Tech Review | http://www.technologyreview.com/articles/05/01/issue/ferguson0105.asp?p=0 January 2005

What's Next for Google

by Charles H. Ferguson

For Eric Schmidt, Google's CEO, 2004 was a very good year. His firm led the search industry, the fastest-growing major sector in technology; it went public, raising $1.67 billion; its stock price soared; and its revenues more than doubled, to $3 billion. But as the search market ripens into something worthy of Microsoft's attention, those familiar with the software business have been wondering whether Google, apparently triumphant, is in fact headed off the cliff.

I've seen it happen before. In September 1995, I had breakfast with Jim Barksdale, then CEO of Netscape Communications, at Il Fornaio in Palo Alto, CA, a restaurant popular with Silicon Valley dealmakers. Netscape had gone public a few months earlier, and Netscape Navigator dominated the browser market. Vermeer Technologies, the company that Randy Forgaard and I had founded 18 months earlier, had just announced the release of FrontPage, a Windows application that let people develop their own websites. Netscape and Microsoft were both preparing to develop competing products. Our choice was to stay independent and die or sell the company to one of them.

At breakfast, and repeatedly over the following months, I tried to persuade Barksdale to take Microsoft seriously. I argued that if it was to survive, Netscape needed to imitate Microsoft's strategy: the creation and control of proprietary industry standards. Serenely, Barksdale explained that Netscape actually invited Microsoft to imitate its products, because they would never catch up. The Internet, he said, rewarded openness and nonproprietary standards. When I heard that, I realized that despite my reservations about the monopolist in Redmond, WA, I had little choice. Four months later, I sold my company to Microsoft for $130 million in Microsoft stock. Four years later, Netscape was effectively dead, while Microsoft's stock had quadrupled.

Google now faces choices as fundamental as those Netscape faced in 1995. Google, whose headquarters in Mountain View, CA—familiarly called the Googleplex—is only five kilometers from Netscape's former home, needn't perish as Netscape did, but it could. Despite everything Google has—the swelling revenues, the cash from its initial public offering, the 300 million users, the brand recognition, the superbly elegant engineering—its position is in fact quite fragile. Google's site is still the best Web search service, and Gmail, its new Web-based e-mail service, Google Desktop, its desktop search tool, and Google Deskbar, its toolbar, are very cool. But that's all they are. As yet, nothing prevents the world from switching (painlessly, instantly) to Microsoft search services and software, particularly if they are integrated with the Microsoft products that people already use.

In November 2004, Microsoft launched a beta, or test, version of a search engine designed to answer questions posed in everyday language and to serve results customized to users' geographical locations. Microsoft has also created additional search software for its Internet Explorer browser and its Office productivity applications. That Microsoft is developing its own Web search engine and desktop search tools is significant in itself. But its competition with Google will have repercussions far beyond the existing search business—or even the software industry itself. Google and Microsoft will be fighting to control the organization, search, and retrieval of all digital information, on all types of digital devices. Collectively, these markets are much larger than the existing market for search services. Over the next several decades, in the view of search industry insiders I've spoken with, they could generate perhaps half a trillion dollars in cumulative revenue.

Microsoft is starting late but has extraordinary resources and powers of persistence—and it joined the browser wars late, too. In contrast, Google is youthful, adventurous, and innovative, and it does some things extremely well. The contest could end in a Cold War standoff, a decisive victory for either side, or even mutual destruction, if the competition frightens away customers and investors.

Peaceful coexistence, however, seems unlikely.

THE PRIZE AND THE CONTESTANTS

Eric Schmidt and Microsoft's Bill Gates will be competing against each other for the third time. For both men, the contest is personal as well as financial.

Gates's philanthropic ambitions depend on Microsoft's continued health. And like a rock star who yearns to be admired for his brains, Gates wants to create new technology. Only by doing so can he overcome his reputation as the college dropout who built his empire by turning other people's ideas into mediocre products. "Bill Gates is desperate to prove that he can innovate," commented a Microsoft executive who prefers to remain anonymous. "And it just might kill us." He pointed to the ambitious goals and long delays that have plagued Longhorn, Microsoft's future (and search-centric) version of Windows.

