

Query

No. 1

No. 2

Styles 18

Designer Michael Cina

About Query is not the typical typeface family. Its members inherit the quirky asymmetrical contrast and spatial rhythm of 1960s wood type and, in varying degrees, the outright geometry of Paul Renner's earlier ideas for Futura. The gradation from black to extra light is not only a change in color, but also character. From the gregarious heavyweights to their even-keeled lightweight siblings, Query carries a range of expressive utility. To take things further, stylistic set 2 squarely mutates certain glyphs, transforming the face into something slightly alien and vaguely futuristic.

Styles

Query No. 1 - Thin
Query No. 1 - ExtraLight
Query No. 1 - Light
Query No. 1 - Regular
Query No. 1 - Medium
Query No. 1 - SemiBold
Query No. 1 - Bold
Query No. 1 - ExtraBold
Query No. 1 - Black

Query No. 2 - Thin
Query No. 2 - ExtraLight
Query No. 2 - Light
Query No. 2 - Regular
Query No. 2 - Medium
Query No. 2 - SemiBold
Query No. 2 - Bold
Query No. 2 - ExtraBold
Query No. 2 - Black

Glyphs per font

557

Language Support

Afrikaans, Albanian, Basque, Bosnian, Breton, Catalan, Croatian, Czech, Danish, Dutch, English, Esperanto, Estonian, Faroese, Fijian, Finnish, Flemish, French, Frisian, German, Greenlandic, Hawaiian, Hungarian, Icelandic, Indonesian, Irish, Italian, Latin, Latvian, Lithuanian, Malay, Maltese, Maori, Moldavian, Norwegian, Pinyin, Polish, Portuguese, Provençal, Romanian, Romany, Sámi (Inari), Sámi (Luli), Sámi (Northern), Sámi (Southern), Samoan, Scottish Gaelic, Slovak, Slovenian, Sorbian, Spanish, Swahili, Swedish, Tagalog, Turkish, Welsh.

A B C D E F G

H I J K L M

N O P Q R S

T U V W

X Y Z

a b c d

e f g h i j k

l m n o p

q r s t u v

w x y z

ÁáĂăĀă ÂâÄäÀà

ĀāĄąÅåÃã ˆ Ææ

ĆćĈĉČčÇç Ďd'Đđ

ÉéĚěÊêËëÈèĒēĘę ĞğĜğĴģ

Ħħ ĪīĴĵ Ķķ Łł'Ł'Łı

ŊŋŃńŅņŇň

ÓóŎňÔôÖöÒòÕõŌōØøÕõ°Œœ

ŔŕŘř ŚśŞşŠšŞş ŢţŤťȚțȚț

ŮůŪūŰű ŰűŲųŪū ÚúŬŭÛù

ŴŵŶŵŴŵŴŵ Ŷŷÿÿÿ

0 0̣ 1 1̣ 1̣ 2 2̣

3 3̣ 3̣ 4 4̣

5 5̣ 6 6̣ 7 7̣

8 8̣ 9 9̣ #

10/10²³/23⁴⁵/45

6 7/8 9/6 7/8 9

0 1 2 3 4 5

6 7 8 9

\$10 ¢23 £45

♥ ♦ ¥67 ♠ ♣

₣88 🍏 €90



SM TEL © P R ® TM




🏠 ○ ⬆️ ● ▲

▴ ▾ ▶ ▼ ◀ ▶

↙ ↘ ↗ ↖ ↕ ↙ ↘

◀ ◁ ▽ ▾ ◻ ◊ ◈ ◼ ▲ ▷ ▶





¡ ¢ \ , . : ; ' ' ' * _ • / ? !

' < " « ({ [- - -] }) » " > '


& § &

†% ‡‰

http://www

we@publictype

¶



Stylistic Set 1

(& 13 il ®®)

103 is a strictly non-palindromic number & 103 is the 27th prime number. The earlier prime is 101 making them both dual primes. It is also a joyful number. [1]®

The 103 series is a DC electric multiple unit (EMU) commuter train type introduced in 1963 [2] by Japanese National Railways (JNR), and currently operated by West Japan Railway Company and Kyushu Railway Company. They were also operated by East Japan Railway Company® and Central Japan Railway Company. In 2004, some sets were shipped to Indonesia, where they operate on the 1H3H1 KRL Jabodetabek system in Jakarta. Those

Stylistic Set 2

(CJMPNSUWY cfhijlmnrstuw)

WALTZ, PYPH, FOR QUICK JIGS VEX BUD.
 SRHPX OF BLACK QUARTZ, JUDGE MY
 VOY. PACK MY BOX WITH FIVE DOZEN
 LIQUOR JUGS. GLIB JOCKS QUIZ PYPH
 TO VEX. JACKDAWS LOVE MY BIG SRHPX
 OF QUARTZ. THE FIVE BOXING WIZARDS
 JUMP QUICKLY. HOW VEXINGLY QUICK
 DAFT ZEBRAS JUMP! Quick zephyrs blow,
 vexing daft Jim. Two driven jocks help fax
 my big quiz. The jay, pig, fox, zebra and
 my wolves quack! Sympathizing would fix
 Quaker objectives. A wizard's job is to vex
 chumps quickly in fog. Watch "Jeopardy!",
 Alex Trebek's fun TV quiz game. By Love, my
 quick study of lexicography won a prize!

Stylistic Set 3

(a)

Accepting any audible arcane and amorous
answer allows arguments. Anyone around
actively abhors analyzing audacious
applications abounding as apeshit actions,
and attracting amendable assholes and
affirming an appropriate avenue always
allows an absolute age acquittal and
adaptation as an articulate acceptance
amplified aboriginal anguish after an
abomination among associations and
allies adhered any and all accusations
accumulating among asteroid appraisers
annual activities, afterward, all astral
ambulation also allowed astronomers an
antiquated artform arranged as an ageless

TYPES!

