Post Mortem - An Assessment of Two Approaches

Tor Stålhane
Torgeir Dingsøyr
NTNU, Norway
Geir Kjetil Hanssen
Nils Brede Moe
SINTEF, Norway

Abstract

Learning from experience is the key to successes for all that develop software. Both the successes and the failures in software projects can help us to improve. Here we discuss two versions of Post Mortem Analysis (PMA) as methods for harvesting experience from completed software projects, which can be part of a larger knowledge management program. The two methods are tailored for use in small and medium size companies and are conceptually easy to apply. In addition, they require few resources compared to other methods in the field. We think that the methods are useful for companies when they need to document their knowledge, find improvement actions and as a start of systematic knowledge harvesting.

Introduction

An obvious way to improve a software development process is to learn from past mistakes. In practice this has turned out to be easier said than done, but the method of post mortem analysis - also called post mortem reviews or PMA - is one way to achieve it. There are several companies already doing this, such as Apple Computers, Rolls Royce and Microsoft. What they all have in common, however, is a rather elaborate and costly process.

We are mostly dealing with small companies, with small projects. The overhead from an approach as ambitious as the one used by, for instance, Apple Computers would be way beyond the reach of most of our customers. We needed something simpler and wanted to try out the two rather simple approaches described in this paper. In order to see if they were useful we tried them out in two companies. Our goals were to see if:

- The approaches worked - i.e. if anything useful came out of them.
- They worked in different ways - i.e. were the information that came out of them different in amount, scope or usefulness.

The research was conducted in close collaboration with industry. We collected experience

from real software projects in a real environment, over which we had only limited control. What we present here is thus not a scientific evaluation but our experience from practical improvement work. Thus, what comes out is not hard science in the strict sense but practical experience for practical people working in the area of software process improvement.

The rest of this paper is organised as follows: First we describe our approaches to PMA. Then we provide a short description of the two companies where we tested our approaches, and what we found out when we did this. Finally, we discuss the usefulness of the two PMA methods and the usefulness of what came out of them.

Harvesting experience

Learning from experience

To improve software development, it is important to help software developers to learn both from positive and negative experience, i.e. to become what is often termed a "learning organisation" . People usually learn best when they can relate what they learn to their ordinary work. Thus, experience collected from a well-known environment will have a greater learning effect than experience collected from unfamiliar environments.

Learning from your own experience requires that you be allowed to reason about tasks that have been completed to see what went well and what did not, and also why things happened the way they did. Thus, it requires observation and reflection .

To cover more material, we can learn from other people's experience as well as our own. Thus, it is a goal to transfer experience from one project to another in order to make individual and project-based learning helpful for, say, a whole organisation. An important part of the problem is how this experience can be collected effectively. We will now turn our attention to this question.

The post mortem analysis - PMA

As a concept, the PMA is simple - gather all participants from a project that is ongoing or just finished and ask them to identify which aspects of the project worked well and should be repeated, which worked badly and should be avoided, and what was merely "OK" but leave room for improvement. In addition, we need techniques that can be used to elicit experience and document it in such a way that they can be made verifiable and reusable in the most effective way. Reuse of experience is, after all, one of the reasons for doing a PMA.
A PMA can be done in many different ways - varying both in goals and degree of formality. A PMA can be focused on harvesting experience that is available from a single activity or process step or it can try to catch all experience available from a project. Data collection can be done as semi-structured interviews, or as a multi-step defined process, for instance consisting of a group process supported by affinity diagrams, followed by a root cause analysis (RCA), see , using an Ishikawa diagram.

Whatever form the PMA might take, one point is important - the consideration of the influence of the project's environment. No experience is reusable without considering these factors. It is often the case that an action undertaken in a project, which led to

unfavourable results, was not bad in itself - it was just not a smart thing to do in the current environment.

Given its loose form and lack of formalities, many people find it almost wrong to dignify PMA with the word "analysis", which they think should be reserved for something more scientific. The loose form and simple concept is, however, a strength and not a weakness. It gives us a method that is flexible, easy to learn and apply and containing a minimum of formalities to take care of. In this way everybody in the organisation can participate and not just as information providers. This is crucial in order to support the TQM philosophy - that quality and quality improvement is everybody's responsibility; it is not the domain of the QA department alone.

