

A Knowledge-Based Approach for Designing Robust Business Processes

Chrysanthos Dellarocas and Mark Klein

Center for Coordination Science
Sloan School of Management
Massachusetts Institute of Technology
Room E53-315, Cambridge, MA 02139, USA
{dell, m_klein}@mit.edu

Abstract. This chapter describes a novel knowledge-based methodology and computer toolset for helping business process designers and participants better manage exceptions (unexpected deviations from a normal sequence of events caused by design errors, resource failures, requirement changes etc.) that can occur during the enactment of a process. This approach is based on an on-line repository exploiting a generic and reusable body of knowledge, which describes what kinds of exceptions can occur in collaborative work processes, how these exceptions can be detected, and how they can be resolved. This work builds upon previous efforts from the MIT Process Handbook project and from research on conflict management in collaborative design.

1 Introduction

Business process models typically describe the “normal” flow of events in an ideal world. For example, the model of a product development process typically includes a “design product” activity, followed by a “build product” activity, which, in turn, is followed by a “deliver product” activity. Reality, however, tends to be more complicated. During the enactment of a business process a lot of *exceptions*, that is, deviations from the normal sequence of events, might occur. For example, product design might prove to be inconsistent with the capabilities of the manufacturing plant. Manufacturing stations might break down in the middle of jobs. Delivery trucks might go on strike. To assure that a process is still able to fulfill its organizational goals, process participants must be able to detect, diagnose and successfully resolve such exceptional conditions as they occur.

Traditionally, managers have relied on their experience and understanding of a process in order to handle deviations from the expected flow of events. However, the increasing complexity of modern business processes and the accelerating pace with which these processes change has made the reliance on individual managers’ experience and intuition an increasingly less satisfactory way to deal with exceptions. There is an increasing need for *systematic* business process *operational risk*

management methodologies. Such methodologies will assist business process designers to anticipate potential exceptions and instrument their processes so that exceptions can either be avoided or be detected in a timely way. Furthermore, when exception manifestations occur during process enactment, these methodologies assist in selecting the best way of resolving them.

Current process modeling methodologies and tools [6, 11, 12] do not make any provision for describing exception handling procedures separately from “main-line” processing. This approach, however, is problematic for a number of reasons. First, it results in cluttered, overly complex, models, which hinder instead of enhancing understanding and communication. Second, the anticipation of possible failure modes once again relies on the experience and intuition of the model designers. Third, the approach cannot help with exceptions that have not been explicitly hard-coded into the model.

This chapter describes a knowledge-based approach for designing robust business processes. Rather than requiring process designers to anticipate all possible exceptions up front and incorporate them into their models, this approach is based on a set of novel computerized process analysis tools, which assist designers in analyzing “normal” process models, *systematically* anticipating possible exceptions and suggesting ways in which the “normal” process can be instrumented in order to detect or even to avoid them. When exception manifestations occur, these tools can help diagnose their underlying causes, and suggest specific interventions for resolving them. The approach is based on an extensible knowledge base of generic strategies for avoiding, detecting, diagnosing and resolving exceptions.

The remainder of this chapter is structured as follows: Section 2 provides an overview of the proposed approach. Section 3 describes how the approach has been successfully applied to analyze operational risks of the Barings Bank trading processes. Section 4 discusses related work. Finally, Section 5 presents some directions for future work.

2 A Knowledge-Based Approach to Exception Handling

2.1 What is an Exception?

We define an exception as any deviation from a “normal” collaborative process that uses the available resources to achieve the task requirements on an optimal way. An exception can thus include errors in enacting a task or distributing results between tasks, inadequate responses to changes in tasks or resources, missed opportunities and so on. To make this more concrete, consider the possible exceptions for the generic coordination process known as “subcontracting”. Subcontracting can be used whenever one wants to share agents who do a service among requestors of that service. The requestor for a service sends out a request for bids (RFB) asking for agents to perform a given task. Interested subcontractors respond with bids. The requestor awards the task to the subcontractor with the best bid (based for example on anticipated cost, quality or timeliness), at which point the subcontractor performs the

task and returns the results to the requestor. This mechanism makes many implicit assumptions; violations of any one of them can lead to exceptions (Figure 1):

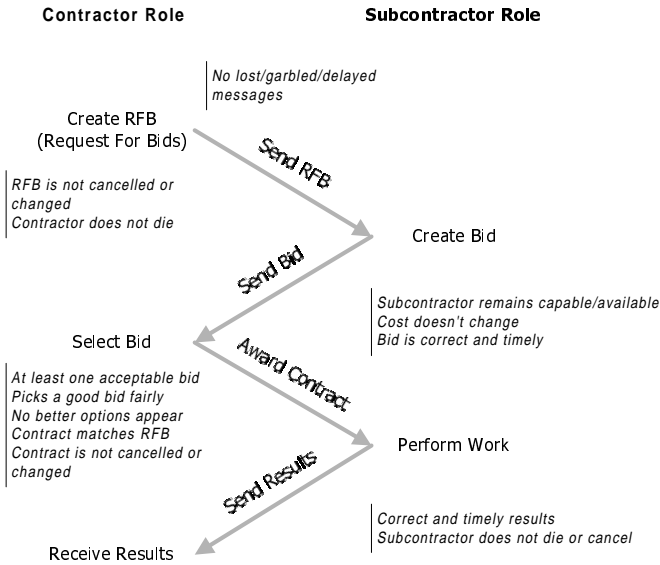


Fig. 1. Implicit Assumptions for Subcontracting.

