Load Equalization on a Two-Armed Robot via Proprioceptive Sensing

Roxana Leontie, Evan Drumwright, Dylan A. Shell, and Rahul Simha

Abstract As humans we use our arms and bodies in addition to our hands to grasp objects. We (and robots) often cannot use caging or closure strategies when manipulating bulky objects. This paper studies manipulating such objects in the context of a particular task: equalizing a load across the arms of a two-armed robot. Our PR2 robot performs this task using only proprioceptive force sensing and a simple, reactive equalization strategy. We demonstrate the robot robustly performing this task using numerous and various objects (*e.g.*, boxes, pipes, broomsticks, backpacks).

1 Introduction

Recent research in robotic manipulation has delivered mobile robots with advanced autonomous capabilities (*e.g.*, [1, 2, 3]). The robot typically manipulates via a gripper or hand, though manipulation using the whole arm and even the body—thereby providing enhanced manipulation capabilities (*e.g.*, grasping when the hands are already loaded)—has been studied [4, 5, 6, 7, 8].

We have been investigating applications of arm and body manipulation for which *caging* and *closure* are difficult or impossible (as in the figure above), such as manipulating bulky objects, manipulating multiple objects simultaneously (*e.g.*, a stack of boxes, a bundle of ski equipment), or manipulating articulated objects (*e.g.*,

Roxana Leontie, Evan Drumwright, Rahul Simha
The George Washington University, Washington, DC, e-mail: {roxana|drum|simha}@gwu.edu

Dylan A. Shell
Texas A&M University, College Station, TX e-mail: dshell@cs.tamu.edu

a squirming child). In applications such as these, the manipulated object or objects are restrained from moving—sometimes only unilaterally—in some directions while being free to move in others.

2 Background and Related Work

We introduce the following subproblem of caging-free and closure-free arm/body manipulation: *given a two-armed robot manipulating an uneven load with its arms, equalize the load such that the arms maintain the load with nearly identical torques.* Here we define a *load* as one or more objects.

Research has established that a load is statically balanced if the projection of the load's center of mass (COM) lies within the polygon of support (defined for our problem as the contact manifold between the load and the robot's arms). It should be clear to the reader that load equalization trivially yields static balancing. We also note that load equalization bestows several benefits apart from stabilizing the load dynamically, including reducing energy consumption by the actuators (if such consumption is approximated using the integral of squared torque, as customary), avoiding lateral postural shifts in humanoids [9], and increasing dynamic stability of the robot [10].

Similar in spirit to our problem of bimanual load equalization is the "devil sticking" task investigated by Schaal and Atkeson [11, 12, 13]. That research also sought to achieve dynamic stability of the object (the "baton"), though by using dynamic movements of the robot rather than precise manipulation. We view both approaches as effective strategies for manipulation without caging or closure; our approach (the latter) is more applicable to larger loads and more powerful on slower robots.

Our work is the first (of which we are aware) to use proprioceptive force sensing in a Robotics application, though our approach—which uses PID control state to sense load on our robot's arms—is identical to that described by Wolpert and Kawato [14] (they also use the error between a motor controller's prediction and the robot's state).

3 Technical Approach

Our experimental platform is the Willow Garage PR2. The dynamics equations for this multibody system would nominally take the form:

$$\mathbf{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \mathbf{C}(\boldsymbol{q},\dot{\boldsymbol{q}})\dot{\boldsymbol{q}} + \boldsymbol{G}(\boldsymbol{q}) = \boldsymbol{\tau} + \mathbf{J}(\boldsymbol{q})^{\mathsf{T}}\boldsymbol{f}_{load} \tag{1}$$

where $\boldsymbol{q}$ are the generalized coordinates of the system, $\mathbf{M}(\boldsymbol{q})$ is the generalized inertia matrix, $\mathbf{C}(\boldsymbol{q},\dot{\boldsymbol{q}})$ is the matrix of fictitious forces, $\boldsymbol{G}(\boldsymbol{q})$ is the vector of gravitational forces, $\boldsymbol{\tau}$ is torque applied by the arms to counteract the load, $\mathbf{J}(\boldsymbol{q})$ is a

Jacobian matrix and $\boldsymbol{f}_{load}$ is a vector of spatial forces (corresponding to the load) acting on one or more of the robot's links. The PR2 employs counter-balancing springs (as a safety mechanism) that nearly eliminate gravitational acceleration of the arms; thus, $\boldsymbol{G}(\boldsymbol{q})$ is approximately zero. Assuming that the robot moves sufficiently slowly (*i.e.*, $\dot{\boldsymbol{q}}, \ddot{\boldsymbol{q}} \approxeq \boldsymbol{0}$) the torques applied by the actuators effectively equal the generalized forces from the load; thus, we can study load equalization with only the PR2's PID control state, albeit at some loss of modeling accuracy (a truly accurate dynamics model of the PR2 does not currently exist, to our knowledge). Thus, we use only proprioceptive sensing (*i.e.*, no vision, LIDAR, *etc.*) to equalize the load.

We formulated our fully reactive (*i.e.*, no planning or intensive computation is employed) equalization strategies from anecdotal experience. For example, humans may move their arms outward to redistribute a load; we tried this simple strategy first. Unfortunately, that strategy failed to account for frictional forces acting between the object and the robot's arms: Coulomb's friction model predicts greater frictional forces will be applied to the arm under greater load. Accordingly, the object (undesirably) moved along with that arm in our experiments, thus failing to shift the COM as desired. We developed two new strategies to address this difficulty: the *Tilting* strategy (Section 3.1) and the *Lock-Move-Release* strategy (Section 3.2).

