Astronomical Data Analysis Software and Systems XII
ASP Conference Series, Vol. 295, 2003
H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, eds.

eSTAR: Building an Observational GRID

Alasdair Allan, Tim Naylor

School of Physics, University of Exeter, Stocker Road, Exeter, EX4 4QL, U.K.

Iain Steele, Dave Carter, Jason Etherton, Chris Motteram

Astrophysics Research Institute, Liverpool John Moores University, Twelve Quays House, Egerton Wharf, Birkenhead CH41 1LD, U.K.

Abstract. The eSTAR Project[1] is a programme to build a prototype robotic telescope network to design and test the infrastructure and software which could be used in larger scale projects. The network consists of a number of autonomous telescopes, and associated rapid data reduction pipelines, connected together using Globus middleware. Intelligent agents carry out resource discovery, submit observing requests, and analyse the reduced data returned by the telescope nodes. The agents are capable of carrying out data mining and cross-correlation tasks using online catalogues and databases and, if necessary, requesting follow-up observations from the telescope nodes. We discuss the design and implications of the eSTAR software and its implications with respect to the GRID.

1. Just Imagine...

Imagine a system which has unified access to archived data, to telescopes and to bibliographic data. In addition the system has intelligent agents (IAs) which can interpret the results.

In this paper I hope to persuade you that such a system, with a seamless interface between telescopes and databases, indeed making telescopes look like databases and visa versa, will bring enormous benefits.

2. The Observational Grid

There are two fundamental ideas behind eSTAR which make it a unique project. The first is to treat both telescopes and databases in as similar a fashion as possible, both being made available as a resource on the 'Observational Grid'. The second is that the main user of that grid should not be humans making observing requests, but should be intelligent agents.

[1] `http://www.estar.org.uk/`

© Copyright 2003 Astronomical Society of the Pacific. All rights reserved.

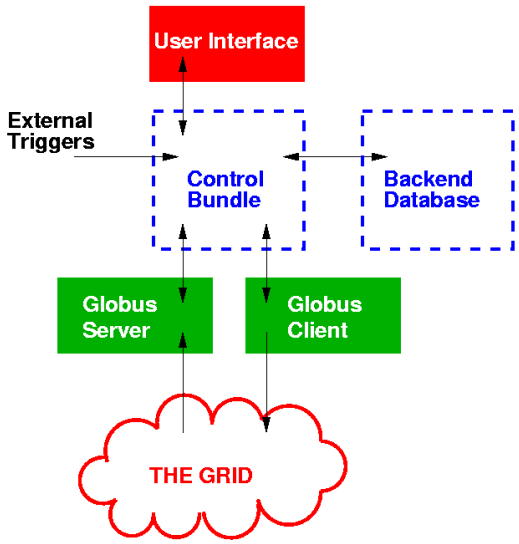

Figure 1. Intelligent Agent architecture diagram. A central control bundle handles both user input and input from external triggers, such as γ-ray burst alerts. The Globus middleware layer is separated from the control bundle to allow easy substitution, e.g., with SOAP.

The design is analogous to computational grids, with no overall supervisor, giving the system scalability with multiple agents talking to multiple nodes.

Intelligent agents submit requests to nodes on the network rather than commands. Each telescope has its own scheduler, which can talk to the agents using the Robotic Telescope Markup Language[2] (RTML) transported over Globus[3] middleware.

A typical sequence begins with the IA carrying out resource discovery, using LDAP to characterise each telescope node and associated instrument suite. Next each telescope is asked by the agent whether it could carry out a particular observation, each node returning a weighted score which encodes the predicted quality of the observation.

The IA will then select the node it wishes to carry out the observation, which will usually be the highest scoring node, and requests that the observation be put in the nodes observing queue by its scheduler.

Once the observation is complete the raw and reduced data are made available to the IA via the grid.

3. Intelligent Agents

One can view the eSTAR network as a unified information grid, within which intelligent agents live.

