

Learning from Failures

Darren Dalcher*.[†] and Colin Tully

*Software Forensics Centre, School of Computing Science,
Middlesex University, Trent Park, Bramley Road,
London N14 4YZ, UK*



Practice Section

This paper addresses the topic of failures in computer-based systems. It does so in three stages: first it describes failure cases in a single domain, that of computer-based ambulance despatch systems; then it discusses some of the features of those cases; finally it considers the importance of reporting and analysing failures, in order to discover their causes and to learn lessons that will help eliminate those causes of failure in the future. Failure, it is argued, provides a valuable learning opportunity that may lead, when recognized, to better practice and future success. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: software development failure; despatch systems; failure causes; learning from failure

1. INTRODUCTION

This paper addresses the topic of failures in computer-based systems. It does so in three stages: first it describes failure cases in a single domain, that of computer-based ambulance despatch systems; then it discusses some of the features of those cases; finally it considers the importance of reporting and analysing failures, in order to discover their causes and to learn lessons that will help eliminate those causes of failure in the future.

The domain from which the failure cases are drawn is that of computer-based ambulance despatch systems. Both medicine and computer systems boast rich records of error and failure, a proportion of which on each side have had fatal results, so it is not surprising that an area of intersection between the two is fertile ground for the purposes of this paper. Both groups of practitioners show strong and chronic symptoms of resisting open enquiry into failure – in contrast, for instance, to the airline industry. The greater demand for

public accountability in the case of health services, however, has meant that there is at least a regular practice of post-mortem investigations and some history of published case studies of the grosser errors. Even so, prevalent characteristics of the medical culture (elitism, blame etc.) conspire to minimize effective learning from failure, and to ensure that errors continue to thrive healthily. Deaths attributable to preventable medical error were recently estimated at 120,000 per annum in the USA, and injuries from preventable medical error are estimated to occur in just under 3% of cases, at an annual cost of \$1 billion just in terms of wasted bed-days.

Computer-based systems have a much better record in terms of deaths: known fatalities are a tiny proportion of medically-induced deaths. In terms of cost and scale, however, the story is different. In the USA in 1995, the cost of effort ploughed into computer projects that were subsequently cancelled, plus the cost of project overruns, was estimated by the Standish Group at a spectacular \$140 billion – and that is only a fraction of the total cost of failure and waste in computer-based systems. Whereas preventable medical errors occurred in fewer than 3% of cases, computer project cancellations and overruns accounted for over 50%

* Correspondence to: Darren Dalcher, Software Forensics Centre School of Computing Science, Middlesex University Trent Park, Bramley Road, London N14 4YZ, UK

[†]E-mail: d.dalcher@mdx.ac.uk



of the total budget for software projects and around 20% of America's total investment in developing and acquiring software (which itself runs at a bit under 2% of national income). Despite this, computer systems practitioners are highly effective in avoiding post-mortem inquiries into their inadequacies, and the 'profession' appears to learn even less, from one to decade to the next, than the medical fraternity.

A number of failure cases in computer-based ambulance despatch systems have been reported. During the 1990s, for instance information about five such cases in the USA was reported in *Software Engineering Notes*: Norfolk, Virginia (Neumann 1991); Chicago (Neumann 1992); San Francisco (Neumann 1996); Iowa (Neumann 1997), and Tacoma, Washington (Neumann 1998). The most striking and best documented cases, however, are probably those that occurred in the UK (London) and Australia (Melbourne). The available information on those two cases is presented in Sections 2 and 3 of this paper. Readers may feel they are already familiar with the stories, but it is believed that the cases are set out here in much greater detail than has previously been attempted in the open literature. The amount of detail is not necessarily of value in itself (though it makes astonishing and compelling reading); the value lies in the breadth of causal analysis which it makes possible, and which is presented in Section 4. Section 5 offers some concluding observations about the importance of learning from failure.

2. FIRST CASE: LONDON AMBULANCE SERVICE

The story of the London Ambulance Service (LAS) began in the mid-1980s. The background to the story had two key features. First, the LAS was falling short of recently established national standards of performance for ambulance mobilization and arrival times. Second, the introduction of the so-called 'internal market' to the National Health Service in the UK, together with chronic shortages arising from long-term under-investment, led to relentless pressure to squeeze budgets while simultaneously improving performance.

The London Ambulance Service (LAS):

- is the largest ambulance service in the world;
- covers an area of just over 600 square miles;

- serves a resident population of about 6.8 million, boosted by commuters and visitors to as many as 10 million;
- at that time had over 300 accident and emergency ambulances (two-thirds of which might be in use at any one time), operating from 70 ambulance stations, with a central control room at LAS headquarters;
- devoted about three-quarters of its budget, and more than half its 2750 staff, to emergency services;
- received a daily average of over 2000 calls, resulting in ambulances attending around 1200 incidents;
- transported a daily average of over 5000 patients, about 1400 of which were emergency cases.

In the early 1980s, a national standard of three minutes was established for ambulance despatch in response to emergency calls. The LAS decided that a computer-based system was essential to enable them to meet that standard. It was further decided to commission a system to handle not just despatch but the whole range of control room functions, involving mobile radio (voice and data) transmission. Primary requirements were to provide:

- a computer-aided despatch system, with incident record-keeping;
- an automatic vehicle location system, with mobile data terminals, able to minimize response time by positioning units optimally and to track long-term asset performance;
- a computer map display system.

In May 1987, after a year of delay, a three-year £2.5 million contract was awarded to IAL (International Aerdio Limited, a subsidiary of British Telecom) and CGS (Cap Gemini Sogeti), to provide a limited despatch system, without any mobile data capability. In 1989, the specification was extended to include mobile data.

In October 1990, the project was terminated, after two tests investigating peak load performance failed. The project was severely behind schedule, and had already accumulated costs of £7.5 million (an overrun of 200%). The LAS sought damages from the prime vendors, claiming they did not understand the requirements. The vendors counter-claimed, alleging that specifications were ambiguous, lacked clarity and were subject to constant change. An out-of-court settlement was eventually reached.



An independent assessment of the failed project was undertaken, sponsored by the LAS chief executive and conducted by Arthur Andersen. It was based on the original statement of requirements, excluding mobile data capability. It recommended acquiring a packaged system on a turnkey basis, estimating a cost of £1.5 million and time-to-field of 19 months. If no suitable packaged solution could be found, it warned that the cost and time for developing a bespoke system would be significantly higher. A number of packages were evaluated, including some operated by other UK ambulance services, but none was judged acceptable. The LAS then embarked on specifying requirements for a new bespoke system.

The proposed system would automate not only manual tasks but also decision-making. Once information from callers had been fed into the system it would take over, allocating and mobilizing ambulances, interacting with crews, and monitoring vehicle positions and performance. The control room would no longer communicate with vehicles by voice radio, or with ambulance stations by telephone, as the system would handle all communications. Intervention by controllers would be required only in exceptional cases, such as failure to despatch an ambulance within eleven minutes, failure of an ambulance crew to acknowledge a message, or an ambulance going 'the wrong way'. In order to maximize benefits, it was decided (a) to achieve them as fast as possible by planning for big-bang rather than phased implementation, switching to full functionality overnight, (b) to achieve economies by using hardware acquired or planned during the previous project.

The specification was completed in early February 1991, less than four months after abandoning the old system. It required the system to be operational in eleven months (8 January 1992). The document was very detailed and prescriptive on how the system should operate, leaving little opportunity for suppliers to incorporate their own ideas and experience or optimize the use of technology. Most of the work (and subsequent bid selection) was carried out by a contract analyst and the systems manager, while a major management restructuring was under way.

