

Silicon Valley: Startups, Strategies, and the Stanford Connection

James F. Gibbons

The following is an edited transcription of the plenary address given by James F. Gibbons, Dean of the Stanford University School of Engineering, at the 1994 MRS Spring Meeting.

My talk today is based on a series of observations I started making about five years ago in response to a question someone asked me: "What role had the Stanford School of Engineering played in the development of Silicon Valley?" Although Stanford, in fact, probably gets more credit for Silicon Valley than it deserves, I thought the question merited some analysis, especially since I had for some time been following some of the Silicon Valley firms, and working with some of them. So I did a small, sort of left-handed study of these companies, involving talks with many of their founders, and out of that came this presentation—which, by the way, I first gave on Wall Street to a Stanford alumni club.

My subject is rather ambitious—a mini history of Silicon Valley—but my purpose is to give you some background on just how the legendary high-tech firms got started, what kinds of firms they are, how their individual economies developed, and what impact they had, as a whole, on the economy of Silicon Valley and the country. First, let's talk about startups—the types of startups, the length of time needed to determine whether or not they'll be viable, and the requirements for their success.

Three Types of Startups

Looking across the Silicon Valley at any given time, you will find, in general, three types of startup firms. The first type is a company that has, already in place, a workable product or idea that does not require extensive product or market development. A firm typical of this category is SUN Microsystems, whose product, the SUN Workstation, was the "solution" to a problem generated about a decade ago in a Stanford Computer Science Department graduate seminar. As a result of that seminar and related activities, Andy Bechtolsheim and his colleagues conceived the workstation, and

recognized that it could be put together with easily available parts. The group built a prototype to test the market for their product, found a significant potential, and proceeded to develop SUN 1, SUN 2, and SUN 3 in rapid succession. In terms of establishing a time scale, you can tell if a firm in this first category is going to be successful or not in a relatively short period of time—something like two to five years.

No successful firm I know of has ever been formed without a highly qualified team.

The second type of firm is one that has a working prototype of a product that requires some development before marketing can begin. An example of such a company is MIPS Computing Systems. MIPS, as some of you may know, is a RISC architecture semiconductor company formed in part by my Stanford colleague, John Hennessy and four of his graduate students. In their PhD work, Hennessy's group was trying to establish benchmarks for a very promising RISC architecture, to determine if it would perform as well as they thought it would. They eventually built the first chip for the architecture at Stanford's Center for Integrated Systems, then left to form the company. What they had at this point was not a working product, but a "proof of concept." They needed to redesign the

chip and develop something that could be manufactured at a reasonable price to serve the needs of a significant industry. The time scale for such development would be about five to ten years—to develop the chips, supply them to computer manufacturers, have them built into systems, and so on. So a "type two" firm would require roughly twice as long to establish its success as a "type one" firm—five to ten years, as opposed to two to five years.

The third type of startup is a company based on a fundamental research breakthrough, one which might, for example, lead to a Nobel Prize. Significant research would be required to develop not just the product, but all the surrounding technologies. One of the best, and no doubt familiar, examples of this type of breakthrough is the transistor. Invented in 1948, the transistor took about 20 years to get off the ground; it was 1968 before anyone was sure the transistor and the resulting integrated circuit would be as ubiquitous as they are today.

My first example of a "type three" startup is Conductus, a firm whose products are based on high-temperature superconductivity. Conductus is already six years old, and has products, but it hasn't at this point developed anything that's going to result in the creation of a new industry in the way the transistor did. At least not yet—you can never tell. My other "type three" example is Genentech, a biotech firm. Many other biotech firms also fall into this category. It takes a long time—10 to 20 years, speaking conservatively—for such companies to realize their potential. If and when that happens, a new industry results.

Four Keys to Success

Now that I've outlined the three types of Silicon Valley startups, I'd like to discuss four of the major requirements for launching a successful venture. I'll preface this by saying that these requirements are neither absolutely necessary nor sufficient; they're simply things I've learned by talking to some of the people who founded Silicon Valley firms.

First, you need an idea or product with significant market potential. For most of the firms I've studied, there are two sources for such an idea or product: research labs (university research labs or industrial research labs such as Xerox PARC, home of the Macintosh computer and laser printer) and market analysis (generating an idea through first-class market research).

The second requirement is a high-quality, highly dedicated team of people. No

Material Matters is a forum for expressing personal points of view on issues of interest to the materials community

successful firm I know of has ever been formed without such a team, which usually comes from one of two sources, university R&D or a corporate spin-off. I'm thinking not only of Xerox PARC spin-offs, but also of firms like Intel or National, both themselves spin-offs of Fairchild.

