MOVING IBM'S TECHNOLOGY FROM RESEARCH TO DEVELOPMENT

Joint programs foster technology transfer by having R&D people plan and work together on next generation technology.

Ralph E. Gomory

Often the greatest challenge facing an industrial research organization is consistently and rapidly moving the results of its work into the product line of the company of which it is a part. Certainly doing good science is not enough today, if indeed it ever was. To survive and prosper, an industrial research group must be a vital part of the company it serves. This means, first and foremost, getting its work accepted elsewhere in the business.

The earliest managers of IBM Research recognized that it was not enough to build a corporate research laboratory on a hill and expect that, somehow, the benefits of its work would trickle down and be exploited by the rest of the company. Even so, and despite a number of major scientific and technical successes since the 1950s, it took two decades of trial and error to evolve a consistently effective way of coupling the work of the Research Division laboratories to IBM's product development community.

A significant mechanism IBM uses, called joint programs, fosters the transfer of technology from research to development by forming, in some areas, continuing joint efforts in which teams of research and development people agree on plans and work together to achieve them. How these joint programs came into being, how they work, and their advantages is the subject of this article. But first, a few words about the fundamental difficulties of transferring technology into development are in order.

Ralph Gomory is president of the Alfred P. Sloan Foundation in New York City, a position he assumed after retiring from the IBM Corporation earlier this year. After receiving his Ph.D. in mathematics from Princeton University in 1954, Gomory served in the U.S. Navy and was Higgins Lecturer at Princeton before joining IBM in 1959 as a research mathematician. In 1964, in recognition for his widespread accomplishments in applied mathematics, he was made an IBM Fellow, and in 1970 he became IBM director of research. He was elected an IBM vice president in 1973 and since 1986 has been IBM senior vice president for science and technology. Gomory is a member of the National Academy of Sciences and the National Academy of Engineering and has served on both their councils. He was a director of the IRI from 1984–88 and has been a member of the White House Science Council and chairman of the Advisory Committee to the President on High Temperature Superconductivity, in addition to serving on numerous other professional and academic boards, committees and panels. The author of more than 50 technical papers, he has received several major awards including the IEEE's Engineering Leadership Award and the National Medal of Science.

Fundamentals of Technology Transfer

One important point is that most "new" products are in fact improved versions of products that were available last year. They are based, not on a brand new idea from science, but on improving an existing product. And the process of repeated incremental improvement (also known as the development and manufacturing cycle) that produces these new versions of the product is inherently resistant to ideas from outside itself.

Many difficulties relate to this existing product orientation. There is an existing set of tools that it is easier, quicker and less costly to use than to get new ones. There are only certain techniques or metallurgies or programming languages that the development team is familiar with, and so on. A new idea or a really new design or part must fit reasonably well into this pre-existing world.

It is only those who work with the product who know these circumstances: what tools they have, what stage the next version is already in—in short, what ideas for improvement are acceptable at a given point in time. It is not impossible for an in-house research (as opposed to development) organization to understand these things, too, but unless it makes the effort to do so, it is severely and almost hopelessly handicapped when it comes to contributing to the normal product improvement process.

Another handicap a corporate research group faces is that it is, in some sense, in competition with its own potential customers. The ideas, new designs and new processes that it proposes can easily compete with, rather than supplement, the product improvements already planned. And if the question is, whose design will development exploit in their next round of products—yours or theirs, other things being equal, the answer is obvious. This set of problems usually goes under the name of "NIH."

The larger the enterprise, the more severe these problems become. For as companies grow, they tend to break up into smaller, more manageable pieces. Research separates from development; technology development may split off from product development, which itself may be divided up between a number of operating units. These organizational boundaries can easily become barriers across which it is very hard to move new technology and ideas.

Part of Research's Responsibility

It is clear from these characteristics of the typical research/development relationship that technology transfer takes work and that someone has to do that work. And it was also clear to us that Research would have to do that work. We could see that sitting back and waiting for others to pick up our results wasn't good enough, and neither was meeting our development counterparts half way. Somehow, by hook or by crook, we had to find a pathway for our efforts that would get the results into the company's product plan.

During the 1970s, this commitment to doing the technology transfer work became one of the guiding principles by which IBM Research operates. Research explicitly took the responsibility for technology transfer, even though it did not have the authority to make technology transfer happen.

