The Consistent Web: A Vision for the Future

Ciaran Ferry
School of Computing
Dublin City University
Dublin 9, Ireland
Tel: +353 1 7005617
Cferry@computing.dcu.ie

Stephen Blott
School of Computing
Dublin City University
Dublin 9, Ireland
+353 1 7005984
Sblott@computing.dcu.ie

ABSTRACT
Over the past decade, the World-Wide-Web has become the preferred vehicle for information dissemination on a global scale. However, no guarantees as to the consistency, or correctness, of the content being provided are made. In this paper, a novel vision of a Consistent Web is proposed. Unlike in the current system, time will have a direct influence on whether requests for specific content should be satisfied. Firstly, the motivation for the work and the limitations of current consistency management techniques are discussed. The concept of a Consistent Web is then presented, and by way of several examples, an illustration of the type of functionality it offers is provided. Finally, an outline of the research strategy to be employed is given.

Categories and Subject Descriptors
C.2.4 [**Computer-Communication Networks**]: Distributed Systems | Client/server, Distributed applications.

General Terms
Algorithms, Performance, Design, Experimentation.

Keywords
Web/Internet, Content Distribution, Consistency Management.

1. INTRODUCTION
Over the past decade, the World-Wide-Web has become the preferred vehicle for information dissemination on a global scale. It is now possible to access anything from the latest news and sports headlines from around the globe, to the most up-to-the-minute stock prices from the world's markets, all at the touch of a button. Recent developments, both technological and commercial, have witnessed the introduction of a range of applications and services that extend the functionality of the Web. Electronic commerce, on-line banking and video on-demand are just a few examples. Highly accessed and frequently updated special event websites, such as those for the Olympic Games and the World Cup, are now extremely common, with performance and reliability issues presenting numerous challenges to the research community. As a consequence of the sheer volume of data that exists, together with its associated access and update patterns, innovative approaches to consistency management are required.

CoNEXT'05, October 24–27, 2005, Toulouse, France.
ACM 1-59593-035-3/05/0010.

Although primarily concerned with improving the overall performance of the web, caching intermediaries and Content Delivery Networks[1] have served to further exacerbate the problem, by storing, replicating and processing content in a highly distributed manner.

2. RELATED WORK
Current consistency management techniques are extremely limited, due largely to the fact that their primary focus has traditionally been on the relationship between origin servers and web intermediaries, with the ultimate goal being to provide a variable level of consistency between both entities. In general, levels range from weak (the data at an intermediary reflects some correct state at an origin server at some time in the past), to strong (the data at both locations is identical). A host of different algorithms and protocols have been proposed, including client polling[2], the use of TTL (time-to-live) values[3], invalidation/update protocols[4], object/volume leasing[5], and Basis Token Consistency[6]. Although certain approaches provide solutions to specific problems, the amount of overhead incurred often makes them unsuitable for deployment in a real-world scenario such as the web. While hosting the websites for the 1998 and 2000 Olympic games, IBM demonstrated Data Update Propagation[7], a method of consistently updating the information presented on the sites. Object Dependence Graphs were used to maintain a record of dependencies among objects, thus when new information entered the system, graph traversal algorithms enabled the efficient propagation of the updates to the relevant objects. However, as is the case with alternative mechanisms, the scalability of the algorithm severely restricts its usefulness.

3. THE CONSISTENT WEB
The vision of a consistent web introduced here is based on the notion of time. It is proposed that in order for a set of web objects to be displayed to a user, the entire group of objects must have co-existed at some arbitrary time in the past. Additionally, once a user accesses a particular version of an individual object, an older copy of the same object must never be seen by that user. This is referred to as consistent browsing, and will become critically important as the amount of distributed data processing increases. In order to demonstrate the ideas more clearly, two real-world examples will be used; the former will be shown in the context of the Olympic Games, and the latter using a regularly updated website detailing the progress of a soccer match.
In figures 1 and 2 below, two separate web objects in the form of two HTML pages are displayed simultaneously. The object on the left contains the medals table for the first day of competition in

track and field, whilst the object on the right provides a listing of the results from recently completed events. Bearing in mind that current consistency management mechanisms consider objects independently, the possibility exists that events listed as having been completed, may not necessarily be reflected in the medals table. The reverse is also true; the medals table may indicate that the overall number of medals won is greater than that suggested by the events list.

