Service Architecture, Prototype Description, and Network Implications of A Personalized Information Grazing Service

Howard E. Bussey   Carmen Egido   Amy Kaplan   Steven L. Rohall   Rosanna Yuan†*

Bell Communications Research
445 South Street
Morristown NJ 07960-1910 USA

ABSTRACT

We present a Passive Information Grazing System (PIGS) which is a prototype information delivery and filtering system for casual users. In PIGS, information articles in multimedia form are electronically transmitted to users over a prototype broadband network. To assist the users in managing the potential flood of articles, PIGS selects information for individual users based on their personalized interest profiles. As users' interests change over time, their profiles could be automatically adapted using implicit and explicit feedback from users about each article. We examine and compare the network implications of several alternative implementations of PIGS, with particular emphasis on transmission and call processing capacity requirements.

1. Introduction

Information technologies in general have been extremely successful in facilitating the creation and dissemination of information. They have been less successful at addressing information management problems. With the advent of high speed electronic networks capable of delivering large amounts of multimedia information to home and office, the deluge of information that we all face is likely to worsen. Our motivation in carrying out this research is twofold. Firstly, we wish to further the understanding of potential uses for public networks by users and vendors. Secondly, we wish to examine the implications of such usage on the network.

This paper describes an information delivery service aimed at ameliorating this problem while at the same time allowing, and actually enhancing, the casual and often accidental access to information. Sections 2 and 3 discuss the motivation for our focus on passive information activities and describe our Passive Information Grazing System (PIGS) service proposal and prototype implementation. Section 4 presents a number of alternative implementations and compares their impact on the network with respect to transmission and call processing requirements.

2. PIGS Service Description

2.1. Motivation

Currently available information services allow users to access information via queries or hierarchical directories. To use these services, users must be willing to invest the time and effort required to learn the query and navigation procedures, as well as the database organization. They must also be willing to engage in active information search. As a result, commercial information services find their audience only among small numbers of users who have a need to find answers to specific, well-defined questions. We refer to this mode of information access as information hunting.

Information grazing is at the other end of the spectrum of information activities. In this mode of information access we receive information in a passive, largely unsolicited, manner from a variety of heterogeneous, multimedia sources such as other people, television, newspapers, direct mailings, and billboards. The only effort that we typically expend to acquire this information is the gross selection of information channels, e.g., selecting a particular TV channel or subscribing to a particular magazine.

This more passive and serendipitous form of information access in fact accounts for a large proportion of our information activities. To accommodate information grazing, on-line information services must be designed as information delivery services rather than as information retrieval services. We adopt this view in the Passive Information Grazing System (PIGS) described below.

2.2. Service Description

In essence, PIGS is a personal "current awareness" information system. The system anticipates a time when a wide range of information sources in a variety of media will be accessible from home and office via broadband networks. To help manage the potential flood of information, PIGS provides personalized filtering, based on each user's own profile, which discards uninteresting or irrelevant articles. PIGS differs from traditional information retrieval systems as well as from Selective Dissemination of Information systems in that it is designed to provide passive access to information that is automatically customized to satisfy individual users' information needs and desires as these evolve over time. During presentation of articles that have passed through the filter, users' workstations or PCs can record explicit or implicit feedback about each article. Based on the feedback data each user's profile can be automatically adapted to track his or her changing information needs.

Conceptually, PIGS resembles standard information retrieval systems in that it must accomplish the following three goals:

1. It must decide what sort of information the user is requesting from the system.
2. It must determine the subject matter of the individual items in the available database or databases.
3. It must appropriately map representations of the user's information goals and interests onto representations of the topics addressed by the items in the database(s) in order to provide the user with satisfactory items.

*Present address: University of Chicago, Department of Music, 5845 South Ellis Avenue, Chicago, Illinois 60637.

†Present address: Shia Systems, 1825 Old Driftway Lane, Lebanon NJ 08833.

Unlike standard information retrieval systems, PIGS cannot rely on partial content specifications of desired subject matter in the form of active queries formulated by the user to achieve the first of these goals. Instead, it determines items the user would like to see by selectively filtering information by referring to a personal interest profile associated with that user. This profile can be automatically adapted to the user's changing information interests and needs through a feedback mechanism that observes system usage and related behavior, as well as through feed-forward mechanisms in which the system occasionally queries the user.

