DR. BERND KULAWIK (EINSIEDELN)

FROM TOP-DOWN TO NETWORK

Abstract

<1>

The internet and also the World Wide Web have been developed to make usage of the advantages of a real network, i.e. a network of communicating aquivalent and self-containing knots. But the structure of both developed into infrastructures of some ›higher priviliged‹ servers as ›sender‹ and a large number of clients or ›receivers‹. This is due to technical restrictions that we are now to overcome in the forthcoming years. At the moment, we are discussing how this infrastructure can be used in the Humanities. Since the development of Wikis started in 1995, more and more interactivity is made available inside this framework of this top-down structure. But recent developments have opened new perspectives to go back to the original structure of really net-working computers being part of a network of ›peers‹. Examples for these developments are online-repositories that can (partly) be accessed to publish data immediately, the technical advancements regarding online access to vast memory space and calculating power; P2P networks for the exchange of data and calculation etc. Usually, these developments are gathered under the metaphors of ›Web. 2.0‹ and the ›cloud‹. – But at the moment, any thinking about the process of exchange and publication of scientific data is limited to the older top-down structures and only does not take into considerations new developments. Therefore, these ideas and customs will be outdated in the very near future, i.e. less than 10 years or so: For instance, it is supposed that central servers will keep the data and will be managed by redactions with higher priviliges than the ›normal users‹ and authors. Instead, I want to argue here, the publication and exchange of data will be done from everybody's own personal computer, tablet, netbook etc. which will be connected online almost all of the time. Dedicated ›servers‹ will take over the role of an indexing system and a repository for the content of computers that are not online at the moment. This demands a radical new thinking about how and what to publish, how to ›control‹ the published data (if at all!) and how to organise a ›reviewing‹ process. I think, the basic structure of the web as we know it, will change radically into a real network of equal participants with equal rights, where real ›peers‹ will decide about the importance of any contribution. The times when special collectives with special rights (and interests!) decided these questions, will be over soon.[1]

Introduction

<2>

While the internet since about 1968 and also the World Wide Web since 1991 have been developed to make usage of the advantages of a real network, the structures of both changed into infrastructures divided into a number of ›higher priviliged‹ servers as ›senders‹ and a far larger number of clients or ›receivers‹. This was mainly due to technical restrictions that are already disappearing or will disappear in the forthcoming years. By ›real network‹ I mean a structure of ›knots‹ (computers) of equal importance and functions, i.e. they (can) serve as servers *and* clients at the same time – and they will be small enough to be situated not only in big server farms but will fit into everyone's pocket. From the user's perspective, one could also say that the structure we have now rather resembles a Master-and-Servant or Top-Down structure: We are transferring data ›down‹ from the server to the client and ›up‹ from the client to the server. Maybe it is no coincidence that we are used to call these processes *Down*- and *Upload*.

At the moment, we as historians in a broader sense are discussing how the old infrastructure could be used more efficiently in the Humanities, but at least since the development of Wikis started in 1995, more and more interactivity is made available even inside this top-down structure.

Now, recent developments allow to change and publish data immediately in a new way: We are not only able to use large Wikis like *Wikipedia* interactively from almost everywhere with our handheld computers and smart phones. We also use wide-spread web-based Content Management Systems, message boards or blogging software etc. which allow to edit and publish data online without having to learn rather complicated special mark-up languages or to use different protocolls (and software) like the File Transport Protocol (FTP).

Other examples are:

Google's several web-based services,

online-repositories synchronising local and distant data immediatels like *Dropbox*,

the *Opera* web-browser extension *Unite*, or

the web-based services *iDisk* and *me.com* by *Apple* and comparable services like *Ubuntu One*.

Other, more or less recent technical advancements extend (online) memory space in quantity and quality into the so-called ›cloud‹. In the ›cloud‹ users do not have to know where their data are kept or processed physically: Their storage server or application simply appears as an additional partition, folder or application in their local system – as long as they are online.

Or they combine the calculating power of thousands of smaller machines anywhere on the internet like SETI and similar projects.

