, and the outside-in bearing sensor.

74. An apparatus comprising.

an estimation module to estimate a pose of an object based on measurement data from sensing elements, the estimation module configured to enable selective performance of

- (a) receiving data from at least one inside-out bearing sensor, and updating an estimated pose of an object based on the data received from the inside-out bearing sensor,
- (b) receiving data from at least one outside-in bearing sensor, and updating an estimated pose of an object based on the data received from the outside-in bearing sensor, and
- (c) receiving data from at least one inside-out bearing sensor and at least one outside-in bearing sensor, and updating an estimated pose of an object based on the data received from the outside-in bearing sensor and the inside-out bearing sensor.

75. An apparatus comprising:


an estimation module to estimate a pose of an object based on measurement data from sensing elements, the estimation module configured to enable selective performance of one of:

(a) updating an estimate of the position or orientation of the object relative to an environment,

(b) updating an estimate of the position or orientation, relative to the object, of at least one sensing element fixed to the object, and

(c) updating an estimate of the position or orientation, relative to the environment, of at least one sensing element fixed in the environment.

FIG. 1






Fig. 2




FIG. 3




FIG. 4




FIG. 5

6/12




FIG. 6




FIG. 7

8/12




FIG. 8

9/12




FIG. 9




Fig. 10




FIG. 11




FIG. 12