Plain text items on the left and right represent tasks done by the Contractor (on the left) on the Subcontractor (on the right). Labeled arcs represent interactions between the Contractor and Subcontractor. Items in italics represent implicit assumptions made by this mechanism (for example that the task required by the Contractor does not change after it has sent out the RFB, that the Subcontractor does not cancel the task it was assigned, become incapable of doing it, make a mistake etc.). Any event that results in the violation of any of these assumptions represents a possible exception for the subcontracting process. In addition to that, some exceptions take the form of dysfunctional systemic behavior that may result even when the mechanism is followed perfectly. Deadlock (where several agents are each waiting for another one to do something) and resource poaching (wherein high-priority tasks are unable to access needed resources because these resources they have already been reserved by lower priority tasks) are all examples of this.

We have developed, as a result of analyses like that shown above, a growing taxonomy of exception types, a subset of which is show below (Figure 2):

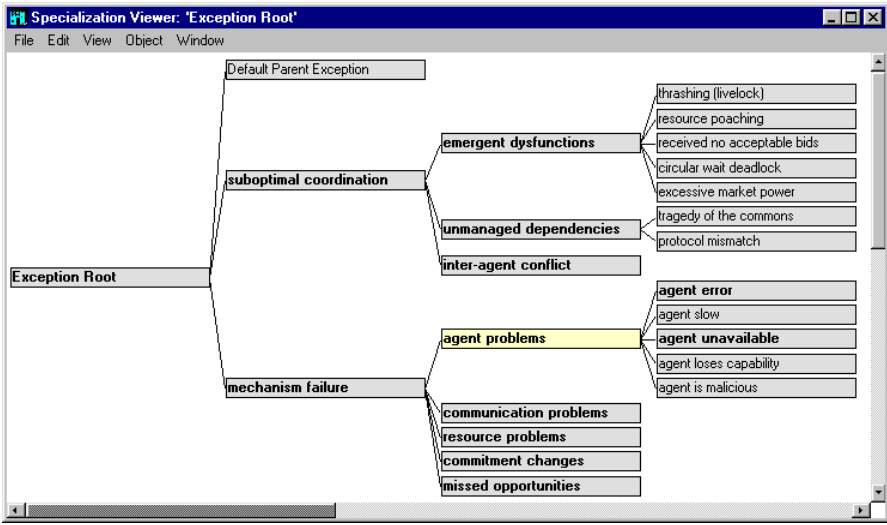


Fig. 2. A Subset of the Exception Type Taxonomy.

As we shall see, the essence of our work is developing a knowledge base that captures such exceptions and relates them to (1) the processes that they can occur in, and (2) the processes that can be used to manage (anticipate, detect, avoid and resolve) them.

2.2 Preparing for Exceptions

The first step in our approach helps process designers to anticipate, for a given “normal” process model, the ways that the process may fail and then instrument the process so that these failures can be detected or avoided. The principal idea here is to compare the process model against a taxonomy of elementary process elements annotated with possible failure modes.

A process element taxonomy can be defined as a hierarchy of process element templates, with very generic elements at the top and increasingly *specialized* elements below. For example, Figure 3 depicts a small activity taxonomy. Each activity can have attributes, e.g. that define the challenges for which it is well-suited. Note that activity *specialization* is different from *decomposition*, which involves breaking an activity down into subactivities. While a subactivity represents a part of a process; a specialization represents a “subtype” or “way of” doing the process [20, 21]. Resource, goal and assumption taxonomies can be defined in a similar manner.

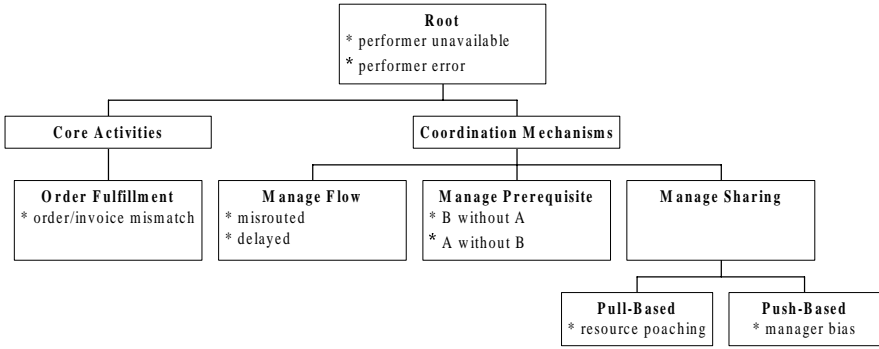


Fig. 3. An Example of a Generic Activity Taxonomy Annotated with Failure Modes.