Navigational URL

**Seek a single website or web page of a single entity
(e.g., PUBLIC TYPE or CINA ASSOCIATES)**

PAGE ENTITY

Engines do not disclose their search logs

QUERY NO. 1 BLACK

Indexing means associating words and other definable tokens found on web pages to their domain names and HTML-based fields. The associations are made in a public database, made available for web search queries. From a user can be a single word, multiple words or a sentence. The index helps find information relating to the query quickly. Some of the techniques for indexing, and caching are trade secrets, whereas web crawling is a straightforward process of visiting all sites on a systematic

The Regular NICNAME/WHOIS Server is an NCP/TCP transaction based query/response server, running on the SRI-NIC machine, that provides nationwide directory service to ARPANET users. It is one of a series of ARPANET/Internet name services maintained by the Network Information Center (NIC) at SRI International on behalf of the Defense Communications Agency (DCA). The server is accessible across the ARPANET from user programs running on local hosts, and it delivers the full name, U.S. mailing address, telephone number, and network mailbox for ARPANET users. This server, together with the corresponding Identification Data Base provides online directory look-up equivalent to the ARPANET Directory. DCA strongly encourages network hosts to provide their users with access to this network service. DCA requests that each individual with a directory on an ARPANET host, who is capable of passing traffic across the ARPANET, be registered in the NIC Identification Data Base. To register, send full name, middle initial, U.S. mailing address (including mail stop and full explanation of abbreviations and acronyms), ZIP code, telephone (including Autovon and FTS, if

Web search engines get their information by web crawling from site to site. The “spider” checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in BLACK HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. “No web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient.

available), and one network mailbox, via electronic mail to NIC@SRI-NIC.A web search engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. One of the first “all text” crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also

SEARCH

"inverted index"

Governed by Strict Syntax rules

"DO, KNOW, GO."

Sidewalk was really aimed at what we now call local search...
Sidewalk is one we should not have gotten out of.

QUERY NO. 1 THIN

Search engines were also known as some of the brightest stars in the Internet investing frenzy that occurred in the late 1990s. Several companies entered the market spectacularly, receiving record gains during their initial thin offerings. Some have taken down their public search engine, and are marketing enterprise-only editions, such as Northern Light. Many search engine companies were caught up in the dot-com bubble, a speculation-driven market boom that peaked in 1999 and ended in 2001. The search engine became so

Between visits by the spider, the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed. The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes. Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing.

In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was indexing, predating the very similar algorithm patent filed by Google two years later in 1998. Larry Page referenced Li's work in some of his U.S. patents for PageRank. Li later used his Rankdex technology for the Baidu search engine, which was founded by Robin Li in China and launched in 2000. Search engines were also known as some of the brightest stars in the Internet investing frenzy that occurred in the late 1990s. Several companies entered the market spectacularly, receiving record gains during their initial public offerings. Some have taken down their public search engine, and are marketing enterprise-only editions, such as Northern Light. Many search engine companies

Between visits by the spider, the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed.[25] The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes.[25] Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing.

EXTRA LG TRANSACTIONAL

Reflect the intent of the user to perform a particular action, purchasing a type family.

Search engines

often support a fourth type, used far less frequently

Query No. 1 ExtraBold

Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek, Inktomi, Northern ExtraBold, and AltaVista. Information seekers could also browse the directory instead of doing a keyword-based search. In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was

It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search. The engine looks for the words or phrases exactly as entered. Some search engines provide an advanced feature called proximity search, which allows users to define the distance between EXTRA BOLD keywords. There is also concept-based searching where the research involves using statistical analysis on pages containing the words or phrases you search for. As well, natural language queries allow the user to type a question in the same form one would ask it to a human. When we say human, we mean a person like Waldo. A site like this would be ask dot com. It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search.

Command–query separation is particularly well suited to a design by contract (DbC) methodology, in which the design of a program is expressed as assertions embedded in the source code, describing the state of the program at certain critical times. In DbC, assertions are considered design annotations – not program logic – and as such, their execution should not affect the program state. CQS is beneficial to DbC because any value-returning method (any query) can be called by any assertion without fear of modifying program state. In theoretical terms, this establishes a measure of sanity, whereby one can reason about a program's state without simultaneously modifying that state. In practical terms, CQS allows all assertion checks to be bypassed in a working system to improve its performance without inadvertently modifying its behaviour. CQS may also prevent the occurrence of certain

The engine looks for the words or phrases exactly as entered. Some search engines provide an advanced feature called proximity search, which allows users to define the distance between keywords. There is also concept-based searching where the research involves using statistical analysis on pages containing the words or phrases you search for. As well, natural language queries allow the user to type a question in the same form one would ask it to a human. A site like this would be ask.com. Microsoft first launched MSN Search in the fall of 1998 using search results from Inktomi. In early 1999 the site began to display listings from Looksmart, blended with results from Inktomi. For a short time in 1999, MSN Search used results from AltaVista instead. In 2004, Microsoft began a transition to its own search technology, powered by its own web crawler (called MSNBOT). Microsoft's rebranded search engine, Bing, was launched on June 1, 2009. On July 29, 2009, Yahoo! and Microsoft finalized a deal in which Yahoo! Search would be powered by Microsoft Bing technology. As of 2019, active search engine crawlers include those of Google, Sogou, Baidu, Bing, Gigablast, Mojeek, DuckDuckGo and Yandex.

CONTENT INFORMATIONAL

Cover a broad topic (e.g., FONTS or LETTERS)
for which there may be thousands of relevant results

Characteristics

broad categories that cover most web

Query No. 1 ExtraLight

Beyond simple keyword lookups, search engines offer their own GUI-EXTRALIGHT or command-driven operators and search parameters to refine the search results. These provide the necessary controls for the user engaged in the feedback loop users create by filtering and weighting while refining the search results, given the initial pages of the first search results. For example, from 2007 the Google.com search engine has allowed one to filter by date by clicking "Show search tools" in the leftmost column of the initial

The Regular NICNAME/WHOIS Server is an NCP/TCP transaction based query/response server, running on the SRI-NIC machine, that provides worldwide directory service to ARPANET users. It is one of a series of ARPANET/Internet name services maintained by the Network Information Center (NIC) at SRI International on behalf of the Defense Communications Agency (DCA). The server is accessible across the ARPANET from user programs running on local hosts, and it delivers the full name, U.S. ailing address, telephone number, and network mailbox for ARPANET users. This server, together with the corresponding Identification Data Base provides online directory look-up equivalent to the ARPANET Directory. DCA strongly encourages network hosts to provide their users with access to this network service. DCA requests that each individual with a directory on an ARPANET host, who is capable of passing traffic across the ARPANET, be registered in the NIC Identification Data Base. To register, send full name, middle initial, U.S. mailing address (including mail stop and full explanation of abbreviations and acronyms), ZIP code, telephone (including Autovon and FTS, if available), and one network mailbox, via electronic

Web search engines get their information by web crawling from site to site. The "spider" checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in BLACK HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. "No web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient. Some websites are crawled exhaustively, while

mail to NIC@SRI-NIC.A web search engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. One of the first "all text" crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also the search engine that was widely known by the public. Also in 1994, Lycos (which started at

CSS, AJAX

TRANSACTIONAL

A little less than a third of users entered three or more shocking unique queries.