The LAS fully realized that the system would require a 'quantum leap' in technology. Combining sophisticated allocation and monitoring with judgement and decision-making tasks involved a major

intellectual and technical effort. They were prepared to attempt this ambitious undertaking despite their stated belief that the previous project, for a simpler system, failed because the vendors did not understand the complexity inherent in the environment. Indeed, insiders report that at this time there was a pervasive sense of optimism, excitement and urgency in the systems team, no doubt fed by the large benefits (operational efficiency and reduced costs) promised by the proposal.

A common factor in successful implementations is that users feel they share the benefits and have a sense of involvement and ownership. Achieving such user attitudes is often critical, and will almost certainly be so in a 'quantum leap' with very high stakes. Based on recent history of industrial conflicts and attitudes towards structural change within the LAS, there was every reason to assume that another major change to the organizational culture would require tact and careful management. The new system would make radical changes to working practices and was bound to introduce anxiety and possibly conflict. As an example, it was proposed that ambulance crews could be allocated to incidents away from their home stations: that would mean that they could spend much of their time covering unfamiliar areas and that, at the end of a shift, it might take them a long time to drive back to their home stations – amounting in effect to unpaid overtime.

There was, however, no attempt to sell the system to the users, or in any other way to reduce resistance to change. The specification, and subsequent guidelines on new working practices, were written without input from ambulance crews, control room staff or unions.

Standing instructions of the Regional Health Authority required all new systems to be put out to tender. An advertisement accordingly appeared in the *Journal of the European Communities* on 7 February 1991. Thirty-five companies expressed interest in bidding and were sent the requirements specification.

Most potential bidders questioned the eleven-month deadline for development and implementation. They were told it was non-negotiable. Half the companies did not proceed. The remaining 17 provided proposals for all or part of the system. Several of them proposed implementing a basic system by January 1992, with full functionality available in early 1993. The LAS's initial screening eliminated all



proposals that did not accept the January deadline for the full system.

The overt criteria for evaluating the handful of remaining contenders included quality and performance factors such as resilience, functionality, flexibility and response times. In practice little attention was paid to these: the effective criterion was price, and moreover there was a secret price target of £1.5 million.

While there is no clear evidence, it is reasonable to assume that the target price was adopted from the Arthur Anderson recommendations. They were made, however, in an entirely different context, assuming a package solution, without a mobile data capability, and with a 19-month implementation schedule. The failed first project, for a simpler system, had lasted over three years before cancellation, and exceeded its £2.5 million budget by a factor of three. Experience, analysis or just plain common sense should have made it obvious that the Arthur Anderson estimate was inappropriate in the new conditions. On the contrary, the early commitment to the £1.5 million price tag seems to have infected the decision-making process, so that it was adopted without challenge and then defended by groupthink. Set out in cold print, it is tempting to condemn the team as mad, but it is a recognized phenomenon that collective rationalization, illusions of invulnerability, self-censorship, and the consequent absence of dissent can lead groups to ignore external opinions and alternative solutions and to adopt positions that are not rationally defensible.

The clear winner was a bid for £937,463 for the complete system to be implemented by the mandated deadline. The nearest competitors came in at £1.6 million and £3 million. The winning bid was from a consortium led by Apricot Computers (a UK hardware supplier owned by Mitsubishi). Apricot would supply networked PCs and a fault-tolerant file server. Systems Options (a small UK software house) would supply the despatch system software, based on their Wings geographical information system. Datatrak would supply the automatic vehicle location system. Solo Electronic Systems would supply radio interface systems and mobile data terminals, which were to have been used in the previous abandoned information systems.

The quoted price for the despatch system software, the main hub of the system, was £35,000. The price represented less than 4% of the overall price of what was clearly a software-intensive system, a

grossly mistaken estimate given the complexity of the requirements. The significance of the software was clearly understood, at least by the LAS and Apricot. The LAS expressed concern by warning Systems Options that the failure of the previous system was largely caused by the supplier's software house not being able to cope with the complexity of the system. Apricot, who led the consortium at the bid stage, refused to assume project leadership because of the dependence of the entire system on the quality of the software. Under strong pressure from the LAS, Systems Options agreed to accept leadership of both the consortium and the project. It later emerged that System Options had been reluctant to bid for the system in the first place, but had been persuaded to do so by Apricot. An earlier bid by the two companies for a far more basic system for the Cambridgeshire Ambulance Service had been rejected because of the lack of technical understanding exhibited in that bid.

The required system was significantly bigger and more complex than anything any of the consortium members had handled previously. The main experience of Systems Options was in developing government administrative systems. LAS felt reassured by the fact that they had delivered systems for police and fire services, but those had been administrative systems, and the company had no experience in developing real-time, safety-critical command and control systems. The letters of reference suggested that Systems Options was overstretched on its current contracts and had been having trouble delivering simpler and less demanding software on time. A letter from the Staffordshire Fire and Rescue Service expressed grave concern over the ability of Systems Options to cope with the LAS project. Those claims were not investigated or even acknowledged by the LAS procurement team.

The procurement team was small and inexperienced in terms of both technical knowledge and acquisition procedures. Procurement guidelines for the Regional Health Authority stated that the lowest tender should be accepted unless there were 'good and sufficient reasons to the contrary'. No attempt was made by the team to investigate whether such reasons existed, why the price was below the internal target price, or why it was so much lower than even the next cheapest. As with authorship of the requirements specification, bid evaluation was undertaken primarily by the systems manager (an ambulanceman who had taken the job temporarily



on the understanding that he would be replaced by a qualified systems manager at some point) and the contract analyst (who had five years experience with the LAS, largely on the failed original project).

The LAS higher management was still in the course of restructuring. Normal practice required it to audit the selection process before accepting the recommendation of the internal procurement team. The audit, performed with the involvement of an external assessor, indicated that the process was risky and required close management attention, but the selection was approved.

The management restructuring was eventually completed in April 1991. Senior and middle management ranks were slimmed by 20%, the four divisions were reduced to three, and a large number of experienced staff left. The rapid changes eliminated any stability within the organization. The remaining managers were unmotivated, stressed, anxious and oppressed by job insecurity. To compensate for the loss of some high-calibre managers, others were promoted or shifted sideways to new positions (rising-to-level-of-incompetence syndrome?). Directors were left with a great span of responsibility and power. The internal climate was influenced by a minimal investment in staff or managerial training and little scope for career advancement. In the absence of consultation with employees and the unions, industrial relations deteriorated, resulting in resentment, lack of trust, low staff morale and increased absenteeism. The reputation of the new CEO for 'sorting out' troublesome employees and interfering trade unions did not help.

On 28 May 1991, with just over seven months to the deadline, the Executive Board ratified the decision on the contract. Learning from the lesson of the previous failure, the LAS required the consortium to provide a complete systems design specification detailing how the final system would operate, as a demonstration that they truly understood the requirements and could provide an adequate solution. A small starter contract was awarded for the development of this design specification, which was carried out during June and July.

