Preface

Logic has been heralded as the basis for the next generation of computer systems. While logic and formal methods are indeed gaining grounds in many areas of computer science and artificial intelligence the expected revolution and breakthrough has not happened as yet. Notwithstanding the object oriented paradigm programming as well as processor design is still done in an imperative way, which has far-reaching consequences for the quality of software and engineering products.

A logical approach instead would offer many advantages such as machine-checked correctness, quick adaptability to design changes, dramatic reduction of maintenance costs, understandability of design, a far-reaching potential for the automation of the synthesis of the product from the design constraints, and so forth. Why then does not everyone follow the logical approach?

In the eighties is was beginning to dawn on the logic community that for most applications logic, as used then, might lack a vital ingredient which, on the other hand, is inherent in imperative languages and which no one would want to miss. What logic lacks is a simple and natural way to describe actions and change without facing inherent problems. In AI these problems center around what is called the *frame problem*. Without a solution to the frame problem—and its cousins—logic would continue to suffer from this shortcoming. Fortunately, the frame problem finally has been solved in such a way that the drawback is disappearing.

By now there is a number of formal variants of the solution to the frame problem. One consists in describing actions and change within the fluent calculus, a first-order Prolog-like formalism. The fluent calculus forms the basis of the contents of this book. It thus sets out from a basis which has overcome the drawback which held logic back for many years.

As mentioned there were the cousins of the frame problem yet to be tackled for a completely satisfactory solution. In particular, it is the so-called ramification and the qualification problem which belong to this family. Professor Thielscher in this book offers a convincing solution to these two accompanying challenges. I am deeply convinced that, as a consequence, there will be a renewed and strong interest in the logical approaches once these solutions will have become more widely known. I therefore wish this book the success it deserves for several reasons.

One of these reasons is the attractive mix of illustrative examples and formal precision, which makes this book easily accessible to a wide-spread readership. Another reason consists in the deep insight provided by the book into a fascinating topic which is central to our human thinking and has been a (mostly philosophical) issue for at least two thousand years. But the main reason is the level reached by the combined solutions to the frame problem and its cousins. Now programming in logic may comfortably include commands which call for logically defined changes without compromising in logical rigorousness. Similarly, engineering design, which always involves change, may now naturally be formalized in a logical setting with all the attractive advantages mentioned above. This includes the logical specification of agents in networks or autonomous robots which exchange information with each other as well as with human users on a most comfortable linguistic level.

It usually takes a number of years until fundamental insights diffuse through the community to a degree that the potential consequences materialize. May this book speed-up this diffusion process by finding many interested readers who spread out the news about another dawn of logic in computer science and artificial intelligence.

Wolfgang Bibel

Prologue

When John woke up he felt an uneasiness as so often these days. An instant later, however, it came to him that there was no reason for worrying anymore. The project he and his colleagues at the lab for months had sacrificed nearly everything for had finally come to a successful end the day before. So John relaxed, closed his eyes, and let his mind wander over the whole course of events again.

At the beginning there was this robot which was capable of performing rudimentary tasks such as moving around obstacles, grasping and handling objects, even climbing stairs (though it looked a bit clumsy to the attentive observer). Yet the robot was completely lacking the ability to solve tasks beyond these primitive ones on his own, that is, without John and his colleagues devising and telling him a minute plan of how to combine elementary actions in order to get the job done. A project was therefore established aiming at providing the robot both with insight into his own capabilities and with the ability to build a cognitive model of his environment. This, John argued, would enable the robot to do planning all by himself by means of reasoning when he has a certain goal in mind, that is, by drawing the right conclusions from what he knows as to the effects of his actions and from what his sensors tells him about his surroundings. A catchy name always being the basis for success of a project, they finally agreed on the acronym ELASER, meaning *Effective Logically Acting and Sensing Robot*.

It was obvious from the beginning that when explaining ELASER the effects of his actions it was impracticable to enumerate all conceivable situations in which an action can be performed and to state the result of its execution separately for each such situation. Actions had rather to be described by some sort of laws which specify the effects in general terms. In this way ELASER learnt, for instance, that grasping an object and carrying it from A to B always causes the object to be at location B and no longer at location A. In this context it turned out vital to provide the robot with a piece of information of universal nature. Namely, whenever he was told an action had such-and-such effects, he should assume this description be complete and so to conclude that moving around printed matter, for instance, does not change its contents. In order to test ELASER whether he had really grasped this crucial point, John asked him one morning to get a copy of the free local newspaper. The robot obediently walked out of the lab, spotted the right paper among different ones lying around in the front yard, correctly concluded that taking the newspaper into the building would not alter its being the local one, and so delivered it to John, who considered this an undeniable success.

