

Kaizen and stochastic networks in support of investigating aircraft failures

Abstract

The investigation of aircraft failures, their causes, and the prevention of their reoccurrence, usually performed by technical failure-analysis teams, are crucial to achieving and maintaining a high standard of flight safety. We developed and applied a twofold improvement framework to the investigation process. In the first phase we reconstructed the investigation process by using a Kaizen project. We then created a simulation model of the resulting stochastic processing network to evaluate alternative configurations. The results indicated a significant improvement in: throughput times and the quality of the investigations which promote flight safety. This unique improvement framework may be suitable for the many organizations that process several concurrent types of jobs (or projects) in a stochastic and dynamic environment.

After a safety incident, it is crucial to the aircraft operators and manufacturers to return the fleet to flying. Technical failure-analysis teams reveal the cause of the failure, and act appropriately to prevent further shortfalls before the decision makers and managers decide to resume flights. These teams routinely operate under time-pressure and perform a critical role in achieving, sustaining and improving the safety of flights. They work in dynamic, stochastic environments where they deal with: incoming investigations that occur randomly; managers who need reliable information quickly to support their decision-making; a high degree of uncertainty regarding the cause of failure during the first stages of an investigation; and scarce resources (e.g., workers, equipment, budget and materials). Such a scenario often results in an environment where more than half of the process's throughput time is typically spent in queues for: resources, parts, information or the decisions of third parties (e.g., Table 3 in Cohen et al., 2004).

We developed an improvement framework that integrates a Kaizen project and stochastic processing networks to improve the performance of an air force failure-analysis team. A Kaizen project is a technique of Japanese origin which is designed to achieve a marked improvement within a week. We then constructed a simulation based on a stochastic processing network model.

We combined the merits and neutralized the pitfalls of both approaches: Kaizen is very effective in reconstructing and implementing new processes, but its short duration and intensive nature prevents making in-depth analyses of complex environments; a stochastic processing networks' approach is ideal to provide the missing in-depth analysis but time consuming with regard to exploring the different improvement possibilities of a process. So by initially reshaping the process on a Kaizen project solves this pitfall.

We endeavor to: acquaint the readers with the failure analysis framework and the failure analysis team management processes; describe the improvement framework and its main components – Kaizen and stochastic processing networks; detail the application process and the model construction; present results and ultimately our conclusions.

A brief introduction to aircraft failure analysis

We commence by describing the major components of aircraft failure analysis. We will then present the failure analysis team and the motivation for the improvement initiatives.

Failure analysis is a complex and wide field of interest that involves many disciplines in the areas of mechanics, physics, metallurgy, chemistry, corrosion, manufacturing processes, stress analysis and accompanying numerical techniques (e.g., finite and boundary element methods), design analysis and fracture mechanics including environmentally induced cracking, non-destructive evaluation, and probabilistic risk evaluation (Decastro and Fernandes, 2004). Owing to the variety of problems and the diverse approaches to their solutions, Thomas suggests that a single general approach cannot encompass this area of interest (Thomas, 1989). We focus upon technical aircraft failure investigations that answer the following questions: what happened? how did it happen? why did it happen? and what should be done to prevent such a failure in the short, medium and long term? There are investigations with legal aspects such as insurance investigations, but these are beyond the scope of this article.

Our methodology for a failure analysis investigation is set by a typical life cycle. This life cycle includes the following components:

- (1) a demand for a new investigation, and collection of available information and evidence
- (2) an initial investigation (documentation and photography, visual inspection – macroscopic and microscopic, characterization of materials)
- (3) an initial report (possible scenarios, initial recommendations, the investigation's proposed plan of action)
- (4) an investigation (laboratories, experiments, analyses, field studies)
- (5) a final report (conclusions and recommendations)

The air force team in charge of these activities is comprised of a commander (the principal investigator), investigators and laboratory technicians. The equipment available enables the preparation of evidence (cleaning and cutting machines), macroscopic examination (e.g., stereoscopes) and microscopic examination (e.g., scanning electron microscope). The team gets laboratory, engineering and analyses services from adjacent organizations in the areas of: materials engineering (metallurgy, non-metallic materials), mechanical testing and experiments, mechanical engineering and analyses (e.g., finite element models), non-destructive methods etc.

An incoming investigation is classified according to its priority. It backlog in an entrance queue until a set of predetermined standard requirements (e.g., information, available evidence and a formal approval to investigate) are met. The commander then assigns it to an investigator. His two main considerations are: matching the needs of an investigation to the skills of its investigator (primary), and balancing the load in the team

(secondary). For example: a complicated accident investigation will be assigned to an experienced investigator regardless of either party's workload. Each investigator has several concurrent investigations that may be at different stages of their life cycles.

The team's main goal is to conduct high quality investigations while minimizing their throughput time. Consequently, the investigators' performances are evaluated according to the throughput time and quality of completed investigations. Quality is evaluated by the principal investigator. Each investigator strives to perform his own work efficiently and reports to the commander at bi-weekly progress meetings.

The improvement initiatives arose owing to the lengthy completion times of investigations (average of 130 days in 2005) and the initial reports (average of 55 days in 2005), and a need to improve the quality of the investigation (i.e., a definitive answer to the 'why did it happen?' question).

High priority investigations usually receive essential resources immediately, so we focused on improving regular priority investigations (about 60% of all investigations). Our objectives were to: reduce the average completion and initial report submission times, and at the same time increase the quality of the investigations.

The improvement framework

We developed a hierarchical twofold improvement framework. The first stage was designed to evaluate the existing air force failure-analysis methodology and improve it by reconstructing the investigation process. In the second stage we developed a dynamic stochastic processing network model. Its aim was to support decision-making, and enable an examination of further adjustments in the team's operating policies. In the rest of this

section we present: the Kaizen framework; then we elaborate on the stochastic processing networks' approach.

What is Kaizen?

