

Searching and Archiving the Web with Tumba!

Mário J. Silva

Universidade de Lisboa, Faculdade de Ciências, Lisboa, Portugal

mjs@di.fc.ul.pt

Searching and Archiving the Web with Tumba!

Abstract

In the past, Internet archives and Web search engines have always been conceived and implemented as independent information systems. This paper shows how Web search and archival functions are supported with tumba!, an information system that can serve national interests in preserving its Web data as cultural heritage, obtaining knowledge about the preferences and interests of its society in the information age and also in intelligence gathering. The tumba! search engine has a new repository architecture and uses innovative ranking and presentation algorithms optimised for this Web.

Keywords: Web search, Web archival, Web data repositories

1. Introduction

The Web has a global nature. However, the behaviour of Web users reflects their social background and their linking patterns reflect how they group in communities [Gibson et al., 1998]. Web communities can be geographically spread across the globe, its members may have many interests and belong to several unrelated communities, but there is a distinct language, vocabulary or other social behaviour characteristic that binds them and motivates the creation of Web links among their Web sites.

Tumba! is a Web search engine specially crafted to provide better results to those searching information on the Portuguese Web. It incorporates knowledge about the usage profile and interests of those who access these pages to improve Web searches. Tumba! is a public service since November 2002 (see <http://www.tumba.pt>). Throughout this paper, I will discuss the Portuguese Web having in mind how it is seen by tumba!. For this search engine, the Portuguese Web is currently defined as the collection of pages from the global Web that satisfy one the following conditions:

- Hosted on a site under a “.PT” domain.
- Having at least one incoming link originating in a Web page hosted under a “.PT” domain, written in Portuguese and hosted in one domain that contains a large set of Portuguese pages (“.ORG”, “.COM”, “.NET”, “.TV”).

I believe that the implementation of a search engine specialised in providing services to community Webs, such as tumba!, is a valuable alternative to global search engines for locating information within the Web pages of these communities. The lack of context or scope information in queries submitted to global engines causes that users in different communities will perform the same query and expect completely different results when a search is made, in particular when these tend to consist of a small number of words which mean different things in different places/languages (see, for instance, the results given by global search engines queries on “passaporte” or “biblioteca”). Tumba! has a similar architecture and adopts many of the algorithms of global search engines, but its configuration data is much richer in its domain of specialisation. It has a better knowledge of the location and organization of Portuguese Web sites (both in qualitative and quantitative terms).

As no other Portuguese organization is systematically crawling and archiving the contents of this Web, the importance of the development of tumba! has more than a simple cultural or

commercial interest: it may become strategic both for government entities and national industries as the resource for locating information and services for Web users communicating primarily in the Portuguese language or interested in locating resources related to Portugal.

Tumba! is now the result of about two years of development with the XLDB group at Universidade de Lisboa. This paper discusses the motivation for building and offering tumba! as a public service, the requirements that have driven its architectural design, which was oriented towards supporting good quality searches and double as an archive for the Portuguese Web. The paper is organized as follows: first, I will quickly review Web search. The next section details the architecture of tumba! and it is followed by a summary of what has been achieved by tumba! Then, I present the arguments for creating tumba! as a search and archive machine. Finally, I derive the conclusions and present some of the directions for future improvements to tumba!

2. Searching and Archiving the Web

Web search engines are almost as old as the Internet (see [Arasu et al., 2001] for a detailed review). The first search engines applied pre-Web information retrieval concepts. The Web was seen as a collection of documents, each abstracted as a simple bag of words extracted from the HTML source. Web search technology eventually evolved to use a new search paradigm. State-of-the-art search engines now rank Web pages based on concepts taken from bibliometrics that provide much better results, such as the measure of documents authority. This is computed as the number of pages that link to a given page from the set of pages that match a query [Kleinberg, 1998]. The most widely used algorithm of this kind is PageRank [Page, 1998]. Unlike the authority measure computed with Kleinberg's algorithm, PageRank computes a measure of the popularity of every page based on the number of pages referencing it. As this measure is computed offline, it can be used in fast ranking computations, making it suitable for global search engines. PageRank builds a graph of the Web, where pages are nodes and edges the links that connect those pages, and then computes the importance of every page as the number of links that recursively point to a page weighted by the importance of each referencing page. PageRank is used in Google¹, the most widely used Web engine today.