Our PIGS proposal assumes that the information is in the form of articles, which we define to be collections of multimedia (text, voice, audio, graphical, image, or video) information items. For instance, an article representing a lecture could have a text information item (the text of the speech), a voice information item (the recording of the lecture), and several graphical information items (the slides or viewgraphs used a visual aids). We also assume a Broadband Integrated Services Digital Network (BISDN) as the infrastructure providing Asynchronous Transfer Mode (ATM) transmission and switching (including multicasting) for this system. It also requires efficient transmission from the publishers, or service providers, to the users, which we shall show can be provided by BISDNs. Finally, we assume that users have Personal Computers (PCs) for storing and presenting articles to users.

2.3 Related work

The problem of floods of information overloading users has been recognized since the early days of computer mail networks [Denning '82, Hiltz '83]. Hiltz and Turoff [Hiltz '82] and Malone, et al., [Malone '86] have proposed *filtering* as a way that users can manage the flood of information. In the Information Lens [Malone '86], a rule-based profile is used in filtering electronic mail messages. Users modify their own profiles using a graphical rule editor or with a rule-creation-by-example facility [Mackay '89]. Mackay, et al., report that users used filtering to identify important messages, to delete irrelevant messages, and to aid in subsequent message retrieval, with some explicit adaptation by a small group of the Information Lens users. In PIGS, we emphasize multiple, multimedia sources with automatic profile adaptation.

3. Prototype Description

3.1. Prototype Software Architecture

The software architecture for the PIGS prototype is shown in Figure 1. The information sources are associated with separate source processes, running on a variety of mini- and mainframe-computers. Each of these sources communicate with the input conversion processes in the prototype "service-provider" minicomputer. Articles from the sources are automatically collected by the conversion processes (which periodically poll their respective sources) and stored in a database. Newly received articles are classified according to subject matter by the classification process, and newly classified articles are transmitted to the PIGS prototype users. The input conversion, classification, and transmission processes run independently of each other, using only information in the database for interprocess communication and synchronization. The transmission process makes "calls" through the Experimental ATM Network and Services Environment (EXPANSE) BISDN prototype[1] to the receive processes that run on each PIGS user's workstation, and transmits each article to each user.

In addition to the receive process, there are filtering and user interface processes running on each user's workstation, along with a database that stores the user profiles and the received articles. As with the service-provider processes, the user-module processes operate independently, using the database for synchronization and communication.

This architecture provides an environment in which processes can be rapidly written, debugged, and integrated because of the process independence and the well-defined database interfaces. It is also easy to add users, since each user's user-module processes communicate only with themselves and the service-provider's transmission module.

3.2. Sources and source processing

Articles are automatically retrieved from various sources, reformatted in some cases, examined, and incorporated into the PIGS database. At present, the majority of these articles are textual, although we will presently integrate video articles from television network stations and other multi-media articles. We are using sources that provide a continuous flow of mostly transient (as opposed to archival) information. These include: Netnews, an international bulletin board comprising over 100 discussion groups on topics ranging from operating systems to vegetarian cooking (we use about 20 such groups in PIGS); Telaris™, a Bellcore database of abstracts of telecommunications articles from over 150 publications; AP newswire, a continuous stream of newspaper articles published by the Associated Press; and a database of abstracts of Bellcore's internal Technical Memoranda. Each of the sources is polled with varying frequency, and articles are placed in our database for varying amounts of time. Information flow is moderately heavy with an average of 1000 articles daily. Articles are processed in different ways depending on their original source. For instance, for the AP wire we take advantage of headers containing information on how to treat individual articles. When a header contains instructions to use the current article to replace or update an earlier version of the same article, the conversion process automatically interprets this information and updates the PIGS database accordingly.

3.3. Service provider classification, editing, transmission

The service-provider processes use a commercial database management system to store or access information about articles and the PIGS prototype users. There are about 20 relations that describe the articles, the information items that articles comprise, the classification information (subjects and keywords), and the PIGS users. The main relation is the article relation, with one record for each article stored in the system. Each article can have multiple information items, and there are several types of information items (e.g., text, voice, video, etc). This structure is reflected in the database, where records in information item relations define the tree of information items that an article comprises. The article records contain a status field, which is updated by each

---

[1] See [Hayward 87] and [Bussey 88] for Broadband prototype descriptions.

™ Telaris is a trademark of Bell Communications Research, Inc.

service provider process indicating the current processing state of the article, and a creation-date field, so that articles can be deleted at the end of their retention period.