Process element templates are annotated with the ways in which they can fail, i.e. with their characteristic *exception types*. Failure modes for a given process template can be uncovered using failure mode analysis [24]. Each process element in a taxonomy inherits all characteristic failure modes of its parent (generalization) and may contain additional failure modes which are specific to it.

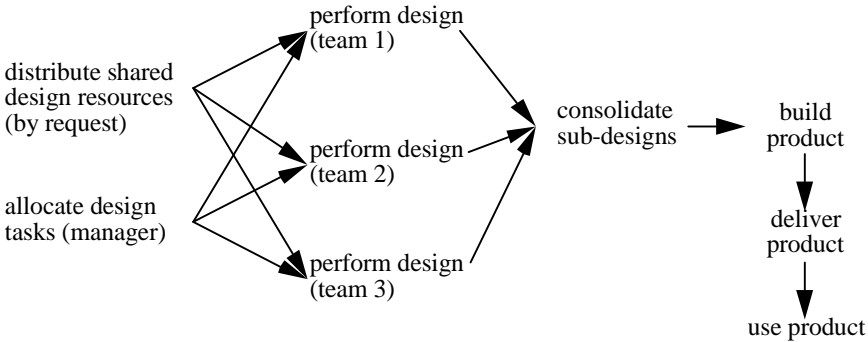


Fig. 4. An Example "Normal" Process Model.

Given a “normal” process model, to identify failure modes we need only identify the generic process element templates that match each component of the model. The potentially applicable exception types will then consist of the union of all failure modes inherited from the matching templates. We can see, for example, that the “distribute shared design resources” activity in Figure 4 is a subtype of the generic “pull-based sharing” process template in Figure 3, since the resources are “pulled” by their consumers rather than “pushed” (i.e. allocated) by their producers. This template includes among its characteristic failure modes the exception called “poaching”, wherein resources go disproportionately to lower priority tasks because agents with lower priority tasks happen to reserve them first. The “deliver product” activity is a

specialization of the “manage flow” template, with characteristic exceptions such as “item delayed”, “item misrouted” and so on. All activities also inherit the characteristic failure modes from the generalizations of these matching templates, such as “responsible agent is unavailable”, and so on.

The process designer can select, from this list of possible exception types, the ones that seem most important in his/her particular context. He/she might know, for example, that the “deliver product” process is already highly robust and that there is no need to augment it with additional exception handling capabilities.

For each exception type of interest, the process designer can then decide how to instrument the process in order to detect these exceptions. While processes can fail in many different ways, such failures have a relatively limited number of different manifestations, including missed deadlines, violations of artifact constraints, exceeding resource limits, and so on. Every exception type includes pointers to *exception detection* process templates in the process taxonomy that specify how to detect the symptoms manifested by that exception type. These templates, once interleaved into the “normal” process model by the workflow designer, play the role of “sentinels” that check for signs of actual or impending failure. The template for detecting the “resource poaching” exception, for example, operates by comparing the average priority of tasks that quickly receive shared resources against the average priority of all tasks. The “item delayed”, “agent unavailable”, and “item misrouted” exceptions can all be detected using time-out mechanisms. Similar pointers exist to *exception avoidance* processes, whose purpose is to try to prevent the exceptional condition from occurring at all.

2.3 Diagnosing Exceptions

When exceptions actually occur during the enactment of a process, our tools can assist process participants in figuring out how to react. Just as in medical domains, selecting an appropriate intervention requires understanding the underlying cause of the problem, i.e. its *diagnosis*. A key challenge here, however, is that the symptoms revealed by the exception detection processes can suggest a wide variety of possible underlying causes. Many different exceptions (e.g. “agent not available”, “item misrouted” etc.) typically manifest themselves, for example, as missed deadlines.

Our approach for diagnosing exception causes is based on heuristic classification [5]. It works by traversing a diagnosis taxonomy. Exception types can be arranged into a taxonomy ranging from highly general failure modes at the top to more specific ones at the bottom (Figure 2). Every exception type includes a set of defining characteristics that need to be true in order to make that diagnosis potentially applicable to the current situation

When an exception is detected, the responsible process participant traverses the exception type taxonomy top-down like a decision tree, starting from the diagnoses implied by the manifest symptoms and iteratively refining the specificity of the diagnoses by eliminating exception types whose defining characteristics are not satisfied. Distinguishing among candidate diagnoses will often require that the user

get additional information about the current exception and its context, just as medical diagnosis often involves performing additional tests.

Imagine, for example, that we have detected a time-out exception in the “deliver product” step (see Figure 4). The diagnoses that can manifest this way include “agent unavailable”, “item misrouted”, and “item delayed”. The defining characteristics of these exceptions are:

- **agent unavailable:** agent responsible for task is unavailable (i.e. sick, on vacation, retired, etc.)
- **item misrouted:** current location and/or destination of item not match original target destination
- **item delayed:** item has correct target destination but is behind original schedule

The user then has a specific set of questions that he/she can ask in order to narrow down the exception diagnosis. If the appropriate information is available on-line, then answering such questions and thereby eliminating some diagnoses can potentially be at least partially automated.