Another dimension of Web search engines lies in the capacity to preserve the successively crawled pages for each Web site in an historic database. The largest effort to archive the Web is an initiative of the Internet Archive [Kahle, 1997]. There is currently no service providing both an interactive search interface and an archival service at the global level. Internet Archive stores the sites it crawls over time but does not offer a search facility or a way for navigating in the pages collected during a certain period of time. There are also several initiatives of national libraries to archive portions of their Webs (see [Day, 2003], for a recent report covering this topic).

The market for search engines is by no means settled. Since the beginning of the explosion of the Internet we have been witnessing the arrival of new players that introduce technical innovations that set new standards for information search. The advent of the Semantic Web [Berners-Lee et al., 1992] and an increasing trend for publishing useful, credible and accurate meta-data by public sector organizations could be the trigger for a new generation of search technologies more useful and reliable than what we have today.

Web search tools aren't limited to engines that crawl the Web, such as the ones presented above. In addition to these, there is a multitude of engines that provide directory services or tend to combine directory listings with Web search results. There is some research in topic distillation, ie, the retrieval of collections of relevant documents for a given information need (not necessarily containing search terms) [Bharat and Henzinger, 1998]. On the other hand,

¹ see <http://google.com>

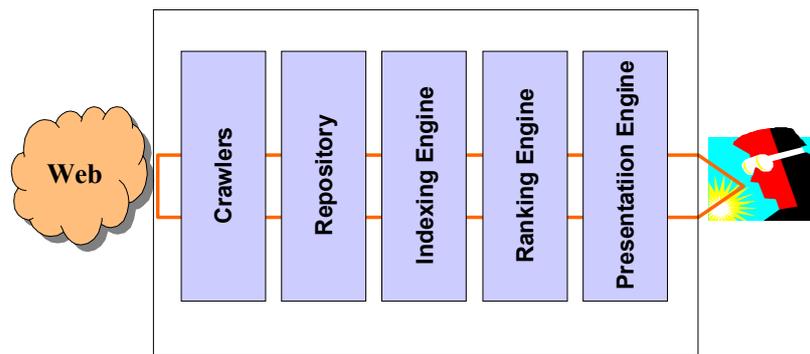


Figure 1 – The main components of tumba! create a pipeline that transforms information collected from the Web in successive stages.

Vivissimo² is one example of a meta-search engine that is able to cluster search results interactively, based on the information provided in the snippets of the results of other search engines.

3. Tumba

Tumba! implements some the best know algorithms for Web search and is pushing the envelope in some domains. Tumba!’s repository is built on a novel framework for parallel and incremental harvesting of Web data, based on version management concepts developed for engineering databases [Katz,1990]. In addition, it can provide better rankings than global search engines, as it makes use of context information based on our local knowledge of the Portuguese Web and our handling of the Portuguese language.

A substantial part of the tumba! software was initially developed for other projects. A selective harvesting system for the digital deposit of Web publications for the National Library of Portugal [Noronha et al., 2001] has provided the environment for building the first prototype of tumba!. In parallel, the development of XMLBase³, a Web data warehousing system began to take shape. Past experience has shown that many of the blocks that compose the current system could be reused and combined to form the basis of a wide range of Internet of applications. Tumba!’s technology could be useful for any software application demanding: i) selective harvesting of Web sites; ii) complex, high-performance, Web data warehouse systems; iii) error-tolerant parsing of badly formed HTML; and iv) new, XML-aware data storage, indexing and search systems.

The concept of a Web crawler as a reusable software component was first proposed in Mercator, the crawler used by the Altavista search engine [Najork and Heydon, 2001]. The Google “anatomy” paper also provided a good source of inspiration for defining the functionality of some of the main software blocks of tumba! However, the APIs of some of tumba!’s components are radically different, in particular the organization of the repository and crawling software.

The remainder of this section details the software architecture of tumba! and its main components.

Architecture

² <http://vivisimo.com>

³ see <http://xldb.fc.ul.pt/xmlbase>

The architecture of tumba! follows the tiered model of high performance information systems. One main difference between tumba! and other search engines is on the emphasis on reusable software components. Some of the main data management software components incorporated in tumba! are available as commercial products, while other are available as open source. Some of the newly developed components of tumba! were conceived to be used independently by other Web data processing applications.