Last but not least: Peer-2-Peer networks and Torrent structures pushed the exchange of large amounts of Data almost into new dimensions, where servers only keep the meta-information, especially about the location of some content, not the content itself anymore.
They all (and there may be some more) have opened new perspectives on the possible uses of the internet.
But, in fact, they are somehow pointing back to the original ideas and structures of really net-working computers being part of real network of peers. Therefore, I think, any discussion about the processes of exchange and publication of scientific information that focuses rather on the established top-down structures is limited – or limits itself – to an infrastructure that will be outdated in the very near future: We should not discuss these processes by supposing that central servers will always keep the data and will always be managed by some sort of authority with higher priviliges than the ›normal users‹ and authors.

<3>
Instead, the publication and exchange of information will be (and in parts is already) done from everybody's own personal computer, netbook, tablet or even mobile phone etc. These devices will be online almost all of the time and will stay connected and interact directly with each other – while dedicated servers will take over the role of an indexing system and a repository for the content of these world-wide inter-connected personal computers in case some of them are not online at a given moment.
For us as (art) historians, researchers, librarians, etc. this demands a radical new way of thinking about what to publish and how – and also: how to control the published data … if control is still wanted or needed at all. This implies the re-thinking of questions like: how will the (old) reviewing processes organised or better: re-organised.
I think, these topics and questions and the basic structure they rely on will soon change radically into a ›real‹ network of equal participants with equal rights, where real ›peers‹ will decide about the importance of any contribution. – They do this, in the long run, already since centuries by citing or forgetting – and sometimes: re-discovering – scientific contributions.
My opinion is, that the era of privileged groups of ›peers‹ who are a little bit ›more equal‹ than their peers in the scientific community and who have more and special rights ... but also: ›interests‹ – that this era will end soon and will be replaced by a more ›democratic‹ structure: »Peer-to-Peer-networks instead of Peer Reviews«, one might say.
After this long introduction that outlined almost everything I wanted to say here, there is still some time left. Therefore, I would like to say a little bit more about the structures and historical developments – ›historical‹ meaning here: the last 42 to 65 years! – that led me to these ›predictions‹ and opinions.

»Back to the future«

<4>

So, I would like to go »back to the future«. I could have called my paper so, but there is already a very interesting and much more relevant paper with this title from 1997, written by – among others – Alan Kay, the co-inventor of the graphical user interfaces we all use today, and of object-oriented programming. To Kay I will return soon.

I can only present some basic ideas here in the form of a draft: But on the other hand and in fact, I do not think that I myself – as a music and art historian ›playing around‹ with IT stuff and computers for a few decades now – that I could be experienced enough and could come to a conclusive paper without the input from others. So, this is rather a sketchy contribution, and a work in progress, too, than a finished paper, to start a discussion.

And, of course, this warning always is true: »Predictions (or forecasts) are always difficult, especially when they regard the future …«. But, as the Alan Kay said already decades ago: »The best way to predict the future is to invent it.«[2]

What does a (real) Net/Network look like?

Introduction

<5>

The starting point I want to use to dive into the history that will be our future, is the meaning of the *net* or *network* metaphor – and its differences from the *web*:

If we talk about networks or ›nets‹ today – especially about *the Net,* by which we now usually mean *the internet* and nothing else –, we do in fact not refer to ›real life‹ nets, and we do not even have real ›networks‹ right now: Neither the internet nor the professional networks among art and other historians are networks in a strict sense. And this is even more true in the context of peer-to-peer publication of scientific content. But my guess is, that this will change soon and that we should not try to block or slow down this development. Instead: we should rather encourage it: Because real networks – as I understand and want to explain them – will bring a significant change in the ways we do science and research – and a positive change, from my point of view.

So let me explain why: To begin with, I would like to very shortly explain the technical and structural background and history of what we now call the internet and its current similarities to the traditional ways of scientific communciation, collaboration and publication.

Examples of Nets

<6>

When we talk of nets or networks in ›real life‹, we may think of a fishernet or a spider web. For the letter one, the German word *Spinnennetz* – meaning spider ›net‹ – would even fit much better than the English ›web‹, because it is no web in the technical sense but rather a net. Other nets or networks are for instance that of geographical co-ordinates. And we could even call a technical static structure a network – like roof steel beem constructions, steel bridges etc.

