

Networks of Innovation

Change and Meaning in the Age of the Internet

Ilkka Tuomi

Feb 20, 2002

List of Figures.....	6
List of Tables.....	7
I. Introduction	8
II. Innovation as multifocal development of social practice	18
Putting the user in focus.....	21
Use as meaningful practice	23
Production as an end	26
Investment and invention of meaning	28
Community as the locus of practice	34
Interpretative flexibility and ecology of social practices	36
Social drivers of innovation.....	39
Individual exploration	42
Spaces of novelty.....	43
Dynamics of networked innovation spaces	48
III. Inventing the Web.....	57
The first WorldWideWeb proposal.....	60
State of the art: KMS.....	63
Architecture of the WorldWideWeb	67
Mobilizing resources	70
The vision of Xanadu	74
Sources of success.....	77
IV. Making of the Internet	82
Laying the infrastructure.....	83
Networking the world	86
Competing technologies.....	91
Message-packets and resilient networks: innovation at RAND	98
Sources of innovation	102
Time-sharing and network society: work at NPL.....	107
Interactive computing: augmenting the human mind.....	116

Memex and the oNLine System	121
Time-sharing and on-line communities	127
IPTO: translating ideas into money and technology	130
The Arpanet	133
V. Analysis of the early phase of Internet development.....	139
Technological frames	140
Augmentation systems frame	143
On-line communities frame.....	144
Communication networks frame	147
Electronic services frame.....	148
Resource mobility in the early phases of Internet history.....	150
VI. Socio-cognitive spaces of innovation and meaning creation	156
Thought collectives	160
Speech genre and chronotope	165
Communities of practice	168
Social learning in communities of practice.....	170
The concept of <i>ba</i>	172
VII. Breaking through a technological frame	180
Two evolutionary paths of communities.....	187
Development of specialization, division of labor, and new technological frames..	188
Combinatorial innovation in an ecology of communities.....	192
Layered <i>ba</i> and combinatorial innovation	197
VIII. Combination and specialization in the evolution of Internet	203
Email as a combinatorial innovation.....	203
ARPANET ecology and the evolution of the Network Working Group	208
RFC's and the four waves of Internet development	212
Organization and community ecology	216
Spin-off and resource combination in the NWG.....	219

IX.	Retrospection and attribution in the history of Arpanet and Internet	225
	“The first paper on packet-switching theory”	227
	Reconstructing the Internet.....	230
X.	Learning from Linux	237
	Evolution of Linux	238
	Basic architectural components of Linux	241
	Growth of Linux source code	243
	Linux developer community	246
	Control and coordination in the Linux community.....	253
	Co-evolution of social and technical structure	254
	Controlling the kernel	257
	Balancing control and innovation in the Linux kernel.....	260
	Sedimentation, translation, and reduction of complexity	266
	Sedimentation of source code.....	268
	Multi-faceted translation and module interfaces	270
	Quality control, Linus’s Law, and the ecology of bugs	272
	Resources in the bug removal process	275
	Rules, regulations, and intellectual property	282
	Forms of open source copyright.....	284
	Social contracts.....	290
	Developer incentives and resource allocation	292
	The Tulip flame war.....	296
XI.	Concluding remarks	304
	Linux as modern economy.....	304
	Hierarchy of innovation	312
	New economy	314
	The road ahead.....	318
XII.	References	322
XIII.	Name Index	322