By contrast, the three men who run Google have impeccable technology credentials. Schmidt has a PhD from the University of California, Berkeley, did research at Xerox PARC, and became chief technology officer of Sun Microsystems, where he oversaw the development of many impressive technologies. In business, however, Schmidt has twice been beaten by Gates. The first time was at Sun; the second was at Novell, where Schmidt was CEO. Both firms made enormous mistakes. Schmidt wasn't entirely responsible, however, because his hands were tied by his superiors at Sun and by his predecessors at Novell. At Google, Schmidt must once again share power—with Larry Page and Sergey Brin, Google's brilliant but young and possibly overconfident founders, both "on leave" from Stanford University's PhD program in computer science. Page and Brin still call many of the shots, and the company's unusual capital structure gives them about 30 percent of the voting shares.

Google seeks to become the gatekeeper for not only the public Web but also the "dark" or hidden Web of private databases, dynamically generated pages, controlled-access sites, and Web servers within organizations (estimated to be tens or even hundreds of times larger than the public Web); the data on personal computer hard drives; and the data on consumer devices ranging from PDAs to cell phones to iPods to digital cameras to TiVo players. Google's founders understand the scale of the opportunity. Larry Page recently said, "Only a fraction of the world's information is indexed on our computers. We are continually working on new ways to index more. . . . Thirty percent [of our engineers] are devoted to emerging businesses." And Sergey Brin once told Technology Review's editor in chief, "The perfect search engine would be like the mind of God."

Until now, competition in the search industry has been limited to the Web and has been conducted algorithm by algorithm, feature by feature, and site by site. This competition has resulted in a Google and Yahoo duopoly. If nothing were to change, the growth of Microsoft's search business would only create a broader oligopoly, similar, perhaps, to those in other media markets. But the search industry will soon serve more than just a Web-based consumer market. It will also include an industrial market for enterprise software products and services, a mass market for personal productivity and communications software, and software and services for a sea of new consumer devices. Search tools will comb through not only Microsoft Office and PDF documents, but also e-mail, instant messages, music, and images; with the spread of voice recognition,

Internet telephony, and broadband, it will also be possible to index and search telephone conversations, voice mail, and video files.

All these new search products and services will have to work with each other and with many other systems. This, in turn, will require standards.

The emergence of search standards would encourage tremendous growth and provide many benefits to users. But standardization would also introduce a new and destabilizing force into the industry. Instead of competing through incremental improvements in the quality and range of their search services, Microsoft, Google, and Yahoo will be forced into a winner-take-all competition for control of industry standards. Steve Jurvetson, a venture capitalist at the firm of Draper Fisher Jurvetson in Menlo Park, CA, says, "This is something of a holy war for Microsoft, and one they can't bear to lose."

In short, the search industry is ready for an architecture war.

PURSUING LOCK-IN

Architecture wars (also known as standards wars) occur because information technology markets require standards in order to manage complexity, communication, and technological change. Historically, proprietary control over a major information technology standard has created more wealth than nearly any other human activity. Architectural dominance mints money; and managed properly, it lasts forever. IBM's mainframe architecture was introduced in 1964; Intel developed its first microprocessor in 1971; Microsoft's first operating system was introduced in 1981; Cisco Systems marketed its first router in 1986. None shows any signs of disappearing, and each has already generated hundreds of billions of dollars in cumulative revenues.

It is only standardization that makes it possible for any browser to display any Web page, or for people to read the documents and e-mail messages they receive from each other. Standards are generally based upon the interfaces that constitute the authorized ways for software systems to communicate with each other. These include application programming interfaces, or APIs, like those Microsoft provides for developing Windows applications; communications protocols such as HTTP (the hypertext transfer protocol), which allows browsers to communicate with websites; and content or document structures, such as the HTML (hypertext markup language) standard for Web pages, or the document structure used by Microsoft Word. These standards are embedded in larger architectures used in the design of general-purpose commercial systems, or platforms, such as the Windows operating system. Platforms, in turn, are used as the starting point for specific applications, such as word processors or accounting systems.

