Linköping University

Department of Computer and Information Science

EMTEK

PM 1998:1

1998-02-19

Erik Sandewall

Scientific Communication on the Internet. The ETAI Experience

Related information can be obtained from the following web pages:

Linköping University:	http://www.liu.se/
Dept of CIS:	http://www.ida.liu.se/
EMTEK:	http://www.ida.liu.se/ext/etai/
This article:	http://www.ida.liu.se/ext/etai/1998/01/
The author:	http://www.ida.liu.se/~erisa/

No abstract.

1 Introduction

The patterns of scientific communication are changing rapidly. The importance of electronic preprint archives is growing rapidly, including both centralized archives like the Ginsparg archive in Los Alamos, and distributed system like the NCSTRL (National Computer Science Technical Report Library). The American Universities' Union (AUU) has recently issued a call for 'unbundling' of the scientific publication, where the peer review is done separately from the publication in the sense of 'making public' or 'putting on-line'. The Association for Computing Machinery has an ambitious electronic publishing policy, and has appointed a committee, chaired by Joseph Halpurn, for working out its detailed strategy.

In this paper I shall describe the experience of the ECCAI (European Coordinating Committee for Artificial Intelligence) in developing an electronic publishing medium called the ETAI (Electronic Transactions on Artificial Intelligence). The evolution of the ETAI since it was announced in May, 1997 seems to indicate that the changes to scientific publication can be much more dramatical and thorough-going than what has been projected so far. In fact, we seem to witness the emergence of new forms of scientific *communication*, and not merely of *publication*, with profound implications for how research is performed and how research is organized, and not only for how its results are reported.

2 ETAI's reviewing and publication scheme

At the outset, ETAI was planned as an electronic journal that would differ from conventional journals in several important respects. The fundamental idea was *open reviewing*, where a large part of the reviewing process should be an open discussion, similar to the question period after a conference presentation, or to a public Ph.D. defense: the questions and the answers as well as the identity of the discussants are open to the public, although in our case they would appear in written form and over the net. Only the final part of the process, leading to a decision of acceptance or non-acceptance should be confidential like in the conventional system. It was expected that feedback to the authors, controversies, and requests for clarification and improvements all take place during the open part, the *reviewing*. Then the final, confidential part, called the *refereeing* could be limited to a simple pass/fail decision.

This choice of approach necessitated a number of other choices:

- *Posteriori reviewing*, where an article is first published, then reviewed. Since the article would be open to the public during the reviewing period, there was no possibility of not considering it as *published* in the true sense of the word.

- *Free availability.* The article should be posted on the net at the expense of the authoring institution, and freely available for anyone to download and to read. This means that most of the distribution cost is shared between the author side (for keeping the article on line) and the reader side (cost of computer system, printer, and paper). The cost of the intermediate connection mechanism can be assumed to be small compared to the costs of subscriptions in the contemporary publication system.

- *Author side retains copyright.* There is no reason why the intermediate connection mechanism should require the author side to release copyright.

The principles of free availability and author-side copyright are in line with all current trends and are by no means unique, but in the ETAI they were placed in a broader perspective. The question of publication vs reviewing is more dramatical, and was motivated as follows. Fifty years ago, there were only two ways of obtaining multiple copies of an article: using a typewriter and carbon paper for preparing a small number of copies, or Gutenberg's method using typesetting and a printing press. In that context it was natural to review the submitted manuscript copies of an article, before going to the expense of printing and distribution. At the same time, there was no big danger that results would leak, since the manuscript was confined to so few hands.

Today, the situation is radically different. While an article is still in its reviewing stage, it may be accessible to the whole world by the click of a mouse. It's only natural to admit that it *has been published* then, with the protection of scientific priority that comes with that concept.

Another necessary consequence of ETAI's publication concept was that we had to *reach the primary readers* of each article. By primary readers, we mean the active researchers in the same speciality: those who work on the same research issues as the article addresses, who are prone to using the article's results, and who are also the ones likely to ask questions and to critique the article. The open discussion or 'defense' of the article will naturally take place in this community of primary readers.