All these nets or networks have something very important in common: These common characteristics of real nets or networks in the true sense of the word are:

- All knots or points of such a network are by definition (almost) of the same importance. For instance: There is no privileged point in a fishernet; not even in a spider web (though it may look like that for biological or constructional reasons) or in a network of co-ordinates. For instance, in geometry, networks have no privileged points: With a simple mathematical transformation you can put your center point, from where you start measuring, wherever you want. In the case of global co-ordinates there seem to be privileged points and meridians – but in fact they can be on the globe wherever we want them to be, only for practical and historical reasons it makes more sense and is more useful for us to put them where they are today. But the important point here is that they do not limit our usage of the net: For instance, if you are using a GPS to drive your car to Acquafredda, you do not have to go to a privileged point like the north pole first and then go from there to the end point of your trip.

- But even if there may be some knots more important than others in a net or network, they should not be of essential importance for the stable functioning of the network, that means: The network should remain stable even if single points or knots ›disappear‹.

- That means: There are many ways through the network from point to point, no privileged ways or privileged knots that have to be used always and cannot be circumvented.

Therefore: The first main point I want to make, is: In a real network there are no privileged points or paths.

Why a Web is Not a ›Net‹

<7>

I mentioned before, that I want to make a strict difference between the notion of a net and that of a web. Why?

In a web, the single threads or paths are not really tied together – the more open such a web is, – that means: the bigger the distances between its threads, – the less stable it is, because there are no ›knots‹ fixing every thread or relation in a more or less stable position. If we remember the fisher net or the spider ›web‹ as examples for real nets with knots, it is easy to see that no-one would use such a ›web‹ structure without knots to catch fish, to catch flies or to erect a load-bearing structure.
Webs are more intended to replace ›walls‹ or any ›separating membrane‹ with a more or less flexible structure.

Why the Internet is Not a ›Net‹

The following picture shows a typical representation of the traffic on the internet in Europe:

See: http://alimorton.posterous.com/europes-internet-traffic-love-a-good-map

As you can see in this representation, nothing resembles the ›nets‹ as we know them from our daily experience, but rather a structure of some central, more important points, that are almost connected with each other like knots in a real network, and many others, that are only connected to these central points. The central points are, of course, the most important servers in the system, why the computers connected to them may be seen as the ›clients‹ for the moment.

What I find interesting here is: This image may also exemplify the structure of scientific publication: The authors are represented by the smaller points or ›clients‹, while the servers resemble the journals with their redactions and reviewers. The readers, on the other hand, are then again in the position of the ›clients‹.
In any way we may interpret these image as a representation of the client-server-client structure of the internet or as a representation of the process of scientific publication and collaboration – in print on paper as well as on the internet: It is obvious that these are not real networks and that they do not own the characteristics or networks mentioned above: Not all knots or members of the network are of equal importance and have the same rights.

<8>
So, why is the internet not a net – at least: until today?
Most of all, there seem to be technical reasons for this structure of the internet as well as of the processes of scientific publication and communication:

- A computer functioning as a server has to have some minimal computational power. It can be replaced by a group of smaller computers that even could scale better under changing load or traffic: which usually is the case ›on the web‹. This is also the reason why virtualization over clusters of many but very simple machines is common today. But: Why should the members of such a cluster be grouped locally together in a physical server centre building at all? It's not necessary anymore, and the *cloud* mentioned above does function exactly in this way, as well as many examples of collaborative computing like SETI. Of course, there will always be very heavy computational tasks that require ›supercomputing‹ in clusters with special machines and special cables etc.: But as long as we or the other normal users on the web do not want to calculate the possible results of an atomic explosion, discover models of nucleids or identify radio signals from outer space, we do not need such a concentrated power in one place or machine. Rather, a distributed structure of small machines sharing tasks and replacing each other if needed, connected by a fast *backbone* structure and accessable anywhere for low costs seems to fit far better the tasks of normal users, including us.

- Another reason, why the ›(inter-) net‹ is not a net yet: There is only a limited number of internet addresses: But this problem will be solved soon (or is already solved) with the new version of the internet protocol, called IPv6.

