ICT-MobileSummit 2009 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2009
ISBN: 978-1-905824-12-0

Guidelines for Mobile Web Content Adaptation by Third-party Proxies

François Daoust[1]
[1]*W3C/ERCIM, 2004 route des Lucioles, Biot, 06410, France*
Tel: +33 4 92 38 79 47, Fax: + 33 4 92 38 78 22, Email: fd@w3.org

Abstract: The Mobile Web is a world of constraints: limited memory, limited screen sizes, limited input methods, limited availability of the user on the go, high latency and low bandwidth networks. Some mobile network operators deployed Content Transformation solutions onto their mobile networks to give their users access to the "long tail" of legacy Web pages that were designed for desktop browsers and that do not render well on mobile devices. Content Transformation is actually envisioned both as an extension set of functionalities to mobile devices browsers and as a way to improve the overall user experience while browsing the Web on a mobile device. In practice, Content Transformation done by third-party proxies creates a number of issues that threaten the Web neutrality and in some cases break the actual delivery of Web content to end users. The W3C is working on a set of Content Transformation Guidelines to ensure that the impact of conforming deployments on the mobile ecosystem stays positive. This paper presents the potential and limitations of content transformation when carried out by third-party proxies, and how the guidelines proposed by the W3C Mobile Web Best Practices working group intend to solve the issues created by existing solutions. The work on the guidelines is on-going and some guidelines may change in the future. The work already helped identify potential topics for future works, in particular the need for explicit semantics that could be used by content providers to tag their content and communicate their intent and expectations with Content Transformation proxies.

Keywords: content adaptation, proxy, mobile, Web, best practices, guidelines, W3C, standardization.

1. Introduction

Mobile devices are increasingly being used to browse the Web. The multiplicity of existing models, each with different capabilities and software, combined with the intrinsic limitations of mobile networks infrastructures and the presence of different generations of devices in the market, turn the Mobile Web ecosystem into a fragmented-and-reserved-for-mobile-experts-mine-field world. The W3C Mobile Web Best Practices Working Group developed a set of Mobile Web Best Practices [1] to help Web content authors take mobile Web constraints into account while developing content, and provide a good user experience on a vast majority of mobile devices.

However, these Best Practices de facto cannot apply to the long tail of legacy Web sites. To enable browsing of Web pages that were not designed with mobile constraints in mind, a few Mobile Networks Operators started to deploy Content Transformation proxies onto their networks [2] [3] as shown in Figure 1, that automatically adapt Web pages for the end user so that they may be rendered on his mobile device.

ICT-MobileSummit 2009 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2009
ISBN: 978-1-905824-12-0

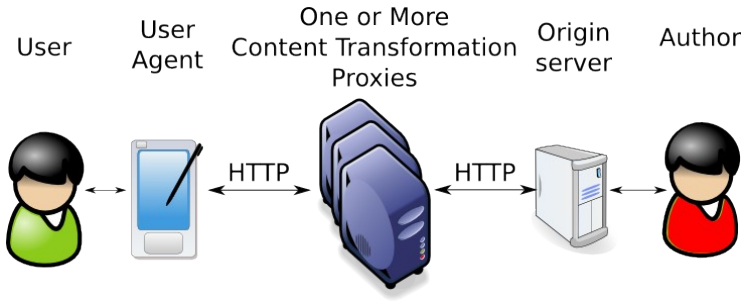

Figure 1: Content Transformation proxies in mobile networks

While this approach certainly increases the percentage of the Web accessible from mobile devices, effectively converting limited devices into first-class Web browsers, it also creates a number of problems (e.g. security, device identity spoofing, disambiguation between browser and non-browser HTTP messages) that may break mobile Web content delivery and mobile Web-based applications.

This paper presents the potential and limitations of content adaptation when performed by third-party transcoding proxies, and focuses on the on-going efforts of the W3C Mobile Web Best Practices Working Group to standardize Content Transformation Guidelines [4] that help address the limitations and preserve the Web neutrality. It is intended to raise awareness on the difficulty to find an acceptable balance between providing an innovative service to end users and respecting the work of content providers. In particular, the paper is not intended to introduce any innovative research in the domain of content adaption.