3.1 Tilting *strategy*

The *Tilting* strategy, depicted in Figure 1, tilts the object by raising one arm and lowering the other until sufficient load has been shifted such that the object will slide along the arm to be moved outward. This strategy is predicated on the object possessing highly similar frictional properties at the points of contact with the robot's arms. Additionally, frictional forces must be sufficiently large such that the object does not slide off the robot's arms when tilted.

The strategy begins by tilting the arms until the imbalance is swapped: if we initially detect x units of force on arm A and y units of force on arm B, where $x - y = \Delta > 0$, we move arm A upward and B downward until the force on B becomes $\varepsilon > 0$ units greater than the force on A. Arm A is then moved outward until Δ load has been transferred, though the object may need to be tilted again during this process to keep the load on B greater than on A. We note that this strategy is sensitive to the frictional properties of manipulated objects: the strategy is not only predicated on nearly equal frictional characteristics at the contact points, but also depends on sufficient frictional effects to keep the manipulated object from sliding off the robot's arms when tilted.

We found this strategy worked poorly. Manipulated objects would slide off of the arms before sufficient weight had shifted; alternatively, joint limits would frequently prevent the object from being tilted enough for sufficient load redistribution. An successful series of steps showing the robot performing this strategy can be seen in Figure 2.

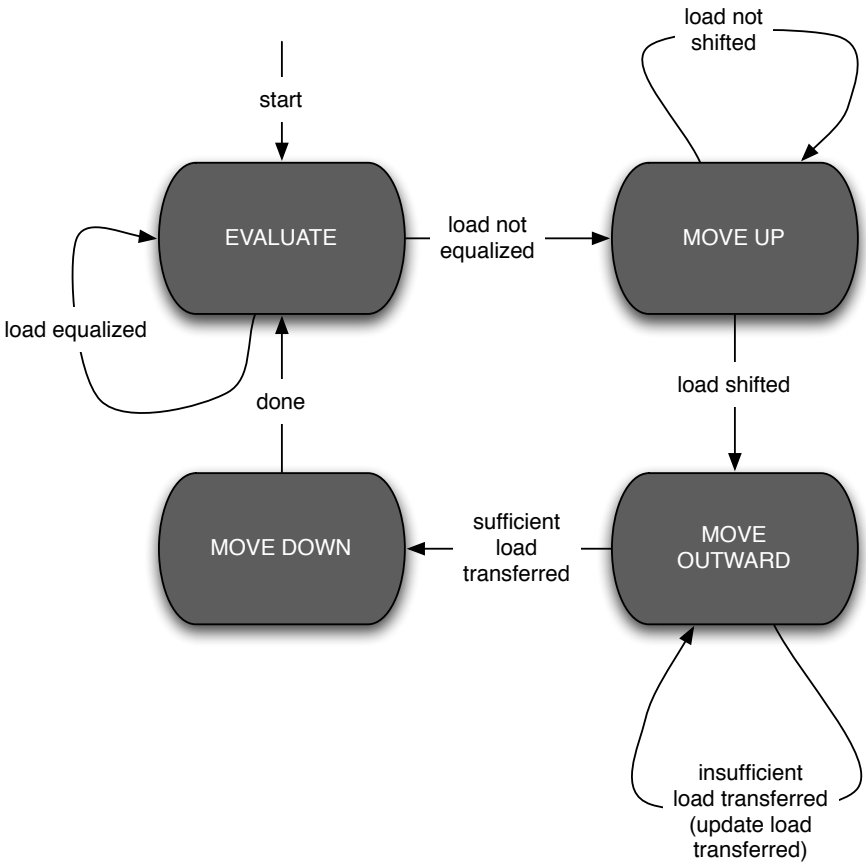

Fig. 1 State transition diagram for the *Tilting* strategy

Our failure to get this strategy working reliably led to the creation of the second strategy, described below.

3.2 Lock-Move-Release *strategy*

In the *Lock-Move-Release* strategy, the robot first grasps the object between its upper arm and forearm (*Lock*). The other arm then moves outward (*Move*), and finally the object is released from the grasp (*Release*). This process repeats until the load is equalized.

In the **EVALUATE** state, the robot analyzes the proprioceptive feedback from the arms (obtained from their PID controllers) to evaluate whether a load (if any) is equalized. To account for settling time (*i.e.*, the time it takes for the system to

