HÖGSKOLAN SKÖVDE

Institutionen för kommunikation och information
Examensarbete i datavetenskap 25 p
D-nivå
Vårterminen 2005

Providing Adaptability in Survivable Systems through Situation Awareness

Daniel Öster

HS-EA-DVA-2005-001

Providing Adaptability in Survivable Systems through Situation Awareness

Submitted by Daniel Öster to the University of Skövde as a dissertation towards the degree of M.Sc. by examination and dissertation in the School of Humanities and Informatics.

2006-09-12

I hereby certify that all material in this dissertation which is not my own work has been identified and that no work is included for which a degree has already been conferred on me.

Signed: _____

Providing Adaptability in Survivable Systems through Situation Awareness

Daniel Öster

Abstract

System integration, interoperability, just in time delivery, window of opportunity, and dust-to-dust optimization are all keywords of our computerized future. Survivability is an important concept that together with dependability and quality of service are key issues in the systems of the future, i.e. infrastructural systems, business applications, and everyday desktop applications. The importance of dependable systems and the widely spread usage of dependable system together with the complexity of those systems makes middleware and frameworks for survivability imperative to the system builder of the future. This thesis presents a simulation approach to investigate the effect on data survival when the defending system uses knowledge of the current situation to protect the data. The results show the importance of situation awareness to avoid wasting recourses. A number of characteristics of the situational information provided and how this information may be used to optimize the system.

Keywords: Survivability, Dependability, Network Based Defense (NBD), Situation Awareness, Simulation

Acknowledgements

Thank you Alice, your love and support made it possible.

Contents

Table of figures

Table of tables

1 Introduction

System integration, interoperability, just in time delivery, window of opportunity, and dust-to-dust optimization are all keywords of our computerized future. As the appreciation of how computerized systems can be used grows, the society's dependence on these systems increases. Imagine a day where the banking systems are inaccessible, the electricity is gone, or the phone system is out of order. The aim of this dissertation is to investigate how to manage these kinds of problems.

The system environment considered is the unbounded network environment (Ellison, Fisher, Linger, Lipson, Longstaff & Mead, 1997). The problems considered in this thesis are those arising from a hostile environment and malicious attacks in the context of Network-Based Defense (NBD) (Warston and Person, 2004). Bendz, Johannisson, Jönsson, and Öhlund (2003) claim that NBD can be seen as a method and that NBD is the method chosen by the Swedish Armed Forces (SwAF) to re-orient the defense in line with "the new warfare". The aim of NBD is to optimize the resource utilization based on a common situational picture (Wik, 2003). This leads to the conclusion that the systems participating in a network-based defense must be adaptive (Warston and Person, 2004). Adaptive and dependable systems are difficult to build, especially when the system is composed of separately developed components.

Survivability is a difficult concept to capture in a single intuitive and usable definition. Ellison et al. (1997) present an informal definition of survivability on which much of the subsequent work on survivability is based. Knight, Strunk, and Sullivan (2003) address the differences between a formal definition and an informal definition. They also propose a framework within which it is possible to outline a definition for survivability, for a specific system. Since the view of survivability differs slightly in these sources, survivability is presented by a short description of each view. The descriptions are followed by a comparison and a discussion of the views. Contrasting dependability and survivability is a task that is becoming increasingly difficult. Avižienis, Laprie, Randell and Landwehr (2004) claim that dependability and survivability are "...essentially equivalent in their goals and address similar threat" while Knight et al. (2003) say that "...define survivability in a way that emphasizes the need to specify systems that can provide different forms of service, each with its own complete set of dependability requirements, under different conditions.". Survivability in distributed system uses much of the taxonomy from dependability and the taxonomy of dependability assimilates much of the discussion around survivability. Consequently, there is no sharp boundary between the two. The reason for discussing the differences is that depending on whether dependability or survivability is emphasized the solutions tend to differ in their basic assumptions and their view of the system environment and the system it self.

The thesis is outlined as follows. First, the background is presented and the problem is defined. This is followed by a description of the simulator design and the simulation setup. The results and the conclusions from the project are then presented.

2 Background

This chapter provides the background and the problem definition of this thesis. The background contains the goals and ideas of network-based defense and the ontology of dependability and survivability. After this, a description of situation awareness is given. Last, two theoretical frameworks are described and contrasted in an effort to find a framework.