Part of the points raised in this paper apply to other in-the-middle content adaptation solutions such as search engine transcoders (e.g. Google's mobile transcoder) and lightweight Web browsers clients that interact with proprietary proxies to render Web pages (e.g. Opera Mini, Skyfire).

In the rest of the paper, Content Transformation is defined as the manipulation, in various ways, by proxies, of requests made to and content delivered by an origin server with a view to making it more suitable for mobile presentation.

Section 2 introduces the motivations behind Content Transformation deployments. Section 3 focuses on the main issues and intrinsic limitations of Content Transformation. Section 4 presents the solutions envisioned by the Content Transformation Guidelines to address the issues. Section 5 takes a look at possible future works in the area.

2. Content Transformation – Motivations

Content Transformation proxies can carry out a wide variety of operations, that may be divided into three different classes:

- Restructuring operations, whereby the original layout is altered so that content is added or removed or where the spatial or navigational relationship of parts of content is altered. It also includes rewriting of links so that subsequent requests route via the proxy. In practice, linearization and pagination of long pages into smaller chunks are commonly performed by Content Transformation proxies. Another example of a restructuring operation is when proxies detect and replace tables that are used to control the layout of the content of the page. Many mobile devices do not render tables as desktop browsers do, or the rendering may result in a poor user experience where the user needs to scroll horizontally as well as vertically.
- Recoding operations, whereby the layout of the content remains the same, but details of its encoding may be altered. Examples include re-encoding HTML as XHTML, correcting invalid markup in HTML, conversion of images between formats, re-

ICT-MobileSummit 2009 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2009
ISBN: 978-1-905824-12-0

encoding of the content using a character set that the mobile browser supports, and language translation. Conversion between formats is not restricted to images but includes audio, video, images, and mere text files.

- Optimizing operations, which include removing redundant white space, re-compressing images (without loss of fidelity) and compressing for transfer. Retrieving 100Kb of data on mobile networks may easily take 1 minute. Optimizing operations reduce the time it takes for a mobile device to render pages.

The three classes of operations serve the same purpose. They enhance and complete the capabilities of a given mobile browser so that it may render as many Web pages as possible, and so that the rendering is optimized.

The benefits for end users are immediate: more Web sites are accessible, Web pages are rendered in less time, the overall browsing experience is better, and final cost is lower.

3. Content Transformation – Limitations

Content Transformation creates a number of issues and concerns. Some are intrinsic technical limitations (e.g. security issues). Some are triggered by the lack of semantics to tell content that may need to be transcoded from content that should not. Most of them are grounded on the usual expectation from endpoints of a communication channel (content providers and end users) that third party components should be transparent.

3.1 Who is in Control?

Figure 2 shows a schematic representation of the HTTP work-flow between an end user and the author of a Web site when one or more Content Transformation proxies are deployed in the network. Examples of existing Content Transformation solutions are listed below the entity that effectively controls the Content Transformation process. Content Transformation proxies are embedded in other services, hardly visible by the end user and the Web content author. There is no established means to control Content Transformation solutions.

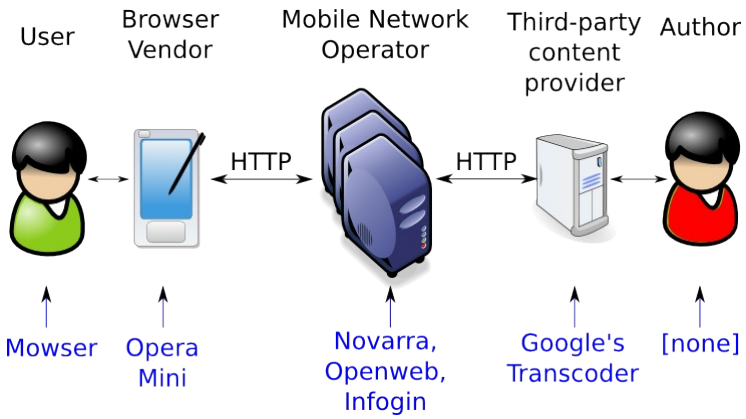

Figure 2: Content Transformation workflow and control

Who should be in control? Users on the Web may already use browsers' extensions on their desktop PCs to disguise their identity using User-Agent switchers and tweak rendered content to their liking using e.g. Greasemonkey [5]. While such tweaks are reserved to a certain class of “power” users for the time being, the question is worth asking: do users have the rights to alter content providers' content? Should the content provider be the only one in control?

