IEEE International Conference on Robotics and Automation, San Diego, CA. May 8-13, 1994:

"RISC " for Industrial Robotics: Recent Results and Open Problems

John F. Canny*
University of California, Berkeley
Berkeley, CA 94720
jfc@ernie.berkeley.edu

Kenneth Y. Goldberg†
University of Southern California
Los Angeles, CA 90089-0273
goldberg@usc.edu

Abstract

At the intersection of robotics, computational geometry, and manufacturing engineering, we have identified a collection of research problems with near-term industrial applications. The common thread is robot systems with Reduced Intricacy in Sensing and Control (RISC), such as light beam sensors and parallel-jaw grippers. We conjecture that such systems, coupled with appropriate algorithms, are capable of recognizing and orienting a broad class of industrial parts. When compared with general-purpose robots, the resulting systems could be: (1) lower in cost, (2) more reliable and (3) easier to reconfigure.

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. By focusing on a small vocabulary of simple hardware, planning become computationally tractable and we can in some cases make guarantees about the existence of solutions.

We borrowed the RISC acronym from computer architecture to acknowledge a common theme: identifying a minimal set of hardware primitives and matching these primitives with highly efficient software. In this paper we review recent algorithms for locating, feeding, inserting and fixturing industrial parts. We discuss related work and propose a set of open problems for future research.

1 Introduction

This paper is concerned with industrial assembly – the layout and programming of workcells – rather than unstructured environments. We start with three assumptions: (1) During each assembly run, all parts have fixed geometry, so models may be used, (2) repetitive motion will occur, and (3) the device being assembled and its components will be redesigned periodically. Thus workcells must be reliable and reconfigurable. Our objectives are: (1) to use the smallest possible set of hardware elements and (2) to develop efficient algorithms to control and when necessary reconfigure these elements based on part geometry. For repetitive operation, such systems can provide an alternative to general-purpose robots. A more complete description of the RISC paradigm is given in [6]. That paper provides a systematic decomposition of hardware into "units" of sensing and control, and then explains how to group the units into functional modules in a task-specific way. Rather than complex manipulators and sensors, RISC favors *instrumented actuators* that incorporate a little of both. Both harware and software are modular, and self-calibration algorithms allow new hardware to be inserted in a workcell and brought up in a matter of minutes. In this paper we focus on recent results and give some concrete open problems.

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. In [17], this was done with a 6dof robot and a parallel-jaw gripper. Sometimes a regrasp is required to insure clearance between the gripper and the assembly. At any given point, only a small subset of the robot's 6 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 is always available although we only require the pose of a known part at fixed points in the sequence.

An alternative would be to use two or more 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. Each motion in this alternative system can be precisely controlled by mechanical alignment of axes and coordinated with a PLC. 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. The latter system describes many existing "hard automation" assembly lines [24].

Researchers have made enormous progress in automatic planning for general-purpose robots, while hard automation continues to rely on manual retooling. In Japan, both systems would be considered robots. Our idea is to apply insights from robotics and computational geometry to the

*Supported by a David and Lucile Packard Foundation Fellowship and by NSF Presidential Young Investigator Grant IRI-8958577

†Supported by the National Science Foundation under Awards IRI-9123747 and DDM-9215362 (Strategic Manufacturing Initiative), and by equipment grants from Adept Technology and Qu-Co.

second type of system. The challenge is to develop software for rapidly reconfiguring a set of simple hardware modules for each new application. By viewing design and layout as dual to sensor-based planning, the result would combine the advantages of general-purpose robots with those of hard automation.

To characterize this hybrid of robotics and automation, we use the term “RISC ” – Reduced Intricacy in Sensing and Control – to refer to systems with a well-defined set of simple hardware elements that can be reconfigured to handle a variety of part geometries.