In this new spirit of carrying our work forward, we began more often to contest disagreements with the product divisions at the level of IBM's Corporate Management Committee. Furthermore, I made the decision that no technology transfer effort in Research could be discontinued without approval at the level of someone reporting to the director of research.

Up until that time, people in IBM Research had been too willing to say, "They won't take it," and simply put their work on the shelf. Requiring the approval of an area director meant that if division management thought the product division was wrong, they could escalate the decision. And if, faced with the possibility of a realistic debate, they decided the product division was right, they could try to adapt or change the project enough to make it acceptable.

One of the things we learned from this process was that we were often wrong. Our project often did not fit well into the real world of product development. Faced with this reality, we learned to do better.

Knowing the Other Side

Another component of the efforts we made to take a more active role in technology transfer was getting to know the other side. Most research people, we realized, had no clear notion of what development people did, the constraints they faced, or how to relate to them.

To help fill this void, we tried a number of things. We invited in the heads of other IBM divisions to speak at a series of well-attended meetings. We launched a sabbatical program in which researchers would spend time at the company's product laboratories and, in the process, get to know how development people think and work. And we organized a management development program, taught by our senior people, which included a section on technology transfer, how to do it, and how it was a responsibility of research managers to make it happen.

Early Attempts at Working Together

Meanwhile, periodic attempts were also made, by both the IBM research and development communities, to

Most research people had no clear notion of what development people did, the constraints they faced, or how to relate to them.

collaborate more closely on specific projects. In the early 1960s, joint work between IBM Research and our components development laboratories had helped in the transfer of our early work on field effect transistors. Later on, the same was true of the IBM-invented thin-film head for magnetic disk files. Because storage systems development had worked with us on that project, they were looking forward to its benefits.

Over the years, both sides made other efforts to establish gateways of collaboration by appointing liaison people, and to formalize somewhat the relationships between the product divisions and Research. But they all shared the same disadvantage: Being mostly *ad hoc* arrangements stimulated by a requirement of the day, they lacked staying power. The problem would be solved, priorities for advanced technology would shift, funding would evaporate, and business as usual would resume. (As we would phrase it today with the benefit of hindsight, we were dealing in individual projects but not in ongoing programs.)

Business as usual, in those days, often meant that what Research did didn't matter greatly to the rest of the company. Research might be a strong advocate of silicon gate technology, for example, but if the components division was practicing metal gate technology, there was little hope of changing their direction. Likewise, whatever Research was working on typically was 10 years ahead of what was in development and manufacturing. That was another perpetual void between us.

One exception, which turned out to be a precursor of things to come, was a joint venture called TRI-One, agreed to in 1978 by the Watson Research Center in Yorktown Heights, New York, and the components laboratories in East Fishkill, New York, and Burlington, Vermont. "TRI" stood for the three locations involved, and "One" for the 1-micron technology on which Research was working at the time. The idea was that by working together, using Research's silicon gate pilot line, we could leapfrog from the 3-micron metal gate technology then in production directly to a produceable 256K DRAM chip.

TRI-One achieved a measure of success. It was instrumental in IBM's being the only contractor in the Defense Department's VHSIC (Very-High-Speed Integrated Circuits) program to meet all its original commitments, on time and at cost. It undoubtedly played a role in IBM's subsequent introduction of silicon gate memories and in being the first producer of 256K, and later megabit, chips. But it pointed up our lack, at the time, of facilities suitable for this kind of advanced technology work, and again there was no continuity; once its job was done, the TRI-One organization ceased to exist.

Two Trips to Japan Set the Stage

If I had to define the sequence of events that led directly to IBM Research's formal joint programs as they exist today, it would probably begin with my second cross-licensing trip to Japan in November 1979. IBM was then engaged in cross-licensing negotiations with a number of Japanese computer manufacturers, and as I had done in 1974, I led a technical team which spent a day at each of the main laboratories of these companies to evaluate what they were doing as input to our upcoming negotiations. We heard a lot about how the Japanese computer companies were organized and what they were doing.

These companies had been very successful at working together in the MITI (Ministry of International Trade and Technology)-coordinated VLSI program to enhance their chip-making capabilities. But more important, we could see that their internal research was more focused on products than ours, their research people had closer ties to their development and manufacturing people, and surprisingly, even in the smaller companies they had outstanding facilities (often government subsidized) to work with.