Sometimes standardization is achieved through nonproprietary efforts managed by governments, standards bodies, or industry coalitions. Examples include the basic Internet protocols, the HDTV broadcasting standard, and most telephone standards. In other cases, like that of the Ethernet protocol invented by Bob Metcalfe while at Xerox PARC, a company donates an architecture to a standards body in the hope of creating or expanding a market. The open-source movement is an interesting variant of nonproprietary standardization based on decentralized control. In the case of open-source software like the Linux operating system, a community of creators and users in effect votes continuously on the direction of a standard.

But in most information technology markets, standardization is achieved via market competition. These contests are extremely complex, but they have a common underlying logic, which Charles Morris and I described a decade ago in our book Computer Wars. The best technology does not always win; superior strategy is often more important. Winners do tend, however, to share several important characteristics. They

provide general-purpose, hardware-independent architectures, like Microsoft's operating systems, rather than bundled hardware and software, like Apple's and Sun's systems. Winning architectures are proprietary and difficult to clone, but they are also externally "open"—that is, they provide publicly accessible interfaces upon which a wide variety of applications can be constructed by independent vendors and users. In this way, an architecture reaches all markets, and also creates "lock-in"—meaning that users become captive to it, unable to switch to rival systems without great pain and expense.

Architecture wars generally begin with a fierce competition for market share. Eventually, the market settles on a de facto standard, a dominant architecture under the proprietary control of one company. Subsequently, only a few rivals survive in the leader's shadow, while the leader expands its empire into neighboring markets.

The search industry is the next place in which a vast architectural empire could be built. Some portions of the emerging search space are now occupied by Google, others by Microsoft, most by nobody. But in the end, there will probably be room for just one architecture. Google's idyllic childhood must therefore give way to a contest much like those Microsoft has fought and won against companies ranging from IBM to Novell to Apple to Netscape. But for several reasons, this architecture war may end differently. First, many of the companies defeated by Microsoft over the past 20 years suffered as much from self-inflicted wounds as from Microsoft's predation. In Eric Schmidt, Google may have a CEO with the technological depth and painfully acquired experience essential to surviving Bill Gates. Second, Google's principal services run on a platform that Microsoft doesn't control—the Web. Third, in some cases (like its fight against Linux, for example), Microsoft's software is now the high-cost incumbent.

Fourth, some analysts believe that Microsoft has lost its edge, that its size and age have bred complacency. Commenting on the collision between Google and Microsoft, Internet industry observer John Battelle recently wrote, "Microsoft is indeed a fearsome competitor, with extraordinary resources (and I don't mean the $50 billion in cash; I mean the ability to leverage Windows). But it's a middle-aged company that moves far more slowly than it did ten years ago, when it first recognized the Web threat." (For John Battelle's views on the future of publishing, see "Megaphone," p. 36.)

Fifth, Microsoft hasn't always won: Adobe and Intuit are doing just fine, MSN hasn't killed AOL or Yahoo, and the Xbox hasn't defeated the Japanese game industry (not yet, anyway). And finally, Microsoft's recent entry in the search wars—the beta version of MSN's search tool—is decidedly unimpressive. (Then again, Windows 1.0 was pretty bad, too.)

So Google's defeat is not a foregone conclusion. Indeed, if it does everything right, it could become an enormously powerful and profitable company, representing the most serious challenge Microsoft has faced since the Apple Macintosh. But if Microsoft gets serious about search—and there is every reason to believe that it will—Google will need brilliant strategy and flawless execution simply to survive.

ARMING SECRETLY

Does Google understand the gravity of the challenges that may confront it? Does it have a strategy for winning an architectural war? The evidence is equivocal.

Google has software developers skilled enough to construct a powerful architectural position. It has hired both newly minted PhDs and experienced technologists from Netscape and even Microsoft. One of its newer employees is Adam Bosworth, famous to software developers for developing the HTML engine in Microsoft's Internet Explorer and for his pioneering work on the Extensible Markup Language, or XML, the standard for machine-to-machine communication on the Web. Other recent hires, significant

for their architectural expertise, include Rob Pike, a pioneer of the Unix operating system at Bell Labs; Joshua Bloch, a leading Java coder from Sun; and Cédric Beust, who developed the Weblogic platform at BEA Systems.