In this paper we describe results for a variety of subproblems such as feeding, fixturing, inserting, and inspecting parts. In each case part geometry is given as input:

- Locating and Inspecting Parts with Sparse Optical Beam Sensors: What is the minimal arrangement of such sensors that will discriminate between a known set of parts?
- Feeding and Sorting a stream of polygonal parts using a set of pneumatically-driven parallel-jaw grippers: At what angles should the grippers be arranged to insure that parts can be distinguished and emerge in a unique orientation?
- High Precision Part Insertion: Given a fixed and a movable beam sensor which are uncalibrated, and a hole with uncertain center and radius, how can a robot perform a precise peg-in-hole insertion in a few seconds?
- Fixturing with Modular Components: Given a square lattice of anchor sites, where should locators and clamps be arranged to hold a part in form closure?

After discussing solutions to each of these problems, we describe related work and pose a number of open problems for future research.

2 Locating and Inspecting Parts with Beam Sensors

We have studied two types of optical beam sensor. The first is called a “cross-beam sensor”, see figure 1. When a part passes through the apparatus, the cross beams perceive a horizontal cross-section of the part. The times when the beams are broken and unbroken are recorded, as shown in figure 2.

With 3 beams, the breakpoints define a hexagon bounding the part cross-section (6 real values). In spite of the coarseness of this information, because the beam measurements are so precise (~25 microns), the pose can almost always be determined unambiguously. Because the measurements are redundant, 6 measurements versus 3 degrees of freedom, the data can actually be used for recognition. A linear-time geometric algorithm to do recognition from beam data is described in [29]. The implementation described there takes a few milliseconds to recognize and compute pose. Also

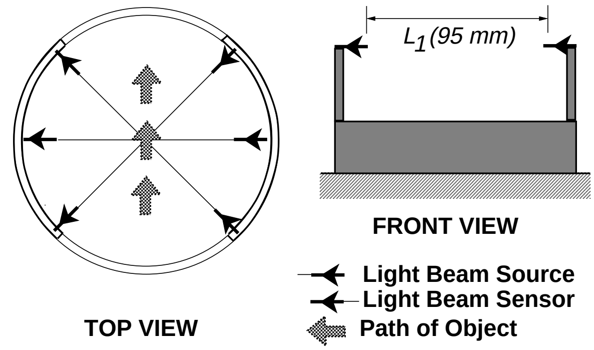

Figure 1: The usual cross-beam sensor configuration

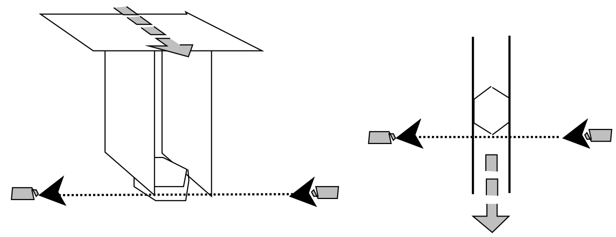

Figure 2: A *critical* point occurs when the part breaks or unbreaks the beam of light

described there is a hash table version which takes a few microseconds to accomplish the same thing. The cross-beam data is particularly well-suited to table lookup because the effective table dimension is only one. So with data quantized to 1000 values, the table takes up a few thousand words of memory.

The cross-beam sensor does not require a-priori part models. The cross-section of a part as seen by the sensor is easily found by passing the part through the sensor at most $2n$ times, where n is the number of sides of the *convex hull* of the cross-section. The part itself need not have convex cross-section, this is simply all that the sensor can sees.

The cross-beam sensor relies on a consistent horizontal cross-section to accomplish its task. It does not work for flat parts. For these we use a parallel-beam sensor, which usually uses reflective elements. A parallel-beam sensor is shown in figure 3.

The scan data from the parallel-beam is particularly difficult to deal with because even using relative measurements, the data still depend on two of the part’s degrees of freedom, unlike the cross-beam sensor which depends only on one. Indexing schemes generate lookup tables whose effective dimension is 2, and are consequently very large.

