



(51) International Patent Classification:
B61L 15/00 (2006.01) *G06Q 10/06* (2012.01)
B61L 17/00 (2006.01)

(21) International Application Number:
PCT/US2016/034715

(22) International Filing Date:
27 May 2016 (27.05.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/167,015 27 May 2015 (27.05.2015) US
62/244,543 21 October 2015 (21.10.2015) US

(71) Applicant: AMSTED RAIL COMPANY, INC. [US/US];
311 South Wacker Drive, Suite 5300, Chicago, Illinois
60606 (US).

(72) Inventors: LEFEBVRE, William; 1145 Saint Finegan
Drive, West Chester, Pennsylvania 19382 (US). BONNES,
Matthew; 4 Woodview Road, Malvern, Pennsylvania
19355 (US). DRAGISH, Darren; 22 Downing Circle,
Downingtown, Pennsylvania 19335 (US). MARTIN, An-
drew; 926 Pineview Drive, West Chester, Pennsylvania
19380 (US). FUHS, Thomas P.; 13 Jonathan Dr.,
Phoenixville, Pennsylvania 19460 (US).

(74) Agents: HECHT, Gary A. et al.; Fox Rothschild LLP,
997 Lenox Drive, Bldg. #3, Lawrenceville, New Jersey
08648 (US).

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

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

Published:

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

(54) Title: SYSTEM AND METHOD FOR BUILDING AND MANAGING A TRAIN CONSIST




Fig. 1

(57) Abstract: Railyard management system for managing, assembling, disassembling and verifying train consists and monitoring railcars in the railyard. The system provides for the collection of data and the movement of data from lower processing levels to higher processing levels, where an inference engine draws inferences regarding the current state of railcars and train consists within the railyard. The inferences are assigned confidence levels based on the methods and available data used to draw the inferences. The system can be used to track the location and orientation of railcars in the railyard and to verify order and orientation of assets in a train consist

System and Method for Building and Managing a Train Consist

Related Applications

[0001] This application claims the benefit of U.S. Provisional Patent Application Serial No. 62/167,015, filed May 27, 2015 and U.S. Provisional Patent Application Serial No. 62/244,543, filed October 21, 2015, which are incorporated herein by reference in their entireties.

Background of the Invention

[0002] It has become increasingly important for railway owners and operators to be able to locate and organize assets, including railcars, locomotives and train consists on a real time basis. From an operational point of view, it is important for railway operators to determine whether a railcar is located within or outside the boundaries of a railyard, is moving or stationary, and whether or not the railcar is part of a train consist or not linked to other railcars.

[0003] The knowledge of the status of railcars allows an operator to determine if railcars are being utilized or idle at any given point in time and provides means to help in the management of railyard operations.

[0004] As current industry practice, the management of train consists and railyards in railroad operations relies on reading, at fixed points in the rail network, passive radio frequency identification (RFID) tags which are affixed to each railcar. While this method provides railroad operators with check-in/check-out list of

assets, it lacks the benefits of a dynamic wireless network capable of transmitting timely information, such as location, status, condition, and/or performance data when not in range of an RFID reader. Additionally, the information typically encoded into an RFID tag is static and therefore, the RFID tag is not capable of providing the current status of the railcar. Further, currently systems do not provide a mechanism to validate a train consist before it leaves the railyard. Mistakes are possible when a train consist is created, and the result of such mistakes can be missing, incorrect or extra railcars in the train consist. There is also a safety risk that can be associated with using human intervention to visually validate a train consist before it departs a railyard.

[0005] It is therefore desirable to provide a train consist management system in a railyard to ease the management of creating and validating train consists. It is intended to eliminate mistakes and to mitigate the safety risks to humans carrying out the manual process of the current systems. Additionally, automating the process improves the efficiency of the management of the railyard, thereby reducing costs.

[0006] Given the demanding and harsh environments in which railroad trains operate, any monitoring system must be rugged, reliable and able to operate for long periods with little or no maintenance. Because there are more than 1.5 million freight railcars in North America alone, and many millions more around the world, a system of monitoring all railcars, both in use and idle in a railyard, is highly desirable and, as such, the system needs to be scalable to handle a very large number of potential devices.

[0007] Train/Rail communication and sensor systems are disclosed in U.S. patent 7,688,218 issued March 30, 2010, U.S. patent 9,026,281 issued May 5, 2015, U.S. patent publication 2013/0342362 published December 26, 2013, PCT application PCT/US2014/067739 filed November 26, 2014, and PCT application PCT/US2014/072380 filed December 24, 2014, the full disclosures of all of these are incorporated herein by reference.

Summary of the Invention

[0008] It is an objective of this invention to provide a comprehensive system which allows the collection of data and the analysis of that data to perform one or more of the following functions:

- detect the presence of railcars within a railyard;
- determine the location and orientation of railcars in the railyard;
- logically monitor the assembly of train consists;
- determine the order and orientation of railcars in a train consist
- validate the order of railcars in a train consist and the orientation of railcars within a train consist
- provide adequate warnings when the railcar order of a train consist is incorrect thus allowing for intervention by humans or automated systems before an operational failure occurs; and
- provide an analysis capability to determine the severity and priority of events and warnings at different levels of processing.

- Determine operational status of railcars in the railyard (loaded, unloaded, handbrake applied, etc.)

- [0009] In one preferred embodiment, and with reference to Fig. 1, the present invention consists of a system and method for building and managing a train consist, and includes the following:
- [0010] A train-based mesh network system **107** using a wireless mesh network to provide bi-directional communication from freight railcars **103(a) or 103(b)** in the train consist **109** to a host or control point.
- [0011] A Powered Wireless Gateway device (PWG) **102** to manage the train-based mesh network **107** and communicate events from individual railcars **103(a) or 103(b)** to the locomotive engineer or to other train management systems.
- [0012] A Powered Wireless Gateway device **102** capable of receiving multiple sensor events from individual railcars and making an inference about the order of the railcars in a train consist **109**.
- [0013] A Powered Wireless Gateway device **102** capable of receiving information from an external control center or data system that specifies the freight railcars **103(a) or 103(b)** that should be in the train consist **109** allowing only those railcars **103(a) or 103(b)** to join and reporting any railcars **103(a) or 103(b)** that are absent.

- [0014] A Communication Management Unit (CMU) **101** on each railcar **103** capable of being a wireless node in the train-based mesh network **107** and being able to send messages to a host or control point.
- [0015] A Communication Management Unit **101** on each railcar capable of using built-in sensors and/or managing a wireless sensor node **104** network on the freight railcar **103** to generate messages that need to be sent to locomotive host or control point.
- [0016] A Communication Management Unit **101** on each railcar **103** capable of supporting a global navigation satellite system (GNSS) sensor to determine location, direction or speed of the freight railcar **103**.
- [0017] A Communication Management Unit 101 on each railcar 103 capable of using a compass.
- [0018] A Communication Management Unit 101 on each railcar 103 capable of using a motion sensor.
- [0019] A Communication Management Unit 101 on each railcar 103 capable of using one or more accelerometers for impact detection.
- [0020] A Communication Management Unit 101 on each railcar 103 capable of using one or more accelerometers for motion sensing.
- [0021] A Communication Management Unit 101 on each railcar 103 capable of supporting one or multiple geo-fences.
- [0022] A Communication Management Unit 101 on each railcar 103 capable of indicating presence of an RFID reader.

- [0023] A Communication Management Unit 101 on each railcar 103 capable of determining presence of mesh network and signal strength.
- [0024] A Wireless Sensor Node 104 containing a temperature sensor and an accelerometer.
- [0025] A Wireless Sensor Node (WSN) containing a motion sensor.
- [0026] A Wireless Sensor Node **104** containing other sensors.
- [0027] A managed railyard or unmanaged location with one or more Powered Wireless Gateway(s) 102 present.
- [0028] A train consist 109 where a train consist is defined as a connected group of railcars 103 and locomotives 108 that form a complete train.
- [0029] The train-based mesh network system 107 used to build and manage a train consist also can be used for event and alert transmission, both during the formation of the train consist 109 (to a control center), as well as after it is complete (to the control center or locomotive 108).

Brief Description of the Drawings

- [0030] Fig. 1 is a diagram illustrating a train consist monitoring system and related hardware components.
- [0031] Fig. 2 is a flowchart illustrating a method of determining the location and orientation of a railcar in a railyard in relation to the rail.

- [0032] Fig. 3 is a flowchart illustrating a method of determining whether a railcar is in a railyard.
- [0033] Fig. 4 is a diagram illustrating how railcars can be linked so that a train consist can be formed.
- [0034] Fig. 5 is a diagram illustrating how data flows from a wireless sensor node, a communication management unit, a powered wireless gateway and to a control center.
- [0035] Fig. 6 is a flowchart illustrating how messages are transmitted based on message priority.
- [0036] Fig. 7 is a diagram illustrating a railyard in which the direction of the railyard is known to be running southwest to northeast with enlargement of railcar showing how the B-end of a railcar with CMU installed can be determined based on the heading of the CMU compared to North.
- [0037] Fig. 8 is a diagram illustrating how to determine if two railcars are on the same rail track or not.
- [0038] Fig. 9 is a diagram illustrating how monitored railcars, not within the presence of a PWG (either in a managed railyard or as part of a managed train consist) can be recognized by a passing locomotive upon which a powered wireless gateway is installed.
- [0039] Fig. 10 shows examples of probability curves for two exemplary sensors.

- [0040] Fig. 11 is a specific example of the use of probability curves for determining the likelihood that two or more railcars are likely to be linked.
- [0041] Fig. 12 shows examples of the use of historical data in lieu of probabilities to determine if two or more railcars are likely to be linked.
- [0042] Fig. 13 is a flow chart showing the process for determining if a coupling event has occurred.

Definitions

- [0043] A train consist, shown in the drawings as reference number 109, is defined as a connected group of railcars and locomotives.
- [0044] A link, shown for example in Fig. 4, is defined as two or more railcars coupled together.
- [0045] A computing device is defined as any machine capable of processing and executing software to perform calculations or otherwise provide functionality. The computing device shall also have data storage and network communication capabilities to perform the functions required by this invention. A computing device includes, but is not limited to, a server, PC, or PWG 102, as described in this document.
- [0046] A manager is defined as any device that is capable of linking together nodes in a mesh network on a time synchronized schedule and maintaining that link schedule such that reliable bi-directional communication is possible between all nodes in the network and with the manager. The manager may also provide a user interface

to another network host for front end communication. A manager includes, but is not limited to, a PWG 102 or CMU 101, as described in this document.

- [0047] A node is defined as any device that is capable of bi-directional wireless communications with another device to transmit and receive data. A node includes, but is not limited to, a CMU 101 or WSN 104, as described in this document
- [0048] A sensor is defined as any device that detects or measures a physical property and records the result, or transmits a resulting signal. One or more sensors may be present on a PWG 102, CMU 101, or WSN 104, as described in this document
- [0049] A wireless sensor node ("WSN"), shown in the drawings as reference number 104, is typically located on a railcar 103(a) or 103(b), is deployed preferably in a self-contained, protective housing, and may include one or more sensors, a power source, circuitry to read the sensor(s) and convert the readings to a digital form, and communication circuitry which allows the WSN to wirelessly transmit the sensor readings to an external receiver. The wireless sensor nodes are used for sensing a parameter to be monitored (e.g. temperature of, for example, bearings or ambient air) or status (e.g., position of a hatch or hand brake). The WSN may also include an intelligence capability, implemented as software running on an embedded microprocessor to analyze the data and determine if the data needs to be transmitted immediately, held for later transmission, or aggregated into an alert. WSNs are typically a member of a wireless mesh network managed by either a CMU or a PWG.

