A Decentralized Domain Name System?
User-Controlled Infrastructure as Alternative Internet Governance

Francesca Musiani

Post-doctoral researcher, MINES ParisTech[1]

Draft. Please contact the author before quoting.

Introduction

> *"The heart of DNS problems aren't [sic] with ICANN. It's with the governments and companies which can control ICANN. The system is centralized."*
>
> - *Peter Sunde, December 2010*

Late 2010. WikiLeaks makes thousands of secret US diplomatic cables public, losing a few days later its web hosting company and the *wikileaks.org* domain. Discussions about a "new competing root-server", able to rival the one administered by the Internet Corporation for Assigned Names and Numbers (ICANN), soon start populating the Web, prompted by well-known Internet "anarchist" Peter Sunde. An alternative domain name registry is envisaged, a decentralized, peer-to-peer (P2P) system in which volunteer users would each run a portion of the Domain Name System (DNS) on their own computers, so that any domain that would be made temporarily inaccessible, because of seizures or blockings, may still be accessible on the alternative registry. Instead of simply adding a number of DNS options to the ones already accepted and administrated by ICANN, this project would supersede the main DNS governance institution – in favor of a distributed, user infrastructure-based model.

The technical and political debates concerning a decentralized, alternative DNS – consequence of domain name seizures for content mediation purposes – are a good illustration of what Laura DeNardis has recently described as the "turn to infrastructure" in Internet governance. They

[1] This paper accounts for work conducted primarily in Washington, DC in 2012-13, with support from the Yahoo! Fund on Communication Technology, International Values and the Global Internet. The author also wishes to acknowledge support from the research program *ADAM-Architectures distribuées et applications multimédias*, funded by the French National Agency for Research (ANR). Earlier versions of this paper have been presented and discussed at the eighth Media in Transition conference (MiT8), Massachusetts Institute of Technology, Cambridge, MA, May 3-5, 2013, and at the 2013 conference of the International Association for Media and Communication Research (IAMCR), Dublin, Ireland, June 25-29, 2013.

show how "Internet governance infrastructures are increasingly being co-opted for political purposes completely irrelevant to their primary Internet governance function" [DeNardis, 2012], and in turn, how developers seek to circumvent this co-optation in disruptive ways – by creating new arrangements of governance-*using*-infrastructure in the process.

This paper draws on perspectives informed by science and technology studies (in particular software studies and critical code studies), and on qualitative interviews with technical and political actors involved in DNS governance, to provide a contribution to the study of the "alternative Internet" [Atton, 2005] and its implications as an *imaginaire* [Flichy, 2007], an organizational principle and a socio-technical artifact.

1. Caring about the (P2P) plumbing

"Peer-to-peer is plumbing, and most people don't care about plumbing," pointed out some years ago Dan Bricklin, the "father" of spreadsheets, in a seminal book about P2P's potential as a *disruptive* technology [Bricklin, 2001: 59]. This assessment of the first file-sharing applications is probably right: likely, their success owes more to the suitability of such tools to rapidly find and download a specific content, than to their underlying architecture in itself. Yet, the *plumbing*, the design of the lower layers – as Susan Leigh Star has effectively put it, the "invisible work" [1999: 385] underlying practices, uses and exchanges in a networked system – informs its adoption and (re)appropriation by users, its regulation, its organizational forms, and should be, as such, an important component of its analysis [Musiani, 2012]. Several bodies of work, crossing Internet studies with science and technology studies (STS), have sought in recent years to explore the social and political qualities of information infrastructures, and to find the 'material' in the virtual of software and code [Star & Bowker, 2002; Monberg, 2005; Manovich, 2001; Fuller, 2008; Marino, 2006; Ribes & Lee, 2010; Kirschenbaum, 2003; DeNardis, 2009, 2010].

Paralleling STS-informed approaches – and conceptualizing network architectures as political, social and legal – some law and economics scholars focus on the relationship between Internet architecture and innovation. It is argued that the Internet's evolution is likely to depend, in the medium-to-long term, on the topology and technical models of Internet-based applications, and on the infrastructure that underpins them [Aigrain, 2011]: thus, emphasis is given to how the "lower layers" of a networked application inform issues and objects that are *de facto* crucial for users, such as the handling and storage of data, computing resources management, information extraction and aggregation; and to the materiality of networked systems is a source of "techno-legal" implications, both for user rights and regulation [Elkin-Koren, 2006]. As the architecture of the Internet, and that of its services, has been a matter of controversy in the past, it is now subjected to manifold tensions. After having recognized its importance as a regulation mechanism, it is increasingly analyzed as a leverage for development and as the origin of market opportunities (and constraints) [van Schewick, 2010], while acknowledging its suitability "by design" to changes and modifications [Braman, 2011].

