

(51) International Patent Classification:
G05B 19/18 (2006.01)(21) International Application Number:
PCT/US2009/036685(22) International Filing Date:
10 March 2009 (10.03.2009)


(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/048,110 13 March 2008 (13.03.2008) US(71) Applicant: **BATTELLE ENERGY ALLIANCE, LLC**
[US/US]; P.O. Box 1625, Idaho Falls, ID 83415-3899
(US).(72) Inventors: **NIELSEN, Curtis**; 343 Talon Drive,
Rexburg, ID 83440 (US). **BRUEMMER, David**; 608
Cedar Ridge Drive, Idaho Falls, ID 83404 (US). **FEW,**
Douglas; 2155 Richards Avenue, Idaho Falls, ID 83404
(US). **WALTON, Miles**; 2756 W. Meadowlark, Lane,
Idaho Falls, ID 83402 (US).(74) Agent: **CHRISTIAN, Stephen, R.**; Battelle Energy Al-
liance, LLC, P.O. Box 1625, Idaho Falls, ID 83415-3899
(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, BR, BW, BY, BZ,
CA, CH, 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, IS, JP, KE, KG, KM, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO,
NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG,
SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA,
UG, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ,
TM), European (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, SE, SI, SK, TR),
OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML,
MR, NE, SN, TD, TG).**Published:**

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments (Rule 48.2(h))

(54) Title: SYSTEM AND METHOD FOR SEAMLESS TASK-DIRECTED AUTONOMY FOR ROBOTS

**FIG. 1**(57) Abstract: Systems, methods, and user inter-
faces are used for controlling a robot. An en-
vironment map and a robot designator are pre-
sented to a user. The user may place, move, and
modify task designators on the environment
map. The task designators indicate a position in
the environment map and indicate a task for the
robot to achieve. A control intermediary links
task designators with robot instructions issued
to the robot. The control intermediary analyzes
a relative position between the task designators
and the robot. The control intermediary uses the
analysis to determine a task-oriented autonomy
level for the robot and communicates target
achievement information to the robot. The tar-
get achievement information may include in-
structions for directly guiding the robot if the
autonomy level indicates low robot initiative
and may include instructions for directing the
robot to determine a robot plan for achieving
the task if the autonomy level indicates high
robot initiative.

TITLE OF THE INVENTION

SYSTEM AND METHOD FOR SEAMLESS TASK-DIRECTED AUTONOMY FOR ROBOTS

RELATED APPLICATION

[0001] This application claims benefit of U.S. Non-provisional application No. 12/048,110, filed March 13, 2008, entitled SYSTEM AND METHOD FOR SEAMLESS TASK-DIRECTED AUTONOMY FOR ROBOTS, which is incorporated herein by reference in its entirety.

CONTRACTUAL ORIGIN OF THE INVENTION

[0002] The United States Government has rights in the following invention pursuant to Contract No. DE-AC07-05-ID14517 between the U.S. Department of Energy and Battelle Energy Alliance, LLC.

TECHNICAL FIELD

[0003] The present invention relates generally to robotics and, more specifically, to software architectures for realizing levels of autonomy in robots.

BACKGROUND

[0004] Historically, robot behaviors have been created for specific tasks and applications. These behaviors have generally been reinvented time and again for different robots and different applications. There has been no sustained attempt to provide a kernel of basic robot competence and decision making that can be used to bootstrap development across many different applications.

[0005] Some architectures have been proposed that provide a generic application programming interface (API) for querying various sensors and commanding various actuators; however, many of these architectures have been limited to raw inputs and outputs rather than provide the intelligence and behavior to a robot. As a result, the behavior functionality created for one robot may not be easily ported to new robots. Other architectures have been proposed to allow limited behaviors to port across different robot platforms, but these have generally been limited to specific low-level control systems.

[0006] The problem with robots today is that they are not very bright. Current robot "intelligence" is really just a grab-bag of programmed behaviors to keep mobile robots from doing stupid things, like getting stuck in corners or running into obstacles. The promise of wireless robots is that they can be sent into remote situations that are too difficult or dangerous for humans. The reality is that today's robots generally lack the ability to make any decisions on their own and rely on continuous guidance by human operators watching live video from on-board cameras.

[0007] Most commercial robots operate on a master/slave principle. A human operator completely controls the movement of the robot from a remote location using robot-based sensors such as video and Global Positioning System (GPS). This setup often requires more than one operator per robot to navigate around obstacles and achieve a goal. As a result very skilled operators may be necessary to reliably direct the robot. Furthermore, the intense concentration needed for controlling the robot can detract from achieving mission goals.

[0008] Although it has been recognized that there is a need for adjustable autonomy, robot architectures currently define levels of autonomy that are only selectable by the user. As a result, performance of the robot and its human controller is often dependent on the human controller's understanding of the robot's autonomy levels and when the various levels are appropriate.

[0009] Therefore, there is a need for systems and methods that provide a more intuitive approach for the user to direct a robot's activities without having to understand the details of available robot behaviors and autonomy levels.

BRIEF SUMMARY OF THE INVENTION

[0010] Embodiments of the present invention comprise methods and systems that provide a more intuitive approach for the user to direct a robot's activities without having to understand the available robot behaviors or autonomy levels. This intuitive approach provides a bridge between a user's intentions and a robot's behaviors by creating a new user interface and adjustments to robot initiative and behavior that does not need to be directly controlled by the user.

[0011] An embodiment of the present invention comprises a Graphical User Interface (GUI) for controlling a robot. The GUI includes an environment map window, a robot designator, one or more task designators, and a control intermediary. The environment map window is configured for displaying a map of an environment proximate the robot. The robot designator is configured for showing a robot position and a robot pose in the environment map

window. The task designators are configured for a user to position them in the environment map window and indicate a task for the robot to achieve. The control intermediary links user defined tasks input by the user with robot instructions issued by the GUI to the robot. The control intermediary analyzes a position of the task designators relative to a position of one or more robot components associated with the task designators. Responsive to the analysis, the control intermediary determines task-oriented robot instructions including defining autonomy level and intelligent behavior for the robot and communicates target achievement information to the robot. The target achievement information may include instructions for guiding the robot to achieve the task if the task-oriented autonomy level comprises low robot initiative. The target achievement information may also include instructions for directing the robot to determine a robot plan for achieving the task if the task-oriented autonomy level comprises high robot initiative.

[0012] Another embodiment of the present invention comprises a method for controlling a robot. The method includes providing a user interface for controlling the robot and positioning at least one task-oriented target in a map on the user interface representing an environment of the robot. The method also includes determining a task-oriented autonomy level of the robot responsive to a change in the at least one task-oriented target. If the change in the task-oriented target is smaller than a change threshold, the method includes instructing the robot to achieve the task-oriented target by intervention instructions from the user interface. If the change in the task-oriented target is larger than the change threshold, the method includes instructing the robot to achieve the task-oriented target using robot initiative to determine a robot plan for achieving the task-oriented target and implementing the robot plan.

[0013] Another embodiment of the present invention also comprises a method for controlling a robot. The method includes receiving instructions for achieving at least one task-oriented target from a user interface. The instructions include at least one of intervention instructions, robot initiative instructions, and instructions for setting a task-oriented autonomy level. If the instructions are robot initiative instructions, the method includes developing a robot plan to achieve the task-oriented target and performing the robot plan. If the instructions are intervention instructions, then the method includes performing the intervention instructions for achieving the task-oriented target and, if present, overriding the robot plan.

[0014] Another embodiment of the present invention comprises a robot platform. The robot platform includes at least one perceptor for perceiving environmental variables of interest, at least one locomotor for providing mobility to the robot platform, and a system controller. The system controller executes a task-oriented autonomy system that receives instructions from a user interface using a communication interface. The instructions are for achieving at least one

task-oriented target and include at least one of intervention instructions, robot initiative instructions, and instructions for setting a task-oriented autonomy level. If the instructions are robot initiative instructions, the system controller develops a robot plan to achieve the at least one task-oriented target and performs the robot plan. If the instructions are intervention instructions, the system controller performs the intervention instructions from the user interface to achieve the task-oriented target and, if present, overrides the robot plan.

[0015] Another embodiment of the present invention comprises a robot control system. The robot control system includes a user computer and a robot platform. The user computer includes a memory and one or more processors for executing a user interface. The user interface is configured for positioning at least one task-oriented target in a map on the user interface representing an environment of the robot. The user interface also determines a task-oriented autonomy level of the robot responsive to a change in the task-oriented target. If the change in the task-oriented target is smaller than a change threshold, the user interface instructs the robot to achieve the task-oriented target using intervention instructions from the user interface. If the change in the task-oriented target is larger than the change threshold, the user interface instructs the robot to achieve the task-oriented target using robot initiative to determine a robot plan for achieving the task-oriented target and implementing the robot plan. The robot platform includes a locomotor and a system controller. If robot initiative instructions are received from the user computer, the system controller develops and performs the robot plan. If intervention instructions are received from the user computer, the robot platform performs the intervention instructions to achieve the task-oriented target and, if appropriate, overrides the robot plan.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] In the drawings, which illustrate what is currently considered to be the best mode for carrying out the invention:

[0017] FIG. 1 illustrates a representative robot platform embodiment of the present invention;

[0018] FIG. 2 illustrates a representative robot control environment including a plurality of robot platforms and a robot controller;

[0019] FIG. 3 is a software architecture diagram illustrating significant components of embodiments of the present invention;

[0020] FIG. 4 illustrates representative hardware abstractions of hardware modules that may be available on robot platforms;

[0021] FIG. 5 illustrates a robot abstraction level including robot attributes that may be available on robot platforms;

[0022] FIG. 6 illustrates a representative embodiment of how a range abstraction may be organized;

[0023] FIG. 7 illustrates an occupancy grid map that may be developed by embodiments of the present invention;

[0024] FIG. 8 illustrates representative robot behavioral components that may be available on robot platforms;

[0025] FIG. 9 illustrates representative cognitive conduct components that may be available on robot platforms;

[0026] FIG. 10A illustrates how tasks may be allocated between an operator and a robot according to embodiments of the present invention;

[0027] FIG. 10B illustrates various cognitive conduct, robot behaviors, robot attributes, and hardware abstractions that may be available at different levels of robot autonomy;

[0028] FIG. 11 illustrates a portion of representative processing that may occur in developing robot attributes and communicating those attributes;

[0029] FIG. 12 illustrates a representative example of communication paths between various hardware abstraction, robot abstraction, and environment abstractions;

[0030] FIG. 13 illustrates a representative example of communication paths between robot abstractions, environment abstractions, robot behaviors, and robot conduct;

[0031] FIG. 14 is a software flow diagram illustrating components of an algorithm for performing a guarded motion behavior;

[0032] FIG. 15 is a software flow diagram illustrating components of an algorithm for performing translational portions of an obstacle avoidance behavior;

[0033] FIG. 16 is a software flow diagram illustrating components of an algorithm for performing rotational portions of the obstacle avoidance behavior;

[0034] FIG. 17 is a software flow diagram illustrating components of an algorithm for performing a get unstuck behavior;

[0035] FIG. 18 is a software flow diagram illustrating components of an algorithm for performing a real-time occupancy change analysis behavior;

[0036] FIG. 19 is a block diagram of a robot system for implementing a virtual track for a robot, in accordance with an embodiment of the present invention;

[0037] FIG. 20 illustrates a user interface for designating a desired path representative of a virtual track for a robot, in accordance with an embodiment of the present invention;

[0038] FIG. 21 is a process diagram for configuring the desired path into a waypoint file for execution by a robot, in accordance with an embodiment of the present invention;

[0039] FIG. 22 illustrates a user interface for further processing the desired path into a program for execution by a robot, in accordance with an embodiment of the present invention;

[0040] FIG. 23 is a diagram illustrating transformation from a drawing file to a program or waypoint file, in accordance with an embodiment of the present invention;

[0041] FIG. 24 is a process diagram of a control process of a robot, in accordance with an embodiment of the present invention;

[0042] FIG. 25 is a flowchart of a method for implementing a virtual track for a robot, in accordance with an embodiment of the present invention;

[0043] FIG. 26 is a software flow diagram illustrating components of an algorithm for handling a waypoint follow behavior;

[0044] FIG. 27 is a software flow diagram illustrating components of an algorithm for performing translational portions of the waypoint follow behavior;

[0045] FIG. 28 is a software flow diagram illustrating components of an algorithm for performing rotational portions of the waypoint follow behavior;

[0046] FIG. 29 is a software flow diagram illustrating components of an algorithm for performing a follow conduct;

[0047] FIG. 30 is a user interface window showing controls that a user may need to operate in order to direct actions of a robot in a conventional user interface;

[0048] FIGS. 31A-31D are user interface windows showing task-oriented targets, a robot position, and a path plan in accordance with one or more embodiments of the present invention;

[0049] FIG. 32 is a user interface window showing task-oriented targets with user interface intervention of the robots behavior in accordance with one or more embodiments of the present invention;

[0050] FIG. 33 is a simplified flow chart illustrating components of a user interface algorithm for controlling seamless task-oriented autonomy adjustments in accordance with one or more embodiments of the present invention; and

[0051] FIG. 34 is a simplified flow chart illustrating components of a robot algorithm for controlling seamless task-oriented autonomy adjustments in accordance with one or more embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0052] In the following description, circuits and functions may be shown in block diagram form in order not to obscure the present invention in unnecessary detail. Conversely, specific circuit implementations shown and described are exemplary only and should not be construed as the only way to implement the present invention unless specified otherwise herein. Additionally, block definitions and partitioning of logic between various blocks is exemplary of a specific implementation. It will be readily apparent to one of ordinary skill in the art that the present invention may be practiced by numerous other partitioning solutions. For the most part, details concerning timing considerations and the like have been omitted where such details are not necessary to obtain a complete understanding of the present invention and are within the abilities of persons of ordinary skill in the relevant art.

[0053] In this description, some drawings may illustrate signals as a single signal for clarity of presentation and description. It will be understood by a person of ordinary skill in the art that the signal may represent a bus of signals, wherein the bus may have a variety of bit widths and the present invention may be implemented on any number of data signals including a single data signal.

[0054] Furthermore, in this description of the invention, reference is made to the accompanying drawings which form a part hereof, and in which is shown, by way of illustration, specific embodiments in which the invention may be practiced. The embodiments are intended to describe aspects of the invention in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized and changes may be made without departing from the scope of the present invention. The following detailed description is not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

[0055] Headings are included herein to aid in locating certain sections of detailed description. These headings should not be considered to limit the scope of the concepts described under any specific heading. Furthermore, concepts described in any specific heading are generally applicable in other sections throughout the entire specification.

1. HARDWARE ENVIRONMENT

[0056] FIG. 1 illustrates a representative robot platform 100 (may also be referred to herein as a robot system) including the present invention. A robot platform 100 may include a system controller 110 including a system bus 150 for operable coupling to one or more communication devices 155 operably coupled to one or more communication channels 160, one or more perceptrors 165, one or more manipulators 170, and one or more locomotors 175.

[0057] The system controller 110 may include a processor 120 operably coupled to other system devices by internal busses (122, 124). By way of example, and not limitation, the processor 120 may be coupled to a memory 125 through a memory bus 122. The system controller 110 may also include an internal bus 124 for coupling the processor 120 to various other devices, such as storage devices 130, local input devices 135, local output devices 140, and local displays 145.

[0058] Local output devices 140 may be devices such as speakers, status lights, and the like. Local input devices 135 may be devices such as keyboards, mice, joysticks, switches, and the like.

[0059] Local displays 145 may be as simple as light-emitting diodes indicating status of functions of interest on the robot platform 100, or may be as complex as a high resolution display terminal.

[0060] The communication channels 160 may be adaptable to both wired and wireless communication, as well as supporting various communication protocols. By way of example, and not limitation, the communication channels may be configured as a serial or parallel communication channel, such as, for example, USB, IEEE-1394, 802.11 a/b/g, cellular telephone, and other wired and wireless communication protocols.

[0061] The perceptrs 165 may include inertial sensors, thermal sensors, tactile sensors, compasses, range sensors, sonar, Global Positioning System (GPS), Ground Penetrating Radar (GPR), lasers for object detection and range sensing, imaging devices, and the like. Furthermore, those of ordinary skill in the art will understand that many of these sensors may include a generator and a sensor to combine sensor inputs into meaningful, actionable perceptions. For example, sonar perceptrs and GPR may generate sound waves or sub-sonic waves and sense reflected waves. Similarly, perceptrs including lasers may include sensors configured for detecting reflected waves from the lasers for determining interruptions or phase shifts in the laser beam.

[0062] Imaging devices may be any suitable device for capturing images, such as, for example, an infrared imager, a video camera, a still camera, a digital camera, a Complementary Metal Oxide Semiconductor (CMOS) imaging device, a charge coupled device (CCD) imager, and the like. In addition, the imaging device may include optical devices for modifying the image to be captured, such as, for example, lenses, collimators, filters, and mirrors. For adjusting the direction at which the imaging device is oriented, a robot platform 100 may also include pan and tilt mechanisms coupled to the imaging device. Furthermore, a robot platform 100 may include a single imaging device or multiple imaging devices.

[0063] The manipulators 170 may include vacuum devices, magnetic pickup devices, arm manipulators, scoops, grippers, camera pan and tilt manipulators, and the like.

[0064] The locomotors 175 may include one or more wheels, tracks, legs, rollers, propellers, and the like. For providing the locomotive power and steering capabilities, the locomotors 175 may be driven by motors, actuators, levers, relays and the like. Furthermore, perceptors 165 may be configured in conjunction with the locomotors 175, such as, for example, odometers and pedometers.

[0065] FIG. 2 illustrates a representative robot control environment including a plurality of robot platforms (100A, 100B, and 100C) and a robot controller 180. The robot controller 180 may be a remote computer executing a software interface from which an operator may control one or more robot platforms (100A, 100B, and 100C) individually or in cooperation. The robot controller 180 may communicate with the robot platforms (100A, 100B, and 100C), and the robot platforms (100A, 100B, and 100C) may communicate with each other, across the communication channels 160. While FIG. 2 illustrates one robot controller 180 and three robot platforms (100A, 100B, and 100C) those of ordinary skill in the art will recognize that a robot control environment may include one or more robot platforms 100 and one or more robot controllers 180. In addition, the robot controller 180 may be a version of a robot platform 100.

[0066] Software processes illustrated herein are intended to illustrate representative processes that may be performed by the robot platform 100 or robot controller 180. Unless specified otherwise, the order in which the processes are described is not intended to be construed as a limitation. Furthermore, the processes may be implemented in any suitable hardware, software, firmware, or combinations thereof. By way of example, software processes may be stored on the storage device 130, transferred to the memory 125 for execution, and executed by the processor 120.

[0067] When executed as firmware or software, the instructions for performing the processes may be stored on a computer readable medium (i.e., storage device 130). A computer readable medium includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact disks), DVDs (digital versatile discs or digital video discs), and semiconductor devices such as RAM, DRAM, ROM, EPROM, and Flash memory.

2. GENERIC ROBOT ABSTRACTION ARCHITECTURE

[0068] Conventionally, robot architectures have been defined for individual robots and generally must be rewritten or modified to work with different sensor suites and robot platforms.

This means that adapting the behavior functionality created for one robot platform to a different robot platform is problematic. Furthermore, even architectures that propose a hardware abstraction layer to create a framework for accepting various hardware components still may not create a robot abstraction layer wherein the abstractions presented for high level behavioral programming are in terms of actionable components or generic robot attributes rather than the hardware present on the robot.

[0069] A notable aspect of the present invention is that it collates the sensor data issued from hardware or other robotic architectures into actionable information in the form of generic precepts. Embodiments of the present invention may include a generic robot architecture (GRA), which comprises an extensible, low-level framework, which can be applied across a variety of different robot hardware platforms, perceptor suites, and low-level proprietary control application programming interfaces (APIs). By way of example, some of these APIs may be Mobility, Aria, Aware, Player, etc.).

[0070] FIG. 3 is a software architecture diagram illustrating significant components of the GRA as a multi-level abstraction. Within the GRA, various levels of abstraction are available for use in developing robot behavior. The object oriented structure of the GRA may be thought of as including two basic levels. As is conventional in object oriented class structures, each subsequent level inherits all of the functionality of the higher levels.

[0071] At the lower level, the GRA includes a hardware abstraction level, which provides for portable, object oriented access to low level hardware perception and control modules that may be present on a robot. The hardware abstraction level is reserved for hardware specific classes and includes, for example, implementations for the actual robot geometry and sensor placement on each robot type.

[0072] Above the hardware abstraction level, the GRA includes a robot abstraction level, which provides atomic elements (i.e., building blocks) of generic robot attributes and develops a membrane between the low level hardware abstractions and controls. This membrane is based on generic robot attributes, or actionable components, which include robot functions, robot perceptions, and robot status. Each generic robot attribute may utilize a variety of hardware abstractions, and possibly other robot attributes, to accomplish its individual function.

[0073] The robot abstraction level may include implementations that are generic to given proprietary low level APIs. Examples of functions in this class level include the interface calls for a variety of atomic level robot behaviors such as, for example, controlling motion and reading sonar data.

[0074] The GRA enables substantially seamless porting of behavioral intelligence to new hardware platforms and control APIs by defining generic robot attributes and actionable components to provide the membrane and translation between behavioral intelligence and the hardware. Once a definition for a robot in terms of platform geometries, sensors, and API calls has been specified, behavior and intelligence may be ported in a substantially seamless manner for future development. In addition, the object oriented structure enables straightforward extension of the architecture for defining new robot platforms as well as defining low-level abstractions for new perceptors, motivators, communications channels, and manipulators.

[0075] The GRA includes an interpreter such that existing and new robot behaviors port in a manner that is transparent to both the operator and the behavior developer. This interpreter may be used to translate commands and queries back and forth between the operator and robot with a common interface, which can then be used to create perceptual abstractions and behaviors. When the "common language" supported by the GRA is used by robot developers, it enables developed behaviors and functionality to be interchangeable across multiple robots. In addition to creating a framework for developing new robot capabilities, the GRA interpreter may be used to translate existing robot capabilities into the common language so that the behavior can then be used on other robots. The GRA is portable across a variety of platforms and proprietary low level APIs. This is done by creating a standard method for commanding and querying robot functionality that exists on top of any particular robot manufacturers control API. Moreover, unlike systems where behavior stems from sensor data the GRA facilitates a consistent or predictable behavior output regardless of robot size or type by categorizing robot and sensor data into perceptual abstractions from which behaviors can be built.

[0076] The Generic Robot Architecture also includes a scripting structure for orchestrating the launch of the different servers and executables that may be used for running the GRA on a particular robot platform. Note that since these servers and executables (e.g., laser server, camera server, and base platform application) will differ from robot to robot, the scripting structure includes the ability to easily specify and coordinate the launch of the files that may be needed for specific applications. In addition, the scripting structure enables automatic launching of the system at boot time so that the robot is able to exhibit functionality without any operator involvement (i.e., no need for a remote shell login).

