

Mapping Large Educational Websites to Interactive DVDs

Swati Ittan
Microsoft Research India
Bangalore, India
t-swati@microsoft.com

Gaurav Paruthi
University of Michigan
Ann Arbor, MI
gparuthi@gmail.com

William Thies
Microsoft Research India
Bangalore, India
thies@microsoft.com

Abstract—In areas lacking computers and Internet connectivity, one promising strategy for delivering electronic educational content is to utilize common household technologies such as TVs and DVD players. Recent work has demonstrated that it is possible to leverage the interactive features of DVD players to provide an (offline) Internet browsing experience, using the remote control for search and navigation. However, until now it has not been documented how to overcome the engineering challenges in mapping large numbers of interactive menus to a single DVD. This paper offers the first description of such a tool. We have applied this tool to burn over 257,000 screens of Wikipedia content to an interactive DVD, for use by low-income communities.

I. INTRODUCTION

Televisions and DVD players remain some of the most ubiquitous household technologies in India. Today, 66% of households in the country have television and 17% have DVD players [1]. Unlike costly technologies such as computers and Internet, TVs and DVD players are routinely acquired and maintained by low-income communities, in both urban and rural areas. In addition, there is a very robust network for distributing DVDs in India. Blockbuster movies are pirated and distributed on DVD even before they leave the theater. However, usually this DVD ecosystem is driven mainly by entertainment. We believe there is an opportunity to leverage the same infrastructure, in concert with the advanced features of DVD players, for the benefit of education.

In pursuit of this goal, we have been developing innovative DVDs that enable large educational corpuses to be distributed at very low cost. A little-known feature of the DVD specification is that DVDs can be very interactive: they enable the viewer to navigate menus by selecting buttons using the remote control. While menus are typically used to choose the chapter of a movie, or to select between different language options, they can also be used to navigate hyperlinked pages, much in the same style as the World Wide Web. Using new tools, we have pushed the capabilities of DVD players to their limits, and have mapped large Internet websites to DVD, for offline browsing on television.

An example of one of our DVDs appears in Figure 1. This double-layer DVD contains the entirety of Schools Wikipedia, a subset of about 5,500 Wikipedia articles for use in education [2]. Our DVD enables the user to search and browse the




Fig. 1. Example screen from Schools Wikipedia on interactive DVD.

disk, which contains about 257,000 menus, using the remote control for navigation. Hyperlinks are numbered in yellow and can be followed by pressing the corresponding key on the remote control. The design and evaluation of this DVD has been described in a prior publication [3]. The use of interactive DVDs for children's books and PowerPoint teaching aids has also been explored previously [4].

However, despite the promise of using interactive DVDs as a platform for education, it remains a significant engineering challenge to map a large number of hyperlinked menus to a single DVD. Prior publications were unable to offer a general solution to this problem; the discs produced were merely prototypes that were often incompatible with off-the-shelf DVD players.

In this paper, we describe a practical tool that enables robust production of interactive DVDs for use in education. Unlike DVDs that are available in stores, the DVDs produced by our tool can contain hundreds of thousands of interactive menus. This presents significant technical challenges for mapping the content to DVD. In the remainder of this paper, we give an abstract description of these challenges, before describing our solution as well as our implementation of a working tool. We have utilized our tool to produce a robust version of Schools Wikipedia on TV-DVD, which we are currently preparing for a public release. In addition, our tool can be applied in other contexts to automatically map any set of hyperlinked web pages to DVD.

II. RELATED WORK

There exists a large literature on the use of video content as an instructional tool; see Goldman et al. for a review [5]. In the Indian context, DVDs have been applied thoughtfully in the classroom as part of the Digital Study Hall project [6], though this project does not emphasize use of interactive features. Non-profit NGOs such as the Children’s Lovycastle Trust and SVYM (in their PremaVidya initiative) have also utilized DVDs in Indian classrooms, including interactive DVD menus to navigate between pages and to do multiple-choice self-assessments. However, the number of interactive menus has been limited (less than 100 screens per disc). Projects such as Digital Green [7] have also found strong applications of DVD videos in educating rural farmers, though have not made heavy use of interactive features to date.

While educational companies have produced interactive TV-DVDs that are widely available in India, they usually focus on video content rather than supporting navigation between a large number of hyperlinked pages. For example, one can buy children’s books and reference materials from BookBox, Pebbles Infotainment, and Vision Interactive Systems; often these discs will read an illustrated book aloud, similarly to a subtitled movie. Though some interactive menus are present, they serve mainly for selecting chapters of the movie as opposed to searching a large database. Conversely, in the United States, companies such as Screenlife have shipped discs that exploit the full power of the DVD specification (including thousands of menus), but focus on the domain of games rather than education.

