

The Nagios Community: An Extended Quantitative Analysis

Jonas Gamalielsson, Björn Lundell, Brian Lings

► **To cite this version:**

Jonas Gamalielsson, Björn Lundell, Brian Lings. The Nagios Community: An Extended Quantitative Analysis. Pär Ågerfalk; Cornelia Boldyreff; Jesús M. González-Barahona; Gregory R. Madey; John Noll. 6th International IFIP WG 2.13 Conference on Open Source Systems,(OSS), May 2010, Notre Dame, United States. Springer, IFIP Advances in Information and Communication Technology, AICT-319, pp.85-96, 2010, Open Source Software: New Horizons. <10.1007/978-3-642-13244-5_7>. <hal-01056022>

HAL Id: hal-01056022

<https://hal.inria.fr/hal-01056022>

Submitted on 14 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



The Nagios community: An extended quantitative analysis

Jonas Gamalielsson, Björn Lundell and Brian Lings
University of Skövde, Sweden
{jonas.gamalielsson, bjorn.lundell, brian.lings}@his.se,
WWW home page: <http://www.his.se>

Abstract. The health of an Open Source ecosystem is an important decision factor when considering the adoption of an Open Source software or when monitoring a seeded Open Source project. In this paper we assess the ecosystem health using approaches involving domain analysis and social network analysis of mailing lists for the Nagios project. We elaborate approaches for how involvement of different roles can be analysed through quantitative analysis, specifically focusing on core developers and professional providers. Our contribution is a step towards a deeper understanding of professional involvement in professional Open Source ecosystems.

1 Introduction

Before an organisation adopts an Open Source project it is important to evaluate its community in order to make sure that it is healthy and that the project is likely to be maintained for a long time (van der Linden et al. 2009), and it can also be the case that a seeded Open Source project needs to be monitored. One important means in such an evaluation is to quantitatively assess the health of an Open Source community (Crowston and Howison 2006). A number of studies have investigated large, well known Open Source projects through quantitative analysis, including the Linux kernel (Moon and Sproull 2000), Apache (Mockus et al. 2002), Mozilla (Mockus et al. 2002), Gnome (German 2004) and KDE (Lopez-Fernandez 2006). Several of these studies focus on social network analysis from different kinds of data sources such as CVS/SVN (Martinez-Romo et al. 2008), bug reports (Crowston and Howison 2005) and mailing lists (Kamei et al. 2008).

We have chosen to analyse social networks derived from the mailing lists of an Open Source project. Such an analysis can reveal how active members in a community are interacting with each other. In earlier work (Gamalielsson et al. 2009) we adopted the approach of Kamei et al. (2008), who studied the community of the Apache web server project, in a study of the Nagios project (www.nagios.org). Kamei et al. (2008) studied the user and developer mailing lists for a period of six months; three months before and after a major release of Apache. In our earlier work, we chose to study a period of six months from January 2009 to June 2009. The aim was to investigate to what extent core developers act as mediators within and between the sub-communities emerging around the user- and developer mailing lists. In this work

we extend the earlier work by studying 22 six-month periods from January 2004 to September 2009, and also suggesting other kinds of analyses.

Important roles in a typical Open Source project are users, developers, core developers and project leaders (Crowston and Howison 2006). The importance of core developers applies to any Open Source project as it is well established in the literature that core developers “contribute most of the code and oversee the design and evolution of the project.” (Crowston et al. 2006). The influence of Nagios core developers was explored in the earlier work (Gamalielsson et al. 2009), and also in this work. Another role to consider, especially with the advent of OSS 2.0 (Fitzgerald 2006), is the provider role. According to Fitzgerald (2006), OSS 2.0 product support is characterised by “customers willing to pay for a professional ‘whole-product’ approach”. The provider role is also emphasised in the OSS stakeholder triangle proposed by Lundell et al. (2009). The model conceptualises an healthy Open Source ecosystem involving different kinds of professional actors. In such ecosystems there are mutualistic relationships between all roles (developer, user and provider). Concerning the provider aspect, it is evident that there are still misconceptions about support of Open Source software (Lundell 2009), something which was identified as a myth ten years earlier by Tim O’Reilly (O’Reilly 1999). This shows that it is important to recognise the professional involvement in Open Source communities and that companies can provide professional support. Of particular interest to our work is that companies can provide different types of support to an Open Source project, including “the participation to online forums in order to keep the community alive by answering to users and customers” (Capra et al. 2009).