In addition to the generally closer ties within the companies, there were some special mechanisms they used—and one that Hitachi called *Tokken,* in particular.

Tokken is a mechanism employed when there is something important to be done (such as introducing a new technology or achieving a major cost reduction in an existing technology or product) that requires no major new invention but must be accomplished quickly. Then senior management assembles a team from whatever parts of the company have the necessary skills, and separates it from the normal management structure. *Tokken* is ongoing; a typical company might always have 3–5 percent of its R&D resource working together on these kinds of projects.

Different members of our team were struck by different things, and every day we compared notes before writing up the day's experiences in a comprehensive report. However, while some of us reacted to the remarkable facilities available for advanced work, some to the lack of internal barriers, and some to the special and effective team structures, we were united on one thing: These competitors were moving more swiftly than we were in transferring new technology into products.

This was especially evident to me and to those members of the team who had made the same trip five years previously. We could see the astonishing progress that had been made. Groups that were far behind us in 1974 were now breathing down our necks.

It was clear to all of us that if we did not improve the way we did advanced technology we would be left behind.

It was clear to all of us that if we did not improve the way we did advanced technology we would be left behind. On the long plane ride back to the United States, we discussed the situation. Although a solution was not obvious to us then, we knew we had to do better, and all of us were determined to find a way.

Actually one member of our team, James McGroddy, was already hatching a plan and sketching out a proposal. It called for a permanent team of researchers and engineers who would work together on the next generation of silicon technology, and for access to facilities and tools as advanced as those available to product developers.

Shortly after our trip, and based on Jim's proposal, I formed a steering committee of key technical executives from both research and development to flesh out the goals and structure of this novel plan.

Proposing the First Joint Program

Part of the problem was that at any given time, research was doing far-out exploratory work; manufacturing was producing the current version of the technology; development was readying the next version for production; no one was really focusing on what the next step after development's version would be.

To fill the gap that existed between research and development, we agreed to propose an interdivisional joint program called the Advanced Silicon Technology Laboratory. The distinction between a program and a project is important here. As a program, the ASTL would be an umbrella organization, through which multiple projects would flow. Its mission, however, would remain the same: the first reduction to practice and definition of the next generation of silicon technology after the one currently being made ready for production.

The ASTL, we agreed, should consist of three groups—one at IBM Research and one at each of the two component laboratories (Fishkill and Burlington)—with a single management and a single technical plan. To shield it from day-to-day crises and shifting priorities would require a long-term (at least three-year) commitment from the corporation. Management out of Research (which tends to be a more protected environment then most in terms of stability) would also help.

And while each division had some existing resources that could be redirected to the ASTL, they would not be enough. In particular, the facilities then available were inadequate and, as we were keenly aware, not comparable with some we had seen abroad. Additional headcount was also needed.

In return for the additional headcount and capital investment the ASTL would require, we were convinced that we could commit to moving up key dates for the availability of new technologies by at least a year. Not only would the definition phase of new technologies be shortened, but the definitions would be based on better information—actual hardware, instead of models and estimates. As a result, there should be fewer slipped schedules, less need for technology redefinition, and a better fit to product needs.

In retrospect all this seems plausible enough, but in real life at the time this was just one more project asking for more people and money to be spent on relatively far-out goals. What were the forces that enabled this program to gain divisional support, and ultimately—with both research and development behind it, and development committed to providing the additional resources—be blessed by the top management of IBM?

The Credibility Factor

It required a considerable leap of faith on the part of the product divisions to invest in the ASTL, instead of spending the money on something they wanted to do themselves. Indeed, I do not think it would have been possible to launch the first joint program if IBM Research did not already have considerable credibility.

Research's scientific and technological achievements, and especially the efforts it had made over the years to reach out to the IBM development community contributed to this credibility. So did joint program precursors like TRI-One. People in development and even in manufacturing were beginning to be convinced that scientists from Research could actually help them solve some of their most pressing down-to-earth problems.

A striking example of this was occurring just as the ASTL proposal was being born. Major problems arose in the early production of a new high-end packaging technology. Without this package, IBM's new high-end computers could not be shipped, and while the problems remained unsolved, the announcement date for those systems was slipping day by day. It was a company-wide crisis and, like everyone else, Research was asked to see if it could help.