Other efforts have focused on low-cost ways to distribute Wikipedia. While offline versions of Wikipedia have been mapped to DVD (see for example [8]), to date these discs have required a computer for playback and are not in the right format for viewing on DVD players. The WikiReader [9] is a special-purpose device that costs \$99 and displays up to 3 million Wikipedia articles (much more than the 5,500 available on our disc). However, the WikiReader omits all images, limits the display to black and white, and has a 240 x 200 resolution (less than that of a TV). The primary advantage of our solution is to leverage existing household technologies to enable anyone to copy and distribute the disc at very low cost.

The distribution and usage of DVDs amongst low-income communities in India has been well-documented [10]. Researchers have also proposed innovative ways to track dissemination of video discs between peers [11]. The potential of interactive TV has also been explored in the Indian context [12].

In prior work, we explored the use of interactive DVDs for displaying PowerPoint teaching aids and libraries of children’s books [4] as well as offline versions of Internet websites such as Schools Wikipedia [3]. However, the authoring tool that we utilized previously had serious limitations, which are resolved by the techniques in this paper. Our prior discs were produced with open-source tools (such as dvdauthor) and proved to

```

Media → Menu+ Statement+
Menu → Image ButtonHandler*
ButtonHandler → Register := Value; jump Label
Statement → show Menu | Label | jump Label |
             if Register = Value then Statement |
             Register := Register op Register |
             Register := Value

```

Fig. 2. Context free grammar of our abstract media format.

be incompatible with many DVD players on the market; our current tool produces discs that are supported by each of the 18 DVD players we have tried. In addition, the number of menus supported by our prior tool (Figure 2 in [4]) is dramatically surpassed by our current tool. In particular, the number of menus per disc no longer has any dependence on the number of buttons per menu.

III. MAPPING LARGE DATABASES TO DVD

The core technical challenge in our work is in mapping a very large number of interactive menus to a single DVD while respecting the constraints of the DVD format. Because the DVD format is very complex, and also is not publicly available, we describe the key ideas behind our technique in terms of an abstract model. Those who are familiar with the DVD specification will understand how to translate our findings into that context. However, our model should **not** be interpreted as a direct representation of the DVD format, which contains many nuances.

A. Abstract Model

The abstract media format that we consider in this paper is specified in terms of a context free grammar (see Figure 2). At the highest level, the format consists of a set of menus and a set of statements. Each menu is composed of an image and a set of button handlers; whenever a button is activated on the menu, the corresponding button handler is invoked.

A button handler, in turn, has a very specific format. It must consist of exactly two statements. The first statement is an assignment of a value to a register. The second statement is a jump to a label. This label must be defined within the set of statements for the media. In addition, only a small number of labels are allowed on a given disc (a constraint that is not captured in the context-free grammar.)

The statements in the media format represent logical code that connects different menus. The statements can perform general-purpose computations, including arithmetic and control flow. There is also a statement to display a new menu on the screen; when this statement is called, the machine pauses until a button handler is invoked.

Using this format, the usual way to encode a transition between two menus is illustrated in Figure 3. First, a button handler is invoked from the starting menu. This handler stores within a register a *unique menu identifier* for the target menu, and then jumps to a label corresponding to a lookup table.

```

Media → (Menu1, Menu2, Menu3) Statements
Menu1 → Image1
Menu2 → Image2
Menu3 → Image3 ButtonHandler1 ButtonHandler2
ButtonHandler1 → R0 := 1; jump L1
ButtonHandler2 → R0 := 2; jump L1
Statements → show Menu3
L1 :
if R0 = 1 then show Menu1
if R0 = 2 then show Menu2

```

Fig. 3. Example media file encoding a jump between menus. When showing the third menu, button 1 transitions to menu 1 while button 2 transitions to menu 2.

The lookup table uses comparison operations to test the stored menu identifier, and when it finds a match, it shows the destination menu.

B. Problem Definition

For small numbers of menus, the model just described is perfectly well-suited for displaying a hyperlinked website using menus and buttons. However, a difficulty arises when mapping a very large number of menus to this model. The problem can be formalized as follows:

Problem 1. Consider a media format with B bits per register. Then the approach in Figure 3 can support at most 2^B menus, since each jump target is encoded in a single register. How can a disc support more than 2^B menus?