Earlier research has had limited focus on professional involvement in Open Source communities (Capra et al. 2009). One notable exception is the study of the Maemo platform (Ghosh 2006), which reported the extent to which different companies contributed to the project.

In this study we have access to informative data in the form of email addresses of mailing list contributors. We also have access to explicit lists of core developers and providers at the website of the project. Given this information, the purpose of our work is to elaborate approaches for studying how these groups of individuals contribute to the community around the mailing lists. As a motivation, it has been suggested that an organisation planning to professionally engage with Open Source needs to assess “the health of an OS product’s ecosystem” as part of their development of an Open Source strategy (Watson et al. 2005).

2 The Nagios project

Nagios is a tool for monitoring IT infrastructure that has been used in many professional organisations and mission critical systems. For example, Toland et al. (2007) found “the Nagios availability tool to be effective in the proactive support of mission-critical radiology and other clinical imaging systems.”. It has received several awards over the years, e.g. InfoWorld’s Best of Open Source Award for both 2008 and 2009. In a recent survey at “thegeekstuff.com” (September 2009), Nagios was voted as by far the most popular Open Source monitoring tool. Nagios has a large base of users, with many active contributors to the project mailing lists. The

core developers of Nagios are explicitly listed at the Nagios website (www.nagios.org/development/teams/core, accessed on 21 December, 2009). Furthermore, Nagios has a number of recognised professional providers. There are two preferred service providers (OP5, and Wuerth Phoenix) and 15 additional service partners for Nagios listed at the project website (www.nagios.org/support/servicepartners, accessed on 21 December, 2009). The explicit lists of core developers and providers are utilised in the result analysis.

Table 1 shows the more recent version history of Nagios, which was derived from the website of the project (www.nagios.org/news, accessed on 21 December, 2009). The first commits on Nagios were performed as early as in the end of 2001.

Table 1. Events in the more recent version history of Nagios.

Event	Date
V2.0	Feb 2006
V2.2	Apr 2006
V2.3, V2.4, V1.4	May 2006
V2.6	Nov 2006
V2.8, V3.0a1	Mar 2007
V2.9	Apr 2007
V3.0b1	Jul 2007
V2.10	Oct 2007
V3.0rc1	Dec 2007
V2.11, V3.0	Mar 2008
V3.1.0	Jan 2009
V3.0.2, V2.12	May 2009
V3.2.0	Aug 2009

In order to illustrate the active community of Nagios, Figure 1 shows number of messages for 22 different and overlapping time windows for the user- and developer mailing lists of Nagios. It is evident that there is much more activity in the user list. In fact, the level of activity in these two mailing lists of Nagios is in the same range as the activity of the corresponding mailing lists for the Apache webserver project for a six month period (Gamalielsson et al. 2009).




Figure 1. Number of messages for different and six-month time windows starting at the month shown at the x-axis. “Use” denotes the user mailing list for Nagios, and “Dev” denotes the developer mailing list.

3 Research Approach

Data was collected from the GMANE (gmane.org) archives of the SourceForge mailing lists for Nagios for the period from January 2004 to September 2009. The GMANE export interface was used to retrieve the raw mbox files for the project. For a message, the “message-id” field was used to get a unique message identifier. The “in-reply-to” field of a message was used, which contains the identifier for the message it was a reply to. Each message also has a “from” field from which the name and email address of the sender were derived. Finally, the “date” field was used to obtain the time and date of the message.

Data cleaning was performed prior to deriving the social networks to make sure that the same person does not appear several times using different identifiers. The approach for doing this is similar to the approach adopted by Kamei et al. (2008). Senders with different names but same email address were treated as the same person. If the same name (or slight variations of the same name) appears with different email addresses it is treated as one person, if the message content confirms this.

Two mailing lists were used: the “Nagios-devel” list, intended for development related issues; and the “Nagios-users” list, intended for problems related to the use of the software. An undirected edge A-B in the network is created if a person B replies to a message earlier sent by a person A. The weight of the edge is defined as the number of times A and B have interacted in the mailing lists used to derive the network, where each edge may account for several interactions between the connected nodes. However, in our analysis the edge weight is currently not used in the calculations.

The betweenness centrality measure has been chosen in several previous studies as a means of assessing the health of a social network representing a community (Kamei et al. 2008, Martinez-Romo et al. 2008). In our study it was calculated for each of the nodes (individuals) in the network. This measure takes on values in the interval $[0,1]$, and quantifies the ability of a node to act as a mediator in the network (Kamei et al. 2008). More precisely, betweenness centrality reflects the number of shortest paths that pass through a specific node. The loss of a node with high betweenness centrality may therefore disconnect parts of the network that the node “glues” together. Betweenness centrality is described in mathematical detail in Kamei et al. (2008).