The input conversion processes poll the sources periodically to obtain new articles and store the articles as individual Unix files, and the information about the articles as records in the database. The classification process polls the database several times each hour to check for newly entered articles. If any are found, it scans through each article to collect title, author, and publication date information. It also determines the subjects and keywords for the articles. Some sources provide header items that specify the article subject and keywords, while articles from other sources require the use of keyword extraction tools.

The transmission process also polls the database several times per hour to find articles to transmit to the PIGS users. When some are available, the process establishes connections with the user stations and transmits the articles to each user through the EXPANSE prototype. Although articles are stored in the database as a hierarchy of information items, they must be transmitted as a flat stream of data. The transmit process "flattens" each article into a sequence of bytes, representing the sections that contain information such as the article ID, title, keywords, and information item contents.

3.4. User reception, filtering, user interface

The key structure in the user module is the personal interest profile. It contains keywords and their weights for each of the various subject classifications that PIGS uses. The user enters any number of keywords to describe his or her interests within a given subject, and keyword weights are programmatically generated, based upon the ordering and number of keywords entered by the user. Changes to the profile must currently be entered manually, but we envision that the profile editor will be used only minimally once profile adaptation mechanisms are in place.[1]

Articles are inserted into the local database as they are received from the network and are filtered on a regular basis, as follows. The list of keywords and their frequencies of occurrence, generated by the classification process in the service provider and the list of keywords and weights from the user profile are retrieved. A *match–index* is calculated by forming the dot product of the frequency and weight vectors. When an article's match–index exceeds a threshold it is kept; otherwise it is discarded.

The architecture of the system, using a central database for software module coordination, allows several different user interfaces to be implemented. The current interface was designed to allow user feedback information to be easily recorded to allow profile adaptation. Figure 2 shows a typical screen. The top region shows subjects for which there are articles waiting to be read. Selecting a subject displays the headlines of the articles for that subject in the second region. Clicking on a headline displays the associated article in the third region. The interface allows multi-media articles to be displayed; it has the capability show raster graphics, animations, and video. Other capabilities allow the user to print, save, delete, and forward (to other PIGS users) articles, or search an article's text.

3.5 PIGS Prototype Experience

The prototype has been in use since June, 1989, by a small group of Bellcore users, including the authors. During that time, PIGS has handled about 800 AP newswire and 200 Netnews articles daily, ≈600 Telaris™ articles per week, and ≈150 Bellcore Technical Memorandum abstracts per month. Users have generally modified their profiles several times, and adjusted the match-index threshold to regulate the number of articles the filter passes according to individual preferences—generally discarding more than 95% of articles. With the current implementation, articles are entered into the service provider's database at a rate of 100s per minute, classified at a rate of about 40 per minute, transmitted to users at a rate of about 30 per minute, and filtered at about 30 articles per minute.

4. Network Implications of PIGS

4.1. PIGS architecture alternatives

The architecture described in section 3 is only one of several architectures that could be used to provide PIGS services using future BISDNs. The major degree of freedom in designing alternative architectures is the locus of filtering relative to the users and the service providers. Locating the filter in the users' PCs—Architecture A—(see Figure 3) provides privacy protection for each user's profile, but it requires that users' PCs have sufficient processing capacity to perform the filtering, and sufficient storage capacity to buffer all articles until the filter can process them.

Since the filtering will discard 90–99% of the articles, it seems wasteful to transmit all the articles to each user. In Architecture B, the filtering is moved to one or more filtering servers, co-located in local exchange switching offices[2], as Figure 4 shows. Because of this the per-user transmission, processing, and storage requirements are reduced, since the users' PCs receive only articles that have passed the filter, and filtering efficiency is increased, since the per-article preprocessing is done once per article instead of once per article per user.

Finally, in Architecture C the filter is moved all the way to the PIGS service providers, as shown in Figure 5. In this example, each user receives his or her own stream of filtered articles, requiring a separate call from the PIGS information provider to each user, and also requiring each article to be transmitted as many times as there are profiles that the article matched.

4.2. Network usage and capacity requirements

Analyzing exactly the full impact of PIGS-like services on the network is not now feasible due to the large number of unknown factors including BISDN architecture, structure, performance, service penetration, usage, etc. Nonetheless it is vital to assess the network implications of PIGS in order to be able to factor the estimated network capabilities and capacity requirements into current BISDN research, planning, and capacity engineering.

[1] Manual profile creation by the user could be eliminated altogether if, say, information vendors provide "seed" profiles for stereotypical users which then grow into individualized profiles through automatic adaptation.

