

Supporting the Writing of Reports in a Hierarchical Organization

Andreas Girgensohn

FX Palo Alto Laboratory
3400 Hillview Avenue
Palo Alto, CA 94304
andreasg@pal.xerox.com

ABSTRACT

In many hierarchical companies, reports from several independent groups must be merged to form a single, company-wide report. This paper describes a process and system for creating and structuring such reports and for propagating contributions up the organization. The system has been in regular use, in-house, by about 30 users for over a year to create monthly status reports. Our experiences indicate that it is possible to change a monthly reporting practice so that the system is easy to use, improves the quality of the written report, fosters collaboration across projects and creates a corporate memory for the company. These results were achieved as a consequence of our design effort to directly support the hierarchical and collaborative process of creating and assembling the report within the organization. User feedback has led to many improvements in the usability and functionality of the system. Further enhancements using information retrieval and text summarization techniques are in progress.

Keywords

Hierarchical organizations, collaborative writing, corporate memory, report generation, user feedback, World Wide Web.

1 INTRODUCTION

In many organizations, workers are required to report periodically on their activities. Such reports are often assembled in a hierarchical fashion, i.e., reports of suborganizations become building blocks for the report of the parent organization. Complicating the problem is the fact that organizational hierarchies are not represented as trees because individuals or suborganizations can contribute to more than one parent organization. This means that reports of individuals or suborganizations have to be sent to several parent organizations so that relevant portions can be extracted. The effort involved in this task often prevents the creation of reports that describe all the important activities.

There are several goals for the report process and system described in this paper: (1) to facilitate the report-writing process by making it easier for the writers, (2) to improve

the quality of the reports by making them more informative and focussed, (3) to increase the awareness among employees about activities outside their immediate environment, (4) to foster collaboration across groups, and (5) to keep a corporate memory of activities.

This paper starts by describing the history and experiences with the process and system. This is followed by a section describing how the system supports the monthly report writing process. The following sections discuss user reactions and related work. The paper concludes with a discussion of future enhancements.

2 BACKGROUND

2.1 Context

At FX Palo Alto Laboratory, ten projects are organized into three research groups. On average, three researchers work on a project. Researchers have “home” groups which are managed by group leaders. In addition to the three research groups, there are two groups for computing support and administration. Before the report system was introduced, each group emailed its monthly activities directly to the corporate secretary. The formats of the reports varied from group to group so that it was difficult to get a coherent picture of the monthly progress of the laboratory. It became clear that a system could be used to facilitate the writing of the reports, and, more importantly, to improve their quality. The new reporting system was supposed to achieve several goals:

- A simple way for writing and combining monthly reports.
- A uniform structure for all group reports with more detail on project status.
- A process that kept management and employees informed about project status and progress.
- An easily accessible archive of past reports to be used as part of a corporate memory.
- A comprehensible report that made it easy to get an overview of ongoing activities.

2.2 History

Throughout the design and deployment of the new system, user feedback has shaped the system and has led to several significant changes. The experiences in that process have been similar to the ones reported, for example, by Okamura *et al.* [14]. Many small adjustments resulted from direct observations of how the users wrote their reports rather than from user feedback alone. Often, the user’s intent could be

inferred from those observations and the system could be adapted to make it easier.

As a first step in designing the new system, a small design team restructured a past monthly report to determine a better format. This sample report consisted of six sections (e.g., “research activities”). Inside each section, paragraphs had run-in headings identifying projects. A Web-based approach was chosen for the report system because an intranet was already widely used. Many users used Hypertext Markup Language (HTML) editors on a regular basis so that it seemed efficient to provide report templates that the users could complete in such an editor. The report made use of HTML formatting features such as color, lists, and tables to distinguish it from the plain format of previous reports. After several rounds of discussion with management, the new system was introduced in May 1997 and has been in regular use by about 30 users.

2.3 Use Experiences

Since the introduction of the system, user feedback and observations of user behavior have led to several design changes in the system. This section discusses some of those changes.

2.3.1 Authoring of Reports

Initially, the system consisted of an HTML template for individual reports and a script for combining individual reports into group reports and group reports into a laboratory report. It became clear right away that most users found the process too cumbersome and error-prone because they had to copy the template file into a directory for the correct month and to rename the copied file. In response, a form was designed in which users could enter their monthly activities that followed the structure of the report. From that form, a corresponding Web page was generated that users could modify using an HTML editor if they wanted to add additional formatting (e.g., links, tables, or images).

In addition to becoming familiar with the new system, users also had to adapt to the process that automatically assembled the reports. Snippets in individual reports had to incorporate enough context so that they could be included in group and lab reports. Users had to remember to write in the third person and to include their full names. They also had to remember that their audience was not just their immediate manager. Therefore, background information had to be included. Similar issues have been reported for collaborative authoring (e.g., [18]).

Early on, users requested support for just marking portions of a report for inclusion in the parent report. This feature enabled the authors to keep additional detail in their reports without having it flow up to the parent report. Only the marked portions were extracted and the remainder of the report was ignored. The marked portions included in the parent report contained cross references back to the source report as hypertext links so that additional details were available on demand. This enabled the reader to see a coherent picture at one level of the hierarchy (e.g., the whole laboratory) and to see more detail for the portions that are of particular interest.

