

Contemporary Research in E-Marketing Volume 1

Sandeep Krishnamurthy
University of Washington, Bothell, USA



IDEA GROUP PUBLISHING

Hershey • London • Melbourne • Singapore

Acquisitions Editor: Mehdi Khosrow-Pour
Senior Managing Editor: Jan Travers
Managing Editor: Amanda Appicello
Development Editor: Michele Rossi
Copy Editor: Bernard J. Kieklak, Jr.
Typesetter: Jennifer Wetzel
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Idea Group Publishing (an imprint of Idea Group Inc.)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@idea-group.com
Web site: <http://www.idea-group.com>

and in the United Kingdom by
Idea Group Publishing (an imprint of Idea Group Inc.)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 3313
Web site: <http://www.eurospan.co.uk>

Copyright © 2005 by Idea Group Inc. All rights reserved. No part of this book may be reproduced in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher.

Library of Congress Cataloging-in-Publication Data

Contemporary research in e-marketing / Sandeep Krishnamurthy, editor.
v. <1 > cm.

Includes bibliographical references and index.

ISBN 1-59140-372-3 (h/c) -- ISBN 1-59140-373-1 (s/c) -- ISBN 1-59140-374-X (eisbn)

1. Internet marketing. I. Krishnamurthy, Sandeep, 1967-
HF5415.1265.C67 2005
658.8'72--dc22

2004016288

British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Chapter XI

Using Server Log Files and Online Experiments to Enhance Internet Marketing

Charles F. Hofacker
Florida State University, USA

Jamie Murphy
University of Western Australia, Australia

Abstract

Unlike most traditional media, the Internet is both digital and interactive. Here we do not simply refer to interactions between consumers and a Web site or e-mail, but also between the marketer and the firm's Web site or e-mail. Furthermore, the digital nature of the Internet records every interaction. These two characteristics — interactivity and digitization — facilitate research possibilities that would be cumbersome and costly using earlier media such as print, radio and television. On the Internet, marketers receive instant feedback on any tactical decision in the form of server log data. We believe that due to technical hurdles, both practitioners and academics under-utilize

this omnipresent data residing in server log files. This is unfortunate for practitioners because their online efforts are far less efficient and effective than they could be. This is also unfortunate for academics because even small sites can generate massive amounts of rich data in relatively short times. This chapter introduces readers to server log files and how the basic information in these files helps management achieve goals for their Web sites and e-mail communication. Next, the chapter uses examples to illustrate how server log files make running online experiments easier than one would expect. The chapter closes with a call for more use of server log files in interdisciplinary research, and collaboration between industry and academia.

Introduction

Industry (Newell, 2000) and academics (Hoffman & Novak, 1996; Kotler, Jain, & Maesincee, 2002) argue that interactivity makes the Internet unique and more powerful than the pre-wired mass media. This interactivity extends beyond consumer interaction with Web sites and incoming e-mail, to include marketers interacting with the firm's Web site and outgoing e-mail. The Internet gives marketers direct and instant feedback on tactical decisions for their Web site and outgoing e-mail. Unlike most marketing, Internet marketing offers a fast-paced cycle of optimizing based on objective feedback.

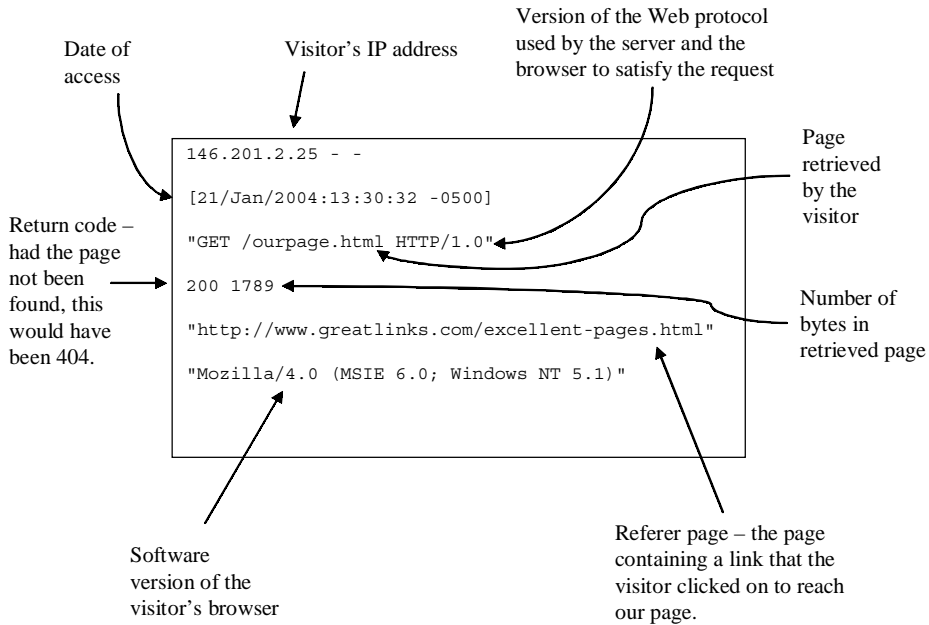
Direct Internet connections, especially Web and e-mail, support this fast-paced cycle and contrast markedly with traditional mass media. Internet channels flow directly between the marketer and the audience — without third party intermediation from a broadcaster, network, entertainer, publisher or billboard owner. In addition, this software-based channel can track and store communication with each consumer. In stark contrast, marketers can attempt to track what happens when a particular mass media advertisement travels towards a particular consumer, but it is an extremely expensive process.

Similarly, Internet sales channels can flow directly between marketers and consumers without intermediary wholesalers or retailers (Alba et al., 1997). In theory, manufacturers can track the activities of individual consumers, but once again, it is an expensive process. In contrast, Web servers automatically and inexpensively track each encounter, even encounters that do not end in a sale. Figure 1 illustrates the typical information contained in a server log file.

The direct Internet marketer-to-consumer connection parallels direct marketer-to-consumer connections (Hofacker & Murphy, 1998; Marinova, Murphy, & Massey, 2002; Novak & Hoffman, 1997; Tizende, Smith, & Murphy, 2002). The direct marketing field has long noted the ease that one can test everything from envelope formatting to offer pricing (Kotler, 2003). The Internet offers millions of opportunities to do essentially the same thing, along with added benefits such as automatic data entry and real-time feedback.

Server log data provide this instant, direct and automatic Internet feedback. While authors promote mining huge quantities of online data (Cooper & Giuffrida, 2000; Newell, 2000; Scharl, 2000), log data are also invaluable for fine grained analysis with little data

Figure 1. Sample log file entry



and a priori hypotheses (Burton & Walther, 2001; Hofacker & Murphy, 1998, 2000; Marinova et al., 2002; Murphy, 1999; Murphy & Hofacker, 2003; Murphy, Hofacker, & Bennett, 2001; Tizende et al., 2002). Despite the potential of log data though, most marketers are content with a cursory glance at simple canned reports of their Web site activity (Sterne, 2002).