- And finally, a political reason: Governments tend to think (and to put into legal praxis) that ›not‹ everybody should be allowed to have uncontrolled internet access and therefore be able to publish any content without restrictions: In Western Germany it was even forbidden to connect a modem to the telephone net in the 1980s!
So, how should a network – among persons and/or their computers – look instead? I think, it rather should look and function like the examples of networks I mentioned above: Like the fisher net, the spider net or the geographical net. Of course, there could be privileged points or persons in this kind of network, too, – but:

- Their authority should not be granted by and should not depend on some higher ›authority‹.

- Rather, their role and authority should depend on the merits they earn among their peers and the connections they established with others.

- We should also keep in mind, that this kind of network is again a rather simple, two-dimensional representation: Usually, there is a third dimension to such a network, be it the time component or the intersection with other networks. Now, imagine, every single person in such a network has his or her own computer connected immediately with each other. Or course, if necessary these single computers could and should be backed-up by others we might still call ›servers‹, but they should not ›depend‹ on any special server, especially not controlled by other, higher authorities.
What would happen?

- Every Person, every ›knot‹ in the net, could reach every other person without having to rely on the servers, or – if we translate this structure into the process of publication –: on the decisions of a small, select group of reviewers.

- Every one can establish new connections immediately and reach another person even by passing-by others who might get in their way.
So what makes me think that the technical structure originally was intended to establish such a network? And why do I think we are close to this point (again).

Some History

History of the Net and of web-based Collaboration

<9>

When we speak of ›the‹ (inter-) net, we usually think of the *World Wide Web*, created by Tim Berners-Lee in 1989/1990, the first successful implementation of a graphical system to use the internet already existing since ca. 1970. The internet itself was originally developed as the DARPAnet for military purposes: And the significant characteristic of this early internet was that there where no privileged computers, and that all relations/connections where (almost) equal. The idea behind this was simple: A tool for military (and political) communication and administration was needed in the case of a nuclear war: If one or more of those computers would be destroyed, there would still be ways for communication over other connections in this network, and, of course, functioning knots, i.e. computers, for administrative purposes.

Tim Berners Lee and the World Wide Web

When Tim Berners-Lee developed the *World Wide Web*, he used only two NeXT computers: Today, only one of them is presented in the show room of the CERN museum – no curator of the museum seems to have thought about the question how a ›single‹ computer could exemplify to the visitors the idea of a ›network‹, especially one that was intended to simplify collaboration and information among different persons via this network. Of course, it was possible to have two parallel systems running as a server and as a client on the same machine for testing purposes, but back in 1990 this would have been a heavy task even for an advanced computer built as a personal workstation, not as a dedicated server.

In fact, there ›is‹ a second NeXT computer used by Berners-Lee, today exhibited in a vitrine in another small show room above the large server hall of CERN's computing center: It looks very much the same like the one in the museum – and even the little note with the inscription attached to the machine and warning not to turn it off looks almost the same.

The main interesting point here is that both computers were running not only a browser that was able to display graphical content from the net, but also a server software: Therefore they both were literally server AND client at the same time to each other: with a 33/40/50 MHz Motorola processor and 8 to 16 MegaByte of RAM. Their operating system was the ›grandfather‹ or ›-mother‹ of what is running on my laptop here and even on my mobile Phone. I would like to mention that in 1993 this little phone here in my hand would have reached place 35 (more or less) in the first list of the *Top500 Supercomputers*. Its power is not only far above the one, Tim Berners Lee had at hand when developing his first version of

the world wide web, but would have made it or its younger brother probably the leading supercomputer in 1990.

Doug Engelbart

<10>

Anyway, this was not the first example of a network of (at least) two almost identical machines, working as server and client for editing and displaying content: The most famous example is even two decades older than the World Wide Web: It is the demonstration of a communicating system with a graphical user interface by Doug Engelbart, inventor of the computer mouse, in 1969, later even called: *The Mother of All Demonstrations*. If you don't know it yet, you should definitely watch in on the web![3]

What Engelbart is showing here is the result of his work trying to realize the ideas of Vannevar Bush from the early 1940s: A concept, Bush called *MemEx*, i.e. *Memory Extension*, and consisting of a large set of paper cards. Bush was working to find a technical solution for this – with textual or image content, interrelated by each other via some sort of footnote referencing system that had to be updated by hand.[4]

In the movie from Engelbarts demonstration you will even see some interesting things working on a large desktop computer in his small network of two connected machines that even today, our daily machines are not able to do – or have not been able until recently – like:

- video conferencing
- text written immediately on the screen.

