A RISC Approach to Robotics

John Canny and Ken Goldberg

Young Rossum invented a worker with the minimum amount of requirements. He had to simplify him. He rejected everything that did not contribute directly to the progress of work...
- Karel Capek, R.U.R.

This article describes a framework that combines simple hardware traditionally used in manufacturing with sensor-based planning and design algorithms from robotics. For repetitive assembly, we argue that this combination can reduce start-up and maintenance costs, increase throughput, and greatly reduce the set-up and changeover times for new products.

Consider the "pick-and-place" operation which is the building block of automated assembly: the part must be picked up off a conveyor or pallet, moved to its destination, and inserted into an assembly. This can of course be accomplished with a 6 dof robot. At any given point, only a small subset of these degrees of freedom are required, yet we continue to pay the overhead for this flexibility in terms of settling time and precision. Similarly, a general-purpose vision system might be used to sense the position of the part. Its full power to provide a rich description of the image at video rates is under-utilized since we only require the pose of a known part at fixed points in the sequence.

An alternative would be to use two or more simple grippers, one for initial grasp, one for final placement etc., and linear pneumatic slides for gross motion. An RCC collar could be used on the insertion gripper for a compliant insert. Binary light beams can be used to measure part pose at the initial and final stages. Furthermore, the stages in this system can be pipelined so that parts at one end are being aligned while parts at the other are being inserted. Thus with a comparable number of degrees of freedom, we get several times the throughput, and all degrees of freedom are working almost all the time. A good testbed for modular workcell design is the RobotWorld system illustrated in Figure 1, which is being used for research at UC Berkeley.

The proposed hardware bears a close resemblance to existing "hard" automation; what is new is the application of computational methods for robust design and control of these systems, and more extensive use of (simple) sensors. Clearly this enhances the capabilities of the hardware. A less-obvious benefit is that software capability is also enhanced—algorithms for fine-motion, grasp planning and some sensing algorithms which would be intractable on a general-purpose robot work in real-time when applied to simple hardware.

To describe this approach we chose the acronym RISC—Reduced Intricacy in Sensing and Control—by analogy with computer architecture. Analogously, we propose to use simple hardware elements that are coordinated by software to perform complex tasks. As a research agenda, these goals encompass computational geometry, mechanics, design, dynamic simulation, and computer vision. Its problems are both more structured and less constrained than traditional robotics. In a robotic planning problem, the robot kinematics and the camera or rangefinder configuration are part of the problem definition. In RISC, one seeks to minimize the complexity of the actuators and sensors, so their configuration is part of the solution of workcell design. A key to doing this is to develop a vocabulary of modular sensing and actuation units, and the design and planning algorithms to support them.

This article is intended as an overview to promote discussion and to incite others to pinpoint research topics in robotics of most likely impact to practitioners. A more detailed description of existing results, related work, and open research problems can be found in [1]. RISC sensing algorithms for binary optical beam arrays are described in [6]. These sensors are extremely well-suited to manufacturing. They are inexpensive, and provide position accuracies of 25 micrometers (μm) and object recognition times of a few milliseconds. Examples of RISC manipulation can be found in [2,5], which describe algorithms for feeding and recognizing industrial parts using a modified parallel-jaw gripper. A sensor-based RISC assembly strategy is described in [4]. The scheme described there achieves peg-in-hole insertions at 25 μm clearance without chamfering at 99% repeatability, without prior calibration.

A RESEARCH FRAMEWORK FOR INDUSTRIAL AUTOMATION

In this section we list several guidelines that we are finding useful in directing our research.

Industrial Assembly is Repetitive.

John Canny is with the University of California, Berkeley, Berkeley CA 94720, fc@cs.berkeley.edu and Ken Goldberg is with the University of Southern California, Los Angeles CA 90089-0273, goldberg@ usc.edu

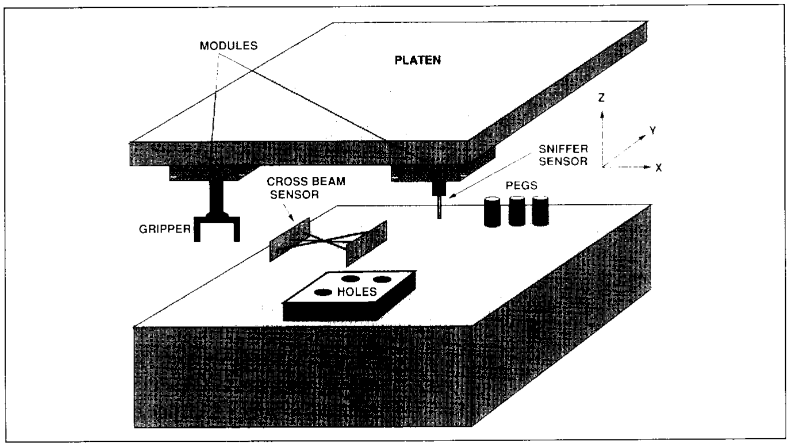

Figure 1. A multiple robot, vision-guided assembly system.