Several technical hurdles, albeit surmountable, mystify practitioners and academics and thereby limit Web log data analysis. This is unfortunate for practitioners because their online efforts are far less efficient. For academics, even small Web sites can generate large amounts of rich data in short periods. Emerging digital and interactive applications, such as mobile phones (Barwise & Strong, 2002; Newell & Lemon, 2001) and interactive TV, highlight the importance of analyzing these new interactive and digital data sources.

This chapter de-mystifies Web site log files and calls for academics and practitioners to perform experimental research using this data. The chapter begins by placing online marketing and log data into context with four classifications of online marketing: communicating, selling, providing content and providing a service (Hofacker, 2001).

Classifications of Online Marketing

Communicating

An important use of Web sites and marketer-initiated e-mail is communication, and more specifically, promotion. As with the mass media, the online marketing expense is justified because it presumably yields increased online and off-line sales. Unlike other media however, log files inexpensively help uncover relationships between communications and sales.

While mass media use unidirectional communication, marketer-to-consumer, software-based Internet communication is more flexible (Hoffman & Novak, 1996; Kotler et al., 2002). Customers communicate with firms and firms communicate with customers. Firms also encourage customer-to-customer dialogue in order to tap the power of online word-of-mouth communication (Chiou & Cheng, 2003; Ha, 2002). This enhanced and expanded communication via the Internet is ideal for nurturing customer relationships (Barwise, Elberse, & Hammond, 2002; Holland & Baker, 2001; Newell, 2000; Reichheld & Shefter, 2000; Romano, 2003; Romano & Fjermestad, 2002).

Due to the Internet's accountability, one can measure how an e-mail or Web site manages the customer relationship, if not the ultimate relationship — sales. Success measures for a communication site, indeed most Web sites, are paradoxical (Johnson, Bellman, & Lohse, 2003). The longer the visitor stays on the site, the better the relationship (Ilfeld & Winer, 2002; Müller & Chandon, 2003). Yet experience and familiarity with a site usually lead to proficient and satisfied visitors spending less time on the site (Johnson et al., 2003). Thus the research focus, given management goals of communication or promotion, includes navigation paths through the site, time on site and site depth, i.e., the number of pages a visitor views.

Selling

Firms also use the Internet for direct sales, especially in the business-to-business arena (Barwise et al., 2002; Hanson, 2000). Compared to communication where a sticky site — visitors enjoy the site and stay a long time — enhances the relationship, selling requires a more focused approach. Selling sites filter visitors.

Assuming that visitors arrive on the home page and need three clicks to purchase a product, some percentage of visitors clicks on the first link while the rest exit. Some of the survivors then click on the second link, and so forth. In effect, a good selling site has a high conversion efficiency (Berthon, Pitt, & Watson, 1996; Hanson, 2000) at each of the key links that filter visitors. Here the research focuses on these key links, and on the pages that form the context for them. The page's context influences the click-through rate on these links. Rather than time on site, for selling sites the researcher concentrates on the click-through rates of key links.

Providing Content

The previous two sections implied that the firm delivers off-line goods or services. In other cases, the site itself — a content provider — is the product. An online news service such as slate.com generates revenue from exposing visitors to ads or having visitors click on the ads. The latter case presents a dilemma to the firm: to earn revenue the visitor must leave the firm's site (Hofacker & Murphy, 2000).

For content sites, the objectives may resemble those of a selling site, specific clicks. The goal is to generate revenue by inducing the visitor to click on ad banners. If the firm receives revenue for exposing visitors to ads, the more pages seen, the more revenue accrued. In this latter case, the goal is to optimize site depth and pages viewed.

Providing a Service

A final category of site is, like the content site, virtual. Rather than content, service-based sites provide networks for users. For example, eBay <www.ebay.com> is a virtual meeting space for buyers and sellers. Similarly, search engines such as Google <www.google.com> provide a mechanism for site owners to register their site in order to attract visitors, and for visitors to find sites. In both examples, the site makes the market. With eBay, the market is between sellers and buyers of myriad, oftentimes hard-to-find objects. Google makes a market between site owners seeking visitors and visitors seeking relevant sites.

A different example of a service-based site is Hotmail <www.hotmail.com>, a firm that provides a platform for visitors to use e-mail and other virtual services. The goals for these three types of sites vary. eBay would optimize the number of transactions, which have a direct relationship to eBay's revenue. The other sites would have goals similar to a content provider; revenue stems from hosting advertising and relates to pages viewed or ads clicked.

In summary, several site-based metrics should interest the Web master, including clicks on particular links, time on site, and how many pages visitors request. Thanks to the Internet's digital nature, the Web server software that dishes out pages in response to user requests automatically stores such information. Marketers can access these data to evaluate visitor behavior on individual links, particular pages, or an entire site.

The Nature of Web Log Data

Measurement industries have sprung up around each mass medium. Traditional mass media's asymmetric and indirect information flow, though, makes measurement difficult. For example, to measure television viewers, diaries or electronic meters track what station is on at what time. Connecting this information to purchases requires expensive and customized experimentation. The Internet, in contrast, is both interactive and digital. The former property implies that information also flows from the consumer to the marketer and

the latter implies easily recording such information. Web site software automatically records these same visits and transactions.

A growing number of authors have written, in glowing terms, about using data from interactive media (Ansari & Mela, 2003; Burton & Walther, 2001; Cutler, 1990; Drèze & Zufryden, 1997; Garofalakis, Kappos, & Makris, 2002; Hofacker & Murphy, 1998, 2000; Johnson, 2001; Murphy et al., 2001; Novak & Hoffman, 1997). Using actual visitors to a site optimizes external validity and the high degree of control afforded by a computer-mediated environment guarantees a high level of internal validity. Furthermore, the use of behavioral data avoids the errors and biases of self-reporting (Alba & Hutchinson, 2000; Blair & Burton, 1987; Lee, Hu, & Toh, 2000; Nisbett & Wilson, 1977).