[0050] A communications management unit (“CMU”), shown in the drawings as reference number 101, is typically located on a railcar 103 and optionally acts as a manager for the railcar-based wireless mesh network 105 overlaid on the railcar. The CMU hardware preferably includes a processor, a power source, for example, a battery, a global positioning system (“GPS”) receiver, Wi-Fi and/or cellular capability, a wireless communications capability for maintaining the mesh network, and, optionally, one or more sensors, such as, but not limited to, an accelerometer or temperature sensor. The CMU may support one or more WSNs in a mesh configuration using the IEEE 2.4 GHz 802.15.4 radio standard. Additionally, the CMU is also a member of either a train-based wireless mesh network, which consists of the CMUs from all enabled railcars in the train consist; controlled by a manager, preferably a powered wireless gateway (PWG), typically located on a powered locomotive; is a member of a railyard-based wireless mesh network, controlled by one or more managers, preferably powered wireless gateways dispersed throughout the railyard; or operating independently outside of a wireless mesh network. The CMU thus supports at least four functions: 1) to support built-in sensors, such as an accelerometer, within the CMU to monitor specific attributes of the railcar such as location, speed, accelerations and more; and 2) to support bi-directional communication to the powered host or control point, such as a locomotive and/or an off-train monitoring and control center; 3) to consolidate data from built-in sensors, and/or any number of WSNs in the railcar-based wireless mesh network and to apply logic to the data gathered to generate warning alerts to a powered host such as a locomotive or remote control

center; and 4) to manage a low-power wireless mesh network overlaid on a railcar.

[0051] The CMU is capable of receiving data and/or alarms from one or more WSNs, or generating data and/or alarms directly, and is capable drawing inferences from this data or alarms regarding the performance of railcar 103, and of transmitting data and alarm information to a remote receiver. The CMU is preferably a single unit that would serve as a communications link to other locations, such as a mobile base station (e.g., the locomotive 108), a land-based base station, etc., and have the capability of processing the data received. The CMU also communicates with, controls and monitors WSNs (when present) in the local railcar-based wireless mesh network. Preferably, the placement of the CMU on each railcar will be consistent, as the placement will be useful in making determinations of the order and orientation of railcars within a train consist, as described later.

[0052] A powered wireless gateway (“PWG”), shown in the drawings as reference number 102, is preferably located either on a locomotive or deployed as part of a railyard-based wireless mesh network. It typically will include a processor, a GNSS receiver, a satellite and or cellular communication system, an Ethernet port and a high capacity network manager. The PWG will have power supplied by the locomotive, if located in the locomotive, or will derive its power from another source. The PWG acts as the manager of a wireless mesh network overlaid on a train consist (a train-based wireless mesh network, as define below), consisting of multiple CMUs from each railcar in a train, or is a member of a wireless mesh network overlaid on a railyard (a railyard-based mesh network, as defined below),

consisting of other PWGs and CMUs from individual railcars not currently associated with a train consist. PWGs can communicate and manage WSNs directly, without requiring the presence of a CMU. The PWG, if located on a powered asset, such as a locomotive 108, will derive power from the powered asset, or will derive its power from another source, for example, from a solar power generator or from a high capacity battery.

[0053] The PWG collects data and draws inferences regarding the performance of the train consist, as opposed to CMUs, which draw inferences regarding the performance of individual railcars.

[0054] A dark railcar is a railcar equipped with a CMU but which is not connected or associated with a train-based wireless network or a railyard-based wireless network, as defined below.

[0055] A railcar-based wireless mesh network shown in the drawings as reference number 105, consists of a CMU on a railcar 103, which is part of and manages a mesh network of a plurality of WSNs, each deployed, preferably, on the same railcar 103.

[0056] A train-based wireless mesh network, shown in the drawings as reference number 107, consists of a powered PWG 102 typically located on a locomotive 108 (but which may be on any moving asset in the train consist), which is part of and manages a mesh network of a plurality of CMUs, each deployed on a railcar, wherein the locomotive and plurality of railcars form a train consist.

- [0057] A railyard-based wireless mesh network, shown in the drawings as reference number 117, consists of one or more land-based, powered PWGs deployed at strategic locations in a railyard. The PWGs form a mesh network which includes one or more CMUs, each deployed on a railcar, and one or more mobile PWGs, each deployed on a powered asset, such as a locomotive, and may optionally include one or more WSNs located on railcars. Under certain circumstances, individual WSNs located on railcars may directly join the railyard-based (or train-based) mesh network, bypassing the CMU on the railcar, by directly communicating with the PWGs located in the railyard. The locomotives and railcars in the railyard-based mesh network are not associated with a train consist, but instead the PWGs, CMUs and, optionally, WSNs located on the railcar are nodes in the railyard-based mesh network.
- [0058] Building off of the IEC 62591 international wireless standard as well as the ISA100.11, a standard from the International Society of Automation, the railyard- and train-based wireless mesh network architectures are developed to these standards.
- [0059] A managed railyard is defined as a railyard having a railyard-based mesh network overlaid thereon.
- [0060] The discussion which follows describes the system in the context of a railcar, however, it will be understood by one of skill in the art that the same methods are applicable to any railroad vehicle or asset. It should also be noted that the definitions above are not meant to be exclusive, in that defined components may have additional components or features not included in the definition. Furthermore, while the description which follows features

a railcar with two trucks (or bogies), it is applicable to any configuration with more or less trucks or axles.

Detailed Description of the Invention

[0061] It is an object of the present invention to provide a train consist management system, where a railyard-based mesh network is overlaid on a railyard, and which includes one or more PWGs present in the railyard which act as communication points and aggregators of data generated and transmitted by the mesh networks of each railcar in the railyard. In addition, the PWGs in the railyard manage train consists and perform analysis of data from multiple monitored railcars and systems. When a railcar is not within a managed railyard, the same data transmission and analysis can be performed in the presence of a powered wireless gateway installed on a locomotive or other moving asset.

[0062] The present invention operates in an environment of a managed railyard, having a topology as shown in Fig. 1. Railcar 103 (shown as both 103(a) and 103(c) in Fig. 1) is typically equipped with multiple WSNs 104 placed at various positions on railcar 103. The positioning of individual WSNs 104 is dependent on the operational parameter(s) of the railcar 103 which are being monitored. CMU 101 is positioned on railcar 103 and forms a railcar-based mesh network 105 being managed by CMU 101 and having the WSNs 104 as nodes in the network. Preferably, CMUs 101 will be positioned and oriented in a consistent manner on each railcar 103. Also preferably, CMU 101 will be positioned toward one end of railcar 103 so as to be useful in determining the orientation of the car within the

train consist and at any location within the railyard. Optionally, railcar 103 may have only a CMU 101, and no WSNs 104, shown as 103(b) in Fig. 1 in which case there will be no railcar-based mesh network associated with that railcar.

[0063] Locomotive 108 is equipped with a PWG 102. PWG 102 also controls a train-based wireless mesh network 107 which is managed by PWG 102 and has CMUs 101 on each railcar in the train as nodes.

[0064] A railcar 103(d) not having a communication management unit 101 or WSNs 104 is considered an unmanaged railcar and is outside of the train-based mesh network 107.

[0065] The present invention also relates to a method of monitoring a railyard wherein, the location and orientation of the railcar within the railyard is determined by the method shown in Fig. 2, the presence of a railcar 103(a) or 103(b) within the railyard is determined by the method shown in Fig. 3, and the building of a train consist proceeds as shown in Fig. 4.

[0066] The order of a railcar in the train consist, the orientation of the railcars and/or the location of the railcar in the railyard may be determined via several methods, discussed below. The orientation of a railcar in the train consist is a critical element in the train consist. As is known in the industry, the ends of a railcar are identified as either “A” or “B”. Readings from a magnetometer or electronic compass and an accelerometer can be used to identify the orientation of the railcar. Additionally, orientation may be determined from the placement of system components on the railcar.

[0067] Fig. 2 is a flowchart showing the method of determining the location and orientation of a railcar within a railyard. The method makes the following assumptions:

- CMUs are installed in a known location and with a known orientation on each railcar.
- There can be one or many CMUs in the railyard.
- The boundaries and orientation of the railyard with respect to magnetic North is known by geo-fences and historical data.
- Time-stamps are associated with all sensor events.
- The orientation of a railcar in a known railyard can be used rather than the position of a device with a compass that is installed on a railcar.

[0068] The method starts with the assumption at 150 that the railcar is in the railyard. At 151, 152 and 153 it is determined whether or not the railcar is moving through use of an accelerometer, a motion sensor and/or a GNSS respectively.

[0069] At decision point 154, if motion was detected control proceeds to 157 where a confidence level is calculated and, at decision point 156, it is determined if the calculated confidence level exceeds the required threshold. The confidence level calculated at 157 is the likelihood that the railcar is actually moving. If, at decision point 156 the threshold is not met or exceeded, control proceeds back to the beginning of the method where various sensors are checked for movement. If it is determined that the railcar is in motion, at 158 a compass heading and GNSS location are periodically obtained at 159 and at 160. Readings from the accelerometer and motion sensor are also periodically obtained. At decision point 163 it is determined if the heading of the B-end of the railcar can be determined. If it can, a confidence level is calculated at 166 and, at decision point 167 it

is determined if the confidence level exceeds the required threshold. If the threshold is exceeded, a message is sent with a direction the B-end the railcar is facing including the confidence level at 169. If the confidence level does not exceed the threshold at decision point 167, then control returns to the beginning of the method where movement is detected at 151, 152 and 153. At decision point 168, the user may optionally configure the system to send the message regardless of the confidence level, in which case the message is sent at 169.

[0070] If, at decision point 154 it is determined that no motion was sensed, the railcar is declared as being stationary at 155 and a compass heading and GNSS location are obtained at 161. At decision point 162 it is determined if the orientation of the railyard is known. If it is unknown, control proceeds to 165 where the GNSS location and compass headings from at least 3 railcars in the train consist are obtained. At 164, the compass heading and GNSS location from the railcar in question is compared to the readings obtained at 165 from at least three other railcars. At decision point 163 it is determined whether or not the heading of the B-end of the railcar can be determined, and, if not, control proceeds as described above. At decision point 162, if the orientation of the railcar is not known, then control proceeds directly to decision point 163 and thereafter proceeds as above.

[0071] Fig. 3 is a flow chart showing a method of determining whether or not a railcar is inside of a railyard. In this case, the method assumes that the railyard is a managed railyard. The method starts at 201 with the railcar. At decision point 202 it is determined if the railcar is a member of the railyard-based wireless mesh network 117. If it is, control proceeds to decision point 205 where it is determined whether or not the location of the

railcar as reported by GNSS is consistent with the railcar being in the railyard. If it is, a confidence level that the railcar is actually in the railyard is calculated at 206.

