Inside iTV- Interactive Commerce Services for Broadcast Digital Television

R. Chernock, P. Dettori, F. Schaffa, and D. Seidman

IBM T.J. Watson Research Center

Abstract Interactive television (iTV), i.e. the enhancement of the television viewing experience by the addition of multimedia content and services beyond the basic video and audio playout of a program, has a great potential for commerce, education and entertainment. By enabling the delivery to the home of rich multimedia content in coordination with broadcast television content, iTV can serve as a powerful platform for widespread home shopping and distance learning, among many other applications. The advent of high-bandwidth channels into the home, such as satellite, terrestrial, two-way digital cable and DSL, enhances iTV's power and the prospects for its acceptance. iTV has the potential to move from a curiosity to a common television enhancement, possibly with the ubiquity of the VCR.

Despite this, iTV has failed, to date, to fulfil its promise. While many fully mature technologies have taken years to find wide acceptance, in the case of iTV some issues can be isolated which can explain its slow approach. Before commercial iTV applications can gain wide acceptance, elements of the iTV network infrastructure must mature from their current state. Among these elements are those relating to the adoption of open standards and the enablement of purchase security and purchase tracking capabilities. For it to become commercially viable for network operators to offer iTV purchasing services to their subscribers, the capability must exist for every purchase transacted by a user to be fully secure and private, to be accounted for, and to be traceable. TV-based systems must be robust, user-friendly, and have reasonable response times.

In this paper, we present an architecture and a suite of commerce-related services for iTV networks designed to fill in some of the gaps in the current iTV software infrastructure, based on our experience in prototyping iTV systems. These services are of three types: operator services, advertiser services, and user services. After a brief review of current iTV commerce

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35584-9_19](https://doi.org/10.1007/978-0-387-35584-9_19)

systems, our proposed approach is described in detail. Finally we relate the details of one implementation of this architecture and these services in a prototype iTV environment. This successful proof of concept instantiated an iTV system with advanced commerce services as developed in the proposed architecture.

Keywords: iTV, Commerce Services, Controller, STB, Broadcast, MPEG-2, Transactions.

1. INTRODUCTION

The broadcasting industry is in the middle of a transition from analog to digital television. In addition to better video and audio quality, digital television allows carrying more standard definition programs in the same bandwidth spectrum, and provides the basic infrastructure for the deployment of new interactive services. There are a number of digital television standards that apply to different geographies, such as DVB [5] [18] in Europe, and ATSC [3] [4] and OpenCable in USA. A common feature among all the different instances of digital television systems is the capability to transport basic audio, video and data services and signalling information on the MPEG-2 Transport protocol [6]. Data services are the common layer which iTV services are built upon.

There are several embodiments of iTV services that require different system capabilities and that can be classified as broadcast-only services or services requiring a two-way communication channel. In the former case, interactive information is injected into the digital broadcast stream, detected and decoded by the Set-Top Box (STB), allowing the user to interact with content in the broadcast. The latter case adds the capability to perform transactions over a two-way network with back end servers, such as retrieving bank account information or sending back a purchase order.

While the deployment of basic digital television services is progressing quite smoothly, the deployment of iTV services has been slowed by a multiplicity of factors. In this paper we will focus on issues related to iTV commerce services, propose an architecture that could accelerate their deployment and present a prototype implementation.

2. ITV COMMERCE SYSTEMS

The iTV commerce infrastructure is built upon a set of building blocks, some of which are defined by standards and some others which are still proprietary, but migrating to open standards. A digital TV head end supports

the insertion of data into the MPEG-2 transport stream. The broadcast network carries the data to a digital STB, which in turn decodes the data and provides users with the capability to interact with the content. Transactions are routed by the network to back end servers, which will process information and will return data either via the broadcast channel or a two-way communication channel (phone line or two-way cable.)

In the past few years several iTV systems relying on proprietary back end systems have been deployed, producing expensive systems that are difficult to maintain and not open to multiple content and service providers. Today the focus is shifting to leveraging as much as possible the Internet infrastructure for interactive and commerce related services, bringing benefits that range from lower costs to more openness and interoperability to better capability to attract merchants and content providers.

