CHAPTER 1

Problem Solving, Simulation, and Computational Red Teaming

1.1 INTRODUCTION

Problem solving is present nowadays virtually everywhere, from industry and research-related activities to daily life, and consists of various methods, grouped in a more or less systematic manner, to find solutions to problems. From a very colloquial understanding, problem solving refers to the process in which one attempts to overcome various real-world difficulties using systematic, including trial and error, steps.

The concept of problem solving is used in many disciplines, where practitioners and researchers treat it from different perspectives and describe it using different terminologies. At one extreme, in psychology, problem solving is mainly concerned with the mental processes involved in finding solutions. At the other end of the spectrum, in computer science and related fields (e.g. computational and artificial intelligence, cognitive science, autonomous systems, and robotics), problem solving is mainly associated with computerised processes, in which computer simulation and computational models are employed to find the desired solutions. Between these two extremes, numerous other fields of activity – such as multiple branches of engineering, industrial design, business and economics – employ various combinations of methods for problem solving, where these methods range from introspection and behaviourism, to prototyping and design of experiments, and further to computer modelling and simulation.

Of these, computer modelling and simulation is currently the most preferred working tool in most, if not all, areas mentioned above [35], due to their low cost, speed, and flexibility. However, computer simulation may not be necessary in every real problem-solving situation. Sometimes the real system in which a problem needs to be solved can be accessed as is, and thereby can be studied easily or economically in the physical world, without transferring the problem into conceptual models that are solved through simulation. However, in most cases, the need for evidence-based

Simulation and Computational Red Teaming for Problem Solving, First Edition.
Jiangjun Tang, George Leu, and Hussein A. Abbass.

decision making and the complexity of real-world problems necessitate the use of modelling and simulation. As a result of the unprecedented development of digital computing, we are able to use simulation on a significantly larger scale than ever before. Today's digital computers offer sufficient computational power to run models of real-world problems of virtually any complexity at very low cost and with very little risks involved, thereby rendering the engagement of actual systems less attractive than in the past. Thus, while computer simulation was once reserved for fields such as science and engineering, it is today pertinent in all aspects of our life and has extended to the humanities, social sciences, and psychology.

Thus, it is natural to start our book about simulation by discussing problem solving to: (i) disambiguate the concept of simulation and (ii) position the concept of simulation within the larger problem-solving concept. In addition, another reason for spending some time discussing problem solving before proceeding with the core discussion on simulation is to explain to the reader why and how computational red teaming (CRT) and simulation are related, given that this relationship only becomes meaningful within the context of problem solving.

Thus, in this very first chapter, we take the reader on a brief yet useful journey through the problem-solving process to emphasise the roles in problem solving of the two concepts that create the topic of this book: computer simulation and CRT. To fulfil this purpose, we structure the remainder of the chapter in two parts. The first part describes the problem-solving process and emphasises the role of simulation within this process, while the second part describes the role of CRT in relation to problem solving and simulation.

1.2 PROBLEM SOLVING

Problem solving can be considered the process of finding solutions to difficult or complex issues. According to this relaxed definition, it is acceptable to consider that solving a problem equates with finding a way out, treating a symptom or simply accomplishing a task that is not trivial. This is an accustomed view of problem solving over many fields of activity, which may be accepted in highly informal contexts. However, from a rigorous perspective, solving a problem refers to finding and eliminating the root cause that generated the problem, where a problem is seen as a risk whose time has come – that is, it has a cause that has manifested. According to this view on problem solving, in the following, we briefly present the major steps of the problem-solving process, with the purpose of positioning simulation in the context of this book. These steps are problem recognition, problem definition, and problem solving.

1.2.1 Recognising the Problem

Problem recognition is a matter of proactive diagnosis. In relation to any activity (where we understand an activity to be any combination of concepts, such as products, services, processes, phenomena, and systems), the stakeholders involved in that activity need to devote continuous attention throughout its entire lifecycle to any potential issue that does or may emerge. To ensure this observant behaviour towards the activity, stakeholders need to be in a continuous state of situational awareness, which can be reached only through acquisition of a considerable amount of knowledge about the past, present, and potential future states of the activity. Problem recognition manifests when one or more stakeholders recognise a difference between the desired and actual state of the activity, which has a significant effect on further decision making related to the activity.

Sometimes the problem is created by shaping the goal state to be different from the current state. For example, a chief executive officer (CEO) of a company may recognise that the performance of the company has been steady for a long time, thereby running the risk of becoming too predictable, which makes the company an easy target for its competitors. Thus, the CEO designs new key performance targets with the primary aim of changing the status quo.

Skills and experience, together with the right knowledge-acquisition tools, are needed for stakeholders to be able to recognise problems that need to be solved and/or trends that can lead to problems. To better understand the involvement of knowledge acquisition in this process, a comprehensive view on knowledge-acquisition methods is offered in Chapter 11.