The third requirement for a successful startup is venture capital, which comes primarily from two sources. Over the past two or three decades, most startup funds have come from small firms that concentrate entirely on venture capital. During the past ten years, some large industrial companies have gotten into the venture capital business, but they've been rather less successful than the smaller firms. They're typically worried about how the outcome will serve their own interests and that usually has a dampening, rather than an encouraging, effect on the startup.

The fourth factor, appropriate infrastructure, is one you have to tease out of a lot of discussions and data, but it is nonetheless extremely important. By appropriate infrastructure I mean primarily cheap, appropriate space and access to basic technology. Between Stanford and the south end of San Jose, there's probably three million square feet of rentable space in which you could locate your business. You don't have to build anything—the structures are already there—and you have easy access to basic technology. If the central idea you're working on doesn't require you to write any software, you can "hire it done." Or you can have a chip made. Or you can have glass blown.

It's a rare place where you can find all the technologies you will need to establish and grow a high-tech company. In fact, I believe there are only two such places in this country: around Boston and in Silicon Valley.

Growth and Development

Next, to get a sense of how Silicon Valley firms have grown and developed, and what their market potential has been, I'd like to present an overview of about 50 firms (Figure 1). This takes the form of a graph, with the vertical axis showing 1988 revenues in millions of dollars on a log scale, and the horizontal axis showing the age of the firm in years. Each firm listed is just a point on the graph; an age and a sales volume, the annual sales volume for 1988. In the top right, you can see the point for Hewlett-Packard, almost 50 years old in 1988, with a revenue of just under \$10 billion.

Below, to the left, is the point for Varian. The broken line going through Varian plots the company's sales history over its 42 years in existence. In the lower left, near the origin of that plot, you can see a line with a slope. Along the top are the annual compound growth rates, ranging from 150% on the left to 25% at the far right. If you construct, with your eye, a sloping line going through the Varian data points, you'll see that when the firm was very young, maybe five or six years old, it had a slope which would have been equivalent to a 50% growth rate.

At the top left, you can see SUN and Tandem, still growing handsomely at more than 100%. Varian, which was 42

years old in 1988, is growing at about 12%, which was roughly the growth rate for its industry. Varian's 1988 sales volume, \$1 billion, was the proportion of the business that they owned. So, in principle, the growth rate for each firm will start fairly high; then, as the firm matures and as the field matures, the growth rate will tail off until it reaches the growth rate for its industry.

Moreover, the point at which a firm's dollar volume begins to peak will be close to its top potential market share. You can see how sales for SUN and Tandem are still growing quite well, while Telesensory Systems (TSI) sales are becoming almost flat. TSI owns the market for electronic aids for the handicapped—about \$10-15 million a year—and has captured almost all of it.

Now, let's go back to Hewlett-Packard. As you can see in Figure 2, H-P (solid line) followed the Varian curve (broken line underneath) for its first 30 years in business, then broke away. H-P broke away at 30 years and a sales volume of about \$300 million. It accomplished this by—anyone want to guess?—by getting into the computer business. Had it stayed in instruments only, it would probably be growing at about the same rate as Varian is. But the computer business gave H-P a dramatic boost, and when you think of the graph as a log scale, you realize that H-P is growing at about the same rate as SUN and Tandem are, in the same time period.

H-P realized that it had to get beyond merely making instruments, that it had to make intelligent instruments that could




Figure 1. Age vs. sales for 50 Silicon Valley firms.

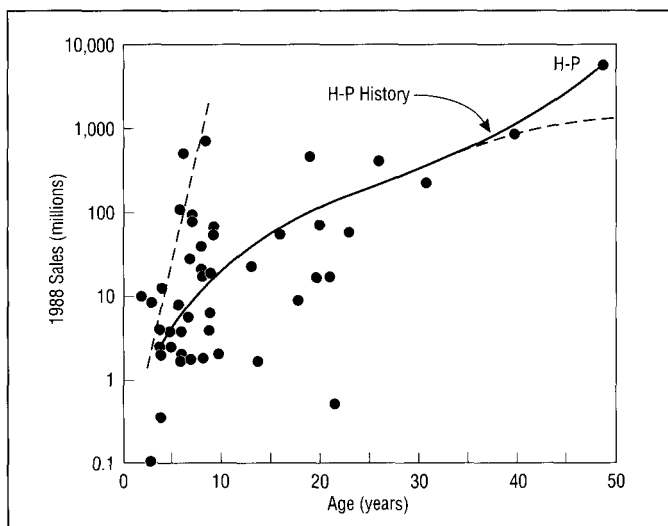


Figure 2. Age vs. sales for Hewlett-Packard.