Was it intuition or mere coincidence that John wanted to double-check the next morning? The night had brought some rain and it was still drizzling when ELASER left the building. Anxiously watching the robot out of his window, John saw that all newspapers were wrapped in transparent protective covers. To his great astonishment, ELASER ripped open the cover of one of the packages, tossed it away, and came back to the lab with a copy of the local newspaper soaked through. Criticizing the robot for his behavior, John and his colleagues learnt that he had no other choice: After all, ELASER explained to them, had he picked up and delivered the cover, the newspaper would still be lying in the front yard. For the single effect of carrying around an object like the cover is that only this very object changes its location, or so they had told ELASER.

Back to the drawing board. Apparently, the description of what happens if objects are moved needs to be split into two cases. Either there is no second item inside, or else there is, in which case both change their location. But what if a third object were placed inside the second one? This seems to require just another rule, which, however, still does not cover the case of four interlocking items, and so on (the alarming picture of an infinite Matryoshka, a nest of innumerable wooden puppets, entered John's mind). Pacing restlessly up and down his office racking his brain over this problem, John's eye fell on a book that had just been mailed to him with a note attached saying that it might be useful for their project. Could it be that a solution to their problem can be found in there, John thought, and so he opened the book and began to read. At least the introductory chapter seemed promising to him. The author first presented a basic theory of actions. He showed how to formally describe actions, including non-deterministic ones like rolling a dice, by specifying their general effects and applicability conditions. Furthermore, it was illustrated how to exploit this knowledge when reasoning about specific situations. The whole theory revolved around the paradigm that each action specification concentrates on what the action potentially affects, so that non-effects are to be inferred rather than being part of the description. The author called this adequacy of action specifications. So far, so good, John thought, but what if it is overly strict to suppose that nothing outside an action law is affected? Soon after starting off reading the second chapter, John realized that the latter was in fact entirely devoted to this question. There may be more to the impact of actions on the environment, so the author argued, than what is specified in action laws, which refer to the *direct* effects only. Actions may, however, have additional, *indirect* effects, which derive from general dependencies, or *constraints*, among the various properties that are used to describe the state of the environment. If John understood correctly, then this means, for instance, that an indirect effect of carrying around the cover is that the newspaper is being relocated, too. This additional effect is triggered by the general fact that two objects being stuck together can never be at different places. Accounting for indirect effects of actions, so the book continued, requires to meet two main challenges. First, the assumption needs to be suitably weakened which says that actions affect nothing but what is mentioned in action laws. Second, the aforementioned constraints often suggest, from the mere formal perspective, indirect effects which would never occur in reality and, hence, need to be sorted out. This was illustrated by an instructive example where toggling a light switch is concluded to have the magical side effect that another switch jumps its position rather than that light turns on, which one should have expected. John learnt that these two aspects together are commonly referred to as the "Ramification Problem."

With growing enthusiasm John kept on reading in hope of encountering a solution. As a matter of fact he found more than one. Unfortunately, all of them seemed perfectly reasonable to him—but only up to the point where the author proved their limited applicability. The author did so by discussing several scenarios for which the respective 'solution' either missed an obvious indirect effect (just like ELASER did in concluding that he had better rip open the cover of the newspaper, John thought), or proposed rather funny effects, which could never occur in reality. Finally, being faced with all these failures, the author introduced the concept of relations each of which directly links a single cause with a single effect. These *causal relationships* form the basis for the generation of indirect effects: Whenever a cause is brought about, then the additional occurrence of the effect is reckoned with. In this way all and also no 'phantom' indirect effects are obtained, so the author argued, provided, of course, the formal causal relationships both soundly and completely reflect causality in reality. Well, thought John, that is convincing, so we just have to provide ELASER with the knowledge that a change of an object's location *causes* any object inside to change its location as well. However, he was not looking forward to telling his colleagues that they have to draw up by hand and feed ELASER all necessary causal relationships. Fortunately the author showed how these relationships can be automatically extracted from much more general knowledge. The chapter concluded with an axiomatization in formal logic of the whole theory of causal relationships as means to solve the Ramification Problem. To John's satisfaction, he noticed that this axiomatization was based on the same principle they had used for designing ELASER—a principle which was called "Fluent Calculus" in the book. So they seem to have made a lucky decision. John remembered how difficult it was in the beginning to convince his colleagues of the advantages of this Fluent Calculus over the other's two favorites, the so-called Situation Calculus and Event Calculus, respectively, when it comes to inferring non-effects of actions. The author of the book shared his conviction, and so John and his colleagues just had to put the proposed axiomatization on top of the one already existing in ELASER.

