

Development and Maintenance of a Web Site for a Bachelor Program

Dr J Rajaram Professor Ph.D ---CSE
E Ravi Assistant Professor M.Tech—CSE
D Navya, Assistant Professor
Donda Prashanth, Assistant Professor
Nagole institute of engineering and technology

Abstract

This paper describes our experiences with the development and maintenance of the BSENG Web site, which provides information about a new Bachelor program for software engineering at the University of Victoria. The site's requirements lead to a design that emphasizes simplicity to ease maintenance. We used Macromedia Dreamweaver to develop the site. During development, we identified two important maintenance tasks—detection of dead links and preservation of the site's navigational structure—that are not well supported by Dreamweaver. We discuss how we use the Rigi reverse engineering environment to aid the Website developer with these maintenance tasks.

1. Motivation

Nowadays, many Web sites are highly complex software systems [17]. At the same time, Web sites are often important assets for an organization. In response to this need, many professional development tools that support the building of sophisticated Web sites have emerged (e.g., Macromedia's Dreamweaver, IBM's WebSphere, and Adobe's GoLive). However, despite tool support, the current practice of Web site development is still immature and does not adhere to software engineering principles [18]. For example, often there is no design or system documentation available [22]. The first release of a Web site is only the beginning of its lifetime. A Web site has to constantly evolve in order to remain useful for its users. Because Web site maintenance is costly and continuous, Web site developers have to employ tools and techniques that also assist them in their maintenance tasks. The Web site's source code (typical artifacts are HTML pages, client-side and server-side scripts, and configuration files) is the most reliable and detailed information available to the developer. But often higher-level abstractions are equally important to facilitate the understanding of the Web site for development and maintenance activities; reverse engineering is concerned with the generation of such documentation. Development tools often focus on the forward engineering aspect of software development, but neglect the reverse engineering aspect. In this paper, we illustrate with a case study how we used reverse engineering technology in addition to a commercial Web development tool in order to improve the development and maintenance of a Web site. Our case study is BSENG, a Web site for a new Bachelor program for software engineering at the University of Victoria. We decided to build BSENG with Macromedia Dreamweaver since it is a mature, commercial tool that fits well into our environment. However, during the development of the BSENG Web site, we identified two important maintenance tasks—detection of dead links and preservation of the site's navigational structure—that are not well supported by Dreamweaver. We customized the Rigi reverse engineering environment [25][24] to visualize dead links as well as BSENG's (navigational) structure. The documentation produced with Rigi is complementary to the information that Dreamweaver provides and gives the BSENG developer valuable, additional information for recurring maintenance tasks. The paper is organized as follows: The next section identifies three different views of a Web site and their relationships. Section 3 gives more background on our case study, BSENG. Section 4 discusses the problem of dead links and how we customized Rigi to visualize them. Section 5 shows how Rigi visualizes Web site structure and Section 6 closes with some observations and future research directions.

2. Web Site Views

For both development and maintenance phases, it is important to distinguish between the different views of a Website. We discern the following views: **client view**: The view of the Web site that a client (typically using a Web browser) sees. **deployment/server view**: The view of the Web site that a Web server (accessing the local file system) sees. **developer view**: The view of the Web site that a developer (using a Web development tool such as Dreamweaver) sees. Reverse engineering tools can operate on any of the above views. The reverse engineering process starts with fact extraction from one (or several) views. Extracted facts are typically stored in a repository, which is then queried by analyses. Analysis results are then visualized to assist in development and

maintenance activities. Depending on the view, fact extraction uses different extraction strategies and extracts different artifacts. In the following, we discuss each view in more detail and give examples of reverse engineering tools and extractors. Traditional stand-alone extractors for Web sites work either on the client view or on the deployment view

.2.1. Client View

Client-side extraction does not require direct access to a Web site's sources, but the Web server has to be treated as a black box; only its output (i.e., served Web pages) can be observed. Examples of client-view extractors are essentially Web crawlers or spiders. These extractors request pages via URLs, communicating with the Web server via HTTP. For example, the strategy of SiteSeer—a tool to collect Web metrics—is typical for client-view extraction [29]: "Beginning at a known URL, SiteSeer spawns a child process to parse the HTML document which will be found there. Each link found within the HTML text will be added to a links file produced by the child, and ultimately returned to the parent. The parent will spawn a new child for each of these new links, so that the whole hyperdocument is eventually examined. . . . This search, which is initially exponentially increasing, is limited by the 'edge' of the website; that is, HTML documents outside the initial website are not parsed, although the links are followed to ensure the URLs are good." SiteSeer's parser is implemented with Lex and Yacc. Similarly, Perlbot is a web crawler written in Perl that gathers metrics for a Web site to help understand its evolution [12]. Ricca and Tonella have developed the ReWeb tool, which consists of an extractor, analyzer, and viewer for static Web sites [19] [20]. The extractor, written in Java, downloads all pages of a certain Web site starting from a given URL. Links in pages that point outside the Web site are ignored. The extracted Web site is represented as a typed, directed graph. The ReWeb tool has several analyses that operate upon the graph structure. Most of these are inspired by traditional compiler (flow) analyses and mapped to the Web site domain (e.g., dominators of a page, shortest path from the start page to a page, and strongly connected components). Results of analyses are visualized with Dotty [16] (a customizable graph editor).

