SMASH – a concept for advanced use of storage in the home

Eric H.J. Persoon *

Abstract

SMASH stands for Storage for Multimedia Applications and Systems in the Home. It is the subject of a European research program as part of the ACTS initiative. The main goal of the project is to explore the future possibilities offered on the one hand by the massive amount of new digital services that enter the consumers home and on the other hand the impressive progress of storage technologies that are suitable for home usage. In this paper we will describe our views on how home storage can be of benefit to future consumers. Although the described ideas could be used in more environments than only the home setting, we have limited ourselves to this environment since it offers many challenges and it is the most demanding setting regarding the variety of users that must be satisfied. The paper describes first the user needs for local storage and the available digital sources that provide the information to be stored. Then the current status of available and expected digital storage technologies for home use are explained. This is followed by the overall architecture of the SMASH system we have realized so far and a description of its basic operations. Finally we end with a chapter that describes on how content and service providers may enhance the usage of local storage at home.

1 - The need for local storage

In many publications during the past years it is often expressed that the ultimate wish of the user would be to have "everything on line". An example of five years ago was the massive effort devoted to "video on demand" systems which were expected to make both video rental shops and video recorders obsolete. It turned out to be quite different. It is important that lessons are learnt from this. One of the lessons is that there should be a good compromise between "how much is technologically possible" and "the price a customer is willing to pay for the service". Nowadays, the Internet seems to promise all that the consumer is waiting for and many people expect that, once the Internet can provide sufficient bandwidth, video on demand will emerge again. It is true that the ever larger amount of information the user will have access to, be it via digital satellite, digital cable or telephone wires will make the necessity of local storage less important. But, there are several fundamental reasons why we think that the amount of local storage at the users home will rather increase than decrease.

The first fact is that the pattern of usage of information by the user at home is very much peaked. This can easily be seen from the "viewing statistics" of people watching TV programs. The hours between 8:00 p.m. and 11:00 p.m. are the most important hours to satisfy the "entertainment and infotainment" demands of consumers. These are simply the hours they are at home and have time for this. Some people have such busy jobs that only the weekend is available for TV watching or Internet surfing. The

* mailing address:
Philips Research Laboratories
Storage & Retrieval (WY-21)
Prof. Holstlaan 4
5656 AA Eindhoven, The Netherlands

variety of the wishes from people differentiated by their different cultural backgrounds, different general interests and different age groups is so large that it is almost impossible to give people what they really want at the right moment. This is where local storage can fulfill an important role. There at 24 hours in a day and it must be possible to use some of the less important hours to download information requested by a user for later viewing. When a VCR is used for this purpose it is not possible to watch a recorded program while recording another one. We will see later how in SMASH this restriction is removed.

Another aspect is that people want to use their time in an optimal way. If one looks at the current broadcast situation it happens that users miss the viewing of a nice program just because they didn't spot it in the TV magazine. Also at some evenings there are interesting programs at several TV stations simultaneously and at other evenings there is really nothing interesting. Our view of the desire of most consumers would be the following situation: each service provider distributes a short description of each program well in advance (only a few hours before the broadcast is sufficient) of the actual broadcast and the storage system at home compares this with an interest profile of the user. The local storage system has sufficient free storage space to record several programs during one day. When the user wants to watch TV he gets on his screen an overview of all on air channels but also of all the recently recorded programs. The user can now get a fast overview of all the recorded material. He has then the following choices: decide that a recorded program is not interesting and delete it, watch a program fully (like a movie) or watch only the interesting parts of a program (in case of talk shows, etc..). This concept gives the user complete control over the time he spends at watching programs. It is clear that to realize this concept not only a lot of storage is required at home but moreover several other technologies must be in place such as intelligent filtering based on user interest profiles and new retrieval technologies for the fast retrieval of interesting information in the stored programs.

When using information from Web sites there are several reasons why local storage can be of advantage. One of the most obvious ones is to have (fast) access regardless of the network load. But it is also useful for keeping certain Web information for later consultation. At Web sites one observes that preference is given to most recent information. For example elections are covered at the moment they are held but after some time this information is not available anymore on the Web. But not only elections, also other interesting information may be deleted unexpectedly from a Web site by the service provider. Some people may be interested to keep such information for later usage.

