

A. Aurum, H. Petersson and C. Wohlin, "State-of-the-Art: Software Inspections after 25 Years", *Software Testing Verification and Reliability*, Vol. 12, No. 3, pp. 133-154, 2002.

State-of-the-Art: Software Inspections after 25 Years

Aybuke Aurum

School of Information Systems,
Technology and Management
University of New South Wales
Sydney NSW 2052
Australia

Phone: +61-2-9385 4418

Fax: +61-2-9662 4061

E-mail: aybuke@unsw.edu.au

Håkan Petersson

Dept. of Communication Syst.
Lund University
Box 118
SE-221 00 Lund
Sweden

Phone: +46-46-222 4910

Fax: +46-46-145823

E-mail: hakanp@telecom.lth.se

Claes Wohlin^a

Dept. of Software Engineering
and Computer Science
Blekinge Institute of Technology
Box 520, SE-372 25 Ronneby
Sweden

Phone: +46-457-385820

Fax: +46-457-27125

E-mail: claes.wohlin@bth.se

a. Contact author

Abstract

Software inspections, which were originally developed by Michael Fagan in 1976, are an important means to verify and achieve sufficient quality in many software projects today. Since Fagan's initial work, the importance of software inspections has been long recognized by software developers and many organizations. Various proposals have been made by researchers in the hope of improving Fagan's inspection. The proposals include structural changes to the process and several types of support for the inspection process. Most of the proposals have been empirically investigated in different studies.

This is a review paper focusing on the software inspection process in the light of Fagan's inspection, and it summarizes and reviews other types of software inspection processes that have emerged in the last 25 years. The paper also addresses important issues related to the inspection process and examines experimental studies and their findings that are of interest with the purpose of identifying future avenues of research in software inspection.

Keywords: software inspections, Fagan inspection, survey, software reviews, reading techniques

1. Introduction

Software inspections are an important means to verify and achieve sufficient quality in software projects today. It is a static verification technique and one of its main benefits is that it can be applied to any artefact produced during software development. It is probably one of the few methods that people actually agree will help in improving software quality, although it is not always applied.

The objective of this paper is to provide a survey of software inspections given that this technique celebrates its 25th anniversary in 2001. Since the first description of software inspections, there have been further improvements, new ways of conducting inspections have been suggested and several methods supporting software inspections have been proposed. This paper summarizes and provides an overview of software inspection processes that have emerged in the last 25 years. The paper identifies the main differences between commonly known inspection methods and Fagan's inspection. It studies the evolution of inspections, highlights articles presenting original ideas, and presents evaluations of new ideas that have since emerged. It also addresses important issues

related to the inspection process that are of interest for the purpose of identifying future avenues of research in software inspection and it provides a list of concrete research questions.

Software inspection is a peer review process led by software developers who are trained in inspection techniques (IEEE, 1998). Michael Fagan originally developed the software inspection process ‘*out of sheer frustration*’ (Wheeler *et al.*, 1996). It has been 25 years since Fagan published the inspection process in his famous article in 1976 (Fagan, 1976). Since then the importance of software inspections has been increasingly recognized by software developers and many organizations. The need to improve the inspection process has become a long-standing concern of software developers and researchers. Fagan's inspection has been studied and presented by many researchers in various forms.

1.1 Terminology

Software inspection is a type of software review process. One common problem is that, the terminology used for different types of software review processes is often imprecise and leads to confusion. In particular, there is no universal agreement on what a *software inspection process* is and how it is different from the other types of software review processes such as a walkthrough, a formal technical review or a management review. In practice, the inspections performed do not necessarily follow the original definition. In other words, an organization can conduct an inspection that may also fit under the category of a walkthrough or a technical review.

Fagan (1976) strongly emphasizes that inspections have formal procedures and produce repeatable results whereas walkthroughs are performed in varying regularity and thoroughness. He also remarks that in some cases walkthroughs may be identical to formal inspections but in many cases they are informal and less efficient (1986). Wheeler *et al.* (1996) point out some principal differences between review processes. Knight and Myers (1993) suggest that walkthroughs are used to examine source code, formal reviews are the presentation of the work product to the rest of the team members, and inspections are error detection techniques and ensure that particular coding standards and issues are enforced. According to these authors Fagan's inspection is a combination of a walkthrough, formal review and inspection. IEEE Std. 1028-1997 (IEEE, 1998) provides the following descriptions:

- An inspection is ‘*a visual examination of a software product to detect and identify software anomalies, including errors and deviations from standards and specifications*’
- A walkthrough is ‘*a static analysis technique in which a designer or programmer leads members of the development team and other interested parties through a software product, and the participants ask questions and make comments about possible errors, violation of development standards, and other problems*’
- A review is ‘*a process or meeting during which a software product is presented to project personnel, managers, users, customers, user representatives, or other interested parties for comment or approval*’

According to the IEEE Std. 1028-1997 standard, the objective of technical reviews and walkthroughs is to evaluate the software product, whereas the objective of an inspection is to detect and identify defects in the software but not to evaluate it. Thus, inspections are viewed as distinctively different from the other techniques, although the document type used is the same. This definition also parallels Fagan's original idea *i.e.* the aim of the inspection process is to detect the defects and this does not involve searching for solutions and reaching consensus over the product implementation.

2. Software Inspection Process

A software inspection is a well-structured technique that originally began on hardware logic and moved to design and code, test plans and documentation (Fagan, 1986). The process itself can be characterized in terms of its objective, number of participants, preparation, participants' roles, meeting duration, work product size, work maturity, output products and the process discipline (Wheeler *et al.*, 1996). The essential criterion to start an inspection is to have a well-defined software process with an exit criterion and a software product that meets with that criterion (Bisant and Lyle, 1989). An inspection team who plays defined roles perform a software inspection. It is important that people performing the inspection are familiar with the product and they have a sound knowledge about the inspection process or otherwise they must be trained. The members of the inspection team examine the material individually to learn the product. Then participants attend a meeting with the intended purpose of effectively and efficiently identifying defects early in the development process. Next, the list of defects is sent to the author of the documents to be repaired and removed in the later stage of the review process. Hence, software inspection provides a powerful way to improve the quality and productivity of the software process (Ackerman *et al.*, 1989; Russell, 1991).

An effective software review process needs to address relationships of all required variables in terms of tasks involved, tools and methods used, and the skill, training and motivation of people. Various researchers have made proposals which attempt to improve upon process of Fagan's inspection. A literature review reveals two major areas of study, as illustrated in Figure 1. There is considerable research that has been carried out on the structure of the inspection process. Researchers have developed several new process models by restructuring the basic processes in Fagan's inspection and these models have been evaluated and validated empirically. Another group of studies focus on particular methods, tools, and models that support the structure of the inspection process. Examples include support to preparation stage, *i.e.* reading techniques, electronic support to both preparation and inspection meeting, and various models that support the re-inspection process, which have again been empirically validated.