.2.2. Deployment View Most reverse engineering research has targeted the client view, but more deployment-view extractors are emerging. A deployment-view extractor has access to the Web site's sources (such as HTML pages, CGI scripts, Java Server Pages, and configuration files). Hassan has developed coarse-grained deployment-view extractors for HTML, JavaScript, VBScript, SQL database access, and Windows binaries [9]. During the extraction process, each file in the directory tree that contains the Website is traversed and depending on the file type the corresponding extractor is invoked. Extracted facts are represented in the Tuple Attribute (TA) format [10]. All extractor output is consolidated into a single file and visualized with the Portable Bookshelf system [7] for the purpose of Website architecture recovery. Di Lucca et al. describe a reverse engineering process for Web sites that employs both static and dynamic analysis [3] [4] [5]. For the static analysis, facts are extracted from HTML pages, client-side scripting languages (JavaScript and VBScript), server-side scripting languages (ASP and PHP), and the directory structure. These facts are represented in a proprietary, XML-based format called IRF, which is translated to a relational database (comprising a schema with 74 tables). Analyses use SQL on the relational database to retrieve facts. Several graph drawing tools (Rigi, VCG, and Dotty) are used for visualization

.2.3. Developer View

To our knowledge, neither extractors nor analyses for the developer view have been developed by the reverse engineering community. Development tools expose the developer view to the Web site developer, but they are geared towards forward engineering as opposed to reverse engineering functionality needed for maintenance activities. A study of three commercial tools (FrontPage 2000, Dreamweaver UltraDev 4, and SmartSite 3) conducted by Tilley and Huang in 2001 came to the conclusion that all tools had limited capabilities with regard to supporting reverse engineering activities [23]. Maintenance activities that are supported by these tools are, for example, validation Proceedings

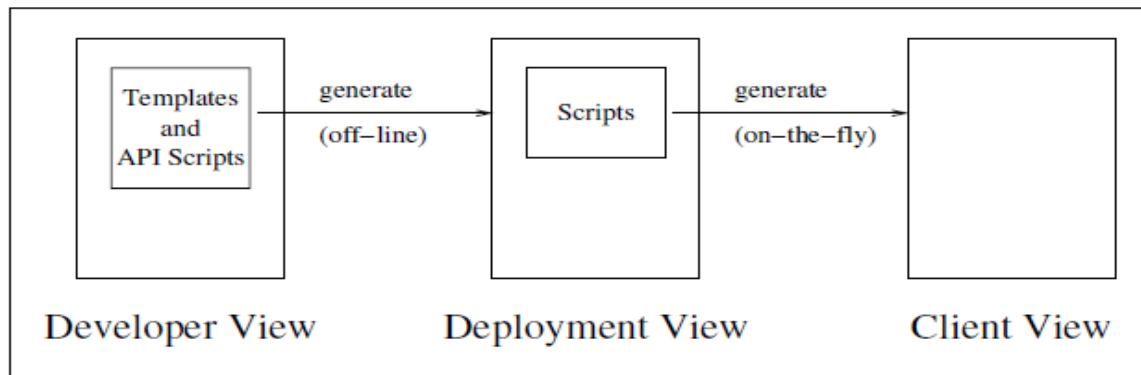


Figure 1. Generative aspect of Web site views

of HTML and XML documents, reports of usage-violations of ALT and META tags, link checking, metrics that summarize characteristics of Web pages, and page download-time estimates. Fortunately, Web site development tools such as Dreamweaver and GoLive have now scripting capabilities and expose tool functionalities with a programmable API. These *API scripts* can be used to automate recurring development activities, but also make it possible to implement dedicated reverse engineering functionality.

2.4. Relationship of Views

All of the three views introduced above are of potential interest to the Web site maintainer. For example, the developer view shows the high-level Web design such as information about templates; the deployment view is the one the Web server uses and thus important for server maintenance; finally, the client view is the one that the user sees and thus is important to assess navigability and structure of the site as well as to detect dead links. The maintainers have to take into account the characteristics of the view that they work with. Analyses that rely on facts extracted from the client or deployment view can give misleading information to the maintainer of a Web site. An example is clone detection of HTML and embedded scripting [11]. If Web pages are generated from templates (which is supported, for example, by GoLive, Dreamweaver, and certain Wikis), a clone detection analysis (operating on the deployment or client view) will report a large amount of cloned material. These clones, however, are of no concern to the maintainers, because they are working in the developer view, which exposes the templates to them. Dreamweaver can export HTML comments that function as meta-data to identify templates in the generated HTML page. A dedicated clone detection algorithm could make use of this information to suppress clones caused by templates. However, the exporting of meta-data information is typically disabled to optimize bandwidth in the deployed version of the Web site. Many of the differences in the views are caused by generative techniques that connect the views. Mechanisms such as templates at the developer view drive the generation of target code (often off-line) for the deployment view. Dynamic, server-side technology (such as JSP and servlets) in turn generates on-the-fly target code for the client view. Another example of differences between deployment and client view are the potentially complex mappings from URLs to file-system paths that are dynamically performed by the Web server (governed by the server's configuration files). These relationships between the views are depicted in Figure 1. Little attention has been paid so far to the particular reverse engineering problem caused by systems that have a generative component; this is the case for Web sites as well as traditional software systems. It is now common to find Web sites that employ generative techniques. Traditional software also uses generators such as compiler compilers (e.g., Yacc), embedded languages (e.g., ESQL in C), or domain-specific languages [27]. To give an example, Dean and Chen describe the design recovery for a traditional software system consisting of a textual domain-specific language (S/SL) with a generator that produces PT Pascal code [2]. In order to analyze such systems, the reverse engineering tool must be targeted to a specific generator. Ideally, an analysis should identify mappings from pre-generation artifacts to post-generation ones and vice versa.

