Paolo Dini and Thanassis Tiropanis

The evolving dynamics of the internet layered archtecture: innovation, net neutrality and the interdependence of structure and function

Book section

Original citation:
Originally published in International Conference on Internet science, Brussels, April 9-11, 2013: conference proceedings. The FP7 European Network of Excellence in Internet Science, pp. 203-208.

The Evolving Dynamics of the Internet Layered Architecture

Innovation, Net Neutrality and the Interdependence of Structure and Function

Paolo Dini
Department of Media and Communications
London School of Economics and Political Science
Houghton Street, London, UK
p.dini@lse.ac.uk

Thanassis Tiropanis
Electronics and Computer Science
University of Southampton
tt2@ecs.soton.ac.uk

In this extended abstract we discuss the infrastructure of the Internet as two distinct layers that are increasingly interdependent: the network layer, below, and the Web, above. The network layer is responsible for the structural properties of the Internet, while the top layer is where the Web content is consumed through Web applications and communications. We argue that the principle of independence between the layers, which guided the early design of the Internet, has enabled an environment conducive to innovation pursued by different communities: networking technologists primarily below, and business, application developers, and humanists/artists above. The innovations in each layer, however, are increasingly influencing design criteria and choices in the other, suggesting that the Internet architecture is evolving towards greater interdependence between the layers. For example, this is explicitly sought in some cases for wireless networks for efficiency optimization. Tight coupling between structural and functional properties is one of the fundamental "architectural" principles of biological organisms, which have evolved to optimize energy efficiency as a requirement for survival and procreation. This view, which mixes strictly functionalist concerns with creative and opportunistic behaviour, suggests that the Internet may be evolving towards an increasingly complex structure and dynamic. The paper argues that an environment in which the two layers are increasingly interdependent can still sustain a high level of innovation as long as no entity has full control of both, and as long as the design principles on each layer, which can be argued to have been fostering innovation, are not changed. We argue that the original Internet has fostered a number of innovations including the Web, P2P applications, and the Cloud and that its potential for innovation could be compromised if the importance of net neutrality and its infrastructural characteristics are undermined.

Keywords: Innovation, net neutrality, layered architecture, tight coupling

I. Introduction

The idea of this paper originated from a conversation between the authors over the course of several meetings of the EINS project.[1] Both authors come from engineering, albeit of different kinds,[2] are still active in the hard sciences, and have been studying social science for the past several years. Our conversations, therefore, have been refreshingly free from the problem usually encountered when speaking with someone whose work is rooted in "the other" disciplinary domain: this is the problem of not being able to find the right words to express an important concept because the other person lacks the ontological, epistemological, and methodological framework to decode and understand the points being made. It is fortunate that since the emergence of Web 2.0 phenomena this problem has been gradually dissipating, as far as the socio-technical interface of interdisciplinary science is concerned, with the result that computer scientists and software engineers are well ahead of the other "hard" sciences in understanding and relating to social phenomena. For example, unlike 10-15 years ago, most computer scientists today are familiar with the concept of social construction. But there are still many opportunities for diverging views in economic and political discussions.

In this extended abstract we begin to sketch the main points of a long-term study of some of the "interdisciplinary entanglements" that increasingly characterize the Internet and an emerging Internet Science. The points raised are meant only as signposts of more in-depth and more nuanced discussions to be pursued during the course of this study. We will analyse the Internet from two viewpoints that for convenience we can associate with the terms "Innovation" and "Net Neutrality".

The first entanglement involves the interaction between the Internet's architecture and socio-technical innovation dynamics. Simplifying the 7-layer OSI stack model or the 4-layer TCP/IP stack model to just 2 layers – network below and Web above – we argue that the original engineering criterion of independence and modularity between the layers facilitated innovation dynamics in the early Internet; but that, ironically, the same innovation dynamics are leading to a progressively *greater interdependence* between the two layers. When interdependence takes the form of interaction only through an agreed interface between otherwise separate modules it is termed 'loose coupling'. This is one of the building blocks of object-oriented software engineering and of its online extension to service-oriented computing or architecture (SOC/SOA) (Papazoglou 2003). Biological systems, by contrast, have evolved opportunistically to optimize their efficiency under scarse energy resources (food), leading to multifunctionality and tight functional interdependence,[3]

[1] http://www.internet-science.eu

[2] Paolo Dini's original background is in aerospace engineering and physics, Thanassis Tiropanis's background is in software engineering and computer science.

