CERN - Computing and Networks Division
CN/91/5

CHEOPS: Really Using a Satellite

J.Altaber, B.Carpenter, S.Cannon, D.Davids, C.Isnard,
J.M.Jouanigot, S.Olofsson, B.Segal, N.Stheneur (CERN),
P.L.Forsström (CSC Helsinki), M.Nordberg (SEFT Helsinki),
S.Simonen, M.Voipio (HUT/IPP Helsinki),
G.Barreira (LIP Lisbon), P.Veiga (INESC Lisbon),
S.Dedoussis, C.Touramanis (Univ. of Salonika),
O.Koudelka (Tech. Univ. Graz), K.Sweeney (ESA/Tecnodata Italia), M.Hine.

30 April 1991

Abstract

CHEOPS is a collaboration between CERN, LIP and INESC (Portugal), SEFT (Helsinki) and four Greek institutes (of which only one will be connected via satellite). CIEMAT (Madrid) and CRS4 (Sardinia) are joining the collaboration. Its objective is a sustained experiment in the use of ESA's OLYMPUS satellite to demonstrate that the properties of satellite transmission can be used to distribute massive quantities of scientific data for remote analysis.

The authors are all active members of the CHEOPS collaboration and are responsible for computing or networking activities in their various institutes. The collaboration may be contacted via B.Carpenter, CN Division, CERN, CH–1211 Geneva 23
(fax +41 22 767 7155, electronic mail brian@priam.cern.ch).

Keywords: satellite, XTP, OLYMPUS.

Paper submitted to 2nd Joint European Networking Conference, Blois, France, May 1991.

Chapter 1

INTRODUCTION

CHEOPS is a collaboration between CERN, LIP and INESC (Portugal), SEFT (Helsinki) and four Greek institutes (of which only one will be connected via satellite). CIEMAT (Madrid) and CRS4 (Sardinia) are joining the collaboration. Its objective is a sustained experiment in the use of ESA's OLYMPUS satellite to demonstrate that the properties of satellite transmission can be used to distribute massive quantities of scientific data for remote analysis.

Previous experiments (STELLA [1], HELIOS [2]) have shown that achieving high-speed point-to-point bit stream transmission by satellite is a well understood problem. CHEOPS aims to benefit from this experience by concentrating on three aspects: efficiency, access to geographically remote regions, and integration of the application of the satellite into the general computing environment.

The CHEOPS earth stations will access OLYMPUS mainly on an overnight schedule when satellite time will be available for our experiment. At each site there will be a UNIX-based server, integrated in a suitable way into the LAN and computing environment. During dayshift, users (i.e. physicists) will make interactive requests via terrestrial links to the CHEOPS system to transfer physics datasets, and the datasets will be staged onto the CHEOPS disc space. At night, connections will be established from CERN to each remote site in turn, and the datasets will be spooled automatically from source to destination. As far as possible, the process will require no operator intervention. Initially standard TCP/IP protocols will be used, with multiple parallel transactions to improve efficiency. Later, optimised protocols will be used in an attempt to maximise throughput.

A successful experiment on the three fronts of geography, efficiency, and integration would clear the way to possible use of commercial services, assuming economic tariffs and a liberal regulatory environment.

Chapter 2

THE CHEOPS APPLICATION

CERN, the European Laboratory for Particle Physics, acts as a source, sink, and distribution centre for many terabytes of data per year. The datasets concerned, whose typical size is 200 Mbytes, are produced either by preliminary processing of raw data from the massive experiments at CERN, or by processing of data from Monte Carlo physics simulations carried out in various European computing centres. In addition, we expect to see large datasets produced by CAD tools during the design of future experiments for the Large Hadron Collider proposed to be built at CERN by the end of the century. Physicists, or designers, need to move hundreds of these data sets in an organised and convenient manner between CERN and the various collaborating academic sites in Europe.