Variables Captured by the Server

When someone requests a Web page, the Web server records information about the computer that requested the page and the files for that page. See Figure 1 and the *Passive Observation of Log Data* section, for a sample log file entry. The information logged includes:

- Internet Protocol (IP) address of the requesting computer
- domain name of the requesting computer
- date and time of the request
- requested URL, or address of the Web page
- return code or status of the request
- bytes returned, or the total amount of data sent to the requesting computer
- referring URL, or previous page that the requesting computer visited
- browser version that the requesting computer used

Which variables appear, and their sequence in the log file, depends on how the Web master configures the server, and whether the site uses Microsoft's Internet Information Server® or the open source and freely available Apache server from Apache.org. Regardless, an important field is the return code or status of the request. Generally, the researcher should keep and analyze only those requests with a return code beginning with "2", meaning that the server successfully sent the file to the visitor.

A good overview of Web log data appears in the book by Stout (1997). The following section highlights measurement issues for interpreting such data, from the marketing practitioner's and academic's points of view. Such issues include using a local cache, identifying individual visitors, and identifying search engine robots and other automated Web page retrieval agents.

Technical Issues

- *Local Cache.* Caching is a clever technique to reduce Internet traffic and help Web pages load faster. The first time a visitor requests a page, the Web server sends that page to the visitor's hard drive. If a visitor requests the same page, the Web browser — such as Internet Explorer or Netscape Navigator — first looks for that page on the computer hard drive. Thus, the Web log data show only the first request for a page by a particular visitor and not subsequent requests, which stem from the cache when the visitor hits the browser's "back arrow."
- *Identifying Visitors.* As the fields captured in Web log data reflect, the online marketer learns more about the visitor's computer than about the visitor. Nevertheless, the IP address and the domain name associated with that IP address help identify individual visitors. Due to various technical and practical reasons however, multiple individuals may use the same IP address.

An example is the student in an Internet café using the computer next to the window to access a particular site. Later, another person sits at the same computer and accesses the same site. Furthermore, some Internet Service Providers (ISPs) allocate IP addresses dynamically. Each dial up request from a visitor using a modem, even broadband users with DSL or cable modems, may have a different IP address. In any case, the same individual may revisit a particular site using a different computer in the Internet café, or using a different IP address from their ISP. The upshot is that one cannot simply map IP addresses and individuals, except within relatively short spans of time. Even this, though, can be tricky.

Proxy servers add to the difficulty of mapping individuals and IP addresses. Internet Service Providers such as America Online (AOL) or large universities centralize the retrieval of Web pages for their users via proxy servers. If an AOL user requests a particular page, the AOL proxy server stores that page for a period of time for use by others with AOL. This speeds up AOL's network traffic, but complicates analyzing log data.

Novak and Hoffman (1997) defined a visitor as a series of logged records from the same domain name or IP address with no more than a 30-minute gap between consecutive Web log records. Reducing this gap to 20 minutes and using only the first three fields of the IP address — i.e., 199.168.1 instead of 199.168.1.292 — helps address identifying visitors and the problem of proxy servers.

Cookies, however controversial and intrusive, help track user visits to a site over days, weeks or months by eliminating the ambiguities discussed above. A cookie is an identifying file that the Web site's server places on the user's hard drive. In the typical case, a CGI Script, that is a computer program that interfaces with the Web server, writes a line such as:

Set-Cookie: Customer=18174

Using this simple mechanism, the server can recognize when the same computer returns by re-reading the cookie that it set previously. Again, server log data track computers, not people. Furthermore, cookies alone cannot extract information such as user e-mail addresses — one can only read the variables previously written by the CGI script. Requiring visitors to log in to a particular site can also identify visitors, but this process goes beyond the simple use of log file data.

- *Robots and Web Crawlers.* Another subtlety in the use of Web log data comes from robots. These automated agents that retrieve and inspect Web pages might traverse a Web site, indexing the pages for a search engine. Each robot generally has a unique “Browser Version” that identifies it. Similar to cleaning survey data or other data, filtering these unwanted Web crawlers from the log data is necessary prior to statistical analysis. According to well-accepted protocols, robots request a file called robots.txt before traversing a site. Visitors that request that file are inhuman. Collecting such identifying characteristics helps filters robots from the log data (Hofacker & Murphy, 1998, 2000; Murphy, 1999; Murphy & Hofacker, 2003; Murphy et al., 2001).

Modeling Issues

This section reviews possible criterion variables to use with log data, and models appropriate to each type of data.

- *Click-Through.* When the focus is on click-throughs for particular links, the data resemble choice data, and in particular, retail scanner data (Hoffman & Novak, 1996). As such, it is natural to use a Logit model (Guadagni & Little, 1983). When one tracks individual visitors and the data are disaggregate, each observation is one visitor’s data with a 1 (clicked) or 0 (not clicked) for the link(s) in question, one can use Maximum Likelihood. Alternatively, when not tracking individual visitors and analyzing aggregated choice probabilities, one can use either Maximum Likelihood or Weighted Least Squares.
- *Site Depth.* Figure 2 shows the frequency that visitors to a hospitality business’ communication site view one page, two pages, and so forth. In 1997, Drèze and Zufryden suggested that such data exhibit a Poisson distribution. The next year, Huberman (1998) proposed the inverse Gaussian distribution based on economic considerations, and obtained a good empirical fit with that distribution. More recently, Bucklin and Sismeiro (2003) suggested that the Type II Tobit model is better as it allows for covariates.
- *Time on Site.* One measures the time on site by calculating the difference between the visitor’s first and last page request. This calculation, though, places a lower bound on the visitor’s time “at” the site. Web servers usually do not maintain state. The visitor neither logs in nor logs off the site. The server simply reacts to individual page requests and has no way of recording if the visitor subsequently requested pages from another site, or stared at the last requested page for 20 minutes before turning to another activity.