Tumba!'s data pipeline is illustrated in Figure 1. Information flows from Web publishing sources to end users through successive stages. At each stage a different transformation is performed on the data:

- In the first stage, *crawlers* harvest Web data referenced from an initial set of domain names and/or Web site addresses, extract references to new URLs contained in that data, and hand the contents of each located URL to the Web repository.
- The Web *repository* is a specialised database management system that provides mechanisms for maximizing concurrency in parallel data processing applications, such as Web crawling and indexing. It can partition the tasks required to crawl a Web space into quasi-independent working units and then resolve conflicts if two URLs are retrieved concurrently. This mechanism is implemented upon a versioning model, derived from the conceptual models of engineering databases.
- The *indexing engine* builds a fast search data structure and provides services that return a list of references to pages in the repository matching user supplied keywords.
- The *ranking engine* sorts the references to pages produced by the indexing engine by their perceived relevance.
- The *presentation engine* receives search results in a device-independent format and formats them to suit multiple output alternatives, including Web browsers, mobile phones, PDAs and Web Services. The presentation engine can also cluster search results in multiple ways.

I will now describe each of these main components in more detail.

Web Crawler and Repository

The Web crawling and data repository sub-system of tumba! has three main components: Versus, a Web meta-data repository; VCR, the Versus Content Repository; and Viúva Negra, a Web crawler built upon Versus.

Versus provides a framework for storing large quantities of information retrieved from the Web throughout time and give access to it to a range of data and knowledge management tools. The objectives of this framework support many of the needs of tumba! [Campos et al., 2002] and [Gomes et al., 2002]:

- harvesting large quantities of data from the Web;
- managing meta-data about large collections of Web resources;
- organizing Web data in a temporal dimension.

Versus performs many Web data management tasks and provides support for parallel execution of many operations, through a versions and workspaces model, data partitioning functions and time support. It is designed to maintain information about Web objects. In Versus, a Web object is an entity that has an associated URL and may contain references to other Web objects. The main classes of the Versus data model are:

- *Workspaces* represent collections of objects. These are shown with 3 different levels of visibility to applications. Private workspaces are exclusive of individual application threads; Group workspaces are shared by application threads and reconcile data received from Private workspaces; finally, Archive workspaces are append-only and can be read by any Versus application. Data can be copied/moved between workspaces through check-out and check-in operations.

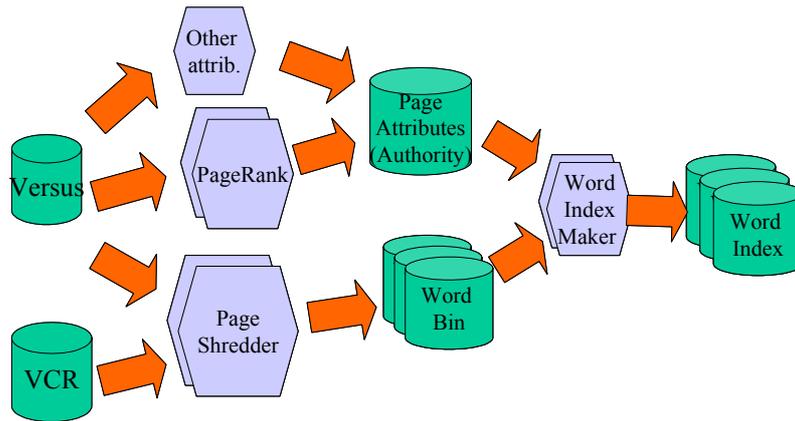


Figure 1 – The off-line process of SIDRA involves parallel shredding of the pages to index, and parallel generation of word indexes. Document references in word indexes are sorted in popularity order.

- **Versions** capture each one of the representations (or “images” – the contents of URLs) that a Web object may have over time. Versions may have multiple **Facets**, which store alternative views of objects: a given Web object may have a facet containing its original publishing format (in HTML or PDF), a text representation consisting of the extracted words, or a well-formed XHTML document generated from the captured data.
- **Layers** represent collections of Versions, with the restriction that they may contain only one version of each object. Versions may have an associated label or time stamp, representing collections such as “the Web pages crawled in March 2003.”
- **Partitions** represent schemes for dividing the objects represented in disjoint groups called Working Units. An example partition would be “the set of hosts on the Web”, where each Working Unit would then represent the set of all Objects with URLs that share the same host component.

