MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Memo 1182 December, 1989

Fast, Cheap and Out of Control

Rodney A. Brooks and Anita M. Flynn

Abstract

Spur-of-the-moment planetary exploration missions are within our reach. Complex systems and complex missions usually take years of planning and force launches to become incredibly expensive. The longer the planning and the more expensive the mission, the more catastrophic if it fails. Always the remedies have been thought to be ever better planning, more redundancy, more thorough testing, and higher-quality components. We argue here for cheap, fast missions using large numbers of mass produced simple autonomous robots that are small by today's standards, perhaps 1 to 2kg. Let loose upon a planet and out of control of ground-based mission planners, we argue that such robots enable the time between mission conception and implementation to be radically reduced, launch mass to be slashed, totally autonomous robots to be more reliable than ground-controlled robots, and large numbers of robots to change the tradeoff between reliability of individual components and overall mission success. Lastly, we suggest that within a few years it will be possible, at modest cost, to invade a planet with millions of tiny robots.

Support for this research was provided in part by the University Research Initiative under Office of Naval Research contract N00014–86–K–0685, in part by the Advanced Research Projects Agency under Office of Naval Research contract N00014–85–K–0124 and in part by a gift from Siemens Corp.

1 Introduction

Based on our experience building ground-based mobile robots, we suggest quickly mounting cheap missions using many mass-produced, simple, small autonomous robots. Over the last four and half years the Mobile Robot Group within the lab has attempted to build completely autonomous mobile robots. We have refined hardware and software tools so that we can build a new one quickly. For instance, Genghis, a six-legged walking robot, was completed in 12 weeks for a Jet Propulsion Lab workshop on microspacecraft **[Jones 88]**. The robot **[Brooks 89]**, **[Angle 89a]**, was principally built and debugged by two people, with occasional help from about half a dozen others. The robot, shown in figure 1, weighs less than a kilogram and can scramble over rough terrain. A follow-on vehicle **[Angle 89b]**, will be able to climb meter high rocks and travel at around three km/hr. Such easy to build, high performance robots suggest some new ways of thinking about planetary exploration.

Two of the principle costs in planetary surface exploration missions are the mass of planetary rover upon launch, and hand construction of a one-of-a-kind vehicle. Both problems can be attacked simultaneously by creating swarms of totally autonomous 1-2kg microrovers. Mass delivered to the planetary surface would be minimized and the large number of rovers would increase the chance of mission success. Mass production of the rovers would lower cost per kilogram.

Total autonomy would actually increase mission reliability. Robots could use force control with tight sensing feedback loops. This is in contrast to the minutes- to hours-long position control feedback loops of teleoperation at planetary distances. By removing all ground-based control of the rovers, their complexity would be reduced drastically, as there would be no need for much of the communications equipment, nor for ground support to maintain communications. Simplicity would increase reliability. And it would allow complete programs to be conceived, researched, developed, and launched in times more reminiscent of the '60s than the '80s.

In the last part of this paper we present some radical ideas on how to scale down the size of planetary rovers even further, to the milligram range inspiring missions which will capitalize on thousands or even millions of rovers let loose on a planetary surface.

One of the keys to such claims is that robots can exhibit clever behavior under the control of small areas of silicon as long as their intelligence is organized the right way.

