SEMI-AUTONOMOUS CONTROL OF AN EMERGENCY RESPONSE ROBOT

Josh Levinger, Andreas Hofmann, Daniel Theobald

Vecna Technologies, 36 Cambridgepark Drive, Cambridge, MA, 02140,
Tel: 617.864.0636 Fax: 617.864.0638 Web: www.vecnarobotics.com
jlevinger@vecna.com, ahofmann@vecna.com, theo@vecna.com

ABSTRACT

The Vecna Battlefield Extraction Assist Robot (BEAR) is designed to locate and rescue people in harm's way. Whether on a battlefield, at a nuclear facility, or inside a structurally-compromised building, the BEAR can rescue those in need without risking additional human life. This type of mission requires the capability to lift heavy objects, including people as well as rubble, and to traverse uneven terrain. Tasks must be performed quickly and safely, while reacting robustly to disturbances in the environment. The operator interface must allow the operator to control the robot in a natural and intuitive way. In particular, the operator should not be burdened with the details of control, since this would interfere with the ability to perform tasks quickly.

The BEAR addresses these demanding requirements through a powerful hydraulically actuated body that features two independent sets of tracked legs, and two six degree of freedom arms. This highly articulated mechanism is controlled using a semi-autonomous approach that combines tele-operated control with autonomous behaviors for basic tasks such as locomotion, grasping, and lifting. This allows the operator to focus on where the robot should go and what it should pick up, rather than the details of how these tasks should be accomplished.

I. INTRODUCTION

The BEAR is humanoid in form, as shown in Fig. 1. Its upper body consists of a steel torso, a pair of six degree of freedom arms, and a head that can pan and tilt. Hydraulic actuation allows the arms to lift over 500 lbs. Flat, soft surfaces are used on the arms to maximize comfort when lifting humans. The head contains a visible-light camera, as well as an infrared camera, allowing for operation in dark environments. The BEAR's lower body features a hybrid tracked/legged design, which allows for a variety of modes of locomotion, depending on terrain conditions. A tank-like mode is used for level to moderately difficult terrain. In this mode, the individually controllable legs provide active suspension. Standing mode is used to walk on very difficult terrain, to step over obstacles, to go up and down stairs, and to climb. A novel hybrid locomotion mode is also possible, where an upright, standing posture is used, as when walking, but the tracks are also turning to maintain balance. This allows for fast traversal of level terrain.

There are significant challenges in controlling such a system. The mechanism is highly nonlinear, and is difficult to balance, especially when carrying a load over uneven terrain. The unstructured environments in which it operates require an unprecedented combination of strength, speed, safety, and robustness to disturbances. The system must also be easy to use. This implies that the BEAR must be able to understand basic task-level commands from humans. Our approach features a library of whole-body maneuvers that provide basic grasping, manipulation and navigation behaviors. Each maneuver contains a control policy

Fig. 1 – The BEAR has a humanoid morphology, allowing it to operate in environments intended for humans.

that guides the robot. Operational limits of the maneuver are also represented, so that the system can tell whether a disturbance will jeopardize successful completion of the maneuver. This information is valuable because it allows the system to abort a maneuver sooner rather than later if there is no chance of successful completion. Maneuvers are assembled into sequences according to task requirements. These sequences are executed by a task executive that monitors state, issues control commands, and switches to the subsequent maneuver when the goals for the current maneuver have been achieved.

The next section specifies operational requirements for the BEAR and describes challenges associated with these requirements. Section III describes the BEAR mechanism, including the torso and the mobility platform. Section IV describes the semi-autonomous control approach. Section V summarizes experimental results and discusses conclusions.

II. OPERATIONAL REQUIREMENTS AND CHALLENGES

The Vecna BEAR robot represents a promising approach to safely extracting combat casualties from urban and wooded terrain or from other areas with numerous obstacles, as shown in Fig. 2. It can operate in environments that are hazardous to humans, such as those containing nuclear, chemical or biological (NBC) hazards, as shown in Fig. 3. It can carry JAUS-compatible mine-sensing payloads for mine sweeping, or a laser for determining if an object is contaminated with nuclear, biological, or chemical materials. It can perform other search and rescue operations, or perform reconnaissance work. It can also be used for a wide range of transport activities, as shown in Fig. 4.