Kaizen is the Japanese word for improvement. Many companies use Kaizen projects (or events) to improve process performance (Bradley and Willett 2004). Unlike other gradual improvement programs, it is designed to achieve a marked improvement within a week. A Kaizen project might impose significant changes such as: the relocation of workers, rearranging equipment, or the introduction of different operational methods and processes. Direct benefits might include lower throughput times, higher quality, lower costs and better products or service levels – qualities that customers and stakeholders appreciate. The improvement process involves using: various tools and methods to characterize the problem, a formal root cause analysis, industrial engineering and methods of operations management. Kaizen is popular in manufacturing settings, service business processes, education and even organizational design (Bradley and Willett 2004; Brunet and New 2003; Swank 2003; Emiliani 2005; Berger 1997).

We now review the Kaizen methodology steps as they apply to the sequel offering specific examples of our work:

Step 0: Problem identification and preparation

Kaizen projects are based on identified problems. Before initiating a Kaizen project specific problems must be defined (e.g., an excessively long throughput time of a regular investigation). The Kaizen approach requires that employees participate. Therefore a

Kaizen team must be comprised of members who are familiar with the problem and its complex processes.

We assembled a team of investigators, customers (mainly from headquarters), suppliers (e.g., workshop technicians; a technician who develops the relevant experiments etc.), and an objective external participant. We gathered information concerning the problem: throughput time, quality data, costs and demand. This information was used to define quantitative goals and evaluate improvement alternatives.

Step 1: Prepare process maps

During this stage the team maps the existing work processes (e.g., from an investigation's arrival to its completion). These maps are very specific, and each of their components is evaluated for its value to the client (when possible, quantitatively and otherwise qualitatively, i.e., high, medium, or low-value). Often this evaluation is simplified by asking if the customer would be willing to pay for only this part of the process. The other attributes of each component are evaluated (i.e., throughput time, cost, required equipment etc.). This allows team members to become familiar with the overall process flow and its specific components. It creates a common understanding of the differentiated value of the process parts enabling the identification and exploitation of improvement opportunities, and reducing process parts deemed less important, or which contain waste.

Our analysis of the process maps indicated that internal processes (e.g., field visit coordination, the approval process of a report distribution, waiting for parts' cutting etc.) are less important to the customer. Whenever possible, these processes were reduced or improved as seen in Step 2.

Step 2: Identify improvement opportunities

In Step 2 a Kaizen team generates ideas to improve the process by examining the mapped process and related data. They then raise and write suggestions for improvements. Ideas may involve process reengineering or local improvements. Each idea is examined thoroughly. The team discusses its applicability, possible improvement and effects. Some ideas are screened at this stage.

An example of one of the ideas was to move a cutting machine to the failure analysis team facilities, and to train team members to operate it. The cutting workshop technician instinctively resented (“you are taking my work”) this initiative. The pleasant atmosphere within the team, and a discussion that demonstrated not only the significance of the process improvement, but also that all sides would benefit (e.g., the technician has enough work without this portion, and the improvement will reduce the number of unexpected, afterhours calls for technicians). It ended in a unanimous agreement to implement the initiative. This initiative simplified the first stages of an investigation when speed is critical. The investigator or laboratory technician became independent in cutting the parts and so reduced bureaucracy and saved time.

Usually at this stage, management grants the Kaizen charter which sets the Kaizen project goals (e.g., reduce the investigation’s average throughput time to 90 days).

Step 3: Improve the process

At this stage the team defines an improved process and simulates its flow. This may lead to the detection of process bottlenecks that need to be solved. Then the team needs to estimate and assess the performance measures and match them to its charter. In complex

processes time constraints prevent in depth simulation, therefore the process flow and its dynamic and stochastic nature would not be realized.

The team documents the actions and resources that are needed to implement the improved process. The accepted procedure is to use existing resources as much as possible (the use of a Kaizen project to manipulate management to obtain budgets and resources is unacceptable).

Step 4: Implement the process

In Step 4 the team makes all preparations to commence work according to the new process. The key principle is to apply all the changes promptly.

At this stage we taught the new investigation process to all those involved people who were not already part of the Kaizen team, moved the cutting equipment and started incoming work in keeping with the improved process.

Step 5: Get management's approval

The final step is to present Kaizen to management. Management offers its guidelines and directions concerning significant gaps that need resources for completion.

In our case management was very enthusiastic about the potential improvement. The new process got the 'go ahead' and was implemented routinely.

Stochastic processing networks

We modeled the investigation process using a stochastic processing network, based on the works of Adler and Cohen (Adler et al., 1996; Cohen et al., 2004). Adler validated the model based on an existent R&D organization, and showed that the model quite

accurately simulated its performance (Adler et al., 1996). We considered investigations to be projects as they were unique in their operational requirements and differed in activity durations. Still, they could be classified as they shared common characteristics such as the precedent relations between investigations' activities. Moreover, some activities were common to all investigations, and for each such activity, the historical realizations of activity times should conform to a common distribution function. Therefore the failure-analysis team operated in a stochastic environment of multiple concurrent investigations which competed for the same set of scarce resources (e.g., investigators, laboratory equipment).

When several activities of an investigation can start being processed contemporaneously, we refer to the phenomenon as a 'fork'; conversely when an activity cannot begin until the precursory activities have been completed, we call it a 'join'. The duration of time required to complete an activity is called its 'processing time', and 'inter-arrival times' refer to the intervals between successive investigation releases.

We use network diagrams to illustrate the activities' precedence requirements. The network is stochastic since inter-arrival times, processing times and precedence requirements are subject to random (stochastic) variability. When an investigation 'arrives' at a node/resource, it either starts its processing immediately or it joins a queue and waits there till its processing starts. Such queues are called 'resource-queues' – they are managed according to priority rules and are subject to resource availability. Another type of waiting occurs in 'synchronization queues', where activities are delayed due to precedence constraints.

The application

According to the twofold framework mentioned, we conducted a Kaizen project, and then developed and ran a simulation of a stochastic processing network model in order to compare and evaluate the impact of various resource allocations and control alternatives. The main objective was to reduce the throughput time of regular investigations (i.e., the average time from the start to the initial report, and the average time from the start to the final report).