[2] `http://alpha.uni-sw.gwdg.de/~hessman/RTML/`

[3] `http://www.Globus.org/`

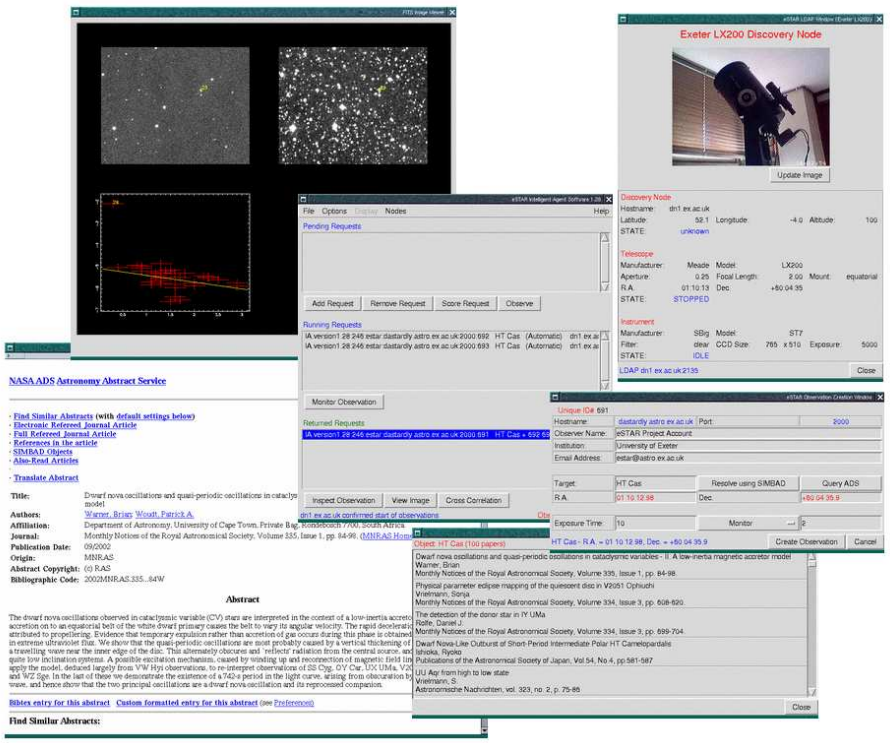

Figure 2. Intelligent Agent prototype developed to monitor Dwarf Novae. The main observation tracker window is shown centre field, with the node window in the upper right. The results of a real time cross correlation of a returned field with the USNO-A2 catalogue are shown in the upper left, while results from datamining SIMBAD and ADS for information about the possible variable the agent has discovered are displayed in the lower part of the figure.

The current prototype agent, see Figure 1, cross correlates the point source catalogue returned by the telescope node with USNO-A2 (in real time) in an attempt to find candidate variable stars, datamining SIMBAD for possible known variables to match the candidate stars.

If the candidate turns out to be a dwarf nova of interest, then the IA will request further followup observations from the network of telescope nodes.

We have been pleasantly surprised by the speed of the cross-correlation stage, which normally completes in 4–10 seconds, including recovering data from three different web sites, indicating that the IA would not be the limiting factor in a system which would be expected to react rapidly to external triggers.

While the prototype agent does lack complexity, it was deliberately developed in a modular fashion to maximise code reuse, and could be trivially re-engineered to handle a very different science goal. For instance, to listen for and response to external triggers such as γ-ray burst alerts.

Indeed, our vision is that agents should be developed by astronomers to address their own science problems, perhaps using some 'agent builder toolkit.'

4. Discovery Nodes

For the prototype we have made use of off the shelf hardware, Meade LX200 and ETX telescopes with SBIG and Apogee CCD cameras, as telescope nodes.

However in general a node is any telescope, or archive, which delivers astronomical data or other resources using a uniform interface. From the point of view of an intelligent agent (or indeed a astronomer) there is little difference between a telescope and a database except the timestamps on the data.

5. Middleware

Currently the project makes use of Globus IO to transport RTML documents between the agents and the discovery nodes, with SSL encryption providing a secure data stream.

Additionally, a new version of RTML (Pennypacker et al. 2002) has been developed by the project, with our improvements being rolled back into the ongoing development of the protocol.

6. Where Now?

We are currently developing the software to make use of emerging web services, and re-engineering the agent to support the next generation Globus Toolkit using Open Grid Service Architecture (OGSA), within which SOAP and WSDL will provide natural language neutral interoperability between the components of the system.

Additionally we are looking to deploy the system on research class telescopes such as the Liverpool Telescope and the United Kingdom Infra-Red Telescope (UKIRT).

7. Summary

We view it as vitally important that federated databases and automated telescopes should share a common interface, as there is little, if any, fundamental differences between them other than the time it takes to return the data and the date stamp.

However if a system similar to this is to become widespread, the idea of requesting observations is equally important. This allows telescopes to join the network and retain their independence, with the intelligent agent contributing, perhaps, only a small fraction of the observations to the telescope scheduler.

Furthermore, telescopes can and must be able to reject requests from an IA if, for example, the astronomer to whom the IA belongs has no time allocated on the telescope.

References

Pennypacker C. et al. 2002, A&A, in press