[2] Perhaps implemented as an extension of today's "enhanced service provider" concept.

In this section, we present estimates of call complexity and transmission bandwidth requirements for PIGS services. Our approach is to use current newspaper and periodical statistics for the USA [US Bureau of the Census, 1989] as a basis for predicting future demand. We have used informal estimates of the number of articles, article sizes, and distribution of probabilities that articles will be read. These assumptions are summarized in Table 1. We further assume a call model [Minzer, 1989] in which a call can include multiple connections; each connection specifies communication over a single channel with one or more parties.

In each of the three architecture alternatives presented in §4.1, the PIGS service providers will use the network differently to place calls. In Architecture A, each service provider sets up a multicast call by placing a single call with a single connection with as many parties as subscribers (to that service provider). The service provider is assumed to choose, according to the tariffs in effect, between setting up each call for a short time (during transmission) or for a long time. In the latter case, the call will be modified, generally at monthly intervals, as subscribers are added or dropped. In Architecture B, the service providers set up multicast calls to the filtering services at the local exchanges, which in turn set up calls with connections to each user. These calls may be of two types: (1) a call to each user for a batch of articles; or (2) a call for each article to all the users who should receive that article. In Architecture C, service providers exhibit the same two call usage patterns as those from the filtering server to the users in Architecture B: a call per user (as shown in Figure 5), or a call per article. In the latter case, the service provider filters each article according to all users' profiles and generates the list of users to receive that article. The service provider then places a multicast call to only those users, transmits the article, and then disconnects the call.

Table 2 qualitatively summarizes the call usage parameters for the three alternatives. With user filtering each service provider places a single extremely large call. In Architecture B (local exchange filtering), each service provider has a single large call to the filter servers, which in turn have either large calls for each article, or a large number of single party calls. In Architecture C, the service providers either have ≈100 calls (one per article), or a large number of single party calls.

Along with call processing capacity, the BISDNs must have sufficient capacity for the article transmission. We can estimate the bandwidth requirements of Architectures A and B by assuming that the BISDN switching centers provide the multicast switching capability, where a single ATM input cell can be replicated, re-addressed, and output to several outgoing trunks. We then define the following transmission plan for all PIGS service providers:

> The articles from each PIGS service provider are transmitted directly to a Quaternary Center/Class 1 Office, from which a spanning tree for article transmission is defined so that each switching office in the tree receives only one copy of each article from its parent in the tree. That office then multicasts the article to each of its children in the spanning tree.[1]

With this plan, the downstream (from the Quaternary Switching Center/Class 1 Office toward the users) transmission capacity is at most the weighted sum of the load from each newspaper and periodical service provider. These calculations are summarized in Table 3, which shows that in the trunk groups that are carrying PIGS traffic, the average PIGS traffic is only 33% of a single STS-3C trunk. Because the transmission plan uses a spanning tree, but the actual BISDN transmission network is likely to be a more richly-connected graph, many of the trunk groups in the BISDN will carry no PIGS article traffic whatever.

The calculations of bandwidth requirements for architecture C with the call-per-article option are identical, since each article is transmitted from the service provider at most once. Since each article results in a call, each service provider makes ≈100 calls, each of which is a multicast to an average of 8 600 subscribers (see Table 4), during the publication interval. In Table 5 we show the calculations for the bandwidth required in architecture C with the call per user option. We calculated the bandwidth required just for daily newspapers (neglecting all periodicals and weekly newspapers); we show that the call per user option for this architecture requires the bandwidth of more than 2 000 STS-3c trunks, because the same set of articles (the top 20%) is separately transmitted to 80% of the subscribers.

4.3. Network Capacity and Service Architecture Trade-offs

There are many more architectures that could be used to implement PIGS than we have discussed, and we have made major simplifications in the architecture we presented. The alternatives generally involve trade-offs between network and service complexity. In this section we discuss two of these trade-offs, which involve network error rates and call complexity.

In our bandwidth estimates, we simplified the model by assuming error-free transmission and multicasting of the data for each article. The PIGS transmission software will have to use techniques such as forward error correction coding, reliable multicast transmission protocols, or a combination of techniques to eliminate the effects of cell loss, corruption, or mis-routing. These techniques even involve trade-offs of their own: bandwidth, processing speed, and storage requirements against robustness for coding alternatives; protocol throughput and complexity against upstream bandwidth of protocol overhead messages.