One Google manager, who preferred not to be named, said his company understands the need for proprietary control, and that future products would prove it. In late 2004, Google did release two important new APIs, for its Deskbar search tool and its advertising systems. But the Google executive declined to comment on future plans, noting that his employer had become secretive to the point of paranoia. (Indeed, Google's senior executives refused to be interviewed for this article.)

The executive then went on to say, "Look, everyone here—right up to our CEO and board of directors—has had the shit kicked out of them over the last five years. A lot of them were at Netscape, or at failed dot coms. Nobody I work with is complacent, and they're all very smart." But there are two important people who haven't had the shit kicked out of them: Google's founders. In a Playboy interview published shortly before Google's IPO, Brin and Page did not mention competitive threats. Rather, they talked about corporate ethics, the creation of foundations, and their efforts to make Google a great place to work.

Google is a great place to work. My friends there absolutely love the place, and in part for that reason, they work very hard. Google allows pets and provides employees with laundry service, drinks, meals, massages, car washes, and (soon) child care. Its corporate motto is "Don't be evil." But long ago, a professor of mine, noting my youthful idealism, remarked that the only successful neutral nations are those, like Switzerland, that are permanently armed to the teeth. And for Google, neutrality is not an option.

But what specifically should Google do? How is Microsoft likely to attack, what will the contest look like, and what will decide its outcome? Let's begin with the current state of search.

THE STATE OF SEARCH

For a long time, search engines were expensive luxuries for those who operated them. They never made money. Market leadership traded hands repeatedly. Sites like AltaVista rose to prominence and fell away. The entirely separate business of selling software products for text indexing and retrieval was a backwater. But then things changed. As the Internet and the Web grew, searchable digital content began to supplant conventional media, and efforts to improve the quality of search results intensified.

Early search engines ranked results largely according to crude criteria such as the number of times a page mentioned the user's chosen keywords. But in a research collaboration that began in 1995, when they were still graduate students, Brin and Page applied a practice called citation ranking to the Web, and it turned out to be a much more reliable way to find relevant information.

For many years, reference publications like the Science Citation Index have ranked scientific papers' "impact" by counting the number of times they were cited in other papers. Brin and Page's insight was that if hyperlinks were viewed as citations, the same thing could be done for the Web. That insight led to the first truly superior search engine. Stanford applied for a patent on Brin and Page's "PageRank" technique in 1998 (it was granted in 2001). Soon afterward, Brin and Page started Google and raised money from top-tier venture capital firms Sequoia Capital and Kleiner, Perkins, Caufield, and Byers.

Today, the search industry has two layers. The leaders, Google and Yahoo, both provide "retail" search services on their own websites. But both firms also license, on a

highly selective basis, their infrastructure and services to other companies in a "wholesale" market. For example, Google provides the underlying search services for AOL and Amazon.com's A9 search subsidiary. Looksmart powered MSN Search for some years. Now, however, Microsoft is developing its own search engine.

Google holds nearly 40 percent of the U.S. retail search market, more than 50 percent of the U.S. wholesale market, and larger shares of the global market. Yahoo enjoys a rough parity with Google in the United States, and Baidu has been expanding in China. Interestingly, while Google operates its own service in China, it also holds an equity stake in Baidu.

Google derives nearly all of its revenues from advertising, of two distinct kinds. First, it places advertisements on pages of search results returned by its own site. Those advertisements are selected according to the words used in the search. Advertisers bid in highly complex auctions for the right to place ads on results pages for searches that use specific terms like "used cars," "SUVs," and so forth. Second, Google manages advertising for a wide network of external websites for which it provides ad placement services. It has combined its search engine with sophisticated text-matching and auction systems to target, price, sell, and evaluate its advertisements, both those placed on its own site and those on its affiliates'.

