Guest Editorial
A Revolution in the Warehouse: A Retrospective on Kiva Systems and the Grand Challenges Ahead

KEN invited me to reflect on the automation and robotics innovations that Kiva Systems developed that allow the company to successfully develop and deploy warehouse systems with hundreds - sometimes thousands - of autonomous mobile robots, and to describe some of the remaining research questions. First, a brief history.[1]

The central idea behind Kiva Systems - to use hundreds of mobile robots to bring inventory to warehouse workers, and thus save them from walking daily marathons to retrieve them - was conceived by Mick Mountz. Convinced that his vision was not only a great business idea, but also technologically viable (albeit barely), Peter Wurman and I took advantage of our upcoming sabbaticals and joined Mick in his quest to revolutionize distribution facilities. Armed with angel investor funds and a well-defined objective, we moved into a small warehouse in the Boston area in January of 2004 to build a team and architect the system that would realize this vision.

By the time Amazon acquired Kiva in May 2012, it was a 300-person company with a long customer list that included Walgreens, Staples, and Saks, and roughly 30 warehouses deployed across Europe and North America.

There were many hurdles that we had to overcome, and only a portion of them were technical. For example, convincing established retailers to embrace a radically different approach to order fulfillment (a competency that is often central to their operation); carefully planning and executing product development with limited cash-flow; maintaining an agile company culture focused on satisfying customer needs and wants that were both varied and changing.

Arguably the biggest challenge we faced was delivering a technical solution that was both economically viable and robust. In 2003, right at the time that Kiva was getting off the ground, Jennifer Carlson and Robin Murphy published their paper "Reliability Analysis of Mobile Robots." (Robotics and Automation, 2003. Proceedings. ICRA'03, IEEE International Conference, 14-19 Sept. 2003.) One of the main conclusions of their research was that the average mean time between failure of mobile robots at the time was 8 hours. For a 1000 mobile robot warehouse operating 24/7, this equates to 3000 robot failures per day - an obviously untenable situation. We thus had to increase mobile robot reliability by many orders of magnitude.

Digital Object Identifier 10.1109/TASE.2012.2214676

[1]Please see also "*Three Engineers, Hundreds of Robots, One Warehouse*," **IEEE Spectrum**, July 2008, for an in-depth history and a discussion of the Kiva business model and technology.

Kiva Systems achieved this by using one of control engineering's most powerful, yet often overlooked, tools: the use of feedback to shift sensitivity. In particular, Kiva uses inexpensive, but very reliable, computation and sensing components to build a highly reliable system from inexpensive electromechanical components and manufacturing processes. Sensitivity is shifted from the subsystems and processes that are costly to make precise, to components whose accuracy and precision can be ensured at a low cost. This allowed us to design extremely dependable robots that can carry loads ranging from a few kilograms to half a ton, with varying mass distributions, at speeds of up to 5 km/h (typical walking speed), in tight quarters, without hitting each other. Furthermore, by having access to large amounts of data, we were able to design model-based adaptation and learning algorithms that improve the performance of the robots over time, while at the same time providing us with important diagnostics and health metrics.

The resulting robust movement infrastructure is leveraged by the multi-agent architecture upon which all high level decisions are made in real time, such as deciding which of the thousands of storage pods should be fetched, which of the hundreds of robots should pick up the pod, and which of the tens of picking stations should be used to fulfill an order. It is in this area where huge benefits can be obtained in a short amount of time: clever algorithms can reduce the number of robots needed in a warehouse by 5 to 10% - a cost savings that is hard to achieve with hardware innovations alone - with only a fraction of the development costs.

The immediate impact of Kiva's success is the adoption of mobile robots in distribution facilities, a multi-billion dollar market. The broader impact, however, is demonstrating that robotics and automation cannot only create new markets, but also revolutionize established ones. Key to our success was that we had first identified a real need based on in-depth market segment knowledge (rather than, like a hammer looking for a nail, trying to find problems to solve with a technology we had invented). We also had the technical skills to create and execute on a product vision, we had the willingness to unequivocally commit to bringing a product to market, and last but not least, we had a little bit of luck and good timing.