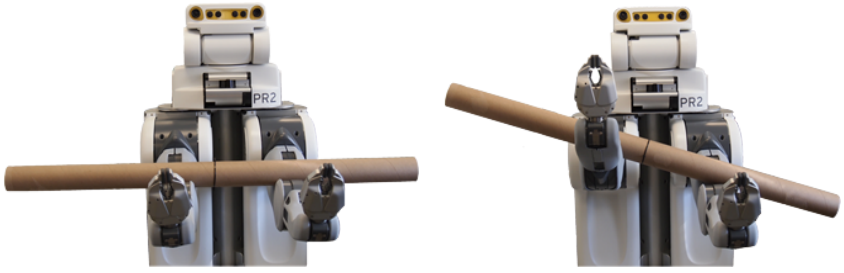

(a) evaluate (b) lifting

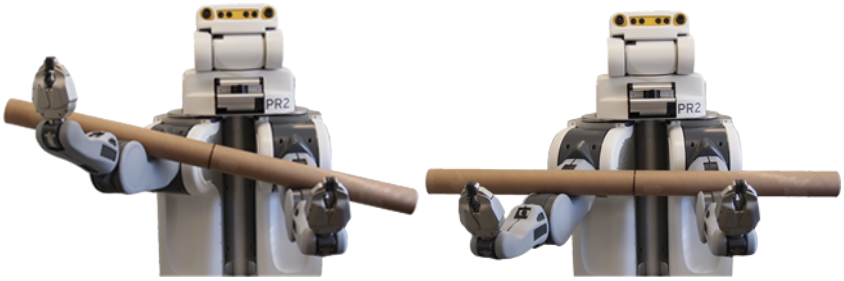

(c) moving outward (d) moving downward

Fig. 2 A series of photos depicting the robot performing the *Tilting* strategy

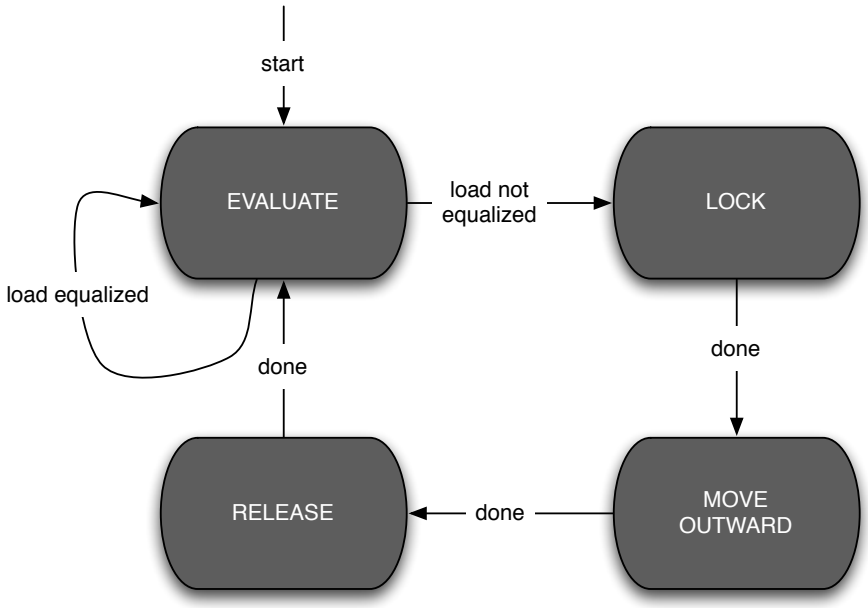

Fig. 3 State transition diagram for the *Lock-Move-Release* strategy

converge to a steady state), we apply a simple filtering strategy: the controller cannot leave the **EVALUATE** state until the sensed change in load on the two arms is smaller than 0.1 units[1] for at least 100ms. If—after the robot is determined to be in a steady state—the sensed difference in load on the two arms is 1.0 units or greater, the controller transitions to the **LOCK** state.

In the **LOCK** state, the robot immobilizes the object on the arm with lesser load by grasping it between the upper arm and the forearm: we assume that the object is placed approximately on the elbow joints. By restraining the object in this manner, we can move the other arm outward without the object also moving. The grasping movement takes place over two seconds; although the movement can generally be performed more quickly, we find that faster movements often reduce the accuracy of the proprioceptive feedback (due to the need for the system to settle).

The **LOCK** state transitions immediately to the **MOVE OUTWARD** state. This state predetermines the distance to move the arm outward toward minimizing repeated adjustments (*i.e.*, repetitions of holding, moving, and releasing). We compute this distance using several steps. First we employ forward kinematics to estimate the location of points of contact between the arms and object. Next we compute the object's initial COM (*i.e.*, its COM before immobilization) as depicted in Figure 4. Given the contact points $\boldsymbol{s}$ and $\boldsymbol{p}$ and the loads $\boldsymbol{x}$ and $\boldsymbol{y}$ on the arms in the **EVALUATE** state, we determine the position of the COM as follows. Noting that the COM is at $\boldsymbol{p}$ if $\boldsymbol{x} = \mathbf{0}$ and the COM is at $\boldsymbol{s}$ if $\boldsymbol{y} = \mathbf{0}$, then the COM is at $\dfrac{\boldsymbol{s}+\boldsymbol{p}}{2}$ if $\boldsymbol{x} = \boldsymbol{y}$. By extension, for arbitrary $\boldsymbol{x}$ and $\boldsymbol{y}$, the COM is located at:

$$\bar{\boldsymbol{x}} = \frac{\boldsymbol{x}s + \boldsymbol{y}p}{\boldsymbol{x}+\boldsymbol{y}} \tag{2}$$

Using the current location of the object's COM, our final step employs inverse kinematics to estimate the distance to move the arm outward to equalize the load (this calculation will assume rigidity of the object, but our experience indicates that this assumption works well in practice even for deformable objects) using the distances between the contact points and COM (d_1 and d_2 in Figure 5). The distance to move outward is calculated as:

$$\boldsymbol{d} = |\boldsymbol{d}_2 - \boldsymbol{d}_1| \tag{3}$$

The robot then executes the outward movement at a constant speed (*i.e.*, dependent on distance moved) to avoid excessive settling time, as described above.

Finally, the robot releases the object from its grasp in the **RELEASE** state, and the forearm returns to its original position. This slow movement occurs over 3.5 seconds, again to avoid excessive settling time. An successful series of steps showing the robot performing this strategy can be seen in Figure 6.

[1] According to Willow Garage, the force units are not calibrated to SI or any other measurement system.