Research, development, and manufacturing teams were formed, and during the weeks that followed, Research Division scientists worked side by side with the development and manufacturing engineers at our plants in East Fishkill and Endicott, New York. The scientists not only helped in an absolutely crucial way but found enormous satisfaction in making this practical and immediate contribution. The engineers discovered that the scientists were practical people like themselves, and that their knowledge and skills could make a vital difference. And both groups got more comfortable working together.

Introduction of new technologies has been speeded up and product cycles shortened.

The Human Factor

Another important fact that enabled agreement on the structure, management and objectives of ASTL was that the actual working groups who participated in it stood to gain in their own terms as well as in terms of the company's objectives. The objectives of this unified program to advance technology were sufficiently compelling to the company that there were going to be better facilities and more people for the groups that participated. This motivation helped overcome the interdivisional barriers and organizational friction.

Getting the Program Underway

This was the background that enabled the steering committee to agree on the funding and resources for ASTL. And with both development and research supporting it, the plan was presented to the Corporate Management Committee in December 1980. IBM Chairman Frank Cary approved it, and soon afterward Data Processing Product Group General Manager Arthur Anderson made available the headcount and the capital needed to construct new pilot lines at Yorktown and East Fishkill.

The next key decision was who would manage the new ASTL. We believed that, although the overall joint program had to be managed from the Research Division for the sake of long-term stability, the director should not be from Research. He or she should be someone the product development people knew, respected and thought of as one of their own.

We found just such a person in the late Robert Henle, an IBM Fellow who had been manager of advanced technology at the East Fishkill components laboratory. Henle directed the ASTL from its inception in 1981 until his death early in 1988. His leadership contributed enormously to its success and thus to the success of the other joint programs that followed.

As agreed by the steering committee, Research committed 15 people whom we felt were applicable to the new organization; 30 came from the Data Systems Division which included the East Fishkill laboratory, and 40 from the General Technology Division which operated the Burlington laboratory. The Data Processing Product Group, to which DSD and GTD belonged, funded the additional headcount ASTL would need to carry out its mission. That was 55 during the first year of operations, increasing to 110 by the third year.

Now, eight years later, the ASTL continues to be supported in the same manner. Research, for example, still provides the original 15 headcount they committed to the program for 1981. IBM United States (today's extended version of the old Data Processing Product Group) and its divisions pay for the rest, and IBM U.S. even funds most of the ASTL's silicon processing facilities in Research.

Shorter Product Cycles

That the IBM product organizations have been willing to continue this kind of investment in the Research-directed ASTL and later joint programs speaks well for the benefits they have received.

The introduction of new technologies has been speeded up and product cycles shortened because development teams, once they have finished their work on a current technology or product, now have a running start on the next generation.

We rarely transfer technology anymore in IBM Research. Instead, we develop it jointly with our colleagues in the product divisions.