We also simplified our architecture presentation by assuming a one- or two-level tree structure of the multicast calls placed by the PIGS service providers. This resulted in single calls that involved 10 000–10 000 000 parties. Call size can be traded off against the number of calls and the processing for lots of small calls. For instance, a five-level hierarchy of processes, each placing a call to 26 lower-level processes (or users) could reach over 11 000 000 subscribers with 18 279 different calling processes.

The use of multiple calling processes to reduce and distribute call processing loads affords a synergistic solution to the problem of providing reliable transmission. By using a reliable multicast protocol between each process and its children (lower-level processes), the effects of cell loss would

[1] This transmission plan is not optimal in that it takes no advantage of locality effects for regional newspapers and periodicals, and sends all traffic to the top of the transmission network. The plan presented, however, affords simple analysis and can be used to bound the transmission bandwidth requirements.

be limited to just the children, and the error-multiplication observed with the single-call multicast tree would be eliminated. As attractive as this solution seems, it requires a change on current views of broadband networks. Few BISDN architectures have proposed including service-provider processes in the local exchange and higher-level trunk switching centers.

5. Conclusions

We have motivated and presented a service description for PIGS—a Passive (and Personalized!) Information Grazing System. We have described the architecture and performance of our PIGS prototype, demonstrating adequate performance and usability with conventional workstations. Finally, we have examined and compared network implications of several possible PIGS architectures. We conclude:

- If network transmission bandwidth is the only criterion, locating filtering in users' PCs or "close to" users, in conjunction with local exchange switching centers is preferable to service-provider based filtering.
- PIGS implementations may use very large calls—multicasting to thousands or millions of users.
- PIGS implementations may decrease call sizes by using multi-level call hierarchies.
- The transmission capacity of future BISDNs that have STS-3c members in each trunk group will likely be adequate for traffic equivalent to today's newspapers and periodicals, with adequate capacity for growth as multi-media articles become commonplace.

Future work in PIGS includes profile adaptation implementation, inclusion of more multi-media sources and articles, and further analysis of network implications of, and alternative architectures for, PIGS services.

References

Bussey, H.E., Porter, F.D., Raitz, J.E. (1988) A second generation BISDN prototype. International Journal of Digital and Analog Cabled Systems 1, 4 (October-December) J.E. Wiley, pp. 275-287.

Denning, P.J., (1982). Electronic Junk. *Communications ACM 23*, 3 (March), pp 163-165.

Hayward, G.A., Linnell, L.R., Mahoney, D., Smoot, L.S. (1987). A Broadband ISDN local access system using emerging-technology components. In *Proceedings 1987 International Switching Symposium* (Phoenix, Arizona, March 15-20), paper A8.1, pp. 597-601.

Hiltz, S.R., Turoff, M., (1983). Avoiding information overload in the electronic office. In *Proceedings of the Sixteenth Hawaii Conference on System Sciences*, (Honolulu, Hawaii, Jan. 5-7), University of Hawaii.

Malone, T.W., Grant, K.R., Turbak, F.A., (1986). The Information Lens: An intelligent system for information sharing in organizations. In *Proceedings of the CHI '86 Conference on Human Factors in Computer Systems* (Boston, Massachusetts, April 13-17), ACM, New York, pp 1-8.

Mackay, W.E., Malone, T.W., Crowston, K., Rao, R., Rosenblitt, D., Card, S.K. (1989). How do experienced Information Lens users use rules? In *Proceedings of the CHI '89 Conference on Human Factors in Computer Systems* (Austin, Texas, April 30-May 4), ACM, New York, pp. 211-216.

Minzer, S.E., Spears, D.R. (1989). New Directions in Signaling for Broadband ISDN. *IEEE Communications Magazine, 27*, 2 (February), pp. 6-14.

U.S. Bureau of the Census. (1989). Statistical Abstract of the United States: 1989. (109th edition.) Washington, D.C. pp. 549-550.

Table 1. Statistics and assumptions used in estimating bandwidth and call complexity for PIGS.