[3] The specific functional interdependence we are referring to here goes well beyond function calls to and from inside given modules. In biology the temporal evolution of a given subunit viewed as a discrete and finite state machine depends entirely on the other subunits it is coupled to *themselves* transitioning to new states (there is no CPU). Depending on the interconnection topology and the number of such interdependent modules one

which we can term 'tight coupling'. In this paper we argue that in some respects the evolution of the Internet towards greater interdependence between the layers reflects some aspects of the functional interdependence or tight coupling of biological systems. Although this is an interesting development from a purely technical viewpoint, it comes with some dangers that we would like to analyse in more depth.

For example, the second entanglement we study, which has been and continues to be widely discussed in the literature, concerns what could be regarded as the most important issue in Internet governance: net neutrality. Net neutrality involves the interaction between the technical management of information and content and the political and market forces that are vying to influence or control the technical design of the Internet (Lessig 2006). The main point the paper argues is that as the interdependence between the two layers of the Internet increases it becomes increasingly important to maintain net neutrality if we wish to retain the ability to innovate.

The points outlined above are linked to some ideas the first author started to work on some years ago but never published beyond the stage of a EU project deliverable (Berdou and Dini 2005), except for Figure 3 which has appeared, through a different argument, also in Dini et al. (2011). In reference to the concepts discussed in Dini & Sartori (2013), the interdependence between the two layers of the Internet can be conceptualized either from a systemic point of view, where an emphasis on language overshadows the role of the individual, or from a more empirical and case study-oriented perspective, which necessarily depends on the analysis of individual initiatives, interests, and motivations.

To argue our point we therefore follow two strategies. First, in Section II we develop a language-based systemic model of the socio-technical Internet phenomenon as a self-reinforcing feedback loop that transcends disciplinary boundaries and offers a possible synthesis of very different disciplinary perspectives. This model spans both the loose coupling and the tight coupling scenarios. Second, in Section III we discuss examples of how innovation in the early Internet was made possible by the modularity and independence between the layers. And in Section IV we discuss examples that show that recent innovation trends are pushing the layers towards ever-greater interdependence. Finally, In Section V we use these different perspectives to argue that net neutrality remains one of the most important architectural principles of the Internet and that the infrastructural nature of the Internet and of the Web needs to be safeguarded to foster further innovation in each of these layers.

II. THE MEDIA STACK

The nested media of the OSI stack are layered in order of increasing abstraction (see Figure 1). Although each layer is not in general a formal transformation of the layers adjacent to it, it certainly can be, as exemplified by software radio or by ASICs (application-specific integrated circuits), which are first implemented as programs and then transposed to logic gates on silicon. This transformational property of ICTs is a consequence of their being formal systems ultimately equivalent to the same abstract machine (the most general for the currently accepted computing paradigm being the Turing machine[4]).

As shown in Figure 2, if we turn the media stack on its side we can arrange different technologies from most concrete to most abstract right to left, culminating with formal languages. As we approach natural languages two interesting things happen: we encounter a boundary beyond which we cannot develop a formal model, and the medium and content converge. Furthermore, as we approach natural language it becomes increasingly difficult to commodify the technology. This becomes clear through the simple observation that society could not function if every spoken utterance were copyrighted and money were exchanged between listener and speaker according to some contract.