In developing the design specification, the consortium opted to use Windows 3.0 for the user interface and Visual Basic for screen dialogues. That decision, which was not specified in the design specification, involved a trade-off against performance. Visual Basic is essentially a prototyping tool, typically used in small, non-safety-critical applications to generate

and test screens rapidly. Code generation, however, is inefficient, and the resulting performance is slow, so that it is often necessary to follow prototyping by fine-tuning to increase speed and reduce memory consumption. In the LAS case that did not happen. Further, the inefficiency of the generated code was compounded by the fact that it was to run on slow 486 machines operating at 25 MHz. The resulting low performance could easily account for the bulk of the mobilization time allowed by the national standard. The allocation algorithm was complex and required high processing speeds to meet performance requirements. An increase in the number of incidents would cause on the one hand an increase in the volume of calculations; on the other hand it would cause a decrease in the number of ambulances available, leading in turn to an increase in the average distance between an incident and the available ambulances, and an increase in the complexity of the calculations. Those increases in the volume and complexity of the calculations caused a slow-down in the system, the severity of which would depend on system processing speed.

While the design specification was being developed, the LAS was considering its management strategy. Among concerns that were noted and minuted, but never followed up, were:

- no full-time LAS staff member, and no project review group, had been assigned to the project;
- the draft project plan left no time for review or revision;
- the six-month schedule was considerably shorter than the 18 months that other services had needed for less complicated despatch systems.

The PRINCE project management method was selected to guide the development effort. As neither LAS staff nor the consortium had much project management experience, or knew anything about PRINCE, a special course was arranged to acquaint them with the methodology. Little use was ever made of PRINCE, however, nor were any alternative methods or tools adopted in its place.

The LAS Director of Support Services proposed inviting one of the losing bidders to perform independent quality assurance on the project. This proposal was rejected, on the grounds that quality assurance was the contractor's responsibility. All quality assurance activities were assigned to Systems Options.



Although there was no formal sign-off of the design specification, and although it contained omissions and undeveloped sections (such as interfaces with other components and systems), the LAS accepted it. The full contract was then formally awarded to the consortium on 8 August 1991, leaving exactly five months for development and implementation. The contract with Solo Electronics for communications hardware, the essence of which was to salvage key features from the earlier abandoned project, was not signed until 16 September, four months before the delivery deadline, and three months after the award of the system design contract.

During early project meetings it became clear that Systems Options was unable to manage the entire project, as it was struggling to manage its own software effort. As already noted, no experienced project manager was available either in the LAS or in the consortium. Despite their inexperience, the LAS Director of Support Services and the contract analyst undertook to add overall project management to their list of duties. The only scant attempt at project management was made by those individuals; none was attempted by any of the consortium members. There was no quality assurance, change control, configuration management or test planning.

Software deliveries were all late. Systems Options blamed the delays on the two months needed to develop the design specification, and on the delay in signing the contract for the communication equipment with which its software needed to interface.

The LAS hired a new systems manager in October. He promptly arranged a formal project review, which reported the following key findings to LAS senior management in November:

- Little time had been allocated for review.
- Increased quality management was needed.
- Troubling technical problems were occurring.
- The January operational date should be maintained if only to keep pressure on the suppliers.

Publicity and external pressure resulted in parliamentary questions. The government expressed the view that the consortium was fully qualified for the task at hand. The LAS management decided to take no action.

By mid-December it was clear that the 8 January deadline was not achievable. The despatch software was incomplete and untested, the radio interface system was yet to be delivered, the design

of the data terminals and their positioning in the ambulances required changing, the vehicle location tracking system was not fully installed, the data provided by the tracking system was reported to be neither accurate nor reliable, and no training of ambulance crews or control room operators had been initiated.

By late December, a new schedule for a three-phase delivery strategy was agreed, with final delivery scheduled for 26 October 1992:

- Phase 1 would deliver the call-taking function. Recorded calls would be printed out for use within the existing manual procedure.
- Phase 2 would deliver an improved allocation function. Allocators would have terminals displaying call information. Ambulance locations would be automatically tracked by the system, which would also issue despatch instructions to vehicle-based data terminals. Decisions on selecting and mobilizing vehicles would still be made by the allocators.
- Phase 3 would deliver the fully automated system. All allocation tasks would be handled by the system, and allocators would no longer be used.

The replacement of big bang by phased implementation gave rise to additional requirements and problems. For example, the need to print call details in phase 1 required the introduction of printers – not part of the original ‘paperless office’ vision. Their introduction led to problems such as screens locking up, server failures, and, on one occasion, the loss of calls when a printer was switched off erasing the contents of its buffer memory.

During phases 1 and 2 different system versions were delivered at different times to the three operating divisions and then enhanced in an uncontrolled fashion with different changes. Thus the systems in use in each division were not functionally or operationally equivalent. This led to great confusion and difficulties in the control room.

Meanwhile, a staff attitude survey, commissioned by the LAS and carried out by Price Waterhouse in January 1992, revealed that only 13% of staff believed that the LAS was providing a quality service, 10% felt they knew what the LAS plans were for the future, and 8% believed that management listened to them.

March 1992 was the phase 2 deadline. It was marked by a major system crash, resulting in a



30-minute delay to emergency calls, delayed ambulances and lost incident reports. Some LAS local area officers ordered a switch back to voice radio communication until the computer was able to match its performance. The union area officer called for an urgent public enquiry. An LAS spokesman reassured the press that these are normal teething problems and no one has anything to worry about.

A new review conducted by the LAS systems manager revealed:

- the radio interface system was failing almost daily;
- no volume testing of the whole communications infrastructure had been done;
- there was no way of tracking changes to the system;
- ambulance crews and control room staff were not trained.

No recommendation on future progress or cancellation of the project was made in the review report, which did not elicit any response from LAS senior management.

A number of letters from computer consultancies and safety experts, warning that the system was 'totally and fatally flawed', reached the LAS and government ministers. In April 1992 hundreds of callers failed to get through to the ambulance service. The LAS management blamed the public for making too many calls and clogging lines. The LAS board was presented with a formal vote of no confidence in the system by staff from one of the three divisions. Additional complaints and expressions of concern throughout the rest of the year were brushed aside by LAS management. The purpose of phased trials, it was explained, was to highlight problems, and there was no need for concern as the final system would work.

Some new problems began cropping up at this stage. These included:

- locking up of terminals (prompting an instruction to staff to simply reboot their screens);
- overload of communications channels;
- inaccurate information provided by the vehicle location system;
- crews using different ambulances from the ones allocated by the system;
- slowness of the system;
- failure of the system to identify the nearest available ambulance;

- failure of the system to identify every 53rd vehicle in the fleet.

As a result, calls were getting lost in the system, while others failed to reach ambulances. Vehicles were being assigned incorrect codes, while others were missing from the system. Call waiting queues and the ever multiplying exception message queues simply scrolled off the top of the screen. At no time during phases 1 and 2 was the system stable or operational across all three divisions.

In the background intense government pressure was being applied on the LAS to reduce its seriously overspent budget. Accordingly it was decided that planned capital improvements and preventive maintenance on emergency vehicles would have to be delayed.

Phase 3 was scheduled to come on line in the early morning of Monday 26 October 1992, at which time the software was known to have two errors that could cause the system not to function, 44 that could cause operational problems likely to affect patient care, and 35 minor problems – a total of 81 errors. No stress testing had been done on the full system. No back-up plans were in place in case of system failure. The LAS did not even have its own network manager, having relied on the contractors to rectify all previous problems.