Applications operate under the following model: first, they define a Partition on a Group Workspace and set a current Layer. Then, they start a group of application threads, where each thread processes a Working Unit at a time in its own Private Workspace and checks the result back in the Group Workspace, associating it with the current Layer. When all Working Units have been processed, the Application may end or instruct Versus to check-in the data in the current layer of the Group Workspace into an Archive Workspace.

When an application needs to store the contents of an URL, it uses Versus services to process the download. Versus is designed to operate above a data store organized as a distributed file server that stores the contents of retrieved objects and any additional objects that may be derived from these during the post-processing tasks that follow. In tumba!, the VCR – Versus Contents Repository, provides that service. The VCR is a distributed storage server for the Web objects loaded and managed by Versus, which hides the physical details of data location and replica management of the archived contents.

Versus detects when contents are replicated in the repository (by computing and maintaining MD5 hashes for each stored contents). This makes it suitable to process successive crawls of Web spaces without the need to maintain complete replicas of all the crawls. However, the repository can always provide Web images containing full copies of any of the past crawls (each crawl is associated with a distinct layer in Versus containing all the references to the pages detected in that crawl).

Indexing and Ranking

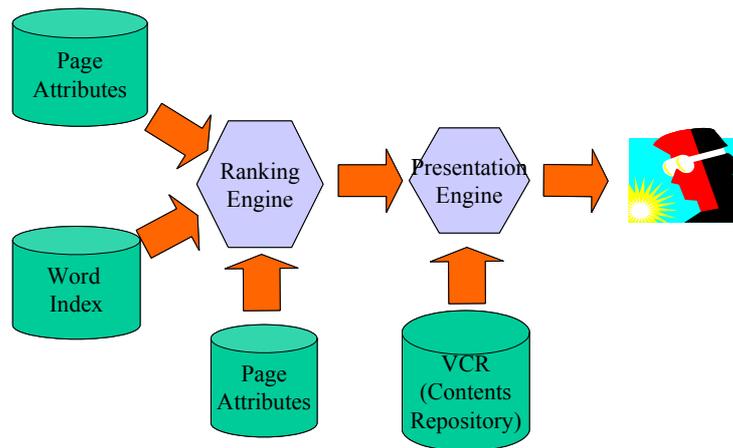


Figure 3 – Architecture of the on-line subsystem of tumba!

Indexing and Ranking in tumba! are performed by SIDRA, a new indexing and ranking system [Costa, 2003]. The design of SIDRA reflects the requirements derived from the experience of building and operating tumba!. Like in any other index, SIDRA runs both off-line and on-line. Most of the expensive computing tasks are performed off-line. These include (see Figure 2):

- retrieving the list of documents and some of their meta data from Versus, and the documents' contents from VCR;
- *shredding* documents into word bins. Shredding makes use of the data partitioning capabilities available in Versus and processes each unit by an independent thread. In the end, each bin is assigned to a specific word-range and will contain the set of document references to associate to each word;
- computing the static ranking (or importance) of each page. Versus provides the linkage data to create the graph of the Web on which the static ranking of each page is computed. In the end, the results are stored in a database that contains the information attributes associated to each page that will be necessary to compute the final rank;
- creating word indexes. This program takes a word bin as input and creates a corresponding word index as a compressed term-document structure. In a word index, each document receives an identifier that reflects its static rank. Each word entry in the word index has the documents sorted by the value of these identifiers.

The online part of SIDRA takes the word indexes as input and computes the final ranking of results to present in response to a user's query. These results are then pipelined to the Presentation Engine (See Figure 3).

The Ranking Engine starts by obtaining the pages with the highest static ranking page matching a query formulated as a Boolean expression of terms. This operation is designed to run with high parallelism. Word indexes are accessed through QueryServers. Each QueryServer is configured to provide the page identifiers contained in one or more word indexes. There is no boundary on the number of Query Servers on the system: they can be freely added to match the load and fault-tolerance levels intended for the system.

In SIDRA, the processing of queries requires, in general, an additional processing step for merging the output of different servers. This is the task of QueryBrokers, which select the QueryServers that will be requested to provide results (based on the input query terms and load-balancing criteria) and combine the results by applying the logic restrictions specified in the input. As the page references in word indexes are pre-sorted, this can be performed very efficiently with a sort-merge-join algorithm.