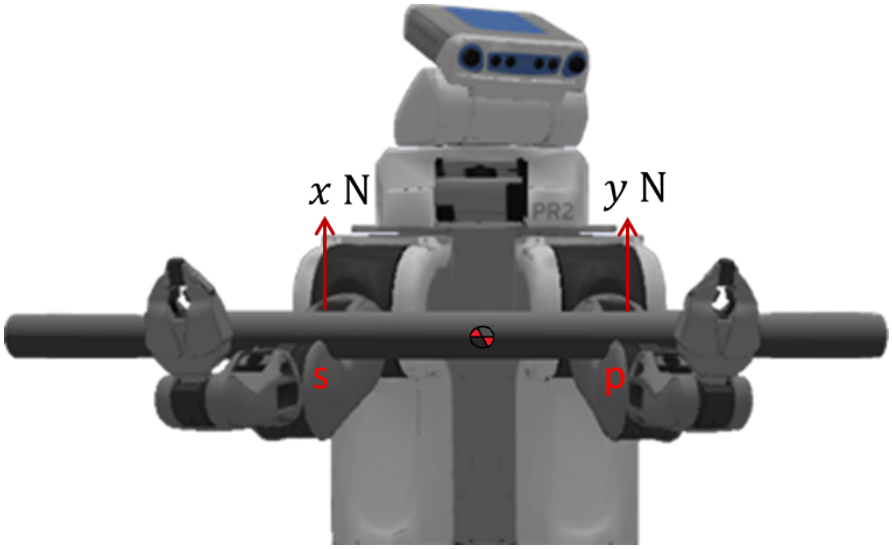

Fig. 4 Knowing (or estimating) the points of contact points ($\boldsymbol{s}, \boldsymbol{p}$) and forces applied to the arms ($\boldsymbol{x}$N and $\boldsymbol{y}$N) allows us to estimate the center-of-mass of manipulated objects (Equation 2).

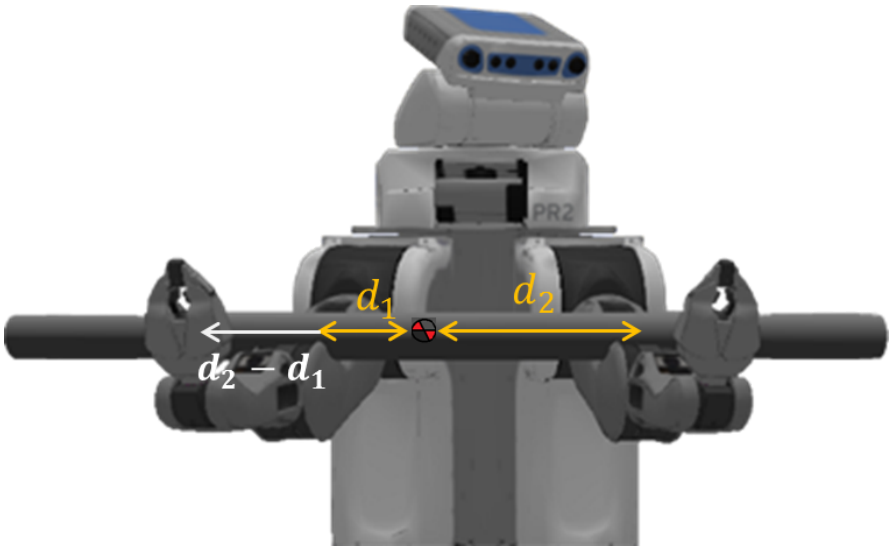

Fig. 5 Computing the distances between the contact points and the estimated center-of-mass allows us to determine how far outward we should move the arm with greater applied load (Equation 3).

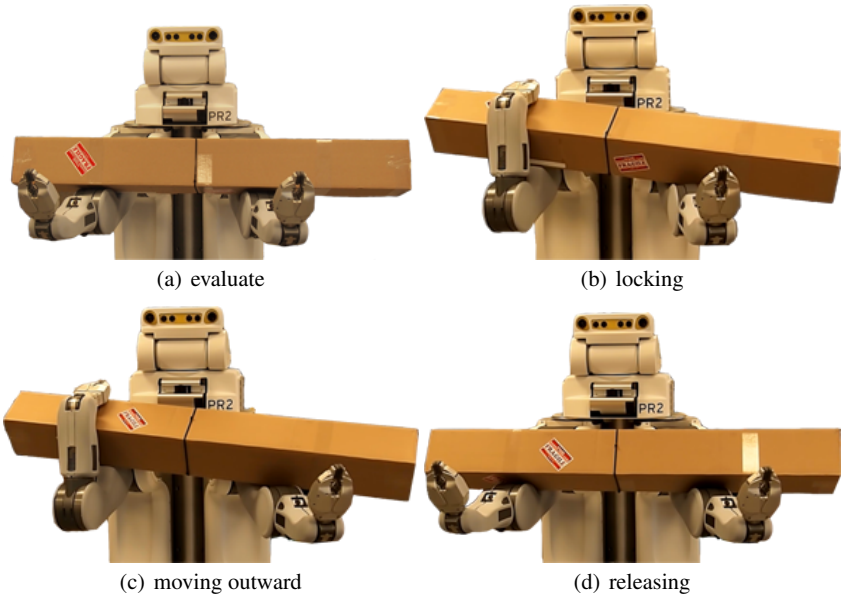

(a) evaluate (b) locking

(c) moving outward (d) releasing

Fig. 6 A series of photos depicting the robot performing *Lock-Move-Release* strategy

4 Experiments and Results

This section discusses a number of experiments conducted using the *Lock-Move-Release* strategy.

