



- (51) International Patent Classification: US
H04W 4/02 (2018.01) PCT/US2018/040269 (CIP)
- (21) International Application Number: Filed on
PCT/US2018/042346 29 June 2018 (29.06.2018)
- (22) International Filing Date: (71) Applicant: KAARTA, INC. [US/US]; Suite 430, 5001
Baum Boulevard, Pittsburgh, Pennsylvania 15213 (US).
16 July 2018 (16.07.2018) (72) Inventors: ZHANG, Ji; Apt. 915, 5 Bayard Road, Pitts-
burgh, Pennsylvania 15213 (US). DOWLING, Kevin,
Joseph; 2139 Le Gran Drive, Gibsonia, Pennsylvania
15044 (US).
- (25) Filing Language: English (74) Agent: AMBROZIAK, Jeffrey, R. et al.; GTC Law Group
PC & Affiliates, One University Avenue, Suite 302B, West-
wood, MA 02090 (US).
- (26) Publication Language: English
- (30) Priority Data: (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,

62/533,261	17 July 2017 (17.07.2017)	US
PCT/US2018/015403	26 January 2018 (26.01.2018)	US
PCT/US2018/040269	29 June 2018 (29.06.2018)	US
- (63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:

(54) Title: ALIGNING MEASURED SIGNAL DATA WITH SLAM LOCALIZATION DATA AND USES THEREOF

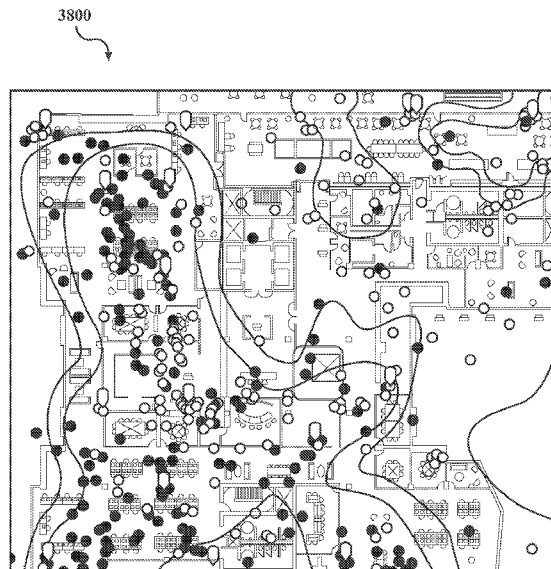


FIG. 38

(57) Abstract: A method includes retrieving a map of a 3D geometry of an environment the map including a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment, receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to at least one of the non-spatial attributes and matching the plurality of received sensor readings to at least one location in the map to produce a determined sensor location.



OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

**ALIGNING MEASURED SIGNAL DATA WITH
SLAM LOCALIZATION DATA AND USES THEREOF**

STATEMENT OF PRIORITY

[0001] This application claims priority to, and is a continuation-in-part of, PCT Application No. PCT/US18/40269 (Atty. Dckt. No. KRTA-0014-WO) entitled " SYSTEMS AND METHODS FOR IMPROVEMENTS IN SCANNING AND MAPPING," filed on June 29, 2018. PCT Application No. PCT/US18/40269 (Atty. Dckt. No. KRTA-0014-WO) claims the benefit of U.S. Provisional Patent Application Serial No. 62/527,341 (Atty. Dckt. No. KRTA-0006-P01), entitled "SYSTEMS AND METHODS FOR IMPROVEMENTS IN SCANNING AND MAPPING," filed on June 30, 2017.

[0002] PCT Application No. PCT/US18/40269 (Atty. Dckt. No. KRTA-0014-WO) claims priority to, and is a continuation-in-part of, PCT Application No. PCT/US2018/015403 (Atty. Dckt. No. KRTA-0010-WO) entitled "LASER SCANNER WITH REAL-TIME, ONLINE EGO-MOTION ESTIMATION," filed on January 26, 2018. PCT Application No. PCT/US2018/015403 (Atty. Dckt. No. KRTA-0010-WO) claims the benefit of U.S. Provisional Patent Application Serial No. 62/451,294 (Atty. Dckt. No. KRTA-0004-P01), entitled "LIDAR AND VISION-BASED EGO-MOTION ESTIMATION AND MAPPING," filed on January 27, 2017.

[0003] PCT Application No. PCT/US2018/015403 (Atty. Dckt. No. KRTA-0010-WO) claims priority to, and is a continuation-in-part of, PCT Application No. PCT/US2017/055938 (Atty. Dckt. No. KRTA-0008-WO) entitled "LASER SCANNER WITH REAL-TIME, ONLINE EGO-MOTION ESTIMATION," filed on October 10, 2017.

[0004] PCT Application No. PCT/US2017/055938 (Atty. Dckt. No. KRTA-0008-WO) claims priority to, and is a continuation-in-part of, PCT Application No. PCT/US2017/021120 (Atty. Dckt. No. KRTA-0005-WO) entitled "LASER SCANNER WITH REAL-TIME, ONLINE EGO-MOTION ESTIMATION," filed on March 7, 2017.

[0005] PCT Application No. PCT/US2018/015403 (Atty. Dckt. No. KRTA-0010-WO) claims priority to PCT Application No. PCT/US2017/021120 (Atty. Dckt. No. KRTA-0005-WO). PCT Application No. PCT/US2017/055938 (Atty. Dckt. No. KRTA-0008-WO) further claims priority to U.S. Provisional No. 62/406,910 (Atty. Dckt. No. KRTA-0002-P02), entitled "LASER SCANNER WITH REAL-TIME, ONLINE EGO-MOTION ESTIMATION," filed on October 11, 2016.

[0006] PCT Application No. PCT/US2017/021120 (Atty. Dckt. No. KRTA-0005-WO) claims the benefit of U.S. Provisional Patent Application Serial No. 62/307,061 (Atty. Dckt. No. KRTA-0001-P01), entitled "LASER SCANNER WITH REAL-TIME, ONLINE EGO-MOTION ESTIMATION," filed on March 11, 2016.

[0007] This application claims the benefit of U.S. Provisional Patent Application Serial No. 62/533,261 (Atty. Dckt. No. KRTA-0007-P01), entitled "ALIGNING MEASURED SIGNAL DATA WITH SLAM LOCALIZATION DATA AND USES THEREOF," filed on July 17, 2017, the disclosure of which is incorporated herein by reference in its entirety and for all purposes.

[0008] This application claims priority to PCT Application No. PCT/US2018/015403 (Atty. Dckt. No. KRTA-0010-WO).

[0009] The disclosures of all patents and applications referred to herein including PCT/US2018/015403, PCT/US2017/055938, PCT/US2017/021120, PCT/US18/40269 and U.S. Provisional Nos. 62/406,910, 62/307,061, 62/533,261, 62/527,341 and 62/451,294 are incorporated herein by reference in their entireties and for all purposes.

BACKGROUND

[0011] A primary means of finding outdoor location are satellite-based radio navigation systems such as the Global Positioning System (GPS) or similar systems such as GLONASS, Galileo, Baidu Navigation Satellite System, and NAVIC. These systems are constellations of satellites that use radio-navigation to provide position on or near the surface of the Earth. However due to attenuation by walls, roofs, and structures determining position is poor or non-existent in indoor spaces. Thus, for indoor navigation other methods have been proposed and implemented using infrastructure, self-estimates such as odometry, and external sources such as passive (barcodes) or active beacons (RF, IR, etc).

[0012] Just as satellite systems such as GPS are used to find one's location and path outdoors, it would be of great practical use to be able to find indoor locations of people, objects, automated vehicles, robots, and the like. Position information can be used in many applications to find a particular location and providing a path to get to/from that location.

SUMMARY

[0013] In accordance with an exemplary and non-limiting embodiment, a method comprises retrieving a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the

environment, receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to at least one of the non-spatial attributes and matching the plurality of received sensor readings to at least one location in the map to produce a determined sensor location.

[0014] In accordance with an exemplary and non-limiting embodiment, a system comprises a device comprising at least one sensor the device adapted to receive a plurality of sensor readings within an environment wherein each of the sensor readings corresponds to at least one non-spatial attribute and a processor adapted to retrieve a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment and further adapted to match the plurality of received sensor readings to at least one location in the map to produce a determined sensor location.

[0015] In accordance with an exemplary and non-limiting embodiment, a method comprises accessing a map comprising a point cloud of an environment comprising a plurality of non-spatial attribute values, receiving, for each of a plurality of non-spatial attributes, a plurality of non-spatial sensor readings within the environment and referencing the received plurality of non-spatial sensor readings to the map resulting in a referenced heat map.

[0016] In accordance with an exemplary and non-limiting embodiment, a system comprises a device adapted to produce a map comprising a point cloud of an environment comprising a plurality of non-spatial attribute values and further adapted to measure, for each of a plurality of non-spatial attributes, a plurality of non-spatial sensor readings and a processor adapted to reference the measured plurality of non-spatial sensor readings to the map resulting in a referenced heat map.

[0017] In accordance with an exemplary and non-limiting embodiment, a method comprises accessing a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters, receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to the non-spatial attribute, creating a heat map of the sensor readings referenced to the map and generating a recommended position of at least one signal emitter based, at least in part, on the heat map.

[0018] In accordance with an exemplary and non-limiting embodiment, a system comprises a device adapted to produce a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters and receive a plurality of sensor readings within the environment wherein each of the sensor readings corresponds to the non-spatial attribute and a processor adapted to create a heat map of the sensor readings referenced to the map and generate a recommended position of at least one signal emitter based, at least in part, on the heat map.

[0019] In accordance with an exemplary and non-limiting embodiment, a method comprises retrieving a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment, receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to at least one of the non-spatial attributes, matching the plurality of received sensor readings to at least one location in the map to produce a determined sensor location, retrieving additional sensor data from the determined sensor location and comparing the additional sensor data to data derived from the map.

[0020] In accordance with an exemplary and non-limiting embodiment, a system comprises a device adapted to produce a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters and receive a plurality of sensor readings within the environment wherein each of the sensor readings corresponds to the non-spatial attribute and a processor adapted to match the plurality of received sensor readings to at least one location in the map to produce a determined sensor location and retrieving additional sensor data from the determined sensor location and compare the additional sensor data to data derived from the map.

BRIEF DESCRIPTION OF THE FIGURES

[0021] FIG. 1 illustrates a block diagram of an embodiment of a mapping system.

[0022] FIG. 2 illustrates an embodiment a block diagram of the three computational modules and their respective feedback features of the mapping system of FIG. 1.

[0023] FIG. 3 illustrates an embodiment of a Kalmann filter model for refining positional information into a map.

[0024] FIG. 4 illustrates an embodiment of a factor graph optimization model for refining positional information into a map.

[0025] FIG. 5 illustrates an embodiment of a visual-inertial odometry subsystem.

[0026] FIG. 6 illustrates an embodiment of a scan matching subsystem.

[0027] FIG. 7A illustrates an embodiment of a large area map having coarse detail resolution.

[0028] FIG. 7B illustrates an embodiment of a small area map having fine detail resolution.

[0029] FIG. 8A illustrates an embodiment of multi-thread scan matching.

[0030] FIG. 8B illustrates an embodiment of single-thread scan matching.

[0031] FIG. 9A illustrates an embodiment of a block diagram of the three computational modules in which feedback data from the visual-inertial odometry unit is suppressed due to data degradation.

[0032] FIG. 9B illustrates an embodiment of the three computational modules in which feedback data from the scan matching unit is suppressed due to data degradation.

[0033] FIG. 10 illustrates an embodiment of the three computational modules in which feedback data from the visual-inertial odometry unit and the scan matching unit are partially suppressed due to data degradation.

[0034] FIG. 11 illustrates an embodiment of estimated trajectories of a mobile mapping device.

[0035] FIG. 12 illustrates bidirectional information flow according to an exemplary and non-limiting embodiment.

[0036] FIGS. 13a and 13b illustrate a dynamically reconfigurable system according to an exemplary and non-limiting embodiment.

[0037] FIG. 14 illustrates priority feedback for IMU bias correction according to an exemplary and non-limiting embodiment.

[0038] FIGS. 15a and 15b illustrate a two-layer voxel representation of a map according to an exemplary and non-limiting embodiment.

[0039] FIGS. 16a and 16b illustrate multi-thread processing of scan matching according to an exemplary and non-limiting embodiment.

[0040] FIGS. 17a and 17b illustrate exemplary and non-limiting embodiments of a SLAM system.

- [0041] FIG. 18 illustrates an exemplary and non-limiting embodiment of a SLAM enclosure.
- [0042] FIGS. 19a, 19b and 19c illustrate exemplary and non-limiting embodiments of a point cloud showing confidence levels.
- [0043] FIG. 20 illustrates an exemplary and non-limiting embodiment of differing confidence level metrics.
- [0044] FIG. 21 illustrates an exemplary and non-limiting embodiment of a SLAM system.
- [0045] FIG. 22 illustrates an exemplary and non-limiting embodiment of timing signals for the SLAM system.
- [0046] FIG. 23 illustrates an exemplary and non-limiting embodiment of timing signals for the SLAM system.
- [0047] FIG. 24 illustrates an exemplary and non-limiting embodiment of SLAM system signal synchronization.
- [0048] FIG. 25 illustrates an exemplary and non-limiting embodiment of air-ground collaborative mapping.
- [0049] FIG. 26 illustrates an exemplary and non-limiting embodiment of a sensor pack.
- [0050] FIG. 27 illustrates an exemplary and non-limiting embodiment of a block diagram of the laser-visual-inertial odometry and mapping software system.
- [0051] FIG. 28 illustrates an exemplary and non-limiting embodiment of a comparison of scans involved in odometry estimation and localization.
- [0052] FIG. 29 illustrates an exemplary and non-limiting embodiment of a comparison of scan matching accuracy in localization.
- [0053] FIG. 30 illustrates an exemplary and non-limiting embodiment of a horizontally orientated sensor test.
- [0054] FIG. 31 illustrates an exemplary and non-limiting embodiment of a vertically orientated sensor test.
- [0055] FIG. 32 illustrates an exemplary and non-limiting embodiment of an accuracy comparison between horizontally orientated and downward tilted sensor tests.
- [0056] FIG. 33 illustrates an exemplary and non-limiting embodiment of an aircraft with a sensor pack.
- [0057] FIG. 34 illustrates an exemplary and non-limiting embodiment of sensor trajectories.

[0058] FIG. 35 illustrates an exemplary and non-limiting embodiment of autonomous flight results.

[0059] FIG. 36 illustrates an exemplary and non-limiting embodiment of an autonomous flight result over a long-run.

[0060] FIGS. 37A and 37B illustrate an exemplary and non-limiting embodiment of a method creating a map that includes both geometry and sensor information and establishing position.

[0061] FIG. 38 is an exemplary and non-limiting embodiment of a heat map style representation of sensed information superimposed on floor plan geometry.

DETAILED DESCRIPTION

[0062] In one general aspect, the present invention is directed to a mobile, computer-based mapping system that estimates changes in position over time (an odometer) and/or generates a three-dimensional map representation, such as a point cloud, of a three-dimensional space. The mapping system may include, without limitation, a plurality of sensors including an inertial measurement unit (IMU), a camera, and/or a 3D laser scanner. It also may comprise a computer system, having at least one processor, in communication with the plurality of sensors, configured to process the outputs from the sensors in order to estimate the change in position of the system over time and/or generate the map representation of the surrounding environment. The mapping system may enable high-frequency, low-latency, on-line, real-time ego-motion estimation, along with dense, accurate 3D map registration. Embodiments of the present disclosure may include a simultaneous location and mapping (SLAM) system. The SLAM system may include a multi-dimensional (e.g., 3D) laser scanning and range measuring system that is GPS-independent and that provides real-time simultaneous location and mapping. The SLAM system may generate and manage data for a very accurate point cloud that results from reflections of laser scanning from objects in an environment. Movements of any of the points in the point cloud are accurately tracked over time, so that the SLAM system can maintain precise understanding of its location and orientation as it travels through an environment, using the points in the point cloud as reference points for the location.

[0063] In one embodiment, the resolution of the position and motion of the mobile mapping system may be sequentially refined in a series of coarse-to-fine updates. In a non-limiting example, discrete computational modules may be used to update the position and

motion of the mobile mapping system from a coarse resolution having a rapid update rate, to a fine resolution having a slower update rate. For example, an IMU device may provide data to a first computational module to predict a motion or position of the mapping system at a high update rate. A visual-inertial odometry system may provide data to a second computational module to improve the motion or position resolution of the mapping system at a lower update rate. Additionally, a laser scanner may provide data to a third computational, scan matching module to further refine the motion estimates and register maps at a still lower update rate. In one non-limiting example, data from a computational module configured to process fine positional and/or motion resolution data may be fed back to computational modules configured to process more coarse positional and/or motion resolution data. In another non-limiting example, the computational modules may incorporate fault tolerance to address issues of sensor degradation by automatically bypassing computational modules associated with sensors sourcing faulty, erroneous, incomplete, or non-existent data. Thus, the mapping system may operate in the presence of highly dynamic motion as well as in dark, texture-less, and structure-less environments.

[0064] In contrast to existing map-generating techniques, which are mostly off-line batch systems, the mapping system disclosed herein can operate in real-time and generate maps while in motion. This capability offers two practical advantages. First, users are not limited to scanners that are fixed on a tripod or other nonstationary mounting. Instead, the mapping system disclosed herein may be associated with a mobile device, thereby increasing the range of the environment that may be mapped in real-time. Second, the real-time feature can give users feedback for currently mapped areas while data are collected. The online generated maps can also assist robots or other devices for autonomous navigation and obstacle avoidance. In some non-limiting embodiments, such navigation capabilities may be incorporated into the mapping system itself. In alternative non-limiting embodiments, the map data may be provided to additional robots having navigation capabilities that may require an externally sourced map.

[0065] There are several potential applications for the sensor, such as 3D modeling, scene mapping, and environment reasoning. The mapping system can provide point cloud maps for other algorithms that take point clouds as input for further processing. Further, the mapping system can work both indoors and outdoors. Such embodiments do not require external lighting and can operate in darkness. Embodiments that have a camera can handle rapid motion, and can colorize laser point clouds with images from the camera, although external lighting may be required. The SLAM system can build and maintain a point cloud in real

time as a user is moving through an environment, such as when walking, biking, driving, flying, and combinations thereof. A map is constructed in real time as the mapper progresses through an environment. The SLAM system can track thousands of features as points. As the mapper moves, the points are tracked to allow estimation of motion. Thus, the SLAM system operates in real time and without dependence on external location technologies, such as GPS. In embodiments, a plurality (in most cases, a very large number) of features of an environment, such as objects, are used as points for triangulation, and the system performs and updates many location and orientation calculations in real time to maintain an accurate, current estimate of position and orientation as the SLAM system moves through an environment. In embodiments, relative motion of features within the environment can be used to differentiate fixed features (such as walls, doors, windows, furniture, fixtures and the like) from moving features (such as people, vehicles, and other moving items), so that the fixed features can be used for position and orientation calculations. Underwater SLAM systems may use blue-green lasers to reduce attenuation.

[0066] The mapping system design follows an observation: drift in egomotion estimation has a lower frequency than a module's own frequency. The three computational modules are therefore arranged in decreasing order of frequency. High-frequency modules are specialized to handle aggressive motion, while low-frequency modules cancel drift for the previous modules. The sequential processing also favors computation: modules in the front take less computation and execute at high frequencies, giving sufficient time to modules in the back for thorough processing. The mapping system is therefore able to achieve a high level of accuracy while running online in real-time.

[0067] Further, the system may be configured to handle sensor degradation. If the camera is non-functional (for example, due to darkness, dramatic lighting changes, or texture-less environments) or if the laser is non-functional (for example due to structure-less environments) the corresponding module may be bypassed and the rest of the system may be staggered to function reliably. Specifically, in some exemplary embodiments, the proposed pipeline automatically determines a degraded subspace in the problem state space, and solves the problem partially in the well-conditioned subspace. Consequently, the final solution is formed by combination of the "healthy" parts from each module. As a result, the resulting combination of modules used to produce an output is neither simply a linear or non-linear combination of module outputs. In some exemplary embodiments, the output reflect a bypass of one or more entire modules in combination with a linear or non-linear combination of the remaining functioning modules. The system was tested through a large number of

experiments and results show that it can produce high accuracy over several kilometers of navigation and robustness with respect to environmental degradation and aggressive motion.

[0068] The modularized mapping system, disclosed below, is configured to process data from range, vision, and inertial sensors for motion estimation and mapping by using a multi-layer optimization structure. The modularized mapping system may achieve high accuracy, robustness, and low drift by incorporating features which may include:

- an ability to dynamically reconfigure the computational modules;
- an ability to fully or partially bypass failure modes in the computational modules, and combine the data from the remaining modules in a manner to handle sensor and/or sensor data degradation, thereby addressing environmentally induced data degradation and the aggressive motion of the mobile mapping system; and
- an ability to integrate the computational module cooperatively to provide real-time performance.

[0069] Disclosed herein is a mapping system for online ego-motion estimation with data from a 3D laser, a camera, and an IMU. The estimated motion further registers laser points to build a map of the traversed environment. In many real-world applications, ego-motion estimation and mapping must be conducted in real-time. In an autonomous navigation system, the map may be crucial for motion planning and obstacle avoidance, while the motion estimation is important for vehicle control and maneuver.

[0070] FIG 1 depicts a simplified block diagram of a mapping system 100 according to one embodiment of the present invention. Although specific components are disclosed below, such components are presented solely as examples and are not limiting with respect to other, equivalent, or similar components. The illustrated system includes an IMU system **102** such as an Xsens® MTi-30 IMU, a camera system **104** such as an IDS® UI-1220SE monochrome camera, and a laser scanner **106** such as a Velodyne PUCK™ VLP-16 laser scanner. The IMU **102** may provide inertial motion data derived from one or more of an x-y-z accelerometer, a roll-pitch-yaw gyroscope, and a magnetometer, and provide inertial data at a first frequency. In some non-limiting examples, the first frequency may be about 200 Hz. The camera system **104** may have a resolution of about 752x480 pixels, a 76° horizontal field of view (FOV), and a frame capture rate at a second frequency. In some non-limiting examples, the frame capture rate may operate at a second frequency of about 50Hz. The laser scanner **106** may have a 360° horizontal FOV, a 30° vertical FOV, and receive 0.3 million points/second at a third frequency representing the laser spinning rate. In some non-limiting

examples, the third frequency may be about 5Hz. As depicted in FIG. 1, the laser scanner **106** may be connected to a motor **108** incorporating an encoder **109** to measure a motor rotation angle. In one non-limiting example, the laser motor encoder **109** may operate with a resolution of about 0.25^0 .

[0071] The IMU **102**, camera **104**, laser scanner **106**, and laser scanner motor encoder **109** may be in data communication with a computer system **110**, which may be any computing device, having one or more processors **134** and associated memory **120**, **160**, having sufficient processing power and memory for performing the desired odometry and/or mapping. For example, a laptop computer with 2.6GHz i7quad-core processor (2 threads on each core and 8 threads overall) and an integrated GPU memory could be used. In addition, the computer system may have one or more types of primary or dynamic memory **120** such as RAM, and one or more types of secondary or storage memory **160** such as a hard disk or a flash ROM. Although specific computational modules (IMU module **122**, visual-inertial odometry module **126**, and laser scanning module **132**) are disclosed above, it should be recognized that such modules are merely exemplary modules having the functions as described above, and are not limiting. Similarly, the type of computing device **110** disclosed above is merely an example of a type of computing device that may be used with such sensors and for the purposes as disclosed herein, and is in no way limiting.