[0077] The Generic Robot Architecture may access configuration files created for each defined robot type. For example, the configuration files may specify what sensors, actuators, and API are being used on a particular robot. Use of the scripting structure together

with the configuration enables easy reconfiguration of the behaviors and functionality of the robot without having to modify source code (i.e., for example, recompile the C / C++ code).

[0078] The GRA keeps track of which capabilities are available (e.g., sensors, actuators, mapping systems, communications) on the specific embodiment and uses virtual and stub functions within the class hierarchy to ensure that commands and queries pertaining to capabilities that an individual robot does not have do not cause data access errors. For example, in a case where a specific capability, such as a manipulator, does not exist, the GRA returns special values indicating to the high-level behavioral control code that the command cannot be completed or that the capability does not exist. This makes it much easier to port seamlessly between different robot types by allowing the behavior code to adapt automatically to different robot configurations.

[0079] The above discussion of GRA capabilities has focused on the robot-oriented aspects of the GRA. However, the robot-oriented class structure is only one of many class structures included in the GRA. For example, the GRA also includes multi-tiered class structures for communication, range-sensing, cameras, and mapping. Each one of these class structures is set up to provide a level of functional modularity and allow different sensors and algorithms to be used interchangeably. By way of example, and not limitation, without changing the behavioral code built on the GRA at the robot behavior level, it may be possible to swap various mapping and localization systems or cameras and yet achieve the same functionality simply by including the proper class modules at the hardware abstraction level and possibly at the robot abstraction level. Additional capabilities and features of each of the levels of the GRA are discussed below.

2.1. HARDWARE ABSTRACTION LEVEL

[0080] FIG. 4 illustrates the hardware abstraction level 210, which includes representative hardware abstractions of hardware modules that may be available on a robot platform. These hardware abstractions create an object oriented interface between the software and hardware that is modular, reconfigurable, and portable across robot platforms. As a result, a software component can create a substantially generic hook to a wide variety of hardware that may perform a similar function. It will be readily apparent to those of ordinary skill in the art that the modules shown in FIG. 4 are a representative, rather than comprehensive example of hardware abstractions. Some of these hardware abstractions include; action abstractions 212(also referred to as manipulation abstractions) for defining and controlling manipulation type devices on the robot, communication abstractions 214 for defining and controlling

communication media and protocols, control abstractions 216 (also referred to as locomotion abstractions) for defining and controlling motion associated with various types of locomotion hardware, and perception abstractions 218 for defining and controlling a variety of hardware modules configured for perception of the robot's surroundings and pose (i.e., position and orientation).

2.1.1. MANIPULATION ABSTRACTIONS

[0081] Action device abstractions 212 may include, for example, vacuum devices, magnetic pickup devices, arm manipulators, scoops, grippers, camera pan and tilt manipulators, and the like.

2.1.2. COMMUNICATIONS ABSTRACTIONS

[0082] The Communication abstractions present substantially common communications interfaces to a variety of communication protocols and physical interfaces. The communication channels 160 may be adaptable to both wired and wireless communication, as well as supporting various communication protocols. By way of example, and not limitation, the communication abstractions may be configured to support serial and parallel communication channels, such as, for example, USB, IEEE-1394, 802.11 a/b/g, cellular telephone, and other wired and wireless communication protocols.

2.1.3. LOCOMOTION ABSTRACTIONS

[0083] Locomotion abstractions 216 may be based on robot motion, not necessarily on specific hardware components. For example, and not limitation, motion control abstractions may include drive, steering, power, speed, force, odometry, and the like. Thus, the motion abstractions can be tailored to individual third party drive controls at the hardware abstraction level and effectively abstracted away from other architectural components. In this manner, support for motion control of a new robot platform may comprise simply supplying the APIs which control the actual motors, actuators and the like, into the locomotion abstraction framework.

2.1.4. PERCEPTION ABSTRACTIONS

[0084] The perception abstractions 218 may include abstractions for a variety of perceptive hardware useful for robots, such as, for example, inertial measurements, imaging devices, sonar measurements, camera pan/tilt abstractions, GPS and iGPS abstractions, thermal sensors, infrared sensors, tactile sensors, laser control and perception abstractions, GPR, compass measurements, EMI measurements, and range abstractions.

2.2. ROBOT ABSTRACTION LEVEL

[0085] While the hardware abstraction level 210 focuses on a software model for a wide variety of hardware that may be useful on robots, the robot abstraction level 230 (as illustrated in FIGS. 3 and 5) focuses on generic robot attributes. The generic robot attributes enable building blocks for defining robot behaviors at the robot behavior level and provide a membrane for separating the definition of robot behaviors from the low level hardware abstractions. Thus, the each robot attributes may utilize one or more hardware abstractions to define its attribute. These robot attributes may be thought of as actionable abstractions. In other words, a given actionable abstraction may fuse multiple hardware abstractions that provide similar information into a data set for a specific robot attribute. For example, and not limitation, the generic robot attribute of “range” may fuse range data from hardware abstractions of an IR sensor and a laser sensor to present a single coherent structure for the range attribute. In this way, the GRA presents robot attributes as building blocks of interest for creating robot behaviors such that, the robot behavior can use the attribute to develop a resulting behavior (e.g., stop, slow down, turn right, turn left, etc).

[0086] Furthermore, a robot attribute may combine information from dissimilar hardware abstractions. By way of example, and not limitation, the position attributes may fuse information from a wide array of hardware abstractions, such as: perception modules like video, compass, GPS, laser, and sonar; along with control modules like drive, speed, and odometry. Similarly, a motion attribute may include information from position, inertial, range, and obstruction abstractions.

[0087] This abstraction of robot attributes frees the developer from dealing with individual hardware elements. In addition, each robot attribute can adapt to the amount, and type of information it incorporates into the abstraction based on what hardware abstractions may be available on the robot platform.

[0088] The robot attributes, as illustrated in FIG. 5, are defined at a relatively low level of atomic elements that include attributes of interest for a robot’s perception, status, and control. Some of these robot attributes include; robot health 232, robot position 234, robot motion 236, robot bounding shape 238, environmental occupancy grid 240, and range 242. It will be readily apparent to those of ordinary skill in the art that the modules shown in FIG. 5 are a representative, rather than comprehensive, example of robot attributes. Note that the term robot attributes is used somewhat loosely, given that robot attributes may include physical attributes such as health 232 and bounding shape 238 as well as how the robot perceives its environment, such as the environmental occupancy grid 240 and range attributes 242.

2.2.1. ROBOT HEALTH

[0089] The robot health abstractions 232 may include, for example, general object models for determining the status and presence of various sensors and hardware modules, determining the status and presence of various communication modules, determining the status of on board computer components.

2.2.2. ROBOT BOUNDING SHAPE

[0090] The robot bounding shape 238 abstractions may include, for example, definitions of the physical size and boundaries of the robot and definitions of various thresholds for movement that define a safety zone or event horizon, as is explained more fully below.

2.2.3. ROBOT MOTION

[0091] The robot motion abstractions 236 may include abstractions for defining robot motion and orientation attributes such as, for example, obstructed motion, velocity, linear and angular accelerations, forces, and bump into obstacle, and orientation attributes such as roll, yaw and pitch.

2.2.4. RANGE

[0092] The range abstractions 242 may include, for example, determination of range to obstacles from lasers, sonar, infra-red, and fused combinations thereof.

[0093] In more detail, FIG. 6 illustrates a representative embodiment of how a range abstraction may be organized. A variety of coordinate systems may be in use by the robot and an operator. By way of example, a local coordinate system may be defined by an operator relative to a space of interest (e.g., a building) or a world coordinate system defined by sensors such as a GPS unit, an iGPS unit, a compass, an altimeter, and the like. A robot coordinate system may be defined in Cartesian coordinates relative to the robots orientation such that, for example, the X axis is to the right, the Y axis is straight ahead, and the Z-axis is up. Another robot coordinate system may be cylindrical coordinates with a range, angle, and height relative to the robot's current orientation.

[0094] The range measurements for the representative embodiment illustrated in FIG. 6 are organized in a cylindrical coordinate system relative to the robot. The angles may be partitioned into regions covering the front, left, right and back of the robot and given names such as, for example, those used in FIG. 6.

[0095] Thus, regions in front may be defined and named as:
 Right_In_Front (310 and 310'), representing an angle between -15° and 15° ;
 Front 312, representing an angle between -45° and 45° ; and
 Min_Font_Dist 314, representing an angle between -90° and 90° .

[0096] Similarly, regions to the left side may be defined as:

Left_Side 320, representing an angle between 100° and 80° ;
 Left_Front 322, representing an angle between 60° and 30° ;
 Front_Left_Side 324, representing an angle between 70° and 50° ; and
 L_Front 326, representing an angle between 45° and 1° .

[0097] For the right side, regions may be defined as:

Right_Side 330, representing an angle between -100° and -80° ;
 Right_Front 332, representing an angle between -60° and -30° ;
 Front_Right_Side 334, representing an angle between -70° and -50° ; and
 R_Front 336, representing an angle between -45° and 0° .

[0098] While not shown, those of ordinary skill in the art will recognize that with the exception of the Left_Side 320 and Right_Side 330 regions, embodiments may include regions in the back, which are a mirror image of those in the front wherein the “Front” portion of the name is replaced with “Rear.”

[0099] Furthermore, the range attributes define a range to the closest object within that range. However, the abstraction of regions relative to the robot, as used in the range abstraction may, also be useful for many other robot attributes and robot behaviors that may require directional readings, such as, for example, defining robot position, robot motion, camera positioning, an occupancy grid map, and the like.

[00100] In practice, the range attributes may be combined to define a more specific direction. For example, directly forward motion may be defined as a geometrically adjusted combination of Right_In_Front 310, L_Front 326, R_Front 336, Front_Left_Side 324, and Front_right_side 334.

2.2.5. ROBOT POSITION AND ENVIRONMENTAL OCCUPANCY GRID MAPS

[00101] Returning to FIG. 5, the robot abstractions may include position attributes 234. Mobile robots may operate effectively only if they, or their operators, know where they are. Conventional robots may rely on real-time video and global positioning systems (GPS) as well as existing maps and floor plans to determine their location. However, GPS may not be reliable indoors and video images may be obscured by smoke or dust, or break up because of poor communications. Maps and floor plans may not be current and often are not readily available, particularly in the chaotic aftermath of natural, accidental or terrorist events. Consequently, real-world conditions on the ground often make conventional robots that rely on a priori maps ineffective.

[00102] Accurate positioning knowledge enables the creation of high-resolution maps and accurate path following, which may be needed for high-level deliberative behavior such as systematically searching or patrolling an area.

[00103] Embodiments of the present invention may utilize various mapping or localization techniques including positioning systems such as indoor GPS, outdoor GPS, differential GPS, theodolite systems, wheel-encoder information, and the like. To make robots more autonomous, embodiments of the present invention may fuse the mapping and localization information to build 3D maps on-the-fly that let robots understand their current position and an estimate of their surroundings. Using existing information, map details may be enhanced as the robot moves through the environment. Ultimately, a complete map containing rooms, hallways, doorways, obstacles and targets may be available for use by the robot and its human operator. These maps also may be shared with other robots or human first responders.

[00104] With the on-board mapping and positioning algorithm that accepts input from a variety of range sensors, the robot may make substantially seamless transitions between indoor and outdoor operations without regard for GPS and video drop-outs that occur during these transitions. Furthermore, embodiments of the present invention provide enhanced fault tolerance because they do not require off-board computing or reliance on potentially inaccurate or non-existent a priori maps.

[00105] Embodiments of the present invention may use localization methods by sampling range readings from scanning lasers and ultrasonic sensors and by reasoning probabilistically about where the robot is within its internal model of the world. The robot localization problem may be divided into two sub-tasks: global position estimation and local position tracking. Global position estimation is the ability to determine the robot's position in an a priori or previously learned map, given no information other than that the robot is somewhere in the region represented by the map. Once a robot's position has been found in the map, local tracking is the problem of keeping track of the robot's position over time and movement.

[00106] The robot's state space may be enhanced by localization methods such as Monte Carlo techniques and Markovian probability grid approaches for position estimation, as are well known by those of ordinary skill in the art. Many of these techniques provide efficient and substantially accurate mobile robot localization.

[00107] With a substantially accurate position for the robot determined, local tracking can maintain the robot's position over time and movement using dead-reckoning, additional global positioning estimation, or combinations thereof. Dead-reckoning is a method of navigation by keeping track of how far you have gone in any particular direction. For example,

dead reckoning would determine that a robot has moved a distance of about five meters at an angle from the current pose of about 37 degrees if the robot moves four meters forward, turns 90 degrees to the right, and moves forward three meters. Dead-reckoning can lead to navigation errors if the distance traveled in a given direction, or the angle through which a robot turns, is interpreted incorrectly. This can happen, for example, if one or more of the wheels on the robot spin in place when the robot encounters an obstacle.

[00108] Therefore, dead reckoning accuracy may be bolstered by sensor information from the environment, new global positioning estimates, or combinations thereof. With some form of a map, the robot can use range measurements to map features to enhance the accuracy of a pose estimate. Furthermore, the accuracy of a pose estimate may be enhanced by new range measurements (e.g., laser scans) into a map that may be growing in size and accuracy. In Simultaneous Localization and Mapping (SLAM), information from the robot's encoders and laser sensors may be represented as a network of probabilistic constraints linking the successive positions (poses) of the robot. The encoders may relate one robot pose to the next via dead-reckoning. To give further constraints between robot poses, the laser scans may be matched with dead-reckoning, including constraints for when a robot returns to a previously-visited area.

[00109] The robot abstractions may include environmental occupancy grid attributes 240. One form of map that may be useful from both the robot's perspective and an operator's perspective is an occupancy grid. An environmental occupancy grid, formed by an occupancy grid abstraction 240 is illustrated in FIG. 7. In forming an occupancy grid 250, a robot coordinate system may be defined in Cartesian coordinates relative to the robots orientation such that, for example, the X axis is to the right, the Y axis is straight ahead, and the Z-axis is up. Another robot coordinate system may be cylindrical coordinates with a range, angle, and height relative to the robot current orientation. Furthermore, occupancy grids may be translated to other coordinate systems for use by an operator.

[00110] An occupancy grid map 390 may be developed by dividing the environment into a discrete grid of occupancy cells 395 and assigning a probability to each grid indicating whether the grid is occupied by an object. Initially, the occupancy grid may be set so that every occupancy cell is set to an initial probability. As the robot scans the environment, range data developed from the scans may be used to update the occupancy grid. For example, based on range data, the robot may detect an object at a specific orientation and range away from the robot. This range data may be converted to a different coordinate system (e.g., local or world Cartesian coordinates). As a result of this detection, the robot may increase the probability that the particular occupancy cell is occupied and decrease the probability that occupancy cells

between the robot and the detected object are occupied. As the robot moves through its environment, new horizons may be exposed to the robots sensors, which enable the occupancy grid to be expanded and enhanced. To enhance map building and localization even further, multiple robots may explore an environment and cooperatively communicate their map information to each other or a robot controller to cooperatively build a map of the area.

[00111] The example occupancy grid map 390 as it might be presented to an operator is illustrated in FIG. 7. The grid cells 395 can be seen as small squares on this occupancy grid 390. A robot path 380 is shown to illustrate how the robot may have moved through the environment in constructing the occupancy grid 390. Of course, those of ordinary skill in the art will recognize that, depending on the application and expected environment, the occupancy grid 390 may be defined in any suitable coordinate system and may vary in resolution (i.e., size of each occupancy cell). In addition, the occupancy grid 390 may include a dynamic resolution such that the resolution may start out quite coarse while the robot discovers the environment, then evolve to a finer resolution as the robot becomes more familiar with its surroundings.

3. ROBOTIC INTELLIGENCE KERNEL

[00112] A robot platform 100 may include a robot intelligence kernel (may also be referred to herein as intelligence kernel), which coalesces hardware components for sensing, motion, manipulation, and actions with software components for perception, communication, behavior, and world modeling into a single cognitive behavior kernel that provides intrinsic intelligence for a wide variety of unmanned robot platforms. The intelligence kernel architecture may be configured to support multiple levels of robot autonomy that may be dynamically modified depending on operating conditions and operator wishes.

[00113] The robot intelligence kernel (RIK) may be used for developing a variety of intelligent robotic capabilities. By way of example, and not limitation, some of these capabilities including visual pursuit, intruder detection and neutralization, security applications, urban reconnaissance, search and rescue, remote contamination survey, and countermeasure operations.

[00114] Referring back to the software architecture diagram of FIG. 3, the RIK comprises a multi-level abstraction including a robot behavior level 250 and a cognitive level 270. The RIK may also include the robot abstraction level 230 and the hardware abstraction level 210 discussed above.

[00115] Above the robot abstraction level 230, the RIK includes the robot behavior level 270, which define specific complex behaviors that a robot, or a robot operator, may want to

accomplish. Each complex robot behavior may utilize a variety of robot attributes, and in some cases a variety of hardware abstractions, to perform the specific robot behavior.

[00116] Above the robot behavior level 250, the RIK includes the cognitive level 270, which provides cognitive conduct modules to blend and orchestrate the asynchronous events from the complex robot behaviors and generic robot behaviors into combinations of functions exhibiting cognitive behaviors, wherein high level decision making may be performed by the robot, the operator, or combinations of the robot and the operator.

[00117] Some embodiments of the RIK may include, at the lowest level, the hardware abstraction level 210, which provides for portable, object oriented access to low level hardware perception and control modules that may be present on a robot. These hardware abstractions have been discussed above in the discussion of the GRA.

[00118] Some embodiments of the RIK may include, above the hardware abstraction level 210, the robot abstraction level 230 including generic robot abstractions, which provide atomic elements (i.e., building blocks) of generic robot attributes and develop a membrane between the low level hardware abstractions and control based on generic robot functions. Each generic robot abstraction may utilize a variety of hardware abstractions to accomplish its individual function. These generic robot abstractions have been discussed above in the discussion of the GRA.

3.1. ROBOT BEHAVIORS

[00119] While the robot abstraction level 230 focuses on generic robot attributes, higher levels of the RIK may focus on; relatively complex robot behaviors at the robot behavior level 250, or on robot intelligence and operator collaboration at the cognitive level 270.

[00120] The robot behavior level 250 includes generic robot classes comprising functionality common to supporting behavior across most robot types. For example, the robot behavior level includes utility functions (e.g., Calculate angle to goal) and data structures that apply across substantially all robot types (e.g., waypoint lists). At the same time, the robot behavior level defines the abstractions to be free from implementation specifics such that the robot behaviors are substantially generic to all robots.

[00121] The robot behavior level 250, as illustrated in FIG. 8 robot, may be loosely separated into reactive behaviors 252 and deliberative behaviors 254. Of course, it will be readily apparent to those of ordinary skill in the art that the modules shown in FIG. 8 are a representative, rather than comprehensive, example of robot behaviors.

[00122] The reactive behaviors 252 may be characterized as behaviors wherein the robot reacts to its perception of the environment based on robot attributes, hardware abstractions, or combinations thereof. Some of these reactive behaviors may include; autonomous navigation, obstacle avoidance, guarded motion, visual tracking, laser tracking, get-unstuck behavior, and reactive planning. As examples, and not limitations, details regarding some of these behaviors are discussed in the section below regarding application specific behaviors.

[00123] In contrast, deliberative behaviors 254 may be characterized as behaviors wherein the robot may need to make decisions on how to proceed based on the results of the reactive behaviors, information from the robot attributes and hardware abstractions, or combinations thereof. Some of these deliberative behaviors may include; waypoint navigation with automatic speed adjustment, global path planning, and occupancy change detection. As examples, and not limitations, details regarding some of these behaviors are discussed in the section below regarding application specific behaviors.

3.2. COGNITIVE CONDUCT

[00124] The cognitive conduct level 270, as illustrated in FIG. 9, represents the highest level of abstraction, wherein significant robot intelligence may be built in to cognitive conduct modules, as well as significant operator-robot collaboration to perform complex tasks requiring enhanced robot initiative 299. Cognitive conduct modules blend and orchestrate asynchronous firings from the reactive behaviors 252, deliberative behaviors 254, and robot attributes 230 into intelligent robot conduct. Cognitive conduct modules may include conduct such as GoTo 272, wherein the operator may simply give a coordinate for the robot to go to and the robot takes the initiative to plan a path and get to the specified location. This GoTo conduct 272 may include a combination of robot behaviors 250, robot attributes 230, and hardware abstractions 210, such as, for example, obstacle avoidance, get-unstuck, reactive path planning, deliberative path planning, and waypoint navigation.

[00125] Another representative cognitive conduct module is human detection and pursuit 274, wherein the robot may react to changes in the environment and pursue those changes. This detection and pursuit conduct 274 may also include pursuit of other objects, such as, for example, another robot. The detection and pursuit 274 conduct may include a combination of robot behaviors 250, robot attributes 230, and hardware abstractions 210, such as, for example, occupancy change detection, laser tracking, visual tracking, deliberative path planning, reactive path planning, and obstacle avoidance.

[00126] Other representative cognitive conduct modules include conduct such as exploration and reconnaissance conduct 276 combined with map building, leader/follower conduct 278, and search and identify conduct 280.

[00127] Of course, it will be readily apparent to those of ordinary skill in the art that the cognitive conduct modules shown in FIG. 9 are a representative, rather than comprehensive example of robot conduct that may be implemented using embodiments of the present invention.

3.3. TIMING AND BEHAVIOR ADAPTATION

[00128] A notable aspect of the RIK is that the cognitive conduct modules 270 and robot behaviors 250 generally operate from a perception of speed of motion in relationship to objects and obstacles. In other words, rather than being concerned with spatial horizons and the distance away from an object, the cognitive conduct 270 and robot behaviors 250 are largely concerned with temporal horizons and how soon the robot may encounter an object. This enables defining the cognitive conduct 270 and robot behaviors 250 in a relativistic sense wherein, for example, the modules interpret motion as an event horizon wherein the robot may only be concerned with obstacles inside the event horizon. For example, a robot behavior 250 is not necessarily concerned with an object that is 10 meters away. Rather, the robot behavior 250 may be concerned that it may reach the object in two seconds. Thus, the object may be within the event horizon when the object is 10 meters away and the robot is moving toward it at 5 meters/sec, whereas if the object is 10 meters away and the robot is moving at 2 meters/second, the object may not be within the event horizon.