be linked together in networks. The company had a lot of concepts in its original business that made it relatively easy to get into computing. We say H-P is in the computer business, but I believe its largest sales volume product is actually the printer, not the computer. So it's still selling a piece of peripheral equipment. But that is a very classy deal—being able to see concepts in your existing business that allow you to enter another line of business that provides opportunities for much greater growth and development. That's smart management.

Picking the Winners

Let's go back now to one of the requirements for a successful startup, a source of venture capital, and see what kind of performance venture capitalists are looking for when they evaluate a startup as a potential investment. Typically, venture capitalists go into a firm, talk to the team, evaluate the business plan, and analyze the market before deciding whether or not to put money into the company. But it's a very rough evaluation, based on the theory that some day the firm is going to go public, or be bought.

Looking at the history of firms that go public, you'll find that the appropriate corporate valuation that NASDAQ puts on a firm when it goes public is something like two times the annual revenues (Figure 3). In other words, they're assuming that you're on a growth curve and they're already giving you credit for the fact that your revenues are going to double in the near future.

Now a venture capitalist will invest a certain amount of money in a firm at $t = 0$, with the hope that in five or six years that money will be multiplied by a factor of 10. It would be great if that could happen in every case but, of course, it doesn't. Realistically, venture capitalists need to realize a return of four to 10 times their invested capital in roughly five years—that translates to something like a 30–60% compound growth rate.

Venture capitalists usually own about half of a firm when it goes public. But let's assume, for the sake of conjecture, that venture capitalists own all of a particular firm when it goes public, and that the public is going to buy their share. The value of that firm, which is two times annual revenues at the time it goes public, has to be equal to the venture capital return. So, to be a desirable investment, a firm should have annual revenues at least two times the invested capital.

Now let's take a look, in Figure 4, at the 1988 revenues (linear scale, vertical axis) of some startup firms, versus their total

Appropriate Corporate Valuation at I.P.O.
 $Value \cong 2 \times Annual\ Revenues$

Desired Payback
 $Return \cong (4\ to\ 10) \times invested\ capital\ in\ \sim 5\ years\ (30 - 60\% CGR)$

Neglecting Time, and Assuming 100% V.C. Ownership
 $Value = Return$
 $2 \times Annual\ Revenues \geq 4 \times I.C.$
 $Annual\ Revenues \geq 2 \times I.C.$

Figure 3. Calculation of venture capital and desired return

invested capital (not including what they raised in their initial public offerings). Since the broken line has a slope of two, or revenues equal to two times capital investment, anything below that line would not be considered a good investment by a venture capitalist. CEMAX, for example, is a terrible investment; someone put \$22 million into that firm and it never saw a cent of revenue. That's why the range of return on investment goes as high as 10—a venture capitalist needs some firms to yield ten times what was put into them.

Above the line, you'll see Collagen, Equatorial, Valid Logic, and others. The "super growth" firms like SUN and Tandem are so far above the line, they're actually off the chart. You'll notice that MIPS is slightly below the line. That was really a fluke, because I had studied MIPS at a particular moment in its accounting history and, three months later, it was up to \$60 million in revenues—above the line. There are other firms, of course, that are below the line, or off the chart completely, in which case they're simply not in existence any more. It's hard to get data for such firms, so the data for this chart is incomplete. But let me assure you that, of the 100 firms used in this sample, most of them lie above the line. Venture capitalists aren't likely to put money in something that lies very far below the line and, if they do, it isn't going to be \$22 million. Bubbles sometimes do occur, but they usually burst.

To be a desirable investment, a firm should have annual revenues at least two times the invested capital.

The Stanford Connection

Assuming that venture capitalists generally manage to pick the winners among high-tech startups, what type of economic impact have these firms actually had? (In asking this question, we've come full circle, to my original point of inquiry.) A look at some of the revenues produced by Silicon Valley during our sample year of 1988 yields some interesting data. The total 1988 Silicon Valley revenues according to Dataquest are \$40 billion. The revenues from 50 Silicon Valley companies that spun off from the Stanford School of Engineering in that same year are \$25 billion. And even though Hewlett-Packard is among those Stanford-based companies, responsible for fully \$9 billion of that \$25 billion, the fact remains that in 1988, 50 Stanford-based firms accounted for more than half the total revenues of Silicon Valley.