In [27] we described an $O(n + A)$ correspondence algorithm for parts with convex polygonal silhouettes, and an $O(n^2 \log n + A)$ algorithm for parts with non-convex silhouettes, where n is the part’s complexity and A is the total number of feasible matches. Typically for convex parts, A, the total number of matches is $O(n)$. The worst case

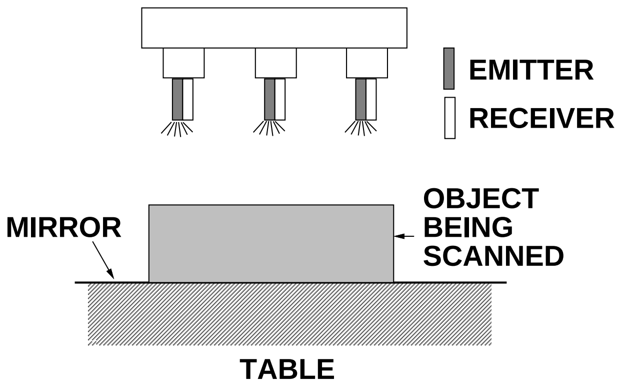

Figure 3: A reflective parallel beam sensor. Relative motion between sensor and part is normal to the page

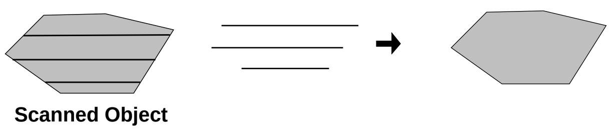

Figure 4: The parallel-beam scan data determined from a part's shadow

for convex parts is $O(n^2)$, and the worst case for general parts remains $O(n^3)$ although it is typically much lower. The model acquisition problem for parallel-beam sensors is harder than for cross-beams and is listed as one of our open problems.

3 Feeding and Sorting Parts with a Parallel-Jaw Gripper

Perhaps the least complex manipulator is the parallel-jaw gripper, having one degree of freedom with binary pneumatic control. Although widely used in industry, conventional wisdom holds that these grippers lack versatility [19]. With a minor modification, however, these grippers can be used to recognize and orient an important class of industrial parts.

3.1 The Modified Gripper

The quality of a grasp configuration depends on many factors including the orientation of the part with respect to the gripper. This orientation may not be known precisely or may be disturbed by the act of grasping. For the parallel-jaw gripper grasping polygonal parts, Brost [3] defined a grasp as stable if at least three vertices of the part are in contact with the gripper jaws and any further closing of the gripper would deform the part; see Figure 5.

Unstable grasp configurations result from friction between the part and the jaws. This suggests that it may be desirable to eliminate friction between the part and the jaws. One approach is to coat the jaws with grease, but this has the disadvantage that the part will slip when the gripper

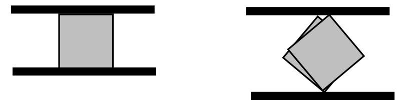

Figure 5: The grasp configuration on the left is stable; those on the right are not.

is lifted out of the plane. We can achieve low friction in the plane of the part but high friction orthogonal to the part by mounting a sliding plate (linear bearing) on one jaw. The inner surface of both jaws is covered with a high-friction material such as rubber. See Figures 6 and 7.

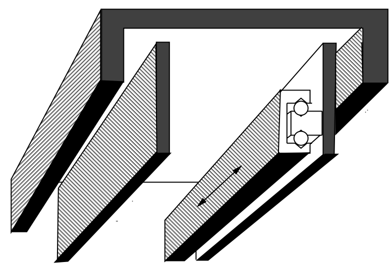

Figure 6: The modified gripper with sliding jaw.

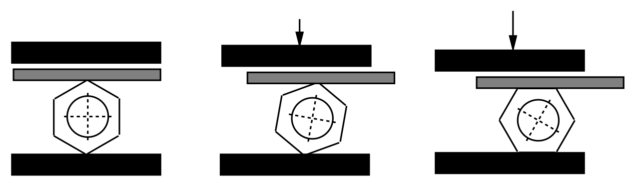

Figure 7: Time-sequence of grasping with the modified gripper. (1) As the two outer jaws close over a typical part (hex nut), horizontal forces cause the sliding jaw to translate to the left (2) until the part is gripped in a stable configuration (3). (Based on drawing by Ben Brown).

