

GroupIntelligence: Automated Support For Capitalizing On GroupKnowledge

Stephan Verveen

GroupSupport.com

Stephan.Verveen@GroupSupport.com

Michiel van Genuchten

GroupSupport.com

Michiel.van.Genuchten@GroupSupport.com

Robert Schuwer

GroupSupport.com

Robert.Schuwer@GroupSupport.com

Theo Bemelmans

Eindhoven University of Technology

T.M.A.Bemelmans@tm.tue.nl

Jules de Waart

Ministry of VROM

jules.dewaart@dio.cs.minvrom.nl

Abstract

Group Support Systems enable harvesting knowledge from groups. The minutes of a typical GSS meeting are mostly textual reports of 10 to 50 pages summarizing the issues discussed, the set priorities and proposed actions. Capitalizing on knowledge generated during GSS meetings appears to be very difficult without additional automated support. In this paper we describe our experiences with several prototypes for computer-aided support for consolidating the results of GSS meetings over time. The development of the prototypes resulted in an online service called GroupIntelligence™. This service transforms accumulated group knowledge from various sources into a dynamic consolidated website.

1. Introduction

This paper describes the results of a R&D project executed over the last year. The goal of the project was to invent a software tool that allows users of GSS software to do more with the results from their electronic meetings. The R&D effort resulted in an Internet service called GroupIntelligence, or GroupI in short. The project was carried out in an industrial environment, in close cooperation with the Eindhoven University of Technology in the Netherlands.

Group Support Systems in general enable collecting large quantities of experience and knowledge available in groups [3],[4]. The result of a typical GSS meeting is mostly a textual report of 10 to 50 pages. The participants and leaders of the groups involved typically use these minutes of individual meetings. From practice it appears to be interesting but also very complicated to put the results of a large number of GSS meetings to use. It is difficult to answer questions such as:

- "Did we not discuss this before?"
- "What were the results of our 20 meetings held during the last twelve months on the topic e-commerce?"

- "When did we discuss the quality problems with our suppliers and what were the most remarkable results?"

One could go through the minutes of all electronic meetings manually to find out the most interesting results, but our experience indicates that users hardly reserve time for this laborious job. It is even more compelling to capitalize multiple sources of group knowledge. For example: a business plan for a company may be represented in a formal text document and be the result of group processes, for instance a series of GSS meetings. One would like to be able to access all the information used in that group process like documents and reports with background information, as well as information available in other data types such as presentations and audio and video files.

GSS as a vacuum cleaner

Group Support Systems are very good for collecting information from a group of people. It acts as a vacuum cleaner that gathers any data, information or knowledge it comes across in a group. After the meeting, a report is created that can be put to use by the participants of that specific meeting. The dust-container is replaced for the next meeting. With each meeting, the organization builds up a collection of hundreds of dust-containers that is hardly accessible as a whole. For real dust-containers this is not a problem, because they contain nothing valuable. For the results of GSS meetings, this is a waste since GSS meetings typically contain very valuable group knowledge that can be of value to an organization in a later stage.

(This metaphor is only to be used by insiders i.e. people that are convinced of the added value of Group Support Systems)

2. The initial idea and customer response

A R&D trajectory was started to develop automated support for improved processing of the results from GSS meetings. The initial idea was to pursue the development

of a database driven website for publishing the results from electronic meetings to the Internet. The working title for the product concept became "Group Intelligence" or "GroupI" in short.

To get early feedback on the R&D project that was about to start, a GroupSystems Survey was deployed among experienced GroupSystems users. The respondents (n=45) had attended 46 electronic meetings on average within a range from 10 to over 100 meetings. In total 73% of the respondents were from the Netherlands. The other respondents were from the UK, US, Switzerland, Portugal, Japan and Hong Kong.

The main conclusions that were drawn from the survey results are presented below.

