

ROBOTICS IN THEORY, ROBOTICS IN PRACTICE: THE 1992 IEEE ROBOTICS AND AUTOMATION CONFERENCE

Background

This conference brought together over 950 participants who heard over 450 papers. Participation from Europe was especially strong since this was the first year the conference has been held outside of the United States.

The meeting was held against the backdrop of NASA's efforts to retrieve the INTELSAT. The counterpoint is particularly ironic since the conference presented a great deal of theory that has as yet seen little or no application; all the while the best robotic technology available in space was unable to retrieve the satellite.

Robotics today has two branches, one dealing with prepared environments (industrial) and the other dealing with unstructured and unpredictable environments (underwater, catastrophes, and space). It was once thought that a great deal of sensing and perception would be needed even for industrial robots but this has turned out not to be true. Industrial robotics is burgeoning in many countries, especially Japan, exactly because time and circumstances are available in which to prepare the robot, its task, and its environment, then to practice the task and perfect it, so that it can be repeated many times. Sensing is used rarely to cover slight unpredictability, but this is sufficient since planning and preparation are the watchwords of manufacturing anyway.

Robotic researchers have had to turn to the unstructured environment in order to justify much of the sensing and dexterity work they have done over the last decade. Yet this environment has eluded them.¹ The fact that astronauts had to perform the rescue lowered the morale of many of the conference's participants, especially those who hoped NASA would increase funding on robotics. The keynote speaker, Prof Jacques Blamont of the Academy of Sciences of France, argued on Tuesday (two days before the satellite was recovered) that the progress of computer hardware and the difficulty of maintaining people in space would eventually drive them out, to be replaced by automation and robotics. "Man is there to smile and vomit," he said.

Obviously, man can do much more but the question remains whether robotics will replace him or not. We cannot go on asking astronauts to risk their lives so that the Olympics can be broadcast. Yet conference participants did not agree with Prof Blamont that the mere power curve of computing hardware was a convincing argument. Software does not follow as easily, and many problems are mechanical and cannot be fixed by software updates.

¹ It is important to note that the Shuttle's difficulties in retrieving satellites are due in large part to the inability to prepare the task environment. The issue is not lack of time to practice but rather documentation of appropriate grasping surfaces or lack of such surfaces on satellites. Regardless of how the task is done, improvisation is necessary from the start.

The issue is one of many that the robotics community has to grapple with: in many areas theory has run years or decades ahead of applications while in other areas either the lack of applicable technology is painfully obvious, or else quite reasonable success can be achieved with much lower technology as long as one is willing to be "unrobotic" and prepare an environment that guarantees task success.

Conference Sessions

I attended two workshops and several technical sessions that bring these points out. The workshops were Assembly Planning (in which I participated as a coauthor) and High Precision Sensors/Actuators and Systems. Sessions I attended included Constrained Motion Control, Mechanical Design of Robots, Fine Manipulation, Force Control, a panel on Large Scale National R&D Projects, Robots and Manufacturing (I was co-chair), and Neural Networks.

Assembly Planning

This is a relatively recent branch of robotics, dealing with fairly high level planning of task sequences rather than detailed planning of gross hand trajectories or fine insertion motions. It was interesting to hear the honorary workshop chairman Prof George Bekey of USC state in his introduction that "Assembly planning is planning and therefore part of Artificial Intelligence." It was interesting because the field originated in the efforts of researchers whose motivation was design of mechanical products and the equipment that would assemble them. The majority of the workshop speakers represented this view and said so explicitly. That is, assembly planning is part of product design and concurrent engineering, not AI.

The talks reflected this in their emphasis on linking assembly planning to CAD models of parts and to the economics and technical challenges of laying out factory floors and choosing appropriate equipment (not always robots!) to accomplish the planned assembly.

The central challenge that this field deals with is determining the thousands of possible assembly sequences and finding a desirable one from among the possible ones. It is challenging because the problem is combinatoric in principle, yet constrained from the worst combinatoric limits by factors that are not understood. Another challenge is to find efficient representations and datastructures for the knowledge that is required in order that efficient automatic or semiautomatic procedures can be used to find the possible sequences.

Two philosophies for representing the knowledge were discussed at the workshop, and the profound differences between them were not evident to the participants until late in the day. These philosophies are 1) geometric and 2) logical or constraint-based. The geometric representation is essentially a CAD model, and the geometric constraints that permit some sequences of assembly and prevent others are found by

computationally manipulating the geometry. The logical representation assumes that the parts were designed by someone who has the opportunity to encode with the geometry most of the information that is relevant to assembly, such as the direction in which parts must approach (screws go into holes along the holes' axes, for example). Such information is there in the design and thus does not have to be discovered by the assembly planning algorithm. This approach is called Feature-based Design and has a ten-year history of research, mostly in supporting the fabrication of parts. Use of FBD to support assembly planning is about three years old. It is the subject of a large new ESPRIT program that will be described in a separate report.