Figure 2. Frequency distribution of site depths for a regional hospitality site

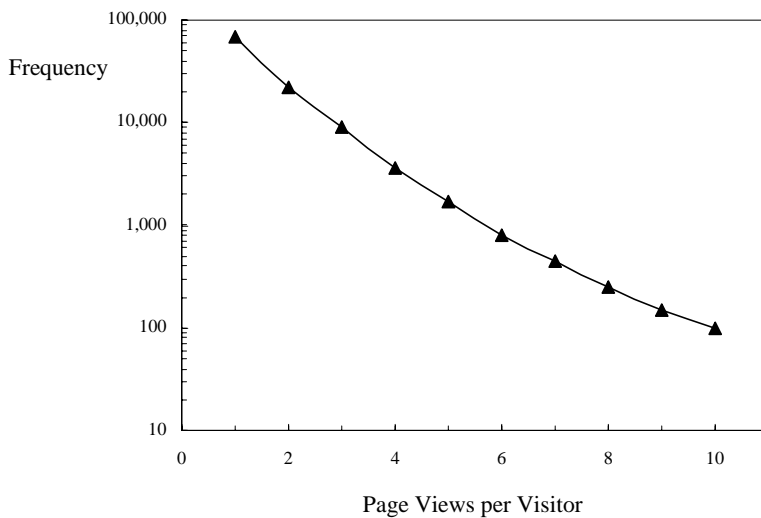
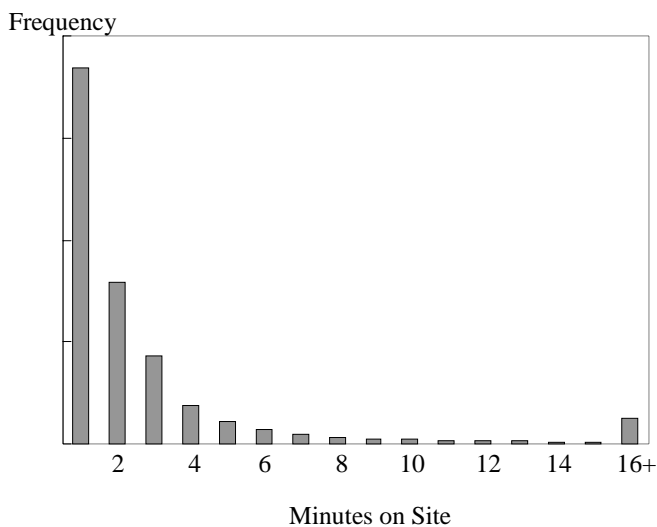


Figure 3. Distribution of time on site for a regional hospitality site



In other words, such data are subject to right censoring. As such, both Site Depth and Time on Site are amenable to modeling using the above mentioned Type II Tobit model (Bucklin & Sismeiro, 2003). Figure 3 shows that like other measures of time, the number of people accessing the site for “n” minutes drops off rapidly.

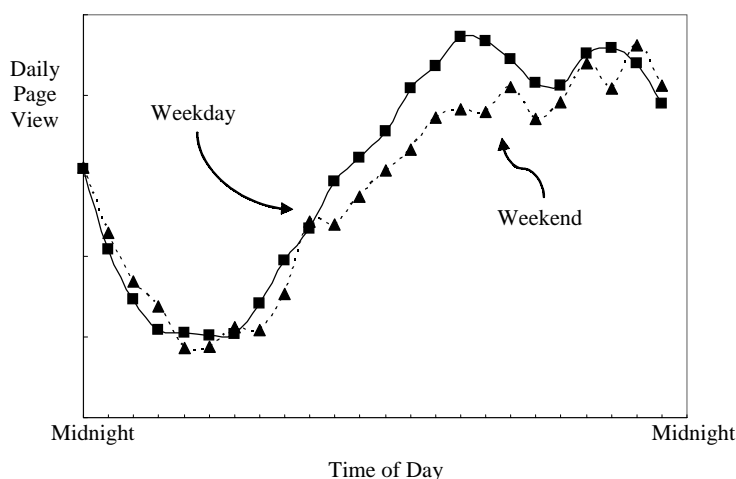

Passive Observation of Log Data

This section shows results from a variety of Web sites that demonstrate the inferences one can make from merely observing the log data. Software products such as WebTrends <www.webtrends.com> produce log file reports and one can work the raw log files as well.

Time of Day

Figure 4 shows average hourly accesses of two Web sites, by weekday vs. weekend. The site access patterns differ with the top graph being the communication site in Figure 2,

Figure 4. Hourly page views for a hospitality site (top) and for a content site (bottom)



and the bottom site being a content provider. The time axis is the time zone of the site itself. Figure 4 suggests that the two sites differ in access from work or home. It is also possible to see a dip around lunch time for both sites. Although neither of these sites is particularly busy, their graphs illustrate the possibility of collecting clean and beautiful data!

Bookmarking

Bookmarking a site (adding the site to the Favorite's list in Microsoft terminology) implies a milestone in the relationship between that visitor and the site (Garofalakis et al., 2002; Murphy & Hofacker, 2003). When this event occurs, Microsoft's Internet Explorer browser requests the file `favicon.ico`, a 16×16-pixel icon that the browser then uses for that site's entry in the visitor's favorites list. According to recent trade publications, Internet Explorer has 96% of the browser market (Murphy & Hofacker, 2003), suggesting that requests for `favicon.ico` indicate that visitors bookmarked the site.

This method works irrespective of whether the file `favicon.ico` is on the server. If the file is on the server, it creates a log entry when the server sends the file to the visitor. Otherwise, it creates a 404 File Not Found record if the `favicon.ico` file does not exist. In either case, the researcher has a valid marker that the user bookmarked the site.

How Did They Get Here?

The referrer URL shows if visitors arrived on the site through links on other Web pages or from a Web search service. If the latter, one can also determine the search terms used as these appear in the log file entry. A typical example might be:

```
192.18.209.138 - - [15/Nov/2003:10:31:29 -0500] "GET /biking.html HTTP/1.0"  
200 2444 "http://www.google.com/search?hl=en&q=fast+cyclists+in+australia"  
"Mozilla/4.0 (compatible; MSIE 6.0; AOL 8.0; Windows NT 5.1)"
```

In this case, a user entered the search term "fast cyclists in Australia" into Google, and requested the Web site's "biking.html" page.

No referring URL in the log entry also provides useful information. In addition to following links, visitors can type the URL directly into the browser, follow a previously made bookmark or have the page as the default opening page on their browser. The former, following links, suggests visitors searching for a page while the latter cases suggest familiar visitors.

A researcher could leverage referral information to assess the impact of off-line materials. For example, the firm could include a unique URL on the material, such as

www.ourcompany.com/special. The log file would show visitors who entered the site by typing that URL.

Site Navigation

On what pages do visitors start, go and finish? After identifying individual visitors, one can follow their navigation through the site. For example, the hospitality site in Figures 2 and 3 discovered that visitors rarely started on the home page (Murphy et al., 2001). Most visitors arrived on a page of maps three levels into the site and then left. Management redesigned this map page and subsequently increased pages visited. This simple example exemplifies the myriad possibilities for using basic log file data to redesign a Web site or Web page to meet management goals, the topic that follows.

Management Goals

With few exceptions (Johnson, 2001), it has been others such as information technology (IT) researchers, and not marketers, at the forefront of studying Web site or e-mail design. This is problematic. The marketing function has an external focus and plans online experiences that further company goals. If marketing academics cede the study of consumer-Web site interaction to IT researchers, and if marketing practitioners cede control of companies' Internet marketing to IT professionals, this is a loss for marketing as a discipline and for companies. This loss grows when one considers new electronic marketing forms such as mobile phones and Interactive TV. Inter-disciplinary books such as this help bridge the gap between disciplines.

