

A Perspective on Computing Research Management

Roy Levin

Microsoft Research Silicon Valley

roy.levin@microsoft.com

ABSTRACT

This paper offers a perspective on a particular set of principles that have guided the development and enabled the success of several noteworthy corporate research labs in computer science. The paper examines the differences between the corporate computing research environment and academia, then describes the model for managing research that Microsoft Research employs, illustrating how it reflects those differences and what the consequences are.

Categories and Subject Descriptors

K.7.1 [The Computing Profession]: Occupations, K.6.1 [Management of Computing and Information Systems] Project and People Management

General Terms

Management

Keywords

Computing research management, technology transfer.

1. PREFACE

This paper offers a perspective on a particular set of principles that have guided the development and enabled the success of several noteworthy corporate research labs in computer science. That perspective is a synthesis acquired from several viewpoints over many years, beginning in 1977. I was the beneficiary of those principles when as a fresh PhD graduate I joined the Computer Science Laboratory at Xerox PARC, where they had been laid down by the founders of the lab, chiefly Bob Taylor. The same principles also served well at DEC's Systems Research Center, where I was a researcher and eventually the director. When I came to Microsoft in 2001, I had the opportunity to apply them in building the Silicon Valley research lab, although many of the same principles had already characterized Microsoft Research since its founding in 1991. Thus, while this paper describes an approach to research and research management that I have lived and practiced for nearly 30 years, I cannot claim credit for its invention.

2. INTRODUCTION

Corporate and academic research in computer science share many similarities and exhibit many differences. However, corporate research settings exhibit more differences among themselves than universities do, both in the way they carry out research and the way they are managed. This should not be surprising. In a broad sense university-based research has two foci, independent of the particular university: to expand human knowledge and to teach the next generation of researchers how to carry out research in their chosen field. By contrast, there is not a common focus for corporation-based research; rather, its character differs markedly depending upon the particular corpo-

rate objectives it is intended to address. (Government-based research, for example at a national laboratory, is yet another kind of animal, but one that I will not discuss further here.)

In this paper, I first attempt to characterize the range of differences that exist among university-based computer science research and the various kinds of corporation-based research, then focus on a particular model for the later. That model, which I have experienced in three successful corporate labs, is perhaps the most "university-like" of corporate research models, yet has important differences that recognize and address the needs of the corporation. It would be wrong to think of this as the "best" model, since corporations have varying needs and what meets the objectives of one company may fail badly in another. But the model does have a proven track record of both expanding knowledge in the field and contributing to the growth of the computing industry, and as such it is, I believe, worth understanding.

A couple of notes on terminology are in order. This paper is about computing research, but for brevity I will henceforth usually omit the word "computing". Also, I will often speak of computing research when I really mean computing research management, that is, the principles and practices for developing and operating a computing research organization. The choice of particular research topics and the means by which they are addressed are of only passing relevance to the discussion.

3. COMPARING RESEARCH MODELS

There are a number of dimensions along which corporate and university research may be compared. I briefly consider several of them here as background for a more detailed examination of the MSR model in the next section.

Funding criteria. In a corporate lab, funding of a particular research activity may depend on the relevance of the hoped-for result to the corporate business. Corporate labs vary considerably in this regard, with some having funding closely tied to business unit activity in a significant portion of their research portfolio and others having a much looser coupling between research topics and business relevance. In a university, funding of a research area is based on criteria established by a funding agency. While agencies vary, in general the space of research funded through government agencies is much broader than a particular corporation funds.

Patent protection. In a corporate lab, patent protection for inventions is generally encouraged. Researchers who (help attorneys to) file patent applications may receive monetary awards when the applications are filed or the patents issue, or both. In a university, patents are less commonly sought, although some universities do try to develop and make money from a patent portfolio.

Publication. Corporate labs vary widely in their view of publication of research results. At one end of the spectrum, some labs strongly encourage publication and evaluate researchers in part

on their publication record. At the other end, some labs discourage publication either explicitly or by creating administrative impediments such as extensive legal review of proposed publications. In a university, publication is mandatory for faculty who seek to be tenured, and the quality of the journals and conferences in which peer-reviewed publication occurs is a significant factor in evaluation of tenure promotion cases.