At this stage, the tumba! ranking reflects the static page ranking (computed off-line). The experience with the evaluation of initial prototypes of tumba! has shown that the set of the top 1000 results of the static ranking contains the set of the top 10 results of the optimal ranking (to be presented at the output). The final stage consists in obtaining these results and changing their sort order to reflect dynamic scores that weight factors such as the TF.IDF score of each page [Baeza-Yates and Ribeiro-Neto, 1999], presence of query terms in the title or description meta-tags of Web pages.

Although the indexing algorithms that we use in tumba! are similar to those described in the literature, our indexing methods differ from those used by global search engines. In tumba!, unlike in other search engines, we consider Portuguese special characters with accents and cedillas. Tumba!'s ranking heuristics are also likely different, but it is not possible to compare them given that commercial search engines keep them secret.

Presentation

The final stage of tumba! involves the generation of the page snippets that show the context of the use of the matching input query terms on the Web pages. This is performed by obtaining from the Versus Content Repository the text previously extracted from each page/document and “greeting” the relevant parts.

By default, tumba! presents results following the style that has been adopted since the first Web search engines were created: output is in ranking order, through a filter that excludes from the listing those pages whose host is the same as in the previous result. A test version of tumba! is already demonstrating an alternative presentation interface using a page clustering algorithm tuned for the Portuguese language [Martins, 2003]. Tumba! provides several user interfaces, supporting different output formats, including HTML and WML (for mobile phones). There is also a simplified HTML version for PDAs available⁴.

4. Implementation, Operation and Usage Details

The tumba! search engine has evolved into a complex software system. It now provides most of the features that are available on global search engines, including related pages search, clustering of search results, a query spell-checker, links to Portuguese language dictionaries and a sub-system that enables users to listen to the pronunciation of query terms (developed by the Spoken Language Systems Lab of INESC⁵).

The software of tumba! can crawl and index not only HTML files, but also the most popular data types containing text information on the Web: Adobe PDF and PostScript, Microsoft Office and Macromedia Flash.

The existing infrastructure is presently composed of 3 medium-scale servers plus 6 older servers that are used for crawling. All the machines run the Linux operating system. A substantial part of tumba! is built on open source software:

- Versus is a software library that implements a Java API that is invoked by its applications. The internal data model for maintaining the metadata is relational, supported in Oracle9 and Hsql database servers.
- The Indexing software is written in C++ and Java; it uses BerkeleyDB databases.
- The Ranking and Presentation engines are written in Java and use the XML manipulation features of the Apache Cocoon publishing framework.

⁴ See <http://movel.tumba.pt>

⁵ See <http://www.l2f.inesc.pt>

It is hard to compare tumba! against other search engines, because there is no information about their details of operation or the data that they have indexed. An analysis of the coverage (i.e., the number of documents indexed by Google, AlltheWeb and tumba!) is an apples-to-oranges comparison: the types of documents indexed differ and there is no data describing how they are distributed, crawls have been obtained at different times, and especially the number of dead links referenced by these engines differs. However, it is possible to observe that tumba! has a deeper coverage of the secondary domains under “.PT”. By querying other search engines and observing the number of documents retrieved, we see that global search engines index about 1 million pages under “.PT”, while tumba! is now indexing above 3 million.

The data archived in the tumba! repository is available for use by other researchers. One of the Versus applications is a dump utility, that can output a stream with the data stored in the repository in XML format. It provides many filtering options for selecting text, links, HTML meta-data attributes, etc. By selecting a specific layer, we can produce a dump file with the contents of all the files harvested during any crawl.

It has been used to provide input to the software used in another project to build a linguistic corpus of the Portuguese Web [Fuchs, 2003]. The Portuguese computational linguistics community is organizing an initiative in evaluation of information retrieval systems processing information in Portuguese [Aires, 2003]. Tumba! will be evaluated and is offering part of its data, in the form of queries and Web document collections for other evaluation tasks [Silva, 2003].

The tumba! access statistics are also available (on demand) to researchers interested in studying the interest profiles of Portuguese users. This data shows that, while sex and related keywords dominate, as in global search engines, the most frequently used keywords reflect the pulse of the Portuguese society: they relate to events in business, environmental issues and national celebrities. On the other hand, the access log data also shows that the user population may still reflect the origins of tumba!: query analysis shows that university students and faculty, the first users, are among the most representative and are using this search tool in their daily activities.