STATISTICS

Type	Frequency	Number
Newspapers	Daily	1745
	Weekly	7400
Periodicals	Weekly	880
	Fortnightly	619
	Monthly	4192
	bimonthly	1558
	quarterly	2245

ASSUMPTIONS

Distribution of the probability that an article is read	
% of articles	P(article read)
20%	0.80
30%	0.15
50%	0.05

AVERAGE PERIODICAL ARTICLE SIZE	
Component	Size
35 sq. in. of image, 90 000 pixels per sq.in., 3 octets per pixel	9450000
Text	50000
Total (octets)	9500000
Number of articles per issue	100

AVERAGE NEWSPAPER ARTICLE SIZE	
Component	Size
15 sq. in. of image, 40 000 pixels per sq.in., 0.75 octet per pixel	450000
Text	10000
Total (octets)	460000
Number of articles per issue	100

Table 2. Qualitative estimates of numbers of calls and call sizes for the different PIGS architecture alternatives considered. Architecture A is user-based filtering, Architecture B assumes filtering servers co-located with local exchange offices, and Architecture C assumes PIGS service providers perform the filtering.

Architecture	Caller (Option)	Number of calls per publication period	Number of parties per call
A	Service Provider	1	100s-10000000s
B	Service Provider	1	10s – 10000s
	Filter Server		
	(call per user)	100s – 10000s	1
	(call per article)	≈100	100s – 100000s
C	Service Provider		
	(call per user)	100s-10000000s	1
	(call per article)	≈100	100s – 10000000s

Table 3. Calculations to determine bandwidth requirements for PIGS service with sources estimated using current newspaper and periodical estimates. (All sizes are given in octets).

	Newspapers		
	Publication Period (days)	Number	daily equivalent
	1	1 700	1 700
	7	7 400	1 057
Newspaper daily equivalent total			2 757
Newspaper average articles per day			100
Newspaper average article size			460 000
Newspaper total size			1.3E+11

	Periodicals		
	Publication Period (days)	Number	daily equivalent
	7	800	114
	14	600	43
	30	4 200	140
	60	1 600	27
	90	2 700	30
Periodical daily equivalent total			354
Periodical average articles per day			100
Periodical average article size			9 500 000
Periodical average daily total size			3.4E+11
Total (newspaper and periodical) average daily size			4.6E+11
SONET and ATM overhead			20%
Total average daily size (bytes, with overhead)			5.6E+11
STS-3c daily capacity (bytes)			1.7E+12
Total average daily size (% of STS-3c capacity)			33.1%

Table 4. Calculation of call size (average number of parties in the call) for Architecture C, call per article option.

Number of Daily Newspapers	1 745
Circulation copies [US Bureau of the Census, 1989]	62 800 000
Average Circulation per Daily	35 989

Article Class	# in class	Call Size
read by 80%	20	28 791
read by 15%	30	5 398
read by 5%	50	1 799
Average Call Size		**8 277**

Table 5. Calculation of Bandwidth required for transmission of daily newspapers only in Architecture C, call per user option.

Number of Daily Newspapers	1 745
Circulation copies [US Bureau of the Census, 1989]	62 800 000
Average Circulation per Daily	35 989
Average Newspaper Article Size	460 000
Average Number of Articles per newspaper	100

Article Class	# in class	# recipients	Bandwidth
read by 80%	20	28 791	1.3E+12
read by 15%	30	5 398	2.5E+11
read by 5%	50	1 799	8.3E+10
Bandwidth for a single newspaper			1.7E+12
Total Bandwidth for all daily newspapers			2.9E+15
Total Bandwidth including SONET & ATM overhead			3.5E+15
Total Daily STS-3c Bandwidth			1.7E+12
# STS-3C trunks required			**2063.9**