2.4 Resolving Exceptions

Once an exception has been detected and at least tentatively diagnosed, one is ready to define a *prescription* that resolves the exception and returns the process to a viable state. This can be achieved, in our approach, by selecting and instantiating one of the generic exception resolution strategies that are associated with the hypothesized diagnosis. These strategies are processes like any other, are captured in a portion of the process taxonomy, and are annotated with attributes defining the *preconditions* that must be satisfied for that strategy to be applicable. We have accumulated roughly 200 such strategies to date, including for example:

- IF a process fails, THEN try a different process for achieving the same goal
- IF a highly serial process is operating too slowly to meet an impending deadline, THEN pipeline (i.e. release partial results to allow later tasks to start earlier) or parallelize to increase concurrency
- IF an agent may be late in producing a time-critical output, THEN see whether the consumer agent will accept a less accurate output in exchange for a quicker response
- IF multiple agents are causing wasteful overhead by frequently trading the use of a scarce shared resource, THEN change the resource sharing policy such that each agent gets to use the resource for a longer time
- IF a new high-performance resource applicable to a time-critical task becomes available, THEN reallocate the task from its current agent to the new agent

Since an exception can have several possible resolutions, each suitable for different situations, we use a procedure identical to that used in diagnosis to find the right one. Imagine, for example, that we want a resolution for the diagnosis “agent unavailable”.

We start at the root of the process resolution taxonomy branch associated with that diagnosis. Three specific strategies are available, with the following preconditions and actions:

- **wait till agent available:** IF the original agent will be available in time to complete the task on the current schedule THEN wait for original agent to start task
- **find new agent with same skills:** IF another agent with the same skills is available, THEN assign task to that agent
- **change task to meet available skills:** IF the task can be performed a different way using agents we have currently available THEN modify and re-assign.

The system user can prune suggested strategies based on which preconditions are satisfied, and enact or customize a strategy selected from the remainder. Note that the substantial input may be needed from the user in some cases in order to refine a generic strategy into specific actions.

2.5 Summary

Figure 5 summarizes the knowledge structure which serves as the basis of the approach described in the previous sections. It consists of two cross-referenced taxonomies: a specialization taxonomy of process templates and a taxonomy of exception types.

During process design time, process models are compared against the process taxonomy in order to identify possible failure modes. Once failure modes are identified, the exception type taxonomy provides links to appropriate detection and avoidance processes. During process enactment time, exception manifestations are compared against the exception type taxonomy in order to identify possible diagnoses. Once plausible diagnoses have been identified, the exception taxonomy provides links to resolution processes.

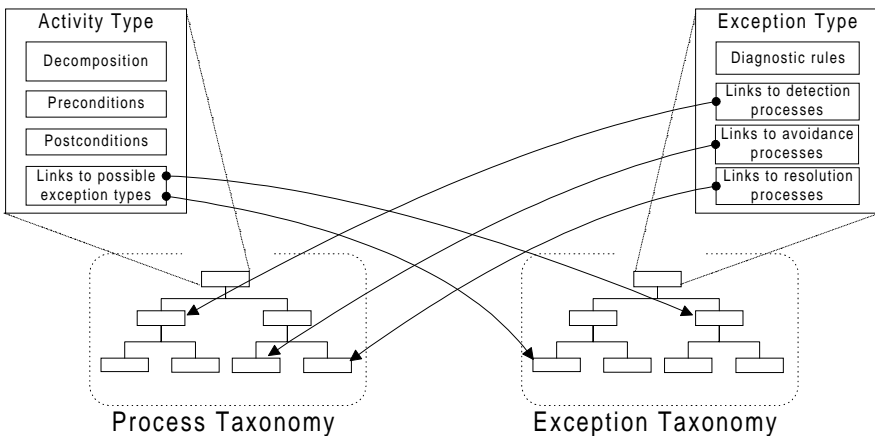


Fig. 5. Overview of Exception Handling Knowledge Structures.

3 Case Study: Barings Bank

The approach described in the previous section can be applied in order to help design robust new processes. It can also be a helpful tool when testing the robustness of existing business processes. This section illustrates how the method has been used in order to systematically expose potential dangers (and suggest possible fixes) in a well-known case of a failed business process.

In February 1995, 233-year old Barings Bank, one of the oldest and most respected investment houses in the United Kingdom, went bankrupt. The entire bank collapsed because of losses of \$1.4 billion incurred in a matter of days by a single young trader, Nicholas Leeson. Nicholas Leeson was a futures trader in the Singapore branch of the bank. For a number of reasons, which are still not entirely clear, Leeson began to engage in unauthorized futures trading in the Singapore exchange. Due to inadequate internal controls and other process failures, Leeson was able to maintain his unauthorized and highly risky activity undetected by the bank headquarters in London until the very end.

The collapse of the Barings Bank is one of the most dramatic and talked about recent disasters in financial markets. There exist several detailed accounts and analyses of why and how it happened (for example, [10, 27]). From our perspective, the Barings disaster is interesting because it was the result of a series of undetected exceptions in one of the bank’s secondary business processes: the futures trading process in Singapore.

In this section, we will demonstrate how the approach described in this paper can be used to systematically point out the gaps in the Barings trading process controls, as well as to suggest ways for closing those gaps.