I'd also like to propose three broader goals for our research community. The first is to collectively define the standards for - and co-develop - an indoor position system (IPS) that robots can use. Its impact will be broader than that of GPS for outdoor environments (and not just for robots), as most human activity takes place indoors. Imagine what types of systems we could build if we had the IPS equivalent of the low-cost, low-weight,

1545-5955/$31.00 © 2012 IEEE

low power GPS receivers available in the market today, additionally capable of providing position information accurate to centimeters or better. This would have two far reaching effects. The first is obvious: it would allow us to deploy robot systems that we could not do otherwise, for reasons of cost, reliability, and performance. The second is that such systems, with minor or no modification, would be able to provide a tremendous amount of ground truth data to researchers developing robust algorithms and architectures that do not require IPS information.

The second activity our community should contribute to is to develop better design tools for robotics and automation systems. By this, I not only mean better visualization tools and integrated, multi-scale and multi-physics simulation environments, although these are important. What is additionally required is a theory and calculus of information flow, decomposition, and hierarchy that takes into account reliability, robustness, and security constraints, to name but a few attributes. Design tools can be built upon such a theory and thus have a firm footing; they can, for example, provide guidelines for how a system can be decomposed, how much information should flow between components, and how much computation should reside in each subsystem. In addition, they should freely borrow from well-established system design principles such as design for verifiability, reconfigurability, and maintainability. These tools will not be fully automated: design is an art, and design without a human element is doomed to fail. Rather, they will empower robot system designers to quickly explore possible solution spaces, with the rote aspects of design delegated to algorithms.

Lastly, we should strive to provide an environment where post-PhDs can become experts at robot system design, a competency that requires both breadth and depth. This will take some explaining. First, let me come to the defense of monodisciplinary research, both experimental and theoretical. The best time to obtain real depth in an area is during a Ph.D. The acquisition of depth has many benefits, beyond the obvious ones of preparing an individual for a career in academia and being up to date on recent advances in an area. Not only is it a great way to develop problem formulation and problem solving skills, but, simply put, it allows one to determine what is difficult and what is not - crucial information if one seeks to eliminate costly, complex, hard to design and maintain subsystems and algorithms. Unfortunately, there is a trend in academia, especially in robotics and related disciplines, to expose doctoral students to more and more breadth at the expense of achieving depth, such that they risk becoming jacks of all trades, masters of none. So one possibility is to give individuals with depth an opportunity to acquire breadth **after** their Ph.D., and in the process establish the important issues that are in their domain of expertise.

I was able to acquire breadth immediately after my PhD by being a founding member of the Cornell University Systems Engineering program and establishing RoboCup as its flagship, multidisciplinary team project. But in general, how do we do this? One option is with man-on-the-moon type projects. I will dream a little. Imagine a multi-year team effort whose objective is to allow humans to fly like birds. Wingsuit pilots already do this to some limited extent, but they cannot maintain altitude, and achieve at best a glide ratio of about 2.5. What would it take to allow humans to take off and land at will, to gain altitude, even to perch, while preserving the intimacy of wingsuit flight? Advances in aerodynamics, materials, structures, actuators, sensors, energy storage, man-machine interfaces, advanced control algorithms and machine learning, and of course their interplay. Or imagine another team effort whose culmination is a demonstration of a robot athlete (not necessarily anthromoporphic, but preserving the beauty and grace of human athletes) coinciding with the Olympics (perhaps even part of the Olympic opening ceremonies), such as a speed skater, a gymnast, or a sprinter. And my final example: the construction of large-scale dynamic sculptures that are unveiled at venues such as the annual Burning Man festival.

Creating a sustainable infrastructure for such endeavors will require some novel business models, probably centered around the intellectual property that is bound to be created from these activities, in addition to the creations themselves if they can be commissioned. Irrespective of how this is done, the individuals that participate in such a program will possess the breadth and depth required to create the next generation of robot systems. They will be magicians and wizards.

RAFFAELLO D'ANDREA, *Guest Editor*
ETH Zurich
Department of Mechanical and Process Engineering
8092 Zurich, Switzerland
rdandrea@ethz.ch
www.raffaello.name
Co-Founder, Kiva Systems