3.2 Copyrights Infringements

What are the transformations allowed on copyrighted material returned by a content provider? From the point of view of Web content authors, slight alteration in the presentation should be expected. Two different Web browsers may not render the same page exactly the same way, depending on their support for various functionalities, including HTML versions and CSS properties. In a way, Web browser clients have always adapted Web content to render it. Where is the limit between transformations that is considered appropriate and transformations that should not take place without the consent of the content provider?

It is easy to go too far and box the user's browsing experience. The result is not the Web as we know it anymore, but a framed representation of the Web. This type of situation is not entirely new in the history of the Web. Back in 1997, the Washington Post, Time, Cable News and Reuters New Media filed a complaint [6] against TotalNews for a similar case: TotalNews was framing the content of the plaintiffs' Web sites and displaying its own ads in the top frame. The case was settled out of court with an agreement to stop the framing, and a general expectation that this would not happen again.

3.3 Browser Requests vs. Non-Browser Requests

When a Content Transformation proxy is deployed in a mobile network, all the HTTP messages go through it. There is nothing that identifies an HTTP message as intended for the Web browser itself. An HTTP response may target a third-party application installed on the device that expects the response in a precise format and structure. In-browser Mobile Web applications may also use an XMLHttpRequest [7] object to send HTTP requests and process HTTP responses in the background. The processing of the HTTP responses usually involves parsing it in search of a particular fragment. Parsing will not work e.g. when only part of the response is received because an intermediary proxy paginated the response.

From the point of view of the proxy, all HTTP messages are alike. HTTP header fields may contain information that can positively be used to determine that a request does not originate from the Web browser, but there is no general solution.

3.4 Identification of the Request Sender

User agent sniffing, where an origin server adapts the content it returns based on the user agent that sent the request, and user agent spoofing, where the user agent pretends it is something else to be returned a "more complete" representation of the content, have been around for about as long as the Web itself. The so-called Web browser war ended up with the need for Web authors to rely on tips and tricks to parse an HTTP User-Agent header field and return versions of the content tuned for specific mainstream desktop Web browsers. Mobile Web developers also heavily rely on the HTTP User-Agent header field to identify the requesting device and adapt the content accordingly.

The main focus of Content Transformation is legacy Web sites. When such Web sites are presented with an unknown User-Agent string (typically most mobile devices User-Agent strings fall in that category when legacy Web sites are being considered), they often reply with an error-like message instead of the real content. Some Content Transformation Proxies fix the issue by replacing the User-Agent string in the HTTP request by a well-known desktop User-Agent string. When this happens in a request sent to a mobile-aware Web site, the Web site is “tricked” into thinking that a desktop browser sent the request and will not serve the mobile optimized version of the content.

3.5 Security Concerns

One of the main security measures on the Web to protect and allow the use of client-side scripting is the same origin policy. It is implemented by all Web browsers and effectively prevents a malicious script on a Web page to access other resources on a different origin (e.g. a script in http://example.org/ cannot retrieve a page on http://example.com/ because "example.org" and "example.com" are two distinct origins).

Some of the Content Transformation restructuring operations that may be carried out by a proxy require that links that appear in the content delivered to a mobile device are re-written so that subsequent requests target the proxy and not the initial origin server. For instance, a link that targets http://example.org/ may have to be re-written into http://ctproxy.example.com/?goto=http://example.org/. Links re-writing change the origin of the page, opening the box for malicious scripts to access sensitive user's information.

This security concern is even stronger in the case of secure exchanges using HTTPS. By definition, HTTPS exchanges use an end-to-end connexion, and a proxy cannot see the content exchanged by the two ends. The only possibility for a proxy to provide Content Transformation services on HTTPS content is to rewrite HTTPS links in the manner mentioned above. When that is done, the HTTPS exchange is relayed by the proxy, and the proxy gains access to highly sensitive user's information such as credentials or credit card numbers. Since this possibility relies on link rewriting, a malicious script could in theory get access to the very same highly sensitive information.

