David A. Patterson

Robots in the Desert A Research Parable for Our Times

DARPA s grand challenge spurred advances in the application of AI and, ironically, acted as a referendum on its de-emphasis on academic research.

To inspire progress in robots, the U.S. government's Department of Defense Advanced Research Project Agency (DARPA) sponsors a grand challenge for autonomous vehicles to complete a 100-plus mile, off-road course in the Mojave Desert (www.grandchallenge.org). In 2004, the first year of the event, Sandstorm from Carnegie Mellon's Red Team traveled the furthest, but that was just 7.36 miles, so no one won the $1M prize.

The publicity surrounding the 2004 competition, plus the doubling of the award to $2M, led to 195 entrants from 36 states and four countries for this year's event. DARPA whittled the list down to 43 in July using a series of two-mile obstacle-course trials.

A final time trial was held at the California Speedway Sept. 27–Oct. 5 to reduce the competition to 23 finalists announced Oct. 6. CMU's Red Team Too edged the Stanford Racing Team for first place in this trial, followed by Sandstorm from CMU's original Red Team. These university entries leverage the latest advances in statistical and machine learning and in vision to automate their vehicles. In particular, the leader of the Stanford Racing Team, Sebastian Thrun, is also one of the developers of FastSLAM, a factored solution to the simulation localization and mapping problem, which was criti-

Stanley, the autonomous ground vehicle from Stanford's Racing Team, won first place and the $2M prize in DARPA's 2005 Grand Challenge competition.

PHOTO COURTESY OF DARPA

President s Letter

If we could reduce the costs of automobile accidents by $30B per year in the U.S., a single year's savings would exceed the government's investment in academic computer science research for the last 50 years.

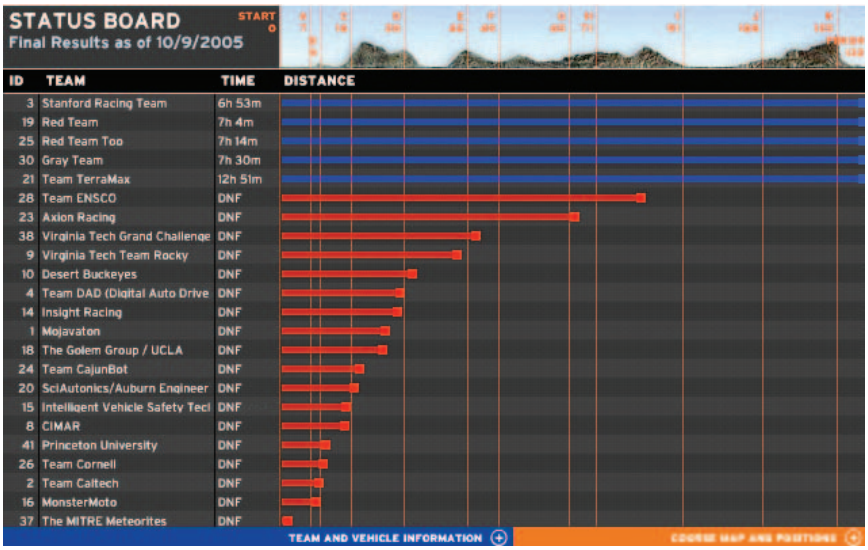

Final results of the 23 bots participating in the last leg of the DARPA race.

cal to Stanley's success [1].

CMU's Red Team Too entry, named H1ghlander, is a 1999 H1 Hummer with a turbocharged diesel engine, traction control, and embedded drive-by-wire technology. Sandstorm, from the CMU Red Team, started life as a 1986 HMMWV. Stanford's entry, Stanley, is a stock Volkswagen Touareg R5 modified with full-body skid plates and a drive-by-wire system controlled by seven Pentium M computers. Other bots that performed well were Team Terramax, a combined effort from Oshkosh Truck and Collins Radio, and the Gray Insurance Ford Hybrid SUV.