2.3.2 Distribution of Individual Contributions

Group leaders discovered potential problems with the new system because it only looked at individual and group reports and did not consider projects as an organizational unit. At the time, it was decided that the group leaders should address this issue by editing the generated report to combine contributions to the same project. In the long run, this provided too little information on the project status, and the process was changed as described below. The implications of the changes were discussed at length in email messages and in meetings. The following text includes excerpts of several email messages from users discussing the proposed changes. Only some of the messages are included. Some issues were never discussed via email so that no written record is available.

Subject: Re: Form for new monthly reports

Date: Tue, 3 Jun 1997 07:12:32 -0700

I also have some procedural questions to raise. In the past my major contribution has involved collating the reports I receive for the three xxx projects, editing them, and inserting my own contributions. The problem seems to be that we have always organized our preparation of material by PROJECT, rather than by the INDIVIDUALS. I am afraid that changing this approach may make our reports harder to read, since it will be less obvious who is contributing to what. How do you propose that we deal with this?

It also became apparent that a strictly hierarchical approach was insufficient. While all researchers had “home” groups, some of them worked on several projects belonging to different groups. Group leaders found it difficult to deal with contributions not belonging to their groups. The mechanism for marking portions of a report was extended to identify the group the contribution belonged to so that the marked portion could be included automatically in the report for that group.

Subject: Re: New script for generating reports ready

Date: Thu, 5 Jun 1997 07:03:29 -0700

I just finished editing the xxx group report. We definitely need a better model of how groups break down into projects. The biggest problem I encountered is that yyy did some work for zzz that has nothing to do with xxx, so I excised it from the group report. However, there is obviously no mechanism by which it is going to get merged with zzz’s material. So this means that, under the current system, the group reports are not going to be particularly accurate, while the individual reports do not reflect the review of the group leaders.

Subject: Re: Monthly reports due tomorrow

Date: Fri, 05 Sep 1997 14:53:06 -0700

I’ve a little nit on the form generator:

> (select a group only if the activity does not belong in your “home” group)

There is a case where an activity doesn't belong to any single group--say spanning two other groups outside of your own group--but under this rule leaving it blank would default to the home group??

2.3.3 Restructuring the Hierarchy and the Report Sections

After seven months of use, a major change was requested by the company president. While the original reports provided several sections for different types of events (e.g., "Research Activities" or "Collaborations"), they did not always capture the essence of the events in a comprehensive manner. Also, reports were structured at the group level so that individual projects could not be tracked. The new version introduced reports at the project level so that there were four different levels (individual, project, group, lab) and restructured the sections to produce more concise and informative reports. The categories in the new version were not completely intuitive and needed some clarification as the following messages show.

Subject: Re: Group report past due
Date: Tue, 13 Jan 1998 10:15:48 -0800

xxx informed me that he just wants 4-8 highlights as a group report for December. I am currently collecting this material from my group (most of whom, like myself, are just returned). I figured I would pass these highlights to xxx by electronic mail; but I can create a version for the Web site as well, just so we have some consistency in the way in which we maintain our records.

Subject: Re: Monthly reports due soon
Date: Thu, 05 Feb 1998 10:14:39 -0800

Some of these categories overlap: Do papers submitted to conferences go under Products or Conferences or both? And when you go to a conference, does that in addition to Conference count as a Presentation and Visit, too?

It was also determined that marking portions for inclusion in the parent report did not improve the ease of writing and the quality of the final report. Instead, it led to additional confusion. The mechanism for marking portions of text for inclusion was replaced by headings indicating the project the contribution belonged to. Such headings were easier to use than the expanded markers mentioned above. After the change, project and group leaders had some extra work in removing unwanted information but they also had a richer selection to choose from. Both the longer experience with the system and the new section structure let individuals write better reports so that the overall report quality has been improved. The better report quality is indicated by reports that are more appropriate for the intended readership at our parent company. Reports now use better English, concentrate on important events, and provide a context for the described events.

Support for collaboration among projects on a single issue required another change. Such reports of all participating

projects should contain the events but there should be an indication that the event is described elsewhere as well. It was decided that such situations are too infrequent to warrant a more complex form interface. Instead, an event was associated with more than one project by listing more than one project in the corresponding heading in the HTML document. Such an event was then included in all referenced project reports and appropriate cross-references were generated.

Subject: monthly report form
Date: Mon, 2 Feb 98 16:45:29 PST

Under the new report form it is impossible to list an activity as being part of two projects. I spoke with Andreas and he said that this change was a result of the way the reports will be filtered as they go up the hierarchy. As part of our first lab goal is integrating technology across projects, it seems unfortunate that it is difficult to report such activities on the new form.

Some users found the incremental changes and announcements confusing. While it was beneficial to have the input of all users, it might have been better to introduce new features to a subset of the user community first to avoid too much confusion.

Subject: Re: More on monthly reports
Date: Mon, 2 Feb 98 11:58:18 -0800

This sequence of bits of explanation is pretty confusing. I hope the website gives the whole story from the top down.

2.3.4 Providing Better Access to the History

The most recent change addressed the needs of users who use the report system to maintain their personal history. A "Next Month" button in the form enabled them to enter information for a month before completing the report for the previous month. Access to the previous month and the combined history of all reports for the current year made it easier for them to refer back to events in the past.