In addition we attempted to increase the quality of the investigations. A good quality investigation delivers good results: that is, prevents future failure. The investigation process and its final conclusions are also parameters that may determine the investigation's quality. It is possible (through no fault of the investigator), for a failure to occur after an investigation that was deemed a good one due to its high quality conclusions. So we defined the quality of an investigation according to its final conclusions. Our aim was to focus future treatment upon the root causes of the failure. A good investigation finds a few root causes. We developed a rating system for the quality of the investigations (Fig. 1). In an ideal situation an investigation indicates with a high degree of certainty, a single root cause of the analyzed failure, thereby pinpointing the treatment and future prevention of failures. The cure for the specific failure type and fleet treatment becomes more complex and labor intensive when either the number of causes increases, or the likelihood of finding the root causes decreases (or a combination of both).

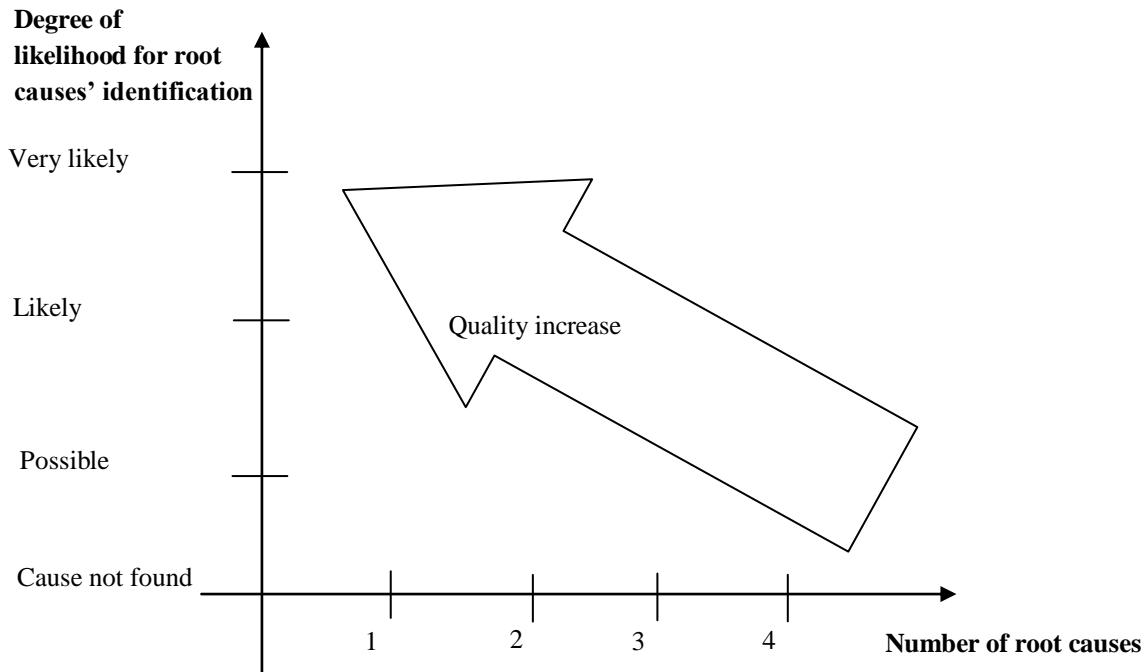


Figure 1. A rating scheme to determine the quality of the investigations.

The Kaizen project application

During the Kaizen charter meeting we set the performance's objective values as: average throughput time 90 days (130 before improvement), average throughput time till initial report 14 days (55 before improvement), and the percentage of investigations (with up to two root causes and likelihood levels of likely or very likely) 50% in comparison to 26% previously. The unique amalgamation of available information, the Kaizen framework and skillful team members from different disciplines enabled an improved investigation process. This process addresses the lack of managerial skills and experience of some of the investigators. In the previous investigation process the performance measures were almost solely dependent upon the investigator's managerial and professional skills - often leading to poor performance. Analyses of the investigation's throughput time indicated that more than half was spent in waiting queues for: resources, parts, data, or third party

decisions (this estimation is coherent with previous research of multi-project stochastic environments: e.g., Cohen et al. 2004). The new process (Figure 2) was designed to rectify such situations, for example: built-in decision points (e.g., nodes A, H, M, N in Figure 2). Accurate decisions at these points should have improved the investigation's quality, reduced throughput time and truncated delays where there was no progress in the investigation. Most of the managerial burden was transferred from the investigator to a professional promoter who collected the required information, coordinated meetings and ensured the availability of information for the investigator. It is important to note that although this was a new function created by the Kaizen team, its tasks were previously carried out by a lab technician or an investigator (e.g., in the original process). When the priority of a task was in question, that of the lab technician or investigator (e.g., polishing a specimen), took precedence over that of the promoter. Clearly, this was not beneficial to reducing the investigation's throughput time.