Real DVD players have 16 bits per register, which means that the encoding in Figure 3 can support at most $2^{16} = 65,536$ menus. In order to support large databases such as Schools Wikipedia, which contain over 250,000 menus, we need to solve Problem 1.

C. Our Solutions

We offer two solutions to the mapping problem. They can be utilized together or in isolation.

Solution 1. Let R denote the number of registers in the machine. Before showing a menu, all registers are initialized to zero. To encode a jump to menu n , the button handler stores the value $n \div R$ into register number $n \bmod R$. Assuming B bits per register, this encoding supports up to $R * 2^B$ menus.

A simple example of Solution 1 (not requiring modulo or division operations) appears in Figure 4.

Solution 2. Let S denote the maximum number of labels on a disc. To encode a jump to menu n , the button handler stores the value $n \div S$ into a given register and jumps to label number $n \bmod S$. Assuming B bits per register, this encoding supports up to $S * 2^B$ menus.

A simple example of Solution 2 (not requiring modulo or division operations) appears in Figure 5.

```

Media → (Menu1, Menu2, Menu3) Statements
Menu1 → Image1
Menu2 → Image2
Menu3 → Image3 ButtonHandler1 ButtonHandler2
ButtonHandler1 → R1 := 1; jump L1
ButtonHandler2 → R2 := 1; jump L1
Statements → R1 := 0
R2 := 0
show Menu3
L1 :
if R1 = 1 then show Menu1
if R2 = 1 then show Menu2

```

Fig. 4. Example media file utilizing Solution 1. The user's experience is the same as in Figure 3, but the implementation scales to large numbers of menus.

```

Media → (Menu1, Menu2, Menu3) Statements
Menu1 → Image1
Menu2 → Image2
Menu3 → Image3 ButtonHandler1 ButtonHandler2
ButtonHandler1 → R0 := 0; jump L1
ButtonHandler2 → R0 := 0; jump L2
Statements → L1 :
show Menu1
L2 :
show Menu2

```

Fig. 5. Example media file utilizing Solution 2. The user's experience is the same as in Figure 3, but the implementation scales to large numbers of menus.

Our solutions can be applied together to support a total of $R * N * 2^B$ menus. In practice, this allows the encoding to address any number of menus that can physically fit on a disc. While there are additional constraints in the DVD specification that eventually limit the number of menus, they do not become relevant until there are tens of millions of menus.

IV. IMPLEMENTATION

To generate the images for use on DVD, we render the HTML pages in a Web browser and take screenshots via .NET controls. These images are processed further (using various third-party tools) and eventually synthesized into a DVD image using Scenarist, the industry-standard DVD authoring tool. An overview of the tool flow appears in Figure 6.

We faced four key challenges in completing this implementation. The first was the simple fact that Scenarist is designed for manual layout of DVD menus, rather than automated generation of large numbers of menus. At the advice of Scenarist customer support, we automated the generation of menus by writing directly to their file format, which is completely undocumented and required considerable trial-and-error.

The second challenge was the runtime of Scenarist. We found that it took about 6 hours for Scenarist to load 4,000 menus via one of our auto-generated save files. This would

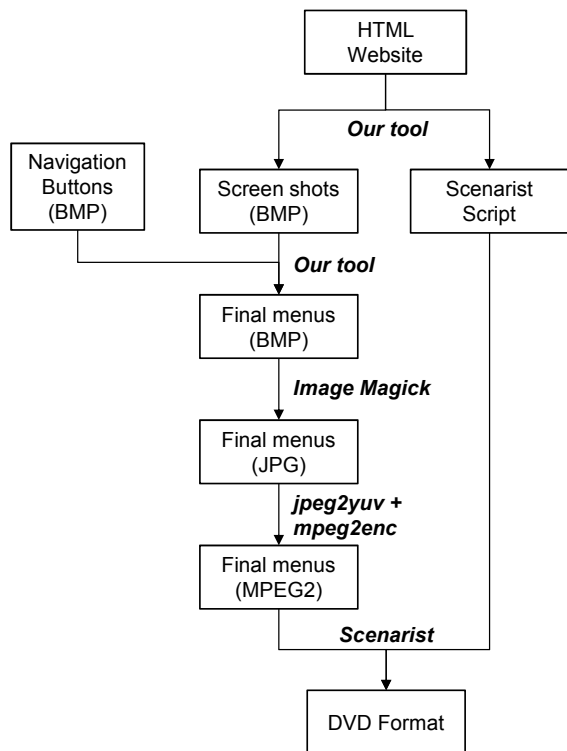


Fig. 6. Toolchain for mapping websites to interactive DVDs.

imply a runtime of at least two weeks to generate the Schools Wikipedia DVD, assuming that the runtime scaled linearly with the number of menus (unfortunately, it didn't). Scenarist customer support acknowledged the problem, but was unaware of a method to speed up the process. When we pushed further, a representative wrote, "I am not familiar with another customer trying to create a disc with the maximum number of menus allowable in the DVD specification."