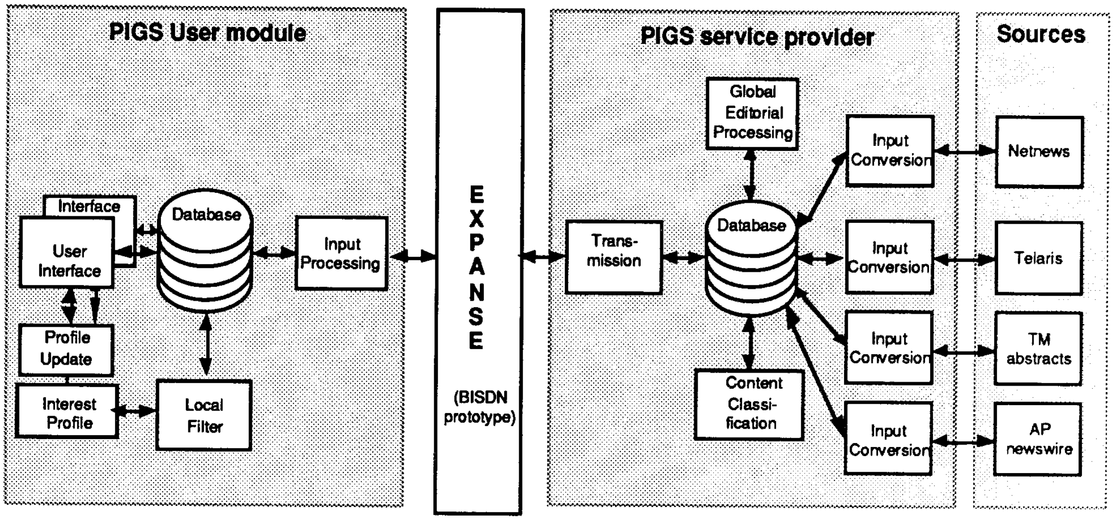

Figure 1. PIGS prototype software architecture

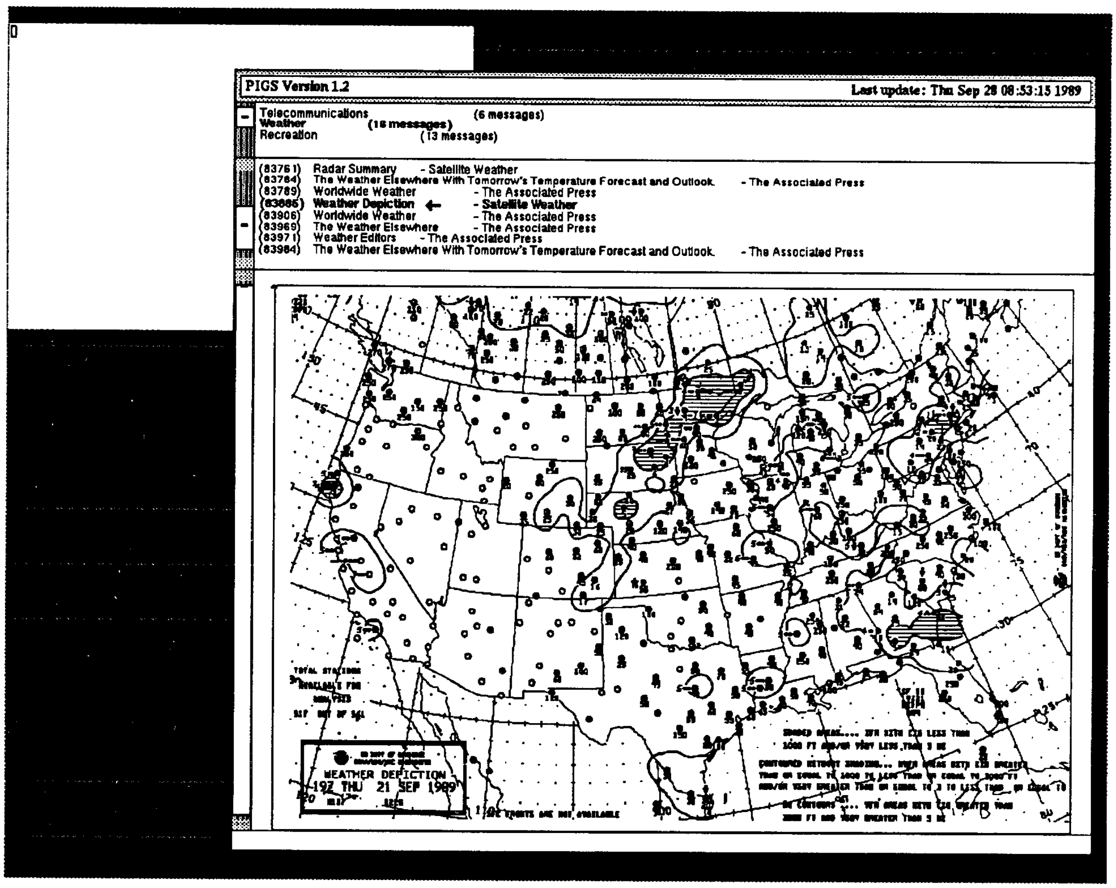

Figure 2. PIGS user interface, showing the subjects, titles for subject "Weather", and a NOAA weather map.