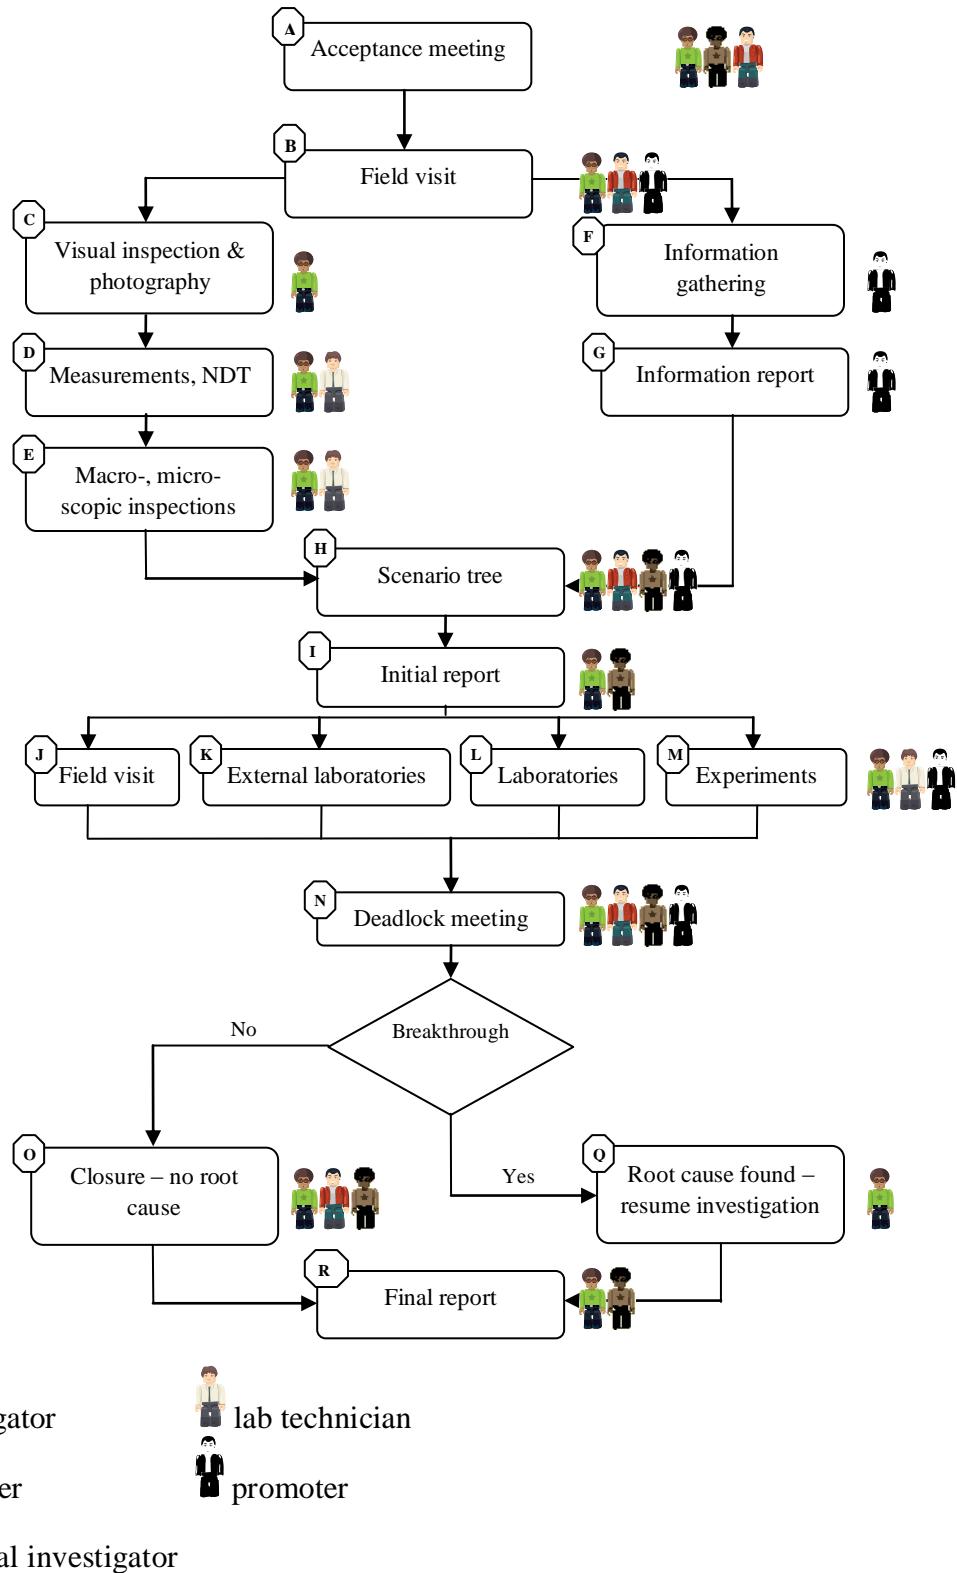


Figure 2. The new investigation process.

Hereunder is an illustrative example to describe the new process (node references found in Figure 2): During takeoff the left-hand main wheel of a fighter aircraft exploded, however the crew managed to keep the plane on the runway. The air force headquarters initiated an investigation. The acceptance procedure was applied: defined the investigation's goals and their closure deadlines and mapped information and evidence gaps (Node A). These gaps were previously identified as one of the main factors which caused delays in on-going investigations. It was found that a field visit proximate to the mishap was crucial for a comprehensive understanding of details and the development of events that led to the failure. The promoter coordinated such a visit. The investigator personally inspected the runway, the aircraft and talked with the air and technical crews (Node B). Necessary visual documentation was photographed, and evidence such as wheel parts was collected. The wheel, along with several other new and used ones, was brought to the laboratories for comparison and a series of examinations (Nodes C, D, E). At the same time, the promoter collected and analyzed information which was important for the investigator to formulate an assessment (Nodes F, G); this included: technical instructions, maintenance manuals and data records that explained the maintenance history of the failed aircraft. The aircraft's database and other operational information were analyzed to find abnormal values or patterns. It is important to note that in the original process concurrent work on Nodes C, D, E and Nodes F, G was not possible as the investigator or lab technician who, in fact, also performed the promoter's tasks, were tied up in their primary work (Nodes C, D, E).

A scenario tree technique (Node H) was used to focus the investigation on the most probable chain of events that led to the failure (interested readers may refer to Kaplan et al. 2005 for a detailed description of the scenario tree technique). An example of such a

scenario could be that the wheel hit a foreign object on the runway causing the initial crack that propagated rapidly owing to the aircraft's weight, until it created an overload of the wheel material and a total failure. However all probable scenarios need to be proven.

Current views, recommendations and future investigation directions were summed up in the initial report (Node I). This report usually enables the customer to make preliminary decisions such as: returning the fleet to flight while limiting the aircraft's weight, inspecting the wheels for damage and keeping the runways clear of foreign objects.

The investigation then followed the directions defined in the preliminary report (Nodes J to M), for example: the wheel underwent thorough laboratory inspections that supported the hypothesis that the initial damage was caused by a foreign object and propagated until failure. Thereafter a series of experiments, designed to repeat the progress of the failure, were conducted on wheels from different manufacturers. The investigator inspected the conditions and maintenance procedures at the various runways.

A deadlock meeting (Node N) would be called either 60 days after the start of the investigation (preset as 75% of the desired investigation throughput time of 90 days), or when the investigator felt that he could finalize it (the shortest of the two). Thereafter the final report (Node R) was written.