One of the reasons for that, of course, is the sheer number of Stanford engineering faculty and students working on various projects, patenting their findings, and taking them into the marketplace. But that's only part of the picture. You can see the connection more clearly by doing a quantitative analysis of the impact of the patents versus the impact of the corporate

In 1988, 50 Stanford-based firms accounted for more than half the total revenues of Silicon Valley.

startups. The Silicon Valley corporate revenues resulting from Stanford licenses in 1988 are \$2.5 billion. We know that figure because the Office of Technology and Licensing at Stanford (or at MIT or Berkeley) gets about one percent of the sales revenues generated by products based on its patents. In 1988 the Office of Technology and Licensing made \$25 million from its licenses, based evidently on \$2.5 billion in product sales.

It's interesting to note that the majority of that \$25 million came from only two patents. One patent, invented by a person in the music department, is the FM synthesizer that forms the basis for Yamaha keyboard instruments. The other patent is for gene splicing. The remaining patents were negligible in their ability to produce revenue for the university. So it is not patents, but the companies started by students, that have the greatest economic impact on Silicon Valley.

In a talk called "The Rise and Fall of

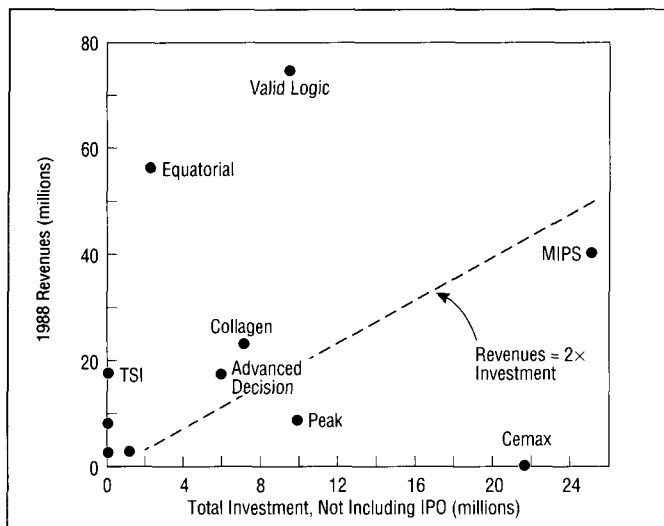


Figure 4. 1988 revenues vs. invested capital for some startup firms.

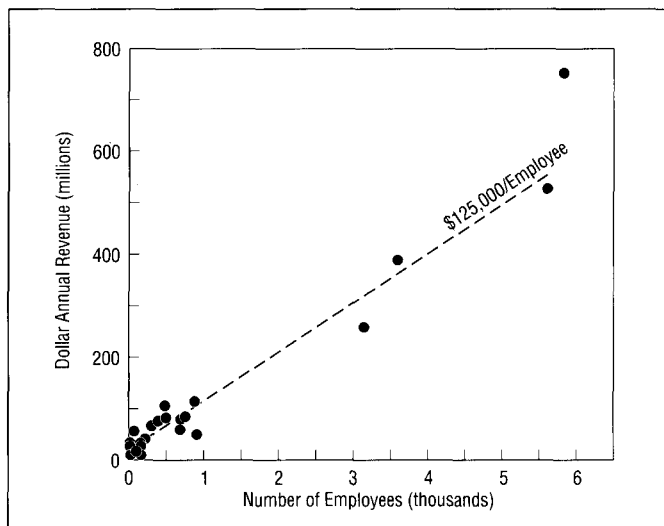


Figure 5. Revenue vs. number of employees, independent of age of company

Silicon Valley," Gordon Moore describes, in order of importance, what he considers to be the three main factors for the success of Silicon Valley. The first is Bill Shockley, not because he invented the transistor, but because he formed an unstable company that eventually spun off into Fairchild, which established a process. The second factor is the pleasure of working and living in what then was called Santa Clara Valley. The third factor is Stanford. According to Gordon, the most important thing Stanford does is to attract, every year, about 800 of the finest master's degree students in the world, educate them, then unleash them on Silicon Valley, thereby replenishing the intellectu-

The most important thing Stanford does is to attract, each year, about 800 of the finest master's degree students in the world, educate them, then unleash them on Silicon Valley.

al pool on a regular basis. I agree that this is exactly what we accomplish. It's also what MIT does, what every major technical school does: train top-level students for the high-tech marketplace.

