

# Computer Science as an experimental science

Testing new concepts & algorithms  
on the DAS-4 Supercomputer

In computer science research, good experimentation facilities for testing new computer system concepts and new algorithms are very important. Dick Epema (Parallel and Distributed Systems) explains the structure and importance of the Distributed ASCI Supercomputer. He is a member of the DAS-4 project Steering Committee.



More than just theory, computer science is also an experimental science. This means that it needs experimental equipment, just like physics or astronomy. The Distributed ASCI Supercomputer (DAS) has played an important role in the Netherlands to make this clear to funding agencies such as NWO. This unique facility, now in its fourth generation, is the only distributed computer system in the world purely meant for computer science research.

## *History*

The real test of new system concepts and new parallel and distributed algorithms is actual deployment. It is pointless to create new massively parallel algorithms for high-resolution weather prediction or to propose energy-efficient scheduling algorithms for data centers if they cannot be tested and analysed on state-of-the-art computer systems. The history of the DAS systems goes back to 1997, when the computer science research school ASCI (Advanced School for Computing and Imaging) obtained funding for the first generation DAS from NWO. Roughly speaking, the DAS is replaced every four to five years. Whereas the first two generations had at least some claim to the term


'supercomputer', this is difficult to maintain for the DAS-4: the fastest supercomputer in the world now has about 550,000 cores versus the 1,600 of the DAS-4. The DAS-4 has only twice as many processors as the DAS-1, but 200 times as much main memory and a wide-area network that is 1500 times as fast.

## *Specs*

The DAS-4 consists of six compute clusters, interconnected by a wide area network. There is one in Delft, one in Leiden, one at ASTRON (Netherlands Institute for Radio Astronomy) and three in Amsterdam. It has a total of 198 nodes with two quad-core 2.4 GHz processors each, and a total of close to 200 TB of storage. The main features of the DAS-4 are the so-called accelerators such as Graphical Processing Units (GPUs) and the 10 Gb/s Infiniband network. GPUs were originally designed for graphical applications, but they can also be used very effectively for many compute-intensive applications.

## *Experimentation*

The different generations of the DAS systems have been instrumental to many dozens of PhD theses in ASCI. Grid resource management systems,



interior-point methods, and web applications are but a few of the topics for which the DAS has been used. The Parallel and Distributed Systems (PDS) group has built the KOALA grid scheduler for systems like the DAS. Whereas the original purpose of KOALA was to provide co-allocation, that is, to allocate processors in multiple clusters to single parallel applications, it can now handle many different types of applications. When testing the scheduling policies of KOALA in the DAS, it is very important to have identical experimental conditions when performing different experiments, just like in physics. Here, this means having exactly the same system configuration available, and having the same system workload. Even in a relatively small and centrally managed system like the DAS, this is difficult to achieve: processors and

network links may fail, and the load imposed on the system by the ordinary users may vary considerably over time.

years to such issues as energy efficiency, data-intensive applications, and clouds. Some of the projects that will

*“DAS-4: the only distributed computer system purely meant for computer science research”*

The PDS group has also used the DAS for measurements of the BitTorrent peer-to-peer file-sharing system. For this purpose, many simultaneous internet connections were required for tracking the download activities of tens of thousands of peers installed at user PCs worldwide.

#### *Present & future*

The focus in parallel and distributed computing has shifted over the last few

use the DAS-4, such as the COMMIT project 'Infrastructure Virtualization for e-Science' (IV-e), look at these topics, and software to make the DAS-4 behave as a cloud has already been installed. However, the DAS-4 offers a very open environment, and researchers in computer science in ASCI can use the DAS-4 for any purpose they wish.

DAS-4 website: [www.cs.vu.nl/das4](http://www.cs.vu.nl/das4)

Als experimentele wetenschap heeft de Informatica experimentele apparatuur nodig. De Distributed Supercomputer van ASCI (DAS-4, '4e generatie') is het enige gedistribueerde computersysteem ter wereld dat uitsluitend voor informatica-onderzoek wordt gebruikt. Eén van de zes rekenclusters staat in Delft. Op de DAS worden nieuwe systeemconcepten en nieuwe parallele en gedistribueerde algoritmen getest. Promovendi van de sectie Parallel and Distributed Systems hebben hem gebruikt bij onderzoek naar o.a. grid resource management systemen en peer-to-peer systemen. Volgens Epema is het lastig experimentele condities als systeemconfiguratie en werklust gelijk te houden. De DAS zal ook een rol spelen bij onderzoek op gebieden als energie-efficiency, data-intensieve applicaties en clouds. Door de zeer open omgeving kan hij echter voor ieder ander gewenst doel in informaticaonderzoek worden gebruikt. ■