1.1. P2P: the search for alternatives built on a cornerstone of Internet history

Peer-to-peer is a computer networking model structured in a decentralized manner, so that communications or exchanges take place between nodes entrusted with equal responsibility in

the system. Participants in the network make some of their computational equipment and resources (power, storage, bandwidth) available to the system; accessible directly by peers, these shared resources are necessary for the proper functioning of the service offered by the network. The dichotomy between a server, provider of resources, and clients, resource-seekers – characteristic of the client-server model – is replaced by a situation where every peer hosts or provides a part of the overall resources, and all peers request it [Schollmeier, 2001].

For a large number of Internet users – since the encounter between P2P and the public, prompted by Napster in 1999 – this technology is a *de facto* synonym for the (illegal) download of cultural content; for others, it represents the ultimate utopia of techno-egalitarianism, or suggests a more sustainable organizational model for the societies of tomorrow. In any analysis of P2P, one cannot completely set aside these strong visions; the history, past and present, of P2P is informed by these visions, discourses, narratives, and informs them in return. However – and while not neglecting the powerful agency of these normative views – this paper does not seek to be a further contribution to the already well-established debate on the sharing/stealing dialectic which P2P is almost "naturally" associated to. Rather, it takes as its starting point the "virtues" of decentralization – effectiveness, stability and resilience – that provide a crucial contribution to the political and technical significance of P2P systems [Elkin-Koren, 2006].

The development of services based on decentralized, distributed, peer-to-peer network architectures has been acknowledged for several years – even, and especially, in today's times of *clouds* and big data – as one of the interesting axes of transformation of our modalities of communication and management of digital content. The concept of decentralization is embedded to some extent at the very core of the Internet, especially in the organization and the circulation of data packets. Yet, today's Internet integrates this principle only partially: while every Internet user has become, at least potentially, not only a consumer but also a distributor and a producer of digital content, a considerable concentration of data takes place in specific regions of the Internet [Minar & Hedlund, 2001; Berners-Lee, 2010]. Recurring to decentralized network architectures and distributed organizational forms for Internet services is thus envisaged by a number of projects, companies, services, in a perspective of effectiveness, resolution of some management difficulties, and digital "sustainable development".

Like several decentralized alternatives to Internet-based services that we have explored elsewhere [Musiani, forthcoming 2013], the P2P DNS project described in this paper integrates a specific design choice: the delegation of the responsibility and the control of data management and flows to the "edges", the margins, or the periphery in the infrastructure of these networking systems. The necessary operations for the proper functioning of these systems, and their ability to correctly provide the services for which they are intended, technically depend on users: their terminals, their computing resources, mobilized in an aggregate manner in order to serve a common purpose. We focus our attention here on the "meeting" between a decision to develop a P2P technical architecture, and a complex and controversial component of Internet infrastructure such as the Domain Name System – which in its current form sets out a clearly identifiable dichotomy between providers of resources – Internet registries and registrars, under the supervision of ICANN – and clients requesting it. Following and trying to clarify the "*ballet* between programmers, software and users" [Abbate, 2012] that builds the project of decentralization for the DNS, this paper contributes to the exploration of the socio-political

implications of the distributed and decentralized approach to the technical architecture of Internet services.

2. Internet governance, a field in the making

Internet governance (IG) today is a lively, emerging field, and the body of research that explores it is no less "in the making" [Latour, 1987]. A "working definition" of IG has been provided in the past, after the United Nations-initiated World Summit on the Information Society (WSIS), by the Working Group on Internet Governance – a definition that has reached wide consensus because of its inclusiveness, but is perhaps too broad to be useful in drawing more precisely the boundaries of the field [Malcolm, 2008]:

> Internet governance is the development and application by Governments, the private sector and civil society, in their respective roles, of shared principles, norms, rules, decision-making procedures, and programmes that shape the evolution and use of the Internet [WGIG, 2005].

This broad definition implies the involvement of a plurality of actors, and the possibility for them to deploy a plurality of governance mechanisms. IG has been described as a mix of technical coordination, standards, and policies [e.g. Malcolm, 2008 and Mueller, 2010]. Technical coordination is conducted, through norms and the market, by the institutions that manage the Internet's technical architecture and resources. Standards development is the set of processes by which, through norms and architecture, technical standards are developed for the operation of the Internet. Public policy governance relates to the development of international public policy for the Internet, and addresses, in particular, matters of regulation of issues such as online privacy. Internet policies are implemented at the domestic and supra-national levels, and discussed at the global level in venues such as the United Nations-promoted Internet Governance Forum. It is in the context of the IGF that the concept of "multi-stakeholderism" was first applied to IG, reflecting the idea that every "holder of stakes" in the Internet should be able to have a voice heard in the shaping of the network of networks [Levinson, 2010].

