A Survey of Software Related Academic Collaborations at Siemens

Matthew Bass
Siemens Corporate Research, Inc
755 College Road East
Princeton, NJ 08540
+1 609 734 3567
Matthew.Bass@Siemens.com

ABSTRACT

As an employer of more than 30,000 software engineers, Siemens has a large financial stake in their software engineering practices. The direct influence that academia has on the software practices within Siemens, however, is relatively small. This paper looks at some of the industrial factors that come into play in industrial/academic collaborations, identifies the collaborative models used, and looks at lessons learned from these efforts.

Categories and Subject Descriptors

D.2.0 [**Software Engineering**]: General

General Terms

Keywords

Technology Transition, Industry Collaboration

1. INTRODUCTION

Siemens employs over 30,000 software engineers world wide and produces large complex software intensive systems for many industries including power generation, automotive, trains, telecommunications, and building automation. Often there are large complex software components to these systems which are developed by large and often geographically distributed teams from a variety of backgrounds. The practices and technologies that govern the development of these systems come from many different sources, but a proportionally small number of them are from academia.

Siemens does engage the academic research community on a variety of topics including software engineering, and while it is not difficult to trace the origins of many of the product innovations to such relationships, it is quite difficult to trace many of our daily software practices or technologies to these collaborations. There are many reasons for this including misaligned expectations, incompatibility with the industrial funding model, and misalignment with individual's goals and objectives. In this paper we will look at the hurdles that are inherent in such collaborations, look at a number of collaborative models within Siemens and explore the lessons learned from these collaborations.

TT'06, May 22, 2006, Shanghai, China.

2. INDUSTRIAL MOTIVATIONS

There are two primary factors that may explain some of these difficulties. The first is the strong profit motive of the stakeholders in large publicly held companies, and the other is the advancement practices within Siemens. These two factors seem to influence the focus of employees at all levels. In this section we will look at the ways in which this is true.

2.1 Profit Motive

It is common knowledge that the fundamental motives for academia and industry are different. Everyone knows that for-profit organizations are by definition motivated by profit. It may not be, however, quite so obvious how this motive is realized in publicly held companies (or at least in Siemens) and what the implications of this are. In this section we will look at the sources of the profit motive and how it may influence the activities of employees at various levels.

Certain financial data for publicly held companies is widely available. Analysts look at factors such as market share, profit margin, and investments into R&D and compare them with the other companies. Based on these factors they rate the organization according to its attractiveness as an investment. These recommendations influence the choices of not just private investors, but also fund managers. The stock price (and resulting value) of an organization is directly impacted by these choices. This means that there is an incredible competition to show increased value for R&D investment compared to the competition (value is determined by the profit that is generated as a result of the R&D investment).

The results mentioned above are viewed on a regular periodic basis (quarterly or annually) and typically directly impact the evaluations of the top level management. As a result of this there is an incentive to focus R&D spending on activities that will positively impact the bottom line directly and reduce spending other areas. These incentives trickle down to employees at all levels. The results are often obvious in policies and programs designed to reduce expenses, reduction of staff, and annual incentives tied directly to things like developing products within budget and schedule.

In addition to these direct incentives there are often mandated programs that may incur costs in the short term and not realize returns for some time (e.g. achieving a particular CMMI level). These programs are often viewed as being harmful to short term reviews and thus are not implemented with as much commitment as one would hope. The same may be true when more centralized strategic efforts (e.g. introducing product line practices) are proposed.

In the case of many large companies including Siemens there are often central research groups. The central research organization within Siemens has a software engineering department focused on software engineering issues. The pressure on these groups has increased as well. It used to be the case that we (the author is part of this department) were funded 100% by a central fund and so was free to conduct research largely as they saw fit. Today, however, we rely on direct funding from the operating companies and so must offer services that are of direct interest to them. This means that we are directly influenced by the same constraints that exist within the development groups.

For groups that are not making their target numbers, it is not difficult to see how these factors lead them to ignore anything that will not directly help their short term issues, or lead to increased sales. In fact it is possible for a career to be cut short if a program is put in place that hurts the short term numbers even if the long term prognosis is good. Groups that don't have problems making their numbers tend not to have the motivation to look for real help with their software engineering practices. They do invest heavily in innovation related R&D, however.

2.2 Advancement Practices

Siemens has a strong record of aggressively promoting from within. The expectation is set early that, successful individuals will change positions every 4 or 5 years. Ideally these positions will allow for a diverse set of experiences from a domain, responsibility, and cultural perspective. The result of this combined with the pressures stated above is that career minded individuals seek results quickly. If it takes three of four years to see noticeable positive results, then the next promotion could be in jeopardy.

3. SURVEY OF COLLABORATIONS

Despite all that was mentioned above, Siemens has in the past and continues to collaborate with academic researchers focused on software engineering topics. In this section I will outline several models of collaboration and describe the typical results of these models. It should be noted that virtually all of the academic collaborations involving software engineering are initiated by individuals from the central research group. It can be assumed in all the models below that unless explicitly noted all company participants are from the central research group.

3.1 Throw Money over the Wall

One approach that is actually not that uncommon is to simply give money to researchers working on problems of interest. This usually is done because a portion of the central research group is engaged in or would like to be engaged in consulting activities in a particular topic area and having a "collaboration" with well known researchers in this area can help bolster an internal reputation. The concrete outcomes from the industrial side are usually in the form of having some access to the researcher, sometimes the researcher will be invited to speak at internal Siemens events, the money will often fund a graduate student, and sometimes there are jointly written papers.

From the researcher's side there are limited benefits of this type of collaboration beyond the funding. The researcher will have some access to the industrial sponsor, but typically there is not any industrial staff time allocated to such a project and so the access to real data or projects is very limited.

