

1. Introduction

This report will cover the concepts of distributed computing and network infrastructure within the context of allowing data to be processed over a system similar to the SETI@Home project (Search for extra terrestrial intelligence).

2. Distributed Processing

A distributed system consists of multiple computers working together on the same task. These computers will each have their own processors, memory and hard disk drives and are able to process information independently; however, they are all interconnected so they are able to share the processing load. (Mullender, 1993) Distributed processing allows a large processing task to be shared across a distributed system. This allows for decentralised processing with centralised control. (Martin, 1981)

The basic infrastructure of the Internet has been around since 1969, when ARPA (Advanced Research Project Agency) connected together several research institutes and universities. (White, 2008) It took about 30 years however for the size of the network and the number of users to reach a point where the Internet could be used as a distributed system. The potential of sharing across the Internet was realised in the late 90's when Napster began offering MP3 downloads using peer-to-peer technology. A similar technology, Grid Computing, is becoming more and more widespread and uses the same principles but allows for the sharing of system resources such as processing power, thus creating a "virtual supercomputer" with massive amounts of processing power, memory and hard disk space which increases as more users become a part of it. (Beekman, 2005)

The SETI@Home project is the best example of how distributed processing across the Internet can be successful. Any user can sign up and need only download a platform independent application that will download "workunits", carry out the necessary processing when the user's computer is idle and then upload the results. (McConnell, 2001) An algorithm allows the vast amount of data collected by SETI (about 39TB a day) to be broken into smaller pieces which can be individually processed. (Englander, 2003) (Englander, 2003)

3. Components of the Technical and Network Infrastructure of SETI@Home

3.1 Digitizing and Preparing Data

The data derived from the telescope is recorded onto removable tapes at a sampling rate of 5mbps, low enough for conversion and transition to occur in a reasonable amount of time and high enough to be scientifically accurate and account for Doppler shifted transmissions. The data is split up into "workunits" by a splitter server which then sends each workunit to both the database server and workunit storage. (Anderson, 2002)

3.2 Storing and Distributing Workunits

The database server uses a RDBMS to store data including: users, workunits and results. The data/results server is a multithreaded Linux system that uses a CGI Backend run by LAMP in order to process requests for workunits and receive results. It also runs 5 daemons written in C++ to process data. (BOINC Client Server Technology, 2009) A client connects to the data/results server over the Internet and downloads a workunit using HTTP, allowing even fire-walled users to do so. The client need only connect again when it is ready to upload the processed workunit. (Anderson, 2002) This

means that the client only communicates with the server to send a finished workunit and request another, and is able to independently process each workunit without a continuous connection to the server.

The feeder daemon retrieves a workunit from the database and stores it in shared memory for the CGI to access. It continually replenishes with new workunits as required.

3.3 Receiving, Processing and Storing Results

To ensure accuracy, the same workunit will be processed by multiple clients, the results of which will be compared by a validating daemon; if they match the workunit is marked complete and a canonical result is created. (BOINC Client Server Technology, 2009)

The assimilator daemon writes the canonical data to a flat file. These files are processed producing both a scientific output and information about the user who processed the workunit. This information is not instantly added to the database, it is instead stored in a buffer which is flushed into the database every few minutes. This method not only reduces pressure on the database it also allows for periods of database overload or outage. Multiple instances of the assimilator daemon will run concurrently to give a higher throughput. (Anderson, 2002)

The file deleter daemon erases the now unneeded flat files and the transitioner daemon performs processing on the scientific data in the database.

3.4 Similar Projects and SETI@Home Servers

The SETI@Home project formed the basis of BOINC (Berkeley Open Infrastructure for Network Computing) which is used by 10's of similar projects including the Large Hadron Collider. (BOINC Manager, 2009)

The SETI@Home project currently uses 10 servers, 1 splitter, 1 master database, 1 replica database, a primary and secondary science database, 1 web server, 2 upload servers, 1 upload server, and 1 feeder/transitioner server. (SETI@Home, 2009)

4. Implications of Failure

4.1 Network Failure

The client server relationship of BOINC is such that a connection between the two is only required when downloading or uploading workunits. This means that the client is able to work independently for the majority of the time. In the event that the server is unreachable when required by the client the only adverse affect will be a delay in result return and acquisition of new workunits. (Anderson, 2002)

Server redundancy and load balancing can help to prevent any periods of network downtime. The SETI@Home project utilises a network of servers many of which all performing the same tasks. (SETI@Home, 2009) As a result it is possible for one or more of the servers to be taken offline intentionally or otherwise without affecting its counterparts. Load balancing also prevents servers from being overloaded or bombarded with requests assisting in the prevention of downtime. A RAID (Redundant Array of Independent Disks) prevents data loss and network downtime if the main hard drive on a server fails a backup drive takes over, pre-loaded with an exact copy of the main hard drive content.