[0072] As illustrated in FIG.1, the mapping system **100** incorporates a computational model comprising individual computational modules that sequentially recover motion in a coarse-to-fine manner (see also FIG. 2). Starting with motion prediction from an IMU **102** (IMU prediction module **122**), a visual-inertial tightly coupled method (visual-inertial odometry module **126**) estimates motion and registers laser points locally. Then, a scan matching method (scan matching refinement module **132**) further refines the estimated motion. The scan matching refinement module **132** also registers point cloud data **165** to build a map (voxel map **134**). The map also may be used by the mapping system as part of an optional navigation system **136**. It may be recognized that the navigation system **136** may be included as a computational module within the onboard computer system, the primary memory, or may comprise a separate system entirely.

[0073] It may be recognized that each computational module may process data from one of each of the sensor systems. Thus, the IMU prediction module **122** produces a coarse map from data derived from the IMU system **102**, the visual-inertial odometry module **126** processes the more refined data from the camera system **104**, and the scan matching refinement module **132** processes the most fine-grained resolution data from the laser scanner

106 and the motor encoder **109**. In addition, each of the finer-grained resolution modules further process data presented from a coarser-grained module. The visual-inertial odometry module **126** refines mapping data received from and calculated by the IMU prediction module **122**. Similarly, the scan matching refinement module **132**, further processes data presented by the visual inertial odometry module **126**. As disclosed above, each of the sensor systems acquires data at a different rate. In one non-limiting example, the IMU **102** may update its data acquisition at a rate of about 200 Hz, the camera **104** may update its data acquisition at a rate of about 50 Hz, and the laser scanner **106** may update its data acquisition at a rate of about 5 Hz. These rates are non-limiting and may, for example, reflect the data acquisition rates of the respective sensors. It may be recognized that coarse-grained data may be acquired at a faster rate than more fine-grained data, and the coarse-grained data may also be processed at a faster rate than the fine-grained data. Although specific frequency values for the data acquisition and processing by the various computation modules are disclosed above, neither the absolute frequencies nor their relative frequencies are limiting.

[0074] The mapping and/or navigational data may also be considered to comprise coarse level data and fine level data. Thus, in the primary memory (dynamic memory **120**), coarse positional data may be stored in a voxel map **134** that may be accessible by any of the computational modules **122**, **126**, **132**. File detailed mapping data, as point cloud data **165** that may be produced by the scan matching refinement module **132**, may be stored via the processor **150** in a secondary memory **160**, such as a hard drive, flash drive, or other more permanent memory.

[0075] Not only are coarse-grained data used by the computational modules for more fine-grained computations, but both the visual-inertial odometry module **126** and the scan matching refinement module **132** (fine-grade positional information and mapping) can feed back their more refined mapping data to the IMU prediction module **122** via respective feedback paths **128** and **138** as a basis for updating the IMU position prediction. In this manner, coarse positional and mapping data may be sequentially refined in resolution, and the refined resolution data serve as feed-back references for the more coarse resolution computations.

[0076] FIG. 2 depicts a block diagram of the three computational modules along with their respective data paths. The IMU prediction module **122** may receive IMU positional data **223** from the IMU (**102**, FIG. 1). The visual-inertial odometry module **126** may receive the model data from the IMU prediction module **122** as well as visual data from one or more individually tracked visual features **227a**, **227b** from the camera (**104**, FIG. 1). The laser

scanner (**106**, FIG. 1) may produce data related to laser determined landmarks **233a**, **233b**, which may be supplied to the scan matching refinement module **132** in addition to the positional data supplied by the visual-inertial odometry module **126**. The positional estimation model from the visual-inertial odometry module **126** may be fed back **128** to refine the positional model calculated by the IMU prediction module **122**. Similarly, the refined map data from the scan matching refinement module **132** may be fed back **138** to provide additional correction to the positional model calculated by the IMU prediction module **122**.

[0077] As depicted in FIG. 2, and as disclosed above, the modularized mapping system may sequentially recover and refine motion related data in a coarse-to-fine manner. Additionally, the data processing of each module may be determined by the data acquisition and processing rate of each of the devices sourcing the data to the modules. Starting with motion prediction from an IMU, a visual-inertial tightly coupled method estimates motion and registers laser points locally. Then, a scan matching method further refines the estimated motion. The scan matching refinement module may also register point clouds to build a map. As a result, the mapping system is time optimized to process each refinement phase as data become available.

[0078] Figure 3 illustrates a standard Kalman filter model based on data derived from the same sensor types as depicted in Figure 1. As illustrated in Figure 3, the Kalman filter model updates positional and/or mapping data upon receipt of any data from any of the sensors regardless of the resolution capabilities of the data. Thus, for example, the positional information may be updated using the visual-inertial odometry data at any time such data become available regardless of the state of the positional information estimate based on the IMU data. The Kalman filter model therefore does not take advantage of the relative resolution of each type of measurement. FIG. 3 depicts a block diagram of a standard Kalman filter based method for optimizing positional data. The Kalman filter updates a positional model **322a** - **322n** sequentially as data are presented. Thus, starting with an initial positional prediction model **322a**, the Kalman filter may predict **324a** the subsequent positional model **322b**, which may be refined based on the receive IMU mechanization data **323**. The positional prediction model may be updated **322b** in response to the IMU mechanization data **323**. in a prediction step **324a** followed by update steps seeded with individual visual features or laser landmarks.

[0079] FIG. 4 depicts positional optimization based on a factor-graph method. In this method, a pose of a mobile mapping system at a first time **410** may be updated upon receipt

of data to a pose at a second time **420**. A factor-graph optimization model combines constraints from all sensors during each refinement calculation. Thus, IMU data **323**, feature data **327a**, **327b**, and similar from the camera, and laser landmark data **333a**, **333b**, and similar, are all used for each update step. It may be appreciated that such a method increases the computational complexity for each positional refinement step due to the large amount of data required. Further, since the sensors may provide data at independent rates that may differ by orders of magnitude, the entire refinement step is time bound to the data acquisition time for the slowest sensor. As a result, such a model may not be suitable for fast real-time mapping. The modularized system depicted in FIGS. 1 and 2 sequentially recovers motion in a coarse-to-fine manner. In this manner, the degree of motion refinement is determined by the availability of each type of data.

[0080] Assumptions, Coordinates, and Problem

Assumptions and Coordinate Systems

[0081] As depicted above in FIG. 1, a sensor system of a mobile mapping system may include a laser **106**, a camera **104**, and an IMU **102**. The camera may be modeled as a pinhole camera model for which the intrinsic parameters are known. The extrinsic parameters among all of the three sensors may be calibrated. The relative pose between the camera and the laser and the relative pose between the laser and the IMU may be determined according to methods known in the art. A single coordinate system may be used for the camera and the laser. In one non-limiting example, the camera coordinate system may be used, and all laser points may be projected into the camera coordinate system in pre-processing. In one non-limiting example, the IMU coordinate system may be parallel to the camera coordinate system and thus the IMU measurements may be rotationally corrected upon acquisition. The coordinate systems may be defined as follows:

- the camera coordinate system $\{C\}$ may originate at the camera optical center, in which the x -axis points to the left, the y -axis points upward, and the z -axis points forward coinciding with the camera principal axis;
- the IMU coordinate system $\{I\}$ may originate at the IMU measurement center, in which the x -, y -, and z - axes are parallel to $\{C\}$ and pointing in the same directions; and
- the world coordinate system $\{W\}$ may be the coordinate system coinciding with $\{C\}$ at the starting pose.

[0082] In accordance with some exemplary embodiments the landmark positions are not necessarily optimized. As a result, there remain six unknowns in the state space thus keeping

computation intensity low. The disclosed method involves laser range measurements to provide precise depth information to features, warranting motion estimation accuracy while further optimizing the features' depth in a bundle. One need only optimize some portion of these measurements as further optimization may result in diminishing returns in certain circumstances.

[0083] In accordance with exemplary and non-limiting embodiments, calibration of the described system may be based, at least in part, on the mechanical geometry of the system. Specifically, the LIDAR may be calibrated relative to the motor shaft using mechanical measurements from the CAD model of the system for geometric relationships between the lidar and the motor shaft. Such calibration as is obtained with reference to the CAD model has been shown to provide high accuracy and drift without the need to perform additional calibration.

MAP Estimation Problem

[0084] A state estimation problem can be formulated as a maximum a posterior (MAP) estimation problem. We may define $\chi = \{x_i\}, i \in \{1; 2; \dots, m\}$, as the set of system states $U = \{u_i\}, i \in \{1; 2; \dots, m\}$, as the set of control inputs, and $Z = \{z_k\}, k \in \{1; 2; \dots, n\}$, as the set of landmark measurements. Given the proposed system, Z may be composed of both visual features and laser landmarks. The joint probability of the system is defined as follows,

$$\text{Eq. 1} \quad P(\chi|U, Z) \propto P(x_0) \prod_{i=1}^m P(x_i|x_{i-1}, u_i) \prod_{k=1}^n P(z_k|x_{i_k}),$$

[0085] where $P(x_0)$ is a prior of the first system state, $P(x_i|x_{i-1}, u_i)$ represents the motion model, and $P(z_k|x_{i_k})$ represents the landmark measurement model. For each problem formulated as (1), there is a corresponding Bayesian belief network representation of the problem. The MAP estimation is to maximize Eq. 1. Under the assumption of zero-mean Gaussian noise, the problem is equivalent to a least-square problem,

$$\text{Eq. 2} \quad \chi^* = \arg \min_{\chi} \sum_{i=1}^m \|r_{x_i}\|^2 + \sum_{k=1}^n \|r_{z_k}\|^2.$$

[0086] Here, r_{x_i} and r_{z_k} are residual errors associated with the motion model and the landmark measurement model, respectively.

[0087] The standard way of solving Eq. 2 is to combine all sensor data, for example visual features, laser landmarks, and IMU measurements, into a large factor-graph optimization problem. The proposed data processing pipeline, instead, formulates multiple

small optimization problems and solves the problems in a coarse-to-fine manner. The optimization problem may be restated as:

Problem: Given data from a laser, a camera, and an IMU, formulate and solve problems as (2) to determine poses of $\{C\}$ with respect to $\{W\}$, then use the estimated poses to register laser points and build a map of the traversed environment in $\{W\}$.

IMU Prediction Subsystem

IMU Mechanization

[0088] This subsection describes the IMU prediction subsystem. Since the system considers $\{C\}$ as the fundamental sensor coordinate system, the IMU may also be characterized with respect to $\{C\}$. As disclosed above in the sub-section entitled Assumptions and Coordinate Systems, $\{I\}$ and $\{C\}$ are parallel coordinate systems. $\omega(t)$ and $a(t)$ may be two 3×1 vectors indicating the angular rates and accelerations, respectively, of $\{C\}$ at time t . The corresponding biases may be denoted as $\mathbf{b}_\omega(t)$ and $\mathbf{b}_a(t)$ and $\mathbf{n}_\omega(t)$ and $\mathbf{n}_a(t)$ be the corresponding noises. The vector, bias, and noise terms are defined in $\{C\}$. Additionally, \mathbf{g} may be denoted as the constant gravity vector in $\{W\}$. The IMU measurement terms are:

$$\text{Eq. 3} \quad \hat{\omega}(t) = \omega(t) + \mathbf{b}_\omega(t) + \mathbf{n}_\omega(t),$$

$$\text{Eq. 4} \quad \hat{\mathbf{a}}(t) = \mathbf{a}(t) - {}^C_W \mathbf{R}(t) \mathbf{g} - {}^I_C \mathbf{t} \|\omega(t)\|^2 + \mathbf{b}_a(t) + \mathbf{n}_a(t),$$

where ${}^C_W \mathbf{R}(t)$ is the rotation matrix from $\{W\}$ to $\{C\}$, and ${}^I_C \mathbf{t}$ is the translation vector between $\{C\}$ and $\{I\}$.

[0089] It is noted that the term ${}^I_C \mathbf{t} \|\omega(t)\|^2$ represents the centrifugal force due to the fact that the rotation center (origin of $\{C\}$) is different from the origin of $\{I\}$. Some examples of visual-inertial navigation methods model the motion in $\{I\}$ to eliminate this centrifugal force term. In the computational method disclosed herein, in which visual features both with and without depth information are used, converting features without depth from $\{C\}$ to $\{I\}$ is not straight forward (see below). As a result, the system disclosed herein models all of the motion in $\{C\}$ instead. Practically, the camera and the IMU are mounted close to each other to maximally reduce effect of the term.

[0090] The IMU biases may be slowly changing variables. Consequently, the most recently updated biases are used for motion integration. First, Eq. 3 is integrated over time.

Then, the resulting orientation is used with Eq. 4 for integration over time twice to obtain translation from the acceleration data.

Bias Correction

[0091] The IMU bias correction can be made by feedback from either the camera or the laser (see **128, 138**, respectively, in FIGS. 1 and 2). Each feedback term contains the estimated incremental motion over a short amount of time. The biases may be modeled to be constant during the incremental motion. Starting with Eq. 3, $\mathbf{b}_\omega(t)$ may be calculated by comparing the estimated orientation with IMU integration. The updated $\mathbf{b}_\omega(t)$ is used in one more round of integration to re-compute the translation, which is compared with the estimated translation to calculate $\mathbf{b}_a(t)$.

[0092] To reduce the effect of high-frequency noises, a sliding window is employed keeping a known number of biases. Non-limiting examples of the number of biases used in the sliding window may include 200 to 1000 biases with a recommended number of 400 biases based on a 200Hz IMU rate. A non-limiting example of the number of biases in the sliding window with an IMU rate of 100Hz is 100 to 500 with a typical value of 200 biases. The averaged biases from the sliding window are used. In this implementation, the length of the sliding window functions as a parameter for determining an update rate of the biases. Although alternative methods to model the biases are known in the art, the disclosed implementation is used in order to keep the IMU processing module as a separate and distinct module. The sliding window method may also allow for dynamic reconfiguration of the system. In this manner, the IMU can be coupled with either the camera, the laser, or both camera and laser as required. For example, if the camera is non-functional, the IMU biases may be corrected only by the laser instead.

[0093] To reduce the effect of high-frequency noises, a sliding window may be employed keeping a certain number of biases. In such instances, the averaged biases from the sliding window may be used. In such an implementation, the length of the sliding window functions as a parameter determining an update rate of the biases. In some instances, the biases may be modeled as random walks and the biases updated through a process of optimization. However, this non-standard implementation is preferred to keep IMU processing in a separate module. The implementation favors dynamic reconfiguration of the system, i.e. the IMU may be coupled with either the camera or the laser. If the camera is non-functional, the IMU biases may be corrected by the laser instead. As described inter-module communication in the

sequential modularized system is utilized to fix the IMU biases. This communication enables IMU biases to be corrected.

[0094] In accordance with an exemplary and non-limiting embodiment, IMU bias correction may be accomplished by utilizing feedback from either the camera or the laser. Each of the camera and the laser contains the estimated incremental motion over a short amount of time. When calculating the biases, the methods and systems described herein model the biases to be constant during the incremental motion. Still starting with Eq. 3, by comparing the estimated orientation with IMU integration, the methods and systems described herein can calculate $b_{\omega}(t)$. The updated $b_{\omega}(t)$ is used in one more round of integration to recompute the translation, which is compared with the estimated translation to calculate $b_d(t)$.

[0095] In some embodiments, IMU output comprises an angular rate having relatively constant errors over time. The resulting IMU bias is related to the fact that the IMU will always have some difference from ground truth. This bias can change over time. It is relatively constant and not high frequency. The sliding window described above is a specified period of time during which the IMU data is evaluated.

[0096] With reference to Fig. 5, there is provided a system diagram of the visual-inertial odometry subsystem. The method couples vision with an IMU. Both vision and the IMU provide constraints to an optimization problem that estimates incremental motion. At the same time, the method associates depth information to visual features. If a feature is located in an area where laser range measurements are available, depth may be obtained from laser points. Otherwise, depth may be calculated from triangulation using the previously estimated motion sequence. As the last option, the method may also use features without any depth by formulating constraints in a different way. This is true for those features which may not necessarily have laser range coverage or cannot be triangulated due to the fact that they are not necessarily tracked long enough or located in the direction of camera motion. These three alternatives may be used alone or in combination for building the registered point cloud. In some embodiments information is discarded. Eigenvalues and eigenvectors may be computed and used to identify and specify degeneracy in a point cloud. If there is degeneracy in a specific direction in the state space then the solution in that direction in that state space can be discarded.

Visual-inertial Odometry Subsystem

[0097] A block system diagram of the visual-inertial odometry subsystem is depicted in Fig. 5. An optimization module 510 uses pose constraints 512 from the IMU prediction module 520 along with camera constraints 515 based on optical feature data having or lacking depth information for motion estimation 550. A depthmap registration module 545 may include depthmap registration and depth association of the tracked camera features 530 with depth information obtained from the laser points 540. The depthmap registration module 545 may also incorporate motion estimation 550 obtained from a previous calculation. The method tightly couples vision with an IMU. Each provides constraints 512, 515, respectively, to an optimization module 510 that estimates incremental motion 550. At the same time, the method associates depth information to visual features as part of the depthmap registration module 545. If a feature is located in an area where laser range measurements are available, depth is obtained from laser points. Otherwise, depth is calculated from triangulation using the previously estimated motion sequence. As the last option, the method can also use features without any depth by formulating constraints in a different way. This is true for those features which do not have laser range coverage or cannot be triangulated because they are not tracked long enough or located in the direction of camera motion.

Camera Constraints

[0098] The visual-inertial odometry is a key-frame based method. A new key-frame is determined 535 if more than a certain number of features lose tracking or the image overlap is below a certain ratio. Here, right superscript l , $l \in \mathbf{Z}^+$ may indicate the last key-frame, and c , $c \in \mathbf{Z}^+$ and $c > k$, may indicate the current frame. As disclosed above, the method combines features with and without depth. A feature that is associated with depth at key-frame l , may be denoted as $\mathbf{X}_l = [x_l, y_l, z_l]^T$ in $\{C_l\}$. Correspondingly, a feature without depth is denoted as $\bar{\mathbf{X}}_l = [\bar{x}_l, \bar{y}_l, l]^T$ using normalized coordinates instead. Note that \mathbf{X}_l , $\bar{\mathbf{X}}_l$, x_l , and \bar{x}_l are different from $\boldsymbol{\chi}$ and \mathbf{x} in Eq.1 which represent the system state. Features at key-frames may be associated with depth for two reasons: 1) depth association takes some amount of processing, and computing depth association only at key-frames may reduce computation intensity; and 2) the depthmap may not be available at frame c and thus laser points may not be registered since registration depends on an established depthmap. A normalized feature in $\{C_c\}$ may be denoted as $\bar{\mathbf{X}}_c = [\bar{x}_c, \bar{y}_c, 1]^T$.

[0099] Let \mathbf{R}_l^c and \mathbf{t}_l^c be the 3×3 rotation matrix and 3×1 translation vector between frames l and c , where $\mathbf{R}_l^c \in \mathbf{SO}(3)$ and $\mathbf{t}_l^c \in \mathbb{R}^3$, \mathbf{R}_l^c and \mathbf{T}_l^c form an $\mathbf{SE}(3)$ transformation. The motion function between frames l and c may be written as

$$\text{Eq. 5} \quad \mathbf{X}_c = \mathbf{R}_l^c \mathbf{X}_l + \mathbf{t}_l^c .$$

[0100] \mathbf{X}_c has an unknown depth. Let d_c be the depth, where $\mathbf{X}_c = d_c \bar{\mathbf{X}}_c$. Substituting \mathbf{X}_c with $d_c \bar{\mathbf{X}}_c$ and combining the 1st and 2nd rows with the 3rd row in Eq. 5 to eliminate d_c , results in

$$\text{Eq. 6} \quad (\mathbf{R}(1) - \bar{x}_c \mathbf{R}(3)) X_l + t_1 - \bar{x}_c t(3) = 0,$$

$$\text{Eq. 7} \quad (\mathbf{R}(2) - \bar{y}_c \mathbf{R}(3)) X_l + t_2 - \bar{y}_c t(3) = 0,$$

[0101] $\mathbf{R}(h)$ and $t(h)$, $h \in \{1, 2, 3\}$, are the h -th rows of \mathbf{R}_l^c and \mathbf{t}_l^c . In the case that depth is unavailable to a feature, let d_l be the unknown depth at key-frame l . Substituting \mathbf{X}_l and \mathbf{X}_c with $d_l \bar{\mathbf{X}}_l$ and $d_c \bar{\mathbf{X}}_c$, respectively, and combining all three rows in Eq. 5 to eliminate d_k and d_c , results in another constraint,

$$\text{Eq. 8} \quad [\bar{y}_c t(3) - t(2)], -\bar{x}_c t(3) + t(1), \bar{x}_c t(2) - \bar{y}_c t(1)] \mathbf{R}_l^c \bar{\mathbf{X}}_l = 0.$$

Motion Estimation

[0102] The motion estimation process 510 is required to solve an optimization problem combining three sets of constraints: 1) from features with known depth as in Eqs. 6-7; 2) from features with unknown depth as in Eq. 8; and 3) from the IMU prediction 520. \mathbf{T}_a^b may be defined as a 4×4 transformation matrix between frames a and b ,

$$\text{Eq. 9} \quad \mathbf{T}_a^b = \begin{bmatrix} \mathbf{R}_a^b & \mathbf{t}_a^b \\ \mathbf{0}^T & 1 \end{bmatrix},$$

[0103] where \mathbf{R}_a^b and \mathbf{t}_a^b are the corresponding rotation matrix and translation vector. Further, let θ_a^b be a 3×1 vector corresponding to \mathbf{R}_a^b through an exponential map, where $\theta_a^b \in \mathfrak{so}(3)$. The normalized term $\theta/\|\theta\|$ represents direction of the rotation and $\|\theta\|$ is the rotation angle. Each \mathbf{T}_a^b corresponds to a set of θ_a^b and \mathbf{t}_a^b containing 6-DOF motion of the camera.

[0104] The solved motion transform between frames l and $c-1$, namely \mathbf{T}_l^{c-1} may be used to formulate the IMU pose constraints. A predicted transform between the last two frames $c-1$ and c , denoted as $\hat{\mathbf{T}}_{c-1}^c$ may be obtained from IMU mechanization. The predicted transform at frame c is calculated as,

$$\text{Eq. 10} \quad \hat{\mathbf{T}}_l^c = \hat{\mathbf{T}}_{c-1}^c \hat{\mathbf{T}}_l^{c-1} .$$

[00105] Let $\hat{\theta}_i^c$ and $\hat{\mathbf{t}}_i^c$ be the 6-DOF motion corresponding to $\hat{\mathbf{T}}_i^c$. It may be understood that the IMU predicted translation, $\hat{\mathbf{t}}_i^c$, is dependent on the orientation. As an example, the orientation may determine a projection of the gravity vector through rotation matrix ${}^c_w\mathbf{R}(t)$ in Eq. 4, and hence the accelerations that are integrated. $\hat{\mathbf{t}}_i^c$ may be formulated as a function of θ_i^c , and may be rewritten as $\hat{\mathbf{t}}_i^c(\theta_i^c)$. It may be understood that the 200Hz pose provided by the IMU prediction module **122** (FIGS. 1 and 2) as well as the 50Hz pose provided by the visual-inertial odometry module **126** (FIGS. 1 and 2) are both pose functions. Calculating $\hat{\mathbf{t}}_i^c(\theta_i^c)$ may begin at frame c and the accelerations may be integrated inversely with respect to time. Let θ_i^c be the rotation vector corresponding to \mathbf{R}_i^c in Eq. 5, and θ_i^c and \mathbf{t}_i^c are the motion to be solved. The constraints may be expressed as,

$$\text{Eq. 11} \quad \Sigma_i^c [(\hat{\theta}_i^c - \theta_i^c)^T, (\hat{\mathbf{t}}_i^c(\theta_i^c) - \mathbf{t}_i^c)^T]^T = \mathbf{0},$$

in which Σ_i^c is a relative covariance matrix scaling the pose constraints appropriately with respect to the camera constraints.