2.1 Network-Based Defense

Network-based defense is the breakdown of the traditional military structure (Wik, 2003). Instead of dividing the forces in different armed services, the classification of the forces is based on their capabilities (Wik, 2003). Depending on the mission, and situation, forces that previously belonged to different armed services can work together to get the most value from the available resource. The defense control structure need to change in order for this new organization to be efficient. Situation adapted control hierarchy with fewer organizational levels needs to replace the static control hierarchies within the defense. The agile control structures will lead to faster decisions and an agile organization. Wik (2003) further claims, “NBD is about managing defence resources as efficiently as possible...” a statement that is a close generalization of the aim of this thesis. Network-based defense is the Swedish counterpart of network-centric warfare (Bendz et al., 2003). Network-centric warfare should not be confused with technical networks or be seen as having a technical focus. The focus in network-based defense is networking and organizational behavior (Tolk and Pullen, 2003). To build a network-based defense, a reengineering of the SwAF is undertaken. The Swedish Armed Forces Enterprise Architecture (FMA) will support the reengineered defense (Bendz et al., 2003). The battle arenas that together describe the reality are the physical, the informational, and the behavioral arena (Wik, 2003). The physical arena constitutes land, sea, air, space, and all other physical things. The information area is all information, intelligence reports, sensor data, etc. The behavioral arena consists of the human behaviors and psychological processes.

2.2 The ontology of Dependability

Dependability dates back to 1980, when the tactical committee on Fault-tolerant computing from the IEEE Computer Society formed a committee on “Fundamental concepts and Terminology”. Their work led to the book “Dependability: Basic Concepts and Terminology” (Laprie, 1992). Laprie (1992) presents the attributes of dependability and the taxonomy of faults. The discussion of dependability is based on Avižienis et al. (2004).

The original definition of dependability was “the ability to deliver service that can justifiably be trusted”. This definition is a good intuitive definition since it describes what a user expects from a dependable system. The definition is not usable from a system builder’s perspective since it is impossible to verify that a system complies, but it gives an informal understanding of the desired system property. An alternative definition of *dependability* is “dependability of a system is the ability to avoid service failures that are more frequent than and more severe than is acceptable” (Avižienis et al., 2004). This definition is more appropriate for developing and evaluating dependable systems. It is more appropriate since it defines dependability in less emotional terms by replacing “justifiable” with a measurable frequency and the difficult concept of trust with an acceptance.

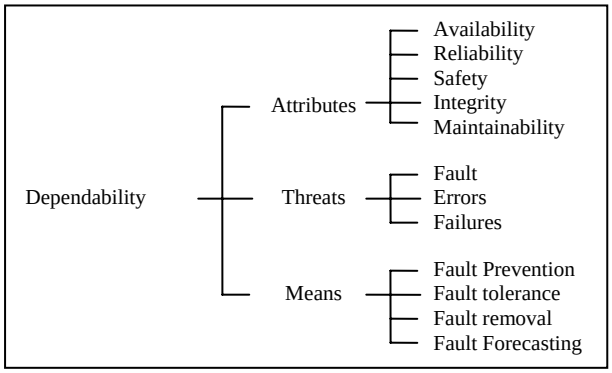

Figure 1. Dependability after Avižienis et al. (2004).

Dependability as a concept is composed of five attributes, threats to these attributes, and means and methods to achieve dependability, as shown in

Figure 1. *Availability* describes the accessibility of the correct service provided by the system. *Reliability* is a measurement of how long a system can provide the service, uninterruptedly. *Safety* is a system's ability to avoid harmful effects on the system environment, including its users. *Integrity* is a system's ability to avoid unintended system changes. *Maintainability* is the systems ability to be repaired and its readiness to undergo modifications.

2.3 The ontology of Survivability

This section describes survivability. The first subsection describes an informal definition of survivability and the domain of survivability, followed by a control theoretic perspective on survivability.

2.3.1 Informal definition of Survivability and its domain

In survivability, we seek to minimize risk and preserving the essential services in unbounded network environments when the system is suffering from attacks, failures, or accidents (Ellison et al., 1997). Ellison et al. (1997) list properties that an unbounded network environment displays:

- An unbounded network environment is composed of multiple participating domains. Each domain has an authentication and an administration policy. There is no central authority.
- There is no global visibility. No participating system has full knowledge of the number of nodes in the network, or their capabilities.
- The involved participants must agree upon the communication and interoperation.
- The difference between legitimate users and malicious attackers is their behavior.