Resources. In a corporate lab, non-personnel resources such as computing equipment are usually available as needed. That is, most computing research projects are not limited by budgetary constraints in acquiring equipment, although certainly there are exceptions in which research requires a large amount of infrastructure that may pose a budgetary challenge. In a university, funds for non-personnel resources are often more limited, which significantly affects the character of the research that can be undertaken in some specialties.

External research collaborations. In the corporate setting, the ability to collaborate with researchers outside the corporation depends substantially on the company's intellectual property strategy. Such collaborations generally require explicit agreements between the two organizations, which involve lawyers, negotiation, and considerable management participation to create. By contrast, university researchers collaborate freely and informally when doing so is beneficial to their work, with little or no administrative impediment. Significantly, some corporations treat collaborations with academic researchers differently (and more flexibly) than collaborations with other corporate researchers.

Research organization structure. Corporate labs vary considerably in the depth of their organizational hierarchy. Of course, this is partly a function of the size of the lab, but some labs emphasize "flatness" much more than others. A university department is typically very flat, with most or all faculty reporting to a department head (a role that frequently rotates). However, each professor is typically the "boss" of a number of graduate students, directing their work on projects that involve varying degrees of collaboration. A corporate lab may be similarly flat, but more commonly divides into a collection of groups each focused on a technical area or business priority of the company.

Management responsibility. A fresh PhD graduate joining a corporate research lab typically has no people-management responsibilities initially and may not acquire any for many years. By contrast, a fresh PhD graduate who becomes an assistant professor typically becomes responsible for advising a collection of graduate students within a few years' time, sometimes in the first year. In universities, research remains part of a professor's job despite management responsibilities, while in a corporate lab, a research manager may or may not pursue an individual research role as well.

Teaching responsibility. Teaching is nearly always a required part of a university faculty member's job, since education is the official primary function of a university. Corporate researchers rarely have teaching responsibilities in the same sense that professors do, although more senior ones may be expected to mentor more junior ones. However, some corporate research labs support teaching in the traditional sense, meaning that they allow or even encourage researchers to spend part of their time teaching a seminar or more formal course, typically at the graduate level, at a local university.

Advancement. Advancement within a corporate research organization may take one of two paths. Progressing up the technical

ladder generally depends on positive impact on the company's products or services, that is, the directly perceivable value of the researcher's work on the company's output. Progressing up the management ladder is likely to be determined by the company's overall view of the role of people managers and is of course affected by the degree of opportunity afforded by the depth of the research organization's structure. In addition, companies vary in the extent to which promotion as an individual technical contributor is possible, although the technical ladder in most companies with highly regarded research organizations parallels the management ladder roughly to the vice-president level. In a university, advancement on the professorial ladder is based on professional standing, which in turn is substantially based on peer review of papers submitted for publication and, to lesser degrees, on service to the professional community and research grants obtained.

Direct impact outside the research community. In a corporation, researchers contribute to the business either through transfer of specific technology to product groups or by sharing with product developers the expertise gained through research. The latter is sometimes termed "internal consulting", and is nearly always an expected part of a researcher's job. For university-based researchers, impact outside the research community comes chiefly through consulting engagements with companies who seek the technical expertise acquired through research. Such consulting is optional but common, especially since it provides an attractive way for a faculty member to supplement the academic salary. Universities often facilitate external consulting by allowing their faculty to sign part-time agreements with companies in which they assign intellectual property rights to the company. In this way, companies can acquire necessary expertise and customized problem-solving on a case-by-case basis.¹

4. MICROSOFT RESEARCH'S MODEL

The preceding section demonstrates the breadth of the space of possibilities for structuring a research organization. I now turn to the particular choices made by Microsoft Research (MSR). These choices follow directly from MSR's strategy, which is:

- to advance the state of the art in the areas in which we conduct research and
- to bring those advances to Microsoft's businesses.

While, for the purposes of this paper, MSR's strategy can be taken as a given, the rationale for that strategy helps to provide an understanding of some of the specifics explained below.²

¹ In the extreme case, a faculty member may temporarily join a company, sometimes as a founder or in a central technical role. However, in such cases the professor generally goes "on leave" and works full-time for the company. While this is certainly a potential way to have impact outside the research community, the professor is effectively leaving that community, albeit temporarily. Thus, it is not really "consulting" in the sense used here.