The simulation model

The construction of the model was based on the stochastic processing-network methodology that was presented earlier. Figure 3 presents the network. We assumed that release times between successive investigations followed an exponential distribution with

a mean time that was calculated, based upon historical data, throughout a two-year period. This assumption seems reasonable as the memorylessness property of exponential distribution is commensurate with the investigations which are independent. Moreover, some empirical evidence has suggested that the exponential distribution provides a reasonable fit to project activity durations. We assumed that the processing times of resources for specific activities (e.g., the duration of the promoter's coordination prior to a field visit, the duration of a field visit by the investigator etc.) were exponentially distributed. Multitasking and preemption were not allowed. It was assumed that each resource unit processed a single activity until completion, and only then would it be allowed to switch to another one. In reality there are times when a higher priority investigation preempts the processing of another one. Nonetheless, the assumption was plausible as we modeled only regular priority investigations (resource capacities, release times etc. were adjusted accordingly). Investigations that entered the processing network had to be completed (i.e., abandons were forbidden). The priority rule for investigations that wait in resource queues is FCFS (first come first served). Based on our familiarity with the process, we defined the scarce resources as: the principal investigator, investigators, laboratory technician, promoter and the SEM (Scanning Electron Microscope).

Our experimental tool was a simulation model written in C and run by a Microsoft Visual Studio C++ 2005 compiler. The simulation runs started with a warm-up period of several hundreds of investigations (equivalent to about 4 years), long enough to ensure a steady state. The transient period was discarded. The model included 47 activities and a maximal number of 20 resource units (that were allocated between the different resource types).

The maximum WIP (work in process) was defined as 800 investigations (i.e. large enough to allow long simulation runs). Each simulation was replicated 50 times.

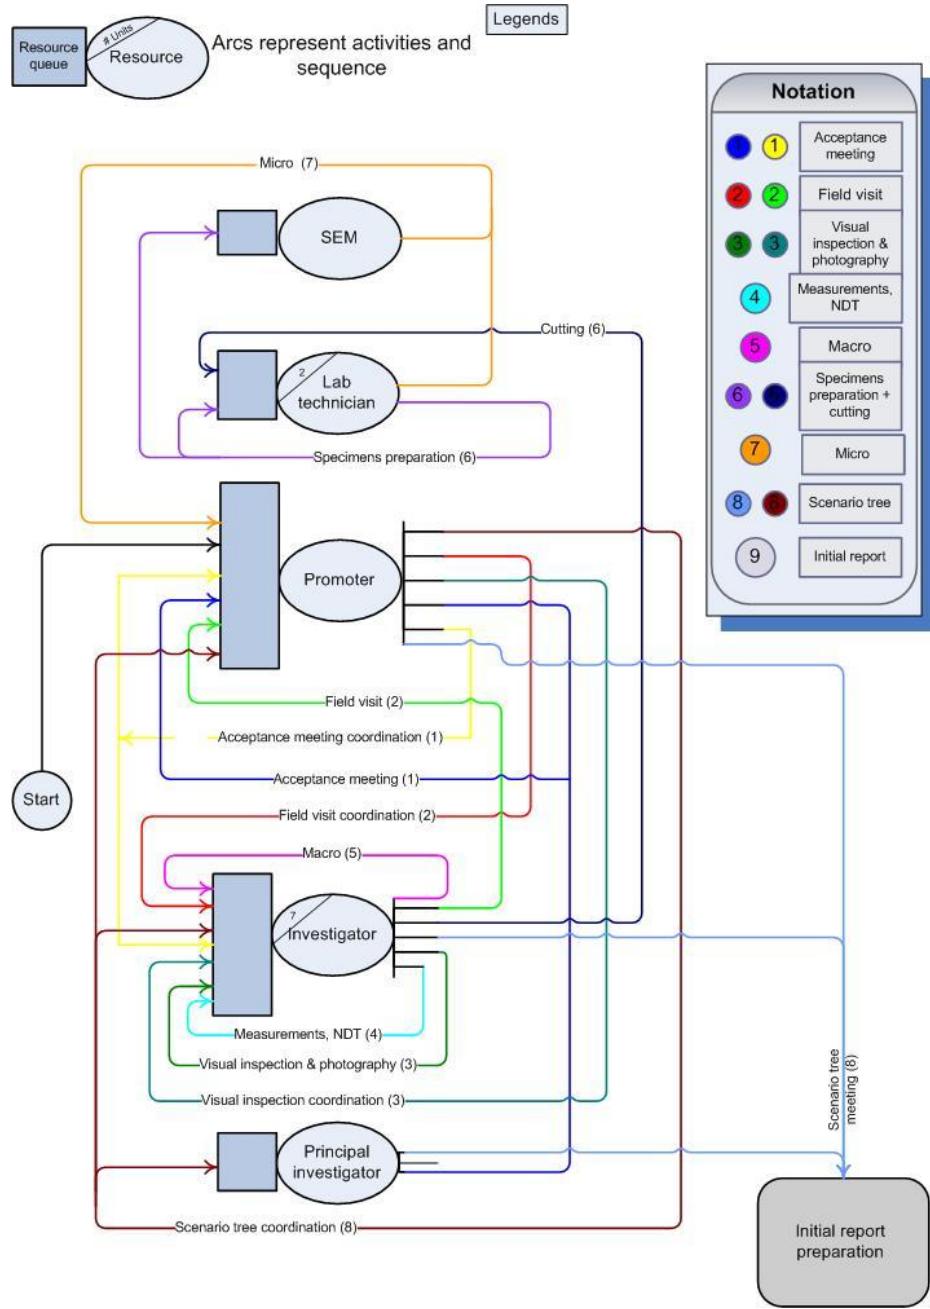


Figure 3a. The stochastic processing-network model of investigations (until the initial report).