4 Results

In this section we present different approaches for establishing ecosystem health that may be of a particular interest to a potential adopter in a professional Open Source context. In so doing, we specifically illustrate professional stakeholders in the Nagios community.

Figure 2 shows an early social network derived from six months of email correspondence (before release of Nagios V2.0, from April to September 2004) in the Nagios developer mailing list, and Figure 3 shows the corresponding network for the most recent six month period (April to September 2009). It can be observed that for both periods, the network has a distributed topology indicating no single point of “failure”. The networks for the user mailing list are also distributed, but have more nodes. The Nagios networks (Figure 2 and 3) have a shape similar to the corresponding networks for the Apache webservice project (Kamei et al. 2008).




Figure 2. Social network derived from Nagios developer mailing list Apr-Sep 2004.




Figure 3. Social network derived from Nagios developer mailing list Apr-Sep 2009.

Simple metrics showing the size of an active community is the number of nodes and edges in the social network. By calculating these metrics for different time windows it is possible to study the community dynamics. Figure 4 illustrates this for the Nagios developer mailing list. Here we can observe a positive trend with a long term growth in both number of nodes and edges. As expected, there is a clear correlation between number of nodes and edges (Pearson correlation of 0.92).



Figure 4. Statistics for social networks derived for different time windows for the Nagios developer mailing list

As earlier mentioned, the mbox files used to derive the social networks contain email addresses for the contributors. This data makes it possible to analyse the domain origins of messages, and assess the involvement of different groups of individuals. As an example, Table 2 shows the top nine email domains for contributed messages in the developer mailing list over time. Each row represents a six-month time window starting at the month specified in the leftmost column. It can be observed in this case that the “com” domain is most dominant, followed by “org” and “se” when studying all messages accumulated from January 2004 to September 2009.

Table 2. Top nine domains for the Nagios developer mailing list. Numbers in brackets represent percentage of messages. The last row shows the top nine domains when accumulated over time.

Win	#1	#2	#3	#4	#5	#6	#7	#8	#9
Jan-04	com(29)	org(17)	net(14)	se(8)	uk(8)	us(4)	au(4)	ru(3)	de(2)
Apr-04	com(28)	se(13)	org(12)	net(11)	us(9)	uk(9)	de(4)	au(3)	ru(3)
Jul-04	com(33)	se(19)	org(13)	us(9)	net(6)	au(5)	de(5)	uk(2)	br(2)
Oct-04	com(30)	se(24)	org(16)	us(7)	de(6)	au(5)	edu(2)	fr(1)	pt(1)
Jan-05	com(27)	se(23)	org(17)	au(6)	de(6)	us(5)	net(4)	pt(2)	edu(2)
Apr-05	com(34)	se(20)	org(14)	net(9)	de(5)	au(5)	fr(2)	hu(2)	no(2)
Jul-05	com(37)	se(18)	org(15)	net(9)	de(7)	fr(2)	no(2)	hu(2)	gov(1)
Oct-05	com(23)	org(17)	se(17)	de(13)	net(10)	br(4)	info(2)	uk(1)	tm(1)
Jan-06	com(22)	org(19)	net(12)	se(11)	de(11)	br(5)	au(2)	nl(2)	no(1)
Apr-06	com(29)	org(20)	net(11)	de(9)	se(6)	edu(4)	au(4)	br(3)	nl(3)
Jul-06	com(39)	org(16)	de(10)	net(10)	se(7)	edu(4)	us(2)	hu(2)	nl(2)
Oct-06	com(34)	org(22)	de(15)	net(8)	se(6)	edu(3)	ca(2)	at(2)	hu(2)
Jan-07	com(32)	org(21)	de(18)	net(7)	se(4)	edu(3)	ca(3)	at(2)	br(1)
Apr-07	com(35)	de(18)	org(17)	se(6)	net(6)	ca(2)	edu(2)	br(2)	at(2)
Jul-07	com(29)	org(21)	se(14)	de(14)	net(4)	dk(2)	edu(2)	ru(2)	ca(1)
Oct-07	com(27)	org(19)	de(16)	se(12)	net(4)	dk(3)	edu(3)	uk(2)	fr(2)
Jan-08	com(38)	de(17)	org(12)	se(6)	fr(4)	net(3)	ca(3)	uk(3)	dk(3)
Apr-08	com(43)	de(16)	se(10)	org(7)	ca(5)	net(4)	fr(4)	br(2)	uk(2)
Jul-08	com(37)	de(19)	se(14)	org(8)	ca(5)	net(3)	fr(3)	br(2)	edu(1)
Oct-08	com(38)	de(18)	se(10)	org(9)	ca(7)	br(5)	fr(3)	net(2)	edu(1)
Jan-09	com(41)	se(14)	de(13)	org(11)	ca(6)	br(5)	net(3)	fr(3)	edu(1)
Apr-09	com(44)	org(15)	se(14)	de(10)	net(4)	ca(4)	br(3)	fr(2)	at(1)
Acc.	com(34)	org(15)	se(12)	de(12)	net(6)	ca(2)	br(2)	edu(2)	fr(2)