AND, OR, + NOT

The average length of a search query was 2.4 terms

Query No. 1 Bold

In February 2010, Mark McCahill was revealed by the philosopher Peter Ludlow (also known by the pseudonym Urizenus Sklar) to be the Internet persona Pixeleen Mistral, a noted "tabloid reporter" covering virtual worlds who was the editor of Ludlow's newspaper The Alphaville Herald. In a 2016 interview with Leo Laporte, McCahill said that his involvement with developing the Croquet Project had led him into contact with Second Life and that he had become interested in the sociology of virtual worlds. As Pixeleen

The Gopher protocol is a communications protocol designed for distributing, searching, and retrieving documents in Internet Protocol networks. The design of the Gopher protocol and user interface is menu-driven, and presented an alternative to the World Wide Web in its early stages, but ultimately fell into disfavor, yielding to the Hypertext Transfer Protocol (HTTP). The Gopher ecosystem is often regarded as the effective predecessor of the World Wide Web.

The protocol was invented by a team led by Mark P. McCahill at the University of Minnesota. It offers some features not natively supported by the Web and imposes a much stronger hierarchy on the documents it stores. Its text menu interface is well-suited to computing environments that rely heavily on remote text-oriented computer terminals, which were still common at the time of its creation in 1991, and the simplicity of its protocol facilitated a wide variety of client implementations. More recent Gopher revisions and graphical clients added support for multimedia. Gopher was preferred by many network administrators for using fewer network resources than Web

The rise of GOPHER (created in 1991 by Mark McCahill at the University of Minnesota) led to two new search programs, Veronica and Jughead. Like Archie, they searched the file names and titles stored in Gopher index systems. Veronica (Very Easy Rodent-Oriented Net-wide Index to Computerized Archives) provided a keyword search of most Gopher menu titles in the entire Gopher listings. Jughead (Jonzy's Universal Gopher Hierarchy Excavation And Display) was a tool for obtaining menu information from specific Gopher servers. While the name of the search engine "Archie Search Engine" was not a reference to the Archie comic book series, "Veronica" and "Jughead" are characters in the series, thus referencing their predecessor. In the late 1980s, McCahill led the team at the University of Minnesota that developed POPmail, one of the first popular Internet e-mail clients. At about the same time as POPmail was being

services. Gopher's hierarchical structure provided a platform for the first large-scale electronic library connections. The Gopher protocol is still in use by enthusiasts, and although it has been almost entirely supplanted by the Web, a small population of actively-maintained servers remains. Gopher system was released in mid-1991 by Mark P. McCahill, Farhad Anklesaria, Paul Lindner, Daniel Torrey, and Bob Alberti of the University of Minnesota in the United States. Its central goals were, as stated in RFC 1436: A file-like hierarchical arrangement that would be familiar to users. A simple syntax. A system that can be created quickly and inexpensively. Extending the file system metaphor, such as searches. Gopher combines document hierarchies with collections of services, including WAIS, the Archie and Veronica search engines, and gateways to other information systems such as File Transfer Protocol (FTP) and Usenet. The general interest in campus-wide information systems (CWISs) in higher education at the time, and the ease of setup of Gopher servers to create an instant CWIS with links to other sites' online directories and resources were the factors contributing to Gopher's rapid adoption.

HTML DOM

Connectivity

Report on the connectivity of the indexed web graph
(e.g., Which links point to this > URL?)

boolean operators

How many pages are indexed from this domain name?

QUERY NO. 1 LIGHT

In the summer of 1993, no search engine existed for the web, though numerous specialized catalogues were maintained by hand. Oscar Nierstrasz at the University of Geneva wrote a series of Perl scripts that periodically mirrored these pages and rewrote them into a standard format. This formed the basis for W3Catalog, the web's first primitive search engine, released on September 2, 1993. In June 1993, Matthew Gray, then at MIT, produced what was probably the first web robot, the Perl-based World

Prior to September 1993, the World Wide Web was entirely indexed by hand. There was a list of web servers edited by Tim Berners-Lee and hosted on the CERN webserver. One snapshot of the list in 1992 remains,[4] but as more and more web servers went online the central list could no longer keep up. On the NCSA site, new servers were announced under the title "What's New!"

The first tool used for searching content (as opposed to users) on the Internet was Archie. The name stands for "archive" without the "v". It was created by Alan Emtage, Bill Heelan and J. Peter Deutsch, computer science students at McGill University in Montreal, Quebec, Canada. The program downloaded the directory listings of all the files located on public anonymous FTP (File Transfer Protocol) sites, creating a searchable database of file names; however, Archie Search Engine did not index the contents of these sites since the amount of data was so limited it could be readily searched manually. The rise of Gopher (created in 1991 by Mark McCahill at the University of Minnesota) led to two new search programs, Veronica and Jughead.

A web search engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. Internet search engines themselves predate

Like Archie, they searched the file names and titles stored in Gopher index systems. Veronica (Very Easy Rodent-Oriented Net-wide Index to Computerized Archives) provided a keyword search of most Gopher menu titles in the entire Gopher listings. Jughead (Jonzy's Universal Gopher Hierarchy Excavation And Display) was a tool for obtaining menu information from specific Gopher servers.

While the name of the search engine "Archie Search Engine" was not a reference to the Archie comic book series, "Veronica" and "Jughead" are characters in the series, thus referencing their predecessor. One of the first "all text" crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also the search engine that was widely known by the public. Also in 1994, Lycos (which started at Carnegie Mellon University) was launched and became a major commercial endeavor. Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek,

statistics

POWER LAW

Where a relative change in one quantity results in a proportional relative change in the other quantity.

**ONE QUANTITY VARIES AS A POWER
OF ANOTHER**

Estimating the exponent from empirical data

Query No. 1 SemiBold

In computing, an event is an action or occurrence recognized by software, often originating asynchronously from the external environment, that may be handled by the software. Computer events can be generated or triggered by the system, by the user or in other ways. Typically, events are handled synchronously with the program flow, that is, the software may have one or more dedicated places where events are handled, frequently an event loop. A source of events includes the user, who may interact

Event driven systems are typically used when there is some asynchronous external activity that needs to be handled by a program; for example, a user who presses a button on their mouse. An event driven system typically runs an event loop, that keeps waiting for such activities, e.g. input from devices or internal alarms. When one of these occurs, it collects data about the event and dispatches the event to the event handler software that will deal with it.