The status of assembly planning, as evidenced by the papers, is as follows:

- 1) The combinatorics have been brought under control, and products with 30 parts or more can be handled with relative ease. It is generally agreed that subassemblies with more than 30 parts are rare, and larger assemblies are usually broken into smaller subassemblies for other reasons. Some approaches require a person to provide intricate geometric reasoning, while others simplify the allowed geometry so that this is not necessary. Very soon this distinction between approaches will disappear since it is easy for users to provide the missing information, and permitting this approach makes the methods applicable to practically every mechanical product imaginable.
- 2) Several criteria and search methods for finding desirable sequences have emerged and have practical relevance: lowest cost, fewest machines, fewest reorientations of the product during assembly, and so on.
- 3) Successful links between these algorithms and CAD models have been demonstrated.
- 4) Good user interfaces exist, making it easy for users to tackle what used to be an impossible problem. In fact, the problem usually was not even recognized by product and process designers since no one gave much thought to finding good assembly sequences systematically until Concurrent Engineering came along and showed that this was an important design step.
- 5) The frontier of unsolved problems includes tolerancing, products with flexible parts, and a stronger link between feasible sequences and desirable factory floor layouts. Several researchers indicated that they were working on these problems.

Paper titles, authors, and brief summaries:

Lega: A Computer-aided Generator of Assembly Plans, by J Henrioud and A Bourjault, Universite de Franche Comte, Besancon, France

Bourjault launched this field with his PhD in 1984. It contains the basic ideas of modeling the product's interconnections with a logical graph called the liaison diagram, and finding the allowed sequences via a Q&A with the designer. The original method required too many questions, and many researchers, including Henrioud, have directed their efforts at making the interaction more efficient. These authors firmly place their work in the domain of product and process design for industry. Recent advances include being able to simplify the search by permitting the designer to define groups of parts to lump temporarily (all share the same approach direction or tool, for example). Other researchers (De Fazio and Whitney, Lee, cited below) have done similar things.

Computer Aids for Finding and Evaluating Assembly Sequences: What is Now Done and What are Some Gaps in Current Application, by T L De Fazio and D E Whitney, Charles Stark Draper Laboratory, Inc, Cambridge, Mass

De Fazio and Whitney were the first to see the need for more efficient user interaction while seeking assembly constraints, and they developed two classes of algorithms. The most recent is called "onion skin" and requires the user only to identify parts remaining on the outside of the assembly that can be removed. The process is strongly aided by the use of assembly features obtained during the design process (or in a concept design implementation, merely typed in by the user). The required information consists of part names, a connection list, and a direction of mating or separation for each part pair. This work has extended to providing users with filters for finding sequences that obey certain industrially useful constraints or that optimize cost, number of reorientations, etc. Output sequences have been connected to assembly system design algorithms. Unsolved problems include linking "facility constraints" into assembly system design: some operations must be done at certain places on the factory floor where space, ventilation, or other resources are located; part supply lines must not interfere with the line, etc. It is similar to circuit board layout. Another area in need of work concerns creating assembly plans that meet quality or strategy requirements. Quality requirements include the need to intersperse in-process test and rework: some sequences are better for this than others. Strategy includes planning for seasonal variations or permitting many models of the same product to be made in unpredictable proportions. Preliminary work has been done in both areas. Tolerances are also starting to be attacked.

Task Decomposition and Planning in an Assembly Workcell Using AND/OR Nets, by A C Sanderson, L S Homem-de-Mello, and T Cao, Rensselaer Polytechnic Institute (Homem-de-Mello is at JPL)

Homem-de-Mello did a seminal PhD on assembly planning in 1989 in which he proved the equivalence of several different datastructures and constraint representations for assembly plans. This presentation summarized that work and follow-on since, all again placed firmly in the context of industrial product design. The algorithms for finding feasible sequences are similar to those described in the

other presentations. The new work includes pruning unwanted sequences by forming groups like those suggested by Henrioud. More interesting is the conversion of a selected assembly plan into a Petri net so that a robot action plan can be obtained. This is the subject of Cao's PhD, in progress. The talk included a video of a two-arm robot system with vision and force sensing planning and executing the assembly of some strut-like parts of a miniature space station. The demo obviously included work by many others besides Cao since it required coordination of many degrees of freedom and sensors, plus scheduling of a complex multiprocessor control system. Other new work includes representing tolerances in the assembly data and adjusting the sequence or the robot's plan to take tolerances into account.