3. BSENG Web Site

The Bachelor of Software Engineering (BSENG) is a new interdisciplinary program offered by the Faculty of Engineering at the University of Victoria (UVic). The program is scheduled to start in autumn 2003 with an initial enrollment capacity of 75 students per term. One of the authors of this paper was in charge of the initial development of the BSENG site. The BSENG Web site is located at the domain <http://www.cs.uvic.ca> (see Figure 2). The BSENG Website is a recruitment tool and information source for potential applicants to the

program. While populating the site with material, it became clear that the site is also an important resource for members of the faculty and committees involved in the program. The site currently provides information about the curriculum, course structure, program requirements, application procedures, contacts, as well as promotional and program-related download material (e.g., class room slides, presentations, and program proposals). As the BSENG program grows and matures, the Web site is expected to evolve accordingly.

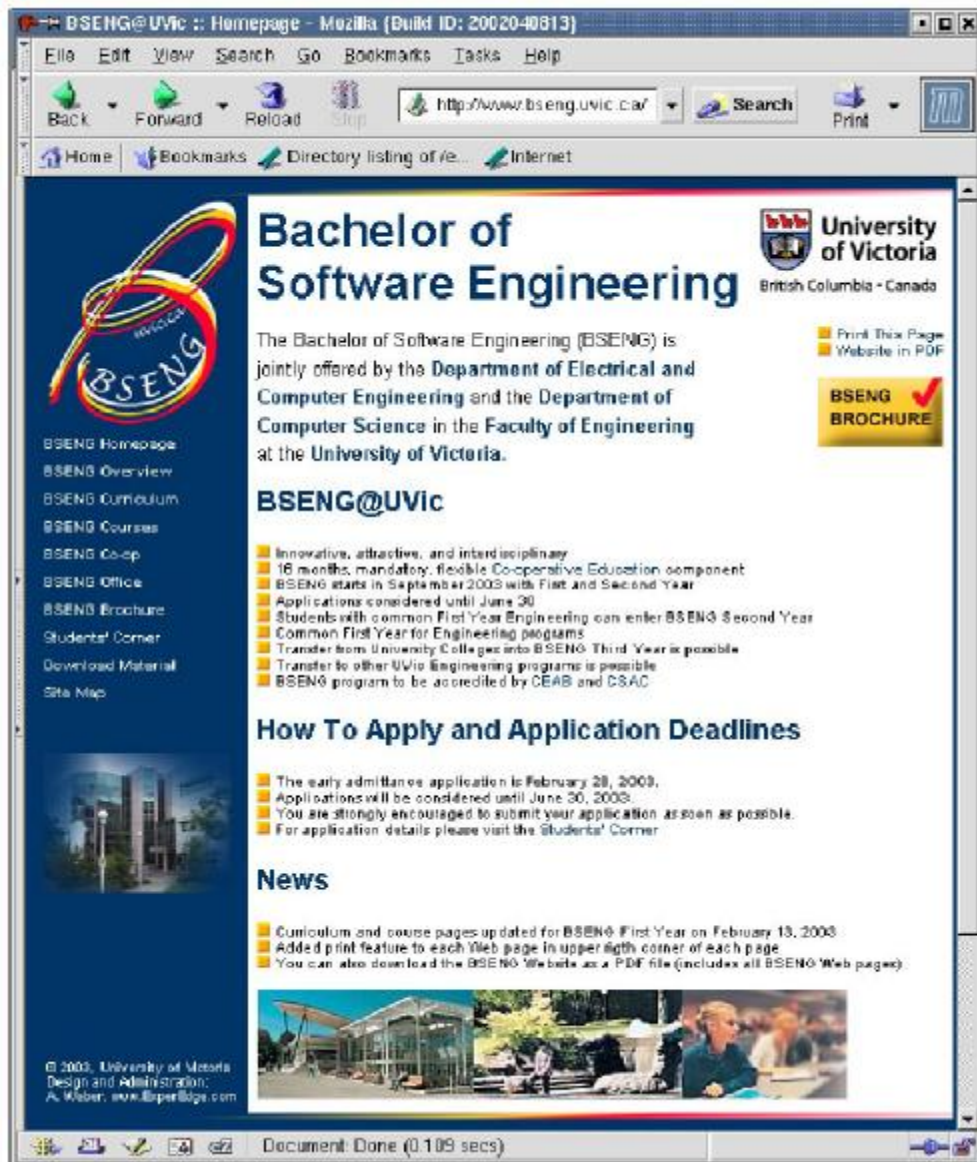


Figure 2. Screenshot of BSENG home page

home page The BSENG site has high textual information content, which will change rather infrequently once the BSENG program is established. However, new information will be added to the Web site continuously. In order to simplify the process of adding new information, we decided on a list structure for each page with a table of contents on top from which the user can jump to the detailed information on the page. Specifically, with about 50 courses and the limited types of changes, we decided that there is no need for database-driven dynamic Web pages. This paper describes the first version of the BSENG site. Since the acceptance of the paper, a new version has been published on June 19, 2003. We further decided on an easy-to-navigate and fast-loading design that is accessible for a wide variety

of browsers. The latter is of special importance in academic environments, which typically exhibit a heterogeneous infrastructure. Another important aspect of keeping the design simple refers to maintenance and costs. In research and education communities turnaround times of Web site administrators are short. We also took the following considerations into account: • On the user side, Nielsen observes that “users have less patience for bleeding-edge technology these days as the Web gets dominated by later adopters and the upgrade speeds for new browsers and plug-ins slowdown” [15]. • On the developer side, Schmeiser advises to subject any Web technology to a rigorous screening before making a decision to adopt [21]. Questions to ask are, for example, “Will this solution save me resources?” and “Will this technology be time-intensive to extract from my system if I decide to replace it?” As a consequence, an overriding principle for the BSENG Web design was the KISS principle: Keep It Simple, Stupid. Despite the simple design outlined above, we aimed for a professional, colorful design that might appeal to students, who have come to expect Web sites that exhibit a high professionalism and quality.