Having redesigned ELASER following the instructions, the project members sent out the robot to fetch a newspaper each and every morning, and the robot did his duty worthily. He always came back with the right paper, plain on sunny days and safely wrapped in a protective cover whenever it had been raining. Until this one day which John will never forget. As usual, ELASER had left the lab sometime in the morning. Eventually, however, John and his colleagues realized that his return was long overdue and still there was no sign of him coming back. Anxiously recalling the disastrous morning when ELASER delivered the soaked newspaper, John followed the robot's path to the front yard. Standing next to a package with the local newspaper inside, ELASER was totally paralyzed; even when John enquired of him what had happened there was no reaction at all. John's last hope was that investigating the package lying nearby would shed some light upon the matter. Indeed he made a surprising discovery. Some rascal, who presumably had watched ELASER picking up a newspaper every morning, had teased the robot by introducing a brick into the package, which thus was too heavy for poor ELASER. Still, however, this did not account for the total blackout. When they had managed to run the robot again, he explained to John and his colleagues that he knew the only precondition of picking up an object is that it must be reachable. Now that was clearly the case when ELASER tried to lift the package, so the formal specification implied, with unerring logic, that success of this action is guaranteed. Nonetheless the expected effect failed to materialize, which entailed a logical contradiction so that ELASER's whole conception of the world broke down instantaneously. By the next morning, they had taught the robot that a second precondition for being able to lift an object is that it is sufficiently light. But then they watched ELASER anxiously ripping open the protective cover around the newspaper again, this time in order to make sure that there be no brick or any other heavy item inside. To John that seemed rather ridiculous. After all, it is highly unlikely that a newspaper package cannot be fetched on account of its weighing too much.

Back to the drawing board again. The problem was to find a suitable way of specifying action preconditions which need not be verified each time prior to assuming that the action in question be executable. John turned to the book which had already served him so well. Indeed the third chapter, entitled "The Qualification Problem," was devoted to exactly their new question. The author started off with arguing that in real-world environments most actions have many more preconditions than one is usually aware of. The reason for this unawareness is that most of these conditions are so likely to be satisfied that they are assumed away unless there is evidence to the contrary. Conditions of this sort are called "abnormal," and if their presence prevents the successful performance of an action, then the latter is considered "abnormally disqualified." Any particular situation, then, is reasonably presupposed to being 'normal' as long as this does not conflict with what is known or has been observed. John vaguely remembered this as a standard technique to deal with assumptions that are made by default because they are usually—but not necessarily always—correct. Yet the author illustrated that making the right assumptions in the context of the Qualification Problem is trickier than usual.

Roughly speaking, this is a consequence of the dynamics inherent in action theories, which implies that abnormalities may naturally arise for reasons of causality. If John understood correctly, the crucial point was the following. Suppose, for instance, ELASER had been told in advance that somebody had planned to add a brick to the newspaper package. Then it would have been reasonable for the robot to assume that this action had been successful and, hence, lifting the package would have been disqualified thereafter. But the application of the aforementioned standard technique would equally well suggest another course of events, namely, that introducing the brick is impossible in the first place due to some mysterious unspecific reason. Therefore any solution to the Qualification Problem, the author argued, requires an account of abnormal preconditions of actions which are brought about as side effects of previous actions. Side effects being nothing else than indirect effects, the preceding solution to the Ramification Problem turned out to furnish a ready fundamental for a solution to the Qualification Problem. The book continued with showing how to seek explanations in case an action surprisingly fails at some point. Even the rare case of inexplicable disqualifications, "miraculous" they were called, had been considered. Like the previous one, the chapter on the Qualification Problem concluded with an axiomatization in formal logic of the entire action theory.

Before John called an assembly of the project group in order to announce that he had found the solution to their new problem, he read through the final, comparably short chapter. The author expanded the connection between the Ramification and Qualification Problem even further. Just like actions may turn out unqualified for some abnormal reason, so the argument went, there may happen exceptions to the occurrence of indirect effects. This time the existing solution to the Qualification Problem in turn furnished a ready approach to this generalization of the Ramification Problem.

Finally, John thought contentedly daydreaming in his bed, with the invaluable assistance of the Book they had brought the project ELASER to a successful end. After having redesigned the robot so as to being capable of coping with the Qualification Problem, they continued making all kinds of tests for weeks. ELASER had passed them all effortlessly. Thinking of that, John happily nodded off again, with a hardly noticeable smile on his face. He dreamed about ELASER strolling around campus when he suddenly bumped into a famous philosopher, who straight away started arguing that the robot lacks conscious understanding of anything he is doing. When ELASER demanded a precise definition of what he meant by conscious understanding, the philosopher finally defined it as being a property only carbon-based brains can possess. I can live with having no understanding of that kind, ELASER thought walking off with a slight shake of his head, leaving the philosopher.

This, however, is not the story of ELASER. Nor is this the story of John.

This is the Book.