By mid-morning callers were often having to wait over 30 minutes to get through. Ambulances arrived late, or not at all, or two at a time. These problems often meant that a single incident generated multiple follow-up calls. The system was slow at logging acknowledgements, and there was therefore a proliferation of warnings that ambulances were not meeting arrival deadlines. These warnings, together with other error and alert messages, and new calls, wiped old calls off the screens, even if they had not been dealt with; the old calls were then lost to the system, again giving rise to follow-up calls. The increasing volumes of system traffic generated by system failures, and the rapidly growing exception queues, were creating overload conditions and slowing the system down. Screens and mobile data terminals began locking up. At one point the exception queue was cleared of its contents, erasing all calls. But that was not all!

The map system failed to recognize certain roads, forcing operators to scramble for maps and despatch ambulances by telephone unbeknown to the system – even though staff were under instruction to



minimize voice communication. Many indeed felt too threatened to make any attempt to rectify problems by voice communication.

The system depended on radio for data feeds on ambulance location and status. The limitations of radio communication in an urban area with tall buildings were not taken into account. No attempt was made to specify or test the system's response to incomplete or incorrect data, and the consequent need for perfect information at all times proved to be a critical millstone. Frustrated staff, unable to communicate with base, switched to their radios, which used the same frequencies as the data communication system, adding to the already clogged airwaves. The number of ambulances that could be traced by the system declined, resulting in more complex allocations, and in delays which further constrained the system. This constituted a positive feedback loop which in the right circumstances would ensure the situation escalated out of control.

A specific aspect of the poor performance of Visual Basic programs is that filling the screens can easily take several seconds. To overcome this, control room staff preloaded all the screens likely to be needed at the beginning of a shift and used Windows to move between them as required. That added to the demand on memory, slowing system performance, and led to extra clutter on controllers' screens, slowing their performance.

The overnight reorganization of the control room eliminated the familiar structure of the room and split up divisional teams who had worked together. Teams with local knowledge and experience of local traffic, shortcuts and hospital locations, who used to share sector desks, were thus separated and placed in unfamiliar positions. These changes, together with the new operating procedures, put control room staff under great stress and made it extremely difficult for them to intervene and correct problems.

Ambulance arrivals within 15 minutes were achieved in only 17% of incidents, against 30% under the phase 2 system, 65% prior to phase 1, and the requirement of the national standard of 95%. The average ring time for each call peaked at 11 minutes, with many callers waiting 15 minutes.

Londoners are lucky that no major incidents capable of stretching the system beyond its limits occurred in the 36-hour period while the system was in a state of total chaos. During that time, over 900 complaints were received, some patients were kept waiting as long as eleven hours, more than

40 patients may have died as a result of delays (depending on which estimates you accept) – and a team of public relations experts were busy promoting the new technological efficiency of the service.

At 14.00 on Tuesday 27 October, a decision was made to take the phase 3 system down and replace it with the semi-manual phase 2 system. More ambulance crews and control room staff were brought on duty in an attempt to cut the backlog. The computer team initiated an attempt to investigate the errors and relaunch the system.

On Wednesday morning, the LAS claimed that no serious disruptions were caused by the computer system. That afternoon, however, the LAS chief executive resigned. It was later argued that he was under extreme pressure to improve performance by the end of that year and had thus been forced to rush the computer implementation. Following intensive pressure from families of patients, the unions, and a stormy session in Parliament on Thursday evening, the Health Secretary announced an official inquiry into the affair. The unions called for the system they described as a 'lethal lottery' to be shut down in the interest of patients. Their plea went unheeded.

The reduced functionality system operated for the next eight days. On 4 November 1992, it in turn started slowing down, and it finally locked up altogether at 02.00. Operators rebooted the system, to discover that it was still frozen. The back-up system (designed to interface with the full-functionality system, but not with the semi-manual system) did not come online. There was then no alternative but to revert to the manual methods of the 1980s. Performance figures for mobilization and arrival recorded dramatic improvements, usually by a factor of around two.

The official inquiry report released on 1 March 1993 was scathing. It described the chaotic management and total lack of planning and technical oversight which led to the disaster, and called for a complete revamping of the LAS and the way it conducted its business. During a press conference marking its publication, a member of the inquiry team commented that the LAS 'went through every mistake in the book'. Following publication, the LAS chairman resigned, saying 'We caused a considerable amount of anguish to the people of London. We failed to deliver the service we could'.

In addition to failings already identified above, the inquiry's findings included the following:



- Despite a claim by LAS management that an excessive number of calls were made on 26 and 27 October, the number was found to be only 6% above the average for October. Owing to the frequent lack of response, there was a higher-than-usual number of call-backs, but even so the total number of calls was within the predicted upper limit. The total number of patients transported on the two days was less than the daily average for October.
 - The LAS management claim that ambulance crews deliberately sabotaged the system was shown to be false. Despite losing the personal touch of having a familiar human voice dispatcher, crews seem to have accepted the initiative, albeit reluctantly, and co-operated. Datatrak, the primary contractor for the vehicle location sub-system, stated that resistance at the LAS was no greater than they experienced at other organizations.
 - The LAS management claim that users and crews had been inadequately trained in how to use the system was supported. Training provision was divided between LAS and Systems Options. In the case of control room staff it comprised just a two-day familiarization exercise. All training was scheduled to be completed by the original implementation date in January 1992. Despite the ten-month delay between that date and the eventual implementation, and the system changes during that period, no further training was given. Operators were thus not in a position to spot problems or be able to override them.
 - Eleven causes for a lack of perfect information (necessary for the allocation algorithms to function correctly) included radio black spots, failure to press correct buttons, noise corruption, wrong call signs, and too few operators. No consideration had been given to the possibility that the assumption of perfect information would not hold.
 - Delayed and duplicated allocations resulted in a build-up of exception messages, which required personal attention from operators and prevented them from answering new calls. As the messages and lists built up, the system slowed down. The linked slowdown of system and operator performance caused an increase in call-backs, and therefore yet further exception messages.
- The positive feedback effect overwhelmed the system and its operators.
- The complete system had never been performance-tested to predict its behaviour under extreme or atypical conditions, such as a disaster scenario, incomplete or contradictory information, operator errors or high volumes of exception messages.
 - The demand to reduce costs and budgets was responsible for more ambulances than normal being out of service due to degradation and lack of repairs, compounding the pressure on the system.
 - An error in Systems Options software meant that file server memory was not released after each ambulance mobilization, leading to the steady consumption of all available memory.
 - While Systems Options 'rapidly found themselves in a situation where they became out of their depth', it was also judged that 'within the time constraints imposed on the project and the scope of the requirements, no software house could have delivered a workable solution'.
- A report from a cross-party committee of MPs described the system as ramshackle and undermined by absenteeism, a culture of blame, cumbersome processes, lack of trust between management and staff, and lack of technology.
- Following the reintroduction of the manual system, and the consequent inability to meet performance standards, further incidents led to deaths over the next two years or so. One in particular gave rise to public outrage and intensive media coverage. Nasima Begum, aged 11, who had a kidney condition known as relapsing nephrotic syndrome, died of pulmonary oedema after her family had telephoned four times for an ambulance and had waited for 53 minutes. Doctors at the Royal London Hospital said that 'Nasima could have been saved had urgent treatment been more forthcoming'. The ambulance that should have saved Nasima had been redirected to treat a patient with a sore toe.
- The third attempt at computerizing the LAS is currently under way. The concept of phased delivery through a well-structured, well-planned series of projects, each delivering some functionality or certain sub-systems, has now been adopted. A series of prototypes refined the call taking, gazetteer, call viewing and management alert system functions. This approach is intended to encourage belief in



the value of change and the possibility of ultimate success.