To achieve a stable grasp, we close the jaws as far as possible without deforming the part. Gripper forces will cause the sliding bearing to translate until the part rotates into a stable orientation. We built a prototype of this gripper using an off-the-shelf linear bearing with a rubber band to provide spring force and a dab of grease to provide damping. We experimented with several part shapes grasped randomly in two sets of 250 trials. Without the bearing, approximately half the grasps were stable. With the bearing, every grasp was stable. Although any physical bearing experiences some friction, we will continue experiments to verify that sticking is very unlikely in practice.

One of the primary advantages is that this modification requires no additional sensors or actuators; a low-cost and lightweight linear bearing can be easily retro-fit to any parallel-jaw gripper without requiring interface software [9].

3.2 Recognizing Parts

The parallel-jaw gripper described above can be used to recognize parts by measuring the distance between the jaws, say with a linear potentiometer. This is similar to using light beams as described in Section 2; the difference is that in this case closing the jaws causes the parts to rotate into a new configuration. For a given set of k parts with constant cross section (2.5D parts), we consider the following two problems: (1) given a set of measurements derived from random grasps of one part, decide which part was grasped. (2) find a sequence of grasp angles for the gripper, conditional on measurements, for efficiently recognizing parts from the given set.

For the first problem, since more than one part may give rise to the same diameter and the diameter sensor may be corrupted by noise due to surface compliance and backlash, we can use a Bayesian decision procedure to estimate the most probable part. Since the set of grasps is random, we can assume that prior to each grasp, the part's orientation with respect to the gripper has a uniform probability distribution on the set of planar orientations. Note that each stable orientation of a given part corresponds to a minimum in the part's diameter function. Thus the prior probability for each measured diameter can be derived in time $O(n)$. This becomes a conditional probability when considering a set of parts. Lacking any information to the contrary, we might assume that initially, each part is equally likely. After each measurement, the posterior probability is computed. After all measurements have been considered, we can decide on the most likely part. This method can also be adapted to allow for sensor noise using a Gaussian error model. For details see [16].

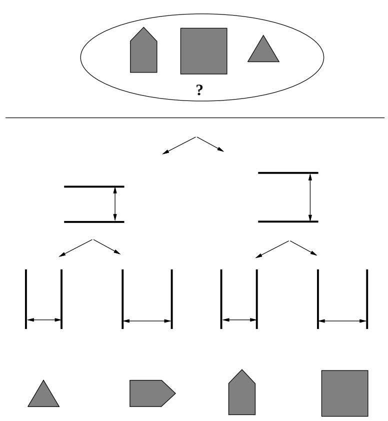

Figure 8: A grasp plan for distinguishing the three parts shown at the top.

Of course we expect to achieve better performance by tailoring the grasp strategy to the geometry of parts in the set, as shown in figure 8. This, the second problem stated above, is the planning problem. As mentioned earlier, [23] showed that some parts cannot be distinguished by measuring diameter alone. But if we restrict attention to parts that are distinguishable, we can find optimal strategies by considering connected components in the following graph. Let $G = (V, E)$ be an undirected graph such that each vertex corresponds to a stable orientation from the given set of parts. Let $n = |V|$. We construct an edge between any two vertices with the same diameter. For each edge, let $R(e)$ be the set of gripper orientations that would disambiguate the neighboring vertices in a subsequent grasp. Let G have m edges. We can construct G in time $O(n^3)$ and it can be partitioned into connected components (which are in fact cliques) in time $O(n + m)$.

Each possible gripper measurement identifies one of the components of G. If the component contains stable orientations from one part, we are done. If it contains stable orientations from exactly two vertices connected by edge e, we pick a gripper angle in $R(e)$, and regrasp to disambiguate between the two associated parts. However, if the clique contains more than two vertices, we look for a gripper angle in the intersection of $R(e)$ for all edges in the clique, $\bigcap_E R(e)$. If this intersection is empty, we require more than one additional grasp to identify the part. In [23] we give two planning algorithms. The first runs in time $O(n^2 \log n)$ but may not generate the shortest plan. The second finds the shortest plan but may require time exponential in n in the worst case because it considers all possible partitions of each clique. In either case, the resulting plan will never require more than n grasps. This algorithm can also be adapted to account for measurement noise.