The building blocks of e-commerce are the technologies of the World Wide Web, its protocols, servers and clients, along with the application development structures such as Java and HTML. Although the introduction of Internet standards and technologies has brought great advances to iTV, a number of issues are still open and slowing the adoption of iTV services. We will discuss some of these issues and proposed solutions in the following sections.

2.1 Current Scenario

The creation of iTV content today is still an expensive and time consuming task, and the general philosophy that should inspire it is "create once, re-use everywhere and many times". The creation of iTV content involves the creation of graphics assets, and the embedding of links either to self-contained content (e.g. content sent also in the broadcast as part of the ad) or to online servers (e.g. commerce servers.) The basic stages for content creation and distribution in an iTV system are authoring, embedding and distribution, ultimately to the viewer with a STB.

Once created, the ad will be delivered to the Broadcaster, together with a "metafile" describing the association of the interactive elements with the video/audio content, where it will be embedded according to a standard (for example, see the SMPTE standard for TV triggers [20]) or a proprietary mechanism. The resulting digital broadcast stream emitted by the broadcaster will include audio, video and data components that will contain the ad's interactive elements, with proper timing relationships with respect to the audio and video components. The STB will receive the regular video and audio program elements plus the interactive data. A STB with proper hardware and software capability will be able to process the data properly and present it in coordination with the video on the TV screen.

The user will interact with the STB to display text and graphics on the screen by means of a remote control. After receiving the initial data from the broadcast channel, the interaction may require loading more data either from the broadcast stream or from a two-way channel. If the user decides to purchase some item presented on TV, he or she will just follow a simple interaction path that typically would require some form of authentication by means of a numeric ID and a PIN or a smart card and a PIN. After authentication, the user will either enter payment and address information (i.e. credit card number) or will access information stored in the STB, in the smart card or in a remote wallet server to finalize the transaction. Finally the purchase request will be sent by the STB to the commerce server, that will process it and will send a confirmation number back to the STB. In some cases an e-mail confirmation with order details can also be sent, either directly to the STB or to some other user's e-mail account. The same flow applies both to physical and digital goods delivery (e.g. purchase of a music CD), but in the latter case the fulfilment happens over the network, by securely delivering the digital content to the STB hard-drive or to some other device (such a MP3 player) directly connected to the STB.

Figure 1 illustrates the current scenario for iTV commerce.

Figure 1. Current scenario for iTV commerce

As mentioned earlier, although this scenario is currently being deployed in some geographies, there are a number of issues that are slowing down its adoption in the market. We will outline these issues in the next section.

2.2 Issues

Network operators need to develop business models and agreements with advertisers and merchants using their networks for iTV commerce services. There is no mechanism today to enforce business rules among the different entities involved in the iTV commerce chain.

iTV content authoring is still an expensive process and ideally the same ad campaign should be re-used at different times in different markets. The same ad campaign may need to be linked to different merchants on different geographies. In the current scenario, users are typically directed to a specific commerce server that does not change over time or with the geography. This would require broadcasters or cable operator to have some mechanism for redirecting commerce related traffic to local merchants that would fulfil the order, or to dynamically update the interactive content related to an ad.

A chief component of the iTV revenue stream comes from advertising. Advertisers are moving from traditional TV ads to more sophisticated ads, and they require ways to track the popularity of an ad. Furthermore, methods to tune interactive ads would be also required. Again, there are no standard mechanisms today to track the lifecycle of an interactive ad and gather statistics on how and how many users interacted with it.

Statistics on commercial transactions need also to be gathered by the merchants and cable operators and reconciled in such a way that each party get its share of the revenue stream.

Finally, there are a number of user requirements to be satisfied, such as providing the capability to securely store information on the network, such as profile information, links to iTV ads and commerce offers. Typical STBs have very limited storage capability, except for those few STBs with hard-drives, making it difficult to utilize local storage for this purpose. Furthermore, links must be maintained and accurately updated. Everyone would agree that a TV user would not accept being presented with a page containing some arcane HTTP code for "page not found", as happens today for PC users on the Internet. Thus a mechanism to store and check consistency of links must be provided in the network.