The first of the staged deliveries won a British Computer Society Information Systems Management award. Allocation decisions are currently still being made by people. Current figures suggest that 89% of activations are within 3 minutes (against a 95% target), 39% of arrivals are within 8 minutes (target 50%) and 89% within 14 minutes (target 95%). However, only 72% (target 95%) of urgent patients reach hospital within 15 minutes. This reflects the fact that only a proportion of the planned system features have yet been delivered. A fully functional replacement system is still some years into the future.

3. SECOND CASE: METROPOLITAN AMBULANCE SERVICE, MELBOURNE, VICTORIA, AUSTRALIA

The story of the Metropolitan Ambulance Service (MAS) began in 1992. The background to the story had two key features. First, the relationship between MAS management and the ambulance workers had always been punctuated by strikes, mistrust and tension. Second, the new Conservative state government elected in 1992 was committed to a cost-cutting war in reaction against their predecessors' 'wasteful financial policies'.

A 1992 review of the MAS proposed significant improvements, the main goal of which should be to generate cost savings in response to concern over the level of government contributions to the service. The government replaced the MAS committee of management with a CEO with no experience in emergency services, whose brief was to cut costs and break the union. His cost-cutting programme focused in part on the utilization of computer technology in the management of emergency calls.

In 1993 the contract for the emergency system, totalling \$A32 million, was awarded to a company called Intergraph by the Victoria State Government. During the negotiation process, the consultant acting for the government was also employed by Intergraph.

The system would represent the state of the art in emergency despatch and communications. The planned components were:

- a despatch system for the automatic despatch of the nearest or most appropriate vehicle;

- a satellite-based vehicle location system, supported by a computerized mapping system;
- mobile data terminals to replace voice radio communication.

Big bang implementation was planned, with switchover to the full system scheduled for 24 August 1995. An independent consultant's report earlier in that year identified more than 50 faults labelled 'critical' or 'high priority'. While this was happening, the contract for providing all emergency services (fire, police, ambulance) state-wide was awarded to Intergraph.

The MAS switchover took place with a number of faults identified in the consultant's report still outstanding. Key problems included:

- the volume of information handled by the automatic vehicle location system was causing it to bog down and report incorrect locations;
- allocation of a vehicle to a job could take up to two minutes, during which time the individual console concerned was locked up;
- responses taking longer than 16 minutes were required to generate exception reports, but the system's measurements of its own response times were not accurate;
- the report generator, responsible for producing exception reports, contained critical faults;
- while no statistical tests had been applied, it was observed that even on simple tasks the automatic route recommendation facility often recommended routes heading in the opposite direction from an accident scene.

Within a few days of switchover MAS officers began complaining to Intergraph, and a heated row developed between the two organizations. The MAS claimed that Intergraph had promised to deliver things that they never had a chance of delivering. In a memo to the MAS, the Intergraph CEO apologized for several issues including service standards, persistent problems and despatch errors. The failure to meet the (already significantly relaxed) performance criteria blossomed into a fully fledged legal dispute, which eventually saw the MAS taking legal advice on the possibility of terminating the contract.

Within two weeks of switchover, the MAS was complaining to the Health Minister. Government advisers suggested that ministers should claim that performance was improving.

In a cost-saving effort, the MAS Regional Training Unit had been closed. Middle management training



for the new system was provided in the form of three-day live-in courses, described by participants as 'mini Camp Wacos'. Participants who expressed opinions critical of new management policies were subjected to emotional and personal abuse which left them exhausted and traumatized. Ambulance officers were coerced into improving their knowledge and skills in their own time.

The MAS internal documents record the view just after switchover that the system often 'teetered on the edge of disaster'. Public pressure led the Victorian Auditor-General to approve an extensive performance audit of the MAS. The Labour opposition released an Intergraph log in October 1995 which showed despatchers repeatedly complaining that ambulances could not be picked up by the system and that it was locking up. In December 1995, the opposition called for an inquiry into the use of public money in the Intergraph deal. The government was eventually forced to withhold contract payments until contracted requirements were met by Intergraph.

In February 1996 the head of the MAS told a ministerial steering committee that 'the services provided to MAS... are progressively getting worse' and 'the service was sub-standard and worse than what was previously provided at the old Communications Centre'.

An incident in August 1996 attracted particular attention. An emergency call to an address in South Street reported that a man was unconscious and thought to be suffering from a drug overdose. The despatch system failed to recognize the street name, and 'corrected' the address from South Street to Sadie Street. The operator failed to notice the error, so that the wrong address was passed to the ambulance crew. The ambulance eventually reached the right address, but too late, and the patient died. It is significant that the system was known to be still confused between the two streets ten months later.

When several emergencies occurred simultaneously, they resulted in severe bottlenecks. During an outbreak of bushfires not long ago, the system became so overloaded that operators had to resort to whiteboards to keep track of events.

The MAS ambulance performance standard stipulated that for serious cases an ambulance should be despatched within 150 seconds in 80% of cases. The figure achieved for November 1996 was 78.7%. The introduction of a new question-and-answer routine for call-takers in December 1996 resulted

in this figure plummeting to 34.2%. Early the following year the figure crept up to 36.7%. The standard also dictated that under no circumstances can the system be allowed to take more than four hours to despatch a non-emergency ambulance. A 1996 ambulance memo recorded that many cases exceeded nine hours.

In January 1997, the head of the MAS complained to the Department of Justice that the MAS was asked to lower its performance standards rather than expect Intergraph to improve the performance of the system. The MAS, he said, 'was of the view that Intergraph do not really understand, or have not come to grips with, operating in an emergency services organisation environment, and are unlikely to do so unless there is a fundamental shift in approach'. A month later, he complained that it was 'no longer acceptable for MAS to rely on Intergraph to eventually 'get it right'. In July he complained to the Bureau of Emergency Services Telecommunications Head about six months of extreme pressure from Intergraph over the reduction of standards. When the use of imported trouble-shooters failed to alleviate the problem, the MAS and Intergraph embarked on a protracted series of negotiations culminating in agreement to lower the benchmark performance standards in return for reduced payment.

In March 1997 the Victorian Public Bodies Review Committee and Department of Finance advised against the government's decision to enter into a further contract with Intergraph. The Auditor-General's special emergency report on the MAS contractual and outsourcing practices was published in April 1997; that was several months earlier than originally scheduled, because of the 'discovery of extremely serious matters identified during the course of a performance audit'. It enumerated serious deficiencies and highly dubious practices, including huge cost blow-outs, entangled relationships between managers in the service and private companies, a biased tender, and inadequate supervision of contractual decisions by MAS management (who were said to have been 'derelict' in their responsibility).

Major flaws identified included:

- reliance on a consultant without a formal contract and with no attempt to establish past experience or association;
- acceptance of the requirements analysis submitted by the consultants in charge of the tendering



process, despite serious reservations about the quality of the analysis;

- serious functional and technical flaws and shortfalls in the original tender document, which was obviously written around the system that would be bid by Intergraph;
- the system specification document, which was hurriedly developed and known to contain major shortcomings, yet was fully accepted by the MAS;
- the absence of documentary evidence to show how the 34 registrations of interest were short-listed to four potential suppliers;
- the inability to produce the evaluation criteria used in selecting Intergraph, and the informal approach used to eliminate the remaining bidders;
- failure to satisfy a key condition set by the government that the system should be capable of being subsequently integrated into a state-wide emergency system, irrespective of who the suppliers of each system would be;
- failure to achieve the projected savings of \$A20 million;
- easy acceptance by MAS management of assurances that all was well.

