59 Technology Transfer

Kevin M. Schofield

CONTENTS

59.1 Introduction 1331
59.2 New Face of Technology Transfer 1331
59.3 Internal Technology Transfer 1332
59.4 Technology Transfer as a Social Process 1332
59.5 Building Trust in Steps 1333
59.6 Thinking for the Future 1334
59.7 Considering HCI in the Tech Transfer Process 1334
59.8 External Technology Transfer 1335
59.9 Know What You Are Selling 1335
59.10 Know What You Are Not Selling 1335
59.11 Know the People Involved 1335
59.12 Crafting a Deal 1336
59.13 Confidentiality and NDAs 1336
59.14 Going to the Negotiating Table 1337
59.15 Embracing the Relationship 1338
59.16 Commonalities for Internal and External Technology Transfer 1338
59.16.1 Intellectual Property 1338
59.17 It All Comes Down to the Relationship 1338
References 1339

59.1 INTRODUCTION

This chapter is intended to provide guidance on how to approach technology transfer of HCI-related research. Admittedly, there are many different perspectives one could take in looking at this often-difficult problem: researcher versus practitioner, industry versus government versus academia. There is also the extra, added dimension of whether the transfer is within the boundaries of a corporation or across corporate lines.

In this chapter, we will discuss tech transfer both internal and external to a company and then discuss commonalities across the two processes. Primarily, we will look at it from the perspective of the researcher, because in practice the burden is more on the researcher to justify the transfer and to make it work. Practitioners, however, will also gain value from reading this, as it will help them to understand the role of the researcher in technology transfer.

Although any technology transfer has its challenges, HCI tends to be a particularly difficult one. This is due to many factors, but the two main factors are that it is often more about abstract ideas than specific implementations and because we in the HCI community are still fighting the (wrong) impression that HCI is an afterthought and not the "real meat" of research and development efforts.

59.2 NEW FACE OF TECHNOLOGY TRANSFER

The last few decades have seen dramatic boom-and-bust cycles in the technology industry that, for a number of reasons, has changed the face of business, commerce, and information flow. Probably, the most earthshaking of those changes has been within the high-tech industry itself, as fitting for such a rabid consumer of its own technologies and ideologies.

The Internet, through its nearly ubiquitous connectivity to information and other people, has reshaped organizational patterns and forged a brand new set of relationships between researchers, designers, developers, manufacturers, marketers, and the other people involved in business and commerce. Those new organizations and relationships respect neither national nor corporate borders and bring us one step closer to the "friction-free" economy that was trumpeted in the mid-1990s.

Technology transfer is not what it used to be, largely because research and development is also not what it used to be. As Friedman (2005) pointed out, the world has flattened, and increasingly research and development can be done anywhere, with anyone, and for anyone. In lockstep, the relationships between corporate and academic R&D have

also evolved; there are fewer and smaller standalone corporate research laboratories as many companies have chosen to "outsource" their basic research to academia. Even those companies that continue to invest in basic research are realizing that with fewer barriers to communication and collaboration, an "open innovation" model of joint R&D, as suggested by Chesbrough (2003), often makes sense to maximize the impact of internally generated ideas as well as to take full advantage of the most relevant externally generated ideas.

Government research, and funding, is also changing. Some of this is due to economic cycles; much is due to shifting political priorities. In Europe, for example, the Lisbon Agenda defines the European Union's desire to make Europe the most innovative economy in the world by 2009, and this drives their investment in research and development. In contrast, in the United States, despite a wide recognition that IT is driving economic advances, the major funding agencies (NSF, DARPA, and NIST) have reduced their support for academic research in computer science, with a built-in assumption that private industry should and will take up the slack. Somewhere in the middle, many countries (e.g., Canada, Australia) have programs where academic researchers who are funded by private companies can apply to the government for matching funding.

Still, as the National Research Council (2003) observed, successful technology transfer is usually a marathon, not a sprint, and may require years to reach completion. It also reflects a complex partnership between government, industry, and academia. What should we take away from this? That, clearly there is no one model for technology transfer from research into industry—there exists a whole spectrum. As mentioned previously, some governments fund research, while others do not, and some are in the middle. Likewise, some companies do their own basic research, while others outsource it, and some are in the middle. Some universities keep a death grip on their faculty's intellectual property; others, such as University of Wisconsin and University of Waterloo, have a policy that IP belongs to the inventor. Many are in the middle.