- [0072] At decision point 208, it is determined if the confidence level exceeds the required threshold for making a determination that the railcar is within the railyard. If the threshold is exceeded, control proceeds to 209 where it is determined that the railcar is in the railyard. If the confidence level is not exceeded, control returns back to decision point 202.
- [0073] If, at decision point 205, the location of the railcar as reported by GNSS is not consistent with the railcar being in the railyard, control proceeds to 207 and the conclusion is drawn that the railcar is not in the railyard.
- [0074] If the railcar is not a member of the railyard-based wireless mesh network 117, control proceeds to decision point 204, where it is determined if the railcar passed an AEI scanner. If the railcar has passed an AEI scanner, control proceeds to decision point 205 and proceeds as above. If, at decision point 204 the railcar has not passed an AEI scanner, it is determined at decision point 203 if the railcar is within a geo-fence defining the boundaries of the railyard. If it is determined that the railcar is within the railyard's defined geo-fence, control proceeds to decision point 205 and proceeds as described above. If, at decision point 203 it is determined that the railcar is external to the railyard's defined geo-fence, it is determined that the railcar is not in the railyard at 207.
- [0075] A collection of links creates a train consist as referenced in Fig. 4. A train consist is built one link at a time. The linking of railcars and links of railcars is a critical part of this process and can be determined by one or more methods, which can be used stand-alone or in combination to provide a level of probability that two or more railcars are linked, or

that two or more links of railcars are linked. The confidence level of the order of the railcars in a train consist is increased if more than one method is used. The sensor readings and process results are associated to an asset, a component of the asset, a phenomenon, and time. The information is stored so that analysis can be performed on both real-time and historical datasets.

[0076] Fig. 13 is a flowchart showing the process for verifying whether two or more railcars have been coupled, or whether two or more links have been coupled. The process starts at 1301 and, at decision point 1302, it is determined if an event has occurred for which a probability curve exists (i.e., an event that may be relevant in determining coupling). If not, control returns back to decision point 1302. If an event of interest was received, the value of the probability for that event is retrieved from the relevant probability curve at 1303. At decision point 1304, it is decided if enough events have occurred such that a coupling can be evaluated. If not, control returns to the decision point 1302. If enough events have occurred, the probabilities from the probability curves for each of the events are retrieved at 1306 and multiplied together to create an overall probability. At decision point 1305 it is determined if the overall probability exceeds the predetermined threshold necessary to declare that a coupling has positively occurred. If not, control returns to decision point 1302. If so, then the coupling event is declared to have occurred at 1308.

[0077] Fig. 4 shows the formation of a train consist built of links of railcars. In Figure 4(a), railcar B impacts railcar A and forms link 401. Likewise, railcar D impacts railcar C and forms link 402. In Fig. 4(b), railcar C impacts railcar B to form larger link 403 shown in Fig. 4(c). In Fig. 4(d) a single railcar E impacts railcar D to form link 404, consisting of railcars A through E, shown in Fig. 4(e).

[0078] CMUs 101 primarily provide data upstream to determine the presence of railcars in a railyard, the location and orientation of railcars in a railyard (Fig. 2), a connecting or linking of railcars as they are prepared to be part of a train consist (Fig. 4), an order of railcars in a train consist, a validation of railcars in a train consist and a direction of travel of a train consist. Additionally, the CMU has an optional means for monitoring the output from a variety of sensors (both internal to the CMU and in WSNs which are in communication with the CMU) as well as attached directly to a railcar and determining the behavior and condition of the railcar and its various components, based on an analysis of the data. The sensors collect, store, analyze and process data, which is then transmitted to the CMU for further transmission to a PWG, where an engineer, control point or automated system can act on the data, for transmission to a remote railroad operations center, or for processing and analysis to build alerts, events or reports.

[0079] The CMU is capable of collecting data from each integrated sensor as well as from WSNs and performing higher-level analysis of the data by applying heuristics and statistical models to data, events and alerts collected from a plurality of WSNs, to determine location, speed, heading, condition and more of a railcar. During such data analysis, heuristics may be applied to determine potential linking of railcars based on statistical models and empirical data. The CMU also is capable of communicating both the data and the results of any analysis to another system remote from the railcar, via any one of a number of communication protocols.

[0080] A PWG may be located, for example, on a locomotive, in a railyard or at an off-train location at a remote railroad operations center. The PWG may also be able to perform higher-level analysis of the condition of an entire train consist by applying heuristics and

statistical models to data, events and alerts collected from a plurality of CMUs, located on different railcars in the train. The analysis of the data collected can be carried out at any one of a plurality of different event engines distributed among the various components in the present invention, including the sensor units, CMU, train-based or land-based PWGs, or other land-based stations. The event engine is used to determine state changes and actions to perform on the device from a plurality of inputs internal or external of the system. The logic used to determine an outcome is based on a set of rules which can be configured and updated remotely.

[0081] Fig. 5 shows a method for managing data as it flows from sensors on WSNs 104 or the CMU 101 and thereafter to various higher-level destinations. The following assumptions are made:

- A method of data analysis is carried out by event engines at each level.
- Logic analysis is pushed out to the lowest level possible to enable more effective management of bandwidth, power consumption and latency.
- Events are only published upstream when necessary.
- Filtering and analysis of data and events is conducted at each level.
- CMUs, PWGs and servers (within the control center) can utilize sensor fusion to better determine the state of larger systems that share events from these different data sources.

[0082] The lowest level of processing 502 includes the optional WSNs 104 disposed on each railcar 103(a) or 103(b), and sensors which may be integrated into CMUs 101 on each railcar. Data collected at lowest level 502 is analyzed by on-board processors included in

each WSN 104 or CMU 101 to determine which data can be discarded and which data needs to be sent to the next higher processing level 504. The next highest processing level 504 includes a CMU 101 on each railcar. CMU 101 on each railcar is capable of making decisions which may require data from multiple WSNs 104 on the railcar. CMU 101 can also determine, based upon this analysis, what data needs to be sent to the highest processing level 506. The highest processing level 506 includes a PWG 102 located on the locomotive, land-based PWGs 116 disposed in the railyard and control center. PWG 102 in the locomotive is capable of making decisions which require information from multiple CMUs 101 or from multiple WSNs 104 on each railcar (i.e., train consist-wide statuses). If a railcar 103(a) or 103(b) is within the confines of a railyard, messages from CMU 101 may be sent to a PWG 116 located in the railyard. This would be a land-based stationary PWG 116. CMU 101 on each railcar at level 506 may also send messages directly to control center. At the highest level of processing, information may be shared between a locomotive-based PWG 102 and railyard-based PWG 116 and control center. Box 506 represents the highest level of processing and decisions at this level typically represent status information regarding an entire train consist or railyard.

[0083] The various levels of processing combine to create a distributed inference engine in which each level of processing can draw inferences requiring data from that level and/or data which has been provided by lower levels of processing and moved to higher levels. As an example, verifying a coupling event requires data from at least two railcars (e.g., detect impact data and location data from each railcar being coupled). As such, the coupling event must be made at the highest level of processing after receiving data from

each railcar. In this case, the highest level of processing is represented by 506 in Fig. 5, which would be a node in the railyard-based wireless mesh network.

[0084] Fig. 6 is a flow chart showing the method of transmitting messages, based on priority, from the lower levels of processing 502 to the higher levels of processing 504 and 506, shown in Fig. 5. The method starts at 501 where an event message is created. At 502 the message is assigned a priority level which is based on a user configuration and, at decision point 503 it is determined if high bandwidth is available to transmit the message. If high bandwidth is available, control proceeds to 510, where the message is transmitted. If high bandwidth is not available, at decision point 505 it is determined if the message has a high priority status. If the message is high priority, control proceeds to decision point 506 where it is determined if there is low bandwidth available. If low bandwidth is available, the message is transmitted at 510. If the low bandwidth is not available or if the message does not have high priority status, control proceeds to decision point 507 where it is determined if the user configuration defines a number of re-transmission attempts over a specified period of time. If so, then control proceeds to decision point 504 where it is determined if the required number of attempts have been exceeded, and if not, control proceeds to decision point 503 and proceeds as described above. If the number of re-transmission attempts has been exceeded, or if the user has not configured the re-transmission option, then the message is stored for a predefined time period before a bandwidth availability check is performed at 508. At decision point 509 it is determined if the bandwidth check time period has been reached, and if so, control proceeds to decision point 503 and proceeds as described above. If the time period has not

been reached then control loops back and the message is stored until the bandwidth check is to be performed again.

- [0085] The following types of methods can be used to determine the linking (or unlinking) of two or more railcars or two or more links, as shown in Fig. 4.
- [0086] **Motion** – If an accelerometer, and or a motion sensor and or GNSS indicate motion on two or more railcars, the time stamps are compared to determine the likelihood that two or more railcars are linked.
- [0087] **Speed and Heading** – When two or more railcars are traveling at the same speed and on the same heading then they are considered linked.
- [0088] **Network Signal Strength** – A link can be determined by comparing the signal strength across two or more railcars and comparing it to the signal strength of other railcars in the railyard-based wireless mesh network. The signal strength is compared to known adjacent railcars, where the railcars are considered linked. The wireless network connection is established when two or more railcars each have installed a CMU 101 that has the ability to communicate with the wireless network. Each CMU 101 has a measurable signal strength where both the presence of the signal and the strength of the signal can be used to determine if two or more railcars are linked.
- [0089] **Impacts** - An impact with time stamp is generated when two or more railcars are coupled. The time stamp across two or more railcars is compared to determine which railcars have time stamps within a specific time period, which is then used to determine if the railcars are linked. Additionally, during an impact, there is a

positive and negative response created, wherein the positive and negative wave profiles are compared and if they are the same or similar the railcars are considered linked.

[0090] **Location** – If two or more railcars have location readings within proximity to the others, it can be assumed they are linked. The confidence level of this type of linking depends upon the complexity of the railyard. Location information may be obtained from a GNSS.

[0091] **Spline Curve Fit** - Knowing at least three railcars in a train consist, utilize location in conjunction with spline curve fit between railcars in a string. As the train consist is assembled, a best fit curve can be applied to the railcars currently in the train consist. Best fit curve must be within constraints of railroad track geometry. This curve can be used to determine if a railcar is incorrectly marked as not within the train consist, based on location position and proximity to the spline.

[0092] **Compass Heading** - Knowing at least three railcars in a train consist, utilize location in conjunction with angle of compass heading between adjacent railcars (Fig. 8) – As the train consist is assembled, angle variation between adjacent railcars can be used to determine potential linked railcars. Angle must be within constraints of railroad track geometry. The difference in angle between railcars can be used to determine if a railcar is incorrectly marked as not within the train consist, based on location position and angle values that match other adjacent railcars within the same known train consist.