3.2 Personnel Exchange

One of the most common ways within Siemens of collaborating with academic institutions is to bring students onsite to either do their research or to work as an intern. While extremely rare, Siemens has on occasion sent people to work onsite at the university as well (to my knowledge, this has only happened with one university).

At any given time there are a reasonably large number of students onsite at the various corporate research facilities. These students are either working on their thesis or working as interns. The European (or at least German) education model allows for the student to spend their research time at an industrial institution (usually 6 months in the case of a Masters degree and 3 years for a PhD). The topic is jointly agreed upon by the academic institution as well as the sponsoring organization. In the case of software engineering topics, this work will sometimes find its way into the activities of the corporate research group.

From the student's perspective they get exposed to an industrial environment, this may help to motivate their research, and in some cases is the primary focus of their research. From Siemens perspective this is an opportunity to have resources that focus on activities that are not typically part of the day to day activities of the technical staff. This is also a good source of future employees. It is not uncommon for students to end up working at Siemens upon graduation.

In three or four cases Siemens has delegated employees to an academic institution. These delegations were all to an applied research institute. This institute had a structured program whereby industry and government agencies could allocate employees to work with the staff of the research institute. This was meant both as a means for transitioning practices and bringing current problems that may influence future work. From a transition perspective this can work well. Because this delegate is an employee of Siemens they end usually actively participate in development efforts and are in a position to bring the work of the research institute into the equation.

The work of this institute is very applied in nature, and it is usually the case that it has moved from the research stage more towards the transition stage by the time Siemens gets involved. This work has been in the area of project management practices, software processes, and software architecture practices. The difficulty with this model is that in order to begin to have a substantial impact on the way Siemens develops software these practices need to move beyond a single project and become institutionalized. This has happened in one area (with SW-CMM and now CMMI), but this is the only such example to date.

Depending on the interest and activities of the delegate, it is possible for the delegate to become aware of and involved with the other activities at the university. Some of the strategic projects mentioned below began that way.

3.3 Strategic Projects

There are several examples of larger more strategic projects that were designed to address specific areas of concern within Siemens. The first was the development and adoption of a real time framework to support the needs of a particular product line activity. The other two were studies and experiments in the area of geographically distributed software development.

In the case of the first project mentioned the collaboration involved three primary entities. There was a Siemens operating company that was building a system, Siemens central research group that was assisting in the design and construction of particular aspects of the system and the academic research partner that had a history of working in the real time middleware arena. This has been a long term relationship that has existed for over 10 years (the relationship is with a particular researcher and has followed him from one institution to another). As a result of this relationship, technology from the academic institution has found its way into a large and very successful product line at Siemens, students have been hired, the researcher has spoken at many Siemens events, and several other joint research projects have been executed (it is not clear what impact if any these other projects have had on Siemens development efforts).

In other case Siemens has initiated and sponsored a large test bed for geographically distributed software development. This project uses student teams from 6 universities world wide (reference the GSP paper) and has a team of researchers from both Siemens and academia collecting data and analyzing the result. It is too early to tell what the long term impact of this project will be (it is currently in its second year), but to date it has spawned papers, Siemens has hired involved students and some of the practices developed are beginning to find their way into ongoing development efforts.

Another effort related to the one above is a collaboration with a researcher in the area of geographically distributed software development (GSD). This researcher has (and is still doing) case studies on GSD projects within Siemens, is working with Siemens to develop a set of best practices in this area and is actively conducting a number of related research activities. It is again too early to tell what the long term impact of these activities will be, but to date both Siemens and the academic collaborator have been fully satisfied with the relationship and the progress.

4. LESSONS LEARNED

In the process of these collaborations many lessons have been learned. Below is a summarization of the primary lessons

Appropriate expectation setting: In order for both sides to have a rewarding collaboration it is important that they both have a reasonable set of expectations. It is not uncommon for industry to not adequately understand the motivations of academic researchers (and vise versa). There is often an expectation that if you (the company) give money for a graduate student (for example) that the graduate student works for you and should act largely as a consultant. As part of this is an understanding of the motives of the *individuals* involved. A high-level understanding of the industrial perspective is one thing, but an understanding of the issues and motives of the individuals involved is important to achieving success. It is not uncommon to hear an academic partner bemoan the actions of an individual thinking that they do not support the long terms objectives of the company when really the academic doesn’t understand the realities of that individual.

Usually a transition agent is needed. In other words the research isn’t ready for prime time (hence it is called research) and having someone that is familiar with both the needs of the developing organization and is familiar with the research can aid is bridging the gap that surely exists. Another way to say this is that there needs to be a champion within the company that is passionate about the research and works hard to make use of it in the organization.

Alignment with corporate initiatives is important for larger efforts. In order to overcome the resistance to dedicating time and resources to anything outside the immediate needs of an organization it helps to have high level support (or pressure). In the cases of all of the strategic projects within Siemens there was either an alignment with a corporate initiative (resulting in political capitol for anyone that participated) or it was supporting a very important project. There are enough corporate initiatives and other pressures that people just don’t have the time to focus on anything else.

Time not money is often the most precious resource of a development organization. This may be the most important factor to keep in mind, time is at least (and often more) as precious as money. There is often a perception that if the researcher offers their services for free in exchange for data or a project to use as context it should be a bargain for industry. This is not the case as these engagements usually will divert key people from their primary responsibilities. This is something that industry can usually not afford. It is also the case that most times the work being done by researchers doesn't fit nicely into the industrial setting as is. If extensive training is required, infrastructure changed, particular languages used and so forth it may be quite a tough sell.