



Source: NSF: www.nsf.gov/discoveries/disc_summ.jsp?cntn_id=100660
Abridged version of original article by David Hart

On the Origins of Google:

Even in the early days of the Internet, people saw the need for better interfaces to growing data collections. A graduate student supported by an NSF digital library project at Stanford University uncovered the missing links in Web page ranking.

In the primordial ooze of Internet content (1993), fewer than 100 Web sites inhabited the planet. Early clans of information seekers hunted for data among the far larger populations of text-only Gopher sites and FTP file-sharing servers. This was the world in the years before Google.

Even in this primitive Internet world, the need for more accessible interfaces to growing data collections had already been recognized. The National Science Foundation led the multi-agency Digital Library Initiative (DLI) that, in 1994, made its first six awards. One of those awards supported a Stanford University project led by professors Hector Garcia-Molina and Terry Winograd.

None of the early DLI proposals -- submitted before the World Wide Web experienced its Cambrian explosion -- explicitly included research into the Web. However, by the time DLI funding began, the information landscape had changed.

In 1994, some of the first Web search tools crawled out of the Internet sea. Two Stanford students started Yahoo!, a manually constructed "table of contents" for Web sites. Other early search engines emerged, such as Lycos and WebCrawler, and began automatically indexing Web pages, focusing on keyword-based techniques to rank search results.

Around the same time, one of the graduate students funded under the NSF-supported DLI project at Stanford took an interest in the Web as a "collection." The student was Larry Page. Page uncovered the missing links, so to speak, in Web page ranking. His evolutionary leap was to recognize that the act of linking one page to another required conscious effort, which in turn was evidence of human judgment about the link's destination. Individually, each link was a simple but effective tool. But collectively, millions of these links provided a key adaptation for the natural selection of search results.

Page was soon joined by Sergey Brin, another Stanford graduate student working on the DLI project. (Brin was supported by an NSF Graduate Student Fellowship.) Together, Page and Brin constructed an ambitious prototype in their Stanford student offices. The equipment for the prototype, called BackRub, was funded by the DLI project and other industrial contributions.

The prototype used well-established technology to crawl from page to page by following links. However, in addition to compiling a standard text index, the prototype also mapped out a vast family tree that reflected the Web links among pages. The pair then developed the PageRank method that, in short, ranks a particular Web page highly if many other highly ranked Web pages link to it. Page and Brin tested the fitness of the approach on live Web data -- initially a test set of 24 million pages.

By late 1997, the BackRub approach proved to be sound, expandable and popular. By the end of the Early DLI Age in 1998, Page and Brin obtained funding and moved their growing facility from the Stanford campus to a friend's garage and incorporated Google, Inc. The rest, as they say, is history.