- GroupSystems users feel their organization can improve the way it capitalizes on the results from its electronic meeting (Strongly agree: 40%, Agree: 49%; Neutral: 7%; Disagree: 4%)
- They think that their organization can benefit from accumulating and publishing the results from electronic meetings (Yes: 76%; No opinion: 18%; No: 4%)
- They believe that searching across results from multiple electronic meetings may reveal valuable relationships (Strongly agree: 16%; Agree: 47%; Neutral: 29%; Disagree: 8%)
- Utilizing results from previous meetings may improve the effectiveness of new electronic meetings (Strongly agree: 23%; Agree: 54%; Neutral: 23%)
- GroupSystems users prefer a website above hyperlinked documents as a medium for electronically publishing e-meeting results.(See Figure 1)




Figure 1. What medium for capitalizing on results do you prefer?

The conceptual foundations for the functionality of the GroupI product are described in an article by Schuwer et al. [6]. In this article, a "knowledge cycle" is proposed, of which the basic elements are given in figure 2. This cycle supplies a framework for classifying software functionality for supporting knowledge management processes.

GroupSystems is a tool that harvests knowledge from a group of people by storing the textual representation of the knowledge in its databases. Based on the framework in figure 2, GroupSystems can be considered as a tool for knowledge acquisition and storage. Up until now, the users of GroupSystems have accumulated masses of meeting data throughout the years, but have no proper means to reflect on their meeting-results in retrospect. Currently there is no functionality in GroupSystems that supports the effective location, dissemination and application of the stored meeting data.

A meeting publishing tool can allow improved knowledge storage and dissemination (see figure 2) by storing additional contextual data and special relations between activities and meetings. The publishing tool can also make the data from GroupSystems easier to disseminate and provide 24 hours a day availability of the results. Special features can enable fast location, selection and retrieval of the available data. Because of these improvements, the data can be reused in new meetings and improve the meeting productivity. Furthermore, additional analyses of the results from previous meetings can be made.

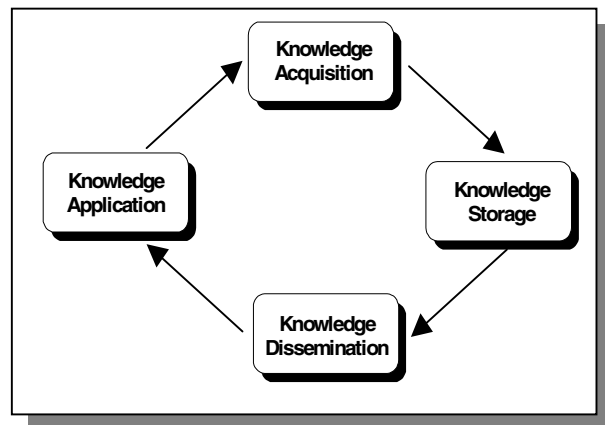


Figure 2. A simplified representation of the "knowledge cycle" by Schuwer et al [6]

3. GroupI development and early applications

A development trajectory was started based on our preliminary ideas and the confirmation we received from the experienced GroupSystems users. Development was structured in small evolutionary steps in two parallel series. The first series of prototypes was aimed at supporting tools for individual meetings, the second was aimed at tools for multiple meetings. The result of each evolutionary step was a working prototype of the product. Both trajectories will be discussed below.

2.1 Support for single meetings

Publishing a single meeting means that the result of one meeting with GroupSystems is converted to GroupI and then published to the Intra/Internet. Several prototypes were developed to publish single meetings. Each prototype on itself was a real life case from customers. With the single meeting publisher, a user can browse the results per activity, access the data through several paths, sort the data on several keys and search for the presence of certain strings in the results. For the first and fourth prototype, a short description of the realization process is given along with the most important lessons learned.

2.1.1 First prototype The first prototype was based on a specific technology (IDC/HTX) that was chosen because of its easy implementation. The constructed prototype was able to display the data from the GroupI database in a flattened list. In a text-field, users could enter a string to perform a simple free-text search on the meeting data.