2.1. A contested definition

Despite the inclusiveness of the above definition, the definition of IG is relentlessly contested by differing groups across political and ideological lines. One of the main debates concerns the authority and participation of specific actors, such as national governments, corporate entities and civil society. In the context of the Internet, relevant actors in governance processes and power arrangements are not only governments. Indeed, under the purview of their sovereignty, governments perform certain IG functions such as regulating abuses, overseeing antitrust measures, and responding to security threats – and they also use content filtering and blocking techniques for surveillance and censorship of citizens. Their role in IG remains central and often ambiguous. However, other areas of IG, such as Internet protocol design and coordination of critical Internet resources, have historically been delegated to transnational, institutional entities, and to private ordering [DeNardis, 2013].

Moreover, one should be careful about subscribing to two opposing ideological positions [Mueller, 2010], an enthusiastic but naïve technological determinism leading the "digital revolution", and the mere reproduction of traditional forms of state sovereignty as applied to the Internet, arguing that "the Internet is nothing new". One should also refrain from reducing the IG debates to one of its components or arenas. A common misconception about IG is, for example, its identification with the United Nations-promoted processes that have led to the establishment of the IGF, a multi-stakeholder dialogue that, even if interesting in its experimentation with innovative governance formats, is not the place where "the practice of internet governance happens" [DeNardis, 2013, *ibid*.]. Similarly, the Internet Corporation for Assigned Names and Numbers, ICANN, while being one of the important institutions of IG managing a delicate part of critical internet infrastructure, has at times been presented in such a way that can lead people to believe that it "runs the internet" on its own – which is not the case.

2.2. The Domain Name System: Spotlight on Internet infrastructure…

The scenario briefly outlined above often leads to neglect or disregard what is, instead, a crucial, albeit discreet and even invisible, aspect of IG: there are a number of components of the Internet's infrastructure and technical architecture in whose design are embedded, to some extent, arrangements of governance. Made of technologies and processes beneath the layer of content and designed to keep the Internet operational, the infrastructure of the "network of networks" is one of today's most critical components of IG.

Among the instances where the "political materiality" of the Internet is revealed, one of the most important is the system of Internet Protocol (IP) addresses. Devices exchanging information over the Internet are identified by unique binary numbers identifying its virtual location, temporary or permanent. Internet routers use these addresses to determine how to route packets over the Internet. The current standard for Internet addresses, IPv4, is in the final stages of exhausting its unallocated address space. A new protocol, IPv6, has been recommended to expand the number of available addresses. However, for a variety of political and technical reasons, the upgrade to IPv6 is still in its infancy and the depletion is getting closer, with important policy implications. Another example of critical Internet infrastructure are Internet Exchange Points (IXPs). They are the physical junctures where different companies' backbone trunks interconnect, exchange packets and route them toward their appropriate destinations. The implications of the management and regulation of IXPs extend to fair competition mechanisms, surveillance and filtering, and stability. There are many more, but let us focus on the central case study for the present paper, the Domain Name System.

In a nutshell, the DNS is a wide database management system, arranged hierarchically but distributed globally, across servers throughout the world; it translates between user-friendly alphanumeric domain names and their associated IP addresses necessary for routing packets of information over the Internet. For this reason, it is oftentimes called the Internet's "phone book"[2]. Nowadays, the number of queries addressed by the DNS is estimated at several billions per day, and by providing a worldwide keyword-based redirection service, the DNS is an essential component of the functionality of the Internet. The Internet's root name servers contain a master file known as the root zone file, listing the IP addresses and associated names of the

[2] The author has organized a recent conference on the topic at the Institute for the Study of Diplomacy, Georgetown University. The conference report is available here (http://isd.georgetown.edu/388419.html).

official DNS servers for all top-level domains (TLDs): generic ones, like .com, .edu, .gov, etc. and country codes such as .us, .uk, .fr. The right to use a domain name is delegated by domain name registrars, accredited by ICANN, the organization charged with overseeing the name and number systems of the Internet, and with controlling the root server system and the root zone file.

This area is rife with controversies, involving institutional and international power struggles over DNS control, and issues of legitimacy, democracy, and jurisdiction. In particular, as an organization of Californian private law and the *de facto* exclusive manager of one of the most delicate infrastructures of IG today, ICANN has been under constant international scrutiny ever since the Internet has become a global, public phenomenon, due to its close ties with the United States government and alleged lack of transparency.

There are additional policy implications to the DNS: it was originally restricted to ASCII characters, precluding domain names in many language scripts such as Arabic, Chinese or Russian. Internationalized domain names (IDNs) have now been introduced. Furthermore, in 2011, ICANN's board voted to end most restrictions on the generic top-level domain names (gTLD) from the 22 currently available. Companies and organizations will now be able to choose essentially arbitrary top-level Internet domains, with implications for consumers' relationships to brands and ways to find information on the internet. Further DNS issues concern the relationship between domain names and freedom of expression, security, and trademark dispute resolution for domain names.