How does marketing optimize firms' online efforts? Customer acquisition and customer retention serve as two organizing principles. The former involves awareness, such as advertising, public relations, online banners, registering the site with appropriate search engines and an intuitive Web site address (Hanson, 2000; Ilfeld & Winer, 2002; Murphy, Raffa, & Mizerski, 2003). The latter, customer retention, is the heart of the forthcoming section.

Web site goals drive the customer retention strategy, whether to induce visitors to click on particular links or click on many pages. E-mail marketing has similar goals as well as a "viral" goal whereby users forward clever e-mails that the firm produces (Godin, 2001). Management goals for the site or the e-mail influence the subsequent design of the site, or e-mail communication. A brief summary of Web site and e-mail goals follows.

Realizing Web Site Goals

- *Entire Sites.* Numerous research issues and practical questions center on the organization of a site (Goldfarb, 2002; Johnson et al., 2003). A Web master has

several options for a site's menu structure. The menu could be deep, with many sublevels, or shallow with more breadth on each page. Menus could also be alphabetical, or according to other semantic principles.

The ideal site structure depends on the Web site goal. For example, should a site have a small number of large pages or a large number of smaller modules? The decision in this case hinges on whether management is trying to get visitors to notice particular links and filter visitors towards these key links, or whether management wants to maximize page views.

Although other scientists investigate these issues from cognitive (Pirolli & Card, 1999) or human-computer interaction (Scheffelaier & Vinsonhaler, 2002, 2003) perspectives, marketers would tend to approach Web site navigation based on the business goals discussed above. If the site sells advertising, the firm can draw upon the Economics of Surfing (Adar & Huberman, 2000) to build a navigational structure that increases page views, and the resulting increased advertising revenue. Yet, if the goal is to funnel visitors to specific pages, other navigational structures might be optimal.

Pirolli and Card (1999) developed a cognitive information model for Web navigation based on the analogy of a foraging animal. For such an animal, the calories spent acquiring food must be less than the calories gained by eating that food. The Web visitor, trying to capture the scent of desired information or entertainment, constantly decides whether to stay on a given site or leave for more fertile searching grounds. A site should provide a scent to the visitor, that is, provide *a priori* information about pages not yet seen. How firms provide this scent — consistent with the site's marketing goals — is an open, empirical question, answerable using the experimental methods later discussed in the chapter.

- *Individual Pages.* At the individual page level, management faces numerous decisions. Social pressures and bandwagon effects often drive management decisions to have a Web site, built on fashion and fad rather than addressing management goals (McBride, 1997; Murphy, Olaru, Schegg, & Frey, 2003). Artists or IT professionals may build pages according to creative whim, blindly deciding the number of links on a page, the type and style of links, format for information, positioning on the page of various elements and use of images. The demise of Boo.com, for example, illustrates fashion rather than function driving the design process (Stockport, Kunnath, & Sedick, 2001).

Marketers should conceptualize a Web page not as a fashion statement but as a supermarket of competing and complementary links. Click-through rates represent the market shares of links on the page. Similar to a store manager, the site manager who controls what goes on the page should optimize pages to meet company goals. In addition to optimizing individual or total clicks, another tactic is strengthening relationships by having the visitor bookmark the page (Garofalakis et al., 2002; Murphy & Hofacker, 2003).

- *Particular Links.* Web log data record clicks to individual links. It is simple to investigate how myriad independent variables pertaining to wording, layout or formatting of specific links or banner ads influence clicks.
- *Personalization.* A hallmark of the Internet is the potential for personalized interaction (Hanson, 2000; Parasuraman & Zinkhan, 2002). Personalization can increase click-throughs (Ansari & Mela, 2003) as well as build and increase loyalty (Newell, 2000; Reichheld & Shefter, 2000; Romano & Fjermestad, 2002).

Realizing E-Mail Goals

Since e-mail oftentimes resembles a Web page, many of the issues are the same. One can investigate the entire e-mail and measure total clicks from the e-mail, or investigate links in the e-mail and measure clicks on these links. In addition, other criteria of interest include actions taken to unsubscribe, and viral effects whereby recipients forward the e-mail to others.

E-mail also resembles direct mail. Unlike Web pages whereby visitors actively seek a page, e-mail comes directly to the recipient's inbox. Thus, researchers can draw upon direct mail analogies to test copy, salutations, serial positioning, number of links and e-mail format (Ansari & Mela, 2003; Marinova et al., 2002; Tizende et al., 2002). The following section illustrates how experiments help realize e-mail and Web site goals.

Active Experimentation

A marketing theory, for example, claims that a blue version of a Web page will outperform a red Web page on a particular criterion. To test this theory, when a visitor requests the Web page the server invokes a random number generator and by chance sends the visitor either the blue or the red version of the page. This methodology extends to two versions of an entire site with dozens of differing sub pages, or reduces to an experiment with two pages differing only in the wording on one link.

Experimenting on actual Web sites creates many advantages for the researcher. For one, using an actual site tends to lead to practitioner implications and managerial relevance. For another, the external validity is high as the sample consists of actual visitors to that site and the experiment is unobtrusive to those visitors (Drèze & Zufryden, 1997). Internal validity can be high as well, because the sequence of events can be computerized and subject to random assignment. To cap it off, even modest sites can generate hundreds of cases in a few days. Once the researcher sets the experiment in motion, the data collect themselves.

The ease with which one can generate large samples of visitors on a site gives researchers time to investigate generalizing across multiple sites. That is, the research investigating

red vs. blue pages might employ a variety of pages sampled from the universe of possible Internet pages. Presumably, the actual page used could be a random effect in the model used.

Technical Methods for E-Mail and Web Experimentation

Several mechanisms help execute experiments similar to those described in the preceding paragraphs. A CGI script (available from the authors) that generates the random number can log the page version sent and the visitor's IP address. After identifying unique visitors to the site, the log generated by the CGI script can then classify those visitors according to the page version they saw.

A second mechanism that does not rely on a CGI script uses a non-existent 1×1-pixel image embedded in the HTML code for each version of the page. The name of this image would vary according to the page version, for example `blue.gif` and `red.gif`. The log records will show a Return Code of 404 for `blue.gif` each time the server sends a blue version of the page, and similarly for the red page version.

Using this nonexistent image, one can forego using a random number generator. Changing the page version based on the time of day also creates experimental control. For example, assuming that the page in question is `x.html`, and that different versions of the page are `a.html` and `b.html`, one could execute a background program that runs, say, every 20 minutes and does this:

```
copy a.html to x.html  
copy b.html to a.html  
copy x.html to b.html
```

In effect, two versions of the page overwrite `x.html` every 20 minutes. One could also execute this code one extra time at midnight. Thus, each version has a different 20-minute interval from the day before, thereby counterbalancing the experimental design. Furthermore, using a different nonexistent image as discussed above — `blue.gif` vs. `red.gif` — helps track which visitor received which page version without worrying about calculating a visitor's particular 20-minute interval.