Think – or look – twice: There is no window around the text, so there is no special application running to display it. In fact,

- Engelbart could point with the mouse anywhere on the screen and start to type or draw.
- Then, he could mark the area with the text/drawings and save it.
- That's it …
- The text would also be visible to his collaborator on the other machine in this little network.

One could see this as an example for a simple form of direct publishing inside a scientific community. That means: Engelbart did not have anything that would resemble our way of scientific communication and publication today:

1. open a special editing program and type some text or draw lines
2. save it as valid HTML
3. copy it to a server

4. make the server software recognize the new content, until recently usually by restarting the server program
5. send his colleague a link pointing to the website where he could find the new information.
6. And, by the way: discussing the text via video chat or making annotations immediately was also possible for the two ›netizens‹.

And that was in 1969!
I think you will have to admit, that this example from 1969 is still somewhere in the future.

Alan Kay: Smalltalk, Squeak and Croquet

<11>
A young colleague of Doug Engelbart was Alan Kay, who was involved in the development of *Smalltalk*, the first graphical user interface not only for programming but for almost any work on a computer. Based on *Smalltalk* he later created Squeak, »a *Smalltalk* environment written in itself« that provides »children of all ages« with the ability to program easily – and test their programs immediately. *Squeak* is an object-oriented programming and computing system running in a virtual machine on almost every hardware architecture long before anyone heard of Java being anything else than an island …
The work of Kay and his colleagues at the *XEROX PARC* (i.e. XEROX *Palo Alto Research Center*) inspired a young person named Steve Jobs to invent a personal computer with a graphical user interface … and, shortly thereafter, inspired another young man named Bill Gates to do the same. Unfortunately, XEROX was not interested in licensing their *Smalltalk* environment to these young ›amateurs‹, most of all, because these developments took place as experiments in the ›playing grounds‹ that XEROX maintained at the *PARC* to eventually get some better graphical interfaces for its copiers.
Squeak was for some years a leading base structure for Wikis, because the *Squeak Wiki*, or: *Swiki*, was running on almost all operating systems, easy to install and easy to use. When I first demonstrated *Swiki* to colleages in the year 2000 they could not believe that I was in real-time editing and changing the content of real websites hosted somewhere in the United States.
During the last years, Kay has led the development of a software derived from Squeak and called OpenCroquet: It is a virtual 3D environment for collaboration and consists of a server architecture running on any computer connected with others over a network. It does not need a special, central server like other 3D collaborative environments do, for instance *Second Life* or many online games. *OpenCroquet* allows participants to not only create three-dimensional spaces where each participant is represented by an avatar and which are

connected with each other through windows. *Croquet* also allows to launch programs on one participant's computer to work together on a file, or to work directly with 3-D-modeling tools in that common space. Any kind of objects can exist as 3-D-models in that space, and then may be annotated by other users. And because everything is an object here, you can set links from any other object to that special space, the object in it or to any of the annotations. There should not be any problem to find these annotations with a text indexing search robot like *Google* or *Bing*. So, it seems to be an ideal tool or, at least: model for scientific collaboration because it does not require participants to own any special software or to maintain any central server – or delegate its maintainance to any higher authority!

Conclusions

<12>

Let me draw some consequences that – in my opinion – will follow sooner or later from the changes in the internet's infrastructure that I described above:

1. One is, that we will first publish and ›then‹ review. That means: Any person can publish anything from his/her own computer – and others, i.e. ›all‹ peers, thereby for the first time becoming real ›peers‹ by doing so, willl discuss this published information, contribute to it, extend it, re-use (i.e.: cite) it in their own contexts, link to it or … even ignore it.

2. Another consequence, following from this first: The published information will not merely consist of ›stable‹ articles or books anymore. And the role of published findings/sources, ideas (what we now call *miscellaneae*) will increase.

3. There may or will still be ›privileged‹ knots in the emerging network(s). But their importance will result rather from agreement among the users then from privileged access to server machines, something we might call ›power‹.