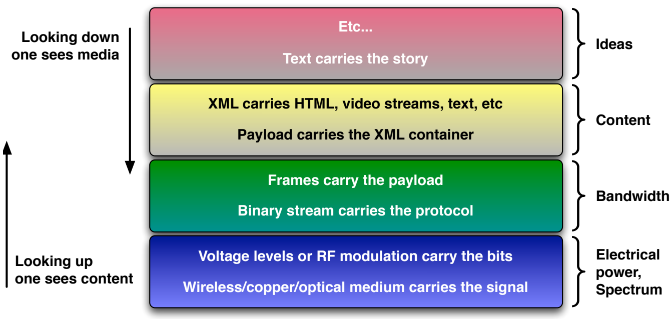

Fig 1. The media stack

Figure 2 also highlights how difficult it is to separate the factors underpinning the co-evolution of technology with socio-economic systems. For example, it is far from clear whether principles such as decentralized architectures or P2P networks were derived from a particular social theory, or whether instead the converse applies. In general, it seems more accurate to state that socio-economic and technical systems are interdependent and tightly intertwined, that socio-technical and socio-economic phenomena appear to emerge spontaneously from their interaction, and that social theory then tries to explain them. This state of affairs can be interpreted as evidence that it is not so easy to make a clear separation between the 'objective' technology we build and our 'subjective' or 'intersubjective' human experience (Ciborra & Hanseth 1998).

As discussed in Feenberg (2005), in Heidegger's early writings 'Aristotle's conception of being in general is derived from the Greek practice of technical making, from τέχνη. τέχνη realizes the inherent potentialities of things rather than violating them as does modern technology' (ibid, xiv). Compatibly with this position, according to Marcuse '…the task of a post-Heideggerian philosophy is to conceive a technology based on respect for nature and incorporating life-affirming values in its very structure, the machines themselves' (ibid, 4). This utopian demand can be understood as 'an implicit recovery of Aristotle's idea of τέχνη in a modern context, freed from the limitations of ancient Greek thought and available as a basis for a reconstructed modernity'

obtains, for example, autocatalytic cycles, which are generally designed out of networking and computer systems as undesirable loops. This form of tight coupling is the basis for emerging models of 'unconventional computing', see for example http://www.biomicsproject.eu.

[4] In his original paper, Turing (1936) introduced also the concept of the 'choice machine' which, unlike what then became known as the Turing machine, could be interrupted during the evaluation of a mathematical function. This is the theoretical starting point of unconventional computing initiatives such as interaction computing (www.biomicsproject.eu), which aim to mimic the tight coupling found within biological systems as well as, at much larger scales, ecosystems.

(ibid, 4). Making things (i.e. engineering) can then be recovered as a life-affirming, deeply human activity, as long as we are not blinded by the myth of the neutrality of technology in an objective world. Feenberg's critical theory of technology (Feenberg 1991, 2002) shows how technology embodies our cultural values and is in fact an extension of our human languages that necessarily generalizes the concept of symbol. It then contributes to the construction of our understanding of reality and in particular of our social reality.