Thus, we have decided to use just a few simple techniques. One or more of these techniques have been used for all the PMAs that we report from in this paper:

- KJ, which is a structured brainstorming technique.
- The Ishikawa diagram – also known as the fishbone diagram.
- Structured interviews, where the structuring factors are:
  - The purpose of the PMA – for instance causes for underestimation of costs
  - The structure of the project – what did we do?
  - The structure of the process – how did we do it?

Spacetec	InfoStream
All project members participated	Only those who did the estimation participated
Interview with the PM to get background information (1)	Study project documents to get background information (a)
Introduce the concept of PMA to the participants (2, b)	
Perform a KJ (3)	Perform structured interviews of project participants (c)
Prioritise the items from the KJ (4)	
Prepare the KJ results for an RCA (5)	-
Perform an RCA by using the Ishikawa diagram (6)	Register cause-effect connections during the interviews (d)
Write final report (7, e)	
Present final report and get feedback from participants (8, f)	

Table 1 : The PMA processes used in the two companies.

The process used in the two projects differed, partly due to company constraints, partly due to each researcher's preferences. The way the PMAs were performed is summed up in table 1 above.

Presenting our results

When presenting our results we have followed the same format for both companies. We start with a short description of the company before describing how we performed the PMA and present some typical results. In addition to the PMAs we had some activities - called support activities -, which improve on the results from the PMAs by supplying some extra facts. In this way, we are not depending only on the participants' memory.
For Spacetec, we did the interviews before the PMA, while for InfoStream a GQM planning session plus data collection and analysis were done after the PMA. In fact, some

of it is still going on.

In order to present the results from the two companies in a uniform way, we have chosen to present the supporting activities after the examples for both companies. In this way, we keep a better focus on the main issues of this presentation, which are the PMAs.

PMA at Spacetec

Description

Spacetec AS is one of the leading producers of receiving stations for data from meteorological and earth observation satellites. They have developed a considerable expertise in delivering turnkey ground station systems, consultancy, feasibility studies, system engineering, training and support.

Spacetec has as an overall goal to increase knowledge transfer between projects, and has chosen estimation of software project costs as its first focus area. The company will build a knowledge repository with cost information on previous projects. One of the goals of the internal improvement project at Spacetec is to improve the estimation accuracy of technical work packages so that the deviation between real and estimated cost will be within 20%.

What we did

Two PMAs were carried out for two typical projects. The first projects had lasted for almost a year, and was 75% finished when we did the analysis. The second project was finished less than a year ago. Both projects built software for analysis of image data from satellites. The PMAs were performed as described in the Spacetec-column in table 1.

Each PMA lasted approximately four hours and we had a feedback session for one hour the following day. Steps 2 and 4 lasted only for one hour each since they were highly creative. In step 3 – see table 1 – we gave each participant a number of post-it notes and told them to write down an item that was either a problem or a success on each post-it note. We then gathered in a meeting room and asked each participant to attach each post-it note on the whiteboard. In addition, they should explain to the other participants what the issue was and why it was important. Post-it notes that were related should be placed close to each other. We used a tape recorder during steps 2 - 4, and the results were transcribed and later inserted into the report. Giving the participants quick feedback was of great value both for the participants and for the researcher. The transcribed material, in particular, created important discussions.

We wrote a post mortem report on the project, containing an introduction which described the process, a short description of the project that we analysed, how the analysis was carried out, and the results of the analysis. The result was a prioritised list of problems and successes in the project. We used statements from the meeting to present what was said about the issues with highest priority, together with an Ishikawa diagram to show their root causes. In an appendix, we included everything that was written down on post-it notes during the KJ session, and a transcription of the presentation of the issues that were used on the post-it notes. In total, this report was about 15 pages long.

The total amount of resources needed for a PMA were three persons selected from the project participants. The total work needed was 35 person hours, distributed as follows:

- 4 person hours for the researchers' preparation.
- 20 person hours for the PMA session - researchers and project personnel.
- 6 person hours for writing the report from the interviews - researchers only.
- 5 person hours for feedback and updating - researchers and project personnel.