A program can choose to ignore events, and there may be libraries to dispatch an event to multiple handlers that may be programmed to listen for a particular event. The data associated with an event at a minimum specifies what type of event it is, but may include other information such as when it occurred, who or what caused it to occur, and extra data provided by the event source to the handler about how the event should be processed. Events are typically used in user interfaces, where actions in the outside world (mouse clicks, window-resizing, keyboard presses, messages from other programs, etc.) are handled by the program as a series of events. Programs written

The usefulness of a search engine depends on the relevance of the result set it gives back. While there may be millions of web pages that include a particular word or phrase, some pages may be more relevant, popular, or authoritative than others. Most search engines employ methods to rank the results to provide the "best" results first. How a search engine decides which pages are the best matches, and what order the results should be shown in, varies widely from one engine to another. The methods also change over time as Internet usage changes and new techniques evolve. There are two main types of search engine that have evolved: one is a system of predefined and hierarchically ordered keywords that humans have programmed extensively. The other is a system that generates an "inverted index" by analyzing texts it locates. This first form relies much more heavily on the computer itself to do the bulk

for many windowing environments consist predominantly of event handlers.

Events can also be used at instruction set level, where they complement interrupts. Compared to interrupts, events are normally handled synchronously: the program explicitly waits for an event to be serviced (typically by calling an instruction that dispatches the next event), whereas an interrupt can demand service at any time.

In computer programming, an event handler is a callback subroutine that handles inputs received in a program (called a listener in Java and JavaScript). Each event is a piece of application-level information from the underlying framework, typically the GUI toolkit. GUI events include key presses, mouse movement, action selections, and timers expiring. On a lower level, events can represent availability of new data for reading a file or network stream. Event handlers are a central concept in event-driven programming. The events are created by the framework based on interpreting lower-level inputs, which may be lower-level events themselves. For example, mouse

EXCITE

(empty search)

Brin, Sergey; Page, Larry.

“The Anatomy of a Large-Scale Hypertextual Web Search Engine”

WITH FEW TERMS

(e.g., place names, zip codes, geographic features, etc.)

QUERY NO. 1 REGULAR

Indexing means associating words and other definable tokens found on web pages to their domain names and HTML-based fields. The associations are made in a public database, made available for web search queries. A query from a user can be a single word, multiple words or a sentence. The index helps find information relating to the query as quickly as possible. Some of the techniques for indexing, and caching are trade secrets, whereas web crawling is a straightforward process of visiting all sites on a systematic

Between visits by the spider, the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed. The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot.

High-level architecture of a standard Web crawler. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes. Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the

Web search engines get their information by web crawling from site to site. The "spider" checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. "[N]o web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient. Some websites are crawled exhaustively, while

keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing.

Beyond simple keyword lookups, search engines offer their own GUI- or command-driven operators and search parameters to refine the search results. These provide the necessary controls for the user engaged in the feedback loop users create by filtering and weighting while refining the search results, given the initial pages of the first search results. For example, from 2007 the Google.com search engine has allowed one to filter by date by clicking "Show search tools" in the leftmost column of the initial search results page, and then selecting the desired date range. It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search. The engine looks for the words or phrases exactly as entered. Some search

HUMMINGBIRD

August, 2035

Boolean operators and parentheses,
a technique traditionally used by
librarians can be applied

NATURAL LANGUAGE PROCESSING

(electronic OR computerized OR DRE)

QUERY NO. 1 MEDIUM

Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek, Inktomi, Northern ExtraBold, and AltaVista. Information seekers could also browse the directory instead of doing a keyword-based search. In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was

Command–query separation is particularly well suited to a design by contract (DbC) methodology, in which the design of a program is expressed as assertions embedded in the source code, describing the state of the program at certain critical times. In DbC, assertions are considered design annotations – not program logic – and as such, their execution should not affect the program state. CQS is beneficial to DbC because any value-returning method (any query) can be called by any assertion without fear of modifying program state.

In theoretical terms, this establishes a measure of sanity, whereby one can reason about a program's state without simultaneously modifying that state. In practical terms, CQS allows all assertion checks to be bypassed in a working system to improve its performance without inadvertently modifying its behaviour. CQS may also prevent the occurrence of certain kinds of heisenbugs.

Broader impact on software engineering. Even beyond the connection with design by contract, CQS is considered by its adherents to have

Event notification is a term used in conjunction with communications software for linking applications that generate small messages (the "events") to applications that monitor the associated conditions and may take actions triggered by events. Event notification is an important feature in modern database systems (used to inform applications when conditions they are watching for have occurred), modern operating systems (used to inform applications when they should take some action, such as refreshing a window), and modern distributed systems, where the producer of an event might be on a different machine than the consumer, or consumers. Event notification platforms are normally designed so that the application producing events does not need to know which applications will consume them, or even how many applications will monitor the event stream. It is sometimes used as a synonym

a simplifying effect on a program, making its states (via queries) and state changes (via commands) more comprehensible.

CQS is well-suited to the object-oriented methodology, but can also be applied outside of object-oriented programming. Since the separation of side effects and return values is not inherently object-oriented, CQS can be profitably applied to any programming paradigm that requires reasoning about side effects. Command–query separation (CQS) is a principle of imperative computer programming. It was devised by Bertrand Meyer as part of his pioneering work on the Eiffel programming language.

It states that every method should either be a command that performs an action, or a query that returns data to the caller, but not both. In other words, Asking a question should not change the answer. More formally, methods should return a value only if they are referentially transparent and hence possess no side effects. This is the last line of the Medium showing. Thank you.