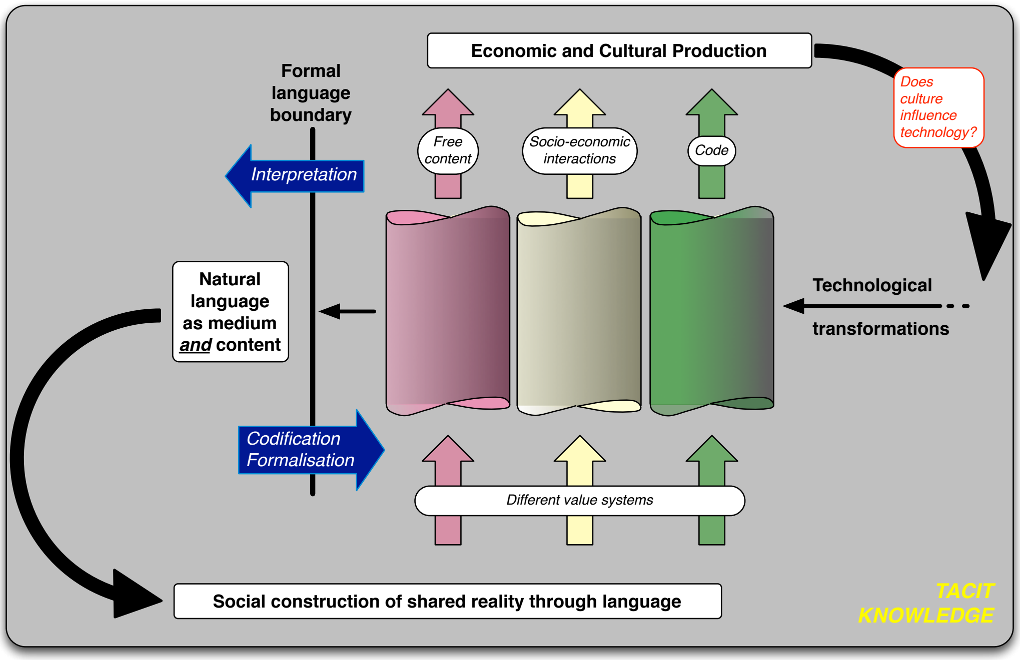

Fig 2. Language, technology and culture

In this panorama of technology recast as an extension of human cultures and languages the Internet plays a unique role because, not only does it share with other kinds of technology this cultural and expressive valence, it *mediates* the very communications that construct the social and cultural systems that created it. It is not clear what the effect of this tight feedback loop might be, but it is pretty clear that it is likely to be a strong one, and perhaps not so easy to control. When looked at through a social science "lens", therefore, the hybrid role of computer science is perhaps best captured by Winograd and Flores's view of computers as communication media (1987). Because communications, in turn, carry commitments (Austin 1962; Flores & Spinosa 1998; Searle 1979), it becomes easier to accept that the Internet has the potential to become a catalyst of social constructivist processes. For example, we can point to the role played by the Web in the formation of the identity of social groups, as discussed by Flores and Spinosa well before Facebook. Figure 3 completes the thought process around the concept of mediation developed in the previous two figures and gives a high-level Escher-like graphical rendition of the feedback loops generated by the interaction of the Internet and media content.

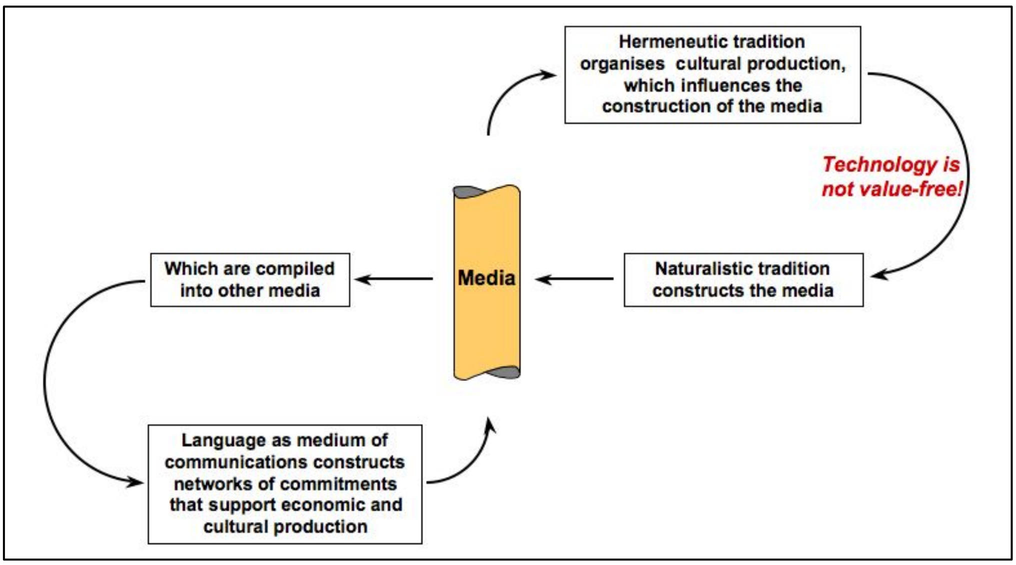

Figure 3: The self-reinforcing cycle at the heart of the Internet (Dini et al. 2011)

This model, although useful for discriminating between how ICTs interact with society and other kinds of technologies (think tractors or container ships), is too abstract to analyse the differences in the innovation dynamics of the Internet under loose coupling vs. tight coupling between its two layers. We need to complement it with an empirical and case study-oriented perspective that examines the Internet as an infrastructure.