4. This will also lead to more ›flexible‹ or rather developing publications: That means, there will not be one, never-changing publication to be cited, but rather a version history: *Wikipedia* and other systems like it are, in fact, already documenting the development of articles. One may think, this would be the end of scientific publications – but the opposite is true. In fact, there is even today nothing like a ›stable publication‹, because – as we all know – the ›interpretation‹ of the same words and letters printed on paper or even engraved in stone changes with time. And this is not only true for poetry, old philosophical or holy scriptures, but

one may encounter this fact in the form of difficulties to understand even classical texts of one's own discipline. – In the (near) future, it will be possible for the author to specify what was meant in the article by adding more information, not only by re-publishing it in a second edition with some new annotations reflecting others' additions, for instance, from a discussion board attached to the original article. This is an opportunity we should welcome, and not bemoan!

5. We will – hopefully – publish more sources in a much faster way than we did before. One might say: For this, we need database servers and their infrastructure. But I don't think so: If you put your transcription of a source or your pre-print article in the ›open‹ folder of your computer, that can be searched and indexed by a text-search engine: Then the provided information could be enough to let others find your contribution over the network. Of course, it might be easier to collect them ›centrally‹ on a big server or so – but in fact this is not necessary: The servers only need to contain the searchable information, the search engine indices and the links to the content. No ›central committee‹ or authority would be needed anymore …

6. If the sentence is true that »Knowledge does not consist of information itself, but of the linking of information« (in the German version: »Wissen besteht nicht in Information, sondern in der Verknüpfung von Information« where the word ›Verknüpfung‹ – i.e., ›knot-tying‹ – points to the essential network-like character of all knowledge within itself!) – so, if this opinion is true, then we will experience a far more intense creation of knowledge.

<13>
And: To connect the information does not mean to create one-way-streets which soon would become dead-end-streets. And it also does not mean to ›channel‹ information, to filter it based on (old-fashioned) scientific models or the latest concepts: It means that ›any‹ information at ›any‹ time should and will be available to ›any‹ interested person – and that this person should also be allowed to connect this information with other information.
For instance: I always made the experience, that one comes to new ideas rather in discussions than by sitting in front of a book or empty sheet of paper – sometimes even in the case that others seemingly did not contribute something new to this discussion. But during the discussion we may have realised that all the significant points were already there. Therefore, the extension of ›networks‹ in the structure of the internet as well as in the internal structural process of scientific work and publication will help to increase the efficiency of our work. One way to achieve this could consist in the publication of all relevant information

before a meeting: In my opinion, it would make much more sense to come together and discuss and develop the basic ideas that were already communicated among each other by publishing them on a web-platform or server (maybe even the small ones in our own pockets) than to come together and ›read‹ them to the others who are usually able to read themselves. Therefore, it would make much more sense to come together for discussions than for listening to other people reading their texts.

Originally I wanted to finish with this sentence: »Information wants to be free.«, i.e. the slogan of the Free Software Foundation. So, then, let's work to help it being free. – But I think the last word should again go to Alan Kay, citing the last paragraph of his article Computers, Networks and Education:

»I have no doubt that as pervasively networked intimate computers become common, many of us will enlarge our points of view. When enough people change, modern culture will once again be transformed, as it was during the Renaissance. But given the current state of educational values, I fear that, just as in the 1500s, great numbers of people will not avail themselves of the opportunity for growth and will be left behind. Can society afford to let that happen again?«

Author's Profile:

Bernd Kulawik studied physics from 1986 to 1988, and musicology, philosophy and art history from 1990 to 1996, master thesis in musicology about the *Seconda prattica* in the Italian madrigal around 1600 (1996) and PhD thesis in art history about drawings regarding the basilica of Saint Peter's in Rome (2002). Working with computers since 1982. Since 1988 several positions, mainly as database developer, at scientific libraries for music and art history (e.g. on Lineamenta). 2006–2008 Coordinator of the *Karman Center for Advanced Studies in the Humanities* at the University of Berne, Switzerland. Since 2010 working at the Bibliothek Werner Oechslin, Luegeten 11, CH-8840 Einsiedeln.

Contact: be_kul@me.com

1 For demonstration purposes, this presentation – including images – was served over a wireless local area network by an *Apache* server from an iPhone.

2 http://en.wikiquote.org/wiki/Alan_Kay

3 http://sloan.stanford.edu/MouseSite/1968Demo.html

4 Art historians may think here of Warburg's Collections