Accesses to tumba! come from all over the world: in the first two weeks of March 2003 we accounted search requests from 74 top-level domains. However, the largest number of accesses are originated from Portugal: 78% of the accesses from resolved IP addresses come from Portuguese service providers (this percentage computed from 65% of the total IP addresses accounted for).

5. The Case for a Search and Archival Machine

As discussed above, state of the art global search engines compute relevance using both word significance and a measure of the documents' importance inferred from the links' structure. However, these make little use of additional knowledge that can be extracted or inferred from the Web. The developers of tumba! can identify the most important Web sites on the Portuguese Web and obtain additional information for each of these that can be useful to improve the quality of search results. We could, for instance, easily characterize the most popular sites of the Portuguese Web in terms of contents, the geographic scope of pages related to specific locations and embed this knowledge in our ranking algorithms.

We could also provide access to the deep and hidden Webs of the most important Web sites of our community, providing references to enter into content-rich databases in results pages. Recently, in a joint project with Biblioteca Nacional de Portugal, we added to the test version of tumba! a demonstration that presents, when the name of a book author is detected on a query

string, a result box with a link to access directly the records about that author in Porbase (the national on-line catalog of the Portuguese public libraries⁶).

Global search engines are commercially motivated and rank higher the sites that pay more to be listed on search results pages. Some claim that ranking is not affected by these payments. However, for smaller communities results will always differ from expected. As the profile of Web starting sites used by the Portuguese is completely different from that of Americans, the popularity of sites that does not take into account this difference will no reflect the ranking that a Portuguese user would expect. That is probably why when we search for “passaporte” or “biblioteca” (“passport” and “library” in English) we see so unexpected results when going to global Web search engines.

In addition to providing information to users about documents that match their queries, search engines have several other functions. Some of these will be important for the Portuguese national interests:

- Advertising of products and services of interest to the Portuguese community
- Conducting sociologic studies based on studying what people look for on the Internet and what sites they are visiting.
- Preserving the language in the digital realm, by providing a tool for searching data on the Web written in Portuguese and data to be used by computational linguistics tools.
- Web Archiving. Who is preserving contents of historical relevance being put on-line for use by history researchers 50 years from today?

Tumba! has a cache feature that gives access to last crawled version of each page in the public site (as available on the Google search engine). The internal search interface of tumba! enables internal users to visualize previously crawled versions of each page (our University research group does not have a mandate to publish this information).

6. Conclusion

This paper presented the case for developing a combined search plus Web archive machine for the Portuguese Web and its strategic importance. Arguments in favour include, among others, the archival of the information published for preservation, ability for providing independent search services for national public authorities, and also providing data for sociologic, linguistic and cultural research studies of the Portuguese community. It also detailed its global architecture and main design decisions, showing that an implementation can be built and operated with moderate resources. I have shown that it is possible to support the needs of various classes of users with diverse information needs with a single infrastructure.

Tumba! is quickly evolving to become an information service that provides a search engine that combines database and information retrieval query processing with context, authority and other knowledge of this Web to provide optimised query results for the Portuguese users.

There is an opportunity for making more specialized information location mediators for the users of thematic, language-specific or cultural Webs. These users have special information needs that cannot be met by an engine that, given its global nature, makes no or little use of contextual information while organizing search results. Tumba! was built from the start having in mind that it would be extensible and should provide support for indexing and searching XML data making use of meta-data in RDF (Resource Description Framework) format. The Semantic Web is the new vision for the next generation of the Web. As this vision unfolds into an extended large-scale database of semantically rich data, search engines will begin to have the foundations for providing more relevant results.

⁶ <http://porbase.bn.pt>

Tumba! does not offer an index to search historical data. The existing index is recreated periodically from the start with the last version of each crawled URL. We have no knowledge of an efficient index structure that could be incrementally updated and support query restrictions based on the crawling dates of each page. This would be an interesting addition for users willing to navigate on the historical data available on the tumba! repository.

7. Acknowledgements

Tumba! is not an individual project, but rather the result of extensive interactions with a group of collaborators. Several members of the XLDB research group have participated in the discussions whose outcome was the design presented in this paper, and developed the system currently in operation. Among these, João Campos, Miguel Costa, Daniel Gomes, Bruno Martins and Norman Noronha developed most of the system currently in operation.

The operation of tumba! and its hardware are supported by the FCCN –Fundação para a Computação Científica Nacional⁷. The XMLBase project is supported by Fundação para a Ciência e Tecnologia, under grant POSI / SRI / 40193 / 2001.