² A historical note: MSR's strategy, established when Rick Rashid came to Microsoft in 1991 to build the organization, follows directly from his experience as a professor in the CMU Computer Science Department and from his observations about the failure of a number of earlier and contemporary corporate labs to transfer their technology effectively. A number of key researchers and research managers whom Rick brought to MSR had worked in these labs. Based on their experiences,

In particular, MSR's top priority is to advance the state of the art. This is not universal among corporate research labs; indeed, it is generally assumed that the first priority of a corporate lab is to transfer technology and/or expertise to the company's businesses. That is the second half of MSR's mission, but it depends on innovation: the creation of new technology. Microsoft's businesses depend on novel technology, which can be obtained either through internal creation or external acquisition. Microsoft does acquire technology externally, but that channel is inherently limited by what other organizations do and are willing to make available. By focusing on advancing the state of the art, MSR creates technology that is not available externally, thereby providing a unique "innovation channel" for Microsoft's businesses. Furthermore, making advancement of the state of the art the top priority emphasizes and fosters the necessary innovation, which otherwise is often limited or stifled by subservience to short-term business contribution.

How does the MSR management create this innovation channel? To a very considerable extent, the principles underlying MSR's operation are concisely described by John Naughton and Robert W. Taylor in their paper "Zen and the Art of Research Management" [1]. I use their pithy prescriptions for research management as a starting point and elaborate on their application in MSR. I have taken the liberty of reordering and grouping their points into half a dozen topics, each treated in a subsection below, although no topic cannot really be treated in isolation from the others. In each subsection, Naughton and Taylor's observations appear in italics, generally at the beginning.

4.1 It's All About the People

HIRE ONLY THE VERY BEST PEOPLE, EVEN IF THEY ARE CURSED. Perhaps especially if they are cursed. Your guiding principle should be to employ people who are smarter than you. One superb researcher is worth dozens of merely good ones.

ONCE YOU'VE GOT THEM, TRUST THEM. Do not attempt to micro-manage talented people ... Set broad goals and leave them to it. Concentrate your own efforts on strategy and nurturing the environment.

The hiring process at MSR is highly selective. In a typical year, Microsoft Research Silicon Valley (MSR-SVC), a 40-researcher group located in Mountain View, California, receives over 100 applications, invites about 20-25 candidates for interviews, and extends offers to 4-6 of them. The selection process focuses on quality of research work within the areas of relevance for the lab, which are the practical and theoretical disciplines related to distributed systems. The process involves all members of the lab, in that all lab members are expected to attend the technical talk given by a candidate, participate in one-on-one interviews of several of each year's candidates, and contribute to a hiring discussion involving the whole lab that precedes the lab director's decision on each candidate. This process ensures that everyone who joins the lab has the support of a substantial fraction of the organization, and everyone in the lab knows who is likely to collaborate with a new arrival.

MSR emphasizes "bottom-up" research. That is, researchers choose their own research agendas and pursue them with minim-

al involvement by management. Most research projects cannot be carried out solo and therefore depend on collaborations, which is why the hiring process at MSR-SVC requires consensus and enthusiasm for bringing particular researchers into the lab. Thus, when a new researcher arrives, the pathway for collaboration within the lab has already been paved.

There is a balance that must be struck between overlap and uniqueness in the researchers' interests. Too much overlap and the opportunity for impact becomes circumscribed; too much uniqueness and collaborations will not form. Every researcher who joins the lab shares some technical strength with some of the existing members, but introduces significant additional technical expertise. The "bottom-up" principle encourages that researcher to engage in work based on that new expertise, which then broadens the lab's research compass and increases its potential for impact on the company.

4.2 Absence of Hierarchy

KEEP THE ORGANISATION CHART SHALLOW. Never let the lab grow beyond the point where you cannot fit everyone comfortably in the same room.