3.3 Feeding Parts

For a known part, it is possible to achieve a desired final grasp configuration *without sensors*. Goldberg [10] describes an algorithm for orienting polygonal parts using the parallel-jaw gripper to grasp and ungrasp at a prespecified sequence of angles depending on part geometry. That is, for any additional stability criterion that prefers one of the stable grasp configurations over the others, we can achieve it using the compliant motion algorithm. The planning algorithm finds the shortest such sequence for any n-sided part in time $O(n^2)$. This algorithm has recently been extended to curved parts [22].

4 High Precision Part Insertion

RobotWorld (See figure 9) is a commercial robot system with multiple 4dof cartesian "placement modules" in a single workspace [26]. Every motion dof of the modules, and also the conveyor and a vice, are treated as actuation units. We refer to unions of these as *virtual robots*.

First, the reflective (sniffer) sensor, which is mounted on a module, is passed through the cross-beam sensor, instantly calibrating the sniffer. Another pass would instantly calibrate the cross-beam sensor itself, that is, determine the angles of its beams. Thus, even a newly-built, uncalibrated, beam sensor can be added to the workspace and run at full accuracy in seconds.

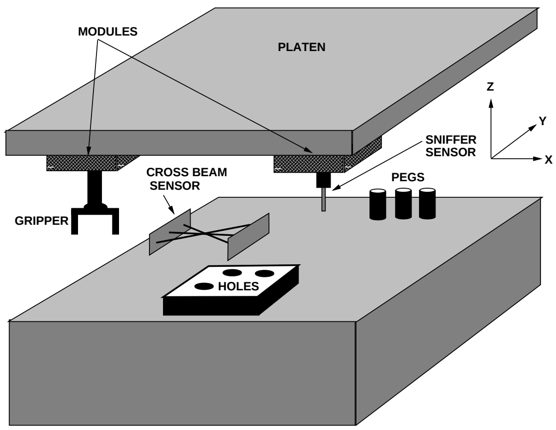

Figure 9: The RobotWorld multi-robot system

Next, the sniffer sensor localizes the hole using a simple "+" pattern. Concurrently with this, another placement module with 3-jaw gripper acquires a peg and moves it through the cross-beam sensor. Once the sniffer has finished and moved clear of the hole, the peg is moved over it, and inserted. With this scheme, we have achieved 25 micron clearance insertions at 99% success rate. More importantly, we can move the sensor or hole plate, or even use a new sensor, and the same results apply.

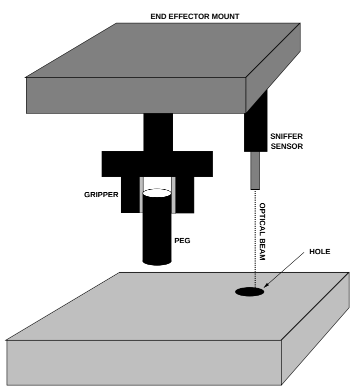

Figure 10: An instrumented gripper

This scheme applies not only to RobotWorld. For other robots which have poorer absolute accuracy, we suggest the instrumented gripper shown in figure 10. In this case, in a single pass through the cross-beam sensor, both the peg and the sniffer sensor are localized. The accuracy of the subsequent peg insertion is determined by the relative position accuracy of the robot over the distance between peg and sniffer.

5 Design of Modular Fixtures

Commercially-available modular fixturing systems typically include a square lattice of tapped and doweled holes with precise spacing and an assortment of precision locating and clamping elements that can be rigidly attached to the lattice using hardened bushings or expanding mandrels. Currently, human expertise is required to synthesize a suitable arrangement of these elements to hold a given part. Besides being time consuming, if the set of alternatives is not systematically explored, the designer may fail to find an acceptable fixture or may settle upon a suboptimal fixture.