[00129] This relativistic perception enables an adaptation to processing power and current task load. If the robot is very busy, for example processing video, it may need to reduce its frequency of processing each task. In other words, the amount of time to loop through all the cognitive conduct 270 and robot behaviors 250 may increase. However, with the RIK, the cognitive conduct 270 and robot behaviors 250 can adapt to this difference in frequency by modifying its behaviors. For example, if the time through a loop reduces from 200 Hz to 100 Hz, the behaviors and conducts will know about this change in loop frequency and may modify the way it makes a speed adjustment to avoid an object. For example, the robot may need a larger change in its speed of motion to account for the fact that the next opportunity to adjust the speed is twice more distant in the future at 100 Hz than it would be at 200 Hz. This becomes more apparent in the discussion below, regarding the guarded motion behavior.

[00130] To enable and control this temporal awareness, the RIK includes a global timing loop in which cognitive conduct 270 and robot behaviors 250 may operate. Using this

global timing loop, each module can be made aware of information such as, for example, average time through a loop minimum and maximum time through a loop, and expected delay for next timing tick.

[00131] With this temporal awareness, the robot tends to modify its behavior by adjusting its motion, and motion of its manipulators, relative to its surroundings rather than adjusting its position relative to a distance to an object. Of course, with the wide array of perceptors, the robot is still very much aware of its pose and position relative to its environment and can modify its behavior based on this positional awareness. However, with the RIK, the temporal awareness is generally more influential on the cognitive conduct modules and robot behaviors than the positional awareness.

3.4. DYNAMIC AUTONOMY

[00132] To enhance the operator/robot tradeoff of control, the intelligence kernel provides a dynamic autonomy structure, which is a decomposition of autonomy levels, allowing methods for shared control to permeate all levels of the multi-level abstraction. Furthermore, the intelligence kernel creates an object-oriented software architecture, which may require little or no source code changes when ported to other platforms and low-level proprietary controllers.

[00133] The dynamic autonomy structure of the RIK provides a multi-level harmonization between human intervention and robot initiative 299 across robot behaviors. As capabilities and limitations change for both the human and the robot due to workload, operator expertise, communication dropout, and other factors, the RIK architecture enables shifts from one level of autonomy to another. Consequently, the ability of the robot to protect itself, make decisions, and accomplish tasks without human assistance may enable increased operator efficiency.

[00134] FIGS. 10A and 10B are depictions of a representative embodiment of a dynamic autonomy structure illustrating different levels of interaction between operator intervention 291 and robot initiative 299. As referred to herein operator, or operator intervention 291, may include human operation via a remote computer in communication with the robot, remote operation by some other form of artificial intelligence operating on a remote computer in communication with the robot, or some combination thereof.

[00135] At the lowest level, referred to as teleoperation mode 293, the robot may operate completely under remote control and take no initiative to perform operations on its own. At the second level, referred to as safe mode 294, robot movement is dependent on manual control from a remote operator. However, in safe mode 294, the robot may be equipped with a

level of initiative that prevents the operator from causing the robot to collide with obstacles. At the third level, referred to as shared mode 295, the robot can relieve the operator from the burden of direct control. For example, the robot may use reactive navigation to find a path based on the robot's perception of the environment. Shared mode 295 provides for a balanced allocation of roles and responsibilities. The robot accepts varying levels of operator intervention 291 and may support dialogue through the use of scripted suggestions (e.g., "Path blocked! Continue left or right?") and other text messages that may appear within a graphical interface. At the fourth level, referred to as collaborative tasking mode 296, a high level of collaborative tasking may be developed between the operator and the robot using a series of high-level tasks such as patrol, search region or follow path. In collaborative tasking mode 296, operator intervention 291 occurs on the tasking level, while the robot manages most decision-making and navigation. At the highest level, referred to as autonomous mode 297, a robot may behave in a substantially autonomous manner, needing nothing more than being enabled by an operator and perhaps given a very high level command such as, for example, survey the area, or search for humans.

[00136] FIG. 10A illustrates a representative embodiment of how tasks may be allocated between the operator and the robot. For example, teleoperation mode 293 may be configured such that the operator defines tasks, supervises direction, motivates motion, and prevents collision, in such a way that the robot takes no initiative and the operator maintains control. In safe mode 294, the operator may still define tasks, supervise direction, and motivate motion, while allowing the robot to take the initiative to prevent collisions. In shared mode 295, the operator may still define tasks and supervise direction, while allowing the robot to motivate motion and prevent collisions. In collaborative tasking mode 296, the robot may possess the initiative to prevent collisions, motivate motion, and supervise direction, while relinquishing operator intervention 291 to define task goals. In autonomous mode 297, the robot's initiative may prevent collisions, motivate motion, supervise direction, and define task goals. Of course, those of ordinary skill in the art will recognize that this allocation of tasks between the operator and the robot is a representative allocation. Many other tasks and behaviors, and allocation of those tasks and behaviors, are contemplated within the scope of the present invention.

[00137] FIG. 10B illustrates various cognitive conduct, robot behaviors, robot attributes, and hardware abstractions that may be available at different levels of robot autonomy. In general, moving from the teleoperation mode 293 toward the autonomous mode 297 represents an increase in the amount of robot initiative 299 and a decrease in the amount of operator intervention 291. Conversely, moving from the autonomous mode 297 toward the teleoperation mode 293 represents a decrease in the amount of robot initiative 299 and an

increase in the amount of operator intervention 291. Of course, those of ordinary skill in the art will recognize that FIG. 10B is a representative sample of available conduct, behaviors, attributes, and hardware, as well as a representative allocation between autonomy levels. The RIK is configured such that many modules may operate across different levels of autonomy by modifying the amount of operator intervention 291, modifying the amount of robot initiative 299, or combinations thereof.

[00138] The autonomy levels are structured in the intelligence kernel such that each new level of autonomy is built on, and encompasses, the subsequent level. For example, a guarded motion mode processing (explained more fully below) may include the behavior and representational framework utilized by the teleoperation mode 293 processing, but also include additional levels of robot initiative 299 based on the various robot attributes (e.g., related to directional motion) created in response to the teleoperation mode 293. Shared mode 295 may include all of the functionality and direct control of safe mode 294, but also allows robot initiative 299 in response to the abstractions produced through the guarded motion mode processing (e.g., fused range abstractions created in response to the directional motion abstractions). In addition, the collaborative tasking mode 296 may initiate robot responses to the abstractions created in shared mode 295 processing such as recognition that a box canyon has been entered or that a communication link has been lost.

[00139] For a robotic system to gracefully accept a full spectrum of intervention possibilities, interaction issues cannot be handled merely as augmentations to a control system. Therefore, opportunities for operator intervention 291 and robot initiative 299 are incorporated as an integral part of the robot's intrinsic intelligence. Moreover, for autonomous capabilities to evolve, the RIK is configured such that a robot is able to recognize when help is needed from an operator, other robot, or combinations thereof and learn from these interactions.

[00140] As an example, in one representative embodiment, the robot includes a Sony CCD camera that can pan, tilt and zoom to provide visual feedback to the operator in the teleoperation mode 293. The robot may also use this camera with increased robot initiative 299 to characterize the environment and even conduct object tracking.

[00141] In this embodiment, the RIK provides a graduated process for the robot to protect itself and the environment. To do so, the RIK may fuse a variety of range sensor information. A laser range finder may be mounted on the front, and sonar perceptors may be located around the mid-section of the robot. The robot also may include highly sensitive bump strips around its perimeter that register whether anything has been touched. To protect the top of the robot, especially the cameras and mission-specific sensors placed on top of the robot,

infrared proximity sensors may be included to indicate when an object is less than a few inches from the robot. Additional infrared proximity sensors may be placed on the bottom of the robot and point ahead of the robot toward the ground in order to prevent the robot from traveling into open space (e.g., traveling off of a landing down a stairway). Together, these sensors provide a substantial field of protection around the robot and allow the operator to command the robot with increased confidence that the robot can take initiative to protect itself or its environment.

[00142] However, avoiding obstacles may be insufficient. Many adverse environments may include forms of uneven terrain such as rubble. The robot should be able to recognize and respond to these obstacles. Inertial sensors may be used to provide acceleration data in three dimensions. This inertial information may be fused with information from the wheel encoders giving velocity and acceleration of the wheels, and electrical current draw from the batteries, to produce a measure of "unexpected" resistance that may be encountered by the robot. As part of the dynamic autonomy, the operator may be able to choose to set a resistance limit that will automatically stop the robot once the specified threshold has been exceeded. The resistance limit may be useful not only for rough terrain, but also in situations when the operator needs to override the "safe motion" capabilities (based on the obstacle avoidance sensors) to do things like push chairs and boxes out of the way and push doors open.

[00143] In addition, the RIK enables operators to collaborate with mobile robots, by defining an appropriate level of discourse, including a shared vocabulary and a shared cognitive work space collaboratively constructed and updated on the fly through interaction with the real world. This cognitive work space could consist of terrain overlaid with semantic abstractions generated through autonomous recognition of environmental features with point-and-click operator validation and iconographic insertion of map entities. Real-time semantic maps constructed collaboratively by humans, ground robots and air vehicles could serve as the basis for a spectrum of mutual human-robot interactions including tasking, situation awareness, human-assisted perception and collaborative environmental "understanding." Thus, the RIK enables human-robot communication within the context of a mission based on shared semantic maps between the robotic system and the operator.

[00144] With reference to FIGS. 10A and 10B, additional details of the dynamic autonomy structure and corresponding operation modes can be discussed.

3.4.1. TELEOPERATION MODE

[00145] In teleoperation mode 293, the operator has full, continuous control of the robot at a low level. The robot takes little or no initiative except, for example, to stop after a specified time if it recognizes that communications have failed. Because the robot takes little or

no initiative in this mode, the dynamic autonomy implementation provides appropriate situation awareness to the operator using perceptual data fused from many different sensors. For example, a tilt sensor may provide data on whether the robot is in danger of overturning. Inertial effects and abnormal torque on the wheels (i.e., forces not associated with acceleration) are fused to produce a measure of resistance as when, for example, the robot is climbing over or pushing against an obstacle. Even in teleoperation mode 293, the operator may be able to choose to activate a resistance limit that permits the robot to respond to high resistance and bump sensors. Also, a specialized interface may provide the operator with abstracted auditory, graphical and textual representations of the environment and task.

[00146] Some representative behaviors and attributes that may be defined for teleoperation mode 293 include joystick operation, perceptor status, power assessment, and system status.

3.4.2. SAFE MODE

[00147] In safe mode 294, the operator directs movements of the robot, but the robot takes initiative to protect itself. In doing so, this mode frees the operator to issue motion commands with less regard to protecting the robot, greatly accelerating the speed and confidence with which the operator can accomplish remote tasks. The robot may assess its own status and surrounding environment to decide whether commands are safe. For example, the robot possesses a substantial self-awareness of its position and will attempt to stop its motion before a collision, placing minimal limits on the operator. In addition, the robot may be configured to notify the operator of environmental features (e.g., box canyon, corner, and hallway), immediate obstacles, tilt, resistance, etc. and also continuously assesses the validity of its diverse sensor readings and communication capabilities. In safe mode 294, the robot may be configured to refuse to undertake a task if it does not have the ability (i.e., sufficient power or perceptual resources) to safely accomplish it.

[00148] Some representative behaviors and attributes that may be defined for safe mode 294 include guarded motion, resistance limits, and bump sensing.

3.4.3. SHARED MODE

[00149] In shared mode 295, the robot may take the initiative to choose its own path, responds autonomously to the environment, and work to accomplish local objectives. This initiative is primarily reactive rather than deliberative. In terms of navigation, shared mode 295 may be configured such that the robot responds only to its local (e.g., a two second event horizon or a six meter radius), sensed environment. Although the robot may handle the low-level navigation and obstacle avoidance, the operator may supply intermittent input, often at the

robot's request, to guide the robot in general directions. For example, a "Get Unstuck" behavior enables the robot to autonomously extricate itself from highly cluttered areas that may be difficult for a remote operator to handle.

[00150] Some representative behaviors and attributes that may be defined for shared mode 295 include reactive planning, get unstuck behavior, and obstacle avoidance.

3.4.4. COLABORATIVE TASKING MODE

[00151] In collaborative tasking mode 297, the robot may perform tasks such as, for example, global path planning to select its own route, requiring no operator input except high-level tasking such as "follow that target" or "search this area" (perhaps specified by drawing a circle around a given area on the map created by the robot). For all these levels, the intelligence resides on the robot itself, such that off-board processing is unnecessary. To permit deployment within shielded structures, a customized communication protocol enables very low bandwidth communications to pass over a serial radio link only when needed. The system may use multiple and separate communications channels with the ability to reroute data when one or more connection is lost.

[00152] Some representative cognitive conduct and robot behaviors, and robot attributes that may be defined for collaborative tasking mode 296 include waypoint navigation, global path planning, go to behavior, retro-traverse behavior, area search behavior, and environment patrol.

3.4.5. AUTONOMOUS MODE

[00153] In autonomous mode 297, the robot may perform with minimal to no operator intervention 291. For behaviors in autonomous mode 297, the operator may simply give a command for the robot to perform. Other than reporting status to the operator, the robot may be free to plan paths, prioritize tasks, and carry out the command using deliberative behaviors defined by the robots initiative.

[00154] Some representative behaviors and attributes that may be defined for autonomous mode 297 include pursuit behaviors, perimeter surveillance, urban reconnaissance, human presence detection, geological surveys, radiation surveys, virtual rail behavior, countermine operations, and seeking improvised explosive devices.

3.5. RIK EXAMPLES AND COMMUNICATION

[00155] Conventionally, robots have been designed as extensions of human mobility and senses. Most seek to keep the human in substantially complete control allowing the operator, through input from video cameras and other on-board sensors, to guide the robot and

view remote locations. In this conventional "master-slave" relationship, the operator provides the intelligence and the robot is a mere mobile platform to extend the operator's senses. The object is for the operator, perched as it were on the robot's back, to complete some desired tasks. As a result, conventional robot architectures may be limited by the need to maintain continuous, high-bandwidth communications links with their operators to supply clear, real-time video images and receive instructions. Operators may find it difficult to visually navigate when conditions are smoky, dusty, poorly lit, completely dark or full of obstacles and when communications are lost because of distance or obstructions.

[00156] The Robot Intelligence Kernel enables a modification to the way humans and robots interact, from master-slave to a collaborative relationship in which the robot can assume varying degrees of autonomy. As the robot initiative 299 increases, the operator can turn his or her attention to the crucial tasks at hand (e.g., locating victims, hazards, dangerous materials; following suspects; measuring radiation and/or contaminant levels) without worrying about moment-to-moment navigation decisions or communications gaps.

[00157] The RIK places the intelligence required for high levels of autonomy within the robot. Unlike conventional designs, off-board processing is not necessary. Furthermore, the RIK includes low bandwidth communication protocols and can adapt to changing connectivity and bandwidth capabilities. By reducing or eliminating the need for high-bandwidth video feeds, the robot's real-world sensor information can be sent as compact data packets over low-bandwidth (<1 Kbs) communication links such as, for example, cell phone modems and long-range radio. The robot controller may then use these low bandwidth data packets to create a comprehensive graphical interface, similar to a computer game display, for monitoring and controlling the robot. Due to the low bandwidth needs enabled by the dynamic autonomy structure of the RIK, it may be possible to maintain communications between the robot and the operator over many miles and through thick concrete, canopy and even the ground itself.

[00158] FIG 11 illustrates a representative embodiment of the RIK processing of robot abstractions 300 and communications operations 350 for communicating information about cognitive conduct, robot behaviors, robot attributes, and hardware abstractions to the robot controller or other robots. The upper portion 300 of FIG. 11 illustrates the robot abstractions, and hardware abstractions that may be fused to develop robot attributes. In the embodiment of FIG. 11, a differential GPS 302, a GPS 304, wheel encoders 306 and inertial data 312 comprise hardware abstractions that may be processed by a Kalman filter 520. The robot attributes for mapping and localization 308 and localized pose 310 may be developed by including information from, among other things, the wheel encoders 306 and inertial data 312.

Furthermore, the localized pose 310 may be a function of the results from mapping and localization 308. As with the hardware abstractions, these robot attributes of mapping and localization 308 and localized pose 310 may be processed by a Kalman filter 520.

[00159] Kalman filters 520 are efficient recursive filters that can estimate the state of a dynamic system from a series of incomplete and noisy measurements. By way of example, and not limitation, many of the perceptors used in the RIK include an emitter/sensor combination, such as, for example, an acoustic emitter and a microphone array as a sensors. These perceptors may exhibit different measurement characteristics depending on the relative pose of the emitter and target and how they interact with the environment. In addition, to one degree or another, the sensors may include noise characteristics relative to the measured values. In robotic applications, Kalman filters 520 may be used in many applications for improving the information available from perceptors. As one example of many applications, when tracking a target, information about the location, speed, and acceleration of the target may include significant corruption due to noise at any given instant of time. However, in dynamic systems that include movement, a Kalman filter 520 may exploit the dynamics of the target, which govern its time progression, to remove the effects of the noise and get a substantially accurate estimate of the target's dynamics. Thus, a Kalman filter 520 can use filtering to assist in estimating the targets location at the present time, as well as prediction to estimate a targets location at a future time.

[00160] As a result of the Kalman filtering, or after being processed by the Kalman filter 520, information from the hardware abstractions and robot attributes may be combined to develop other robot attributes. As examples, the robot attributes illustrated in FIG. 11, include position 332, movement 334, obstruction 336, occupancy 338, and other abstractions 340.

[00161] With the robot attributes developed, information from these robot attributes may be available for other modules within the RIK at the cognitive level 270, the robot behavior level 250, and the robot abstraction level 230.

[00162] In addition, information from these robot attributes may be processed by the RIK and communicated to the robot controller or other robots, as illustrated by the lower portion of FIG. 11. Processing information from the robot conduct, behavior, and attributes, as well as information from hardware abstractions serves to reduce the required bandwidth and latency such that the proper information may be communicated quickly and concisely. Processing steps performed by the RIK may include a significance filter 352, a timing module, 354, prioritization 356, and bandwidth control 358.

[00163] The significance filter 352 may be used as a temporal filter to compare a time varying data stream from a given RIK module. By comparing current data to previous data, the

current data may not need to be sent at all or may be compressed using conventional data compression techniques such as, for example, run length encoding and Huffman encoding. Another example would be imaging data, which may use data compression algorithms such as Joint Photographic Experts Group (JPEG) compression and Moving Picture Experts Group (MPEG) compression to significantly reduce the needed bandwidth to communicate the information.

[00164] The timing module 354 may be used to monitor information from each RIK module to optimize the periodicity at which it may be needed. Some information may require periodic updates at a faster rate than others. In other words, timing modulation may be used to customize the periodicity of transmissions of different types of information based on how important it may be to receive high frequency updates for that information. For example, it may be more important to notify an operator, or other robot, of the robot's position more often than it would be to update the occupancy grid map.

[00165] The prioritization 356 operation may be used to determine which information to send ahead of other information based on how important it may be to minimize latency from when data is available to when it is received by an operator or another robot. For example, it may be more important to reduce latency on control commands and control queries relative to map data. As another example, in some cognitive conduct modules where there may be significant collaboration between the robot and an operator, or in teleoperation mode where the operator is in control, it may be important to minimize the latency of video information so that the operator does not perceive a significant time delay between what the robot is perceiving and when it is presented to the operator.

[00166] These examples illustrate that for prioritization 356, as well as the significance filter 352, the timing modulation 354, and the bandwidth control 358, communication may be task dependent and autonomy mode dependent. As a result, information that may be a high priority in one autonomy mode may receive a lower priority in another autonomy mode.

[00167] The bandwidth control operation may be used to limit bandwidth based on the communication channel's bandwidth and how much of that bandwidth may be allocated to the robot. An example here might include progressive JPEG wherein a less detailed (i.e., coarser) version of an image may be transmitted if limited bandwidth is available. For video, an example may be to transmit at a lower frame rate.

[00168] After the communication processing is complete, the resultant information may be communicated to, or from, the robot controller, or another robot. For example, the information may be sent from the robot's communication device 155, across the communication

link 160, to a communication device 185 on a robot controller, which includes a multi-robot interface 190.

[00169] FIGS. 12 and 13 illustrate a more general interaction between hardware abstractions, robot abstractions, environment abstractions, robot behaviors, and robot conduct. FIG. 12 illustrates general communication between the hardware abstractions associated with sensor data servers 210 (also referred to as hardware abstractions), the robot abstractions 230 (also referred to as robot attributes), and environment abstractions 239. Those of ordinary skill in the art will recognize that FIG. 12 is intended to show general interactions between abstractions in a representative embodiment and is not intended to show every interaction possible within the GRA and RIK. Furthermore, it is not necessary to discuss every line between every module. Some example interactions are discussed to show general issues involved and describe some items from FIG. 12 that may not be readily apparent from simply examining the drawing. Generally, the robot abstractions 230 may receive and fuse information from a variety of sensor data servers 210. For example, in forming a general abstraction about the robot's current movement attributes, the movement abstraction may include information from bump sensors, GPS sensors, wheel encoders, compass sensors, gyroscopic sensors, tilt sensors, and the current brake state.

[00170] Some robot attributes 230, such as the mapping and localization attribute 231 may use information from a variety of hardware abstractions 210 as well as other robot attributes 230. The mapping and localization attribute 231 may use sonar and laser information from hardware abstractions 210 together with position information and local position information to assist in defining maps of the environment, and the position of the robot on those maps. Line 360 is bold to indicate that the mapping and localization attribute 231 may be used by any or all of the environment abstractions 239. For example, the occupancy grid abstraction uses information from the mapping and localization attribute 231 to build an occupancy grid as is explained, among other places, above with respect to FIG. 7. Additionally, the robot map position attribute may use the mapping and localization attribute 231 and the occupancy grid attribute to determine the robot's current position within the occupancy grid.

[00171] Bold line 362 indicates that any or all of the robot abstractions 230 and environment abstractions 239 may be used at higher levels of the RIK such as the communications layer 350, explained above with respect to FIG. 11, and the behavior modulation 260, explained below with respect to FIG. 13.

[00172] FIG. 13 illustrates general communication between the robot abstractions 230 and environment abstractions 239 with higher level robot behaviors and cognitive conduct. As

with FIG. 12, those of ordinary skill in the art will recognize that FIG. 13 is intended to show general interactions between abstractions, behaviors, and conduct in a representative embodiment and is not intended to show every interaction possible within the GRA and RIK. Furthermore, it is not necessary to discuss every line between every module. Some example interactions are discussed to show general issues involved and describe some items from FIG. 13 that may not be readily apparent from simply examining the drawing.

[00173] As an example, the event horizon attribute 362 may utilize and fuse information from robot abstractions 230 such as range and movement. Information from the event horizon attribute 362 may be used by behaviors, such as, for example, the guarded motion behavior 500 and the obstacle avoidance behavior 600. Bold line 370 illustrates that the guarded motion behavior 500 and the obstacle avoidance behavior 600 may be used by a variety of other robot behaviors and cognitive conduct, such as, for example, follow/pursuit conduct, virtual rail conduct, countermine conduct, area search behavior, and remote survey conduct.