4.2 Database Failure

To accommodate for any database unavailability due to overload or outage, the assimilator daemon buffers results in a flat file before writing them to the database every few minutes. (BOINC Client Server Technology, 2009) Loss of data is a very serious implication, be it caused by failure of hardware or software, a malicious user, power outage, fire or a natural disaster like flooding. A daily backup stored in a secure alternate location solves this issue, as well as utilising the ancestral system in which a back log of backups are stored. This means that in the event of data loss no more than 24 hours of data will be lost.

4.3 Security and Validity

There are instances where results calculated during the processing of a workunit could be inaccurate. This could be caused by a malicious user or interception and tampering of results whilst they are being uploaded by the client. Inaccuracy could also be caused by differing operating systems, processor types or even a malfunction in the hardware or software of the client computer. To prevent this from happening, the same workunit is distributed to multiple clients running the same operating system with processors of the same type (For example Intel Pentium) meaning that the results returned should be the same. A validation daemon compares the results and based on a tolerance produces a final canonical result. (BOINC Client Server Technology, 2009)

5. Expected Contributors

The SETI@Home (and similar projects) relies completely on users voluntarily sharing their system resources. (Foster, 2004)

To get participants to signup, allowing drug test results to be processed on their computers, they must be told the benefits that their contribution will have to give them an incentive. SETI@Home started with about 400,000 user registering in the first year, the number has now grown to over 3.83 million. (Anderson, 2002) Any user who is interested in helping better medical science or even just users looking to be involved in distributing processing projects are suitable participants for this project.

The online community provides users with the ability to view information about anyone who is part of the programme. This is the reason why information about users is collected when uploading workunits (See 3.3). Individuals can see how well their processors are performing, how many workunits they have completed and other non private statistics. Users also have the ability to join teams meaning their individual statistics contribute to a team total. A system of credits and awards is utilised rewarding users for processing. The credits awarded are used to generate a highest contributors "leader board" showing the users who have completed the most processing. This encourages friendly competition between both individuals and teams to process more data in an attempt to process more than their rivals. The upshot of which results in more workunits being completed which is beneficial to the project. (SETI@Home, 2009)

The technical staff required should be quite minimal as they need only maintain the servers and the database. SETI@Home is run by only 8 people, 2 of which are system administrators and software developers, 1 is a database administrator, a director and 4 scientists who work on analysing the results and are not involved in running the servers in any way. As BOINC already exists creating a whole new distributed processing application is unnecessary so no development personal will be required other than implementing and altering the software, where necessary, to allow it to accept

drug test input and produce the required output. A similar staffing of 2 server administrators and 1 database administrator should be sufficient.

6. Expected Technical Issues

Databases, software and servers crash or start behaving oddly from time to time and will often take days to find and fix the problem causing it. [AND02] During these periods, servers may not communicate with each other correctly or at all, and clients may be unable to connect to the servers causing “chaos”. The SETI@Home project experienced an issue similar to this, which was causing validation to fail on completed workunits.

Periods of server availability are sometimes caused intentionally, for example when backing up the database or rebooting a server. The SETI@Home project backs up the database daily (Technical News, 2009) which makes it unavailable. There have been numerous issues caused by a server’s inability to mount a hard drive, which is solved by rebooting. (Technical News, 2009) Also, the SETI@Home project is always looking to expand and develop its network to improve it in terms of speed, processing power or simply to replace hardware which has worn out. This means that the entire system will be affected as components are taken offline for work.

Software does not always perform as it is supposed. To prevent database overload, results are stored in flat files and flushed to the database every few minutes (see 3.3 and 4.2). Sometime however, the flat files are not completely flushed to the database meaning that there may be results in the database that are also being stored in flat files and results still in flat files that are not in the database. As a result, directories can become very large and unmanageable and custom written program are required in order to clean them up and place the missing information into the database whilst removing all redundant flat files. (Technical News, 2009)

Whilst the aim of the project is to get as many users processing data as possible, this causes severer backlogs and bandwidth issues as more and more users sign up to a network which was not designed to handle an ever increasing amount of traffic. Servers can literally get gridlocked where they are receiving so many requests from clients their responsiveness decreases until eventually reaching almost a standstill. The amount of bandwidth on the Internet connection, the routers and Ethernet connections on the servers is finite and as more requests sent by an ever increasing number of users will cause the network to slow down. This problem can be rectified by improving the Internet connection, replacing routers with faster ones able to handle more traffic and upgrading server Ethernet hardware. (Technical News, 2009)

Due to the budget, it is only possible to have so many people working during so many hours of the day – despite the fact that servers could require rebooting or other work at any time of day. SETI@Home staff have remote control, however, meaning any staff member can reboot a server from home.