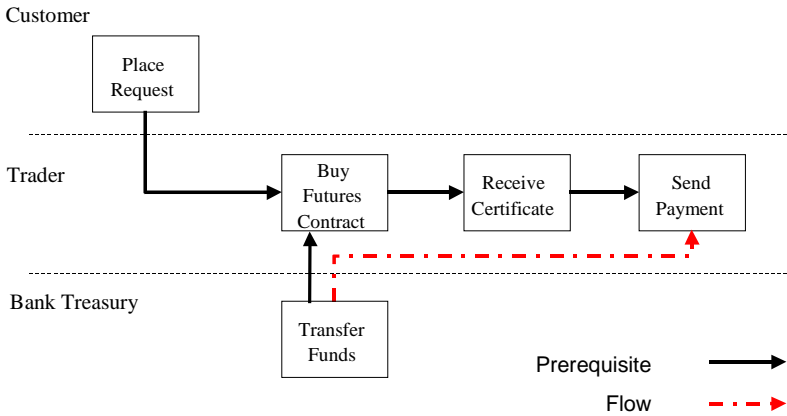


Fig. 6. The Barings Futures Trading Process

As described in the previous section, the approach begins with a “normal” model of the process. Figure 6 depicts a simplified but accurate model of the futures trading process, based on the descriptions contained in [10] and [27]. The model consists of boxes, which describe process activities, and lines, which describe various dependency relationships, that is, constraints that must hold true in order for the process to succeed. The following is a brief description of the process: When a customer requests a futures trade, the trader asks the bank headquarters for advances of funds in order to cover the customer’s margin account¹. Once the funds have arrived, the trader performs the trade, waits to receive the corresponding security certificate and finally pays the exchange. In an “ideal” world, a trader only performs trades when authorized to do so by customers, correct certificates are always received, and payment for trades exactly match the funds forwarded to the trader by the bank headquarters. These conditions are implied by the “prerequisite” and “flow” relationships, which are part of the “normal” process model.

The first step in our exception handling methodology consists of identifying the possible exceptions that are associated with each element of the “normal” process model. For simplicity we will only consider here exceptions associated with dependency relationships in the model.

According to the failure mode taxonomy shown in Figure 3, one possible exception of any prerequisite relationship is a prerequisite violation (“B without A”), that is, the possibility of activity B happening without a prior occurrence of activity A. In the context of the Barings trade process such violations would translate into unauthorized trading, unwanted security receipts and unnecessary payment (Figure 7).

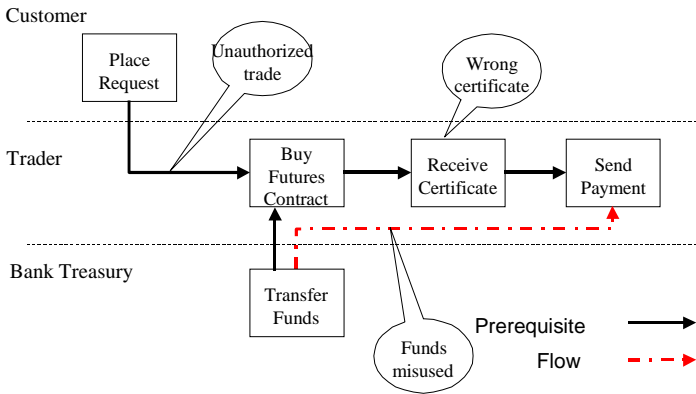


Fig. 7. The Barings Futures Trading Process with Associated Exceptions

¹ To find out more about derivatives trading and the meaning of margin accounts, the interested reader is referred to Zvi Bodie, Alex Kane, Alan J. Marcus, *Investments* (4th Edition), Irwin, 1998 (Part IV).

Likewise, one possible exception of a “flow” process is mismatch between the amount produced and the amount consumed. In the context of the Barings process this would translate into a misuse of headquarter funds.

After possible exceptions have been identified, the next step is to use the information stored in the exception type taxonomy (Figure 2) in order to find ways for avoiding or detecting the exceptions. It turns out that, because the trading process at Barings involves several independent entities (customer, bank, exchange) and requires some initiative from the part of the trader, there are no practical mechanisms for avoiding the exceptions. There were, however, several mechanisms for detecting them.

