

Using Java for distributed computing in the Gaia satellite data processing

William O'Mullane · Xavier Luri · Paul Parsons ·
Uwe Lammers · John Hoar · Jose Hernandez

Received: 27 October 2010 / Accepted: 4 July 2011 / Published online: 29 July 2011
© Springer Science+Business Media B.V. 2011

Abstract In recent years Java has matured to a stable easy-to-use language with the flexibility of an interpreter (for reflection etc.) but the performance and type checking of a compiled language. When we started using Java for astronomical applications around 1999 they were the first of their kind in astronomy. Now a great deal of astronomy software is written in Java as are many business applications. We discuss the current environment and trends concerning the language and present an actual example of scientific use of Java for high-performance distributed computing: ESA's mission Gaia. The Gaia scanning satellite will perform a galactic census of about 1,000 million objects in our galaxy. The Gaia community has chosen to write its processing software in Java. We explore the manifold reasons for choosing Java for this large science collaboration. Gaia processing is numerically complex but highly distributable, some parts being embarrassingly parallel. We describe the Gaia processing architecture and its realisation in Java. We delve into the astrometric solution which is the most advanced and most complex part of the processing. The Gaia simulator is also written in Java and is the most mature code in the system. This has been successfully running since about 2005 on the

W. O'Mullane (✉) · U. Lammers · J. Hoar · J. Hernandez
Space Robotic Exploration, European Space Astronomy Centre, Madrid, Spain
e-mail: womullan@sciops.esa.int

X. Luri
ICC-UB/IEEC, Departament d'Astronomia i Meteorologia,
Universitat of Barcelona, Barcelona, Spain

P. Parsons
The Server Labs, Madrid, Spain