Query No. 2

A B C D E F G

H I J K L M

N O P Q R S

T U V W

X Y Z

a b c d

e f g h i j k

l m n o p

q r s t u v

w x y z

ÁáĂăĀă ÂâÄäÀà

ĀāĄąÅåÃã ˆ Ææ

ĆćČčÇç ĎdĐđ

ÉéĚěÊêËëÈèĒēĘę ĞğĜğĴģ

Ħħ ĪīĴĵİıǪǫ Łł Ł'ł'L'ł

ŊŋŅņŦŧŃń

ÓóŎňÔôÖöÒòÕõŌōØøŨũ°Œœ

ŔŕŘř ŚśŞşŠšŞş ŢţŤťȚțȚț

ŮůŪūŰű ŰűŲųŪū ÚúŬŭÛù

ŴŵŶŷŴŵŶŷ Ŷŷÿÿ

0 0̣ 1 1̣ 1̣ 2 2̣

3 3̣ 3̣ 4 4̣

5 5̣ 6 6̣ 7 7̣

8 8̣ 9 9̣ #

10/10²³/23⁴⁵/45

6 7/8⁹/67/8 9

0 1 2 3 4 5

6 7 8 9

\$10¢23£45

♥ ♦ ¥67 ♠ ♣

₣88 🍏 €90




SM TEL © P R ® TM






🏠 ○ ⬆️ ● ▲

▶️ ◀️ ▴ ▾

↙ ↘ ↗ ↖ ↕ ↔️ ↠ ↡

◀️ ◁ ▽ ▿ ◻ ◊ ◈ ◼ ▲ ▷ ▶️





¡ ¢ £ ¤ ¥ ¦ § ¨ © ª « ¬ ® ¯ ° ± ² ³ ´ µ ¶ · ¸ ¹ º » ¼ ½ ¾

¿ ˆ ˇ ˘ ˙ ˚ ˛ ˜ ˝ ˞ ˟ ˠ ˡ ˢ ˣ ˤ ˥ ˦ ˧ ˨ ˩ ˪ ˫ ˬ ˭ ˮ ˯ ˰ ˱ ˲ ˳ ˴ ˵ ˶ ˷ ˸ ˹ ˺ ˻ ˼ ˽ ˾ ˿


& § &

†% ‡%00

http://www

we@publictype

¶



Stylistic Set 1

(& 13 il ®®)

103 is a strictly non-palindromic number & 103 is the 27th prime number. The earlier prime is 101 making them both dual primes. It is also a joyful number. [1]®

The 103 series is a DC electric multiple unit (EMU) commuter train type introduced in 1963 by Japanese National Railways, and currently operated by West Japan Railway Company and Kyushu Railway Company. They were also operated by East Japan Railway Company® and Central Japan Railway Company. In 2004, some sets were shipped to Indonesia, where they operate on the 1H3H1 KRL Jabodetabek system in Jakarta. Those

Stylistic Set 2

(CJMNPSUWY cfhijlmnrstuw)

WALTZ, PYPPIH, FOR QUICK JIGS VEX BUD.
 SPRINK OF BLACK QUARTZ, JUDGE MY
 VOW. PACK MY BOX WITH FIVE DOZEN
 LIQUOR JUGS. GLIB JOCKS QUIZ PYPPIH
 TO VEX. JACKDAWS LOVE MY BIG SPRINK
 OF QUARTZ. THE FIVE BOXING WIZARDS
 JUMP QUICKLY. HOW VEXINGLY QUICK
 DAFT ZEBRAS JUMP! Quick zephyrs blow,
 vexing daft Jim. Two driven jocks help fax
 my big quiz. The jay, pig, fox, zebra and my
 wolves quack! Sympathizing would fix
 Quaker objectives. A wizard's job is to vex
 chumps quickly in fog. Watch "Jeopardy!",
 Alex Trebek's fun TV quiz game. By Love, my
 quick study of lexicography won a prize!

Stylistic Set 3

(a)

**Accepting any audible arcane and amorous
answer allows arguments. Anyone around
actively abhors analyzing audacious
applications abounding as apeshit actions,
and attracting amendable assholes and
affirming an appropriate avenue always
allows an absolute age acquittal and
adaptation as an articulate acceptance
amplified aboriginal anguish after an
abomination among associations and
allies adhered any and all accusations
accumulating among asteroid appraisers
annual activities, afterward, all astral
ambulation also allowed astronomers an
antiquated artform arranged as an ageless**

TYPES!

Deepwater Horizon

Single website or web page

PAGE ENTITY

Engines do not disclose their search logs

QUERY NO. 2 BLACK

Indexing means associating words and other definable tokens found on web pages to their domain names and HTML-based fields. The associations are made in a public database, made available for web search queries. From a user can be a single word, multiple words or a sentence. The index helps find information relating to the query quickly. Some of the techniques for indexing, and caching are trade secrets, whereas web crawling is a straightforward process of visiting all sites on a systematic

The Regular NICNAME/WHOIS Server is an NCP/TCP transaction based query/response server, running on the SRI-NIC machine, that provides nationwide directory service to ARPANET users. It is one of a series of ARPANET/Internet name services maintained by the Network Information Center (NIC) at SRI International on behalf of the Defense Communications Agency (DCA). The server is accessible across the ARPANET from user programs running on local hosts, and it delivers the full name, U.S. mailing address, telephone number, and network mailbox for ARPANET users. This server, together with the corresponding Identification Data Base provides online directory look-up equivalent to the ARPANET Directory. DCA strongly encourages network hosts to provide their users with access to this network service. DCA requests that each individual with a directory on an ARPANET host, who is capable of passing traffic across the ARPANET, be registered in the NIC Identification Data Base. To register, send full name, middle initial, U.S. mailing address (including mail stop and full explanation of abbreviations and acronyms), ZIP code, telephone (including Autovon and FTS, if available), and one network

Web search engines get their information by web crawling from site to site. The “spider” checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in BLACK HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. “No web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient. Some

mailbox, via electronic mail to NIC@SRI-NIC.A web search engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. One of the first “all text” crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also the search engine that was widely

SEARCH

conspiracy theories

This leads to an effect called a filter bubble

"giggling"

A phenomenon in which websites use algorithms to selectively guess what information a user would like to see

QUERY NO. 2 THIN

Search engines were also known as some of the brightest stars in the Internet investing frenzy that occurred in the late 1990s. Several companies entered the market spectacularly, receiving record gains during their initial thin offerings. Some have taken down their public search engine, and are marketing enterprise-only editions, such as Northern Light. Many search engine companies were caught up in the dot-com bubble, a speculation-driven market boom that peaked in 1999 and ended in 2001. The search engine became so

Between visits by the spider, the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed. The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes. Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing. Between visits by the spider,

In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was indexing, predating the very similar algorithm patent filed by Google two years later in 1998. Larry Page referenced Li's work in some of his U.S. patents for PageRank. Li later used his Rankdex technology for the Baidu search engine, which was founded by Robin Li in China and launched in 2000. Search engines were also known as some of the brightest stars in the Internet investing frenzy that occurred in the late 1990s. Several companies entered the market spectacularly, receiving record gains during their initial public offerings. Some have taken down their public search engine, and are marketing enterprise-only editions, such as Northern Light. Many search engine companies were caught up in the dot-com

the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed.[25] The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes.[25] Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing.