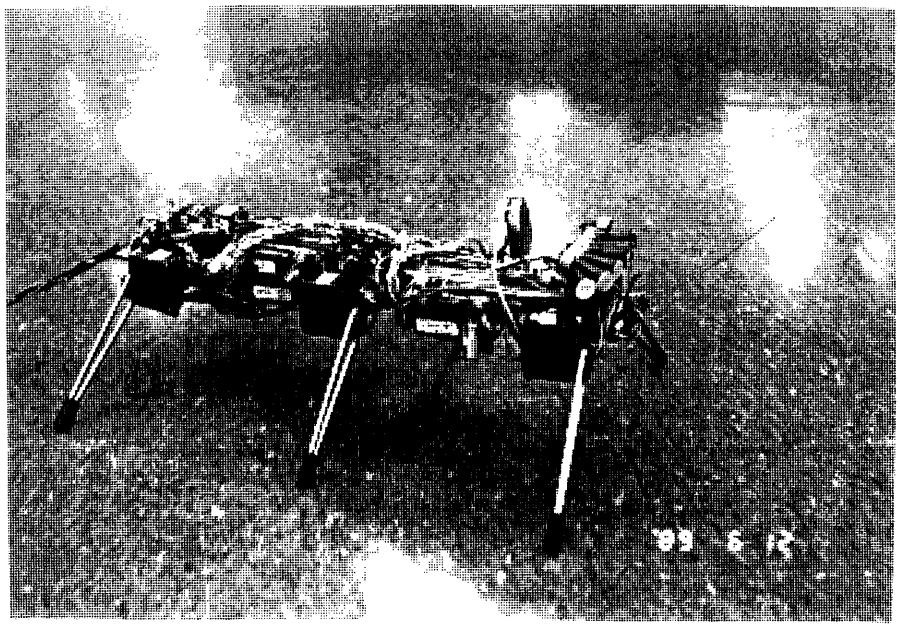

Figure 1: Genghis is a 1kg six legged robot. It can walk and climb over rough terrain carrying a variety of sensors. It has four onboard processors, twelve actuators with force feedback, six pyroelectric sensors, two whiskers, and pitch and roll inclinometers. Total time for the project between initial conception and completion of the robot was twelve weeks.

2 Intelligence for Mobile Robots

At first, the Mobile Robot Group tackled the questions of what the essential components of an intelligent robot are and how they should be put together. Influenced by experiments with actual robots, we took a different approach from the traditional one in artificial intelligence. We decided to drop most mission planning, do away with a computerized world model against which perceptions and actions of the robot are judged, use biology and evolution as models in design, and build complete systems and test them against the real world so that we would not trick ourselves into skipping hard problems.

In practice, these ideas lead to a general layering methodology for organizing the intelligence system. Simple behaviors were built first connecting sensing to actuation. Then higher level task behaviors were layered on in parallel. Because, when necessary, a more sophisticated higher layer takes over, or subsumes, a task a simpler lower one would do, we called this framework the subsumption architecture **[Brooks 86]**. A variety of robots have been built which carry out different tasks. The family portrait of all our robots is shown in figure 2.

One writes a subsumption program by specifying layers of networks of augmented finite state machines. Finite state machines are logic circuits with a small number (often under 10) states. Augmented finite state machines have a timer added. As an example of their use, a rover may be told to leave the mother ship and search for something. If a rover were unsuccessful, after a set time it would be told to return to the mother ship. Or if the leg of a walker were swung forward and did not encounter a force after a set time – perhaps because it is dangling over a dip or crevasse – it would be told to reposition itself. With a timer, rover action does not have to rely entirely on outside events.

These principles have been built into Genghis and Squirt. Genghis walks under subsumption control and has an extremely distributed control system. It successfully walks over rough terrain using 12 motors, 12 force sensors, six pyroelectric sensors, two inclinometers and two whiskers. A pyroelectric sensor registers a change in temperature within its field of view. With its array of pyroelectric sensors, Genghis can follow cooperative humans.

Genghis has no central controller. Instead, each leg is granted a few simple behaviors and each leg independently knows what to do under various circumstances. For instance, one of the most basic behaviors can be thought of as, "If I'm a leg, and I'm up, put myself down." Additionally, there are behaviors such as, "If I'm forward, put the other five legs back a little," and

Figure 2: The MIT Mobots come in a variety of shapes and sizes. Toto, Allen, Herbert, Seymour and Tito can be seen in the back row, left to right. Genghis, Tom and Jerry, and Labnav are in the middle row and tiny Squirt is in the foreground.

"If I'm up, then swing forward." These processes exist independently, run at all times, and fire whenever the sensory preconditions are true.

The only scrap of central control necessary for walking is sequencing the legs. As soon as a leg is raised, it automatically swings itself forward. But the act of swinging forward causes all the other legs to move back a little. Since those legs touch the ground, the body moves forward. Now the leg notices it is up in the air, and so puts itself down. Then the next leg lifts, and so on.

Additional layers pay attention to new sensors such as pitch and roll inclinometers or force sensors on the legs. With these sensory triggers, new behaviors can be composed that make the robot improve its walking performance. But there is no need to modify the original basic layers. New higher level behaviors just suppress the original layers whenever the higher levels get triggered. So if Genghis is climbing over a pile of rocks and one leg detects a high force before it has reached its set position, it triggers a behavior to move the set position closer to the current position. This is done by suppressing the original layer's command to the leg. Code for the lower layer has not been altered, just ignored in appropriate circumstances. As a result, the robot's legs comply with rough terrain. Back on flat terrain, higher level behavior is not triggered.

