UNIVERSITY
OF TRENTO

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

38050 Povo – Trento (Italy), Via Sommarive 14
http://www.dit.unitn.it

LIQUID PUBLICATIONS: SCIENTIFIC PUBLICATIONS MEET THE WEB

Fabio Casati, Fausto Giunchiglia and Maurizio Marchese

December 2007

Technical Report DIT-07-073

Liquid Publications: Scientific Publications meet the Web

Changing the way scientific knowledge is produced, disseminated, evaluated and consumed[1]

Fabio Casati, Fausto Giunchiglia, Maurizio Marchese
University of Trento
firstname.lastname@dit.unitn.it
Version 2.3, October 1 ,2007

Abstract

The world of scientific publications has been largely oblivious to the advent of the Web and to advances in ICT. Scientific knowledge dissemination is still based on the traditional notion of "paper" publication and on peer review as quality assessment method. The current approach encourages authors to write many (possibly incremental) papers to get more "tokens of credit", generating often unnecessary dissemination overhead for themselves and for the community of reviewers. Furthermore, it does not encourage or support reuse and evolution of publications: whenever a (possibly small) progress is made on a certain subject, a new paper is written, reviewed, and published, often after several months.

We propose a paradigm shift in the way scientific knowledge is created, disseminated, evaluated and maintained. This shift is enabled by the notion of Liquid publications, which are evolutionary, collaborative, and composable scientific contributions. Many Liquid Publication concepts in this document are based on a parallel between scientific knowledge artifacts and software artifacts, and hence on lessons learned in (agile, collaborative, open source) software development. Liquid Publications concepts are reified by a model based on i) Scientific Knowledge Objects (SKOs), which are the digital instantiation of liquid publications, by ii) the processes involved in their creation, evolution, and quality assessment, and by iii) the people and roles that contribute to knowledge creation (authors, reviewers, bloggers..). Various models (including social reputation models) are developed to analyze and improve publication quality assessment and the process for attributing credit to and measuring reputation for individuals.

1. "The Web changes everything". But what about scientific publications?

The current Web and advances in ICT have created new scenarios that radically change the knowledge production process: we have virtually unbounded storage capabilities and essentially no limits in our ability to interact with peers. This new network- and storage-mediated knowledge production process is impacting on all aspects of knowledge creation on all types of knowledge. The Web is becoming the most extensive knowledge repository that ever existed. As examples of

[1] This is ongoing work. It is an evolution of the paper *"**Publish and perish: why the current publication and review model is killing research and wasting your money"**ACM Ubiquity 8(3), Feb 2007. http://www.acm.org/ubiquity/views/v8i03_fabio.html*

this trend, it is sufficient to consider the success of Wikipedia, or the upcoming initiatives such as Web Science Research Initiative (WSRI)[2], and the Faculty of 1000 initiatives[3].

Surprisingly, it is easy to observe that innovative forms of scientific publications are still lagging behind, and that the world of scientific publication has been largely oblivious to the advent of the Web and to advances in ICT. Even more surprisingly, this is the case even for research in the ICT area: ICT researchers have been able to exploit the Web to improve the (production) process in almost all areas, but not their own! We are producing scientific knowledge (and publications in particular) essentially following the very same approach we followed before the Web. Indeed, dissemination of scientific knowledge is still based on the traditional notion of "paper" publication and on peer review as quality assessment method. The current approach encourages authors to write many (possibly incremental) papers to get more tokens of credit, generating often unnecessary dissemination overhead for themselves and for the community of reviewers. Furthermore, it does not encourage or support reuse and evolution of publications: whenever a (possibly small) progress is made on a certain subject, a new paper is written, reviewed, and published.

The problem is analogous - and even worse - for scientific books in general and textbooks in particular. Textbooks today are written, reviewed, printed, and distributed. For a given topic, professors have at their disposal several textbooks to choose from. However, it is rarely the case any single one of these textbooks matches the professor's intended class topics and level of depth, and it is even more rare that the books are up to date, especially in a quickly evolving area such as IT. In the best cases, the authors "maintain" the book more or less up to date every few years, often with minor, poorly done extensions. As a consequence, a professor wanting to teach a class that is up to date is somehow forced to use parts of several books along with supplementary material taken from the Web. Again, there is little reuse, collaboration, and evolution in the generation of textbooks. Indeed, it is not surprising that some professors lately teach by using Wikipedia as a reference for the various topics, rather than textbooks.

This paper explores how ICT and the lessons learned from i) *software engineering* and ii) *the social Web* can be applied to provide a radical paradigm shift in the way scientific knowledge is created, disseminated, evaluated, and maintained.

Indeed, *software and scientific knowledge are more similar than they may look at a first sight. They are both complex "artistic" creation of the human mind which are "malleable" in that they can be modified and adapted (evolved) over time in several (sometimes wrong) directions and possibly in a collaborative fashion*. This malleability property implies that, unlike other "artifacts" characteristics of other engineering disciplines (such as all kinds of hardware, e.g., electronic or mechanical), both software and knowledge can be easily taken and evolved by anybody based on their ideas, preferences, or needs. Many of the ideas developed for producing and handling software, such as evolutionary and agile software process models and open source software development, can be applied to the production and management of scientific knowledge.

Another area from which we seek input is that of the social Web. Web technologies are providing us with ways to facilitate collaborative authoring and, equally interesting, **collaborative evaluation** of knowledge and artistic artifacts. Prime examples are measures of "interestingness" for pictures (see, e.g., flickr[4]), or processes for evaluation of Wikipidia entries and Wikipedia-like content. The social Web is also showing us the way in which social networks are created and how

---

[2] http://www.webscience.org/

[3] http://www.facultyof1000.com/

[4] www.flickr.com

they evolve, as well as how people establish trust and reputation in various communities spanning all sort of domains, from travel to art. While, in the past, quality assessment has always been based on “peer reviews” provided by a small set of “knowledgeable” individuals (movie critics, travel guidebooks) who took time and effort to provide their “expert opinion”, today various forms of feedback from the community are often used to establish absolute or relative quality of artifacts, products, and services (e.g., number of sales, feedback from hotel guests, social bookmarkings, or *associations* - “people who like your favorite book also like…..”). These are all lessons to be learned for how networks are formed and how reputation can be established in the community of interest of LiquidPub, the scientific community, where similar methods – and variations thereof - may be used to complement peer reviews and citation counts as the traditional forms of quality assessment in scientific knowledge production.

The goal to exploit these novel technologies to enable a transition of the “scientific paper” from its traditional “solid” form, (i.e., a crystallization in space and time of a scientific knowledge artifact) to a Liquid Publication (or LiquidPub for short), that can take multiple shapes, evolves continuously in time, and is enriched by multiple sources. The intended benefits of this approach are:

- To increase the early circulation of innovative ideas (and hence foster a more effective dissemination);
- To optimize the time spent by researchers in creating, assessing and disseminating knowledge, while improving the quality of the paper selection processes for conferences and journals;
- To facilitate collaborative research efforts that builds upon previously developed knowledge;
- To develop a new way of credit attribution process based on social networks, team/community work, collaborative problem solving, social reputation, and distribution of knowledge;
- To deliver innovative services and products for publishers of Liquid Publications, in order to add value to their traditional businesses.

In a nutshell, this paper provides the following main contributions:

1. It introduces the notion of *Liquid Publications* (and, analogously, *Liquid Textbooks*) as evolutionary, collaborative, multi-faceted knowledge objects that can be composed and consumed at different levels of detail.
2. It abstracts the notions of journals and conferences into *collections*, which are groupings of publications that can be based on topic and time but also on arbitrary rules in terms of what is included and how the quality of publications is assessed for them to be included in the collection. Collections can themselves be liquid. We believe that journals as they are conceived today (a periodic snapshot of papers on a given topic, selected by a restricted group of experts and based on submissions) will soon become obsolete both in their printed and electronic forms, and replaced by this higher level and more generic notion. This in itself is a major shift in the way knowledge is disseminated, and also in its creation and evaluation.
3. It proposes a radically different evaluation method for publications and for authors, based on the *interest* they generate in the community and on their innovative contributions and that is maintained in real time and possibly without reviewing effort (peer reviews can be used as a complement). The method also encourages early dissemination of innovative results.