2.3 Proposed Architecture

We have designed an architecture that capitalizes on the current iTV infrastructure and plugs in a new element that will allow a smooth

introduction of iTV services. A controller is introduced for the interactive television commerce environment, acting as a central point of control and information for enabling and facilitating information distribution and transactions among the entities involved in the iTV commerce value chain. The controller offers a low-cost solution for the introduction of iTV commerce services, in contrast with alternative systems, which involve the redesign of the network's IT infrastructure. It is designed to be retrofitted into an existing network.

The controller allows network operators to keep track of all flows in the network transactions that take place over their networks between clients and any servers.

The controller manages, adds to, and modifies the content of interactive ads and programs that are broadcast to their viewing audiences. It supports redirection of network traffic to the appropriate servers for local fulfilment of goods and services which are advertised in nationally broadcast ads, and which are purchased through interactivity with those ads.

The controller reports to network operators, advertisers, and viewers, on purchasing activities that are conducted on iTV networks, based on consensual information gathering and analysis of this data.

Figure 2 illustrates our approach. The controller is inserted in the digital TV network, and it is typically located in the cable operator's head end.

An additional element, the Service Control Point (SCP) communicates with multiple controllers to centrally manage them. The SCP updates database tables in the controllers, gather aggregate statistics and enables a single point of control and interface for distributed controllers.

Figure 2. Proposed scenario for iTV commerce

3. SUITE OF ITV COMMERCE SERVICES

We will describe in this section the benefits and services supported by the architecture we discussed in section 2.3.

3.1 Enablement of Interactivity

Interactive content, in the typical scenario considered by this paper, is added to the basic video and audio content prior to broadcast or at broadcast time. However, it may be useful to selectively enable the utilization of interactive content by client applications, for example if an interactive commercial is aired after the time period in which the added content is valid, such as a limited-time offer. This selective enablement can be located at the controller (see section 2.3) or at each individual set-top box.

3.2 Transaction Capture

Typically, some display of multimedia content to the viewer during a program or commercial begins (and motivates) the interaction process.

If (as is the case with purchases) the interactivity proceeds to the point where connection to a server is necessary, this connection is facilitated by use of URIs (Universal Resource Identifiers) [17] that are included in the added content. At this point, the interaction could proceed between the STB and the server; however, this would leave the network operator out of the loop, not being aware of the type and amount of commerce traffic being handled on the network.

It is necessary for the controller to intermediate between the STB and the server in order for the network operator to be kept in this loop. Information on the items being purchased, their quantity, price, etc., can be extracted from the HTTP messages moving between the STB and the server, subject to privacy and security considerations (see sections 3.7 and 3.8.) Each transaction conducted over the network should be logged by the controller and made accessible (in the aggregate) to the network operator, who may receive a fraction of the profit. Statistics may be generated for the individual or in the aggregate, regarding purchase patterns, interests, etc. Again, this is subject to the approval of the clients, and must align with their privacy and security concerns.

3.3 URI Translation

As described above, an interactive TV advertisement is created by a content creator, typically the same entity which created the audio and video content. Alternatively, the interactive content may be added to a broadcast program. The enriched content (audio, video, and data) is broadcast and received by the client's STB that is running an application to detect and extract this data. Typical data includes URIs, text describing sale items, graphics such as icons, and "metadata", data that describes the data that has been added to the broadcast. The advertisement may be broadcast nationally or locally.

A service control point (SCP) (see section 2.3) has prepared the controllers, which reside in the head ends of network operators offering TV e-commerce services, to prepare for the broadcast. "Redirect URIs" are specified by the SCP and stored in the controller, in order to map the URI from the broadcast ad to the appropriate local server. When the user interacts with the advertisement (typically by pressing a key on the remote control), the controller recognizes the advertisement (by use of a unique ad identifier) and maps it to a local server, as specified by the SCP.

The transaction then proceeds between the client and the local server, with the controller intermediating to record user activity and transaction details. This mechanism allows localization of interactive content for different markets and spreading the transaction loads between different servers.

3.4 Activity Tracking

At any time, the client may request a review of his/her purchase activity. The controller maintains records for all clients' purchases in the database, and responds with a display of all purchases for this client that were transacted over this network. The client may also request the status of a given purchase. The controller then assists in the tracking of a purchase, based on the confirmation number, which it maintains for each purchase, and the purchase tracking mechanisms available for the item.