The most interesting research work often occurs when two technical areas grow close enough together that two researchers, one in each, can work together. Researchers naturally cluster in established areas, but if the organization chart institutionalizes those areas, cross-connections become more difficult since researchers naturally look within their cluster for collaborators and, to some extent, set their research agenda accordingly. A shallow organization chart – in particular, one that doesn't group researchers by topic area – increases the opportunity for cross-connection and unexpected research results.

The principle is practiced in MSR-SVC by having no fixed organizational structure within the group. Instead, researchers create informal associations around technical work on a project-by-project basis. Such associations are inherently of limited duration; when the project ends, the "group" no longer exists. Because of technical affinities, the same researchers may work together over a number of projects, but the project leader often changes. Moreover, the role of project leader is largely a technical one that does not include people management. Thus, project leaders do not take on budget responsibility or performance review responsibilities as a side-effect of technical leadership. The membership of a project is determined collectively (and often spontaneously) with researchers being included based on the relevance of their technical interests and skills.

At first blush, this sort of structure, in which the lab director has 40 or more direct reports, appears unworkable. In practice, it works well provided that the members of the lab value it and are willing to invest some of their time to help the management keep it functioning. Obviously, a single manager cannot provide individual mentoring for 40+ researchers, so this role has to be distributed, with senior researchers helping to advise more junior ones on their research directions and professional development. At MSR-SVC, approximately a third of the researchers operate in this mentoring role with no individual advising more than four others. These mentors also help to draft their advisees' annual performance reviews, although the reviews themselves and attendant compensation actions are handled by the lab director. Thus, the relationship is that of a more senior researcher and a more junior professional colleague, and it operates without the complications that a organization chart link introduces.

they shared a philosophy about the best way to conduct computing research in a corporate setting. As a result, many of the principles and practices mentioned in this paper have spread naturally throughout MSR's labs around the world.

The chief limitation of this informal structure is scalability. It works effectively as long as everyone in the lab can keep track of what everyone else is doing, at least to the extent that each researcher, off the cuff, can say a sentence or two about what nearly every other researcher in the lab is working on. When this is the case, researchers are able to form projects and choose appropriate collaborators within the lab. Obviously, some investment of time is required to maintain the necessary level of familiarity with others' work. At MSR-SVC, a weekly meeting of the entire lab at which researchers take turns giving talks on current technical work provides a natural and low-overhead mechanism to keep informed. (As Naughton and Taylor note, this requires a room into which the whole lab can comfortably fit.) To a lesser extent, other activities help as well, including lunch-time discussions, group off-site lunches, and the hiring discussions already mentioned. The physical arrangement of the lab is relevant too, since it can encourage informal discussion and maximize chance encounters with lab members outside a researcher's current projects. Finally, to increase the management bandwidth, the lab has an assistant director who works closely with the director but without introducing another layer in the management hierarchy. Experience suggests that the informal structure of MSR-SVC will scale to a lab of 50-60 researchers, and possibly somewhat more, before the size and breadth of technical interests overwhelm the mechanisms that make it work.

Because research projects are collaborative, they are formed by attraction of like-minded researchers, not by management fiat. This is in sharp contrast to a product group, in which a hierarchical structure is the norm and individuals are generally assigned based on their skills and project resource needs. A research project that is unable to recruit³ sufficient collaborators simply doesn't happen.

4.3 Life in the Office

REMEMBER THAT CREATIVE PEOPLE ARE LIKE HEARTS: they go where they are appreciated. They can be inspired or led, but not managed.

PROTECT YOUR RESEARCHERS FROM EXTERNAL INTERFERENCE, whether from company personnel officers, senior executives, or security personnel. Remember that your job is to create a supportive and protective space within which they can work.

MAKE YOUR RESEARCHERS DEBATE WITH ONE ANOTHER REGULARLY. Let them tear one another's ideas to pieces. Ensure frank communication among them. Observe the strengths and weaknesses which emerge in the process.