Some of these affiliates use Google's search services, but most do not. In fact, almost half of Google's revenue and profits come from its external advertising network, a business where its superior indexing and search capabilities play a less critical role. Google also sells a "search appliance," a Linux server running its indexing and search software, to organizations wishing to provide search services for their internal Web servers. This business, however, is quite small.

Yahoo's search business is similar. Like Google, Yahoo earns a substantial fraction of its total revenue through search-related advertising, both on its own site and on a network of affiliates. Yahoo's portal offers a wider variety of information services than Google, including news, dating, chat, and shopping. But Google is rapidly diversifying: in addition to allowing users to download its free personal search tool, Google's website has news, shopping, e-mail, and photo storage services in various stages of development.

Today, the wholesale search market has significant barriers to entry. Economies of scale have asserted themselves, secondary competitors have folded, and the creation of new search engines by startups is becoming prohibitively expensive. Consider: to crawl, index, and search more than eight billion pages—still only a fraction of the Web—Google now operates a global infrastructure of more than 250,000 Linux-based servers of its own design, according to one Google executive I spoke with, and it is becoming a major consumer of electrical power, computer hardware, and telecommunications bandwidth.

But the consolidation of the wholesale market does not mean that the search industry is mature. Quite the contrary.

First, there is no lack of new competition. This comes from any number of sources: large firms, like Amazon and its A9 subsidiary, with sufficient resources to enter the market; startups commercializing a wide variety of new search functions; information retrieval and filtering firms such as LexisNexis or Viv'simo, whose products are competitive with or complementary to Web-based search services; and, in a class by itself, Microsoft. Moreover, while basic Web crawling is a mature technology with high barriers to entry, many other search-related functions are not. Secondly, services that have thus far been confined almost exclusively to the public Web are now expanding to personal computers, the dark Web, and other platforms. Finally, the search arena is still unstructured and without standards. Search sites are self-contained islands. They do not interoperate, and independent developers cannot use search sites as platforms upon

which to offer specialized products and services, because, with minor exceptions, the search industry lacks open APIs. For the most part, each service is confined to what it can do on its own.

But the search industry cannot resist APIs, standards, and open architectures much longer. No single company can offer users all the functions they want. Users will demand search products and services that work across many different platforms. And Microsoft will almost certainly exploit both its ownership of the Windows platform and its search engine. Indeed, Microsoft has already announced that it intends to provide third-party developers with APIs to its new search engine, enabling them to construct applications based on it.

TRENDS IN SEARCH: TECHNOLOGY

The advantage conferred on Google by its PageRank algorithm, once overwhelming, is gradually disappearing. Many other clever algorithms have been developed; indexing and searching are being applied to more data sources and data types; and ever more nuanced user interfaces and functions are being offered to users.

Some of these efforts seem quite promising. Amazon has scanned more than 100,000 books and made their text searchable for Amazon users. Google Print provides a similar service and also offers direct links to bookselling sites. PubSub, a small startup in New York City, has developed a high-performance "matching engine" that monitors online information: if you subscribe to a topic, PubSub will scan data in real time and notify you whenever there is news. For the sorting and clustering of search results, the leader is Viv'simo, a Carnegie Mellon University spinoff in Pittsburgh, with its new Clusty website. Software from Blinkx, of San Francisco, lets users search multiple information sources, including their desktops, websites, and blogs. X1 Technologies of Pasadena, CA, also provides a popular desktop search tool.

As these examples suggest, many new search functions are being introduced by startups rather than by Google or established companies. A few of these startups may become large, independent firms. But most will remain small vendors, will be acquired, or will simply fail, depending on what Google, Yahoo, or Microsoft choose to do. Many offer products that would be natural additions or complements to existing search services, since their utility depends upon access to a search engine. But Google and Yahoo do not usually provide such access, even though it would benefit users. Google's sole Web API is laughably limited, offering little functionality and contractually restricting users to 1,000 queries per day.

Just what services could be built upon a fully open Google architecture? They could take many forms, but some of the most obvious would make indexing and searching processes on the desktop, on Web servers, and on Google's own website work together better. A single search could then span not just Google's index of the public Web but whatever other sources might be appropriate: a newspaper archive, a medical database, an antique-car parts catalogue, or your own hard drive. Google, or others building upon its APIs, would unify the results, explain any access restrictions on particular sources, and facilitate purchases of information. At the same time, independent firms could create services that call on Google's search and indexing functions to retrieve information, but present that information in new and creative ways.