Eight incremental layers make up the control system: stand up, simple walk, force balancing, leg lifting, whiskers, pitch stabilization, prowling, and steered prowling. Prowling and steered prowling pay attention to pyroelectric sensors that detect people. They suppress walking unless triggered. Thus the net effect is that Genghis stands up and stays still until a person walks through its field of view. Then it attacks.

Fifty-seven finite state machines comprise the control system, 48 of which are organized as six copies of an eight-machine control system for each leg, two of which are associated with local behaviors connecting whiskers to the front legs and two of which are associated with inhibition of balance behaviors in the front and back pairs of legs. This leaves only five finite state machines with any sort of central role, of which two coordinate walking, one coordinates steering and two produce following behavior using pyroelectric sensors.

This manner of organization does away with any need for sensor fusion or having to make judgements about which sensors to believe and when. One fallout of this approach is that there is no need to calculate footfalls or safe places for the robot to place each leg. Genghis does not bother to try to place a rear foot on a footprint from an earlier leg either. It just puts its

foot down, but if that causes the body to pitch, then the pitch stabilization behavior kicks in and adjusts the stiffness of the force-balancing behavior on the load-bearing legs. Failure to make this adjustment would allow those load-bearing legs to collapse in a misguided effort to shift weight to other legs.

By having lots of tight, real-time feedback loops that run in parallel and respond to sensory input, we can get around the bottlenecks of long contemplative thought about what to do with input from a multitude of sensors.

We are building a new version of Genghis, called Attila, that will be a much stronger climber and able to scramble at around 3 km/hr. Each leg (shown in figure 3) **[Angle 89b]** has three degrees of freedom and three force sensors mounted on load-bearing beams. Each leg has a single-chip microprocessor with onboard random-access memory and a program stored in EEPROM. Attila will weigh 1.3kg, including batteries to power about 30 minutes of walking. After that, it will have to recharge from solar cells for about 4.5 hours in Earth sunlight.

Another robot, Squirt **[Flynn, Brooks, Wells and Barrett 89]**, is the smallest we have ever built. It weighs less than 50g and occupies about one and one-quarter cubic inches. Most of a robot is usually made up of motors and batteries, whereas sensors and computers take up only a small amount of space. We built Squirt as an exercise in shrinking brawn down to the scale of the brain using strictly off-the-shelf components.

Even at its modest dimensions however, Squirt incorporates an 8-bit computer, an onboard power supply, three sensors and a propulsion system, which can be seen in figure 4. Its normal mode of operation is to act like a bug, hiding in dark corners, venturing out in the direction of noises, only after the noises are long gone, and looking for a new place to hide near the origin of the previous noises.

3 Planetary Rover Scenarios

Genghis and Squirt show that it is possible to build small autonomous mobile robots that would make possible new ways of exploring planets. But could small vehicles traverse as rough terrain as large vehicles? The answer depends on the means of locomotion. On Earth, ants can traverse much wider varieties of terrain than humans or machines. Admittedly, they cannot jump over large fissures, but over most of the Earth's surface these are