[00106] In embodiments of the visual-inertial odometry subsystem, the pose constraints fulfill the motion model and the camera constraints fulfill the landmark measurement model in Eq. 2. The optimization problem may be solved by using the Newton gradient-descent method adapted to a robust fitting framework for outlier feature removal. In this problem, the state space contains θ_i^c and \mathbf{t}_i^c . Thus, a full-scale MAP estimation is not performed, but is used only to solve a marginalized problem. The landmark positions are not optimized, and thus only six unknowns in the state space are used, thereby keeping computation intensity low. The method thus involves laser range measurements to provide precise depth information to features, warranting motion estimation accuracy. As a result, further optimization of the features' depth through a bundle adjustment may not be necessary.

Depth Association

[00107] The depthmap registration module **545** registers laser points on a depthmap using previously estimated motion. Laser points **540** within the camera field of view are kept for a certain amount of time. The depthmap is down-sampled to keep a constant density and stored in a 2D KD-tree for fast indexing. In the KD-tree, all laser points are projected onto a unit sphere around the camera center. A point is represented by its two angular coordinates. When associating depth to features, features may be projected onto the sphere. The three closest laser points are found on the sphere for each feature. Then, their validity may be by calculating distances among the three points in Cartesian space. If a distance is larger than a

threshold, the chance that the points are from different objects, e.g. a wall and an object in front of the wall, is high and the validity check fails. Finally, the depth is interpolated from the three points assuming a local planar patch in Cartesian space.

[00108] Those features without laser range coverage, if they are tracked over a certain distance and not located in the direction of camera motion, may be triangulated using the image sequences where the features are tracked. In such a procedure, the depth may be updated at each frame based on a Bayesian probabilistic mode.

Scan Matching Subsystem

[00109] This subsystem further refines motion estimates from the previous module by laser scan matching. FIG. 6 depicts a block diagram of the scan matching subsystem. The subsystem receives laser points **540** in a local point cloud and registers them **620** using provided odometry estimation **550**. Then, geometric features are detected **640** from the point cloud and matched to the map. The scan matching minimizes the feature-to-map distances, similar to many methods known in the art. However, the odometry estimation **550** also provides pose constraints **612** in the optimization **610**. The optimization comprises processing pose constraints with feature correspondences **615** that are found and further processed with laser constraints **617** to produce a device pose **650**. This pose **650** is processed through a map registration process **655** that facilitates finding the feature correspondences **615**. The implementation uses voxel representation of the map. Further, it can dynamically configure to run on one to multiple CPU threads in parallel.

Laser Constraints

[00110] When receiving laser scans, the method first registers points from a scan **620** into a common coordinate system. $m, m \in \mathbb{Z}^+$ may be used to indicate the scan number. It is understood that the camera coordinate system may be used for both the camera and the laser. Scan m may be associated with the camera coordinate system at the beginning of the scan, denoted as $\{C_m\}$. To locally register **620** the laser points **540**, the odometry estimation **550** from the visual-inertial odometry may be taken as key-points, and the IMU measurements may be used to interpolate in between the key-points.

[00111] Let P_m be the locally registered point cloud from scan m . Two sets of geometric features from P_m may be extracted: one on sharp edges, namely edge points and denoted as ε_m , and the other on local planar surfaces, namely planar points and denoted as \mathcal{H}_m . This is through computation of curvature in the local scans. Points whose neighbor points are already selected are avoided such as points on boundaries of occluded regions and points whose local

surfaces are close to be parallel to laser beams. These points are likely to contain large noises or change positions over time as the sensor moves.

[00112] The geometric features are then matched to the current map built. Let Q_{m-1} be the map point cloud after processing the last scan, Q_{m-1} is defined in $\{W\}$. The points in Q_{m-1} are separated into two sets containing edge points and planar points, respectively. Voxels may be used to store the map truncated at a certain distance around the sensor. For each voxel, two 3D KD-trees may be constructed, one for edge points and the other for planar points. Using KD-trees for individual voxels accelerates point searching since given a query point, a specific KD-tree associated with a single voxel needs to be searched (see below).

[00113] When matching scans, \mathcal{E}_m and \mathcal{H}_m into $\{W\}$ are first projected using the best guess of motion available, then for each point in \mathcal{E}_m and \mathcal{H}_m , a cluster of closest points are found from the corresponding set on the map. To verify geometric distributions of the point clusters, the associated eigenvalues and eigenvectors may be examined. Specifically, one large and two small eigenvalues indicate an edge line segment, and two large and one small eigenvalues indicate a local planar patch. If the matching is valid, an equation is formulated regarding the distance from a point to the corresponding point cluster,

$$\text{Eq. 12} \quad f(\mathbf{X}_m, \theta_m, \mathbf{t}_m) = d,$$

where \mathbf{X}_m is a point in \mathcal{E}_m or \mathcal{H}_m , $\theta_m, \theta_m \in \mathbf{so}(3)$, and $\mathbf{t}_m, \mathbf{t}_m \in \mathbb{R}^3$, indicate the 6-DOF pose of $\{C_m\}$ in $\{W\}$.

Motion Estimation

[00114] The scan matching is formulated into an optimization problem **610** minimizing the overall distances described by Eq. 12. The optimization also involves pose constraints **612** from prior motion. Let \mathbf{T}_{m-1} be the 4×4 transformation matrix regarding the pose of $\{C_{m-1}\}$ in $\{W\}$, \mathbf{T}_{m-1} is generated by processing the last scan. Let $\hat{\mathbf{T}}_{m-1}^m$ be the pose transform from $\{C_{m-1}\}$ to $\{C_m\}$, as provided by the odometry estimation. Similar to Eq. 10, the predicted pose transform of $\{C_m\}$ in $\{W\}$ is,

$$\text{Eq. 13} \quad \hat{\mathbf{T}}_m = \hat{\mathbf{T}}_{m-1}^m \mathbf{T}_{m-1}.$$

[00115] Let $\hat{\theta}_m$ and $\hat{\mathbf{t}}_m$ be the 6-DOF pose corresponding to $\hat{\mathbf{T}}_m$, and let Σ_m be a relative covariance matrix. The constraints are,

$$\text{Eq. 14 } \Sigma_m \left[(\hat{\theta}_m - \theta_m)^T, (\hat{\mathbf{t}}_m - \mathbf{t}_m)^T \right]^T = \mathbf{0}.$$

[00116] Eq. 14 refers to the case that the prior motion is from the visual-inertial odometry, assuming the camera is functional. Otherwise, the constraints are from the IMU prediction. $\hat{\theta}'_m$ and $\hat{\mathbf{t}}'_m(\theta_m)$ may be used to denote the same terms by IMU mechanization. $\hat{\mathbf{t}}'_m(\theta_m)$ is a function of θ_m because integration of accelerations is dependent on the orientation (same with $\hat{t}_i^c(\theta_i^c)$ in Eq. 11). The IMU pose constraints are,

$$\text{Eq. 15 } \Sigma'_m \left[(\hat{\theta}'_m - \theta_m)^T, (\hat{\mathbf{t}}'_m(\theta_m) - \mathbf{t}_m)^T \right]^T = \mathbf{0},$$

where Σ'_m is the corresponding relative covariance matrix. In the optimization problem, Eqs. 14 and 15 are linearly combined into one set of constraints. The linear combination is determined by working mode of the visual-inertial odometry. The optimization problem refines θ_m and \mathbf{t}_m , which is solved by the Newton gradient-descent method adapted to a robust fitting framework.

Map in Voxels

[00117] The points on the map are kept in voxels. A 2-level voxel implementation as illustrated in FIGS. 7A and 7B. M_{m-1} denotes the set of voxels **702**, **704** on the first level map **700** after processing the last scan. Voxels **704** surrounding the sensor **706** form a subset M_{m-1} , denoted as S_{m-1} . Given a 6-DOF sensor pose, $\hat{\theta}_m$ and $\hat{\mathbf{t}}_m$, there is a corresponding S_{m-1} which moves with the sensor on the map. When the sensor approaches the boundary of the map, voxels on the opposite side **725** of the boundary are moved over to extend the map boundary **730**. Points in moved voxels are cleared resulting in truncation of the map.

[00118] As illustrated in FIG. 7B, each voxel $j, j \in S_{m-1}$ of the second level map **750** is formed by a set of voxels that are a magnitude smaller, denoted as S_{m-1}^j than those of the first level map **700**. Before matching scans, points in \mathcal{E}_m and \mathcal{H}_m are projected onto the map using the best guess of motion, and fill them into $\{S_{m-1}^j\}, j \in S_{m-1}$. Voxels **708** occupied by points from \mathcal{E}_m and \mathcal{H}_m are extracted to form Q_{m-1} and stored in 3D KD-trees for scan matching. Voxels **710** are those not occupied by points from \mathcal{E}_m or \mathcal{H}_m . Upon completion of scan matching, the scan is merged into the voxels **708** with the map. After that, the map points are downsized to maintain a constant density. It may be recognized that each voxel of the first level map **700** corresponds to a volume of space that is larger than a sub-voxel of the second level map **750**. Thus, each voxel of the first level map **700** comprises a plurality of

sub-voxels in the second level map **750** and can be mapped onto the plurality of sub-voxels in the second level map **750**.

[00119] As noted above with respect to FIGS. 7A and 7B, two levels of voxels (first level map **700** and second level map **750**) are used to store map information. Voxels corresponding to M_{m-1} are used to maintain the first level map **700** and voxels corresponding to $\{S_{m-1}^j\}$, $j \in S_{m-1}$ in the second level map **750** are used to retrieve the map around the sensor for scan matching. The map is truncated only when the sensor approaches the map boundary. Thus, if the sensor navigates inside the map, no truncation is needed. Another consideration is that two KD-trees are used for each individual voxel in S_{m-1} - one for edge points and the other for planar points. As noted above, such a data structure may accelerate point searching. In this manner, searching among multiple KD-trees is avoided as opposed to using two KD-trees for each individual voxel in $\{S_{m-1}^j\}$, $j \in S_{m-1}$. The later requires more resources for KD-tree building and maintenance.

[00120] Table 1 compares CPU processing time using different voxel and KD-tree configurations. The time is averaged from multiple datasets collected from different types of environments covering confined and open, structured and vegetated areas. We see that using only one level of voxels, M_{m-1} , results in about twice of processing time for KD-tree building and querying. This is because the second level of voxels, $\{S_{m-1}^j\}$, $j \in S_{m-1}$, help retrieve the map precisely. Without these voxel, more points are contained in Q_{m-1} and built into the KD-trees. Also, by using KD-trees for each voxel, processing time is reduced slightly in comparison to using KD-trees for all voxels in M_{m-1} .

[00121] **Table 1.** Comparison of average CPU processing time on KD-tree operation

Task	1-level voxels		2-level voxels	
	KD-trees for all voxels	KD-trees for each voxel	KD-trees for all voxels	KD-trees for each voxel
Build (time per KD-tree)	54ms	47ms	24ms	21ms
Query (time per point)	4.2ns	4.1ns	2.4ns	2.3ns

Parallel Processing

[00122] The scan matching involves building KD-trees and repetitively finding feature correspondences. The process is time-consuming and takes major computation in the system.

While one CPU thread cannot guarantee the desired update frequency, a multi-thread implementation may address issues of complex processing. FIG. 8A illustrates the case where two matcher programs **812**, **815** run in parallel. Upon receiving of a scan, a manager program **810** arranges it to match with the latest map available. In one example, composed of a clustered environment with multiple structures and multiple visual features, matching is slow and may not complete before arrival of the next scan. The two matchers **812** and **815** are called alternatively. In one matcher **812**, P_m , **813a** P_{m-2} , **813b** and additional P_{m-k} (for k =an even integer) **813n**, are matched with Q_{m-2} **813a**, Q_{m-4} **813a**, and additional Q_{m-k} (for k =an even integer) **813n**, respectively. Similarly, in a second matcher **815**, P_{m+1} , **816a** P_{m-1} , **816b** and additional P_{m-k} (for k = an odd integer) **816n**, are matched with Q_{m-1} **816a**, Q_{m-3} **816a**, and additional Q_{m-k} (for k =an odd integer) **816n**, respectively. The use of this interleaving process may provide twice the amount of time for processing. In an alternative example, composed of a clean environment with few structures or visual features, computation is light. In such an example (FIG. 8B), only a single matcher **820** may be called. Because interleaving is not required P_m, P_{m-1}, \dots , are sequentially matched with Q_{m-1}, Q_{m-2}, \dots , respectively (see **827a**, **827b**, **827n**). The implementation may be configured to use a maximum of four threads, although typically only two threads may be needed.

Transform Integration

[00123] The final motion estimation is integration of outputs from the three modules depicted in FIG. 2. The 5Hz scan matching output produces the most accurate map, while the 50Hz visual-inertial odometry output and the 200Hz IMU prediction are integrated for high-frequency motion estimates.

On Robustness

[00124] The robustness of the system is determined by its ability to handle sensor degradation. The IMU is always assumed to be reliable functioning as the backbone in the system. The camera is sensitive to dramatic lighting changes and may also fail in a dark/texture-less environment or when significant motion blur is present (thereby causing a loss of visual features tracking). The laser cannot handle structure-less environments, for example a scene that is dominant by a single plane. Alternatively, laser data degradation can be caused by sparsity of the data due to aggressive motion. Such aggressive motion comprises highly dynamic motion. As used herein, "highly dynamic motion" refers to substantially abrupt rotational or linear displacement of the system or continuous rotational or translational motion having a substantially large magnitude. The disclosed self-motion

determining system may operate in the presence of highly dynamic motion as well as in dark, texture-less, and structure-less environments. In some exemplary embodiments, the system may operate while experiencing angular rates of rotation as high as 360 deg per second. In other embodiments, the system may operate at linear velocities up to and including at 110kph. In addition, these motions can be coupled angular and linear motions.

[00125] Both the visual-inertial odometry and the scan matching modules formulate and solve optimization problems according to EQ. 2. When a failure happens, it corresponds to a degraded optimization problem, i.e. constraints in some directions of the problem are ill-conditioned and noise dominates in determining the solution. In one non-limiting method, eigenvalues, denoted as $\lambda_1, \lambda_2, \dots, \lambda_6$, and eigenvectors, denoted as $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_6$, associated with the problem may be computed. Six eigenvalues/eigenvectors are present because the state space of the sensor contains 6-DOF (6 degrees of freedom). Without losing generality, $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_6$ may be sorted in decreasing order. Each eigenvalue describes how well the solution is conditioned in the direction of its corresponding eigenvector. By comparing the eigenvalues to a threshold, well-conditioned directions may be separated from degraded directions in the state space. Let $h, h = 0; 1, \dots, 6$, be the number of well-conditioned directions. Two matrices may be defined as:

$$\text{Eq. 16 } \mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_6]^T, \bar{\mathbf{V}} = [\mathbf{v}_1, \dots, \mathbf{v}_h, 0, \dots, 0]^T.$$

[00126] When solving an optimization problem, the nonlinear iteration may start with an initial guess. With the sequential pipeline depicted in FIG. 2, the IMU prediction provides the initial guess for the visual-inertial odometry, whose output is taken as the initial guess for the scan matching. For the additional two modules (visual-inertial odometry and scan matching modules), let \mathbf{x} be a solution and $\Delta\mathbf{x}$ be an update of \mathbf{x} in a nonlinear iteration, in which $\Delta\mathbf{x}$ is calculated by solving the linearized system equations. During the optimization process, instead of updating \mathbf{x} in all directions, \mathbf{x} may be updated only in well-conditioned directions, keeping the initial guess in degraded directions instead,

$$\text{Eq. 17 } \mathbf{x} \leftarrow \mathbf{x} + \mathbf{V}^{-1}\bar{\mathbf{V}} \Delta\mathbf{x}.$$

[00127] In Eq. 17, the system solves for motion in a coarse-to-fine order, starting with the IMU prediction, the additional two modules further solving/refining the motion as much as possible. If the problem is well-conditioned, the refinement may include all 6-DOF. Otherwise, if the problem is only partially well-conditioned, the refinement may include 0 to

5-DOF. If the problem is completely degraded, $\bar{\mathbf{V}}$ becomes a zero matrix and the previous module's output is kept.

[00128] Returning to the pose constraints described in Eqs. 14 and 15, it may be understood that the two equations are linearly combined in the scan matching problem. As defined in Eq. 16, \mathbf{V}_V and $\bar{\mathbf{V}}_V$ denote the matrices containing eigenvectors from the visual-inertial odometry module, $\bar{\mathbf{V}}_V$ represents well-conditioned directions in the subsystem, and $\mathbf{V}_V - \bar{\mathbf{V}}_V$ represents degraded directions. The combined constraints are,

$$\Sigma'_m \mathbf{V}_V^{-1} \bar{\mathbf{V}}_V \left[(\hat{\theta}_m - \theta_m)^T, (\hat{\mathbf{t}}_m - \mathbf{t}_m)^T \right]^T + \Sigma'_m \mathbf{V}_V^{-1}$$

$$\text{Eq. 18 } (\mathbf{V}_V - \bar{\mathbf{V}}_V) \left[(\hat{\theta}'_m - \theta_m)^T, (\hat{\mathbf{t}}'_m(\theta_m) - \mathbf{t}_m)^T \right]^T = \mathbf{0}.$$

[00129] In a normal case where the camera is functional, $\bar{\mathbf{V}}_V = \mathbf{V}_V$ and Eq.18 is composed of pose constraints from the visual-inertial odometry as in Eq. 14. However, if the camera data are completely degraded, $\bar{\mathbf{V}}_V$ is a zero matrix and Eq. 18 is composed of pose constraints from the IMU prediction according to Eq. 15.

Case Study of Camera Degradation

[00130] As depicted in FIG. 9A, if visual features are insufficiently available for the visual-inertial odometry, the IMU prediction **122** bypasses the visual-inertial odometry module **126** fully or partially **924** - denoted by the dotted line -- depending on the number of well-conditioned directions in the visual-inertial odometry problem. The scan matching module **132** may then locally register laser points for the scan matching. The bypassing IMU prediction is subject to drift. The laser feedback **138** compensates for the camera feedback **128** correcting velocity drift and biases of the IMU, only in directions where the camera feedback **128** is unavailable. Thus, the camera feedback has a higher priority, due to the higher frequency making it more suitable when the camera data are not degraded. When sufficient visual features are found, the laser feedback is not used.

Case Study of Laser Degradation

[00131] As shown in FIG. 9B, if environmental structures are insufficient for the scan matching **132** to refine motion estimates, the visual-inertial odometry module **126** output fully or partially bypasses the scan matching module to register laser points on the map **930** as noted by the dotted line. If well-conditioned directions exist in the scan matching problem, the laser feedback contains refined motion estimates in those directions. Otherwise, the laser feedback becomes empty **138**.

Case Study of Camera and Laser Degradation

[00132] In a more complex example, both the camera and the laser are degraded at least to some extent. FIG. 10 depicts such an example. A vertical bar with six rows represents a 6-DOF pose where each row is a DOF (degree of freedom), corresponding to an eigenvector in EQ. 16. In this example, the visual-inertial odometry and the scan matching each updates a 3-DOF of motion, leaving the motion unchanged in the other 3-DOF. The IMU prediction **1022a-f** may include initial IMU predicted values **1002**. The visual-inertial odometry updates **1004** some 3-DOF (**1026c**, **1026e**, **1026f**) resulting in a refined prediction **1026a-1026f**. The scan matching updates **1006** some 3-DOF (**1032b**, **1032d**, **1032f**) resulting in a further refined prediction **1032a-1032f**. The camera feedback **128** contains camera updates **1028a-1028f** and the laser feedback **138** contains laser updates **1038a-1038f**, respectively. In reference to FIG. 10, cells having no shading (**1028a**, **1028b**, **1028d**, **1038a**, **1038c**, **1038e**) do not contain any updating information from the respective modules. The total update **1080a-1080f** to the IMU prediction modules is a combination of the updates **1028a-1028f** from the camera feedback **128** and the updates **1038a-1038f** from the laser feedback **138**. In one or more of the degrees of freedom in which feedback is available from both the camera (for example **1028f**) and the laser (for example **1038f**), the camera updates (for example **1028f**) may have priority over the laser updates (for example **1038f**).

[00133] In practice, however, the visual-inertial odometry module and the scan matching module may execute at different frequencies and each may have its own degraded directions. IMU messages may be used to interpolate between the poses from the scan matching output. In this manner, an incremental motion that is time aligned with the visual-inertial odometry output may be created. Let θ_{c-1}^c and \mathbf{t}_{c-1}^c be the 6-DOF motion estimated by the visual-inertial odometry between frames $c-1$ and c , where $\theta_{c-1}^c \in \mathfrak{so}(3)$ and $\mathbf{t}_{c-1}^c \in \mathbb{R}^3$. Let θ'_{c-1}^c and \mathbf{t}'_{c-1}^c be the corresponding terms estimated by the scan matching after time interpolation. \mathbf{V}_V and $\bar{\mathbf{V}}_V$ may be the matrices defined in Eq. 16 containing eigenvectors from the visual-inertial odometry module, in which $\bar{\mathbf{V}}_V$ represents well-conditioned directions, and $\mathbf{V}_V - \bar{\mathbf{V}}_V$ represents degraded directions. Let \mathbf{V}_S and $\bar{\mathbf{V}}_S$ be the same matrices from the scan matching module. The following equation calculates the combined feedback, f_c ,

$$\text{Eq. 19 } f_c = f_V + \mathbf{V}_V^{-1}(\mathbf{V}_V - \bar{\mathbf{V}}_V) f_s,$$

where f_V and f_s represent the camera and the laser feedback,

$$\text{Eq. 20 } f_V = \mathbf{V}_V^{-1}(\bar{\mathbf{V}}_V)[(\theta_{c-1}^c)^T, (\mathbf{t}_{c-1}^c)^T]^T,$$

$$\text{Eq. 21} \quad f_s = \mathbf{V}_S^{-1} \bar{\mathbf{V}}_S [(\theta'_{c-1})^T, (\mathbf{t}'_{c-1})^T]^T.$$

[00134] note that f_C only contains solved motion in a subspace of the state space. The motion from the IMU prediction, namely $\hat{\theta}_{c-1}^c$ and $\hat{\mathbf{t}}_{c-1}^c$, may be projected to the null space of f_C ,

$$\text{Eq. 22} \quad f_I = \mathbf{V}_V^{-1} (\mathbf{V}_V - \bar{\mathbf{V}}_V) \mathbf{V}_S^{-1} (\mathbf{V}_S - \bar{\mathbf{V}}_S) \left[(\hat{\theta}_{c-1}^c)^T, (\hat{\mathbf{t}}_{c-1}^c)^T \right]^T$$

[00135] $\tilde{\theta}_{c-1}^c(\mathbf{b}_\omega(t))$ and $\tilde{\mathbf{t}}_{c-1}^c(\mathbf{b}_\omega(t), \mathbf{b}_a(t))$ may be used to denote the IMU predicted motion formulated as functions of $\mathbf{b}_\omega(t)$ and $\mathbf{b}_a(t)$ through integration of Eqs. 3 and 4. The orientation $\tilde{\theta}_{c-1}^c(\mathbf{b}_\omega(t))$ is only relevant to $\mathbf{b}_\omega(t)$, but the translation $\tilde{\mathbf{t}}_{c-1}^c(\mathbf{b}_\omega(t), \mathbf{b}_a(t))$ is dependent on both $\mathbf{b}_\omega(t)$ and $\mathbf{b}_a(t)$. The biases can be calculated by solving the following equation,

$$\text{Eq. 23} \quad \left[(\tilde{\theta}_{c-1}^c(\mathbf{b}_\omega(t)))^T, (\tilde{\mathbf{t}}_{c-1}^c(\mathbf{b}_\omega(t), \mathbf{b}_a(t)))^T \right]^T = f_C + f_I.$$

[00136] When the system functions normally, f_C spans the state space, and $\mathbf{V}_V - \bar{\mathbf{V}}_V$ and $\mathbf{V}_S - \bar{\mathbf{V}}_S$ in Eq. 22 are zero matrices. Correspondingly, $\mathbf{b}_\omega(t)$ and $\mathbf{b}_a(t)$ are calculated from f_C . In a degraded case, the IMU predicted motion, $\hat{\theta}_{c-1}^c$ and $\hat{\mathbf{t}}_{c-1}^c$, is used in directions where the motion is unsolvable (e.g. white row **1080a** of the combined feedback in FIG. 10). The result is that the previously calculated biases are kept in these directions.