7. Benefits of Distributed Processing

Distributed data processing reduces cost as an organisation is able to utilise processing power on volunteers’ computers without having to invest in their own mainframe. (Blurtit, 2009) The disadvantage of this is actually getting users to volunteer their processing power for free.

Providing that enough users volunteer, the speed at which data can be processed will grow to much more than anything that could be realistically provided by the organisation running its own mainframe. This unfortunately has the disadvantage that as more and more users volunteer the servers, available bandwidth and network infrastructure need to be able handle the distribution and receipt of workunits and need to be constantly monitored as any downtime can cause loss or delays with receiving results or aggravate users who may discontinue their participation if they are continually unable to connect to download new workunits.

The biggest problem with relinquishing data to potentially anyone is security. Any sensitive information could not be processed in this manor as it would not be secure. Similarly if sensitive data was sent out for processing it could potentially be read by anyone. There is also the possibility that a malicious user may attempt to send back results that are in accurate or tampered with.

2,203 words

Glossary

Advanced Research Project Agency

An agency of the American government responsible for the development of new technologies during the 1970's. One such technology was a vast computer network which became the Internet. (Defence Advanced Research Projects, 2009)

Backend

The name given to a function of a piece of software that displays results derived from user input. (Backend, 2009)

Bandwidth

A measurement of the amount of data that can be transferred by a network, measured in kilobytes or megabytes per second. (Wikipedia, 2009)

BOINC (Berkley Open Infrastructure for Network Computing)

A grid computing application which was developed by the SETI@Home project. It allows volunteers to sign up to participate in distributed data processing. It is now used by many projects in many fields such as medicine, mathematics and astrophysics. (Berkeley Open Infrastructure for Network Computing, 2009)

Buffer

A section of memory in a computer that temporarily holds data before it is transferred to another more permanent location. (Wikipedia, 2009)

Canonical

Describes data that has been reduced to its simplistic and most significant form without loss of generality. (Wikipedia, 2009)

CGI

Common gateway interface. The method for interfacing external application software with an information server – usually a web server. (Wikipedia, 2009)

Client

A client is a computer that accesses a remote service on another computer.

C++

A computer programming language. (Wikipedia, 2009)

Database

A shared collection of logically related data, and a description of the data, designed to meet the information needs of an organisation. (Connolly et al, 2005)

Daemon

A piece of software that runs in the background on UNIX based computers. (Wikipedia, 2009)

Doppler Effect

The name given to the change that occurs in the frequency and wavelength of a wave from the point of view of an observer moving relative to the source. (Wikipedia, 2009)

Flat File

A plain text or binary file on a computer that usually contains one record per line.

Hyper Text Transfer Protocol (HTTP)

A computer protocol (a set of pre determined rules) for transferring data between to computers usually over the Internet.

Internet

A vast computer network that follows the mesh topology. It allows user to access the World Wide Web, share files, video conference and many more.

LAMP

Linux, Apache, MySQL, PHP. A computer or server can use all four of these services to provide a web page. Linux is an operating system (see Linux). Apache is web server software that listens for HTTP connection and returns a web page. MySQL is database software and PHP is a server side scripting language.

Large Hadron Collider

The world's largest particle accelerator, built by CERN in an attempt to prove the existence of the Graviton particle. (Wikipedia, 2009)

Load Balancing

A technique in which the work load is spread across one or more computers locally networked in order the share the processing load and make optimal use of available resources.

Linux

A UNIX based operating system run primarily on servers (such as a web server) and also used by developers.

Multithreaded

A computer program often runs more than one "thread" which is perceived as multiple tasks. A single processor computer will switch between threads giving the appearance that it is running them concurrently. A computer with multiple processors will often run one thread on each processor. (Wikipedia, 2009)

Napster

A music sharing website which indexes files available for download using P2P (peer-to-peer) technology. (Beekman, 2005)

Peer-to-Peer

A file sharing technique in which files are download by a user from many other users who have some or all of that file.

Relational Database Management System (RDBMS)

A database which joins together records in one table with records in another using a primary key to foreign key relationship. (Wikipedia, 2009)

Sampling

The term used when translating analogue into digital data. The analogue waveform is “sampled” at a predetermined frequency and the amplitude of the wave is recorded digitally.

Server

A computer that provides services to other computers (clients) over a network.