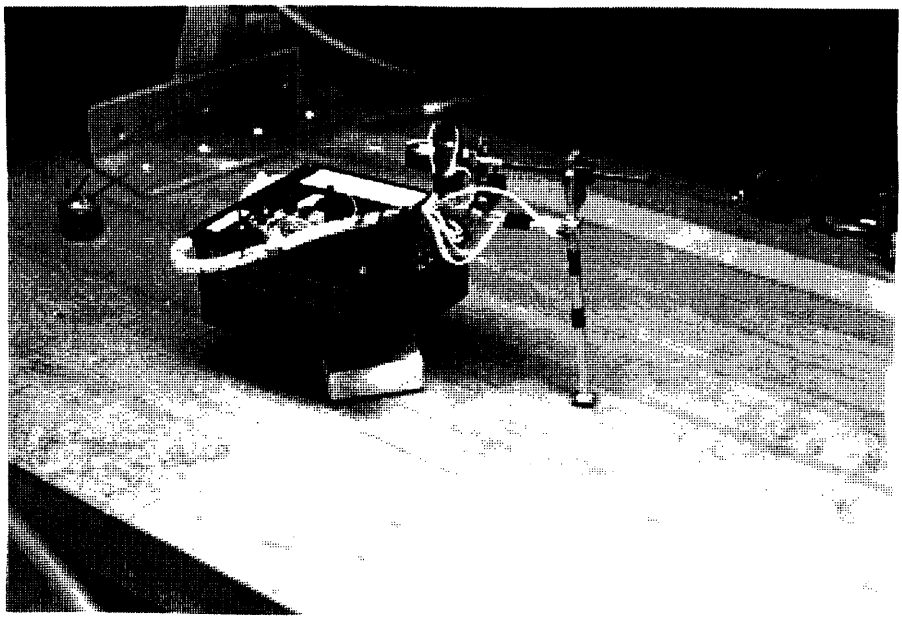

Figure 3: A three-axis force controlled leg is the basis for a new six-legged rover. Each leg will have its own microprocessor for force servoing. Microprocessors will be connected together in a ring network. Total weight for the six legged walker will be 1.3kg, including batteries and solar cells for recharging. On Earth the robot will have a 10% working duty cycle during the day.

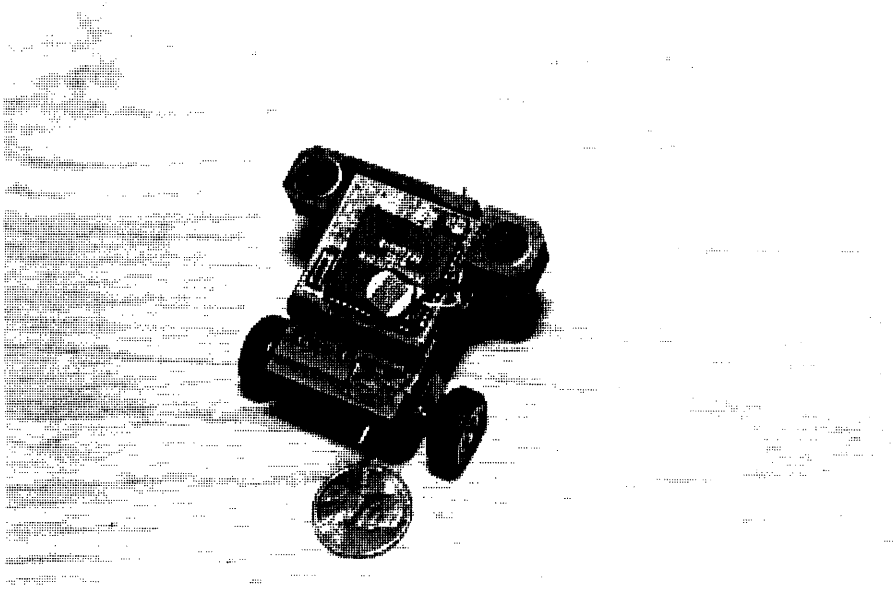

Figure 4: Squirt, the smallest robot we've built to date, packs motor, batteries, microcomputer, interface electronics and three sensors into a volume slightly larger than one cubic inch.

rather rare.

In fact, it was easier to make Genghis walk well than it is to make larger robots walk well. At a smaller scale, the strength-to-weight ratio increases dramatically, since mass goes down by a cube law, while cross sections go down only by a square law. Getting a leg stuck in a crack temporarily does not mean disaster as it might for a horse-sized rover. Furthermore, if a foot placement is missed, the distance to fall is short and the impact velocity low. With Genghis we have found that we can simply ignore foot placement issues and rely on persistent oscillation of the leg to get the robot out of troublesome situations.

3.1 Augmenting a Large Rover

There are major problems with planning a space mission which relies solely on one large planetary rover. If a mission is restricted to such a single large robot, there is a tremendous cost associated with losing the rover and thus a rash of conservatism will develop among the mission planners. There could be great trepidation in sending the vehicle into terrain that was unknown, rough, sloped with loose gravel or otherwise apparently dangerous, even though the area could be scientifically very interesting.

However, if the large rover carried a set of small potentially disposable 1kg rovers along, it could open up options available to the mission planners. There is much lower cost associated with losing one of the small rovers, and for sufficiently interesting sites one could be sent off to carry out scientific tasks.

One-kilogram rovers could carry out a variety of tasks. They could relay TV images back to the large rover. A camera and transmitter for such a purpose can be made to weigh less than 50g. Or they could collect small, loose samples of soil, run simple chemical analyses using solid state silicon sensors, or determine soil characteristics by measuring forces on a leg as it swings back and forth in the dust. Adding small rovers to an existing large rover would add little cost but vastly increase scientific payback.