3.5 Bookmaking Purchase Opportunities

Because of the dynamic nature of television, purchase opportunities may be seen by clients who wish to interact with them but are unable to do so in the time allotted. The controller enables the client to set a "bookmark" during an iTV advertisement to indicate an interest in the advertised item. Each client's bookmarks are stored at the STB or at the head end. The client may request to be reminded, periodically, of the bookmarks he/she has set by way of an icon or audio beep. The controller takes care of purging bookmarks when the purchase opportunities they reference expire.

3.6 Offline Purchasing

At any time, the client may request to view the "channel" of purchase opportunities. This may be in response to an interactive ad which the client saw but did not interact with, or for items which the client did not see. The controller maintains a database of all interactive ads and displays the icons related to these ads to the client. Ads are characterized by a number of keywords, and these keywords are used to display the purchase opportunities in categories such as Food, Sports, Clothes, etc.

3.7 Data Mining

Subject to the consent of the viewers, information collected by the controller (see section 3.2) can be used to generate statistics on purchasing and viewing habits, in the aggregate or for individuals. The latter can be

useful in the targeting of specific opportunities to viewers with particular interests. Commercial packages like IBM's Business Intelligence can be utilized to generate statistics from the raw data thus collected.

Among the statistics of interest can be the effectiveness of an iTV ad for a particular demographic, at a particular time, on a particular channel or during a particular program. Statistics can reflect the number of "clicks" (i.e. remote control keys pressed in response to offered iTV content), whether or not a purchase was completed. Advertisers may be interested in the interest generated by an iTV ad and how highly the viewers' response to interactivity is correlated to purchases at that time or in the future.

3.8 Privacy

One issue, which may slow the adoption of iTV purchasing, is concern over privacy. Viewers do not want anyone monitoring the channels that they choose to watch. All information which is being gathered on viewer activity (purchases, channel changes, etc.) must be collected strictly with the consent and prior approval of the viewer.

The controller, intermediating between the STB and the servers, needs to "snoop" on only a very limited amount of the traffic between the two, e.g. total price of items to be purchased, amount of sales tax, etc. The remaining traffic can pass through the controller in encrypted form.

Information on individual activity can be discarded, and retained only in the aggregate. Incentives can be offered to viewers who wish to volunteer more individual information.

3.9 Security

The controller represents a vulnerable node in the guarantee to the viewer of secure transactions. It is of paramount importance that every precaution be taken to ensure security of information and funds. Since the STB may not have the capability of reading a smart card and the entering of detailed viewer information (such as credit card number, address, etc.) may be difficult, the controller may need to know this information in advance and provide access to it as a service. Access to this information for purchases can be secured via a PIN number, dedicated STB software or hardware, or other means.

3.10 Hand-offs

It is to be expected that while viewers might be interested in and attracted by the dynamic nature of multimedia iTV purchase offers, many will pursue

such offers only to the point of getting more information, and stop short of making a purchase, for several reasons. Despite the best security technology, many users (especially those who are not computer-savvy) will be reluctant to make a purchase using their television as the medium, for fear of privacy and security breaches. In order to coordinate the distance and in-person shopping experiences, "hand-offs" may be used, i.e. the initiation of a purchase (or purchase inquiry) in one medium which is followed by its continuation in another medium.

The controller would facilitate the management of shopping across media boundaries, by maintaining status information for each user on each initiated transaction, and enabling transfer of shopping data between user devices.

4. IMPLEMENTATION

Our research team has implemented a prototype controller system demonstrating several of the features described above. The system consists of three main components:

1. Transaction Capture Engine
2. Application Server
3. Persistent Storage Server.

The controller is inserted in the iTV network to enable advanced interactive services. It provides services to the iTV users, captures transactions and redirects commerce-related traffic to external commerce servers.

We have implemented and tested the controller in a pilot iTV commerce system built in our lab, consisting of a scaled-down cable head end, an iTV authoring suite, a commercial Internet commerce server, and an off-the-shelf STB, demonstrating the viability of our architecture as enhancement for current iTV systems.