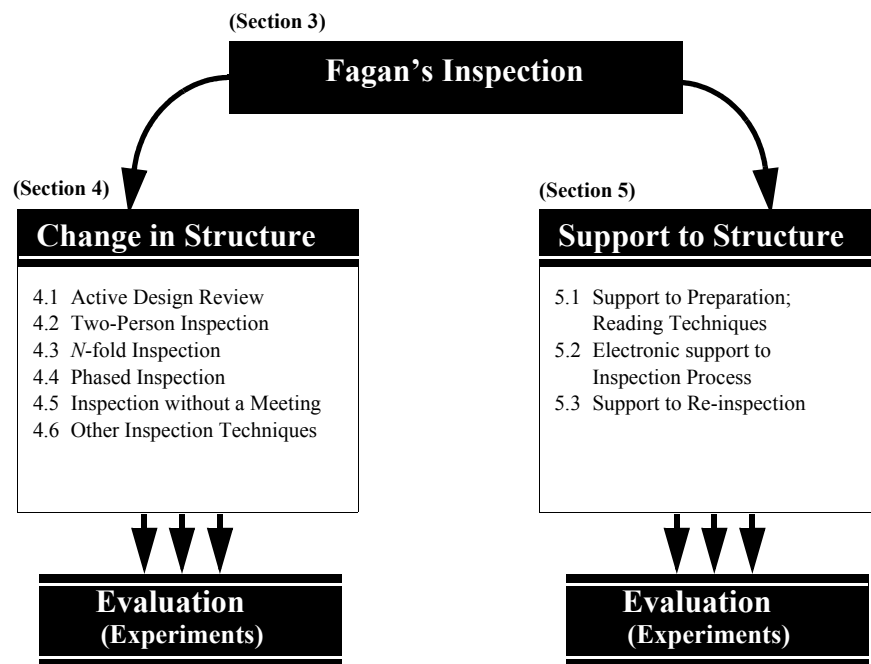


FIGURE 1. Evolution of the Inspection Process with Two Major Areas: Change in Structure and Support to Structure.

This paper focuses on the software inspection process in the light of Fagan's inspection. The structure of this paper is drawn from Figure 1. This section addresses the review process in terms of its objectives and benefits, roles of the participants, documents, inspection pace and the team size. Section 3 describes Fagan's inspection process and Section 4 reviews the literature on structural changes made to this process and experimental findings. Section 5 focuses on methods and models that support the structure of the inspection process and the experiments that validate these studies.

2.1 Objectives and Benefits

The basic objective of software inspections is to improve the quality of the product by analysing the product and detecting defects in time, and removing them before the product is released. According to O'Neill (1997) 42% of all defects result from lack of traceability from the code to the design, to the requirements and to the business case. Researchers all agree that inspections appear to be an effective technique in error detection and productivity improvement (Fagan, 1986; Ackerman *et al.*, 1989; Russell, 1991). First of all, inspections reduce the number of defects in the software throughout the development process. The more defects that can be found earlier, the easier and much less expensive they are to correct. Secondly, they can also uncover defects that would be difficult or impossible to discover in later stages of the life cycle. Thirdly, the inspection process improves learning and communication within the software team (Briand *et al.*, 2000b). Overall, inspections reduce the software development cost, increase the quality of software, and improve the productivity as well as the quality of the decision making process for management (Fagan, 1986; Ackerman *et al.*, 1989; Gilb and Graham, 1993).

Despite the documented benefits of inspection, in many organizations productivity is still measured in terms of lines of code per effort (Briand, 2000b), and therefore the benefits of inspection are not perceived well by management and thus inspections are not used widely in software industry. Shirey (1992) points out that one of the problems in the inspection process is that the information collected from the process may not necessarily be used. Grady (1994) reports that to get inspections incorporated into a large company *e.g.* Hewlett-Packard, may take a long time *e.g.* 15 years. Since inspections are labour intensive, the return from this process is not immediate and clear to the management (Shirey, 1992). Therefore there is a perception that inspections cost more than they are worth (Russell, 1991).

2.2 Roles

Software inspections start with a request from the author of the work product. Following this request a team of reviewers is formed and roles assigned to each member of the team. The team members may play multiple roles. Each role requires specific skills and knowledge. There are roles related to the inspection process, and to the product being inspected. The process roles defined by Fagan are: author, moderator, reader and reviewer(s). According to some researchers, when reviewing requirements documents, each reviewer may be assigned a product role to represent the viewpoint of that role during the review process (Bisant and Lyle, 1989; Basili *et al.*, 1996). Researchers seem to all agree that the manager (project leader) should not be involved with the review process of technical documents. In some cases, the manager is the person who plans the inspection, assigns resources and roles. But on the other hand, a moderator (not being the manager) chairs the software inspection activities and facilitates the interaction between reviewers (O'Neill, 1997). The moderator may also collect the defects from the reviewers and creates a list of defects (Russell, 1991). In other cases, the moderator is a leader who is drawn from other development projects and has many responsibilities and may act as administrator, manager and supporter (Gilb and Graham, 1993). In each case, the moderator must be trained to carry out the review process. Although Russell (1991) recommends that the author participates as a silent observer in the review process, Gilb and Graham (1993) stresses the active participation of the author; in fact they claim that the author may be the best defect finder.

2.3 Documents

The document type can be a design, a requirements specification, code, a test plan or any type of document that is related to the product. One of the strengths of inspections is that a similar process can be applied to a wide range of documents. The document format can be in textual, graphical or tabular form or program code.

The size of a document of work product can vary in length, starting from a couple of pages to dozen of pages or thousands of lines of code. In order to prevent the information-overloading problem on reviewers, a document is generally divided into small chunks and then presented to reviewers. Hence, inspection of a document may take several cycles.

2.4 Inspection Pace

Researchers agree that a slower review process yields more impressive results (Russell, 1991; Gilb and Graham, 1993). Russell (1991) recommends that in code reading, maximum reading speed should not exceed 150 lines per hour. Gilb and Graham (1993) also support this idea by suggesting an optimum-reviewing rate as one page per hour per participant. Laitenberger and DeBaud (1997) report that 200-300 lines of C code is a right size of document to inspect within two hours frame.

2.5 Team Size in Software Inspection

Software inspections involve teamwork where a group of individuals work together to analyse the product in order to identify and remove defects. Teamwork is essential for software developers. There are some major assumptions behind teamwork. First of all, the software products in today's market are too complex to be designed by individuals. The complexity of the products has become the driving force behind the creation of teams. It has been found that the quality of the product improves as it is inspected from multiple viewpoints (Laitenberger and DeBaud, 1997). Secondly there is a belief that people are more committed to their work if they have a voice in the design of the product or use of the development process (Smart and Thompson, 1998).

Inspection teams are formed from small groups of developers. Fagan (1976), in his article suggests that four people constitute as a good-sized inspection team. This suggestion also ties with the industrial work done by Weller (1993), who reports that four person teams are more effective and efficient than three person teams. IEEE STD 1028-1997 suggests teams of 3 to 6 people (IEEE, 1998). Teams of 4 or 5 people are common in practice (Wheeler *et al.*, 1996). Owens (1997) reports that although it is expensive to have more reviewers, it is more effective to have more points of view for requirements inspection (*e.g.* 5-6 inspectors) than for design, and more inspectors are needed for design than coding (*e.g.* 1-2 inspectors for coding). Bisant and Lyle (1989) also support the idea of having 2 people in coding inspection. Their findings illustrate that overall programming speed increases significantly as a result of two-person inspections. Porter *et al.* (1997) performed a detailed study on team size and found that the performance of one-person teams is significantly lower than two or four person teams. Porter *et al.* (1997) also did not find a significant difference in performance between two and four person teams for code inspection.

3. Fagans Inspection

Fagan (1976), in his influential article describing an inspection process developed at IBM, defines inspections as '*a formal, efficient and economical method of finding errors in design and code*'. Fagan's inspection is now known to be applicable to any phase of the software development lifecycle including requirements analysis, test documents, or management documents as well as to design and code. Below is a summary of the key elements from the 1976 paper.

Software inspections involve teamwork and are organized by a moderator. Essentially, a number of participants review a document with the aim of discovering defects. Fagan states that the inspection will be successful if all team members play out their roles fully in the process. In summary the following inspection process roles are described:

- *Moderator*: This is the key person in the inspection who manages the team and coordinates the inspection process.
- *Author (designer or coder)*: The programmer who produces the program design and who may also translate the design into code.
- *Reader*: The programmer who paraphrases the design or code during the meeting.
- *Tester*: The programmer who reviews the product from testing point of view.

Fagan's inspection also consists of six major steps to be followed:

1. *Planning*: An inspection team is formed and the roles are assigned to team members.
2. *Overview*: An optional stage where the inspection team is informed about the product.