Some examples

One result from one of the KJ sessions was four post-it notes grouped together and named "preparatory work." They are shown in the lower left corner of the results from the KJ process – see figure 1. The arrows indicate relationships between the classes and tell which success factors that influence other success factors.

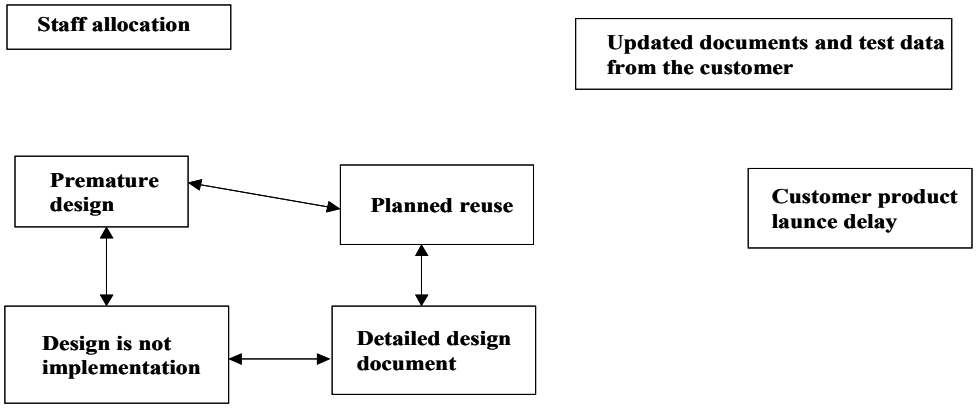

During the PMA meeting, the developers explained the problems related to "preparatory work" in the following way:

- *Design before "full insight". I dare say this as well: It is OK that we are supposed to reuse, but if we had known – now I was not participating in the design, so it is wrong of me to say it - but if we had full control we could have made an architecture that was in fact possible to use later on. ("Premature design")*
- *We put down a lot of effort into the design of the processor [a software module for analysing satellite data], and then it turned out when we were to implement it that we did not at all use the design that we had planned of in the document, so it was maybe a bit stupid to use a lot of time to… and then we don't use it afterwards. ("Design is not implementation")*

When we later tried to find the root causes for the "preparatory work" problem, we ended up with the following Ishikawa diagram:

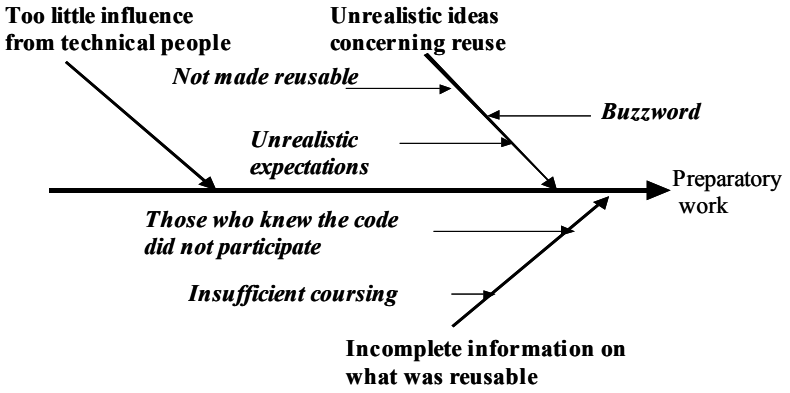

Fig. 2: Ishikawa diagram for "Preparatory work".

The root causes for this problem, as the developers saw it, was that people in the company in general had an idealised conception of reuse. Other root causes were that the company did not have a good enough overview of what was reusable from earlier projects and that the technical personnel did not have influence on the project activities where issues on reuse were discussed.

We then asked people to come up with suggestions for how to reduce this problem:

- Get an overview of reusable components in the company.
- Teach people higher up in the organisation about reusability.
- Update or write a new detailed design document for this project.

These suggestions on short- and long-term improvement actions were documented in the report. At the end of the meeting, we asked people how they felt about the results, and got feedback like: *The issues about contracts with the customer – I think we are too kind. There is so much that is mentioned in the contract that… it is not supposed to be like that.* In general, people said they were made aware of new issues, and were able to see smaller issues more in a larger context.

