

# SeFS: Unleashing the Power of Full-text Search on File Systems\*

USENIX FAST '07(WiP)

George Margaritis

Stergios V. Anastasiadis

Department of Computer Science

University of Ioannina, Ioannina 45110, GREECE

**Dropping prices in storage space capacity combined with a constant increase in the accumulated amount of digitized content create an unprecedented need for full-text search automation in local and distributed file systems. The current state of the art in what is usually called desktop search is mainly (or seems to be) an adaptation from technology previously developed for the needs of information retrieval and web search applications. Indeed, search facilities for large-scale storage systems have not received the attention that they deserve from system designers. We point out particular deficiencies of existing systems and outline approaches likely to overcome them.**

Existing content search technology was originally developed for the needs of the information retrieval applications. Essentially, the availability of news articles in digital text format provided the basis for designing systems that can store and index large amounts of text repositories. Typically, individual articles stored in such systems were not supposed to be changed over time. However, large numbers of new articles were constantly added to the repository on a daily basis.

As a result, both the storage management and the indexing methods used in such systems were mainly customized for the needs of append-only updates. Quite surprisingly, the indexing structure that maps each searched term to the list of documents that contain the term is assumed to be contiguously stored on the disk. To a large extent, this fundamental decision dominates the structure of the entire system by enforcing indexing information constantly relocated to guarantee contiguity as new articles are added to the system [BCL06].

In fact, alternative approaches that organize indexing information in blocks were not previously exhaustively evaluated because older studies restricted the block size to a few tens of kilobytes [TGS94]. Therefore, the design of modern file system search facilities should be reconsidered from scratch taking into account both the particular needs of desktop search and the features of the latest storage technology.

Another source of expertise for developing search facilities for file systems originates in the web search community [SG05]. However, the web is assumed to be an unreliable storage medium. As a result, the corresponding search facilities update the indexing structure only periodically at time intervals that may span several days. On the contrary, file systems are regarded as highly reliable storage media and users' expectations from them are totally different. For example, stored data should become visible immediately and searchable within a small interval of time that can hardly exceed a few tens of minutes.

We believe that existing expertise in full-text search is highly valuable. Today, it is not uncommon to search over the web for content that is somewhere available in our local hard disks but we don't know where. One reason for that is the fact that current desktop search tools sometimes prove heavy-weight in their update processing needs and may be turned-off by frustrated users. This is especially true in file systems that exceed a few gigabytes and are updated frequently. Very large files or enormous numbers of small files [AGC05] may involve additional restrictions on the operation of existing desktop search tools, such as indexing only part of a large file or delays in handling normal queries.

Overall, we consider file system search an important open issue that should be highly regarded by the systems research community. More importantly, we point out the need for studying it from scratch but taking into account all the useful ideas that have been accumulated during the last decades in similar but not exactly the same contexts.

[TGS94] A. Tomasic, H. Garcia-Molina, K. Shoens, *Incremental updates of inverted lists for text document retrieval*. ACM SIGMOD, 1994, pp. 289-300.

[BCL06] S. Butcher, C. L. A. Clarke, B. Lushman, *Hybrid index maintenance for growing text collections*, ACM SIGIR, 2006.

[SG05] C. A. N. Soules, G. R. Ganger, *Connections: Using context to enhance file search*. SOSP 2005, pp. 119-132.

[AGC05] S. V. Anastasiadis, S. Gadde, J. S. Chase, *Scale and performance in semantic storage management of data grids*, Intl. Journal on Digital Libraries, 5(2), April 2005, pg. 84-98.

\* Supported in part by an Interreg IIIA grant (I2101005).