4. REPRESENTATIVE BEHAVIORS AND CONDUCT

[00174] The descriptions in this section illustrate representative embodiments of robot behaviors and cognitive conduct that may be included in embodiments of the present invention. Of course, those of ordinary skill in the art will recognize these robot behaviors and cognitive conduct are illustrative embodiments and are not intended to be a complete list or complete description of the robot behaviors and cognitive conduct that may be implemented in embodiments of the present invention.

[00175] In general, in the flow diagrams illustrated herein, T indicates an angular velocity of either the robot or a manipulator and V indicates a linear velocity. Also, generally, T and V are indicated as a percentage of a predetermined maximum. Thus V=20% indicates 20% of the presently specified maximum velocity (which may be modified depending on the situation) of the robot or manipulator. Similarly, T=20% indicates 20% of the presently specified maximum angular velocity of the robot or manipulator. It will be understood that the presently specified maximums may be modified over time depending on the situations encountered. In addition, those of ordinary skill in the art will recognize that the values of linear and angular velocities used for the robot behaviors and cognitive conduct described herein are representative of a specific embodiment. While this specific embodiment may be useful in a wide variety of robot platform configurations, other linear and angular velocities are contemplated within the scope of the present invention.

[00176] Furthermore, those of ordinary skill in the art will recognize that the use of velocities, rather than absolute directions, is enabled largely by the temporal awareness of the robot behaviors and cognitive conduct in combination with the global timing loop. This gives the robot behaviors and cognitive conduct an opportunity to adjust velocities on each timing loop, enabling smoother accelerations and decelerations. Furthermore, the temporal awareness creates a behavior of constantly moving toward a target in a relative sense, rather than attempting to move toward an absolute spatial point.

4.1. AUTONOMOUS NAVIGATION

[00177] Autonomous navigation may be a significant component for many mobile autonomous robot applications. Using autonomous navigation, a robot may effectively handle the task of traversing varied terrain while responding to positive and negative obstacles, uneven terrain, and other hazards. Embodiments of the present invention enable the basic intelligence necessary to allow a broad range of robotic vehicles to navigate effectively both indoors and outdoors.

[00178] Many proposed autonomous navigation systems simply provide GPS waypoint navigation. However, GPS can be jammed and may be unavailable indoors or under forest canopy. A more autonomous navigation system includes the intrinsic intelligence to handle navigation even when external assistance (including GPS and communications) has been lost. Embodiments of the present invention include a portable, domain-general autonomous navigation system which blends the responsiveness of reactive, sensor based control with the cognitive approach found through waypoint following and path planning. Through its use of the perceptual abstractions within the robot attributes of the GRA, the autonomous navigation system can be used with a diverse range of available sensors (e.g., range, inertial, attitude, bump) and available positioning systems (e.g., GPS, laser, RF, etc.).

[00179] The autonomous navigation capability may scale automatically to different operational speeds, may be configured easily for different perceptor suites and may be easily parameterized to be portable across different robot geometries and locomotion devices. Two notable aspects of autonomous navigation are a guarded motion behavior wherein the robot may gracefully adjust its speed and direction near obstacles without needing to come to a full stop and an obstacle avoidance behavior wherein the robot may successfully navigate around known obstacles in its environment. Guarded motion and obstacle avoidance may work in synergy to create an autonomous navigation capability that adapts to the robots currently perceived environment. Moreover, the behavior structure that governs autonomous navigation allows the

entire assembly of behaviors to be used not only for obstacles but for other aspects of the environment which require careful maneuvering such as Landmine detection.

[00180] The robot's obstacle avoidance and navigation behaviors are derived from a number of robot attributes that enable the robot to avoid collisions and find paths through dense obstacles. The reactive behaviors may be configured as nested decision trees comprising rules which "fire" based on combinations of these perceptual abstractions.

[00181] The first level of behaviors, which may be referred to as action primitives, provide the basic capabilities important to most robot activity. The behavior framework enables these primitives to be coupled and orchestrated to produce more complex navigational behaviors. In other words, combining action primitives may involve switching from one behavior to another, subsuming the outputs of another behavior or layering multiple behaviors. For example, when encountering a dense field of obstacles that constrain motion in several directions, the standard confluence of obstacle avoidance behaviors may give way to the high level navigational behavior "Get-Unstuck," as is explained more fully below. This behavior involves rules which, when activated in response to combinations of perceptual abstractions, switch between several lower level behaviors including "Turn-till-head-is-clear" and "Backout."

4.1.1. GUARDED MOTION BEHAVIOR

[00182] FIG. 14 is a software flow diagram illustrating components of an algorithm for the guarded motion behavior 500 according to embodiments of the present invention. Guarded motion may fuse information from a variety of robot attributes and hardware abstractions, such as, for example, motions attributes, range attributes, and bump abstractions. The guarded motion behavior 500 uses these attributes and abstractions in each direction (i.e., front, left, right, and back) around the robot to determine the distance to obstacles in all directions around the robot.

[00183] The need for guarded motion has been well documented in the literature regarding unmanned ground vehicles. A goal of guarded motion is for the robot to be able to drive at high speeds, either in response to the operator or software directed control through one of the other robot behaviors or cognitive conduct modules, while maintaining a safe distance between the vehicle and obstacles in its path. The conventional approach usually involves calculating this safe distance as a product of the robot's speed. However, this means that the deceleration and the distance from the obstacle at which the robot will actually stop may vary based on the low level controller responsiveness of the low level locomotor controls and the physical attributes of the robot itself (e.g., wheels, weight, etc.). This variation in stopping speed

and distance may contribute to confusion on the part of the operator who may perceive inconsistency in the behavior of the robot.

[00184] The guarded motion behavior according to embodiments of the present invention enables the robot to come to a stop at a substantially precise, specified distance from an obstacle regardless of the robot's initial speed, its physical characteristics, and the responsiveness of the low-level locomotor control schema. As a result, the robot can take initiative to avoid collisions in a safe and consistent manner.

[00185] In general, the guarded motion behavior uses range sensing (e.g., from laser, sonar, infrared, or combinations thereof) of nearby obstacles to scale down its speed using an event horizon calculation. The event horizon determines the maximum speed the robot can safely travel and still come to a stop, if needed, at a specified distance from the obstacle. By scaling down the speed by many small increments, perhaps hundreds of times per second, it is possible to insure that regardless of the commanded translational or rotational velocity, guarded motion will stop the robot at substantially the same distance from an obstacle. As an example, if the robot is being driven near an obstacle rather than directly towards it, guarded motion will not stop the robot, but may slow its speed according to the event horizon calculation. This improves the operator's ability to traverse cluttered areas and limits the potential for operators to be frustrated by robot initiative.

[00186] The guarded motion algorithm is generally described for one direction, however, in actuality it is executed for each direction. In addition, it should be emphasized that the process shown in FIG. 14 operates within the RIK framework of the global timing loop. Therefore, the guarded motion behavior 500 is re-entered, and executes again, for each timing loop.

[00187] To begin, decision block 510 determines if guarded motion is enabled. If not, control transitions to the end of the guarded motion behavior.

[00188] If guarded motion is enabled, control transfers to decision block 520 to test whether sensors indicate that the robot may have bumped into an obstacle. The robot may include tactile type sensors that detect contact with obstacles. If these sensors are present, their hardware abstractions may be queried to determine if they sense any contact. If a bump is sensed, it is too late to perform guarded motion. As a result, operation block 525 causes the robot to move in a direction opposite to the bump at a reduced speed that is 20% of a predefined maximum speed without turning, and then exits. This motion is indicated in operation block 525 as no turn (i.e., $T = 0$) and a speed in the opposite direction (i.e., $V = -20\%$).

[00189] If no bump is detected, control transfers to decision block 530 where a resistance limit determination is performed. This resistance limit measures impedance to motion that may be incongruous with normal unimpeded motion. In this representative embodiment, the resistance limit evaluates true if; the wheel acceleration equals zero, the force on the wheels is greater than zero, the robot has an inertial acceleration that is less than 0.15, and the resulting impedance to motion is greater than a predefined resistance limit. If this resistance limit evaluation is true, operation block 535 halts motion in the impeded direction, then exits. Of course, those of ordinary skill in the art will recognize that this is a specific implementation for an embodiment with wheels and a specific inertial acceleration threshold. Other embodiments, within the scope of the present invention, may include different sensors and thresholds to determine if motion is being impeded in any given direction based on that embodiment's physical configuration and method of locomotion.

[00190] If motion is not being impeded, control transfers to decision block 540 to determine if any obstacles are within an event horizon. An event horizon is calculated as a predetermined temporal threshold plus a speed adjustment. In other words, obstacles inside of the event horizon are obstacles that the robot may collide with at the present speed and direction. Once again, this calculation is performed in all directions around the robot. As a result, even if an obstacle is not directly in the robot's current path, which may include translational and rotational movement, it may be close enough to create a potential for a collision. As a result, the event horizon calculation may be used to decide whether the robot's current rotational and translational velocity will allow the robot time to stop before encroaching the predetermined threshold distance. If there are no objects sensed within the event horizon, there is no need to modify the robot's current motion and the algorithm exits.

[00191] If an obstacle is sensed within the event horizon, operation block 550 begins a "safety glide" as part of the overall timing loop to reduce the robot's speed. As the robot's speed is reduced, the event horizon, proportional to that the speed, is reduced. If the reduction is sufficient, the next time through the timing loop, the obstacle may no longer be within the event horizon even though it may be closer to the robot. This combination of the event horizon and timing loop enables smooth deceleration because each loop iteration where the event horizon calculation exceeds the safety threshold, the speed of the robot (either translational, rotational, or both) may be curtailed by a small percentage. This enables a smooth slow down and also enables the robot to proceed at the fastest speed that is safe. The new speed may be determined as a combination of the current speed and a loop speed adjustment. For example, and not limitation, $\text{New_speed} = \text{current_speed} * (0.75 - \text{loop_speed_adjust})$. The loop_speed_adjust variable may

be modified to compensate for how often the timing loop is executed and the desired maximum rate of deceleration. Of course, those of ordinary skill in the art will recognize that this is a specific implementation. While this implementation may encompass a large array of robot configurations, other embodiments within the scope of the present invention may include different scale factors for determining the new speed based on a robot's tasks, locomotion methods, physical attributes, and the like.

[00192] Next, decision block 560 determines whether an obstacle is within a danger zone. This may include a spatial measurement wherein the range to the obstacle in a given direction is less than a predetermined threshold. If not, there are likely no obstacles in the danger zone and the process exits.

[00193] If an obstacle is detected in the danger zone, operation block 570 stops motion in the current direction and sets a flag indicating a motion obstruction, which may be used by other attributes, behaviors or conduct.

[00194] As mentioned earlier, the guarded motion behavior 500 operates on a global timing loop. Consequently, the guarded motion behavior 500 will be re-entered and the process repeated on the next time tick of the global timing loop.

4.1.2. OBSTACLE AVOIDANCE BEHAVIOR

[00195] FIG. 15 is a software flow diagram illustrating components of an algorithm for the obstacle avoidance behavior that governs translational velocity 600 of the robot according to embodiments of the present invention. Similarly, FIG. 16 is a software flow diagram illustrating components of an algorithm for the obstacle avoidance behavior that governs rotational velocity 650 of the robot. Obstacle avoidance may fuse information from a variety of robot attributes and hardware abstractions, such as, for example, motion attributes, range attributes, and bump abstractions. In addition, the obstacle avoidance behavior may use information from other robot behaviors such as, for example, the guarded motion behavior and a get unstuck behavior. The obstacle avoidance behavior uses these attributes, abstractions, and behaviors to determine a translational velocity and a rotational velocity for the robot such that it can safely avoid known obstacles.

[00196] In general, the obstacle avoidance behavior uses range sensing (e.g., from laser, sonar, infrared, or combinations thereof) of nearby obstacles to adapt its translational velocity and rotation velocity using the event horizon determinations explained earlier with respect to the guarded motion behavior. As stated earlier, the obstacle avoidance behavior works with the guarded motion behavior as building blocks for full autonomous navigation. In

addition, it should be emphasized that the processes shown in FIGS. 15 and 16 operate within the RIK framework of the global timing loop. Therefore, the obstacle avoidance behavior is re-entered, and executes again, for each timing loop.

[00197] To begin the translational velocity portion of FIG. 15, decision block 602 determines if waypoint following is enabled. If so, control transfers out of the obstacle avoidance behavior to a waypoint following behavior, which is explained form fully below.

[00198] If waypoint following is not enabled, control transfers to decision block 604 to first test to see if the robot is blocked directly in front. If so, control transfers to operation block 606 to set the robot's translational speed to zero. Then, control transfers out of the translational velocity behavior and into the rotational velocity behavior so the robot can attempt to turn around the object. This test at decision block 604 checks for objects directly in front of the robot. To reiterate, the obstacle avoidance behavior, like most behaviors and conducts in the RIK, is temporally based. In other words, the robot is most aware of its velocity and whether objects are within an event horizon related to time until it may encounter an object. In the case of being blocked in front, the robot may not be able to gracefully slow down through the guarded motion behavior. Perhaps because the object simply appeared in front of the robot, without an opportunity to follow typical slow down procedures that may be used if an object is within an event horizon. For example, the object may be another robot or a human that has quickly moved in front of the robot so that the guarded motion behavior has not had an opportunity to be effective.

[00199] If nothing is blocking the robot in front, decision block 608 tests to see if a detection behavior is in progress. A detection behavior may be a behavior where the robot is using a sensor in an attempt to find something. For example, the countermines conduct is a detection behavior that is searching for landmines. In these types of detection behaviors, obstacle avoidance may want to approach much closer to objects, or may want to approach objects with a much slower speed to allow time for the detection function to operate. Thus, if a detection behavior is active, operation block 610 sets a desired speed variable based on detection parameters that may be important. By way of example, and not limitation, in the case of the countermines conduct this desired speed may be set as: $\text{Desired_Speed} = \text{Max_passover_rate} - (\text{Scan_amplitude}/\text{Scan_Speed})$. In this countermines conduct example, the Max_passover_rate may indicate a maximum desired speed for passing over the landmine. This speed may be reduced by other factors. For example, the $(\text{Scan_amplitude}/\text{Scan_Speed})$ term reduces the desired speed based on a factor of how fast the mine sensor sweeps an area. Thus, the Scan_amplitude term defines a term of the extent of the scan sweep and the Scan_Speed defines

the rate at which the scan happens. For example, with a large Scan_amplitude and a small Scan_Speed, the Desired_Speed will be reduced significantly relative to the Max_passover_rate to generate a slow speed for performing the scan. While countermeasure conduct is used as an example of a detection behavior, those of ordinary skill in the art will recognize that embodiments of the present invention may include a wide variety of detection behaviors, such as, for example, radiation detection, chemical detection, and the like.

[00200] If a detection behavior is not in progress, decision block 612 tests to see if a velocity limit is set. In some embodiments of the invention, it may be possible for the operator to set a velocity limit that the robot should not exceed, even if the robot believes it may be able to safely go faster. For example, if the operator is performing a detailed visual search, the robot may be performing autonomous navigation, while the operator is controlling a camera. The operator may wish to keep the robot going slow to have time to perform the visual search.

[00201] If a velocity limit is set, operation block 614 sets the desired speed variable relative to the velocity limit. The equation illustrated in operation block 614 is a representative equation that may be used. The 0.1 term is a term used to ensure that the robot continues to make very slow progress, which may be useful to many of the robot attributes, behaviors, and conduct. In this equation, the Speed_Factor term is a number from one to ten, which may be set by other software modules, for example the guarded motion behavior, to indicate a relative speed at which the robot should proceed. Thus, the desired speed is set as a fractional amount (between zero and one in 0.1 increments) of the Max_Limit_Speed.

[00202] If a velocity limit is not set, operation block 616 sets the desired speed variable relative to the maximum speed set for the robot (i.e., Max_Speed) with an equation similar to that for operation block 614 except Max_Speed is used rather than Max_Limit_Speed.

[00203] After the desired speed variable is set by block 610, 614, or 616, decision block 618 tests to see if anything is within the event horizon. This test may be based on the robot's physical dimensions, including protrusions from the robot such as an arm, relative to the robot's current speed. As an example using an arm extension, something inside the event horizon may be determined by the equation:

$$\text{Min_Front_Range} < 1.0 + \text{Arm_Extension} + (1.75 * \text{Abs}(\text{Current_Velocity}))$$

[00204] Where the Min_Front_Range indicates a range to an obstacle in front, 1.0 is a safety factor, Arm_Extension indicates the distance beyond the robot that the arm currently extends, and Current_Velocity indicates the robot's current translational velocity.

[00205] If there is something detected within the event horizon, operation block 620 sets the current speed based on the distance to the obstacle. Thus, the example equation in block

620 sets the speed based on the range to the object less a Forward_Threshold set as a safety factor. With this speed, guarded motion has an opportunity to be effective and the speed may be reduced further on the next iteration of the timing loop if the object is still within the event horizon. After setting the speed, control transfers out of the translational velocity behavior, and into the rotational velocity behavior.

[00206] If there is nothing detected within the event horizon, operation block sets the speed to the desired speed variable that was set previously by operation block 614, 616, or 618. After setting the speed, control transfers out of the translational velocity behavior, and into the rotational velocity behavior

[00207] To begin the rotation velocity portion of FIG. 16, decision block 602 determines if waypoint following is enabled. If so, control transfers out of the obstacle avoidance behavior to a waypoint following behavior, which is explained form fully below.

[00208] If waypoint following is not enabled, control transfers to decision block 604 to first test to see if the robot is blocked directly in front. If so, control transfers to operation block 606 to set the robot's translational speed to zero. Then, control transfers out of the translational velocity behavior and into the rotational velocity behavior so the robot can attempt to turn around the object. This test at decision block 604 checks for objects directly in front of the robot. To reiterate, the obstacle avoidance behavior, like most behaviors and conducts in the RIK, is temporally based. In other words, the robot is most aware of its velocity and whether objects are within an event horizon related to time until it may encounter an object. In the case of being blocked in front, the robot may not be able to gracefully slow down through the guarded motion behavior. Perhaps because the object simply appeared in front of the robot, without an opportunity to follow typical slow down procedures that may be used if an object is within an event horizon. For example, the object may be another robot or a human that has quickly moved in front of the robot so that the guarded motion behavior has not had an opportunity to be effective.

[00209] If nothing is blocking the robot in front, decision block 608 tests to see if a detection behavior is in progress. A detection behavior may be a behavior where the robot is using a sensor in an attempt to find something. For example, the countermeasure conduct is a detection behavior that is searching for landmines. In these types of detection behaviors, obstacle avoidance may want to approach much closer to objects, or may want to approach objects with a much slower speed to allow time for the detection function to operate. Thus, if a detection behavior is active, operation block 610 sets a desired speed variable based on detection parameters that may be important. By way of example, and not limitation, in the case of the

countermine conduct this desired speed may be set as: $\text{Desired_Speed} = \text{Max_passover_rate} - (\text{Scan_amplitude}/\text{Scan_Speed})$. In this countermine conduct example, the Max_passover_rate may indicate a maximum desired speed for passing over the landmine. This speed may be reduced by other factors. For example, the $(\text{Scan_amplitude}/\text{Scan_Speed})$ term reduces the desired speed based on a factor of how fast the mine sensor sweeps an area. Thus, the Scan_amplitude term defines a term of the extent of the scan sweep and the Scan_Speed defines the rate at which the scan happens. For example, with a large Scan_amplitude and a small Scan_Speed , the Desired_Speed will be reduced significantly relative to the Max_passover_rate to generate a slow speed for performing the scan. While countermine conduct is used as an example of a detection behavior, those of ordinary skill in the art will recognize that embodiments of the present invention may include a wide variety of detection behaviors, such as, for example, radiation detection, chemical detection, ground penetrating radar, and the like.

[00210] If a detection behavior is not in progress, decision block 612 tests to see if a velocity limit is set. In some embodiments of the invention, it may be possible for the operator to set a velocity limit that the robot should not exceed, even if the robot believes it may be able to safely go faster. For example, if the operator is performing a detailed visual search, the robot may be performing autonomous navigation, while the operator is controlling a camera. The operator may wish to keep the robot going slow to have time to perform the visual search.

[00211] If a velocity limit is set, operation block 614 sets the desired speed variable relative to the velocity limit. The equation illustrated in operation block 614 is a representative equation that may be used. The 0.1 term is a term used to ensure that the robot continues to make very slow progress, which may be useful to many of the robot attributes, behaviors, and conduct. In this equation, the Speed_Factor term is a number from one to ten, which may be set by other software modules, for example the guarded motion behavior, to indicate a relative speed at which the robot should proceed. Thus, the desired speed is set as a fractional amount (between zero and one in 0.1 increments) of the Max_Limit_Speed .

[00212] If a velocity limit is not set, operation block 616 sets the desired speed variable relative to the maximum speed set for the robot (i.e., Max_Speed) with an equation similar to that for operation block 614 except Max_Speed is used rather than Max_Limit_Speed .

[00213] After the desired speed variable is set by block 610, 614, or 616, decision block 618 tests to see if anything is within the event horizon. This test may be based on the robot's physical dimensions, including protrusions from the robot such as an arm, relative to the robot's current speed. As an example using an arm extension, something inside the event horizon may be determined by the equation:

$$\text{Min_Front_Range} < 1.0 + \text{Arm_Extension} + (1.75 * \text{Abs}(\text{Current_Velocity}))$$

[00214] Where the Min_Front_Range indicates a range to an obstacle in front, 1.0 is a safety factor, Arm_Extension indicates the distance beyond the robot that the arm currently extends, and Current_Velocity indicates the robot's current translational velocity.

[00215] If there is something detected within the event horizon, operation block 620 sets the current speed based on the distance to the obstacle. Thus, the example equation in block 620 sets the speed based on the range to the object less a Forward_Threshold set as a safety factor. With this speed, guarded motion has an opportunity to be effective and the speed may be reduced further on the next iteration of the timing loop if the object is still within the event horizon. After setting the speed, control transfers out of the translational velocity behavior, and into the rotational velocity behavior.

[00216] If there is nothing detected within the event horizon, operation block 622 sets the robot's current speed to the desired speed variable that was set previously by operation block 614, 616, or 618. After setting the speed, control transfers out of the translational velocity behavior 600, and into the rotational velocity behavior 650.

[00217] FIG. 16 illustrates a representative software flow diagram illustrating components of an algorithm for the obstacle avoidance behavior that governs rotational velocity behavior 650 of the robot. To begin the rotational velocity behavior of FIG. 16, decision block 652 determines if waypoint following is enabled. If so, control transfers to decision block 654 to determine if the angle to a target exceeds a predefined threshold.