- [0093] **Brake Events** – During a braking event, a pressure change occurs to modify the braking state on each railcar. This event of a pressure change will be perceived by each connected railcar in series from the locomotive to the last connected railcar. The time of this event is used to determine connected railcar order in the train consist.
- [0094] One example of this would be the brake test. A brake test must occur before a train consist can leave a railyard. In this case, brake lines in connected railcars will be pressurized to a standard pressure. This ensures the brakes are released. During a brake test, a sudden drop in pressure occurs to actuate the brakes on each railcar. This event of a sudden pressure drop will be perceived by each connected railcar in series from the locomotive to the last connected railcar. The time of this event is used to determine connected railcar order in the train consist.
- [0095] **AEI Tags** - If two or more railcars are scanned by the same AEI (Automatic Equipment Identification) reader, use the time of the scan, the time difference or offset between the scan of each railcar and the speed of each railcar to determine if the railcars are linked.
- [0096] When an “event” occurs, either asynchronously triggered by external phenomenon (e.g. motion starts) or on a timed basis, the event is recorded and transmitted to a CMU or PWG within the railyard or train consist. The sensors are installed on different components of an asset, recording the asset, time, and details of the event. Some examples of sensors and methods are listed below (but not limited to):

- Asset impact – measured in g-force

- Railcar coupler impact – measured in g-force (this is a more specific form of asset impact)
- Asset GNSS location – latitude and longitude
- Asset speed and heading – measured in mph & direction of travel in degrees
- Brake line pressure change – measured in psi
- Asset AEI tag scan – presence of scan (true/false)

[0097] Fig. 7 shows the method whereby the orientation of a railcar within a railyard is determined utilizing the on-board compass. This is a method that is performed in at 161, 159 and 165 of Fig. 2. This method makes several assumptions. First, the orientation of the railcar can be determined by assuming that the CMU is installed in a known place and orientation on the railcar. It is also assumed that the orientation of the tracks within the railyard with respect to North are known, as shown in Fig. 7(a).

[0098] If the asset is in motion, the orientation of the railcar can be determined by comparing the changes in compass heading, or the lack thereof, over time parallel to the direction of travel as determined by the GNSS location updates. If the vector of the compass matches the vector created by the difference between two or more GNSS points, then the railcar is moving towards the B-end (if the CMU is installed/oriented in that way). This is shown in Figure 7(b). If the vectors are opposite, then the railcar is moving towards the A-end. This is shown in Fig. 7(c).

- [0099] If the asset is stationary, the compass and location can be used to compare to a known railyard layout and orientation stored within the system as shown at 162 of Fig. 2. The compass orientation and GNSS location will be used to compare against the railyard location and orientation to determine the railcar heading. If the asset is stationary and the railyard location is not known, then the orientation of a railcar in question can be compared with other assets in a known group of linked railcars. This is shown at 165 of Fig. 2.
- [00100] Because the rail track can only curve at a small and defined rate, if three or more railcars are known as being linked, the variation in compass heading is small (when accounting for the 180 degree difference if facing opposite directions). If the asset in question is in close proximity to the railcars used for the baseline, or linked as part of the same train consist, a compass reading of the asset can be compared to the other assets to determine heading. As with other methods discussed herein, a confidence level can be assigned to the result, as shown at 166 and 167 of Fig. 2.
- [00101] Fig. 8 shows a method to determine whether two railcars are on the same rail track or not. This method uses a spline curve fit to apply a best fit curve to the assets in the train consist. Any best fit curve that is not within the constraints of the railroad track geometry can indicate railcars on different tracks. As with previous methods, CMUs 101 on each railcar must be installed in a known location and orientation on the railcar. These locations are used to pair assets with the closest proximity to each other. The angle is calculated between railcars in close proximity (within the configurable distance of the maximum railcar gap)

to determine the relative angle differences between railcars in close proximity. A GNSS reading of two railcars is used to determine the vector between each. This vector direction is compared to the compass heading of the railcar (against North). When angles between the GNSS vector and the compass heading are small, then the likelihood of the assets being on the same track is very high. If a difference in vector between the GNSS vector and the compass is high, then it is unlikely that the assets are linked and on the same track. The difference in angles becomes worse as problems cascade down the track.

[00102] As an example, with reference to Fig. 8, if the angle between A and B is small these are likely linked. If the angle between B and C is large, then these are likely not linked. The angle between C and D is also high and are also not likely linked. The maximum angle threshold can be used to determine if assets are likely linked or not. In Fig. 8, angle AB is the angle of railcar A relative to railcar B, and an example of an angle within the bounds of “Z” degrees (i.e., degrees indicating that track geometry has not been violated). Angle BC is the angle of the heading of railcar B with respect to railcar C, and angle CD is the angle of railcar C with respect to railcar D. Angle BD represents the difference between angle BC and angle CD. If angle BD exceeds “Z” degrees then it can be determined that railcar C is on a different rack than railcars B and D. If not, then railcar C is likely on the same track as railcars B and D. The threshold “Z” degrees is determined by geometry of the rail tracks.

[00103] A statistical logic engine is used to determine the confidence level of various determinations that may be inferred from the data that is collected from each

railcar, including, for example, which assets are linked. Conditional probability is used to combine several different inputs, of different phenomenon types and units of measure, to provide a single output based on the knowledge of those other events.

[00104] For each method, component, and phenomenon, a probability chart is supplied to determine the difference between events occurring on two separate assets. Depending on the method used, the X axis represents the difference between the events or data collected from sensors on two (or more) assets.

[00105] Each sensor (component and phenomenon pairing) and method has a probability curve showing the likelihood of a coupling event between two assets, wherein the X-axis can be based on the phenomenon that is measured, the time between events, or both (as a three-dimensional graph), as observed between two assets, and the Y-axis represents the probability of a coupling event. A coupling event is not guaranteed, to occur at any particular X measurement, but the measurement represents the opportunity for the coupling event to occur. A 1.0 on the graph indicates a coupling event is possible, for this sensor type or method. A 0.0 on the graph precludes a coupling event, invalidating all other sensor input curves in combination. Examples of probability charts are shown in Fig. 10, where Fig. 10(a) shows a probability curve for time between an impact event across two railcars and Fig. 10(b) shows a probability curve for the distance between two assets.

[00106] When events are received from multiple assets, the probability result is generated based on available data at the time. If the analysis of events across assets does not

result in a coupling (or railcar linking) event, the events are saved, and can be reprocessed again when other events occur between the asset pair.

- [00107] An example is shown in Fig. 11. Fig. 11(a) shows information is obtained regarding the impact times, showing the difference in time between two impacts, as measured on two railcars, is 0.19 seconds, resulting in an output value of 0.85, which represents an 85% probability that a linking has occurred. Fig. 11(b) shows a difference in distance between two railcars as 55 meters, resulting in an output value of 0.62, representing a 62% probability that a linking has occurred.
- [00108] It is important to account for inaccuracies and imprecision in different sensors and methods when generating probability curves and assigning weighting to different methods. A curve should not have a probability level above the accuracy provided. Preferably, more accurate and precise methods are weighted higher than other methods.
- [00109] In the simplest manifestation of the algorithm, the individual probabilities are multiplied together to get a combined probability, which, in this example, results in a .527 probability that a linking has occurred. This calculation does not utilize other sensor inputs, historical data, or apply a configurable weighted average, but all of these possibilities are within the scope of the invention.
- [00110] The output value is compared to the user-defined threshold of what constitutes a linking event. If, for example, the threshold was set to 0.75, then this instance would be marked as “not linked”, but an analysis can be executed again when new data is received for the assets in question.

- [00111] There is a minimum threshold value which must be equaled or exceeded for the system to declare that a coupling event has occurred. The link state between an asset pair is defined as linked, not linked, or no data. Linked indicates that the calculated result is above the minimum threshold. Not linked indicates a calculation was executed, but fell below the minimum threshold – these asset pairings can be re-calculated when new event data is received for the assets and their respective components. No data indicates that there are no sensor readings for the asset pairing in question
- [00112] In addition to the pre-defined probability curves, historical metrics can be used for the same X and Y graphs, to compare results against a histogram of instances and verified results. The sensor histograms can optionally be used in place of the pre-defined probability curves, or in combination with the pre-defined probability curves (multiply the two results together per sensor), to show a confidence interval in a valid asset-coupling result (and quantity of events). An example of this is shown in Fig 12, wherein Fig 12(a) shows an historical histogram for the difference in times of impact and Fig 12(b) shows difference in distance.
- [00113] In another embodiment, a version of the histogram method shown in in Fig. 12 could use used to identify the accuracy of the asset link assumption itself. In other words, the histogram would show how often the result was correct (linked or not-linked) instead of only showing how often the X value resulted in an actual asset linking event.
- [00114] Using this method, many different parameters and inputs can be used to generate the conditional probability of a linking event. As an example, two railcars are

coupled together in a railyard, using a locomotive travelling at roughly 3 mph. An event is recorded on two separate railcar coupler accelerometers, both indicating peak impact events of 7 g's, within 1 millisecond of each other. A three-dimensional probability graph for a railcar coupler accelerometer uses the difference in time for the X axis, the difference in g-force as the Z axis, and the probability (0.0 to 1.0) as the result in the Y axis. After the event occurs, the PWG requests a location and speed of both assets, and the result is transmitted back to the PWG, indicating both assets are now stationary. The graph for difference in speed is used in combination with difference in time and the difference in g-forces to provide a secondary input, resulting in a value above the threshold used to mark the assets as being linked.

- [00115] In one embodiment of the invention, the probability curves that associate to sensors and methods can be dynamically added, modified, and removed from the system. Machine learning algorithms can be used to automatically generate curves based on historical data when the final train manifests are provided.
- [00116] In another embodiment, the system can be user-configurable. Method and sensor selections can be marked as enabled, ignored, or required. Additionally, the minimum number of distinct methods required to perform analysis (e.g. 2 or more needed or a result is not generated) can be specified.
- [00117] In another embodiment, the system also has the capability of proving probability curves for each method, component, and phenomenon. A hierarchy of curves can exist for each sensor, mapping to more specific measurements, if available. For example, there may be an overall probability curve for impact, but if an asset has

an impact sensor mounted on the coupler on a railcar, that more distinct probability curve for a coupler impact event can be applied in place of the higher-level impact curve. In the event that one asset has a more specific sensor mapping and the other has the higher-level mapping for the same phenomenon, the association between the assets can be configured to be allowed or rejected

- [00118] In another embodiment, the ability to provide a relative weighting metric for different methods is provided. For example, GNSS location between two linked railcars may be determined to be 4 times as important as compass heading to determine if a linking has occurred.
- [00119] The system also has the ability to utilize historical data and a final result, provided externally, to validate linking events against known outcomes. This feedback is used to enhance probability curves and confidence intervals for different method, component, and phenomenon inputs. For example, if a railroad provides a final manifest for trains created, the actual data could be used as a check against predicted assumptions of railcar links, and mark each as valid or invalid.
- [00120] The system also has a user-configurable window of time indicating when historical events are valid for analysis. The window indicates how long existing data can be used for analysis, based on each sensor type or method.
- [00121] In another aspect of the invention, the system is capable of determining the order of railcars within a train consist. Any combination of the following can be used to determine the order of train.

- [00122] Using historical data, and any combination of the “linking” algorithms previously described, the orientation and order of railcars within the train consist can be determined based on the time of the event, and the railcars involved for each link.
- [00123] The system also utilizes physical constraints to accept or reject events that result in a link. For example, a single asset can only be linked to, at most, two other assets because there are physically only two couplers per railcar.
- [00124] The time scan of the AEI tag plus elapsed time provides the position of a railcar within a train consist and optionally railcar heading and railcar speed, and can be used to validate the order and orientation of the railcars within the train consist as the train passes by the AEI reader (typically as the train is leaving the railyard).
- [00125] The railcar’s location can be used, however, direction of travel will not be determined and the confidence level will be low. The railcar’s location plus the compass heading of the same railcar can be used, however the direction of travel will not be determined.
- [00126] Using the “accordion effect” or push/pull, an accelerometer in each railcar’s CMU records impact force as the railcar is pushed and pulled when the train moves. The impact force is recorded with a time stamp and offset and compared with other railcars in the train. Such movement creates a cascading events through the train, in which the event time stamps can be compared to determine in what order two or more railcars are moving. If the impacts and time stamps from two or more railcars show a time gap it is assumed there is a number of unmonitored railcar in the train consist.