As an example of how the top domains can be analysed in more detail, Table 3 shows the top five subdomains for the top five domains accumulated for all windows over time. It can be noted that for several of the top domains there are some very dominant subdomains like “nagios” for the “org” domain and “op5” for the “se” domain.

Table 3. Top five subdomains for the accumulated top five top domains in the Nagios developer mailing list.

Top dom	#1	#2	#3	#4	#5
com (34)	gmail(8)	opservices(2)	zango(2)	altinity(1)	ena(1)
org(15)	nagios(11)	ldschurch(2)	clewett(<1)	gmane(<1)	debian(<1)
se(12)	op5(12)	forumsq(1)	cendio(<1)	dokad(<1)	iis(<1)
de(12)	process-zero(4)	netways(1)	gmx(<1)	consol(<1)	ederdrom(<1)
net(6)	seanius(1)	lordsfam(<1)	gmx(<1)	netshel(<1)	elan(<1)

Similarly, it is possible to derive domain trees using the email addresses in the messages for a mailing list. This can be particularly informative if there are several levels of subdomains in an email address.

The availability of mailing list data containing domain information makes it possible to assess involvement of a certain group of stakeholders in an OSS ecosystem. As an example of this, Figure 5 shows the percentage of messages contributed over time by the group of explicitly listed core developers in the Nagios project. The results show that core developers in this project communicate more on development related issues than on user related issues, and that they with some fluctuation on average contribute to about one third of the development related correspondence.




Figure 5. Percentage of messages contributed by the core developers for the developer- and user mailing lists of Nagios.

Another specific group of stakeholders are the providers. The percentage of messages contributed by providers for Nagios is illustrated in Figure 6. It is evident that the providers contribute more to the developer mailing list, and it can also be noted that the shape of the curves is quite similar (Pearson correlation of 0.59). This is partly due to the fact that one individual of the most influential provider (OP5), who contributes to both lists, is the most active contributor in the set of providers.




Figure 6. Percentage of messages contributed by providers in the developer- and user mailing lists.

At a more intricate level, we can rank all nodes with respect to betweenness centrality for the social networks derived from messages contributed during the 22 different time windows from January 2004 to September 2009. This is exemplified in Table 4 which shows what individuals that are the top 10 mediators in the social network derived from the developer mailing list for different six-month time windows. The number in each table cell (except for those cells in the rightmost column) represents a unique identifier for an individual. As an example, individual 1 is top ranked in the Jan-04 window, ranked second from Apr-04 to Jul-05, and ranked first from Oct-05 to Jan-08. It can be argued that this kind of table reflects the sustainability of the community that actively participates in mailing lists. The rightmost column shows the number of new mediators in the top 10 of a window, which have not been in the top 10 before in the table. One general observation is that there is a constant addition of new mediators over time and that there are only a few individuals that contribute over all time windows (mediator 1 and 3). Mediators at lower ranks (towards 10) also tend to exhibit lower endurance over time. For example, mediator 12 is ranked as the 8th most important mediator for the two time windows starting in April 2004 and January 2005, and is thereafter not amongst the ten most important mediators.

Table 4. Mediator index at ranks 1 to 10 for the developer mailing list.