4.1 Balancing a load with COM projection within the support polygon

One of our first experiments with the *Lock-Move-Release* strategy tested whether the robot could equalize loads with centers-of-mass located non-equidistantly between the robot's arms yet with projections remaining within the support polygons. The transparent photo of the robot in Figure 7 shows the manipulated object used in the experiments in Sections 4.1–4.3: a 2.7m, 1.4kg iron pipe. That figure also depicts the object's geometric center (marked in pink near the robot's left arm); the geometric center is located near the pipe's COM. The data plotted in Figure 7 correspond to immobilizing the object with the robot's right arm (0–2.5) seconds), shifting the load (2.5–6 seconds), and releasing the object (6–9.5 seconds); the plot shows that the actuator force decreases when the object is grasped and increases on release. We have yet to observe a failure using the *Lock-Move-Release* strategy over numerous experimental trials with variable initial condition (location of the pipe's geometric center within the support polygon).

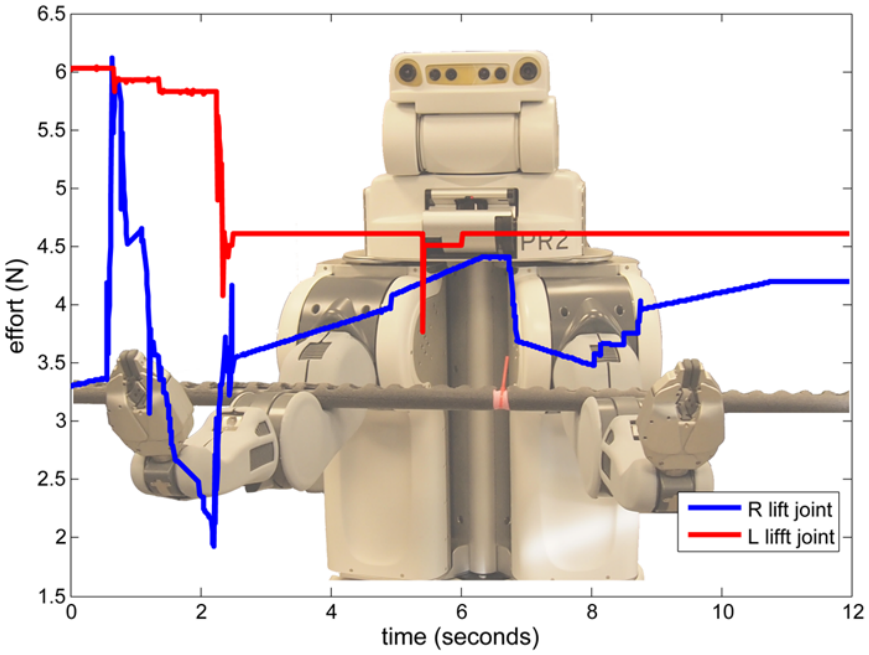

Fig. 7 Plot of the "shoulder lift" motor torques during the load equalization process corresponding to a trial of the first experiment: as seen in the photo, the geometric center of the pipe (marked in pink near the robot's left arm) lies within the object's support polygon.

4.2 Equalizing time varying loads

A second series of experiments tested the ability of our strategy to handle time varying loads (as if objects were continually added to the robot). The first ten seconds of the plots in Figure 8 correspond to the robot equalizing loads from only the iron pipe (identically to the conditions described in the previous section). After each load was equalized, we added an additional load (0.6–0.85kg) to the pipe at a point with projection near or at the inside edge of the support polygon (depicted by the photos in Figure 8). The plots at approximately fifteen seconds correspond to the robot beginning to re-equalize the loads. The re-equalization process requires fewer than ten seconds to complete.

4.3 Equalizing time varying loads added outside the support polygon

A third series of experiments (depicted in Figure 9) was designed to stress the ability of our strategy to handle loads near the margin of dynamic instability. The experimental setup for these trials was identical to those described above but with the