- [00127] The railyard-based wireless mesh network or the train-based wireless mesh network can determine if a railcar is in the network and if so, can compare the signal strength of the railcar with the signal strength of other railcars in the network. There is a low confidence level using this method.
- [00128] There are multiple ways to validate the order of a train consist as it leaves a railyard. Data can be collected regarding location, speed, heading, movement, network signal strength and paths. Using these data points increases the confidence level regarding the order and orientation of the railcars within the train consist, when they are consistent with the pre-supposed configuration of the train consist.
- [00129] In another aspect of the invention, the direction in which a train is traveling can be determined by employing one or more of the methods described below and as referenced in Fig. 7.
- [00130] In aspects of the invention, the heading and orientation of a railcar can be determined. Regarding orientation, it is desirable to know whether the “A” end or the “B” end of a railcar is facing the head end of the train. This is important to railroads and to shippers to know the “A” and “B” end orientation because a railcar may be required to be positioned at its final destination such that the “A” or “B” end it facing a specific direction. In Fig. 2, the data from sensors and an algorithm to process the data provide a confidence level that the correct end of the railcar will be known. The CMU must be installed in a known orientation, for example, positioned on the B-end of railcar. The heading of the CMU is compared with North to determine the orientation of the railcar. Also, it is

preferred that the direction of the railyard be known based on historical or geographic data such as rail track is in a Southwest to Northeast direction (See Fig. 7).

[00131] If the orientation of the railyard is not known, location data and compass heading of at least three linked railcars can be used to determine railcar heading by comparing compass heading of a railcar versus the direction of the track inferred by three or more linked railcars. If the orientation of at least one railcar is known, the heading of other railcars that are linked can be derived by comparing the compass heading of a railcar versus the known heading of the other linked railcars. If the orientation of at least one railcar is known, the heading of other railcars that are linked can be derived by comparing the timing of the impact during the coupling event as measured at the “A” and “B” of railcar. This impact information combined with the known orientation of one railcar will determine the orientation of the other railcar.

[00132] In another aspect of the invention, the system can be used to determine when assets are removed from a train consist or set of assets linked together. Similar to determining if the assets linked as described above, the removal of one or more assets can be inferred by the reciprocal event. Assets are assumed to be linked until otherwise determined by any number of the methods below:

[00133] **Motion** – If an accelerometer, and or a motion sensor and or GNSS indicate motion on two or more railcars with different values, the time stamps are compared to determine if the two or more railcars are unlinked.

- [00134] **Speed and Heading** – When two or more railcars are not traveling at the same speed or on a different heading then they are considered unlinked.
- [00135] **Network Signal Strength** – Unlinking can be determined by comparing the signal strength across two or more railcars and comparing it to the signal strength of other railcars in the railyard wireless mesh network. Where the signal strength is comparable to known unlinked railcars, the railcars are considered unlinked.
- [00136] **Location** – If the location readings of two or more linked railcars are not within proximity to each other within a specified time interval, it is likely they are unlinked. The confidence level of this type of linking depends upon the complexity of the railyard.
- [00137] **Spline Curve Fit** - Knowing at least three railcars in a train consist, location can be utilized in conjunction with spline curve fit between railcars in a string. A best fit curve can be applied to the assets currently in the train consist. Any best fit curve not within the constraints of railroad track geometry can indicate unlinked railcars.
- [00138] **Compass Angle** - Knowing at least three railcars in a train consist, utilize location in conjunction with angle of compass heading between adjacent railcars (Fig. 7). Divergence in the angle variation between adjacent railcars can be used to determine potential un-linked railcars. In other words, the change in heading between consecutive railcars. Angle must be within constraints of railroad track geometry.

- [00139] **Brake Events** – During a braking event, a pressure change occurs to modify the braking state on each railcar. This event of a pressure change will be perceived by each connected railcar in series from the locomotive to the last connected railcar. The time of this event is used to determine connected railcar order in the train consist. If there is no similar pressure change for a railcar, it is less probable to be part of the train consist.
- [00140] **AEI scans** – If two or more railcars are scanned by the same AEI reader, the differences in the time of the scan, or offset between the scan of each railcar and the speed of each railcar can be utilized to determine if the railcars are not linked.
- [00141] The system also utilizes physical constraints to further invalidate links between assets. For example, two railcars heading north in a railyard that only has tracks in the east/west direction invalidates the GNSS sensor method for the calculation.
- [00142] In another aspect of the invention, the presence of a dark railcar can be determined and reported. Dark railcars can be identified by a PWG on the locomotive directly, or the presence of a dark railcar can be passed through the wireless network from the CMU on one or more railcars in the train consist. This process is shown in Fig. 9.
- [00143] Locomotive 108 has a PWG 102 and a railcar 103(a) or 103(b) has a CMU 101, which may be in a state that listens for radio broadcasts from other railcars 103(a) or 103(b) that are not connected to a train-based network, not connected to a managed railyard, or are sitting in an unmanaged railyard.

[00144] As locomotive 108 or a CMU 101 passes a railroad siding upon which at least one monitored railcar 103(a) or 103(b) are sitting, locomotive 108 will listen for radio broadcast identification information from monitored railcars 103(a) or 103(b). If a broadcast is detected, the PWG on locomotive 108 will transmit the identification information about the railcar 103(a) or 103(b) to the remote operations center.

[00145] In a second embodiment, a dark railcar will be in listen mode for other networks. When a railcar 103(a) or 103(b) within a train-based or yard-based wireless mesh network is in range proximity to the dark railcar, the dark railcar will hear “advertisements” from the railcar 103(a) or 103(b) in network. The dark railcar will reply to the advertisement from the railcar, with its identification and settings, which will be passed to the PWG 102. The PWG 102 will have the option of allowing the dark railcar to join the train-based or railyard-based wireless mesh network, passing the information down through the other CMUs to the dark railcar. If the dark railcar is blacklisted, it will not be allowed to join the train-based wireless mesh network. Once the railcar is in the network, it changes to the normal operating profile, and is no longer a dark railcar.

[00146] An important aspect of the invention is the capability of measuring certain parameters on vehicles in the train and relating the measurements or events to a common time base. This enables inferences to be made based on the relative measures. This same capability is important within railyards, to correlate events for train consist creation or facility operations. An example might include being able to sample vehicle acceleration on every railcar in the train consist and using

the relative acceleration (or deceleration) to detect run in and run out at any point in the train. Another example is relating wheel impact events to individual track anomalies, where all wheels on one side of a train may detect it, and we want to associate all events to a single track feature. A railyard example would utilize this functionality to determine the cascading of coupling events as the force of impact translates through several railcars during train consist creation.

[00147] Assets within a railyard or train consist, which are managed, are synchronized to a precise network clock, with time accuracy synchronized across all devices. In the preferred embodiment of the invention, for example better than 1 millisecond time accuracy synchronization is used. This enables direct correlation of events across all assets.

[00148] In a train-based or railyard-based network, where a multitude of CMUs or WSNs having microcontrollers or microprocessors are used, to take a measurement or detect an event, clock drift becomes a limiting factor in the confidence placed on the time base of any measurement. In wired or permanently powered wireless systems with high bandwidth, regular synchronization of the clocks to a master time is an established practice. However, wireless, self-contained and self-powered CMUs and WSNs would use too much bandwidth and consume too much power to maintain the tight time synchronization needed to differentiate between certain types of events or provide a set of instantaneous measures from across the train. Clock drift becomes particularly limiting at temperature extremes or when the temperature changes rapidly over a relatively short period. It is further exacerbated when multiple discrete networks are used (a railcar-based

mesh network connecting with a train-based mesh network for instance) and a mesh topology is employed versus a point-to-point network.

[00149] The present invention overcomes this constraint through the use of a very high accuracy network time base running over a time synchronized mesh network which is used to periodically (based on the desired accuracy) correct the microcontroller's timing mechanism to a predetermined accuracy. In the preferred embodiment of the invention, for example, 1 millisecond accuracy is desired. The system also has the ability to use a broadcast or scheduled event to trigger time-synchronized sampling across the entire train and/or railyard. CMUs are corrected to PWG time and WSNs are corrected to CMU time. This enables simultaneous sampling of data across all components (PWGs, CMUs, and WSNs) to within the predetermined accuracy, with no impact to network bandwidth capacity or power use.

We claim:

1. A system for managing assets in a railyard comprising:

one or more powered wireless gateways disposed in a railyard; and

one or more railcar-based communication management units;

wherein said powered wireless gateways and said communication management units form a railyard-based network;

a computing device, having access to said railyard-based network, said computing device running software for performing the functions of:

collecting data from said railcar-based communication management units regarding events occurring on or the status of their respective railcars;

drawing inferences from said data regarding the state of said railcars;

and

reporting said inferences.

2. The system of claim 1 wherein said drawn inferences are assigned a confidence level.

3. The system of claim 2 wherein said confidence level represents a probability that said inference is true.

4. The system of claim 3 wherein said confidence level is a combination of probabilities from one or more of said events.

5. The system of claim 4 wherein an inference is declared to be true when said confidence level exceeds a pre-defined value.
6. The system of claim 1 wherein said events include one or more of detected impacts, motion, acceleration, GNSS location, speed, compass heading, brake line pressure change and AEI scan.
7. The system of claim 1 wherein said inference is the presence of a railcar within said railyard.
8. The system of claim 7 wherein said data collected includes AEI scan and location.
9. The system of claim 6 wherein said inference is the location and orientation of a railcar within said railyard.
10. The system of claim 9 wherein said data collected includes acceleration information, motion information, GNSS location, compass heading.
11. The system of claim 6 wherein said inference is that two or more railcars are linked.
12. The system of claim 6 wherein said inference is that two or more railcars are not linked.
13. The system of claim 11 wherein said data collected is detected impacts, motion, acceleration and location from each of said two or more railcars.

14. The system of claim 4 further comprising the step of assigning probabilities that an inference is true for each event used in drawing said inference, and

wherein said confidence level that said inference is true is a combination of the probabilities of each event used in drawing said inference.

15. The system of claim 1 wherein said software performs the function of logically building, and validating train consists.

16. The system of claim 1 wherein said software performs the function of logically unlinking railcars from a train consist and validating that said railcars are not in a train consist.

17. The system of claim 16 wherein said inference that a railcar is not in a train consist is validated when data collected from two or more railcars supports a confidence level that exceeds a predetermined threshold.

18. The system of claim 15 wherein said function of building of a train consist comprises the steps of:

(a) determining multiple couplings of two or more railcars, said multiple couplings resulting in a link containing all railcars in said train consist;

(b) determining uncouplings of railcars or links required to form said train consist;

(c) determining the coupling of a locomotive to said link containing all railcars in said train consist; and

(d) forming a train-based wireless network consisting of a manager and at least one node from each railcar.

19. The system of claim 18 wherein said couplings are determined when data collected from each railcar having at least one node thereon supports a confidence level that exceeds a predetermined threshold.