Win	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	N _{new}
Jan-04	1	2	3	4	5	6	7	8	9	10	10
Apr-04	3	1	6	10	11	5	4	12	13	14	4
Jul-04	3	1	11	15	6	10	5	16	13	17	3
Oct-04	3	1	11	15	16	18	17	19	5	20	3
Jan-05	3	1	7	11	18	20	15	12	5	21	1
Apr-05	3	1	7	5	20	22	23	18	24	21	3
Jul-05	3	1	25	22	23	26	27	28	24	29	5
Oct-05	1	3	20	30	31	32	33	26	34	29	5
Jan-06	1	3	20	11	35	31	33	36	37	32	3
Apr-06	1	3	38	39	40	36	11	41	31	42	5
Jul-06	1	43	31	24	3	40	44	45	46	38	4
Oct-06	1	3	44	47	45	46	37	33	48	49	3
Jan-07	1	33	3	45	44	47	50	37	51	31	2
Apr-07	1	33	3	45	52	50	48	53	54	31	3
Jul-07	1	3	33	52	45	55	24	56	57	53	3
Oct-07	1	3	33	45	55	24	58	59	44	60	3
Jan-08	1	3	33	45	61	24	62	31	63	59	3
Apr-08	3	1	33	45	47	61	44	16	31	64	1
Jul-08	3	33	1	45	64	44	16	58	31	65	1
Oct-08	3	45	33	1	61	66	31	65	58	67	2
Jan-09	3	33	45	1	65	68	61	69	70	63	3
Apr-09	3	1	33	71	65	44	72	73	74	16	4

When studying a particular group of stakeholders and the ability to act as a mediator, it may be better to study the ranking of each of the individuals in the group. Table 5 illustrates this for the six core developers in the developer mailing list over time. Like in Table 4, this kind of information gives an indication of sustainability, but here with respect to a particular group of interest. In this particular case it can for example be noted that CD1 and CD2 have been very important mediators in the developer mailing list community since it started, and that CD4 has become equally important since January 2007. Results of this kind can also be derived for the provider list.

Table 5. Mediator rank for the six core developers (CD1-CD6) in the developer mailing list. N denotes total number of contributors in a specific time window

Win	N	CD1	CD2	CD3	CD4	CD5	CD6
Jan-04	69	1	3	57	-	-	-
Apr-04	76	2	1	66	-	-	-
Jul-04	83	2	1	-	-	-	-
Oct-04	93	2	1	-	-	-	-
Jan-05	88	2	1	-	-	-	-
Apr-05	74	2	1	-	-	-	-
Jul-05	71	2	1	-	-	-	-
Oct-05	96	1	2	25	7	-	-
Jan-06	104	1	2	25	7	-	-
Apr-06	87	1	2	14	18	-	-
Jul-06	69	1	5	7	16	-	-
Oct-06	84	1	2	3	8	-	-
Jan-07	112	1	3	5	2	9	-
Apr-07	121	1	3	13	2	26	-
Jul-07	131	1	2	39	3	126	-
Oct-07	135	1	2	9	3	-	125
Jan-08	118	1	2	11	3	-	114
Apr-08	108	2	1	7	3	104	-
Jul-08	116	3	1	6	2	108	-
Oct-08	114	4	1	17	3	106	-
Jan-09	105	4	1	16	2	41	25
Apr-09	111	2	1	6	3	35	20

5 Conclusion and discussion

In this paper we have presented an extended quantitative analysis of the Nagios community by performing domain- and social network analysis of mailing lists. It can be concluded from our study that there is an increasing ability to develop useful approaches and metrics to cover the broader aspects of the OSS ecosystem. This means that we have shown that not only the user and developer role can be analysed, but also the provider role. This is an important contribution that can offer a more comprehensive understanding of professional involvement in professional Open Source.

Specifically, the way domain analysis has been performed in the context of mailing lists contains elements of novelty. We have proposed approaches for showing how a specific group of stakeholders contribute to a mailing list community. In particular, we have contributed approaches for analysing the sustainability of a mailing list community or a specific group of stakeholders in a community in terms of mediation of information when studying social networks derived from mailing list data over time.

One limitation of the study is that no account was taken of the content of email responses, so that each response was equally weighted. A clearer indication of the

nature and value of responses would add value to such analyses, but would require techniques beyond the scope of this study.

Since quantitative research on OSS to a large extent is data driven, the results can only be as informative as the available data sources permit. OSS projects that would like their projects to be taken up more broadly should consider providing as informative data as possible. In fact, the community as a whole would benefit from a continuous joint effort to make available informative data sources in order to be able to cover broader aspects of the OSS ecosystem.

In healthy Open Source communities people are active and responsive to questions during the life cycle of a software system. It is important to consider such indicators of health prior to any organisational adoption or during the seeding of a community. The kind of analyses elaborated in this paper serve as important means for establishing the health of an Open Source community.