2.3. **… and its *détournements***

While these manifold controversies play a prominent role in the shaping of today's Internet – and have been addressed at length in literature – this paper focuses on a slightly different, albeit related, take on Internet infrastructure. In recent years, we witness a number of (more or less successful) attempts, by political and private entities, to co-opt infrastructures of internet governance for purposes other than the ones they were initially designed for. Not only is there governance *of* infrastructure – as outlined up to this point – but governance is carried out *by* leveraging infrastructure in "creative ways". This particularly applies to content mediation: conflicts over the ways in which information is carried and circulates in the Internet are, increasingly often, fought in its lower layers. Forces of globalization and technological change have diminished the capacity of sovereign nation states and media content producers to directly control information flows through laws and policies – and these actors are acknowledging infrastructure as a mechanism for regaining this control. What Laura DeNardis [2012] calls the "turn to infrastructure for Internet governance" entails not only issues of economic freedom – but of information and communication freedoms.

Examples of how content mediation controversies have shifted into the realm of Internet governance infrastructure can be found, for example, in the intentional outages of basic telecommunications and Internet infrastructures, enacted by governments via private actors, whether via protocols, application blocking, or termination of access services. The government-initiated Internet outages in Egypt and Libya, in the face of revolution and uprisings, have illustrated this and may have set a dangerous precedent.

However, the DNS is perhaps, nowadays, the best illustration of this "governance by infrastructure" tendency. Domain name seizures, using the domain name system to redirect queries away from an entire web site rather than just the infringing content, have recently been considered as a suitable means of intellectual property rights enforcement. DNS-based enforcement was at the heart of controversies over the legislative projects Protect IP Act (PIPA) and Stop Online Privacy Act (SOPA). Governance by infrastructure enacted by private actors was also visible during the WikiLeaks saga, when Amazon and EveryDNS blocked Wikileaks' web hosting and domain name resolution services. And finally, attempts at governance by infrastructure are at the core of the project of an alternative, decentralized or P2P DNS, in which volunteer users would each run a portion of the DNS on their own computers. Faced with Internet infrastructure co-optation for content mediation functions that eventually restricts their freedom of expression and access, user/developers seek, in turn, to circumvent this co-optation in disruptive ways. And through this *détournement* [Akrich, 1998] they create new arrangements of governance-*using*-infrastructure.

3. Towards a decentralized alternative for DNS? Debates, balances, concerns

At the end of 2010, the organization WikiLeaks makes thousands of secret US diplomatic cables public, losing a few days later its web hosting company and the *wikileaks.org* domain. Discussions about a "new competing root-server", able to rival the one administered by ICANN, raise a new wave of interest on the Web, prompted by well-known Internet "anarchist" Peter Sunde. An alternative domain name registry is envisaged, a decentralized, peer-to-peer (P2P) system in which users would each host a part of the DNS on their computers, so that any domain made temporarily inaccessible by a registry, for whatever purpose, may still be accessible on the alternative registry. Instead of simply adding a number of DNS options to the ones already managed by ICANN (like OpenNic or NewNet had done before), this project is aimed at superseding or circumventing the main DNS governance institution – in favor of a distributed, user infrastructure-based model. The remainder of this section explores the debates on a decentralized, user-controlled infrastructure response to DNS co-optation.[3]

3.1. A history of dissatisfactions, a history of attempts

Dissatisfactions about the ways in which the DNS is managed are hardly recent. Due to its being hierarchical, and initially not built with security in mind, unsavory people have made it the center of their attention in the past; previously mentioned controversies, and in particular, the control over the root system – which, it is argued, is *de facto* managed by the government of the United States via ICANN – have made it the subject of several heated international and intergovernmental meetings.

[3] The source material for this section is made of in-depth qualitative interviews conducted with technical developers – working as academics and/or in the private sector – and Internet governance specialists, most of whom have expressed the wish to remain anonymous (thus, I have decided not to explicitly name the non-anonymous minority, as well). Interviews have been conducted in Washington, D.C., New York and Boston, and remotely in Italy and France, in the September 2012-April 2013 period.

In recent years, the *Combating Online Infringement and Counterfeits Act* (*COICA*) law project, in 2010, and its re-written, but no less controversial, 2011 version (*Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property Act*, or *PIPA*), both in the United States, have drawn public opinion's attention to the risk of a control of the Internet exercised *via* the DNS. Without waiting for the eventual passing of the COICA, the U.S. government did, in November 2010, proceed to order the cancellation of several domain names. According to the member of a D.C.-based ICT governance think-tank, this was

> "completely on behalf of the entertainment industry."