A broad set of research areas (the MSR web site lists more than fifty) and the bottom-up principle establish the context for researchers to pursue their work. Without that work, there can be no innovation channel and MSR cannot fulfill its mission. Thus, it is essential to the success of MSR that researchers be impeded as little as possible by administrative burdens. The creative

process requires uninterrupted time, and ideas and insights occur at unexpected moments. If the environment does not recognize this, it compromises creativity. There are many ways in which a research manager can shape the environment to respect and foster creativity. Here are a few:

- Researchers must travel, both to participate in conferences that keep them up-to-date with the state of the art and to interact with product groups in locations remote from the research lab. When management makes travel easy – from authorization to booking reservations to completing expense reports – it minimizes the operational and psychic overhead associated with this necessary part of a researcher's job. In Microsoft, online systems enable researchers to book travel and complete expense reports efficiently, and sensible reporting requirements reduce the amount of bookkeeping that travelers need to do in order to be reimbursed. (It is amazing how petty many corporations are about such matters, wasting relatively expensive employee time and sending the message that those employees aren't trusted to report their expenses honestly.) At MSR-SVC, researchers are told to exercise their best judgment as to which trips are necessary (there is no fixed approval process), with the result that nearly all desired travel can be supported.
- Some areas of computing research require moderate amounts of equipment, and most require occasional modest purchases of equipment and software. By streamlining simple purchases and enabling both online ordering and invoice approval, Microsoft again minimizing the administrative overhead to acquire materials researchers need to carry out their work. MSR management further reduces the burden by setting purchasing authorization limits high enough that most purchases incur no management intervention.
- Research is not a 9-to-5 activity, nor is it conducted solely in the workplace. MSR management recognizes this reality and supports creativity by allowing researchers flexible hours and not requiring an accounting of time spent outside the office. It is worth noting that this works in part because much of the business of research can be effectively conducted by email. However, excessive telecommuting can undermine effective collaboration and prevent the serendipitous interactions that often give birth to new ideas, so researchers who are not traveling are expected to be in the workplace a significant portion of most work days.
- As previously noted, collaboration is essential for most successful research. Since no lab has a corner on good ideas, external collaborations offer a mechanism for extending the breadth and depth of a lab's work. There are challenges to creating external collaborations in a corporate setting, with intellectual property ownership being the chief one. By creating processes that make such collaborations easy, from the creation of suitable consulting agreements to the administration of expense reimbursement for consultants, MSR management enables and encourages researchers to collaborate with experts outside Microsoft. (MSR-SVC, for example, has 20+ such arrangements every year, an average of one for every two researchers.)
- *REMEMBER TO INITIATE AND SPONSOR CELEBRATIONS when merited.* Managers in any organization, not just research, should heed this principle, and budget accordingly. The qualifier "when merited" should not be overlooked – celebrations should recognize significant events,

³ The recruiting of colleagues to collaborate on a research project is traditionally called "Tom Sawyering" by analogy with that character's approach to enlisting fence white-washers. While colorful, the analogy is imperfect, since Tom sought to avoid doing the white-washing himself, whereas researchers are eager participants in the projects for which they recruit colleagues.

so that they do not become devalued, but management should also avoid setting the bar too high. In a well-functioning research organization, there should be good reasons to celebrate annually and possibly somewhat more often. In MSR, celebrations tend to recognize team- or lab-wide achievements more frequently than individual achievements, though the latter are certainly important too.

- *INSTALL A WORLD-CLASS COFFEE MACHINE and provide plenty of free soft drinks.* It has been observed that a mathematician is a machine for turning coffee into theorems.⁴ An analogous remark applies to computing researchers. When MSR-SVC was founded in 2001, the first capital purchase was an espresso machine (and Microsoft provides free soft drinks throughout its facilities).
- *BUY AERON CHAIRS. Remember that most computer science research is done sitting down.* MSR pays heed to this remark and expands on it to include any ergonomic equipment that assists researchers to maintain their physical well-being in the office.
- *INSTITUTE A “TOY” BUDGET, enabling anyone in the lab to buy anything costing less than a specified amount on their own authority ...* Effective researchers are curious about many things, and a wise manager provides opportunities in the workplace for them to go where their curiosity leads them. A budget for “toys” is an example. Microsoft provides extensive online purchasing available to most employees and the purchasing limit for researchers enables them to satisfy their curiosity about many things.
- It is perhaps obvious that the physical space occupied by a research lab can help or hinder creativity. The space needs to support simultaneously both individual work in a quiet environment and group discussion, with group discussion space being available most of the time without pre-planning. The space should also encourage the chance meeting. The aforementioned coffee machine is one such mechanism, but management should also consider arranging the office assignments so that likely collaborators are not too near each other, thereby forcing them to encounter others as they move through the lab.