INTERNET

TRANSACTIONAL

Reflect the intent of the user to perform a particular action, purchasing a type family.

Search Engines

OFTEN SUPPORT A FOURTH TYPE, USED FAR LESS FREQUENTLY

Query No. 2 ExtraBold

Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek, Inktomi, Northern ExtraBold, and AltaVista. Information seekers could also browse the directory instead of doing a keyword-based search. In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was

It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search. The engine looks for the words or phrases exactly as entered. Some search engines provide an advanced feature called proximity search, which allows users to define the distance between EXTRA BOLD keywords. There is also concept-based searching where the research involves using statistical analysis on pages containing the words or phrases you search for. As well, natural language queries allow the user to type a question in the same form one would ask it to a human. When we say human, we mean a person like Waldo. A site like this would be ask dot com. It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search. The engine looks for the

Command–query separation is particularly well suited to a design by contract (DbC) methodology, in which the design of a program is expressed as assertions embedded in the source code, describing the state of the program at certain critical times. In DbC, assertions are considered design annotations – not program logic – and as such, their execution should not affect the program state. CQS is beneficial to DbC because any value-returning method (any query) can be called by any assertion without fear of modifying program state. In theoretical terms, this establishes a measure of sanity, whereby one can reason about a program's state without simultaneously modifying that state. In practical terms, CQS allows all assertion checks to be bypassed in a working system to improve its performance without inadvertently modifying its behaviour. CQS may also prevent the occurrence of certain

words or phrases exactly as entered. Some search engines provide an advanced feature called proximity search, which allows users to define the distance between keywords. There is also concept-based searching where the research involves using statistical analysis on pages containing the words or phrases you search for. As well, natural language queries allow the user to type a question in the same form one would ask it to a human. A site like this would be ask.com. Microsoft first launched MSN Search in the fall of 1998 using search results from Inktomi. In early 1999 the site began to display listings from Looksmart, blended with results from Inktomi. For a short time in 1999, MSN Search used results from AltaVista instead. In 2004, Microsoft began a transition to its own search technology, powered by its own web crawler (called MSNBOT). Microsoft's rebranded search engine, Bing, was launched on June 1, 2009. On July 29, 2009, Yahoo! and Microsoft finalized a deal in which Yahoo! Search would be powered by Microsoft Bing technology. As of 2019, active search engine crawlers include those of Google, Sogou, Baidu, Bing, Gigablast, Mojeek, DuckDuckGo and Yandex.

CONTENT INFORMATIONAL

Cover a broad topic (e.g., FONTS or DOGS)
for which there may be thousands of relevant results

Characteristics

broad categories that cover most web

Query No. 2 ExtraLight

Beyond simple keyword lookups, search engines offer their own GUI-EXTRALIGHT or command-driven operators and search parameters to refine the search results. These provide the necessary controls for the user engaged in the feedback loop users create by filtering and weighting while refining the search results, given the initial pages of the first search results. For example, from 2007 the Google.com search engine has allowed one to filter by date by clicking "Show search tools" in the leftmost column of the initial

The Regular NICNAME/WHOIS Server is an NCP/TCP transaction based query/response server, running on the SRI-NIC machine, that provides worldwide directory service to ARPANET users. It is one of a series of ARPANET/Internet name services maintained by the Network Information Center (NIC) at SRI International on behalf of the Defense Communications Agency (DCA). The server is accessible across the ARPANET from user programs running on local hosts, and it delivers the full name, U.S. mailing address, telephone number, and network mailbox for ARPANET users. This server, together with the corresponding Identification Data Base provides online directory look-up equivalent to the ARPANET Directory. DCA strongly encourages network hosts to provide their users with access to this network service. DCA requests that each individual with a directory on an ARPANET host, who is capable of passing traffic across the ARPANET, be registered in the NIC Identification Data Base. To register, send full name, middle initial, U.S. mailing address (including mail stop and full explanation of abbreviations and acronyms), ZIP code, telephone (including Autovon and FTS, if available), and one network mailbox, via electronic mail to NIC@SRI-NIC.A web search

Web search engines get their information by web crawling from site to site. The "spider" checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in BLACK HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. "No web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient. Some websites are crawled exhaustively, while

engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. One of the first "all text" crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also the search engine that was widely known by the public. Also in 1994, Lycos (which started at Carnegie Mellon University) was launched

CSS, AJAX

TRANSACTIONAL

A little less than a third of users entered three or more shocking unique queries.

AND, OR, + NOT

The average length of a search query was 2.4 terms

Query No. 2 Bold

In February 2010, Mark McCahill was revealed by the philosopher Peter Ludlow (also known by the pseudonym Urizenus Sklar) to be the Internet persona Pixeleen Mistral, a noted "tabloid reporter" covering virtual worlds who was the editor of Ludlow's newspaper The Alphaville Herald. In a 2016 interview with Leo Laporte, McCahill said that his involvement with developing the Croquet Project had led him into contact with Second Life and that he had become interested in the sociology of virtual worlds. As Pixeleen

The Gopher protocol is a communications protocol designed for distributing, searching, and retrieving documents in Internet Protocol networks. The design of the Gopher protocol and user interface is menu-driven, and presented an alternative to the World Wide Web in its early stages, but ultimately fell into disfavor, yielding to the Hypertext Transfer Protocol (HTTP). The Gopher ecosystem is often regarded as the effective predecessor of the World Wide Web.

The protocol was invented by a team led by Mark P. McCahill at the University of Minnesota. It offers some features not natively supported by the Web and imposes a much stronger hierarchy on the documents it stores. Its text menu interface is well-suited to computing environments that rely heavily on remote text-oriented computer terminals, which were still common at the time of its creation in 1991, and the simplicity of its protocol facilitated a wide variety of client implementations. More recent Gopher revisions and graphical clients added support for multimedia. Gopher was preferred by many network administrators for using fewer network resources than Web