Subject: More nits for the activities log
Date: Fri, 10 Jul 1998 14:15:08 -0700

When I follow a link to a month which I haven't filled out yet, it should bring up the form. And there should be a link from the report to the form, to next/previous month.

We continue to solicit user feedback, and to observe problems with the system as it evolves with the user community and the needs of the use.

3 THE REPORTING SYSTEM

The experiences described in the previous section have shaped the design of the reporting system. This section provides an overview of the system functionality and a description of the technical details.

Figure 1 shows the flow of reports in the organizational hierarchy. As suggested by the user feedback described in

Editing generated HTML document

Combining reports and generating new document

Editing generated HTML document

Combining reports and generating new document

Editing generated HTML document (optional)

Generating HTML from the form data

Filling out the report form

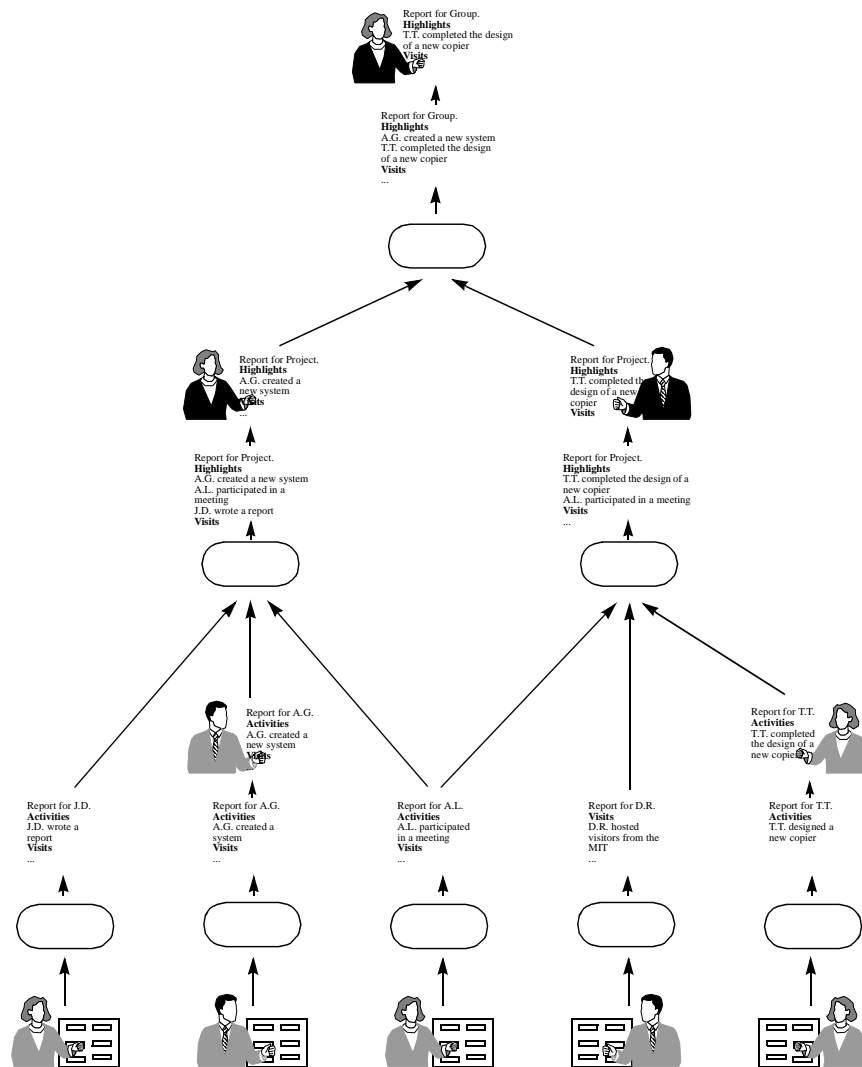


Figure 1: Report Flow in the Organizational Hierarchy

the previous section, contributions from a report can flow up to several different parent organizations. The automatic combination of reports helps managers concentrate on editing reports to tell a better story rather than on menial tasks such as copying and pasting text blocks. Contributors can learn from the example of others because all reports are easily accessible. In addition, managers interact with contributors directly to ask for additional detail or to suggest changes.

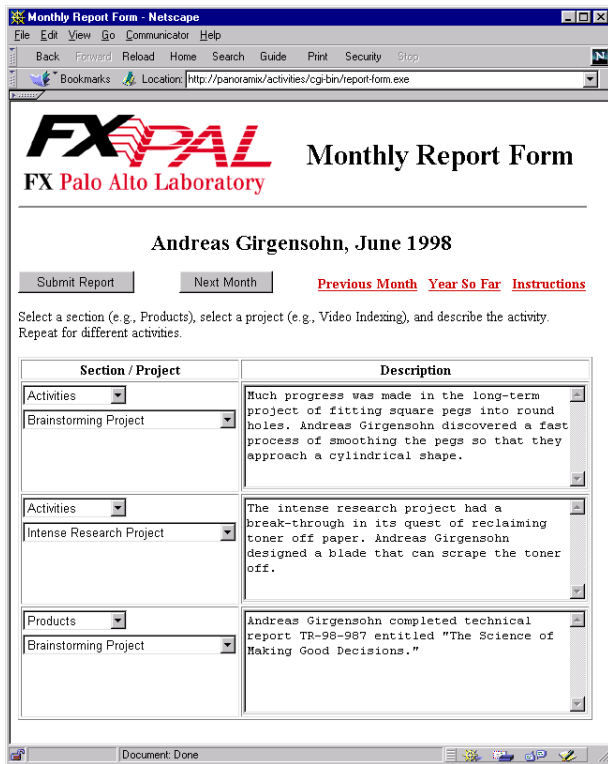
The system described here is by no means restricted to our organization. Because the organizational hierarchy and the report sections are both described in tables, the system can be easily adapted to different organizations and their needs. Any organization that produces reports in a hierarchical fashion could make use of the system.