In addition, the joint programs allow the development

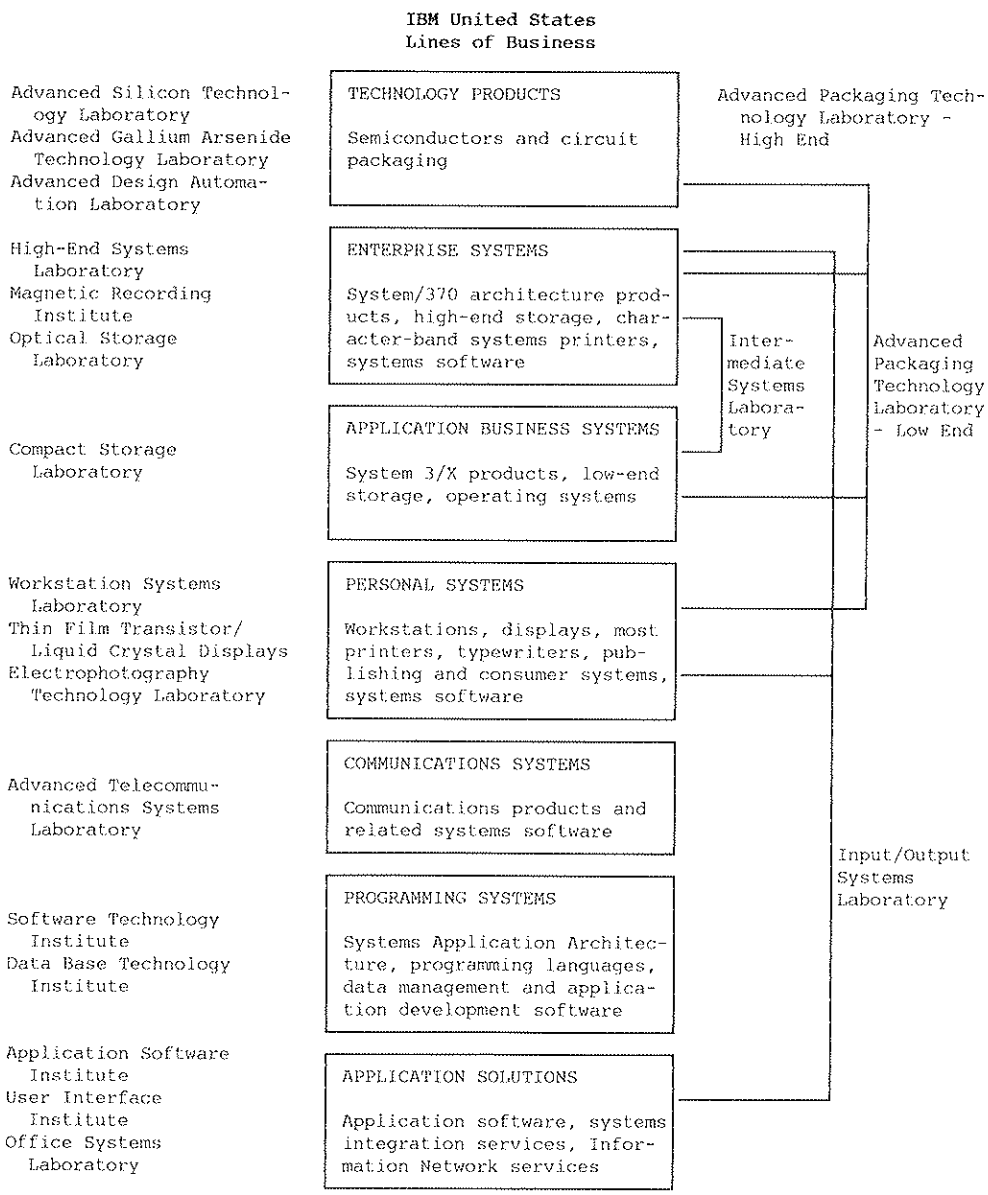

Joint programs between IBM Research and IBM Lines of Business have increased from one in 1981 to 19 today.

organizations to "pull" on what is going on in Research. Now the product divisions help set the goals Research aims for, rather than being presented with something Research has created and is trying to sell them.

Also, because part of the joint programs' job is the reduction of advanced technologies to practice, elements of their next generation work tend to feed forward and enhance committed technologies already in development.

Today's 19 Joint Programs

Perhaps the best indication of how useful joint programs can be is that their number has grown from one in 1981 to 19 today. They are listed in the diagram on the preceding page,which also indicates the IBM Line(s) of Business which are Research's partners in each of these programs.

Among the earliest of these joint programs (after the ASTL) were the Magnetic Recording Institute formed in 1982 to focus on storage technology and the High- and Low-End Packaging Technology Laboratories started in 1983. They are essentially analogs of the ASTL and have produced the same kinds of benefits in their equally important fields of technology.

More recently, the joint program concept has been applied to software technology and to several systems areas as well—to the point where virtually every field in which IBM Research works (except for the physical and mathematical sciences) includes at least one joint program.

In each case, the key elements are:

- Focus on the reduction to practice of the next generation technology.
- Resources (people as well as dollars) provided by a development organization and by Research.
- A single program director, based in Research.
- A single comprehensive plan, agreed to by both parties, defining the work required to bring the new technology into the product plan.
- A long-term management commitment to continue the work.

As should be evident by now, we rarely transfer technology anymore in IBM Research. Instead, we develop it jointly with our colleagues in the product divisions.

This is what works. ⓐ

Acknowledgments

The author gratefully acknowledges the contributions to this article by several IBM colleagues, including Hirsh Cohen, consultant to the director of research; the late Robert Henle, who directed IBM Research's first joint program, the Advanced Silicon Technology Laboratory, from its inception until his untimely death in January 1988; and James McGroddy, former Research Division director of semiconductor science and technology, who is now IBM director of research.