XIV. Subject Index.....	355
last page:	390

List of Figures

Figure 1. WWW servers connected to the Internet.	10
Figure 2. Changing production-consumption linkages.	50
Figure 3. How the "World Wide Web" became "the web."	80
Figure 4. Investment cost of transatlantic cable per minute of use.	90
Figure 5. Total number of transatlantic voice circuits.....	91
Figure 6. Revenues from International Service in the U.S., 1950-1997.	94
Figure 7. Total share of international revenues in the U.S., 1950-1997.	95
Figure 8. U.S. Postal Service, pieces of mail handled.....	96
Figure 9. Some key loci of innovation in the evolution of Internet, 1940-1975.	142
Figure 10. The general structure of thought collectives.....	164
Figure 11. The basic SECI-model of knowledge creation.	173
Figure 12. Bakelite as a structural coupler.....	195
Figure 13. A layered <i>ba</i>	200
Figure 14. Layers of a systematization <i>ba</i>	201
Figure 15. The waves of Internet technology development.....	214
Figure 16. Division of labor in the early phases of Internet development.	219
Figure 17. Production of fractal community structure.	223
Figure 18. The "Morris worm" in newspapers 1988-1990.	232
Figure 19. The Morris Internet worm in major newspapers, 1988-1999.	233
Figure 20. Early Internet awareness.	234
Figure 21. Internet articles in Financial Times, 1994.....	235
Figure 22. Layers of a Unix operating system.....	242
Figure 23. Linux kernel distributions.	243
Figure 24. New source code directories.	245
Figure 25. Removed directories as a function of time.	246
Figure 26. Number of people in the Linux CREDITS file, 1994-2002.	248
Figure 27. Location of active contributors, March 1994.....	250
Figure 28. Geographical expansion of development activity, March 1994-March 2000.	251
Figure 29. Linux architecture.....	258
Figure 30. Growth of Linux core kernel components.	263

Figure 31. Growth of architecture dependent and extensible components.	265
Figure 32. Resource and object development paths.	269
Figure 33. Standard interface as a multi-faceted translation mechanism.....	271
Figure 34. The basic bug removal cycle.	278
Figure 35. Mediated interactions in the bug removal process.	281
Figure 36, License type distribution in the Linux Software Map database, 1992-2002.	289

List of Tables

Table 1. Augmentation frame, 1965.....	144
Table 2. On-line communities frame, 1965	145
Table 3. Communication networks frame, 1965.	147
Table 4. Electronic services frame, 1965.	149
Table 5. Kernel bug management resources.	280
Table 6. License types in the Linux Software Map database, 2002.	288

I. Introduction

According to user surveys, the Linux operating system is rated as the best operating system available. It is considered to be more reliable than its main competitors. Its functionality is claimed to be better, and according to many experts, new releases of Linux implement innovative ideas faster than its competitors. In other words, it is argued that Linux development creates complex new technology better and faster than the biggest firms in software industry.¹

Yet, Linux also seems to break many conventional assumptions that underlie research on innovation and technological change. Linux is developed by an informal self-organizing social community. There is no well-defined market or hierarchy associated with it. Most of Linux development occurs without economic transactions. Instead of getting paid for their efforts, the developers often spend a lot of money and effort to be able to contribute to the advancement of the development project.

The open source development model, which underlies Linux, has attracted increasing attention in the last years. Today, Linux is considered to be a serious threat to Microsoft's market dominance in operating systems. More generally, open source development projects have in recent years had major impact in software and internet-based industries. For example, almost 60 per cent of Internet connected Web servers were open source Apache servers in October 2000. As can be seen from Figure 1, the second most popular Microsoft servers were about one third as popular with 20 per cent. Although Microsoft has gained market share with its Internet Information Server, at the end of 2001 about 63 per cent of active web sites were running Apache. The most common operating system in the web server machines was Linux.² Some open source projects, such as Sendmail, Perl, and Emacs, have achieved large user bases, making it difficult for commercial enterprises to enter the market.

¹ <http://www.uk.linux.org/LxReport.html>

² Source: Netcraft, <http://www.netcraft.com/survey/>. For a discussion on server market shares, see Netcraft and Peeling and Satchell (Peeling & Satchell, 2001).

Linux has been developed in the open source mode to a large extent because the Internet itself was to a large extent developed in this same mode. The collaborative and participatory development model gained visibility in the mid-1960's, when the early users of time-shared computers realized that collaboration often produced unexpected benefits. The predecessor of the modern Internet, ARPANET, was created in this mode, and many critical contributions, such as Internet email, Usenet news, and the World Wide Web emerged as a result of open collaboration. The Internet Engineering Task Force, which defines standards for the Internet, has also used an open source approach since its formation in 1986 (Bradner, 1999).