The report system is integrated into the corporate intranet by including hyperlinks from portions of a report to additional information on the intranet and by linking back from the intranet documents to the reports. The collection of reports


represents a part of a corporate memory. Our future plans are to further develop the corporate memory concept.

3.1 Types of Reports

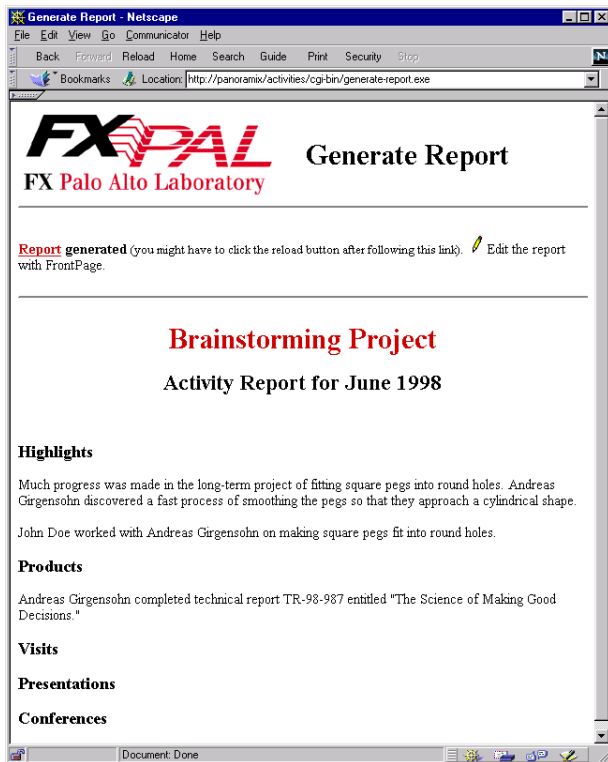
Reports are divided into different categories to structure the reported activities. Such structure helps the contributors focus on the important accomplishments. For each category, all the corresponding contributions are collected from the input reports and combined into a single section. This process keeps related accomplishments together. In the combined report, there can be subheadings for the contributing entities within a category or all the events can be combined as a flat list. Figures 2c and 2d show examples for both styles (actual reports cannot be shown because of the proprietary nature of the information). The categories are defined in a table so that they can be adapted easily for different needs of an organization. Categories can have different names at different levels of the hierarchy (e.g., “Activities” maps to “Highlights”) or several categories can be combined into a single category. Defining appropriate categories and explaining their intended uses are very important for producing good reports. As described in Section 2, the ini-



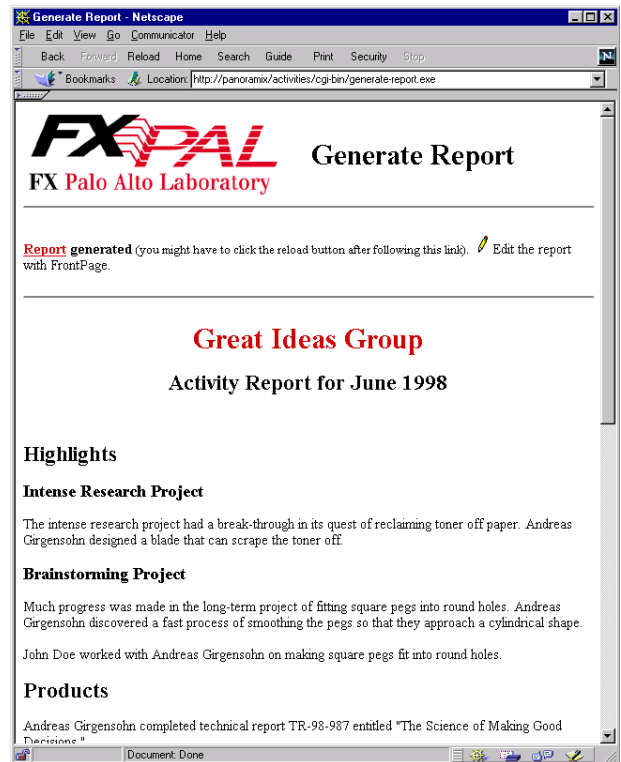
(a) Form for entering reports



(b) Report generated from the form input



(c) Project report combining individual reports



(d) Group report combining project reports

Figure 2: Web Interface for Report System

tial report structure was changed after observing its use for several months. The new structure required several clarifications for the purpose of different sections.