The next day, we went through the report from the meeting together with the people who attended, as well as the project manager. Some of the participants disliked that we gave a complete transcript of part of the meeting, but after some discussion, they agreed that this is something that could trigger new discussion on relevant topics. We got positive feedback on the way we organised the review, which differed from the project completion reports that the project manager would normally write alone:

One of the problems with the project completion report is that the project manager sits down in the end and sum up what went well and what went badly. And it is one man who remembers what […] has happened the last two years. Then you have lost a lot. You do not collect a lot of experience. It is the large problems that are already discussed – things that have gone really badly. But if you do this kind of an analysis several times during the project, then we agree – it shows how… you have to think things through!

Support activities

As a basis for the PMA analysis we first had interviews with another eight persons from the company (developers, project managers, management and marketing personnel). Focus in these interviews was the estimation process and how the projects

were accomplished compared to the initial plans.

We also conducted an analysis of the project completion report from 12 projects. This resulted in the following improvement suggestions:

- Introduce a 28% contingency budget on each work package to decrease the probability of cost over-runs.
- Do a focused PMA to examine why the work packages ”management” and ”QA” were strongly underestimated in one project, and why ”coding” and ”unit testing” are strongly underestimated in general.
- Do PMAs on all future projects in order to increase learning in general.
- Use a standardized structure of work packages to make projects more comparable – thus improving the possibility for comparative data analysis.

PMA at InfoStream

Description

InfoStream is a medium sized Internet consulting company that recently became part of the international Integra group. Most of the customers are in the financial, energy, media, telecom, and manufacturing sectors. InfoStream has no product line, but develops tailor made solutions for their customers.

Many of the projects of the recent past have missed on the estimates. There have been several delays in delivery and a tendency to overuse resources in the projects. This has serious economic consequences since the estimates are used to set a fixed price. Estimation precision is considered to be an important success factor, and the company is investigating their projects to find what causes these problems.

What we did

In order to perform PMAs in this company, we asked our contact person to select three typical projects that had been finished less than a year ago. This limit was imposed in order to make sure that the project experience is still available with a reasonable certainty. The three PMAs were focused on estimation and were performed as described in the InfoStream-column in table 1.

Each project PMA lasted approximately five hours, including a lunch break in the middle. During the interviews and discussions, we used a whiteboard and a flip-over to document the points as they surfaced. The flip-over was used as a collective short-term memory for the group. Except for the structure imposed by the project’s work breakdown structure and the focus provided by the goal of understanding the reasons for incorrect estimates, no further structure was imposed. Thus, the notes from each PMA differs somewhat when it comes to internal structure. For instance: for one project, the participants identified the most important cause for cost overruns for each activity, while this was not done for the other two.

When the interview was finished, the researchers took all the material back to their office and structured and wrote down all the points raised. The points where numbered so that cause-and-effect relationships could be documented. This was done by inserting statements like “The estimate was increased by 40 person hours. This was caused by

point 53". See also the chapter "Some Examples" below. The resulting report was sent to the participants for feedback.

The researchers went through all the three PMA summaries and extracted all the registered information that was related to the PMA focus - improving the estimation process. Afterwards we defined a slogan for each information point and wrote them down on post-it notes. This was used as an input to a KJ process undertaken by the two researchers in order to analyse the information. This way to use KJ is in line with the original way to use this method, see .

Figure 3 shows all the results from this KJ process. Each box contains several experiences but these are not shown in order to keep the diagram simple. The categories contained in each box are supposed to be the key success factors, and we believe that if the company paid enough attention to them we would not have experienced the poor quality of the estimates that we actually did.