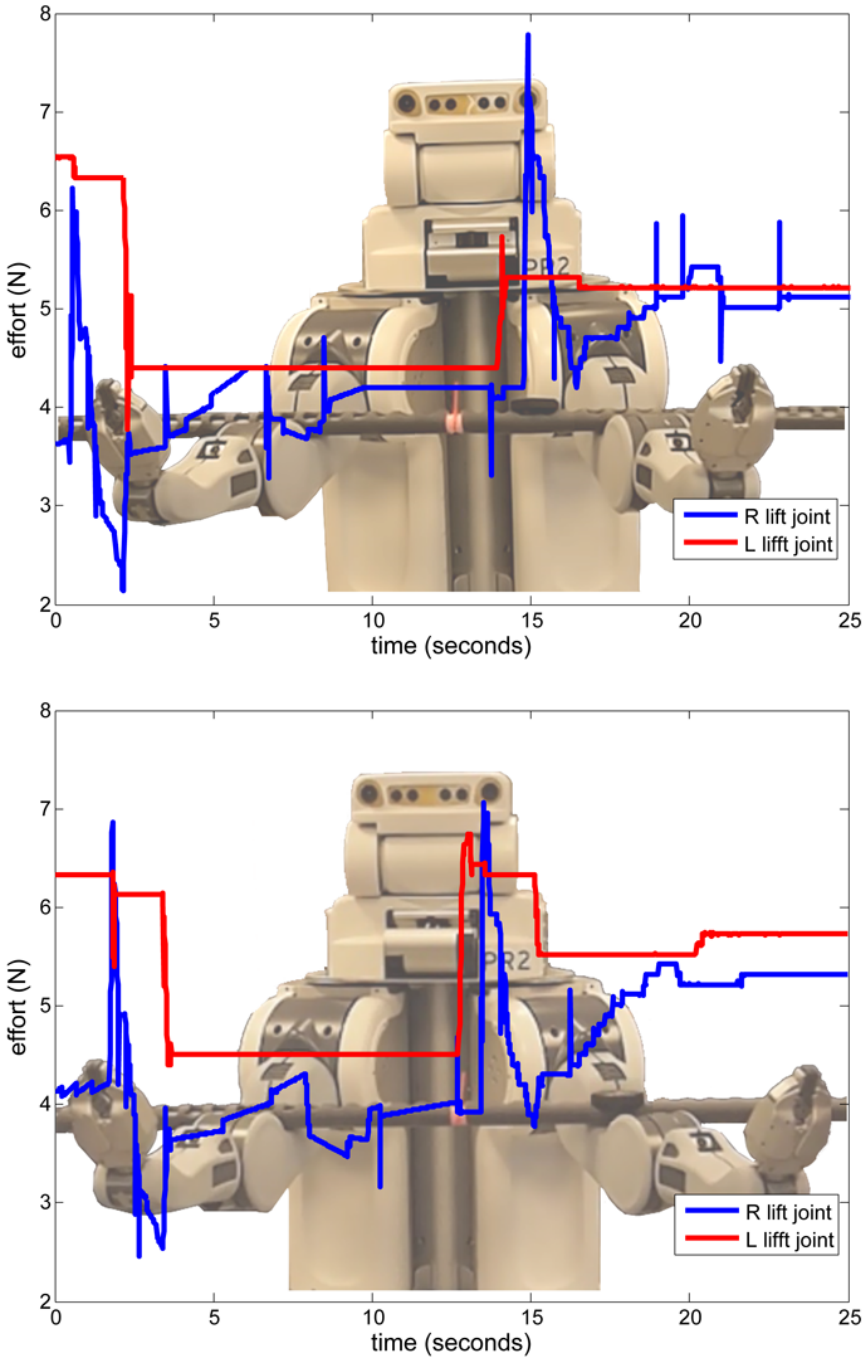

Fig. 8 Plots of the "shoulder lift" motor torques during the load equalization process corresponding to two trials of the second experiment: as shown in the photos, magnetic weights have been added to the iron pipe near (top) and directly above (bottom) the robot's left elbow to effect a change to the load.

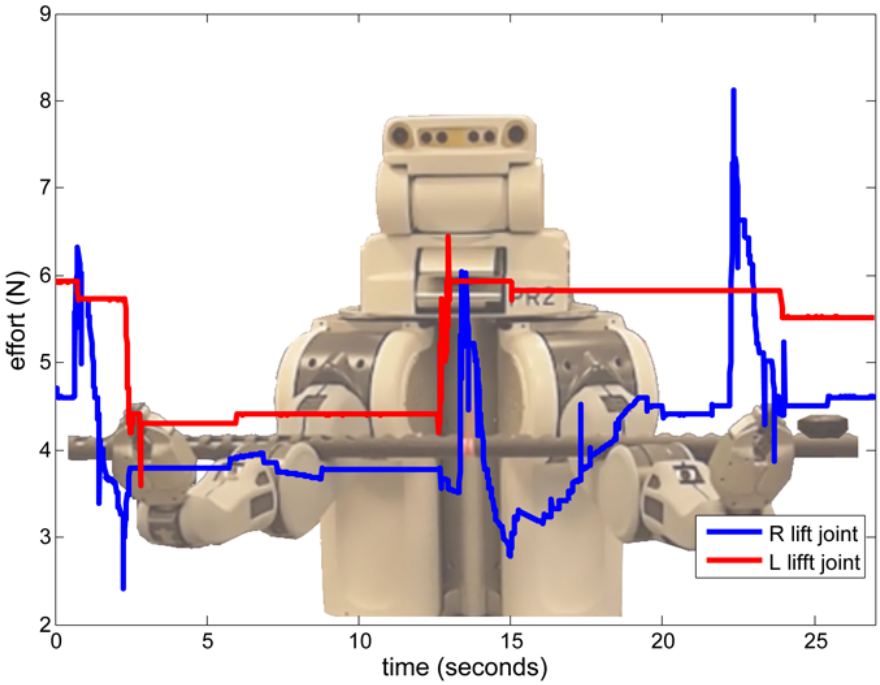

Fig. 9 Plot of the "shoulder lift" motor torques during the load equalization process corresponding to a trial of the third series of experiments: as seen in the transparent photo (near the right elbow), a magnetic weight has been added well outside the support polygon of the iron pipe to simulate a modification to an equalized load.

additional load (a 0.6kg weight) placed well outside of the support polygon. We tested various positions of the weight—including at the end of the 2.7m iron bar—and still observed complete success of our strategy. The equalization process for the added load generally required two iterations (this was the case with the trial depicted in Figure 9): the first iteration (13–20 seconds) nearly equalized the load, while the second iteration yielded full equalization. As with the plots in Figure 8, the plot in Figure 9 shows that both the detection and equalization of the added load was conducted quickly.