Consequently, there is no "playbook" that will explain exactly how to transfer a research result. Both researchers and practitioners will need to be very flexible and willing to adjust to the situation in which they find themselves. Tech transfer is, in the end, a business, and the best way to succeed at it is to think like a businessperson.

59.3 INTERNAL TECHNOLOGY TRANSFER

I am constantly surprised by the number of people who believe that technology transfer is some sort of Rube Goldberg machine,* where technology is inserted in one end of the contraption, strange things happen in the middle that usually involve people in uncomfortable and contorted positions, and then magically it pops out on the other end. Countless books and articles have been written (Lesko, Nicolai, and Steve 1998; Mock, Kenkeremath, and Janis 1993) in an attempt to document the perfect mechanical process for technology transfer. And yet, despite the fact that nearly everyone has had painful experiences trying to define a mechanical process for technology transfer, they still try to do it and complain bitterly when it fails (Butler 1990; Hiltzik 1999; Isaacs and Tang 1996; Singh 1993).

59.4 TECHNOLOGY TRANSFER AS A SOCIAL PROCESS

Tech transfer is not a mechanical or logistical process; it is fundamentally a social process. It succeeds when people build a relationship between the provider and the recipient that fosters trust and communication. Manning (1974) recognized that successful technology transfer centers on viewpoints and perspectives and fundamentally on communications. Foley (1996) spoke to this point most directly, that technology transfer is a "full-contact sport" that centers on the people.

Successful product organizations understand that risk is their mortal enemy. They work hard to proactively manage the risk in their development process or to remove the risk factors altogether. One of the most prevalent and difficult-to-manage risk factors is an external dependency, and let us be honest: an external dependency from a research organization looks about as risky as it gets. As long as your counterparts in product organizations think of a research organization that way, technology transfer is difficult at best and often outright impossible.

To succeed with technology transfer, we need to mitigate the risk or at least the perceived risk in the minds of the people we wish to receive our technology.

Up to this point, none of this is particularly controversial, but this is where the paths diverge. Many people will tell you that you succeed in mitigating the risk by creating well-defined, step-by-step processes through which you and your industry partner will enact the technology transfer. I argue that this approach fails more often than it succeeds, for two main reasons:

- The only experience that people in industry have with external dependencies is the occasional dependence on an external contractor or supplier to deliver a finished component ready for integration. They inevitably use this as the model for defining their tech-transfer process from research, and it is fundamentally incompatible. Research technologies are not finished components, and any product organization that expects a research group to deliver a finished component fundamentally misunderstands the role, expertise, and hiring practices of a research organization. Research groups almost never understand the development and test practices of a commercial product organization; even if they did, those practices vary so widely between organizations that past

* For those in the United Kingdom, that would be "Heath Robinson machine."

experience does not guarantee that they could successfully deliver a finished component. Moreover, it is not the goal of a research organization to develop technologies into finished components; its goal is to discover and prove solutions to previously unsolved problems. To do so, a research organization requires different skills and expertise and frankly different development and testing methodologies. Both research and product organizations need to understand this fundamental difference and embrace it as a way to complement each other's strengths, rather than ignore it and delude them about a theoretical capability that, practically speaking, is not there.

- People who do not understand each other cannot communicate and do not trust each other and cannot be expected to work cooperatively toward a shared goal, even within the best-defined process. The trust and communication must come first. If the two people on opposite sides of a table trust each other, then the two of them can accomplish anything; if they do not, they will never accomplish anything of value.

59.5 BUILDING TRUST IN STEPS

So then comes the catch: how does one build trust? By working side by side, of course! This means that one must start with the kind of activities that are initially low in risk, but high in communication, and build on one's successes to build more trust and overcome successively higher levels of risk.

Step 1 is to establish trust that one is an expert in the domain. Technical people, whether in research or in industry, are almost universally avid readers and understand the importance of staying up to date in their fields. But we all suffer from a lack of time to weed through the volumes of less-than-useful information to find the truly valuable nuggets of wisdom. If someone in a product organization expresses an interest in one's field, an offer to forward them a set of papers, articles, and books that represent the seminal reading is a great first response. Granted, doing a literature search is not glamorous work, but it fundamentally demonstrates a working knowledge of the domain and an ability to provide guidance. Equally important, it shows a healthy respect for the people's intellect and a flattering assumption that they will be able to read and digest the material. One of two things will happen. Either they will actually read the materials sent to them, in which case they have not only made an initial investment in seeing a technology transfer happen but have also been provided with great topics for further conversation, or they will not read the material and most likely conclude that it is simply easier to rely on the researcher as their expert rather than to become experts themselves. Either way, it is a foot in the door. They will ask endless questions as they try to decide for themselves what is within the realm of possibility and, more importantly, practicality. It is essential to ask them just as many questions to understand as completely as possible their constraints and to make clear recommendations on what they can expect to build.