Besides the above mentioned reasons for storage the user at home will always have the desire to keep a record of some important TV broadcasts or other delivered information for longer storage. Maybe he or she will never watch those programs again but it is a natural desire to keep things. Without any doubt, as more digital information will enter the home, the consequence is that users will also want to keep more of this information.

2 - The digital sources

To a large extent storage of digital information in the home will be dependent on information entering the home in a digital way. The amount of information a user creates by himself is only a very small fragment of what he consumes. There are several technologies available now to carry digital information into the homes each one

however at very different speeds and costs. The digital satellite receiver, the digital cable modem, ISDN lines and the telephone modems are the most used ones. The digital satellite system is by far the most cost effective way of distributing information to many people and the telephone modem the most expensive way especially over long distances. But not only cost, also bandwidth is an issue. A good quality TV program requires at least 3 to 4 Mbps (Megabits per second) of transfer rate. It will take a long time before such speeds are offered through the copper telephone wires in a cost effective way.

Besides cost of connection and speed, very important are also the standards that are used to distribute the information. This is especially important for consumer equipment which is very cost sensitive. It is very difficult to make cost effective implementations if no standards exist. In principle one could for every service download first the application software such as is currently done with Web browsers. However for high speed real time services this would require expensive software programmable decoding hardware. At this moment in the SMASH project we are concentrating on the standards which are agreed and being developed in DVB and DAVIC.

An important standard for a SMASH device is a standard that specifies how the information in a digital source is transmitted in a serial fashion. We have found that the current MPEG-TS (Transport Stream) concept is very suitable as a general framework to distribute and store real time services. This is not only because it allows for the packetized transmission of audio and video but also for the simultaneous broadcast of related information such as teletext and program specific SI information. Currently we are investigating how the SI information could be extended for the purpose of storage at home. However, to make the TS system suitable for storage, we had to work in the past on the "partial TS" concept. This is needed because the original DVB TS concept was specified for the transmission over a satellite, cable or terrestrial channel which could carry bit rates up to 40 Mbps and carry several programs. For local storage those are too high bit rates and it is necessary to first demultiplex the complete stream and then make a selection in the available programs. The resulting stream is called partial TS because it contains only part of the original multiplex.

This leads us to another topic, although we will not cover it in a lot of detail here. A SMASH storage device should be connectable in a digital fashion to all kinds of other equipment in the home. For example the STB for both the reception but also the display of MPEG programs, but to other A/V equipment and the PC. It is important that the necessary digital interface is well constructed, standardized and accepted by many peripheral makers. Such an interface will have to carry, amongst others, partial MPEG TS streams as the preferred standard of communicating Audio and Video to and from the storage device.

An important other topic is the way the service is delivered to the user. Will it be scrambled or will it be copy protected? Many digital broadcast services are transmitted scrambled. In principle this is not a problem for the SMASH device because it can easily store scrambled information. The problem arises however when the user wants to retrieve the information for viewing. In section 5 it is explained how we want to ease retrieval by presenting the user with "visual abstracts" of the stored material. The SMASH device can only do this if the received information is not scrambled. But there is a more down to earth problem. Most scrambled signals must be played back through a conditional access system (present in the

STB) that is time sensitive. This means that after some hours or days the information will not be decodable anymore by the descrambler. It is obvious that broadcasters will have to adopt a new or modified conditional access system if the recording of scrambled information is desired. We think that this should be made possible. It would allow a new way of distributing content to a large community. The service provider can broadcast in parallel for example all the recent movies to all people having a storage device at home. This is done in a scrambled way and free of charge to the user. Only when a user decides to watch a stored movie will he have to pay for decoding the movie. This offers by the way a very cost effective realization of a video on demand system where the service provider does not have to invest in expensive video server technology.

A last point is the speed of the service. Already was mentioned that 40 Mbps is too high for many storage devices to handle. In section 5 a proposal is formulated for the maximum transfer bit rate of a service.

3 - The storage technologies

There are two basic ways in which information can be stored: magnetically and optically. Besides that there is the possibility to either store information on a tape or a disk. Both the magnetic and optical storage density has grown dramatically over the past years. The most obvious one is the magnetic hard disk present in every computer. Most people may remember that the first PC's came with a 20 Mbytes hard disk drive. At this moment a 2 Gigabyte hard disc is regarded as the minimum one should buy with PC's. The preferred way of storing information is on discs because this enables fast access times. However if one wants to store a lot of information then magnetic tape has certainly cost advantages.