SETI (Search for Extra Terrestrial Intelligence)

A project which uses a deep space radio telescope in an attempt to detect radio signals of an extra-terrestrial origin. (Wikipedia, 2009)

Shared Memory

Memory on a computer which can be simultaneously accessed by more than one program with the intent to provide communication above them. (Wikipedia, 2009)

Terabyte (TB)

A measurement of data, consisting of 1000 gigabytes. (Wikipedia, 2009)

Throughput

The ratio of data input to data output. The more output there is compared to input, the higher the throughput. (Google, 2009)

References

- White. R. (2008). How Computers Work. 9th Ed. US: Que Publishing
- Englander. I. (2003) The Architecture of Computer Hardware and Software Systems. 3rd Ed. US: Leyh Publishing
- Beekman. G, Quinn. M (2006) Computer Confluence: Tomorrow's Technology and you. International Edition. US: Prentice Hall
- Foster. I. And Kesselman. C. Eds. (2004) The Grid. 2nd Ed. US: Morgan Kaufmann Publishers
- Mullender. S. (1993). Distributed Systems. 2nd Ed. US: ACM Press
- Martin. J. (1981). Design and Strategy For Distributed Data Processing. 1st Ed. US: Prentice Hall
- McConnell. B. (2001). Beyond Contact: A Guide to SETI and Communication with Alien Civilisations. 1st Ed. US: O'Reilly & Associates Inc
- SETI@Home [Online] Available From :<http://setiathome.ssl.berkeley.edu> [Accessed 18th March 2009]
- Anderson. D. (2002) SETI@Home: An Experiment in Public Resource Computing
- BOINC Client Server Technology [Online] Available From: http://en.wikipedia.org/wiki/BOINC_client-server_technology [Accessed 18th March 2009]
- BOINC Manager [Application] (Version 6.4.5) Available From: <http://boinc.berkeley.edu> [Accessed 18th March 2009]
- Technical News [Online] Available From: http://setiathome.berkeley.edu/tech_news.php [Accessed 21st March 2009]
- Defence Advanced Research Projects Agency [Online] Available from: http://en.wikipedia.org/wiki/Defense_Advanced_Research_Projects_Agency [Accessed 22nd March 2009]
- Backend [Online] Available from: <http://en.wikipedia.org/wiki/Backend> [Accessed 22nd March 2009]
- Berkeley Open Infrastructure for Network Computing [Online] Available from: <http://en.wikipedia.org/wiki/Boinc> [Accessed 22nd March 2009]
- Wikipedia Data Buffer [Online] Available from: http://en.wikipedia.org/wiki/Data_buffer [Accessed 22nd March 2009]
- Canonical [Online] Available from: <http://en.wikipedia.org/wiki/Canonical> [Accessed 22nd March 2009]
- Bandwidth (Computing) [Online] Available from: [http://en.wikipedia.org/wiki/Bandwidth_\(computing\)](http://en.wikipedia.org/wiki/Bandwidth_(computing)) [Accessed 22nd March 2009]
- Shared Memory [Online] Available from: http://en.wikipedia.org/wiki/Shared_memory [Accessed 22nd March 2009]
- SETI [Online] Available from: <http://en.wikipedia.org/wiki/SETI> [Accessed 22nd March 2009]

Terabyte [Online] Available from: <http://en.wikipedia.org/wiki/Terabyte> [Accessed 22nd March 2009]

Large Hadron Collider [Online] Available from: <http://en.wikipedia.org/wiki/Lhc> [Accessed 22nd March 2009]

C++ [Online] Available from: <http://en.wikipedia.org/wiki/C%2B%2B> [Accessed 22nd March 2009]

Define: Throughput [Online] Available from:
<http://www.google.co.uk/search?hl=en&safe=off&client=firefox-a&rls=org.mozilla%3Aen-GB%3Aofficial&hs=dWI&q=define%3A+throughput&btnG=Search&meta=> [Accessed 22nd March 2009]

Common Gateway Interface [Online] Available from:
http://en.wikipedia.org/wiki/Common_Gateway_Interface [Accessed 22nd March 2009]

Connolly, T. Begg, C. (2005) Database Systems: A practical Approach to Design, Implementation, and Management. 4th Ed. US: Addison Wesley.

Daemon (Computer Software) [Online] Available from:
[http://en.wikipedia.org/wiki/Daemon_\(computer_software\)](http://en.wikipedia.org/wiki/Daemon_(computer_software)) [Accessed 22nd March 2009]

Doppler Effect [Online] Available from: http://en.wikipedia.org/wiki/Doppler_Effect [Accessed 22nd March 2009]

Load Balancing (Computing) [Online] Available from:
[http://en.wikipedia.org/wiki/Load_balancing_\(computing\)](http://en.wikipedia.org/wiki/Load_balancing_(computing)) [Accessed 22nd March 2009]

Thread (Computer Science) [Online] Available from: <http://en.wikipedia.org/wiki/Multithreaded> [Accessed 22nd March 2009]

Can you brief some advantages of distributed data processing system? [Online] Available from:
<http://www.blurtit.com/q471790.html> [Accessed 22nd March 2009]