4.4 Equalizing a load from a box

We also experimented with the *Lock-Move-Release* strategy on objects of different shapes, such as boxes (one of which is depicted in Figure 10). For the data plotted in Figure 10, the box's position shifted during the lock and release phases of the strategy, thereby introducing significant noise into the process (as evident in the plot). Nevertheless, no failures were encountered; Figure 10 corresponds to an initially

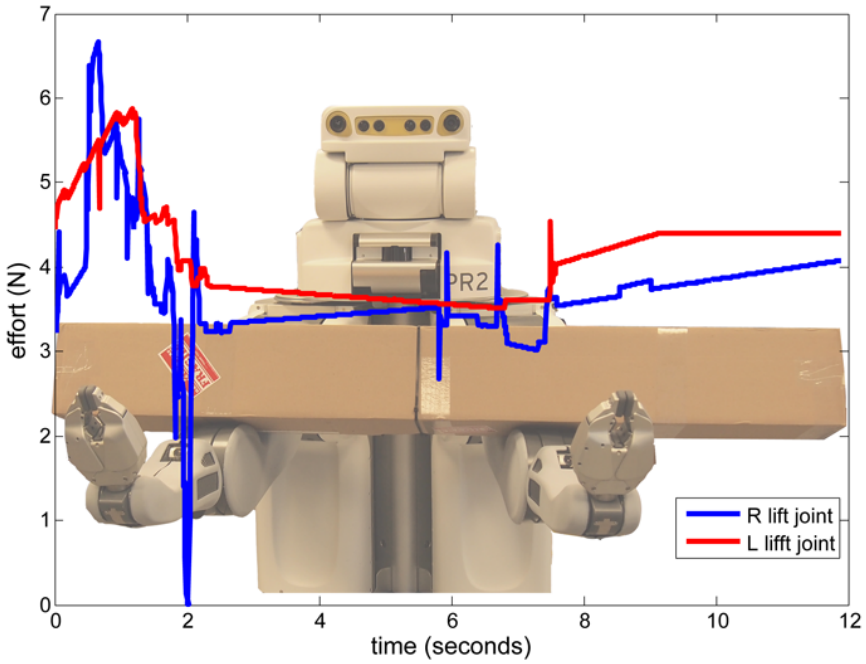

Fig. 10 Plot of the "shoulder lift" motor torques during the load equalization process corresponding to the robot equalizing a load from a cardboard box.

unequal load from a box that is equalized in fewer than ten seconds using only a single iteration of our strategy.

4.5 Load equalization with non-rigid objects

We conducted a final series of experiments on bulky non-rigid objects (camping backpacks with frames removed; see Figure 11) with non-uniformly distributed loads. These experiments entailed having the robot equalize one backpack placed on the robot's arms followed by a second backpack placed on top of the first one. Again, no failures were encountered, though the equalization process required three iterations.

4.6 Miscellaneous testing

We also performed a number of tests of our strategy using very light objects (0.25 kg)—though one would expect that load equalization would not be in great need for

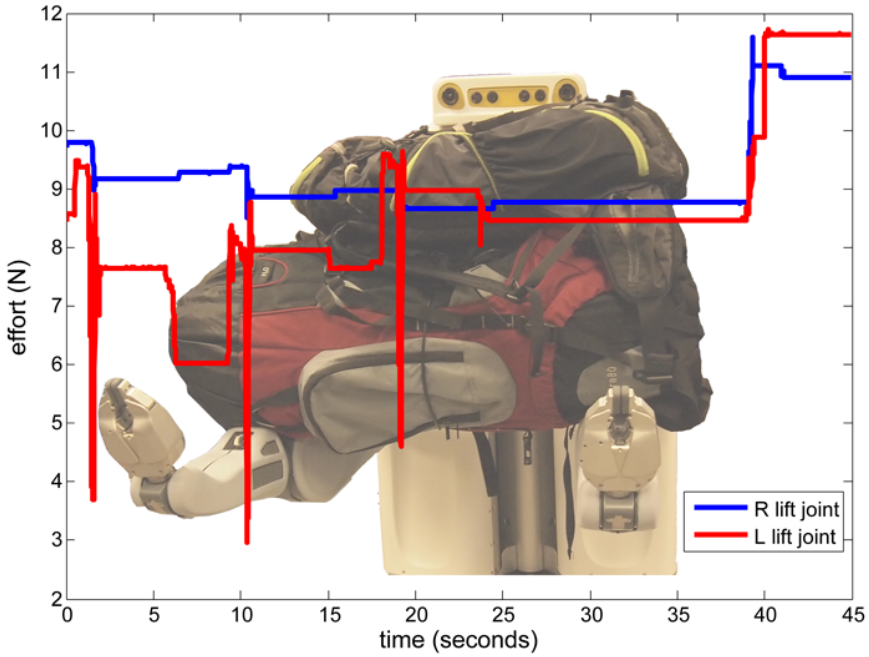

Fig. 11 Plot of the "shoulder lift" motor torques during the load equalization process conducted on non-rigid bodies (camping backpacks with frames removed and non-uniformly distributed loads).

such objects on current robots with bimanual manipulation capabilities. Again, we never observed a failure of the strategy.

5 Main Experimental Insights

This section discusses distinct insights gleaned from the development of the *Lock-Move-Release* strategy and the performance of the robotic platform during the experiments with respect to both primary (*i.e.*, load equalization) and ancillary (*i.e.*, grasping) tasks.