Brost and Goldberg [2] and Wallack and Canny [28] considered classes of modular fixtures that prevent a part from translating and rotating in the plane by providing 4 frictionless point contacts with the part's polygonal boundary. Both methods use an extra degree of freedom to supplement a set of fixed round peg restraints that lie on a regular lattice, since fixed constraints on a grid cannot constrain general shapes. In [2], the extra degree of freedom moves a linear clamp. In [28] a subset of the pegs move relative to the others. Hardware prototypes of both configurations are being built.

Algorithms are given in [2] and [28] that accept part geometry as input and synthesize the set of all fixture designs that achieve form closure for the given part. Brost implemented the algorithm of [2] on a Lisp Machine and demonstrated non-intuitive fixture designs. Variations on this fixturing model are also being studied by Bud Mishra at NYU.

6 Previous Work

Many others have argued for simplicity in the factory. As Dan Whitney argued his 1993 ICRA Keynote Address, it is vital to consider the robot in the context of the assembly environment [31]. Nevins and Whitney [20] stressed the distinction between structured and unstructured environments, and Whitney [30] pointed out the what is often an opposition between flexibility and efficiency in manufacturing. In structured environments such as factories, where repetition is the rule, the emphasis is on efficiency. The idea of modular manufacturing systems is gaining in popularity [25, 8, 14].

There are many examples of research that can be viewed as examples of RISC ; space does not permit us to include a survey. One example is the RCC wrist that achieves compliant peg-in-hole assembly using passive mechanical elements rather than active feedback in a general-purpose manipulator [7]. Recently, Goswami and Peshkin showed how to "program" such a wrist to achieve desired behavior by changing the damping constants of its passive components [12].

One of the recurring themes in our approach is the role of mechanical compliance in lieu of sensing. Mason [18] and his students have developed a science of part pose control through sliding motion, demonstrating that it is often possible to replace sensors with mechanical solutions. Others, such as Erdmann and Donald are studying ways to reduce the complexity of sensors. Other applications of mechanical compliance are explored in [4, 21].

An important principle in computer vision is the exploitation of *domain constraints* to simplify algorithms. Rather than representing all conceivable images, using domain constraints one describes the simpler space of images that can occur that satisfy various physical and structural constraints. RISC is the natural extension of domain constraints to hardware. A part on a table has only 3 degrees of freedom, so a sensor that provides 6 numbers is sufficient.

Kanade used the term "KISS" (Keep it Simple) to describe a collection of recent results in machine vision where simple processors at each pixel permit extremely fast update rates. Here, the correspondence problem can be avoided since motion between frames is greatly reduced [15]. While there is some relation to our use of simple elements, the primary difference is that Kanade applies fixed arrays of simple elements to unstructured scenes rather than planning for repetitive operations.

7 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. At UC Berkeley, Canny is now implementing a very general algebraic constraint satisfaction system to perform mechanical design. This system is the product of 5 years of work on practical algebraic algorithms. A preview of the system was presented in an ESPRIT workshop on motion planning in Rodez, France in March of this year [5]. Most of the algorithms that comprise the system had not been implemented before. In early tests, we have found that each contributes one to several orders of magnitude of speedup over other methods. Overall, the system should provide a qualitatively higher level of design automation than previous systems. If efficient, the automated design of workcells can provide feedback to the part designer during the design stage, in turn facilitating design of parts for manufacturability.

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.

In short, we propose that "hard" automation be reconsidered in light of the last 20 years of research in robotics and computer science. We argue that the flexibility observed in robotic systems is as much due to good algorithms as it is to hardware.