3. *Preparation*: Individual reviewers inspect the material independently. The inspection material may be a product of analysis, design (such as entity-relationship diagrams, data-flow diagrams, state-transition diagrams or a text specification) or coding. The aim of preparation is to learn the material and to fulfil their assigned roles.
4. *Examination*: It is also called *inspection meeting*. The aim of the meeting is not to discuss or evaluate the solution but to find and to pool the defects together. The moderator, who guides the meeting, takes notes and prepares a list of defects. The meeting should not last more than two hours. After the meeting the moderator produces a written report of the meeting to ensure that all issues identified in the meeting will be addressed in the rework and follow-up.
5. *Rework*: The defects, to be verified by the moderator, are corrected by the author. Some products must be reworked and re-inspected several times if necessary.
6. *Follow-up*: The moderator checks and verifies each correction.

Although the main steps are preparation, examination and rework in the review process, Fagan states that all steps are necessary and skipping or combining steps is not recommended. One of the benefits of this approach is the feedback provided to the programmer that may take place within a few days of writing a program. Fagan also recommends that, since the inspection process from overview to follow-up may take up to 100 people hours, it is best to have a two hour session when performing an inspection meeting.

Fagan (1986) points out that software inspections are less expensive if defects are detected and repaired at the stage where they are inserted as opposed to when they are found and removed during the testing and operational stage. He presents a detailed cost-benefit analysis showing its cost-effectiveness and provides descriptions to allow software developers to better understand the inspection process. He describes a 1976 project which, when inspection was applied, resulted in 38% less defects in the final product than would have been accomplished if walkthroughs were used (1976). He also reports that IBM inspections found 90% of all defects detected over the life-cycle of a product and IBM had 9% reduction in average project cost compared to when walkthroughs were applied (Fagan, 1986).

4. Changes in the Inspection Process

Figure 2 compares the seven well known inspection methods based on the main process steps involved. The grey and black parts on each bar illustrate the relative emphasis put on the various steps in each method. For instance, in the Active Design Review (Parnas and Weiss, 1985) preparation and inspection meeting steps are handled differently from Fagan's inspection process. On the other hand, in Gilb inspections (Gilb and Graham, 1993) an extra step is added to the inspection process to improve the software development process that produced the document being inspected. The following section further describes the seven methods and addresses experimental

findings on the inspection team size, the number of teams and the coordination strategy within the team or the multiple teams.

		Planning	Overview	Preparation	Inspection Meeting	Rework	Follow-up
1976	Fagan's Inspection	[Solid black bar]					
1985	Active Design Review	[Grey bar]		[Solid black bar]		[Grey bar]	
1989	Two-Person Inspection	[Solid black bar]					
1990	N-Fold Inspection	[Grey bar]		[Solid black bar]		[Grey bar]	
1993	Phased Inspections	[Grey bar]		[Solid black bar]	[Grey bar]		
1993	Inspection without meeting	[Grey bar]		[Solid black bar]	Collection	[Solid black bar]	[Grey bar]
1993	Gilb Inspection	[Grey bar]		[Grey bar]	Process Brainstorming	[Grey bar]	[Grey bar]

FIGURE 2. Commonly Known Inspection Processes with Comparison to Fagan's Inspection

4.1 Active Design Review

Parnas and Weiss (1985) presented active design reviews in the mid 80s. According to the authors the conventional inspection methods fail to find defects in the product development because (1) reviewers are generally overloaded with information during the preparation stage and finding the relevant information is difficult, (2) in most cases the reviewers are not familiar with the goals of the design and they may not be competent in the task, (3) social interaction in large meetings brings a few drawbacks such as production blocking and evaluation apprehension. In order to achieve a full review process they redesigned the review process.

The Active Design Review is conducted as several brief reviews as opposed to one large review, in which each review focuses on certain part of the product. Different reviewers who have particular expertise and specific responsibilities perform each review. In order to guide the reviewers the errors are classified in terms of inconsistency, inefficiency and ambiguity. Reviewers are also provided with a list of questions to be used as a guideline during the review process. Reviews are designed according to the properties of the product that the reviewer is examining. Each review consists of three stages. In the first stage, the reviewers are presented with a brief description of the main product. In the next stage, reviewers study the material following the guideline. Then, in the third stage the issues raised by the reviewers are discussed in small meetings where only the designer and the particular reviewers meet.

Parnas and Weiss successfully applied this approach to the design of military flight navigation system, but did not provide any quantitative measurements.

4.2 Two-Person Inspection

This approach uses a formal method for software inspection. It removes the requirement of the initial resources of Fagan's inspection by eliminating the role of the moderator. The inspection team consists of two persons: an author and a reviewer.

Bisant and Lyle (1989) by applying this approach, conducted experiments to study the programmers' productivity and to determine whether their productivity improves with this technique during the design or coding phase of product development. They found that two-person inspections had immediate benefits in the program quality and program productivity. They pointed out the significant improvements in individual productivity, in particular with novice programmers and less productive programmers. They suggested using this method in small organizations where the team sizes are relatively small, and as a transition to the larger team technique which will have the extra benefits of establishing consistent data to be used in process control.

4.3 N -fold Inspection

This approach uses N independent small efficient teams. The aim in using multiple teams is to identify defects that may not be found by a single team. It is also expected that different teams will detect different faults. Therefore, having a larger number of teams should result in the detection of more defects. The team size varies between 3 and 4 people. The participants of this approach are the author, reviewers and one single moderator. The moderator is assigned for coordinating teams and collecting data from teams. Each team follows the six steps of the Fagan inspection. In this approach multiple teams work in parallel sessions. First each team member individually reviews the document for about two hours. Then, they meet as a team to discuss the defects that may also last up to two hours. After the inspection only one single team becomes responsible for the rest of the development process.

Martin and Tsai (1990) who originally developed this technique, suggested using N -fold inspection for the initial phase of software development for mission-critical systems. They conducted their experiments based on this approach and found that (a) N -fold teams out-performed single teams: even one pair of teams found more defects during the inspection than a single team discovered after inspection, specification and design. Tripp *et al.*, (1991) who applied multiple team inspection to improve the technical review and quality of a safety-critical software standard have also reported favourable results. They found out that there was a low redundancy in defects between teams. Schneider *et al.*, (1992) found that a single team was able to detect 35% of defects present whereas N teams ($N=9$) was able to detect 78% of defects present, (b) there was no significant overlapping between the faults found by different teams. Martin and Tsai (1990) raised a point that although increasing the number of teams result in finding more defects, this increment may not be significant. The best value for N depends on (a) availability of teams, (b) cost of additional teams, and (c) potential cost for not finding a defect during the inspection. Tripp *et al.*, (1991) also described that multiple team inspection is costly, but provides substantial benefits.

Kantorowitz *et al.*, (1997) who conducted experiments with both information and real-time systems identified that N -fold inspection is not dependent on the type and the size of the system. Their analysis suggested that the performance of N -fold inspection depends on the level of expertise of the team members and the number of teams. Based on their findings they developed a probabilistic model for N -fold inspections, in which the model captures the performance of N -fold and it can also be used to determine the optimal number of inspection teams.

4.4 Phased Inspection

Knight and Myers (1993) developed Phased inspection. Some of the ideas in active design reviews, Fagan inspection and *N*-fold inspection are adopted for phased inspection. In this approach the software product is reviewed in a series of partial inspections called *phases*. A simple inspection may have up to six phases and phases are conducted in sequential order. Each phase has a specific goal and is intellectually manageable. During each phase, reviewers fully examine and validate the product for fulfilment with a specific property. The inspection does not progress to the next phase until corrections are completed. To achieve a higher degree of confidence, the approach uses two types of phases: *single-inspector* phase and *multiple-inspector* phase. In the first case, a single person studies the product to determine whether it complies with a checklist. It is recommended to use a technical writer for this process. After this, in the multiple-inspector phase, several reviewers individually study the product by using a different checklist and then attend a meeting session to compare their findings. The multiple-inspector phase may take around four hours. According to Knight and Myers phase inspections can be used not only for fault detection but also to examine the other desirable characteristics of the product *e.g.* portability, reusability or maintainability.