As the search industry evolves, it also touches upon—and often competes with—a widening array of other industries, from publishing to software, in both business and consumer markets. The search industry wants to become the starting point for a larger proportion of digital activities. Some companies are happy to oblige: Amazon, for instance, opens its databases to search services, so that search results can point directly to

relevant Amazon products, bypassing the need to navigate Amazon's own site. Others are less welcoming. Microsoft will be displeased, to put it mildly, if Google Desktop begins to supplant the traditional Windows desktop interface and file systems.

However, the most important trend in the search industry is the proliferation of new computing platforms—and the increasing cross-pollination of data between these devices, PCs, and Web services. These emerging—and merging—markets represent Google and Microsoft's greatest opportunity for future growth and the greatest threat they pose to each other. In the absence of a common architecture, the information on these systems is almost unsearchable. Today, a user cannot possibly conduct a search such as "Show me everything about the Chinese economy that has appeared in the last month in my e-mail attachments, Word documents, bookmarked websites, corporate portal, voice mail, or Bloomberg subscription." Many computing platforms, old and new, have no useful search facilities at all. Most existing search tools are available on only one or at most a few platforms; and due to their lack of standardization, they cannot talk to each other.

Thus, while Google provides an excellent service for searching the public Web and has made a good start on PCs with Google Desktop (the hard-drive search tool) and Google Deskbar (which performs searches without launching a browser), the search universe as a whole remains a mess, full of unexplored territories and mutually exclusive zones that a common architecture would unify. Given the size and growth rate of the markets involved, the dominant search provider a decade from now could easily have revenues of $20 or $30 billion per year.

GOOGLE VERSUS MICROSOFT

Who will win? Google certainly has impressive assets. Moreover, Microsoft does not own the server side of the Web and probably never will. Nor does it control the architectures of the newer computing platforms, whose markets are growing much faster than the PC's. And in these newer markets, Microsoft faces a painful choice: either provide search technology that will run on, and thereby support, competing platforms such as Linux machines, or let others take the lead.

Yet Microsoft's control of Windows, Internet Explorer, and Office is a real advantage. For instance, if desktop search tools enjoyed deeper access to the internal document structures of Word and Excel, they would be much more useful. Similarly, operating systems can potentially collect information about user behavior that could improve search tools substantially. Other recent search innovations are really enhancements to the Web browser. Google, Ask Jeeves, A9, Blinkx, Yahoo, and Microsoft are all providing search toolbars that can be downloaded into the browser, and independent developers have created many search-related enhancements to the open-source Firefox browser.

But we know who really owns the browser. Ramez Naam, group program manager for MSN Search, declined to say whether or not search functions would be integrated directly into Microsoft's Internet Explorer. But a Microsoft executive, who asked to remain unnamed, told me that his company had recently reconstituted its browser development organization. "Microsoft effectively disbanded the Internet Explorer group after killing Netscape," he said. "But recently, they realized that Firefox was starting to gain share and that browser enhancements would be useful in the search market." He agreed that if Microsoft got "hard-core" about search (as Bill Gates has promised), then, yes, Google would be in for a very rough time.

Why? Because in contrast to Microsoft, Google doesn't yet control standards for any of the platforms on which this contest will be waged—not even for its own website. Although Google has released noncommercial APIs—which programmers may use for their

own purposes, but not in commercial products—until recently, it avoided the creation of commercial APIs. In late 2004, however, Google announced APIs for its advertising systems and for the Google Deskbar. The advertising APIs could help create an infrastructure of firms dependent on Google's platform and specializing in the management of automated, Web-based advertising strategies. This could protect Google's advertising revenues against future price competition from Microsoft. The Google Deskbar APIs, likewise, should encourage third parties to create search functions for the Windows desktop.

