Doctoral Studies and Research Proposition

Diversity in Peer-to-Peer Networks

Mikko Pervilä

Helsinki 24 November 2008

UNIVERSITY OF HELSINKI

Department of Computer Science
Supervisor: prof. Jussi Kangasharju

1 Diversity in Peer-to-Peer Networks

Current peer-to-peer systems (P2P) have been identified to hold the property of being *heterogeneous* [StW06], which means that the clients participating in the network cannot be assumed to manifest similar features, e.g., operating systems, hardware, or network connections. This property is also advantageous when modelling reliability, for the nodes can now be assumed to fail independently of each other. A statistical model that contains the assumption of uncorrelated failures is far more simple than one that tries to model every feature separately. The proposed research questions this simplicity, in how far the elegant models [Sto01, KRT03] deviate from the diversity [RoD01] of existing P2P systems.

When the total number of participating nodes scales upwards to hundreds of even thousands of peers, the assumption that no common properties would manifest may become unsound. Global manufacturing, shipping, and sales processes as well as clearly defined installation and maintenance procedures give cause to believe that at least the computing hardware can encounter simultaneous failures.

A motivational example concerns RAID arrays, whose administrators are instructed by manufacturers to only install their own disk subsystems, ideally of the same model for performance reasons. If the administrator then conveniently purchases the disks from a single outlet, shipping procedures ensure that with a high probability, the disks originate from the same manufacturing plant with [almost] sequential serial numbers. Thus, the resulting full array can no longer be exempt from correlated failures. Correlated disk failures were noted by Chen *et al.* [Che94] as early as 1994 – yet RAID arrays are still built with homogeneous devices.

1.1 Research Questions and Methodology

While RAID arrays are arguably distributed systems of the previous generation, anecdotal evidence [IBM05, Ses08, KuN08] suggests that the example carries over to many, if not all, systems employing COTS hardware. Even if the hardware of a P2P network is not administered by any single organisation, similar scenarios may be caused by the required client software [Avi84, Per06] when programming defects are replicated into every participating node.

Whereas the severity of hardware failures depends largely on their position (hosts versus routers) and mean time to recovery, software failures may be considerably more significant and logically system-wide. Thus, further research into the reliability and performance of P2P systems is proposed. The conjecture is that these properties are related by a balance decision, wherein performance and administration benefit from the composition of homogeneous nodes, while reliability benefits from heightened diversity in the network.

Performance and reliability studies combine well with other vectors of interest into peer-to-peer systems. The necessary implementations could be simultaneously studied by members of a project team with similar research projects. Small scale networks should be feasible by using surplus (nightly) processing hours of the local workstations at the Department of Computer Science. For larger measurements, the proposed project will probably require the aid of volunteers in order to gather sufficient research data.

Research into this field must by necessity be of practical nature, i.e., by observing existing P2P networks and reference implementations, for modelling every relevant property may be infeasible. Existing models can be compared with the data obtained by measurements in order to examine the amount of entropy present, and if the models consequently hold for real systems at all. One possibility might be to categorise the clients into *classes* based on their similar properties: hardware and software components, location, ownership, origin, usage patterns, and so on. An open research question concerns the depth of such a taxonomy, for the number of combinations may be so vast that it can only be quantifiable by an automated system.

1.2 Expected Research Outcomes

The measurements of current systems are expected to yield a great amount of real-world data, which shall be suitably formatted and made available for additional research and analysis. Some of the inherent properties of P2P systems will probably be of greater interest initially than others. These points of interest are thought to be similar to performance bottlenecks, but in the domain of reliability. Immediate optimisations will then target the most interesting properties first. For example, it is expected that faults in client software will be of greater consequence than faults in hardware, because of the smaller set of software available for a given P2P application, when compared to compatible hardware.

Ultimately, the proposed research seeks to prove the conjecture that there exists a balance question between administration and reliability when it comes to the diversity of peer-to-peer networks. For now, completely diverse systems are expected to remain infeasible to totally control by any single point of authority. These systems will perhaps become plausible in the near-to-mid-term future, when the tasks of specification, development, deployment, and administration can be automated to a greater extent.

1.3 Schedule and Funding

The research is proposed to take place in a time line of five years, to be funded by a doctoral student position in the field of Distributed Systems and Telecommunication, at the Department of Computer Science, University of Helsinki. Visiting other universities may be required if research or relevant studies are organised abroad. The goal is a doctoral dissertation from the University of Helsinki, as supervised by professor Jussi Kangasharju.

2 Plan of Doctoral Studies

The following courses are selected to support my doctoral studies from three different viewpoints. First, the courses may concern P2P Networks directly. Second, they may support the analysis of P2P reliability or performance. Third, they may assist in teaching courses at the Department of Computer Science.

Spring 2008

- 4 credits, *Peer-to-Peer Networks* (582615)
- 8 credits, *Tietokoneavusteiset oppimisympäristöt* (58037)
- 8 credits, *Natural Language Processing* (582602)
- 3 credits, *Participation in Scientific Conferences*

Autumn 2008

- 3 credits, Seminar: Future Internet and Hot Topics in Networking (58308305)
- 10 credits, Mathematics

Spring 2009

- 10 credits, Statistics (Probability theory)
- 4 credits of additional CS courses, e.g., Suorituskykyanalyysi (if lectured)

Autumn 2009

- 6 credits, Seminar in Doctoral Studies (582710)
- 4 credits of additional CS courses

Yielding a total of 60 credits, not including the doctor's thesis.

References

AvK84 Avižienis, A. and Kelly J., Fault tolerance by design diversity - Concepts and experiments. *Computer* Vol. 17 (Aug. 1984), 67-80.

Che94 Chen, P., et al. RAID: High-Performance, Reliable Secondary Storage. *ACM Computing Surveys* (CSUR) Vol. 26, Issue 2 (June 1994), 145-185.

IBM05 Internet Archive Wayback Machine, IBM Desktar 75GXP Class Action Website, http://web.archive.org/web/20060315210819/http://www.ibm-deskstar75gxplitigation.com/. [9.6.2008]

KRT03 Kangasharju, J., Ross, K. W., Turner, D. A., Secure and Resilient Peer-to-Peer Email: Design and Implementation. *Proc. International Conference on Peer-to-Peer Computing*, Linköping, Sweden, Sept. 1-3, 2003.

KuN08 Interview with Kutvonen, P. (IT Manager), and Niklander, P. (IT Specialist), Computing Facilities Staff, Department of Computer Science, University of Helsinki, 9.6.2008.

Ses08 Private communication with Seskar, Ivan, Associate Director at Wireless Information Network Laboratory (WINLAB), Rutgers University, New Jersey, 6.5.2008.

Sto01 Stoica, I., et al., Chord: A Scalable peer-to-peer lookup service for Internet applications. *Proc. SIGCOMM*, San Diego, CA, Aug. 27-21, 2001, pages 149-160.

Per06 Pervilä, M., Luotettavuus vikasietoisin toisinnoin (in Finnish). Bachelor's Thesis, Department of Computer Science, University of Helsinki, December 2006.

RoD01 Rowstron, D. and Druschel, P., Pastry: Scalable, Distributed Object Location and Routing for large-scale peer-to-peer systems. *Proc. IFIP/ACM International Conference on Distributed Systems Platforms* (Heidelberg, Germany, Nov. 2001), London, UK, 2001, pages 329-350.

StW06 Steinmetz, R. and Wehrle, K. (eds.), *Peer-to-Peer Systems and Applications* (Lecture Notes in Computer Science). Springer-Verlag New York Inc., NJ, USA, 2005.