Figure 2 – Combat casualty extraction in an urban environment.

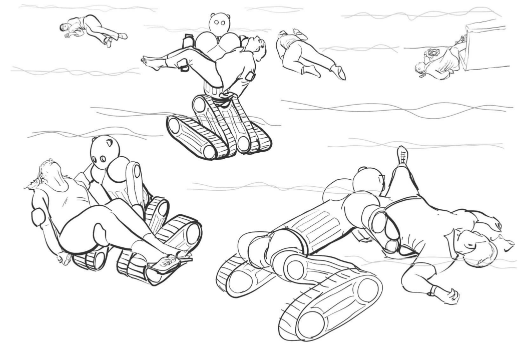

Figure 3 – Operation in NBC environments.

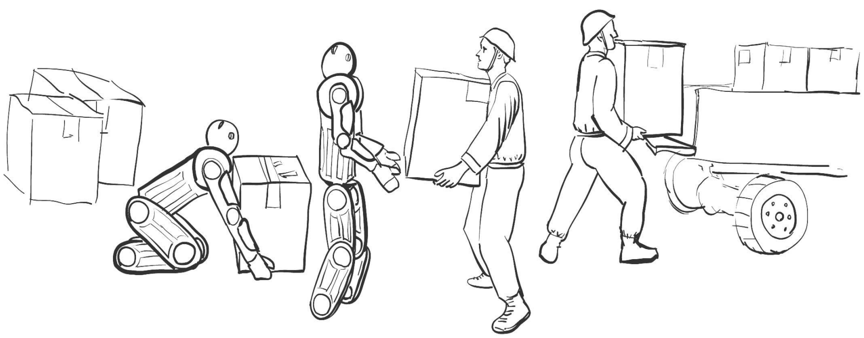

Figure 4 – Loading and transport activities.

To accomplish these kinds of missions, the BEAR must be capable of lifting objects of varying size, shape, and firmness, carrying such objects over a variety of terrain, and depositing them safely in a desired location. Such tasks must be performed safely, quickly, and efficiently. Among other things, this implies that the robot must be easy to operate by a human. The human operator should specify what the robot should do, but not how. Ideally, a human would interact with the robot as a coach would with an inexperienced assistant.

Human-like autonomy will require decades of more research, and is therefore not currently a feasible goal. However, we believe that a level of autonomy that is less than that of a human being, but significantly greater than the simple remote-control driving and tele-operation approachescurrently used in most mobile robots is both achievable, and sufficient for BEAR missions. In particular, we expect the basic mode of interaction to be one of designating an object and issuing a command. Object designation may be accomplished by laser illumination, touching a screen showing the robot's camera view, or simply by pointing at the object. Issuing commands may be accomplished by a button or keyboard interface, or by speech. For example, the operator would point to an object like a box, and would say “pick this up”. The operator might then say carry it over there, pointing to the location where it should be carried. The robot determines how to accomplish the task; the operator doesn't have to worry about the details of the robot's movements. This raises the level of interaction from one where the operator drives the robot, to one where the operator directs the robot. This approach makes the robot easier to use by the operator, and therefore more efficient.

Although this semi-autonomous capability is far below that of humans, it nevertheless implies a significant level of capability in terms of local navigation, object detection and recognition, maneuver control, and task understanding capabilities. For example, in order to maneuver to a location designated by the operator, the robot must analyze the terrain before it, decide on a locomotion mode, and plan and execute a sequence of control actions, automatically compensating for disturbances.

In the next section we describe the robot mechanism, including its morphology, and hydraulic actuators. In the subsequent section, we describe the BEAR's semi-autonomous capabilities in more detail.