[00218] At decision block 658, the process checks to see if the robot is blocked in front. If so, the process performs a series of checks to see where other obstacles may be to determine a desired rotational velocity and direction. This obstacle checking process begins with decision block 660 testing to see if the robot is blocked on the left side. If the robot is blocked on the left side, and also in front, operation block 662 sets a new value for a turn velocity to the right. In the representative embodiment illustrated in FIG. 16 a positive rotational velocity is defined as a turn to the left and a negative rotational velocity is defined as a turn to the right. Thus, generally, Turn_left is a positive value indicating a rotational velocity to the left and Turn_right is a negative value indicating a rotational velocity to the right. Thus, operation block 662 reduces the rotational velocity in the current direction by about one half plus a small offset used to ensure that the rotational velocity does not reach zero. After setting the new rotation velocity, the process exits.

[00219] If the robot is not blocked on the left, decision block 664 tests to see if the robot is blocked on the right. If so, operation block 666 sets a new value for a turn velocity to

the right similar to that velocity setting in operation block 662. In other words, set the rotational velocity to the left to about one half plus a small offset used to ensure that the rotational velocity does not reach zero. After setting the new rotation velocity, the process exits.

[00220] If the robot is blocked in the front, but not on the left or right, the process then decides which way to turn to get around the blockage by checking to see whether the nearest obstacle in a measurable range is to the right or left and adjusting the rotational velocity to be away from the obstacle. Operation block 668 checks to see if the nearest obstacle is to the left. If so, operation block 670 sets the rotational velocity to the right (i.e., away from the obstacle) at a velocity of 30% of a maximum defined rotational velocity. If the nearest obstacle is not to the left, operation block 672 sets the rotational velocity to the left at a velocity of 30% of a maximum defined rotational velocity. After setting the new rotation velocity by either operation block 670 or 672, the process exits.

[00221] If the robot was not blocked in front, based on decision block 658, then decision block 674 performs a “threading the needle process.” This starts with decision block 674 determining a range to obstacles that may still be in front of the robot but not directly blocking the robot. To do this, decision block 674 test to see if Min_Front_Range is greater than two times a predefined threshold for the front direction, and to see if Min_Narrow_Front is greater than two times the robot’s length. If both these tests are true, it may be relatively clear in front and the process decides to reduce the rotational velocity in the current direction to make the direction more straight ahead until the next global timing loop. Therefore, decision block 676 tests to see if the current rotational direction is left. If so, decision block 678 tests to see if the magnitude of the left rotational velocity is greater than twice a turn threshold. If so, operation block 680 reduces the rotational velocity in the left direction by one half, and the process exits. If the current rotational direction is not left, decision block 682 tests to see if the magnitude of the right rotational velocity is greater than twice a turn threshold. If so, operation block 684 reduces the rotational velocity in the right direction by one half, and the process exits.

[00222] If decision block 674, 678, or 682 evaluates false, decision block tests to see if anything is currently within the event horizon.

[00223] This test may be based on the robot’s physical dimensions, including protrusions from the robot such as an arm, relative to the robot’s current speed. In addition, this test is likely the same as the event horizon described above for the translational velocity when discussing decision block 618 on FIG. 15. In other words, is the Minimum_Front_Range less than an Event_Range. Wherein the $\text{Event_Range} = 1.0 + \text{Arm_Extension} + (1.75 * \text{Abs}(\text{Current_Velocity}))$.

[00224] If there is nothing within the event horizon (i.e., decision block 690 evaluates false), there is likely no need to change the current rotational velocity so the process exits. If there is something within the event horizon, but not within the threading the needle process or blocking the robot in front, the rotational velocity may be adjusted at a more gradual rate. Thus, if operation block 690 evaluates true, decision block 692 tests to see if the closest object is on the left side. If so, operation block 694 sets a new rotation velocity to the right. If the closest object is not on the left, operation block 696 sets a new rotation velocity to the left. The rotational velocity that is set in operation blocks 694 and 696 are similar except for direction. In this representative embodiment, the rotational velocity may be set as a function of the Event Range from the event horizon test of decision block 690. Thus, the rotational velocity may be set as:

$$(\text{Event_Range} - \text{Min_Front_Range}) / 4.$$

[00225] After setting the rotational velocity in either operation block 694 or 696, the process exits.

[00226] As mentioned earlier, the obstacle avoidance behavior 600 operates on the global timing loop. Consequently, both the translational velocity and rotational velocity may be adjusted again on the next time tick of the global timing loop, allowing for relatively quick periodic adjustments to the velocities.

4.2. GET UNSTUCK BEHAVIOR

[00227] A get unstuck behavior 700, as illustrated in FIG. 17, includes significant robot initiative to extricate itself from the stuck position with little or no help from the operator. Sometimes, when a robot is operating under its own initiative, or even under operator control, the robot may get stuck and have difficulty getting free from that position. Often times, the operator may have limited understanding of the robots position relative to the robots understanding with its wide variety of perceptors. In general, the get unstuck behavior 700 may use range sensing (e.g., from laser, sonar, infrared, or combinations thereof) to determine nearby obstacles and their position relative to the robot.

[00228] The get unstuck behavior 700 begins at decision block 710 by determining if the current path is blocked. This blocked situation may be defined as obstacle present in front, on the front-right side, and on the front-left side. If the path is blocked, control transfers to operation block 740, which is explained below. For an example, and using the range definitions defined above under the description of the range attribute, a blocked path may be defined by the Boolean equation:

[00229] Blocked =
 ((right_in_front < (robot->forward_thresh + 0.2)) || FRONT_BLOCKED) &&
 (l_front < (robot->forward_thresh * 2)) &&
 (r_front < (robot->forward_thresh * 2)) &&
 (left_front < (robot->forward_thresh * 2)) &&
 (right_front < (robot->forward_thresh * 2))

[00230] Wherein: (robot->forward_thresh) is a predetermined threshold parameter, that may be robot specific, to define a safety distance, or maneuverability distance, away from the robot.

[00231] If the path is not blocked, decision block 720 determines if forward motion and turning motion is obstructed. If motion is obstructed, control transfers to operation block 740, which is explained below. For an example, this motion obstruction may be determined by the Boolean equation:

[00232] Obstructed_motion =
 (FR_LEFT_BLOCKED || R_RIGHT_BLOCKED) &&
 (FR_RIGHT_BLOCKED || R_LEFT_BLOCKED) &&
 FRONT_BLOCKED

[00233] If motion is not obstructed, decision block 730 determines if the robot is in a box canyon. If the robot is not in a box canyon, the get unstuck behavior exits because it appears the robot is not in a stuck situation. If the robot is in a box canyon, control transfers to operation block 740. For an example, this box canyon situation may be defined by the Boolean equation:

[00234] Box_canyon =
 (right_in_front < (robot->forward_thresh+.2)) &&
 (right_front < (robot->forward_thresh * 2.0)) &&
 (left_front < (robot->forward_thresh * 2.0)) &&
 ((right_side + left_side) < (robot->turn_thresh * 3.0)) &&
 (BACK_BLOCKED=0)

[00235] Wherein: (robot-> turn _thresh) is a predetermined threshold parameter, which may be robot specific, to define a maneuverability distance that enables the robot to turn around.

[00236] Once the determination has been made that the robot may be stuck, operation block 740 begins the process of attempting to get unstuck. Operation block 740 performs a back-out behavior. This back-out behavior causes the robot to backup from its present position while following the contours of obstacles near the rear sides of the robot. In general, the back-

out behavior uses range sensing (e.g., from laser, sonar, infrared, or combinations thereof) of nearby obstacles near the rear sides to determine distance to the obstacles and provide assistance in following the contours of the obstacles. However, the back-out behavior may also include many robot attributes, including perception, position, bounding shape, and motion, to enable the robot to turn and back up while continuously responding to nearby obstacles. Using this fusion of attributes, the back-out behavior doesn't merely back the robot up, but rather allows the robot to closely follow the contours of whatever obstacles are around the robot.

[00237] For example movements, the robot may attempt to equalize the distance between obstacles on both sides, keep a substantially fixed distance from obstacles on the right side, or keep a substantially fixed distance between obstacles on the right side. As the back-out behavior progresses, decision block 750 determines if there is sufficient space on a side to perform a maneuver other than backing out. If there is not sufficient spaces, control transfers back to operation block 740 to continue the back-out behavior. If there is sufficient space on a side, control transfers to operation block 760. As an example, the sufficient space on a side decision may be defined by the Boolean equation:

[00238] $\text{Space_on_side} = \text{space_on_left} \parallel \text{space_on_right}$, wherein:

[00239] $\text{Space_on_left} =$
 $(l_front > (\text{robot} \rightarrow \text{forward_thresh} + .2)) \ \&\&$
 $(\text{turn_left} > (\text{robot} \rightarrow \text{arm_length} + \text{robot} \rightarrow \text{turn_thresh} + .2)) \ \&\&$
 $(\text{turn_left} \geq \text{turn_right})$

[00240] $\text{Space_on_right} =$
 $(r_front > (\text{robot} \rightarrow \text{forward_thresh} + .2)) \ \&\&$
 $(\text{turn_right} > (\text{robot} \rightarrow \text{arm_length} + \text{robot} \rightarrow \text{turn_thresh} + .2)) \ \&\&$
 $(\text{turn_right} \geq \text{turn_left})$

[00241] Once sufficient space has been perceived on the right or left, operation block 760 performs a turn-until-head-is-clear behavior. This behavior causes the robot to rotate in the sufficient space direction while avoiding obstacles on the front sides. As the turn-until-head-is-clear behavior progresses, decision block 770 determines if, and when, the head is actually clear. If the head is not clear, control transfers back to the operation block 760 to continue the turn-until-head-is-clear behavior. If the head is clear, control transfers to operation block 760.

[00242] Once the head is clear, decision block 780 determines whether an acceptable egress route has been found. This egress route may be defined as an acceptable window of open space that exists for the robot to move forward. To avoid potential cyclical behavior, the acceptable window may be adjusted such that the robot does not head back toward the blocked

path or box canyon. If an acceptable egress route has not been found, control transfers back to operation block 740 to attempt the back-out behavior again. If an acceptable egress route is found, the unstuck behavior exits. As a specific example, the window may be defined by the equation:

window = 1.25 meters - (seconds-in-behavior / 10.0); and the egress route may be defined as true if the window < (robot->forward_thresh * 2.5).

[00243] As with the guarded motion behavior, the get-unstuck behavior 700 operates on a global timing loop. Consequently, the get-unstuck behavior 700 will be re-entered and the process repeated on the next time tick.

4.3. REAL-TIME OCCUPANCY CHANGE ANALYSIS

[00244] FIG. 18 is a software flow diagram illustrating representative components of an algorithm for performing a real-time occupancy change analysis behavior 800. Despite the much discussed potential for robots to play a critical role in security applications, the reality is that many human presence and motion tracking techniques require that the sensor used in tracking be stationary, removing the possibility for placement on a mobile robot platform. In addition, there is a need to determine substantially accurate positions for changes to recognized environmental features within a map. In other words, it may not be enough to know that something has moved or even the direction of movement. For effective change detection, a system should provide a substantially accurate position of the new location.

[00245] The Real-Time Occupancy Change Analyzer (ROCA) algorithm compares the state of the environment to its understanding of the world and reports to an operator, or supporting robotic sensor, the position of and the vector to any change in the environment. The ROCA robot behavior 800 includes laser-based tracking and positioning capability which enables the robot to precisely locate and track static and mobile features of the environment using a change detection algorithm that continuously compares current laser scans to an occupancy grid map. Depending on the laser's range, the ROCA system may be used to detect changes up to 80 meters from the current position of the laser range finder. The occupancy grid may be given a priori by an operator, built on the fly by the robot as it moves through its environment, or built by a combination of robot and operator collaboration. Changes in the occupancy grid may be reported in near real-time to support a number of tracking capabilities, such as camera tracking or a robotic follow capability wherein one or more robots are sent to the map location of the most recent change. Yet another possible use for the ROCA behavior is for target acquisition.

[00246] A notable aspect of the ROCA behavior is that rather than only providing a vector to the detected change, it provides the actual X, Y position of the change. Furthermore, the ROCA behavior can operate "on-the-move" meaning that unlike most human presence detection systems which must be stationary to work properly, it can detect changes in the features of the environment around it apart from of its own motion. This position identification and on-the-move capability enable tracking systems to predict future movement of the target and effectively search for a target even if it becomes occluded.

[00247] In general, once the robot has identified a change, the change may be processed by several algorithms to filter the change data to remove noise and cluster the possible changes. Of the clustered changes identified, the largest continuous cluster of detected changes (i.e., "hits") may be defined as locations of a change (e.g., possible intruder) within either the global coordinate space, as a vector from the current pose of the robot, other useful coordinate systems, or combinations thereof. This information then may be communicated to other robot attributes, robot behaviors, and cognitive conduct within the RIK as well as to other robots or an operator on a remote system.

[00248] As discussed earlier when discussing the range attribute, a variety of coordinate systems may be in use by the robot and an operator. By way of example, a local coordinate system may be defined by an operator relative to a space of interest (e.g., a building) or a world coordinate system defined by sensors such as a GPS unit, an iGPS unit, a compass, an altimeter, and the like. A robot coordinate system may be defined in Cartesian coordinates relative to the robots orientation such that, for example, the X axis is to the right, the Y axis is straight ahead, and the Z-axis is up. Another robot coordinate system may be cylindrical coordinates with a range, angle, and height relative to the robot current orientation.

[00249] The software flow diagram shown in FIG. 18 includes representative components of an algorithm for performing the ROCA behavior 800. As stated earlier, the ROCA process 800 assumes that at least some form of occupancy grid has been established. However, due to the global timing loop execution model, details, probabilities, and new frontiers of the occupancy grid may be built in parallel with the ROCA process 800. The ROCA process 800 begins at decision block 810 by testing to determine if the robot includes lasers, the laser data is valid, an occupancy grid is available, and the ROCA process is enabled. If not, the ROCA process 800 ends.

[00250] If decision block 810 evaluates true, process block 820 performs a new laser scan, which includes obtaining a raw laser scan, calculating world coordinates for data included in the raw laser scan, and converting the world coordinates to the current occupancy grid. The

raw laser scan includes an array of data points from one or more lasers sweeps with range data to objects encountered by the laser scan at various points along the laser sweep. Using the present occupancy grid and present robot pose, the array of range data may be converted to an occupancy grid (referred to as laser-return occupancy grid) similar to the present occupancy grid map.

[00251] Next, decision block 830 tests to see if the current element of the array of range data shows an occupancy element that is the same as the occupancy element for the occupancy grid map. If so, control passes to decision block 860 at the bottom of the range data processing loop, which is discussed later.

[00252] If there is a difference between the laser-return occupancy cell and the corresponding cell for the occupancy grid map, decision block 840 tests the laser-return occupancy cell to see if it is part of an existing change occurrence. In other words, if this cell is adjacent to another cell that was flagged as containing a change, it may be part of the same change. This may occur, for example, for an intruder, that is large enough to be present in more than one occupancy grid. Of course, this test may vary depending on, for example, the granularity of the occupancy grid, accuracy of the laser scans, and size of the objects of concern. If decision block 840 evaluates true, operation block 842 clusters this presently evaluated change with other change occurrences that may be adjacent to this change. Then control will transfer to operation block 848.

[00253] If decision block 840 evaluates false, the presently evaluated change is likely due to a new change from a different object. As a result, operation block 844 increments a change occurrence counter to indicate that there may be an additional change in the occupancy grid.

[00254] Operation block 848 records the current change occurrences and change clusters whether from an existing cluster or a new cluster, then control transfers to decision block 850.

[00255] Decision block 850 tests to see if the change occurrence counter is still below a predetermined threshold. If there are a large number of changes, the changes may be due to inaccuracies in the robot's current pose estimate. For example, if the pose estimate indicates that the robot has turned two degrees to the left, but in reality, the robot has turned five degrees to the left, there may be a large number of differences between the laser-return occupancy grid and the occupancy grid map. These large differences may be caused by the inaccuracies in the pose estimate, which would cause inaccuracies in the conversion of the laser scans to the laser-return occupancy grid. In other words, skew in the alignment of the laser scan onto the occupancy grid

map due to errors in the robot's pose estimation, from rotation or translation, may cause a large number of differences. If this is the case, control transfers to operation block 880 to update the position abstraction in an attempt to get a more accurate pose estimate. After receiving a new pose estimate from the position abstraction, the ROCA process begins again at decision block 810.

[00256] If decision block 850 evaluates true or decision block 860 was entered from decision block 830, decision block 860 test to see if there are more data points in the laser scan to process. If so, control transfers back to decision block 830 to process the next element in the laser scan array.

[00257] If decision block 850 evaluates false, all the data in the laser scan array has been processed and decision block 870 again tests to see if the change occurrence counter is still below a predetermined threshold. As discussed earlier, if the change occurrence counter is not below the predetermined threshold, operation block 880 updates the position abstraction in an attempt to get a more accurate pose estimate, the ROCA process begins again at decision block 810.

[00258] If decision block 870 evaluates true, then processing for this laser scan is complete and operation block 890 updates a change vector and information regarding change occurrences and change clusters is made available to other robot attributes, robot behaviors, and cognitive conduct modules.

[00259] By way of example, and not limitation, the ROCA results may be sent to the user interface, used by a tracking behavior, and combinations thereof. For example, ROCA results may be used with additional geometric calculations to pan a visual camera, a thermal camera, or combination thereof to fixate on one or more of the identified changes. Similarly, a manipulator, such as, for example, a weapon may be panned to acquire a target identified as one of the changes. If the detected change is moving, tracking position updates may arrive in near real time (the actual rate may depend on the speed and latency of the communication channel), allowing various sensors to continuously track the target. If desired, the robot may also continuously move to the new location identified by the change detection system to provide a mobile tracking capability.

[00260] When coupled with an operator interface, the tracked entity's movements may be indicated to an operator in near real time and visual data from a camera can be used by the operator to identify the tracked entity.

[00261] As with other behaviors, the ROCA behavior 800 operates on the global timing loop. Consequently, the ROCA behavior 800 will be re-entered and the process repeated on the next time tick.

4.4. VIRTUAL RAIL CONDUCT

[00262] One representative cognitive conduct module enabled by the RIK and GRA is a virtual rail system for robots. Many industrial and research applications involve moving a vehicle or target at varying speeds along a designated path. There is a need to follow physical paths repeatably either for purposes of transport, security applications or in order to accurately record and analyze information such as component wear and tear (e.g., automotive testing), sensor responsiveness (e.g., sensor characterization), or environmental data (e.g., monitoring). Such applications require both accuracy and repeatability.

[00263] Conventional practice methods have required the building of physical or actual tracks along which a vehicle can be moved. Drawbacks of such an approach include the significant limitations of the configuration of paths that may be created and the feasibility of building permanent tracks. Also, for characterization and other readily modifiable tasks, reconfiguration of physical track networks quickly becomes cost and time prohibitive.

[00264] Although it has long been known that physical tracks or rails are problematic, mobile robots have not had a means by which to maintain accurate positioning apart from such fixed-track methods. For some tasks, absolute positioning can be achieved by various instrumented solutions such as visual, laser-based tracking systems or radio frequency positioning systems that triangulate distance based on beacons placed in the environment. Each of these systems is costly to implement; in fact, the cost for purchasing and installing such a positioning system is often more than the total cost of the robot itself.

[00265] Moreover, the utility of visual or laser tracking systems is limited by occlusions within the environment. For example, RF beacons are only appropriate for environments where the beacons can be fixed in a static, known location. The physical properties of a remote sensing environment are constantly changing. In fact, walls are often shifted within the building to model different operational environments. Accordingly, absolute positioning is sometimes less feasible, impractical and frequently impossible to implement. Therefore, there is a need to provide a method and system for configuring a virtual track or rail system for use by a robot.

[00266] The present invention includes various embodiments including a robot system configured to follow pre-planned routes forming a “virtual rail” or “virtual track” and may

include defined speeds for traversing various segments of the pre-planned routes. One application of a virtual rail system includes the repeated testing of a sensor or system to characterize the device. Due to the accuracy and repeatability of the virtual rail system, sensors and systems may be tested with data collected that conforms to a “sufficient comparable data” standard. Such a data collection standard requires acceptance of data only when consistent and comparable data is generated in response to repeatable tests carried out under the same conditions. For example, the virtual rail system may be used in a laboratory, research facility, or manufacturing environment to characterize a vast number of sensors. Accordingly, characterization tests that may previously have required a significant amount of time for execution, may now be characterized in a fraction of the time.

[00267] Sensor characterization is only one example of a specific application. Other applications include automated mail carts and other delivery systems, security and surveillance systems, manufacturing and monitoring systems. In particular, the technology is useful for parts handling, as well as replacement of current railed robotic systems especially within the manufacturing and defense industries.

[00268] FIG. 19 is a block diagram of a robot system for implementing a virtual track for a robot, in accordance with an embodiment of the present invention. A robot system 2100 includes a robot 2102 and a control generation system 2104. In robot system 2100, a user interfaces with control generation system 2104 to implement a virtual track for tracing or following by robot 2102. Robot 2102 is responsive to programming commands generated by control generation system 2104 and further conveys feedback and sensor information to control generation system 2104 over communication interface 2106. A user through a user interface of control generation system 2104 designates a desired path comprised of one or more representative path segments. In the various embodiments of the present invention, robot 2102 is configured or programmed to follow a virtual track or a virtual rail similar in resulting operation to a robot following a fixed track or physical rail. In the various embodiments of the present invention, however, the shortcomings of a fixed physical rail configuration are overcome by enabling the formation of a virtual track or rail system without the appreciated physical and economical limitations associated therewith.

[00269] FIG. 20 illustrates a user interface for generating a desired path representative of a virtual track or virtual rail, in accordance with an embodiment of the present invention. A user interface 2120 operating on a conventional computer or other hosting interface provides an environment wherein a user may configure and readily reconfigure a virtual track or rail configuration for execution and following by a robot.

[00270] The user interface 2120 provides an environment for the generation of a desired path comprised of at least one segment representative of the virtual track for the robot. The user interface 2120 may take the form of a Computer Aided Design (CAD) program for the formation of the desired path. The desired path comprised of one or more segments representative of the virtual track for the robot may take the form of lines, arcs or any of a number of design shapes known by those of ordinary skill in the art, and are collectively referred to herein as “segments.” By way of example, a desired path 2122 includes a plurality of line segments 2124-2132 with line segment 2132 illustrated as being selected. Line segments 2124-2132 may be generated using any of a number of commercially available CAD system that may generate file formats that are readily convertible and parsable. By way of example and not limitation, the CAD file format may be directly saved or converted into a file format such as Drawing Exchange Format (.dxf).