Step 2 is to move from simply giving domain guidance on the state of the art to provide feedback on product-design plans. This involves offering to review specifications and providing timely feedback. Timeliness is critical—schedules are the rules of the game, and an ability to stay within their stated schedule reflects an understanding of the rules, an appreciation for their importance, and a commitment to the success of their project. This is also a critical test of a researcher's ability to think practically; in their distrust, they might expect suggestions of wildly impractical things that would have a negative impact on their schedule or require resources out of proportion the relative importance of the technology to the overall product. It is the researcher's job to show once again an understanding of their constraints and the value added to their team effort. Success will be apparent when a subtle but important shift happens: instead of the researcher asking to review their design documents, they will ask the researcher to review them.

Step 3 is a significant one indeed: when the clients ask the researcher to help write the specification for the product. Do not expect this to happen until there has been a clear success at Step 2 that has established credibility. But when it does happen, the product organization is making a loud and unambiguous statement that they now think of the researcher as part of the product team. This is an enormous step for a product group to take in their relationship, and it is a heavy responsibility to take on. At first, they will probably only delegate small parts, and often they might ask the researcher to co-write design documents. But regardless of the size of the assignment, the key to success is the same: whatever is designed must be easily buildable and testable. If there is any significant disagreement on whether the design can be built or tested, the product organization will not take the risk. Development organizations (at least the successful ones) are inherently conservative and will overstate the costs to build new technologies.* This is not only another test of whether the researcher understands their constraints but also equally whether he or she understands their development process. I encourage "aiming low" initially and looking for indications from the team that they would like to work together to design something more aggressive. If there is success in co-designing components, they will loosen the reins and delegate more responsibility (with more autonomy).

Step 4 is where one (finally) gets involved in implementation. It has taken enormous patience on the part of the researcher to get here, and there are still landmines everywhere. No two development organizations are alike; they all have different practices for creating, documenting, integrating, accepting, testing, and deploying new products. It is impossible to understand all of their processes, and it is

* Ironically, it has been my experience that development organizations tend to estimate incorrectly not in the new technologies, but rather in the incremental improvements to legacy components, and particularly in the "integration" work in making multiple components work together.

very unlikely that they will all be written down; yet, every one of them has an opportunity to break form and cause a rift. I would encourage asking the group manager how they bring a new employee into the group and what training and mentoring that person would go through; further, see if there are opportunities to take advantage of such a process to help to get up to speed. If one has made it this far, the product team wants to see a success as much as the researcher does. Because one's success is the other's success, the product team will be very reasonable about doing things to help themselves be understood by the researcher, especially if it is clear to them that that is the goal. All development groups fall in love with their own processes, and one can earn their cooperation by showing equal respect for those processes, no matter how silly they might seem to an outsider.

The key to success in a development process is to realistically promise and over deliver. Set rational expectations; they should think that you can carry your own weight, but not that you are God's gift to engineering. Be honest about the readiness of your technology in as crisp terms as possible; Speser (2006) suggests terminology for Technology Readiness Levels that represent how far away from market introduction the technology is currently. Promise metrics for work and for the technologies created that are achievable but not overly aggressive, and then exceed those metrics. By doing that, it is possible to fully gain their trust and move on to discuss with them more aggressive technology transfers.

It is important to note that in this "pyramid" of sorts that we are building with increasing levels of risk and corresponding trust, it is possible to peak at any level. For instance, if the researcher does not have the development skills to co-develop components with a real product team, then do not try to do it! By all measures, every step in this process can be considered technology transfer. Product organizations need knowledge, understanding, and ideas about technology just as much as they need finished technology components; they need to understand what cannot be built just as much as what can. And, most importantly, they need researchers to tell them honestly what they are and are not capable of doing for them. Even without ever delivering a finished component to a product organization, one can still have a litany of technology-transfer successes for which the product organization will sing praises. It is more important to proceed in measured steps built upon past successes and build the trust and the lines of communication that will guarantee future successes.