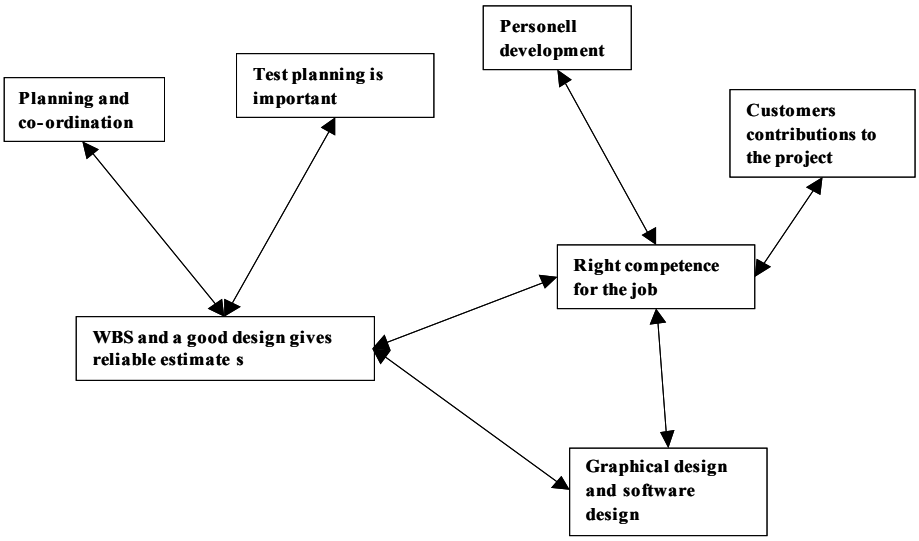

Fig. 3: Post-it notes showing the summary of problems related to estimation

The total effort for each PMA was 45 person hours or approximately one person-week per PMA. This should not frighten anyone if the work results in a substantial improvement in the way the company do their project estimates. The costs were distributed as follows:

- 5 person hours for the researchers' preparation.
- 20 person hours for the interviews - researchers and project personnel.
- 12 person hours for writing the report from the interviews and performing the final KJ - researchers only.
- 8 person hours for commenting and correcting the report - researchers and project personnel.

Some examples

The following are examples of the notes taken during two of the PMAs. All items are numbered so that it was possible to cross-reference them and thus to construct cause-and-effect chains.

PMA 1
1. The developers controlled the concept development themselves (because of 2).
2. The ideas for the product were too abstract for the customer (caused 1 and 4).
3. The developers got little feedback from the customers (because of 2).
4. The customers felt that they could change the concept later (caused 3).
5. The customer representatives were few and positive. These were the right people (caused 6 and 7).
6. The customer's participation was controlled by a subsidiary of InfoStream (because of 5).
7. The developers needed few meetings with the customers (because of 5).

PMA 2.
15. Too little time was allocated to testing.
16. InfoStream was under heavy pressure from their customer concerning time of delivery (caused 15. see also 10).
17. The customer was supposed to take full responsibility for the system's test.
18. More errors than expected – confirmed by the test reports.
19. InfoStream has later tried to improve their test planning for this customer (see 17).

Already during the first reading of the PMA reports, the participants were able to identify several improvement possibilities. The following is just a small sample:
- The developers must interact with just a few customer representatives - preferably only one. Otherwise, they will receive a large amount of uncoordinated, often conflicting comments on details in the system under development. To handle this consumes considerable - mostly unproductive - resources.
- Once set, it is next too impossible to change the customer's expectations. Thus, we must be careful when suggesting solutions.
- The graphical designers must co-operate with one or more software developers during the design of all web pages.
- The number of correction / feed back meetings with the customer must be fixed and stated before the estimates are made.

Support activities

Experience and experiments have shown that most persons have a tendency to be selective in what they remember. The reasons for this are legion and will not be discussed here. An interesting, albeit provocative, discussion on this topic can be found in . In order to have a broader basis for our final conclusion and recommendations, we decided early in the process to collect measurement data for a project that was about to start at that point in time. We ran a GQM (Goal Question Metrics) process in order to develop a measurement plan. Another reason for combining PMA with GQM is that from a data analysis point of view, PMA is mainly a bottom up technique while GQM is a technique that is mainly top-down. Thus, one should expect the two methods to give a more useful and complete picture of the estimation process than either one of them would achieve when applied alone. We discovered later that others have also arrived at this conclusion, albeit via another route - see .

Discussions

PMA experience

Our most important experience from doing PMAs is that the participants later told us that it was a highly positive experience. Since the concept of PMA is easy to grasp, everybody could participate. In addition, the developers all felt that they got something back – better insight, ways to improve their job and so on.

It is our experience that PMA is a method that is easy to use. There are no complex routines or tasks that have to be explained in advance, the method relies on the participant's intuitive understanding. A PMA will, however, create a large amount of information. Some of it is important and some is not. The job of separating the two can be both time consuming and challenging.