An embedded image can also track individual visitors in order to analyze aggregate choice probabilities. For a study that focuses on the click-through rate of one or more links on the experimental pages, one divides the number of people who clicked on a link by the number of people exposed to the links. A unique (nonexistent) image on each page version shows the denominator for the page versions via the log count of 404 return codes.

For the numerator, one sets up each version of the page with unique URLs. For example, to compare red and blue versions of a page, with the criterion variable being the number

of clicks on the link to “www.experiment.com/get-more-info.html,” the coded URL embedded in the red and blue versions of the page could be

www.experiment.com/get-more-info.html?red
www.experiment.com/get-more-info.html?blue

The Web server ignores anything after the question mark, as the hypertext transfer protocol ignores that part of the URL. But the Web server does log *get-more-info.html?red* and *get-more-info.html?blue* as two different pages, thereby recording the number of clicks on that link made from each page version. Dividing these numbers by the count of the non-existent images — from any standard Web log report or one available from the authors — yields the choice probabilities.

One studies e-mail messages similarly, giving the URLs embedded in the e-mail a unique code matched to the experimental condition. To test a blue page e-mail against a red page e-mail, the criterion variable is how many recipients click on a URL embedded in both e-mails, which points to the Web page, www.mycompany.com/newsletter-followup.html. The red and blue e-mails go out with these URLs, respectively:

www.mycompany.com/newsletter-followup.html?red
www.mycompany.com/newsletter-followup.html?blue

The words “red” and “blue” after the question mark permit simple log file analysis of which version yielded the most clicks. As with the Web page experiments, each e-mail version points to different nonexistent 1×1-pixel images in order to estimate how many people in each experimental condition opened the e-mail. This number represents a lower limit as some mail programs let users turn off images or preview e-mail without opening the e-mail.

Possible Pitfalls

Occasionally a visitor may observe multiple experimental conditions by using the reload button on the experimental page. Drezé and Zufryden (1997) suggested another randomization mechanism to avoid this possibility, namely to assign visitors to conditions using the last digit of the visitor’s IP address. This approach helps assure that each visitor sees only one version of the site during the experiment. The down side is that proxy servers, mentioned earlier, make it problematic to assume that all users have the same IP address — even in a short site visit. One could avoid this pitfall, however, using cookies.

Sample Empirical Results

- *Link Design and Copy Testing.* An unpublished experiment illustrates this technique. The authors compared two banner ads for a national dating service, but also tracked what the visitors who clicked on those banners did on the advertiser's site. The banners were identical in all design elements, including the name of the firm and a simple logo, except for the copy. Message A simply said "Click Here" while the longer message B said "Done with bars and time wasting classifieds? Serious about finding your soulmate?"

Table 1 shows that while message A won the click-through battle ($\chi^2 = 147.353$, $p = .001$), it lost the battle for subsequent page requests on the dating service's Web site (Poisson Regression, $|t| = 8.321$, $p < .001$). While it is tempting to pick the banner with the highest click-through probability, that banner performed worse in an absolute sense; it yielded fewer page views on the advertiser's site.

The Elaboration Likelihood Model may explain these results (Petty & Cacioppo, 1986). The short copy and more peripheral route processing achieved higher click-throughs, but the long copy and central route processing lead to significantly more clicks on the advertiser's site.

- *Opt-In E-mail Design.* E-mail newsletters and other outgoing e-mail messages resemble mass media. They are marketer, rather than customer, initiated. Yet even though e-mail is the most popular Internet application in business and at home (Ramsey, 2001), for young (Pastore, 2002) and for old (BBCNews, 2002), there seems to be little published research on this useful customer relationship management tool (Krishnamurthy, 2001).

An early experiment tested personalized vs. non-personalized greetings in a hotel e-mail campaign. Against expectations, there were significantly better results with the non-personalized greeting (Marinova et al., 2002). The authors suggested that as the hotel obtained most names at registration, there was a weak relationship and recipients may have objected to the personalization.

In 2002 unpublished study, Hofacker and Voorhees of the Florida State University performed a simple experiment that compared three versions of an e-mail newsletter for a leisure and recreation company. One version was plain text, one used basic HTML, and

Table 1. Results of ad banner click-throughs and follow up page requests

	N	Banner Click-through	Follow Up Page Requests on the Site
Message A	2498	.1137	.0246
Message B	2361	.0779	.2134

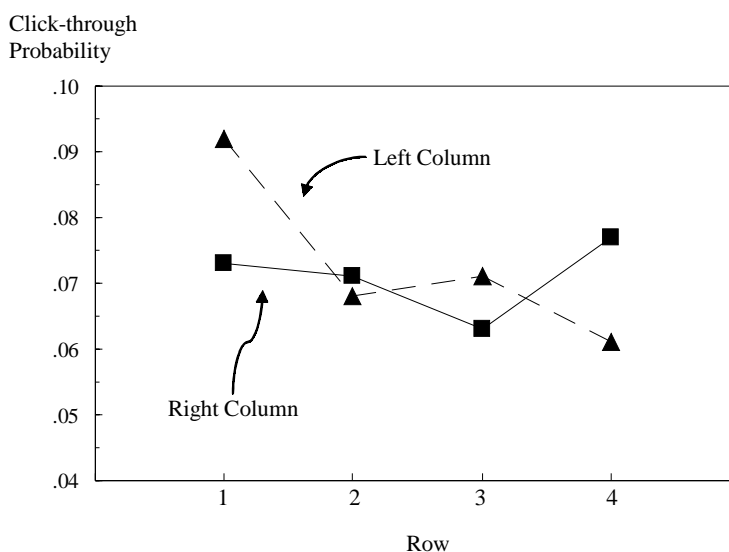
another HTML version included graphical banner ads rather than simple text links. Contrary to expectations, the text condition out-performed both HTML versions, although the effect was slight.

- *Page Design.* The authors executed several studies of click-throughs on links in simple Web menus. One of the first studies utilized a marketing practitioner site that had, on its home page, eight links arranged in a four-row by two-column menu. The first hypothesis was a monotonic relationship between row and click-through probability. In other words, the higher on the screen, the higher the click-through rate would be. Secondly, given English reading habits, left to right and top to bottom, links in the left column should garner higher click-through rates than links in the right column.

Should a Web master put the two most important links as the first two items either in the left column or as the two items in the top row?