These steps, however, are at best half-measures. Google has not yet faced the need for full architectural competition and in some respects has arguably been moving in the wrong direction. It still has not provided open APIs for its core search engine. (Raúl Valdés-Pérez, Viv'simo's CEO, says that he tried to license Google's search engine services but was refused.) Furthermore, it sells its search software to enterprises only in the form of a bundled, Linux-based hardware system. This alienates other hardware and software vendors, leaves most of the non-Linux market unserved, and presents a huge opportunity for Microsoft.

Google may feel that APIs are of secondary importance in its coming war with Microsoft. Two Google employees (both of whom prefer not to be named) told me that Google's leaders believe that the company's expertise in infrastructure—knowing how to build and operate those 250,000 servers—constitutes a competitive advantage more important than APIs or standards. This could be a major, even fatal, error. Microsoft can certainly obtain or cultivate the skills necessary to operate large-scale computing infrastructures; indeed, it already operates MSN, with nearly 10 million users.

Worse, Google may feel that APIs can wait. Peter Norvig, the company's director of search quality, told Technology Review, "We've had the API project for a few years now. Historically, it's not been that important: it's had one person, sometimes none. But we do think that this will be one important way to create additional search functions. Our mission is to make information available, and to that end we will create a search ecology. We know we need to provide a way for third parties to work with us. You'll see us release APIs as they are needed."

Those words do not convey much sense of urgency. There is, however, another possibility: Google understands that an architecture war is coming, but it wants to delay the battle. One Google executive told me that the company is well aware of the possibility of an all-out platform war with Microsoft. According to this executive, Google would like to avoid such a conflict for as long as possible and is therefore hesitant to provide APIs that would open up its core search engine services, which might be interpreted as an opening salvo. The release of APIs for the Google Deskbar may awaken Microsoft's retaliatory instincts nonetheless. For Google to challenge Microsoft on the desktop before establishing a secure position on the Web or on enterprise servers could be unwise.

STRATEGIES AND PRESCRIPTIONS

In all of Microsoft's successful battles, it has used the same strategies. It undercuts its competitors in pricing, unifies previously separate markets, provides open but proprietary APIs, and bundles new functions into platforms it already dominates. Once it has acquired control over an industry standard, it invades neighboring markets.

In contrast, the losers in these contests have usually made one or more common mistakes. They fail to deliver architectures that cover the entire market, to provide products that work on multiple platforms from multiple companies, to release well-engineered products, or to create barriers against cloning. For example, IBM failed to retain proprietary control over its PC architecture and then, in belatedly attempting to recover

it, fatally broke with established industry standards. Apple and Sun restricted their operating systems to their own hardware, alienating other hardware vendors. Netscape declined to create proprietary APIs because it thought Microsoft would never catch up. Google—and Yahoo—would do well to take note.

What will Microsoft do? Publicly, it doesn't care about building a broad search architecture reaching across many platforms. "There will be a lot of innovation and competition around search by a broad number of vendors, but it is wishful thinking to believe it is a platform tidal wave like the initial emergence of the browser and the Web," says Charles Fitzgerald, Microsoft's general manager of platform strategy. And indeed, Microsoft has begun innocently enough: a decent though unspectacular search site, some software, no bundling—nothing, you know, violent. But the company will provide APIs to its Web search engine, and its long-term strategy could be brutal. If it acts logically, it will bundle better search facilities into Internet Explorer and Office; it will build advanced indexing and searching tools into both its PC and server operating systems; and it will alter its own products to make searches of many kinds more fruitful. Search tools could tailor results to a user's interests, based upon data collected by the operating system. Microsoft could even deliberately cause failures in Google's products—for example, altering its file formats so that Google's crawlers could not properly index Word or Excel files. Microsoft has been accused of such conduct repeatedly in the past, notably in its battles against the DR-DOS operating system (an attempted clone of MS-DOS) and Lotus spreadsheet software.