A "mail meeting owner" button was implemented to elicit feedback on the meeting results from the attendants. This allowed extending the time in which participants could give feedback, without bothering them with complete collaborative functionality. Even though the first R&D efforts looked promising, the displayed information was hard to interpret. (Winograd 1983) A shortcoming was that the original agenda and the activity structures were not displayed. Furthermore, contextual data regarding the meeting goal, the used approach and the attending persons was absent.

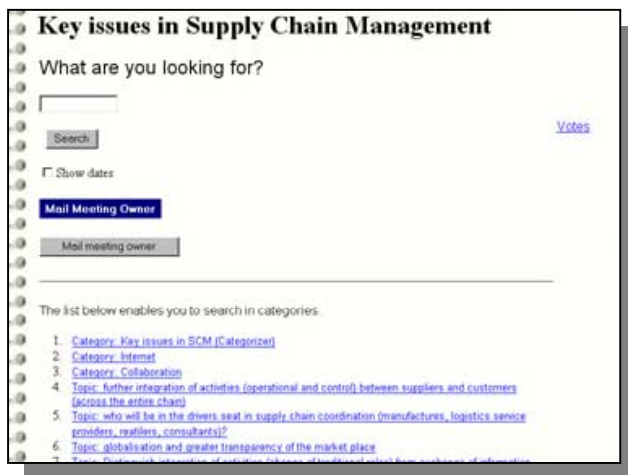


Figure 3. First prototype

2.1.2 Fourth prototype Between the development of the first and fourth prototype, Active Server Pages (ASP) technology replaced the IDC/HTX technology because this opened up the opportunity to make the output application completely independent from the content in the GroupI database. The output now allowed browsing

through and performing searches on all the data in the activities from the GroupSystems agenda. The search results now also presented hyperlinks to the originating activities. The search functionality proved to be useful for finding the numerous remarks regarding certain subjects from all contributors to the meeting.

Especially the quantitative tools required special attention because of the more complex database operations that must be made to calculate aggregate statistics.

The presentation of the results closely followed the GroupSystems format, since participants became familiar with this format during the meetings. In the presentation also dates, time, activity descriptions, etc. were included. All the representation logic was fed with data from the GroupI database, without any required programming efforts for each new dataset.

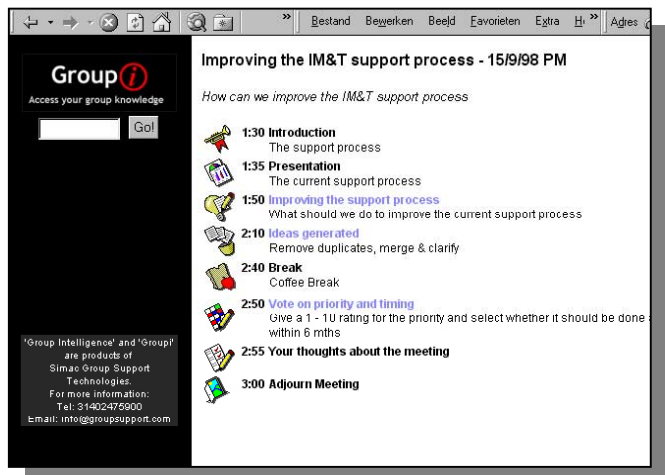


Figure 4. The fourth prototype

2.1.3. Developed functionality The key functionality that evolved from the developed prototypes for publishing single meetings is given below:

- **Dynamic, database driven output generation.**
Databases with results from different meetings can be accessed through exactly the same output application.
- **Content retrieval through browsing the original meeting structure**
Based on the additional meeting structure information in the GroupI database, hyperlinks are created that give access to specific parts of the results of a meeting.
- **Mail meeting owner**
Making it easy to contact the original meeting "owner" may help resolving questions regarding the meeting results and promotes continuing discussions based on the meeting results.
- **Show and exploit relations between activities**
The GroupI database stores process data that was not available in the original GroupSystems database.

Based on this data, it is possible to show relations that clarify the data manipulations during a meeting (e.g. a shift from a brainstorm to a vote). This enables linking an idea from a brainstorm activity to its vote results in a later activity.