Naughton and Taylor’s advice about encouraging debate among researchers (the third quotation at the start of this subsection) must be implemented carefully, especially the second sentence. The quality of research ideas indubitably improves with discussion and debate. “Tearing one another’s ideas to pieces” must be understood as a constructive activity; that is, management should encourage lively, exploratory discussion whose purpose is to get to the best solution for a problem. At the same time, management must not permit the tearing apart of ideas to stray into the tearing apart of individuals. Productive results depend on collegial respect, and if a researcher cannot try out a half-baked idea on colleagues for fear of losing their esteem, the lab’s collaborative environment will collapse.

4.4 Technology Transfer

MUCH OF WHAT YOU DO WILL FALL INTO THE CATEGORY OF ABSORBING THE UNCERTAINTY OF YOUR RESEARCHERS.

DO NOT PAY TOO MUCH ATTENTION TO “RELEVANCE,” “DELIVERABLES,” and other concepts beloved of senior management.

MSR’s strategy depends on eventual contributions to Microsoft’s businesses. Although the research areas in which the contributions will occur can generally be anticipated (and, therefore, are areas of focus within the research labs), the timing of those contributions frequently cannot be predicted with confidence. This makes research unique within the corporation, for much of a company’s activity follows from the heartbeat of the external financial world. That is, the expectations of the marketplace, and financial analysts, often strongly influence the actions of the company’s senior management, either on an annual or a quarterly basis. Research results cannot fit comfortably within these rigid timing constraints. For this reason, the senior management in Microsoft does not impose short-term quantitative metrics on MSR. Instead, MSR’s deliverables are evaluated over the long term, meaning many years, and the evaluation is fundamentally qualitative, meaning that no attempt is made to attach a specific financial value to individual research innovations. Rather, the senior management looks at the depth and breadth of MSR’s contributions to Microsoft businesses in a cumulative way without trying to shape or direct it in the short term.

While measuring individual deliverables is not practical, the research management needs to judge how well the innovation channel is working overall. How good is the work? How useful is it to the business?

Fortunately, it is relatively easy to assess the quality of the research. MSR emphasizes publication by researchers in the professional literature, that is, the conferences and journals in relevant areas of specialization. Since these are peer-reviewed venues, the acceptance of a paper constitutes a judgment by the leaders of the field that the work represents an advance in the state of the art and is worth communicating to other practitioners. A significant presence of MSR work in the top-tier conferences and journals therefore represents an independent assessment of the quality of the research. Management is careful not to attach too much weight to results in any particular year, since there are many reasons why a particular piece of work may not be accepted for publication in a particular conference. Nevertheless, trends over several years are indicative. MSR has had a consistently strong presence in top publications in graphics, information retrieval, systems, and computing theory, to name just a few.

For novel work to be useful to Microsoft’s business in most cases, it is necessary (but not sufficient) for the central innovation to be patented. Suitable patents permit publication of the work while preserving Microsoft’s ability to choose how to use the technology. For this reason, MSR management strongly encourages researchers to seek patent protection for their innovations. Researchers work with MSR’s internal intellectual property legal counsel and outside attorneys who prepare the actual patent filings. MSR doesn’t use patents awarded as a specific metric of novelty; the professional peer-review process is more effective. Nevertheless, patents are an essential contributor to the effectiveness of the innovation channel.

A more direct measure of the utility of the research work is the extent to which Microsoft’s businesses adopt it. Management

⁴ This characterization is generally attributed to Paul Erdős, although some sources cite various of his colleagues.

doesn't find year-by-year metrics to be particularly useful, since many factors affect the "take-up rate" of MSR technology, including in particular the timing in the development cycle of a product. But a multi-year examination of the breadth of MSR's influence on Microsoft's businesses helps management to assess how well the innovation channel is working. Since 2001 when MSR-SVC was founded, during which time the lab averaged approximately 25 researchers, there were three major technical engagements with product groups in which researchers played leading architect roles that leveraged the technical expertise derived from their research work. There were also at least a dozen specific technology transfers varying in scope from algorithmic "nuggets" to complete (prototype) implementations of eventual products. These transfers spanned three of the seven main business units of the company. Looking at MSR overall during its 15-year lifetime, research technologies have affected virtually every business and product of the company.