3.2 A Single Small Rover

Another intriguing possibility is to send just one 1kg rover to the Moon, an asteroid, or Mars. If on-surface payload were kept small enough, the vehicle mass needed in low Earth orbit could be low enough to be piggybacked into orbit on a satellite that did not use all of its launch vehicle's payload capacity. Risk would be high with only one rover, but cost would be low,

and since we have had no mobile surface exploration of any body other than the Moon, payoff would be enormous.

3.3 A Herd of Small Rovers

More radically, one can consider replacing a large rover by a collection of small ones. Economies of scale would considerably reduce cost per kilogram of the rovers. Each could be less reliable than the single large rover, since an individual failure would not jeopardize the whole mission. Upon landing, the rovers would disperse over a wide area. Not all need be alike. Having many well-separated rovers that can nevertheless communicate with each other might lead to new and better measurement techniques in some fields.

3.4 Micro Rovers and Manned Missions

A different sort of mission could prepare the way for manned activity. A manned lunar colony might be planned using lunar soil as radiation protection. To avoid high-cost, fast-paced soil moving right after the first manned landing, a troop of say 100 1kg microrovers could be landed many months or even years ahead. They would be totally autonomous and not even communicate with Earth, though a central station might send television pictures of the area to Earth. Like ants, they would mine the soil, perhaps tunneling, or just piling loose collections of soil for later manned use. Because time would not be a pressing factor, the rovers could use very slow techniques. Surely some would expire, but their total number would be large. For power, they would have small solar cells recharging batteries.

4 Emerging Technologies – Gnat Robots

Though small, these rovers can be made even smaller. Most of a mobile robot's weight and bulk consist of motors and batteries, low-tech items that have never experienced the drastic cost and size reductions made in integrated circuits. As many jobs for planetary explorers consist primarily of collecting data, extra bulk and weight in a rover yield no benefit. Often, a robot that begins as a chassis with a few motors and batteries grows in an ascending spiral merely because large motors draw hefty amounts of power, which calls for large batteries, which calls for a sturdier chassis, and so on.

We would move in the opposite direction. By scaling down and using smaller motors, which could make do with tiny power supplies, we could

gain a tremendous advantage. Most of the components we are interested in for our rovers, the computers and sensors, can fit on a small silicon chip. Why not put the entire robot on a chip?

Recently, several groups **[Bart, Lober, Howe, Lang and Schlecht 88]**, **[Fujita and Omodaka 87]**, **[Tai, Fan and Muller 89]**, **[Jacobsen, Price, Wood, Rytting and Rafaelof 89]** and **[Trimmer and Jebens 89]**, have begun to design and build micromotors. Through a technique known as silicon micromachining originally developed for microsensors, it is possible to etch freely movable structures onto silicon wafers. These actuators are primarily electrostatic and often no more than a few hundred microns in diameter. If electronics and actuators were integrated on the same substrate, we could print robots like we print integrated circuits, by the thousands.

We like to call these robots gnat robots **[Flynn 87]** and **[Flynn, Brooks and Tavrow 89]**. Thousands would fit in one payload. Such redundancy would increase the likelihood of acquiring large amounts of data and, in general, of mission success. Also, an integrated robot would have no connectors, where most problems occur.

Putting gnat robots to work requires new perspectives. A good analogy is the relation of parallel processing computers to traditional sequential uniprocessors. Programming an algorithm for a parallel computer requires standing on your head and thinking sideways in comparison to traditional ways, but if the algorithm is well matched to the parallel computer, it can deliver tremendous gains in speed. Lot and lots of very simple processors work together to outperform a goliath uniprocessor.

Similarly, we can match gnat robots to many planetary exploration tasks and solve problems in better, albeit different ways. Gnat robots introduce two new concepts to robotics: massive parallelism and disposability. Each robot is cheap and dumb, but very little is required of a single individual.

Although such tiny robots cannot maintain two-way communication, they could provide one-way low-bandwidth signaling. For instance, a tiny corner reflector could be rotated or uncovered. An orbiter would scan the planetary surface with a laser. With its corner reflector, a gnat could signal its position and its desire to communicate. Each additional corner reflector tuned to a different frequency would enable it to signal another bit.