In this section we briefly describe the building blocks for the prototype controller, and the test system we have adopted for functional and interoperability testing.

4.1 Transaction Capture Engine

The Transaction Capture Engine is a programmable proxy responsible for capturing HTTP transactions. This data is used for later analysis of system traffic and user activity. The Transaction Capture Engine is built on top of the IBM WeB Intermediaries (WBI) [8]-[10]. The WBI Development Kit provides a convenient and flexible API for programming intermediaries on the Web. In simple terms, an intermediary is a programmable proxy. More

precisely, intermediaries are computational entities that lie along the HTTP stream and are programmed to tailor, customize, personalize, or otherwise enhance data as it flows along the stream.

From the perspective of an HTTP client, WBI acts as an HTTP request processor, receiving a request and sending a response. The processing that happens in between is completely controlled by the functions WBI is set up to provide. The WBI architecture provides a simple and powerful programming framework for developing any such Web applications.

A WBI transaction flows through three basic stages: request editors, generators, and document editors. Request editors receive a request and may modify it before passing it along. Generators receive a request and produce a corresponding response (i.e., a document.) Document editors receive a response and may modify it before passing it along. When all steps are completed, the response is sent to the originating client. A fourth type of processing element, a monitor, can be designated to receive a copy of the request and response but cannot otherwise modify the data flow. Collectively, we refer to Monitors, Editors, and Generators as MEGs.

In our implementation with WBI, a special monitor has been built to gather data from the HTTP session and store it in a data storage facility for later processing and analysis (see sections 3.2, 3.4 and 3.7.)

4.2 Application Server

The Application Server runs the applications that support services such as those described in section 3. The Application Server interoperates with the Persistent Storage Server to maintain metadata related to the TV programming and user-related data. In our implementation the Application Server has been built on top of the IBM's WebSphere Application Server [14].

WebSphere Application Server provides an environment for open distributed computing. Users and processes on a wide variety of platforms can interact by using the facilities provided by WebSphere. The WebSphere software platform delivers the first commercial server that provides J2EE integrated with tools that support Web services standards, such as Simple Object Access Protocol (SOAP) [13] and Universal Description, Discovery and Integration (UDDI) [11]. By implementing these open standards in ways that preserve and enhance the server's scalability and reliability, the WebSphere software platform delivers a highly productive environment for developing, deploying and integrating e-business applications. The incorporation of Web services opens a new door to building WebSphere software applications that dynamically interact with Web services using

SOAP, UDDI, Web Services Description Language (WSDL) [12] and XML [19].

The Persistent Storage is provided by Enterprise Java Beans (EJB) [15] Entity Beans, which provide the Object mapping to the database tables maintained by the Persistent Storage Server. The database tables store information such as user profiles, re-mapping of URIs to local servers, bookmarks, etc.

Business logic has been implemented with EJB Session Beans. One of the advantages of using EJBs for transaction-based services is the capability to provide ACID (Atomic, Consistent, Isolated, Durable) transactions. This supports the requirements of high reliability and consistency that are necessary to run transactional services.

User interface and management functions have been implemented using Servlets and Java Server Page (JSP) technology. JSP and Servlets interact with the EJB Session and Entity Beans to run the applications that support the services listed in section 3.

4.3 Persistent Storage Server

The Persistent Storage Server maintains all the tables used by the controller for its operation, such as metadata related to the TV programming and user-related data. It is based on the IBM's DB2 Relational Database [16].

4.4 Test System

The test system for the prototype controller is illustrated in Figure 3. iTV commerce content is created with the IBM HotMedia TV authoring package [21], and then is transferred to a broadcast server that sends the MPEG-2 Transport multiplex composed of video, audio and interactive data to a Scientific Atlanta QAM modulator.

The digital broadcast signal is received by a Scientific Atlanta Explorer 2000 STB, which runs a PowerTV HTML engine. A user can interact with iTV-enhanced content, and purchase items through the STB. When a commerce transaction is started, the controller tracks it and routes it to the IBM WebSphere Commerce Suite (WCS) server [22]. Digital media files can be purchased from the STB and transferred to a PDA. A subset of the services described in section 3 have been deployed and tested in this environment.