3.1. Site Development There is a staggering variety of technologies, standards, and tools to choose from when developing a Web site. Often Web sites try to employ the latest cutting-edge techniques without regards to the site's clientele as well as development and future maintenance implications. According to our considerations for maintenance outlined in Section 3.2, we consciously tried to avoid advanced features of the tool in favor of simpler, proven technologies. The BSENG site is static, which means that pages are computed at application definition time and remain immutable during application usage [8]. This is in contrast to dynamic sites where pages are computed on-the-fly and can change during application usage. Our site is Level 0 (i.e., “only HTML pages without frames”) according to the classification introduced by Ricca and Tonella [19]. The site uses plain HTML 4.01 without frames. The page layout is controlled with nested tables. We exploit cascading style sheets (CSS) for accessibility purposes. They are mainly used to set font sizes and colors without affecting the navigation of the Web site. In the first version, we did not use client-side scripts, but we will use these for future versions, for example, to adapt the layout depending on the detected browser type. The BSENG site has been developed with Macromedia Dreamweaver MX, an integrated development environment for Web site development. Dreamweaver is available for both Windows and Mac, which is an important consideration in our university environment where the site is being maintained.



Figure 3. Developer views of the BSENG site

in Macromedia Dreamweave

Dreamweaver supports multiple editing views for a page. The WYSIWYG view (cf. large window of Figure 3) allows editing of the page without an in-depth knowledge of HTML. If a page is based on a template, only certain parts of the page can be edited. The code view (cf. top window of Figure 3) shows the deployment view of the page (optionally pretty-printed). Similar to the WYSIWYG view, code that defines the template cannot be modified. Either view can be edited during development and both views are synchronized. An advanced feature of Dreamweaver that we decided off-line with an API script (cf. Figure 1) that is invoked by the developer before deploying the site.

total files	98
HTML files	28
image files	9
all links	1968
external links	345
broken links	294
orphaned links	29

Figure 4. Statistics of the BSENG

site Figure 4 shows some statistical information for the BSENG site taken from Dreamweaver's developer view. The site has only 28 pages (half of which are pages for the print view), but a rather large number of links

3.2. Site Maintenance

The maintenance aspect of a Web site must be taken into account right from the start of the Web site development effort. In fact, the maintenance aspect of Web sites is often neglected. Nielsen lists "Forgetting to Budget for Maintenance" among his *Top Ten Mistakes of Web Management* [14]. The choice of technologies and tools has a significant impact on the future maintenance effort. For the BSENG site we chose to simplify maintenance by employing a small set of simple, well-established technologies. We consciously avoided approaches that have a negative impact on maintenance; specifically:

- No advanced HTML features such as frames and client-side image maps. These are often hard to understand for maintainers who lack a Web development background and can cause problems when migrating to another Web development tool.
- No JavaScript for important functionality such as link navigation. If a JavaScript cannot be properly executed (e.g., because it is broken or has been disabled by the client) it should not degrade the usefulness of the site (e.g., in terms of navigability).
- No server-side scripting (such as JSP, PHP, or ColdFusion). This simplifies configuration and updating of the Web server as well as eases migration to another Web server. It also simplifies maintenance of the Web site itself. For these reasons we decided to keep the BSENG site static by generating the print view of the pages off-line.
- Limited use of technology and features that are specific to the development tool. This mitigates vendor lock-in in case there is the desire to change the development tool. Worse, sometimes a needed update of the development tool forces the developer to migrate code—the consequences and costs of such a migration are hard to predict. As an example, Toeter describes the migration of several Websites at the University of Amsterdam to a new version of the ColdFusion Markup Language (CFML) [26]. 12000 lines of code had to be migrated from CFML 1.0 to 4.0 because the ColdFusion development tool ceased to support the old CFML version. The migration had to deal with technological changes, changes in the tool, and non-backwards compatible changes in the CFML. We hope

to avoid such migration scenarios for the BSENG site. During the development of the BSENG site, we noticed two important tasks that we expect will remain critical during the maintenance and evolution of the site: 1. Dead links both within the site and as the result of link rot. 2. Ensuring the integrity of the site's navigational structure. Unfortunately, both of these tasks are not well supported by Dreamweaver. As a response to this gap in tool support, we customized the Rigi reverse engineering environment. We discuss the tool support that Rigi provides for the maintainer in more detail in Sections 4 and 5. During development and maintenance activities, Dreamweaver and Rigi are used alternately. Typically, the Web developer first modifies the Web site with Dreamweaver and then deploys the new site on the server. Once deployed, the Web site is crawled and its structure visualized in Rigi. The visualization allows the maintainer to identify problems such as dead links, which leads to a subsequent maintenance activity with Dreamweaver.

4. Dead Links Maintenance is constantly necessary even if the Web site itself does not change—the reason being that the environment in which the Web site operates changes. This often causes links to pages that point outside of the Web site to become unreachable. (This phenomenon has been dubbed link rot [6].) In a study conducted by Linos et al., four out of nine Web sites had more than one percent of broken links; one site had as many as 7.5% dead links [12]. Another study conducted in 1999 found that 5.7% of links on the Web are broken.³³ Available at <http://www.pantos.org/atw/35654.html>.

Figure 5. Dreamweaver's link checker Dreamweaver has a link checker that reports broken links. However, the checker operates in the developer view only. This means that the report is accurate for broken intrapage links (i.e., anchors), but can give false positives for links between pages. False positives occur if mappings between the deployment view and the client view are not properly resolved. The mapping of a link between the two views is governed by the Web server, which is not taken into account by Dreamweaver. Figure 5 shows part of the broken links reported for the BSENG site. In this run of the checker, broken intra-page and inter-page links are detected. During development, broken links often act as to-do items indicating unfinished work. While Dreamweaver's ability to detect broken intra-site links is useful for the Web site developer, it is too limited because dead external links are missed. Dreamweaver lists all external links, but does not check them. Thus, the maintainer needs an additional tool that checks the client-view for dead links. Web2Rsf, which has been previously developed by one of the authors of this paper, extracts the link structure (client view) of a Web site for subsequent analysis [13]. The result of the extraction is visualized with the Rigi graph editor [25]. Rigi uses directed, typed graphs to process and display the data to be analyzed. Rigi can be retargeted to different domains by defining a suitable *schema* that expresses the types and attributes of the arcs and nodes that constitute the domain. In the Web schema that we defined, URLs are represented by nodes and links between URLs by arcs. Different node and arc types are represented with different colors (cf. Figure 7). There are different node types corresponding to different URLs: HTML pages (HTML, blue), image files (Image, blue), email addresses (Email, green), and other files (Other, purple). There are also two distinct node types to indicate dead links (DeadURL, red) and syntactically incorrect URLs (MalformedURL, pink). Web2Rsf is

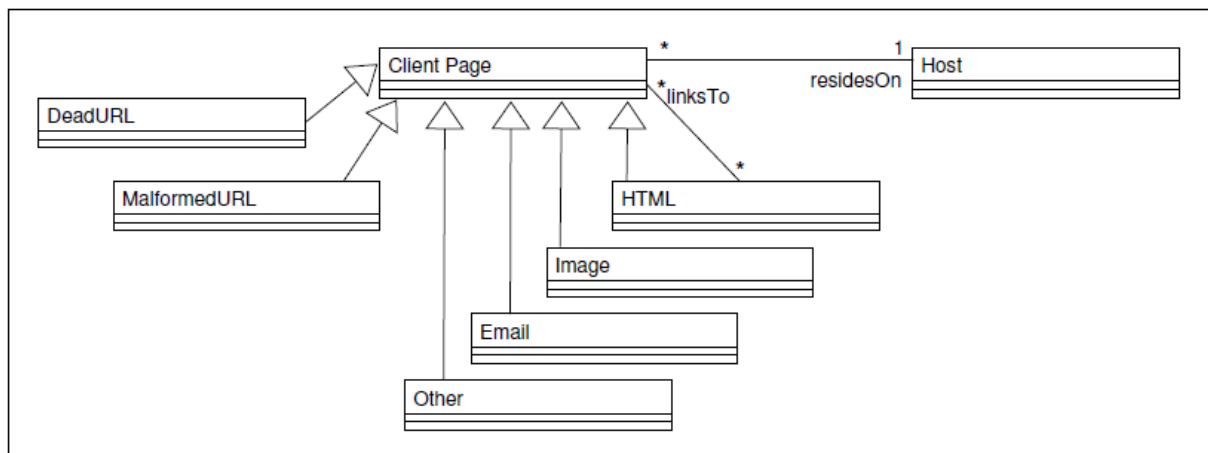


Figure 6. Rigi's Web schema for BSENG (brown),

mailto: (Email, green) etc. Files for downloading (e.g, PDF documents, ZIP archives, and Power-Point presentations) are grouped in a single type (Other, purple). There are also two distinct node types to indicate dead links (DeadURL, red) and syntactically incorrect URLs (MalformedURL, pink). Web2Rsf is