3 The publication layer

For the reasons that have now been described, ETAI was initially defined for a two-level publication scheme, consisting of a *publication layer* and a *reviewing layer*. These are different layers in the sense that they are organized in completely different ways.

The purpose of the publication layer is easy to summarize: to make the article easily available, independently of space and time. The basic service is of course one of making each submitted article available on the net, or more precisely, of making articles available on the net so that they can be submitted. The reviewing layer *will only receive articles that have already been published* using the publication layer. This principle is diametrically opposite to the policy of conventional journals who maintain that they only accept "previously unpublished work".

Besides making articles freely available on the net, the publication layer must also guarantee that articles can not be changed. Unlike in ordinary file systems, it must in particular guarantee that the author himself or herself is not able to change the article after the stated publication date. What is more, the publication layer must be organized in such a way that *everyone is convinced* that noone is able to tamper with the article after it was published. This calls for both technical and administrative safeguards.

The reviewing layer contains both the open reviewing and the subsequent, confidential refereeing, as outlined above. In fact, there are three points where the two layers interact:

1. Before the article is first submitted to the ETAI, it must be published using the publication layer

2. After the open review period and before the article is sent to the referees, the authors may wish to revise it based on the feedback that they have received. The revised paper must again be published using the publication layer.

3. In order to compare well with traditional journals, accepted articles should be integrated into well-defined journal issues with a traditional, journal-style "look and feel".

We refer to these as *first publication*, *revised publication*, and *journal republication*, respectively. Note that technically speaking, the ETAI issues containing those articles that have been accepted by the referees do in fact *republish* articles, since they have been individually published before.

First publication can not consist merely of the author putting the article on her or his own web page, since then the integrity of the article over time could hardly be guaranteed. For similar reasons, we could not say that any article that has been published as a departmental technical report qualifies as published for the purpose of ETAI reviewing. On the other hand, it was not possible (especially not on the European scene) to assume a centralized repository for all articles. We therefore decided to define a concept of *First Publication Archives* or FPA's, that is, organizations that are able to accept vthe required responsibilities.

The publication layer is thus organized as a family of First Publication Archives. We would expect every major university to run its own FPA, just like it has a university library. In fact, Linköping University Electronic Press, which was the first FPA, is organized in close cooperation between the university library and the computing center of its university. The ETAI itself does not participate in the FPA's: it specifies a number of *requirements* for authorization of FPA's, but that is all. The basic requirements are:

- The FPA must be organized to put research articles on-line on the Internet and to protect them against tampering, using a credible protection scheme.
- The FPA must be committed to keeping articles on-line for a sufficiently long period of time. 25 years from the date of publication is a recommended period. It must be based on a sufficiently strong decision, for example by the Board of its university, for the future retention of the articles.
- The FPA must take responsibility for the graphical and linguistic quality of published articles. The ETAI will not intervene for the language quality of individual articles, but it will object to the responsible for the FPA if there should be problems in this respect.
- The FPA must have correct routines for publication agreements with authors, so that the right of the FPA to keep articles on-line for the required period of time is not jeopardized. Usually, this means to let the authors sign a contract whereby the FPA obtains a non-exclusive right to publish the article.
- The FPA must do whatever is required published articles to be considered as published in the legal sense in its country.

If the responsibility for publication is taken by each university, it means in fact a return to a classical system where each academy or university published its own works, but then 'publish' was the same as 'print'. In the electronic world it may be convenient to resume that practice.

FPA publication works fine for first publication and revised publication as defined above. But what about journal republication? We chose the following solution: The ETAI prescribes a standard style (in Latex) for submitted articles. This style is chosen so that the journal version of the article can be obtained by removing a few pages at the beginning and replacing them by a few others. These operations can be easily done on the level of the postscript version, and are in fact almost entirely automated. Therefore, the reviewing layer is able to create the journal republication version of the article by an easily managed transformation on the FPA published version. The reader is invited to visit

http://www.ida.liu.se/ext/etai/...

in order to see the appearance of the resulting issues of the ETAI.