Several commercial software firms have recently tried to adopt aspects of the open source model. For example, Netscape announced in 1998 that it would distribute the source code of Netscape Communicator with open source license. IBM decided to use the open source Apache server as the core of its Web server offers. Red Hat, SuSE, Caldera, and other new economy firms, in turn, thrive on packaging Linux distributions and by producing added value for Linux users. Sun Microsystems has used a version of the open source model to support development of its Java and Jini platforms. After launching an attack on Linux in 2001, also Microsoft declared that it will have its own Shared Source Philosophy, which was aimed at making open source development possible without losing intellectual property rights. In all these cases, business firms are experimenting with ways to benefit from innovation that occurs in the open source communities. Instead of traditional economic competition, such initiatives rely on symbiotic relationships, and on the willingness of developer communities to collaborate.




Figure 1. WWW servers connected to the Internet.³

In much of the innovation literature, innovation is defined as something that has economic impact. Linux and other open source initiatives show that this definition is problematic and possibly misleading in important practical cases. For example, during its history, most Linux development has occurred independent of direct economic concerns. It would be tempting to argue that Linux development is different from “economic activity” and something that, strictly speaking, should not be called innovation. Indeed, in its early history Linux development was not in any obvious way associated with changes in production functions, market competition, or appropriation of economic investment and surplus. Yet, Linux developers obviously collectively produce new technology. If economy is about collective production, this is it.

Linux, therefore, is an interesting test case for economic theories of innovation and technology development. For example, the history of Linux allows one to question to what extent existing economic models of innovation and technological development capture phenomena that underlie collective production of new technologies.

In very practical terms, Linux is an economically important phenomenon. Indirectly, the success of many new businesses, venture capitalists, investment funds, and individual investors critically depends on the productive activities of the Linux community. Today,

³ Source: Netcraft, <http://www.netcraft.com/survey/>.

many corporations, governments, public sector organizations, and individual developers are starting to deploy Linux to cut costs, promote interoperability, and avoid lock-in to proprietary systems. Yet, when we consider the entire history of Linux, the economic impact seems to appear almost as an afterthought and as a side effect of a long period of technology creation. Linux, therefore, provides an interesting history of globally networked innovation, illustrating the substance that underlies the discussions on the “new economy.” If the “new economy” is about global Internet-enabled and software-driven production, this is it.

More generally, the history of Internet-related innovations enable us to discuss those social and cognitive phenomena that underlie technological change. By studying such innovations, we can open some black boxes of innovation theory, including such widely used concepts as learning, capability, utility, and consumption. By observing the development of the Internet, we can describe the microstructure of innovation, and transcend the boundary between invention and innovation.

Although such studies have obvious consequences for innovation research in general, Internet related innovations are, however, also special. On the Internet the products of innovative activity are externalized as technological artifacts and documents that can be studied relatively easily. Never before has innovation and its results been recorded in such historical detail. On the net we live in dog years, but our memory is that of an elephant. There exists sufficient documentation so that we can—at least tentatively—describe some key principles that underlie the development of Internet related innovations. For a researcher on technological change, this is an exciting opportunity.

Internet related innovations are obviously important as the Internet has become a key technology in many areas of our everyday life. Below I will argue, however, that these innovations reveal important aspects of all innovative activity. Indeed, my key message is that the traditional models of innovation are often misleading, and that they are becoming increasingly misleading in the future. In practice, we have to move beyond abstract descriptions and ask what makes novelty meaningful. This leads to social and cognitive theories of innovation.

From a practical point of view, Internet related innovations also provide test cases for analyzing product development models and proposals for organizing for innovation. For example, the extensive use of modern communication and collaboration technologies in Linux development highlights some aspects of technology development that were not easy to see in earlier studies on innovation. Although I will not explicitly discuss organizational or policy implications below, I believe that the following chapters highlight several points which have such implications.

Linux, open source projects, and internet-related innovations may have developmental histories where collaboration and networking are more visible than in some earlier innovations. The open source model, however, obviously goes beyond software programming projects. As many commentators have observed, the process of science itself is very much based on peer-review, incremental development, non-economic motives, and geographically distributed collaboration. Indeed, traditional models of innovation often assumed that basic research generates ideas and technologies that are appropriated by entrepreneurs who turn them to products and money. The history of Linux and internet-related innovations enable us to see how the boundaries between basic and applied research are being transformed. Indeed, I will argue below that the distinction between basic and applied research needs to be reconsidered.