The audit also uncovered a memorandum that had been sent to the MAS CEO by the manager of information systems, prior to the awarding of the contract, expressing the depth of his concerns and reservations about the proposed system and the selected contractors. It called for:

- withdrawal of the specification, as it did not cover MAS requirements, and a delay of four weeks to enable a team of specialist staff, with expertise in communications, information systems, technical services, ambulance procedures, and despatch systems, to review and redraft it;
- review and redrafting of the schedule, and especially of project milestones and phases, to obtain a realistic schedule with 'practical implementable stages';
- review of the short-listed suppliers;
- appointment of a project manager, ideally with skills in computer-aided despatch and related technical areas.

Minutes from subsequent meetings reveal that these concerns had been discussed, and were shared by all key personnel with the exception of the chairman.

Copyright © 2003 John Wiley & Sons, Ltd.

They were brushed aside, however, and the tender process had continued unchanged.

Following the publication of the Auditor-General's report, the systems support manager resigned. The former MAS boss was heavily criticized for failing to take a more proactive stance. Ambulance officials who evaluated the Intergraph bid called for a full-scale Royal Commission, claiming the tender was rigged and the process a sham; later evidence confirmed both claims.

An investigation shortly afterwards by *The Sunday Age* newspaper unearthed further evidence:

- The head of the consultants overseeing the tendering process ended up as CEO of Intergraph. (It was also subsequently revealed that the former MAS CEO had connections with the consultancy firm.)
- Members of the MAS tender evaluation team each ranked all the bids as part of their final evaluation. The averaged rankings gave Intergraph the lowest score. Two days later, the team was redirected by a consultant representing MAS management to change the evaluation criteria and rewrite their final report. As the Intergraph bid was particularly weak on backup and security aspects, team members were asked to change these and several other items that did not favour Intergraph.
- Members of the evaluation team were directed not to inspect the operating sites of Intergraph's competitors
- When serious problems became apparent during testing, team members were instructed to abandon testing. To avoid further embarrassing crashes, the team was only allowed to ask questions but was denied access to the equipment.
- Warnings by the head of the MAS communications department about the safety risks in changing over to the Intergraph system were ignored.
- Additional concerns over lack of medical knowledge among the Intergraph staff (including operators) were identified as a major safety risk.
- 'The timeframes were deliberately short so the due diligence couldn't be done properly. They got a system that will never work, that is flawed in its basic design.'

In March 1998, Internal Health Department documents released under Freedom of Information

Softw. Process Improve. Pract., 2002; 7: 71–89



legislation revealed that the consulting group overseeing the tender was hired on a \$A45K contract but was ultimately paid \$A1.4 million, without providing appropriate documentation of costs.

In response to the Auditor-General's findings the government initiated a police and Queen's Counsel inquiry into the contracts. The government's solicitor managed to block the release of sensitive documents on the ground that they 'would be reasonably likely to prejudice a Victorian police investigation of a possible breach of law'. In May 1998, a report by the Major Fraud Squad was submitted to the Office of Public Prosecutions. The report identified possible criminal offences in the MAS procurement and contracting process. The homes of the former chief of the MAS, his deputy, and the former government consultant were raided by the police following the disappearance of documents containing warnings about the Intergraph deal. The volume of missing documents seriously hampered police investigations throughout. The Minister for Police resigned. The Director of Public Prosecutions subsequently decided there was insufficient remaining evidence to bring criminal charges.

Allegations were made that Intergraph staff were placing emergency calls which were answered immediately, thereby reducing average delay times and improving performance. Intergraph claimed the calls were simply to test the system. An independent auditor's report confirmed that Intergraph artificially manipulated response times through fake calls.

The opposition health spokesman said in parliament that 'the Government had got into bed with shysters and crooks and been played for a sucker, and it had been robbed blind by contractors and consultants'.

In August 1998, giving evidence before a Victorian Civil and Administrative Appeals Tribunal, a former member of the Intergraph bidding team said that the despatch system had no chance of working because of a major software problem; that it was not compatible with the automatic location and communications software, thus forcing ambulances to use the existing communication system; that the team ignored these obvious faults; and that the team knew in advance they would win the bid, and were supplied with 'excellent intelligence' about competing bids. 'The ambulance service was never informed of the problem and was also misled about the cost of the contract'. The witness

had lost his job after expressing his concerns to his superior, who was planning to become the future Minister for Police and Emergency Services. In December 1998, the Tribunal ordered the State to release all documents on the MAS despatch system, including a secret report about Intergraph's role, stipulating that the release would not prejudice future litigation against companies involved in the contract.

A two-year police investigation found 'a *prima facie* case of misconduct in public office', that a State minister misled parliament, that government officials hampered police investigations, and that the government illegally covered up improprieties by suppressing key documents on the eve of the 1996 election. State lawyers subsequently argued that key documents should be exempt from release on the grounds that they 'would be reasonably likely to prejudice a Victoria police investigation of a possible breach of law'. In an additional bid to block the release, the State financed private legal action by ministerial advisers and officials implicated in the documents. A Supreme Court judge accused State lawyers of 'committing one of the worst abuses of the court process'.

The MAS computerization programme resulted in a less efficient ambulance service plagued with technical hitches, loss of experienced staff and low morale. The performance of the system caused much public concern as faults in the system have been implicated in a series of misadventures in which people have died while waiting for delayed or misdirected ambulances. To date, there have been at least nine Coroner's Court inquests into deaths associated with serious delays. In one case the Coroner proposed that the MAS should conduct a detailed examination of call taking, communication and despatch systems. The Victorian secretary of the ambulance employees union is on record as saying 'The worst aspect of the Intergraph affair is not the cost, but the fact that the Government lost control of essential services, and as a result people died unnecessarily'.

Following a recent electoral victory by the Labour party, a Royal Commission was appointed, with wide-ranging terms of reference that enable it to investigate (among other things) the political links that secured the contract and the illegal release of details of competing bids. It appears likely that criminal charges against senior political figures may follow the inquiry.



4. CAUSES OF FAILURE

Just a straight-through reading of the above tales makes it clear that no single or simple causes of failure can be proposed. Causes are multiple, complex and inter-related. This section sets out the more obvious, almost all of which apply (albeit sometimes with differences in detail or emphasis) to both cases. Indeed it is striking that, though the stories unwind in very different ways, those underlying causes are so similar – and so basic; so basic that one would have expected them to be well understood by any ‘professional’ practitioner, even a decade ago.

In undertaking the analysis for this paper, each of the following failure causes was validated by traces back to the specific passages in the case histories which provide the supporting evidence, although it is not possible to show those traces in the limited space available here. It should also be noted (a) that the case histories themselves are severely truncated accounts of what actually happened, meaning that there are undoubtedly causal factors that are not revealed, (b) that the analysis has not yet gone as far as to investigate fully the root causes and systemic relationships between causes (for instance mutually reinforcing feedback loops), although it requires little analysis to guess at some of them.

Each failure cause below is followed by a brief discussion. The causes are presented in no significant sequence.

4.1. Blind Optimism that Novel Technology Solutions are Achievable

Developers (whether in-house or contractors) tend to innate technological optimism, and this is often not counterbalanced by management realism. High levels of system complexity, ambition and innovation are not recognized, and necessary steps to manage them are not taken. The project sets out to create in a single leap the ultimate in its domain, beyond that achieved in any existing system, and is infected with technical naivety, machismo and hubris. The prospect of enormous benefits and the splendour of the utopian achievement blinds participants to the scale of the challenge and the risks of failure, which are wilfully and irresponsibly overlooked. It is assumed that experience with smaller simpler systems, or in other domains, qualifies clients and contractors to take the leap of faith.