Note that link rewriting does not need to happen in the case of Mobile Web browsers clients like Opera Mini that are composed of a heavy server part that uses a proprietary protocol to exchange information with a lightweight client part on the mobile device. HTTPS exchanges are de facto decrypted on the server part in these cases.

3.6 Who knows better? Who is responsible?

When a Web author carefully crafted Web content so that it may be rendered by most mobile devices, should a Content Transformation proxy be allowed to further adapt the content because it detected that there could be a way to improve the user experience on a given mobile device? The author might have considered this option and concluded that it did not improve the user experience in his view. Who knows better?

Content Transformation is not an exact science. Some Web content that actually would be rendered correctly by a mobile device may still be adapted by a proxy and the transcoded result may break the user experience on the mobile device. How many users know that a Content Transformation proxy is in place? How many users know what a Content Transformation proxy actually is? Most users would blame the final Web site owner for the broken experience, negatively impacting their trust in the Web site brand.

4. Content Transformation Guidelines

The different stakeholders in Content Transformation apparently have diverging needs:

- Mobile Web authors do not welcome the introduction of yet another layer of complexity in an already complicated and fragmented world.
- Mobile network operators want their users to access the Web at large with as good a user experience as possible, but also wants to keep control on security and on network loads.

- End users basically want things to work. They want things to work as efficiently as on regular desktop browsers. They do not want to have to worry about security matters or potential impacts on their monthly bill.

Reconciliation between diverging views is an area where standardization helps. The W3C Mobile Web Best Practices Working Group provides a forum for discussion on Content Transformation issues. It lead to the publication of Content Transformation Guidelines [4] that define a set of rules to be followed by Content Transformation proxies to ensure their impact on the mobile Web ecosystem stays positive. The work on the guidelines is on-going, triggers passionate discussions among key players of the Content Transformation scene, and some of the guidelines may be adapted in the near future.

The guidelines are based, as much as is practical, on existing technologies, and in particular on the HTTP/1.1 protocol [8]. Guidelines are as follows.

4.1 Content Transformation deactivation

Existing mechanisms do not define fine-grained means to control potential Content Transformation parameters over HTTP. The only existing control thus plays as a global switch. The HTTP "Cache-Control: no-transform" directive may be used in HTTP requests and responses to prevent any Content Transformation.

The directive needs to be used with care since it prevents all kinds of transformation, including optimization, that is often essential in mobile networks with limited bandwidths.

4.2 Content tasting

Content Transformation proxies should not change the HTTP headers in an HTTP Request unless they receive a rejected response from the server where rejected includes human-oriented messages such as "Sorry, your Web browser is not supported". In such cases, they may alter the HTTP header fields and send a second HTTP request to the origin server. This content tasting approach guarantees that the origin server receives the original HTTP Request sent by the user agent and is given a chance to return mobile-friendly content tailored for that particular user agent.

Content Transformation proxies may send altered HTTP header fields if they determine that it is likely that the origin server would return a rejected response, but should, on receipt of a Vary HTTP header field from the origin server indicating that different representations of the content are available depending on the values of modified HTTP header fields, re-issue the unaltered HTTP Request.

4.3 Transformation of Mobile Web content

Content that identifies itself as mobile should be left untouched by Content Transformation proxies, e.g. through the use of a mobile doctype such as XHTML Basic 1.1 [9]. It reveals that the Web author is at least aware of some of the mobile constraints, and took a step towards mobile friendliness.

Of course, the term "mobile" is vague to begin with. Does an iPhone version of a Web site constitute a good mobile representation for other classes of mobile devices? Content Transformation could probably be used to further improve the user experience on these other mobile devices, but it should not be done in the absence of better semantics for the Web author to flag his content as something more precise than "made for mobile".

4.4 Security guidelines

Client-Side scripting is hardly supported by mobile devices for the time being. To avoid the problems raised by changes of origin, the Content Transformation proxy could remove all traces of script from transcoded pages, effectively stopping the malicious scripting menace. However, this solution merely controls the threat, but does not suppress the initial weakness of Content Transformation deployments in that area.

Security guidelines are still being heavily discussed.