Assembly Research at Sandia Lab, by David Strip, presented by Randy Brost

Sandia is attempting to find ways to automate low volume production of complex mechanical assemblies, such as are found in the triggering mechanisms of nuclear weapons. Cleanliness and care are very important, but the volume is so low that people have a hard time learning the right methods before all assemblies have been built. On the other hand, conventional robotics usually requires higher volume than DOE has in order to be economically attractive.² Strip identified difficulty of teaching and of designing fixtures, and created an assembly planner that can do both. The planner works only on a very restricted class of objects, however, and he has not extended it beyond these limits since the original work about two years ago. However, it is the only assembly planner that includes fixture design in the process. It operates on simple CAD descriptions of the parts and is rule-based.

Backward Assembly Planning with DFA Analysis, by Sukhan Lee, USC

Lee's approach is similar to Homem-de-Mello's and they have collaborated for two years. (They also organized the workshop and recently published a book of papers on the topic which contains reports by all the presenters here plus others.) Lee's main contribution has been to identify need to break large assemblies into subassemblies in order to help reduce the combinatoric load. While this is often done routinely in industry, Lee has been seeking automatic methods. These include ones similar to Henrioud's, such as defining common directions or tools. Also on Lee's list are common operations that do not accomplish assembly, such as testing and painting. Several of his methods are heuristic and are difficult to understand. This is true of other papers in this field when the need to include desirable functions exceeds the ability to represent them quantitatively. The needs of real designers are still beyond the abilities of algorithms here, and this is why several researchers have fallen back on simply querying the designer at key points in the process.

² Actually, the economics are complex and a good case can often be made for low volume automation in precisely this environment. The argument is based on the percentage of successful assemblies that can be obtained, and the cost of failed ones. Automation usually has higher yield.

An Integrated Assembly System, by C S G Lee and Y F Huang, Purdue University

This paper also places assembly planning in the industrial context. The work was done in conjunction with Purdue's Engineering Research Center, which already has facilities for "quick turnaround" CAD/CAM of single parts. Lee's work adds assembly to this capability. An important component of this work is the attempt to include facility capabilities such as the constraints of tools and fixtures. Another is the wish to provide designers with direct feedback on assembly problems and suggestions for redesigns. This last is one of the most difficult goals to achieve in this field, and the wish has again run ahead of the ability to deliver. The approach has been largely implemented by asking the user to comment, with very simple rules being used by the computer to point out areas where problems exist. The group discussion after this paper indicated that people doubt whether serious redesign suggestions can be expected from computer-based methods due to the need to be creative and to understand many unstated design constraints and detailed engineering knowledge.

Methods of Knowledge Representation for Assembly Planning, by Jan Wolter, Texas A&M University

This was a useful paper because it summarized most of the existing methods for representing geometric assembly constraints and the resulting plans. It also carefully compared the assumptions and limitations of each method. There is more similarity than differences, so the result was to show where the field is and what problems remain. For example, all approaches are limited to one approach direction at a time and no temporary positions for parts during assembly. Only one method permits three or more disjoint part sets to be mated during one operation (Wilson, see below). He has proven, as part of his PhD, that the main types of constraint used by different methods are in fact equivalent.

Following these main presentations were several minipresentations for which there are no papers.

Carlos Ramos, Oporto University, Portugal described a five-agent software system that plans and executes assembly. Tasks are described symbolically, and the user provides some of the geometric reasoning.

Randy Wilson, Stanford University, described his recent PhD in which he developed what is probably the most efficient algorithm for finding the set of feasible sequences. He works with CAD models and finds separation constraints by finding contact planes between parts and determining separation directions. The set of allowed directions is then reduced to its intersection, which determines how the part can separate from its neighbors. As in other methods discussed today that take a similar approach, the fact that a part is trapped is easily discovered once the set of

escape directions is known. Wilson then exploits this knowledge very efficiently to produce the final result. However, Wilson assumes no design information is available, in contrast to the feature-based-design approach. Thus his method has to recreate a lot of information that is not exclusive to assembly planning. The combination of the efficiency of the rest of his algorithm with the FBD approach would create a dynamite method.

Luis Homem-de-Mello of JPL described his recent work planning the assembly of the space station. This structure is regular and has regular substructures. Luis showed that he could represent these structures in a hierarchical graph in which all the geometric constraints were encoded as symbols of different shapes. The search for allowed assembly operations is very much simplified and resembles the FBD approach in many ways. It also places the two knowledge representation philosophies (geometry vs logic) in the starkest contrast and shows the power of the latter most strongly of all the papers presented.

