A Short History of the Australian Centre for Field Robotics (ACFR)

By Hugh Durrant-Whyte

'Fortune favours the Brave', Virgil, Aeneid

Coming to Sydney

In mid 1994 I came across a job advert in the IEEE Spectrum for a Professor of Mechatronics at the University of Sydney. Australia: Big, empty and a long way from anywhere; the perfect place to do field robotics! I showed my wife who more pragmatically saw sun, surf and a sister also living in Sydney. I sent in an application by e-mail. John Kent, then Head of School called me the very next day, I flew out the following week. John took me to the beach at Manly, introduced me to an Alumnus, John Young, who had just started a senior management position at Patrick, and I met the group in mechatronics who I could see had great enthusiasm and potential. The deal was done, I moved to Sydney with my family in June 1995.

The new-look robotics group hit the ground running. It seemed like Australia was ready for robotics. As a group, Eduardo Nebot, Gamini (Dissa) Dissanayake and David Rye, led by the new Professor of Mechatronics Hugh Durrant-Whyte resolved to work closely together on 'field robotics'. We established a very important 'team' ethos that has persisted to this day and underpins the success of the ACFR.

Field robotics is distinguished from more traditional automation research by its focus on large-scale outdoor autonomous systems in applications that are characterised by relatively unstructured, difficult and often hazardous environments. It draws together the most advanced research areas in robotics, including navigation and control, sensing and data fusion, safety and reliability, and planning and logistics. It aims to develop autonomous robotic systems capable of operating in highly challenging applications sectors including mining, construction, cargo handling, agriculture, undersea and aerospace systems. Field robotics is an area notable for focussing the ideas from robotics and mechatronics research in order to solve demanding real-world engineering applications.

In the first year, 1996, the robotics group successfully secured three ARC funded projects. One of these was to result in the key SLAM algorithms which have become arguably the most important single piece of international robotics research in the past 15 years and which established the international reputation of the group at Sydney. A second project was a Linkage grant with Patrick aimed at developing new high-speed crane technology. This also was very important as the first industry partnership for the ACFR and appeared at a critical juncture as Patrick, under the leadership of Chris Corrigan, began investing in new technology with a view to improving efficiency on the Australian waterfront. Chris came to the University to see the scale-model crane that had been developed and to demonstrate the control and operation of the system. He asked what else we could do and he was shown a video of some robotic vehicles. He asked if we could automate a straddle carrier. 'No problem', we said.

The AutoStrad Project

The AutoStrad project (as it was to become) had a great many things going for it. First, it was essentially an opportunity to re-visit some earlier work in port automation undertaken by Hugh with all the knowledge and experience gained. Second, it benefitted from the direct support of senior management in Patrick which proved to be essential when the going got tough. Third, we assembled a hugely talented team, not just the group at Sydney, but two critical people who had previously worked with Hugh at Oxford, Mike Stevens and Phil Avery, and also a great management team from Patrick, Grahame Nelmes and John Young. Our first straddle carrier arrived in Sydney in late 1997 and work began at a site at St Mary's. At this time Patrick also began their dispute with the Unions, which added a certain clandestine spice to the proceedings.

AutoStrad involved much technical innovation, especially in the areas of navigation and systems engineering. Many ideas were the subject of research and thinking by David, Dissa, Salah Sukkarieh and others. The first demonstration to the Patrick board occurred at Port Botany in early 2000. This was followed by Patrick acquiring a new terminal at Fisherman's Island in Brisbane which was scheduled to be fully automated. A new company called Patrick Technology and Systems was established and two leading engineers from ACFR, Ben Rogers and Daniel Pagac, moved to the new company. By 2003, four AutoStrad systems were operating at the new terminal unloading commercial ships and demonstrating that the complete systems concepts worked. By 2005, with the addition of new berths, eighteen AutoStrad systems were in operation. The Brisbane terminal became the most advanced automation system anywhere in the world, a far cry from where Patrick started in 1995. Of course, this description glosses over many traumatic moments, the hard work of a great team of engineers over a sustained period and the commitment and investment made by Patrick senior management.

An interesting spin-out from the work with Patrick was the establishment of another new company, NavTech Engineering, which commercialised the mm-wave radar technology used for the navigation system. This company has gone on to be a major player in this niche area, building and supplying radar technology not just for container terminals, but also for security, debris-monitoring on airport runways, traffic monitoring and many other applications.

