

Garbage Collecting the World Wide Web

Stephen M. Watt
Western University
London, Ontario, Canada N6A 5B7
Stephen.Watt@uwo.ca

Abstract

The World Wide Web has grown over the past decade and a half from being a collection of modest resources to an essential part of our every day lives. While the curation of early web sites could be managed by hand, present web portals are backed by collections of numerous hosts with complex inter-connections. The issues arising in the management of present-day web resources bears many similarities to the issues encountered with explicit storage management in the first few decades with programming languages. For example, the owner of a web site may remove a page when it is still referenced by others, or a page may be kept long after it is no longer accessible. Lisp, a programming language for symbolic computation, introduced the notion of garbage collection for automatic storage management. We show how the ideas of modern distributed garbage collection can be applied to the world wide web and present the results of some experiments. This presentation is based on joint work with my doctoral student Yannis Chicha [1].

References

- [1] Yannis Chicha. Practical Aspects of Interacting Garbage Collectors. PhD Thesis, University of Western Ontario, 2002.