In contrast to unstructured environments where new objects are often encountered, factories are characterized by repetition. This requires that assembly systems be efficient but not necessarily flexible in the short run. When part geometry does change, the system must be reconfigured. Speeding this process is a fundamental research challenge. To develop a scientific basis for automatically reconfiguring such systems, we need to identify a basic vocabulary of modular elements, and the sensing and control algorithms to support them.

Decompose Complex Tasks into Small Steps.

Dextrous insertion operations can be simplified by simplifying the hardware. For example, a pick-reorient-insert sequence may employ a parallel-jaw gripper for the initial grasp and a second three-fingered gripper with RCC collar for the final insertion. In general, many benefits accrue from breaking a complex manipulation or sensing step into several, such that each step can be performed by a simple sensor or actuator. Hardware is simplified, degrees of freedom are more effectively used, and pipelining is possible.

Decompose Sensors and Actuators into "Units."

A sensor can be viewed as comprising some number of "units of sensing." For the beam sensors, a unit of sensing is most naturally a single beam. For other sensing technologies, a unit of sensing should correspond to a single real value provided by the sensor. So array sensors like cameras and tactile sensors comprise roughly 256k and several hundred units respectively. We propose this point of view for several reasons: (i) It makes explicit the amount of information that the sensor interpretation algorithms must process, in relation to the number needed for pose determination. (ii) It indicates the approximate initial and maintenance cost of the sensor. (iii) It supports a task-specific design of a sensor that provides enough information for the task at hand without overkill.

For similar reasons it makes sense to break actuators down into "units of actuation". These will normally correspond to the degrees of freedom of the actuators. This allows every device that causes or constrains part motion to be considered. Not just robot arms and grippers, but fixtures, conveyors, AGVs, and various types of feeders. All these devices affect the 3 to 6 degrees of freedom of a part.

The advantages of this point of view are: (i) It indicates the approximate complexity of controlling the actuator. (ii) It is a good guide to setup and maintenance cost. (iii) It allows a measure of the "efficiency" of the actuator, i.e., how many actuator degrees of freedom are used, and how many part degrees of freedom are constrained. (iv) It provides a uniform vocabulary for the trading off options when designing assembly cells. We note that this view resonates with the (often derided) Japanese view of robots as any device with degrees of freedom that affects part motion or shape.

Distribute Sensing throughout the Workcell.

The accuracy of a sensor is typically measured by the accuracy of the measurements it returns. But at least as important is the proximity of the sensor to the task. Good sensor values are useless without an actuator with comparable accuracy. If the sensor is kinematically close (e.g., attached to the last link) to the end effector, overall accuracy depends on the repeatability of the actuator, not its absolute accuracy. E.g. the configuration of Figure 2 allows a standard industrial SCARA robot to perform reliable unchamfered peg-in-hole insertions at 25 μm tolerance, even though its absolute position accuracy is about 2 mm. By using local sensing, we get optimum position accuracy from the actuators and bypass the need for accurate global calibration.

There are many other examples of local sensing, several of which are used in the UC Berkeley Workcell. Cross-beam sensors can be mounted on conveyor belts to determine the very uncertain pose of objects coming down the conveyor. Cross-beam and reflective sensors can be mounted on grippers. The cross-beam sensor allows the gripper to accurately center over a part to be grasped, assisting in part acquisition. The reflective sensor allows the end-effector to accurately locate a feature for an insertion step, as shown in Figure 2, assisting in part placement.

Avoiding global calibration is a great aid to rapid workcell editing. Local sensing makes calibration unnecessary in almost all cases. Because of the proximity of sensor and actuator, only accurate relative displacements are needed.

Integrate Hardware and Software into Soft-Hard Objects (SHOs).

Once one steps back from the view of robots as 6-axis universal positioners, a huge variety of possibilities opens up for forming novel liasons between actuators, fixtures and other types of passive elements to effect part pose. Matt Mason and his students at Carnegie-Mellon University have created a science of part pose control through sliding motion. It is difficult to say whether the tilting trays, sliding fences and barriers they use should be called robots, feeders or fixtures, but they clearly have aspects of all three.

Critical to fast workcell editing is rapid incorporation of support software when a piece of hardware is added to the cell. To facilitate this, we propose a twist on the usual object-oriented program-