The WikiLeaks case has also well illustrated the pressures exercised *via* the DNS against freedom of expression, and the risks of concentration: *wikileaks.org* was down for days because there was only one DNS hosting for that domain. Although illustrated by many episodes originating with American institutions and companies, this issue is not specific to the United States. In France, the *Loi d'orientation et de programmation pour la performance de la sécurité intérieure* (*LOPPSI*) provides for a mandatory filtering of domain names that the government deems as threatening, a measure whose implementation could be done via the DNS.

The frustrations concerning DNS management, present and future, are – as outlined until this point – many and legitimate. Peter Sunde's uproarious "call to arms" is perhaps the most recent and media-savvy[4], but historically, these frustrations have already led to a number of alternative DNS-related projects, aimed at creating alternative root servers so as to bypass ICANN or existing registrars, or at developing name resolution systems not using the DNS hierarchy, for example based on Distributed Hash Tables (DHTs)[5]. Such projects include the Cornell-based, DNS safety-net project CoDoNS[6] or the Italian-based ANDNA, part of the decentralized routing system project Netsukuku[7]. There appears to be a consensus among developers that new, alternative DNS projects should therefore begin by pondering why these projects, some of which (such as CoDoNS or the alternative root Open Root Server Network, ORSN, terminated in 2008) are technically sound and innovative, have never known significant deployment. Otherwise, they are likely to encounter, sooner or later, the same fate.

3.2. Replacing what?

From a technical viewpoint, developers make several different remarks on the feasibility of a decentralized or P2P DNS. The DNS serves two fundamental operations: name registration (the management of the reservation of Internet domain names by different entities) and name resolution (the behind-the-scenes task of converting domain names to their corresponding IP address). Historically, the term "Domain Name System" has referred to both as if they were necessarily linked. But this is not the case, even if the name registration service and the

[4] E.g. http://digitizor.com/2010/12/01/the-pirate-bay-co-founder-starting-a-p2p-based-dns-to-take-on-icann/
[5] Systems providing a lookup service in which responsibility for maintaining mapping information is distributed among the nodes, so that a change in the set of participants in the system causes a minimal amount of disruption.
[6] http://www.cs.cornell.edu/people/egs/beehive/codons.php
[7] http://netsukuku.freaknet.org/

resolution protocol[8] have interactions – both have a tree-like structure, for example. The registration mechanism ensures the *uniqueness* of names, one of the DNS's most important functions, and the resolution mechanism allows a machine to *obtain information* – for example, IP addresses – in exchange for a domain name. One could, thus, foresee to replace only one of them; this is precisely what the CoDoNS project was doing – replacing the resolution function by a DHT, while maintaining the registration mechanism intact in its previous form.

The replacement of the resolution mechanism, although a daunting task (it should modify hundreds of thousands of machines) is possible: today, alternative mechanisms already exist. Change the system of naming and registration seems much more unrealistic to a number of technical developers, for a reason that is primarily not technical: the fact that it is already known and adopted by so many users. One of the developers describes users as

> "more complicated to update than software",

referring to the cumulative, "snowballing" effect that a critical mass of individuals using an information system has on other individuals.

A crucial, controversial issue is thus to clarify the function that a P2P DNS project should tackle. In its initial call, Peter Sunde mentions the creation of an alternative root, a move that would entail fundamental evolutions in the domain name registration mechanism. Others have spoken in the past about creating a new top-level domain name, .p2p; and yet others seem to be seeking to replace the DNS with a BitTorrent-informed mechanism. Several different projects coexist among developers, each with different specifications, sharing only a shared technical and political dissatisfaction with the current system – and a P2P-informed culture. Developers are aware of this "broader picture" and of the *fils rouges* that may weave them together, but so far, these projects have remained relatively isolated from one another.

Another issue discussed in the context of alternative DNS projects is the extent to which current DNS governance arrangements could be replaced or erased altogether. Several developers involved in P2P DNS projects mention the possibility that organizations that are still running a root alternative, like OpenNic[9], participate in such projects as alternative instances of governance. Such an organization could be the registrar of the .p2p domain, and a web page exists on its wiki that describes the project[10]. In this case, though, it is argued that the problems now represented by instances such as ICANN, VeriSign[11] or national registries would likely simply be shifted to OpenNic:

> "Power would not dissolve, but shift or transfer from one actor to another, and this would not, in itself, entail any solution to the problem."

---

[8] Standardized in two Requests for Comments of the Internet Engineering Task Force, the memoranda published by the IETF describing methods, behaviors, research, or innovations applicable to the working of the Internet and Internet-connected systems: RFC 1034 and RFC 1035.

[9] http://www.opennicproject.org/

[10] http://wiki.opennicproject.org/dotP2PTLD

[11] The Virginia-based company that operates a diverse array of network infrastructure, including two of the Internet's thirteen root name servers, and the authoritative registry for the .com and .net TLDs.

3.3. What would be left behind?

This is the point when, most often, discussions among developers (within single projects and occasionally, more transversally) shift into the crucial question, one that is both technical and deeply political: what services does the DNS provide, and what interests does it serve?