Experiments

Tests with Scanners

[00137] The odometry and mapping software system was validated on two sensor suites. In a first sensor suite, a Velodyne LIDAR™ HDL-32E laser scanner is attached to a UI-1220SE monochrome camera and an Xsens® MTi-30 IMU. The laser scanner has 360° horizontal FOV, 40° vertical FOV, and receives 0.7 million points/second at 5Hz spinning rate. The camera is configured at the resolution of 752 x 480 pixels, 76° horizontal FOV, and 50Hz frame rate. The IMU frequency is set at 200Hz. In a second sensor suite, a Velodyne LIDAR™ VLP-16 laser scanner is attached to the same camera and IMU. This laser scanner has 360° horizontal FOV, 30° vertical FOV, and receives 0.3 million points/second at 5Hz spinning rate. Each sensor suite is attached to a vehicle for data collection, which are driven on streets and in off-road terrains, respectively.

[00138] For both sensor suites, a maximum of 300 Harris corners were tracked. To evenly distribute the visual features, an image is separated into 5 x 6 identical sub-regions, each sub-

region providing up to 10 features. When a feature loses tracking, a new feature is generated to maintain the feature number in each sub region.

[00139] The software runs on a laptop computer with a 2.6GHz i7 quad-core processor (2 threads on each core and 8 threads overall) and an integrated GPU, in a Linux[®] system running Robot Operating System (ROS). Two versions of the software were implemented with visual feature tracking running on GPU and CPU, respectively. The processing time is shown in Table 2. The time used by the visual-inertial odometry (**126** in FIG. 2) does not vary much with respect to the environment or sensor configuration. For the GPU version, it consumes about 25% of a CPU thread executing at 50Hz. For the CPU version, it takes about 75% of a thread. The sensor first suite results in slightly more processing time than the second sensor suite. This is because the scanner receives more points and the program needs more time to maintain the depthmap and associate depth to the visual features.

[00140] The scan matching (**132** in FIG. 2) consumes more processing time which also varies with respect to the environment and sensor configuration. With the first sensor suite, the scan matching takes about 75% of a thread executing at 5Hz if operated in structured environments. In vegetated environments, however, more points are registered on the map and the program typically consumes about 135% of a thread. With the second sensor suite, the scanner receives fewer points. The scan matching module **132** uses about 50-95% of a thread depending on the environment. The time used by the IMU prediction (**132** in FIG. 2) is negligible compared to the other two modules.

Accuracy Tests

[00141] Tests were conducted to evaluate accuracy of the proposed system. In these tests, the first sensor suite was used. The sensors were mounted on an off-road vehicle driving around a university campus. After 2.7km of driving within 16 minutes, a campus map was built. The average speed over the test was 2.8m/s.

[00142] **Table 2.** Average CPU processing time using the first and second sensor suites

Environment	Sensor suite	Visual-inertial odometry (time per image frame)		Scan Matching (time per laser scan)
		GPU Tracking	CPU Tracking	
	First suite	4.8ms	14.3ms	148ms
Structured	Second suite	4.2ms	12.9ms	103ms
	First suite	5.5ms	15.2ms	267ms
Vegetated	Second suite	5.1ms	14.7ms	191ms

[00143] To evaluate motion estimation drift over the test, the estimated trajectory and registered laser points were aligned on a satellite image. Here, laser points on the ground are manually removed. It was determined, by matching the trajectory with streets on the satellite image, that an upper bound of the horizontal error was < 1:0m. It was also determined, by comparing buildings on the same floor, that the vertical error was < 2:0m. This gives an overall relative position drift at the end to be <0:09% of the distance traveled. It may be understood that precision cannot be guaranteed for the measurements, hence only an upper bound of the positional drift was calculated.

[00144] Further, a more comprehensive test was conducted having the same sensors mounted on a passenger vehicle. The passenger vehicle was driven on structured roads for 9.3km of travel. The path traverses vegetated environments, bridges, hilly terrains, and streets with heavy traffic, and finally returns to the starting position. The elevation changes over 70m along the path. Except waiting for traffic lights, the vehicle speed is between 9-18m/s during the test. It was determined that a building found at both the start and the end of the path was registered into two. The two registrations occur because of motion estimation drift over the length of the path. Thus, the first registration corresponds to the vehicle at the start of the test and the second registration corresponds to the vehicle at the end of the test. The gap was measured to be < 20m, resulting in a relative position error at the end of < 0:22% of the distance traveled.

[00145] Each module in the system contributes to the overall accuracy. FIG. 11 depicts estimated trajectories in an accuracy test. A first trajectory plot **1102** of the trajectory of a mobile sensor generated by the visual-inertial odometry system uses the IMU module **122** and the visual-inertial odometry module **126** (see FIG. 2). The configuration used in the first trajectory plot **1102** is similar to that depicted in FIG. 9B. A second trajectory plot **1104** is based on directly forwarding the IMU prediction from the IMU module **122** to the scan matching module, **132** (see FIG. 2) bypassing the visual-inertial odometry. This configuration

is similar to that depicted in FIG. 9A. A third trajectory plot **1108** of the complete pipeline is based on the combination of the IMU module **122**, the visual inertial odometry module **126**, and the scan matching module **132** (see FIG. 2) has the least amount of drift. The position errors of the first two configurations, trajectory plot **1102** and **1104**, are about four and two times larger, respectively.

[00146] The first trajectory plot **1102** and the second trajectory plot **1104** can be viewed as the expected system performance when encountering individual sensor degradation. If scan matching is degraded (see FIG. 9B), the system reduces to a mode indicated by the first trajectory plot **1102**. If vision is degraded, (see FIG. 9A), the system reduces to a mode indicated by the second trajectory plot **1104**. If none of the sensors is degraded, (see FIG 2) the system incorporates all of the optimization functions resulting in the trajectory plot **1108**. In another example, the system may take the IMU prediction as the initial guess and but run at the laser frequency (5Hz). The system produces a fourth trajectory plot **1106**. The resulting accuracy is only little better in comparison to the second trajectory plot **1104** which uses the IMU directly coupled with the laser, passing the visual-inertial odometry. The result indicates that functionality of the camera is not sufficiently explored if solving the problem with all constraints stacked together.

[00147] Another accuracy test of the system included running mobile sensor at the original 1x speed and an accelerated 2x speed. When running at 2x speed, every other data frame for all three sensors is omitted, resulting in much more aggressive motion through the test. The results are listed in Table 3. At each speed, the three configurations were evaluated. At 2x speed, the accuracy of the visual-inertial odometry and the IMU + scan matching configurations reduce significantly, by 0.54% and 0.38% of the distance traveled in comparison to the accuracy at 1x speed. However, the complete pipeline reduces accuracy very little, only by 0.04%. The results indicate that the camera and the laser compensate for each other keeping the overall accuracy. This is especially true when the motion is aggressive.

[00148] **Table 3.** Relative position errors as percentages of the distance traveled

[00149] (Errors at 1x speed correspond to the trajectories in Fig. 11)

Configuration	1x speed	2x speed
Visual-inertial odometry	0.93%	1.47%
IMU + scan matching	0.51%	0.89%
Complete pipeline	0.22%	0.26%

[00150] With reference to Fig. 12, there is illustrated an exemplary and non-limiting embodiment of bidirectional information flow. As illustrated, three modules comprising an IMU prediction module, a visual-inertial odometry module and a scan-matching refinement module solve the problem step by step from coarse to fine. Data processing flow is from left to right passing the three modules respectively, while feedback flow is from right to left to correct the biases of the IMU.

[00151] With reference to Figs. 13a and 13b, there is illustrated an exemplary and non-limiting embodiment of a dynamically reconfigurable system. As illustrated in Fig. 13a, if visual features are insufficient for the visual-inertial odometry, the IMU prediction (partially) bypasses the visual-inertial odometry module to register laser points locally. On the other hand, if, as illustrated in Fig. 13b, environmental structures are insufficient for the scan matching, the visual-inertial odometry output (partially) bypasses the scan matching refinement module to register laser points on the map.

[00152] With reference to Fig. 14, there is illustrated an exemplary and non-limiting embodiment of priority feedback for IMU bias correction. As illustrated, a vertical bar represents a 6-DOF pose and each row is a DOF. In a degraded case, starting with the IMU prediction on the left where all six rows designated are "IMU", the visual-inertial odometry updates in 3-DOF where the rows become designated "camera", then the scan matching updates in another 3-DOF where the rows turn designated "laser". The camera and the laser feedback is combined as the vertical bar on the left. The camera feedback has a higher priority – "laser" rows from the laser feedback are only filled in if "camera" rows from the camera feedback are not present.

[00153] With reference to Figs. 15a and 15b, there is illustrated an exemplary and non-limiting embodiment of two-layer voxel representation of a map. There is illustrated voxels on the map M_{m-1} (all voxels in Fig. 15a), and voxels surrounding the sensor S_{m-1} (dot filled voxels). S_{m-1} is a subset of M_{m-1} . If the sensor approaches the boundary of the map, voxels on the opposite side of the boundary (bottom row) are moved over to extend the map boundary. Points in moved voxels are cleared and the map is truncated. As illustrated in Fig. 15b each voxel j , $j \in S_{m-1}$ (a dot filled voxel in Fig. 15a) is formed by a set of voxels S_{m-1}^j that are a magnitude smaller (all voxels in (Fig. 15b) $\in S_{m-1}^j$). Before scan matching, the laser scan may be projected onto the map using the best guess of motion. Voxels in $\{S_{m-1}^j\}$, $j \in S_{m-1}$ occupied by points from the scan are labeled in cross-hatch. Then, map points in cross-hatched voxels are extracted and stored in 3D KD-trees for scan matching.

[00154] With reference to Fig. 16, there is illustrated an exemplary and non-limiting embodiment of multi-thread processing of scan matching. As illustrated, a manager program calls multiple matcher programs running on separate CPU threads and matches scans to the latest map available. Fig. 16a shows a two-thread case. Scans P_m, P_{m-1}, \dots , are matched with map Q_m, Q_{m-1}, \dots , on each matcher, giving twice amount of time for processing. In comparison, Fig. 16b shows a one-thread case, where P_m, P_{m-1}, \dots , are matched with Q_m, Q_{m-1}, \dots . The implementation is dynamically configurable using up to four threads.

[00155] In embodiments, a real time SLAM system may be used in combination with a real time navigation system. In embodiments, the SLAM system may be used in combination with an obstacle detection system, such as a LIDAR- or RADAR-based obstacle detection system, a vision-based obstacle detection system, a thermal-based system, or the like. This may include detecting live obstacles, such as people, pets, or the like, such as by motion detection, thermal detection, electrical or magnetic field detection, or other mechanisms..

[00156] In embodiments, the point cloud that is established by scanning the features of an environment may be displayed, such as on a screen forming a part of the SLAM, to show a mapping of a space, which may include mapping of near field features, such as objects providing nearby reflections to the SLAM system, as well as far field features, such as items that can be scanned through spaces between or apertures in the near field features. For example, items in an adjacent hallway may be scanned through a window or door as the mapper moves through the interior of a room, because at different points in the interior of the room different outside elements can be scanned through such spaces or apertures. The resulting point cloud may then comprise comprehensive mapping data of the immediate near field environment and partial mapping of far field elements that are outside the environment. Thus, the SLAM system may include mapping of a space through a “picket fence” effect by identification of far-field pieces through spaces or apertures (i.e., gaps in the fence) in the near field. The far field data may be used to help the system orient the SLAM as the mapper moves from space to space, such as maintaining consistent estimation of location as the mapper moves from a comprehensively mapped space (where orientation and position are well known due to the density of the point cloud) to a sparsely mapped space (such as a new room). As the user moves from the near field to a far field location, the relative density or sparseness of the point cloud can be used by the SLAM system to guide the mapper via, for example, a user interface forming a part of the SLAM, such as directing the mapper to the parts of the far field that could not be seen through the apertures from another space.

[00157] In embodiments, the point cloud map from a SLAM system can be combined with mapping from other inputs such as cameras, sensors, and the like. For example, in a flight or spacecraft example, an airplane, drone, or other airborne mobile platform may already be equipped with other distance measuring and geo-location equipment that can be used as reference data for the SLAM system (such as linking the point cloud resulting from a scan to a GPS-referenced location) or that can take reference data from a scan, such as for displaying additional scan data as an overlay on the output from the other system. For example, conventional camera output can be shown with point cloud data as an overlay, or vice versa.

[00158] In embodiments, the SLAM system can provide a point cloud that includes data indicating the reflective intensity of the return signal from each feature. This reflective intensity can be used to help determine the efficacy of the signal for the system, to determine how features relate to each other, to determine surface IR reflectivity, and the like. In embodiments, the reflective intensity can be used as a basis for manipulating the display of the point cloud in a map. For example, the SLAM system can introduce (automatically, or under user control) some degree of color contrast to highlight the reflectivity of the signal for a given feature, material, structure, or the like. In addition, the system can be married with other systems for augmenting color and reflectance information. In embodiments, one or more of the points in the point cloud may be displayed with a color corresponding to a parameter of the acquired data, such as an intensity parameter, a density parameter, a time parameter and a geospatial location parameter. Colorization of the point cloud may help users understand and analyze elements or features of the environment in which the SLAM system is operating and/or elements of features of the process of acquisition of the point cloud itself. For example, a density parameter, indicating the number of points acquired in a geospatial area, may be used to determine a color that represents areas where many points of data are acquired and another color where data is sparse, perhaps suggesting the presence of artifacts, rather than “real” data. Color may also indicate time, such as progressing through a series of colors as the scan is undertaken, resulting in clear indication of the path by which the SLAM scan was performed. Colorization may also be undertaken for display purposes, such as to provide differentiation among different features (such as items of furniture in a space, as compared to walls), to provide aesthetic effects, to highlight areas of interest (such as highlighting a relevant piece of equipment for attention of a viewer of a scan), and many others.

[00159] In embodiments, the SLAM system can identify “shadows” (areas where the point cloud has relatively few data points from the scan) and can (such as through a user interface) highlight areas that need additional scanning. For example, such areas may blink or be rendered

in a particular color in a visual interface of a SLAM system that displays the point cloud until the shadowed area is sufficiently "painted," or covered, by laser scanning. Such an interface may include any indicator (visual, text-based, voice-based, or the like) to the user that highlights areas in the field that have not yet been scanned, and any such indicator may be used to get the attention of the user either directly or through an external device (such as a mobile phone of the user). In accordance with other exemplary and non-limiting embodiments, the system may make reference to external data of data stored on the SLAM, such as previously constructed point clouds, maps, and the like, for comparison with current scan to identify unscanned regions.

[00160] In embodiments, the methods and systems disclosed herein include a SLAM system that provides real-time positioning output at the point of work, without requiring processing or calculation by external systems in order to determine accurate position and orientation information or to generate a map that consists of point cloud data showing features of an environment based on the reflected signals from a laser scan. In embodiments, the methods and systems disclosed herein may also include a SLAM system that provides real time positioning information without requiring post-processing of the data collected from a laser scan.

[00161] In embodiments, a SLAM system may be integrated with various external systems, such as vehicle navigation systems (such as for unmanned aerial vehicles, drones, mobile robots, unmanned underwater vehicles, self-driving vehicles, semi-automatic vehicles, and many others). In embodiments, the SLAM system may be used to allow a vehicle to navigate within its environments, without reliance on external systems like GPS.

[00162] In embodiments, a SLAM system may determine a level of confidence as to its current estimation of position, orientation, or the like. A level of confidence may be based on the density of points that are available in a scan, the orthogonality of points available in a scan, environmental geometries or other factors, or a combination thereof. The level of confidence may be ascribed to position and orientation estimates at each point along the route of a scan, so that segments of the scan can be referenced as low-confidence segments, high-confidence segments, or the like. Low-confidence segments can be highlighted for additional scanning, for use of other techniques (such as making adjustments based on external data), or the like. For example, where a scan is undertaken in a closed loop (where the end point of the scan is the same as the starting point, at a known origin location), any discrepancy between the calculated end location and the starting location may be resolved by preferentially adjusting location estimates for certain segments of the scan to restore consistency of the start- and end-

locations. Location and position information in low-confidence segments may be preferentially adjusted as compared to high-confidence segments. Thus, the SLAM system may use confidence-based error correction for closed loop scans.

[00163] In an exemplary and non-limiting embodiment of this confidence-based adjustment, a derivation of the incremental smoothing and mapping (iSAM) algorithm originally developed by Michael Kaess and Frank Dellaert at Georgia Tech ("iSAM: Incremental Smoothing and Mapping" by M. Kaess, A. Ranganathan, and F. Dellaert, IEEE Trans. on Robotics, TRO, vol. 24, no. 6, Dec. 2008, pp. 1365-1378, [PDF](#)) may be used. This algorithm processes map data in "segments" and iteratively refines the relative position of those segments to optimize the residual errors of matches between the segments. This enables closing loops by adjusting all data within the closed loop. More segmentation points allows the algorithm to move the data more significantly while fewer segmentation points generates more rigidity.

[00164] If one selects the points for segmentation based on a matching confidence metric, one may utilize this fact to make the map flexible in regions with low confidence and rigid in areas with high confidence so that the loop closure processing does not distribute local errors through sections of the map that are accurate. This can be further enhanced by weighting the segmentation points to assign "flexibility" to each point and distribute error based on this factor.

[00165] In embodiments, confidence measures with respect to areas or segments of a point cloud may be used to guide a user to undertake additional scanning, such as to provide an improved SLAM scan. In embodiments, a confidence measure can be based on a combination of density of points, orthogonality of points and the like, which can be used to guide the user to enable a better scan. It is noted that scan attributes, such as density of points and orthogonality of points, may be determined in real time as the scan progresses. Likewise, the system may sense geometries of the scanning environment that are likely to result in low confidence measures. For example, long hallways with smooth walls may not present any irregularities to differentiate one scan segment from the next. In such instances, the system may assign a lower confidence measure to scan data acquired in such environments. The system can use various inputs such as LIDAR, camera, and perhaps other sensors to determine diminishing confidence and guide the user through a scan with instructions (such as "slow down," "turn left" or the like). In other embodiments, the system may display areas of lower than desired confidence to a user, such as via a user interface, while providing assistance in allowing the user to further scan the area, volume or region of low confidence.

[00166] In embodiments, a SLAM output may be fused with other content, such as outputs from cameras, outputs from other mapping technologies, and the like. In embodiments, a SLAM scan may be conducted along with capture of an audio track, such as via a companion application (optionally a mobile application) that captures time-coded audio notes that correspond to a scan. In embodiments, the SLAM system provides time-coding of data collection during scanning, so that the mapping system can pinpoint when and where the scan took place, including when and where the mapper took audio and/or notes. In embodiments, the time coding can be used to locate the notes in the area of the map where they are relevant, such as by inserting data into a map or scan that can be accessed by a user, such as by clicking on an indicator on the map that audio is available. In embodiments, other media formats may be captured and synchronized with a scan, such as photography, HD video, or the like. These can be accessed separately based on time information, or can be inserted at appropriate places in a map itself based on the time synchronization of the scan output with time information for the other media. In embodiments, a user may use time data to go back in time and see what has changed over time, such as based on multiple scans with different time-encoded data. Scans may be further enhanced with other information, such as date- or time-stamped service record data. Thus, a scan may be part of a multi-dimensional database of a scene or space, where point cloud data is associated with other data or media related to that scene, including time-based data or media. In embodiments, calculations are maintained through a sequence of steps or segments in a manner that allows a scan to be backed up, such as to return to a given point in the scan and re-initiate at that point, rather than having to re-initiate a new scan starting at the origin. This allows use of partial scan information as a starting point for continuing a scan, such as when a problem occurs at a later point in a scan that was initially producing good output. Thus, a user can “unzip” or “rewind” a scan back to a point, and then recommence scanning from that point. The system can maintain accurate position and location information based on the point cloud features and can maintain time information to allow sequencing with other time-based data. Time-based data can also allow editing of a scan or other media to synchronize them, such as where a scan was completed over time intervals and needs to be synchronized with other media that was captured over different time intervals. Data in a point cloud may be tagged with timestamps, so that data with timestamps occurring after a point in time to which a rewind is undertaken can be erased, such that the scan can re-commence from a designated point. In embodiments, a rewind may be undertaken to a point in time and/or to a physical location, such as rewinding to a geospatial coordinate.

[00167] In embodiments, the output from a SLAM-based map can be fused with other content, such as HD video, including by colorizing the point cloud and using it as an overlay. This may include time-synchronization between the SLAM system and other media capture system. Content may be fused with video, still images of a space, a CAD model of a space, audio content captured during a scan, metadata associated with a location, or other data or media.

[00168] In embodiments, a SLAM system may be integrated with other technologies and platforms, such as tools that may be used to manipulate point clouds (e.g., CAD). This may include combining scans with features that are modeled in CAD modeling tools, rapid prototyping systems, 3D printing systems, and other systems that can use point cloud or solid model data. Scans can be provided as inputs to post-processing tools, such as colorization tools. Scans can be provided to mapping tools, such as for adding points of interest, metadata, media content, annotations, navigation information, semantic analysis to distinguish particular shapes and/or identify objects, and the like.

[00169] Outputs can be combined with outputs from other scanning and image-capture systems, such as ground penetrating radar, X-ray imaging, magnetic resonance imaging, computed tomography imaging, thermal imaging, photography, video, SONAR, RADAR, LIDAR and the like. This may include integrating outputs of scans with displays for navigation and mapping systems, such as in-vehicle navigation systems, handheld mapping systems, mobile phone navigation systems, and others. Data from scans can be used to provide position and orientation data to other systems, including X, Y and Z position information, as well as pitch, roll and yaw information.

[00170] The data obtained from a real time SLAM system can be used for many different purposes, including for 3D motion capture systems, for acoustics engineering applications, for biomass measurements, for aircraft construction, for archeology, for architecture, engineering and construction, for augmented reality (AR), for autonomous cars, for autonomous mobile robot applications, for cleaning and treatment, for CAD/CAM applications, for construction site management (e.g., for validation of progress), for entertainment, for exploration (space, mining, underwater and the like), for forestry (including for logging and other forestry products like maple sugar management), for franchise management and compliance (e.g., for stores and restaurants), for imaging applications for validation and compliance, for indoor location, for interior design, for inventory checking, for landscape architecture, for mapping industrial spaces for maintenance, for mapping trucking routes, for military/intelligence applications, for mobile mapping, for monitoring oil pipelines and drilling, for property evaluation and other

real estate applications, for retail indoor location (such as marrying real time maps to inventory maps, and the like), for security applications, for stockpile monitoring (ore, logs, goods, etc.), for surveying (including doing pre-scans to do better quoting), for UAVs/drones, for mobile robots, for underground mapping, for 3D modeling applications, for virtual reality (including colorizing spaces), for warehouse management, for zoning/planning applications, for autonomous missions, for inspection applications, for docking applications (including spacecraft and watercraft), for cave exploration, and for video games/augmented reality applications, among many others. In all use scenarios, the SLAM system may operate as described herein in the noted areas of use.

[00171] In accordance with an exemplary and non-limiting embodiment, the unit comprises hardware synchronization of the IMU, camera (vision) and the LiDAR sensor. The unit may be operated in darkness or structureless environments for a duration of time. The processing pipeline may be comprised of modules. In darkness, the vision module may be bypassed. In structureless environments, the LiDAR module may be bypassed or partially bypassed. In exemplary and non-limiting embodiments, the IMU, Lidar and camera data are all time stamped and capable of being temporally matched and synchronized. As a result, the system can act in an automated fashion to synchronize image data and point cloud data. In some instances, color data from synchronized camera images may be used to color clod data pixels for display to the user.