1.2.2 Defining the Problem

Once the stakeholders identify the existence of a problem, the natural step towards solving the problem is to define it. This means that they must structure, state, and represent the problem in an intelligible way, so that the problem can be communicated to peers for further consideration and solving. As with problem recognition, this part of problem solving is also a process that involves acquiring a substantial amount of knowledge about the context, so that it can be properly stated, structured, and represented. In addition, similar to problem recognition, Chapter 11 is a useful reading for understanding how knowledge about problems can be acquired.

It is critical to understand that these three aspects – structure, statement, and representation – are useful for facilitating further actions, such as development and evaluation of potential solutions; thus, problem definition must capture the underlying causes, not (only) the symptoms of the problem. For example, a road management authority observes that roads in a particular area are substantially jammed, and concludes that there is a traffic problem. Thus, when the problem is defined, the traffic jam is not what defines the problem – it is only a symptom. To facilitate solving the problem, it needs to be further quantified in terms of the amount of traffic that exceeds the road capacity. In this way, potential root causes are implied (e.g. insufficient roads, unexpected demand or both), so the problem is defined in a way that can guide analysts to choose appropriate methodologies.

1.2.3 Solving the Problem

1.2.3.1 Establish System Boundary A problem does not exist in a vacuum. A problem is associated with a certain system (or phenomenon) through its causes, its effects, interactions with it or a combination of these factors. In turn, this system is situated within an environment. Thus, when solving a problem, it is necessary to establish the boundaries between the system to which the problem is related and the rest of the environment. Whether these boundaries are soft or strong is a matter of the type of problem, the type of system, the type of environment, and ultimately the objectives stated by the system's stakeholders. However, a boundary needs to be considered to ensure further traceability and solvability of the problem.

1.2.3.2 Formalisation After a system boundary is established, the problem can be articulated in a language that is universally understandable by various stakeholders involved in its solving. Thus, the problem, which is real and attached to a real system, can be further abstracted and formalised in a language that allows analysts to discuss it with peers and solve it. To achieve this, the problem is transferred in a model, which is a convenient representation of its real-world version, by using a certain formalism. There are numerous ways to build models that represent real-world problems, such as natural language, prototypes, diagrams, logical constructs or mathematical equations. Each of these has advantages and disadvantages, depending on the purpose of the modeller. A detailed discussion on these is offered later in the book, in Chapter 5. The resultant model is then the foundation on which a computer simulation is designed and performed.

1.2.3.3 Solve the Model Simulation is one of the tools used to both understand the problem and system under consideration, and, when coupled with an appropriate search algorithm, to solve the conceptual model of the problem at hand. An attempt to understand this view is a necessary effort to disambiguate the concept of simulation.

It has become normal across various disciplines for simulation to be viewed as the process of emulating or imitating a real-world phenomenon [35], which is a rather colloquial view that does not sufficiently emphasise the role of a model in a simulation environment. While simulation can indeed be associated with the idea of emulating a real-world phenomenon, it is actually the model representing that phenomenon that is simulated, with the purpose of understanding the model space of behaviours to further understand the phenomenon's space of behaviours.

Thus, depending on how the model represents the real problem, various types of simulation can be considered. For example, let us consider a very simple real-world situation of a boat moving on a river from Point A to Point B. The problem we would like to solve here is to investigate how long it takes to move from A to B, given different speeds of the boat. We need to solve this problem in a way that does not involve the boat itself or the river because finding the solution in this way would not be beneficial for multiple reasons, such as cost, logistics, time, and effort. Thus, we need to solve it in a way that does not involve the real thing, so we build a model to obtain a representation of the real situation, which is good enough to serve our purpose.

Let us consider three cases, where we build three models that represent the boat problem at an increasing level of abstraction: (1) we build a small toy-like wooden boat and place it in a water recipient; (2) on a piece of paper, we draw a line representing the river and a bubble representing the boat; and (3) we represent boat movement using Newton's equations of motion. Each of these three models abstract quite well the real boat problem, and we can use various settings to explore the behaviours of these models. Then, assuming that the models are sufficient representations of the real-world problem, we can infer the behaviours of the boat on the river under different circumstances. Simulation is when we actually run these models under the conditions that we consider relevant to the investigation. In (1), we may move the toy boat by hand along the water recipient at different speeds, conveniently scaled to resemble the spatial/physical constraints present in the real world. In (2), we may add to the drawing several vectors with different directions and amplitudes, representing different velocities of the boat. In (3), we may give different values to variables and parameters that are part of the laws of motion. By using the models to do this, we actually simulate the models, and create the premises for further analysis of their behaviour under different conditions.

