Astronomical Data Analysis Software and Systems VI
ASP Conference Series, Vol. 125, 1997
Gareth Hunt and H. E. Payne, eds.

A Configuration Control and Software Management System for Distributed Multiplatform Software Development

E. Huygen,[1] B. Vandenbussche, and G. Bex

Katholieke Universiteit Leuven, Instituut voor Sterrenkunde, Celestijnenlaan 200B, B-3001 Heverlee.

P. R. Roelfsema, D. R. Boxhoorn, and N. J. M. Sym

Stichting Ruimteonderzoek Nederland, SRON, The Netherlands.

Abstract. A system to efficiently control and manage the software development of a large project is presented. The overall design, implementation and use of the configuration control system (CoCo) and the calibration database (CalDB) used for off-line software development of the Short Wavelength Spectrometer (SWS) in the Infrared Space Observatory (*ISO*) is described.

1. Introduction

CoCo is a software configuration control system especially designed for the development of the *ISO*-SWS interactive analysis data reduction software (IA^3). Since this software is being developed by a team spread out over several institutes, on machines with completely different architectures, we needed a configuration control system that was easily portable and applied a client/server technology.

The CoCo system is capable of efficiently controlling the development of the IA^3 software. The server runs at the satellite tracking station of the European Space Agency (ESA) in VILSPA, Spain, and controls the main repository of what we call CoCo objects, i.e., all source code (FORTRAN, C, IDL, Perl), calibration files, comparison spectra and photometry, documentation, etc. Changes to these objects are saved as new generations and can be requested individually by any CoCo client program.

The calibration database CalDB is constantly and automatically fed with information about all these changes and dependencies of calibration files, algorithms, and data. This information is accessible through a WWW interface.

The software management for IA^3 is greatly simplified by a Software Problem and Modification Report (SPR) system which is built into CoCo. These reports are handled as special objects by CoCo and special commands are implemented to work with them.

---

[1] E-mail: rik@ster.kuleuven.ac.be

2. Using CoCo

CoCo uses standard configuration control commands to work with its objects, i.e., create/reserve/insert/extract, etc. All changes to objects must be accompanied by a software problem report (SPR) which fully describes the problem, the proposed solution, and the actual change to the object. Changing objects with CoCo requires a mandatory sequence of CoCo commands: a user raises a SPR, software management assigns an action to a developer, the developer reserves the object (which is automatically extracted to his/her local machine), and changes are tested and inserted back into the CoCo system. This will automatically create a new generation for the affected object. The reservation mechanism avoids developers making conflicting versions of the same code, and the SPR system streamlines the reporting of problems from the users to the developers.

All developers and users of the IA^3 data reduction software have individual access and change permission on the CoCo server, which can be granted or revoked by the configuration control manager.

At the consortium institutes, an automatic process starts up a CoCo client every night to update the local system. This means that every site has a consistent version of the software system every morning, including the changes made at any institute the day before.

3. The CoCo Server Implementation

The CoCo server is a dæmon program that accepts commands from any CoCo client by electronic mail. This protocol was chosen because of its simplicity, and because it was, at that time, the only robust way to communicate between the different institutes and the different computer architectures. It is now, for security reasons, the only possible way to communicate with the server, located at VILSPA. Requests from CoCo clients are handled as they appear in the mail queue.

The server is organized in several different command modules (see Figure 1). The command manager parses the requests sent in by a client, and delegates the commands to the proper command modules. Configuration control commands are those which create, extract, reserve, or insert objects in the main repository. They communicate with the Code Management System (CMS) running on a VMS machine and simultaneously feed the calibration database. Info commands provide the user with status information about the system and the individual objects. Administrative commands are used to maintain and debug the CoCo server.

4. The CoCo Client Implementation

The CoCo client is a stand-alone program that accepts user input via a command line interface or a batch file. The user request is converted into a syntactically correct CoCo command and sent to the command manager. The command manager distributes these commands to the proper command modules.

Local commands are handled by the client program itself. These commands do not require interaction with the server, and consist mainly of information commands and commands to check the structure and contents of CoCo objects.