- **Results sorting**
Data can be sorted on several attribute values such as date and mean vote result.
- **Free-text search**
The results can be searched for occurrences of a certain string of characters. For example: if a text string is found in a comment, GroupI will indicate to which idea this comment belongs, in which category and which activity the idea was generated and if the idea was put to a vote.

2.1.4. Lessons learned Based on the developed prototypes and the interaction with the users, the following lessons were learned:

- **Background and reasons for a meeting must be available in the meeting result presentation**
In many meetings the facilitator explains the meeting objective verbally, sometimes supported by a slideshow. This information has to be presented to the GroupI users in order to enable them to interpret the meeting results correctly.
- **Meeting data often needs additional processing**
The names, descriptions and instructions from most meetings are unclear, missing and often too brief.
- **Display the results in their original structure for people that attended the meeting**
Users that attended the meeting can browse the original meeting structure to review results. Users that did not attend the meeting do not recognize the presented structure and may be served better through a different presentation format.

2.2 Publishing multiple meetings

The first development of the prototypes for publishing multiple meetings originated from a project for the Dutch ministry of VROM. A project was started to develop an application that would provide easy access to the results from previous meetings on a server at the ministry. This application is called the meeting manager.

2.1.1 An extension of the single meeting publisher The first prototype for the VROM project was an adaptation of the application that was used to publish single meetings. The adaptation required changes in the database structure and additional support for searching across multiple meetings. The first prototype was used to elicit feedback from the users at VROM regarding missing functionality and the screen layout. Based on the feedback on this first

prototype, it was decided to extend the stored data with additional clarifying data to make a correct interpretation of the presented results possible.

2.1.2 Applying superstructures and keywords With the decision to recreate the single meeting publisher with ASP code, the meeting manager had to be rebuilt from scratch also. After building the new single meeting output functionality, new functionality for displaying and browsing superstructures was created. These superstructures (series) allow the user to access the data from a higher hierarchical level, which reduces the required efforts for locating results. The series are views that are defined by the manager of the data that is stored within the GroupI database. In the VROM case, meetings with related subjects were shared under the same series. Figure 5 shows four series that contain over thirty meetings in total.

By attributing descriptions and keywords to the series and the meetings within a series, locating data was made easier. Finding data within meetings based on keywords is a powerful search mechanism. The prototype was fitted with an option that shows an (alphabetical) list of all the used keywords. In this view, the user can view what keywords are used and the number of meetings or series that are related to them.

The search function was still based on a simple free-text query. In the results-screen however, a number of improvements were made. During the retrieval of the results, the number of found results is presented to the user instantly. This way, the user can tell right away whether or not the search statement was too generic. Each found result is presented with its date of creation. The user can sort the results based on these dates to find the results from a certain date more quickly.



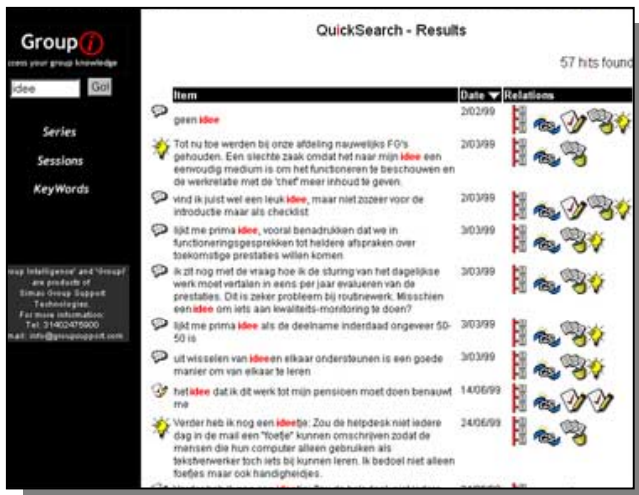
Figure 5. Combining meetings in series

Next to each of the results, the relationships with all the hierarchical levels in the meeting data are displayed. By clicking on the icons that represent these relationships,

the user is taken to the series, agenda, activity or idea that is related to the result.