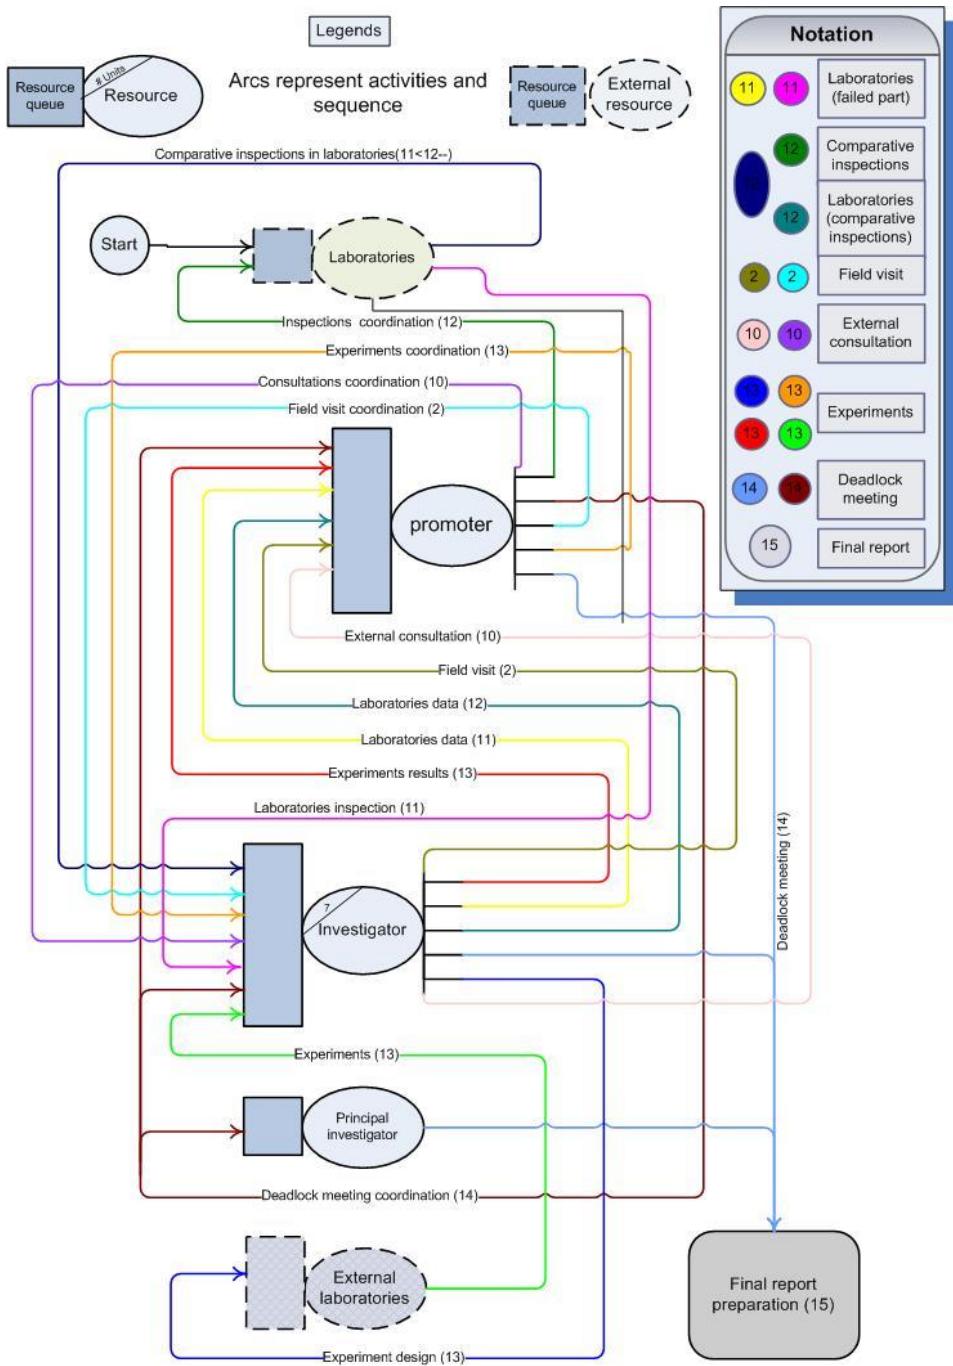


Figure 3b. The stochastic processing-network model of investigations (after initial report).

Notation and formulas

We defined x as a vector representing the resource allocation x_1, \dots, x_5 , where $x_k \in \{1, 2, 3, \dots\}$, $\forall k = 1, \dots, 5$. We needed to decide upon the allocation of all the resources where x_1 represented the number of principal investigators, and x_2, x_3, x_4, x_5 represented the number of investigators, promoters, lab technicians and SEM, respectively.

As in Cohen et al. (2004), we considered n simulation replications, each one's length was m . We defined the duration of the i^{th} investigation from the j^{th} replication as Y_{ij} ($i = 1, 2, \dots, m$; $j = 1, 2, \dots, n$). The average investigation throughput time for replication j was $T_j = \sum_{i=l+1}^m Y_{ij}/(m-l)$, where l was the length of the warm-up period. As replications were independent of each other, we assumed that the T_j 's were independent random variables with an expectation $E(T_j)$, approximately equal to the true steady state average. The overall average throughput time was then computed as $T_{mean} = \sum_{j=1}^n T_j/n$.

The mean standard deviation was defined as $\sigma = \sqrt{\sum_{j=1}^n (T_j - T_{mean})^2 / (n-1)}$. Similarly we defined $T_{ir,mean}$ as the average throughput time till the initial report. Resource utilization (i.e. the loading of a resource type k) was $\rho_k = \sum_{j=1}^n (\sum_{i=l+1}^m T_{ijk} / U_j) / (n * N_k)$ where T_{ijk} was the processing time of investigation i by resource k in replication j , N_k was the number of k -type resource units and U_j represented the steady state duration of replication j . An investigation was either being serviced or was waiting in line to receive service. The portion of this waiting time was defined as WT.