8. References

- Aires, R. and Aloísio, S. and Quaresma, P. and Santos, D. and Silva, Mário J. (2003). An initial proposal for cooperative evaluation on information retrieval in Portuguese. PROPOR'2003 - 6th Workshop on Computational Processing of the Portuguese Language - Written and Spoken, Faro, Portugal.
- Arasu, Arvind and Cho, Junghoo and Garcia-Molina, Hector and Paepcke, Andreas and Raghavan, Sriram, (2001). Searching the Web, ACM Transactions on Internet Technology.
- Baeza-Yates, Ricardo and Ribeiro-Neto, Berthier, (1999). Modern Information Retrieval, Addison-Wesley.
- Bharat, K. and Henzinger, M.R. (1998). Improved Algorithms for Topic Distillation in a Hyperlinked Environment. In Proceedings of SIGIR - Conference on Research and Development in Information Retrieval.
- Brin S. and Page L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. Proceedings of the Seventh World Wide Web Conference (WWW7), Brisbane, also in Computer Networks and ISDN Systems, Volume 30, issues 1-7.
- Campos, João, (2002). Versus: a Web Repository, Master Dissertation, Faculdade de Ciências da Universidade de Lisboa. Also available as DI/FCUL TR-03-08, <http://www.di.fc.ul.pt/tech-reports/03-8.pdf>.
- Costa, Miguel and Silva, Mário J. (2001). Ranking no Motor de Busca TUMBA, CRC'01 - 4ª Conferência de Redes de Computadores, Covilhã, Portugal (in Portuguese).
- Costa, Miguel (2003). SIDRA: Web Indexing and Ranking System, Master Dissertation, Faculdade de Ciências da Universidade de Lisboa, (in preparation).
- Day, Michael, (2003). Collecting and Preserving the World Wide Web. http://www.jisc.ac.uk/uploaded_documents/archiving_feasibility.pdf

⁷ see <http://www.fccn.pt>

- Fuchs, R., and Martins, B. and Silva, Mário J. (2003). Statistics of the Tumba! Corpus, Technical Report. Universidade de Lisboa, Faculdade de Ciências, Departamento de Informática (in preparation).
- Gibson, D., Kleinberg, J., and Raghavan, P. (1998). Inferring Web communities from link topology. In Proceedings, 9th ACM Conference on Hypertext and Hypermedia, New York, NY. ACM, ACM.
- Gillmor, David, (2002). 'Google effect' reduces need for many domains. SiliconValley.com, January, 12, 2002.
- Hearst, M. A. and Pedersen J., (1996). O. Reexamining the Cluster Hypothesis: Scatter/Gather on Retrieval Results. Proceedings of 19th ACM SIGIR Conference on Research and Development in Information Retrieval.
- Kahle, B. (1997). Preserving the Internet. *Scientific American*, 276 (3), 72-73.
- Henzinger, Monika R. and Motwani, Rajeev and Silverstein, Craig (2002). Challenges in Web Search Engines, SIGIR Forum, 36(2).
- Katz, Randy, (1990). Towards a Unified Framework for Version Modeling in Engineering Databases. *ACM Computing Surveys*, 22(4):375—408.
- Kleinberg, J. (1998), Authoritative sources in a hyperlinked environment, in 'Proceedings ACM-SIAM Symposium on Discrete Algorithms', San Francisco, California, pp. 668--677
- Martins, Bruno, (2003). Web Search using Inter-document Similarities, Master Dissertation, Faculdade de Ciências da Universidade de Lisboa, (in preparation).
- Noronha, Norman, and Campos, João P. And Gomes, Daniel and Silva, Mário J., and Borbinha, José Luís, (2001). A Deposit for Digital Collections. Proceedings of the European Conference on Digital Libraries, ECDL, pp. 200-212.
- Najork, M., and Heydon, A (2001). High-performance Web crawling. Tech. Rep. Research Report 173, Compaq SRC.
- Silva, Mário J and Martins, Bruno and Costa, Miguel (2003). Avaliação de Sistemas de Recuperação da Web em Português: Uma Proposta Inicial à Comunidade. Avalon'2003, Faro, Jun de 2003. Available as DI-FCUL Technical Report TR-03-11. <http://www.di.fc.ul.pt/tech-reports/03-11.pdf>.