2.1.3 Developed functionality The last prototype was implemented at the ministry of VROM. It is now possible to answer the question ‘did we not discuss this before’ right away in any GSS meeting. GroupIntelligence is installed standalone on the leader station and the facilitator can switch between GroupSystems and GroupIntelligence at any time during a meeting. The GroupI database currently contains 40 meetings and the intent is to grow the database over time. The key functionality for publishing multiple meetings is given below:

- **Series**
Series are superstructures that combine the results from different meetings and make these accessible from a higher hierarchical level. The series functionality is the first steps towards integrating results from multiple meetings.
- **Keywords**
Keywords provide a brief taxonomy of the content of a series or meeting. The used keywords can be used as a point of entry to meeting results. This function also allows increased integration of meeting results across multiple meetings.
- **Sorting mechanism**
By sorting results on a certain attribute, a user can browse the results quickly.
- **Hit highlighting in search results**
By highlighting the search string in the search results, the relevance of the found results can be assessed quicker.
- **Cross structural relations in search results**
For each search result, it is possible to browse the structure it belongs to on all hierarchical levels. In the future, more advanced search technology will be considered. [1],[2]



2.1.4 Lessons learned Experiences in processing multiple meetings from the past for this project learned the following lessons:

- **Enriching data from multiple meetings afterwards is a laborious job**
Recalling contextual data from previous meetings from memory appears to be difficult. In this specific case, the expected benefits of the publication of many meetings with GroupI afterwards did not outweigh the required efforts to clean up and enrich the data. There is an analogy with reuse of software: trying to make software reusable as an afterthought after the development proves to be very difficult. It has to be designed and the goal to reuse software has to be taken into account from the start.
- **Applying superstructures to meetings can be difficult**
Without a preconceived plan for conducting related meetings, it is difficult to apply superstructures to the data afterwards. This might be compared to sifting waste without separation at the source.
- **Superstructures help in locating results from different meetings**
By reducing the number of options to browse, the user can browse the available results top down from a higher hierarchical level.

2.1.5 Feedback from users In the last prototype version of the single meeting publisher, a link to an online survey was included to elicit feedback with regard to the use of the GroupI application. The intermediate results from this survey (n=17) are included to illustrate the results so far with remarks from experienced GroupSystems users. The main preliminary conclusions, illustrated with remarks from open-ended questions, are given below:

- **A majority judges the opportunities for disseminating GroupSystems results to be better than a traditional GroupSystems report.**
(Much better: 23.5%, Better: 70.5%, The same: 6%)
- **A majority judges the opportunities for finding results with the GroupI application to be better than a traditional GroupSystems report.**
(Much better: 58.8%, Better 29.4%, The same: 6%, Worse: 6%)
- **A majority considers GroupI to be a valuable supplement to GroupSystems.**
(Yes: 88.2%, No Opinion: 11.7%)
- **A majority thinks GroupI can support for knowledge management processes.**
(Yes: 82.4%, No Opinion: 17.6%)

One respondent mentioned: "It is easier to make cross references to knowledge that was generated in other meetings. The accumulated knowledge is preserved better in this way. Furthermore, results can be easily disseminated among people that did not attend the

meeting." But: "GroupI requires additional functionality. Knowledge management is not about storing everything, but about making good choices regarding content and representation mechanisms."

Advantages of GroupI that were mentioned were:

- "It has become much easier to 'zoom in and out'. For example when looking at the vote results, you can easily go deeper and look at the comments behind the topic." "The smaller amount of information that is presented to the reader makes the information more accessible."
- "Specific searches are possible and the result is not an interpretation, but an exact representation."
- "Electronic publication opens up opportunities for future functionality."

Notes of criticism on the current version were:

- "We are still left with a large amount of information which is difficult to communicate to those who didn't participate in the meeting."
- "The results of some sessions can not be viewed in the right context when leaving the oral discussions out."