We circulated an early version of this paper and received some comments that we would like to address:

- **"RISC sounds like just another acronym, where are the new ideas?"**
  Robotics and manufacting are rife with buzzwords; one hesitates before introducing a new term. In this case it seemed useful to identify a common theme in a growing body of research. We acknowledge that on one hand, simple hardware is not new in manufacturing; on the other hand, computational planning algorithms are not new in robotics. What is new, we feel, is the combination of the two. This paper represents our best attempt to articulate a new direction for research.
- **"The problems are not new; vision researchers have solved 2D model-based recognition long ago."**
  Certainly it is true that commercial machine vision systems have matured and dropped in cost in recent years; These systems may continue to find applications in manufacturing. However, RISC sensing can be distributed throughout a workcell at very low cost, and in many locations cameras cannot reach. Processing some RISC sensor data is non-trivial. Conventional model-based vision techniques perform poorly on parallel-beam data. We have developed compression techniques that greatly reduce the size of indexing tables for feature lookup. 2D parallel-beam recognition has the same table dimension as the recognition of 3D objects from 2D edge data. We are currently implementing this compression scheme for the 3D model-based vision problem, and discovering similar reductions in lookup table size. So the study of RISC has in fact lead to improvements in the state-of-the-art for 3D vision.
- **"The approach should be validated experimentally before claims are made about its usefulness."**
  Rather than making claims, we are trying to generate interest in a new area of research. Of course specific algorithms should be carefully tested in labs and in industry. We have implemented almost all of the algorithms described in this paper and tested them in the lab. At UC Berkeley, we are now designing a complete workcell for the assembly of a model-aircraft engine. In the course of developing this workcell, we will test the capability of our existing hardware and see if new mechanisms or sensors are justified. Of course the ultimate test will on the factory floor.

The RISC approach suggests theoretical questions with short-term practical consequences. We believe this work holds potential for significant scientific progress during the next five years. Related projects are being initiated at Stanford, Sandia Labs, Carnegie Mellon, New York University,

University of Padua and with Adept Technology. We close with ten open problems.

8 Open Problems

1. Can we orient any planar part up to symmetry using parallel-jaw grippers (ie, parts with piecewise algebraic contours)?
2. Is there a polynomial-time algorithm to find the shortest plan for sorting parts with an instrumented parallel-jaw gripper?
3. Parts can be oriented with a sequence of fences as they pass on a conveyor belt [1]. Are frictionless fences complete for the class of Polygonal Parts? (Is there a part we cannot orient with fences?)
4. What is a lower bound on the complexity of designing modular fixtures? For what class of parts are modular fixtures complete?
5. Given a known set of parts, can we locate a registration mark on each part to efficiently distinguish them?
6. Stable poses of 3D curved parts: Given a CSG part (constructed by negation, union, and intersection operations on n primitive solids) with its center of mass, what is the complexity of finding all stable poses of the part on a flat surface?
7. Given a family of parts, choose a beam layout to minimize the probability of mis-identification of parts and mis-calculation of pose.
8. Model generation: Using a sequence of probes with a moving beam sensor, plan a strategy for determining the shadow of a part for recognition by a parallel-beam sensor.
9. Pose determination from sparse depth probes. Given k fixed depth probes, determine part identity and pose given part models. (We are experimenting with simple beam arrays that will provide this data.)
10. Given polyhedral part shape, design a "pallet" such that parts flowing over the pallet will fall into the pallet in a unique orientation and are prevented from jamming [13].

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.

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 [11]. We thank the participants for helpful feedback, 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, Pradeep Khosla, Damian Lyons, Peter Allen, Daniela Rus, Todd Rockoff, and Howard Moraff for insightful feedback on earlier drafts of this report.

References

[1] Mike Brokowski, Michael A. Peshkin, and Ken Goldberg. Curved fences for part alignment. *ASME Journal of Mechanical Design*, 1992. (Accepted September 1993. A preliminary version of this paper appeared in the 1993 IEEE International Conference on Robotics and Automation, Atlanta, GA.).

[2] Randy Brost and Ken Goldberg. A complete algorithm for synthesizing modular fixtures for polygonal parts. Technical Report SAND 93-2028, Sandia National Laboratories, Oct 1993.

[3] Randy C. Brost. Automatic grasp planning in the presence of uncertainty. *The International Journal of Robotics Research*, December 1988. Also appeared in Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, April, 1986.