Existing academic data networks, including those dedicated to particle physics, are clearly inadequate to this task. Today, with rare exceptions, bulk data are transported by sending copies of magnetic tape cartridges by air. The process of copying, labelling, shipping, customs clearance, etc., is clumsy and takes several days. It creates a manual workload in keeping track of the tapes, entering their details in multiple databases, and so on. The ability to transport bulk data automatically and economically between sites would greatly simplify these logistic and administrative problems.

A related question is that of how users "see" their mass of data. CERN has developed a distributed data management system called FATMEN [3], based on the IEEE Mass Storage Reference Model, to support a distributed dataset catalogue, and uniform access to datasets for users and for software.

The CHEOPS application can therefore be defined as the construction of an automat, which will accept requests via FATMEN to move datasets between CERN and other sites, and will execute these requests when network capacity (in the first instance, OLYMPUS) is available. This automat has the job of hiding the entire network process from the user, and reporting back only in the case of success or irrecoverable failure. Constructing this automat, rather than implementing basic transmission, is the main challenge of CHEOPS. Nevertheless, the remainder of this paper takes a bottom–up approach, but starting with the satellite some 36,000 km above us.

Chapter 3

SATELLITE TRANSMISSION

The OLYMPUS experimental communications satellite was successfully launched in mid 1990 and is working well despite a partial failure in its solar arrays early in 1991. The satellite itself indeed has built–in redundancy in case of partial failure, but no backup satellite is available in the case of total failure. It is in geo–stationary orbit at 19 W with a technical lifetime of about 5 years. It carries several payloads including the "Specialised Services" payload, which is the only one considered here in detail. This operates in the 12/14 GHz Ku–band with a steerable cluster of five spot beams operating at various frequencies and polarisations. There are four transponders which can be patched by remote control between the five beams. Not all combinations are possible, due to physical constraints in the satellite and due to interference between adjacent beams. For example, CERN cannot talk simultaneously to Finland and Greece. However, by steering the beam cluster slightly from its nominal position any site in Western Europe can be reached from CERN. In fact the payload has been designed such that quasi–simultaneous connection between the areas covered by the five spots is only possible using SS–TDMA (satellite switched time division multiple access), a substantial complication which has been avoided in the initial design of CHEOPS. Thus a suitable schedule will have to be used to establish links successively between various pairs of sites.

Time on OLYMPUS is allocated to different experiments according to a daily schedule. ESA have indicated that CHEOPS can expect two one–hour sessions each working day during prime time, and several hours per night and at weekends. There is no chance of more than this, due to the many approved experiments.

The four initial CHEOPS earth stations, supplied by Nokia Telecom, will have 3m antennae, 8 W radio–frequency amplifiers, and G.703 modems capable of bit rates up to 8.442 Mbit/s with a bit error rate better than 10^{-9}. Such earth stations are cheap and robust, and can be installed almost anywhere with little technical difficulty, subject to official approval. They will access the Ku–band transponders of the Specialised Services payload, allowing point–to–point links to be established between CERN and each other site in sequence.

Thus we hope that by the end of 1991, OLYMPUS and these earth stations will deliver scheduled 8 Mbit/s G.703 links between CERN and each CHEOPS site for fixed periods, with the 250 ms satellite transit delay as the only technical disadvantage.

It should be added that obtaining official approval for small send/receive earth stations, even in a purely scientific and experimental context, is not simple. Although most countries involved now have a telecommunications regulatory body which is theoretically independent of the telecommunications service provider(s), the collective influence of the PTOs is still a brake on satellite technology in Europe.

Chapter 4

SERVERS, INTERFACES, AND CONTROL

Data rate calculations show that to use an 8 Mbit/s duplex link effectively, the datasets must be available on a dedicated disc server at the source, and space for them must be available on a dedicated server at the destination. The I/O rates require servers of at least 20 RISC MIPS. These servers will be Sun SPARCservers fitted with high–performance IPI discs, with the initial goal that the CERN server can deliver up to 5 Gbytes of data per night, and each remote site accept up to 1 Gbyte per night at a nominal rate of 1 Gbyte/hour. These quantities of data correspond to estimates of how much physics data a significant computer centre can usefully absorb, and to 25% effective utilisation of the 8 Mbit/s channel.