Figure 1 Auto-Straddle Carrier

Mining

In parallel with our work with Patrick, in 1996 Hugh was commissioned to undertake a detailed study and cost-benefit analysis of automation for CRA. This was a major formative project as we had never previously been involved with the mining industry. It resulted in a report that had major ramifications, but a decade later. At the conclusion of this work, CRA merged with RTZ to become Rio Tinto. At the time, Rio Tinto saw themselves as “fastest followers” rather than innovators and so work in automation was put firmly off the agenda.

In the meantime, we joined the CRC for Mining Technology and Equipment (CMTE). We undertook a significant project to automate underground mining vehicles and another project with Komatsu in support of their automated mine haul trucks. Neither of these were ultimately successful in a commercial sense, although they introduced and developed a number of new automation concepts. It is interesting to speculate as to why these projects were not successful. The commercial failure of these projects made us reconsider our position with the mining industry and to re-focus on developing systems which would add safety or value to existing manned operations.

This proved to be a good move. Led by Eduardo Nebot, we developed concepts such as HaulCheck, a system using lasers to maintain large haul trucks on a roadway, mm-wave radars, developed by Dr Graham Brooker, for monitoring stopes and ore-passes in underground mining, and proximity monitoring systems for mine operations to ensure safety. All of these projects were developed in partnership with CMTE and its successor CRC Mining, and commercialised through a new company Acumine.

In 2007, at the height of the mining boom, Hugh was approached by Andy Stokes from Rio Tinto, who had also worked with Hugh on the mine automation work for CRA a decade earlier. Rio Tinto now wished to innovate in the area of automation and wished to establish a long-term research and development program aiming to automate surface mining. This was a seismic shift from the stance of any previous endeavour. Rather than focusing on short-term gains and automation of individual platforms, Rio Tinto saw that long term investment was needed and that automation had to be understood as an integrated system.

The work of the ACFR Rio Tinto Centre for Mine Automation (RTCMA) aims to automate surface mining as a process, especially focusing on issues of data fusion, systems architecture and integration of platforms and information into the mining operation. A large group of researchers at Sydney, currently led by Dr Steve Scheding (a University of Sydney Mechatronics graduate), and engineering staff at Rio Tinto are supporting the development. It is a very exciting place to be at this juncture with key new systems being developed and deployed including automated drills, new sensing and data fusion technology as well as the infrastructure for the complete “mine of the future”. At over $35m, the RTCMA is the largest ever investment by a company in university research in Australia.

Figure 2 Eduardo Nebot and Auto-Truck

Defence

One of the key research ideas developed in the ACFR was the idea of decentralised robotics systems; make all components both physically and algorithmically modular; essentially building a robot from a network of interchangeable parts. To do this we developed a number of new ideas in the theory of network sensing and decentralised data fusion (DDF). In the late 1990's sensor networks and (in particular) ideas such as "network-centric warfare" had come of age. BAE Systems was a company with which Hugh had collaborated in the UK and which had supported earlier work in DDF. They came with the idea of demonstrating DDF on a fleet of cooperative unmanned air vehicles (UAVs) to show network-centric operations in a military-relevant context. The resulting ANSER project, initiated in 1999, was a massive undertaking which we seriously underestimated.

The ANSER project was led by Salah Sukkarieh, starting initially as a fresh-minted PhD student. The project aimed to fly a fleet of autonomous UAVs, all sharing and fusing data in a network in real time. At its peak, the project engaged 50 staff spread across three continents. While challenging, the ANSER project produced a long-list of world-first operations including the use of multiple UAVs and demonstration of real-time network-centric data exploitation. Major outcomes were achieved in terms of applications in defence systems for BAE Systems and established a major presence for BAE Systems Australia in the international UAV and autonomous systems domains.

The ANSER project was a tremendously exciting program that really pushed the state-of-the-art in autonomy and has had major ramifications for developments in these areas. ANSER was followed by a number of projects demonstrating combined fusion of information with ground systems, including humans, and work in cooperative control. These programs were supported by both BAE Systems and a number of US defence agencies.