III. Infrastructures and Innovation

So far we have been discussing the network layer and the Web as two parts of the same overall infrastructure, the Internet. Whereas this is conceptually correct, the everyday use of these terms by practitioners from each layer is slightly different: the network layer itself tends to be referred to as the Internet – consistently with how things stood before the Web – while use of the Web as the top layer is the same as we have used in this extended abstract so far. As from this section we are shifting the discussion from a model-based, systemic and deductive approach to an empirical, individualist, and inductive perspective, we see no harm in switching the convention for the associated terminology as well.

It appears that both the Internet and the Web can be seen as two distinct infrastructures which have been fostering innovation. However, many people perceive the Web and the Internet infrastructure as a single artefact, especially considering that their individual roles are not easily distinguishable when examined in the context of social construction through language, as discussed in the previous section. For this reason, that view of the Internet and the Web could be complemented by approaching them as two distinct layered and interdependent infrastructures. The Internet itself is based on the telecommunication infrastructure, and supports the Web. More recently, the global deployment of Cloud services could be seen as another infrastructural layer on top of the Internet.

In economic terms, infrastructures have a number of characteristics (Frischmann 2012): (i) governments played an important role in the deployment of those networks, (ii) they were managed on an egalitarian, non-discriminatory basis, and (iii) they generated spillovers (positive externalities) that resulted in social gains. Telecommunication networks had all the characteristics of infrastructure, as summarised by Frischmann (2012), and they led to significant spillovers, e.g. telephone sales or communication among family and social groups. It was on this telecom infrastructure that the Internet was deployed as part of those externalities. One could say that, initially, it may have been viewed as another service, but a retrospective account of its evolution could classify it as another layer of infrastructure deployed on a global scale over those telecommunication networks and as a consequence of the increasingly deregulated telecoms market. Many of the tussles that are currently fought on the Internet, as discussed by Clark et al. (2005), especially those around principles such as the end-to-end argument and openness, are exactly about ensuring that the character of the Internet remain that of an infrastructure instead of just an application. Discussions on Internet governance and on equitable access to it are indeed emerging because the Internet has been established as infrastructure. But when it comes to examining the positive externalities that the Internet has to show, that's when the topic of discussion shifts to the Web.

In a similar fashion, the Web can be viewed as something that initially appeared to be a service based on the Internet. In technical terms, it is based on a protocol that was deployed in the *Application Layer* of the Internet. In the 90s one could view the Web as yet another application offered with a package of Internet applications including *EMAIL* and *FTP*. Nevertheless, the Web was able to leverage direct network effects on a large scale and grew to become a global resource for information publication and discovery in the first instance, and, later, for communication, collaboration and knowledge construction. The peer organisation of the Internet and its support for easy information sharing has led to the success of the Web but also to surprising outcomes ranging from the support for social movements to trading in Wall street (Johnson 2012). Thus, one can now see that the Web evolved to become an infrastructure itself, which is based on the Internet, and which enabled new social architectures and human interaction, such as Wikipedia (Johnson 2012).

In that light, the Internet appears to provide the means for access to Web servers, databases, people and devices, while the Web provides ways of publishing and linking information based on new innovative social structures. Ensuring equitable and fast access to the Internet and the Web is in the programme of many governments in the world, given the promise of the Web as an infrastructure that will lead to significant spillovers and new waves of innovation.