III. BEAR MECHANISM

The BEAR is strong enough for lifting and transportation tasks in commercial environments, such as construction, manufacturing, and delivery, as well as for the lifting tasks described previously for military missions. The BEAR's humanoid form (Fig. 1) allows it to operate in environments intended for humans, in a manner similar to the way humans do. We now describe the upper body and mobility platform in more detail.

III A. Upper Body

The BEAR upper body is shown in Fig. 5. It consists of a torso, two arms, and a head with an articulated neck. The use of two arms is crucial for lifting large objects such as large boxes using whole-body grasps, as shown in Fig. 4; a comparably sized robot with one arm would not be able to perform this task. Two arms are also particularly useful for lifting and carrying humans, and other objects with flexibility and articulation, as shown in Figs. 2 and 3, and also in Fig. 6.

Figure 6 – The BEAR's two arms are particularly useful for carrying people.

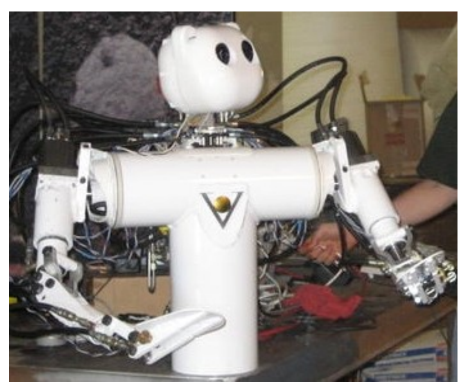

Figure 5 – BEAR Upper Body

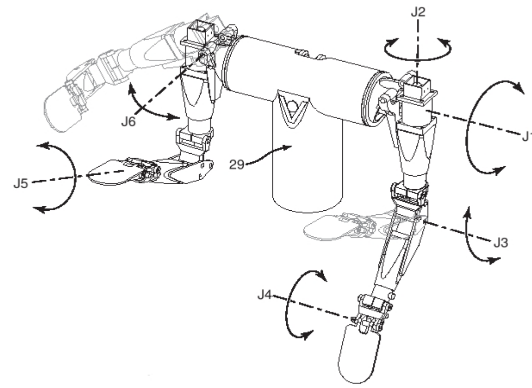

Figure 7 – The BEAR torso and arms.

The BEAR's arms each have six degrees of freedom, as shown in Fig. 7. Three revolute degrees of freedom at each shoulder, a single revolute degree of freedom at the elbow, and two revolute degrees of freedom at the wrist. Although six degrees of freedom per arm are necessary to arbitrarily choose the position and orientation of the end effector in the workspace, the human arm has seven degrees of freedom, Experience has shown that six degrees of freedom are adequate for lifting and maneuvering of humans and a wide variety of objects. Holding one's hand palm down, the yaw degree of freedom, that which allows the hand to move back and forth, is not implemented. The pitch (up and down) and roll (twisting) degrees of freedom are necessary and sufficient for lifting and transfer tasks.

The BEAR's arms have a flat, broad, and soft surface to maximize comfort when lifting humans, and potential surface contact area when lifting objects. We have tested a variety of arm lengths, widths, and materials. The arms may be contoured to optimize pressure distribution. The end effectors are currently simple paddles, but more capable end effectors, with basic grasping capability are currently being designed.

Control of the hydraulic joint actuators is accomplished through Pulse-Width Modulation (PWM) of valves associated with each actuator. A single hydraulic pump is used, which connects to the actuator valves via a hydraulic manifold. An advanced pump flow and pressure control system helps to ensure smooth movement and accurate position control. This system also uses advanced state estimation techniques to estimate actuator force, in order to support force and impedance control modes. Future development plans in this area include more fully instrumented actuators in order to further improve force control capabilities.

The head contains a visible-light camera as one "eye", and an infrared camera as the other. The latter provides for operation in dark environments. The cameras support both tele-operated and semi-autonomous control modes. In the tele-operated mode, the operator "drives" the robot based on video from the cameras. The head is attached to the torso by an electrically actuated neck that provides pan and tilt motions, allowing the cameras to sweep over a wide section of the environment in the vicinity of the robot.