5.1 Ability to equalize loads

The PR2 can very robustly equalize loads applied by various objects using a simple reactive strategy (this is likely explains much of its robustness), minimal sensory data, few assumptions on object geometry or mass distribution, and little modeling. More sophisticated strategies might yield faster equalization, though we sense that we are near the PR2's performance envelope:dynamically equalizing shifting loads

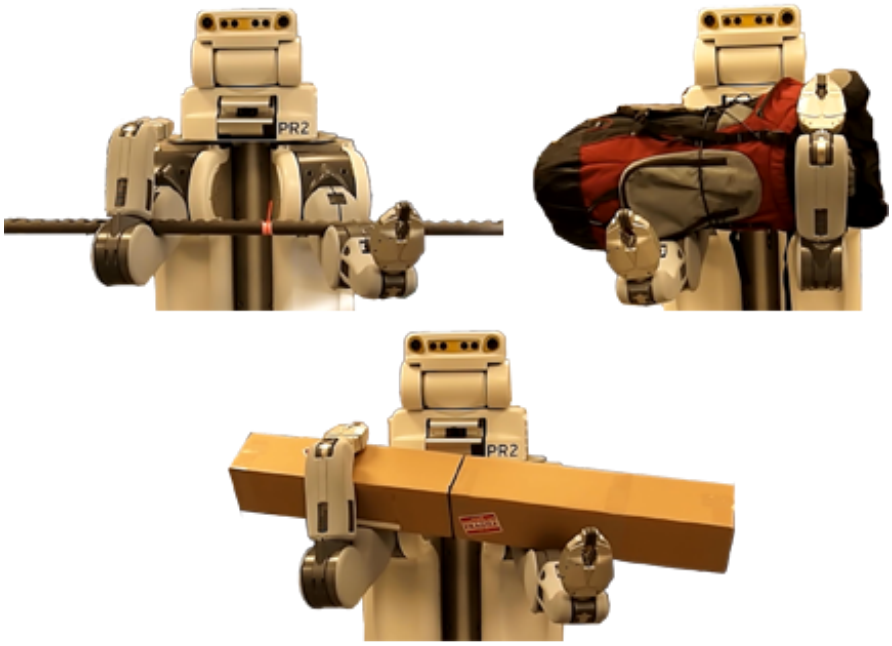

Fig. 12 Sample grasps of various objects manipulated by the robot's arms during the course of our experiments.

does not seem possible with this robot. Due to the minimal assumptions used in our work and the simplicity of our strategy, we predict that our approach will transfer to other manipulation platforms without difficulty.

5.2 Insights from use of purely proprioceptive sensory data

The PR2 is slightly better able to detect uneven loading for heavier objects than for lighter ones; it can be challenging to distinguish an uneven load from noise (hence our unit threshold for triggering the equalization strategy). This finding corresponds with an informal experiment conducted by the authors: each of us was able to accurately compare masses for only relatively heavy objects (greater than several kilograms). Thus, while proprioceptive sensory data is likely sufficient for heavier objects, lighter objects may require the addition of visual or tactile sensory data.

5.3 Insights from arm grasping

The PR2 seems quite capable of grasping objects placed onto various points on its arms. We tested assorted objects, including planks, broomsticks, pipes, boxes, and backpacks (among others), during the course of our experiments. The PR2 never failed to grasp an object with its arms and seemed fairly insensitive to the positioning of the objects. We would be curious to see a study gauging effectiveness of grasping with the arms against grasping with the grippers (at least, but not limited to, the PR2). Sample arm grasps realized during performances of our strategy are depicted in Figure 12.

References

1. W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote, J. Hsu, R. B. Rusu, B. Marthi, G. Bradski, K. Konolige, B. Gerkey, and E. Berger, "Autonomous door opening and plugging in with a personal robot," in *Proc. IEEE Intl. Conf. Robotics and Automation (ICRA)*, Anchorage, 2010.
2. J. van den Berg, S. Miller, K. Goldberg, and P. Abbeel, "Gravity-based robotic cloth folding," in *Springer Tracts in Advanced Robotics: Algorithmic Foundations of Robotics IX: Selected Contributions of the Ninth International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2010.
3. T. L. Chen and C. C. Kemp, "A direct physical interface for navigation and positioning of a robotic nursing assistant," *Advanced Robotics*, vol. 25, pp. 605–627, 2011.
4. K. Salisbury, "Whole arm manipulation," in *Proc. Intl. Symposium Robotics Research (ISRR)*, 1987, pp. 183–189.
5. S. Yeap and J. C. Trinkle, "Dynamic whole-arm dextrous manipulation in the plane," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 1995.
6. P. Song, M. Yashima, and V. Kumar, "Dynamics and control of whole arm grasps," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Seoul, 2001, pp. 2229–2234.
7. R. P. Jr., A. H. Fagg, and R. A. Grupen, "Extending fingertip grasping to whole body grasping," in *Proc. of IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2003, pp. 2677–2682.
8. K. Hsiao and T. Lozano-Perez, "Imitation learning of whole-body grasps," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Beijing, 2006, pp. 5657–5662.
9. J. Watson, R. Payne, A. Chamberlain, R. Jones, and W. I. Sellers, "The kinematics of load carrying in humans and great apes: Implications for the evolution of human bipedalism," *Folia Primatologica*, vol. 80, no. 5, pp. 309–328, Nov 2009.
10. J. Liu, T. E. Lockhart, and K. Granata, "Effect of load carrying on local dynamic stability," in *Proc. Human Factors and Ergonomic Society Annual Meeting*, vol. 51, no. 15, Oct 2007, pp. 909–913.
11. S. Schaal and C. G. Atkeson, "Open loop stable control strategies for robot juggling," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 3, Atlanta, 1993, pp. 913–918.
12. ——, "Robot juggling: An implementation of memory-based learning," *Control Systems Magazine*, vol. 14, no. 1, pp. 15–71, 1994.
13. ——, "Assessing the quality of local linear models," in *Proc. of Advances in Neural Information Processing Systems*, J. D. Cowan, G. Tesauro, and J. Alspector, Eds., 1994, pp. 160–167.
14. D. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Networks*, vol. 11, pp. 1317–1329, 1998.