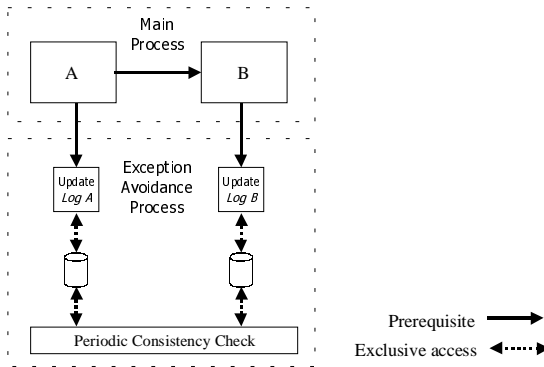


Fig. 8. Logging is a Generic Process for Detecting Prerequisite Violations

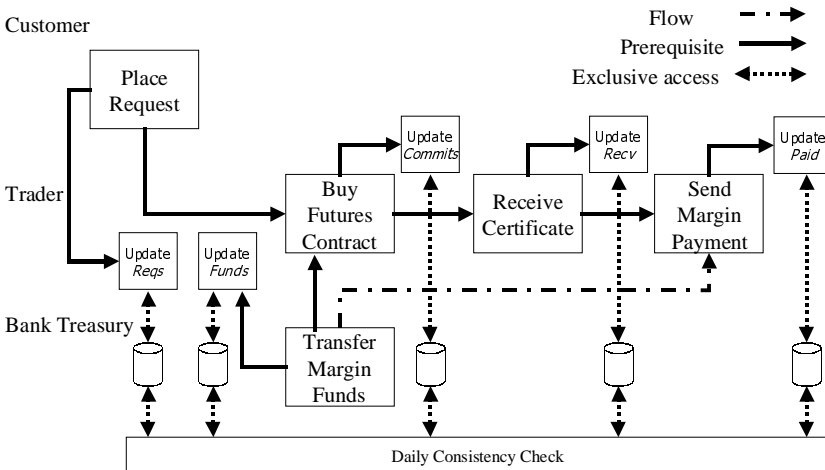


Fig. 9. Barings Process Properly Instrumented with Logging Processes.

For example, *logging* is one (out of several) generic mechanism for detecting prerequisite relationship violations (Figure 8). Logging involves recording all occurrences of activities A and B in some reliable storage medium and periodically conducting checks for prerequisite violations. In order for logging to be successful it is, in turn, required that (a) all occurrences of A and B are reliably logged and (b) the log can only be modified by the processes that do the logging.

If we insert a logging process for all dependencies listed in Figure 8 we get a model of a properly instrumented trading process (Figure 9).

At this point, we can compare the process derived using our approach with the actual Barings described in [10, 27]. It can immediately be seen that, although Barings did log some information about trades, it had two crucial gaps relative to the properly instrumented process of Figure 9 (see Figure 10):

First, it failed to log and compare the amount of funds forwarded by headquarters to the trader to the amounts actually paid by the trader for customer trades (in other words, the log labeled “Funds” in Figures 9-10 was missing from the Barings process). Second, Nick Leeson, in addition to being a trader, was also in charge of the back room operations in the Singapore branch. This gave him the authorization to modify the trades logs (and thus violated requirement (b) above of the logging process).

Nick Leeson was able to use these two gaps to his advantage as follows: Whenever he received a trade request from a customer, he requested an amount of funds far greater than what was required for the customer trade. He then performed the customer trade, as well as some additional unauthorized trades on his behalf. All of these trades were automatically logged into logs “Commits”, “Received” and “Paid” (see Figures 9-10). Leeson then erased the records of his unauthorized trades from logs “Commits”, “Received” and “Paid”. Therefore, at the end of each day, the log of “Requests” matched perfectly the other three logs. By not checking for discrepancies between the funds forwarded to Leeson and the total funds recorded at the “Paid” log, headquarters remained unaware of Leeson’s activities until it was too late.

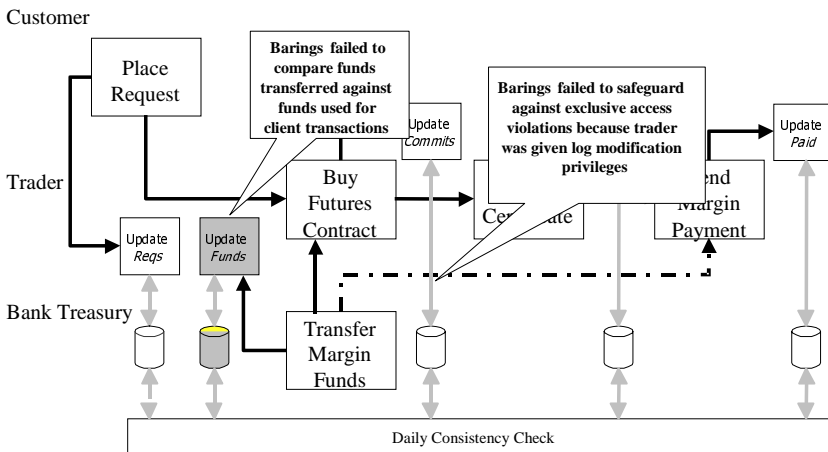


Fig. 10. Comparison between Ideal and Actual Barings Process

It is probably too simplistic to claim that the Barings disaster would have been avoided if the management of Barings had at their disposal knowledge-based exception handling methodologies, such as the ones described in this paper. Nevertheless, this exercise demonstrates that these methodologies and tools can be used in real-life cases to alert management of potential weaknesses and suggest ways for making vital business processes more robust.