5. Outlook

Existing mechanisms do not allow Web authors to express a more fine-grained position than merely stating "do not transform". Mobile devices are unlikely to be formatted along the same lines in the future, meaning they will continue to show different input methods, different sizes and different capabilities. The possibility to describe precisely what a given content was designed for could help the interaction between Web authors and Content Transformation proxies. The Protocol for Web Description Resources (POWDER) [10] could be used in the future to express such semantics. A Content Transformation vocabulary would need to be standardized though.

More generally, content negotiation possibilities in HTTP communications may need to be improved. Transparent Content Negotiation [11] could be completed and deployed to take Content Transformation requirements into account.

On top of providing and improving ways to control Content Transformation done by proxies, the democratization of ubiquitous Web content designed from scratch to provide at least a functional user experience on as many devices as possible and limit the amount of versions to maintain should be encouraged. The Mobile Web Best Practices [1] and the associated mobileOK mark were created to bring the Mobile Web to authors without requiring that they become experts in the field. The mobileOK Checker is an easy-to-use tool to check the level of mobile-friendliness of a Web page:

http://validator.w3.org/mobile

6. Conclusion

Designing ubiquitous Web content is obviously the best way for Web authors to target as many Web-enabled devices as possible at once. Content adaptation, done directly on the server, can then be used to enhance the end user's browsing experience on specific classes of devices. However, the existence of a “long tail” of legacy Web sites that are unaware of mobile constraints cannot be dismissed. Content adaptation done by third party services may help bridge the gap and bring the entire Web to mobile devices, provided such proxies are implemented and deployed with great care.

If not, the presence of Content Transformation proxies in networks threatens the Web neutrality. Former walled gardens approaches would be replaced by an open and general Web access only in appearance.

The Content Transformation Guidelines document developed by the W3C Mobile Web Best Practices Working Group, and in particular by main actors of the Content Transformation scene, defines rules to enable the deployment of Content Transformation solutions that respect the mobile Web ecosystem and opens the path to further work in the area. The work on the guidelines should be completed by the end of 2009.

Acknowledgement

This work is part of the MobiWeb 2.0 project supported by the European Union's 7th Research Framework Programme (FP7).

References

[1] Mobile Web Best Practices 1.0, Jo Rabin, Charles McCathieNevile, W3C Recommendation, 29 July 2008, http://www.w3.org/TR/2008/REC-mobile-bp-20080729/
[2] Vodafone Mobile Internet available today, Vodafone press release, 06 June 2007, http://www.vodafone.com/start/media_relations/news/local_press_releases/uk_press_releases/2007/vodafone_mobile_internet.html
[3] Sprint Customers to Get a PC-Like Experience of the Open Internet on Their Phones, Sprint press release, 18 March 2008, http://newsreleases.sprint.com/phoenix.zhtml?c=127149&p=irol-newsArticle_newsroom&ID=1119727
[4] Content Transformation Guidelines 1.0, Jo Rabin, W3C Working Draft, 01 August 2008, http://www.w3.org/TR/2008/WD-ct-guidelines-20080801/
[5] Greasemonkey, Firefox extension that allows to customize the way Web pages look, http://www.greasespot.net/
[6] The Washington Post, et als. v. TotalNews, Inc., et als., Southern District of New York, Civil Action Number 97-1190, 20 February 1997, http://legal.web.aol.com/decisions/dlip/washcomp.html
[7] The XMLHttpRequest Object, Anne Van Kesteren, W3C Working Draft, 15 April 2008, http://www.w3.org/TR/2008/WD-XMLHttpRequest-20080415/
[8] Hypertext Transfer Protocol -- HTTP/1.1 Request for Comments: 2616, R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, June 1999, http://tools.ietf.org/html/rfc2616
[9] XHTML Basic 1.1, Shane McCarron, Masayasu Ishikawa, Editors, W3C Recommendation, 29 July 2008, http://www.w3.org/TR/2008/REC-xhtml-basic-20080729/
[10] Protocol for Web Description Resources (POWDER): Description Resources, Phil Archer, Kevin Smith, Andrea Perego, W3C Working Draft, 14 November 2008, http://www.w3.org/TR/2008/WD-powder-dr-20081114/
[11] Transparent Content Negotiation in HTTP, K. Holtman, A. Mutz, March 1998, http://tools.ietf.org/html/rfc2295