From the very beginning, the Internet has been used to distribute work and its results. Division of labor is the foundation of all societies; the Internet, however, makes it possible in qualitatively new ways. A study on internet-related innovations, therefore, has implications when we try to understand the ongoing social transformation towards the network society. To give just one example: when NASA run its Clickworkers pilot where volunteer Internet users could mark craters on pictures of Mars, between December 2000 and June 2001 people marked over 1.9 million craters. Although each volunteer only marked a few craters, collectively their results were indistinguishable from those of a well-trained expert.⁴ This example is interesting as it shows⁴ that a trivial individual effort

⁴ <http://clickworkers.arc.nasa.gov/documents/crater-marking.pdf>: “Clickworkers results: crater marking activity.” July 3, 2001.

may lead to high-quality collective outcome. In a very concise form it shows one way by which a new balance may emerge in the network society between increasing specialization and network-enabled participatory decisionmaking. Internet-related innovations, therefore, have relevance both when we try to understand how new technologies are developed but also when we try to understand how technological development and social change could be linked in the future.

History is always constructed from the perspective and for the purposes of the present. A useful history, however, provides opportunities for more than one interpretation. Historical description, therefore, has to be rich enough in detail and it has to give room for multiple voices. Yet, a balance has to be found between details and conciseness. Reality is always richer than any of its descriptions. I have tried to solve this problem by combining relatively general conceptual arguments with outlines of specific innovation histories and more detailed in-depth case studies. Some chapters make rather controversial theoretical claims without extensive empirical support for these claims. Subsequent chapters, hopefully, fill in some of the details.

The first chapter introduces some main concepts and assumptions that underlie the present work. In effect, it tries to set the reader in a position where the subsequent discussion can make sense. It points out that innovation is fundamentally about social change, and that innovations emerge and become articulated when they are taken into meaningful use in social practice. It argues that meaningful use—as well as the meaning of technology itself—is grounded on social groups that can be called practice-related communities. As a result, innovation and technological change can be studied as phenomena that occur within an ecology of such communities. Construction of technology requires construction of meaning, and new technology is much more than improved functionality. Instead of the “upstream” of the traditional linear model of innovation, we have to focus on the “downstream” where social communication and change occurs. All innovation is social innovation. Innovation does not happen “out there” in the world of objects, but in the society and in the minds. More particularly, it happens in the minds of the users, which are intrinsically integrated with the activities of

the users. Those cultural and material resources that are available for the users, therefore, become key resources in the innovation process.

The second chapter is a quick first take on making these concepts more concrete. It illustrates the nature of innovation by outlining the history of the World Wide Web. It asks who invented the Web, what were the resources used in its invention, and what actually was invented in the process. Many of the details of this history are well known. Many accounts of the history of the World Wide Web, however, also show that some details of the story are often missed. These details become important when we try to understand innovations such as the World Wide Web.

The third chapter moves from the recent history back in time, describing the early phases of the evolution of the Internet. More exactly, the focus is on that point of time when computer networking was only an idea. The chapter introduces the historical data that will be used in subsequent chapters. Although there now exists excellent histories of the Internet, such as those written by Abbate (1999) and Naughton (2000), it is necessary to provide enough historical detail to make the origins of the Internet understandable. In the process, I will also make some notes that hopefully complement existing histories in interesting ways. The chapter describes how electronic communication systems evolved and laid conceptual and material foundations for computer networks. It also introduces leading actors who played key roles in the early phases of computer networking.

The fourth chapter summarizes the early history of the Internet and describes the various technological frames that generated the basic innovations of computer networking. In other words, it puts history in the context of technology and innovation studies. It also discusses resource mobility in the early phases of the Internet development. One main claim in the book is that innovation occurs when social practice changes. The mobility of resources, therefore, is a key factor in enabling and constraining innovation.

The fifth chapter returns to the topic of communities. It discusses several alternative theoretical traditions that have described the social basis of meaning, knowing, and knowledge creation. It starts by introducing the concept of thought community that was originally introduced by Ludwik Fleck (1979) in the 1930's. Fleck's historical study

described many of those social processes that underlie the emergence of new scientific knowledge and new technologies. The chapter further discusses Bakhtin's speech genres, cultural-historical activity theory, social learning in communities of practice, and the concept of *ba*. Ikujiro Nonaka and his colleagues have argued that innovation and knowledge creation occur in knowledge creation spaces, or *ba*'s. The chapter discusses the nature of *ba*'s, and links this concept back to its origins in the epistemological theory of Kitaro Nishida and the Kyoto School. The fifth chapter, therefore, introduces a set of alternative theoretical views that can be used to understand the cognitive and social basis of innovation.