> "The DNS has survived and scaled well, in a somewhat surprising way,"

points out a developer. It provides unique names, which can be memorized by a human in a relatively easy way, and can be resolved by a program. Moreover, it has been working for more than twenty years, despite the significant changes experienced by the Internet during its evolution from a quiet, "symmetrical" utopia of passionate intellectuals to an agitated mass-medium.

Ultimately – argues a Boston-based developer – before moving to another system, all interested stakeholders will have to consider the system they would have to give up in order to do so. This necessity of balancing pros and cons, particularly the fact that no solution will solve all problems with no inconvenience or side-effect, is well-known in the community of P2P developers, although several declarations, especially following Sunde's project announcement, have let the enthusiasm for the decentralized utopia prevail, somewhat a-critically[12]. Two alternatives are possible to facilitate file search and recovery in a P2P system: either a hierarchical system is used – this is the case for classical BitTorrent, where the recovery of a .torrent file is done through a Uniform Resource Locator (URL)[13], thus, a domain name; or, the system works in a completely peer-to-peer and decentralized manner, and in this case, there is no uniqueness, no single root[14]. The same name can refer to two completely different files, it can be recorded by two different entities, and refer to completely different content.

As a policy analyst for a D.C. firm underlines, whether this implication of decentralization is an issue worth facing in all its technical and political complexity, is

> "bound to subjective assessment of the different stakeholders of the DNS system."

Several of the interviewed *techies*, though, are not optimistic about the practical implications of these assessments. If some actors may consider that

> "the fact of circumventing or eliminating ICANN, the currently heavily hierarchical registrars, and the like, is well worth bearing some disadvantages,"

the important shifts in security and authentication patterns that the P2P DNS would entail seem severely under-represented in primarily non-technical online and off-line debates, if compared to their relevance: they are mentioned explicitly and clearly only in a handful of related articles, that note: "And yes, it's not going to be secure and authenticated like the present system. We're

[12] A *Wired* article is cited in support: http://www.wired.co.uk/news/archive/2010-12/02/peter-sunde-p2p-dns
[13] The character string that constitutes a reference to an Internet resource.
[14] RFC 2826, Internet Architecture Board's technical comment on Unique DNS Root and http://www.bortzmeyer.org/2826.html

just going to have to deal with that[15]". Unique name registration in a P2P environment, with no need of a central registry, has already been theorized and coded into algorithms. However, its correct functioning is based upon a premise very difficult to achieve in real-life P2P contexts, as decades of history of this technology have shown: all parties must cooperate.

If the name resolution system is changed, what is to be gained, and what lost? Here, developers insist again on the fact that the current DNS is based on over twenty years of experience and interaction with the "real world." Any other mechanism – those based on DHT are technically sound, and certainly worth the attention of any ambitious developer – is certainly going to take years before a sufficiently mature development stage, and a long coexistence should be expected; a developer insists that

> "claims of being able to replace current DNS operators and governance entities in three months are hardly anything more than unrealistic boasting."

3.4. Engineering, adoption, and governance: the *triple challenge* of DNS alternatives

According to developers' comments, DNS alternatives face a triple challenge. The first one has to do with "good engineering": the security of the name resolution mechanism. Currently, the user's confidence in the result of the resolution comes from the fact that a machine has queried a known server. In a P2P environment, this "one-directional" validation mechanism disappears, resulting in a scenario where any participant node in the system can contribute anything and everything to the DHT, without a server acting as a "legitimizing" authority of the validity of that information. The CoDoNS project solved the problem by applying DNSSEC[16] to its system, a technically correct way to address the problem. But in doing so, one has simply changed the resolution system, while the registration infrastructure and its governance remain the same, with their flaws[17]. More generally, there is increasing evidence of the fact that achieving complete security may be impossible in a "pure" P2P environment[18].

In the case that any of the decentralized DNS projects matures to the stage of a relevant user appropriation, the crucial issue may become trust of users in other users:

> "With the current setup, we are putting our trust in the DNS servers like OpenDNS, Google DNS etc. to point us to the right direction when we want to access a website. With the scheme that P2P-DNS is proposing, we will have to rely on others in the network to direct us. It is one thing to trust OpenDNS, Google etc. but completely another thing to do the same with a random computer[19]".

---

[15] http://techcrunch.com/2010/11/29/peter-sunde-seconds-the-idea-of-an-alternative-root-dns/

[16] The Domain Name System Security Extensions (DNSSEC) is a suite of IETF specifications for securing certain kinds of information provided by DNS. It is a set of extensions to DNS which provide origin authentication of DNS data, authenticated denial of existence, and data integrity. The DNSSEC basically attempts to add security, while maintaining backwards compatibility, to a system – the DNS – whose original design as a scalable distributed system did not include security.

[17] DNSSEC uses the DNS tree structure for validating signatures.