However, apart from the Web, the Internet has fostered further innovation in terms of teleconferencing applications, the Cloud, and interaction with devices. At the same time, the Web has fostered significant innovation with a wave of collaboration and online social networking services. The emergence of the Internet and subsequently of the Web as infrastructures on which ecosystems of innovation flourished has established a precedent that has been followed in online services such as online social networks; these services facilitated the emergence of ecosystems by leveraging network effects and global reach. For example, Facebook fostered the evolution of an application ecosystem that is specific to Facebook. Could these emerging services be seen as an extra layer of infrastructure on top of the Web and the Internet? Do they have the characteristics of an infrastructure in terms of governance, equitable access and the potential for positive externalities? The fact that Google and Facebook are not open platforms in the same sense of the Internet or the Web suggests that they may evolve differently.[5] These are topics for discussion in the future. Nevertheless, it appears that this two-layered infrastructure of the Internet and the Web is characterized by a high and increasing degree of interdependence between the layers, which presents a challenge when considering changes on a technological or policy level.

IV. Interdependence between the Internet and the Web

The growth of the Web has had an impact on the Internet in both technological and policy terms. The growth of the Web as a global infrastructure based on the Internet has had an impact on the Internet itself. ISPs had to optimize their routers to cope with Web traffic. Initially, this involved small-sized Web objects which gradually became larger (Brownlee et al. 2002). In its early years the Web contributed significantly to Internet traffic, followed, later, by peer-to-peer (P2P) applications (Odlyzko 2003). Recently, Web, P2P and Web-

[5] We are indebted to an anonymous referee for pointing this out.

based video-on-demand services were estimated to comprise 85% of Internet traffic (Berl et al. 2010), with Cloud computing adding to that.

Applying law enforcement for content on offshore websites has led to considerations of changes in the domain name system (DNS) in the US under the SOPA act; however, it was argued in a petition to the White House that such changes in the DNS could compromise the openness of the Internet and its potential for innovation.[6] The discussion on net neutrality could be seen as a reflection of the tussle between the stakeholders involved in the deployment of services on the Web and in the provision of Internet access. The significance of open standards as a driver for innovation, the risks of 'Web islands' built around online social networks and application ecosystems, and the need to keep the Web and the Internet as two separate layers have been argued by the inventor of the Web, Tim Berners-Lee (Berners-Lee 2010).

Keeping the Internet and the Web as separate, identifiable (although interdependent) infrastructures can indeed continue to foster innovation. Since the emergence of the Web, the governance model of the Internet and its end-to-end design made new innovation possible. For example, cloud services could be seen as innovation on the Internet driven partly by thriving Web-based services and partly by the widespread use of a variety of networked devices for domestic and industrial applications (including home routers, smartphones, utility usage monitors and media devices). Could it be that the pressure on Internet Service Providers (ISPs) for Cloud access will be increasing, leading to a "flatter" structure of the Internet and making it gradually indistinguishable from the Cloud? Could such developments pose a threat to the potential of the Internet for innovation?

V. Conclusion and Outlook: Net Neutrality and Internet Innovation

The relative roles of infrastructures and applications discussed in the second part of the paper mirror the relationship between media and content discussed in the first part, but add a dynamic element whereby applications and services – under certain conditions that we have only hinted at here – have the potential to evolve gradually into infrastructures which, in turn, can engender more applications, and so on in an apparently never-ending process. Further, the empirical perspective on infrastructures captures elements such as political economy and an analysis of individual interests that are absent from the functionalist and systemic language-based media-stack model. It suggests that the viability of the latter may depend to a significant extent upon characteristics such as openness, democratic values, and market vs. monopolistic behaviour. In other words, embracing both the systemic and individualist epistemological perspectives captures a richer picture of the various layers of the Internet and of some of the factors that drive their evolving dynamics.

In this extended abstract, so far we have argued that the pressure of Web traffic, then P2P traffic and, more recently, of Cloud traffic on the Internet as infrastructure suggests that the Internet with its governance model and its design principles has been fostering innovation with a high impact. The Internet is increasingly supporting interaction among devices, and the distribution of entertainment content is bringing increasing demands for quality of service. The multitude of users engaging with the Internet via devices, the Cloud, or the Web is shaping its evolution. To understand how people engage with the Internet could help us ensure its sustainability and improve it further. Mixed research methods can provide us with new insights on how to improve the Internet – and its continually emerging new layers of applications – as infrastructure to ensure its continuing contribution to innovation. Internet Science will be providing the interdisciplinary methods, the best practices and the research momentum in the study of the Internet as a techno-social phenomenon in this context.