If it acts logically, Microsoft would also perform a "cashectomy" on Google—just as it did in the browser wars when it gave away Internet Explorer. Even with nearly $2 billion in cash, Google is vulnerable to this tactic. For instance, Microsoft could offer free wholesale access to its search engine. Then it could attack Google's advertising networks by offering free or subsidized advertising placement. These businesses are based primarily upon agreements with third-party websites, most of which have no long-term allegiance to Google. (Google's forthcoming advertising APIs could, however, change this.) Finally, Microsoft will try to play competitors off against each other, as is its custom. Microsoft thrives when its opponents are fragmented and possess no alternative common standard.

So what should Google do? Given Microsoft's ferocity in the past, panic might be a productive first step. Google should understand that it faces an architecture war and act accordingly. Its most urgent task must be to turn its website into a major platform, as some other firms have already done. Amazon, as we have noted, does not merely operate a retail website. It has developed proprietary but open APIs that have made it the capital of an electronic economy (see "Amazon: Giving Away the Store," p. 28). Other merchants set up stores under the Amazon umbrella, and other websites can offer direct links to Amazon's product pages. Recently, Amazon has gone even further, creating ways for consumers to search and find products without visiting Amazon at all.

Thus, Google should first create APIs for Web search services and make sure they become the industry standard. It should do everything it can to achieve that end—including, if necessary, merging with Yahoo. Second, it should spread those standards and APIs, through some combination of technology licensing, alliances, and software products, over all of the major server software platforms, in order to cover the dark Web and the enterprise market. Third, Google should develop services, software, and standards for search functions on platforms that Microsoft does not control, such as the new consumer devices. Fourth, it must use PC software like Google Desktop to its advantage: the program should be a beachhead on the desktop, integrated with Google's broader architecture, APIs, and services. And finally, Google shouldn't compete with Microsoft in browsers, except for developing toolbars based upon public APIs. Remember Netscape.

When Google's Peter Norvig was read this list—presented not as recommendations,

but as things that Google would do—he did not deny any of it. When Technology Review asked, "If we reported any of this, would we be wrong?", Norvig answered, "We don't like the word 'beachhead.' That implies a war, and we don't want to go there." Pressed, he said, "Factually, nothing wrong"—although he stressed that APIs were only one way Google might create a "search ecology." But historically, proprietary APIs have been the only way to create a loyal customer base—one that can't readily switch to a competitor.

BIG QUESTIONS

Would such an architectural strategy work? I'm not sure, but I think so. I also suspect that if Google doesn't do something like this fast, and Microsoft attacks, Google will go down. Its decline would take longer than Netscape's precipitous descent, but it would be no less final. And at least during the second term of the George W. Bush administration, it is highly unlikely that antitrust policy would come to the rescue.

Whether Google or Microsoft wins, the implications of a single firm's controlling an enormous, unified search industry are troubling. First, this firm would have access to an unparalleled quantity of personal information, which could represent a major erosion of privacy. Already, one can learn a surprising amount about people simply by "googling" them. A decade from now, search providers and users (not to mention those armed with subpoenas) will be able to gather far more personal information than even financial institutions and intelligence agencies can collect today. Second, the emergence of a dominant firm in the search market would aggravate the ongoing concentration of media ownership in a global oligopoly of firms such as Time Warner, Bertelsmann, and Rupert Murdoch's News Corporation.

If the firm dominating the search industry turned out to be Microsoft, the implications might be more disturbing still. The company that supplies a substantial fraction of the world's software would then become the same company that sorts and filters most of the world's news and information, including the news about software, antitrust policy, and intellectual property. Moreover, Microsoft could reach a stage at which its grip on the market remains strong, but its productivity falls prey to complacency and internal politics. Dominant firms sometimes do more damage through incompetence than through predation.

Indeed, as so many have noted, much of Microsoft's software is just plain bad. In contrast, Google's work is often beautiful. One of the best reasons to hope that Google survives is simply that quality improves more reliably when markets are competitive. If Google dominated the search industry, Microsoft would still be a disciplining presence; whereas if Microsoft dominated everything, there would be fewer checks upon its mediocrity.

Disclosure: *As the result of selling Vermeer Technologies to Microsoft in 1996, Charles Ferguson still holds a substantial quantity of Microsoft stock, a position which is partially but not completely hedged. He has no other financial interest relevant to this article.*