Now I'd like to present some data for

MIT and its impact on the Massachusetts economy (similar to the Stanford/Silicon Valley data). The gross domestic product of the State of Massachusetts for 1988 was \$80 billion. MIT-based companies from all sources, not only from the School of Engineering, were responsible for \$40 billion of that. So MIT accounts for roughly half of Massachusetts revenues, while Stanford-based companies generated revenues that represent a little more than half of Silicon Valley's. That's real economic clout!

Economic Impact: Jobs and Growth

Continuing our study of the economic impact of Silicon Valley, let's look at how well the Valley does job creation. We'll focus again on 1988, keeping in mind that back then Silicon Valley was involved not so much in software as in manufacturing. On this graph (Figure 5), each point represents a Silicon Valley firm, the vertical axis showing annual revenues in millions of dollars, and the horizontal axis showing the number of employees. The firms seem to be widely scattered, but actually they fall on a fairly straight line, whose slope is, of course, the revenue per employee. As you can see, in 1988 that revenue was about \$125,000 per employee. An interesting point. SUN, Tandem, and Varian all had revenues of about \$1 billion and, according to our chart, also had nearly the same number of employees. Now SUN and Tandem got to the \$125,000 per employee level in about six or seven years, but it took Varian 42

years, so the average job creation rate for Varian is approximately one sixth the rate for SUN or Tandem. My point is that a startup that moves aggressively and gets its product to market quickly is creating jobs like crazy. It's really true that startups, if they grow to become significant companies, are going to generate jobs at a very high rate, especially during their first decade.

The next chart (Figure 6) will illustrate another interesting point about the job-creation rates of Silicon Valley firms. For three firms—SUN Microsystems, Hewlett-Packard, and Varian—I've listed the 1988 sales, then the 1989 sales, adjusted for inflation and expressed in 1988 dollars. In the 1988-89 year, SUN went from \$1.05 to \$1.68 billion, a growth of \$630 million. If you divide that growth by the \$125,000-per-employee figure from our previous chart, you could assume that the firm added 5,000 employees. Actually it added 3,000 employees. Keeping that in mind, let's look at Hewlett-Packard. H-P grew from 9.8 to 11.33 billion during that same period. At our average productivity rate, it should have added 11,000 employees. In fact, it added 8,000. In other words, the added revenue per employee for SUN and H-P was more than the average, a sign that their businesses were growing. That's good for them and for the economy.

Varian, on the other hand, went from \$1.17 billion to an adjusted \$1.14 billion, down by \$30 million, or 3%. At the average productivity rate, it would have had

to release 240 employees. But it actually lost 1,000 employees to arrive at the \$30 million, a rather depressing outcome. The fact is, when your business is growing, you are probably adding people who are more productive than the average, but when business is decreasing, you're releasing people who are less productive than the average. You may take that scenario in stride from a management point of view, while, say, President Clinton, looking at it from a jobs point of view, may be very unhappy. A lot of this type of "right sizing" is involved in jobs-creation data across the country, and it's troubling if you're trying to create jobs.

I want to make two points here. First, highly successful new firms create jobs at a very high rate. Second, established firms need some kind of breakaway strategy if they're going to outperform their sector economy, in which case they will also generate jobs at a significantly high rate.

Strategic Alliances: A Great Way to Grow

Now I'd like to talk about one of the ways high-tech firms can grow their businesses—by forming strategic alliances. By way of definition, a strategic alliance is a business relationship based on technological synergy between an established company and a high-technology firm. A strategic alliance can function in two ways: as an alternative to internal devel-

	'88 Sales	'89 Sales	No. of New Employees (est.)
SUN Microsystems	\$1.05 bil	\$ 1.68 bil	3,000 (5,000)
H-P	\$9.8 bil	\$11.33 bil	8,000 (11,000)
Varian	\$1.17 bil	\$ 1.14 bil	-1,000 (-240)

- Highly successful new firms create jobs at high rate.
- Established firms need "break-away" strategy to outperform their sector economy.

Figure 6. Job creation rates for high-tech "startups."

opment and as an alternative to traditional acquisitions. For example, when Hewlett-Packard got into the computer business, it accomplished the entire task internally; when H-P wanted to enter the workstation business, or increase its stature in this field, it acquired Apollo. So H-P has done both, internal development and acquisition.