The rise of GOPHER (created in 1991 by Mark McCahill at the University of Minnesota) led to two new search programs, Veronica and Jughead. Like Archie, they searched the file names and titles stored in Gopher index systems. Veronica (Very Easy Rodent-Oriented Net-wide Index to Computerized Archives) provided a keyword search of most Gopher menu titles in the entire Gopher listings. Jughead (Jonzy's Universal Gopher Hierarchy Excavation And Display) was a tool for obtaining menu information from specific Gopher servers. While the name of the search engine "Archie Search Engine" was not a reference to the Archie comic book series, "Veronica" and "Jughead" are characters in the series, thus referencing their predecessor. In the late 1980s, McCahill led the team at the University of Minnesota that developed POPmail, one of the first popular Internet e-mail clients. At about the same time as POPmail was being

services. Gopher's hierarchical structure provided a platform for the first large-scale electronic library connections. The Gopher protocol is still in use by enthusiasts, and although it has been almost entirely supplanted by the Web, a small population of actively-maintained servers remains. Gopher system was released in mid-1991 by Mark P. McCahill, Farhad Anklesaria, Paul Lindner, Daniel Torrey, and Bob Alberti of the University of Minnesota in the United States. Its central goals were, as stated in RFC 1436: A file-like hierarchical arrangement that would be familiar to users. A simple syntax. A system that can be created quickly and inexpensively. Extending the file system metaphor, such as searches. Gopher combines document hierarchies with collections of services, including WAIS, the Archie and Veronica search engines, and gateways to other information systems such as File Transfer Protocol (FTP) and Usenet. The general interest in campus-wide information systems (CWISs) in higher education at the time, and the ease of setup of Gopher servers to create an instant CWIS with links to other sites' online directories and resources were the factors contributing to Gopher's rapid adoption.

HTML DOM Connectivity

Report on the connectivity of the indexed web graph
(e.g., Which links point to this > URL?)

BOOLEAN OPERATORS

How many pages are indexed from this domain name?

QUERY NO. 2 LIGHT

In the summer of 1993, no search engine existed for the web, though numerous specialized catalogues were maintained by hand. Oscar Nierstrasz at the University of Geneva wrote a series of Perl scripts that periodically mirrored these pages and rewrote them into a standard format. This formed the basis for W3Catalog, the web's first primitive search engine, released on September 2, 1993. In June 1993, Matthew Gray, then at MIT, produced what was probably the first web robot, the Perl-based World Wide Web Wanderer, and

Prior to September 1993, the World Wide Web was entirely indexed by hand. There was a list of webservers edited by Tim Berners-Lee and hosted on the CERN webserver. One snapshot of the list in 1992 remains,[4] but as more and more web servers went online the central list could no longer keep up. On the NCSA site, new servers were announced under the title "What's New!"

The first tool used for searching content (as opposed to users) on the Internet was Archie. The name stands for "archive" without the "v". It was created by Alan Emtage, Bill Heelan and J. Peter Deutsch, computer science students at McGill University in Montreal, Quebec, Canada. The program downloaded the directory listings of all the files located on public anonymous FTP (File Transfer Protocol) sites, creating a searchable database of file names; however, Archie Search Engine did not index the contents of these sites since the amount of data was so limited it could be readily searched manually. The rise of Gopher (created in 1991 by Mark McCahill at the University of Minnesota) led to two new search programs, Veronica and Jughead.

A web search engine or Internet search engine is a software system that is designed to carry out web search (Internet search), which means to search the World Wide Web in a systematic way for particular information specified in a textual web search query. The search results are generally presented in a line of results, often referred to as search engine results pages (SERPs). The information may be a mix of links to web pages, images, videos, infographics, articles, research papers, and other types of files. Some search engines also mine data available in databases or open directories. Unlike web directories, which are maintained only by human editors, search engines also maintain real-time information by running an algorithm on a web crawler. Internet content that is not capable of being searched by a web search engine is generally described as the deep web. Internet search engines themselves predate the debut of the Web

Like Archie, they searched the file names and titles stored in Gopher index systems. Veronica (Very Easy Rodent-Oriented Net-wide Index to Computerized Archives) provided a keyword search of most Gopher menu titles in the entire Gopher listings. Jughead (Jonzy's Universal Gopher Hierarchy Excavation And Display) was a tool for obtaining menu information from specific Gopher servers.

While the name of the search engine "Archie Search Engine" was not a reference to the Archie comic book series, "Veronica" and "Jughead" are characters in the series, thus referencing their predecessor. One of the first "all text" crawler-based search engines was WebCrawler, which came out in 1994. Unlike its predecessors, it allowed users to search for any word in any webpage, which has become the standard for all major search engines since. It was also the search engine that was widely known by the public. Also in 1994, Lycos (which started at Carnegie Mellon University) was launched and became a major commercial endeavor. Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek,

statistics

POWER LAW

Where a relative change in one quantity results in a proportional relative change in the other quantity.

ONE QUANTITY VARIES AS A POWER
OF ANOTHER

Estimating the exponent from empirical data

Query No. 2 SemiBold

In computing, an event is an action or occurrence recognized by software, often originating asynchronously from the external environment, that may be handled by the software. Computer events can be generated or triggered by the system, by the user or in other ways. Typically, events are handled synchronously with the program flow, that is, the software may have one or more dedicated places where events are handled, frequently an event loop. A source of events includes the user, who may interact

Event driven systems are typically used when there is some asynchronous external activity that needs to be handled by a program; for example, a user who presses a button on their mouse. An event driven system typically runs an event loop, that keeps waiting for such activities, e.g. input from devices or internal alarms. When one of these occurs, it collects data about the event and dispatches the event to the event handler software that will deal with it.

A program can choose to ignore events, and there may be libraries to dispatch an event to multiple handlers that may be programmed to listen for a particular event. The data associated with an event at a minimum specifies what type of event it is, but may include other information such as when it occurred, who or what caused it to occur, and extra data provided by the event source to the handler about how the event should be processed. Events are typically used in user interfaces, where actions in the outside world (mouse clicks, window-resizing, keyboard presses, messages from other programs, etc.) are handled by the program as a series of events. Programs written

The usefulness of a search engine depends on the relevance of the result set it gives back. While there may be millions of web pages that include a particular word or phrase, some pages may be more relevant, popular, or authoritative than others. Most search engines employ methods to rank the results to provide the "best" results first. How a search engine decides which pages are the best matches, and what order the results should be shown in, varies widely from one engine to another. The methods also change over time as Internet usage changes and new techniques evolve. There are two main types of search engine that have evolved: one is a system of predefined and hierarchically ordered keywords that humans have programmed extensively. The other is a system that generates an "inverted index" by analyzing texts it locates. This first form relies much more heavily on the computer itself to do the bulk of the

for many windowing environments consist predominantly of event handlers.

Events can also be used at instruction set level, where they complement interrupts. Compared to interrupts, events are normally handled synchronously: the program explicitly waits for an event to be serviced (typically by calling an instruction that dispatches the next event), whereas an interrupt can demand service at any time.