After a report has been generated, it is intended to be used as a starting point for an edited version of that report that combines duplicate items, excludes portions only of limited interest at this level in the organizational hierarchy, and summarizes parts with much detail. To support the user in that task, the generated report contains a link that starts an HTML editor and saves the changes back to the report. Users can concentrate on the authoring task and are not distracted by clerical aspects.

Organizations and individuals are grouped in levels with all individuals at the first level. In collecting activities for reports, all the reports at the previous level are checked whether they contribute anything to the organization in question. It is also possible for contributions to skip a level if there are gaps in the hierarchy. Individuals can contribute to several projects so that project names are used as headings in the individual reports to direct the described contributions to the appropriate project report. For all the organizational units above the individual level, the hierarchy is assumed to be a tree. Experiences at our company showed that inserting an additional level (project reports) in the hierarchy improved the reports because the material was better grouped and more focussed.

For managerial purposes, it is also possible to combine a report from all the individual reports contributing directly or indirectly to an organization, and thus skip levels in the hierarchy. Such a report provides a manager of an organization with a quick overview of all the individual contributions.

To provide widespread and cross-platform access to the reporting system, Web browsers and HTML are used for the user interface and the internal report format. Data entered into the report form is automatically converted into an HTML document with sections representing the different categories of activities and subsections for the different organizations each activity is associated with. This document can be edited further using an editor if additional formatting is desired.

To support individual users in writing their reports, the annual history is available from the report form (see Figure 3). This history combines on demand past reports by section and inserts subsections for each month. The product section is of particular interest because it provides the readers with an overview of all the products such as technical reports they have already produced for the year. This particular feature was introduced in response to requests from users who wanted to refer back to past events while writing their reports. Users also wanted to see all their accomplishments at the end of the year.

Annual reports are also interesting at higher levels in the organizational hierarchy. Annual project reports show all the accomplishments of a project. The annual history for the whole laboratory starts with the history of all executive

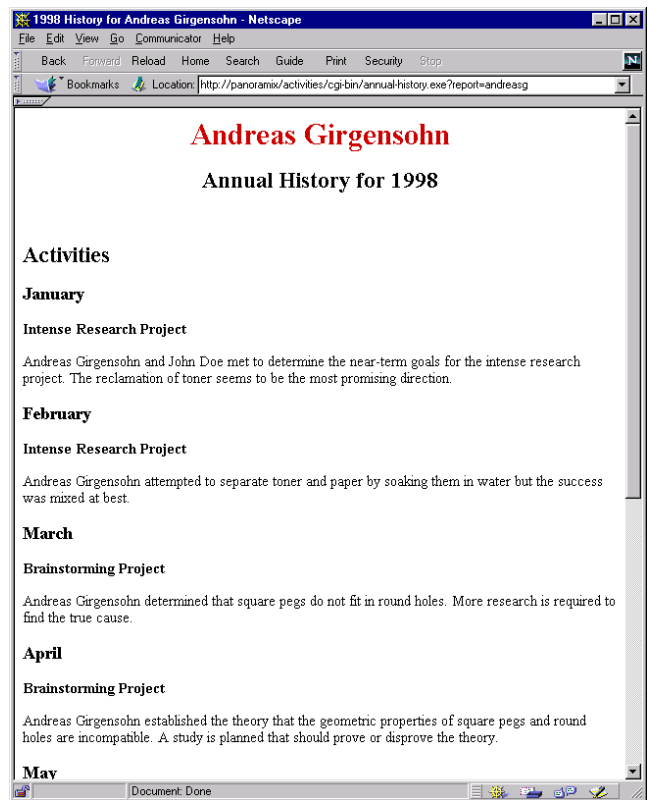


Figure 3: Annual Report History

highlights which provides a good overview of all the activities for the year.

3.2 Technical Details

The system consists of two main components, a form interface for entering individual reports and a report generator for combining several reports into a single report. The component for creating annual histories is a variant of the report generator. In addition, there are utilities for maintaining the organizational hierarchy and for browsing collections of reports.

The form interface (see Figure 2a) provides users with a simple means to enter event descriptions and to associate them with categories and organizations. Form input is saved so that users do not have to complete the form in a single session. An HTML document (see Figure 2b) is generated from the form. The document groups events by categories (e.g., “Activities”), and by organizations (e.g., “Intense Research Project”) within each category. A link is provided for making changes to the generated document with an editor.

The component for combining reports, the report generator (see Figures 2c and 2d), is the most complex part of the system. Because users can make free-form changes to their reports, the document parser has to use heuristics for determining the structure of a document. Users might have used different heading levels or just bold text to indicate a heading, or they might use line breaks instead of paragraph breaks. While such choices look almost the same in the displayed document, they are very different in the underlying

HTML representation. In searching for boundaries, differences in style (e.g., font size, emphasis, and additional whitespace) are ignored because they do not add much to the content.

The parser determines the boundaries of sections describing different categories of events as well as the boundaries of subsections associating events with organizations by collecting all heading environments in the HTML document. Heading levels are ignored because the proper nesting of the recognized headings is known and users sometimes inadvertently change the level to something inappropriate. Slight variations in section titles are handled by keeping a synonym list and checking for transposed characters as well as likely abbreviations. If marks for including portions of a report (e.g., square brackets on separate lines) are used, the parser has to identify the marks that again can be in a variety of styles (e.g., different font size). Once the boundaries of all headings have been identified, leading and trailing whitespace including blank lines is skipped in all the portions in between (again while ignoring differences in style). This provides the boundaries of all portions that can be collected in generated HTML documents.