There are some ground rules that should be applied in order to get a successful PMA:

- The process needs to be structured – it is not a free-for-all happening. On the other hand, too much structuring – for instance extensive use of time boxing – was considered to be negative.
- The PMA will work better with an external facilitator since internal personnel may hesitate to bring up sensitive issues.
- If we are doing a general, catch-all PMA, people from all parts of the project need to participate. It is for instance not enough to have participants from those who did the coding if someone else did the design.
- If we are doing a focused PMA, the process needs to be steered – sometimes quite strongly - in order not to go astray and wind up as a general PMA or general lamenting over everything that is bad in the company.

The first decision needed if we want to do a PMA is whether it should be focused or not. Both focused and unfocused PMAs have their strong and weak points that should be considered before we make our choice.

- An unfocused PMA will usually give more surprises since it covers a broader area of experience. Thus, it is always a good idea to do a PMA as a two-stage process - first a general PMA in order to identify all the important issues and then a focused PMA for each issue afterwards.
- A focused PMA is best if we want to understand or improve a single activity. It should be done with just a few participants since the discussion must be steered in order to keep the focus. Few participants will favour structured interviews combined with discussion sessions. A focused PMA will in general require that the facilitator have in-depth knowledge of the field under discussion. This is consistent with the PROFES experience – see .

We started out by believing that PMAs would work best in a project-oriented organisation without a rigid control structure or an excessive hierarchy. We later learned, however, that the US Army – and maybe other military organisations – uses their own version of the PMA – called After Action Review (AAR) – with great success . These authors have a more narrow definition of a PMA than we do. In their opinion the AAR is not a PMA since it is done on a regular bases during the activity's lifetime. Anyhow, the important point is that it seems that the major factor is whether the organisation want to improve itself by learning from past mistakes, not how it is organised.

For small companies - less than five developers, say - it might be beneficial to let all employees participate actively in the PMA whether they were involved in the project or not. This will make everybody aware of quality problems and the ways to solve them.

It will thus strengthen the understanding that quality is everybody's business.

At least two problems are still open:

- The PMA way of learning from experience might not work well when we have several cultures present in the group. A case in point is a PMA where both software developers and graphical designers participate. On the other hand, divergent groups can also be the basis for constructive discussions.
- How can we perform PMAs when the participants are geographically distributed? This can be a problem today and will be even more so as we get more and more distributed or virtual organisations.

What did we learn

What has been the benefit for the two companies? What have they learned and what will they do different the next time? The feedback after the sessions tells us that the participants felt that they got useful information through the PMA sessions and that they discovered new aspects, or at least was able to express knowledge that earlier had only been tacit.

- *Documenting knowledge* - Documenting a company's knowledge is always a difficult task. By using PMA it became much easier for the companies to document knowledge. The resulting documentation also gave valuable insight into the development process. The knowledge harvested in the process came from several persons in the projects. This made the conclusions more deeply rooted in the organisation than for example an experience report.
- *Source for improvement actions* - The results of the PMA sessions lead to several improvement actions. In Spacetec for example, after discovering in a PMA session that the customer delayed the project, and did not provide test data according to the contract, they sent a Contract Change Note to the customer regarding the amount of hours to be used in the project. Since this method gives fast feedback it is possible to get some "quick wins". You can have a PMA meeting one day and a list of improvement possibilities the following day.
- *A system for harvesting experience* - PMA helped the companies in the process of developing their own system for harvesting experience. Because of the strong involvement from the developers, there was an increasing focus on "learning from experience" in the companies. The motivating factor of the PMA also makes it easier for the companies to implement a system for knowledge management. PMA will be one of the central techniques in such a system.

What would we do different the next time?

- *Preparation Work* – If using interviews, it is important to have good knowledge of the project in advance, in order to ask specific and relevant questions. In the KJ phase of the Post Mortem, it is important to be open and leave the word to the participants. Here, being able to moderate discussions is the main virtue.
- *PMA Process* – How the PMA is carried out depends on the number of participants, the complexity of the issues under investigation, and what the focus of the PMA is. Using the KJ method requires at least three participants in order to have enough material to discuss. If the issue under investigation is complex, it might be better to use interviews because it is more difficult to keep people on track using an open process such as the KJ. To get everyone to participate, the KJ method is working very well.
- *Participant Activation* – Some of the developers did not participate much in the

RCA phase. It could help people to be more active if one of the participants were moderating this session. Another possibility is to have participants working in groups first, and then present results in a plenary session afterwards.