Around these main concepts, we advocate the need for services that benefit authors, readers, reviewers, conference organizers, editorial boards, and even evaluation committees. Examples of such services are an analysis center for helping committees to assess the scientific quality of people and publications, ways for people to bookmark papers or people of interest and to define collections, and an authoring/sharing/versioning environment for maintaining and evolving liquid publications and for the fruition of their content.

2. Scientific Publications Today

2.1. Why we publish

The research world, and specifically the academic world, is centered around the notion of *publication* as the basic mean to disseminate results, foster interaction among communities, and achieve international recognition (and career advancement). Publications are done in conferences or journals, and are usually reviewed by a committee of experts, also referred as "peers". Typically, each paper is reviewed by 3 or 4 reviewers. The "best" papers among all the submitted ones are then accepted for publication in the journal or in the conference proceedings. In the computer science area, people typically publishes a dozen paper per year, and submit a little more than that (not all papers are accepted the first time around). Acceptance rates for conferences are often around 20% or lower[5]. Journals are published and printed (or, today, put online), and run in similar ways as conferences, only with continuous submission process, slightly different requirements (at least in theory) in terms of paper quality, and possibility of revision by authors after the peer review. Journal issues are published periodically.

There are three drivers behind this model:

1. **Disseminate ideas and make them visible**. Through publication and review, papers are made known to colleagues, and the review process is supposed to ensure that the best papers are more visible, so that researchers know where to go (good journals and conferences) if they want to read literature on certain topics. Publications also have legal implications as they "timestamp" work and ideas.
2. **Get credit, recognition.** Having papers accepted at prestigious conferences and journals is a way to prove (in theory) that the work is valuable. This in turn is a major criterion to determine career advancement.
3. **Meeting and networking.** Publications and conference participation leads to exchange of ideas with colleagues, and to networking. Conferences are also very useful for students to come and learn how the research community operates.

In the past, this model was working reasonably well, or, rather, there was not much of an alternative. Early dissemination as well as social or collaborative creation and evaluation of knowledge is unpractical without the Web. Papers, just like software, were produced using the waterfall model. Furthermore, until recently, scale was not a big problem. The number of papers submitted to conferences was under control. Without the Web, the pace of collaboration was a little slower, printed publication was the real mean of dissemination, the implementation of models other than peer review by a small group of experts was the only viable option: The

---

[5] http://www.adaptivebox.net/research/bookmark/CICON_stat.html

number of conferences was also much smaller, review cycles longer, and researchers were not forced to work in Web time.

2.2. Why this model is both unsustainable and unreasonable in the Web era

This model is incredibly inefficient under every perspective, results in a colossal waste of public funding, and forces researchers worldwide to waste countless hours that could be devoted to better research (or to have fun with family and friends). It is a system deeply rooted in the past, oblivious to the advent of the Web and related new forms of communication, information sharing, social networking and reputation. Here are some problems with the current state of affairs:

– **Too much time is spent writing papers rather than developing research**. Dissemination of results is important, and writing problem statements and results in a clear manner is also important. It is in integral part of the research work. This being said, one thing is to write papers with the purpose of making results available, and another is struggle to package and "sell" the work to try to get the highest number of papers published in the best conferences (or, in those conferences that guarantee career advancement in a certain institution). The latter is a huge effort and often results in papers that are incremental work with respect to previous research by the same authors.

– **The reviewing process kills good papers and is inherently flawed**. In general, reviewing a paper is not easy, and it is rarely done properly. There are many problems with the peer review process today:

  1. Judging the impact of a paper is very hard, in general. Even smart people and great researcher have a hard time assessing whether a topic is interesting and relevant and likely to have an impact.
  2. Sometimes good papers are cut because of bad reviews. It is not unheard of to have a paper rejected by a conference and win the best paper award at the next one. The main reason is that only one bad review is often enough to kill a paper. Reviews are often inconsistent, sometimes an author gets reviews criticizing the paper and saying opposite things.
  3. There are reviewers who are generally more negative and some that are more positive. So it is often a matter of luck to a certain extent whether your paper gets accepted. Clearly good papers eventually go through, but sometimes late and after a lot of reworks.
  4. Reviewing takes time, and is not necessarily time that results in better papers. Reviewers, especially scrupulous ones, spend a lot of time in doing reviews, and authors spend a lot of time adapting and tuning the paper not so much for the sake of making the best possible explanation, but to please reviewers and the conference style. While improving papers following comments is a good thing, very often one has to fight with meaningless or contrasting comments as well as space limitations that make the whole work cumbersome. Furthermore, sometimes there are certain styles of writing papers that is better accepted by reviewers, or that reviewers feel particularly bad in rejecting.
  5. A common effect of this review process is that many conferences tend to accept very detailed papers resulting from very detailed studies, rather than more innovative and creative papers.

– **Limited dissemination.** The entire review process itself limits dissemination (unless people post the papers on the web, which is a different kind of "publication", and likely a more appropriate one): reviewing introduces delays and if the paper is rejected then 6 more months will pass till the work has the chance to be published. Moreover, and very curiously indeed, research sponsored with public money is given to private publishing companies that profit from it and that sell papers. Furthermore, although it is nice to have papers in front when hearing presentations, printed proceedings by institutions tend to increase the cost of conferences.

- **Sudoku research.** In the mind of many researchers, the goal is not to make contributions, but to do publications. These researchers tend to work on problems if they are "publishable", not if they are useful. Unfortunately the two things don't always go hand in hand, to put it mildly. The result is that people make up problems (often creating nice models and languages supported by formalizations behind them and along with mappings between models and languages….) and then solve them. This is much like what happens with the now popular Sudoku games. In research we have Sudoku researchers, Sudoku papers, and even Sudoku communities. A side effect of this is that recently we in the academia tend to lag behind industry in innovation in computer science, and this should be a cause of concerns for all of us.

Furthermore, the current publication model, and even the notion of "publication", are rooted in the past. If academic research was born after the Web, we would not even be talking about publications as they are today. With a printed paper model, typical of journals, one needs to have the notion of publication, which happens periodically. If the authors do some extra work or have new findings, they need to write another paper, they cannot update or extend the current one. If people want to comment or discuss on the paper, they need to do this via email and via private discussions with the authors. Of course there is the issue of how to evaluate and give credits to people, but that is a separate matter (see below). With the Web, this is not the case, and there is no reason for the "publication" model to go on unchanged.

Taking the criticism further, one could easily argue that academic journals are not needed any longer. We don't read journals today, *we read papers in journals,* and we don't really read them from the printed journal, but we download and print them from the Web. The importance of being in a journal is credit attribution from the perspective of the authors, and quality measure from the perspective of the reader, meaning that we take the time to download a paper published in a journal because we assume it has been properly reviewed and approved. So the point is not about having journals with papers, the point is that of ensuring trusted quality evaluation and credit attribution for each scientific contribution. Once this is done, whether the paper is in this or that journal or this or that issue is irrelevant. In the end, the fact that in the past there used to be a hardcopy publication is what drives this whole model that we (community of bright researchers and bold innovators) seem to be incredibly slow in changing.

Notice that our goal in this paper is not to say that the goals are to be changed, or that people should collaborate in the interest of science and not worry about credit recognition. We do assume that one of the main, if not the main goal of researcher is to achieve reputation and career advancement. The goal of this work is (also) to help make the credit attribution process more fair and drive it towards rewarding (early) dissemination while reducing the dissemination overhead.

2.3. Previous attempts for changes

Despite these very significant shortcomings, the research community has been unable to come up with a better model. This is certainly also because the problem is hard in itself, but we suspect a significant reason is that people respected in the community are successful in the current system, and hence are not very interested in changing it. Besides, people are always so busy writing papers that it is hard to take a break and think about creating and pushing for a better system.