III B. MOBILITY PLATFORM

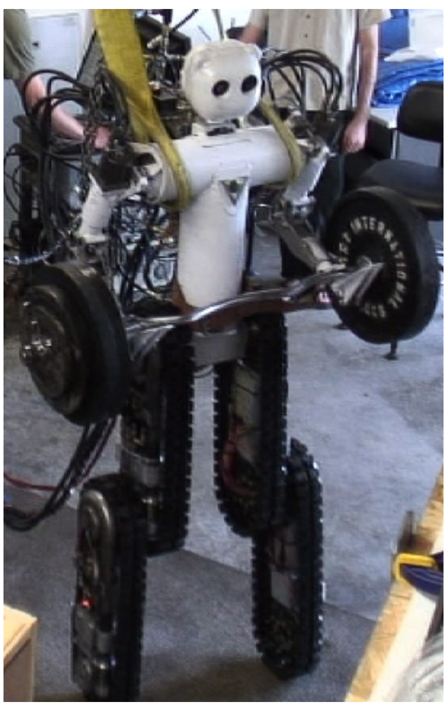

Figure 8 – The BEAR balancing in an upright position while carrying a load of 275 pounds.

A unique feature of the BEAR mobility platform are the two independent legs, providing significant position flexibility. Each upper and lower leg link is identical, and features a tread driven by an electrical motor, as well as a hinge joint driven hydraulically. For the lower leg links, the hinge joint is the knee. For the upper leg links, the hinge joint provides causing the leg to move forward and backward.

Figure 9 – The BEAR Mobility Platform

The main objective of this design is to provide independently articulated legs to improve the robot's ability to balance in both the side to side, as well as forward and backward directions. This will allow the robot to perform new movements such as bending one knee to balance on inclined planes, or placing one lower link forward of the other wheel to increase the dynamic stability of the robot climbing and descending inclined surfaces. Previous mobility platforms were restricted in lateral movement because a Segway base was used where the wheels are mechanically connected together. The new mobility platform will include a hip joint that allows the legs to be moved side to side, dramatically improving lateral balance and maneuver capabilities.

IV. BEAR SEMI-AUTONOMOUS CONTROL

Effective use of autonomous robots in unstructured, human environments requires that the robots have: 1) sufficient autonomy to understand task-level commands from humans, 2) sufficient size and strength to perform useful tasks in the environment, 3) sufficient speed to accomplish tasks in a timely manner, and 4) sufficient operating safety and robustness to disturbances while utilizing their size, strength, and speed. Note that requirements 2, 3, and 4 are often in conflict; operating with sufficient safety may require a reduction in speed or energy of movement. Ideally, such compromises are made by taking into account the full capabilities of the robot, and fully understanding the task requirements.

An important characteristic of this type of application is that there is often significant spatial and temporal flexibility in the task specifications. For example, the task of moving the robot to a position where it can pick up an object may have significant flexibility in timing, as well as the final pose of the robot; there may be many possible poses that allow the robot to reach and pick up the object. Note that this is different from typical factory manipulator "pick and place" applications, where exact timing and position repeatability are crucial.

Successful operation of robots in unstructured environments requires taking advantage of task flexibility in order to react appropriately to disturbances. Most currently existing robots do not do this. The proper exploitation of plan flexibility represents a significant gap in current capabilities.

The approach we describe here addresses this gap. We view this problem as one of dynamic plan execution, where the plan representation must capture the flexibility inherent in the goal specification. Our approach is based on techniques for dynamic execution of temporally flexible plans, but extends this using recently developed algorithms for state reachability analysis and optimal controller synthesis. The resulting plan execution system is able to take advantage of spatial and temporal flexibility in the plan specification to improve handling of disturbances while the plan is executing. This approach is superior to traditional robotic planning and control approaches that focus on generating and following individual state trajectories, and therefore, are unaware of the more complete set of execution options when a disturbance occurs.