4 The reviewing layer

The publication layer and the reviewing layer are partitioned along opposite axes: each university may have its own FPA, the present ETAI can review articles from any of its approved FPA's, and when there are several journals similar to ETAI, each FPA may serve any of them. In addition, however, the ETAI is organized as a number of *research areas*, such as "planning and scheduling", "reasoning about action and change", "intelligent human-computer interaction", and so on. Articles are submitted to research areas, discussed within them, and refereed within them. This is the only way an article can be accepted by the ETAI. If it does not fit into any of the existing areas, then it can't be received.

Each ETAI area has an area editor who organizes the work, and an area editorial committee which is in charge of the refereeing. However, the area must also have a "constituency", that is, a list of researchers who identify with its choice of research problems and research methods. The active participation in the open reviewing must come from the constituency.

First of all, the discussion is *posted* and *preserved*. It would not be sufficient to only run it as a conventional mailgroup. This is for several reasons:

- If the discussion is preserved, then it is possible to look back at it later on and to see what arguments were put forward, what critique did the author receive, and how well did he or she answer to it.
- If the discussion is preserved, then a well-written contribution to it will also give credit to the discussant. By contrast, writing an insightful review of a journal article in the confidential reviewing system does not count for anything.
- Making and keeping the discussion public is a protection against nonserious contributions to the discussion.

The most obvious solution is then to set up a structure of web pages containing the successive contributions to the discussion; one discussion page for each submitted article. (The alternative solution of using a teleconferencing system such as First Class (R) was not considered because we felt that we would then not be able to attract sufficiently many participants). Initially, we arranged to set up a mailing list for the entire constituency in each of the ETAI areas; submission of articles were reported by direct E-mail to the constituency, but then everyone had to visit the web page in order to see how the discussion had developed since his or her last visit.

In our experience, it was very difficult to get off the ground with that approach: we ran into a vicious circle of few discussion contributions, few visits to the web page, and accordingly a low motivation to contribute to the discussion. The solution that evolved fairly soon was to create a *Newsletter*, which can be described as a somewhat edited variant of a mailgroup. It works as follows: contributions are sent to the area editor, presently using a plain E-mail message. One or more contributions are collected into a newsletter, which is sent out to the constituency. As far as possible, newsletters are sent out the same day as the contributions come in, but never more than once a day. (This provides a protection against overly intense interactions). All newsletters are accumulated without change to a webpage structure, so a subscriber never needs to save arriving newsletters - they are retrievable from the webpage of the ETAI area in question.

Besides retaining the plain newsletter in .txt format, there are also some other operations. The text is reformatted into .html for more convenient reading. Also, the constituent messages are organized into the pages for the respective articles that they are commenting about, and in a structure where one can easily follow what is the answer to what question. Finally, the article discussions are also reformatted into latex and used for generating a postscript file that reads conveniently on paper.

Now we have described everything that is required in the reviewing layer. The area editor has to organize a web page for the area, and he or she must have a mailing list for sending out newsletters. When an article is submitted to the area, an article interaction page must be set up, and the constituency must be informed in a newsletter issue. When discussion contributions come in, they must both be circulated using the newsletter, and accumulated to the interaction page. After three months of discussion, the area editor asks the authors whether they wish to revise the article, awaits the URL of the revised article, then nominates three referees who decide whether to pass or fail the article, posts the verdict, and informs the editor-in-chief. The latter maintains a web page for the ETAI as a whole, where accepted articles are posted in journal look-and-feel format.

In principle this is all. In practice, there are a few more aspects to it that must be mentioned:

- What happens if there are no contributions, or insufficiently penetrating contributions to the discussion? Our present policy is as follows. The discussion period is three months. If four-six weeks have passed without much discussion, the area editor invites a few colleagues to read the article and to ask questions.

  In a longer perspective, one may consider extending this scheme so that there is *always* one or two "opponents" for each submitted article, similar to how Ph.D. theses are defended in some countries, or even similar (in some ways) to a court proceedings. It is important, in these cases, that designated opponents shall never be asked to be referees for the same article.