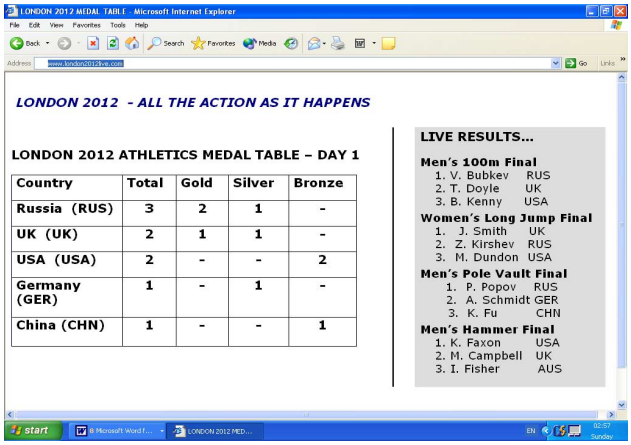

Figure 1 Illustrating inconsistency where total number of medals awarded is greater than that recorded in medals table

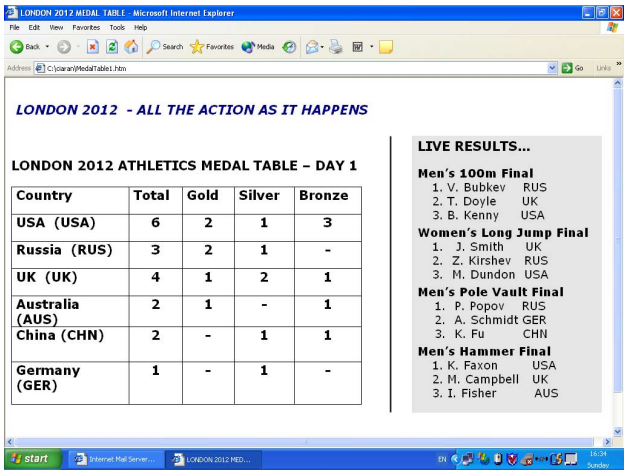

Figure 2 Illustrating inconsistency where total number of medals recorded in medals table exceeds total number awarded

The consistent web proposes a mechanism whereby inter-object consistency may be guaranteed. Objects must have existed together at a fixed point in time in order for their consistency to be guaranteed. In the case of the above example, consider multiple versions of both objects existing over a 10-hour period, with updates to one object taking place on the hour, and updates to the other taking place on the half hour. Capturing a snapshot of both objects at a specific point in time would ensure that the information presented is correct, thus avoiding data discrepancies such as those already presented.

Now, consider a single page, consisting mainly of text-based commentary of a soccer match. Every two minutes, the page is updated, with a further description of the action being added. Suppose that after twenty minutes, the score was observed to be two goals each, and the user decided to check the latest news headlines. Then, after three minutes, returned, only to notice that the score was now one goal each, and that only fourteen minutes had elapsed in the match. In The Consistent Web, thanks to the introduction of the requirement that access to older versions of objects than those already accessed is not permitted, the inconsistency demonstrated above could be avoided.

4. RESEARCH STRATEGY AND CONCLUSION

In this paper, a novel vision of a consistent web has been proposed. It introduces a number of consistency-based guarantees, which aim to make the WWW a more reliable and dependable environment for the deployment of next-generation web applications. Going forward, the research strategy to be adopted will consider the intrinsic links that exist between Internet and web related technologies, and attempt to provide a solution that may be useful across a range of application areas. Initially, emphasis will be placed on the construction of a consistency model, through which data correctness may be formally specified. Focus will then turn to the design and development of an algorithm that will facilitate the implementation of the consistency requirements set out above. A study of the various optimisation issues involved in implementing the algorithm will follow, and a working prototype will be produced. Finally, a detailed evaluation of the approach will be undertaken, and an analysis of its effectiveness performed.

REFERENCES

[1] Akamai: http://www.akamai.com

[2] Arun Iyengar, Erich Nahum, Anees Shaikh, and Renu Tewari. Web Caching, Consistency and Content Distribution, in *The Practical Handbook of Internet Computing*, Copyright 2005, Chapman & Hall/CRC Press.

[3] James Gwertzman and Margo Seltzer. World-wide web cache consistency. In *Proceedings of the USENIX Technical Conference*, pages 141-152. USENIX Association, January 1996.

[4] Dan Li and Pei Cao. WCIP: Web cache invalidation protocol. In *5th International Web Caching and Content Delivery Workshop*, Lisbon, Portugal, May 2000.

[5] A. Ninan, P. Kulkarni, P. Shenoy, K. Ramamritham, R. Tewari. Cooperative Leases: Scalable Consistency Maintenance in Content Distribution Networks. In *Proceedings of the WWW2002 conference*, (Honolulu, Hawaii), May 2002.

[6] Adam D. Bradley and Azer Bestavros, Basis Token Consistency: Supporting Strong Web Cache Consistency, *Global Internet 2002 (GI2002)*, Taipei, Taiwan, 2002.

[7] Jim Challenger, Paul Dantzig, and Arun Iyengar. A Scalable and Highly Available System for Serving Dynamic Data at Frequently Accessed Web Sites. In Proceedings of *Supercomputing'98,* Orlando, Florida.