4 Related Work

The approach described here integrates and extends two long-standing lines of research: one addressing coordination science principles about how to represent and utilize process knowledge, another addressing how artificial intelligence techniques can be applied to detecting and resolving conflicts in collaborative design settings:

One component is a body of work pursued over the past six years by the Process Handbook project at the MIT Center for Coordination Science [8, 20, 21]. The goal of this project is to produce a repository of process knowledge and associated tools that help people to better redesign organizational processes, learn about organizations, and automatically generate software. The Handbook database continues to grow and currently includes over 4500 models covering a broad range of business processes. A mature Windows-based tool for editing the Handbook database contents, as well as a Web-based tool for read-only access have been developed. A key insight from this work is that a repository of business process templates, structured as a specialization taxonomy, can assist people to design innovative business processes more quickly by allowing them to retrieve, contrast and customize interesting examples, make “distant analogies”, and utilize “recombinant” (mix-and-match) design techniques.

The other key component of this work is nearly a decade of development and evaluation of systems for handling multi-agent conflicts in collaborative design [15, 16] and collaborative requirements capture [17]. This work resulted in principles and technology for automatically detecting, diagnosing and resolving design conflicts between both human and computational agents, building upon a knowledge base of roughly 300 conflict types and resolution strategies. This technology has been applied successfully in several domains including architectural, local area network and fluid sensor design. A key insight from this work is that design conflicts can be detected and resolved using a knowledge base of generic and highly reusable conflict management strategies, structured using diagnostic principles originally applied to medical expert systems. Our experience to date suggests that this knowledge is relatively easy to acquire and can be applied unchanged to multiple domains.

The work described in this paper integrates and extends these two lines of research in an innovative and, we believe, powerful way. The central insights underlying this integration are that (1) business process exceptions can be handled by generalizing the diagnostic algorithms and knowledge base underlying design conflict and (2) the exception handling knowledge base can be captured as a set of *process templates* that can be retrieved, compared and customized using the principles embodied in the Process Handbook.

This work also constitutes, we believe, a substantive and novel contribution to previous efforts on exception handling, which have been pursued in the context of workflow [1, 9, 13, 18, 22, 25, 28-30] manufacturing control [14, 23, 26], model-based fault diagnosis [3, 7, 19], planning [3, 4], and failure mode analysis research [24]. Most workflow research has focused on languages for expressing correctness-preserving transforms on workflow models, providing no guidance however concerning *which* transforms to use for a given situation. There has been some manufacturing and workflow research on providing guidance for how to handle exceptions, but this has been applied to few domains (mainly software engineering and flexible manufacturing cell control) and/or has addressed a small handful of exception types. The planning work, by contrast, has developed a range of computational models but they are only applicable if the planning technology was used to develop the original work process. This is typically not the case for workflow settings where processes are defined by people rather than planning tools. Model-based fault diagnosis approaches use a single generic algorithm to uncover the causes of faults in a system without the need for a knowledge base of failure modes and resolution heuristics. This approach is predicated, however, on the availability of a complete and correct model of the system's behavior. This is possible for some domains (e.g. the analysis of electrical circuits) but not for many others including, we would argue, most collaborative work settings that include human beings and/or complex computer systems as participants. Model-based fault diagnosis also typically assumes that resolution, once a fault has been diagnosed, is trivial (e.g. just replace the faulty component) and thus does not provide context-specific suggestions for how to resolve the problem. Current work on failure mode analysis describes a systematic process. However, the actual work must be done by people based on their experience and intuitions. This is potentially quite expensive, to the extent that this analysis is rarely done, and can miss important failure modes due to limitations in the experience of the analyst [24].

5 Future Work

This chapter has emphasized the use of our exception handling knowledge base as a decision support tool for humans. Our ongoing work is also focused on connecting our technology with automated process enactment systems, such as workflow controllers and software agent systems. It is widely recognized that state-of-the art workflow technology provides only rudimentary support for exception handling [2, 9]. The result of our work will be a prototype implementation of a domain-independent *exception handling engine*, which oversees the enactment of a workflow script, monitors for exceptions and decides (automatically for the most part) how to intervene in order to resolve them. Given a "normal" workflow script, the engine first uses the exception handling knowledge base in order to anticipate potential exceptions and augment the system with additional actions that play the role of *software sentinels*. During enactment time, these sentinels automatically trigger the diagnostic services of the engine when they detect symptoms of exceptional

conditions. The diagnostic services traverse the taxonomy of exception types, select (possibly with human assistance) a diagnosis and then select and instantiate a resolution plan. The resolution plan is eventually translated into a set of workflow modification operations (e.g. add tool, remove tool, modify connection, etc.), which are dynamically applied to the executing workflow.

For further information about our work, please see the Adaptive Systems and Evolutionary Software web site at <http://ccs.mit.edu/ases/>. For further information on the Process Handbook, see <http://ccs.mit.edu/>

Acknowledgment

The authors gratefully acknowledge the support of the DARPA CoABS Program (contract F30602-98-2-0099) while preparing this paper.