At the time of writing, two options are still open for the connection between the servers and the satellite modems. The first option is to multiplex four 2 Mbit/s Sun interfaces into the 8 Mbit/s channel, and the second is to use a dedicated Ethernet interface and a CONWARE Ethernet–G.703 bridge. Both of these options have implications for performance which are under study.

Each server will also be connected at the highest speed available to the local site infrastructure, so that datasets can be spooled into and out of the server during daytime. At CERN, the interface will be Ultranet [8] with FDDI as an alternative.

The earth station network will be controlled by a dedicated Transmission Management PC, itself linked to the server at CERN.

Chapter 5

PROTOCOL EFFICIENCY

For practical reasons we have only considered TCP/IP family protocols. It is well known that "raw" FTP/TCP/IP will perform very badly over an 8 Mbit/s link with 500 ms round–trip time, even with maximum tuning. Using an artificial delay on Ethernet, we have so far measured single–stream TCP performance up to 256 kbit/s using 16kByte windows. Other investigators have reported measurements in the range 400 kbit/s to 1 Mbit/s with careful tuning [4][5]. Running several parallel streams, we have shown that performance collapses above four or five streams, a result also observed elsewhere [5].

We therefore expect that use of a transport protocol designed for "long fat networks" will be required, and we are evaluating XTP [6] and Extended TCP [7] for this purpose.

The choice of hardware interface mentioned above may also lead to work on low level IP implementation, especially if the option of four 2 Mbit/s interfaces is chosen.

Chapter 6

THE CHEOPS AUTOMAT

The heart of CHEOPS is the software automat which schedules and controls the following processes:

- During dayshift, physicists or other users specify interactively the datasets they wish to be transferred. The main data source will be cartridge tapes, mounted robotically or manually, but disc files may also be specified.

- Overnight, the datasets will be transferred via OLYMPUS, between CERN and each site in succession.

- The next day, the user will receive a success/failure report.

User requests at the various sites must be forwarded to the CHEOPS server at CERN outside the satellite transmission window. The servers will therefore communicate via terrestrial links at conventional speeds, in principle via existing or planned TCP/IP connections.

Figure 1 illustrates the general architecture. The software design is now being completed using the SA/SD methodology, which has been found an excellent tool for this project. About 35 pages of diagrams and data dictionary suffice to define the CHEOPS automat. Very little of the design actually concerns satellite transmission as such, and the CHEOPS automat is expected to be re-usable to exploit any form of circuit-switched capacity (e.g. BISDN at off-peak tariffs) or even leased lines when they are empty overnight.

Chapter 7

STATUS AND PLANNING

The CHEOPS design is essentially complete, some of the servers have been delivered, and delivery of the first two earth stations is imminent. We expect that by the end of 1991, transmission tests using unmodified FTP/TCP/IP will have been completed between CERN and three other sites, that an optimised protocol such as XTP will be on test, and that the CHEOPS software automat will be ready for testing. Commissioning and tests of the full system, and the addition of two extra sites, would occur in early 1992, followed by an extended pseudo-operational test.

REFERENCES

[1] M.Hine, The STELLA experiment – final report, EUR 9090, CEC, 1984.

[2] A.Cohen, T.Massart, B.Sales, P.Van Binst, The HELIOS Project: OSI Conformant High Speed Data Communication by Satellite in Europe, Proc. 9th International Conf. on Computer Communication, Tel Aviv, 1988, ed. J.Raviv, North Holland, p.94.

[3] J.Shiers, FATMEN CERN Program Library entry Q123 Long Writeup, CERN internal report, 1990.

