Carnegie Mellon

Hans Moravec
Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891
USA
(412) 268-3829 - FAX: (412) 268-5895
hpm@cmu.edu www.ri.cmu.edu/~hpm

November 1, 2002

This letter is to announce commercialization of my work on free range robot navigation, and to solicit interest in participating in the opportunity.

The point of the work is robot perception good enough to let mobile machines free range reliably indoors for long distances without route-specific preparation. This seemingly straightforward functionality has eluded everybody to this day: I've been at it for three decades. Lessons from those 30 years (summarized in the accompanying illustrated sheets), along with a 1000x increase in computer power at 1/1000 the price, have finally put it in reach. It's surely an "enabling technology."

Today the parts cost is over $5,000, for high-end computing and stereoscopic cameras to do dense 3D statistical perception and mapping. For this reason, high-value AGVs (Automatic Guided Vehicles for factories and warehouses) may be the most plausible early application. We are contacting major suppliers of varied AGV navigation systems (AGV Electronics in Sweden, Siemens Dematic in Michigan) to explore that option. For AGVs, a camera-based mapping module, resembling a laser navigation unit, could provide significant advantages over existing guidance methods. Without need of bar code targets or floor embedded wires, magnets or patterns, a mapping AGV could be installed in new locations or rerouted with small effort, perhaps just led through a new route by a worker. Mapping is potentially more accurate and reliable than laser navigation, because a dense 3D sense gives a firmer statistical grip on the surroundings than the three or four points of a laser localization. The rich map opens the possibility of extra functionality, for instance long-range obstacle negotiation, locating movable destinations and large object recognition.

By my numbers, the cost of computation has recently been halving each year, a combination of increasing computer performance and decreasing unit price. Camera costs are falling almost as rapidly. Within five years the parts cost should be well below $1,000, with significantly increased performance. By developing expensive products now for the near term, we achieve a head start in the creation of algorithms for hardware expected to be available in 5 years. Additionally, the development cost will be largely amortized within the early (2 to 3 year) target market. The cost reductions will make it possible to use the approach with smaller, less-expensive, vehicles, where the advantages are much greater and the market orders of magnitude larger. While million-dollar AGV systems must be carefully preplanned, inexpensive small transport vehicles could be used casually. A mapping vehicle might be taught new routes at any time by being led through once, and remember several destinations. It could then function as a junior employee, transporting where and when required on command. The comprehensive 3D sense enables straightforward programming to deal with unexpected route hazards, and to locate destinations that move unpredictably. Easily installed security robots are a related application.

Another is industrial floor cleaning robots—there are a few today, but most require specialist installation and routing. Denning, a now-defunct company I was affiliated with in the 1980s, made the navigation system for one such. Siemens recently began to offer a navigator called Sinas based on a 2D mapping laser from Sick, AG (many robot research projects are using it also) that is more nearly self installing—it doesn't have the potential of 3D mapping and probably won't drop as fast in cost, but it shows that things are heating up. We're contacting a cleaning robot manufacturer (Kärcher in Germany—they make the radar-guided BR 700 Robot floor scrubber) who inquired about collaborating with me a few years ago. A machine able to automatically map a new space could be programmed to locate the boundaries of a room and the major obstacles, and to plan and execute a systematic cleaning trajectory on the spot. A supervisor might be able to shepherd a group of such machines down an industrial corridor, and drop them off one by one in rooms to be cleaned, like human workers, trusting each to do its job automatically and reliably, then directing them to new rooms, and collecting them at the end of the shift. On subsequent nights the machines might repeat the entire routine without any external guidance.

Further cost reductions enabled by the economics of a growing mid scale market for industrial machines should then open even larger possibilities in consumer markets. Kärcher has a prototype consumer robot vacuum cleaner that's very small and can self-charge and empty its dust contents at a docking station. Electrolux is having initial success in Sweden with their less ambitious model and iRobot, Dyson, Hoover and others may not be far behind. But the simple-minded navigation of these machines is a serious impediment. Severe cost constraints and limitations in their technology prevent them from understanding their surroundings or even knowing their location. They move randomly, miss areas and easily become lost or stuck. A mapping machine could keep track of exactly where it had been, what remained, identify and work around navigation hazards and reliably find its docking station, thus work completely autonomously, recharging and emptying itself repeatedly while moving from room to room.

In turn, useful single-purpose machines lead naturally to a truly exciting, enormous market for more advanced utility robots with arms, programmable for multiple tasks. Then robots will truly begin to live up to their promise, as their range of potential application grows to cover almost every physical task. (The rest of the story is for books, not business plan.)

My research on the mapping and navigation problem has been conducted in the open for thirty years. Finally the results are sufficient to support long-term reliable free ranging. It will take a few years more to develop them into a complete demonstration that convinces every casual observer. Unfortunately, doing this last stage of the work in the open would likely compromise the commercial value of the result, and it is not the fastest route to the goal. Various branches of robotics research are experiencing great ferment as functionalities long out of reach rapidly approach practicality. Competition will soon become great and fast-moving. For these reasons I've decided to start the commercial effort now, to build such a prototype in a focused, accelerated effort, followed immediately by a product.

I've linked up with a small company (incorporated as "Botfactory") of 6 individuals with technical backgrounds pursuing this purpose. We hope to raise about $5 million to develop a prototype navigation unit in two years, and a first sellable product a year after

that. The product would be a retrofittable unit something like AGV laser navigation devices (which scan a laser horizontally at about 10 Hz, detecting retroreflective bar code targets on walls and pillars. Three or more targets allow the machine to triangulate its position and orientation), but with a wide-angle stereo camera head scanning 360 degrees at about 1 Hz. Onboard processing of 1000 MIPS or more will permit dense stereoscopic range images to be generated at several Hz and digested into probabilistic 3D maps used for localization, possibly several times per scan, and available for more advanced functions. (We would prefer to use four fixed camera sets instead of a scanner, but cost at present favors a scan. That is likely to change within a few years as inexpensive CMOS cameras advance. The approach would also work well with an imaging rangefinder instead of stereoscopic cameras, but cameras are today more compact and less expensive. Stereoscopy requires more processing, but only about a quarter of the total needed for mapmaking and other functions, and the fraction will decline as the system computer power increases.)

Sincerely

Hans Moravec
Principal Research Scientist
Robotics Institute
Carnegie Mellon University

Chief Technology Officer
Botfactory, Inc.

Animated versions of a lot of the illustrations in the attachments can be found on the first row of the table gotten by clicking **Presentations** on the right of my home page

`http://www.ri.cmu.edu/~hpm`

PDF files of this letter and illustrated attachment sheets can be found in

`http://www.ri.cmu.edu/~hpm/mrl`