- In extreme cases where the discussion ever takes off, one can always go back to the conventional system, that is, to appoint confidential reviewers whose tasks is not only make a pass/fail decision, but also to give feedback to the author and propose changes to the article.

- What happens if the author thinks of some improvements to the results while the article is in the review discussion, either because of feedback from discussants, or on her own? As a matter of policy, we do not allow the author to improve the *results* in the revised version of the article, only to improve the *presentation* of the results. This is in order for the publication date of first publication to be relevant for priority purposes for the revised version as well. Instead, the author is invited to write up the additional results as a separate research note. Unlike in a paper-based journal, the electronic medium makes it possible to link the amending note to the base article for easy cross-reference.

Being an area editor in this sense need not be very much work, in particular if one has adequate software support for the various operations involved. (A forthcoming article will describe the software support system that has been developed for this purpose). It does however require responsiveness: the arrangement is only credible for the constituency if incoming messages do get forwarded and administered the same day, with no or minor exceptions. Anyway it is certainly possible for this task to be performed by the research community itself. The only function that is performed by journal publishers today and that the ETAI system is not set up to cope with, is to handle articles from authors that do not have access to contemporary computer and Internet technology. In our context, some special solution has to be found for them.

The following is one example of how the ETAI reviewing layer has changed the reviewing practice. (... example ...)

5 The colloquium layer

The two-layer system consisting of a publication layer and a reviewing layer corresponds fairly well to our initial design hypothesis at the time the ETAI was started. Only the emergence of the newsletter with direct E-mail of daily issues (when needed) was an extension of the original concept. However, once this basic structure was in place, we began to see other activities that could also be run over the same structure. To repeat, the review layer provided an ETAI area with:

- A constituency, represented by an E-mail mailing list;
- A newsletter that the constituency learned to recognize and to appreciate, and which therefore became a channel for reaching the constituency;

- A webpage structure, containing information of interest for the constituency, which they also started to use;
- Among many members of the constituency, an interest and an ability to use the newsletter and the webpage structure for appearing, engaging in debate, and making their voice heard;
- Software support for administrating a relatively large flow of short texts ("messages") that are to be sorted into different structures (by date, topic, etc), and provided with headings and formatted in different ways (text, html, latex).

The following are some of the activities that have developed on that basis:

- *Panel discussions* on topics of current interest in the community, analogous to panels at ordinary conferences.
- *Bibliography support.* We maintain a database of articles of current interest, including a slot for storing the URL for where the full text of the article can be accessed. Links to this database can be included in all submitted messages, which means that the discussion contributions in HTML format get to contain "hot" links to other articles.
- *Topic oriented bibliographies.* The bibliography support is also a tool for setting up and maintaining bibliographies of articles on selected topics. Unlike the classical bibliographies that had to be printed in order to be distributed, ours can be updated continously in response to both new arrived articles, and to reminders about articles that had inadvertently been missed at first.
- *Follow-up of discussions at workshops.* Once the software for administrating discussions about submitted articles was in place, we could use it as well for presenting resumes of the discussion periods at a regular workshop. In those cases where a regular book contained a summary of the discussion after each article, it could take several years to complete the book. In our case, the discussion resume came on-line only three weeks after the end of the workshop. (We had to allow some time for the workshop participants to check and correct the resume). Also, of course, authors are able to correct and extend what they may have said at the workshop, and the discussion can continue on-line.
- *Calendarium information.* It is natural to maintain a list of forthcoming conferences and workshops, and to integrate that list seamlessly with the list of accepted articles (when it becomes available), links to the full text of the articles (when available), and later on to add on-line discussions about the articles.

- *Publishing software and other nonstandard entities.* The newsletter and the web page contain descriptions of current software, links to them so that they can be run across the net, and links to full articles reporting on their salient properties.
- *Monthly news journal as a regular publication.* Although the daily newsletter is perfect for the active members of the constituency, there are others who wish to get an overview of what is going on, but not as often. For this reason, newsletter contributions are also accumulated into monthly *news journals*, which at first only appeared in html form. Soon, however, it became apparent that a latex/postscript version of the News Journal would also be appreciated, and it was straightforward to automate its production. The News Journal for one of the ETAI areas has been turned into a regular periodical.