But the strategic alliance is a completely different way to grow a business, one which, if done correctly, should be a lot more efficient than the two more traditional ways. Let's look at the competitive advantages a strategic alliance holds for both of its partners, the large company and the small high-tech firm. The large company (Figure 7) can use a strategic alliance, first of all, to gain rapid access to market (the small firm already has the

product). Next, the alliance gives the large firm a low-risk access to new technology; an acquisition tends to kill the young firm, an alliance is not supposed to. The strategic alliance also increases productivity, technological diversification, product and market diversification, and access to entrepreneurial talent.

Figure 8 shows the competitive advantages from the small high-tech's point of view.

Obvious advantages are access to additional markets and access to new marketing channels, both highly important. Next, an alliance provides the small firm with capital inflows, crucial if it's hanging by the fingernails most of the time. Then, long-term business focus, risk reduction, and enhanced image and stability are what the small firm hopes to achieve from a strategic alliance—when it works.

Let's go back to the larger firms. Thinking about how Hewlett-Packard got into the computer business (internal development plus acquisition), I've outlined in Figure 9 some of the pros and cons of both the strategic alliance and internal development (acquisition has been left out because few companies these days are thinking about acquiring a business; if anything, they're thinking about disassembling one). The pros for an alliance are short time to market, low risk, and diversification. On the negative side is less control (the large firm can't control the small one), external expertise (the

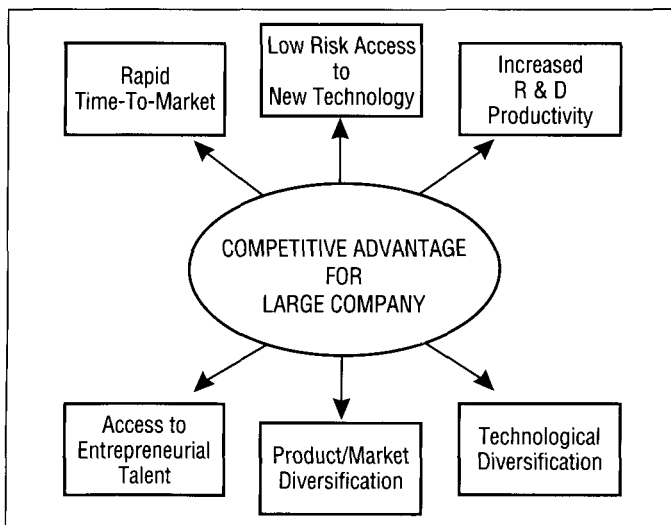


Figure 7. Competitive advantage of a strategic alliance for a large company.

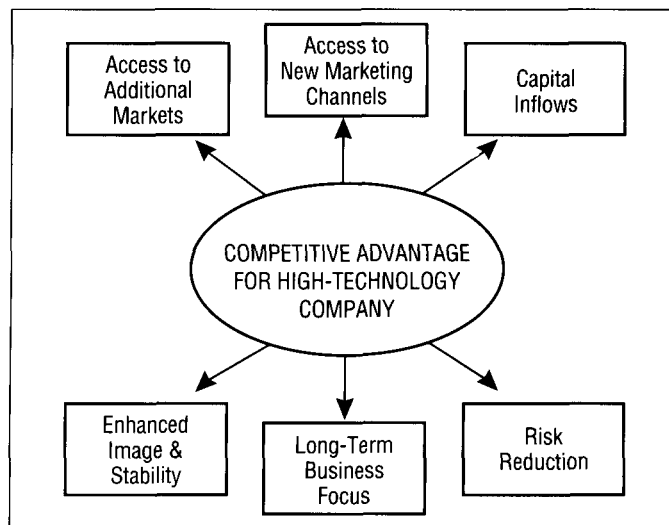


Figure 8. Competitive advantage of a strategic alliance for a high-technology company.

ALLIANCE		INTERNAL DEVELOPMENT	
PRO	CON	PRO	CON
• Short time-to-market	• Less control	• Complete control	• Long time-to-market
• Diversification	• Expertise is external	• Expertise is internal	• Bureaucratic inertia
• Low risk	• Financial exposure	• 100% ownership	• No assurance of R&D success
• New business enhances innovation			• Entrenched business deters innovation
• Low cost			• High cost

Figure 9. Pros and cons of the strategic alliance vs. internal development.