Once the document has been divided into portions associated with a single category and a single organization, the appropriate portions can be collected into the report for an organization. To accommodate events that relate to more than one organization, a section heading can list several organizations. Such events are included in reports for all the listed organization and cross-references to the other reports are added.

The combined report is rendered as another HTML document according to the chosen options for the report structure such as the inclusion of subheadings of suborganizations. In turn, this report can be used as input for another report; either immediately after it has been generated or after it has been edited.

Both the form interface and the report generator are implemented as CGI scripts in C++. While this is not an issue for our company, this approach provides good support for geographically distributed organizations. Also, no software needs to be installed on the users' workstations (assuming that they already have a Web browser) so that they can get started easily and always have the current version of the software. For editing reports, a Web interface for invoking Microsoft FrontPage was already in use in our company at the introduction of the report system. It consists of a CGI script that sends the name of the document to be edited and a helper application that launches FrontPage and opens the desired document. Users have to click on the pencil icon on the Web page (see Figure 2b) to launch FrontPage to start editing the document.

4 USER REACTIONS

Some user reactions have already been illustrated by the email messages included in Section 2. Users sent many more messages but most feedback was delivered verbally. Several meetings were held with groups of users to explore the possibilities for making improvements and for adding

new features. All researchers at the company (about 30) still use the system every month to write their monthly reports. Project and group leaders spend more time with the system because they have to combine individual reports. Not surprisingly, many of the suggestions have come from them.

Overall, the social environment of the system is very interesting. Monthly reports are important deliverables to our parent company overseas, and thus are of high importance to the company president. The writing of monthly reports was a requirement even before the introduction of the system described in the paper so that some user resentments are unrelated to the current process or system. Some individual researchers strongly object to the fact that they have to write monthly reports. Given a choice, they probably would not write any reports no matter how easy it is to use a support system. Others see monthly reports as a good vehicle to keep management informed of progress and to keep a history of their own accomplishments. All in all, researchers only gain small benefits from writing reports so that a design goal for the system was to make the process as effortless as possible for them while substantially increasing the quality of the reports. The individual report quality has improved greatly from the reports written shortly after the introduction of the system. The better report quality is indicated by reports that are more appropriate for the intended readership at our parent company. Reports now use better English, concentrate on important events, and provide a context for the described events. The example of others and discussions about the purpose of the reports have had the desired effects.

Initially, users were skeptical. For the individual researchers, it was fairly easy to send an email message with the monthly accomplishments so that they were reluctant to switch to another system. Unfortunately, the old system placed all the burden on the group leaders and did not produce reports of sufficient quality. Also, the members of the laboratory had no simple means to stay informed about other projects. These shortcomings convinced the users to give the new system a try.

After using the system for a little while, the users did not find it any more difficult to use than the old system and in some cases much easier. Several users kept making suggestions for improvement which led, for example, to the introduction of a form for entering individual reports. Quick implementation of the suggestions convinced the users that their input mattered, and this in turn increased the acceptance of the system.

Users were pleased with the fact that they could read about the activities in the other projects. It provided them with a bigger picture of the overall direction of the laboratory and helped identify areas of possible collaboration with other researchers. The company president really liked the system because it provided him with a comprehensive monthly report rather than a collection of disconnected individual activities. The new report system also fosters collaboration between our research laboratory and our parent company. On several occasions, business units of our parent company have contacted project members at our laboratory after see-

ing the description of a project accomplishment in the monthly report. These initial contacts have led to further explorations and discussions of how project results could be incorporated into products.

For most people, the automatic reminder sent via email at the beginning of each month is sufficient, but some require individual reminders from time to time. The example of other reports has led people slowly away from writing “laundry lists” of contributions towards describing a few highlights that situate the importance of the accomplishments in a context. This trend simplifies the job of the people higher up in the hierarchy who have to do less synthesis work so that they can concentrate on editing reports for consumption up the hierarchy.

Currently, the main complaint of the users is that they occasionally cannot remember all their activities for a month. The current design does not address that issue but future additions have been envisioned that automatically check sources such as a user’s mailbox and calendar to retrieve activity information. Some users keep adding to their reports over the course of the month to deal with that problem.

5 RELATED WORK

Intranets [17, 19] usually perform some aspect of the reporting system described here. In most intranets, there are project-related pages that provide overviews of project goals and status. Unfortunately, such pages are only updated infrequently. Also, each project tends to use its own style for presenting information, making it difficult to get a quick overview of the status of several projects. Our reporting system uses Web-based technologies and integrates well with existing intranets by using hyperlinks from monthly reports to additional information stored in the intranet and back from intranet documents to monthly reports.

The monthly report system also serves as a corporate memory tool that provides a history of the accomplishments of all research projects. Monthly reports at all levels, as well as combined annual reports, are available via the Web. Some organizational memory systems such as Ackerman’s Answer Garden [1] are concerned mostly with finding the right expert. The reporting system described in this paper provides more of a genuine memory function in which the history of projects and larger organizational units can be found. In the current incarnation, access is less structured than in a typical corporate memory system, but some of the techniques discussed in the next section could overcome that potential limitation.