4.5 Inspection without a Meeting

In Fagan's inspection the main focus is on inspection meetings where the members of the inspection team meet to identify and to discuss the defects. The goal of inspection meetings is to bring synergy to the software team. It is believed that the combination of different viewpoints, skills and knowledge from many inspectors creates this synergy (process gains). Synergy can overcome many difficulties encountered in organizational life, including those in the software development process. Fagan (1976) strongly emphasized that the inspection meeting is crucial and reported that most defects will be found during the inspection meeting. He suggested that reviewers who have meetings will be more successful in finding a larger quantity of defects, as opposed to those who work individually. Johnson and Tjahjono (1998) pointed out that the reason that teams who have meetings will be more successful than individuals working alone, is that in Fagan's inspection defect detection only becomes an explicit goal only during the meeting phase.

Recently, it has been reported that in many organizations the structure of Fagan's inspection is radically changed, and performed in three stages: preparation, collection (with or without an inspection meeting) and rework (Votta, 1993; Porter *et al.*, 1995). Furthermore, the expected outcome from preparation and inspection meeting has also considerably changed, *i.e.* the preparation stage is used to identify defects, whereas the collection stage is used to pool out the defects from the reviewers (McCarthy *et al.*, 1995). Experimental findings in Eick *et al.*, (1992) showed that when this structure was applied, reviewers were able to identify 90% of the defects during preparation stage, whereas only 10% of the defects were found during the inspection meeting.

After observing several meetings and collecting data from the programmers of the work products, Votta (1993) identified five reasons for holding meetings: synergy, education, schedule deadline, competition, and requirement. His experiments in AT&T Bell Labs showed that on the positive side holding meetings tend to minimize 'false positives', but on the negative side they may not necessarily bring the synergy. He claimed that holding meetings after the preparation stage does not have a significant effect on the inspection. This finding also ties with the experimental work done by Johnson and Tjahjono (1998). These researchers reported that meeting based software inspections are more costly and do not find significantly more defects than non-meeting based methods, but the former approach is significantly better at reducing the number of 'false positives' and the reviewers actually prefer inspection meetings over a non-meeting approach. Mashayekhi

et al., (1993) argued that face-to-face meetings can be quite expensive because they involve 3-6 people. It is labour intensive and takes many meetings to completely inspect a product since each inspection covers only a small part of the product. Porter *et al.*, (1995) also indicated that holding meetings is time consuming, because bringing several people into one room on a specific day and time requires significant amount of effort.

Votta (1993) raised a point that in a meeting, although n number of reviewers participate, only two of the reviewers can actually interact at any time, because in face-to-face meetings communication is established in sequential order. He also found that, most of the time, only 30-80% of $n-2$ reviewers were actually listening to the conversation. Hence, in each meeting on average $n-4$ reviewer hours per meeting were wasted. Votta found that most defects were also identified before the meeting. He concluded that having a meeting is not cost effective and made a few suggestions: (1) using a nominal¹ group approach or, (2) using correspondence instead of face-to-face meetings or, (3) replacing meetings with all reviewers with two or three person meetings, called *depositions*, which consists of an author, a reviewer and a moderator (optional). According to his experimental findings in many cases depositions were more efficient than inspections.

4.6 Other Inspection Techniques

There are several more inspection techniques, similar to Fagan's inspection, that have several distinct features in relation to the number of people involved in the process, the role of the reviewers and the goal of inspections. For instance, Yourdon (1989) defines a specific walkthrough that is similar to Fagan's inspection. Yourdon calls this approach 'structured walkthroughs', where a walkthrough is described as being simply a peer group review of any product. In Yourdon's inspection, the preparation stage and the inspection meeting does not take more than two hours. In the preparation stage the reviewers are encouraged to note positive aspects of the document. Gilb and Graham (1993) suggest adding an extra step to improve the inspection process itself, right after the inspection meeting, called 'process brainstorming meeting'. The reader can find a literature review on various software inspection processes in Porter *et al.* (1996), a detailed taxonomy of review processes in Wheeler *et al.* (1996) and Johnson (1998). Recently, Laitenberger and DeBaud (2000) also provided a survey and developed a different taxonomy which characterized the nature of software inspection in various dimensions and then they illustrated the differences between software inspection processes during different phases of the software development lifecycle.

5. Support for the Structure of the Inspection Process

There are a considerable number of studies that focus on methods and tools to support the structure of the inspection process. Section 5.1 reviews different reading techniques and the results of experimental findings. Section 5.2 presents the research on electronic support tools and the experimental studies. Finally, Section 5.3 is concerned with various methods and models that are developed to support the re-inspection process and the experiments that validate these studies.

5.1 Support for Preparation: Reading Techniques

Multiple reviewers identifying potential defects in the material using a particular reading technique perform the preparation stage. There are few techniques available to support this activity

1. A nominal group is an artificial group that is created based on combining the results of the performance of the individuals. This means that no group meeting is held.

that are proven to be more effective. Researchers all agree that the choice of reading techniques has potential impact upon inspection performance. Reading techniques are classified as systematic techniques and non-systematic techniques (Porter and Votta, 1994; Porter *et al.*, 1995). The systematic reading techniques such as Perspective-Based Reading, further described in Section 5.1.5., apply a very highly explicit and structural approach to the process. It provides a set of instructions to reviewers and explains how to read the software document and what they should look for (Shull *et al.*, 2000) The non-systematic reading techniques, such as Ad hoc reading or Checklists, apply an intuitive approach and offer little or no support to the reviewer. A few empirical studies have also made comparisons between reading techniques by measuring the number of defects detected by each technique. The following sections describe the most commonly used forms of reading techniques and report the results of a few empirical studies.

5.1.1 Ad hoc Reading

This technique takes a very general viewpoint of reviewers. There is no clear procedure to follow for the reviewer and no training is needed. Reviewers use their own knowledge and experience to identify defects in the documents. Ad hoc reading does not offer any support to the reviewer.

5.1.2 Checklist

Checklist reading is considered to be more systematic than Ad hoc reading. The proposed procedure by Fagan (1976) included the use of checklists. The reviewer answers a list of questions or ticks a number of predefined issues that needs to be checked. It is expected that the questions will guide the reviewer throughout the inspection process. The aim is to define the responsibilities of the reviewers and provide guidance to them to identify the defects. According to Gilb and Graham (1993) checklists are developed from the project itself and they must be prepared for each individual type of documentation and for each different type of product or process role. A checklist may concentrate on questions that help reviewers identify major defects or prioritise different defects. A checklist should be no more than one single page for each type of documentation (Gilb and Graham, 1993).

5.1.3 Stepwise Abstraction

This reading technique was developed in the late 70s as a code reading technique and has been experimentally validated (Linger *et al.*, 1979). In code reading by Stepwise Abstraction, each reviewer identifies subprograms in the software and determines their functions. Next, each reviewer determines the function of the entire program by combining sub-functions together and constructing their own specification for the program. Then, the reviewer compares this derived specification with official specification. Based on this comparison, inconsistencies are identified as defects and studied in the next stage.

Basili and Selby (1987) conducted some experiments to compare code reading by Stepwise Abstraction with Functional Testing and Structural Testing in four small programs. They reported that code reading by Stepwise Abstraction detected more defects than the other two techniques did. They pointed out that the number of defects observed, defect detection rate and the total effort in detection depended on the type of software tested.