An unbounded network environment originates from organizational integration where system communication crosses organizational boundaries to improve the effectiveness of the involved organizations.

Ellison et al. (1997) discuss survivability as maintaining a balance between different quality-attributes, attributes that are similar to those of dependability. Further, Ellison et al. (1997) define *survivability* as "…the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents".

The first part of the definition places the focus on the mission. The *system mission* describes the current purpose or assignment of the system. The system mission must be fulfilled even if system elements are lost or replaced. The *system value* is a measurement of the system contribution, a measurement of how the service provided by the system is valued by those served. The evaluation of the system value must be performed in the context of the system environment and with attention to the system mission. A complete evaluation of the system value is therefore not possible at design time. The definition of survivability explicitly recognizes the timely completion of the mission as a requirement for mission success. Most systems have requirements on their timeliness, either explicit as the in the case of real-time systems, or implicit arguably as in the case of all other systems. The last part of the definition recognizes attacks as a threat to survivability. This is important since it implies intent and intelligence, as opposed to chance and random failures. The chance of malicious attacks orchestrated by an intelligent adversary puts harsh requirements on the design, implementation, and composition of these systems. Within survivability, the effects from attacks or failures are not separated. It is important to make the possibility of an intelligently planned attack explicit to avoid assumptions of independent faults.

It is crucial that a survivable system continues to deliver the essential services even when subject to attacks, failures, or accidents (Ellison, Linger R and Longstaff, 1999). The essential services are the subset of services that are required to fulfill the system mission. The classification and protection of essential services are central to building, and maintaining, survivable systems (Ellison et al., 1997).)

2.3.2 Control theoretic approach to Survivability

Sullivan, Knight, Du and Geist (1999) approach survivability through a control theoretic perspective. They suggest that the design of survivable systems should be according to a decentralized, hierarchical, discrete-state, adaptive architectural style. Knight and Sullivan (2000) list system and environment characteristics that affect survivability: *System size;* critical information systems are large, both geographically and in number of computing nodes. *Externally observable damage;* the system may be so extensively damaged that the damage is visible to the user, preferably as degraded service or degraded quality of service. *Damage and repair sequences;* damaging events may be dependent and not mutually exclusive. *Time-dependant damage effects;* the effect of the damage to the system increases with time. *Heterogeneous criticality;* the requirements for dependability change with time and system needs. *Complex operational environments;* the operational environments are complex and unpredictable. They carry risks of natural, accidental, and malicious nature. *Time-varying operational environments;* the operational environments of the systems are changing with time. Knight et al. (2003) suggest a definition of survivability based on a survivability specification. A survivability specification specifies all foreseeable forms of services together with their requirements and environmental conditions. A set of valid system transitions describes the transitions between the different survivability specifications and is included in the survivability specification.

2.4 Situation awareness

Situation awareness is a wide concept that essentially means that the system has information about the current situation. The detail of situation awareness can vary but even a low detail of knowledge on the situation can prove valuable when adapting a system.

2.5 Design and architecture considerations of survivable systems

Dependability is the traditional terminology of dependable systems. Survivability is a more recent research field that has emerged as a response to threat scenarios traditionally associated with Internet and network security vulnerabilities (Anderson, Hearn and Hundley, 1997). Dependability and survivability has the same end goal; to build highly dependable software to support critical infrastructural systems. Even with similar goals dependability and survivability approaches the problem from different angles. Within dependability, researchers define a rigorous taxonomy and different types of threats to the system are classified and categorized. Survivability focuses on the characteristics of the system in its environment and tries to establish the properties of the execution environment to the system of interest. Survivability focuses on the mission of the system and the environment in which the system needs to survive. Dependability focuses on the measured performance of the solutions. The two approaches lead to different angles on the problem that complement each other. The dependability approach provides valid good solutions on real problems that can be used within survivability in different survivability specifications.

A problem with the engineering definition of survivability proposed in Knight et al. (2003) is that it requires an enumeration of all possible environmental states that the system may encounter. This results in a twofold problem. First, it is difficult to partition and enumerate all possible environmental states. Second, a partitioning of the environmental states at design time leads to a design based on operational modes, which is appealing since it is a well-studied subject but drifts away from the adaptive dynamic systems approach.