[18] I.e., one that does not entrust any of its components, for example super-nodes, with any special oversight or management function over the system.

[19] http://digitizor.com/2010/12/01/the-pirate-bay-co-founder-starting-a-p2p-based-dns-to-take-on-icann/

Beyond choices of design and innovation, comes the issue of political governance, of which developers are acutely (and perhaps surprisingly) aware. The original questions that cause P2P DNS proposals to proliferate are deeply political: they are about control, freedom, and censorship. Says a developer:

> "This 'governance using infrastructure', as you call it, if it does not happen via the DNS, will happen via the Border Gateway Protocol[20], or via some of the many IP filtering mechanisms[21] that are out there on the Internet."

So, technical solutions to controversial issues that have a political component to them should, at some point, be accompanied by evolutions of institutions, lest the governance of the Internet be reduced to a war of surveillance and counter-surveillance technologies, of infrastructure cooptation and counter-cooptation. Ultimately, from technical and political actors alike, a common concern seems to be raised. The Internet may indeed find ways to "treat censorship as damage and route around it", as Internet pioneer David Clark once pointed out: technical design choices are as political as any law laid down on paper. However, in the long run, more sustainable solutions to restrictions of Internet freedoms are likely to be achieved by Internet governance institutions' capacity to engage in reflexivity and review: of their means, their aims, and their delicate roles in the management of today's foremost "global facility".

Conclusions

What do stories of infrastructure cooptation and "creative disruption", such as the decentralized DNS debates, tell us for the close future of Internet governance? A frequent critique of this interdisciplinary field is that it often tends to focus on a limited number of international institutions and debates about the global politics of the Internet. The "Internet governance" qualification does not generally apply to the study of a large number of activities and daily practices, on and with the Internet, that play a very important role in the shaping and the regulation of the "network of networks" [van Eeten, 2009].

In this context, approaches informed by an STS perspective on infrastructures, like the one adopted in this paper, can contribute to a disengagement from a conception of the Internet as an *a priori* identifiable and rigidly bounded space, either a stranger to the institutional forces of the off-line "reality", or on the contrary, entirely entrenched within the codified spaces of traditional politics [Cheniti, 2009]. This perspective allows to unveil the set of mechanisms that lead different participants in the technical, political and economic management of the "network of networks" to build common knowledge, legitimize some of it as "facts" of the Internet, and shape limits and boundaries able to reconcile the concerns of both experts and users.

[20] The Border Gateway Protocol (BGP) is the protocol which is used to make core routing decisions on the Internet; it involves a table of IP networks or "prefixes" which designate network reachability among autonomous systems.
[21] Techniques that control the incoming and outgoing network traffic by analyzing the data packets and determining whether they should be allowed through or not, based on a predetermined rule set.

This perspective also suggests that a different spin can be put on the "alternative Internet" debates that have thrived in recent years, particularly in the wake of the so-called Arab Spring. The label usually refers to social movements' appropriation of Internet-based communication and social media tools, with the aim of promoting "bottom-up" goals of reform and change of the political and social order [Atton, 2005]. But it may also acquire a deeper meaning: what is at stake is also a different architecture of the network of networks in its lower layers, with potentially far-reaching implications for the extent and the quality of users' control on their machines, data and exchanges, and ultimately for the very values underlying the global Internet.

In projects such as a P2P DNS, or the alternative P2P electronic currency Bitcoin, Internet users not only entrust the rest of the network with a portion of their software and hardware resources, but also rely on other users and machines in the network to manage information, communications, transactions. If users are accustomed to putting their trust in a classic DNS server, or in a central bank, to point them in the right direction when they want to access a website or legitimize the value of their currency, what does it take for them to do the same with a random, domestic computer? What values need to underlie the network's conception and implementation for users to be willing to turn their computational equipment into a part of the Internet's "phone book", or in a generator of virtual, decentralized currency, for the sake of a global, alternative Internet? How can political actors of Internet governance harness the potential for "disruptive change" [Rejeski, 2003] of the alternative Internet's recent and diverse manifestations, rather than being overwhelmed by it or choosing to oppose it by default?

With this paper, we have proposed some directions and tools to answer these questions, and contributed to the exploration of those new-born systems that – focused on users and self-organization, characterized by decentralized development and control – propose alternatives to the centralized infrastructures of today: not disposing of hierarchy, but reconfiguring it. In doing so, we have attempted to show how Internet infrastructure and technical architecture are today at the core of Internet governance debates and arrangements – not only as an *object of* governance, but as a set of *tools for* governance. This shift has important implications: these tools need to be used in ways that, albeit creative and potentially disruptive, will not pose a threat to the Internet's stability and security. Increased awareness, by all relevant actors, of what happens in the lower layers of the "network of networks" is needed, if the co-optation of Internet infrastructure, for functions its design is not suited to serve, is not to have unintended and gravely prejudicial consequences.