This does not mean to say that no attempts have been made or that the problem has not been studied. Over the last decades, there have been a few attempts to experiment with different models as well as to study in a scientific way the effectiveness of the current approach to paper evaluation and publication. In terms of conference models, variations include:

- **Peer-review with rebuttal** (e.g., ICSOC'05) **or double blind review** (e.g., Sigmod): unlike traditional conference review models where authors cannot reply, some conferences are experimenting today with rebuttal,, where authors have a few days to reply, in a few lines, to the reviewers to correct errors in the review. In theory, this is used as input in the discussion among PC members. In practice, rebuttal rarely leads to reviewers changing their minds, but it affects PC chairs when making decisions and, most importantly, leads to better reviews in the first place. Double blind reviews occurs when reviewers do not know the name of authors. There is contradicting research on whether double blind improves the fairness of the selection process.
- **Community review** (e.g., eclipseCon 2006): the community can vote on papers or on abstracts. There is no restricted program committee, the community decides what they want to be presented. This approach had very little success, for reasons yet to be fully studied and understood.
- **Open** (e.g., INFORMS): There is little to no selection, everybody can go to present. Participants can read abstract and exercise their own judgment with respect to what presentation they will listen to. Open conferences do not assign credit to the papers, though they are great for dissemination and networking.
- **By invitation** (e.g., in physics): the conference organizers invite people to come and give presentations. This appears to be good as it is a freeform way for the community to select top researchers to come to conferences. However it is not clear how to distinguish good conferences/meeting from average ones and at times, if people are not serious, it may be more based on friendships rather than scientific merit.

Journals also experimented with alternative models. The most significant one is ETAI, where papers are first put online and then reviewed, with comments openly posted on the pages (open reviewing) before a review process begin. For reasons that are still unclear, but probably related to the fact that people were posting comments in the open, this approach did not succeed and ETAI stopped publishing in 2002.

In terms of research on this topic, a few papers have been published on various aspects of the reviewing process, sometimes with contradicting results (see e.g., papers on double blind reviewing or repeatability of the review process [Tung, 2006; Madden, 2006; Fisher, 1994; Rothwell, 2000]. The conclusions are sometimes contradictory. There are no indications on which review process and model works best and no clear evaluation of benefits and shortcomings of each, so that program chairs and journal editors are still left in the dark and, in the absence of a clearly stated "better way", proceed with the status quo. This is often the approach that generates the least discussions: even if most people want a different model, they disagree on which one, so in the end it is sometimes just "easier" to keep going with the same old approach. However a large-scale study is still missing, and contributions mostly focus on small samples of reviews.

3. Steps towards a solution: Knowledge, software, and the social Web

Although the problems highlighted above are by no means limited to the IT area, in IT the magnitude of this problem is greater, due to the high number of researchers, conferences and journals, and to the heavy dissemination overhead (researchers write many rather long papers every year, and is not uncommon for researchers to do over a 100 reviews per year). However, we believe that the IT community is the one which is more suited to deal with these issues and to start a profound paradigm shift in the way knowledge is produced, as the IT community understands the technology which enables a paradigm shift (viciously, IT itself, the Web, distributed systems, …). We next discuss lessons that can be learned from software engineering and the social Web and then discuss how these can be applied to Liquid Publications.

3.1. Scientific knowledge and software

Software and scientific knowledge are more similar than they may look at a first sight. They are both complex "artistic" creation of the human mind which are "malleable" in that they can be modified and adapted (evolved) over time in several (sometimes wrong) directions and possibly in a collaborative fashion. This malleability property implies that, unlike other "artifacts" characteristics of other engineering disciplines (such as all kinds of hardware, e.g., electronic or mechanical), both software and knowledge can be easily taken and evolved by anybody based on their ideas, preferences, or needs.

We can therefore make a parallel between paper writing and software development. In software, the code is developed and then improved. New functionality is added with time, and the artifact is released and then improved. In extreme programming approaches [Beck, 1999], the code is also "evaluated" quickly in the process, rather than waiting till development is complete. Taking into account differences that do exist, one can borrow ideas from software development and try to apply them to writing. In software development, we do not change the name of a class each time we make a change to a function. We just release a new version of the class. Once a certain amount of functionality is developed, then the code is released for "testing". Similarly, with scientific papers, an approach that seems sensible is to publish versions of the paper when the work is sufficiently mature and clear so that somebody can read and gain insights from it, and then improve it. More importantly, minor changes (delta contributions) should not result in yet another paper (class) and yet another set of peer reviews as it is always the case today, but in variations or extensions to (versioning of) an existing work.

The development of a large program is a cooperative effort, while researchers compete more than cooperate, so this has to be taken into account. One sometimes does not want to release initial ideas for fear that they are copied, but usually this does not happen and whoever posts a version of a work has a significant lead on others. Besides, early posting, coupled with a secure and community trusted timestamp mechanism, gives people the right to claim that they have been the "first" to a certain discovery. Furthermore, the researchers keep the control on when they want to release the new version of a paper, and who to make it visible to, much like what happens today on photo sharing sites. Needless to say (but we say it nonetheless, as it is a goal the community should strive for), early releases contribute to science more than late releases.

Open source development can also provide interesting insights for the way people cooperate to provide feedback and improve the development. Again this is challenged by the fact that researchers are not very cooperative while open source development is often led by enthusiast that really use the results of what they develop. Still, it is a very effective way to improve and extend an artifact and it would be interesting to see what can be "reused" for paper evaluation and even improvement.

This similarity between knowledge and software is becoming more and more manifest as knowledge is increasingly available in electronic form. Indeed, we argue that (scientific) knowledge will follow a similar path to the one opened up by software since its decoupling from hardware over 50 years ago. The main difference is that it will follow this path much faster, as we can learn from lessons in software development. Indeed, as the rest of the proposal will make clear, many of the ideas developed for producing and handling software, such as evolutionary and agile software process models and open source software development, can be applied to the production and management of scientific knowledge. Again, this is more evidence that an ICT perspective (suitably integrated with inter-disciplinary competences to properly understand and model the knowledge generation process, as it is the case in this consortium), is a good standpoint from which to tackle this problem, and that the results produced by this project should have long lasting effects.

We do realize that this parallel cannot be taken to the extreme. Software and scientific knowledge do differ, if not else for the fact that one is to be consumed by machine and the other by humans. Hence we are only borrowing lessons from the progress in software development (especially over the last five years) as long as they help us reach our goal of a better knowledge generation and dissemination process. Said this, we also believe that this distinction will become narrower and narrower the more humans will consume knowledge mediated by computers. The extent to which computers will need to understand the semantics of the knowledge presented to users, will define the boundary between *knowledge-for-humans* and *knowledge-for machines*.

3.2. Scientific knowledge and the (Social) Web

Another area from which we seek input is that of the social Web. Web technologies are providing us with ways to facilitate collaborative authoring and, equally interesting, collaborative evaluation of knowledge and artistic artifacts. Prime examples are measures of "interestingness" for pictures (see, e.g., flickr[6]), or processes for evaluation of Wikipidia entries and Wikipedia-like content. The social Web is also showing us the way in which social networks are created and how they evolve, as well as how people establish trust and reputation in various communities spanning all sort of domains, from travel to art. While, in the past, quality assessment has always been based on "peer reviews" provided by a small set of "knowledgeable" individuals (movie critics, travel guidebooks) who took time and effort to provide their "expert opinion", today various forms of feedback from the community are often used to establish absolute or relative quality of artifacts, products, and services (e.g., number of sales, feedback from hotel guests, social bookmarkings, or *associations* - "people who like your favorite book also like…..."). These are all lessons to be learned for how networks are formed and how reputation can be established in the community of interest of LiquidPub, the scientific community, where similar methods – and variations thereof - may be used to complement peer reviews and citation counts as the traditional forms of quality assessment in scientific knowledge production.

Web search gives an almost instantaneous way to identify significant documents. One wonders how much of this can be applied to evaluate posted versions of papers. Today's approaches use page rank to rate documents [Brin, 1998] and citation/impact factors to evaluate papers (research document). More details are provided in the related work section. The problem here is how much of these can be leveraged to either "automatically" evaluate papers, or at least to assist reviewers or perform a preliminary screening.