Of course, technology transfer doesn't just happen. It requires substantial effort, both by researchers and management, and even then it frequently doesn't succeed. Books have been written on the subject, and there are many business case studies that document the difficulties of successful technology transfer. This paper cannot hope to cover all the reasons why it works at Microsoft, but a couple of facets deserve attention.

MSR has a small group of "program managers" (MSR-PM) that plays an essential role in the technology transfer process. At Microsoft, program management is a role whose specifics vary widely from business to business, so the job title is not particularly informative. However, many program managers in product groups have responsibility for a particular component or feature of a product. That responsibility includes specification and design, implementation and test, and sometimes documentation and/or marketing as well. The individuals in MSR-PM usually have served in that role, so they know very well what it entails, but now operate as a "connector" between others in that role and researchers with relevant technology and expertise.

This connection function is bi-directional. When a researcher has produced something of potential value to a product group, (s)he may often not know where that potential recipient might be in the company. Microsoft is a big place – over 70,000 people – and like most large companies it can be difficult to navigate. The MSR-PM team makes it their collective job to know in significant depth what is going on in the individual product groups, especially when those groups are in a phase of their development cycles in which they are receptive to new technology. Thus, the researcher looking for a technology outlet can go to the MSR-PM team and enlist their help in making appropriate connections. By approaching the right people at the right time, researchers substantially increase their chances of initiating a successful transfer. MSR-PM also assists with connections in the opposite direction by being the "go to" organization for program managers in product groups who seek technical expertise to solve a particular problem. The MSR-PM team makes it their collective job to know what is going on inside MSR, especially including the expertise of individual researchers. With the help of MSR-PM, the product group program manager can efficiently find the right researcher to help solve a pressing technical problem. This has the additional benefit of establishing a connection that can become the basis for future engagements and may even stimulate a new research direction influenced by the product group's longer-term needs.

In a company of 70,000+ employees, even the well-informed MSR-PM team may not spot some technology transfer opportunities. To provide a more direct though less "managed" mechanism for connection between researchers and product groups, MSR puts on an annual internal trade show called TechFest. At this multi-day event, the results of research projects from MSR's labs around the world are presented in nearly 200 demonstration booths and lectures. Thousands of employees walk around the show floor visiting the booths that interest them and talking directly with the researchers who carried out the work. In this way, unanticipated connections can be made, enabling research technology to be applied in novel ways.

Both MSR-PM and TechFest are mechanisms designed to "impedance match" researchers and product developers. Technology transfer requires technically knowledgeable individuals on both sides, but technical knowledge is not sufficient. Researchers and product developers think about problems very differently. Their objectives, their constraints, and even their vocabularies differ. These differences can create an impermeable barrier that blocks the flow of research work to the rest of the company unless explicit "impedance matching" efforts occur. The marked decrease in the number of computing research labs in corporations can be partly attributed to a lack of mechanisms for matching the impedance of researchers and product developers.

Technology transfer is not the only way in which MSR contributes to Microsoft's business, although it is the primary one. Not everything that MSR invents can be used by the businesses. Creativity, by definition, cannot be predicted. While management can make a reasonable guess about the areas in which a particular researcher's innovation will occur, surprises often occur. Sometimes those creative surprises don't fit well with the company's businesses. The lack of fit may be technical or may be a matter of timing; that is, sometimes a researcher's idea is potentially relevant, but the company is simply not organized in a way that permits the idea to be exploited promptly. In such cases, the technology may be licensed externally. This is a relatively new but expanding means for MSR's work to benefit the company's business. It is also one that generally requires a lower time investment by the researcher than internal transfers do.

4.5 Quiet Leadership

REMEMBER THAT YOU [the manager or a research lab] ARE A CONDUCTOR, NOT A SOLOIST ...