There are several types of magnetic tapes in the market and it is not the purpose of this paper to give an in depth overview. Currently the largest capacity which can be stored on a single cassette is around 50 Gigabytes. And there are possibilities to still increase this capacity significantly. Optical tapes are non-existing for consumer use.

Magnetic disks come in two kinds, the hard disk drive which is non removable and the removable ones. Hard disk capacities will keep rising and currently a 4 Gigabyte hard drive is common. Removable magnetic disks have less capacity amongst others because they have only a single storage platter whereas most hard disc drives have 4 platters for storage. Storage capacities for removable discs of up to one Gigabyte are achievable already.

Optical discs which are rewritable have capacities at this moment of 3 Gigabytes per side. This capacity is expected to grow also over the next years to 10 Gigabytes and in the more distant future to even higher capacities.

From this short overview it can easily be seen that a storage capacity at the home of one Terabyte is not at all impossible over a couple of years. It means at this moment 20 tape cassettes. When optical discs will contain 10 Gigabytes per disk, a disc changer containing 100 disks also amounts to 1 Terabyte. To have a reference how much this all is, a good quality movie of 2 hours requires around 4 Gigabytes of storage. So 1 Terabyte enables the storage of 250 movies (each of 2 hours) or in total 500 hours of recording.

4 - The architecture of SMASH

How should a home storage system look like? We will begin with an analysis on how the user may want to use such a system.

The first desirable element would be that, in contrast to current VCR's, the user

would not have to program explicitly each program he wants to get stored. This requires the analysis of the received information and the comparison of this with a prespecified interest profile of the user. Of course the user is still entitled to program specific broadcasts to be stored and those will be given preference. To make this possible some intelligence needs to be built in the SMASH device.

Once the programs are stored the user will want to have a quick overview of the stored material and moreover, inside a stored program he would like, much as is the case in a VCR, to scan fast in a forward or reverse manner in order to locate interesting parts and skip non-interesting material. To be able to perform this task, the storage device will have to make a "condensed" version of the recorded material. This should be done at the same time the recording is done so that, immediately thereafter (or even while the recording is happening) the user can search through the material. This makes it necessary to analyze the incoming video and audio in real time and make intelligent abstracts of this. In the SMASH project, some quite interesting results have been made in this area.

The diagram below sketches the basic architecture. You can see that it consists of two storage elements: the Hard Disc Drive and the Mass Storage part. The HDD is partitioned in two parts: one part contains the Data Base and the other part a pool of temporary buffers.

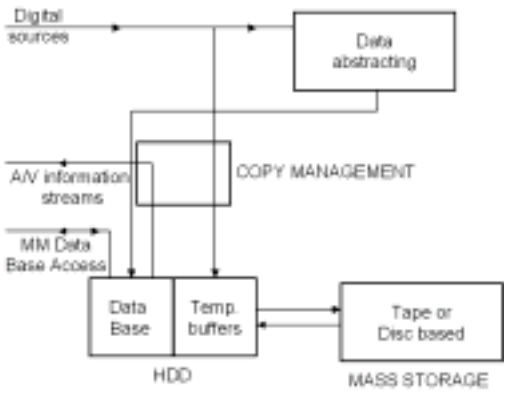

The SMASH device obtains its main information input from the digital sources. In our system we derive them either from a DVB Set Top Box or from Internet. The information which enters the SMASH device (digital sources) is sent directly to the hard disc where it is buffered on one of the available buffer spaces. At the same time this information is sent to a "data abstracting" module. The task of this module is to extract in real time information which later will be used in the Data Base. In our system we use for this purpose single images extracted from the incoming video stream, information out of teletext and information provided by the DVB - SI stream. On the hard disk there is a special data base system implemented which can be separately accessed by the user (Data Base access). The pool of temporary buffers on the hard disk can be used for several purposes: either for the reading or writing to the tape device, but also buffers which can implement in real time a pause function for incoming streams. In the current experimental system we use a fast 2 Gigabyte hard disk and a digital tape drive with a transfer rate of 12 Mbps.