One of the main arguments below will be that innovation can properly be understood only by studying the social basis of innovation. The heroic individual innovator is not a good model when we try to understand the evolution and development of technology. If knowledge and the meaning of technology is grounded in communities that reproduce existing social practice, as this book argues, it may seem, however, that innovation is a contradiction in terms. How is it possible that new social practices emerge when communities more or less by definition reproduce their current practices? How do we break technological frames and how new technological frames are created? Chapter six argues that there are two distinctive ways that new communities and new technological practices can emerge. One is based on increasing specialization, and the other on combination of existing resources. In other words, there exists two qualitatively different dynamics of innovation, and their analysis requires two different theoretical approaches. As a result of these two different modes of socio-technical evolution, the concept of *ba* can therefore be redefined. The chapter links the concept of *ba* to the sociocultural basis of knowledge, and proposes a new interpretation of Nonaka's knowledge creation model.

Using these theoretical concepts, chapter seven then returns to the history of the Internet. It briefly discusses email as an example of combinatorial innovation, and describes the evolution of the social structure that provided the basis for the creation of ARPANET and the Internet. It shows, for example, that both resource combination and evolution of specialization have played important roles in the development of social structure of Internet-related innovation communities. The current Internet community is in many

ways rooted to the Network Working Group, which started in 1968 as an informal group of computer students. Internet, itself, however, would not have been possible without combination of resources that came from outside this nucleus or the Internet culture.

Chapter eight picks up one aspect of this history, which is an interesting topic for both innovation studies and policy. This is the question of retrospection and attribution of authorship. If innovations are to an important part created by their users and the meaning of innovation is reconstructed from the present position, how should we read historical accounts that describe evolution of technology? And to whom should the credit go? Did Al Gore really invent the Internet? Or was he just doing what Rembrandt did: signing off works that, strictly speaking, were produced by others, but which could not have existed without him? Should Linus Torvalds get a patent on Linux? What, indeed, intellectual property means when technology development uses resources that are networked, cumulative, often unintended, and when adaptation of new technological opportunities depends on institutional change and competence development in the downstream? Should we reconsider the author, or is the confusion created by a wrong conceptualization of the products themselves? By analyzing newspaper articles that have discussed the Internet during the last fifteen years, we show how the common understanding of “the Internet” has evolved. As chapter eight shows, the heroes of innovation are mental reconstructions, but so is the technology itself.

Chapter nine, finally, returns to the case of Linux. It describes both social and technological evolution of Linux and its development community. For example, it shows how technological architecture and social structure co-evolve as technical problems are solved in the social domain and social problems are solved in the technical domain. By analyzing in detail the evolution of the structure of Linux source code over a period of years, it shows how social control and coordination become embedded in a technological artifact. It also shows how social interaction can be “translated” into resources by “black-boxing” some of the underlying complexity behind technological interfaces. The chapter argues that one reason why the open source development model has been successful is that the social translation mechanisms it uses allow several communities to simultaneously interface to a common technological artifact. Moreover, the open source

model guarantees that when software fails, it fails gracefully, at least in the social sense. In open source, black boxes have transparent and penetrable walls. The chapter also discusses the bug removal process in Linux and highlights some trade-offs that are needed to make distributed innovation and technology development effective.

The last chapter puts the open source model of technology development in a broader perspective, and discusses the cultural and value system that underlies open source. Indeed, it argues that a study on socio-cognitive basis of innovation leads to a new approach in economic theory, where the concept of value has to accommodate the idea that in innovation processes new meaning is created and new domains of social practice are generated. Such “expansive” theory of economics may lead to new insights when we formulate and study technology and innovation policy. The chapter also points out that the networked mode of production that underlies open source may lead to new dynamics in the socio-economic development as the social institutions that usually provide stability in socio-economic systems are constantly renegotiated in the network mode of development. The chapter also discusses the differences and similarities between the open source model and the Silicon Valley innovation system. The chapter finally points out some areas for further study, and ends with some concluding remarks.