20. The system of claim 19 wherein said step of determining couplings comprises drawing said inferences based at data provided by said two or more railcars being coupled.

21. The system of claim 20 wherein said provided data includes at least data on detected impacts and location data.

22. The system of claim 21 wherein said detected impact and location data is collected by said computing device and processed by said inference engine to create said inference that railcars have been coupled.

23 . The system of claim 22 wherein said inference engine is running on a node in said railyard-based wireless mesh network or on said computing device, or partially on said railyard-based wireless mesh network and on said computing device.

24. The system of claim 23 wherein said distributed inference engine draws inferences regarding which railcars are in said links base on inferences regarding the coupling of railcars.

25. The system of claim 24 wherein said inference engine uses one or more of the following types of data collected from each railcar involved in a coupling event to raise or lower the confidence level that said two or more railcars are coupled:

(a) signal strength of said railyard-based wireless mesh network received by each railcar;

(b) motion data;

(c) speed and heading data;

(d) spline curve fit data;

(e) compass angle data;

(f) brake event data; and

(g) AEI data;

26. The system of claim 12 wherein said function of validating a train consist comprises the steps of:

(a) collecting data from each railcar having at least one node thereon, said data including at least speed, location and compass heading data;

(b) drawing inferences based on said collected data;

(c) verifying that the speed, location and compass headings of each railcar having at least one node thereon is consistent with the overall motion of said train consist.

27. A method of managing assets in a railyard comprising:

software, running on a computing device having access to a railyard-based network, said software performing the functions of:

(a) collecting data from one or more railcar-based communication management units regarding events occurring on or the status of their respective railcars;

(b) drawing inferences from said data regarding the state of said railcars;

and

(c) reporting said inferences.

28. The method of claim 27 wherein said software further performs the function of assigning a confidence level to each of said drawn inferences.

29. The method of claim 28 wherein said confidence level represents a probability that said inference is true.

30. The method of claim 29 wherein said confidence level is a combination of probabilities from one or more of said events.