written in Java and reports a dead link if an input/output exception is thrown while connecting. Figure 6 shows the Web schema that is sufficient to model the BSENG site. The schema is described with UML and is similar to the one developed by Conallen [1] and DiLucca et al. [5]. Web pages are modeled as classes; an association is used to represent hyperlinks. The complete RigiWeb schema is discussed in detail in a previous WSE paper [13]. The Rigi views of the Web site allow Web site authors to readily locate internal and external dead links. Web2Rsf does not check anchors, but these are reliably checked with Dreamweaver. Figure 7 shows a rendering of a previous version of the BSENG site. This Rigi view shows 54 DeadURL nodes, which are rendered in red. No syntactically incorrect URLs were found. The site maintainer can now filter out all the working URLs to focus on the dead ones. When inspecting the dead links, we found external ones (referring to UVic's online calendar) as well as internal ones. Many broken internal links were caused by problems with wrong relative links in the print view. Furthermore, not every page had a correct URL to its print view.

5. Navigational Structure

A well designed Web site has an appropriate navigational structure that is apparent to its users and thus helps them to effectively find desired information. Ricca and Tonella discuss several recurring navigational structures of Web sites [20]. A tree structure is acyclic and each node has exactly one parent. A user navigates the tree structure from top to bottom, making a choice at each level. A fully-connected structure means that each page can be reached by following a single link. An indexed sequence means that pages are arranged into a single/double linked list. Pages within the list allow navigation to their previous/next page. Optionally, a table-of-content page—from which all pages in the list are directly accessible—can super-impose a flat tree structure on the list. The BSENG site's navigational structure is essentially a flat tree. Web sites can be composed of regions, each one having a different navigational structure. For example, a future version of the BSENG site might introduce a virtual tour of the UVic campus. This could be accomplished by adding an indexed sequence that is anchored with a table-of-content page to the current tree structure. A well-engineered Web site should be designed with a certain navigational structure and the structure should be clearly documented. Changes to a page due to maintenance activities and evolution of the site should not (unintentionally) change or violate the navigational structure. The navigational structure is determined by the linkage of client-view Web elements (e.g., HTML pages and `mailto:`). The Rigi Web schema describes these elements, and they are visualized in the Rigi graph editor as nodes and arcs. Thus, Rigi graphs can help the maintainer to assess the site's navigational structure. Rigi graphs can be laid out with a spring algorithm to expose clusters in the site's structure (cf. Figure 7). This allows the developer to assess the usability of the Web site by evaluating the reachability of the most important Web pages. Data gathered from Web server log files can be used to investigate usage patterns of Web sites, i.e. which of the available links users actually followed and which of several possible paths users took to locate the Web pages they were looking for. This allows the Web site authors to optimize the structure of their Web site to increase user satisfaction.