[00172] The unit may comprise four CPU threads for scan matching and may run at, for example, 5Hz with Velodyne data. The motion of the unit when operating may be relatively fast. For example, the unit may operate at angular speeds of approximately 360 degree/second and linear speeds of approximately 30m/s.

[00173] The unit may localize to a prior generated map. For example, instead of building a map from scratch, the unit's software may refer to a previously built map and produce sensor poses and a new map within the framework (e.g., geospatial or other coordinates) of the old map. The unit can further extend a map using localization. By developing a new map in the old map frame, the new map can go further on and out of the old map. This enables different modes of use including branching and chaining, in which an initial "backbone" scan is generated first and potentially post-processed to reduce drift and/or other errors before resuming from the map to add local details, such as side rooms in building or increased point density in a region of interest. By taking this approach, the backbone model may be generated with extra care to limit the global drift and the follow-on scans may be generated with the focus

on capturing local detail. It is also possible for multiple devices to perform the detailed scanning off of the same base map for faster capture of a large region.

[00174] It is also possible to resume off of a model generated by a different type of device or even generated from a CAD drawing. For example, a higher global accuracy stationary device could build a base map and a mobile scanner could resume from that map and fill in details. In an alternate embodiment, a longer range device may scan the outside and large inside areas of a building and a shorter range device may resume from that scan to add in smaller corridors and rooms and required finer details. Resuming from CAD drawings could have significant advantages for detecting differences between CAD and as-built rapidly.

[00175] Resuming may also provide location registered temporal data. For example, multiple scans may be taken of a construction site over time to see the progress visually. In other embodiments multiple scans of a factory may help with tracking for asset management.

[00176] Resuming may alternately be used to purely provide localization data within the prior map. This may be useful for guiding a robotic vehicle or localizing new sensor data, such as images, thermal maps, acoustics, etc within an existing map.

[00177] In some embodiments, the unit employs relatively high CPU usage in a mapping mode and relatively low CPU usage in a localization mode, suitable for long-time localization/navigation. In some embodiments, the unit supports long-time operations by executing an internal reset every once in a while. This is advantageous as some of the values generated during internal processing increase over time. Over a long period of operation (e.g., a few days), the values may reach a limit, such as a logical or physical limit of storage for the value in a computer, causing the processing, absent a reset, to potentially fail. In some instances, the system may automatically flush RAM memory to improve performance. In other embodiments, the system may selectively down sample older scanned data as might be necessary when performing a real time comparison of newly acquired data with older and/or archived data.

[00178] In other exemplary embodiments, the unit may support a flying application and aerial-ground map merging. In other embodiments, the unit may compute a pose output at the IMU frequency, e.g., 100Hz. In such instances, the software may produce maps as well as sensor poses. The sensor poses tell the sensor position and pointing with respect to the map being developed. High frequency and accurate pose output helps in mobile autonomy because vehicle motion control requires such data. The unit further employs covariance and estimation confidence and may lock a pose when the sensor is static.

[00179] With reference to Figs. 17(a)-17(b) there is illustrated exemplary and non-limiting embodiments of a SLAM. LIDAR is rotated to create a substantially hemispherical scan. This is performed by a mechanism with a DC motor driving a spur gear reduction to LIDAR mount via a spur gear assembly 1704. The spur gear reduction assembly 1704 enables the LIDAR to be offset from the motor 1708. There is a slip ring in line with the rotation of the LIDAR to provide power and receive data from the spinning LIDAR. An encoder 1706 is also in line with the rotation of the LIDAR to record the orientation of the mechanism during scanning. A thin section contact bearing supports the LIDAR rotation shaft. Counterweights on the LIDAR rotation plate balance the weight about the axis of rotation making the rotation smooth and constant. As depicted in the LIDAR drive mechanism and attachment figures below, the mechanism is designed with minimal slop and backlash to enable maintenance of a constant speed for interpolation of scan point locations. Note that a motor shaft 1710 is in physical communication with a LIDAR connector 1712.

[00180] With reference to Fig. 18, there is illustrated an exemplary and non-limiting embodiment of a SLAM enclosure 1802. The SLAM enclosure 1802 is depicted in a variety of views and perspectives. Dimensions are representative of an embodiment and non-limiting as the size may be similar or different, while maintaining the general character and orientation of the major components, such as the LIDAR, odometry camera, colorization camera, user interface screen, and the like.

[00181] In some embodiments, the unit may employ a neck brace, shoulder brace, carrier, or other wearable element or device (not shown), such as to help an individual hold the unit while walking around. The unit or a supporting element or device may include one or more stabilizing elements to reduce shaking or vibration during the scan. In other embodiments, the unit may employ a remote battery that is carried in a shoulder bag or the like to reduce the weight of the handheld unit, whereby the scanning device has an external power source.

[00182] In other embodiments, the cameras and LIDAR are arranged to maximize a field of view. The camera-laser arrangement poses a tradeoff. On one side, the camera blocks the laser FOV and on the other side, the laser blocks the camera. In such an arrangement, both are blocked slightly but the blocking does not significantly sacrifice the mapping quality. In some embodiments, the camera points in the same direction as the laser because the vision processing is assisted by laser data. Laser range measurements provide depth information to the image features in the processing.

[00183] In some embodiments, there may be employed a confidence metric representing a confidence of spatial data. Such a confidence metric measurement may include, but is not

limited to, number of points, distribution of points, orthogonality of points, environmental geometry, and the like. One or more confidence metrics may be computed for laser data processing (e.g., scan matching) and for image processing. With reference to Figs. 19(a)-19(c), there are illustrated exemplary and non-limiting example images showing point clouds differentiated by laser match estimation confidence. While in practice, such images may be color coded, as illustrated, both the trajectory and the points are rendered as solid or dotted in the cloud based on a last confidence value at the time of recording. In the examples, dark gray is bad and light gray is good. The values are thresholded such that everything with a value >10 is solid. Through experimentation it has been found that with a Velodyne <1 is unreliable, <10 is less reliable, >10 is very good.

[00184] Using these metrics enables automated testing to resolve model issues and offline model correction such as when utilizing a loop-closure tool as discussed elsewhere herein. Use of these metrics further enables alerting the user when matches are bad and possibly auto-pausing, throwing out low confidence data, and alerting the user when scanning. Fig. 19(a) illustrates a scan of a building floor performed at a relatively slow pace. Fig. 19(b) illustrates a scan of the same building floor performed at a relatively quicker pace. Note the prevalence of light gray when compared to the scan acquired from a slower scan pace arising, in part, from the speed at which the scan is conducted. Fig. 19(c) illustrates a display zoomed in on a potential trouble spot of relatively low confidence.

[00185] With reference to Fig. 20, there is illustrated an exemplary and non-limiting embodiment of scan-to-scan match confidence metric processing and an average number of visual features that track between a full laser scan and a map being built from the prior full laser scans may be computed and presented visually. This metric may present useful, but different confidence measures. In Fig. 20, a laser scan confidence metric view is presented in the left frame while an average number of visual features metric is presented in the right frame for the same data. Again, dark gray line indicates lower confidence and/or fewer average number of visual features.

[00186] In some embodiments, there may be employed loop closure. For example, the unit may be operated as one walks around a room, cubicle, in and out of offices, and then back to a starting point. Ideally the mesh of data from the start and end point should mesh exactly. In the reality, there may be some drift. The algorithms described herein greatly minimize such drift. Typical reduction is on the order of 10x versus conventional methods (0.2% v 2%). This ratio reflects the error in distance between the start point and end point divided by the total distance traversed during the loop. In some embodiments, the software recognizes that it is

back to a starting point and it can relock to the origin. Once done, one may take the variation and spread it over all of the collected data. In other embodiments, one may lock in certain point cloud data where a confidence metric indicates that the data confidence was poor and one may apply the adjustments to the areas with low confidence.

[00187] In general, the system may employ both explicit and implicit loop closure. In some instances, a user may indicate, such as via a user interface forming a part of the SLAM, that a loop is to be closed. This explicit loop closure may result in the SLAM executing software that operates to match recently scanned data to data acquired at the beginning of the loop in order to snap the beginning and end acquired data together and close the loop. In other embodiments, the system may perform implicit loop closure. In such instances, the system may operate in an automated fashion to recognize that the system is actively rescanning a location that that comprises a point or region of origin for the scan loop.

[00188] In some embodiments, there may be employed confidence-based loop closure. First, one may determine a start and end point of a loop scan of an area that includes multiple segments. Then, one may determine a confidence of a plurality of the multiple segments. One may then make an adjustment to the lower quality segments rather than the higher quality segments in order for the beginning and end of the loop to be coincident.

[00189] In other exemplary embodiments, there may be performed multi-loop confidence-based loop closure. In yet other embodiments, there may be employed semantically adjusted confidence-based loop closure. For example, structural information may be derived from the attribution of a scanned element, i.e., floors are flat, corridors are straight, etc.

[00190] In some instances, there may be employed colorization of LIDAR point cloud data. In some embodiments, coarse colorization may be employed in real time to collected points to help identify what has been captured. In other embodiments, off-line photorealistic colorization may be employed. In other embodiments, each pixel in the camera can be mapped to a unique LIDAR pixel. For example, one may take color data from a pixel in the colorization camera corresponding to LIDAR data in the point cloud, and add the color data to the LIDAR data.

[00191] In accordance with exemplary and non-limiting embodiments, the unit may employ a sequential, multi-layer processing pipeline, solving for motion from coarse to fine. In each layer of the optimization, the prior coarser result is used as an initial guess to the optimization problem. The steps in the pipeline are:

[00192] 1. Start with IMU mechanization for motion prediction, which provides high frequency updates (on order of 200 Hz), but is subject to high levels of drift.

[00193] 2. Then this estimate is refined by a visual-inertial odometry optimization at the frame rate of the cameras (30-40Hz), the optimization problem uses the IMU motion estimate as an initial guess of pose change and adjusts that pose change in an attempt to minimize residual squared errors in motion between several features tracked from the current camera frame to a key frame.

[00194] 3. Then this estimate is further refined by a laser odometry optimization at a lower rate determined by the "scan frame" rate. Scan data comes in continuously, and software segments that data into frames, similar to image frames, at a regular rate, currently that rate is the time it takes for one rotation of the LIDAR rotary mechanism to make each scan frame a full hemisphere of data. That data is stitched together using visual-inertial estimates of position change as the points within the same scan frame are gathered. In the LIDAR odometry pose optimization step the visual odometry estimate is taken as an initial guess and the optimization attempts to reduce residual error in tracked features in the current scan frame matched to the prior scan frame.

[00195] 4. In the final step, the current scan frame is matched to the entire map so far. The laser odometry estimate is taken as the initial guess and the optimization minimizes residual squared errors between features in the current scan frame and features in the map so far.

[00196] The resulting system enables high-frequency, low-latency ego-motion estimation, along with dense, accurate 3D map registration. Further, the system is capable of handling sensor degradation by automatic reconfiguration bypassing failure modules since each step can correct errors in the prior step. Therefore, it can operate in the presence of highly dynamic motion as well as in dark, texture-less, and structure-less environments. During experiments, the system demonstrates 0.22% of relative position drift over 9.3km of navigation and robustness with respect to running, jumping and even highway speed driving (up to 33m/s).

[00197] Other key features of such a system may include:

[00198] **Visual feature optimization w/ and w/o depth:** The software may attempt to determine a depth of tracked visual features, first by attempting to associate them with the laser data and secondly by attempting to triangulate depth between camera frames. The feature optimization software may then utilize all features with two different error calculations, one for features with depth and one for features without depth.

[00199] **Laser feature determination:** The software may extract laser scan features as the scan line data comes in rather than in the entire scan frame. This is much easier and is done by looking at the smoothness at each point, which is defined by the relative distance between that point and the K nearest points on either side of that point then labeling the smoothest

points as planar features and the sharpest as edge features. It also allows for the deletion of some points that may be bad features.

[00200] Map matching and voxelization: A part of how the laser matching works in real-time is how the map and feature data is stored. Tracking the processor load of this stored data is critical to long term scanning and selectively voxelizing, or down-sampling into three-dimensional basic units in order to minimize the data stored while keeping what is needed for accurate matching. Adjusting the voxel sizes, or basic units, of this down-sampling on the fly based on processor load may improve the ability to maintaining real-time performance in large maps.

[00201] Parallel Processing: The software may be setup in such a way that it can utilize parallel processing to maintain real-time performance if data comes in faster than the processor can handle it. This is more relevant with faster point/second LIDARS like the velodyne.

[00202] Robustness: The way the system uses separate optimization steps without including the prior steps estimates as part of the next estimate (aside from being the initial guess) creates some inherent robustness.

[00203] Confidence Metrics: Each optimization step in this process may provide information on the confidence in its own results. In particular, in each step, the following can be evaluated to provide a measure of confidence in results: the remaining residual squared error after the optimization, the number of features tracked between frames, and the like.

[00204] The user may be presented a down scaled, (e.g., sub sampled) version of the multi-spectral model being prepared with data being acquired by the device. In an example, each measured 3cm x 3cm x 3cm cube of model data may be represented in the scaled down version presented on the user interface as a single pixel. The pixel selected for display may be the pixel that is closest to the center of the cube. A representative down-scaled display being generated during operation of the SLAM is shown below. As described, the decision to display a single pixel in a volume represents a binary result indicative of either the presence of one or more points in a point cloud occupying a spatial cube of defined dimensions or the absence of any such points. In other exemplary embodiments, the selected pixel may be attributed, such as with a value indicating the number of pixels inside the defined cube represented by the selected pixel. This attribute may be utilized when displaying the sub sampled point cloud such as by displaying each selected pixel utilizing color and/or intensity to reflect the value of the attribute.

[00205] A visual frame comprises a single 2D color image from the colorization camera. A LIDAR segment comprises a full 360 degree revolution of the LIDAR scanner. The visual

frame and LIDAR segment are synchronized so that they can be combined and aligned with the existing model data based on the unit positional data captured from the IMU and related sensors, such as the odometry (e.g., a high speed black/white) camera.

[00206] Problems giving rise to lower confidence metrics include, but are not limited to, reflections off of glossy walls, glass, transparent glass and narrow hallways. In some embodiments, a user of the unit may pause and resume a scan such as by, for example, hitting a pause button and/or requesting a rewind to a point that is a predetermined or requested number of seconds in the past.

[00207] In accordance with an exemplary and non-limiting embodiment, rewinding during a scan may proceed as follows. First, the user of the system indicates a desire to rewind. This may be achieved through the manipulation of user interface forming a part of the SLAM. As a result of indicating a desire to rewind, the system deletes or otherwise removes a portion of scanned data points corresponding to a duration of time. As all scanned data points are time stamped, the system can effectively remove data points after a predetermined time, thus, "rewinding" back to a previous point in a scan. As discussed herein, images from the camera are gathered and time stamped during each scan. As a result, after removing data points after a predetermined point in time, the system may provide the user with a display of an image recorded at the predetermined point in time while displaying the scanned point cloud rewound to the predetermined point in time. The image acts as a guide to help the user of the system reorient the SLAM into a position closely matching the orientation and pose of the SLAM at the previous predetermined point in time. Once the user is oriented close to the previous orientation of the SLAM at the predetermined point in time, the user may indicate a desire to resume scanning such as by engaging a "Go" button on a user interface of the SLAM. In response to the command to resume scanning, the SLAM may proceed execute a processing pipeline utilizing newly scanned data to form an initial estimation of the SLAMs position and orientation. During this process, the SLAM may not add new data to the scan but, rather, may use the newly scanned data to determine and display an instantaneous confidence level of the user's position as well as a visual representation of the extent to which newly acquired data corresponds to the previous scan data. Lastly, once it is established that the SLAM's location and orientation are sufficiently determined with respect to the previously scanned data, scanning may continue.

[00208] As described above, this ability to rewind is enabled, in part, by the data being stored. One may estimate how many points are brought in per second and then estimate how much to "rewind". The unit may inform the user where he was x seconds ago and allow the

user to move to that location and take a few scans to confirm that the user is at the appropriate place. For example, the user may be told an approximate place to go to (or the user indicate where they want to restart). If the user is close enough, the unit may figure it out where the user is and tell the user if they are close enough.

[00209] In other exemplary embodiments, the unit may operate in transitions between spaces. For example, If a user walks very quickly through a narrow doorway there may not be enough data and time to determine the user's place in the new space. Specifically, in this example, the boundaries of a door frame may, prior to proceeding through it, block the LIDAR from imaging a portion of the environment beyond the door sufficient to establish a user's location. One option is to detect this lowering of confidence metric and signal to the operator to modify his behavior upon approaching a narrow passage to slow down, such as by a flashing a visual indicator or a changing the color of the screen, and the like.

[00210] With reference to Fig. 21, there is illustrated an exemplary and non-limiting embodiment of a schematic of the SLAM unit 2100. The SLAM unit 2100 may include a timing server to generate multiple signals derived from the IMU's 2106 pulse-per-second (PPS) signal. The generated signals may be used to synchronize the data collected from the different sensors in the unit. A microcontroller 2102 may be used to generate the signals and communicate with the CPU 2104. The quadrature decoder 2108 may either be built into the microcontroller or on an external IC.

[00211] In some exemplary embodiments, the IMU 2206 supplies a rising edge PPS signal that is used to generate the timing pulses for other parts of the system. The camera may receive three signals generated from the IMU PPS signal including one rising edge signal as described above and two falling edge signals, GPIO1 (lasting one frame) and GPIO2 (lasting two frames as illustrated with reference to Fig. 22.

[00212] As illustrated, each camera receives a trigger signal synchronized with the IMU PPS having a high frame rate of approximately 30 Hz or 40 Hz and a high resolution of approximately 0.5 Hz - 5 Hz.

[00213] Each IMS PPS pulse may zero a counter internal to the microcontroller 2202. The LIDAR's synchronous output may trigger the following events:

[00214] • read the current encoder value through the quadrature decoder, and

[00215] • read the current counter value.

[00216] The encoder and the counter values may be saved together and sent to the CPU. This may happen every 40 Hz, dictated by the LIDAR synchronous output as illustrated with reference to Fig. 23.

[00217] An alternate time synchronization technique may include IMU based pulse-per-second synchronization that facilitates synchronizing the sensors and the computer processor. An exemplary and non-limiting embodiment of this type of synchronization is depicted with reference to Fig. 24.

[00218] The IMU 2400 may be configured to send a Pulse Per Second (PPS) signal 2406 to a LIDAR 2402. Every time a PPS is sent, the computer 2404 is notified by recognizing a flag in the IMU data stream. Then, the computer 2404 follows up and sends a time string to the LIDAR 2402. The LIDAR 2402 synchronizes to the PPS 2406 and encodes time stamps in the LIDAR data stream based on the received time strings.

[00219] Upon receiving the first PPS 2406, the computer 2404 records its system time. Starting from the second PPS, the computer 2404 increases the recorded time by one second, sends the resulting time string to the LIDAR 2402, and then corrects its own system time to track the PPS 2506.

[00220] In this time synchronization scheme, the IMU 2400 functions as the time server, while the initial time is obtained from the computer system time. The IMU 2400 data stream is associated with time stamps based on its own clock, and initialized with the computer system time when the first PPS 2406 is sent. Therefore, the IMU2400, LIDAR 2402, and computer 2404 are all time synchronized. In embodiments, the LIDAR 2402 may be a Velodyne LIDAR.

[00221] In accordance with exemplary and non-limiting embodiments, the unit includes a COM express board and a single button interface for scanning.

[00222] In accordance with exemplary and non-limiting embodiments, the process IMU, vision and laser data sensors may be coupled. The unit may work in darkness or structureless environments for long periods of time. In some embodiments, four CPU threads may be employed for scan matching, each running at 5Hz with Velodyne data. As noted above, motion of the unit may be fast and the unit may localize to a prior map and can extend a map using localization. The unit exhibits relatively high CPU usage in mapping mode and relatively low CPU usage in localization mode thus rendering it suitable for long-time.

[00223] In accordance with exemplary and non-limiting embodiments, methods and systems described herein enable collaboration between mapping from the ground and mapping from the air according to the characteristics of each. Ground-based mapping is not necessarily prone to limitations of space or time. Typically, a mapping device carried by a ground vehicle is suitable for mapping in large scale and can move at a high speed. On the other hand, a tight area can be mapped in a hand-held deployment. However, ground-based mapping is limited by the sensor's altitude making it difficult to realize a top-down looking configuration. As

illustrated in Fig. 25, the ground-based experiment produces a detailed map of the surroundings of a building, while the roof has to be mapped from the air. If a small aerial vehicle is used, aerial mapping is limited by time due to the short lifespan of batteries. Space also needs to be open enough for aerial vehicles to operate safely.

[00224] The collaborative mapping as described herein may utilize a laser scanner, a camera, and a low-grade IMU to process data through multi-layer optimization. The resulting motion estimates may be at a high rate (~200Hz) with a low drift (typically <0.1% of the distance traveled).

[00225] The high-accuracy processing pipeline described herein may be utilized to merge maps generated from the ground with maps generated from the air in real-time or near real-time. This is achieved, in part, by localization of one output from a ground derived map with respect to an output from an air derived map.

[00226] While the method disclosed herein fulfills collaborative mapping, it further reduces the complexity of aerial deployments. With a ground-based map, flight paths are defined and an aerial vehicle conducts mapping in autonomous missions. In some embodiments, the aerial vehicle is able to accomplish challenging flight tasks autonomously.

[00227] With reference to Fig. 26, there is illustrated an exemplary and non-limiting embodiment of a sensor/computer pack that may be utilized to enable collaborative mapping. The processing software is not necessarily limited to a particular sensor configuration. With reference to Fig. 26(a), the sensor pack 2601 is comprised of a laser scanner 2603 generating 0.3 million points/second, a camera 2605 at 640×360 pixels resolution and 50Hz frame rate, and a low-grade IMU 2607 at 200Hz. An onboard computer processes data from the sensors in real-time for ego-motion estimation and mapping. Fig. 26(b) and Fig. 26(c) illustrate the sensor field of view. An overlap is shared by the laser and camera, with which, the processing software associates depth information from the laser to image features as described more fully below.

[00228] In accordance with exemplary and non-limiting embodiments, the software processes data from a range sensor such as a laser scanner, a camera, and an inertial sensor. Instead of combining data from all sensors in a large, full-blown problem, the methods and systems described herein parse the problem as multiple small problems, solve them sequentially in a coarse-to-fine manner. Fig. 27 illustrates a block diagram of the software system. In such a system, modules in the front conduct light processing, ensuring high-frequency motion estimation robust to aggressive motion. Modules in the back take sufficient

processing, run at low frequencies to warrant accuracy of the resulting motion estimates and maps.

[00229] The software starts with IMU data processing 2701. This module runs at the IMU frequency to predict the motion based on IMU mechanization. The result is further processed by a visual-inertial coupled module 2703. The module 2703 tracks distinctive image features through the image sequence and solves for the motion in an optimization problem. Here, laser range measurements are registered on a depthmap, with which, depth information is associated to the tracked image features. Since the sensor pack contains a single camera, depth from the laser helps solve scale ambiguity during motion estimation. The estimated motion is used to register laser scans locally. In the third module 2705, these scans are matched to further refine the motion estimates. The matched scans are registered on a map while scans are matched to the map. To accelerate the processing, scan matching utilizes multiple CPU threads in parallel. The map is stored in voxels to accelerate point query during scan matching. Because the motion is estimated at different frequencies, a fourth module 2707 in the system takes these motion estimates for integration. The output holds both high accuracy and low latency beneficial for vehicle control.