- *Documentation Techniques* – The techniques we used, KJ and RCA together with detailed minute writing and transcription worked well. The results were easy to understand for people reading the report afterwards, and we got a lot of information in a short time.

Conclusions

We started out with two questions – did our approaches work and did they differ in what we got out of them?

First of all – both the approaches worked quite well. The approaches were conceptually simple and had a low cost (the participants did for instance not need to prepare them selves) but even so; quite a lot of useful information came out of them. In Spacetec the results from the sessions led to a Contract Change Note to the customer regarding the amount of hours to be used in the project and useful input to the experience database. In InfoStream the results gave important input to a set of checklists that were under development. Thus, at least for SMEs, we do not need the extensive approached used by for instance Microsoft or Apple Computers. The companies involved found it easy to identify improvement opportunities and the participants found the approaches easy to use.

The two approaches differed mainly with respect to the types of information and improvement opportunities identified. We found that semi structured interviews worked well for a focused PMA while the other approach quickly lost focus. The other, open approach – KJ plus RCA – worked well in a catch-all situation and gave more surprises. An optimum solution would be to start out with an open PMA to identify all the issues and then follow up with focused PMAs on the most important ones. This will bring us close to a simplified version of the approach used by Apple Computers.

Acknowledgement

The authors want to thank the Norwegian research Council that has supported this work through the PROFIT (PROcess Improvement For the IT industry) program and our colleagues at SINTEF and NTNU who, through discussions and comments, have provided many important inputs to this paper.

References

[1] Collier, B. DeMarco, Tom and Fearey, P. A Defined Process For Project Postmortem Review, IEEE Software, July 1996.

[2] Nolan, Andrew J, Learning form Success, IEEE Software, January / February 1999.

[3] Cusumano, Michael A and Selby, Richard W. Microsoft Secrets, The Free Press, 1995, ISBN 0-02-874048-3.

[4] Garvin, David. Building a Learning Organisation, Harvard Business Review, July / August, 1993

[5] Kolb, David. Management and the learning process, in How Organizations learn (Ed. Starkey, Ken), Thomson Business Press, London, 1996

[6] Straker, David. A Toolbook for Quality Improvement and Problem Solving, Prentice Hall International (UK) Limited, 1995

[7] Scupin, R. (1997) The KJ Method: A Technique for Analyzing Data Derived from Japanese Ethnology, Human Organization, Vol. 56, No. 2, pp. 233-237

[8] Jørgensen, Magne and Sjøberg, Dag. The Importance of NOT Learning from Experience. Proceedings of the EuroSPI 2000 Conference.

[9] van Solingen, Rini, and Berghout, Egon. The Goal/Question/Metric Method, McGraw Hill, 1999. Japanese Ethnology, Human Organization, Vol. 56, No. 2, pp. 233-237

[10] Mendoca, Manoel G. and Basili, Victor, Validation of an Approach for Improving Existing Measurement Frameworks, IEEE Transactions on Software Engineering, vol. 26,no. 6, June 2000.

[11] Oivo, Marrku, Andreas Birk, Seija Komi-Sirviø, Pasi Kuvaja and Rini van Solingen: Establishing Product Process Dependencies in SPI, Proceedings of the fourth annual SEPG conference, June 7 - 10, 1999, Amsterdam, The Netherlands

[12] Baird, L. Holland, P. and Deacon, S. Learning from Action: Imbedding More Learning into the Performance Fast Enough to Make a Difference. Organizational Dynamics, vol. 27, no. 4, 1999.

The following information is not part of the paper:

Appendices

CV for Tor Stålhane

Tor Stålhane was born in 1944 and became a M.Sc. at the Norwegian Institute of Technology, University of Trondheim (NTNU) in 1969. During 1969 to 1985 he worked at SINTEF - RUNIT, department for languages and compilers. From 1985 he worked on his Ph.D. studies and finished his thesis on software reliability in 1988. From 1988 he was back at SINTEF where he mainly worked with quality assurance and software safety and reliability. In 1997 he became professor in Computer Science at the Stavanger Polytechnic. In 2000 he became professor at the Norwegian University of Technology and Science in Trondheim where he today works full-time. During the latest decade he has been mainly been working with safety analyses of software intensive systems and measurement based process improvement.