Undeterred, we set out to generate the menus more quickly by parallelizing the process. We discovered that the VOB components of a DVD disc can be independently generated and copied to the destination folder, while retaining the capability to jump from one component to another. Thus we broke the task into 64 units, each requiring about 6 hours to generate. Generating 10 units in parallel, we could complete the assembly of the Schools Wikipedia disc (starting from 257,000 encoded screenshots) in about 38 hours; since this process required intermittent manual attention, this translated to 3 working days. Including the time to generate the screenshots, the entire process took approximately a week.

The third challenge was shrinking the assets of the disc to fit within the 8.55 GB limit of a double-layer DVD. Encoding the images at full quality resulted in a considerably larger file. Thus, we reduced the bitrate of mpeg2enc to 2500 kbps; this caused compression artifacts that are visible on a high-end monitor, but not on an analog TV (the most common device for our target users). Also, we discovered that unnecessary space was being consumed by the "highlight stream", a video

layer which highlights buttons when pressed. Since we rely on numeric key presses instead of graphical button selection, we removed this stream and saved approximately 500MB of space.

The final challenge was in ensuring that the generated disc was compatible with all DVD players. Originally we utilized the open-source dvdauthor program instead of Scenarist; however, we discovered that the resulting DVDs were rejected by up to half of the DVD players we tested in retail outlets in India. Even after moving to Scenarist, we encountered unexpected difficulties with a handful of unbranded (locally-manufactured) players, that we conjecture are implementing only a subset of the complete DVD specification. To overcome these compatibility challenges, we made minor adjustments to our tool, including avoiding use of a special element called the "language unit", as well as the "VTSM" (menu) domain. We tested our final disc in 18 brands of commercial DVD players, including locally manufactured models displayed in second-hand shops in Bangalore, and found that it worked in every case.

V. CONCLUSIONS

This paper describes a practical tool that overcomes the technical barriers in mapping a large number of interactive menus to a single DVD. While we described this tool in the context of Schools Wikipedia, it is equally applicable for mapping any set of Web pages to the TV-DVD platform. We believe that this tool could have broad applications for bringing hyperlinked educational content to regions lacking computers and Internet connectivity.

VI. ACKNOWLEDGMENTS

We are grateful to Scenarist's technical support team (especially Matt Barr and Dave Smith), rOIZ from IfoEdit, and Trai Forrester for their generous time and help with our detailed questions.

REFERENCES

- [1] Screen Digest, "Market Research Report," 2012.
- [2] "Schools Wikipedia," <http://schools-wikipedia.org>, 2010.
- [3] G. Paruthi and W. Thies, "Utilizing DVD Players as Low-Cost Offline Internet Browsers," in *CHI*, 2011.
- [4] K. Gaikwad, G. Paruthi, and W. Thies, "Interactive DVDs as a Platform for Education," in *ICTD*, 2010.
- [5] R. Goldman, R. Pea, B. Barron, and S. J. Derry, Eds., *Video Research in the Learning Sciences*. Lawrence Erlbaum, 2007.
- [6] R. Anderson, C. Robertson, U. Sahni, E. Nabi, and T. Setia, "Facilitated video instruction in low resource schools," in *ICTD*, 2012.
- [7] R. Gandhi, R. Veeraraghavan, K. Toyama, and V. Ramprasad, "Digital Green: Participatory Video for Agricultural Extension," in *ICTD*, 2007.
- [8] "Wikipedia on DVD," <http://www.wikipediaondvd.com/>, 2012.
- [9] "WikiReader," <http://thewikireader.com/>, 2012.
- [10] N. Sambasivan, N. Rangaswamy, E. Cutrell, and B. Nardi, "UbiComp4D: Infrastructure and Interaction for International Development - the Case of Urban Indian Slums," in *UbiComp*, 2009.
- [11] N. Sambasivan, E. Cutrell, and K. Toyama, "ViralVCD: Tracing Information-Diffusion Paths with Low Cost Media in Developing Communities," in *CHI*, 2010.
- [12] N. Rangaswamy, S. Nair, and K. Toyama, "My TV is the family Oven/Toaster/Grill': Personalizing TV for the Indian Audience," in *uxTV*, 2008.