3.3. Summary of ingredients for publications in a liquid world

In summary, the key learning and key ingredients of a novel approach to scientific publications are:

- *Evolutionary and collaborative approach to writing papers*. Borrowing concepts from collaborative software development and agile software development methods, we believe papers should be developed in incremental fashion and facilitating collaboration among groups. This is not just a matter of providing an authoring environment. We do have collaborative authoring tools at our disposal. It is a matter of changing the way credit is attributed on the one hand, and the mentality and habits on the other. Incidentally, successful groups are not shy or conservative in publishing their work early (e.g., Stanford posts technical reports way before the work is published).

[6] www.flickr.com

- *Separation of the dissemination, evaluation/recognition, and retrieval aspects*. Today, with a publication, researchers achieve all of them. A publication disseminates the work, causes recognition for the authors (the peer evaluation recognizes it as quality work), and makes the paper "visible" in that people can look on papers published in "good" conferences or journals if they want to find "good" work in a certain area. However, there is no reason for these three aspects to be tied now that dissemination is not necessarily related to the physical, paper printing of the scientific contribution in a journal.
- *Definition of fair, efficient, and innovation-rewarding quality assessment methods.* This objective can be achieved in different ways and has different facets. First, we need to understand and remove biases in the current peer review model. This is true even if no change is done to the publication model. Second, we need to define and experiment alternative, and possibly social evaluation methods where publications are rated or ranked based on interest from the community and especially (but not only) from reputed individuals in the community. This goes also past (or complements) the time-consuming peer review approach and allows for an almost effortless evaluation, much like what happens today on the Web. Third, we need to develop models that, while not discouraging the still important delta contributions, adequately reward early dissemination of innovative ideas. Key to all this will be the acceptance of these innovative evaluation models by the community, including the bodies in charge of deciding on careers of individuals.
- *Sustainable business models for publishers and service centers*. It is now recognized that the current business model scientific publisher are following is not sustainable [EU pub market, 2006]. People are simply not going to pay to access and download papers, and rightly so, in our opinion. Yet, it is not easy to imagine scientific dissemination without somebody picking up the baton of making them available along with additional services. In a liquid world, today's publishers (or new players) will likely offer more advanced services such as maintenance of scientific social networks, semi-automated generation of related work, automated notifications of new contributions in a certain area, social bookmarkings, and the like, along with collaborative authoring, blogging, and reviewing. They will be the analogous of the yahoo, flickr, digg, and delicious of the publication world.

4. Liquid publications

We aim at achieving the above objectives by modeling "papers", publications, and scientific knowledge in general as a complex (software) system characterized by three main interacting components: **Scientific Knowledge Objects, People,** and **Processes.**

4.1. Scientific Knowledge Objects

Scientific Knowledge Objects (**SKOs** for short) are the digital counterparts of the traditional notion of scientific paper, which we view as complex, social, evolving software objects. We consider SKOs as being:

- *Evolutionary*: SKOs and their constituents evolve (possibly in a continuous fashion) over time as people contribute knowledge to them. Hence, they exist in multiple versions, at different and hopefully increasing degrees of maturity. This is somewhat analogous to what happens in agile software development methods.
- *Collaborative*: They enable the collaboration and contribution of a number of interested researchers on a specific SKO, each with different levels of "ownership" and control on the SKO, and each able to claim credit and responsibility for the contribution. The combined effect of collaboration (possibly open to the entire community, where different people may

have different perspectives, opinions, or interests) indicates that SKOs may evolve in a tree-like fashion as new branches are created by groups or individuals wanting to explore a certain aspect or to follow a certain approach on the research addressed by a SKO. This resembles open source software development where people collaborate and where in case of disagreement or in case of desire to explore different development directions, different branches are created in the codebase.

- *Multi-faceted* and *Multi-purpose*: SKOs are complex objects in that they contain different kinds of content that have different purposes, but all of which contribute to the knowledge embedded in the SKO. Such content may include text (the "paper"), associated images, videos, and slides, experimental datasets, but also reviews and feedback by the community, which are also a form of knowledge and contribution. Similarly, a software application is composed of code, specifications, manuals, and, from a developer's perspective, of test cases, bug lists, code reviews and the like.
- *Composite* and *composable*: They support the creation of new SKOs by composing (and extending) existing ones. This is applicable to paper but even more so to textbooks, where professors could compose and create a customized (and possibly liquid, hence evolving) textbook for their class by putting together composable content from other "liquid" textbooks.

An example of SKO is this very paper. This version is an evolution of a previous one, it has been developed collaboratively and posted on the Web so that it could get comments and feedback, and indeed we did get *a lot* of it. Early versions have been continuously released. Right now it does not yet have other material for the simple reason that such material does not exist, but once available, material such as results of review analysis, presentations, and the like will be available to complement this paper.

Another example is a textbook. Ideally, a textbook should have companion material (as is the case today with companion Web sites), but most importantly it should be evolutionary and collaborative (with the community empowered to complement it or create versions and branches) and, ideally, even composable (to allow the creation of novel textbooks out of many existing ones). We believe that static, printed, and entirely author-controlled textbooks in rapidly evolving areas such as most of the ones in computer science are obsolete, and will disappear very rapidly as more and more professors share content (today in a somewhat unstructured way) and refer to online sources such as Wikipedia, white papers, online presentations or videos, and the like.

In the following we use the term SKO to refer to the digital object itself, and hence to the IT aspect of the knowledge creation and dissemination problem. We use instead the term **liquid publication** to refer to the concept of an evolutionary and possibly collaborative publication which, in this project, is reified and made possible by SKOs, its processes, and the roles that enact these processes. In contrast, we use the term **solid publication** to refer to the current notion of publication as a contribution that is written by a closed circle of authors, reviewed, and published (typically in print), and that is then "set in stone" in the sense that it never changes from that point on. In a way, solid publications are analogous to waterfall software process models.

4.2. People

People are the agents involved in the scientific knowledge processes, playing various cooperating and competing *roles*. Some of these roles come from the current practice (e.g., *authors, readers, reviewers*), some other come from a refinement of these traditional roles (e.g., one can be *an author, a reader, or a reviewer of a paper, of a draft, of a revision of a paper, of an incremental result on top of a previous paper, of a review of a paper, of a comment of a paper, of a comment* *on a comment on a paper*, … and so on), while others are totally new roles deriving, for instance from the new collaboration paradigm of the Social Web *(e.g., taggers, bloggers, bookmarkers, content aggregators, classifiers, SKO quality certifiers, credit certifiers …*and so on) and more. Each of these roles contributes to the creation and evolution of a SKO. For example, reviewers are a specific kind of readers, and if they provide feedback then they also become, in our paradigm, a specific kind of authors. In our approach, the scientific community at large and to some extent also people as a whole have the possibility to play a central role towards the creation of knowledge and towards the credit attribution process, as well as in ensuring quality and trust on authors and reviewers.

4.3. Processes

Different **processes** are adopted to create, evaluate, and evolve SKOs and hence to manage their lifecycle. In particular, different processes can be supported for allowing different ways of:

- Editing and evolving SKOs and their components;
- Enforcing various degrees of control by SKO owners while still allowing and encouraging the community to create knowledge and to disseminate;
- Managing intellectual property and other legal aspects;
- Supporting the assessment of the quality of both contributions (SKOs) and contributors (People), also ensuring that proper credit is given to novel ideas and to their proponents, so that early dissemination is encouraged.

The description above confirms the close parallelism between the lifecycle of liquid publications and software process models. This is rooted in the considerations reported at the beginning of the section. Indeed, the paradigm shift we advocate for scientific publications is similar to what started to happen a few years ago in software engineering with the progressive adoption of more agile and iterative development processes, from the spiral model to extreme programming up to the "social", open source development. The current software engineering practice can teach us a lot in our move towards liquid publications. Thus, for instance, the open source and Web communities can inspire us on how to perform a "validation" of the people's work and "credit attribution" (which is key to people's careers and goals) that is fair, relatively accurate, that allows for the production of high quality artifacts, but that is lightweight (relative to the current practice in peer review) in terms of requirements in reviewing time.