Our approach features a representation for temporally and spatially flexible tasks, called a *Qualitative State Plan* (QSP), and a *Flexible Executive* that executes QSP's. We use the QSP representation to elevate the level of command of hybrid systems to the task level, while allowing the controller full latitude in responding to disturbances safely. The *Flexible Executive* generates control actions that achieve the QSP goals, even if disturbances occur. Key features of this executive are a plan compiler that transforms the QSP into an easily executable form called a *Qualitative Control Plan*

(QCP), and a plan dispatcher that executes the QCP.

The plan compiler incorporates compilation techniques used for temporally flexible plans[1], but extends these using recently developed algorithms for state reachability analysis and optimal controller synthesis[2]. The BEAR is a highly-nonlinear, tightly coupled system, so computing control actions that achieve a desired state is a challenging problem. To solve this problem, we linearize and decouple the robot plant into a set of loosely coupled linear plants, resulting in an *abstracted plant* that is easier to control. We accomplish this through use of a whole-body controller[3,4]. Thus, the plan compiler produces a QCP that contains feasible state and control input trajectory sets called *flow tubes*. The flow tube representation prunes infeasible trajectories from consideration at runtime, allowing the dispatcher to focus only on control actions that are feasible. Because computation of flow tubes is time consuming, and because the executive must run in real time, we perform this step off-line, as a compilation[5,6].

At execution time, the dispatcher attempts to execute tasks successfully by issuing control actions that keep the system within the allowable states defined by the flow tubes. A disturbance that displaces the system state away from the nominal trajectory can be handled successfully as long as the displaced state is within the flow tube. In this case, the executive may adjust control actions to bring the system back towards the nominal trajectory within the flow tube. If a disturbance pushes the system state outside the flow

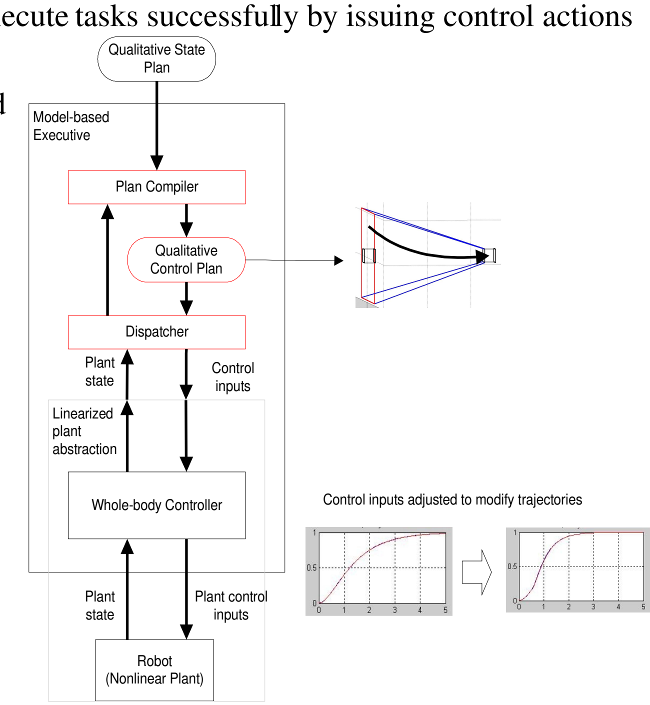

Figure 10 – Flexible Executive consisting of compiler, dispatcher, and whole-body controller

tube, then plan execution fails. At this point, a higher-level control authority must switch to a different plan. Note that we assume here that the robot plant's state can be estimated accurately. This type of estimation is performed by a hybrid mode estimator[7].

V. CONCLUSIONS AND NEXT STEPS

We now summarize results from some experiments performed with the BEAR, and discuss future development plans. Several key capabilities have been demonstrated on the BEAR robot. These include:

1. The ability to lift 500 lbs., more than twice the weight of the robot;
2. The ability to lift a fully weighted mannequin and move it to another location, setting it down on a table;
3. The ability to carry the mannequin while standing erect using dynamic balancing;
4. Driving up and down stairs while carrying a 200 lb. load;
5. Carrying one casualty while towing another;
6. Moving quickly, up to 10 – 15 mph, while carrying a 200 lb. load.