[00271] FIG. 21 is a process diagram for configuring the desired path into a waypoint file for implementing a virtual track or rail and for execution by a robot, in accordance with an embodiment of the present invention. A virtual track or rail is specified in the form of a desired path 2122 (FIG. 20) including at least one segment representative of the virtual track as input through the user interface 2120 (FIG. 20). The graphical input of the desired path is converted or stored in a form that is capable of further processing or manipulation by the control generation system 2104 (FIG. 19) which generates programming commands destined for execution by robot 2102 (FIG. 19). By way of example and not limitation, the stored format for the desired path of the one or more segments representative of the virtual track may be a drawing file 2202. The format of drawing file 2202, among others, includes file formats (e.g., .dxf) configured to represent various line segments, arcs and other drawing elements as expressed by a user through a graphical user interface 2120 (FIG. 20).

[00272] A path plan process 2204 receives the CAD-generated drawing file 2202 and processes the one or more segments of the desired path into a waypoint file 2206 that includes instructions that are capable of being executed by robot 2102 (FIG. 19). The processing of drawing file 2202 includes the assignment process 2208 of input velocities 2200 to the segments or vertices of the desired path segments or elements. A verification process 2210 analyzes the desired input velocities 2200 by comparing the velocities with the mobility capabilities of robot 2102 (FIG. 19). Discrepancies or incompatibilities between the desired path and input velocities as compared with the execution capabilities of robot 2102 are reported and/or resolved.

[00273] Path plan process 2204 further includes a waypoint generation process 2212 for generating waypoint file 2206 which precipitate from the original drawing file 2202

undergoing assignment process 2208 followed by verification process 2210 for determining the compatibilities of the desired path and the robot capabilities. Waypoint file 2206 includes a listing of waypoints as well as any modified velocities 2214 which may be different than the originally specified input velocities 2200.

[00274] FIG. 22 illustrates a user interface for further processing the desired path into a program for execution by a robot, in accordance with an embodiment of the present invention. A path plan process user interface 2420 provides an environment for the rendering of a previously defined drawing file 2202 (FIG. 21) and further enables the generation of a waypoint file 2206 (FIG. 21) through the assignment of start and end points 2440, 2442 to the desired path 2422 as well as the association of velocities with such paths. The desired path 2422, comprised of one or more segments 2424-2432 representative of the virtual track for the robot, thereafter includes assigned motion qualities and characteristics including start and end points 2440, 2442 to the desired path 2442 as well as assigned input velocities 2200 (FIG. 21) or speeds that should be executed by robot 2102.

[00275] As stated, the one or more line segments 2424-2432 with the assigned motion qualities is compared or verified through verification process 2210 (FIG. 21) with the performance capabilities of a specific robot 2102 which compares the requested desired path with mobility limitations and capabilities of robot 2102. In one embodiment of the present invention, an algorithm analyzes the path including traversal of the segments at various speeds, including velocity transitions between line and arc segments and determines the turn gain to insure minimal oscillations during traversal of the line segments. Furthermore, the algorithm is capable of carving smooth arcs by adjusting the turn gain based on an analysis of the arc shape and the commanded forward velocity. This algorithm provides the ability to arbitrate between waypoint following and motor schema control as speed and point types change.

[00276] After resolution of any inconsistencies or incompatibilities, a waypoint file 2206 (FIG. 21) is generated by path plan process 2204 with waypoint file 2206 (FIG. 21) being transferred over communication interface 2106 (FIG. 19) to robot 2102 (FIG. 19) for execution. Robot 2102, executing the various waypoints and specified velocities 2214 (FIG. 21) associated therewith, traces out or follows a virtual track or virtual rail as specified and/or modified by a user through the control generation system 2104 (FIG. 19).

[00277] The user interface 2420 for controlling path plan process 2204 (FIG. 21) enables a user to generate commands in the form of waypoint file 2206 (FIG. 21) for execution by robot 2102 which results in the formation of a virtual rail or track that is followed or traced by robot 2102. The virtual track or rail may be created from an abstraction or may be generated

with reference to an available map or other boundary designations of the operating environment. Furthermore, accurate positioning of the robot 2102 (FIG. 19) may be maintained by application of Markov localization techniques that may combat problems such as odometry drift. Generation of waypoint file 2206 (FIG. 21) allows a robot 2102 (FIG. 19), given accurate position data, to traverse a trace of arcs and lines at various speeds. The various embodiments of the present invention may utilize various mapping or localization techniques including positioning systems such as indoor GPS, outdoor GPS and DGPS, a theodolite system as well as others which may be devised in the future.

[00278] As stated in FIG. 22, desired path 2422 includes a plurality of line segments 2424-2432. Through the use of the user interface 2420 start point 2440 and end point 2442 may be selected with each individual line segment 2424-2432 being individually selected thereby allowing the association of a velocity therewith. By way of example, line segment 2432 is illustrated as being selected with a representative speed of 0.5 meters per second being associated therewith. The path plan process 2204 (FIG. 21) through user interface 2420 uses the properties of each segment within drawing file 2202 (FIG. 21) to spatially locate each segment (e.g., line or arc) and then creates a default path based on the initial order of segments found in the drawing file 2202.

[00279] Path plan process 2204 (FIG. 21), through user interface 2420, can be used to manipulate various properties of the initial desired path 2422. For example, when segment 2432 is selected, the segment is highlighted in the user interface 2420. Once a segment is highlighted, its properties are displayed and can be edited, if desired. The order of segments can be changed, for example, either by using the “Move Up” and “Move Down” buttons or by selecting and dragging a segment to its new position. Each segment can be included or excluded, for example, from the path by appropriately marking the “Include this entity in the path” checkbox. This allows additional features that are not part of the path to be included in the drawing file without the requirement that they be a part of the virtual track or rail. Additional input boxes may be provided to set the initial speed, the final speed or constant acceleration, and provide for comments for each segment.

[00280] Once motion characteristics, such as velocity, have been associated with each of the line segments 2424-2432, other processing may be performed such as an estimation of run time as well as verification of velocity transitions 2210 (FIG. 21). Once velocities have been associated therewith and verification of compatibility with the capabilities of the target robot have been performed, a waypoint file 2206 (FIG. 21) may be generated by activating generate waypoint process 2212 (FIG. 21) within the user interface 2420.

[00281] FIG. 23 is a diagram illustrating transformation according to path plan process 2204 (FIG. 21) from a drawing file to a waypoint file, in accordance with an embodiment of the present invention. A drawing file 2202 as received and generated in a user interface 2120 is transformed as stated above, with respect to FIG. 21, from a drawing file 2202 to a waypoint file 2206 according to path plan process 2204 (FIG. 21). As stated, drawing file 2202 is a limited expression of graphical segments and must be augmented through path plan process 2204 to include, among other things, motion characteristics such as velocities as well as execution ordering of the segments. Additionally, for the generation of waypoint file 2206, the information in drawing file 2202 also undergoes verifications to determine if input velocities 2200 (FIG. 21) are within the capabilities of the robot.

[00282] By way of example and not limitation, waypoint file 2206 assumes one or more formats, an example of which is illustrated with respect to FIG. 23. Waypoint file 2206 may include an estimated traversal time 2402 identifying a summation of the traversal times of each segment of the virtual track. By way of example, waypoint file 2206 includes a listing of ordered vertices identifying the waypoints 2404-2414 for traversal by the robot 102 (FIG. 19). Each waypoint 2404-2414 includes a waypoint number indexed according to order as previously described, X and Y coordinate values, a velocity value, and an arc continuation flag for associating a set of waypoints for short line segments that comprise an arc traversal.

[00283] FIG. 24 is a functional block diagram of a control process of a robot, in accordance with an embodiment of the present invention. Robot control process 2300 executes a waypoint file to trace-out or follow a virtual track or rail first defined and processed within the control generation system 2104 (FIG. 19). Robot control process 2300 includes a localization process 2302 wherein the robot processes environmental parameters including physical boundaries to determine a present frame of reference for use in alignment or referencing the virtual track or rail. Localization is a continuous process that the robot uses to determine its present location with reference to its internal map. For example, when the robot first starts executing the waypoint file 2304, the robot's start point is the origin (0,0) with positive X in the forward direction for referencing the internal map. As the robot moves around, the robot locates items in the operating environment and then places those items in the robot's internal map. As the robot continues to move, the robot may encounter familiar features with an expectation that the recognized features are located in the same relative position. However, if the features have moved relative to where the robot believes the features should be, then the robot assumes that the robot may not be in the right place on the internal map. The continual correction of the robot's position in the internal map may be described as "localization."

[00284] The localization process of the robot allows the robot to accurately and repeatedly trace the waypoints forming the virtual rail or track. The navigation process 2306 responds to the localization process 2302 and sensor data from sensor process 2310 to generate controls to the robot motion process 2308. Additionally, the robot uses sensor data from sensor process 2310 to determine surrounding features. The robot control process 2300 does not need to necessarily identify the composition or identity of the features, but only the fact that they are part of the environment which forms boundaries for the robot. Robot 2102 may utilize one or more sensors for providing feedback to the localization process. Sensors may include wheel measuring devices, laser sensors, ultrasonic sensors, and the like.

[00285] Waypoint navigation process 2306 generates commands or control signals to a robot motion process 2308. Robot motion process 2308 generates controls to actuators for generating motion, rotation, etc., as well as velocities associated therewith. Waypoint navigation process 2306 further receives from sensor process 2310 sensor information in the form of feedback for determining when traversal of one or more segments of the virtual rail has been accomplished. Sensor process 2310 may also provide information to waypoint navigation process 2306 in the form of changes to environmental parameters which enables waypoint navigation process 2306 to protect or guard against unforeseen changes to the environment. Additional details with respect to waypoint navigation are described below with respect to FIGS. 26-28.

[00286] FIG. 25 is a flow chart 2500 of a method for implementing a virtual track for a robot, in accordance with an embodiment of the present invention. An execution of a robot traversing a virtual rail is initiated by a user developing a desired path for the robot to follow using a drawing package to generate 2510 a drawing file.

[00287] As stated, a drawing file or other illustration of a desired path is generated 2510 and includes at least one segment representative of the virtual track to be configured for the virtual track which the robot will traverse. Generation of a desired path results in the creation of a specific file format representing the illustrated segments of the desired path. The file format, in one embodiment of the present invention, is converted 2520 into a standardized file format, an example of which is the .dxf format. Generation 2510 and converting 2520 steps may be accomplished through the use of one or more applications which are made usable through a user interface, such as user interface 120 (FIG. 20).

[00288] Through path plan process 2204 (FIG. 21) and as further illustrated with respect to a user interface 2420 (FIG. 22), the drawing file 2202 (FIG. 23) is imported 2530 and start points 2440 (FIG. 22), end points 2442 (FIG. 22) and segment ordering may be assigned

2540 to the various segments of the desired path 2422 (FIG. 22). Through verification process 2210 (FIG. 21), continuity may be checked or verified 2550 and input velocities 2200 (FIG. 21) may be assigned 2560 to the various segments 2424-2432 (FIG. 22) of the desired path 2422 (FIG. 22). Further checking and reporting 2570 of inconsistencies or incompatibilities may also be performed.

[00289] Once the desired path has been illustrated and start and end points as well as velocities have been associated therewith as well as a successful completion of verification processes, a waypoint list 2206 (FIG. 23) is generated 2580 and stored in a waypoint file. Upon completion of the generation of waypoint file 2206 (FIG. 21) by control generation system 104 (FIG. 19), the waypoint file 2206 is sent 2590 via a communication interface 106 (FIG. 19) to a robot 102 (FIG. 19). Thereafter, robot 102 may execute 2600 a first waypoint from waypoint file 2206 and subsequently execute 2610 a second and subsequent waypoints using waypoint navigation process 2306 (FIG. 24).

4.5. WAYPOINT FOLLOWING BEHAVIOR

[00290] FIGS. 26, 27, and 28 are software flow diagrams illustrating representative algorithms for performing waypoint following according to embodiments of the present invention. The waypoints may come from an algorithm such as the virtual robot rail system described above, interaction with an operator, interaction with other robots, internally generated, or combinations thereof. FIG. 26 illustrates components of a handler algorithm for handling transitions between waypoints, FIG. 27 illustrates handling of translational velocities during waypoint following, and FIG. 27 illustrates handling of rotational velocities during waypoint following.

[00291] The waypoint handler, illustrated in FIG. 26, starts with decision block 902 to test whether path planning is active and the time since the achieving the last waypoint is greater than a threshold. In the representative embodiment of FIG. 26, the threshold is set at three seconds. If sufficient progress has not been made toward a waypoint within the threshold, there may be a barrier blocking the robot's progress toward the waypoint. For example, perhaps a door was closed that the waypoint planning had assumed was open, or perhaps a new obstacle was placed in the environment such that the robot cannot find a way around the obstacle to achieve the next way point. In these types of circumstances, it may be appropriate to plan a new path with a new waypoint list. Thus, if path planning is active and the threshold is exceeded, operation block 904 performs a routine to delete the current waypoint list and plan a new waypoint list, then control transfers to decision block 906.

[00292] If decision block 902 evaluates false, or operation block 904 completes, decision block 906 test to see if the current waypoint is defined as part of an arc. If the current waypoint is part of an arc, operation block 908 sets a variable named Waypoint_Radius as the current speed times one half of the robot's length. This Waypoint_Radius variable is used later as a test threshold when determining how close the robot is to the waypoint. If the current waypoint is not part of an arc, operation block 910 sets Waypoint_Radius to one half the robot's length plus one half the length of the arm extension. Thus, the waypoint radius is defined as the physical extent of the robot from the Robot's center.

[00293] With the Waypoint_Radius variable set, decision block 912 tests to see if the angle to the target waypoint is currently less than 90 degrees to the left or right. If so, operation block 914 sets the range to the target as the closest range within plus or minus 15 degrees of the current angle to the target. If the waypoint is not less than 90 degrees away, the range to target is set as Min_Front_Distance, which, as explained earlier, is the range to the nearest object within plus or minus 90 degrees of the robot's forward direction. The current angle to the target defines the angle towards the target relative to straight ahead. However, Range_To_Target defines a range (i.e., distance) from the robot to an obstacle in the direction of the waypoint.

[00294] After setting the Range_To_Target variable, decision block 918 tests to see if the distance to the current waypoint is less than the waypoint radius defined previously. If so, the waypoint is considered to be achieved, so operation block 920 iterates to the next waypoint in the waypoint list, and the process exits.

[00295] If decision block 918 evaluates false, a more specific test is performed to see if the waypoint has been achieved. In some instances, it may not be possible to actually place the center of the robot over the waypoint. For example, the waypoint may have been placed too close to a wall, or perhaps even behind the wall. However, if the robot can get close enough, it may be sufficient to say the waypoint has been achieved. Thus, if decision block 922 evaluates true, operation block 920 iterates to the next waypoint in the waypoint list, and the process exits. However, if decision block 922 evaluates false the process exits and continues on with the current waypoint.

[00296] A representative evaluation of a test for close enough to a waypoint is illustrated in block 924. Of course, those of ordinary skill in the art will recognize that other parameters, distances, and decisions may be made within the scope of the present invention to define whether a waypoint has been achieved. In block 924, the first test checks to see if the Range_to_Target variable is less than the arm extension plus the larger of a forward threshold or a side threshold. If not, there may still be room to move forward or rotate toward the waypoint,

so the process may exit and continue on with the current waypoint. Otherwise, the second test checks to see if the distance to the waypoint is less than the sum of the ram extension and the robot length. If not, there may still be room to move forward or rotate toward the waypoint, so the process may exit and continue on with the current waypoint. Otherwise, the third test checks to see if the distance to the closest object on the front side of the robot (i.e., Min_Front_Distance) is less than the arm extension plus twice the forward threshold. If not, there may still be room to move forward or rotate toward the waypoint, so the process may exit and continue on with the current waypoint. Otherwise, the final check tests to see if the angle to the target is less than 45 degrees or the range to the nearest obstacle is less than the turn threshold. If not, there may still be room to move or rotate toward the waypoint, so the process may exit and continue on with the current waypoint. Otherwise, operation block 920 iterates to the next waypoint in the waypoint list, and the process exits.

[00297] FIG. 27 is a software flow diagram illustrating components of an algorithm for adjusting translational velocity 930 during the waypoint follow behavior. First, operation block 932 tests to see if the distance to the next waypoint is less than one tenth of the robot's length. If so, operation block 934 performs an update of the robot's current pose to be certain that the pose is precise relative to the waypoint location.

[00298] If operation block 932 evaluates false, or after the pose is updated, decision block 936 tests to see if the range to the closest object in front is less than twice a predefined threshold. If not, control transfers to decision block 944. However, if the range to the closest object in front is less than twice a predefined threshold, the robot may be approaching close to an obstacle, so decision block 938 tests to see if the robot is blocked in the direction of the target. If so, operation block 940 performs a backup procedure and the process exits. If the robot is not blocked in the target direction, decision block 942 tests to see if the angle to the target is greater than 60 degrees. If so, the robot may not be able to achieve the target without backing up, so operation block 940 performs a backup procedure and the process exits. If the angle to the target is not greater than 60 degrees, a backup procedure may not be needed and control transfers to decision block 944.

[00299] Decision block 944 tests to see if the angle to the target is greater than 45 degrees. If so, operation block 946 sets the translational speed to zero enabling the robot to stop making forward progress while it rotates to face more directly toward the target. After setting the speed to zero, the process exits.

[00300] If the angle to the target is not greater than 45 degrees, new translational velocity determination continues by decision block 948 testing to see if a detection behavior is in

progress. As stated earlier when describing the obstacle avoidance behavior, a detection behavior may be a behavior where the robot is using a sensor in an attempt to find something. For example, the countermine conduct is a detection behavior that is searching for landmines. In these types of detection behaviors, it may be desirable to approach much closer to objects, or to approach objects with a much slower speed to allow time for the detection function to operate. Thus, if a detection behavior is active, operation block 950 sets a desired speed variable based on detection parameters that may be important. By way of example, and not limitation, in the case of the countermine conduct this desired speed may be set as:

$\text{Desired_Speed} = \text{Max_passover_rate} - (\text{Scan_amplitude}/\text{Scan_Speed})$. In this countermine conduct example, the Max_passover_rate may indicate a maximum desired speed for passing over the landmine. This speed may be reduced by other factors. For example, the $(\text{Scan_amplitude}/\text{Scan_Speed})$ term reduces the desired speed based on a factor of how fast the mine sensor sweeps an area. Thus, the Scan_amplitude term defines a term of the extent of the scan sweep and the Scan_Speed defines the rate at which the scan happens. For example, with a large Scan_amplitude and a small Scan_Speed, the Desired_Speed will be reduced significantly relative to the Max_passover_rate to generate a slow speed for performing the scan. While countermine conduct is used as an example of a detection behavior, those of ordinary skill in the art will recognize that embodiments of the present invention may include a wide variety of detection behaviors, such as, for example, radiation detection, chemical detection, and the like.

[00301] If a detection behavior is not in progress, decision block 952 tests to see if a velocity limit is set. In some embodiments of the invention, it may be possible for the operator to set a velocity limit that the robot should not exceed, even if the robot believes it may be able to safely go faster. For example, if the operator is performing a detailed visual search, the robot may be performing autonomous navigation, while the operator is controlling a camera. The operator may wish to keep the robot going slow to have time to perform the visual search.

[00302] If a velocity limit is set, operation block 954 sets the desired speed variable relative to the velocity limit. The equation illustrated in operation block 954 is a representative equation that may be used. The 0.1 term is a term used to ensure that the robot continues to make very slow progress, which may be useful to many of the robot attributes, behaviors, and conduct. In this equation, the Speed_Factor term is a number from one to ten, which may be set by other software modules, for example the guarded motion behavior, to indicate a relative speed at which the robot should proceed. Thus, the desired speed is set as a fractional amount of the Max_Limit_Speed.

[00303] If a velocity limit is not set, operation block 956 sets the desired speed variable relative to the maximum speed set for the robot (i.e., Max_Speed) with an equation similar to that for operation block 614 except Max_Speed is used rather than Max_Limit_Speed.

[00304] After the Desired_Speed variable is set by block 950, 954, or 956, decision block 958 determines if the distance to the current waypoint is less than the current velocity plus a small safety factor. If not, operation block 968 sets the new translational speed for the robot to the Desired_Speed variable and the process exits. However, if the current waypoint is getting close, as determined by decision block 958 evaluating true, decision block 960 determines if the current waypoint is part of an arc. If so, operation block 962 sets the translational speed such that the robot can smoothly traverse the arc. Thus, operation block 962 is a representative equation that sets the new translational speed as a function of the larger of either the angle to the target, or the turn angle to the next waypoint. In other words, the translation velocity will be reduced by setting the new speed to the current speed multiplied by a fractional change factor. This fractional change factor may be defined as the cosine of the larger of either the angle to the target, or the turn angle to the next waypoint.

[00305] If the current waypoint is not part of an arc, it may still be desirable to slow the robot's translational speed down in preparation for turning toward the next waypoint. Thus, operation block 964 is a representative equation for setting the new translational speed for the robot by multiplying the current speed by a different fractional change factor. This fractional change factor may be set as about $(0.7 + (0.3 * \cos(\text{Next_Turn_Angle}))$. In other words, the new speed will be set somewhere between 70% and 100% of the current speed based on the angle towards the next waypoint. If the angle is small, for example zero degrees, there may be no need to slow down and the new speed can be set at 100% of the current speed. Conversely, if the angle is large, for example 90 degrees, it may be desirable to slow down significantly in preparation for a turn to the new waypoint. Thus, the new translational velocity is set at 70% of the current speed. Of course, the next time through the global timing loop presents another chance to adjust the translational speed if the angle to the next waypoint is still large.

[00306] This sets the translational speed based on the severity of the turn that will be negotiated to achieve the next waypoint.

[00307] After setting the current speed, from operation block 968, 964, or 962, the translational velocity process 930 ends.

[00308] FIG. 28 is a software flow diagram illustrating components of an algorithm for performing rotational velocity adjustments 970 of the waypoint follow behavior. First, decision block 972, checks to see if waypoint following is enabled. If not, the process exits, because

rotational velocity changes will be handled by another behavior, such as, for example, the obstacle avoidance behavior.

[00309] If waypoint following is enabled, decision block 974 tests to see if the robot is blocked in front. If not, rotational velocity determination can continue at decision block 986. However if the robot is blocked in front, decision block 976 determines whether the current waypoint is to the left of the robot. If so, decision block 978 tests the area to the left of the robot where the robot may want to turn towards and finds the range to the nearest object in that area. If the range is larger than a turning threshold, as tested by decision block 982, there is room to turn, so operation block 980 sets the rotational velocity to the left at 30% of a predefined maximum rotational velocity. After setting the rotational velocity, the process exits.