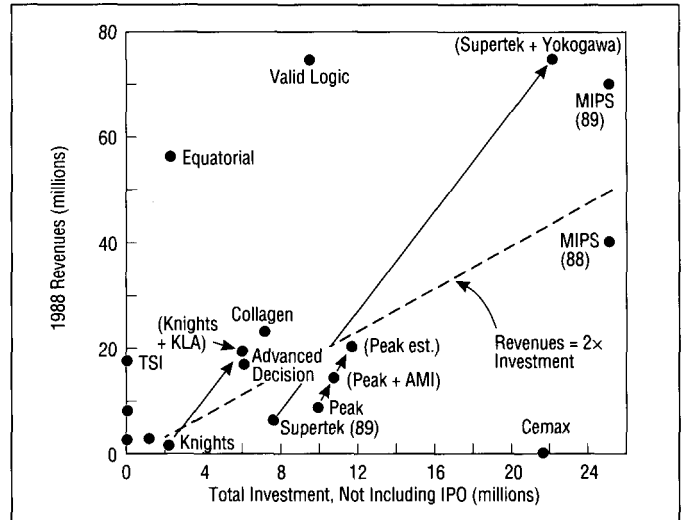


Figure 10. Revenues vs. investment when a strategic alliance is successful.

expertise is in the small firm, not the big one), and financial exposure (the small firm could still go belly up because the large firm is not buying it, only supporting it).

For internal development, the pros are complete control, internal expertise, and 100% ownership. If the venture is successful, the rewards stay completely in-house. On the down side are a long time to market, bureaucratic inertia, and high costs (the large firm has to pay its own overhead).

Strategic Alliances: Economic Efficiency

Figure 10 shows what can happen economically when a strategic alliance is successful. This graph is similar to the one on venture capital, with total investment on the horizontal axis, 1988 revenues in millions on the vertical axis. Take a look at Knights, in the lower left corner. As some of you may know, Knights is a small firm that made equipment for measuring very thin films and oxide films, among other things. Knights formed a strategic alliance with KLA, with about \$2 million in invested capital and no revenues. KLA put in what looks like \$4 million. It wasn't actually \$4 million, but the equivalent of that: KLA took over manufacturing and distribution, advertised the products, etc. So it put in \$4 million worth of value. Now, move up the chart to the right, to the point marked Knights + KLA. Just two years later, Knights was making \$20 million in revenue. From a \$4 million investment, it

reaped \$20 million in revenue, which is a lot larger than a slope of two. When a large firm forms the right kind of alliance with a small firm, it can significantly increase its leverage in the marketplace. Had Knights acquired \$4 million from venture capitalists and tried to accomplish the growth itself, it more than likely would have ended up on a slope similar to the broken line. A firm starting below that line, as Knights was, never gets above it without becoming very lucky. A strategic alliance is a good way to get lucky.

Another interesting example is Supertek, which was making mini Crays, useful for weather prediction and other purposes. Supertek tried to identify a U.S. firm for a strategic alliance and couldn't, but finally found one in Japan. It formed an alliance with Yokogawa, a firm that made radios and other products but wanted to get into the field of high-speed computing and mini mainframes. Supertek had about \$8 million in U.S. invested capital and about \$10 million in sales at the time the alliance was formed. Yokogawa put in an equivalent amount, roughly \$12 million, and the sales went promptly to \$75 million. So again you can see that the slope of their line is dramatically higher than one would have expected had Supertek merely put in another \$12 million in venture capital.

When strategic alliances work, that's how they're supposed to work. Sometimes an agreement is made, then turns out to be not as good, not as futuristic as it should have been. An example of a less

than successful alliance is one formed by Peak and AMI (Figure 10, bottom center), which was supposed to increase their revenues to \$20 million. But sometime after the agreement was signed, AMI and Peak disagreed on the terms and just about everything else, and the partners ultimately failed to honor the terms of the document. You shouldn't conclude, from what I'm saying, that spending a lot of time with lawyers would help such a situation. If you want a strategic alliance to work, you've got to have a common understanding and be willing to let the small firm develop its market.

There's never been an active alliance market in the United States to match the one in Japan.

So that's the theory behind strategic alliances. We don't have many such alliances in the United States and in most of the successful ones I've known, the partner has been Japanese. That was true, at least, from 1985 to 1990, after which few strategic alliances were formed here. There's never been an active alliance market in the United States to match the one in Japan; we have no major drive to take advantage of the technology developed by the young high techs of Silicon Valley or Route 128.

Conclusion

While I don't propose to be an economist, I do believe that successful startups are based on a marketable idea or product, a high-quality team, venture capital, and appropriate infrastructure. I hope I've demonstrated to your satisfaction that the economic impact of successful startups is, and can be, very important to the economy as a whole, both in terms of the revenues they generate and the jobs they create.