We have tested key components of the Flexible Executive architecture described previously, using high-fidelity dynamic simulations. Capabilities tested include robust balance and posture control in the presence of disturbances, and basic object lifting and balancing, as shown in Fig. 12.

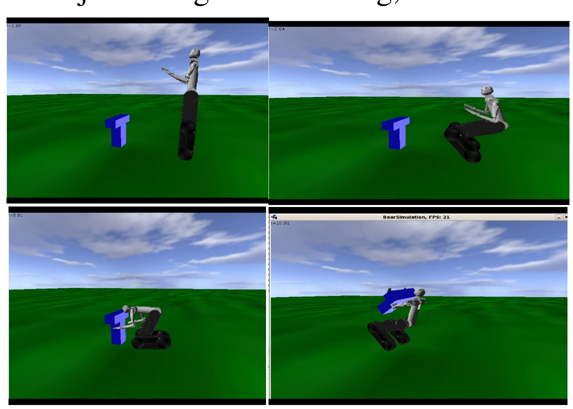

Fig. 12– The BEAR performing a lifting maneuver.

NEXT STEPS

The Flexible Executive shown in Fig. 10 is a key component in the BEAR's semi-autonomous control system, but it is supported by a number of other components currently under development. A combined vision and hybrid mode estimation system provides information about both the state of the robot and the state of the environment in the vicinity of the robot. The latter includes a terrain map in the area immediately surrounding the robot as well as identification of relevant objects in this range. A QSP assembler component is used to continually interpret task goals from the operator, and to select appropriate QSP's from a large database representing a wide range of maneuvers. Augmentation of this database of maneuvers is an on-going project.

Our goal is to produce a robot with significant local navigation and operation, capable of whole-body lifting and carrying maneuvers requiring coarse motor skills. The architecture described here will provide a level of autonomy that is less than that of a human being, but significantly greater than the simple remote-control driving and tele-operation approaches currently used in most mobile robots. In particular, it will support an object designation/command mode of interaction. In this way, the operator specifies what task should be performed, and the robot figures out how; the operator doesn't have to worry about the details of the robot's movements. This raises the level of interaction from one where the operator drives the robot, to one where the operator coaches the robot. This approach makes the robot easier to use by the operator, and therefore, more efficient.

ACKNOWLEDGEMENTS

We wish to thank the Telemedicine and Advanced Technology Research Center (TATRC), and the Tank Automotive Research and Development Center (TARDEC) for supporting development of the BEAR robot. Additionally, we are grateful for congressional funding, and for internal funding from Vecna Technologies.

REFERENCES

1. N. Muscettola, P. Morris, and L. Tsamardinos, "Reformulating temporal plans for efficient execution," *Proc. of Sixth Int. Conf. On Principles of Knowledge Representation and Reasoning*, (1998).
2. A. Bemporad, M. Morari, V. Dua, and E.N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," Automatica, Vol. 38, no. 1, pp. 3-20 (2002).
3. A. Hofmann, S. Massaquoi, M. Popovic, and H. Herr, "A sliding controller for bipedal balancing using integrated movement of contact and non-contact limbs," Proc. Int. Conf. on Intelligent Robots and Systems, (2004).
4. O. Khatib, L. Sentis, J. Park, J. Warren, *International Journal of Humanoid Robotics*, 1(1):1-15 (March 2004)
5. A. Hofmann, "Robust Execution of Bipedal Walking Tasks from Biomechanical Principles", Ph.D. Thesis, MIT (2005)
6. A. Hofmann and B. Williams, "Exploiting spatial and temporal flexibility for plan execution of hybrid, under-actuated systems" *Proc. AAAI* (2006)
7. M. Hofbaur and B. Williams, "Hybrid Estimation of Complex Systems" *IEEE Transactions on Systems, Man and Cybernetics* – Part B: (2004).