[4] Randy C. Brost. *Analysis and Planning of Planar Manipulation Tasks*. PhD thesis, CMU, January 1991.

[5] John Canny. A toolkit for non-linear algebra. UCB Technical Report, 1993.

[6] John Canny and Ken Goldberg. A RISC approach to sensing and manipulation. In *International Symposium on Robotics Research*, Hidden Valley, Penn., 1993.

[7] S. H. Drake. *Using Compliance in Lieu of Sensory Feedback for Automatic Assembly*. PhD thesis, MIT, 1977.

[8] H. Tsukune et al. Modular manufacturing. *Journal of Intelligent Manufacturing*, 4, 1993.

[9] Ken Goldberg. A kinematically-yielding gripper. In *22nd International Symposium on Industrial Automation*, October 1991. See U. S. Patent 5,098,145, granted March 26, 1992.

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

[11] Ken Goldberg, Matt Mason, and Ari Requicha. Nsf workshop on geometric uncertainty in robot motion planning: Summary report and bibliography. Technical Report 297, IRIS, August 1992. (The Workshop was held on Catalina Island, June 14-17, 1992 and supported by NSF Award IRI-9208161. A summary of this report appeared in the IEEE Robotics and Automation Newsletter, Vol. 6, Number 4, October, 1992. It is also available via anonymous ftp from 128.125.51.19.).

[12] Ambarish Goswami and Mike Peshkin. Mechanical computation for passive force control. In *American Control Conference*, 1990.

[13] Hajime Hitakawa. Advanced parts orientation system has wide application. *Assembly Automation*, 8(3), 1988.

[14] Debra Hoitomt. Paradigm shifts in manufacturing. *IEEE Robotics and Automation Society Newsletter*, 7(3), August 1993.

[15] Takeo Kanade. Keynote talk: The kiss principle in machine vision. In *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, July 1992.

[16] Duk Kang and Ken Goldberg. Shape recognition by random grasping. In *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, July 1992. To appear in the *IEEE Transactions on Robotics and Automation*.

[17] Tomas Lozano-Perez, Joe Jones, Emmanuel Mazer, and Pat O'Donnell. *Handey: A Robot Task Planner*. MIT Press, 1992.

[18] Matthew T. Mason. Kicking the sensing habit. In *Asilomar Winter Workshop*. AAAI, November 1991.

[19] P. J. McKerrow. *Introduction to Robotics*. Addison-Wesley, 1991.

[20] James L. Nevins and Daniel E. Whitney. Computer-controlled assembly. *Scientific American*, 1978.

[21] Michael A. Peshkin and Art C. Sanderson. Planning robotic manipulation strategies for workpieces that slide. *IEEE Journal of Robotics and Automation*, 4(5), October 1988.

[22] Anil Rao and Ken Goldberg. Grasping curved planar parts with a parallel-jaw gripper. Technical Report 299, IRIS, August 1992.

[23] 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).

[24] Frank J. Riley. *Assembly Automation, A Management Handbook*. Industrial Press, New York, 1983.

[25] G. G. Rogers. Modular production systems: A concurrent manufacturing philosophy. In *IEEE International Conference on Robotics and Automation*, 1993.

[26] Victor Scheinman. A multiple robot vision guided assembly system. In Robert Bolles and Bernard Roth, editors, *Robotics Research: The Fourth International Symposium*, 1987.

[27] A. Wallack and J. Canny. A geometric matching algorithm for beam scanning. In *SPIE Symposium on Vision Geometry II*, 1993. Boston, Massachusetts.

[28] A. Wallack and J. Canny. Planning for modular and hybrid fixtures. In *IEEE Conference on Robotics and Automation*, 1994. (submitted).

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

[30] Daniel E. Whitney. 'real robots don't need jigs'. In *International Conference on Robotics and Automation*. IEEE, May 1986.

[31] Daniel E. Whitney. From robots to design. *IEEE Robotics and Automation Society Newsletter*, 7(3), August 1993.