Part of the SMASH system must be concerned with copy management. It operates both on the incoming stream as well as on the outgoing streams. This operation is implemented in its simplest form by a few copy bits which are associated with each program. We are investigating also more robust techniques which use watermarking on the received MPEG bit streams.

The output connection "A/V information streams" can contain one or more simultaneous Audio Video streams and is usually connected to MPEG decoders to convert the compressed streams to baseband signals.

The connection "MM Data Base Access" is bidirectional. It communicates with the application device (e.g. a STB) and allows for

the retrieval of information from the data base. But it is also used for example to delete files from the Disk or Tape or to edit the data base directly from the application.

5 - The basic operations

Recording

For recording the received data stream, it is first sent to the hard disk where there is a sufficiently large circular buffer present. When the buffer is nearly full transfer of the data to the mass storage device will be started. The transfer of the data to the hard disk happens at the rate it is received but the transfer to the mass storage device preferably happens at the maximum bit rate possible. Moreover for tape drives, the transfer to the tape drive should be with constant bit rate since most tape drives do not support variable transfer rates. For digital consumer tape drives the maximum transfer speed is at this moment 14 Mbps (for D-VHS) and 25 Mbps (for DV). Higher bit rates are possible but most current products (Standard Mode) will support the above bit rates. Since those transfer rates are quite high the SMASH device may in principle support the recording of two or three programs in parallel provided that the bit rate of each program is not too high. For current DVB transmissions broadcasters use between 4 and 5 Mbps as a transmission rate for each program. The recording of two programs at the same time may be an attractive feature since, especially during peak evening hours, it often occurs that interesting programs are broadcasted at the same time. The speed of the hard disc drive will be in general not an issue since the data rates are already quite high and will increase in the future. However the speed of the mass storage devices may put restrictions on the recording speed. In the DVB project a guideline has been proposed to limit the total bandwidth of any single program to 9 Mbps. This seems to be a bit rate which is quite acceptable for any digital video broadcast and which is currently not exceeded by any DVB broadcaster in Europe as far as we know. The 9 Mbps maximum limit has been accepted by the DVB group as a guideline for implementation.

Retrieval

The retrieval task consists of two parts. The retrieval of information regarding which programs have been stored and the retrieval of information inside a chosen program.

The retrieval of information on the stored programs is organized using several sources of information. Preferably we use the SI data which can give details of a program such as the title, genre, etc. But also teletext offers a source of information where some teletext pages give more detail information of the program being broadcasted. We are also looking at program information which is available at the Web sites of several broadcasters although this is not a preferred solution. The information can also be extended by short video abstracts of the program. All this information is stored in a database together with the information where this program is stored (tape number and tape position).

The retrieval of information inside a program is a quite different task. It can be compared with the search functions on a VCR. However in SMASH we prefer that the search functions are not done on the tape drive but on the hard disk. The hard disc offers much more flexibility and also higher virtual search speeds. In order to do this, we need to analyze the incoming A/V streams and must identify episodes inside a given program. This is done based on the concept of key frames and intelligent clustering of key frames into larger logical units. All this information is also stored in the database again accompanied by the location where the corresponding information is located on tape.

Playback

The SMASH device can, due to its high bandwidth, support in principle the playback of more than one program at the same time. This can be useful when it is used by more than one person at the same time and as such it can function as a home server. But it can also playback one program and at the same time record another one provided that the sum of the bit rates of the two programs is less than the maximum transfer between the disk and the tape. This could be a useful feature compared to the possibilities of current VCR's. A VCR can only record or playback at any given moment and cannot perform both functions simultaneously.

6 - On the organization of information . The role of the service provider

If we believe in the advantageous role of local storage then the people who provide the digital services should start thinking of organizing their information accordingly. The two main current examples of information services are Radio/TV broadcasts and Web sites. The first services are organized for linear non-interactive use and have to be planned at least one week in advance to be able to print all the TV magazines. Very limited support is given by those service providers to support the role of local storage. The WWW providers organize their information for interactive use but again they do not take yet advantage of local storage possibilities.