These additional uses of the mechanisms that were created by and for the reviewing layer, have now become so varied that they are best viewed as a third layer of communication activity, the *colloquium layer*. We talk of an electronic colloquium, keeping in mind that the traditional academic colloquium is a forum where researchers meet in order to hear the latest news in their area of interest, to listen to results, to ask for advise about where to find what information, and so on. The electronic colloquium serves many of the same roles, but in a way that defies distance and that is also preserved over time.

This is already enough for talking about a new form of communication between researchers. Distance education is already a fact on the undergraduate studies level; the ETAI experience suggests that a similar development may be within reach for graduate study and for research. With the ETAI electronic colloquia, grad students anywhere can follow the discussions between senior researchers in their field of interest, and they can also participate in those discussions.

However, the scenario offers additional possibilities, which we have barely started to tap into. One such possibility concerns the size of written contributions; this is a key to new ways of publishing systems research, in particular. Until now, articles in computer science have tended to be fairly long: long in comparison with some other fields, such as physics, and long in the sense that reading an article often requires a considerable investment of time. At the same time, authors often wanted to write even longer articles; journals editors are known to require authors to cut down the size of the articles.

In the ETAI scheme, we can extend the size range both upwards and downwards. Long articles are not a problem in the sense that it's the author institution that pays the cost anyway; it is possible to cover the entire spectrum from article-length to book-length publications. Conversely, short notes have never been widely used in our literature, presumably for a number of reasons:

- Articles are supposed to be self-contained. (This requirement is valid for printed articles, but much less necessary in the on-line environment where it's easy to link to other papers providing background, definitions, etc)
- Very short articles are difficult to administrate. You don't really want to make a tech report of a two-page paper. Bothering to send it to referees is also a nuisance.

In the ETAI system we have two sizes below article size. There are *messages* into discussion sessions; messages are edited into discussion protocols, and a message is typically between four and fourty lines. Then, there are *newsletter notes* that can be two pages and up. Newsletter notes have individual existence, with an author and a title; they obtain published status by appearing in a monthly News Journal. Thus, just as full articles are first published by FPA's and as separate publications, newsletter notes can be bundled and appear in a News Journal, in order to reduce the handling costs.

6 New patterns for communicating research results

In its present form, the ETAI provides a system for handling *conventional research articles and information about those articles*: questions to the author, commentary by the author, and so forth. It defines a communication channel between reseachers in an area. This channel was initially set up for the purpose of the open reviewing process but, once in place, it is also being used for many other communication purposes.

This experiment has only existed for less than a year, but even during this short time it has developed some features and facilities that were not anticipated at the outset. Therefore it is likely that additional developments will emerge during the coming years, and for the same reason it is very difficult to predict the direction of those developments. However, we shall anyway venture to speculate on some of the obvious possibilities.

Developments in this area will only happen in response to a perceived need: a perceived failure of the existing system. Consider, therefore, the question of how systems oriented research publishes its results. For a long time, there have been complaints about how this works. Conference participants request more systems papers and fewer theory papers. Conference organizers answer that they really try to get systems papers, but few are submitted, and many of them have insufficient quality. Authors of these papers complain that their work is only weighed for its theoretical contents, and that it is therefore too harshly treated in the reviewing process.

It may be appropriate, therefore to question whether the conventional-size article is a natural way of communicating the results of a systems oriented research project. Is it reasonable to squeeze the description of a large project into a short article? Is it reasonable to require the system description to also contain an analysis of, and comparison with alternative approaches? And of course: is it reasonable to have articles that do not allow the reader to test-run the system being described?

For the sake of the argument, consider the following alternative scenario for publishing the results of systems-oriented reseach, in an ETAI-style constituency where several of the participating groups build comparable systems. First, each group develops a publication structure that describes their system, and maintains it on-line. Thus, they do not necessarily generate a sequence of papers with more or less overlapping contents; instead, the group makes a serious job of *maintaining* an on-line documentation structure consisting of a number of parts that can be exchanged asynchronously, and which at each point in time is a fair description of their system in terms of motivation, principles, and architecture as well as details.