Based on the idea that different reading techniques may identify different defects, Roper *et al.* (1997) investigated the relative effectiveness of different reading techniques. They examined the effect of combining three techniques: code reading by Stepwise Abstraction, Functional Testing

and Structural Testing. They found that individual techniques were similar to each other in terms of finding defects. They also pointed out that the reading techniques would be more effective when used in conjunction with other techniques.

5.1.4 Scenario-Based Reading (Defect-Based Reading)

Scenario-Based reading is proposed as a systematic reading technique with specific responsibilities for reviewers (Porter and Votta, 1994; Porter *et al.*, 1995). This technique was originally developed to identify defects from requirements documents. In this approach, defects are classified and a set of questions is developed for each defect class. Scenarios, which are a collection of procedures for detecting particular types of defects, are also developed and focused on a particular viewpoint. The reviewer answers the questions by following a specific scenario.

Several researchers compared Scenario-Based reading with Ad hoc and Checklist reading techniques and studied the potential benefits of Scenario-Based reading on requirements documents by conducting multiple experiments (Porter and Votta, 1994; Porter *et al.*, 1995). In these experiments documents were written in both plain English and tabular form. For scenario-based reviewers, they developed a few scenarios that were derived from checklists and aimed to detect particular classes of defects. Then, each reviewer was assigned to a scenario and was asked to search for different classes of defects. For checklist reviewers, they developed a single checklist that contained several questions. Ad hoc reviewers did not receive any assistance. Experimental findings showed that Scenario-Based reading had a higher detection rate than Ad hoc and Checklist methods. In particular, the Scenario-Based approach found 35% more defects than the other two groups. Checklist reviewers performed no better than Ad hoc reviewers.

Gough *et al.*, (1995) performed a large-scale study using scenarios in a market driven industrial environment (a medical system) and the use of Fagan's inspection with stakeholders, *i.e.* customers and developers. Each team consisted of a moderator, a scribe and an average of five reviewers who were assigned to particular roles *e.g.* requirements engineer or customer representative. In this large case study reviewers identified typically 2-4 defects per hour. The authors pointed out that scenarios are the prime means of elicitation of requirements.

Fusaro *et al.*, (1997) who also conducted experiments using requirements documents on real-time systems, and compared Scenario-Based reading with Ad hoc and Checklist techniques, discovered that the average defect detection rate for Scenario-Based reading was not significantly different from those obtained from the other two techniques.

5.1.5 Perspective-Based Reading

Perspective-Based Reading (PBR) is originally developed and experimentally validated at NASA (Basili *et al.*, 1996). PBR is an enhanced version of scenarios. The technique focuses on the point of view or needs of the stakeholders. Each scenario consists of a set of questions. Scenarios are developed based on the viewpoint of stakeholders. During the review process, reviewers read the document from a particular viewpoint and produce a physical model that can be analysed to answer the questions based on the viewpoint. It is expected that this structural approach will reduce any existing gaps or overlaps between reviewers during the inspection process.

Basili *et al.*, (1996) conducted several experiments at NASA to study the effectiveness of PBR on requirements documents. They found no significant difference in the performance of reviewers who were using PBR and those using their own usual technique when finding defects. However,

PBR reviewers performed significantly better on the generic documents. Laitenberger and DeBaud (1997), also found no significant performance differences when they ran a more detailed experiment using PBR on code documents at Robert Bosch GmbH. Shull *et al.*, (2000) pointed out that PBR is suited for reviewers with a certain range of experience. According to these authors, reviewers who use PBR to inspect requirements documents tend to detect more defects than those who use less-structured approaches. They also emphasized that PBR has beneficial qualities because it is systematic, focused, goal-oriented and customisable, and transferable via training.

5.2 Electronic Support for the Inspection Process

In recent years, there have been a number of attempts to further increase inspection efficiency by the introduction of electronic tool support that has resulted in a number of prototype systems. These systems are designed to support either synchronous, asynchronous or both of types of communications between team members. Some systems are designed to fully support the inspection process whereas others partially support the process *e.g.* the preparation or the inspection meeting. The researchers who conducted experiments studied the effectiveness of these tools and the performance of meetings with and without electronic support. Below several examples of these tools, which are evaluated in empirical studies, are presented.

Mashayekhi *et al.*, (1993) developed a tool for reviewers who are geographically distributed. The tool provided support for the preparation and inspection meeting. This system was structured for software inspections on all types of products including requirements, design and coding and it was designed to facilitate both asynchronous and synchronous activities in software inspection. The authors conducted experiments where each team consisted of six members: an author, a moderator, a recorder and three reviewers. Their findings indicated that meetings using electronic support were as effective as face-to-face meetings.

Johnson (1994) developed a tool to support asynchronous communication. The inspection mainly consists of two stages: *private review* and *public review*. During the private review, reviewers study the material individually and cannot see other participants' comments and issues. During the public review, reviewers can access all comments and issues. Then they vote on whether they agree on issues. Next, the moderator analyses the private and public review and decides whether to go a head with a face-to-face group review meeting or not.

Murphy and Miller (1997) also developed a prototype to support asynchronous communication that used electronic mail as the basis of communication. In this approach some of the tasks of team members were modified: the author did not participate in the process whereas the moderator was responsible for planning the inspection, monitoring the communication, and broadcasting the list of defects to the inspection team to be re-inspected and compiling the list of defects to the author.

Vermunt *et al.*, (1998), who conducted experiments using electronic support only for the inspection meeting, found that 95% of the defects were found during the preparation stage and groups using electronic support performed as well as the groups who did not have electronic support.

5.3 Support for Re-Inspection

This section is concerned with methods and models that have been developed to support the re-inspection process.

5.3.1 Inspection Metrics

Inspection metrics provide important information for the software project manager. Gilb and Graham (1993) state that '*The metrics themselves are the lifeblood of inspections*'. Inspection metrics are useful in many ways; It helps the project manager to see the economics of the software review process, to judge the quality of inspection and decide whether to re-inspect the material or not, to better understand the potential problems within the software process and also provides feedback for the review process.

Although Fagan (1976) does not explicitly discuss the importance of collecting metrics from the inspection process, some metrics on the software inspection process are presented in the original paper. Fagan mentions the classification of defects and gives examples from the inspection reports and forms such as the number of defects found, date of inspection, size of product and the length of inspection. Ackerman *et al.* (1989) list twelve basic measures that should be collected for each inspection, *e.g.* date of inspection, identity of product to be inspected, number of reviewers and number of defects. Barnard and Price (1994) present nine derived key metrics that are used to answer questions concerning different things such as inspection cost, product quality and the effectiveness of the inspection process. They also show various analyses and process control techniques that can be used in inspections. Detailed information about inspection metrics and descriptions of measures to be collected from an inspection process can be found in Gilb and Graham (1993), and Ebenau and Strauss (1994).

Inspection metrics measure various aspects of inspection itself directly as well as indirectly (Gilb and Graham, 1993). There are many useful metrics that can be collected from inspections for analysis to monitor, control and improve the review process. Examples of these inspection metrics include total number of defects found, elapsed time for preparation and inspection meeting, and number of reviewers. By using the inspection metrics several statistics can be computed *e.g.* inspection rate, the number of defects found per reviewer, and the number of remaining defects can be estimated. By collecting inspection metrics, developers can establish a baseline of measurement about their development process. This historical data may help them to determine the quality of the current inspection of the product and to control the process (Christenson *et al.*, 1990; Gilb and Graham, 1993; Burr and Owen, 1996). Chillarege *et al.* (1992), introduced the concept of Orthogonal Defect Classification in order to provide a cause-effect relationship that can be used to monitor the software process. In this approach, by keeping the same classification of defects throughout the phases of the development, and computing the total number of found defects, it is possible to obtain information about the status of the product and the process.

Collecting metrics can be tedious, in particular when dealing with various sizes and types of products. Furthermore, the developers may avoid measurement because of a fear of data misuse. Shirey (1992) reported from a survey at Hewlett-Packard that in most cases the data from the inspection process had never been used for any of the intended purposes.