There is a belief that ‘it will be all right in the end’. The required knowledge for attempting such a breakthrough in technology does not exist, and the need to extend budgets and schedules to allow for research and experiment to acquire that knowledge, and to proceed towards the end-goal by a series of controlled-risk steps, is ignored.

One commentator on the LAS (Charette 1995) noted that, even given perfect conditions (such as proper management, unlimited timescale and full resources), it would not have been possible to achieve the functionality and quality called for in the requirements. Even using, for instance, global positioning systems, inertia navigation systems and satellite-based communications would not have provided the specified level of performance. Further, allocation and mobilization algorithms are known to be difficult to design even for less challenging environments. Given a metropolis the size of London, given the contingencies that may arise (storms, major accidents, road works, festivals, processions, sports matches, demonstrations, broken-down vehicles, staff holidays, missing drivers, malfunctioning equipment, radio blackspots), given the fallibility of human operators, the algorithm may belong to the class of intractable problems.

Lesson: either do not venture into unexplored territory or, if you do, take proper precautions.

4.2. Impossible and Inflexible Deadlines, Relentless Rush

This is another form of blind optimism, that ‘it will be all right in the end’. Obviously for any important system, other things being equal, the sooner it can be delivered, and its benefits realized, the better. Sometimes circumstances dictate an absolute deadline (e.g. a Y2K project), but that is often not the case. Frequently some essentially arbitrary deadline (often motivated by political or PR considerations, or by personal ambition) is proposed, accepted without proper estimating or planning to establish its viability, and then becomes a fixed part of the landscape for managers and developers alike, to be defended at all costs. Where there is a client–contractor relationship, contractors are all too willing to collude with their clients’ delusions. Even as the deadline gets nearer and nearer, and common sense would seem to dictate that it is increasingly unachievable, there is no review and



no attempt to change either the deadline or the solution. Short cuts are taken, and essential processes (such as testing, reviews, problem resolution, training) are ignored in the rush to complete essential technical tasks. Senior managers rarely take action to prevent this – on the contrary they are often prime sources of pressure.

4.3. The Fatal Attraction of Big Bang Implementation

This is closely associated with the two previous causes. The attempt to leapfrog to a system that delivers all the expected functionality and benefits in a single-shot implementation is fraught with risk. With complex systems, phased implementation with each phase delivering some useful functionality and building on what already exists offers a controlled approach to dealing with risk. The overall effort can be assessed based on results to date, and it becomes easier to spot trouble earlier and to adjust overall effort if necessary. This approach also gives users a flavour of what the system can do for them, can generate enthusiasm and support, and can reduce resistance. The presence of one or more effective, tried and tested back-up systems, with the necessary cutover procedures, and with staff trained in their use, is always important, but even more so where big bang is being attempted.

4.4. Impossible and Inflexible Budgets, Obsession with Costs

This failure cause is the other side of the 'impossible deadline' coin, and the two normally go together. In an atmosphere of internal or external pressure to cut or control costs, cost reduction is regarded as the prime goal of the system, and low development cost is seen to be essential in gaining project approval or the prime criterion in bid selection. Just as there is resistance to changing deadlines, there is often a natural reluctance to cancel projects with runaway costs. Investment already made in the project feeds the escalation cycle and results in throwing more money after bad, rather than a re-evaluation of the return on continuing investment.

Copyright © 2003 John Wiley & Sons, Ltd.

4.5. Deafness to Alerts

Alerts to potential problems may come from a variety of sources: past experience (in-house or public-domain); warnings from knowledgeable sources (own staff, consultants or other outside observers); feedback from reviews; observation of events; or just plain common sense. Deafness to such alerts may afflict any stakeholder group – senior management, project management, technical developers etc. If those capable of sounding the alert observe that deafness is the usual outcome (or worse, threats not to rock the boat), then the alert will be sounded less and less frequently.

4.6. Groupthink

This is a stronger version of the previous failure cause. Rather than mere deafness, groupthink is a phenomenon in which a magic circle of individuals collectively adopt a set of *idées fixes* and display active resistance to any outside influences that may threaten that mindset. It is a defence of the insecure against external reality, a circling of the waggon against the Indian hordes – a vicious circle, indeed, of self-reinforcing mutual assurance. The positive feedback loop can occur between individuals in a development team, between developers and managers, or between client and contractor. It can sustain the blind optimism in technological magic or impossible deadlines and budgets that has been discussed above. It is not limited to IT-based systems, but has been noted in such high-profile disasters as Bhopal, Three-Mile Island, Challenger, Chernobyl and even Pearl Harbour.

4.7. Dysfunctional Relationships among Stakeholder Groups

Key stakeholder groups include senior management (of the enterprise for whom the system is being developed), the system developers/suppliers (either in-house or external), the commissioning group (where the developers/suppliers are external), and the system users. It is essential that those (and other) stakeholder groups are integrated into the system process, with each having an effective voice so as to be able to express both their needs and their knowledge. Poor inter-group communication increases risk, and success relies on establishing stakeholder ownership and commitment.

Softw. Process Improve. Pract., 2002; 7: 71–89



Examples of dysfunctional relationships between groups include:

- failure of senior management to understand issues, get involved and committed, organize for the systems effort, and undertake properly informed and joined-up decision-making;
- failure of senior management to relate the systems effort to other concurrent problems/changes within the enterprise (such as management restructuring or BPR), and to understand and deal with their interactions;
- unreasonable pressure by senior management on development teams;
- lack of attention by senior management to staff concerns (both developers and users) exacerbated by major overspending of scarce resources on glossy brochures, management consultants and corporate image, or by perceptions that it is lying to the public;
- a climate of fear, blame and low morale among developers and/or users, of which senior management is ignorant or tolerant, or at worst which it encourages;
- scapegoating of one group by another;
- failure of development/commissioning teams to seek senior management involvement, and to gain their support in resolving key problems;
- failure of development/commissioning teams and senior management to consult and involve user groups and unions;
- failure of development teams to adopt user-centred design practices;
- a culture of strong internal politics;
- improper relationships and vested interests between commissioning groups and external developers/suppliers.

4.8. Inadequate Skills, Experience and Training Provision

In almost all projects there will be individuals who will be required to undertake tasks that take them beyond the limits of their past experience and existing skills, and in some projects that will apply to a whole team; relevant training delivered at the right time is essential. In almost all cases systems will demand changes in working practice from users, who will also need relevant and well-timed training. Too often training is ill-planned and delivered

too early, resulting in skills decay. It is unfair to condemn people as incompetent when the fault is with poor staff assignment and development processes.