As this is an extended abstract, we now include a critique in the form of a list of points that will be explored in greater depth in our future work:[7]

- As it is known that the Internet net neutrality does not always hold today, its relationship to innovation needs to be investigated further.
- The increasing interdependence between the two layers discussed deserves a more formal analysis: Why? How? What is the impact if any? Why does this limit innovation?
- Is the principle about the separation of layers a necessary condition to support open innovation, or is this just related to the necessity to expose APIs naturally required to support new services from third parties? What does it mean that two layers are more dependent without either having control of both?
- The concept of architecture could usefully be clarified to draw out the roles of design and praxis, especially given the genesis of the Internet from an experiment among a closed and cohesive group with largely shared perspectives and objectives and the distance from 'then' till 'now'.
- The concept of independence needs greater definition; there are many senses in which the Internet and the Web can be said to be independent, but in relation to choice and behaviour, or the interpretation of observations or stylised facts, or the construction of an analytic framework for deriving and testing hypotheses (in the face of endogeneity and unobservability) something more concrete is needed.
- It is not clear why the feedback loop between innovation and independence/modularity should be seen as ironic, but it points out a possibly useful standpoint: the Internet as a complex system, characterized by emergent behaviours and self-organization. The existence of transitory, context- and path-dependent, multiply-directed and (inter-)subjective feedback loops is expected and may represent the true 'genius' of the system as a whole. It also allows the artificially simplified view of the linked systems to be enlarged to track the way human beings individually and in networked association have evolved with the Internet.

[6] whitehouse.gov. Available from: http://www.whitehouse.gov/blog/2012/01/13/obama-administration-responds-we-people-petitions-sopa-and-online-piracy

[7] These points and observations were provided by the anonymous referees, to whom we are grateful for the constructive input.

- The tight vs. loose characterisation is a bit forced. Common knowledge – and typically formalised – hardly seem loose in the intuitive sense, and linkages in even modestly complex food webs (at least under conditions of relative abundance or evolutionary slack) are hardly tight in the manner defined. This is the source of resilience in such systems.
- The definition of net neutrality will need to be made more explicit, and its virtues will need to be derived more carefully. The issue is not innovation per se, but which kind of innovation. There are both good and bad types of innovation, just as there are good (even essential) and bad departures from neutrality (e.g. Ramsey pricing).
- The discussion of language and communication will need to be extended: not all exchanges involve a single listener and a single hearer and not all impediments to the free, easy and potentially pre-emptive flow of communication bear positive fruit. Transaction costs do get in the way, but the communication observed in computer-based financial trading shows that it may be useful to slow things down enough to permit reflection (e.g. learning together how to use what we have 'innovated') or to introduce impediments that we have to innovate around, producing profound improvements and disrupting entrenched power and control (tipping and persistence of power being problems that the current Internet has yet to surmount). An example relevant to this context is structural holes (e.g. extensive literature from theory to empirics).
- The potential of complexity to produce useful plasticity and variable geometries that usurp and refresh the presumed layered structure – if only by emergence – should be addressed (e.g. Arthur 2009).
- The infrastructure vs. service discussion could be improved by relating it to the model of intensive vs. extensive competition, and to the metaphor of services competing on platforms, using relationships between service providers and users on one side and the platform or infrastructure on the other, or end-to-end relations as suggested by the 2-sided market model.

VI. Acknowledgement

The work reported in this article was funded in part by the EU projects EINS, contract number CNECT-288021, and BIOMICS, contract number CNECT-318202. Their support is gratefully acknowledged.

VII. References

[1] Austin, J L (1962). *How to do Things with Words*, Oxford: Oxford University Press.