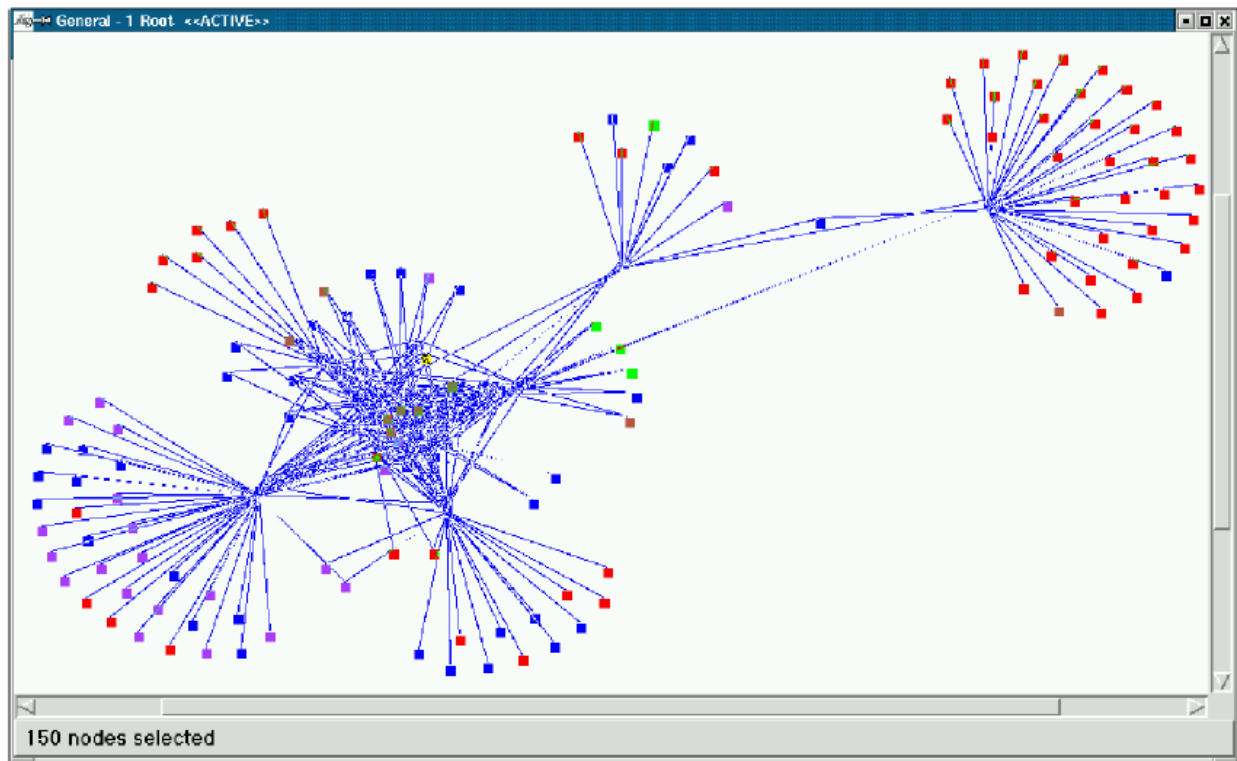


Figure 7. Visualization of dead links in Rigi and thus the potential success of their Web site. Since the BSENG site was still quite young and relatively unknown at the time of our study, we were not able to detect browsing patterns of Web site users. The data Web2Rsf gathered from the log files still revealed some interesting information when viewed in the Rigi graph editor. Web2Rsf records the number of accesses for every successful and unsuccessful page request. In the log file we analyzed, there were about 300 requests for the BSENG startpage and about 130 requests for a file called `robots.txt`. By convention, this file is used to inform web spiders and search engines, whether they are allowed to index a website. The numbers suggest that about 40 percent of traffic on the Web site is caused by spiders. The main style sheet of the BSENG site was only requested about 250 times, suggesting that almost 20 percent of clients disregard stylesheets. Though these results can only be regarded as preliminary because of the small time span the log file covered, they give some useful hints about the users of the Web site that should be considered during future site maintenance. In Rigi, groups of related nodes can be collapsed into a single “super” node (Collapse, cyan). Collapsing of nodes can be recursive, leading to a hierarchy of Collapse nodes. Such groupings are useful to document the site’s organization. When exploring a Rigi graph, a Collapse node can be expanded to reveal the elements that it contains. The maintainer can do groupings manually in the Rigi editor. Alternatively, Rigi’s scripting capabilities make it possible to develop a script that (semi-)automatically performs the grouping. Figure 8 shows BSENG’s top-level graph after running the script that we developed for the BSENG site. After a new version of the BSENG site has been deployed, the maintainer first runs Web2Rsf and then executes the Rigi script to visualize and assess the site’s new organization.

6. Conclusions and Future Work

In this paper, we described our experiences with the development and maintenance of the BSENG Web site, which is an important asset for UVic’s BSENG program. We identified the different views (development, deployment, and client) that a Web developer has to work with. Different tools and technologies operate on different views. It is important for the developer to understand the differences between these views in order to interpret the information of the views correctly. Our experiences with the BSENG site show that Macromedia Dreamweaver is an effective tool for site development but lacks in support for important maintenance tasks. We discussed how the Web2Rsf extractor and the Rigi reverse engineering environment can

be used to generate client-view documentation for assessing dead links and site navigability. Current reverse engineering approaches focus on the deployment and/or client view, but neglect the development view. In the future, we plan to investigate reverse engineering support for the developer view by extending existing commercial Web development tools such as Dreamweaver, GoLive, and WebSphere with reverse engineering functionality. Reverse engineering analyses often generate documentation that represents information at a higher level of abstraction. An important criterion for documentation is that mappings between the different abstraction levels must be preserved [28]. Mappings between abstraction levels are *vertical mappings*. The views that we introduced show that documentation needs to be available for all three views and that *horizontal mappings* between these documents are of equal importance. Horizontal mappings can be quite complex and unintuitive for the Web site developer. We plan to investigate tool support that will allow Web site developers to effectively navigate such horizontal mappings.

Acknowledgments

Thanks to Keith Edwards and Crina-Alexandra Vasiliu for proofreading. This work has been supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Consortium for Software Engineering (CSER), and the Center for Advanced Studies (CAS), IBM Canada Ltd