After the trials, the finalists had one day to repair their robots. The race began Oct. 8, with each entrant given a CD defining a 132-mile, obstacle-laden course; they were positioned according to their lap times in the final time trial. The results of the finals are listed in the table here.

Of the 23 entrants, 18 were disabled, but all but one exceeded the maximum distance from 2004. Five finished the course; although only four did it under the 10-hour limit needed to win the award. Stanford's Racing Team took home the $2M prize. Given last year's results, I was certainly surprised that many teams did so well. Many AI experts consider this milestone a more significant achievement than when IBM's Deep Blue computer beat the reigning world chess champion in 1997.

The success of the challenge could help accelerate the deployment of driver-assistance devices, potentially reducing the loss of life and property from traffic accidents. In just the U.S. alone, traffic accidents cost 40,000 lives and $55B. Worldwide, nearly 1.2 million die each year, 20 to 50 million more are injured or disabled, and more than $500B is spent annually on traffic accidents [4]. Indeed, ACM A.M. Turing Award winner Butler Lampson suggests a worthy grand challenge for our field is to create the technology to eliminate all automobile accidents. Even if the technology prevented *half* the accidents, it would save $25B to $30B and prevent tens of thousands of deaths and injuries in the U.S.

Upon further reflection, it's clear the DARPA challenge tested the importance of academics participating in a research mission in an open competition against all comers. The agency recently shifted the emphasis toward more classified research and research with short go/no-go deadlines that discour-

PHOTO COURTESY OF DARPA

age participation from universities. The result is a 50% decline of DARPA funding of contracts with academics as principal investigators [3].

University teams were 30% to 40% (depending if you count joint industry-academic teams) of the 43 semifinalists and 40% to 50% of the 23 finalists. Not only did Stanford University's team win first place, university teams swept the top three places, and only one other entrant completed the course within the 10-hour window. These results jibe with the NRC report illustrating the important role of academic research in helping create 17 multibillion-dollar industries [2]. Whether you look to history or to this competition, what conclusion can you reach other than academic research is vital to technological advances in IT and to the creation of new IT industries?

If we could reduce the costs of automobile accidents by $30B per year in the U.S., a single year's savings would exceed the government's investment in academic computer science research for the last 50 years. Given the tremendous progress over the last year as a result of this competition, doesn't it seem plausible that if we invested, say, 5% of the potential savings in computer science research for five to seven years, that we could at least halve traffic accidents down the road? Wouldn't almost all drivers urge such an investment? If so, what are we waiting for? C

References

1. Montemerlo, M., Thrun, S., Koller, D., and Wegbreit, B. FastSLAM: A factored solution to the simulation localization and mapping problem. In *Proceedings of the AAAI National Conference on Artificial Intelligence,* 2002; robots.stanford.edu/papers/montemerlo.fastslam-tr.html.
2. National Research Council. *Rising Above The Gathering Storm: Energizing and Employing America for a Brighter Economic Future.* National Academies Press, Washington D.C.
3. Patterson, D.A. President's letter: The state of funding for new initiatives in computer science and engineering. *Commun. ACM 48,* 4 (Apr. 2005), 21–25.
4. Peden, M., Scurfield, R., Sleet, D., et al. *World Report on Road Traffic Injury Prevention.* World Health Organization, Geneva, 2004.

David A. Patterson (pattrsn@eecs.berkeley.edu) is president of ACM and the Pardee Professor of Computer Science at the University of California at Berkeley.

The 2006 Editorial Calendar for *Communications*

JANUARY
Personal Information Management

FEBRUARY
Computer Forensics

MARCH
Self-Managed Systems and Services

APRIL
Supporting Exploratory Search

MAY
Two Decades of Language/Action

JUNE
Hacking

JULY
Services Sciences

AUGUST
Music Information Retrieval

SEPTEMBER
Security in Highly Dynamic Systems

OCTOBER
Flexible and Distributed Software Development Processes

NOVEMBER
Entertainment Networking

DECEMBER
Software Product Line