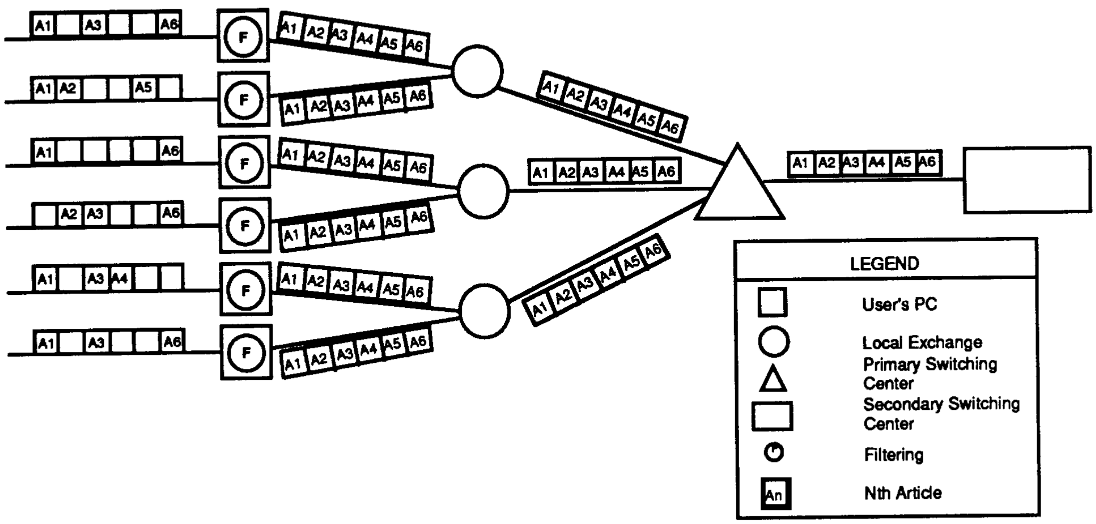

Figure 3. System/service architecture option A.

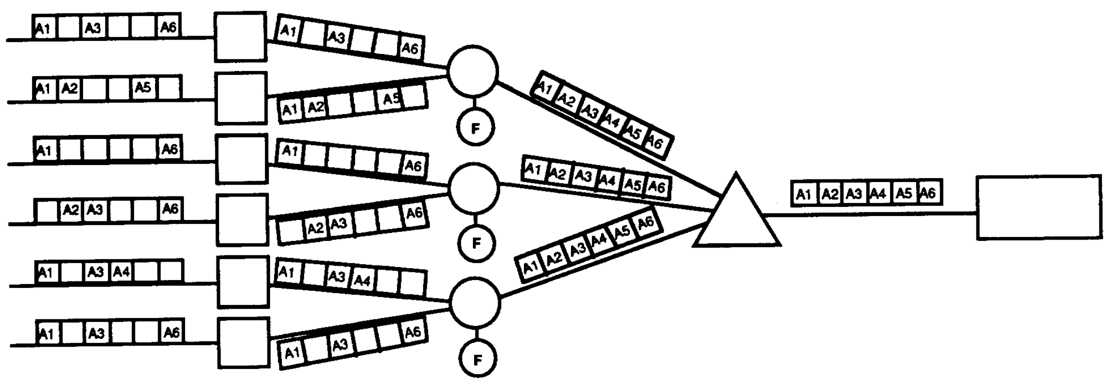
Figure 4. System/service architecture option B.

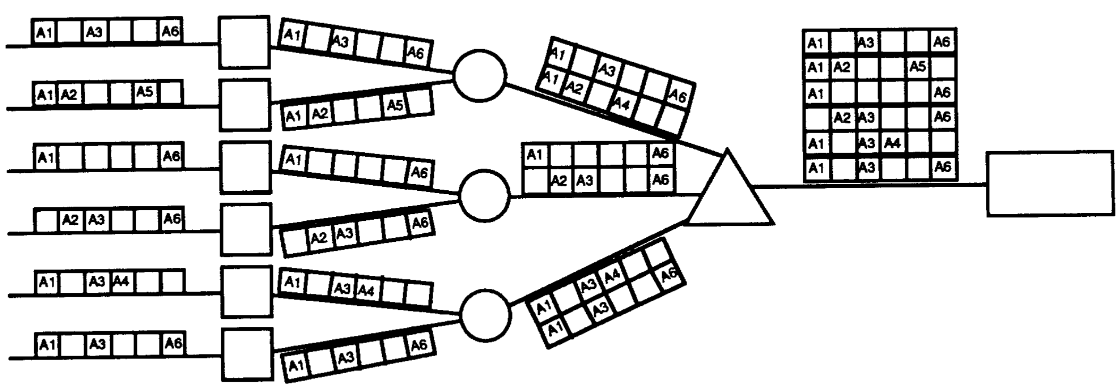
Figure 5. System/service architecture option C.