References

- [1] J. Conallen. Modeling Web application architectures with UML. *Communications of the ACM*, 42(10):63–70, Oct. 1999.
- [2] T. Dean and Y. Chen. Design recovery of a two level system. *11th International Workshop on Program Comprehension (IWPC 2003)*, pages 23–32, May 2003.
- [3] G. A. Di Lucca, M. Di Penta, G. Antoniol, and G. Casazza. An approach for reverse engineering of web-based applications. *Eighth Working Conference on Reverse Engineering (WCRE '01)*, pages 231–240, Oct. 2001.
- [4] G. A. Di Lucca, A. R. Fasolino, F. Pace, P. Tramontana, and U. De Carlini. WARE: a tool for the reverse engineering of web applications. *Sixth European Conference on Software Maintenance and Reengineering (CSMR '02)*, 2002.
- [5] G. A. Di Lucca, A. R. Fasolino, and P. Tramontana. Towards a better comprehensibility of web applications: Lessons learned from reverse engineering experiments. *4th International Workshop on Web Site Evolution (WSE 2002)*, pages 33–42, Oct. 2002.
- [6] D. Eichmann. Evolving an engineered web. *1st International Workshop on Web Site Evolution (WSE '99)*, Oct. 1999.
- [7] P. J. Finnigan, R. C. Holt, I. Kalas, S. Kerr, K. Kontogiannis, H. A. M'uller, J. Mylopolous, S. G. Perlegut, M. Stanley, and K. Wong. The software bookshelf. *IBM Systems Journal*, 36(4), 1997.
- [8] F. Fraternali. Tools and approaches for developing data intensive web applications: A survey. *ACM Computing Surveys*, 31(3):227–263, March 1999.
- [9] A. E. Hassan and R. C. Holt. Towards a better understanding of web applications. *3rd International Workshop on Web Site Evolution (WSE 2001)*, pages 112–116, Nov. 2001.
- [10] R. Holt. TA: The tuple-attribute language. <http://plg2.math.uwaterloo.ca/~holt/papers/ta-intro.html>, 1997.
- [11] F. Lanubile and T. Mallardo. Finding function clones in web applications. *Seventh Conference on Software Maintenance and Reengineering (CSMR 2003)*, pages 379–386, Mar. 2003.
- [12] P. K. Linos, E. T. Ososanya, and H. Natarajan. Maintenance support for web sites: A case study. *3rd International Workshop on Web Site Evolution (WSE 2001)*, pages 70–76, Nov. 2001.
- [13] J. Martin and L. Martin. Web site maintenance with software-engineering tools. *3rd International Workshop on Web Site Evolution (WSE 2001)*, pages 126–131, Nov. 2001.
- [14] J. Nielsen. Alertbox for June 15: Top ten mistakes of web management. *useit.com*, 1997. <http://www.useit.com/alertbox/9706b.html>.
- [15] J. Nielsen. Alertbox for May 2: "top ten mistakes" revisited three years later. *useit.com*, 1999. <http://www.useit.com/alertbox/990502.html>.
- [16] S. C. North and E. Koutsofios. Applications of graph visualization. *Graphics Interface '94*, pages 235–245, 1994.
- [17] J. Offutt. Quality attributes of web software applications. *IEEE Software*, 19(2):25–32, Mar./Apr. 2002.
- [18] R. S. Pressman. What a tangled web we weave. *IEEE Software*, 17(1):18–21, Jan./Feb. 2000.
- [19] F. Ricca and P. Tonella. Web site analysis: Structure and evolution. *International Conference on Software Maintenance (ICSM '00)*, pages 76–86, Oct. 2000.
- [20] F. Ricca and P. Tonella. Understanding and restructuring web sites with ReWeb. *IEEE MultiMedia*, 8(2):40–51, Apr.–June 2001.
- [21] L. Schmeiser. Web site evolution: Design and developing for the future. *1st International Workshop on Web Site Evolution (WSE '99)*, Oct. 1999.
- [22] M. Taylor, J. McWilliam, J. Sheehan, and A. Mulhaney. Maintenance issues in the Web site development process. *Journal of Software Maintenance and Evolution: Research and Practice*, 14(2):109–122, Mar./Apr. 2002.
- [23] S. Tilley and S. Huang. Evaluating the reverse engineering capabilities of web tools for understanding site content and structure: A case study. *23rd International Conference on Software Engineering (ICSE 2001)*, pages 514–523, May 2001.
- [24] S. R. Tilley. Domain-retargetable reverse engineering II: Personalized user interfaces. *1994 International Conference on Software Maintenance (ICSM '94)*, pages 336–342, Sept. 1994.
- [25] S. R. Tilley, H. A. M'uller, M. J. Whitney, and K. Wong. Domain-retargetable reverse engineering. *Conference on Software Maintenance (CSM '93)*, pages 142–151, Sept. 1993.
- [26] B. Toeter. Lexical scanners for 4GL-source maintenance of a corporate web site. *2nd International Workshop on Web Site Evolution (WSE 2000)*, pages 51–56, Mar. 2000.
- [27] A. van Deursen, P. Klint, and J. Visser. Domain-specific languages: An annotated bibliography. *ACM SIGPLAN Notices*, 35(6):26–36, June 2000.
- [28] A. van Deursen and T. Kuipers. Building documentation generators.

International Conference on Software Maintenance(ICSM '99), pages 40–49, Aug. 1999.[29] P. Warren, C. Boldyreff, and M. Munro. The evolution of websites. *7th International Workshop on Program Comprehension(IWPC '99)*, pages 178–185, 1999.Proceedings of the Fifth IEEE International Workshop on Web Site Evolution (WSE'03)0-7695-2016-2/03 \$17.00 © 2003 IEEE View publication