Gnats could be spread over a large area to signal their positions if and only if they found some condition was met locally. The orbiter would get a map of the occurrence of that condition. Small chemical field effect transistors could be used to detect specific compounds, for instance. On Mars,

gnats could be spread on the wind. Elsewhere, they could disperse by hopping. Solar cells would collect energy and store it in a silicon spring. After a certain compression, a catch would release the spring, and the robot would go flying.

Similarly, millions of seismographic sensors could be distributed over the surface of a planet at regular intervals. They would incorporate micro-accelerometers and -vibration sensors and would communicate very crude tremor magnitudes to the orbiter. Millions of such sensors could be placed all over the planetary surface at the same cost as one more traditional large sensor.[1]

Sensors of many other sorts could be given mobility, for useful autonomous microrovers could be designed, built, and tested quickly. And they would be cheap and reliable enough to be sent out to explore the planets.

5 Conclusion

Exploration of the Earth proceeded by many small spontaneous sorties into the unknown. Small autonomous rovers give us the same opportunity for the rest of the solar system.

Useful autonomous robots can be designed, built and tested on fast timescales. They are cheap because of their size and the ability to mass produce them. They are cheap because they reduce the required launch mass. They are cheap and reliable because they are out of control of a large ground-based mission organization. With imagination and nerve we can let them loose to invade the whole solar system.

6 References

[Angle 89a] "Genghis, A Six Legged Autonomous Walking Robot", Colin M. Angle, *MIT S.B. Thesis in Electrical Engineering and Computer Science*, March.

[Angle 89b] "Attila's Crippled Brother Marvin", Colin M. Angle, *MIT Term Paper in Electrical Engineering and Computer Science*, May.

[1] Recall that we are only talking about a few kilograms of total mass here, so any impact on other planets in terms of pollution, is minimal.

[Bart, Lober, Howe, Lang and Schlect 88] "Design Considerations for Microfabricated Electric Actuators", S.F. Bart, T.A. Lober, R.T. Howe, J.H. Lang and M.F. Schlect, *Sensors and Actuators*, 14(3):269-292, July.

[Brooks 86] "A Robust Layered Control System for a Mobile Robot", Rodney A. Brooks, *IEEE Journal of Robotics and Automation*, RA-2, April, 14–23.

[Brooks 89] "A Robot that Walks: Emergent Behavior from a Carefully Evolved Network", Rodney A. Brooks, *Neural Computation 1:2*, summer.

[Flynn 87] "Gnat Robots (And How They Will Change Robotics)", Anita M. Flynn, *IEEE Micro Robots and Teleoperators Workshop*, Hyannis, MA, November.

[Flynn, Brooks and Tavrow 89] "Anita M. Flynn, Rodney A. Brooks and Lee S. Tavrow", Twilight Zones and Cornerstones: A Gnat Robot Double Feature, *MIT AI Memo 1126*, July.

[Flynn, Brooks, Wells and Barrett 89] "The World's Largest One Cubic Inch Robot", Anita M. Flynn, Rodney A. Brooks, William M. Wells and David S. Barrett, *Proceedings IEEE Micro Electro Mechanical Systems*, Salt Lake City, Utah, February, 98–101.

[Fujita and Omodaka 87] "Electrostatic Actuators for Micromechatronics", H. Fujita and A. Omodaka, *IEEE Micro Robotics and Teleoperators Workshop*, Hyannis, MA, November.

[Jacobsen, Price, Wood, Rytting and Rafaelof 89] "The Wobble Motor: An Electrostatic, Planetary-Armature, Microactuator", S.C. Jacobsen, R.H. Price, J.E. Wood, T.H. Rytting and M. Rafaelof, *Proceedings IEEE Micro Electro Mechanical Systems*, Salt Lake City, UT, 20-22, February.

[Jones 88] "Micro Spacecraft for Space Science", Ross M. Jones, *Workshop, Jet Propulsion Laboratory*, July.

[Tai, Fan and Muller 89] "IC-Processed Micro-Motors: Design, Technology and Testing", Y-C. Tai, L-S. Fan and R.S. Muller, *Proceedings IEEE*

Micro Electro Mechanical Systems, Salt Lake City, UT, 1-6, February.

[Trimmer and Jebens 89] "An Operational Harmonic Electrostatic Motor", W. Trimmer and R. Jebens, *Proceedings IEEE Micro Electro Mechanical Systems*, Salt Lake City, UT, 13-16, February.