[00230] The modularized system also ensures robustness with respect to sensor degradation, by selecting “healthy” modes of the sensors when forming the final solution. For example, when a camera is in a low-light or texture-less environment such as pointing to a clean and white wall, or a laser is in a symmetric or extruded environment such as a long and straight corridor, processing typically fails to generate valid motion estimates. The system may automatically determine a degraded subspace in the problem state space. When degradation happens, the system only solves the problem partially in the well-conditioned subspace of each module. The result is that the “healthy” parts are combined to produce the final, valid motion estimates.

[00231] When a map is available, the method described above can be extended to utilize the map for localization. This is accomplished using a scan matching method. The method extracts two types of geometric features, specifically, points on edges and planar surfaces, based on the curvature in local scans. Feature points are matched to the map. An edge point is matched to an edge line segment, and a planar point is matched to a local planar patch. On the map, the edge line segments and local planar patches are determined by examining the eigenvalues and eigenvectors associated with local point clusters. The map is stored in voxels to accelerate processing. The localization solves an optimization problem minimizing the overall distances

between the feature points and their correspondences. Due to the fact that the high accuracy odometry estimation is used to provide initial guess to the localization, the optimization usually converges in 2-3 iterations.

[00232] The localization does not necessarily process individual scans but, rather, stacks a number of scans for batch processing. Thanks to the high-accuracy odometry estimation, scans are registered precisely in a local coordinate frame where drift is negligible over a short period of time (e.g., a few seconds). A comparison is illustrated with reference to Fig. 28 (8.4) where Fig. 28(a) is a single scan that is matched in the previous section (scan matching executes at 5Hz), and Fig. 28(b) shows stacked scans over two seconds, which are matched during localization (scan matching runs at 0.5Hz). One can see the stacked scans contain significantly more structural details, contributing to the localization accuracy and robustness with respect to environmental changes. Additionally, low-frequency execution keeps the CPU usage to be minimal for onboard processing (localization consumes about 10% of a CPU thread). Stacked scans contain significantly more structural details, contributing to the localization accuracy and robustness with respect to environmental changes. Roughly, about 25% of the environment can change or be dynamic and the system will continue to operate well.

[00233] The localization is compared to a particle filter based implementation. The odometry estimation provides the motion model to the particle filter. It uses a number of 50 particles. At each update step, the particles are resampled based on low-variance resampling. Comparison results are shown in Fig. 29 and Table 8.1. Here, errors are defined as the absolute distances from localized scans to the map. During the evaluation, the methods and systems described herein choose a number of planar surfaces and use the distances between points in localized scans to the corresponding planar patches on the map. Fig. 29 illustrates an exemplary and non-limiting embodiment of an error distribution. When running the particle filter at the same frequency as the described method (0.5Hz), the resulting error is five times as large. While in Table 8.1, the CPU processing time is more than double. In another test, running the particle filter at 5Hz helps reduce the error to be slightly larger than the disclosed method. However, the corresponding CPU processing time increases to over 22 times greater. These results imply a particle filter based method does not necessarily take full advantage of the high-accuracy odometry estimation.

Table 8.1 Comparison of CPU processing time in localization. When running the particle filter at 5Hz, denser data is processed at 25% of real-time speed due to high CPU demand.

Method	Particle filter		Ours
Frequency	0.5Hz	5Hz	0.5Hz
Time per execution	493ms	478ms	214ms
Time per second	247ms	2390ms	107ms

[00234] With reference to Fig. 30, there is illustrated an exemplary non-limiting embodiment of a sensor study wherein the sensor pack is carried horizontally in a garage building. Fig. 30(a) shows the map built and sensor trajectory. Fig. 30(b) is a single scan. In this scenario, the scan contains sufficient structural information. When bypassing the camera processing module, the system produces the same trajectory as the full pipeline. On the other hand, the methods and systems described herein run another test with the sensor pack tilted vertically down toward the ground. The results are shown in Fig. 31. In this scenario, structural information in a scan is much sparser as shown in Fig. 31(b)). The processing fails without usage of the camera and succeeds with the full pipeline. The results indicate the camera is critical for high-altitude flights where tilting of the sensor pack is required.

[00235] With reference to Fig. 32, there is illustrated an exemplary and non-limiting embodiment wherein the sensor pack is held by an operator walking through a circle at 1-2m/s speed with an overall traveling distance of 410m. Fig. 32(a) shows the resulting map and sensor trajectory with a horizontally orientated sensor configuration. The sensor is started and stopped at the same position. The test produces 0.18m of drift through the path, resulting in 0.04% of relative position error in comparison to the distance traveled. Then, the operator repeats the path with two sensor packs held at 45° and 90° angles, respectively. The resulting sensor trajectories are shown in Fig. 32(b). Clearly, tilting introduces more drift, where the relative position errors are 0.6% at 45° (blue dash curve) and 1.4% at 90° (red dash-dot curve). Finally, by localizing on the map in Fig. 32(a), the drift is canceled and both configurations result in the black solid curve.

[00236] An exemplary and non-limiting embodiment of a drone platform 3301 is illustrated at Fig. 33. The aircraft weighs approximately 6.8kg (including batteries) and may carry a maximum of 4.2kg payload. The sensor/computer pack is mounted to the bottom of the aircraft, weighting 1.7kg. The bottom right of the figure shows the remote controller. During autonomous missions, the remote controller is operated by a safety pilot to override the autonomy if necessary. Note that the aircraft is built with a GPS receiver (on top of the aircraft). GPS data is not necessarily used in mapping or autonomous.

[00237] In the first collaborative mapping experiment, an operator holds the sensor pack and walks around a building. Results are shown in Fig. 25. In Fig. 25(a), the ground-based mapping covers surroundings of the building in detail, conducted at 1-2m/s over 914m of travel. As expected, the roof of the building is empty on the map. Second, the drone is teleoperated to fly over the building. In Fig. 25(b), the flight is conducted at 2-3m/s with a traveling distance of 269m. The processing uses localization w.r.t. the map in Fig. 25(a). That way, the aerial map is merged with the ground-based map (white points). After the ground-based map is built, the take-off position of the drone is determined on the map. The sensor starting pose for the aerial mapping is known, and from which, the localization starts. Fig. 34 presents the aerial and ground-based sensor trajectories, in top-down and side views.

[00238] Further, the methods and systems described herein conduct autonomous flights to realize aerial mapping. With reference to Fig. 35, a ground-based map is built first by hand-held mapping at 1-2m/s for 672m of travel around the flight area. The map and sensor trajectory are shown in Fig. 35(a). Then based on the map, way-points are defined and the drone follows the way-points to conduct aerial mapping. As shown in Fig. 35(b), the curve is the flight path, the large points on the curve are the way-points, and the points form the aerial map. In this experiment, the drone takes off inside a shed on the left side of the figure, flies across the site and passes through another shed on the right side, then returns to the first shed to land. The speed is 4m/s crossing the site and 2m/s passing through the shed. Fig. 35(c) and Fig. 35(d) are two images taken by an onboard camera when the drone flies toward the shed on the right and is about to enter the shed. Fig. 35(e) shows the estimated speed during the mission.

[00239] Finally, the methods and systems described herein conduct another experiment over a longer distance. As shown in Fig. 36, the ground-based mapping involves an off-road vehicle driven at 10m/s from the left end to the right end, over 1463m of travel. With the ground-based map and way-points, the autonomous flight crosses the site. Upon take-off, the drone ascends to 20m high above the ground at 15m/s. Then, it descends to 2m above the ground to fly through a line of trees at 10m/s. The flight path is 1118m long as indicated by the curve 3601 in Fig. 36(b). Two images are taken as the drone flies high above the trees (see Fig. 36(c)) and low underneath the trees (see Fig. 36(d)).

[00240] In accordance with an exemplary and non-limiting embodiment, global positioning data such as from a GPS may be incorporated into the processing pipeline. Global positioning data can be helpful to cancel ego-motion estimation drift over a long distance of travel and register maps in a global coordinate frame.

[00241] In some embodiments, GPS data may be recorded simultaneously with mapping activity. As the system moves there is some level of drift that causes an error to grow over distance. One may typically experience only 0.2% drift rate but when traveling 1000 meters that is still 2 meters for every 1000 meters of travel. At 10km this grows to 20 meters, etc. Without closing the loop (in the traditional sense of coming back to the beginning of the route) this error cannot be corrected. By providing external information in the form of GPS coordinates, the system can know where it is and correct the current position estimate. While this is typically done in a post-processing effort, the present system is able to accomplish such a correction in real-time or near real-time.

[00242] Thus, utilizing GPS provides a method by which one may close the loop. GPS, of course, has some amount of error as well, but it is usually consistent in a given area and many GPS systems today can provide better than 30cm accuracy in position X and Y on the surface. Other more expensive and sophisticated systems can provide cm level positioning.

[00243] This use of GPS provides important capabilities: 1. The location of the point cloud on the planet. 2. the ability to use the course-corrected information to align and "fix" the map so that the map becomes even more accurate since one knows one's position and any data taken at that position may now be referenced to the series of GPS points that are also collected. 3. The ability for the system to act as an IMU when GPS is lost.

[00244] In accordance with an exemplary and non-limiting embodiment, dynamic vision sensors may be utilized to further improve estimation robustness with respect to aggressive motion. A dynamic vision sensor reports data only on pixels with illumination changes, delivering both a high rate and a low latency.

[00245] This high rate (typically defined as more than approximately 10Hz) may provide rapid information quickly to the ego-motion and estimation system thus improving values for localization and, subsequently, mapping. If the system is able to capture more data with fewer delays, the system will be more accurate, and more robust. The features that are identified and tracked by the dynamic vision sensor enable better estimates since more features, and faster updates enable more accurate tracking and motion estimation.

[00246] Direct methods may be used to realize image matching with a dynamic vision sensor for ego-motion estimation. Specifically, direct methods match sequential images for feature tracking from image to image. In contrast, the feature tracking method disclosed herein is superior to the direct method.

[00247] In accordance with an exemplary and non-limiting embodiment, parallel processing may be implemented to execute on a general purpose GPU or FPGA and therefore enable data processing in larger amount and higher frequencies.

[00248] Using parallel processing techniques, a number of feature may be tracked simultaneously or data in defined areas or directions may be partitioned into multiple smaller problems. Parallel architectures may take the form of multiple cores, thread, processors or even specialized forms such as Graphics Processing Units (GPUs). By employing a divide and conquer approach, there may be a significant speed-up in overall processing as each node in the parallel architecture works on its portion of the problem at the same time. Typically such speed-ups are not linear. That is, by dividing the problem into 16 parts and processing each part separately does not necessarily speed up the processing by 16 times. There is overhead in making the division, communications, and then re-assembling the results.

[00249] In contrast, with linear processing, each sub-problem or calculation is worked on once and data is shuttled in and out of memory, the CPU, and then back into memory or storage. This is slower although a pipeline architecture, if set up correctly, such as in systolic arrays can improve things as long as the pipeline is kept filled.

[00250] In accordance with an exemplary and non-limiting embodiment, loop closures may be introduced to remove ego-motion estimation drift by global smoothing.

[00251] For example, when you return to your starting point you know that the position estimate can be corrected since you were there already. The starting point is your origin and there must have been an accumulated error, which, at the end of the traverse, becomes evident. By taking the difference between the origin (0,0,0) and the position you think you are at when you return to the beginning (e.g. (1,2,3)) you now know that error value. But the value accumulated, perhaps evenly, over the traverse.

[00252] For example, if you traveled 100 meters and returned to your starting point and found you were off by 1 meter in xyz, you now have the problem of determining where to apply that error over the whole distance. One way would be to spread that error over the entire distance then you might correct the travel by 1cm at each meter increment to stay on the path and, importantly, be back where you started with no error. This is global smoothing.

[00253] Another way may be to use some other measure of error during the traverse and applying corrections for only those places where the measure of error is high. For example, the covariance matrix provides a convenient metric for map quality and may be used to proactively distribute this error over the full traverse. In other words, locations with low-quality matching that show up in the covariance matrix could be used to unevenly spread the error across the

traverse in accordance with the proportion of the error over that traverse. This will act to correct the error in the precise areas where low-quality values were associated with particular locations.

[00254] In accordance with exemplary and non-limiting embodiments, there is disclosed one or more means to determine position, velocity, and other information of a sensing device such as a commonly-carried cellular phone, smart phone, or other device that picks up, for example, a variety of wireless signals and uses a camera or other sensors to provide visual or geometric information of the environment such as may be utilized to produce a 3D model or map of said environment.

[00255] Once a position is known, then other separate measurements, such as with other sensors, may be made and assessed at the known location. By moving around either in random or pre-determined patterns such as a raster pattern to provide complete or sufficient coverage, a map of the environment may be visualized and/or stored that will provide the separately sensed data values coordinated with the mapped locations, such as within a facility.

[00256] This calculation of position or, more generally, pose, can be combined with data from the other sensors to produce maps of any of the sensed qualities in multiple dimensions, including the same dimensions being captured over time to, for example, model time-varying qualities by location. This combined data may be used to model time-varying qualities, and the like. Time is a useful aspect of dynamic data and adds another dimension of interest. Variations can be just as important. One may either provide time-based measurements assuming some periodicity to the signals or extract salient information about distribution using slower traverses.

[00257] With reference to Fig. 37A, there is illustrated an exemplary and non-limiting embodiment of mapping and data embedding. First, at step 3700, a data gathering device, such as a SLAM, is moved through an environment to create a map recording both point cloud data and sensor pose information. Simultaneously, at step 3702, non-geospatial sensor data is recorded and time information is used to synchronize the map position with the sensor information. Lastly, at step 3704, an overall map that includes geometry and additional non-geospatial sensor information is created. Non-geospatial sensors may be one or more of RF detection, thermal, visual, acoustic, electrical signal, magnetic, radioactive, or other sensing modalities.

[00258] With reference to Fig. 37B, there is illustrated an exemplary and non-limiting embodiment of deriving sensor position from the map of Fig. 37A. First, at step 3706, the

map embedded with non-geospatial sensor data is retrieved. At step 3708, newly acquired non-geospatial sensor data is collected from a device. Then, at step 3710, localization software matches the received sensor readings with one or more positions in the map that exhibits similar values corresponding to the sensor readings. Lastly, one may optionally embed the newly acquired sensor pose data into the map.

[00259] As used herein, a "heat map" refers to a map comprised of a plurality of non-spatial attributes each having a value indicative of an intensity of measurement. In some instances, a heat map may comprise non-spatial attribute values for a region or sub-region of the extent comprising the map. In other instances, the heat map may be formed of a plurality of discreet non-spatial attribute values from which additional heat map values may be interpolated.

Location-based examples and sensing modalities may include:

- Mapping of buried ordnance, or other objects through the use magnetometers. Wherein magnetic fields, or other quantities are indicative of the presence or absence of ferrous materials.
- Mapping of magnetic field along high-voltage power lines. Still using magnetometers, to evaluate the magnetic interference to other electrical devices close to the power lines.
- Mapping of radioactivity levels in a facility such as a nuclear power plant or processing plants.
- Mapping of temperature, thermal information, in all locations – a heat map.
- Light level mapping to capture spectral power distribution or photometric values from which a variety of other measurements can be calculated.
- Temporal changes - If done at different times of the day or over a season or less or greater frequency than a temporal and spatial map of such values can be created to provide useful time-varying information. Measurements done at different times of the day or over a season or less or greater frequency than a temporal and spatial map of such values can be created to provide useful time-varying information.
- Measures of RF sources to provide signal strength as a function of distance.
- Audio measurements in 3D, such as to map the acoustics of a performing arts hall or other arena where sound quality matters.

[00260] With reference to Fig. 38, there is illustrated an exemplary and non-limiting embodiment of a heat map style representation 3800 of sensed information superimposed on floor plan geometry. Such data could represent one or more of many sensing modalities and/or sensor attributes. As illustrated, contour lines reflect areas of similar attribute value.

[00261] In these examples, the colocation of a localization means, such as a SLAM device with that of an additional sensor, or a wireless communication device to provide another measurement may provide an accurate map of the sensed quantity or quantities of the information including the information collected by the additional sensor in space and/or time.

[00262] When scanning an environment, movement of the device provides variation over a much larger set of inputs and calculations. In some embodiments, one may utilize advanced SLAM-based real-time mapping techniques thus providing much greater information. For example, 100HZ information versus 0.1Hz information would supply several orders of magnitude more information and better tie information at any instant to information before and after a given measurement.

[00263] In this way, an accurate map may be made with a sensor set, such as a SLAM device and a mobile phone or other signal detection device as described above and elsewhere herein, that may be impractical to deploy frequently. This set may be too expensive, too large, too heavy, or too unwieldy to use or be limited in duration and accuracy. But by using a more sophisticated mapping system to create an initial accurate map that may include multiple modes of sensing (e.g., SLAM plus WiFi signal, heat, radioactivity, sunlight, and others), a lesser system can later be enabled to facilitate determination of a position that is more accurate than the lesser system by itself would provide.

[00264] For example, a robot or a person might carry a mapping and localization device, such as a SLAM device, to produce a registered 3D point cloud that can later be converted into a map of the 3D geometry of the environment. Other signals such as RF sources such as Wi-Fi and Bluetooth can be integrated into that map and database so that when a device that uses one or more of these other signals is introduced into the environment, such as at an unknown position, the integrated map and database may be used to provide a position fix for the newly introduced device.

[00265] The density and quality of the resultant maps can be improved by interpolating between measured points within a given sensor type (e.g., RF) or across sensor types (e.g., RF and magnetic) to provide a more continuous representation of the various different sensor data in the environment. Such an interpolation may or may not be linear as some sensed quantities may fall off as a power function with distance. Magnetic fields are one example where the level may drop off exponentially.

[00266] The use of this multi-modal data set may include determining placement of RF or other signaling devices to ensure adequate coverage.

[00267] One example commercial need is the placement of transmitting and receiving micro antennas for RF sources including, but not limited to, 5G, LTE, Wi-Fi, or Bluetooth to provide adequate and complete coverage. In essence, this may be the opposite of ‘fingerprinting’ because the problem is determining the best placement of the source rather than recording current placement.

[00268] The superposition of wireless communication reception and 3D mapping tools allow a detailed analysis of the signal strength across large areas. This enables initial correspondence and mapping of those signal sources and this may, in turn, be used to determine one’s position. For telecommunications, Long-Term Evolution (LTE) is one example of a standard for high-speed wireless communication for cell phones and other data sources. This process can also be used to evaluate the reception of signals from such devices to ensure they provide full coverage or specialized coverage of a site. This is an important application because if this placement is not optimized then either gaps exist, or there is inefficient overlapping coverage.

[00269] Location-aware services continue to expand beyond outside GPS-location services. Interiors, which are usually GPS-denied, offer a new generation of location needs and services. Many alternatives have been suggested using infrastructure, beacons, RF signals, bar codes, and so forth. Wi-Fi, one such RF signal, is attractive because smartphones are equipped to receive such signals. In spaces such as shopping malls, offices, airports, factories, warehouses, and service industry venues such as restaurants and hotels, Wi-Fi and Bluetooth signals are widely used and a combination of features such as signal strength and relative geometry can be used to help determine position of a phone and thus, the person carrying it. A signal propagation model using signal strength may be used as a basis for determining position.

[00270] An exemplary approach to generating and using a multi-modal map of an environment includes determining position. A user may create real-time and interactive 3D models using a device such as a SLAM device while simultaneously gathering information on signal strength from a variety of sources such as Bluetooth, Wi-Fi, cell tower, and the like. By identifying these and other RF-based information sources, for example, the 3D model generated by the SLAM device may be used to determine signal strength, which may translate into position of RF or other beacon sources. In addition, smartphone or other cameras may also be used to help identify places within the environment. By knowing the camera position (note that the camera may be part of a smartphone or other device) and orientation, the image and subsequent analysis may be used to further determine the position

and even be used to layer graphics and identifying information such as direction of travel, places of interest, and directions. Examples include: (i) finding one's position; (ii) formulating a path to a destination, such as an exit, fire extinguisher, human rescue, other interior location (e.g., another store in a mall); (iii) identifying a place of interest; (iv) forming a basis for augmented reality of the image being captured through the phone camera or the like.

[00271] Indoor positioning may be based on signal strength from several sources including, but not limited to, Wi-Fi. One approach includes using localization algorithms to find a position that corresponds to a particular set of signals. Another approach may include using those algorithms based on a fingerprinted position. The former uses the signal propagation model to convert measured signal strength into distance information. Then, the target coordinates can be calculated according to the distance between the moving target and the multiple access points with known coordinates.

[00272] In most buildings, the actual deployment of Wi-Fi access points is very irregular, such as due to a lack of centralized planning and because signal propagation is seriously influenced by some factors such as the building structure. Wi-Fi signals are susceptible to environmental factors including walls, windows, equipment, and windows, even people. The attenuation level is related to the shape, size, and material of the materials. Therefore, for indoor environments with complex structure, a generalized signal propagation model may not directly and accurately correspond between the actual distance and the signal strength, even given a good understanding of the position of the Wi-Fi signal sources (e.g., the routers, repeaters, and the like).

[00273] Fingerprint localization consists of offline location-fingerprint database generation and online positioning. Typically, to build the database, some appointed locations in the building are sampled for signal strength. A collection of Wi-Fi signals and their intensity will be recorded and considered as signal fingerprint.

[00274] In some embodiments, signals may be normalized and statistical measures including regression, Chi-squared, or Bayesian formulations may all be used to classify data, determining if sensed quantities are independent or not. Other classifiers include neural networks (CNN, ANN, etc) and this would require training sets to successfully classify inputs and sensor reading. Over time certain signals that show less variation may be accorded greater confidence and thus greater confidence implies greater trust in those values and resultant calculations of position.

[00275] The collection and fingerprinting is done in coordination with the high-resolution mapping (e.g., SLAM based) so that the fingerprints can be aligned to the map. To perform online positioning after the off-line multi-modal database is available, signal information is collected around the position to be localized and used to generate a currently sensed fingerprint. These currently sensed fingerprints may be compared with fingerprints in the offline database by, for example, determining a best match to the fingerprints in the offline database. The corresponding map position matching the position whose fingerprint can attain the best match is chosen as the estimated position. In contrast, although the Wi-Fi fingerprint-based positioning technology needs to make a fingerprint database, or map, at the early stage, it may effectively avoid the influence of building structure. Furthermore, fingerprint-based methods do not require that Wi-Fi access points are known beforehand. Therefore, it has a higher practicability.

[00276] One problem that can be simplified using the methods described herein include determining how best to provide good LTE or the like signal coverage indoors. A warehouse-type store or a multi-facility environment is a typical environment. Current solutions include installing micro antennas, often in the ceilings at some model-based distribution algorithm. In contrast, walking around the environment with a 3D real-time mapping device and recording the signal strength at each position can be used to optimize antenna positioning. This multi-modal data could be processed to create, among other things a color-coded volume showing changes in signal strength. Antenna placement could then be optimized by knowing visually where the signal is strong and where it is weak.

[00277] Uses of this data could also help with placement of ultra-high bandwidth coverage, such as 5G networks to provide high bandwidth data. Some applications could greatly benefit from these systems that have bandwidths that are up to three orders of magnitude greater bandwidth than Wi-Fi rates. For example, medical imaging and the like could use high-resolution 4K, 8K, and the like cameras, microphones, EKG machines, ultrasound machines, and the like. Ensuring that bandwidth coverage for such devices, while ensuring lower bandwidth coverage elsewhere in an environment may allow real-time remote monitoring and control.

[00278] Another use may include providing a localization system utilizing a combination of active and passive sensing modalities to measures local geometry to provide a map that is based on the simultaneous gathering of RF sources that also provides readings of signal strength of RF sources in the environment.

[00279] An exemplary multi-modal sensing system may use one or more of each of the following: LIDAR, camera, depth camera, infrared camera, IMU, ultrasonic, Wi-Fi, Bluetooth, temperature sensor, barometer. The LIDAR may be a 1D (single point distance) or a 2D (line) or a 3D LIDAR with mechanical or solid-state means to control the beam. The camera may operate at a high or low frame rate. The camera may be high (e.g., 4K) or low resolution. The depth camera may have long or short range. Such a device operated with the methods and systems described herein may provide accurate measurement with noisy data sources.