Louis Whitcomb of Yale described his recently completed PhD in which he uses potential functions and dynamic equations of motion to generate assembly plans. Parts repel each other and are attracted to their final desired locations in the assembly. He has implemented this only for spherical objects. He showed a video of these parts literally dancing into their final positions.

Workshop on Microsensors and Actuators:

This all-day workshop is in an area unfamiliar to me, and I came to learn. My impression of micromachines is that the field is in the exciting discovery period in which the possibilities are just starting to be discovered. The limitations are not being considered at all, and no one can predict whether really useful techniques and devices will emerge. In some cases, better alternatives exist. An example is "cell sorting," a possible application cited by several speakers. The micro-machine approach is a scale-down of a macro-machine approach, as if cells were little boxes at the biological post office. At the post office one would identify each box and mechanically push it in the right direction. So the micromachine people envision little micropushers. However, cell sorters are starting to be developed that work on fluid flow principles, using valves or gates to switch the flow and direct cells in different directions. Since cell sorting is likely to be a high volume job, a flow approach makes sense.

Another clue that the micromachiners are still in discovery mode is that they are generating new capabilities in isolation from possible classes of applications, and in particular have not analyzed applications to obtain design requirements on their machines. When I asked what force the actuators could produce I was told "enough to support their weight." When I asked what forces were needed for target tasks, there was no answer.

Construction of these devices has advanced to the point where they are getting complex, often consisting of a half dozen or more individual parts that must be "assembled." Assembly takes two forms. One is well known and consists of various semiconductor processing steps that add or remove material in a complex sequence. The other is new and involves making parts in the plane and folding them up to make 3D structures, rather like setting up a flagpole or stepping a mast.

The point is that design of such structures and their assembly processes requires a lot of skill, experience, and 3D geometric reasoning. I determined from asking both US and German speakers that there are as yet no design tools to aid this process. (One paper described using CAD to visualize how the folded-up structures might fit together.) Here is an important research need, one that the assembly planning community might contribute to.

Papers, authors, and brief summaries:

New Microactuators, by Prof T Fukuda, Nagoya University

He described several electrostatic devices made by layering methods. An example is a parallel jaw gripper. It is made from a monolith by semiconductor processes. The moving jaws are suspended by slender leaf springs that permit relative jaw motion in one lateral and two angular directions. The different directions are obtained by turning on and off different combinations of segments on the face of the actuator. He can get motions as large as 0.18 degrees and 50 microns. The devices must operate in a vacuum because squeeze film damping would bring their motions practically to a halt. He uses capacitive sensors to close control loops around the jaws to obtain controlled motions. For some reason, however, he has not designed the devices with actuation symmetry (push-pull arrangement that eliminates most nonlinearities), with the result that his controller is complex and results are only fair. He has extended the principle to a complete 6 degree of freedom actuator by combining three of the simpler ones. His most interesting device was an "optical actuator." It is made from material that changes length when UV is shone on it. The principle is one of polarization, not heating, although other speakers used the latter. Fukuda has built little untethered walking horses and showed a video of one stubbing its way along. These are the first micro autonomous vehicles. Right now the motion is very slow.

Microstructures and Microactuators for Submillimeter Robots, by Ron Fearing of UC Berkeley

Fearing has constructed 3D devices by folding them up from 2D shapes. Folding is done either manually, painfully, or somewhat destructively by blowing fluid over the flat structures. Many do not survive, but he will undoubtedly improve on this. Folding takes place around hinge joints that are made by a selective plating and etching process. Some hinges are made with square pins to provide a bistable design: the parts want to snap upright once they are pushed up a little from the flat

orientation. Among the structures he has built are multi-part hot-wire anemometers, complete with built-in signal amplifiers. These fold up in several directions and have to interlink in the process. Tab-slot assembly is used, which is done manually with microprobes under a microscope. It's incredible. The structures have integral hot-wire sensors made from laying down 3 micron thick silicon layers. The amplifiers are on the structures and get their power from "ribbon cables" made from 50 micron wide 3 micron thick layers. The ribbon cables span the space at the foot of the structure where it folds up from the substrate. He plans to chain hinged structures together to make little robot arms. Actuation would be provided by electrostatic forces acting across adjacent joints. An important aspect of his approach is that it would need only 10 volts, in contrast to Fukuda, who needs 300. He has anticipated problems from air friction, and feels that he can live with the slow speeds it would cause.