CV for Torgeir Dingsøyr

Torgeir Dingsøyr was born in 1973 and is working with knowledge management as a way to improve software development as a doctoral candidate at the Department of Computer and Information Science at the Norwegian University of Science and Technology. He is analysing the usage of knowledge management systems in two medium-sized software-developing companies in Norway, using a combination of qualitative and quantitative research methods.
He received his Master degree in computer science at the Norwegian University of Science and Technology in 1998, with a stay at Université Dauphine in Paris. He has published papers on knowledge management in software engineering, software engineering education, and on combining the artificial intelligence techniques case-based reasoning and data mining.

He is currently staying at the Fraunhofer Institute for Experimental Software Engineering in Kaiserslautern, Germany.

CV for Geir Kjetil Hanssen

Geir Kjetil Hanssen (born 1969) is a Research Scientist in SINTEF Telecom and Informatics. He is a M.Sc. (Cand.Scient) from the Norwegian University of Science and Technology (1996). He works within the computer science area covering various aspects of information technology. System architecture is one of his focus areas, and he has experience in systems design and systems engineering methodology and processes. Other focus areas are transport telematics and process improvement where he assists several Norwegian software companies in improving their internal software development processes. He is working with messaging systems using XML based technology and a system for generating XML messages based on graphical message models (UML). He has experience from the IT industry, both as a systems engineer and developer and as a project leader.

CV for Nils Brede Moe

Nils Brede Moe was born in 1972 and became a M.Sc. at the Norwegian University of Science and Technology (NTNU) in 1998. His main research areas in the field of Software Processes Improvement include: Measurement based improvement, assessments and improvement on an organisational level. Other research areas are Human-Computer Interaction (HCI) and home care (health informatics).

He has been working on the research projects SPIQ and PROFIT, supported by the Norwegian Research Council.

Company description for SINTEF

SINTEF is an independent, non-profit research foundation based in Trondheim and Oslo, Norway. Our role is to encourage innovation and improve competitiveness in Norwegian industry and public administration. In doing so, we maintain close links with the technical Universities in Trondheim and Oslo, collaborating on projects, and sharing equipment and other resources.

With over 1800 employees and a turnover of NOK 1.4 billion, SINTEF is Scandinavia's largest independent research organization. It is organized into eight separate research institutes, covering all major scientific areas and industrial sectors. Refer to our web site www.sintef.no for further information.
SINTEF has over the years been a leading company in the area of software engineering and have broad and deep experience in this area. This experience serves as a sound basis for our Software Process Improvement (SPI) work. Our major SPI activities are:

- The SPIQ program – Software Process Improvement for better Quality. This is a national Norwegian program that is partly a national ESSI type project and partly a co-operation with Norwegian industry to increase the use of process improvement methods in Norwegian software industry.
- The QIS project – Quality Improvement in Scandinavia. This is an EU sponsored project that shall give help and assistance to PIEs in Norway and Sweden and market the concept of SPI to software industry in these two countries.
- The PROFIT program – PROcess improvement For the IT industry, which is a follow-up to the SPIQ program.

Company description for NTNU

As the name states the Norwegian University of Science and Technology, NTNU, is a centre for technological education and research in Norway, with a solid foundation in the natural sciences. This tradition is interwoven with broadly based expertise in the classical university disciplines of the humanities, medicine and the social sciences.
The Department of Computer and Information Science is one of three departments at the Faculty of Physics, Informatics and Mathematics, NTNU. The department is currently located partly at Gløshaugen in Trondheim, Norway. Although modest in size (approximately 150,000 inhabitants), Trondheim is nevertheless the only Norwegian city found in Wired Magazine's recent overview of the world's 46 hottest spots of the global high tech network.

The department offers a broad selection of courses, covering most areas of computer and information science. Currently, the department employs approximately 90 faculty,

staff, scientists and doctoral fellows. Approximately 30 of these are assistant, associate or full professors in permanent positions.