[00280] With one or more sensors, one may localize on a map that is embedded with data acquired by the same sensors. Once localization is determined, one can retrieve the old data embedded on the map, compare it with the new data to detect environmental changes, e.g. temperature change and color change in the infrared domain. One can also embed the new data on the map, such that a map can have multiple layers, while each layer contains data logged at a different time. Such a procedure provides a record of environmental changes over time, all embedded on a single map.

[00281] By performing localization on a map, one can also retrieve the color, texture, and 3D information from the map. Information as such may be helpful for visual assistance to people in determining their self-locations. The information may also be used by a computer software to generate a path to a destination, e.g. a restaurant or café, for guidance of people in GPS occluded areas.

[00282] Sensors such as cameras and LIDAR produce signals that can vary considerably in capability and accuracy. Range measurement for example, in some systems can range from meters to sub-millimeters depending on the design, construction, and calibration of these units. Very often there is a trade-off between accuracy and cost.

[00283] Other aspects of the methods and systems of multi-modal SLAM-based 3D mapping with additional sensing may include a user interface to facilitate measurements between the means. A user interface tool that allows the measurements by taking the spatial median value of point cloud sub-groups such as walls and providing a measure to another such point cloud structure. The mean-to-mean distance is thus more accurate than any one set of measurements between the two surfaces. The user interface may take the form of a measurement by drawing a line or a curve. The software may then cluster the overlaid points and provide a single line corresponding to the median value of the line.

[00284] Another aspect may include combining low-resolution with high-resolution to produce high quality maps. The combination of a lower resolution LIDAR to provide good-

enough positioning information with a higher resolution LIDAR to provide high-resolution geometry of the environment combined with high fidelity LIDAR imagery. This combines the use of a lower resolution device to provide good positioning by referencing the maps produced by the higher resolution devices while using the lower resolution device.

[00285] While the disclosure has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the present disclosure is not to be limited by the foregoing examples, but is to be understood in the broadest sense allowable by law.

[00286] In accordance with some exemplary embodiments, combinations of maps from multiple mapping units may be combined by wireless transmission or file transfers so many mappers update a complete map in real-time. In other embodiments, multiple consumers may be led towards each other to find someone else or for two people to efficiently find each other in the shortest time possible. In other embodiments, paths may be stored for first responders to find people quickly.

[00287] While only a few embodiments of the present disclosure have been shown and described, it will be obvious to those skilled in the art that many changes and modifications may be made thereunto without departing from the spirit and scope of the present disclosure as described in the following claims. All patent applications and patents, both foreign and domestic, and all other publications referenced herein are incorporated herein in their entireties to the full extent permitted by law.

[00288] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software, program codes, and/or instructions on a processor. The present disclosure may be implemented as a method on the machine, as a system or apparatus as part of or in relation to the machine, or as a computer program product embodied in a computer readable medium executing on one or more of the machines. In embodiments, the processor may be part of a server, cloud server, client, network infrastructure, mobile computing platform, stationary computing platform, or other computing platform. A processor may be any kind of computational or processing device capable of executing program instructions, codes, binary instructions and the like. The processor may be or may include a signal processor, digital processor, embedded processor, microprocessor or any variant such as a co-processor (math co-processor, graphic co-processor, communication co-processor and the like) and the like that may directly or indirectly facilitate execution of program code or program instructions stored thereon. In addition, the processor may enable

execution of multiple programs, threads, and codes. The threads may be executed simultaneously to enhance the performance of the processor and to facilitate simultaneous operations of the application. By way of implementation, methods, program codes, program instructions and the like described herein may be implemented in one or more thread. The thread may spawn other threads that may have assigned priorities associated with them; the processor may execute these threads based on priority or any other order based on instructions provided in the program code. The processor, or any machine utilizing one, may include non-transitory memory that stores methods, codes, instructions and programs as described herein and elsewhere. The processor may access a non-transitory storage medium through an interface that may store methods, codes, and instructions as described herein and elsewhere. The storage medium associated with the processor for storing methods, programs, codes, program instructions or other type of instructions capable of being executed by the computing or processing device may include but may not be limited to one or more of a CD-ROM, DVD, memory, hard disk, flash drive, RAM, ROM, cache and the like.

[00289] A processor may include one or more cores that may enhance speed and performance of a multiprocessor. In embodiments, the process may be a dual core processor, quad core processors, other chip-level multiprocessor and the like that combine two or more independent cores (called a die).

[00290] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software on a server, client, firewall, gateway, hub, router, or other such computer and/or networking hardware. The software program may be associated with a server that may include a file server, print server, domain server, internet server, intranet server, cloud server, and other variants such as secondary server, host server, distributed server and the like. The server may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other servers, clients, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the server. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the server.

[00291] The server may provide an interface to other devices including, without limitation, clients, other servers, printers, database servers, print servers, file servers, communication servers, distributed servers, social networks, and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking

of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the server through an interface may include at least one storage medium capable of storing methods, programs, code and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[00292] The software program may be associated with a client that may include a file client, print client, domain client, internet client, intranet client and other variants such as secondary client, host client, distributed client and the like. The client may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other clients, servers, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the client. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the client.

[00293] The client may provide an interface to other devices including, without limitation, servers, other clients, printers, database servers, print servers, file servers, communication servers, distributed servers and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the client through an interface may include at least one storage medium capable of storing methods, programs, applications, code and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[00294] The methods and systems described herein may be deployed in part or in whole through network infrastructures. The network infrastructure may include elements such as computing devices, servers, routers, hubs, firewalls, clients, personal computers, communication devices, routing devices and other active and passive devices, modules and/or components as known in the art. The computing and/or non-computing device(s) associated with the network infrastructure may include, apart from other components, a storage medium such as flash memory, buffer, stack, RAM, ROM and the like. The processes, methods, program codes, instructions described herein and elsewhere may be executed by one or more

of the network infrastructural elements. The methods and systems described herein may be adapted for use with any kind of private, community, or hybrid cloud computing network or cloud computing environment, including those which involve features of software as a service (SaaS), platform as a service (PaaS), and/or infrastructure as a service (IaaS).

[00295] The methods, program codes, and instructions described herein and elsewhere may be implemented on a cellular network has sender-controlled contact media content item multiple cells. The cellular network may either be frequency division multiple access (FDMA) network or code division multiple access (CDMA) network. The cellular network may include mobile devices, cell sites, base stations, repeaters, antennas, towers, and the like. The cell network may be a GSM, GPRS, 3G, EVDO, mesh, or other networks types.

[00296] The methods, program codes, and instructions described herein and elsewhere may be implemented on or through mobile devices. The mobile devices may include navigation devices, cell phones, mobile phones, mobile personal digital assistants, laptops, palmtops, netbooks, pagers, electronic books readers, music players and the like. These devices may include, apart from other components, a storage medium such as a flash memory, buffer, RAM, ROM and one or more computing devices. The computing devices associated with mobile devices may be enabled to execute program codes, methods, and instructions stored thereon. Alternatively, the mobile devices may be configured to execute instructions in collaboration with other devices. The mobile devices may communicate with base stations interfaced with servers and configured to execute program codes. The mobile devices may communicate on a peer-to-peer network, mesh network, or other communications network. The program code may be stored on the storage medium associated with the server and executed by a computing device embedded within the server. The base station may include a computing device and a storage medium. The storage device may store program codes and instructions executed by the computing devices associated with the base station.

[00297] The computer software, program codes, and/or instructions may be stored and/or accessed on machine readable media that may include: computer components, devices, and recording media that retain digital data used for computing for some interval of time; semiconductor storage known as random access memory (RAM); mass storage typically for more permanent storage, such as optical discs, forms of magnetic storage like hard disks, tapes, drums, cards and other types; processor registers, cache memory, volatile memory, non-volatile memory; optical storage such as CD, DVD; removable media such as flash memory (e.g. USB sticks or keys), floppy disks, magnetic tape, paper tape, punch cards,

standalone RAM disks, Zip drives, removable mass storage, off-line, and the like; other computer memory such as dynamic memory, static memory, read/write storage, mutable storage, read only, random access, sequential access, location addressable, file addressable, content addressable, network attached storage, storage area network, bar codes, magnetic ink, and the like.

[00298] The methods and systems described herein may transform physical and/or intangible items from one state to another. The methods and systems described herein may also transform data representing physical and/or intangible items from one state to another.

[00299] The elements described and depicted herein, including in flow charts and block diagrams throughout the figures, imply logical boundaries between the elements. However, according to software or hardware engineering practices, the depicted elements and the functions thereof may be implemented on machines through computer executable media has sender-controlled contact media content item a processor capable of executing program instructions stored thereon as a monolithic software structure, as standalone software modules, or as modules that employ external routines, code, services, and so forth, or any combination of these, and all such implementations may be within the scope of the present disclosure. Examples of such machines may include, but may not be limited to, personal digital assistants, laptops, personal computers, mobile phones, other handheld computing devices, medical equipment, wired or wireless communication devices, transducers, chips, calculators, satellites, tablet PCs, electronic books, gadgets, electronic devices, devices has sender-controlled contact media content item artificial intelligence, computing devices, networking equipment, servers, routers and the like. Furthermore, the elements depicted in the flow chart and block diagrams or any other logical component may be implemented on a machine capable of executing program instructions. Thus, while the foregoing drawings and descriptions set forth functional aspects of the disclosed systems, no particular arrangement of software for implementing these functional aspects should be inferred from these descriptions unless explicitly stated or otherwise clear from the context. Similarly, it will be appreciated that the various steps identified and described above may be varied, and that the order of steps may be adapted to particular applications of the techniques disclosed herein. All such variations and modifications are intended to fall within the scope of this disclosure. As such, the depiction and/or description of an order for various steps should not be understood to require a particular order of execution for those steps, unless required by a particular application, or explicitly stated or otherwise clear from the context.

[00300] The methods and/or processes described above, and steps associated therewith, may be realized in hardware, software or any combination of hardware and software suitable for a particular application. The hardware may include a general- purpose computer and/or dedicated computing device or specific computing device or particular aspect or component of a specific computing device. The processes may be realized in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable device, along with internal and/or external memory. The processes may also, or instead, be embodied in an application specific integrated circuit, a programmable gate array, programmable array logic, or any other device or combination of devices that may be configured to process electronic signals. It will further be appreciated that one or more of the processes may be realized as a computer executable code capable of being executed on a machine-readable medium.

[00301] The computer executable code may be created using a structured programming language such as C, an object oriented programming language such as C++, or any other high-level or low-level programming language (including assembly languages, hardware description languages, and database programming languages and technologies) that may be stored, compiled or interpreted to run on one of the above devices, as well as heterogeneous combinations of processors, processor architectures, or combinations of different hardware and software, or any other machine capable of executing program instructions.

[00302] Thus, in one aspect, methods described above and combinations thereof may be embodied in computer executable code that, when executing on one or more computing devices, performs the steps thereof. In another aspect, the methods may be embodied in systems that perform the steps thereof, and may be distributed across devices in a number of ways, or all of the functionality may be integrated into a dedicated, standalone device or other hardware. In another aspect, the means for performing the steps associated with the processes described above may include any of the hardware and/or software described above. All such permutations and combinations are intended to fall within the scope of the present disclosure.

[00303] While the disclosure has been disclosed in connection with the preferred embodiments shown and described in detail, various modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the present disclosure is not to be limited by the foregoing examples, but is to be understood in the broadest sense allowable by law.

[00304] The use of the terms "a" and "an" and "the" and similar referents in the context of describing the disclosure (especially in the context of the following claims) is to be construed

to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms "comprising," "has sender-controlled contact media content item," "including," and "containing" are to be construed as open-ended terms (i.e., meaning "including, but not limited to,") unless otherwise noted. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., "such as") provided herein, is intended merely to better illuminate the disclosure and does not pose a limitation on the scope of the disclosure unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the disclosure.

[00305] While the foregoing written description enables one of ordinary skill to make and use what is considered presently to be the best mode thereof, those of ordinary skill will understand and appreciate the existence of variations, combinations, and equivalents of the specific embodiment, method, and examples herein. The disclosure should therefore not be limited by the above described embodiment, method, and examples, but by all embodiments and methods within the scope and spirit of the disclosure.

[00306] All documents referenced herein are hereby incorporated by reference.

CLAIMS

What is claimed is:

1. A method comprising:

retrieving a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment;

receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to at least one of the non-spatial attributes; and
matching the plurality of received sensor readings to at least one location in the map to produce a determined sensor location.

2. The method of claim 1, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of the non-spatial sensor readings.

3. The method of claim 2, wherein the interpolation is non-linear.

4. The method of claim 1, wherein the received plurality of sensor readings correspond to a generally stationary position of the device.

5. The method of claim 1, wherein the received plurality of sensor readings correspond to a movement of the device through the environment.

6. The method of claim 1, wherein the matching comprises utilizing a plurality of lower resolution sensor readings to determine an approximate sensor location and utilizing a plurality of higher resolution sensor readings to determine the determined sensor location.

7. The method of claim 1, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.

8. A system comprising:

a device comprising at least one sensor the device adapted to receive a plurality of sensor readings within an environment wherein each of the sensor readings corresponds to at least one non-spatial attribute; and

a processor adapted to retrieve a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment and further adapted to match the plurality of received sensor readings to at least one location in the map to produce a determined sensor location.

9. The system of claim 8, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of the non-spatial sensor readings.
10. The system of claim 9, wherein the interpolation is non-linear.
11. The system of claim 8, wherein the received plurality of sensor readings correspond to a generally stationary position of the device.
12. The system of claim 8, wherein the received plurality of sensor readings correspond to a movement of the device through the environment.
13. The system of claim 8, wherein the matching comprises utilizing a plurality of lower resolution sensor readings to determine an approximate sensor location and utilizing a plurality of higher resolution sensor readings to determine the determined sensor location.
14. The system of claim 8, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
15. A method comprising:
 - accessing a map comprising a point cloud of an environment comprising a plurality of non-spatial attribute values;
 - receiving, for each of a plurality of non-spatial attributes, a plurality of non-spatial sensor readings within the environment; and
 - referencing the received plurality of non-spatial sensor readings to the map resulting in a referenced heat map.
16. The method of claim 15, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of the non-spatial sensor readings.
17. The method of claim 16, wherein the interpolation is non-linear.
18. The method of claim 15, wherein the received plurality of sensor readings correspond to a generally stationary position of a device.
19. The method of claim 15, wherein the received plurality of sensor readings correspond to a movement of a device through the environment.
20. The method of claim 15, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
21. The method of claim 15, wherein the referenced heat map is a georeferenced heat map.

22. A system comprising:
a device adapted to produce a map comprising a point cloud of an environment comprising a plurality of non-spatial attribute values and further adapted to measure, for each of a plurality of non-spatial attributes, a plurality of non-spatial sensor readings; and
a processor adapted to reference the measured plurality of non-spatial sensor readings to the map resulting in a referenced heat map.
23. The system of claim 22, wherein the referenced heat map comprises both non-spatial sensor readings and interpolated non-spatial sensor readings.
24. The system of claim 22, wherein a plurality of the non-spatial attributes are signal strengths.
25. The system of claim 22, wherein the referenced heat map is a georeferenced heat map.
26. The system of claim 22, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of a plurality of non-spatial sensor readings.
27. The system of claim 26, wherein the interpolation is non-linear.
28. The system of claim 22, wherein the measured plurality of sensor readings correspond to a generally stationary position of the device.
29. The system of claim 22, wherein the measured plurality of sensor readings correspond to a movement of the device through the environment.
30. The system of claim 22, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
31. The system of claim 22, wherein the referenced heat map is a georeferenced heat map.
32. A method comprising:
accessing a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters;
receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to the non-spatial attribute;
creating a heat map of the sensor readings referenced to the map; and
generating a recommended position of at least one signal emitter based, at least in part, on the heat map.
33. The method of claim 32, wherein the recommended position of the at least one signal emitter results in a more uniform heat map.

34. The method of claim 32, wherein a plurality of the non-spatial attributes are signal strengths.
35. The method of claim 32, wherein the heat map is a georeferenced heat map.
36. The method of claim 32, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of a plurality of non-spatial sensor readings.
37. The method of claim 36, wherein the interpolation is non-linear.
38. The method of claim 32, wherein the measured plurality of sensor readings correspond to a generally stationary position of the device.
39. The method of claim 32, wherein the measured plurality of sensor readings correspond to a movement of the device through the environment.
40. The method of claim 32, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
41. A system comprising:
a device adapted to produce a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters and receive a plurality of sensor readings within the environment wherein each of the sensor readings corresponds to the non-spatial attribute; and
a processor adapted to create a heat map of the sensor readings referenced to the map and generate a recommended position of at least one signal emitter based, at least in part, on the heat map.
42. The system of claim 41, wherein the recommended position of the at least one signal emitter results in a more uniform heat map.
43. The system of claim 41, wherein a plurality of the non-spatial attributes are signal strengths.
44. The system of claim 41, wherein the heat map is a georeferenced heat map.
45. The system of claim 41, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of a plurality of non-spatial sensor readings.
46. The system of claim 45, wherein the interpolation is non-linear.
47. The system of claim 41, wherein the plurality of sensor readings correspond to a generally stationary position of the device.

48. The system of claim 41, wherein the plurality of sensor readings correspond to a movement of the device through the environment.
49. The system of claim 41, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
50. A method comprising:
retrieving a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to one of a plurality of non-spatial attributes and indicative of a plurality of non-spatial sensor readings acquired throughout the environment;
receiving a plurality of sensor readings from a device within the environment wherein each of the sensor readings corresponds to at least one of the non-spatial attributes;
matching the plurality of received sensor readings to at least one location in the map to produce a determined sensor location; and
retrieving additional sensor data from the determined sensor location and comparing the additional sensor data to data derived from the map.
51. The method of claim 50, further comprising updating the map data with the retrieved additional sensor data.
52. The method of claim 50, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of the non-spatial sensor readings.
53. The method of claim 52, wherein the interpolation is non-linear.
54. The method of claim 50, wherein the received plurality of sensor readings correspond to a generally stationary position of the device.
55. The method of claim 50, wherein the received plurality of sensor readings correspond to a movement of the device through the environment.
56. The method of claim 50, wherein the matching comprises utilizing a plurality of lower resolution sensor readings to determine an approximate sensor location and utilizing a plurality of higher resolution sensor readings to determine the determined sensor location.
57. The method of claim 50, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.
58. A system comprising:
a device adapted to produce a map of a 3D geometry of an environment the map comprising a plurality of non-spatial attribute values each corresponding to a non-spatial attribute

and indicative of a plurality of non-spatial sensor readings acquired throughout the environment from a plurality of signal emitters and receive a plurality of sensor readings within the environment wherein each of the sensor readings corresponds to the non-spatial attribute; and

a processor adapted to match the plurality of received sensor readings to at least one location in the map to produce a determined sensor location and retrieving additional sensor data from the determined sensor location and compare the additional sensor data to data derived from the map.

59. The system of claim 58, wherein the map data is updated with the retrieved additional sensor data.
60. The system of claim 58, wherein a portion of the plurality of non-spatial attribute values are determined via an interpolation of the non-spatial sensor readings.
61. The system of claim 60, wherein the interpolation is non-linear.
62. The system of claim 58, wherein the received plurality of sensor readings correspond to a generally stationary position of the device.
63. The system of claim 58, wherein the received plurality of sensor readings correspond to a movement of the device through the environment.
64. The system of claim 58, wherein the matching comprises utilizing a plurality of lower resolution sensor readings to determine an approximate sensor location and utilizing a plurality of higher resolution sensor readings to determine the determined sensor location.
65. The system of claim 58, wherein each of the sensor readings is derived from a sensor comprising at least one of LIDAR, a camera, a depth camera, an infrared camera, an IMU, an ultrasonic sensor, a Wi-Fi sensor, a Bluetooth sensor, a temperature sensor, and a barometer.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6