On the ground, the group also developed major outcomes in areas of perception and data fusion for unmanned ground vehicles (UGVs), as part of the DSTO Centre for Autonomous and Uninhabited Vehicles. One interesting outcome of this work was a project to develop a robot for sniper-target training. This is a system in which multiple robots act as targets on which snipers can train. The idea for this came from the Australian Army and was developed in the Centre into a now operational system. A new start-up company, Marathon Robotics, has since commercialised the system and has recently sold its first system to the US Marines. This is one of the great benefits of having a wealth of high-quality researchers who are able to adapt their ideas to an outcome and who are motivated to build new companies on the back of these ideas.

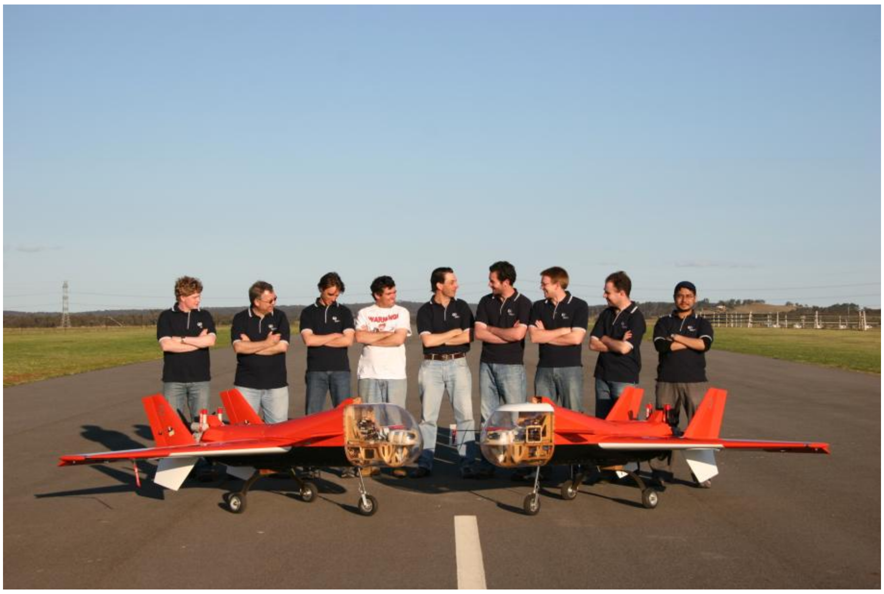

Figure 3 Unmanned Air Vehicles and team

Land, Sea and Air

Following on from ANSER, the group developed an excellent set of research and operational expertise that could be applied to a range of civilian applications, including commercial aviation where Salah has a number of large projects with Qantas Airways. Salah has also been leading the development of UAVs to find and eradicate weeds, with precision, plant-by-plant. The systems have been demonstrated for detecting aquatic and terrestrial weeds, which are estimated to cost the economy up to four billion dollars in lost agriculture and land use. This has been extended to also include the development of UGVs for farm operations, in particular the development of systems that can autonomously estimate fruit yield, offer precision spray of pesticides, and ultimately harvest crops.

Under the leadership of Stefan Williams, the group has substantial capability in marine robotics. Here, many of the same navigation and data fusion principles developed in ground and air systems are applied to the underwater domain. The group runs some of the most sophisticated autonomous underwater vehicles (AUVs) now in operation, able to undertake precise mapping and data interpretation tasks for scientists and industry.

David Rye is also leading a group in robotics and media-art which has exhibited robotic art forms in museums across the globe, and used robotic systems for education in promoting the sciences and engineering.