[00310] If the waypoint is not on the left, decision block 82 tests the area to the right of the robot where the robot may want to turn towards and finds the range to the nearest object in that area. If the range is larger than a turning threshold, as tested by decision block 982, there is room to turn, so operation block 984 sets the rotational velocity to the right at 30% of a predefined maximum rotational velocity. After setting the rotational velocity, the process exits.

[00311] If the robot is blocked in front and there is not room to turn (i.e., either decision block 978 or 982 evaluates false), then the process exits to a get unstuck behavior in an effort to find away to get around the obstacle in front so that the robot can continue to pursue the current waypoint.

[00312] If the robot is not blocked in front, decision block 986 tests to see if the angle to the waypoint target is less than ten degrees. If so, the robot is close to pointed in the correct direction and only minor corrections may be useful. Thus, in a representative method for determining an appropriate change to the rotational velocity, operation block 988 sets a Waypoint_Turn_Gain as the angle to the target divided by 100. Conversely, if the waypoint target is equal to or greater than ten degrees, a larger correction to the rotational velocity may be appropriate to get the robot pointed toward the current waypoint. Thus, in a representative method for determining an appropriate change to the rotational velocity, operation block 990 sets a Waypoint_Turn_Gain as the base 10 logarithm of the angle to the target minus one. As a result, the larger the angle to the waypoint target, the larger the value will be for the Waypoint_Turn_Gain.

[00313] With the Waypoint_Turn_Gain set, decision block 992 test to see if the waypoint is on the left. If so, operation block 994 sets the turn velocity to the left by multiplying the current turn velocity by the Waypoint_Turn_Gain, and the process exits. If the waypoint is

not on the left, operation block 996 sets the turn velocity to the right by multiplying the current turn velocity by the Waypoint_Turn_Gain, and the process exits.

[00314] As with other behaviors, the waypoint algorithms 900, 930, and 970, in FIGS. 26, 27, and 28, respectively operate on the global timing loop. Consequently, the decision of whether a waypoint has been achieved to move on to the next waypoint, adjustments to the translational velocity, and adjustments to the rotational velocity, may be repeated on each time tick of the global timing loop.

4.6. ROBOTIC FOLLOW CONDUCT

[00315] One representative cognitive conduct module enabled by the RIK is a robotic follow capability wherein one or more robots are sent to a map location of the most recent change in the environment or directed to follow a specific moving object. FIG. 29 is a software flow diagram illustrating components of an algorithm for performing the follow conduct 1000.

[00316] This relatively autonomous conduct may be useful for a fast-moving robot with perceptors that may be used by robot attributes and robot behaviors to detect and track changes in the environment. It would be difficult for conventional robots under direct operator control to avoid obstacles, track where the robot is going, and track the object of pursuit at the same time. However, with the relatively autonomous conduct and collaborative tasking enabled by the RIK, a high-speed chase may be possible.

[00317] The RIK may include a tracking behavior that allows the robot to track and follow an object specified by the operator with the camera or other tracking sensors, such as thermal, infrared, sonar, and laser. Consequently, the tracking behavior is not limited to visual tracking, but can be used with any tracking system including a thermal imaging system for tracking human heat signatures.

[00318] In visual tracking, for example, the operator may specify an object to be tracked within the operator's video display by selecting a pursuit button on the interface and then manipulating the camera view so that the object to be tracked is within a bounding box. The camera can then track the object based on a combination of, for example, edge detection, motion tracking, and color blob tracking. Furthermore, the camera can track the motion of the target independently from the motion of the robot, which allows the robot to follow the optimal path around obstacles even if this path at first may take the robot in a direction different from the direction of the target.

[00319] Thus, the robotic follow conduct 1000 effectively blends robot behaviors, such as, for example, tracking, obstacle avoidance, reactive path planning, and pursuit behaviors. To

begin the follow conduct 1000, operation block 1010 illustrates that the conduct queries or receives information regarding the present bearing to a target. This present bearing may be generated by a tracking behavior, such as, for example, the ROCA behavior discussed above or from an updated map location from the operator or other robot. In addition, the bearing may be converted to a robot relative coordinate system, if needed. Both the tracking behavior and follow conduct 1000 operate on the global timing loop. As a result, the follow conduct 1000 will be re-entered each timing tick and be able to receive an updated bearing to the target, from the tracking behavior or other defined location, each timing tick.

[00320] Decision block 1015 tests to see if the robot has reached the target. If so, the follow conduct 1000 exits. If not, the follow conduct 1000 transitions to decision block 1020. In this representative embodiment, reaching the target is defined as: 1) The Closest obstacle in a 30° region in which the tracked object lies is closer than the closest obstacle in a 30° region on the opposite side; 2) Both L-front and R-front are obstructed; 3) the angle to the object lies in the front region; and 4) the distance to the object in front is less than the distance on the right and left.

[00321] Decision block 1020 tests to see if the front is blocked. If so, control transfers to operation block 1030 to attempt to get around the obstacle. If not, control transfers to decision block 1070. The front blocked decision may be based, for example, on a flag from the guarded motion behavior discussed previously.

[00322] Decision block 1030 begins a process of attempting to get around a perceived obstacle. To begin this process, decision block 1030 checks the current speed. If the speed is not greater than zero, control transfers to decision block 1040. If the speed is greater than zero, operation block 1035 sets the speed to zero before continuing with decision block 1040.

[00323] Decision block 1040 tests to see if the robot is blocked on the left or right. If not, control transfers to decision block 1050. If the robot is blocked on the left or right, the robot may not have an area sufficient to make a turn, so the robot is set to begin backing up with an angular velocity of zero and a linear velocity that is 20% of the presently specified maximum, then the follow conduct 1000 exits.

[00324] Decision block 1050 tests to see if the robot is blocked in the direction of the target. If so, control transfers to decision block 1060. If the robot is not blocked in the direction of the target, the robot is set to turn toward the target with a linear velocity of zero and an angular velocity that is 60% of the presently specified maximum, then the follow conduct 1000 exits.

[00325] Decision block 1060 tests to see if the target is positioned substantially in front of the target. If so, the target is in front of the robot, but the robot is also blocked by an obstacle. Thus, operation block 1062 attempts to move forward slowly but also turn around the obstacle by setting the linear velocity to 10% of the presently specified maximum and the angular velocity to 60% of the presently specified maximum and away from the obstacle. Then, the follow conduct 1000 exits.

[00326] If decision block 1060 evaluates false, then the direction directly in front of the robot is blocked, and the direction toward the target is blocked. Thus, operation block 1064 attempts to find a clear path to the target by setting the linear velocity to -20% of the presently specified maximum (i.e., backing up) and the angular velocity to 30% of the presently specified maximum and in the direction of the target. Then, the follow conduct 1000 exits.

[00327] Returning to decision block 1020, if decision block 1020 evaluates false, then decision block 1070 begins a process of attempting to progress toward the target since the front is not blocked. Thus, decision block 1070 tests to see if the robot is blocked in the direction of the target. If so, operation block 1075 attempts to move forward while gradually turning away from the target in an effort to try to find a clear path to the target by setting the linear velocity to 20% of the presently specified maximum and the angular velocity to 20% of the presently specified maximum. Then, the follow conduct 1000 exits.

[00328] If decision block 1070 evaluates false, then the target is not in front of the robot and the robot is free to move forward. Thus, operation block 1080 attempts to move forward and turn toward the target. In this representative embodiment, the robot is set with an angular velocity toward the target that is determined by the current bearing toward the target divided by a predetermined turn factor. Consequently, the speed at which the robot attempts to turn directly toward the target may be adjusted by the turn factor. In addition, the robot is set to move forward at a safe speed, which may be set as 10% of the maximum, to ensure the robot keeps moving, plus a safe speed adjustment. The safe speed adjustment may be defined as $\text{Front_forward_threshold}/2$. Wherein Front defines the distance to the nearest object in the vicinity of directly in front as defined by the range attribute discussed earlier, and forward_threshold defines a distance to which the robot may be relatively certain that objects are outside of its time horizon. Thus, the robot makes fast, but safe, forward progress while turning toward the target, and the speed may be adjusted on the next time tick based on new event horizon information.

[00329] As with other robot behaviors and cognitive conduct, the follow conduct 1000 operates on the global timing loop. Consequently, the ROCA behavior 700 will be re-entered and the process repeated on the next time tick.

5. SEAMLESS AUTONOMY

[00330] Embodiments of the present invention include methods and systems that provide a more intuitive approach for the user to direct a robot's activities without having to understand the available robot behaviors and autonomy levels. This intuitive approach provides a bridge between a user's intentions and a robot's behaviors by creating a new user interface and adjustments to robot initiative and behavior that do not need to be directly controlled by the user.

[00331] This approach provides a means to control a remote robot via a continuous range of autonomy levels without requiring the operator to explicitly change the autonomy level on the robot. Conventional approaches have distinct autonomy modes or levels and the operator generally chooses a mode and then the robot navigates according to the dictated level of autonomy.

[00332] In contrast, in embodiments of the present invention, the operator is no longer responsible for directly choosing the mode of autonomy. Rather, the operator places a navigational "target" and the robot moves towards the target. This target may also be referred to herein as a hotspot, a task designator, and a task-oriented target.

[00333] Efficient use of the robot often depends on the interactions between the operator and the robot as interacting members of a human-robot team. As tasks change, the requirements of the user and the robot should also change in order to ensure efficiency. The responsibilities of the robot are often referred to as the autonomy of the robot, and when multiple levels of autonomy are available, conventionally, the operator has the responsibility of choosing the appropriate level of autonomy. The challenge with this user-selected approach is that operators often do not realize when they are in a situation where the autonomy of the robot should be changed in order to maintain high performance. Furthermore, as new behaviors and intelligence are added to the robot, the traditional approach requires the operator to maintain appropriate mental models of what all the modes are and how they should all be used.

[00334] Leaving the autonomy switching and mode switching responsibility to the operator means that mistakes will be made from not recognizing the need to switch, or switching to the wrong autonomy level. In general, requiring the operator to have an understanding of the appropriate context for each autonomy mode is cognitively challenging and prone to error because of the requirement for the operator to make the decision about autonomy modes. Despite this, most solutions to adjustable autonomy are to create various "levels" of autonomy from which the user must choose. Even when the robot is given the ability to automatically switch between the discrete autonomy modes based on its own assessment of the situation, the

operator often feels frustrated and confused about what the robot is doing and why it usurped control, which usually results in a fight for control between the operator and the robot. This fight for control generally emanates from poor communications between the robot and the human as to why the decision was made and how the decision is being put into action.

[00335] Embodiments of the present invention solve many of these problems by supporting seamless levels of interaction based on the operator indirectly controlling the robot. This indirect control is performed by allowing the operator to specify the intent of the user from a human-centered reference frame. Instead of selecting discrete levels of autonomy, the operator uses defined intentions in the form of the task-oriented targets. By using indirect control through specifying intent, the “fight-for-control” problem present in direct control approaches may be eliminated because issues with communications latency and bandwidth requirements are mitigated and the user may experience a more intuitive feel for what the robot is doing.

[00336] FIG. 30 is a user interface window showing controls that a user may need to operate in order to direct actions of a robot in a conventional user interface 1200. In this conventional user interface 1200, the user may need to control where a camera is looking using a set of camera controls 1210 to perform operations such as zoom, tilt, and pan. The user may also have another set of controls 1220 for navigating the robot when the robot is operating in a low initiative mode such as, for example, teleoperation mode or safe mode. The user may also have autonomy controls 1230 to directly set autonomy levels. For example, the user may need to set the autonomy level to teleoperation mode or safe mode in order for the navigation controls 1220 to be effective. As another example, the user may also have another set of controls 1240 for selecting intelligent behaviors. These intelligent behaviors may include go-to points, waypoints, path-planning, search region, patrol region, and retro-traverse.

[00337] FIGS. 31A-31D are user interface windows 1300 in accordance with one or more embodiments of the present invention. The user interface windows 1300 include an environment map 1390, a robot position (1310A, 1310B, 1310C, and 1310D), task-oriented targets (1320 and 1330), a path plan 1315, and an imaging window (1325A, 1325B, 1325C, and 1325D).

As a non-limiting example, seamless autonomy changes using task-oriented targets may be incorporated within the GRA, RIK, robot behaviors, and robot conduct described herein. However, embodiments of the present invention are not so limited. Task-oriented targets and seamless autonomy according to embodiments of the present invention may be practiced in other robot architectures that define specific behaviors for the robot to accomplish and include a general rubric of varying autonomy levels.

[00338] Using the seamless autonomy and task-oriented targets, embodiments of the present invention allow the operator to specify a goal for the robot to achieve with the use of a target icon. In some instances, the robot is then free to determine, on its own, how to get to the destination. However, in other instances, such as, for example, when the task-oriented target is placed closer to the robot, when the task-oriented target is moving relatively frequently, or when the task-oriented target is inside some event horizons, the robot is more directly controlled. In these direct control situations, where the robot has a low robot initiative, the user more directly controls the robot's operation by placement, movement, or modification of the task-oriented target.

[00339] In a very general sense, when the target is placed farther from the robot, it is as if there is lots of play in a leash between the user and the robot and the navigational responsibilities fall more on the robot. In this approach, the operator is less concerned with how the robot will get to the destination. Consequently, the user's attention can be focused more on managing the overall task or mission as opposed to the actual movement of the robot. The end effect is that the operator's cognitive capabilities are freed from the navigation responsibility and allowed to concentrate on tasks more appropriate for the human to consider.

[00340] The environment map 1390 depicts the portions of an interior space that the robot comprehends up to this point in time. As the robot progresses through the environment, it may fill out more details of the environment map 1390. In addition, details of the environment map 1390 may be filled in from other sources, such as, for example, GPS information, or other information that may be supplied by the user interface. Occupancy grid-based maps may be a good choice of map because they provide a higher resolution for tasking, however, any map representation can work as long as the robot is localized with respect to the map. Furthermore, the map may depict any environment that the robot is in, not just an interior space as depicted in FIGS. 31A-31D, and 32.

[00341] In FIG. 31A, the robot's current position and pose 1310A is shown in the upper-right corner of the environment map 1390. A task-oriented target 1330 is shown in the lower-left corner of the environment map 1390. The task-oriented target is placed by the user to designate a specific task that the user wants the robot to achieve. In this case, the task-oriented target 1330 is a navigation target 1330 that the robot should go to. With this placement of the navigation target 1330, the distance between the robot and the navigation target is large enough that the user interface directs the robot to achieve the target with a high degree of autonomy. In other words, the user interface may send instructions to develop a robot plan. The robot plan may include planning a path to the navigation target and to begin moving toward the target

along the planned path 1315. The robot reports the planned path 1315 back to the user interface and the user interface displays the planned path as a series of waypoints in the environment map 1390. Once planned and displayed, the user may adjust the waypoints if the user believes there may be a better way to approach the navigation target 1330. Thus, the robot can still achieve the navigation target 1330 with a high degree of autonomy and the user is free to concentrate on other operations.

[00342] Also illustrated in FIG. 31A is another task-oriented target 1320 positioned near the center of the environment map 1390. This task-oriented target is an imaging target 1320. Thus, the user interface sends instructions to the robot to carry out robot behaviors to point an imaging device (e.g., a camera) toward the imaging target 1320 and to keep pointing at the imaging target 1320. As the robot moves, the imaging device will automatically develop a robot plan to adjust the imaging device (e.g., pan, tilt, focus) to stay on the imaging target with a high degree of autonomy requiring little or no intervention from the user. Also illustrated in FIG. 31A is an imaging window 1325A. The imaging window 1325A shows the current view of the camera, which is pointed at, or moving to achieve, the imaging target 1320. FIGS. 31B, 31C, and 31D illustrate the same user interface window 1300 with the environment map 1390, navigation target 1330, planned path 1315, and imaging target 1320.

[00343] The robot plans for achieving the various task-oriented targets may include using a variety of robot conduct and robot behaviors available within the robot software architecture. As non-limiting examples, the conduct and behaviors may include focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

[00344] FIGS. 31B, 31C, and 31D also illustrate various robot position and pose icons 1310B, 1310C, and 1310D, respectively, as the robot traverses along the planned path 1315. Finally, FIGS. 31B, 31C, and 31D also illustrate various images perceived by the camera 1325B, 13125C, and 13125D, respectively, as the robot traverses along the planned path 1315 while the camera autonomously adjusts to point at the imaging target 1320.

[00345] FIG. 32 is a user interface window 1300 showing the environment map 1390, a robot position and pose 1310E, an imaging target 1320, and a planned path 1315. However, in FIG. 32, the user has moved the navigation target from a first position 1330A to second position 1330B. This new navigation position 1330B is much closer to the current position of the robot, as is indicated by vector 1318. With this movement, the user interface will direct the robot to abandon the currently planned path 1315. Since the second target is much closer to the robot, rather than telling the robot to plan a path, the user interface will instruct the robot to enter a

lower initiative mode and follow instructions from the user interface. Note that the instructions are still not coming directly from the user as in instructions that may be generated by a user using a joystick to control the robot's movements. Rather, the instructions are generated by the user interface based on the vector 1318 between the robot and the navigation target 1330B. In this low initiative mode for the robot, one can view the navigation target 1330B from the user's perspective as a carrot for the robot to follow. The user may drag the navigation target 1330B around and as the navigation target 1330B moves, the user interface determines a new vector 1318 and sends new instructions to the robot on how to move to achieve the new position of the navigation target 1330B. This moving of a target for the robot to chase after may be much more intuitive for a user and the robot can proceed more directly to the intended point in the environment map. Thus, the navigation target 1330B may eliminate, or greatly reduce, the possibility of missing the intended point by being off to one side, going too far, or stopping short, all of which can easily happen if the user is in direct control of driving the robot with a joystick. In other words, the performance in achieving a target is better using the navigation target 1330B because the users perception is map centric (i.e., where on the map the target is), rather than view centric (i.e., where in my field of view am I trying to get the robot to go).

[00346] The instruction from the user interface for achieving the task-oriented target in these lower initiative modes may include using a variety of robot conduct and robot behaviors available within the robot software architecture. As non-limiting examples, the conduct and behaviors may include focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

[00347] With these task-oriented targets, a user with less trust in the robot can keep direct control of the robot by keeping the target relatively close to the robot and moving the target as the robot moves, similar to leading the robot on a tight leash. If the user has more trust in the robot's ability to achieve the target on its own, or the user does not care how the robot achieves the target, the user can place the target farther away from the robot and move the target less often. Consequently, the changes in autonomy level for the robot are indirectly controlled by the user based on where a task-oriented target is placed, how often it moves, and whether it is inside an event horizon. Event horizons were explained earlier and also will be discussed later when discussing the robot's actions in the seamless autonomy system.

[00348] Navigation targets 1330 and visual targets 1320 have been illustrated and discussed. However, other task-oriented targets are contemplated within the scope of the present invention. As non-limiting examples, the task-oriented target may be a navigation target, an imaging target, an artillery target, a sensor target, a manipulator target, and the like.

[00349] In addition, as illustrated in FIGS. 31A-31D and 32, multiple task-oriented targets may be used at any given time. As a non-limiting example of multiple task-oriented targets, a user may use a navigation target to direct the robot toward a certain point in the environment, a sensor target to direct the robot to sense for a parameter of interest near the robot, and an imaging target to direct the robot to look behind so the user can determine if anything is approaching near the robot.

[00350] Task-oriented targets may be placed in a two-dimensional map or a three-dimensional map and include two-dimensional or three-dimensional coordinates. In addition, task-oriented targets may include task attribute information beyond just location within an environment map. As a non-limiting example, a task-oriented target may be associated with a manipulator as a manipulator target. This manipulator target may be placed, for example, on a door handle. Thus, the user interface will direct the robot to use the manipulator to achieve the manipulator target on the door handle. Other task attribute information assigned to the manipulator target may include information such as which way to turn the door handle, which way to push the door, and similar attributes associated with the manipulator in question. As another non-limiting example, a bomb sniffer element may be capable of detecting different types of bombs based on metal content, chemical content, and the like. The specific types of attributes that the bomb sniffer is searching for may be defined as attributes attached to a bomb-sniffer target. As another non-limiting example, a navigation target may be defined as a center-point around which the robot should circle. Thus, task attributes for a navigation target may include an orbit attribute and a radius attribute for the orbit. As yet another example, offsets and pose information may be defined relative to a navigation target. The offset may define where the robot should stop when it has achieved the target. For example, an offset attribute may define that the robot should stop one meter in front of the target, to the side of the target, or centered on the target. As an example, the robot may include a manipulator or sensor that extends one meter in front of the robot. Thus, if the offset attribute is set to one meter in front of the robot, the robot stops in an optimal position for using the manipulator without additional requirements from the user to fine-tune the robot's position before using the manipulator.

[00351] In addition, task-oriented targets may be used in a wide variety of robot platforms, such as, for example, Unmanned Ground Vehicles (UGVs), Unmanned Air Vehicles (UAVs), Unattended Ground Sensors (UGSs), and Unmanned Water Vehicles such as unmanned surface vehicles (USVs) and unmanned Underwater vehicles (UUVs).

[00352] As an example of the strength of task-oriented targets, in a UAV an image target may be placed on a specific location on the ground. Then the user can direct the UAV to

circle, or directly navigate the UAV using a navigation target. All the while the user is directing navigation, possibly with high user intervention and low robot initiative, the user interface and robot are cooperating to keep the camera directed at and focused on the image target with low user intervention and high robot initiative.

5.1. USER INTERFACE

[00353] FIG. 33 is a simplified flow chart illustrating components of a user interface algorithm 1400 for controlling seamless task-oriented autonomy adjustments in accordance with one or more embodiments of the present invention.

[00354] A common operating picture block 1420 may receive information from a variety of sources to develop an understanding of the robot and its environment. Non-limiting examples of these sources are, UGV data 1402, UAV data 1404, UGS data 1406, Geospatial Information System (GIS) information 1412, and Sensor data 1410, either directly or via a sensor map 1408. The common operating picture 1420 includes the GUI that is presented to the user as well as all the collected information from the various data sources and information about the robot and the various task designators (i.e., task-oriented targets).

[00355] The common operating picture 1420 also receives information about the task designators in the form of intended tasks 1418, commanded tasks 1416, and confirmed tasks 1414, as is explained more fully below.

[00356] Operation block 1422 indicates user input in the form of a mouse, a joystick, a keyboard, haptic devices, and other suitable input devices for conveying a user's intentions to the user interface. The user's intentions may be actions for the robot to perform, such as, for example, where to go, what to do, what to look for, what to detect, and the like. The user interface algorithm 1400 illustrated herein is a simplified flow chart intended to show general actions performed as a result of the user placing, moving, and modifying task-oriented targets. Those of ordinary skill in the art will recognize that there may be many other inputs from the user not directly related to achieving task-oriented targets.