How can information be better organized by the service providers to support local storage? In one of the previous chapters it has been explained that a possible interesting concept for SMASH is the feature of "interest profile recording". To make this possible the SMASH device must be able to select amongst the broad offering of programs only those that are of interest to the specific user. This means that each program must be accompanied by a suitable description. It is not sure how extensive this description must be. It is obvious that, for example for movies, just describing it as "a movie" is not sufficient since then every movie will have to be stored. Together with each movie attributes such as movie type, main actor(s) and year of production will probably also be required. In the SMASH project we are investigating what type of information is needed for each program. It must be clear that the SMASH device need not always to make 100 percent correct decisions since the user can easily decide not to watch a recorded program. However not too many programs should be recorded and there is another restriction: depending on the data transfer rate of the SMASH device, only 2 or 3 programs may be recorded in parallel making it necessary to make a selection in case more than 2 or 3 programs fit the interest profile of the user. Once the program has been recorded, the user is confronted with finding back the stored information. As explained before, there he has two tasks: retrieve a program and finding inside a program interesting parts. For the first task, the description which the broadcaster has prepared can be very useful since this can be stored on the hard disk for easy retrieval.

But to enable the second task, extra information is needed about the structure of a program or possible highlights in a program. For different types of programs this means different things. Just one example: in case of a talk show it will be interesting to know which people are speaking at which moment and what subjects are being discussed at what moment so that the user can easily jump to the section that interests him. This type of extra information can be supplied for free but digital broadcast offers the possibility to send this information encrypted and make it only

available to the user for a fee. All this extra information should preferably be simulcasted on the same transponder as on which the actual broadcast will happen so that no extra receivers or tuners are necessary to receive this information.

Another area of interest is the information which is available from Internet. At this moment Internet is not so attractive for consumers when a lot of pictorial information or video has to be downloaded. It is predicted by many people that the speed of Internet will continue to grow but only when the transfer rate will go up to more than 1 Mbps will it become attractive to view and download movies (be it at not optimal quality). However it can be expected that for many services the speed of downloading will remain a very big problem for a long time, since this also depends on the power of the Web server and the amount of parallel hits it can support. Therefore also for Internet services we think that local storage may offer some great benefits. Internet is however used for many on line services and users expect that the information they receive is up to date. When using local storage, advantage can only be taken when information has been downloaded ahead of time. This can easily make the locally stored information out of date. It would be nice if Internet services could construct their information in a way which easily supports local storage. When examining the build up of a typical Web site the amount of new information which appears each day is only a minor fraction of the total information available of the Web site. Even inside a given a Web page photographs may stay the same but only prices may change. Advantage can be taken of this fact. Suppose a user works a lot during the day with the information on some specific Web sites. He could instruct his computer to download each day the complete Web site during say the early morning hours. It becomes important however that during the next day, when he will contact the Web site information, he will know which information is still valid and which of the Web pages have been changed since the last update. Internet protocols provide the possibility to perform such tasks but the information on a Web site must also be organized accordingly. Currently to perform this task one must scan the complete Web site and check the dates of each page. It could be more convenient if the Web site provides a separate file which contains the list of pages that have been modified, added or deleted since for example the day before. If the content of a Web site may become too large to download it every day, it is conceivable to make also weekly downloads possible and that each day only the updates which have happened during that week are updated on the local storage system.

7 – Conclusion

It must be obvious that the combination of sufficient local storage and the support thereof by service providers will be a big benefit for users in the future. Storage devices with sufficient capacity and at sufficient low prices will appear on the market in the near future. They will offer capacities from 10 to 50 Gigabyte on a single carrier (magnetic tape or optical disc) and a collection of 20 tapes (or 100 discs) offers a Terabyte of storage in the home. Several interesting new ways have been presented on how this can be used and it has been explained how the service providers can provide extra information to get the most benefits of the system. We hope that service providers start to consider now the new possibilities we are offered by low cost digital storage at home and that they start thinking on how to organize their services to take full advantage of this. In the end the consumer will appreciate the nice possibilities offered to him. To conclude this paper I want to

acknowledge the work of the many people who are contributing to the SMASH project from the different participating institutes. More information regarding the partners of the project but also references to other SMASH publications can be found at the following Web address:
http://www.extra.research.philips.com/euprojects/smash

presentation at IMAGINA'98, March 4-6, 1998