The experiment used a Latin Square design that rotated each of the links through each of the menu positions. Figure 5 shows the empirical results. Click-throughs for the individual positions show that the upper left location was the most clicked position, as predicted, but that the lower right location was second. English language reading habits suggest reading the upper left first, and the lower right, last. These results, however, reflect a primacy and recency effect. We replicated these results on other sites, with

Figure 5. Click-through probability as a function of item position within a menu



different length menus and different menu structures. The results suggest that given a goal of increased clicks on particular links, a Web master should place the second most important link last, not second.

- *Site Design.* Hofacker and Murphy (2000) also studied pages containing one or two links in order to test two propositions concerning online choice. They used two pages, i and j, with page i containing a superset of the links on page j. In other words, page i had all the links that page j had, plus an additional link or links. The first proposition was that the total clicks for page i would not be less than the total clicks for page j. The second proposition was that individual clicks on any link included in both pages i and j would not be less for page j than for page i. The data confirmed these simple notions.

Conclusions and Call for Future Research

The number of unique domain names with hosts, almost two hundred million in January of 2003 <www.isc.org/ds/>, suggests a lower estimate of the number of Web sites. Similarly, the popular search engine Google's index of over four billion Web pages in 2004 serves as a lower estimate of Web pages. These online repositories have at least two things in common. One, minute by minute a Web server is automatically logging visits and thereby creating a valuable data set for improving that site or page. Two, each site owner is a potential partner for collecting data that might be of theoretical interest to marketing academics and of practical interest to the site owner.

In our opinion, much of the current electronic marketing research attempts to recreate familiar mass media situations. For example, Internet user panels track how a particular individual navigates online, going from one site to another. Other research uses surveys to capture the attitudes of Web visitors. Such data are certainly valuable, albeit expensive.

We argue however, that most marketing researchers are missing an important Internet component, namely its direct and interactive nature. The marginal cost of acquiring and analyzing Web log data from individual sites is essentially zero. These data are valuable enough as is, describing what individuals do on that site. Combined with experimentation on the site, the data become markedly more valuable to both the practitioner and academic.

Rather than argue or guess what might work best online, practitioners, along with academics armed with the appropriate theory, could run experiments. Similarly, industry already running these experiments could discover additional insights by sharing the results with academia. There are numerous research possibilities.

At the level of individual links, experiments could test wording, the icon or image that anchors the link, text vs. image anchors, the type of appeal, and a host of other psycholinguistic and format issues. At the page level, we know little about the effectiveness of simple format variables and the appearance of pages, supplementary images and text that might provide a context for the links on the page, the position, number and type of links, and how one link might cannibalize another.

Finally, at the level of the whole site, productive research could investigate site structure, and how to best present that structure to the visitor. Issues pertaining to the former include modularization vs. larger pages, semantic distinctions to define sub-pages, the information space (hierarchical, linear, serial, factorial, or semantic network) and menu structure. Issues pertaining to the latter include navigation bar content and format, representing relationships between pages, and how visitors receive feedback on where they are, where they have been and where they can go.

Live Web or e-mail experiments can test virtually every consumer marketing theory, as well as theories in other disciplines. Possible theories could deal with exposure, perception, cognition, scent, learning, persuasion, elaboration, reasoning, memory, choice, heuristics, risk, flow, influence and compliance, motivation, culture, communication, hedonic shopping, economic rationality, pricing, effort and commitment, and judgment.

With even a modest Internet site, and a partnership with the site owner, academics can create multiple versions of the site, sit back, and let the data roll in. The winning experimental condition gives the practitioner a better site and gives the academic a valid test of the theory that suggested it.

References

- Adar, E., & Huberman, B. A. (2000). The economics of surfing. *Quarterly Journal of Electronic Commerce*, 1(2), 203-214.
- Alba, J. W. et al. (1997, July). Interactive Home Shopping: Incentives for Consumers, Retailers, and Manufacturers to Participate in Electronic Markets. *Journal of Marketing*, 61, 38-53.
- Alba, J. W., & Hutchinson, J. W. (2000, September). Knowledge calibration: What consumers know and what they think they know. *Journal of Consumer Research*, 27, 123-156.
- Ansari, A., & Mela, C. F. (2003, May). E-customization. *Journal of Marketing Research*, XL, 131-145.
- Barwise, P., Elberse, A., & Hammond, K. (2002). Marketing and the Internet: A research review. Retrieved 30 January, 2004 from: www.london.edu/marketing/Future/future_media_research_projects/Untitled/internet/MATI_1.3.pdf
- Barwise, P., & Strong, C. (2002). Permission-based mobile advertising. *Journal of Interactive Marketing*, 16(1), 14-24.

- BBCNews. (2002). Elderly take up surfing. Retrieved 30 January, 2004 from: news.bbc.co.uk/hi/english/sci/tech/newsid_1977000/1977823.stm
- Berthon, P., Pitt, L. F., & Watson, R. T. (1996). The World Wide Web as an advertising medium: Toward an understanding of conversion efficiency. *Journal of Advertising Research*, 36(1), 43-53.
- Blair, E., & Burton, S. (1987, September). Cognitive processes used by survey respondents to answer behavioral frequency questions. *Journal of Consumer Research*, 14, 280-288.
- Bucklin, R. E., & Sismeiro, C. (2003, August). A model of Web site browsing behavior estimated on clickstream data. *Journal of Marketing Research*, XL, 249-267.
- Burton, M. C., & Walther, J. B. (2001). The value of Web log data in use-based design and testing. *Journal of Computer Mediated Communication*, 6(3). Retrieved from: www.ascusc.org/jcmc/vol6/issue3/burton.html
- Chiou, J.-S. C., & Cheng, C. (2003). Should a company have message boards on its Web sites? *Journal of Interactive Marketing*, 17(3), 50-61.
- Cooper, L. G., & Giuffrida, G. (2000). Turning datamining into a management science tool. *Management Science*, 46(2), 249-264.
- Cutler, B. (1990). The fifth medium. *American Demographics*, 12(6), 24-49.
- Drèze, X., & Zufryden, F. (1997). Testing Web site design and promotional content. *Journal of Advertising Research*, 37(2), 77-91.
- Garofalakis, J. G., Kappos, P., & Makris, C. (2002). Improving the performance of Web access by bridging global ranking with local page popularity metrics. *Internet Research: Electronic Networking Applications and Policy*, 12(1), 43-54.
- Godin, S. (2001). *Unleashing the idea virus*. New York: Hyperion.
- Goldfarb, A. (2002). Analyzing Website choice using clickstream data. In M. R. Baye (Ed.), *Advances in applied microeconomics, the economics of the Internet and e-commerce* (vol. 11, pp. 209-230): Elsevier Science Ltd.
- Guadagni, P. M., & Little, J. D. C. (1983). A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3), 203-238.
- Ha, H. Y. (2002). The effects of consumer risk perception on pre-purchase information in online auctions: Brand, word-of-mouth, and customized information. *Journal of Computer-Mediated Communication*, 8(1). Retrieved from: www.ascusc.org/jcmc/vol8/issue1/ha.html
- Hanson, W. (2000). *Principles of Internet marketing*. Cincinnati, OH: Southwest College Publishing.
- Hofacker, C. F. (2001). *Internet marketing* (3rd ed.). New York: John Wiley & Sons.
- Hofacker, C. F., & Murphy, J. (1998). World Wide Web banner ad copy testing. *European Journal of Marketing*, 21(7/8), 703-712.
- Hofacker, C. F., & Murphy, J. (2000). Clickable World Wide Web banner ads and content sites. *Journal of Interactive Marketing*, 14(1), 49-59.