The hierarchical assembly of reports shares some properties with hierarchical workflow [12]. In both cases, documents float up in a hierarchy and are modified along the way. The additional feature of our report system is the fact that several documents are combined into a single document before further editing. In general, Web-based workflow systems [2, 4, 7] share the property with our system that documents become available via the Web. In addition, workflow systems focus on the control of the document flow whereas our

system supports the automatic combination of many documents.

The shared writing of reports raises issues similar to the ones found in collaborative authoring. Ede and Lunsford [5] and Tammaro *et al.* [18] discuss different patterns of social interaction such as different styles of authoring. Neuwirth *et al.* [13] propose several approaches for supporting such interactions. Part of the interaction pattern is the allocation of roles among contributors and editors. In our report system, the same effect can be observed with the researchers as contributors and the managers as editors. Many systems that support collaborative authoring (e.g., [3]) focus on synchronous aspects. But even when concentrating on asynchronous authoring, the observed interactions in our reporting systems tended to be less interactive. Exceptions are the interactions between managers and researcher where the researchers are reminded to complete their reports or asked to provide more detail on events. Authors learned from the example of others and from the changes their contributions went through on their way up the hierarchy. Further changes in the authors’ behavior were caused by occasional announcements that clarified the purpose of aspects of reports.

6 FUTURE WORK

We are currently exploring several issues. Contributions by different people on the same topic should be grouped together. The writing style of individuals sometimes prevents the verbatim incorporation of their contributions because they write in the first person or use only first names. As the descriptions of accomplishments move up the hierarchy, only the more important ones move up so that the identification of such descriptions should be supported. Automatic summaries would help in the writing of executive highlights.

Larger organizations sometimes find it particularly difficult to combine related contributions. In the current system, contributions are already automatically grouped by project but within a project they are displayed in a random order. A long-standing request by the project and group leaders has been the automatic clustering of individual contributions with respect to the topic. Usually, several people report on the same event or accomplishment but they provide different facets so that the contributions should be merged appropriately by a leader. Placing related contributions next to each other or merging related paragraphs would reduce the amount of work required of the leaders. We are currently in the process of exploring several ideas, and should have some early results soon. We plan to apply a hierarchical agglomerative clustering technique (e.g., [15]) to all contributions to the same project for a given month. As a similarity measure, we will use the inner product of a standard term frequency - inverse document frequency score (TF-IDF) with word stemming [8]. IDF will be computed from all reports in the past to overcome the relatively small amount of data in a single month. Hearst’s TextTiling system [9] addresses the opposite problem of finding topic boundaries in text. In her system, terms are filtered through a stop list and reduced to their roots by a morphological analysis function. Term sequences of fixed length are used for compari-

son because normalizations cannot deal with sentences of widely different lengths. For the actual comparison, Hearst uses the normalized inner product of the term frequencies in the sequences. She reports that TF-IDF turned out to be less robust than just using term frequencies. We expect that TF-IDF should work better for our domain because of its relative homogeneity.

At the same time, to address the problem of people writing in the first person or just using first names, the system will recognize names from a database of employees, a collection of first names [11], and a list of common titles (e.g., “Ms.,” “Dr.,” “-san”). In case of duplicate first names (e.g., “John”), the person closest to the author in the organizational hierarchy should be used. Names of outside collaborators and employees of our parent organization should be automatically learned from reports and added to the name database. The name database could also be used to correct misspelled names. Normalized names should provide much better results when used as terms in measuring the similarity of paragraphs. Furthermore, the system can automatically change sentences to the third person and expand names to full names.

Another goal is to support the task of identifying the most important accomplishments for the monthly highlights. Topics described by several people are good candidates for that. Alternatively, individuals can prioritize project accomplishments and the system can automatically combine the priorities and determine the most important accomplishments.

For the corporate memory function of the monthly reports, tools for additional summaries are desirable. Text summarization techniques can extract an executive summary from an annual project report. Most domain-independent approaches to text summarization are based on statistical techniques (e.g., [16]). It is also possible to incorporate a commercial system such as InXight’s LinguistX [10]. Text summarization techniques do not always produce correct and easily readable summaries, so that it might be more appropriate to use the automatically generated summaries as starting points for a human editor. Key phrases can be used to automatically insert hyperlinks to background information such as milestones in the annual project plan. Golovchinsky [6] discusses the automatic generation of links in detail.

7 CONCLUSIONS

The experiences described in this paper show that it is possible to change a monthly reporting practice so that both ease of use and the quality of the reports are improved. This is mostly due to the fact that the system took advantage of the existing organizational hierarchy and supported the collaborative assembly of reports within that hierarchy. The automatic combination of reports while preserving the section structure made it much easier for project and group leaders to incorporate all the relevant accomplishments and to concentrate on telling a coherent story. In addition, many of the improvements have been made in response to user feedback or to observed problems. The constant involvement of users has both increased the acceptance and improved the system.