Then, the on-line discussion mechanism of the ETAI system is used for comparing notes about these systems. Their performance for specific examples or specific classes of problems is compared in debate mode. Also, a few researchers from different groups may co-operate in making a detailed comparison of their respective systems from some specific point of view, and publishing the result, possibly sparking additional discussion. News Journal notes may be used for contributions that are too long or too substantial to be edited into a plain discussion protocol. Throughout the discussion, the primary publications of the respective groups are used as framework and reference for the discussion. However, those readers who are not themselves actively involved in exactly the same kind of research are likely to profit much more from reading the discussions than from trying to read the full articles from each project. - Let us use the term *comparative publication* for this way of communicating research results.

It may well be argued that comparative publication is more adequate for systems oriented research than the traditional paper format is. But if we think so, there are again a number of issues that must be addressed. In particular, how do we assign merit for research that has been done and communicated in this way? The traditional merit scheme is based on identifying influential articles, often also *counting* the number of articles, taking into account the prestige of the journal where the article was published, and counting the number of citations to the article from other, later articles. Most of these criteria become inadequate and misleading in the comparative publication style.

So how will reseachers that subscribe to comparative publication obtain their promotion? First of all, the members of a promotion committee are likely to find the discussion protocols for the debates between the research groups more informative, and easier to understand and evaluate, than the orginal articles of each of the groups.

An ambitious evaluation of research quality should maybe look less at the praise made by peers in confidential letters of recommendation, and more at how the project was able to present itself and to assert itself in open, technical discussions.

In addition, the open "colloquium" environment might organize special events where different systems compete with each other. An obvious kind of competition event is the *championship*, where participants are asked to use their systems for solving some concrete tasks. Another possibility is to have *duels*, where the group that developed system A challenges the group of system B with a concrete task or set of tasks, proposing that a closed competition be set up with only systems A and B, and for solving the problem(s) proposed by the A system team.

It will be necessary to add a few recommendations to the rules for these competitions. For example, in the interest of fairness one may rule that for every task proposed by the challenging team in a duel situation, the challenged team shall be able to responed with another task of their choice.

The use of competitions for evaluating research results differs in important ways from the classical criteria, where research is given credit for the number and the popularity of its original results, that is, results where the author or the authoring team are able to claim priority. In the competitive scheme, priority is not directly relevant. If team A has developed a powerful new concept and incorporated it into their system, and team B picked up the idea and were able to introduce it into their system in time for the next competition, then the historical priority of team A does not matter. On the other hand, team A had a head start, and if they are able to use it well they can proceed more rapidly to the next innovation. This arrangement may seem very unsatisfactory if you take priority of results as the ultimate touchstone of scientific success. On the other hand, if you feel that priority is an ill-defined concept since new ideas tend to come up concurrently in several laboratories, then you may argue that competitive evaluation makes more sense than publication counts and priority rules.

Competitive evaluation also has its problems. In its extreme form, it might lead to excessive secrecy, since a research project will not only strive to keep its results secret until they are properly published; they may have an interest in not telling at all about some of their developments, in the hope of being able to use them during forthcoming competitions during an extended period of time. This risk may be reduced by a combined system, where competition and conventional publication are used together. Then, a research project that hides some of their results may win in competitions, but they may also lose in publications since someone else will publish the method sooner or later. Yet another idea may be to impose explainability requirements on the competing teams: the merit value of a victory is qualified by the extent to which the winning team is able to explain why and how they won. These issues will no doubt have to be worked out in terms of the special characteristics of the discipline and the type of system that is concerned.

7 Conclusion

The ETAI scheme of publication has developed over less than a year. Similar approaches have reportedly been tried in some other disciplines, although not yet in computer science, to the best of our knowledge. This is actually a bit surprising, since the operation of ETAI depends quite heavily on software support, which means that it ought to be particularly easy to realize it in a computer science environment. This applies in particular for the third, colloquium layer, but also for the first two layers. Anyway, now that a concrete example exists, we hope and believe that similar systems will be set up for many other specialized areas as well.