59.6 THINKING FOR THE FUTURE

It is also important to be thinking to the future—to be thinking about what comes next. There is always a desire to simply throw a technology over the wall and then to move on to the next research project, but this is unrealistic. It never really works that way, and even if it did, "throwing it over the wall" would end the relationship and any opportunities for future technology transfer; the ongoing relationship after the transfer is an opportunity to carry on a dialogue about the next great technology.

From the product–organization perspective, there is rarely a "clean" way to integrate a component. The overwhelming majority of development work is revisions to existing products; very rarely are new products started. Revision work means that new components need to be integrated into an existing legacy framework; this usually requires development work on both sides of the integration to ensure the optimal match.

59.7 CONSIDERING HCI IN THE TECH TRANSFER PROCESS

As if this was not difficult enough, applying this approach to HCI-related technology transfer introduces its own challenges. One can read The Psychology of Human–Computer Interaction (Card, Moran, and Newell 1983) and learn that at a very fundamental level, a set of scientific principles holds very broadly. However, we in the HCI community have also learned that interactive systems must be designed within the context of a particular task and human and that this very fact makes it tricky at best, and misleading or impossible at worst, to try to generalize specific designs to other contexts. Even with the best of intentions and the most thorough usability testing, there are no clear guidelines about how much of an HCI-related research technology can actually be transferred and particularly for integration into an existing product. So "throwing over the wall" is especially difficult, as it calls for potential redesign, as well as further development and integration, test, localization, support, and operations (it is increasingly a service world, after all).

In the traditional view of technology transfer, lack of a clean handoff would be fatal. In the "relationship" view, however, this is an opportunity to build a working model that lets you overcome the challenges and work side-by-side with a product organization to guide the transfer of your work.

Beyond simply moving up the pyramid, one can do other important things to deepen the relationship. Scheduling regular "maintenance" conversations can help to maintain communication channels and to keep abreast of activity in the product organization; it is also an opportunity to continue to update them on progress on new work. The relationship can also be used to improve the researcher's own work by learning about critical real-world issues. Good product organizations have a wealth of information about their customers; by using their access to real customers (and aggregate information about them), the relevance of research activities can be improved. It is an opportunity to ask key questions, learn about critical product strategic direction, conduct user studies, and find out what difficult HCI and technical issues are about to become critical issues for real customers. This is the golden opportunity to get out of the ivory tower.

As an aside, it is worth pointing out that HCI has its own value-add in the tech transfer process helping to quantify improvements. Often it is important, when there are competing technologies, to demonstrate the superiority of your technology. HCI's processes to measure quantitative differences in ease of use and "time on task" can be very helpful in these cases.

59.8 EXTERNAL TECHNOLOGY TRANSFER

Technology transfers outside the boundaries of a legal entity, regardless of whether they originate in a government agency, academia, or an industrial lab, are almost by definition cleaner types of transfers. This means that one will be participating in a transaction involving the sale or licensing of technology, or contracting to provide some service, or both. So, the first order of business is negotiating the "deal."

59.9 KNOW WHAT YOU ARE SELLING

It is critical to know and understand what is being sold: outright ownership of intellectual property or a license to it? Is it a complete solution that has been developed or just pieces? Those pieces might include any or all of the following:

- User interface design
- System specifications
- A working prototype
- An actual implementation, tested to some level of quality
- The source code for a software implementation
- Copyrights
- Patents
- Working time as a commitment to support ongoing productization
- A running service that you host

59.10 KNOW WHAT YOU ARE NOT SELLING

It is equally important to know what is not being sold. If one would like to continue this work, he or she will need to make sure to preserve rights and ownership to continue that work. Otherwise, one could very well put oneself out of business by selling complete ownership to a valuable asset or by signing an exclusive license, which precludes licenses with any other company.

59.11 KNOW THE PEOPLE INVOLVED

Deals inevitably involve lawyers as well as what are known as "business development" people: those whose job are to negotiate deals that further the business interests of their employers. One can assume that the company negotiating will have both business development people and lawyers; it makes good sense for the researcher to have them too. This is a situation where we need to put our pride and high-minded notions about doing business "on a handshake" aside. While the overwhelming majority of companies are not in the business of stealing from people like us, and will not try to do so, to get the most value in return for what is being offered one needs someone on his or her side who understands what is customary in intellectual property deals and how to negotiate for it. Even a brilliant and excellent debater who does not know what to reasonably ask for is at a serious disadvantage. The bottom line: find someone with good business development skills and experience to negotiate the deal.