5.3.2 Estimation of Remaining Defects

The main objective of an inspection is to find and correct defects. The correction of defects leads to increased quality, hence eliminating defects from the product is an important issue. One way of measuring the quality is estimating the number of remaining defects after an inspection. If the number of remaining defects is estimated, it becomes easier to calculate an estimation of the risk of failure. This would also give a hint of the amount of effort that has to be spent further on the

project. Thus the knowledge of the number of remaining defects gives an estimation of the quality of the inspected artefact.

During the review process, the moderator records the number of defects the reviewers find. The list of defects is then sent to the author to be corrected. Meanwhile, by using various methods the number of remaining defects is calculated. Based on the results, the developers decide whether to re-inspect the document. Re-inspection is performed after the author revises the software product (IEEE, 1998). The most commonly used estimation techniques which are applied to software inspection data are summarized below (Eick, *et al.*, 1992; Ebenau and Strauss, 1994; Briand *et al.*, 2000a). Other methods also exist to estimate defect density, however, this is not in the scope of this paper. The focus here solely on estimations from inspections.

(a) **Defect Density:** This is expressed in terms of the defects detected per unit reviewed (Ebenau and Strauss, 1994). A reasonably high defect density may indicate that the review process is working well or that the document contains an unreasonably high number of defects. Since the aim is to eliminate the causes of defects, by re-inspecting the document it is expected that defect density will continually become lower throughout the development. Using upper and lower thresholds on the number of defects per unit of size is also a common approach (Vander Wiel and Votta, 1993). The drawback in this approach is that reviewers tend to reach the upper limit regardless of the document's quality (Briand *et al.*, 1997).

(b) **Subjective Assessment:** Reviewers make their own estimation at the end of the review process. The approach is not fully reliable, thus it is useful to combine this approach with an objective re-inspection criterion (Briand, *et al.*, 2000a).

(c) **Using Historical Data:** The remaining number of defects is estimated based on the historical data from the previous projects. The document is re-inspected if the number of undetected defects is significantly different from the historical average. Too many defects may indicate that the document is poorly written. Too few defects may indicate that the document is poorly reviewed. According to Eick *et al.*, (1992) this approach assumes that the variation between reviewers is larger than the variation between the documents. The drawback in this approach is that, if this assumption is not true then a high quality document may be re-inspected, and a poor quality document may not be re-inspected if the inspection is poorly performed. Gilb and Graham (1993) also suggest using historical data as an estimation method.

(d) **Capture-Recapture Method:** This approach is based on the amount of overlap in defects found by various reviewers and uses estimation models to calculate the number of remaining defects in the product. The capture-recapture method was originally used to estimate the number of animals in wilderness (animal population). Recently, researchers started using the same method to estimate the number of defects in software products. The approach is based on applying statistical method to collected defects. Three methods have been applied for this purpose: Maximum Likelihood Estimator (MLE) and Jackknife Estimator (JE), and Chao's Estimator. Capture-recapture models applied to software inspections make use of models based on assumption of two factors (1) the probability of a specific defect being found (2) the ability of the reviewers.

The capture-recapture method for estimating the number of remaining defects has been empirically studied by many researchers during the last decade. Eick *et al.*, (1992) performed the first application of capture-recapture in software inspections. They applied this method to estimate the defect content in the design documents in AT&T's International Switching Division. Vander Wiel and Votta (1993) used Monte Carlo simulation to investigate the assumptions made by capture-recapture methods, in particular MLE and JE. They found that if defects are classified into a small

number of groups, MLE performs better than JE and if there is no grouping then MLE performs poorly. Based on this Wohlin *et al.*, (1995) carried out the first controlled experiment on capture-recapture in software engineering. They proposed the application of a filtering process that grouped the defects into small classes. The classes are defined in such way that some of the assumptions of the capture-recapture method are better fulfilled, hence the estimate is improved.

Briand *et al.*, (1997) studied the impact of the number of reviewers on the performance of capture-recapture models. They commented that the capture-recapture models generally underestimate the remaining defects. They suggested that the estimates will not be accurate if the group size is less than four people. They recommended using the Jackknife estimator for group size four, and the Chao estimator for group size five. The experimental findings from Miller (1999) also supported this recommendation. Miller found that the Jackknife estimator is the best estimator for team size 3-5 and recommended using Chao's estimator for larger group sizes or the Jackknife estimator with combination of other estimators. Briand *et al.*, (2000a) conclude that models are strongly affected by the number of reviewers and, therefore, developers must consider this factor before using capture-recapture models.

(e) **Curve-fitting Methods:** Wohlin and Runeson (1998) proposed estimating the number of defects using a curve fitting approach. They plotted the inspection data and fitted a curve to the plot to produce an estimate. They pointed out that curve fitting is more useful than capture-recapture estimates as the latter is based on statistical analysis with underlying assumptions that are largely not fulfilled. The curve-fitting approach was integrated with capture-recapture methods by Briand *et al.* (1998), where they also applied a linear function to perform the estimation. The function applied by Wohlin and Runeson (1998) was an exponentially decreasing function.

The estimation methods listed above, from (a) to (e), are validated in empirical studies. It is important to note that most of the evaluations are from laboratory environments (controlled experiments), whereas only a few of them are based on industrial case studies.

6. Summary and Conclusion

Software inspections are important for improving the quality of software. Although software inspections are widely used among developers, the application of the process varies. This paper reviewed existing literature on software inspections and addressed issues related to the inspection process. Reviewing the literature in this field made us aware of the differences and overlaps between studies. Figure 1 illustrates the summary of the evolution of software inspections.

The following two major areas of study is identified in the literature.

- In the last 25 years several variations of Fagan's inspection have been proposed to improve the performance of software inspections. There are some distinctive structural differences between models. Restructuring the basic processes in Fagan's inspection creates the different models. This includes changing the activities in each preparation or inspection meeting, emphasizing different goals in each stage, changing the number of participants in teams, employing single or multiple teams, changing the coordination strategy between teams or participants of the teams. A summary of these inspection methods is illustrated in Figure 2. Most of the methods in Figure 2 have been empirically evaluated.
- There is a considerable amount of research on various models, methods and tools that support the structure of software inspections. Support approaches include various reading techniques,

electronic support tools, and various models that support the decision making process in relation to re-inspection.

Table 1 provides a brief summary of the two major areas of research and corresponding empirical studies in the literature. The first two columns show the main approaches or suggested changes in software inspections. The first column includes changes to the structure of software inspections and the second column presents new ideas of how to support software inspections. The objective, in the first two columns, has been to give credit to the authors that proposed the new ideas. Thus, the references provided in the first two columns are intended to reference the original source. The new ideas are annotated so that cross-references to empirical studies in columns three and four can easily be identified. The references in columns three and four are to papers where primarily a new idea has been evaluated, although some minor changes may also have been introduced. The table is divided into three rows to classify the research into the three major phases of software inspections, *i.e.* preparation, meeting and re-inspection. The table is by no means exhaustive. The intention has been to capture some of the main contributions in the area.

Table 1: A Summary of the Literature.