Whether facilitating the creation of technology or its transfer, the research manager is enabled by the work done by the researchers. When a recognizable success occurs, publicity often centers on the research manager. It is a wise manager who keeps this in proper perspective and turns the spotlight on the individuals who carried out the real work: the researchers. Proper attribution of work forms a cornerstone of professional research conduct, and the manager who takes pains to highlight the work done by his/her researchers pays proper respect to that principle. When the conductor takes a bow, it is on behalf of the entire orchestra. The respected conductor ensures that individual performers receive the recognition due them.

MSR includes a small communications group whose role includes ensuring that individual researchers' work is properly recognized externally. This includes arranging interviews with the technical press, publishing researcher profiles on the Web, and staging technical "road shows" in which the researchers

present their work to selected external audiences. Since all of these activities compete with research for the researchers' time, they must be carefully vetted by research management to achieve a proper balance, one in which the recognition increases the researcher's future impact.

For audiences within Microsoft, researchers, not management, generally present their work. This includes technology transfer engagements and TechFest (discussed above) as well as more general settings. For example, Microsoft has an annual all-employee meeting in which, among other things, technologies from across the company are demonstrated. Researchers demonstrate their work in this meeting; their managers are merely spectators.

4.6 Keep in Touch

PAY ATTENTION TO WHAT GOES ON IN UNIVERSITIES. Every significant breakthrough in computing in the last four decades has involved both the university and corporate sectors at some point in its evolution.

BE NICE TO GRADUATE STUDENTS. One day they may keep you, even if only as a mascot. (Moreover, they are a lot of fun!)

The inherent differences between corporate and academic research, described in Section 3 of this paper, make it clear that a corporate organization focused on broad innovation must stay well-connected with university research. MSR does this in many ways. Perhaps the most obvious connection results from the emphasis on publication and peer review (described previously), which ensures that researchers stay up-to-date on the latest developments in their specialties by reading the professional literature in which academics chiefly publish. But publication lags research results by months to years, so reading the literature alone does not provide sufficient currency.

To stay at the leading edge of research work, MSR engages with universities on several fronts. Every year, more than 250 graduate student interns come to MSR's US labs to participate in research projects. They bring deep knowledge of work going on at their universities and, in the course of their internships, they share that knowledge with MSR researchers through presentations and informal discussions. These students inject a fresh perspective and new ideas into MSR projects and research directions. In addition to strengthening work in MSR, the students often acquire a broader perspective on their own work, one that they take back to their universities and that influences the work

there. (And, as Naughton and Taylor remind us, having the students around is a lot of fun.) MSR also encourages its researchers to teach courses and seminars in local universities, which provides them a qualitatively different opportunity than the internship program to interact with graduate students, to their mutual benefit.

MSR also engages deeply with individual faculty members on topics of mutual interest. These engagements extend from short-term "consulting" in which a faculty member works with MSR researchers for a week or two, to year-long "visiting researcher" appointments in which a faculty member on sabbatical works at MSR full-time. These collaborations nearly always produce jointly authored papers and frequently extend MSR's reach into areas that would not otherwise have been explored.

Though MSR is a large organization (over 750 researchers world-wide) with a broad agenda, it does not attempt to cover all computing-related research areas. There are many topics in which MSR does not invest internally but which it deems worthy of investigation. Accordingly, MSR supports university research through unconstrained grants, fellowships, and grants for specific work (generally determined by through a proposal process). MSR also operates a faculty fellowship program that identifies and supports innovative university researchers in the early stages of their careers when they may find government funding difficult to obtain.

5. ACKNOWLEDGMENTS

Many people have contributed to the elaboration and application of the research model practiced in Microsoft Research. However, the origin of the principles that underlie the model is clear: Bob Taylor. Those principles influenced academic research beginning in the 1960's when Taylor guided much groundbreaking work from his leadership position in ARPA. His influence has reached MSR through many people who lived and worked in the environments that he created. Taylor was my boss for 19 years at Xerox PARC/CSL and DEC/SRC, and it is my privilege to apply his research principles in the lab I presently manage.

6. REFERENCES

- [1] Naughton, J., and Taylor, R. W. Zen and the Art of Research Management. In *Computer Systems: Theory, Technology, and Applications* Andrew Herbert and Karen Spärck-Jones, ed., Springer Monographs in Computer Science, New York, NY, 2004, 223-4.