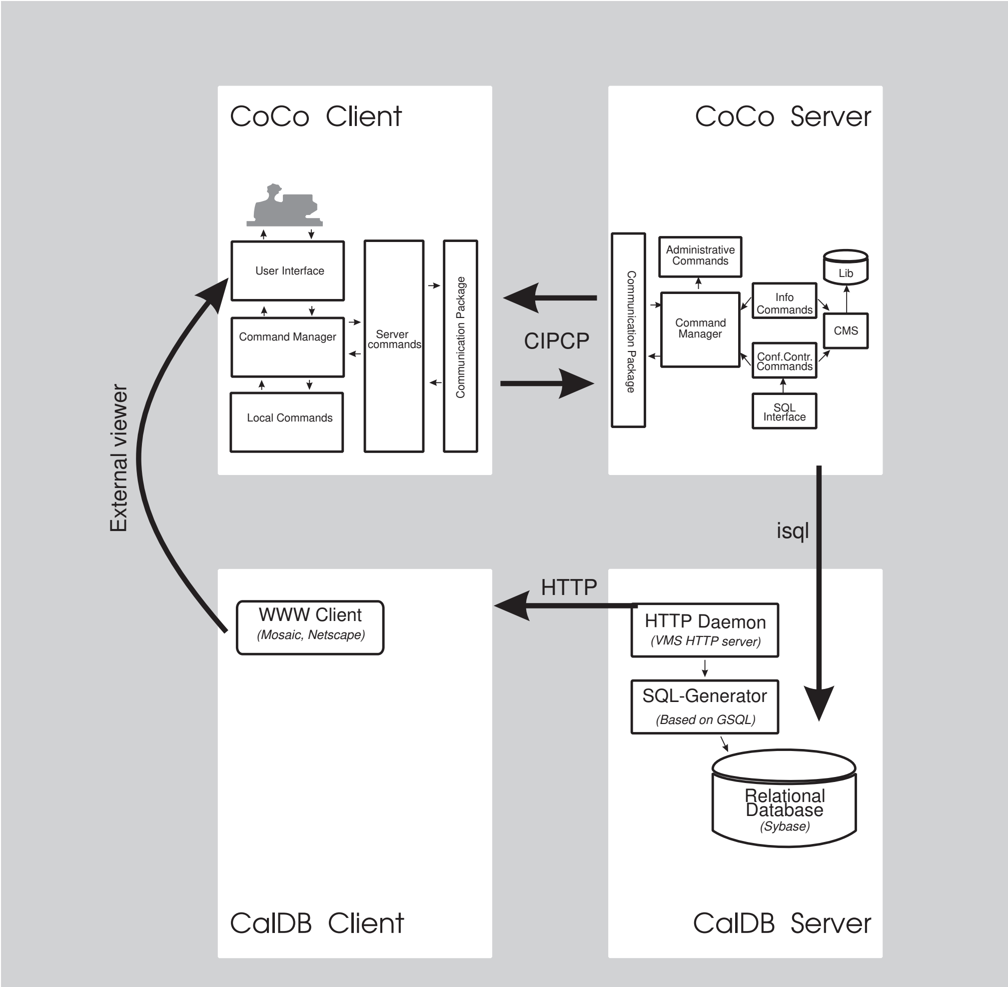

Figure 1. This picture shows the essential design of the complete configuration control and database system.

All commands that require server interaction are sent to the CoCo server. The client program then waits for a response from the server. If this response requires further action from the client, this is done before returning control to the user.

Apart from communication with the server, the client is also responsible for maintaining the local copy of the IA^3 software and calibration data. This is accomplished by an automatic full-update mechanism, whereby clients will send a request to the server to retrieve a list of all the CoCo objects with their last modification time. From this list, the client can decide which objects are out-of-date on the local system, and send the server an update request for each such object.

5. The Calibration Database CalDB

Many different types of objects are managed by CoCo. These objects have not only their own properties (e.g., creation date, scientist who created them, reasons why they changed, purpose, ...), the objects are also related to each other (e.g., calibration files are derived with special algorithms, the comparison spectra are used to derive calibration files). The properties of all the objects in CoCo and the relationships between them are stored in the Calibration Database, CalDB, a relational database implemented in Sybase.

Every creation or change of the state of an object on the CoCo server is automatically stored in CalDB. The interface between the CoCo server and Sybase is isql, the interactive sql interface of Sybase which can also handle batch files of sql statements. This approach makes it easy to switch to any other relational database engine with an interactive sql interface that can handle sql input files.

Users can query the database via the WWW. The selection criteria and the output format are specified by filling out a form and submitting it to the HTTP server on the CalDB host. The script translating the form to sql is based on the NCSA GSQL gateway.

If the result of the query contains the name of CoCo generations, these will be hyperlinks. Following such a hyperlink will start a CoCo client with the external "viewer" mechanism. The CoCo client will then extract the generation from the CoCo server. The CoCo server does not depend on CalDB. CalDB provides the user with extensive query possibilities and extra information about the objects and generations in CoCo.

6. Conclusion

This article briefly describes the CoCo software control and management system designed and implemented for the development of the *ISO*-SWS interactive analysis and data reduction software. The CoCo system has now been used extensively for almost two years. World wide, ten clients with a total of twenty-six developers are currently registered at the server. Since June 1995, when the server was transfered to the ESA satellite tracking station at VILSPA, more than 60,000 client requests were handled by the server, which is more than 100 transactions per day. For more information on CoCo consult the CoCo Homepage.[2]

Acknowledgments. We like to thank all those who supported us during development and maintenance of this system, especially the SWS Instrument Dedicated Team (SIDT) that constantly uses the system and continuously gives feedback on its use and possible improvements.

[2]http://www.ster.kuleuven.ac.be/coco