The software engineering analogous to sticking with the current publication model would be to have software development only done by very small groups, to base development on the pure waterfall model, to create applications with character-based interface and no additional explanatory material, and to redevelop the application from scratch each time an extension is made. This is as unreasonable for software development as it is for scientific publications, but it is what is happening in publications today, although we are seeing some initial attempts to break this model (such as *wiki-papers*). IP management is also another aspect that is showing increasing convergence between knowledge and software, so that licenses such as the GPL (GNU Public License) or BSD, or the various Creative Commons licensing schemes provide inspiration for IP management in scientific knowledge.

5. Conferences, Journals, and Collections in a Liquid World

The traditional notions of *journal* and *conference* will still apply in the world of liquid publications, though they will most likely change, possibly in a very significant way, with respect to what they are today. For example, assuming that all contributions are available in SKO repositories, as we foresee in this project, “solid” journals could be a “collection” of snapshots of (liquid) publications related in some way (e.g., by time, by topic, etc…), while “liquid” journals could be an analogous collection but evolving along with the liquid publications they collect. Similarly, conferences could operate by selecting and inviting “best (liquid) papers” in a certain area, where the evaluation of what is “best” can be done with traditional peer review as today or can leverage continuous evaluation mechanisms developed as part of this line of research and that can for example include community feedback, downloads, and citations (which, in a liquid world, are available in “real time” as dissemination occurs continuously). Finally, besides conferences or journals, it is likely that novel forms of collection of papers will arise, as in the end all that matters is the ability for readers to find good papers in a certain area and for authors to be recognized for their contributions.

5.1. Solid vs Liquid collections and social bookmarking

Today, conferences and solid journals do not differ much from a publication perspective (this is true for computer science, we are aware that in other disciplines they are indeed very different). They are both essentially based on taking a snapshot of what's available at a given time. They are indeed characterized by *time* when being referred to (e.g. ICSE 2007, or IEEE Computer, May 2007) and, as discussed, inclusion in an issue or in proceedings is based on peer review. Besides conferences and journal, we can also however think of *liquid collections*, which are groups of contributions that evolve with time, possibly in a continuous fashion, and where the scientific topic and the rules for inclusion or exclusion take precedence over time, though time is a dimension. There are many variations of these models, both for journals/conferences as well as for generic “collections”:

1. *Liquid Journals:* Once a set of papers has been identified as worthy of inclusion in a journal, their liquid version as opposed to their static snapshot is added to the journal. So if a new version of a paper in a journal is published by the authors, it is now the new version that belongs to the journal (possibly after some reviewing, the review process is not our concern in this section - see later). We call these as liquid journals. The number and kind of papers are static, but the journal evolves with the included papers. We may still have the notion of “issue”, though it begins to be decoupled with time.
2. *(Solid or Liquid) Topic-based Collections*: As another, quite different and likely more useful model, we may think of a variation of a journal which is in fact topic-based collections, such as a *papers on "data warehousing in outsourcing contexts"*. This collection will contain the most important papers in the area. As a new paper which is "good enough" or "worth reading" becomes available, it is added to the collection. Again there are many (possibly subjective) ways in which one may define if a paper is good enough or worth reading, but that's not the point here. The point is that we have a topic-based aggregation of contributions which evolves with time and is extended as new contributions appear. Yes, this means that this "journal" can grow indefinitely (so ratings and rankings become important), and also that it may go out of fashion as maybe 10 years from now a journal on "data warehousing in outsourcing contexts" becomes irrelevant. (side note: this happens also today, it is just that journals and conferences keep the name but change topics. for example very few paper on the VLDB journal are actually on very large databases). When this collection will go out of fashion people will simply stop maintaining it or looking at it. Notice that we do not say anything here about who maintains the collection. It can be individuals, institutions such as ACM, committees such as those that run conferences today, and so on. Also, the collection can be structured into subcollections, based on subtopics or even time. A journal today is an example of a solid collection that has subcollections based on time.

3. *Generic Collections*: The generalization of the model above is the case in which we have a set of "rules" or criteria for selecting papers to be put in some kind of collection. here is where creativity comes in :). So, for example, we can think of a collection such as "papers on knowledge management by Trento", or "vision papers on the future of software engineering", or "experimental papers by the group of John Smith highly rated by the community", or "most influential papers of the last 12 months, plus papers published in the last ICSE", or "papers that are liked (downloaded, bookmarked) by Vanessa and George".

The bottom line of the last and more general category is that papers in principle may join and leave the collection at any time (so the collection is liquid and versioned), possibly based on a combination of automated rules and reviews. Or yes, the collection can also be branched: I can take a defined collection done by Bob, or defined by ACM, and create a variation for my own use, based on what i think is important for me. And, of course, i can be notified when something is added or removed from the collection if I want, and go read the new paper, or perhaps override the addition or removal from the collection.

Observations:

- Depending on the approach there can be PUSH (continuous) or PULL (discrete) collections. PULL collections are analogous to conferences: they are initiated by some action (call for papers, and review) and end up at some stage with the new version of the collection. PUSH or continuous collections may change, possibly automatically, whenever there is a new contribution (or version thereof).
- If this looks very much flickr-ish or facebook-ish or delicious-ish, well, that is indeed one of the directions we may end up going..... note that this is effortless reviewing in a way, but none of the above excludes a partial or exclusive use of peer reviews. You can implement all the above collection types with peer review if you want. That's one particular rule.
- As mentioned earlier, one begins to wonder what's the point of journals in this new world. Indeed, we do wonder. Nobody reads "journals" anyways. We read *papers* in journals. Which in most cases we get from the Web. We care about the fact that they are in a journal only because we believe this supposedly means it is a good paper. The constraint that forced us to have solid journals in the past are now gone. What we want is ranked search and dynamic collections.

5.2. DEFINING AND GENERATING COLLECTIONS

Discrete/Pull-based collections

The definition of pull collections, such as conferences or journals, is based on two key aspects: **selection processes and metrics**. A selection process is characterized by a set of steps where at each step an action is taken, typically to select, exclude, or rank papers. Metrics are qualitative or quantitative indexes that measure the "intrestingness" of a paper.

Actions (process steps)

We next describe examples of actions related to the review process, grouped in macro areas.

- Call for "submission"/submission
- Peer review
  - With x number of reviewers per paper
  - Visible, hidden, double blind
  - Breath or depth driven, or mixed. depth driven is somewhat traditional: a small

number of reviewers per paper reads the whole paper. breath driven means that a large number of reviewers reads an extended abstract. or even an unextended one. or just takes a decision by browsing a paper. mixed approach means a couple of reviewers read all, many other browse.
    - Closed or open: closed to a committee or open to anybody
- Rebuttal (author replies)
    - based on comments, or also with score visible
    - shorter vs longer replies allowed by authors
- Discussion
- Marking (Selection or exclusion)
    - some papers are excluded from the candidate ones, or some papers are marked for acceptance
- Export / import
    - make reviews of rejected papers avail for the next "submission" analogous to paper rollover models

Examples of peer review processes

Below are some examples of description of processes that recent conferences have followed.