Fig. 1

2/13

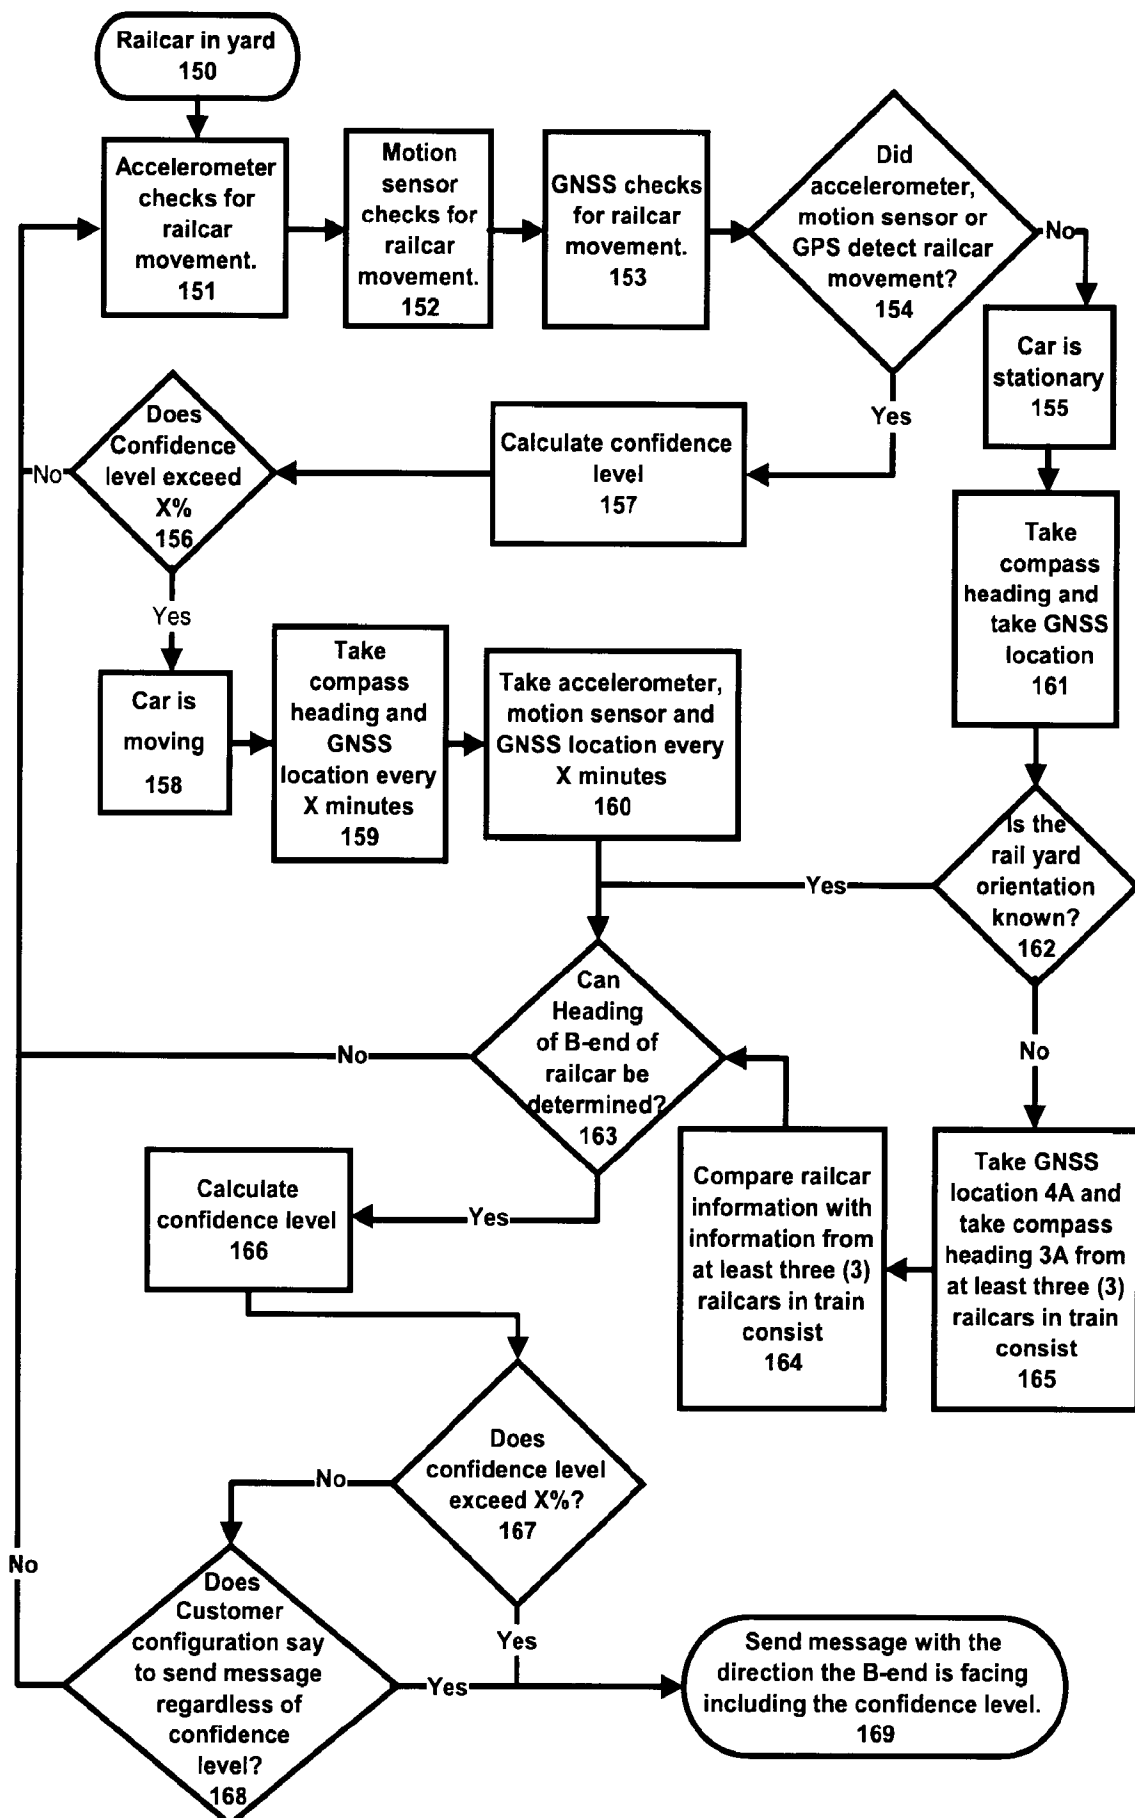


Fig. 2

3/13

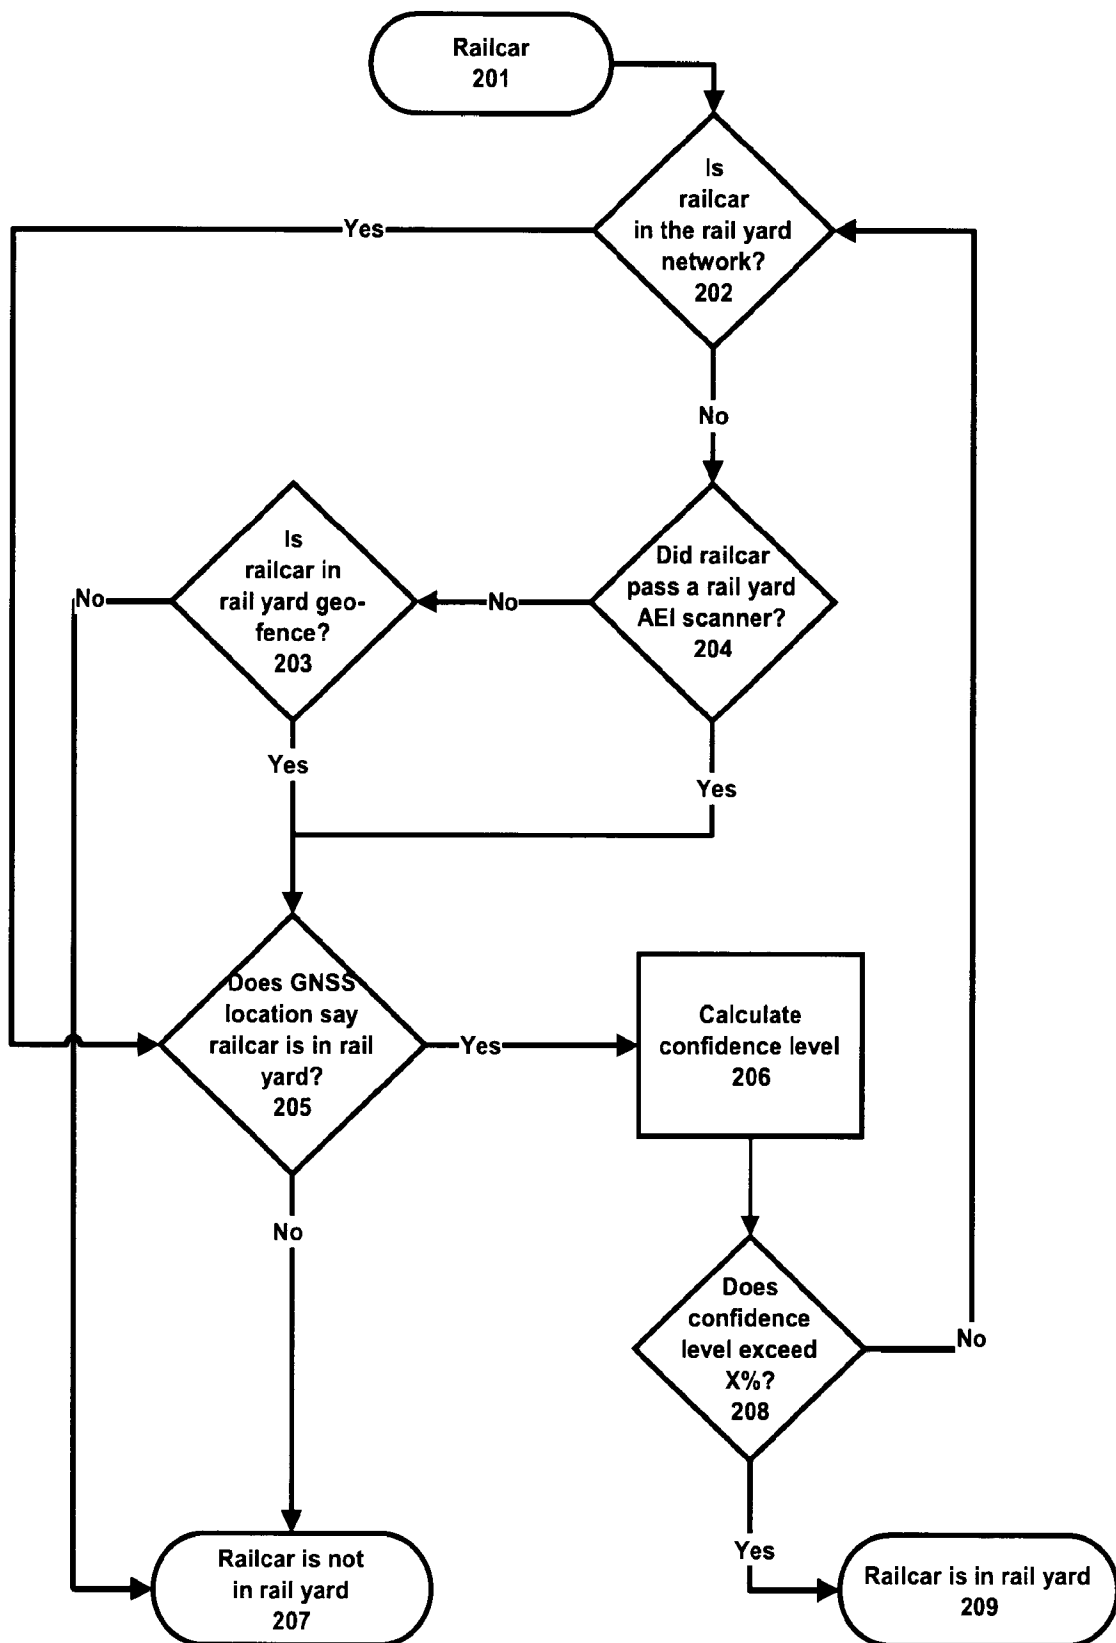


Fig. 3




Fig. 4

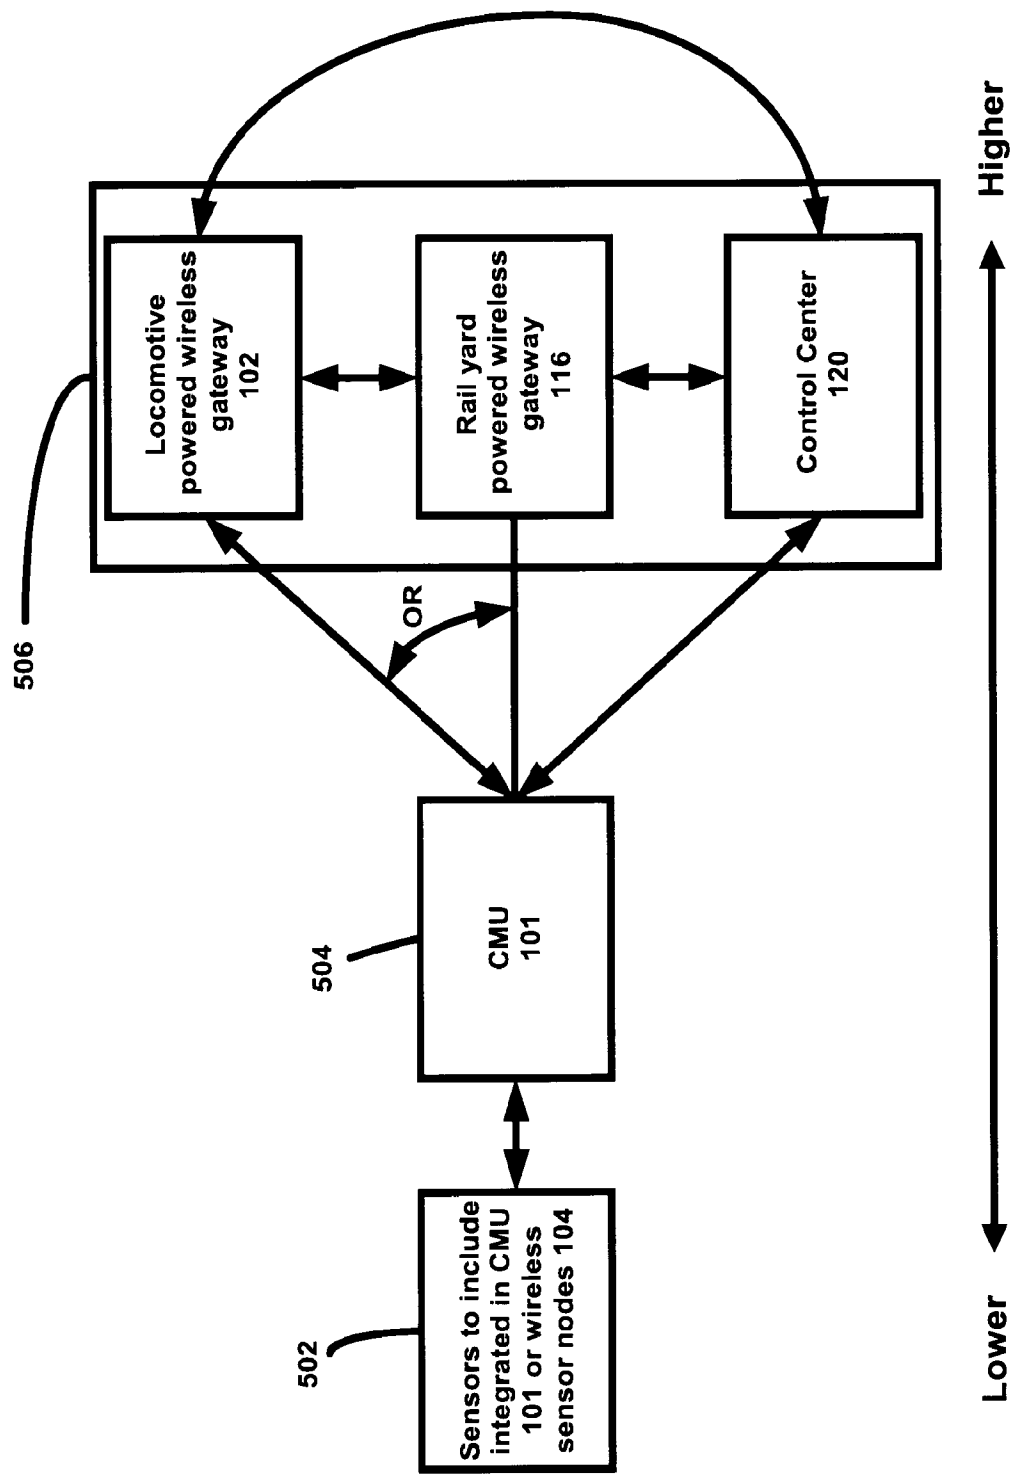


Fig. 5

6/13

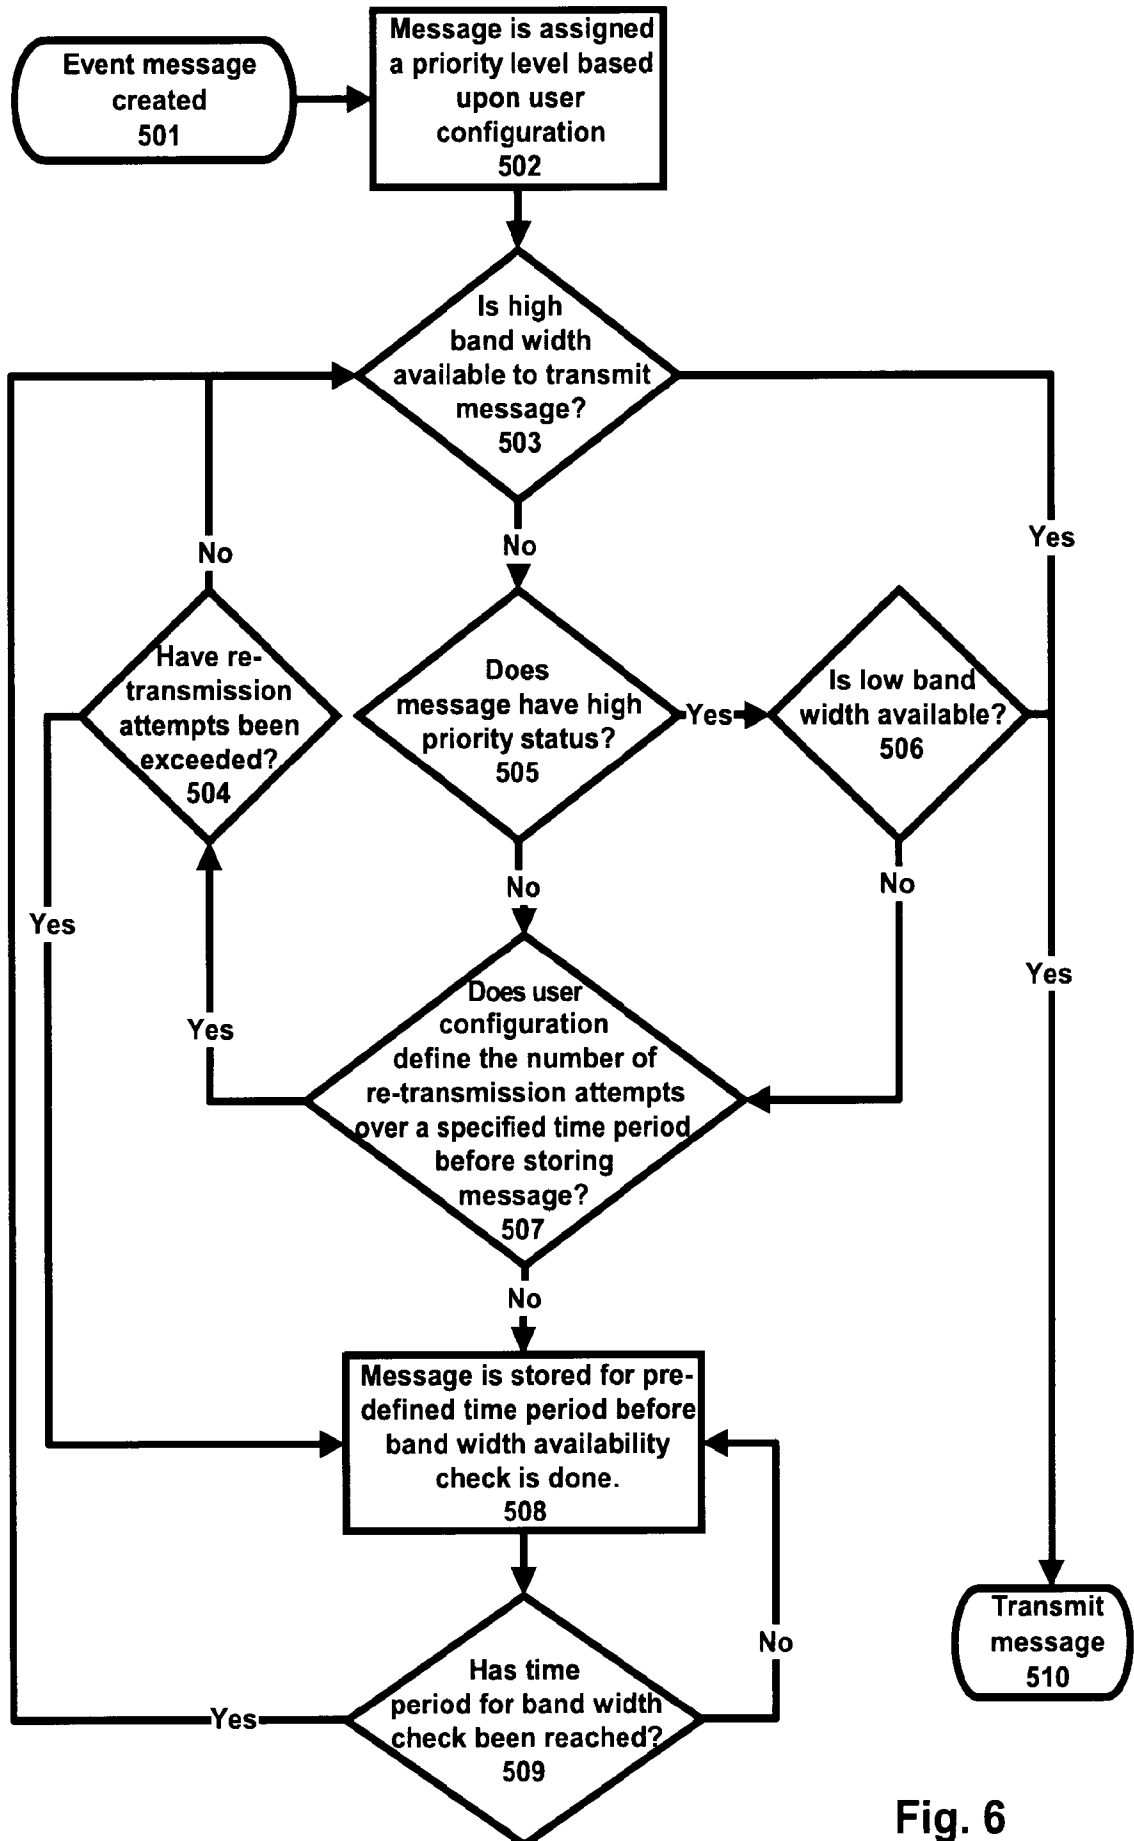


Fig. 6

7/13




Fig. 7




Fig. 8




Fig. 9





Fig. 10



(a)

(b)

Fig. 11


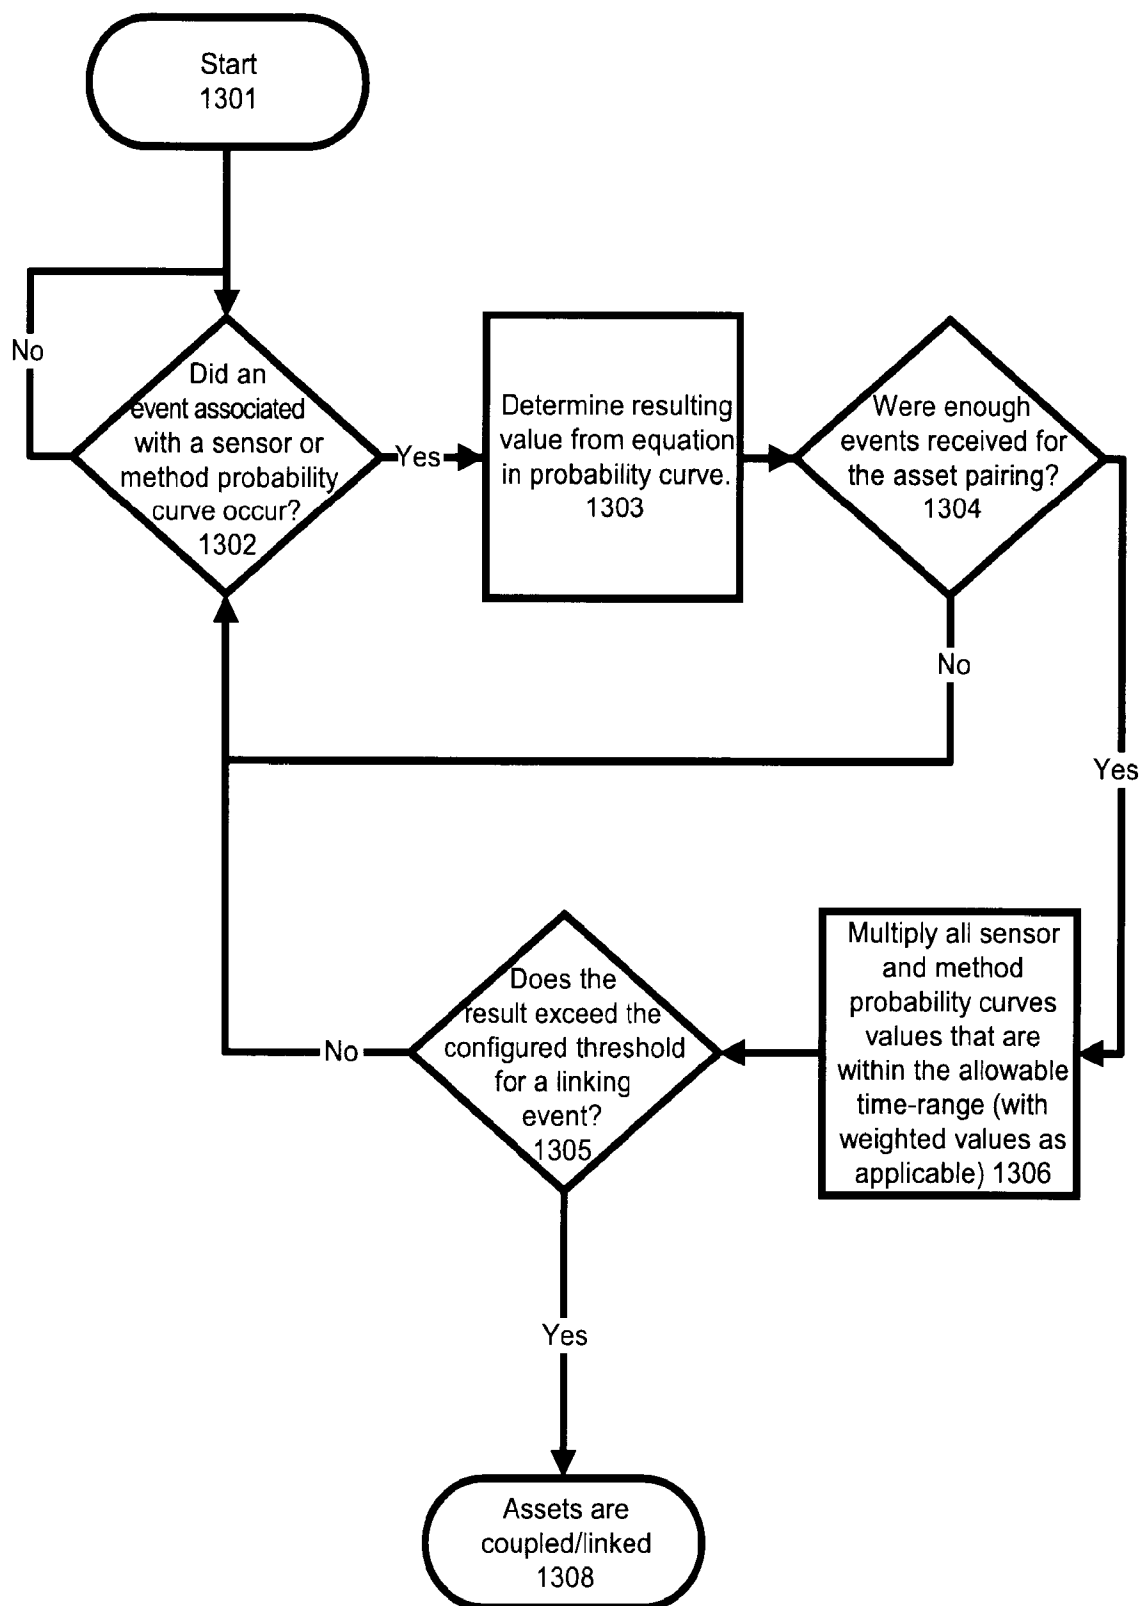


Fig. 12

13/13

**Fig. 13**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US16/34715

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - B61L 15/00, 17/00; G06Q 10/06 (2016.1)

CPC - B61L 15/0018, 15/0027 25/02, 27/0077; G06Q 10/06

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)

IPC(8) Classifications: B61L 15/00, 17/00; G06Q 10/06, 10/0631 (2016.1)

CPC Classifications: B61L 15/0018, 15/0027, 15/0081, 25/02, 27/0077; G06Q 10/06, 10/0631

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

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

PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, IN, CA, INPADOC Data); Google; Google Scholar; EBSCO; IP.com; keywords: statistical, simulation, probability, yard, depot, hump, railyard, hump-yard, sensor, detect, measure, impact, collision, switch, wireless

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X — Y	US 2010/0032529 A1 (KISS J et al.) February 11, 2010; figure 1; paragraphs [0027], [0028], [0031], [0032], [0035], [0038]	1, 6-9, 11, 12, 16, 27 10, 13, 15, 18
Y	US 5,986,547 A (KORVER K et al) November 16, 1999; column 5, lines 30-40; column 6, lines 15-20	10
Y	US 2003/0182030 A1 (KRAELING MB et al.) September 25, 2003; paragraphs [0023], [0024], [0026], [0028]	13
Y	US 2008/0097659 A1 (HAWTHORNE MJ) April 24, 2008; paragraphs [0028]-[0031]	15, 18
A	US 2007/0005200 A1 (WILLS MS et al.) January 04, 2007; entire document	1-30
A	US 2013/0116865 A1 (COOPER J et al.) May 09, 2013; entire document	1-30
A	US 2014/0372498 A1 (INTERNATIONAL ELECTRONIC MACHINES CORPORATION) December 18, 2014; entire document	1-30

☐ 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

05 August 2016 (05.08.2016)

Date of mailing of the international search report

02 SEP 2016

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