Likewise, once the terms of a deal have been negotiated, one needs a lawyer to write it up and make it legally binding. Do not even think about self-representation in the drafting of an intellectual property agreement; the laws are changing too quickly (which is not the fault of the lawyers) for a researcher to understand which ones apply to the situation and should be factored into the drafting of an agreement. Mock, Kenkeremath, and Janis (1993) set out the basics of existing laws and how they relate to technology transfer, though the details have changed substantially since then. Many good books exist to help one to get educated on current intellectual property laws, although that is a poor substitute for a competent attorney skilled in the current practice. The Association of University Technology Managers (http://www.autm.net) also provides a wealth of resources to its members on a number of issues related to technology transfer and intellectual property.

The key to success is to understand the defined role of each of the three people on a negotiating team: oneself, the business-development person, and the lawyer. The researcher's role is to be the technical expert and to place a value on the work as well as what the people on the other side of the table are offering in return; one's role as "client" is to decide what is needed to be successful and what additionally is desired but negotiable. The role of the business development person is to take the articulated needs and desires and try to structure the terms of an agreement that will work for both parties. He or she understands business risks and will help the team members to understand them and make informed decisions about how much risk can be tolerated. The role of the lawyer is to take the terms and write them down in words that both parties understand and that can be interpreted under the law to protect the client's interests. The lawyer also understands and can articulate the legal risks; laws are often subtle and ambiguous things that can be interpreted in many ways (in fact, nations have an entire branch of government that does nothing but interpret the laws). Any contractual obligation runs the risk of being interpreted in a way other than how it was intended, and a lawyer can help the team members to understand how likely that is, based on the language of the law and similar previous cases where the law has been interpreted by the courts. Just as there is always business risk, there is always legal risk, and in the end, it will fall upon the researcher to make the decision as to whether the risk is acceptable.

It is also important to understand who plays these roles on the "buying" team. Speser argues that the most important role to understand (and in particular who is playing that role) is that of the decision maker, who at the end of the day will actually decide whether to go ahead and acquire a technology. It most likely will not be someone in the legal department or someone in their internal research department, who might even see you as unwanted competition. It most likely will be the person in the development team who actually plans to use the technology, as that person is the one who will need to find the money to purchase it as a part of running their business.

Where do deals go wrong? In my personal experience, they often go wrong when these three roles become confused and when the business development person and the lawyer start to make the key decisions. The researcher must live with the result of the deal, not them. On one end of the spectrum, there is no such thing as a risk-free deal; on the other end, even a high-risk deal could be worth doing if the reward is also high enough. Those decisions are the client's, not theirs, and the researcher as client should insist on making them.

There are many good sources for both business development people and lawyers. Venture capitalists will often have a "short list" of ones they trust to do their business transactions. The Chamber of Commerce for a local area can also provide recommendations and often will track complaints registered against specific ones. Many regions also have associations of entrepreneurs, inventors, and small-business owners, with great resources to draw from.

Of course, if one already works in a research lab, the institution most likely has a technology-licensing office that will negotiate and draft deals on the researcher's behalf (and are likely required to do so if the researcher's employment agreement assigns ownership of inventions to the employer). In that case, one will still need to stay involved to make sure that the researcher's needs are met in whatever deal is negotiated.

59.12 CRAFTING A DEAL

The most difficult process of negotiating a deal is crafting an arrangement that meets the needs of both sides. I have seen that many negotiations take much longer to conclude, and in many cases fail to conclude successfully, because either or both sides did not bother to try to understand the other side's business needs. Business partnerships are always about finding a way to help both parties be more successful. Speser goes as far as to suggest that deal making is the search for a Nash equilibrium whereby all parties have more incentive to stay with the deal than to change their tactics. The best way to do that is to understand what one's prospective partner's business is about and likewise to share enough information about one's own business, openly and honestly, so that together a combination can be found that works for both sides. Find out everything available about a partner's current business situation:

- Their revenues and profits
- Their competition
- Their most important customers
- What customers are saying about their product
- Where they say they want to take their business in the future
- The problems and challenges they are facing

This essential information will guide you to deeper insights on how to offer terms that will be seen as valuable to a potential buyer or licensee. Steinberg (1998) described his experiences in negotiating deals and his own well-known and well-respected philosophy for how to structure deals makes good business sense for all parties.