WISE'07 (standard peer review)

1. get submissions (papers)
2. peer review (hidden, depth-oriented, 4 reviewer per paper)
3. discuss (mainly reviewers to agree on marks, and PC members)
4. marking (final eval)

SIGMOD'06 (peer review with rebuttal)

1. get submissions (papers)
2. peer review (hidden, depth-oriented, 2 reviewer per paper)
3. marking (reject papers with 2 reject marks)
4. peer review (assign the paper to one or two more reviewers)
5. rebuttal
6. discuss (mainly reviewers to agree on marks, and PC members)
7. marking (final eval)

ICSOC'07 (peer review with area chairs)

1. get submissions (papers)
2. peer review (hidden, depth-oriented, 4 reviewer per paper)
3. discuss (mainly reviewers to agree on marks, and area chairs/PC chairs)
4. marking (final eval)

Metrics (support process decision and the marking steps)

Orthogonal to the process issue is the issue of marking and ranking, which is in fact a typical and crucial step in the processes described above. The metrics below are intended to provide selection committees or even individuals defining rules for their collections with the building blocks of a kind of **analysis center** that can support the decision, but without relying reviewers in answering the dozens of scores that are sometimes asked in conference review form, rarely filled with care anyways and rarely used by the chairs. Most of the below metrics and bias compensation techniques can be computed automatically and complement the overall score given by the reviewers.

- De-biasing (bias compensation)
  - reviewer positive/negative bias (some reviewers tend to be considsently more positive or negative than others, e.g. some always reject)
  - reviewer scale bias (some are more extreme than other)
  - confidence bias (compensate for the fact that low confidence reviewer tend to give more accepting marks, or to flatten marks in the middle - if that's the case)
  - gender and geography or institution bias
- Metrics based on reputation of reviewers
- Metrics based on reputation of authors, institutions (this is in a way the opposite or removing institution bias)
- Metrics based on interestingness of the topics
- Post-posting metrics: these are applicable when the contribution is available (and has been avail for a while on some web site) before it is "evaluated". It is a common approach we expect in liquidPubs
  - count citations, count bookmarkings/social tagging
  - downloads
  - accesses w/o downloads
  - public ratings
  - number of comments entered, possibly consider "reputation" of bloggers
  - the above with compensations for attacks, or for self-sustaining communities
  - metrics that take into account agreements between open/social and peer reviews

Push-based collections

The above metrics can be the starting point for defining push-based, automated collections of paper, where quality is dynamically evaluated and papers enter or exit the collection, and in any case are ranked according to different criteria thereby helping people 1) search for interesting papers and 2) be notified when interesting paper in a given topic are published. Examples of criteria for search and ranking are classic approaches based on e.g., papers by authors who are often cited or included in pull-based collections. However, more interesting techniques are based on social bookmarkings for ranking papers, as briefly discussed below.

PubRank, PeopleRank, and encouraging dissemination

PageRank is at the heart of google’s search technology, and it proved to be very effective. The interesting aspect of *both* social bookmarking and push-based collections in scientific publication is that variations of this method can be used for

1. contributing to measuring the quality of scientific publications (and hence to ranking them)
2. contributing to measuring the quality of scientists
3. encouraging the early release of innovative content, also reducing the fear of having credit stolen by having the content stolen.

Indeed, inclusion of a publication among one’s bookmark is a sign of interest in a paper. In a way, this is analogous, as we said, to a citation, with a twist or two: first, it is not a publication who makes the link, but a person. Second, one can bookmark a specific version of a SKO, or of any version of a SKO, or of any branch. Bookmarks of SKOs and versions are also, in most cases, an indirect declaration of interest in the scientific production of the authors of those SKOs.

Furthermore, the bookmarking can be *qualified*, so that for example I may mark some papers are particularly interesting, or as landmarks, or as groundbreaking in an area. Hence, it is much richer than a citation. It may also fade with time, so that for example with time I may find that a paper I thought was interesting is in fact incorrect or irrelevant and remove it from the bookmarks.

Collections (including in particular push-based collections and subscriptions) are also declaration of interest. They are not however necessarily directed at a publication. They may be directed at a topic, but, more interestingly, they can be directed at people, as in "notify me of new SKOs or new *major* versions of SKOs by John". Notice that a subscription from content by a person is stronger than a bookmark on a pub by that person. By subscribing, I am stating that I am interested in all content produced by that person. Each of us has a few researchers we think of very highly for which this is the case. This is implicitly a self-regulating mechanism: I am not just going to subscribe to content from a lot of people simply because I do not want to be flooded with emails or pubs. Also, I am not going to subscribe to content from people who do tons of delta pubs, for the same reason. Finally, and for the same reason, I am likely to give priority to people who are innovative and release results early.

The above observations suggest a possible way to measure the "reputation" of pubs and people, via PubRank and PeopleRank algorithms that treats bookmarks and subscriptions similarly to how PageRank classifies pages based on hyperlinks.

Hence, we rank higher pubs that have i) a high number of "declaration of interests" (via bookmarks or inclusion in collections), ii) by giving more weight to declarations coming from highly reputed people and highly accessed/reputed collections, and iii) by giving more weight to declarations coming from people who have few bookmarks. Of course, many variants can be designed (changing weights, discounting the time of citation, and including the "type" of bookmarkings and personal (private) ratings as part of the algorithm. Notice that the fact that bookmarks are used for ranking does not mean that they have to be public. They just have to be known to the system.

An analogous metric can be used for people, but this time based on subscriptions to contribution by a person in addition to bookmarks on pubs by a person. We rank higher people for which not only the community is interested in their pubs via citations or bookmarks or collections, but also, and mainly, people for which the community is writing a blank check by saying "I want to know about any contribution you *will* publish".

Again, there are tons of variations of the basic algorithm framework, e.g. one can take into account subscriptions more than bookmarks, or can consider the specific kind of subscription (e.g. I want to know ALL that this person writes, vs I want to know all that this person writes but only on a specific topic, etc), or can consider the breadth of the communities that are inertested in the person's work, etc.

6. Success measures and support from the community

Right now it is hard to predict the extent to which this paradigm shift will fly. The success criteria are based on:

1. **Acceptance by the community of the SKO concept.** This can be qualitatively and quantitatively measured by observing the usage of the platform by scientists. It is the ultimate measure of the success of the approach. A preliminary measure will be the interest of conferences, events, on-line journals and communities in using the LiquidPub approach and tools;

2. **Level of reuse.** Measures the extent in which SKOs are reused to create "publications" as opposed to creating new SKOs from scratch. Measured as percentage of SKOs which are evolutions of previous ones or that reuse/integrate parts of other SKOs;
3. **Increase in the quality of the assessment process.** This metric can be measured in the short term by computing the reviewers' consensus and the reduction in gender, geography, or other kinds of biases, as well as with the review quality metrics developed within the project. In the long term, the quality of assessment will be determined by looking at which institutions consider the LiquidPublication assessment process as the basis for career advancement and other credit attributions. This aspect will strongly depend on the type of scientific community that will use SKOs as building blocks of the members' career tracks;
4. **Ability to support the existing publication and review processes,** measured in terms of number/percentage of covered processes;
5. **Measurable reduction in the time spent in dissemination (writing papers) and paper review time**, measured via surveys among researchers.
6. **Improved fruition and consumption of scientific knowledge** thanks to the collaborative and multimedia nature of the artifacts composing SKOs and to search and navigation features that take advantage from using the quality assessment processes. A series of surveys analysis among researchers will be organized to understand how effective SKO services are. Also, a set of indicators will be defined to evaluate the trend of innovation and knowledge creation into SKOs. The underlying knowledge management theories (called Intellectual Capital Theories [Stewart, 2001]) tend to make explicit the value of knowledge and are already experimented in business scenarios.

The most critical aspect is getting buy-in from the community. To get things started, we have prepared an EU FET project proposal that has the support and commitment of the main bodies/committees in computer science.

These conferences will assist us in providing guidelines and testing the results of this research.

7. Impact on other areas and long-term vision

The impact of liquid publications on publication models and on the scientific community immediately descends from the description above and the objectives we set out forth in the introduction. We have also discussed the need for publishers to drive towards new business models. We next discuss other kinds of impact and the long term vision

7.1. Education

We expect that if LiquidPub is successful, it will impact the world of education – and in particular high-level education, such as undergraduate and graduate courses - in a significant way. University education is characterized by:

– A certain degree of customization in the sense that each professor teaches the class a little bit differently in terms of content and style, sometimes because the goals of the class are a little bit different, the level of the audience is different, or the professor has specific ideas on what is the best way to teach a class.
– Frequent evolution in the subject matter, due to innovations in science.