Model validation was necessary to establish its reliability. We simulated the model with existing resources after the Kaizen project (i.e. base case: one principal investigator, seven investigators, two lab technicians, one promoter and one SEM) and compared the results with real-life information (Figure 4, Case #2). Although the original process did not formally employ a promoter, the investigators and laboratory technicians acted as part-time promoters. In our base case we included one promoter to perform these tasks. This was reasonable, as one of the targets was to compare and evaluate the impact of different resource allocations and control methodologies in the new process. In the absence of documented information, the validation was based on an expert's opinion. Simulated results showed compatibility with the real ones, for example: the simulated average throughput time to complete the initial report was 42 days compared to 55 days taken for real results (23% difference); the average simulated throughput time of an investigation was 130 days compared to 129 taken for real results (0.8% difference). The average number of incoming and completed simulated investigations was 74 and 68 per year respectively, (74 and 70 respectively using real data). Simulation results indicated resource utilization levels of 97%, 73%, 55%, 52% and 48% for the promoter, SEM, principal investigator, lab technicians and investigators respectively. When we analyzed the average time volume (work hours) that each resource spent on an individual investigation, we found that 52% was performed by the investigator and 14% by the promoter. Initially, when we compared the resources' time volume investment to their utilization levels, this might have seemed surprising. As there were seven investigators and one promoter who processed several concurrent investigations, the results become self-explanatory. The investigation's throughput time analysis showed that 45% of it was spent in resource queues. This was consistent with previous research, and manifested the

improvement potential of reducing the amount of time spent in queues (Cohen et al. 2004).

Our simulation compared the impact of alternative scenarios and control methodologies. First we analyzed and learnt from the case where resources were virtually unlimited (Figure 4, Case #3). We then applied a semi-closed control: Constant Number of Projects in Process (CONPIP). According to this control new investigations are started based on a predetermined number of investigations in process (NPIP) (Anavi-Isakow and Golany, 2003). A new investigation is allowed to enter and is assigned to an investigator who has less than the stipulated NPIP concurrent investigations in process. When all investigators have been assigned NPIP investigations, the investigation waits in an external queue till it can be processed. Thus NPIP is defined as the maximal number of investigations allowed to be processed concurrently by an investigator. Multiplying NPIP by the number of investigators defines the maximum number of investigations in process.

In Figure 4, Case #4 we reduced the amount of work in process by setting NPIP to 5 (35 concurrent investigations were allowed compared to 49 in Case #2). Figure 4, Case #5 demonstrated the effect of increasing the overall number of resources, and Cases #6 and #7 explored different resource allocations.

Simulation results and insights

The results for the simulated cases are summarized in Figure 4. Increasing the number of resources (Case #3) changed the investigations' throughput time to 77 instead of 130 as in the base case. The time taken till the completion of the initial report was 22 days as opposed to the original 42, thus shortened by 48%. Resources were virtually unlimited, as

now investigations wait in queues only 8% of their stay-time compared to 45% in the base case. The optimum results that we could hope for are illustrated in this case.

Reducing the maximal number of concurrent investigations from 49 (base case), to 35 (Case #4) enabled a throughput time of 99 days. This was in accordance with Little's Law stating that reducing the amount of work in process, while keeping the same throughput rate, results in lower throughput times. However the average waiting time of an incoming investigation until the acceptance meeting was 21 days which was unacceptable. The promoter utilization was 92%, which was relatively high loading. We investigated the option of increasing the overall resource number by adding an additional resource unit in the form of another promoter (Case #5); performance improved in comparison to Case #4 and the base case. The throughput time was 87, and the time till the completion of the initial report was 22 days (similar to the result using unlimited resources). The promoters' utilization decreased to 52% (compared to 92% in Case #4), and the portion of time that an investigation spent in queues was 19%. If increasing the number of resources was unreasonable, it would have been logical to cut out the less utilized resource type which was the lab technician (51%). Simulating the system with only one lab technician led to a 130-day investigation's throughput time, and a mean time of 43 days until the completion of initial reports. This did not meet our objectives. Clearly the single lab technician became the bottleneck resource with a utilization of 99%. The Queueing Theory has manifested that being close to 100% utilization leads to a poor performance and long waiting times in queues (network congestion), in this case 46% of the time. In Case #7 we decreased the number of investigators to 6 and increased NPIP to 8 to maintain the approximate number of concurrent investigations (48) similar to the base case (49). Results indicated that on average an investigation took 88 days to

complete, and the initial report took on average 22 days. Resource utilization ranged from a high of 74% for the SEM, to a low of 49% for the promoters and lab technicians.

Case #	x_1, x_2, x_3, x_4, x_5	NPIP	$T_{mean} (\sigma)$	$T_{ir,mean}$	$\rho_1, \rho_2, \rho_3, \rho_4, \rho_5$	WT
1 – Historical data	(1,7,1,2,1)	7	129	55	Unknown	Unknown
2 - Base case	(1,7,1,2,1)	7	130 (30)	42	(0.55,0.52,0.97,0.48,0.73)	45%
3	(1,9,3,3,3)	-	77 (28)	22	(0.60,0.45,0.35,0.29,0.28)	8%
4	(1,7,1,2,1)	5	99 [#] (27)	27	(0.51,0.49,0.92,0.47,0.72)	28%
5	(1,7,2,2,1)	7	87 (29)	22	(0.59,0.58,0.52,0.51,0.79)	19%
6	(1,7,2,1,1)	7	130 (30)	43	(0.53,0.50,0.46,0.99,0.74)	46%
7	(1,6,2,2,1)	8	88 (29)	22	(0.57,0.62,0.49,0.49,0.74)	18%

Long and unrealistic waiting time before investigation's initiation

Figure 4. Summary of results for the different cases.

Simulation results showed that even under optimum conditions it took an average of 22 days to complete the initial report. Conversely a total investigation's throughput time of 90 days was a realistic goal. Although cost considerations were not defined as a constraint to begin with, it was preferable to avoid unnecessary expenses. Adding more resources amounted to raising the costs, while changing their mix (e.g., adding a promoter instead of an investigator) was insignificant due to small differences between the salaries. When the overall number of resources was kept at 12, it was advisable to reinforce the promoter at the expense of investigators (Case # 7).

Results

We improved the processes and focus in keeping with the insights gained from both the Kaizen project and the simulation model. The new investigation process followed Figure 2. We established a promoter function and manned it, and prioritized the promoter's work (she frequently worked overtime). When that did not suffice, an investigator or a lab technician was taken from his work to help. We estimated that the monthly number of man-hours invested in the promoter's job approximately equaled that of two positions. Six months after implementation, the average investigation's throughput time was found to be 92.5 days, and the initial report was completed on average within 31 days. The quality of investigations (measured by the percentage of investigations with up to two root causes, and with likelihood levels of likely or very likely), was 68% (26% before the improvement). We believe that a combination of the new quality rating scheme (Figure 1) and the improved investigation process were keeping the team more focused resulting in such a significant improvement. Results stayed relatively stable after a year (90 days throughput time, 36 days to the initial report and a quality measure of 55%).