Companies will pay for the value that can be delivered to them. They will pay a certain amount of money to make an even larger reduction in their costs (because in the end they save money). They will of course pay to help themselves make even more money. Finally, they will pay if one solves a problem for them. As a deal is structured, try to cast it in terms of what it does for them; those are terms that they can understand and, more importantly, quantify in a valuation.

Take it further. Help the buyers in any way possible to place a value on what is brought to the table. For example, conduct a user study on their existing product and another showing how the technology being offered will improve their product (if it happens to address a key customer complaint, all the better, and that should definitely be brought to their attention). During the negotiations, show them a smart, effective professional who can work with them. This does not mean that you need to negotiate hard to the very last item; contrary to popular belief, the tough negotiators are not always the most respected, and in fact, they are often the ones that create their own reputation for being difficult to work with. It is much more important to demonstrate an understanding of the buyers and an ability to speak their language, as well as willingness to help them make the case to the decision makers in their organization's senior management.

59.13 CONFIDENTIALITY AND NDAs

One large challenge in negotiating deals is the issue of confidentiality, which usually rears its head first in the often-dreaded nondisclosure agreement. Nondisclosure agreements are signed before revealing confidential information to ensure that the information would not be disclosed to competitors. That part is a good business and a natural, noncontroversial part of good-faith bargaining—we all need to be able to keep secrets. The difficult part of nondisclosures is the issue of residuals: by looking at confidential information, I learn things and then I carry that learning around in my head for the rest of my career. What am I allowed to do with that information in my head, and who in fact owns it? From the discloser's point of view, one wants to make sure that someone cannot use his or her own confidential information to compete. From the other side, it is impossible to know exactly what will be disclosed, or what business opportunities are going to come one's way tomorrow, so it is deeply problematic to sign away the ability to enter certain businesses simply for the privilege of looking at confidential information. Both sides sound very reasonable, and they are, which is why NDAs are no trivial matter and often become the stopping point in negotiations.

Whenever possible, try to complete as much of the negotiations as possible without entering into an NDA, because it simplifies matters and prevents the trust issues from rising to the surface too early. The downside is that this makes

the early negotiations a precarious dance, where a researcher needs to show the other party enough to convince them that the technology is real and solves their problem, without giving away key secrets. There are things that still can be done: tell them "what" it does, instead of "how" it does it, and show them the system working. The goal is to make them crave it enough that they will want to sign an NDA under the researcher's terms to complete the technical due-diligence required for them to close the deal.

It is critical to think this all through before getting to the negotiation table—these are never decisions to make under time and social pressure. It is also critical to realize that the whole issue of confidentiality and NDAs is one more business risk; admit it and decide for yourself whether (and when) the potential reward outweighs the risk. This is one of the clear cases in which a lawyer will be extremely conservative and protective and describe in great detail everything that could be lost by entering into an NDA (or by showing technology without an NDA). But, in the end, the decision is the researcher's.

59.14 GOING TO THE NEGOTIATING TABLE

Negotiating a deal is probably the most hyped and feared part of this whole process; perhaps, we have all had too many nightmares about slick car salesmen tricking us into paying too much for too little in return. The reason we fear car salesmen is that the salesman has all the information about what the car is really worth and shares none of it with us; we are forced to blindly trust him, and many of us do not.

Steinberg (1998) once again shared his wisdom on a sound and ethical approach to negotiation in his "twelve essential rules for negotiating." Not nearly as ambitious as Steinberg, I have only three basic rules for negotiating:

1. I obtain as much information about each side's position as possible before arriving at the table.
2. I have a list of what I really need to succeed and a separate list of what I want to have in addition. I hold firm on my needs, and I am willing to compromise on my wants.
3. I always negotiate a deal in which both sides win.

Understanding a potential business partner's position is critical to negotiating success for a number of reasons. First, it tells what they are looking for. Ask the same list of questions we discussed with internal technology transfer: Who are their customers? What are those customers saying about their products? Who are their competitors? What is the company looking for that will help them to be more successful? What are their strengths and weaknesses? Second, and very much related, it tells how the buyer will value what they are offered. Successful business deals involve an exchange of value that both sides view as fair and equitable, but value is of course relative to the company and its context. Understand that in order to find an equal trade. Reading the company's annual report is an excellent source of information about a company (if it is a publicly traded company). Reading news articles and competitive reviews also provides invaluable information about the business pressures the company is under, as well as the assets that they bring to the table.