Figure 3. iTV commerce test system

5. CONCLUSIONS

The deployment of open iTV commerce systems has been slowed down by a multiplicity of factors. Business rules need to be defined and enforced among the entities involved in the commerce scenario. Furthermore, a number of operators, advertiser and user services need to be supported. In order to overcome some of these issues, we have designed an architecture that leverages the current iTV infrastructure by introducing a new plug-in element, the controller. The controller enforces business agreements between the parties involved in the iTV commerce flow, acts as a central point of control and information for data exchange, and provides the platform for the introduction of operator, advertiser and user services. The concept of controller as a service platform enables the introduction of new services beyond those described in this paper.

Our team has implemented a successful prototype demonstration system that embodies the concepts discussed above and interoperates with off-the-shelf iTV components.

REFERENCES

- [1] J. Paraszczak, F. Schaffa, D. Seidman, R. Chernock, P. Dettori and S. Depp. The Convergence of TV and PC - An Arranged Marriage. In 1999 Broadcast Engineering Conference Proceedings, NAB, 506-515, April 17-22, 1999.
- [2] R. Chernock, P. Dettori, X. Dong, J. Paraszczak and F. Schaffa, D. Seidman. Approaches in Enabling Electronic Commerce Services over Digital Television. In Proceedings of the 2nd International Conference on Telecommunications and Electronic Commerce, Nashville, TN, Oct. 6, 1999.
- [3] ATSC Digital Television Standard. In ATSC Standard, ATSC A/53 September 1995.
- [4] ATSC Data Broadcast Standard. In ATSC Standard, ATSC A/90, July 26, 2000.
- [5] DVB - Specification for Data Broadcast. EN 301 192, Edition 1.2.1.
- [6] ISO / IEC 13818-1 Generic Coding of Moving Pictures and Associated Audio: Systems.
- [7] ISO / IEC 13818-6 Generic Coding of Moving Pictures and Associated Audio: Digital Storage Media Command and Control.
- [8] IBM WBI Toolkit, available at <http://www.alphaworks.ibm.com/tech/wbidk>
- [9] Barrett, R. & Maglio, P. P. (1999.) Intermediaries: An approach to manipulating information streams. IBM Systems Journal, 38, 629-641.
- [10] Barrett, R., Maglio, P. P., & Kellem, D. C. (1997.) How to personalize the web . In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97), Atlanta, GA
- [11] "UDDI Version 2.0 Data Structure Reference", June 8, 2001. Available at <http://www.uddi.org/pubs/DataStructure-V2.00-Open-20010608.pdf>.
- [12] Web Services Description Language (WSDL) 1.1, March 15, 2000. Available at <http://www.w3.org/TR/wsdl>
- [13] SOAP Version 1.2, W3C Working Draft 9, July 2001. Available at <http://www.w3.org/TR/soap12>.
- [14] IBM WebSphere Application Server. See <http://www-4.ibm.com/software/webservers/appserv/>
- [15] Enterprise Java Bean 2.0 Specifications. Available at <http://java.sun.com/products/ejb/docs.html>
- [16] IBM DB2 Relational Database: <http://www-4.ibm.com/software/data/db2/>
- [17] Uniform Resource Identifiers (URI): Generic Syntax", IETF Draft Standard, August 1998. Available at <http://www.ietf.org/rfc/rfc2396.txt>
- [18] DVB - Multimedia Home Platform, Revision 1.5, February 2000, SB280007.

- [19]XML - Extensible Mark Up Language. See <http://www.w3.org/TR/PR-xml-971208>.
- [20]Declarative Data Essence, Content Level 1, Interactive Television Triggers. SMPTE Draft 06/11/01
- [21]"Synchronized Data Services over MPEG-2 Digital Television Broadcast Systems", J. Brunheroto, R. Chernock, P. Dettori, X. Dong, J. Paraszczak, F. Schaffa, D. Seidman, European Transactions In Telecommunications Journal, Vol. 12, Issue No. 3 (2001.)
- [22]IBM WebSphere Commerce Suite. See <http://www-4.ibm.com/software/webservers/commerce>