It is important to note that in the military environment personnel changes are frequent, especially among compulsory service soldiers and officers. Therefore, during this year the failure-analysis team was subjected to changes. Two of the main influencing factors were: several urgent time-consuming investigations appeared and deferred existing ones; and two of the investigators (one of them was a senior investigator) left their positions and were subsequently replaced. Therefore the improvement is all the more impressive. The average improvement rates after 6 and then after 12 months were: throughput time

improved by 30% and the goal was met; the time to the initial report decreased by 20%; and the investigation's quality measure improved by about 136%.

Conclusions

We have presented a unique hierarchical, twofold improvement framework. We found that this improvement framework was very natural. The integration of a Kaizen project and stochastic networks modeling were found to be supplementary of each other, and very useful in improving the performance of an air force failure-analysis team that investigates aircraft failures.

The intensive effort made by the participants of the Kaizen project enabled a thorough examination of the investigation process. It was mapped, improvement opportunities and inefficient procedures were identified, and a new process was constructed. This new process is different from the previous one. For example, it was restructured to include built-in decision points for treating problematic investigations. Machines and responsibilities were relocated and new performance indicators were defined. It draws from the common industrial engineering principles of limiting the amount of work in process and balancing resource utilization.

We demonstrated that controlling the number of investigations in process was a simple, yet effective management policy. Its implementation required minimal effort and it prevented high loading and the effects of congestion.

Our improvement framework applies not only to military environments and to failure analysis teams but also to: job shops, service organizations and multi-project environments where the projects may be classified into several types. These

environments are found in organizations that process software maintenance projects, product development in the chemical industry, and maintenance or retrofit projects in aerospace companies (Leung 2002; Adler et al. 1996; Gemmill and Edwards, 1999).

We conclude with two observations that provide insight for those who plan a similar improvement effort:

1) A Kaizen project is a very efficient tool for achieving change in a short time. Its insights should be applied immediately as they usually produce better processes and improve the results. Nevertheless in complex systems there is a trade-off between the short-term effort and the ability to make a thorough analysis. In order to gain more, sometimes much more, we recommend performing steady-state modeling and analysis.

2) There is a trade-off between the loading of a system and its steady-state performance. In loaded systems the time-based performance decreases with the load. It is a managerial decision to allocate resources and set the loading level in such a way that the system operates efficiently (that is resources are not underutilized), and the overall performance is satisfactory. In heavily loaded systems it might be wise to limit the amount of work in process. This comes at the price of delaying incoming work, but the reward is a superior performance.

Our model has several limitations. It is difficult to capture quantitatively the impact of improvements such as moving equipment, a professional promoter or built in decision points. The simulation was designed to compare between different resource allocations, loading and control methodologies which it did satisfactorily. It is worth applying common sense changes that are likely to improve the process (such as equipment relocation) and with time, the effect of such improvements on the process throughput

could be measured quantitatively. The model did not take into consideration certain factors such as the experience of the investigators, and the external resources needed to carry out the investigations. While results are satisfactory, it would be interesting to make a more accurate model that includes these factors. Such a model may provide more insight into aspects of the team's members and the effect that the external resources has on the process. Our model also assumes that activities' durations follow exponential distribution. Some empirical evidence and the application's results have shown that the exponential distribution is a reasonable assumption for project activities. However, the model can easily be adapted to other distributions including empirical ones. Our observations and insights in this section will still be valid for any distribution selected.

We envision the use of the simulation tool in the future for the examination of: process improvements, the impact of resource allocations and the impact of adaptations in activity durations.

Acknowledgements

Many were involved in the work described in this article. I would like to acknowledge the contribution of Rafael Kimchi, David Kazir, Jacob Shmerler and Ofer Levy. I thank also Professor Avraham Shtub and Professor Avishai Mandelbaum for their valuable inputs to this work. Of course, the responsibility for any errors found in the article rests solely with the author.

References

Adler, P.S., A. Mandelbaum, V. Nguyen, E. Schwerer. 1996. Getting the most out of your product development process. *Harvard Business Review* March-April 134-152.

Anavi-Isakow, S., B. Golany. 2003. Managing multi-project environments through constant work-in-process. *International Journal of Project Management* 21 (1) 9-18.

Berger, A. 1997. Continuous improvement and Kaizen: standardization and organizational designs. *Journal of Integrated Manufacturing Systems* 8 (2) 110-117.

Bradley, J.R., J. Willett. 2004. Cornell students participate in Lord corporation's kaizen projects. *Interfaces* 34 (6) 451-459.

Brunet, A.P., S. New. 2003. Kaizen in Japan: an empirical study. *International Journal of Operations & Production Management* 23 (12) 1426-1446.

Cohen, I., A. Mandelbaum, A. Shtub. 2004. Multi-project scheduling and control: a process based comparative study of the critical chain methodology and some alternatives. *Project Management Journal* 35 (2) 39-50.

De Castro, Paulo M.S.T., A.A. Fernandes. 2004. Methodologies for failure analysis: a critical survey. *Material and Design* 25 117-123.

Emiliani, M.L. 2005. Using kaizen to improve graduate business school degree programs. *Quality Assurance in Education* 13 (1) 37-52.

Gemmill, D.D., M.L. Edwards. 1999. Improving resource-constrained project schedules with look-ahead techniques. *Project Management Journal* 30 (3) 44-55.

Kaplan, S., S. Visnepolschi, B. Zlotin, A. Zusman. 2005. *New tools for risk analysis: Anticipatory Failure Determination (AFD) and the theory of scenario structuring*. USA: Ideation International Inc

Leung, H.K.N. 2002. Estimating maintenance effort by analogy. *Empirical Software Engineering* 7 157-175.

Swank, C.K. 2003. The lean service machine. *Harvard Business Review* 81 (10) 123-129.

Thomas, J. M. 1989. Industrial experience with structural failure. T. Wierzbicki, N. Jones, eds. *Structural Failure*. John Wiley & Sons, New York, 511–545.