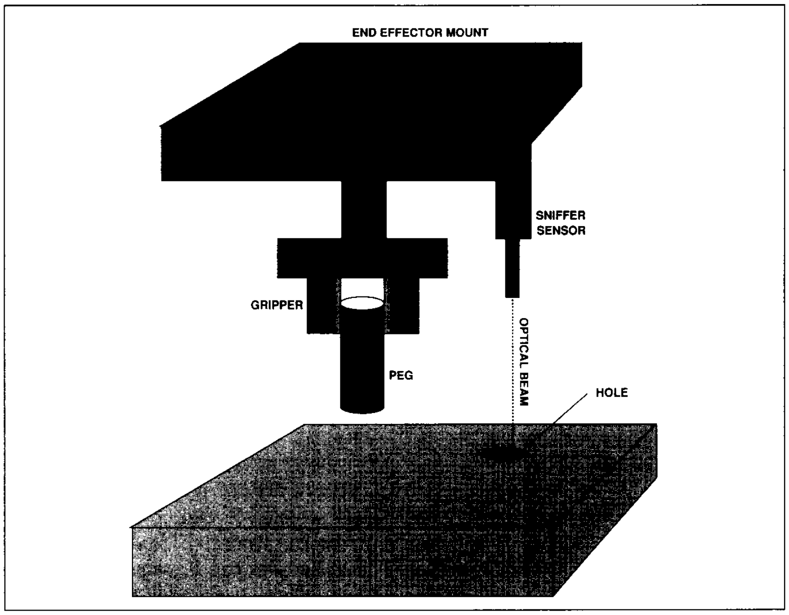

Figure 2. Optical sensor mounted directly on gripper platform accurately locates holes prior to peg insertion.

ming model. Each piece of physical hardware has associated with it a software object that presents a public interface of abstract operations, such as "sense pose" or "move to position", and which hides the details of how the hardware implements this operation.

We propose in addition that the software object instance contain specific geometric data for a "hardware instance" of given type, e.g. for a particular cross-beam sensor, this information would be beam angles and relative positions. Editing the physical workcell involves adding or removing object instances from the software workcell, in an isomorphic way.

For the future, we argue that it will be critical for vendors of manufacturing hardware such as actuators, sensors, grippers, fixtures and feeders to provide software objects to drive and present an abstract interface to their hardware, just as is done now for hardware accessories for personal computers. These routines would be most widely usable as C libraries.

Integrating Actuators, Feeders, Sensors into "Virtual Robots."

At a slightly higher level, since we have removed the traditional grouping of units of actuation and sensing, we must choose an alternative. In the RISC context, we are free to do this in a task dependent way, again using an object-oriented programming model. The result is "virtual robots," which perform more abstract tasks, such as "acquire part" or "peg-in-hole insert," and which use many soft-hard objects. [3]

DISCUSSION

One of the most intriguing aspects of RISC is that it blurs the distinction between planning and design. A configuration of RISC elements can be thought of as a compiled version of an assembly plan. For example, rather than planning motions of a multi-fingered hand to hold a part, we can configure a modular fixture to hold the part. In effect, the grasp plan is reduced to hardware.

Our initial motivation for considering simple hardware elements was to reduce the complexity of planning for general purpose robots. Simple elements also have the advantage of:

- Increased Reliability. RISC sensors and actuators have fewer components so less can go wrong.
- Lower Start-up and Maintenance Costs. Many of the hardware elements are available off-the-shelf, and are easily repaired or replaced.
- Increased Speed. Simple sensor data can be processed very fast. Simple actuators with decoupled dynamics can move very fast without losing accuracy.
- Rapid Reconfigurability, critical for future manufacturing systems. RISC sensors self-calibrate, and the modular design of RISC actuators and feeders supports easy "editing" of the workcell.

The beauty of RISC is that it suggests theoretical questions with short-term practical consequences. Related projects are being initiated at Stanford, Sandia Labs, Carnegie Mellon, New York University, University of Padua and with Adept Technology and Silma, Inc. We believe this work holds potential for significant scientific progress during the next five years.

ACKNOWLEDGEMENTS

We presented an early version of these ideas at the NSF Workshop on Geometric Uncertainty in Robot Motion Planning in the summer of 1992. We thank the participants for their input, in particular Brian Carlisle, Jean-Claude LaTombe, Randy Brost, Matt Mason, Mike Erdmann, Bud Mishra and Anil Rao. Also we thank Dan Whitney, Richard Wallace, Todd Rockoff, Pradeep Khosla, Damian Lyons, Peter Allen, Daniela Rus, and Howard Moraff for insightful feedback on earlier drafts of this report.

BIBLIOGRAPHY

[1.] John Canny and Ken Goldberg. " 'RISC' for industrial robotics: Recent results and open problems." Technical Report IRIS-93-315, USC Institute for Robotics and Intelligent Systems, September 1993. Available via anonymous ftp from 128.125.51.19.

[2.] Ken Goldberg. "Orienting polygonal parts without sensors." *Algorithmica*, 10(2):201-225, August 1993. Special Issue on Computational Robotics.

[3.] Jean-Claude Latombe. "Robot algorithms." In Takeo Kanade and Richard Paul, editors, *Robotics Research: The Sixth International Symposium*, 1993.

[4.] Eric Paulos and John Canny. "Informed peg-in-hole insertion using optical sensors." In SPIE Conference on Sensor Fusion VI, 1993. Boston Massachusetts.

[5.] Anil Rao and Ken Goldberg. "Shape from diameter: Recognizing polygonal parts with a parallel-jaw gripper." *International Journal of Robotics Research*, 13(1), February 1994 (to appear).

[6.] A. Wallack, J. Canny, and D. Manocha. "Object localization using crossbeam sensing." In IEEE Conference on Robotics and Automation, pages 692-699, 1993.