Education support material in this area has exactly the characteristics of liquid publications. Indeed, it is a prime candidate for liquefaction and collaborative development. Today, education material is characterized mainly by printed textbooks and online slideware. We already mentioned some of the drawbacks of this approach in section 1: too many textbooks, none being a good fit for the professor's preferences in terms of content and style, often outdated and with rather poor support for additional material. Sometimes, they are also quite expensive.

The end goal we have in mind is to go towards a composable and collaborative model for the preparation of education material, where content is structured and multi-faceted (includes text, exercises, slides, along with online communities of classes doing, e.g., joint exercises, or collaborating in software development), collaborative (many professors contribute to the educational material, again sometimes agreeing and providing the same perspective, sometimes disagreeing and perhaps branching the material into separate versions), certainly evolutionary, and also composable as each professor may want to pick and choose elements to form its own educational material.

These are exactly the characteristics of liquid publication, with an increased emphasis on composability aspects as we expect that many professors, even when not contributing to the content of a liquid textbook, will want to compose and create their own educational material out of components (which ideally are "modular" to the possible extent).

Incidentally, note that even when playing the role of "users" (e.g., not adding content to the educational material itself), the mere act of composing material into a "textbook" or even of selecting a specific version is in fact a quite significant contribution: it creates modules that other professors may find interesting and want to reuse, as well as provide input to the evaluation models ("using" content is a strong form of "citation") and hence facilitate ranked search and navigation through the maze of material available.

7.2. Projects

Everyone involved in project preparation, evaluation, and execution knows how much work it is to perform these activities, not only for the R&D activities, but also for the necessary preparation and evaluation overhead. Liquid publication concepts aim at both reducing this overhead and improving the selection and evaluation process. For example:

- *Projects can be prepared and evaluated in an evolving fashion*. This is in fact what to a certain extent is done with EU FET proposals, and it is clear that this is advantageous. Each year a very significant amount of (mostly public) money and a great deal of time is spent in preparing project proposals that in most cases do not get accepted. One could argue that this is good anyways as it makes people connect and think, but ideally we should try to identify proposals to be rejected as soon as possible minimizing both the preparation and the evaluation effort. Evolutionary and possibly collaborative evaluation techniques can support this aspect.
- *Improved evaluation process, in terms of quality of results and effort*: we expect that techniques and services developed for evaluating scientific publications can be of use to evaluate proposals. This may include variations of the current peer review techniques with methods and visual tools that help make sure biases are removed, that reputation of proposers and evaluators (if so desired) is taken into account, and the like.
- *Reduced execution overhead and increased visibility*: preparation of deliverables, a necessary step, is often cited as a major overhead. Deliverables are however needed for interaction among partners and visibility from the evaluating body. Deliverables are just like liquid publications (or, they should be like liquid publications). First, they are certainly collaborative. Second, they are (or should be) evolutionary: it is hard to imagine deliverables set in stone. All artifacts can indeed change and evolve until the very end of the project. This is true for software development as it is for publications, and the same applies to deliverables. Ideally, deliverables should proceed from M1 to the end month of a project, in a continuously evolving fashion but of course with snapshots, just like for publications, where other project

members or project evaluators are invited (perhaps automatically upon “publication”) to look and the work done and results achieved, also based on their access rights. Furthermore, deliverables are in principle multi-faceted as they are conceptually a combination of documents, software, and presentation. Note that preparing versions (and snapshots) is less effort with respect to preparing different deliverables. Sometimes the evolution to a deliverable may be very small and quick to do, but conceptually significant.

7.3. Long term vision

We are going towards a world where knowledge will be entirely available in electronic form. A lot of evidence of this process is already available, e.g., the progressive digitalization of all kinds of knowledge including cultural heritage, spatial and geographic knowledge (e.g., Google Earth), scientific papers (e.g., CiteSeer and Google scholar), common sense and encyclopedic knowledge (e.g., Wikipedia), the availability of more and more indexed knowledge and data available for fast search, and so on.

We have already noticed that software and scientific knowledge are more similar than they may look at a first sight in that they are both creation of the human mind and are “malleable”, since they can be modified and adapted (evolved) over time in several (sometimes wrong) directions and possibly in a collaborative fashion. We also observed that the similarity between knowledge and software is becoming more and more manifest as knowledge is increasingly available in electronic form. Analogously, we can notice that the process from knowledge on papers to knowledge on computers (in electronic form) is similar to the process where software was progressively decoupled from hardware. Thus, initially, software was in the form of machine code and assembler, and there was a clear mapping to the underlying computer architecture. Then, fourth generation languages with suitable compilers and interpreters came about, providing an abstract view of the underlying hardware. Progressively, other abstract, more or less successful, programming languages were invented, e.g., object oriented programming, functional programming, logic programming, agent oriented programming, mobile agents, workflow languages, data base languages, languages for service oriented computing, and so on. We are taking a similar path. These are only the first steps and the impact, consequences and further developments of what we are witnessing will be far more pervasive and radical than we can imagine, similarly to what has happened with software.

This paper provides a substantial contribution in the process highlighted above in that it suggests and provides guidelines for exploiting the malleability property of electronic knowledge. As stated in the introduction, a goal of this paper is to leverage the process of production of scientific knowledge and the huge possibilities provided by ICT as well as by the social web. This paper will provide us with scientific knowledge which is evolutionary, collaborative, multi-faceted and multi-purpose, composite and composable. As it has been customary in the past, the progress made in scientific knowledge will then extend to all the other forms of knowledge thus providing an important contribution to the vision described above.

8. Related work

Our approach goes beyond current R&D directions in collaborative authoring approaches such as Wikipedia (http://www.wikipedia.org/), Swiki (http://wiki.squeak.org/swiki/), and Open Access initiatives, from which we want to harvest experience, success stories and evaluation. Current initiatives do not address the complex interconnections among knowledge objects (SKOs), actors (people) and related processes. For instance, in the various Open Access initiatives, the main focus is on the accessibility and usability of the knowledge object, but nothing is said with respect to its evolution, social reviewing process, and quality assessment. On the contrary, Wikipedia and similar initiatives are focused on the collaborative aspects and on the dynamicity of the knowledge objects, but at present they are lacking in regard to trust, reputation, and credit attribution processes. In this project we learn from these and other related research, as well as from experiences in failures (such as the case of the Electronic Transactions on Artificial Intelligence). We next briefly discuss work related to the areas relevant to this project.

8.1. Writing papers as developing collaborative complex artifacts

As stressed, we think that an interesting parallelism can be drawn between paper writing and software development. In software, the code is developed and then improved. New functionalities are added over time, and the artifact is subject to frequent cycles of improvement and release. In agile methodologies, such as extreme programming [Beck, 1999], the code is also "evaluated" quickly in the process, rather than waiting until the development is complete. We do not change the name of a class each time we make a change to a function. We just release a new version of the class. Once a certain amount of functionality is developed, then the code is released for "testing". Similarly, with scientific papers, an approach that seems sensible is to publish versions of the paper when the work is sufficiently mature and clear so that somebody can read and gain insights from it, and then improve it. More importantly, minor changes (delta contributions) should not result in yet another paper (class) and yet another set of peer reviews as it is always the case today, but in variations or extensions to (versioning of) an existing work.

Of course the development of a large program is a cooperative effort, while researchers compete more than cooperate, so this has to be taken into account. One sometimes does not want to release initial ideas for fear that they are copied, but usually this does not happen and whoever posts a version of a work has a significant lead on others. Besides, early posting, coupled with a secure and community trusted timestamp mechanism, gives people the right to claim that they have been the "first" to a certain discovery. Furthermore, the researchers keep the control on when they want to release the new version of a paper. Needless to say, early releases contribute to science more than late releases.