We are currently working on developing approaches for obtaining valuable information about professional involvement in Open Source communities. Such information is essential for any professional organisation wishing to better understand the broader Open Source ecosystem.

6 Acknowledgement

The authors would like to thank their colleagues in the OSA-project for their encouragement and support. The OSA-project is financially supported by the Knowledge Foundation (“KK-stiftelsen”).

References

- Capra, E., Francalanci, C., Merlo, F. (2009). A Survey on Firms’ Participation in Open Source Community Projects, In Proceedings of the Fourth Conference on Open Source Systems (OSS 2009), pages 225–236.
- Crowston, K. and Howison, J. (2005). The social structure of Free and Open Source software development. *First Monday*, 10(2).
- Crowston, K. and Howison, J. (2006). Assessing the Health of Open Source Communities. *IEEE Computer*, 39(5):89-91.
- Crowston, K., Wei, K., Li, Q. and Howison, J. (2006). Core and Periphery in free/Libre and Open Source software team communications. In Proceedings of the 39th Hawaii International Conference on System Sciences, page 118.1.
- Fitzgerald, B. (2006). The Transformation of Open Source Software, *MIS Quarterly*, 30 (3)
- Gamalielsson, J., Lundell, B. and Lings, B. (2009). Social Network Analysis of the Nagios project. Open Source Workshop (OSW 2009), Skövde, Sweden, October 15-16 2009.
- German, D. (2004). The GNOME project: a case study of open source global software development. *Journal of Software Process: Improvement and Practice*, 8(4): 201–215.
- Ghosh, R. A. (2006). Study on the: Economic impact of open source software on innovation and the competitiveness of the Information and Communication Technologies (ICT) sector

- in the EU, http://ec.europa.eu/enterprise/sectors/ict/files/2006-11-20-flossimpact_en.pdf, accessed 23 December 2009.
- Kamei, Y., Matsumoto, S., Maeshima, H., Onishi, Y., Ohira, M. and Matsumoto, K. (2008). Analysis of Coordination Between Developers and Users in the Apache Community. In Proceedings of the Fourth Conference on Open Source Systems (OSS 2008), pages 81-92.
- van der Linden, F., Lundell, B. and Marttiin, P. (2009). Commodification of Industrial Software: A Case for Open Source. *IEEE Software*, 26 (4): 77-83.
- Lopez-Fernandez, L., Robles, G., Gonzalez-Barahona, J. M. and Herraiz, I. (2006). Applying Social Network Analysis Techniques to Community-driven Libre Software Projects. *International Journal of Information Technology and Web Engineering*, 1:27-48.
- Lundell, B. (2009). Being Open about the ten Open Source myths: Myth #2 is still alive!, Lightning talk presented at OpenMind 2009, Tampere, Finland, September 30th 2009.
- Lundell, B., Forssten, B., Gamalielsson, J., Gustavsson, H., Karlsson, R., Lennerholt, C., Lings, B., Mattsson, A. and Olsson, E. (2009). Exploring Health within OSS Ecosystems, First International Workshop on Building Sustainable Open Source Communities (OSCOMM 2009), Skövde, Sweden, June 6th 2009.
- Martinez-Romo, J., Robles, G., Ortuño-Perez, M. and Gonzalez-Barahona, J. M. (2008). Using Social Network Analysis Techniques to Study Collaboration between a FLOSS Community and a Company. In Proceedings of the Fourth Conference on Open Source Systems (OSS 2008), pages 171-186.
- Mockus, A., Fielding, R. T. and Herbsleb, J. D. (2002). Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology*, 11 (3): 309-346.
- Moon, Y. J. and Sproull, L. (2000). Essence of distributed work: The case of the Linux kernel. *First Monday*, 5 (11).
- O'Reilly, T. (1999). Ten myths about open source software. Transcript of talk given to a group of Fortune 500 executives in October 1999, <http://www.oreillynet.com/lpt/a/2019>, accessed 22 December 2009.
- Toland, C., Meenan, C., Warnock, M. and Nagy, P. (2007). Proactively Monitoring Departmental Clinical IT Systems with an Open Source Availability System. *Journal of Digital Imaging*, 20:119-124.
- Watson, R.T., Wynn, D. and Boudreau, M.-C. (2005). JBOSS: THE EVOLUTION OF PROFESSIONAL OPEN SOURCE SOFTWARE, *MIS Quarterly Executive*, 4 (3):329-341