[00357] Based on the user input, and other information gathered by the common operating picture 1420, the user interface determines the appropriate tasks and autonomy level for the robot. As was explained above, in a very general sense, the autonomy level will have a low robot initiative if the task-oriented target is close to the robot, moving relatively frequently, or within an event horizon.

[00358] Operation block 1426 tests to see if the task-oriented target meets certain filtering requirements and filters information about the task-oriented target to meet those

filtering requirements. These filtering requirements may be time based and distance based. As a non-limiting example, if the user is dragging a task-oriented target around in the environment map, it may not be advisable to flood the communication channel between the user interface and robot with many small changes in position. Rather, the user interface may attempt to filter out many of these small changes to present fewer position changes to the robot based on communication capabilities. After creating the filtered information, operation block 1418 indicates that the intended task, based on the autonomy and task switching 1424 and filtering 1426 is communicated back to the common operating picture 1420 for display to the user. The intended task indicates what the user interface intends to instruct the robot to do.

[00359] In parallel with generating the intended task 1418, decision block 1428 tests to see if the map is geo-referenced. In other words, can the map be tied to Global Positioning System (GPS) coordinates, Universal Transverse Mercator (UTM) coordinates, or other suitable geo-referenced coordinate system? If the map is geo-referenced operation block 1430 determines the current geo-referenced coordinates for the task-oriented target. If the map is not geo-referenced operation block 1432 determines the current coordinates for the task-oriented target in terms of another suitable map, such as environment map coordinates.

[00360] Operation block 1434 determines the task coordinates from whichever map coordinates are provided. As non-limiting examples, these coordinates may be relative to the robot's current position and pose, relative to a specific sensor, relative to a specific manipulator, or relative to the robot's understanding of the environment.

[00361] The decision block string of 1440, 1460, and 1470 are example decisions that may be made for various task-oriented targets. Along with operation blocks 1442, 1462, and 1472, these decisions and operations are intended to be exemplary and not to cover the details of each possible task-oriented target or all possible operations for each task-oriented target.

[00362] Decision block 1440 determines if a navigation task needs to be addressed. If so, operation block 1442 determines what type of navigation commands should be sent to the robot. As a non-limiting example, details of the type of decisions and commands that may be performed are illustrated for operation block 1442. Decision block 1444 tests to see if the robot is a UGV. If so, decision block 1446 tests to see if the current task is a high intervention task or a high initiative task. As explained earlier, if the target is close to the robot, or the target is moving relatively frequently, the user interface decides that high intervention should be used.

[00363] In addition, if the target is within an event horizon high intervention may be appropriate. As a non-limiting example, the distance between the target and the robot may be close enough that the time to communicate instructions to the robot, perform a path plan on the

robot, and communicate the path plan back to the user interface may be too long. In other words, at its present speed, the robot may reach the target before the planning operations can be performed. In these scenarios, high intervention is more appropriate because there is not enough time to allow high initiative.

[00364] If high intervention is appropriate, operation block 1448 sends a waypoint command to the robot. The waypoint command indicates to the robot where in its environment it should immediately head toward.

[00365] If high initiative is appropriate, operation block 1450 sends a “goto” command to the robot. The robot can then use its initiative to attempt to achieve the navigation target. The robot then plans a path to the navigation target and communicates the planned path back to the user interface for display to the user.

[00366] If the robot is not a UGV, decision block 1452 tests to see if the robot is a UAV. If not, operation block 1458 indicates that other vehicle navigation commands should be sent to the robot. As a non-limiting example, the robot may be an underwater vehicle and decision should be made about the underwater vehicle’s capabilities relative to the navigation target. Then, appropriate navigation commands at various autonomy levels can be sent to the underwater vehicle.

[00367] If the robot is a UAV, decision block 1454 tests to see if the current navigation task is to orbit around the navigation target. If the task is not an orbit task, operation block 1450 sends a “goto” command as was described above for the UGV. If the task is an orbit task, operation block 1456 sends an orbit command.

[00368] Those of ordinary skill in the art will recognize that operation block 1442 for handling navigation targets is a simplified block diagram. Many other types of robot vehicles may be involved as well as many other navigation decisions and navigation commands.

[00369] Returning to decision block 1440, if the current task is not a navigation task, decision block 1460 tests to see if the current task is a camera task. If so, operation block 1462 sends appropriate high robot initiative or low robot initiative commands to control a camera to the robot. As non-limiting examples, commands may be related to operations such as pan, tilt, focus, and zoom.

[00370] If the current task is not a camera task, decision block 1470 tests to see if the current task is a payload task. If the current task is a payload task, operation block 1472 sends appropriate high robot initiative or low robot initiative commands to control a payload. Non-limiting examples of payloads are bomb sniffers, manipulators, chemical sensors and the like.

[00371] If the current task does not appear to relate to tasks for the robot to perform, operation block 1480 tests to see if the user wishes to change the map perspective and, if so, changes map views based on user instructions.

[00372] Operation blocks 1462, 1472, 1458, 1456, 1450, and 1448 all are operations sending instruction to the robot and are indicated as such by using a bold box. As part of sending the instruction to the robot, the user interface also defines the commanded task 1416 and sends it to the common operating picture 1420 for display to the user.

[00373] Operation block 1490 sets the commanded task indicating that the robot has received the commanded task and has acknowledged the instructions. The acknowledgement may be a simple reply or a complex reply that may include information about what the robot intends to do, such as, for example, a waypoint list for a planned path. As part of the acknowledgement, the user interface reports a confirmed task 1414 to the common operating picture 1420 for display to the user. Thus, the common operating picture 1420, and the user if appropriate, is aware of tasks to be performed at three levels; an intended task 1418 not yet sent to the robot, a commanded task 1416 indicating what the user interface wants the robot to perform, and a confirmed task 1414 indicating that the robot acknowledges that the task is being attempted.

5.2. ROBOT ACTIONS

[00374] FIG. 33 is a simplified flow chart illustrating components of a robot algorithm for controlling seamless task-oriented autonomy adjustments on the robot in accordance with one or more embodiments of the present invention. Operations 1500, shown above the dashed line, are operations performed by the user interface. Operations 1520, shown below the dashed line, are operations performed by the robot platform. A change to the current state of the seamless autonomy related to task-oriented targets generally begins with a hotspot (i.e. task-oriented target) movement 1502, which may include hotspot placement and changes to task attributes. The hotspot movement along with other information is communicated to the common operating picture 1504 (designated as 1420 in FIG. 33) as was explained above with respect to FIG. 33. Movement filtering 1506 (designated as 1426 in FIG. 33) with respect to time and distance may be performed.

[00375] As non-limiting examples, task-oriented targets may be related to manipulators 1510, artillery 1512, navigation 1516, and cameras 1514. As non-limiting examples, cameras may include still and video images generated for visual images, thermal images, and gamma images.

[00376] Robot instructions related to task-oriented targets may call a focus entity abstraction 1525 on the robot. Other robot instructions, such as, for example, navigation targets, and possibly manipulator targets may go directly to operation block 1530.

[00377] The focus entity abstraction 1525 generally may be used in connection with visual hotspots for various visual tracking behaviors, such as, for example, determining line of sight, tracking, determining distance, and the like.

[00378] Operation block 1530 determines the hotspot position relative to the robot's local pose and position. This may include conversion to a robot coordinate system or an environment map coordinate system. In other words, the current hotspot position is put in robot understandable coordinates.

[00379] Decision block 1532 tests for whether the current hotspot being analyzed is a navigation target. If so, decision block 1550 tests to see if the hotspot is within a planning event horizon. As discussed earlier, the distance between the target and the robot may be close enough that the robot may not have time to perform a path-planning algorithm. In other words, at its present speed, the robot may reach the target before the planning operations can be performed. If the navigation target is within the planning event horizon, the robot may use a real-time reactive pursuit behavior 1596, such as, for example, the robotic follow conduct discussed above. During performance of the pursuit behavior 1596, dynamic re-planning 1552 may be requested because the navigation target may move outside of the event horizon enabling the possibility to generate a path plan. After dynamic re-planning 1552, control returns to decision block 1550 to see if the target is still within the event horizon.

[00380] If the navigation target is outside the event horizon, decision block 1554 tests to see if the navigation target is within the map boundaries. The navigation target placed by the user may actually be outside the boundaries of the map the robot has of its environment. If the navigation target is not within the map boundary, operation block 1556 finds the nearest point on the map boundary based on where a vector between the robot and the target crosses the map boundary. As the robot moves toward this new point on the map boundary, it may discover more about its environment on its own or through the user interface enabling the robot to re-plan toward the navigation target. After finding the perimeter point 1556, control returns to operation block 1530 to determine robot position relative to the perimeter point.

[00381] If the navigation target is within the map boundary, operation block 1558 determines a path plan and operation block 1560 filters out some of the waypoints in the path plan if appropriate. For example, the filter may remove multiple waypoints in a substantially straight line or reduce the number of waypoints on a curve. With the path plan determined and

filtered, a waypoint following conduct 1598, such as, for example, the virtual rail conduct discussed above, is performed to achieve the navigation target.

[00382] Returning to decision block 1532, if the current hotspot is not a navigation target, the target may be related to a payload such as a sensor, manipulator, or camera. Operation block 1534 calculates the angle of error relative to a payload offset 1536. A payload offset may indicate a payloads position relative to the robot. As a non-limiting example, a camera may be mounted on the back of a robot and compensation may be needed relative to the center of the robot.

[00383] Operation block 1538 calculates the needed pan/tilt/zoom movement, or other appropriate movements for manipulators or sensors, needed to achieve the hotspot. Information about the payload state 1540 may be used in the calculations of operation block 1538. Decision block 1542 tests to see if the calculated movements are possible. For example, a manipulator may extend beyond the robot and movement of the manipulator in the desired manner may cause it to meet an obstacle. If the payload movement is possible, an appropriate payload actuation behavior 1544 is executed.

[00384] If the payload movement is not possible, operation block 1546 activates a navigation task 1546 in an effort to move the robot to a position, pose, or combination thereof that will allow the desired payload movement.

[00385] In some instances, the robot may have better knowledge of the environment than the user. In these cases, the robot may move the navigation target, or modify waypoints in a path plan. As a non-limiting example, the hotspot may be on an opposite side of a wall that the user was not aware of or did not see. In such cases, the robot may move the hotspot, or create a new path plan to get around the wall. Similarly, a new obstacle may be placed in the robot's current path, causing the robot to re-plan a path or move the hotspot.

As with FIG. 33, the simplified flowchart of FIG. 34 is intended to illustrate only some of the operations and decisions that the robot may need to perform to achieve a task-oriented target with an emphasis on navigation decisions and operations. Those of ordinary skill in the art will recognize that many other navigation decisions and operation are included within the scope of the present invention. Similarly, those of ordinary skill in the art will recognize that for the sake of brevity, decisions and operation related to other task-oriented targets would be similar to those for navigation and need not be described in detail herein.

[00386] Although this invention has been described with reference to particular embodiments, the invention is not limited to these described embodiments. Rather, the invention is encompassed by the appended claims and their legal equivalents

CLAIMS

What is claimed is:

5

1. A graphical user interface for controlling a robot, comprising:
an environment map window configured for displaying a map of an environment proximate the
robot;
a robot designator configured for showing a robot position and a robot pose in the environment
10 map window;
at least one task designator configured for positioning by a user in the environment map window
and indicating a task for the robot to achieve; and
a control intermediary configured for linking user defined tasks with robot instructions by:
analyzing a position of the at least one task designator relative to a position of at least
15 one robot component;
determining a task-oriented autonomy level for the robot responsive to the analysis; and
communicating target achievement information to the robot, wherein the target
achievement information comprises:
instructions from the control intermediary for guiding the robot to achieve the
20 task if the task-oriented autonomy level comprises low robot initiative;
and
instructions from the control intermediary directing the robot to determine a
robot plan for achieving the task if the task-oriented autonomy level
comprises high robot initiative.

25

2. The graphical user interface of claim 1, wherein the map is selected from the group
consisting of a two-dimensional map and a three-dimensional map.

30

3. The graphical user interface of claim 1, wherein the at least one task designator includes
position information selected from the group consisting of two-dimensional information and
three-dimensional information.

4. The graphical user interface of claim 1, wherein the at least one task designator includes
task attribute information.

5. The graphical user interface of claim 1, wherein the at least one task designator is selected from the group consisting of a navigation target, an imaging target, an artillery target, a sensor target, and a manipulator target.

5

6. The method of claim 1, wherein analyzing a position further comprises analyzing a change in the at least one task designator, wherein the change is selected from the group consisting of:

a time interval between movements of the at least one task designator;

10 a distance between the robot and the at least one task designator;

an event horizon relative to the at least one task designator;

a change from one task designator to another task designator; and

combinations thereof.

15 7. The method of claim 1, wherein instructions from the control intermediary for guiding the robot include instructions for robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

20 8. A method for controlling a robot, comprising:

providing a user interface for controlling the robot;

positioning at least one task-oriented target in a map on the user interface representing an environment of the robot;

determining a task-oriented autonomy level of the robot correlated to a change in the at least one

25 task-oriented target;

if the change in the at least one task-oriented target is smaller than a change threshold,

instructing the robot to achieve the at least one task-oriented target by intervention instructions from the user interface; and

if the change in the at least one task-oriented target is larger than the change threshold,

30 instructing the robot to achieve the at least one task-oriented target using robot initiative to determine a robot plan for achieving the at least one task-oriented target and implementing the robot plan.

9. The method of claim 8, further comprising selecting the at least one task-oriented target from the group consisting of a navigation target, an imaging target, an artillery target, a sensor target, and a manipulator target.

5 10. The method of claim 8, wherein the change in the at least one task-oriented target is selected from the group consisting of:

a time interval between movements of the at least one task-oriented target;

a distance between the robot and the at least one task-oriented target;

an event horizon relative to the at least one task-oriented target;

10 a change from one task-oriented target to another task-oriented target; and combinations thereof.

11. The method of claim 8, wherein the user interface adjusts the task-oriented autonomy level between a level selected from the group consisting of:

15 a teleoperation mode configured to maximize a user interface intervention and minimize the robot initiative;

a safe mode configured to include less of the user interface intervention and more of the robot initiative relative to the teleoperation mode;

20 a shared mode configured to include less of the user interface intervention and more of the robot initiative relative to the safe mode;

a collaborative tasking mode configured to include less of the user interface intervention and more of the robot initiative relative to the shared mode; and

an autonomous mode configured to minimize the user interface intervention and maximize the robot initiative.

25 12. The method of claim 8, wherein the intervention instructions include instructions for robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

30 13. The method of claim 8, wherein the robot plan includes robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

14. The method of claim 8, wherein positioning the at least one task-oriented target in the map is performed by a user and the intervention instructions are determined by the user interface responsive to the change in the at least one task-oriented target.

5 15. A method for controlling a robot, comprising:

receiving instructions for achieving at least one task-oriented target from a user interface, the instructions comprising at least one of intervention instructions, robot initiative instructions, and instructions for setting a task-oriented autonomy level;

10 if the instructions are the robot initiative instructions, then developing a robot plan to achieve the at least one task-oriented target and performing the robot plan; and

if the instructions are the intervention instructions, then performing the intervention instructions for achieving the at least one task-oriented target and, if present, overriding the robot plan to achieve the at least one task-oriented target.

15 16. The method of claim 15, further comprising reporting the robot plan to the user interface.

17. The method of claim 15, wherein the intervention instructions include instructions for robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

18. The method of claim 15, wherein the robot plan includes robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

25 19. The method of claim 15, further comprising returning to the performing the robot plan after performing the intervention instructions if the user interface modifies the task-oriented autonomy level to provide additional robot initiative.

30 20. The method of claim 15, further comprising selecting the at least one task-oriented target from the group consisting of a navigation target, an imaging target, an artillery target, a sensor target, and a manipulator target.

21. The method of claim 15, wherein the user interface adjusts the task-oriented autonomy level between a level selected from the group consisting of:

a teleoperation mode configured to maximize a user interface intervention and minimize the robot initiative;

5 a safe mode configured to include less of the user interface intervention and more of the robot initiative relative to the teleoperation mode;

a shared mode configured to include less of the user interface intervention and more of the robot initiative relative to the safe mode;

10 a collaborative tasking mode configured to include less of the user interface intervention and more of the robot initiative relative to the shared mode; and

an autonomous mode configured to minimize the user interface intervention and maximize the robot initiative.

22. A robot platform, comprising:

15 at least one perceptor configured for perceiving environmental variables of interest;

at least one locomotor configured for providing mobility to the robot platform; and

a system controller configured for executing a task-oriented autonomy system, comprising:

receiving instructions from a user interface using a communication interface, the

20 instruction for achieving at least one task-oriented target and comprising at least one of intervention instructions, robot initiative instructions, and instructions for setting a task-oriented autonomy level;

if the instructions are the robot initiative instructions, developing a robot plan to achieve the at least one task-oriented target and performing the robot plan; and

25 if the instructions are the intervention instructions, then performing the intervention instructions for achieving the at least one task-oriented target from the user interface and, if present, overriding the robot plan to achieve the at least one task-oriented target.

30 23. The robot platform of claim 22, further configured for reporting the robot plan through the communication interface to the user interface.

24. The robot platform of claim 22, wherein the system controller is further configured for returning to performing the robot plan after performing the intervention instructions if the user interface modifies the task-oriented autonomy level to provide additional robot initiative.

25. The robot platform of claim 22, wherein the system controller is further configured for returning to performing the robot plan after performing the intervention instructions if the robot determines that additional robot initiative is appropriate and the user interface enables a change in autonomy initiated by the robot.

26. The robot platform of claim 22, wherein the at least one task-oriented target is selected from the group consisting of a navigation target, an imaging target, an artillery target, a sensor target, and a manipulator target.

27. The robot platform of claim 22, wherein the robot plan includes robot conduct and robot behavior selected from the group consisting of focus behaviors, manipulation behaviors, go-to points, waypoints, path-planning, search region, patrol region, retro-traverse, and laser tracking.

28. A robot control system, comprising:
a user computer comprising a memory and at least one processor configured for executing a user interface, the user interface configured for:
positioning at least one task-oriented target in a map on the user interface representing an environment of the robot;
determining a task-oriented autonomy level of the robot correlated to a change in the at least one task-oriented target;
if the change in the at least one task-oriented target is smaller than a change threshold, instructing the robot to achieve the at least one task-oriented target by intervention instructions from the user interface; and
if the change in the at least one task-oriented target is larger than the change threshold, instructing the robot to achieve the at least one task-oriented target using robot initiative to determine a robot plan for achieving the at least one task-oriented target and implementing the robot plan; and
a robot platform comprising a locomotor and a system controller, the system controller configured for:
developing the robot plan to achieve the at least one task-oriented target and performing the robot plan if robot initiative instructions are received from the user computer; and

performing the intervention instructions for achieving the at least one task-oriented target
and, if appropriate, overriding the robot plan to achieve the at least one task-
oriented target if the intervention instructions are received from the user
computer.

5

29. The system of claim 28, wherein the task-oriented target is selected from the group
consisting of a navigation target, an imaging target, an artillery target, a sensor target, and a
manipulator target.

10

30. The system of claim 28, wherein the change in the at least one task-oriented target is
selected from the group consisting of:

a time interval between movements of the at least one task-oriented target;

a distance between the robot and the at least one task-oriented target;

an event horizon relative to the at least one task-oriented target;

15



a change from one task-oriented target to another task-oriented target; and
combinations thereof.

31. The system of claim 28, wherein positioning the at least one task-oriented target in the
map is performed by a user and the intervention instructions are derived by the user interface

20

responsive to the change in the at least one task-oriented target.

1/34

**FIG. 1****FIG. 2**

200 →




FIG. 3

3/34





FIG. 4

4/34

**FIG. 5**






FIG. 6


6/34

**FIG. 7**


7/34

**FIG. 8**

8/34

**FIG. 9**

Autonomy Mode	Defines Task Goals	Supervises Direction	Motivates Motion	Prevents Collision
<i>Teleoperation mode 293</i>	Operator	Operator	Operator	Operator
<i>Safe Mode 294</i>	Operator	Operator	Operator	Robot
<i>Shared Mode 295</i>	Operator	Operator	Robot	Robot
<i>Collaborative Tasking Mode 296</i>	Operator	Robot	Robot	Robot
<i>Autonomous Mode 297</i>	Robot	Robot	Robot	Robot

FIG. 10A**FIG. 10B**

Robot Intelligence Kernel Communication






FIG. 11

Robot Intelligence Kernel Abstractions

**FIG. 12**

Intelligence Kernel Behavior Modulation

**FIG. 13**

13/34

Guarded Motion Behavior

500

Execute for each direction of motion using resistance, bump, and range regions for each direction (forward, left, right, back)




FIG. 14

Obstacle Avoidance Behavior Translational Velocity




FIG. 15

Obstacle Avoidance Behavior Rotational Velocity






FIG. 16

Get-Unstuck Behavior**FIG. 17**

17/34

ROCA Behavior**FIG. 18**

18/34




FIG. 19

19/34





FIG. 20

20/34

**FIG. 21**

21/34






FIG. 22


22/34

**FIG. 23**


23/34

**FIG. 24**


24/34

**FIG. 25**

25/34

Waypoint Handler**FIG. 26**

Waypoint Following Behavior Translational Velocity



Waypoint Following Behavior Rotational Velocity




FIG. 28

28/34




FIG. 29




FIG. 30




FIG. 31A




FIG. 31B




FIG. 31C




FIG. 31D





FIG. 32



34 / 34




FIG. 34

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 09/36685

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G05B 19/18 (2009.01)
USPC - 700/250

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)
USPC: 700/250

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 700/245 (keyword limited - see terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PubWEST(PGPB, USPT, USOC, EPAB, JPAB); GOOGLE
Search Terms Used: Graphical user interface, environment, map, initiative, intervention, autonomous, robot, task, target, level, instruction, command, receiving, autonomy.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2004/0073360 A1 (Foxlin) 15 April 2004 (15.04.2004), entire document, especially; abstract, para [0194], [0164], [0170].	1-14, 22-31
Y	US 2008/0009967 A1 (Brummer) 10 January 2008 (10.01.2008), entire document, especially; abstract, para [0013]-[0015], [0122], [0138], [0141]-[0148], [0153]- [0169], [0176], [0182], [0186], [0193], [0212], [0213], [0285], [0295]-[0299], [0382].	1-31
Y	US 2006/0095160 A1 (Orita et al.) 04 May 2006 (04.05.2006), entire document, especially; abstract, para [0075], [0089].	15-21

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search
22 October 2009 (22.10.2009)

Date of mailing of the international search report

02 NOV 2009

Name and mailing address of the ISA/US
Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

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