	Change in Structure	Support to Structure	Evaluation of Structure	Evaluation of Support to Structure
Preparation	A1) Fagan Inspection Fagan'76 A2) Active Design Review, Parnas and Weiss'85 A3) Two-Person Inspection, Bisant and Lyle'89 A4) N-fold Inspection, Martin and Tsai'90 A5) Phased Inspection, Knight and Myers'93 A6) Inspection without Meet- ing, Votta'93 A7) Gilb Inspection, Gilb and Graham'93	B1) Ad hoc Reading B2) Checklists B3) Stepwise Abstraction, Linger et al.'79 B4) Scenario Based Reading, Porter and Votta'94 B5) Perspective Based Read- ing, Basili et al.'96 B6) Electronic Support	<ul style="list-style-type: none"> • Wheeler et al.'96, A1 • Russell 91', A1 • Tripp et al.'91, A4 • Schneider et al.'92, A4 • Kantorowitz et al.'97, A4 	<ul style="list-style-type: none"> • Basili et al.'87, B3 • Roper et al.'97, B3 • Porter et al.'95, B4 • Gough et al.'95, B4 • Fusaro et al.'97, B4 • Laitenberger and Debaud'97, B5 • Shull et al.'00, B5 • Mashayekhi et al.'93, B6 • Johnson'94, B6 • Murphy and Miller'97, B6
	Meeting	C1) Active Design Review, Parnas and Weiss'85 C2) N-fold Inspection, Martin and Tsai'90 C3) Inspection without Meet- ing, Votta'93 C4) Gilb Inspection, Gilb and Graham93	D1) Electronic Support	
Re-Inspection		E1) Defect Density E2) Subjective Assessment E3) Historical Data E4) Capture-Recapture, Eick et al.'92 E5) Curve Fitting, Wohlin and Runeson'98		<ul style="list-style-type: none"> • Briand et al. '00b, E2 • Vander Wiel and Votta'93, E4 • Wohlin et al.'95, E4 • Briand et al.'97, E4 • Briand et al.'00a, E4 • Miller'99, E4 • Briand et al.'98, E4 and E5

Software inspections are widely used and recognized as a cost-effective way of removing defects. Despite this, there is still room for improvement in this field. The extensive research into software inspections has provided valuable knowledge and experience. There have been contributions to areas such as the inspection process, reading techniques and defect estimation techniques. It is not possible to identify any single study or proposal that has provided a major breakthrough in how software inspections are researched or practiced. The studies have, however, jointly contributed to the evolution of software inspections. This means that they have contributed to the general body of knowledge in software inspections. It should however be noted that no definitive answers have

been found and that many research questions still are open, although inspections have matured considerably during the last 25 years. Some of the open research questions are:

- What is an efficient reading technique? The research is inconclusive on this issue and more work is needed.
- Which reading technique should be used for which type of artefacts? The inability to find a superior way of reading may be due to different reading techniques being better suited for different artefacts or even applications.
- How should models, in for example UML, be inspected? Object orientation, UML and Java have been major changes in the last couple of years, but it is not clear how to efficiently inspect these models and notations.
- What is the best way to support inspectors in their work? There is little work done on how to develop checklists in an efficient way and also how to best support different scenarios when using scenario-based techniques.
- What is the best way to plan and staff software inspections? Research is needed toward supporting the person who takes the decisions prior to an inspection. Support is required so that a person in a cost-effective and efficient way can determine the number of reviewers needed and the competence of the reviewers.
- Is it possible to develop a benchmark that can provide a general measure of the performance of software inspections which can also be further tailored for specific applications? Software inspections are carried out in a number of ways and at a large number of places, but so far it was not possible to make comparisons between different ways of conducting software inspections. Research along this line should aim at establishing methods for benchmarking software inspection processes.
- In what way can the decision making process after an inspection be supported? Work has been conducted regarding estimation of remaining software defects, but little work has been carried out of how to use the information. When is it suitable to add reviewers, make a re-inspections or continue with the development process?
- What is the relationship between software inspections and software testing? Both these techniques aim at identifying software defects, but which defects are best found by which technique? What is the best way to ensure that the techniques complement each other in the best possible way?
- What is the best way to pool the comments from individual reviewers? Some work has been done regarding inspection meetings, but using electronic support to view each others comments may be helpful and possibly it is more efficient to only include a sub-group of the reviewers for these meetings.
- Are the estimation methods good estimators of the actual number of defects? This research question includes evaluation and validation of defect estimation methods in industrial case studies, where the actual data may be compared with the estimated value.
- Is it possible to identify relationships between different parameters in software inspections? This includes empirical studies related to different parameter in software inspections, for example, inspection pace versus effectiveness and defect density versus inspection pace.
- Are inspections different between different application domains such as real-time systems versus information systems or web-based systems?

These open research questions illustrate that although a lot of research has been devoted to software inspections during the last 25 years, there are still many unresolved issues. The research has come far since 1976, but it is important to continue to increase our understanding of software inspections to allow for improvements in the area. Given the potential, in terms of increasing software quality throughout the software development process, research into software inspections is still very relevant and important for the success of software development.

Acknowledgment

This research was partially conducted during a sabbatical visit by Dr. Aurum to the Department of Communication Systems, Lund University. Dr. Aurum was partially supported by the Swedish Institute Research Grant and a Research Grant from Department of Communication Systems, Lund University.

References

- Ackerman, F. A.; Buchwald, L. S.; Lewski, F. H. (1989): 'Software Inspections: An Effective Verification Process'. *IEEE Software*, 31-36, May.
- Barnard, J. and Price, A. (1994): 'Managing Code Inspection Information'. *IEEE Software*, 59-69, March.
- Basili, V. R.; Selby, R. W. (1987): 'Comparing the Effectiveness of Software Testing Strategy'. *IEEE Transactions on Software Engineering*, 13(12), 1278-1296, Dec.
- Basili, V. R.; Green, S.; Laitenberger, O.; Lanubile, F.; Shull, F.; Sörumgård, S., Zelkowitz, M. (1996): 'The Empirical Investigation of Perspective-Based Reading'. *International Journal on Empirical Software Engineering*, 1(12), 133-144.
- Bisant, D. B. and Lyle, J. R. (1989): 'A Two Person Inspection Method to Improve Programming Productivity'. *IEEE Transactions on Software Engineering*, 15(10), 1294-1304, October.
- Briand, L. C.; El Emam, K.; Freimut, B. G.; Laitenberger, O. (1997): 'Quantitative Evaluation of Capture-Recapture Models to Control Software Inspections'. *Proceedings of the 8th International Symposium on Software Reliability Engineering*, 234-244.
- Briand, L. C.; El Emam, K.; Freimut, B. G. (1998): 'A Comparison and Integration of Capture-Recapture Models and the Detection Profile Method'. *Proceedings of the 9th International Symposium on Software Reliability Engineering*, 32-41.
- Briand, L. C.; El Emam, K.; Freimut, B. G.; Laitenberger, O. (2000a): 'A Comprehensive Evaluation of Capture-Recapture Models for Estimating Software Defect Content'. *IEEE Transactions on Software Engineering*, 26(6), 518-539, June.
- Briand, L. C.; Freimut, B.; Vollei, F. (2000b): 'Assessing the Cost-Effectiveness of Inspections by Combining Project Data and Expert Opinion'. *Proceedings of the 11th International Symposium on Software Reliability Engineering*, 124-135.
- Burr, A. and Owen, M. (1996): *Statistical Methods for Software Quality*. Thomson Computer Press, ISBN 1-85032-171-X.