In computer programming, an event handler is a callback subroutine that handles inputs received in a program (called a listener in Java and JavaScript). Each event is a piece of application-level information from the underlying framework, typically the GUI toolkit. GUI events include key presses, mouse movement, action selections, and timers expiring. On a lower level, events can represent availability of new data for reading a file or network stream. Event handlers are a central concept in event-driven programming. The events are created by the framework based on interpreting lower-level inputs, which may be lower-level events themselves. For example, mouse movements and

EXCITE

(empty search)

Engines use web crawlers that will eventually find most web sites on the Internet without assistance

WITH FEW TERMS

(e.g., place names, zip codes, geographic features, etc.)

QUERY NO. 2 REGULAR

Indexing means associating words and other definable tokens found on web pages to their domain names and HTML-based fields. The associations are made in a public database, made available for web search queries. A query from a user can be a single word, multiple words or a sentence. The index helps find information relating to the query as quickly as possible. Some of the techniques for indexing, and caching are trade secrets, whereas web crawling is a straightforward process of visiting all sites on a systematic

Between visits by the spider, the cached version of page (some or all the content needed to render it) stored in the search engine working memory is quickly sent to an inquirer. If a visit is overdue, the search engine can just act as a web proxy instead. In this case the page may differ from the search terms indexed. The cached page holds the appearance of the version whose words were previously indexed, so a cached version of a page can be useful to the web site when the actual page has been lost, but this problem is also considered a mild form of linkrot.

High-level architecture of a standard Web crawler. Typically when a user enters a query into a search engine it is a few keywords. The index already has the names of the sites containing the keywords, and these are instantly obtained from the index. The real processing load is in generating the web pages that are the search results list: Every page in the entire list must be weighted according to information in the indexes. Then the top search result item requires the lookup, reconstruction, and markup of the snippets showing the context of the

Web search engines get their information by web crawling from site to site. The "spider" checks for the standard filename robots.txt, addressed to it. The robots.txt file contains directives for search spiders, telling it which pages to crawl. After checking for robots.txt and either finding it or not, the spider sends certain information back to be indexed depending on many factors, such as the titles, page content, JavaScript, Cascading Style Sheets (CSS), headings, or its metadata in HTML meta tags. After a certain number of pages crawled, amount of data indexed, or time spent on the website, the spider stops crawling and moves on. "[N]o web crawler may actually crawl the entire reachable web. Due to infinite websites, spider traps, spam, and other exigencies of the real web, crawlers instead apply a crawl policy to determine when the crawling of a site should be deemed sufficient. Some websites are crawled exhaustively, while others are

keywords matched. These are only part of the processing each search results web page requires, and further pages (next to the top) require more of this post processing.

Beyond simple keyword lookups, search engines offer their own GUI- or command-driven operators and search parameters to refine the search results. These provide the necessary controls for the user engaged in the feedback loop users create by filtering and weighting while refining the search results, given the initial pages of the first search results. For example, from 2007 the Google.com search engine has allowed one to filter by date by clicking "Show search tools" in the leftmost column of the initial search results page, and then selecting the desired date range. It's also possible to weight by date because each page has a modification time. Most search engines support the use of the boolean operators AND, OR and NOT to help end users refine the search query. Boolean operators are for literal searches that allow the user to refine and extend the terms of the search. The engine looks for the words or phrases exactly as entered. Some search

HUMMINGBIRD

August, 2035

Boolean operators and parentheses, a technique traditionally used by librarians can be applied

NATURAL LANGUAGE PROCESSING

(Electronic OR Computerized OR DRE)

QUERY NO. 2 MEDIUM

Soon after, a number of search engines appeared and vied for popularity. These included Magellan, Excite, Infoseek, Inktomi, Northern ExtraBold, and AltaVista. Information seekers could also browse the directory instead of doing a keyword-based search. In 1996, Robin Li developed the RankDex site-scoring algorithm for search engines results page ranking and received a US patent for the technology. It was the first search engine that used hyperlinks to measure the quality of websites it was

Command–query separation is particularly well suited to a design by contract (DbC) methodology, in which the design of a program is expressed as assertions embedded in the source code, describing the state of the program at certain critical times. In DbC, assertions are considered design annotations – not program logic – and as such, their execution should not affect the program state. CQS is beneficial to DbC because any value-returning method (any query) can be called by any assertion without fear of modifying program state.

In theoretical terms, this establishes a measure of sanity, whereby one can reason about a program's state without simultaneously modifying that state. In practical terms, CQS allows all assertion checks to be bypassed in a working system to improve its performance without inadvertently modifying its behaviour. CQS may also prevent the occurrence of certain kinds of heisenbugs.

Broader impact on software engineering. Even beyond the connection with design by contract, CQS is considered by its adherents to have

Event notification is a term used in conjunction with communications software for linking applications that generate small messages (the "events") to applications that monitor the associated conditions and may take actions triggered by events. Event notification is an important feature in modern database systems (used to inform applications when conditions they are watching for have occurred), modern operating systems (used to inform applications when they should take some action, such as refreshing a window), and modern distributed systems, where the producer of an event might be on a different machine than the consumer, or consumers. Event notification platforms are normally designed so that the application producing events does not need to know which applications will consume them, or even how many applications will monitor the event stream. It is sometimes used as a

a simplifying effect on a program, making its states (via queries) and state changes (via commands) more comprehensible.

CQS is well-suited to the object-oriented methodology, but can also be applied outside of object-oriented programming. Since the separation of side effects and return values is not inherently object-oriented, CQS can be profitably applied to any programming paradigm that requires reasoning about side effects. Command–query separation (CQS) is a principle of imperative computer programming. It was devised by Bertrand Meyer as part of his pioneering work on the Eiffel programming language.

It states that every method should either be a command that performs an action, or a query that returns data to the caller, but not both. In other words, Asking a question should not change the answer. More formally, methods should return a value only if they are referentially transparent and hence possess no side effects. This is the last line of the Medium showing. Thank you.

We Are Public

About

Public is a collective of typeface and graphic designers. We have a boundless interest in constructing form to communicate. We are for curious discovery. We are advocates for type design. We work on custom typefaces for companies of all shapes and sizes. We contribute to a local and worldwide community. We are for sharing knowledge. We facilitate workshops. We use some of our profits to fund scholarships for aspiring designers who need a leg up. We are for tearing down walls. We count on a better tomorrow. We are Public.