References

Abbate, J. (2012). L'histoire de l'Internet au prisme des STS. *Le temps des médias*, 18: 170-180.
Aigrain, P. (2011). Another Narrative. Addressing Research Challenges and Other Open Issues session, *PARADISO Conference*, Brussels, 7–9 Sept. 2011.
Akrich, M. (1998). Les utilisateurs, acteurs de l'innovation, *Education permanente*, 134 : 78-89.
Atton, C. (2005). *An Alternative Internet*. Edinburgh, UK: Edinburgh University Press.
Berners-Lee, T. (2010). Long Live the Web: A Call for Continued Open Standards and Neutrality, *Scientific American*, November 2010.

Braman, S. (2011). Designing for Instability: Internet Architecture and Constant Change. *Media In Transition 7 (MIT7) Unstable Platforms: the Promise and Peril of Transition*, Cambridge, MA, May 13-15, 2011.
Bricklin, D. (2001). The Cornucopia of the Commons. In A. Oram (Ed.), *Peer-to-Peer: Harnessing the Power of Disruptive Technologie*s (pp. 59-63). Sebastopol, CA : O'Reilly.
Cheniti, T. (2009). *Global Internet Governance in Practice. Mundane Encounters and Multiple Enactments*. Unpublished DPhil Thesis, University of Oxford.
DeNardis, L. (2013). The Emerging Field of Internet Governance, in William Dutton (ed.) *Oxford Handbook of Internet Studies*. Oxford: Oxford University Press.
DeNardis, L. (2012). Hidden Levers of Internet Control: An Infrastructure-Based Theory of Internet Governance, *Journal of Information, Communication and Society*, 15 (3).
DeNardis, L. (2009). *Protocol Politics : The Globalization of Internet Governance*. Cambridge, MA : The MIT Press.
Elkin-Koren, N. (2006). Making Technology Visible: Liability of Internet Service Providers for Peer-to-Peer Traffic. *New York University Journal of Legislation & Public Policy*, 9 (15), 15-76.
Flichy, P. (2007). *The Internet Imaginaire*. Cambridge, MA: The MIT Press.
Fuller, M. (2008, Eds.). *Software Studies: A Lexicon*. Cambridge, MA: The MIT Press.
Kirschenbaum, M. (2003). Virtuality and VRML: Software Studies after Manovich. *Electronic Book Review*.
Latour, B. (1987). *Science in Action : How to follow scientists and engineers through society*. Cambridge, MA : Harvard University Press.
Levinson, N. (2010). Co-creating Processes in Global Governance: the Case of the Internet Governance Forum. *Fifth Annual Global Internet Governance Academic Network Conference*, Vilnius, Lithuania.
Malcolm, J. (2008). *Multi-Stakeholder Governance and the Internet Governance Forum*. Wembley, WA : Terminus Press.
Manovich, L. (2001). *The Language of New Media*. Cambridge, MA: The MIT Press.
Marino, M. C. (2006). Critical Code Studies. *Electronic Book Review*.
Minar, N. et Hedlund, M. (2001). A network of peers – Peer-to-peer models through the history of the Internet. In A. Oram (Ed.), *Peer-to-peer: Harnessing the Power of Disruptive Technologies*, 9-20. Sebastopol, CA: O'Reilly.
Monberg, J. (2005). Science and Technology Studies Approaches to Internet Research. *The Information Society*, 21 (4): 281-284.
Mueller, M. (2010). *Networks and States: The Global Politics of Internet Governance*. Cambridge, MA: The MIT Press.
Musiani, F. (2013). *Nains sans géants. Architecture décentralisée et services Internet.* Paris : Presses des Mines.
Musiani, F. (2012). Caring About the Plumbing: On the Importance of Architectures in Social Studies of (Peer-to-Peer) Technology. *Journal of Peer Production*, 1.
Rejeski, D. (2003). Making Policy in a Moore's Law World. *Ubiquity*, December 2003.
Ribes, D. & Lee, C. P. (2010). Sociotechnical Studies of Cyberinfrastructure and e-Research: Current Themes and Future Trajectories. *Computer Supported Cooperative Work*, 19, 231-244.
Schollmeier, R. (2002). A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications. *Proceedings of the First International Conference on Peer-to-Peer Computing*, 27–29.

Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, 43 (3), 377-391.
Star, S. L. & Bowker, G. (2006). How To Infrastructure. In Lievrouw, L. A. (Ed.), *Handbook of New Media* (pp. 151-162), London: Sage.
van Eeten, M. (2009). Where is the Governance in Internet Governance? *GigaNet Annual Symposium*, Sharm-el-Sheikh, Egypt, November 14th, 2009.
van Schewick, B. (2010). *Internet Architecture and Innovation*. Cambridge, MA: The MIT Press.