- Hoffman, D. L., & Novak, T. P. (1996, July). Marketing in hypermedia computer-mediated environments: Conceptual foundations. *Journal of Marketing*, 60, 50-68.
- Holland, J., & Baker, S. M. (2001). Customer participation in creating site brand loyalty. *Journal of Interactive Marketing*, 15(4), 34-45.
- Huberman, B. A., Pirolli, P. L. T., Pitkow, J. E., & Lukose, R. M. (1998). Strong regularities in World Wide Web surfing. *Science*, 280, 95-97.
- Ilfeld, J. S., & Winer, R. S. (2002). Generating Website traffic. *Journal of Advertising Research*, 42(5), 49-61.
- Johnson, E. J. (2001). Digitizing consumer research. *Journal of Consumer Research*, 28(2), 331-336.
- Johnson, E. J., Bellman, S., & Lohse, G. L. (2003, April). Cognitive lock-in and the power law of practice. *Journal of Marketing*, 67, 62-75.
- Kotler, P. (2003). *Marketing Management* (11th ed.). Upper Saddle River, NJ: Prentice Hall.
- Kotler, P., Jain, D. C., & Maesincee, S. (2002). *Marketing moves*. Boston: Harvard Business School Press.
- Krishnamurthy, S. (2001). A comprehensive analysis of permission marketing. *Journal of Computer Mediated Communication*, 6(2). Retrieved from: www.ascusc.org/jcmc/vol6/issue2/krishnamurthy.html
- Lee, E., Hu, M. Y., & Toh, R. S. (2000, February). Are consumer survey results distorted? Systematic impact of behavioral frequency and duration on survey response errors. *Journal of Marketing Research*, 37, 125-133.
- Marinova, A., Murphy, J., & Massey, B. (2002). Permission e-mail marketing as a means of targeted promotion. *Cornell Hotel and Restaurant Administration Quarterly*, 43(1), 61-69.
- McBride, N. (1997). Business use of the Internet: Strategic decision or another bandwagon? *European Management Journal*, 15(1), 58-67.
- Müller, B., & Chandon, J.-L. (2003). The impact of visiting a brand Website on brand personality. *Electronic Markets*, 13(3), 210-221.
- Murphy, J. (1999). Surfer and searchers: An examination of Web-site visitor's clicking behavior. *Cornell Hotel and Restaurant Administration Quarterly*, 40(2), 84-95.
- Murphy, J., & Hofacker, C. F. (2003). A preliminary investigation of an e-loyalty metric, bookmarking. Paper presented at the *World Marketing Conference*, Perth, Australia.
- Murphy, J., Hofacker, C. F., & Bennett, M. (2001). Website generated market-research data: Tracing the tracks left behind by visitors. *Cornell Hotel and Restaurant Administration Quarterly*, 42(1), 82-91.
- Murphy, J., Oлару, D., Schegg, R., & Frey, S. (2003). The bandwagon effect: Swiss hotels' Website and e-mail management. *Cornell Hotel and Restaurant Administration Quarterly*, 44(1), 71-87.

- Murphy, J., Raffa, L., & Mizerski, R. (2003). The use of domain names in e-branding by the World's top brands. *Electronic Markets*, 13(3), 30-40.
- Newell, F. (2000). *Loyalty.com: Customer relationship management in the new era of Internet marketing*. New York: McGraw Hill.
- Newell, F., & Lemon, K. N. (2001). *Wireless rules - New marketing strategies for customer relationship management*. New York: McGraw-Hill.
- Nisbett, R., & Wilson, T. (1977, May). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231-259.
- Novak, T. P., & Hoffman, D. L. (1997, Winter). New metrics for new media: Toward the development of Web measurement standards. *World Wide Web Journal*, 2, 213-246.
- Parasuraman, A., & Zinkhan, G. M. (2002). Marketing to and serving customers through the Internet: Conceptual frameworks, practical insights, and research directions. *Journal of the Academy of Marketing Science*, 30(4), 286-534.
- Pastore, M. (2002). Internet key to communication among young. Retrieved 30 January, 2004 from: cyberatlas.internet.com/big_picture/demographics/article/0,,5901_961881,00.html
- Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. New York: Springer-Verlag.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106(4), 643-675.
- Ramsey, G. (2001). The e-mail marketing report, executive summary. Retrieved 30 January, 2004 from: www.the-dma.org/cgi/registered/whitepapers/eMail_Market_Exec_Sum_SS.pdf
- Reichheld, F., & Shefter, P. (2000). E-loyalty: Your secret weapon on the Web. *Harvard Business Review*, 78(4), 105-114.
- Romano, N. C., Jr. (2003). Customer relationship management for the Web-access challenged: Inaccessibility of *Fortune* 250 business Web sites. *International Journal of Electronic Commerce*, 7(1), 81-117.
- Romano, N. C., Jr., & Fjermestad, J. (2002). Electronic commerce customer relationship management: An assessment of research. *International Journal of Electronic Commerce*, 6(2), 61-113.
- Scharl, A. (2000). *Evolutionary Web development*. London: Springer.
- Scheffelmaier, G. W., & Vinsonhaler, J. F. (2002-2003). A synthesis of research on the properties of effective Internet commerce Web sites. *Journal of Computer Information Systems*, 43(2), 23-30.
- Sterne, J. (2002). *Web metrics: Proven methods for measuring Web site success*. New York: John Wiley and Sons.
- Stockport, G. J., Kunnath, G., & Sedick, R. (2001). Boo.com - The path to failure. *Journal of Interactive Marketing*, 15(4), 56-70.
- Stout, R. (1997). *Web site stats: Tracking hits and analyzing traffic*. Berkeley, CA: Osborne McGraw-Hill.

Tizende, T., Smith, B., & Murphy, J. (2002). Getting permission: Exploring factors affecting permission marketing. *Journal of Interactive Marketing*, 16(4), 28-36.