3 Future plans

The future plans with GroupI include the following activities:

- **Become more platform independent**
The focus of GroupI is the knowledge that is elicited from groups. One data type in which this kind of knowledge is captured is in the results of GSS meetings. Other sources of group knowledge may be audio and video clips, text documents, presentations or e-mail. GroupI will incorporate group knowledge available in these other datasources. On the output side there is also a need to become more platform independent. A migration to XML is likely.
- **Pro-active knowledge harvesting**
The development of GroupI started with the goal to do more with the results of GSS meetings as a reactive activity. During the last year, it became clear that the most powerful application may be pro-active knowledge collection where the group knowledge is the goal and the GSS is one of the means to collect it. The scenario is the following: one wants to collect the knowledge on a specific subject available within an organization. The following schedule is likely:
 - interview experts one-on-one to get their in-depth feedback
 - gather information on the subject from the Internet

- conduct a survey among members of organizations that have experience with the subject on hand
- plan a sametime GSS meeting with those respondents to the survey with most useful feedback
- execute distributed GSS meetings to collect information from experts, both within and outside the organization.

The results of all sources of group knowledge are incorporated in a GroupI database that is made accessible to all those relevant. In March and April 2000 we conducted an effort to collect information on virtual teams. The results can be found on <http://www.groupsupport.com>. It incorporates both the results from an electronic meeting on virtual teams, as well as a HICSS paper [7] on experiences with virtual teams.

4. Conclusions and recommendations

Goal of the R&D project was to do more with the results from GSS meetings. The first experiences with publishing single meetings proved to add value to the use of GroupSystems. We believe that the real power of using GroupI lies in its ability to merge the results of multiple sources of group knowledge, such as documents, GSS results and audio and video clips. Merging these results makes it possible to compare the results from different groups and analyzing the development of opinions over a period of time.

An important lesson that we learned over the course of this project is that the structure in which group knowledge is presented should in many cases be different from the way in which it is collected. One example to illustrate this: we used a distributed GroupSystems meeting to collect all the possible problems a user of a specific software product could encounter. With the Group outliner tool, the defects were classified as hardware, network, software and usage errors. 200 possible defects were classified in 20 classes. This worked very well for the collection phase. The results of that GroupSystems meeting were converted to a website that is accessible to the users. When a user is confronted with an error while using the software, he does not care about this structure. If the error message on his screen shows: 'error2361', the user only wants to do a search through the database for this string. In case of a more complicated problem, the user may need the structure used for collection or another structure to find the cause and a solution.

Nunamaker stated in IEEE Computer in 1999 [5]: 'Ultimately, I envision group support systems not only as facilitators of decision making and collaborative work,

but as gateways to a knowledge repository that archives an organization's collective expertise.' If Group Support Systems as we know them are regarded as the gateway in, successors of tools such as GroupIntelligence may give a first indication on improving the gateways out of the knowledge repository.

5. References

- [1] Fayyad. U. M., ed. (1996), "*Advances in Knowledge Discovery and Data Mining*", Boston: MIS Press
- [2] Gray, P. and Watson, H. (1998) "*Decision Support in the Data Warehouse*", PrenticeHall
- [3] Morrison, J.; Morrison, M.; and Vogel, D. (1992) "*Software to Support Business Teams*," Group Decision and Negotiation, September issue, pp. 91-115
- [4] Nunamaker, J.F. et al. (1997) "*Lessons from Ten Years of Group Support Systems Research: A Discussion of Lab and Field Findings*", Journal of Management Information Systems Vol. 13, issue 3
- [5] Nunamaker, J.F. (1999) "*Collaborative Computing: The Next Millennium*", IEEE Computer September issue
- [6] Schuwer, R. et al. (1999) "*Kennismanagement en de rol van informatietechnologie*", Management en Informatie
- [7] Vogel, D (2000). "*Distributed Experiential Learning: the Hong Kong - Netherlands Project*," Proceedings of the 33rd Annual Hawaii International Conference on System Sciences, CDRom
- [8] Winograd, T. (1983) "*Language as a cognitive progress (Volume I: Syntax)*", Addison-Wesley Publishing Company