- Chillarege, R.; Bhandari, I. S.; Chaar, J. K.; Halliday, M. J., Moebus, D. S.; Ray, B. K.; Wong, M. (1992): 'Orthogonal Defect Classification: A Concept for in-process Measurements'. *IEEE Transactions on Software Engineering*, 18(11), 943-956.
- Christenson, D. A.; Huang, S. T.; Lamparez, A. J. (1990): 'Statistical Quality Control Applied to Code Inspections'. *IEEE Journal on Selected Areas in Communications*, 8(2), 196-200.
- Ebenau, R. G. and Strauss, S. H. (1994): *Software Inspection Process*. McGraw Hill (System Design and Implementation Series), ISBN 0-07-062166-7.
- Eick, S. G.; Loader, C. R.; Long, M. D.; Votta, L. G.; Vander Wiel, S. (1992): 'Estimating Software Fault Content Before Coding'. *Proceedings of the 14th International Conference on Software Engineering*, 49-65.
- Fagan, M. E. (1976): 'Design and Code Inspections to Reduce Errors in Program Development'. *IBM Systems Journal*, 15(3), 182-211.
- Fagan, M. E. (1986): 'Advances in Software Inspections'. *IEEE Transactions on Software Engineering*, 12(7) July.
- Fusaro, P.; Lanubile, F.; Visaggio, G. (1997): 'A Replicated Experiment to Assess Requirements Inspection Techniques'. *International Journal on Empirical Software Engineering*, 2(1), 39-57.
- Gilb, T. and Graham, D. (1993): *Software Inspection*. Addison Wesley Publishing Company. ISBN 0-201-63181-4
- Gough, P. A.; Fodemski, F. T.; Higgins, S. A.; Ray, S. J. (1995): 'Scenarios- an Industrial Case Study and Hypermedia Enhancements'. *Proceedings of the 2nd. IEEE International Symposium on Requirements Engineering*, 10-17.
- Grady, R. B. (1994): 'Successfully Applying Software Metrics'. *IEEE Computer*, 27(9), 18-25, September.
- IEEE Std., 1028-1997 (1998): *Standard for Software Reviews*. The Institute of Electrical and Electronics Engineering, Inc. ISBN 1-55937-987-1.
- Johnson, P. M. (1994): 'An Instrumented Approach to Improving Software Quality through Formal Technical Review'. *Proceedings of 16th International Conference on Software Engineering*, 113-122.
- Johnson, P. M. (1998): 'Reengineering Inspection: The Future of Formal Technical Review'. *Communications of the ACM*, 41(2), 49-52, February.
- Johnson, P. M. and Tjahjono, D. (1998): 'Does Every Inspection Really Need a Meeting?'. *Empirical Software Engineering*, 3, 9-35.
- Kantorowitz, E.; Guttman, A.; Arzi, L (1997): 'The Performance of the N-Fold Requirements Inspection Method'. *Requirements Engineering Journal*, 2, 152-164.
- Knight, J. C. and Myers, A. E. (1993): 'An Improved Inspection Technique'. *Communications of ACM*, 36(11), 50-69, November.

- Laitenberger, O. and DeBaud, J. N. (2000): 'An Encompassing Life Cycle Centric Survey of Software Inspection'. *Journal of Systems and Software*, 50(1), 5-31.
- Laitenberger, O. and DeBaud, J. N. (1997): 'Perspective-based Reading of Code Documents at Robert Bosch GmbH'. *Information and Software Technology*, 39, 781-791.
- Linger, R. C.; Mills, H. D.; Witt, B. I. (1979): *Structured Programming: Theory and Practice*. Addison Wesley.
- Martin, J. and Tsai, W. T. (1990): 'N-Fold Inspection: A Requirements Analysis Technique'. *Communications of ACM*, 33(2), 225-232, February.
- Mashayekhi, V.; Drake, J.; Tsai, W. T.; Riedl, J. (1993): 'Distributed Collaborative Software Inspection'. *IEEE Software*, 66-75, September.
- McCarthy, P.; Porter, R.; Riedl, J. (1995): 'An Experiment to Assess Cost-Benefit of Inspection Meetings and Their Alternatives'. *Technical Report*, Computer Science, Dept., University of Maryland.
- Miller, J. (1999): 'Estimating the Number of Defects after Inspection'. *Software Testing Verification and Reliability*, 9(4), 167-189.
- Murphy, P. and Miller, J. (1997): 'A Process for Asynchronous Software Inspection'. *Proceedings on 18th. IEEE International Workshop on Software Technology and Engineering*, 96-104.
- O'Neill, D. (1997): 'Issues in Software Inspection'. *IEEE Software*, 14, 18-19, Jan./Feb.
- Owens, K. (1997): 'Software Detailed Technical Reviews: Findings and Using Defects'. *Wescon'97, Conference Proceedings*, 128-133.
- Parnas, D. L. and Weiss, D. M. (1985): 'Active Design Reviews: Principles and Practices'. *Proceedings of ICSE'85*, (London, England, Aug 28-30), IEEE Computer Society, 132-136.
- Porter, A. A. and Votta, L. G. (1994): 'An Experiment to Assess Different Defect Detection Methods for Software Requirements Inspections'. *Proceedings on 16th International Conference on Software Engineering, ICSE-16*, 103-112.
- Porter, A. A.; Votta, L. G.; Basili, V. (1995): 'Comparing Detection Methods for Software Requirements Inspection: A Replicated Experiment'. *IEEE Transaction on Software Engineering*, 21(6), 563-575.
- Porter, A. A.; Siy, H. P.; Votta, L. G. (1996): 'A Survey of Software Inspection'. *Advances in Computers*, 42, 40-76, November.
- Porter, A. A.; Siy, H. P.; Toman, C. A.; Votta, L. G. (1997): 'An Experiment to Assess the Cost-Benefit of Code Inspections in Large Scale Software Development', *IEEE Transactions on Software Engineering*, 23(6), 329-346.
- Roper, M.; Wood, M.; Miller, J. (1997): 'An Empirical Evaluation of Defect Detection Techniques'. *Information and Software Technology*, 39(11), 763-775.
- Russell, G. W. (1991): 'Experience with Inspection in Ultralarge-Scale Developments'. *IEEE Software*, 25-31, January.

Schneider, G. M.; Martin, J.; Tsai, W. T. (1992): 'An Experimental Study of Fault Detection in User Requirements Documents'. *ACM Transactions on Software Engineering and Methodology*, 1(2), 188-204, April.

Shirey, G. C. (1992): 'How Inspections Fail'. *Proceedings of 9th. International Conference on Testing Computer Software*, 151-159.

Shull, F.; Rus, I.; Basili, V. (2000): 'How Perspective-Based Reading Can Improve Requirements Inspection'. *IEEE, Computer*, 73-79, July.

Smart, K. L. and Thompson, M. (1998): 'Changing the Way We Work: Fundamentals of Effective Teams'. IPCC'98. *Proceedings of IEEE International Communication Conference*. 2, 383-390.

Tripp, L. L.; Struck, W. F.; Pflug, B. K. (1991): 'The Application of Multiple Team Inspections on Safety Critical Software Standard'. *IEEE, Proceedings of 4th. Software Engineering Standards Application Workshop*, 106-111.

Vander Wiel, S. and Votta, L. (1993): 'Assessing Software Designs Using Capture-Recapture Methods'. *IEEE Transactions on Software Engineering*, 19(11), 1045-1054.

Vermunt, A.; Smits, M.; Van der Pijl, G. (1998): 'Using GSS to Support Error Detection in Software Specifications'. *IEEE, Proceedings 31st Annual Hawaii International Conference on System Sciences*, 566-574.

Votta, L. G. (1993): 'Does Every Inspection Need a Meeting?' *Proceedings of ACM Symposium on Software Development Engineering*. Reprinted in *Software Inspection: An Industry Best Practice*. *IEEE, Computer Society Press, USA*, (1996). ISBN 0-8186-7340-0.

Weller, E. F. (1993): 'Lessons from Three Years of Inspection Data'. *IEEE Software*, 10(5), 38-45, September.

Wheeler, D. A.; Brykczynski, B.; Meeson, R. N. (1996): 'Peer Review Process Similar to Inspection'. *Software Inspection: An Industry Best Practice*. *IEEE Computer Society Press, USA*. ISBN 0-8186-7340-0.

Wohlin, C.; Runeson, P.; Brantestam, J. (1995): 'An Experimental Evaluation of Capture-Recapture in Software Inspection'. *Software Testing, Verification, and Reliability*, 5, 213-232.

Wohlin, C. and Runeson, P. (1998): 'Defect Content Estimation from Review Data'. *Proceedings of 20th. International Conference on Software Engineering*, 400-409.

Yourdon, E. (1989): *Structured Walkthroughs*. 4th Edition. Prentice-Hall, Englewood Cliffs, N.J., 1989.