1.2.3.4 Sensitivity Analysis After a model is simulated under different conditions of interest, we may perform an analysis on the resultant outputs to infer useful conclusions about its behaviour. In the classic view, where simulation is seen as just an emulation of reality of some kind, the analysis process may not be even relevant or pertinent to the purpose of simulation. However, when we consider simulation from a problem-solving perspective, then we clearly see a motivation, utility, and meaning in performing analysis on the output resulting from simulating the model.

Apart from the analysis of the simulation output, which is needed to infer conclusions on the space of behaviours, sensitivity analysis also needs to be performed, once the model (and its computer implementation, in the case of computer simulation) is in place. Sensitivity analysis investigates how the uncertainty in the output of a model can be related to different sources of uncertainty in its inputs. It usually involves calculating the outcomes under various assumptions to determine the effect of deviations in the inputs on the outputs. Sensitivity analysis is used as part of verification and validation of the model – which we describe in the next paragraph – leading to a number of benefits to the model and the resultant simulation process, such as showing the robustness of the results in the presence of uncertainty, enabling a deeper understanding of the relationship between inputs and outputs, unveiling errors in the model, simplifying the model by fixing those inputs with no or insignificant/irrelevant effect on the output, and identifying redundant parts of the model structure.

1.2.3.5 Verification and Validation However, since simulation refers to running the model, then what we obtain is outputs of the model. Consequently, the analysis of these outputs shows the behaviour space of the model, not of the real-world problem. Thus, only if the model is a valid representation of the real problem, and the simulations were properly run on the model, will we be able to infer and import the simulation conclusions (generated by the model) onto the real problem. These two aspects refer to the validation and verification processes that we need to perform to ensure that the conclusions in the behavioural space of the model apply to the behavioural space of the real problem as well.

In general, verification and validation are the processes that evaluate the answers to the questions *did we build the product correctly?* and *did we build the correct product?*, respectively. In the context of problem solving and simulation, these questions become *did we solve the model correctly?* and *did we solve the correct model?*, respectively. Therefore, through verification we check if the way the simulator was implemented based on the model (and consequently, the way the simulation produces outputs) is correct. Therefore, through verification, we check if the way the simulator was implemented based on the model (and, consequently, the way the simulation produces outputs) is correct. Further, through validation, we check if the conceptual model we created actually represents the real problem in a necessary and sufficient manner, so that it can serve our purpose.

1.3 COMPUTATIONAL RED TEAMING AND SELF-'VERIFICATION AND VALIDATION'

While this is the introductory chapter of this book and should be accessible to all readers, we do recommend a brief consultation of Chapter 13 before continuing reading this section for those readers who are not at all familiar with red teaming concepts. This will help understand the involvement of red teaming in the topic of simulation.

In brief, red teaming is a practice in which an entity, denoted the blue team, employs another entity, denoted the red team, to act as its hypothetical opponent. By employing a red team that emulates the opponent's intentions and challenging actions towards the blue team's system, the blue entity can test and evaluate its own course of actions in relation to specific objectives. There may be many types of blue, which we call red teaming objects, to illustrate the object to which the red teaming process is applied. For example, we can apply red teaming to plans, decisions, ideas, strategies, models, simulations, systems, the operation of systems, and the performance of tasks. In addition, red teaming may be used by blue for a variety of reasons, which we call red teaming purposes. Some examples of these purposes can be unveiling vulnerabilities, discovering blind spots, verifying and validating assumptions, and exploring the space of possibilities.

Let us now emphasise a certain combination of these objects and purposes. Let us assume a situation in which the red teaming objects are a model of a real system and the resultant simulation, and the red teaming purposes are verification and validation of assumptions. Now, we start seeing a certain connection between simulation and red teaming in the context of problem solving. If we express the above situation in other words, we obtain the following.

We have a real-world problem that we need to solve. We transfer the real problem into a conceptual model that represents the real problem at an assumed necessary and sufficient degree of fidelity, and hope they are equivalent with respect to our purpose of finding the solution. Further, we transfer the conceptual model into a computer (numeric) simulator, and then simulate the model under different conditions, and reach conclusions about the behaviour of the model – that is, we find the solution to the model. Then, assuming that (1) the simulation reflects the conceptual model and (2) the model reflects the real-world problem, we infer that the solution for the model can also be applied to the real-world problem. Thus, we solve the problem using simulation. However, we can never be sure whether (1) and (2) are true – we can only hope that our assumptions were correct. We may have made an implementation mistake, and the simulation did not run the model correctly. Also, we may have made incorrect assumptions when creating the conceptual model, so the model did not suit the real-world problem. By employing a red teaming exercise to test these two aspects, we use red teaming as a verification and validation mechanism for our simulation and problem solving, respectively.