References

1. E. Auramaki and M. Leppanen. Exceptions and office information systems. In B. Pernici and A.A. Verrijn-Stuart, editors: *Office Information Systems: The Design Process*, pp.167-182, North Holland Publishing Co., 1989.
2. P. Barthelmess and J. Wainer. Workflow Systems: a few Definitions and a few Suggestions. *Proceeding of the Conf. On Organizational Computing Systems (COOCS'95)*, pp. 138-147, 1995.
3. L. Birnbaum, G. Collins, M. Freed and B. Krulwich. Model-Based Diagnosis of Planning Failures. *Proceedings of the 8th National Conf. on Artificial Intelligence (AAAI-90)*, pp.318-23, 1990.
4. C.A. Broverman and W.B. Croft. Reasoning About Exceptions During Plan Execution Monitoring. *Proceedings of the 6th National Conf. on Artificial Intelligence (AAAI-87)*, pp. 190-195, 1987.
5. W. J. Clancey. Heuristic Classification. *Artificial Intelligence* 27(3), pp. 289-350, 1985.
6. T. Davenport. *Process Innovation: Reengineering Work through Information Technology*. Harvard Business School Press, 1993.
7. J. deKleer and B. Williams. Reasoning About Multiple Faults. *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86)*, pp. 132-9, 1986.
8. C. Dellarocas, J. Lee, T.W. Malone, K. Crowston and B. Pentland. Using a Process Handbook to Design Organizational Processes. *Proceedings of the AAAI 1994 Spring Symposium on Computational Organization Design*, pp. 50-56, 1994.
9. C.A. Ellis, K. Keddera and G. Rozenberg. Dynamic Change Within Workflow Systems. *Proceedings of the Conf. On Organizational Computing Systems, (COOCS'95)*, pp. 10-21, 1995.
10. S. Fay. *The collapse of Barings*. W.W. Norton, New York, 1997.
11. V. Grover and W. J. Kettinger, editors. *Business Process Change: Concepts, Methodologies and Technologies*. Idea Group Publishing, 1995.
12. M. Hammer and J. Champy. *Reengineering the Corporation: A Manifesto for Business Revolution*. Harper Business, 1994.

13. B.H. Karbe and N. G. Ramsberger. Influence of Exception Handling on the Support of Cooperative Office Work. In S. Gibbs and A. A. Verrijin-Stuart, editors: *Multi-User Interfaces and Applications*, Elsevier Science Publishers, pp. 355-370, 1990.
14. D. Katz and S. Manivannan. Exception management on a shop floor using online simulation. *Proceedings of the 1993 Winter Simulation Conference*, pp.888-96. 1993.
15. M. Klein. Conflict resolution in cooperative design. University of Illinois at Urbana-Champaign Technical Report UIUCDCS-R-89-1557.
16. M. Klein. Supporting Conflict Resolution in Cooperative Design Systems. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6), pp. 1379-1390, 1991.
17. M. Klein. An Exception Handling Approach to Enhancing Consistency, Completeness and Correctness in Collaborative Requirements Capture. *Concurrent Engineering: Research and Applications*, 5 (1), pp. 37-46, 1997.
18. T. Kreifelts and G. Woetzel. Distribution and Error Handling in an Office Procedure System. *Proceedings of IFIP WF 8.4 Working Conference on Methods and Tools for Office Systems*, Pisa, Italy, 1987.
19. M. Krishnamurthi and D.T. Phillips. An expert system framework for machine fault diagnosis. *Computers & Industrial Engineering* 22 (1), Jan. 1992, pp.67-84.
20. T.W. Malone, K. Crowston, J. Lee and B. Pentland, Tools for Inventing Organizations: Toward a Handbook of Organizational Processes, Proceedings of 2nd IEEE Workshop on Enabling Tech. Infrastructure for Collaborative Enterprises (1993) 72-82.
21. T.W. Malone, K. Crowston, J. Lee, B. Pentland, C. Dellarocas, G. Wyner, J. Quimby, C. Osborne, and A. Bernstein. Tools for inventing organizations: Toward a handbook of organizational processes. *Management Science*, in print.
22. P. Mi and W. Scacchi. Articulation: An Integrated Approach to the Diagnosis, Replanning and Rescheduling of Software Process Failures. *Proceedings of the Eighth Knowledge-Based Software Engineering Conference*, IEEE Comput. Soc. Press. 1993, pp.77-84.
23. S. Parthasarathy. Generalised process exceptions-a knowledge representation paradigm for expert control. *Proceedings of the Fourth International Conference on the Applications of Artificial Intelligence in Engineering*, 1989, pp.241-56.
24. D. Raheja. Software system failure mode and effects analysis (SSFMEA)-a tool for reliability growth. *Proceedings of the Int'l Symp. on Reliability and Maintainability (ISRM'90)*, Tokyo, Japan, pp. 271-77, 1990.
25. D.M. Strong. Decision support for exception handling and quality control in office operations. *Decision Support Systems* 8(3), June 1992, pp. 217-27.
26. A. Visser. An exception-handling framework. *International Journal of Computer Integrated Manufacturing* 8(3), May-June 1995, pp.197-203.
27. G. Zhang. *Barings bankruptcy and financial derivatives*. World Scientific Publishing Co, Singapore, 1995.
28. M. Klein, C. Dellarocas and A. Bernstein, editors. Special Issue on Adaptive Workflow Systems, *Computer-Supported Collaborative Work*, January 2000.
29. S. Ellis and K. Keddera. A workflow change is a workflow. Part II, Chapter 6 in this volume.
30. A. Agostini and G. de Michelis. Improving Flexibility of Workflow Management Systems. Part II, Chapter 7 in this volume.