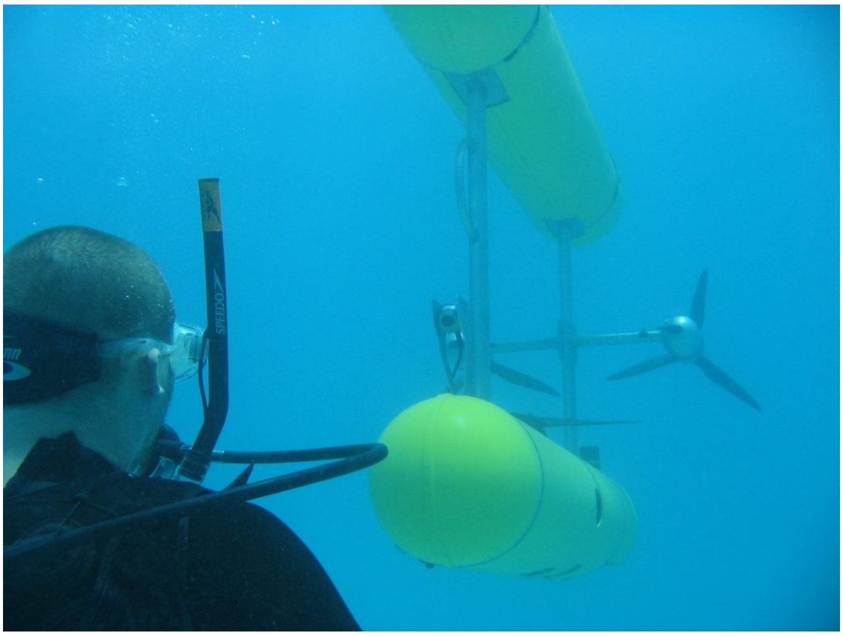

Figure 4 Autonomous Underwater Vehicle

Meanwhile back in the Research World

All of these different commercial activities were being undertaken against a backdrop of major successes in the research sphere. The ACFR grew rapidly both in size and in terms of fundamental research outcomes. Researchers and students were attracted to the ACFR both by our commercial work and by the success we were having in fundamental research. The ACFR become an ARC Key Centre in 1999 and an ARC Centre of Excellence in 2003. Staff and student numbers rose from ten to two hundred and fifty. The Centre grew from almost nothing to being the second largest robotics group in the world. It is interesting also to see the synergistic role played between our fundamental research program and industry-supported work. Each feeds on the other, successively challenging, implementing and learning from each other. An important success in the ACFR has been to marry both fundamental and applied research in this manner. This marriage has also been essential in being able to attract the very best researchers and students to work in the group.

The ACFR benefitted hugely from attracting some brilliant graduate students. Paul Newman started in 1996, working on underwater robotics for his PhD, then went on to a postdoc at MIT and in 2012 was appointed the Professor of Information Engineering at the University of Oxford. Salah Sukkarieh started in 1997 and Stefan Williams in 1998. Both have since gone on to be Professors at the University of Sydney. Steve Clarke (started 1997) went on to found NavTech Engineering, a very successful radar company, Alex Makerenko (2001), Alex Brooks (2002) and Toby

Kaupp (2005) went on to found Marathon Robotics; Eric Nettleton (1998) now leads the Rio Tinto Automation program. These are not isolated stories: many other ACFR graduates have gone on to work at all the leading robotics groups around the world.

Although the ACFR made numerous research breakthroughs, the stand-out was undoubtedly the development of the Simultaneous Localisation and Mapping (SLAM) method. The SLAM method was first pioneered at the ACFR and many of the most successful implementations were undertaken by the ACFR. This work, initially undertaken in the period 1995-2001, is arguably the most important advance in robotics in the past two decades. The papers published by members of the group in SLAM are amongst the most highly cited in robotics world-wide.

Equally, members of the ACFR won a large number of national and international awards for their work including at least a dozen best paper prizes. Hugh Durrant-Whyte was elected to Fellowships of the Academy of Technological Sciences in 2002 (FTSE); the Institute of Electrical and Electronic Engineers in 2005 (FIEEE); the Australian Academy of Science in 2009 (FAA); and the Royal Society in 2010 (FRS). He was also named Engineer of the Year by the Australian Institute of Engineers in 2008 and NSW Scientist of the Year in 2010. The walls of the ACFR were filled with other similar awards won by all members of the group.

Figure 5 ACFR 2005

Into the Future

We have only just begun the journey that will see robotics becoming a major Australian industry. In the near-term, mining, intelligent transport systems, and agriculture will be key drivers in the development of robotics and automation technology with some of the largest programs now underway in Australia. In the further future, one of the next big applications for robotics will be in the stewardship of our environment, both in the marine domain and for terrestrial ecosystems.

Australia has the skills and opportunity to lead the way in these endeavours. Yet further, there are great opportunities for using robotics in remote health care, infrastructure maintenance and management of disasters including bush fires.

In the past fifteen years Australia has come to lead the world in the development and application of robotics in large-scale outdoor field applications. Robotics and autonomous systems will be one of the most important and transformational technologies in the future of this country. It is an exciting future with enormous opportunities.

Figure 6 Hugh and Robots