FIG. 7A






FIG. 8A




FIG. 8B




FIG. 9A




FIG. 9B




FIG. 10




FIG. 11

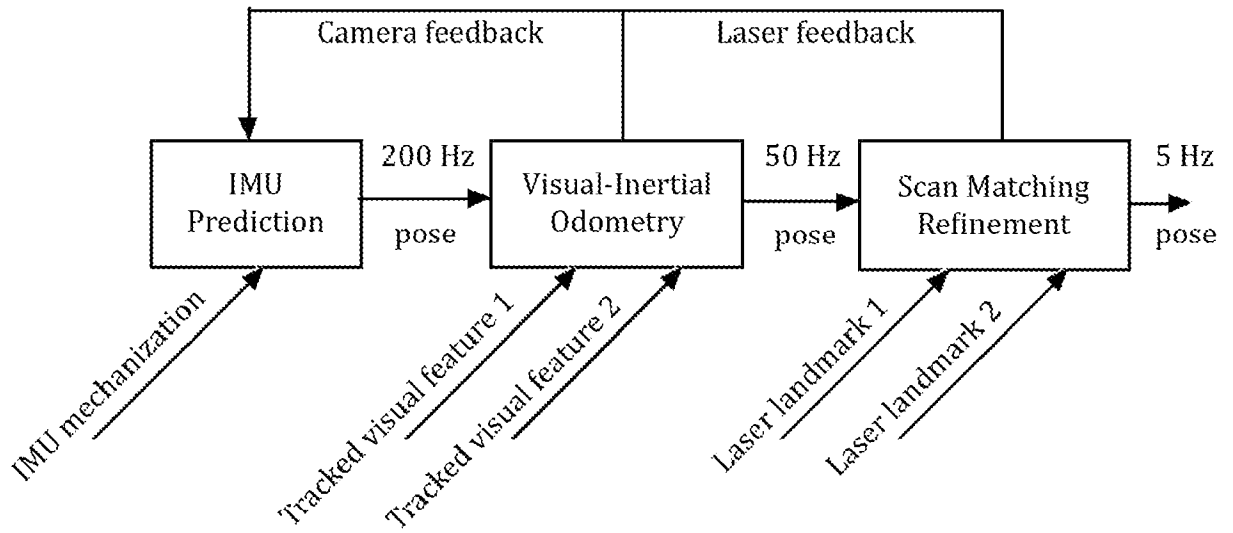


FIG. 12

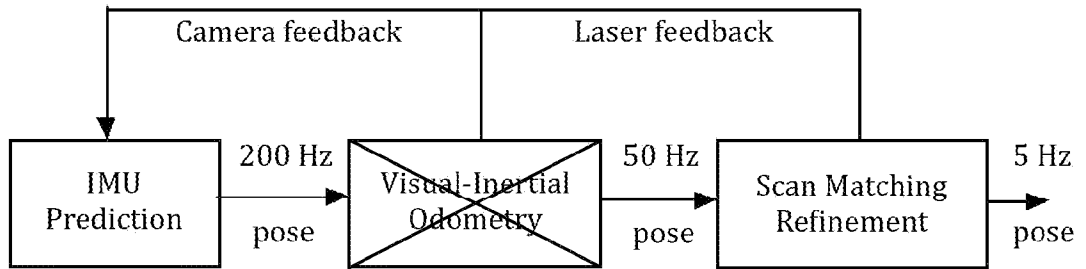


FIG. 13a

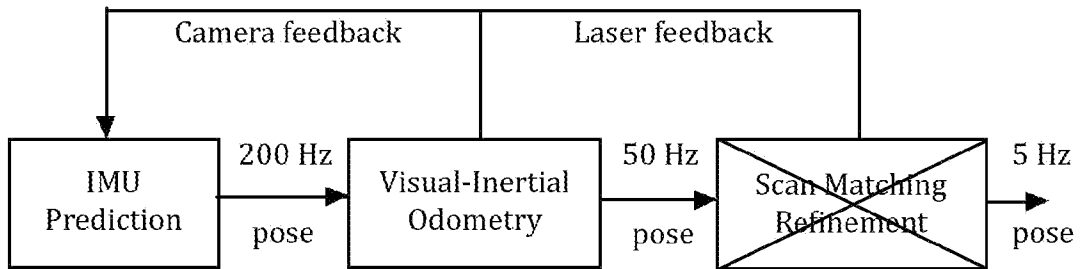


FIG. 13b

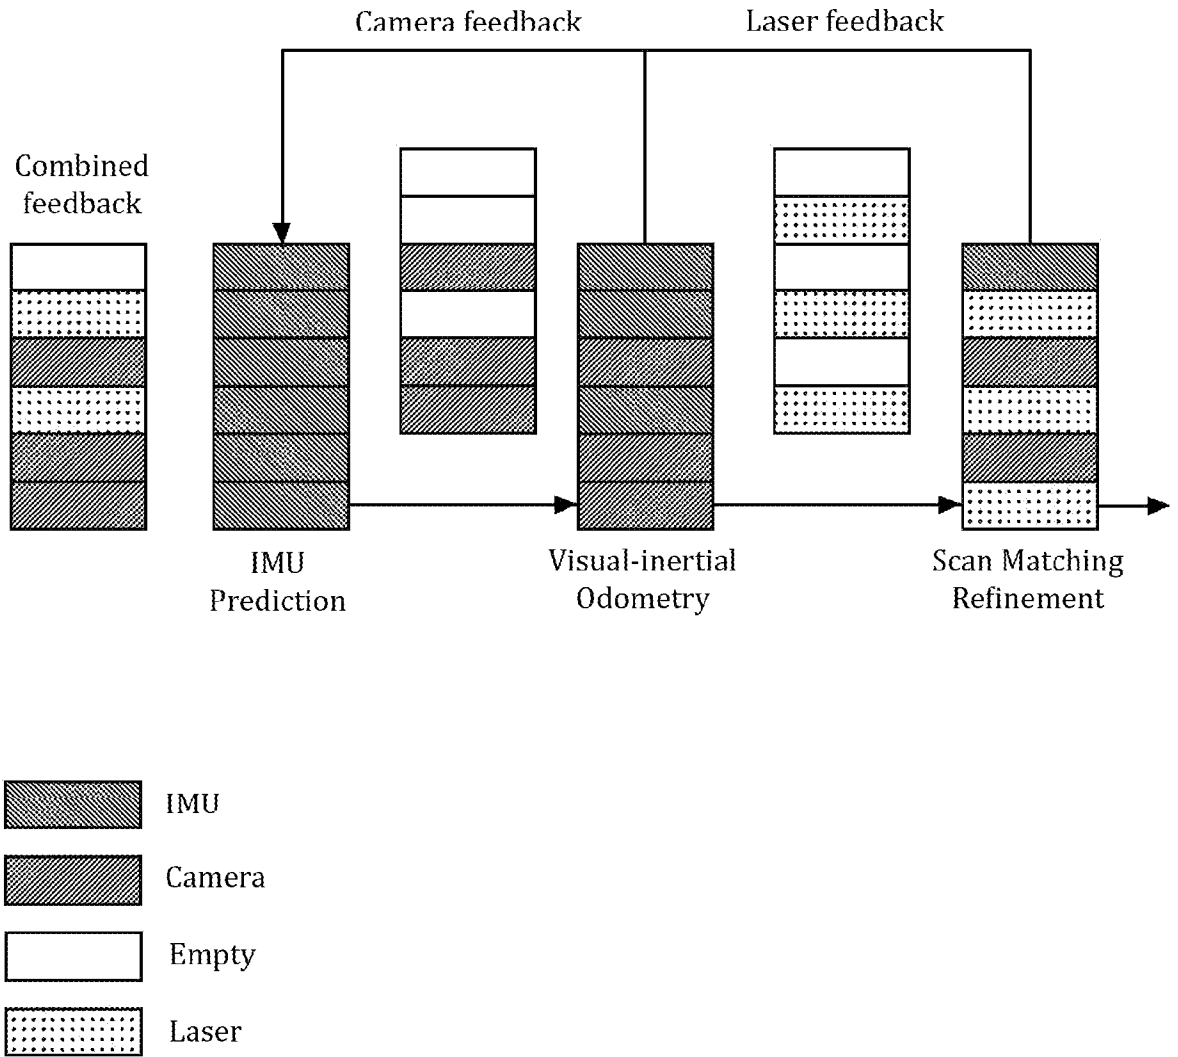


FIG. 14




FIG. 15a




FIG. 15b

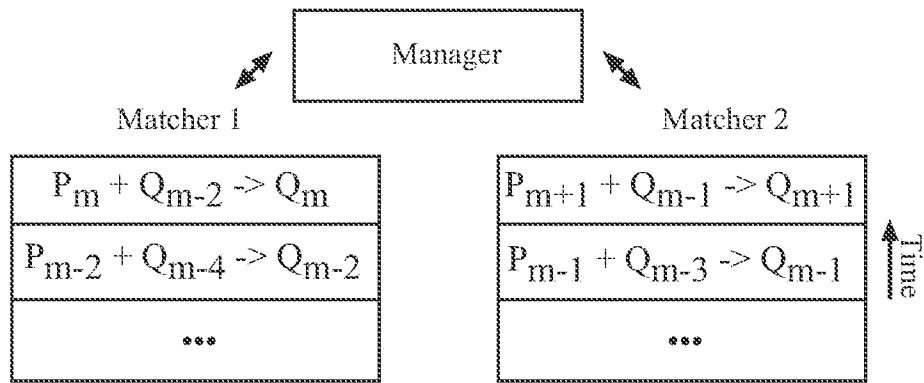


FIG. 16A

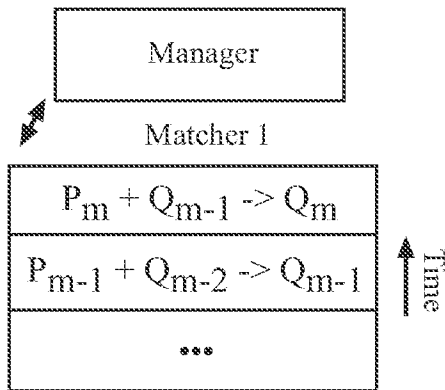


FIG. 16B




FIG. 17A




FIG. 17B




FIG. 18

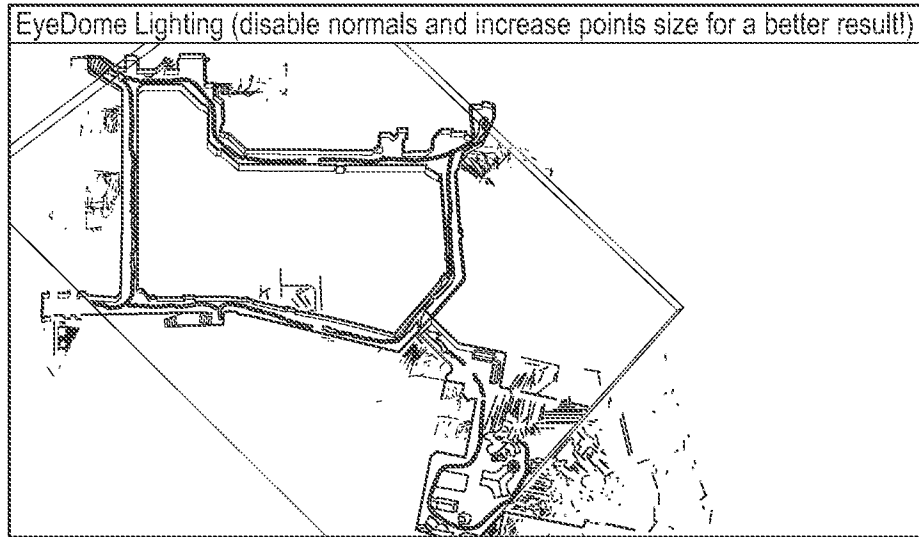


FIG. 19A




FIG. 19B




FIG. 19C

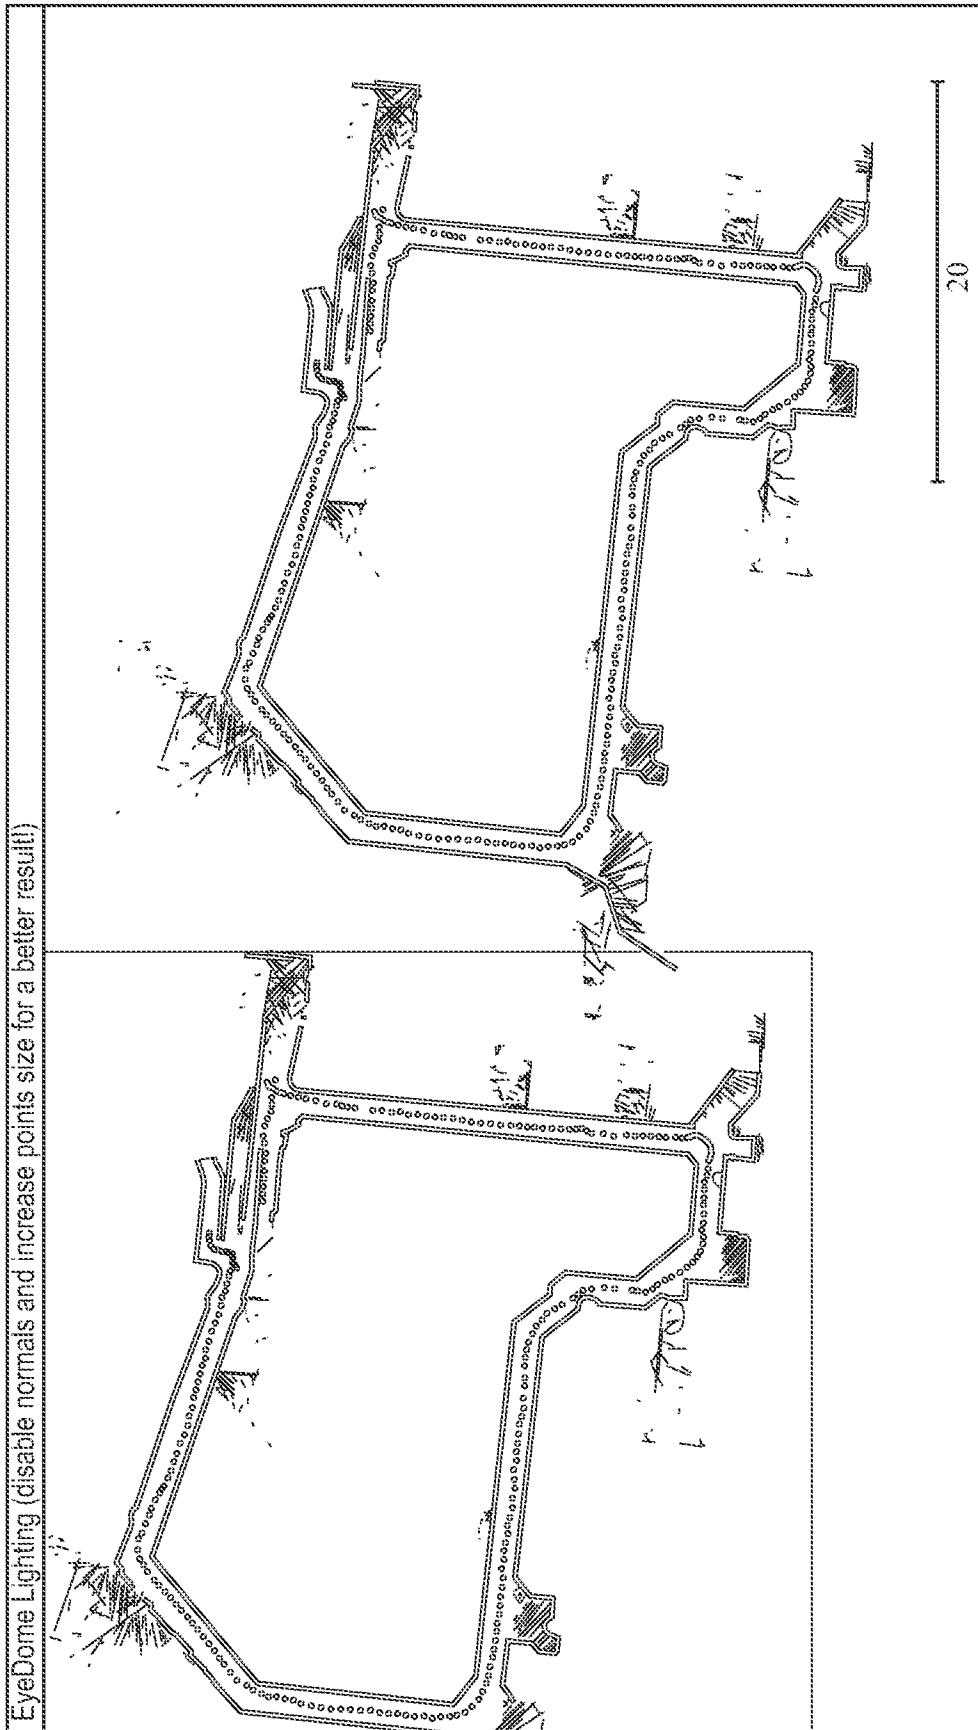


FIG. 20




FIG. 21




FIG. 22




FIG. 23






FIG. 24

22/37



(a)



(b)

FIG. 25




FIG. 26

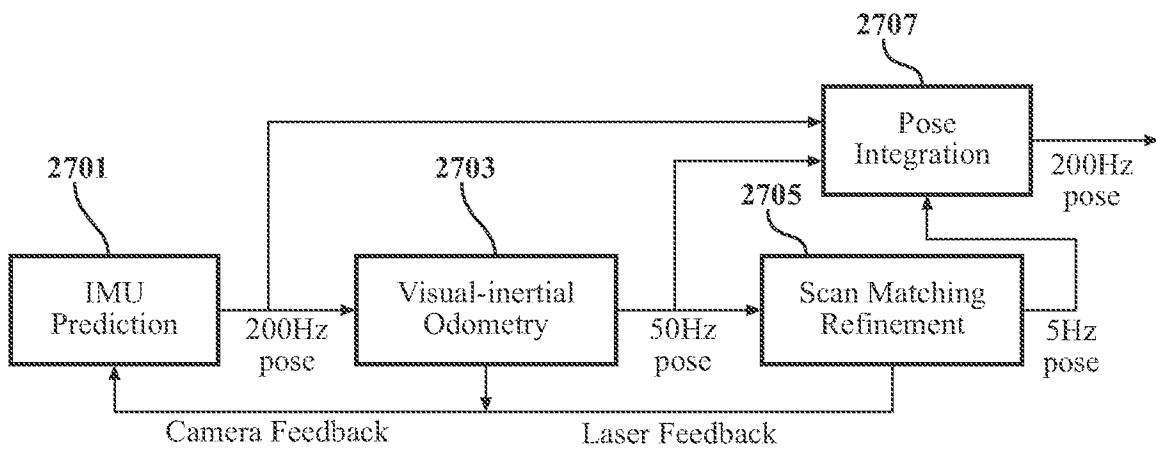




FIG. 27



(a)



(b)

FIG. 28






FIG. 29

27/37




(a)




(b)

FIG. 30

28/37



(a)




(b)

FIG. 31



(a)



(b)

FIG. 32




FIG. 33

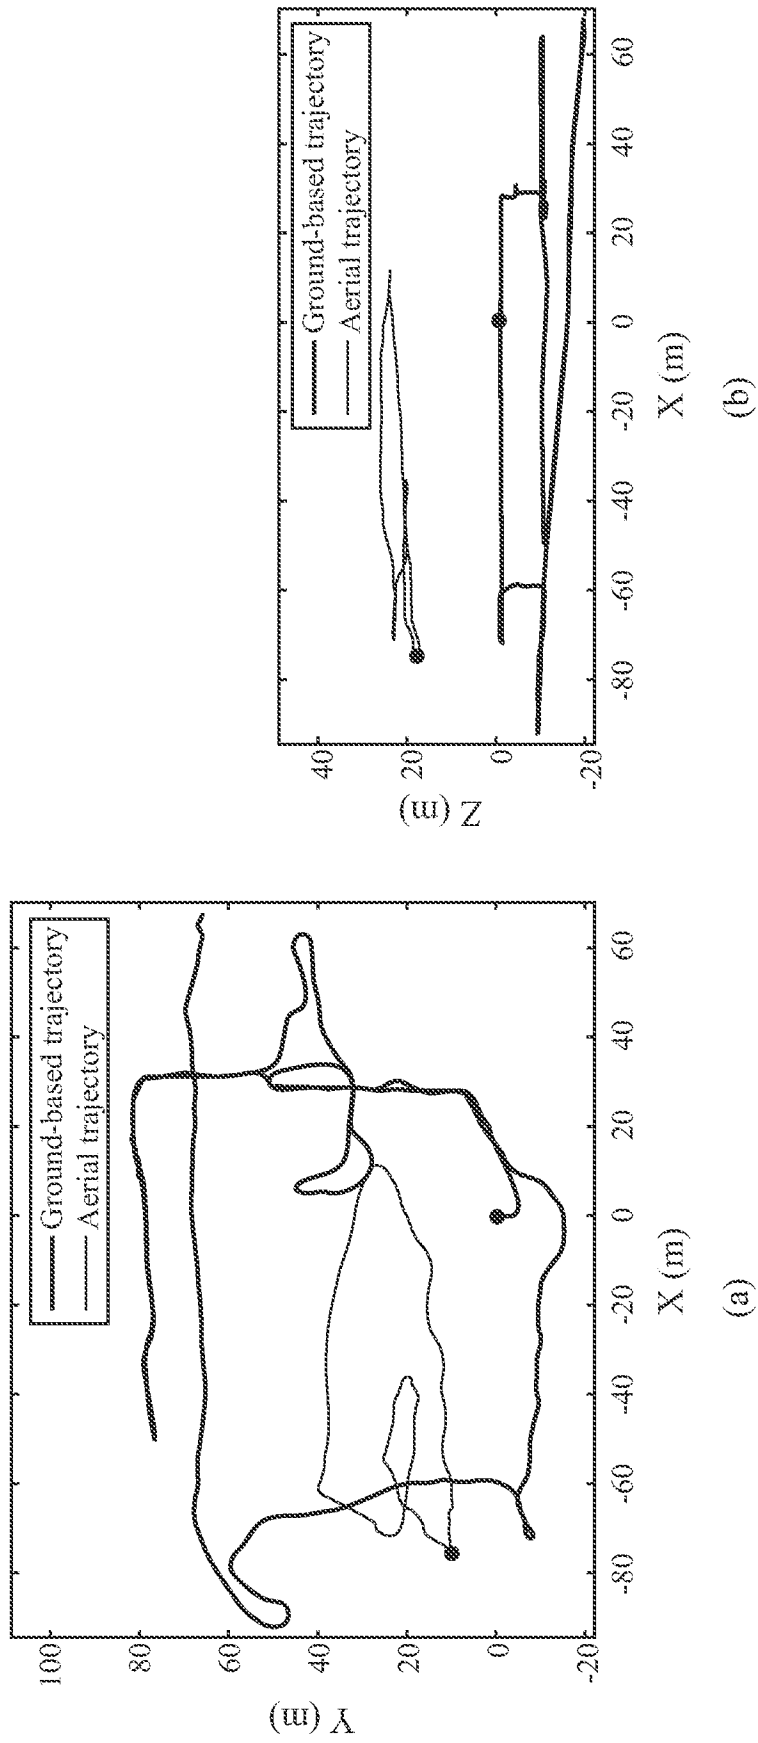



FIG. 34

32/37



(a)




(b)

FIG. 35 (1 of 2)




(c)

(d)




(e)

FIG. 35 (2 of 2)




(a)




(b)

364

FIG. 36 (1 of 2)



(c) (d)



(e)

FIG. 36 (2 of 2)

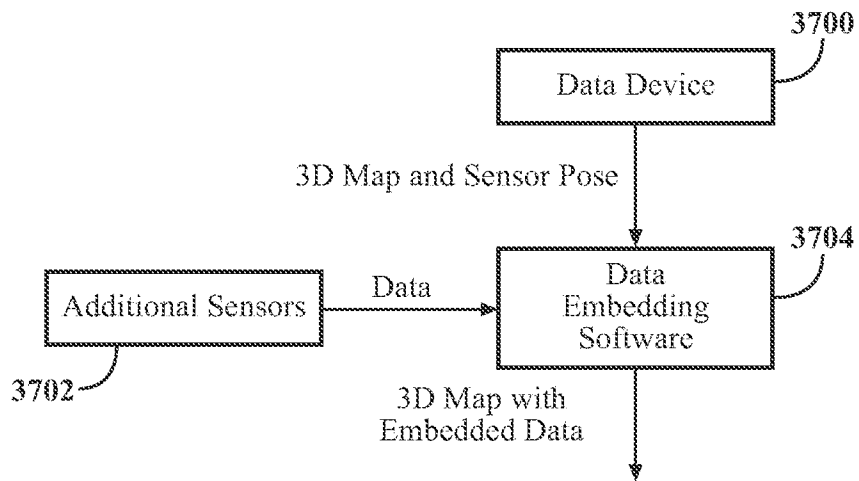


FIG. 37A

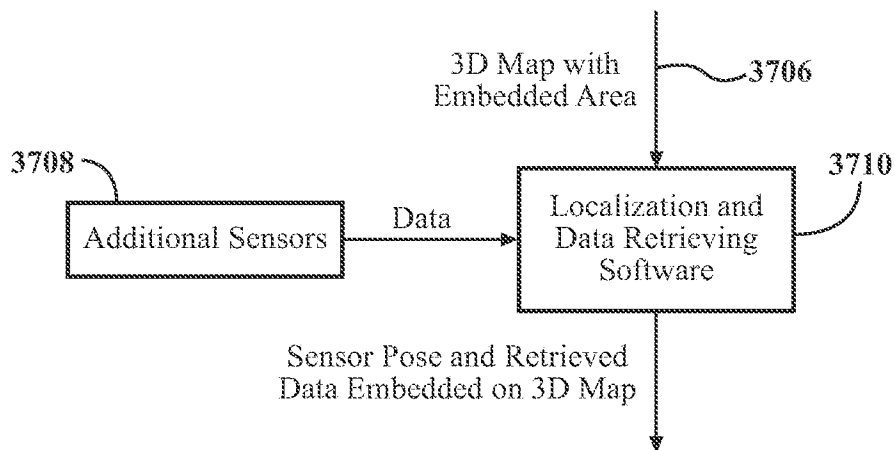


FIG. 37B

3800




FIG. 38

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US18/42346

A. CLASSIFICATION OF SUBJECT MATTER

IPC - H04W 4/02 (2018.01)

CPC - H04W 4/023, 4/025; G01C 21/005

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

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

See Search History document

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/0193040 A1 (OMIIMII LTD.) 10 July 2014; abstract; paragraphs [0048], [0049], [0090], [0105]-[0107], [0115]	1, 4, 5, 7, 8, 11, 12, 14, 15, 18-22, 24, 25, 28-31, 50, 51, 54, 55, 57-59, 62, 63, 65
---		---
Y		2, 3, 6, 9, 10, 13, 16, 17, 23, 26, 27, 32-49, 52, 53, 56, 60, 61, 64
Y	US 2013/0085378 A1 (IMRICOR MEDICAL SYSTEMS, INC.) 4 April 2013; paragraphs [0033], [0036], [0043]	2, 3, 9, 10, 16, 17, 23, 26, 27, 36, 37, 45, 46, 52, 53, 60, 61
Y	US 2010/0090899 A1 (ZHAO, J et al.) 15 April 2010; abstract; paragraphs [0017], [0043]	6, 13, 56, 64
Y	US 2014/0297092 A1 (TOYOTA MOTOR ENGINEERING & MANUFACTURING NORTH AMERICA, INC.) 2 October 2014; abstract; paragraph [0024]	32-49

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

13 September 2018 (13.09.2018)

Date of mailing of the international search report

01 OCT 2018

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774