4.9. Lack or Inadequacy of Essential Processes

If you are going to ensure that you have the individuals with the necessary skills, then you must know what skills you want – in other words, you must know what processes need to be carried out. The following are among processes normally recognized as essential that were missing or inadequately undertaken in one or both of the ambulance despatch cases. Each of these process deficiencies could be a major cause of failure in its own right. Note that they almost all figure as key processes or practices in the Capability Maturity Model and in related models and standards for high capability:

- Project management: sizing and estimating, resource and schedule planning, tracking and oversight, budgetary control, problem identification and resolution. Ensure plans provide for all processes, and include contingency allowances (e.g. rework). If metrics from past projects are not in place to support estimating, ensure that they are at least put in place in this project.
- Risk management: risk identification (including scenario development), risk analysis, risk response planning, risk resolution and monitoring. No project is risk-free; high-novelty projects involve high risk. Plan for all contingencies, throughout the system lifetime.
- Tender, bid and contract management. Ensure transparency in bid evaluation and selection.
- Requirements management: elicitation; analysis; negotiation; description/modelling; change management; tracking. Avoid requirements creep. Do everything to ensure common understanding of requirements between client and contractor (even when they are within the same enterprise), acceptance requirements, and clear separation between problem statement and design solution.
- Configuration management.
- Quality management: inspections; design reviews; milestone reviews; V & V; testing (see below); sign-off of key work products; quality assurance; quality planning and control.
- Appropriate and mature technical practices: architecture and design methods; performance



estimating/simulation; prototyping; analysis of design alternatives and trade-offs etc., together with other technical issues discussed in the following bullets.

- Testing practice, including test plans and designs, testing under ranges of operational loading and other conditions, testing of back-up and other contingency measures, and means of predicting test completeness, is an especially critical technical issue. It impinges on the issue of release management – judging when the system is fit to be released. In the case histories, release decisions were made in the known presence of serious defects. It could be said that both systems were riddled with design and technical flaws that collectively triggered their collapse.
- Reuse can also be a key issue. Systematic and well managed reuse of elements of previous systems/solutions can bring great benefits, but without proper understanding and analysis they can be high-risk (refer not only to the cases above but also, for instance, to Ariane 5).
- Wrong decisions on technology infrastructure can have massively damaging effects. This is well illustrated in the LAS case by two such decisions. One was the use of Visual Basic, already discussed. Another was the adoption of Apricot PCs and servers which, however reliable for the purposes for which they were designed, lacked the power offered by the mid-range computers commonly used for such applications at the time. Further, using a network of PCs meant they had to be programmed to share the workload – an addition to the already massively overstretched software activity.
- Certain classes of systems (such as safety-critical, which includes the sub-class of ambulance systems) present requirements for certain more rigorous technical practices and design approaches: there is a considerable literature on these, and they are not pursued further here. Suffice it to say that those practices add to development costs: in a climate of cost-cutting there are severe risks that they will not be followed.

4.10. Failure to Recognize Systems Dimension

The need to incorporate component subsystems designed and optimized by various agencies and vendors underscores the critical role of effective systems engineering. The complete system depends on

the interactions between its components. No single supplier can offer all the systems and components required for complex systems. Aside from assuring the timely delivery of all components, the central co-ordinators need to ensure that all packages and hardware communicate and interconnect in a way which optimizes the overall system. Often multiple sourcing implies multiple concurrent development projects and the consequent need for programme management capability.

4.11. Inappropriate Human–Machine Interface

The technological optimism identified in the first failure cause often leads to overreliance on technology at the expense of what is properly the domain of human skill and judgement. Machine computation may be used to supplant human decision makers when the contextual system is relatively closed; open systems, such as ambulance despatch systems, however, with their inherent imperfections and unknown factors, need to retain a degree of reliance on human judgement. Rule-based analytical approaches, as attempted in both case histories, cannot deal as well as experienced operators with special cases. These systems wrongly reduce operators' influence. There is no evidence that machine capability can ever match the human mind's ability to deal with unexpected or unpredictable situations. The more investigators understand the complex and changing nature of information and knowledge, the more clear it becomes that trained operators will always be required as elements in a co-operative relationship between technology and humans. Overreliance on technology often results in ignoring the human element. The most valuable asset is knowledgeable staff, especially those with informal knowledge and keen awareness of the subtleties of the system. Systems should be built to utilize the skills of the people working within an organization, rather than overriding or marginalizing them.

Successful despatch systems act as decision aids and not decision makers. They offer the human decision maker the luxury of having additional information and alternative channels for obtaining information, while retaining the flexibility and contingency offered by this redundancy. The human–machine interface must reduce the likelihood of errors and present all the information required for rapid response to exceptional



situations. Operators must be empowered to make decisions in short time spans with fully adequate information. Visualization, displaying this information in a clear way, is a key to enabling rapid decision making and reaction to feedback. Usability engineering is becoming an essential part of systems development. When human capabilities are ignored, the versatility they can offer the system is lost.

5. LEARNING FROM FAILURE: FORENSIC SYSTEMS ENGINEERING

'Experience is a dear teacher, and only fools will learn from no other.' (*Benjamin Franklin*)

'The most important of my discoveries have been suggested by my failures.' (*Sir Humphrey Davy*)

'The history of engineering may be told by its failures as well as its triumphs.' (*Henry Petroski*)

Failures offer invaluable, if often painful, lessons. Having made an investment, it is important to reap some reward, even if it not the one that was planned. For the organizations that suffer them, they offer opportunities for organizational learning, improvement and the enhancement of corporate knowledge base. For the individuals involved in them, learning from failure should be a required part of professional ethics.

A group of authors (Pinkus *et al.* 1997) have identified a set of basic ethical principles applicable to engineering, the first of which is the obligation to technical competence. It includes the obligation to admit when expertise does not exist and to seek it by updating knowledge and skills. Being willing to participate in post-mortems, and studying the documented history of failures, is one of the most effective ways of doing that. Developers should not be encouraged, even allowed, to bury their mistakes. One of the greatest resource available is the knowledge that can be gleaned out of what did not work.

Forensic systems engineering is the post-mortem analysis and study of project failure case histories. The work involves a systematic investigation of a project, its environment, the decisions taken,

the relationship between sub-systems, and the relationships between the processes, people, technology and organization involved in the project. It draws on a multidisciplinary body of knowledge and method, including decision analysis, systems science, and chaos and complexity theories, and assesses projects from several directions and viewpoints. The concept of systems is central for understanding the complex relationships and their implications in the overall project environment. Forensic case histories provide important inputs to risk management and mitigation.

There is no ideal word for the scientific or systematic investigation of systems failures and their causes. Without inventing one, *forensics* is the best available and mirrors its use, for instance, in *forensic science* and *medical forensics*. This usage has gained acceptance within several engineering disciplines. Its Latin root, *forensis*, has the dual meaning of *pertaining to the courts* (which gives rise to its more common English usage), and also *public*. It is this latter meaning on which we draw, emphasizing the importance of bringing failure information into the public domain rather than burying it.

In a nutshell, the message of this paper is: understand that engineering (like science, art and indeed life) advances by trial and error, and that failures are unavoidable; acknowledge and value failures; study them to gain the knowledge that is embedded in them, and use that knowledge so that you and your organization do better next time.

ACKNOWLEDGEMENTS

The authors wish to thank the numerous contributors and correspondents who volunteered information in various formats.

REFERENCES

- Charette RN. 1995. No one could have done better. *American Programmer* 8(7): 21–28.
- Pinkus RLB, Shuman LJ, Hummon NP, Wolfe H. 1997. *Engineering Ethics*. Cambridge University Press: New York.
- Neumann PG. 1991. Risks to the public in computers and related systems. *Software Engineering Notes* 16(1): 10.



Practice Section

Learning from Failures

Neumann PG. 1992. Risks to the public in computers and related systems. *Software Engineering Notes* **17**(1): 11–12.

Neumann PG. 1996. Risks to the public in computers and related systems. *Software Engineering Notes* **21**(2): 19.

Neumann PG. 1997. Risks to the public in computers and related systems. *Software Engineering Notes* **22**(2): 21.

Neumann PG. 1998. Risks to the public in computers and related systems. *Software Engineering Notes* **23**(4): 22.