[4] C. Garrido, J.Stjernevi, Computer Networking using High Speed Satellite Link, ESA Bulletin, October 1990; C.Garrido, private communication, 1991.

[5] J.Lekashman, Type of Service Wide Area Networking, Proc. Supercomputing '89, November 1989, Reno, Nevada.

[6] R.Sanders, A.Weaver, The Xpress Transfer Protocol (XTP) – a Tutorial, ACM Computer Communication Review, Vol. 20 No. 5 p. 67 (1990).

[7] V.Jacobson, R.Braden, TCP Extensions for Long Delay Paths, RFC 1072, Network Information Centre, 1988; V.Jacobson, R.Braden, L.Zhang, TCP Extension for High-Speed Paths, RFC 1185, Network Information Centre, 1990.

[8] L.Robertson et al, SHIFT, the Scalable Heterogeneous Integrated Facility for HEP Computing, Proc. Computing for High Energy Physics, March 1991, Tsukuba, Japan.

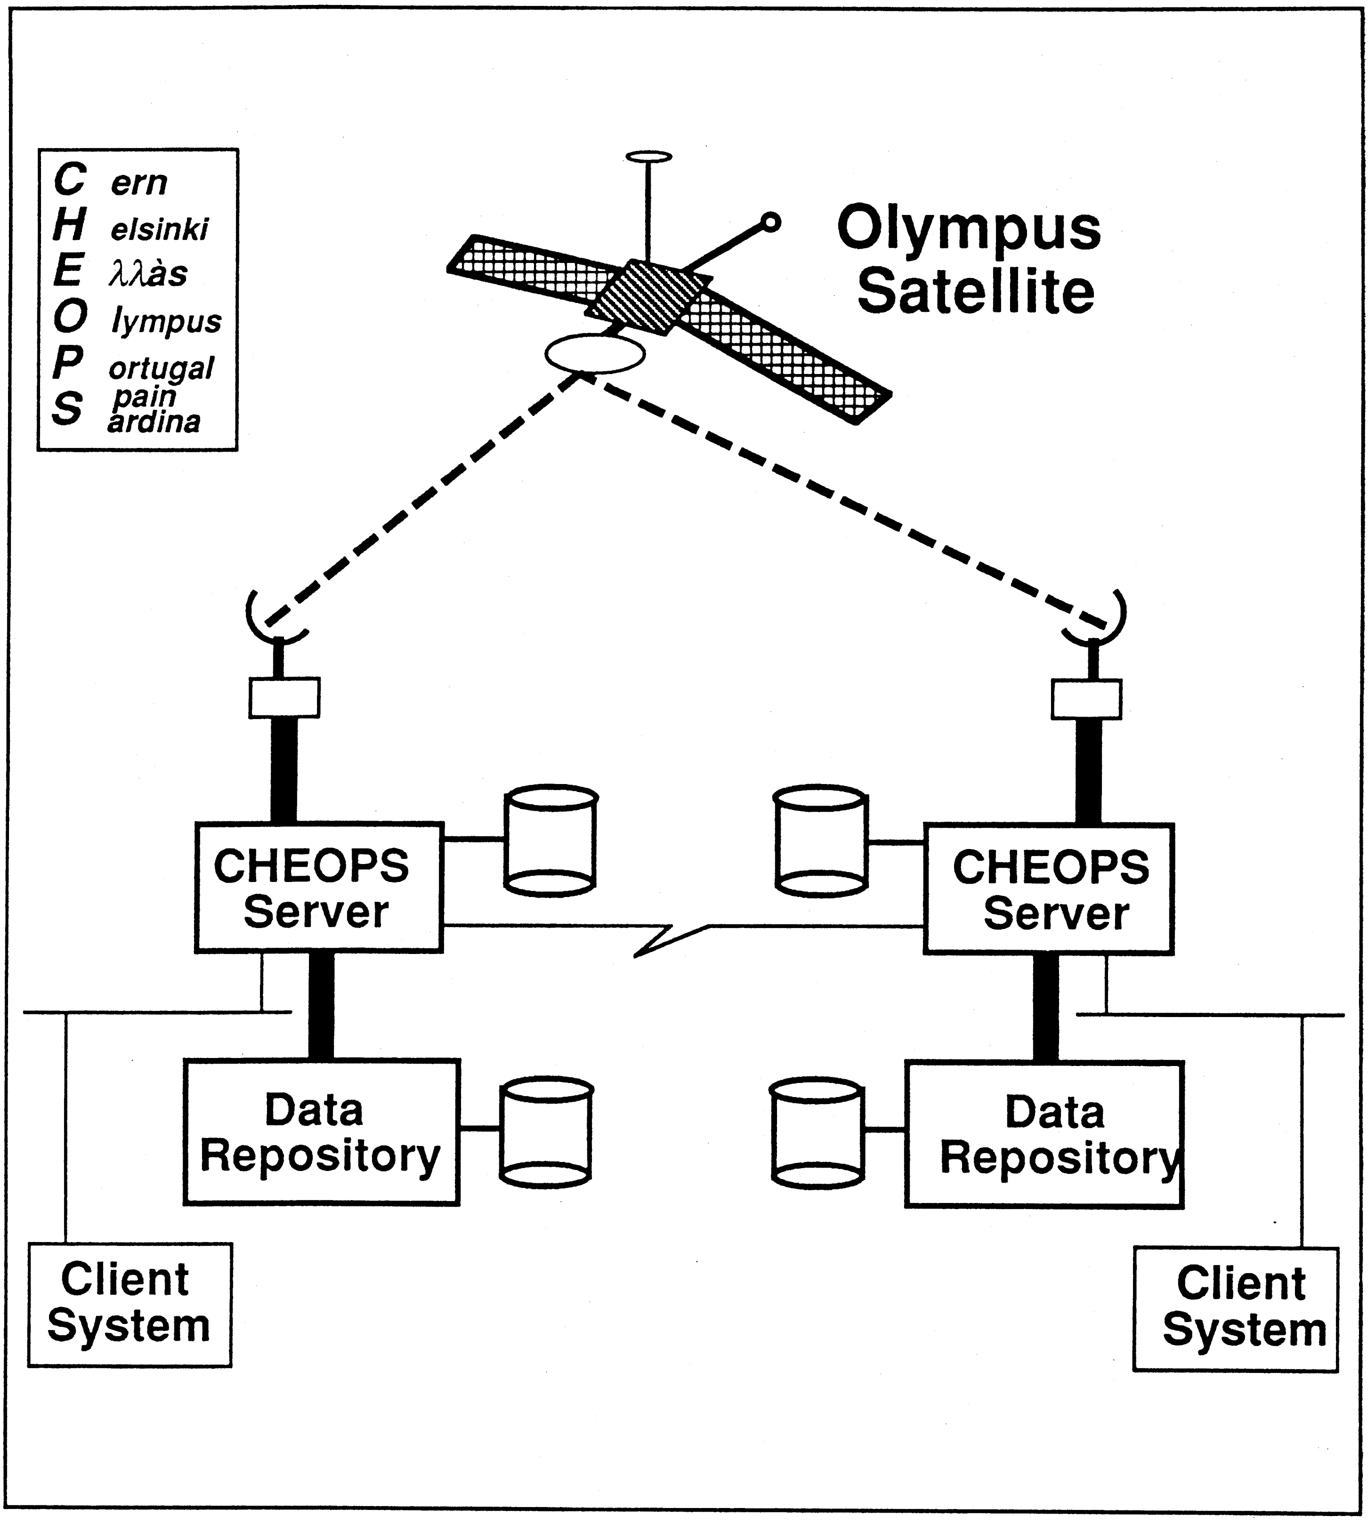

Fig.1. CHEOPS Architecture

Imprimé au CERN. Prix : 0.50 Fr.