In using the system, the users have become more aware of activities in other research projects so that collaboration across projects has been fostered. The archive of monthly reports has served as a corporate memory of research activities that helps new employees to get started and provides an overview of research activities over a longer period. On several occasions, contacts with our parent company were initiated by a description in a monthly report. These initial contacts have led to further explorations and discussions of how project results could be incorporated in products.

The applicability of this process is not restricted to our organization. Because both organizational hierarchies and sections are customizable via tables, the system could be adapted easily to another organization with a different report structure.

Future enhancements are expected to simplify the reporting process even further. Some of these enhancements are mainly beneficial for projects and other organizational units with many contributors but others are in direct response to observed problems with the current system. Some problems reflect on people’s preferences (e.g., it is more natural to write about accomplishments in the first person). We continue to solicit user feedback, and to observe problems with the system to suit the needs of the user community.

ACKNOWLEDGMENTS

I would like to thank the users of the system for their valuable feedback. I would also like to thank Lynn Wilcox, Alison Lee, Morgan Price, Patrick Chiu, Gene Golovchinsky, and Elizabeth Churchill for suggesting improvements to this paper.

REFERENCES

1. M.S. Ackerman. *Augmenting the Organizational Memory: A Field Study of Answer Garden*. In Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW’94), New York: ACM, pp. 243-252, 1994.
2. C.K. Ames, S.C. Burleigh, and S.J. Mitchell. *A Web Based Enterprise Workflow System*. In Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work, GROUP’97 (Phoenix, AZ), New York: ACM, pp. 214-220, 1997.
3. R.M. Baecker, D. Nastos, I.R. Posner, and K.L. Mawby. *The User-centred Iterative Design Of Collaborative Writing Software*. In Proceedings of INTERCHI 1993 Human Factors in Computing Systems, New York: ACM, pp. 399-405, 1993.
4. P. Dourish, J. Holmes, A. MacLean, P. Marquardsen, and A. Zbyslaw. *Freeflow: Mediating Between Presentation and Action in Workflow Systems*. In Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW ’96), New York, NY: ACM Press, pp. 190-198, 1996.
5. L. Ede and A. Lunsford. *Singular Texts / Plural Authors: Perspectives on Collaborative Writing*. Carbondale, IL: Southern Illinois University Press, 1990.

6. G. Golovchinsky. *What the Query Told the Link: The Integration of Hypertext and Information Retrieval*. In Proceedings of Hypertext 97, New York, NY: ACM Press, pp. 67-74, 1997.
7. W. Gräther, W. Prinz, and S. Kolvenbach. *Enhancing Workflows by Web Technology*. In Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work, GROUP'97 (Phoenix, AZ), New York: ACM, pp. 271-280, 1997.
8. D. Harman. *Ranking Algorithms*, in *Information Retrieval Data Structures & Algorithms*, W.B. Frakes and R. Baeza-Yates (Eds.). Prentice Hall: Upper Saddle River, NJ, pp. 363-392, 1992.
9. M.A. Hearst. *TextTiling: Segmenting Text into Multi-paragraph Subtopic Passages*. Computational Linguistics. **23**(1): pp. 33-64, 1997.
10. InXight, *LinguistX*, <http://www.inxight.com>: 1998.
11. M. Kantrowitz, *Name Corpus*, <http://www-cgi.cs.cmu.edu/afs/cs.cmu.edu/project/ai-repository/ai/areas/nlp/corpora/names/0.html>, 1994.
12. P. Mambrey and M. Robinson. *Understanding the Role of Documents in a Hierarchical Flow of Work*. In Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work, GROUP'97 (Phoenix, AZ), New York: ACM, pp. 119-127, 1997.
13. C.M. Neuwirth, D.S. Kaufer, R. Chandhok, and J.H. Morris. *Computer Support for Distributed Collaborative Writing: Defining Parameters of Interaction*. In Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94), New York: ACM, pp. 145-152, 1994.
14. K. Okamura, M. Fujimoto, W.J. Orlikowski, and J. Yates. *Helping CSCW Applications Succeed: The Role of Mediators in the Context of Use*. In Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94), New York: ACM, pp. 55-65, 1994.
15. E. Rasmussen. *Clustering Algorithms*, in *Information Retrieval Data Structures & Algorithms*, W.B. Frakes and R. Baeza-Yates (Eds.). Prentice Hall: Upper Saddle River, NJ, pp. 419-442, 1992.
16. G. Salton, A. Singhal, C. Buckley, and M. Mitra. *Automatic Text Decomposition Using Text Segments and Text Themes*. In Hypertext'96: Seventh ACM Conference on Hypertext, New York: ACM, pp. 53-65, 1996.
17. R. Scheepers and J. Damsgaard. *Using Internet Technology within the Organization: a Structural Analysis of Intranets*. In Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work, GROUP'97 (Phoenix, AZ), New York: ACM, pp. 9-18, 1997.
18. S.G. Tamaro, J.N. Mosier, N.C. Goodwin, and G. Spitz. *Collaborative Writing Is Hard to Support: A Field Study of Collaborative Writing*. Computer Supported Cooperative Work: The Journal of Collaborative Computing. **6**: pp. 19-51, 1997.
19. C.M. Vishik. *Internal Information Brokering and Patterns of Usage on Corporate Intranets*. In Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work, GROUP'97 (Phoenix, AZ), New York: ACM, pp. 111-118, 1997.