Having the list of the things that one really needs and the things that one wants in addition is a valuable step in preparing for negotiation. I have seen many people come to the table unwilling to compromise on anything; they ask for too much in the beginning and believe strongly that compromising on anything is a sign of weakness. Negotiations like those always take longer than they should and are very frustrating. In some cases, they spend more money in lawyers' fees for fruitless negotiations than the value of the small items on which they refuse to compromise. I recommend starting with a basic negotiation on the core needs of both sides; not only does that keep you focused on the heart of the deal but it also tends to simplify things just by taking all the peripheral items off the table. Once the heart of the deal is done, and both sides feel comfortable that they can be successful because they are getting what they need, adding additional pieces is much easier with a lower stress level and a structure in place.

Remember to be honest with yourself as you detail all this information, because lying to yourself is the surest path to failure. Assume that they will have an accurate valuation of what is brought to the table (regardless of whether they are willing to tell you what it is), so insisting that something is worth more than its real value is foolish. Understand one's own strengths and weaknesses, be honest about what is needed to be successful, and do not promise things that cannot be delivered.

All this brings us to the last rule: always negotiate a win–win. Always negotiate a deal in which both sides feel that they are receiving what they needed and can be successful. Beyond the obvious ethical reasons for doing this, there is also the very practical consideration that the two parties will need to continue to work together. Many people go to the negotiating table believing that signing a deal is the end of the process, when in fact it is just the beginning. Once the paperwork is signed, a relationship that essentially lasts forever begins between the parties. This relationship often makes itself known in unanticipated ways; for example, if one is licensing a patent to a company, the company has a vested interest in ensuring that the patent maintenance fees continue to be paid to the PTO to keep the patent valid and will want to have regular information to confirm that the payments are being made. Almost every deal one can imagine, no matter how cut-and-dry, has some aspect that will require communication between the parties on an ongoing basis after the deal is done. If one licenses a technology to a company, the company may be sending royalty checks, and the licensor may want some way to audit their sales to ensure that they are accurately paying. They, in turn, may want technical support, including important

bug fixes and updates. They may additionally have negotiated the option to license future upgrades, and as a paying customer, they will likely want to provide their input on features and enhancements to the technology that would be of most help to them.

59.15 EMBRACING THE RELATIONSHIP

By admitting from the beginning that there is an ongoing relationship, one can embrace this notion and turn it to an advantage. In fact, I encourage its use to build a future revenue stream—and in light of the increasing attention to Chesbrough's "open innovation" model, this is a likely outcome indeed. A researcher can build design and consulting services into the deal, which is particularly helpful for HCI-related technology transfer since, as we discussed earlier, they often need reworking to fit into a larger context. One can use the ongoing relationship as a "foot in the door" to be able to offer future sales and deals as new technologies are developed. In fact, viewed as a "strategic partner," one might even want to offer them the right of first refusal on future offerings, as a way of demonstrating a commitment to them. Finally, as in contemplating future growth of the researcher's business, one may need additional sources of funding, and a partner who has a vested interest in the researcher's success can be a great source for funding. Even if none of this is true, assume that one day there will be a need for a good reference or recommendation from them; that alone is reason enough to want to have a great ongoing relationship.

59.16 COMMONALITIES FOR INTERNAL AND EXTERNAL TECHNOLOGY TRANSFER

59.16.1 Intellectual Property

One of the issues common to both internal and external technology transfer is intellectual property (IP). I am not a lawyer, and so I obviously cannot give legal advice on how to protect intellectual property or how to treat others' IP. What I can do is point out some places where the IP issues get thorny and make some business recommendations about how to deal with them.