I haven't tried to demonstrate that successful growth depends on high-quality management and a consistent technical orientation, but I think those are crucial factors. Finally, I believe strategic alliances can be a significant part of the high-tech picture by supplying missing factors such as manufacturing processes and marketing channels. Strategic alliances may not always succeed, but when they do, the results can be magnificent.

Questions and Answers

Q: Are the four principles you outline for a successful startup exclusive to the United States, or are they relevant all over the world as well? Or do other places have different criteria for starting a high-tech business?

A: We have had senior ministers from more than 10 different countries visit Stanford to find out how Silicon Valley works and what Stanford had to do with it. Several years ago French President Mitterand and his staff paid us a visit. Now the French minister of science didn't believe that infrastructure was important. But I would be very surprised if anything like Silicon Valley could develop in France without appropriate infrastructure.

Bill Hewlett listed the four principles before I did, and most of the people you talk to from Silicon Valley would, I think, tell you that all four prerequisites, particularly infrastructure, are very important. Infrastructure, however, is very hard to reproduce. It would be difficult to find in Austin, Texas the kind of logistical and technology advantages that exist in Silicon Valley or around Route 128. The same firms may be in Austin, but they didn't start there.

If I had to prioritize the four principles, I would say the single most important one is a high-quality team. Venture capitalists bet as much, or more, on the team than on anything else because if the first

business plan doesn't work, the second one probably will. The idea ought to look right—they don't want you starting off on something you don't believe in—but many firms have ended up doing something different from what they had originally planned. So I would rank them in this way: first, the team; second, money; and third, infrastructure.

Q: In studying the various firms you've mentioned, have you noticed whether or not government R&D funding at the beginning, or even a little beyond that, had any positive or negative effect on the eventual commercialization of any products? At the APS meeting in March, John Rowell gave a very depressing talk in which he told us how many jobs have been created by the high T_c superconductors. After eliminating factors which he thought weren't legitimate, like jobs created by government and R&D contracts, the result was something like zero. How many of these success stories were launched because of, or in spite of, that kind of funding?

A: If you mean direct funding, then some firms certainly were not. Hewlett-Packard was not, Varian was not, probably most of those companies were not started from research supported by the Department of Defense (DOD). On the other hand, you've got to consider that Stanford's Computer Science Department is heavily involved in such research. Also, most of the money that comes into the School of Engineering comes from the faculty, from the work they do to fulfill their research contracts. Engineering's budget is about \$125 million, about \$80 million of which is raised by the faculty from the DOD. Probably 70 percent of the research is defense-based, and 95 percent of it is federal government research of one sort or another.

So it's a little hard to answer that question. MIPS, for example, was created from a government contract. In fact, MIPS and Silicon Graphics were started from the same government contract. Jim Clarke and John Hennessy were its respective managers. So sometimes that happens, but it's not the common practice.

Q: As a quick follow up, did you include any government R&D revenue in the revenues you were plotting?

A: No. The revenues I listed were solely what the companies reported on their

10Ks. That's it. The answer, I think, is that defense-supported research had very little to do with getting these companies to where they are. And it probably had little to do with their startups, although the defense-related environment that spawned them may have had some influence. So the government is not without impact.

Q: In the area of computers, which seems to be one of the hot fields now, how much of the business is devoted to hardware, how much to software? What balance would you predict for the future?

A: That's a good question. When you look at how Hewlett-Packard is growing, you see that it's growing because of hardware; but that hardware has software embedded in it, and is acquiring an increasing amount of software as time goes on. In most of the firms that are growing—in which the average revenue per employee is higher than average—software is usually the reason for that growth, because the average revenue per employee in the software industry is two to three times higher than that for the manufacturing industry. When you see a firm growing, more than likely it's because that firm is adding software components and technology. So the software side of the business is creating most of the growth.

Q: Relating to infrastructure, there's a lot of talk these days about the possible impact of enhanced communication—the information highway, that type of thing. The implication is that, with such advanced communication, you could have a more diffuse infrastructure and still achieve the same results. Do you think that will be the case?

A: I don't know. I suppose we're going to find out. But I don't think the situation you describe will spin itself out too rapidly. And I would be very surprised if working with an infrastructure that's spread out over a wide area will be as effective as working with one that's concentrated in a single place. It has something to do with rubbing elbows in a lab. That's hard to reproduce when you have people scattered around the globe—whatever the electronic connection among them might be. Perhaps I'm being a little defensive—I don't know what is going to happen. □