[2] Arthur, B (2009). *The Nature of Technology: What it is and How it Evolves*. Penguin.

[3] Berdou, E and Dini, P (2005). D18.3: Report on the Socio-Economics of Free/Open Source, DBE Deliverable, European Commission. Available at http://files.opaals.eu/DBE/

[4] Berl, A, Gelenbe, E, Di Girolamo, M, Giuliani, G, De Meer, H, Dang, MQ & Pentikousis, K (2010). Energy-efficient cloud computing. *The Computer Journal*, **53**(7):1045–1051.

[5] Berners-Lee, T (2010). Long Live the Web. *Scientific American Magazine*.

[6] Brownlee, N & Claffy, KCCMI (2002). Understanding Internet traffic streams: dragonflies and tortoises. *Communications Magazine*, IEEE. **40**(10).

[7] Ciborra, CU & O Hanseth (1998). From tool to Gestell: agendas for managing the information infrastructure. *Information Technology & People*, **11**(4):305-327.

[8] Clark, DD, Wroclawski, J, Sollins, KR & Braden, R (2005). Tussle in cyberspace: defining tomorrow's Internet. Networking, *IEEE/ACM Transactions*. **13**(3):462–475.

[9] Clark, DD, Wroclawski, J, Sollins, KR & Braden, R (2002). Tussle in cyberspace: defining tomorrow's internet. In: August 2002, ACM.

[10] Dini, P, Iqani, M, and Mansell, RE (2011). The (Im)Possibility of Interdisciplinarity: Lessons from Constructing a Theoretical Framework for Digital Ecosystems, *Journal of Culture, Theory and Critique*, **52**(1): 3-27.

[11] Dini, P & Sartori, L (2013). Science as Social Construction: An Inter-Epistemological Dialogue between Two Internet Scientists on the Inter-Epistemological Structure of Internet Science, Part 1, in EINS2013: 1st International Conference on Internet Science, 10-11 April.

[12] Espinel, V, Chopra, A & Schmidt, H (n.d.). Combating Online Piracy while Protecting an Open and Innovative Internet | We the People: Your Voice in Our Government. [Online]. petitions.whitehouse.gov. Available from: https://petitions.whitehouse.gov/response/combating-online-piracy-while-protecting-open-and-innovative-internet. [Accessed: 18 /02/13].

[13] Feenberg, A (1991). *Critical Theory of Technology*, Oxford: Oxford University Press.

[14] Feenberg, A (2002). *Transforming Technology: A Critical Theory Revisited*, Oxford: Oxford University Press.

[15] Feenberg, A (2005). *Heidegger and Marcuse: The Catastrophe and Redemption of History*, Routledge.

[16] Flores, F and C Spinosa (1998). Information technology and the institution of identity: reflections since *Understanding Computers and Cognition*. *Information Technology & People*, 11(4): 351-372.

[17] Frischmann, BM (2012). *Infrastructure*. Oxford: Oxford University Press.

[18] Johnson, S (2012). *Future Perfect*. Penguin.

[19] Lessig, L (2006). *Code Version 2.0*, Basic Books.

[20] Odlyzko, AM (2003). Internet traffic growth: Sources and implications, in *Optical Transmission Systems and Equipment for WDM Networking II*, B. B. Dingel, W. Weiershausen, A. K. Dutta, and K.-I. Sato, eds., Proc. SPIE, **5247**:1-15.

[21] Papazoglou, M (2003). 'Service-oriented computing: concepts, characteristics and directions', *Proceedings WISE'03*, IEEE, pp. 3–12.

[22] Searle, JR (1979). *Expression and Meaning*, Cambridge, UK: Cambridge University Press.

[23] Turing, A (1936). On Computable Numbers, with an Application to the Entscheidungsproblem, *Proceedings of the London Mathematical Society*, **42**(2):230-265; A correction, ibid., 1937, **43**:544-546.

[24] Winograd, T, & Flores, T (1987). *Understanding Computers and Cognition*, Addison Wesley.