Anything received from a third party may carry restrictions on how it may be used and, more importantly, whether it can be redistributed in its original or modified form or combined with some other components. This includes libraries of software routines, data, copyrighted works and designs, and patents. These restrictions can come from explicit license agreements that accompany the third-party components, or they could come from any of a number of different laws, including patent, copyright, trademark, trade secret, and export. Any time a third-party component is used in one's work, a business risk of constraining the ability to transfer the work to another party arises, because either one does not possess the right to do so or the rights that are possessed are not sufficient to the needs of the party that wants to license it. This question fits into the larger scheme of what it traditionally called the "build or buy" decision: whether it makes more business sense to build something oneself or to buy or license it from a third party.*

I strongly recommend that, whenever possible, the IP issues should be dealt with at the time a third-party component is acquired, rather than waiting until an opportunity to transfer it. This accomplishes two things. First, it allows one to negotiate and make business decisions about acquiring the component before there is a commitment and dependence on the component built into your technology; once the dependency is there, the "switching cost" is much higher for moving to an alternate and one could be forced to pay a much larger licensing fee than before. Second, it simplifies the tech-transfer process. Any company worth its salt will perform a "due diligence" on the technology before it closes a licensing deal. Part of that will be an analysis of who really owns the technology or whether the licensor has acquired the right to further license it. In essence, one will need to prove the right to license one's work to the company. And do not be surprised if the company also asks the researcher to "warrant" the work—to guarantee that he or she has the right to license all of it to the company and that right will be defended in court if necessary. The bottom line: clear it up front, and save a lot of trouble later.

59.17 IT ALL COMES DOWN TO THE RELATIONSHIP

The most significant common aspect of internal and external transfer comes back to the notion that we began with: technology transfer is a social process that succeeds or fails based on the relationships that have been built. Tech-transfer partners need to trust each other and that trust is built with communication and follow-through. Researchers should understand clearly and honestly what value they bring to the table at the various stages of the relationship, and they should make it their business to know how their partners see and value what they bring. After all, business is fundamentally the exchange of value between partners who need each other's competencies, and the truly successful companies are the ones that build relationships that last across a continuous series of business transactions.

When beginning a technology-transfer effort, assume that what is started that day is the beginning of a working relationship that will last forever. Build the relationship from the ground up with the expectation of ever increasing levels of cooperation and trust that will allow the partnership to take on ever more challenging technology transfers, in whatever manner is most appropriate to the needs of the business.

There are far too many stories of companies who have struggled with the transfer of technologies that could change

* It is important to point out, though, that building from scratch does not necessarily mean an automatic escape from third parties; for example, one can still violate someone else's patent even with code written from scratch.

the world—and failed. As Buderi (2000) and Freidman (2005) described, the next chapter of this story is being written now. By rewriting the rules to focus on the social side of the process, we can ensure that our best work will see the light of day and this story will have a happy ending. The good news is that there are more companies open to technology transfer today than at any time in the history of the industry, and the opportunities are there—on both sides—for us to take.

REFERENCES

Buderi, R. 2000. *Engines of Tomorrow: How the World's Best Companies are Using their Research Labs to Win the Future.* New York: Simon & Schuster.

Butler, K. 1990. Collaboration for technology transfer—or "How do so many promising ideas get lost?" In *Proceedings of the CHI '90 Conference on Human Factors in Computing* Systems, 349–51. New York: ACM.

Card, S. K., T. P. Moran, and A. Newell. 1983. *The Psychology of Human-Computer Interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Chesbrough, H. 2003. *Open Innovation.* Boston, MA: Harvard Business School Press.

Foley, J. 1996. Technology transfer from university to industry. *Commun ACM* 39(9):30–1.

Friedman, T. 2005. *The World is Flat.* New York: Farrar, Straus, and Giroux.

Hiltzik, M. 1999. *Dealers of Lightning.* New York: Harper Collins Publishers.

Isaacs, E. A., and J. C. Tang. 1996. Technology transfer: So much research, so few good products. In *Proceedings of the CHI '96 Conference on Human Factors in Computing Systems*, (2), 155–6. New York: ACM.

Lesko, J., P. Nicolai, and M. Steve. 1998. *Technology Exchange in the Information Age.* Columbus, OH: Battelle Press.

Manning, G. K., ed. 1974. *Technology Transfer: Successes and Failures.* San Francisco, CA: San Francisco Press.

Mock, J. E., D. C. Kenkeremath, and F. T. Janis. 1993. *Moving R&D to the Marketplace: A Guidebook for Technology Transfer Managers.* Falls Church, VA: Technology Prospects, Inc.

National Research Council. 2003. *Innovation in Information Technology.* Washington, DC: National Academies Press.

Singh, G. 1993. From research prototypes to usable, useful systems: Lessons learned in the trenches. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 139–43. New York: ACM.

Speser, P. 2006. *The Art & Science of Technology Transfer.* Hoboken, NJ: John Wiley & Sons.

Steinberg, L. 1998. *Winning with Integrity*. New York: Villard Books.