Open Source development can provide interesting insights in the way people coordinate their efforts and cooperate for the production of an innovative complex artifact. In this respect, there is a growing evidence which show how in the Open Source community both development practices and the management of intellectual property rights (e.g., software licenses) allow interested developers or even (experienced) final users of the software to contribute to a project during its development in several ways, such as: offering extension or modification to the code, highlighting the existence of bugs, discussing technical features or requirements, suggesting patches to correct specific problems, developing a different version of the software (also known as "forking"), etc [Fogel, 2005]. Again, this is challenged by the fact that scientific scholars are usually not very cooperative while open source development is often led by enthusiasts that really use the results of what they develop. On the other side, reputation issues and credit attribution concerns seems to be high also within the Open Source community and it seems that both current practices and licensing specifications can most of the times take into account these "competitive concerns" without hampering cooperative efforts from the developers. Overall, it seems that Open Source development suggests innovative and effective ways to improve and extend a joint artifact and it would be interesting which practices or IPR schemes can be adapted from in our framework for paper production, evaluation and improvement.

Finally, starting from theoretical frameworks such as network theory [Cross and Parker, 2004], collaborative problem solving [Nelson, 1999], open learning environments [Hannafin, Land and Oliver, 1999], and constructivist learning environments [Jonassen, 1999], the project will focus on the role on SKOs as social and virtual platform which might enable scholars to address their efforts towards joint research issues.

8.2. Prior art in publications and the (social) Web

Electronic publishing, digital libraries, electronic proceedings, on-line patents repositories and, more recently, blogs and scientific news streaming are rapidly expanding the amount of available scientific/scholarly digital content. Designing effective platforms for *creating*, *integrating* and *retrieving* such a variety and quantity of scientific knowledge will be crucial in the coming years.

At present, scientific/scholarly "power users" have an increasing number of useful tools in order to support their demands. For instance, they can rely on specialized, vertical, high priced commercial products/portals like (among many others):

- SpringerLink (www.springerlink.de): one of the world's leading interactive databases for high-quality journals, book series, books, reference works and the Online Archives Collection.
- ScienceDirect (http://www.sciencedirect.com/): offers more than a quarter of the world's scientific, medical and technical information online. Over 2,000 peer-reviewed journals and hundreds of book series, handbooks and reference works
- Chemical Abstracts Service (http://www.cas.org/) for chemistry-related articles: this service provides a pathways to published research in the world's journal and patent literature - virtually everything relevant to chemistry plus a wealth of information in the life sciences
- Web of Knowledge from ISI-Thomson (http://scientific.thomson.com): a dynamic, fully integrated research environment base on selected established content sources. The Thomson Scientific editorial process serves as a filter, sorting out the many different content sources and presenting only the most influential resources.

A different breed of scientific/scholarly content providers has also emerged, starting from communities services for particular sectors, such as: CiteSeer (http://citeseer.ist.psu.edu), DBLP (http://dblp.uni-trier.de/), and more recently the wider service of Google Scholar (http://scholar.google.com). They offer free, useful and powerful tools to search, navigate and find scientific content. Furthermore, emerging standard, like the DOI® (Digital Object Identifier, http://www.doi.org/) are appearing and acquiring momentum to provide a system for persistent and actionable identification and interoperable exchange of managed information on digital networks.

Building on these kinds of tools and environments, researchers are beginning to develop applications capable of using these repositories to assist the scientific community above and beyond the pure dissemination of information. In [Rodriguez, 2006] a deconstructed publication model is presented in which the peer-review process is mediated by an Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) peer-review service. This peer-review service uses a social-network algorithm to determine potential reviewers for a submitted manuscript and for weighting the influence of each participating reviewer's evaluations. Such work can be quite beneficial as input to our research on evaluation methods.

Chen et al [Chen, 2007] studied alternative metrics of paper quality and impact. They have applied a variant of the PageRank algorithm [Rodriguez, 2006; Ball, 2006] to assess the relative importance of all publications in the Physical Review family of journals from 1893–2003. PageRank number and the number of citations for each publication are in fact positively correlated. Furthermore, outliers from this linear relation identify other exceptional papers or "gems" that are not easily found with traditional citation/impact factors. The reasoning behind this approach is that the situation in citation networks is not that dissimilar from that in WWW links: scientists commonly discover relevant publications by simply following chains of citation links from other papers. Thus it is reasonable to assume that the popularity or "citability" of papers may be well approximated by the random surfer model that underlies the PageRank algorithm.

One meaningful difference between the WWW and citation networks is that citation links cannot be updated after publication, while WWW hyperlinks keep evolving together with the webpage containing them. Another limitation of citations is that in the current publication models they cannot be used directly for evaluation in the extreme writing model as they assume that a paper is published, visible, and with an "identifier" (published in a journal/conference or at least as a technical report), because before the paper has a high citation count it has to be above the noise level among all documents, and because this is a slow process (you need for many referring papers to be released before you can assess the quality of a paper).

Pre-print repositories, such as e-Prints[7] and academic digital libraries and academic web search services, like CiteSeer.IST[8], Google Scholar[9] and Windows Academic Live[10], have also seen a significant increase in use over the past years across multiple research domains. Furthermore, emerging standard, like the DOI®[11] (Digital Object Identifier) are appearing and acquiring momentum to provide a system for persistent and actionable identification and interoperable exchange of managed information on digital networks. On this basis, researchers are beginning to develop applications capable of using these repositories to assist the scientific community above and beyond the pure dissemination of information. In [Rodriguez, 2006] a deconstructed publication model is presented in which the peer-review process is mediated by an Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) peer-review service. This peer-review service uses a social-network algorithm to determine potential reviewers for a submitted manuscript and for weighting the influence of each participating reviewer's evaluations.

9. Acknowledgements

We thank Vincenzo D'Andrea, Alessandro Rossi, Roberta Cuel, Gloria Origgi, Roberto Casati, Zhang Li, Jordi Sabater, and Carles Sierra for the fruitful discussions on the topic.

---

[7] http://www.eprints.org/

[8] http://citeseer.ist.psu.edu/

[9] http://scholar.google.com/

[10] http://academic.live.com/

[11] http://www.doi.org/

10. References

[Blog, 2005] Academia's Conflicted Reaction To Blogging. 2005 http://acrlblog.org/2005/11/28/academias-conflicted-reaction-to-blogging/

[Beck, 1999] K. Beck. *Extreme Programming Explained: Embrace Change*. Addison-Wesley Professional, Reading. 1999.

[Chen, 2007] P. Chen, H. Xie, S. Maslov, and S. Redner. *Finding scientific gems with Google. Journal of Informetrics*, to appear (2007); (http://physics.bu.edu/~redner/pubs/ps/google.ps).

[Brin, 1998] S. Brin and L. Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Computer Networks and ISDN Systems, 30, 107 (1998).

[Ball, 2006] Philip Ball. *Prestige is factored into journal ratings*, Nature 439, 770–771 (2006).

[EU pub market, 2006] European Commission. Study on the Economic and Technical Evolution of the Scientific Publication Market in Europe. 2006

[Fisher, 1994] Martin Fisher, MD; Stanford B. Friedman, MD; Barbara Strauss. *The Effects of Blinding on Acceptance of Research Papers by Peer Review*. *JAMA.* 1994. http://www.ama-assn.org/public/peer/7_13_94/pv3058x.htm

[Madden, 2006] Madden and DeWitt. *Impact of Double-Blind Reviewing on SIGMOD Publication Rates*. Sigmod Record, Sept 2006.

[Rodriguez, 2006] Marko A. Rodriguez, Johan Bollen, Herbert Van de Sompel. *The convergence of digital libraries and the peer-review process*. Journal of Information Science, Vol. 32, No. 2, 149-159 (2006)

[Rothwell, 2000] P. Rothwell and C. Martyn. *Reproducibility of peer review in clinical neuroscience*. Brain, Vol. 123, No. 9, 1964-1969, Sept. 2000 http://brain.oxfordjournals.org/cgi/content/abstract/123/9/1964

[Santini, 2005] Simone Santini. *We are sorry to inform you…* IEEE Computer, 38(12). Dec 2005.

[Tung, 2006] A. Tung. *Impact of Double Blind Reviewing on SIGMOD Publication: A More Detail Analysis*. 2006. http://www.comp.nus.edu.sg/~atung/