The two main algorithmic strategies to build survivable systems are structurally hierarchical algorithms (Sullivan et al., 1999) and local emergent algorithms (Fisher and Lipson, 1999). Using a hierarchical structure and a hierarchical communication pattern improves the scalability of the system (Sullivan et al., 1999). Local algorithms only communicate with a bounded neighborhood. In an unbounded network environment, it is important that the number of systems to communicate with does not depend on the total number of integrating systems. An example of a local replication algorithm from the field of mobile ad hoc networks is an algorithm that forms associations between systems over longer distances to enable the faster long-range communication (Hemly, Garg and Nahata, 2005). Local does not necessarily mean geographically local, an important fact when building survivable systems.

2.6 Problem definition

The increasing sophistication of attacks against computerized systems together with the vulnerability of networked infrastructure systems makes survivability a top priority (Ellison, Linger, Lipson, Mead and Moore, 2002). In the NBD domain, the high requirement on rapid integration and reuse in hostile and changing environments necessitates the use of a stable and robust framework (Bendz et al., 2003). No simulations investigating how the situation awareness affects data survival in volatile and hostile environments have been found during the course of this project. The aim of the simulations performed in this project is to identify and investigate important factors of the situation awareness, and of the systems aided by the situation awareness.

2.6.1 Aim

The aim is to investigate basic data replication strategies in a replicated database that utilize situational awareness to minimize the number of data replicas used and to minimize the amount of lost data. Lost data means that the data do no longer exist anywhere in the system. The investigation will be carried-out by executing a number of simulations in a custom-made simulator. The investigated environment is volatile and destructive.

2.6.2 Motivation

In the domain of network-based defense, the requirement on rapid integration and reuse in hostile and changing environments necessitates the use of a stable and robust framework (Bendz et al., 2003). System used in NBD must use the available resources efficiently (Wik, 2003). In this thesis, the resources are replicas of data. It is resource demanding to create and maintain replicas of data. Maintaining consistent data replicas requires large amount of data sent over the network and it is difficult to provide real-time guarantees of the data freshness, as desired in NBD. There are proposed algorithms and frameworks for survivable systems, for example Cukier, Ren, Sabnis, Henke, Pistole, Sanders, Bakken, Berman, Karr, and Schantz (1998) and Narasimhan, Kihlstrom, Moser and Melliar-Smith (1999). There is a lack of models and techniques for describing situational information and of investigations of the benefits and pitfalls of having access to situational information.

2.6.3 Objectives

Develop a replication strategy that replicates data in order to save the data, where save means that at least one replica of the data remains in the replicated database. Investigate how the adaptive replication strategies compare to a deploy-time configured replication strategy in different configurations of resource utilization and environments characteristics. Each configuration is investigated through a simulation scenario.

2.6.3.1 Hypothesis

The following overall hypothesis is investigated through the developed strategies.

A replication strategy that uses situation awareness to decide where to replicate data, avoids loosing data with fewer replication targets than a replication strategy that selects replication targets randomly.

2.6.4 Expected results

The expected result is replication management strategies suited for hostile and volatile environments. The experiments investigate the benefits of situation-aware replication strategies. In addition, the simulations investigate how the parameters enumerated in section 3.2 affect the performance of the adaptive situation-aware replication strategies.

3 Method

The base of the investigation is experiments in a simulated environment developed in this project. This investigation compares situation-aware data replication strategies with a baseline strategy. The project uses a random configuration as baseline since it has not been possible to identify a published benchmarking scenario or test data. The random configuration is also chosen to avoid the unconscious introduction of unfavorable static replication configuration in the experimentation set.

3.1 Data survival simulation setup

The thesis investigates the eventual benefits and drawbacks of using an adaptive replication strategy. Its performance is tested against a static random replication configuration in a discrete event simulation.

3.1.1 Data collection and model definition

Based on recommendations in Law and Kelton (1991) the simulation model contains few details.

A *system node* represents a computer on which the distributed database is replicated. The distributed database consists of data storage and a *replication manager*. The system nodes only fail in a fail-halt manner in the simulation. A halted system node and all its data are destroyed. Data items in the simulation represent the stored and replicated data in the distributed database.

A *data item* may represent an object in an object oriented database or a table in a relational database, etc. In the simulation, the data items are binary data updated in the database and replicated to a set of system nodes. A *data item replica* is a copy of a data item stored at a system node. A *replication configuration* describes on which system nodes each data item and its data item replicas are located. The simulation does not make a distinction of the exact form or type of data since that is not a focus of this thesis. A data item survives if there is a replica of the data item on any of the surviving system nodes.

The *adversary* initiates and executes attacks on the system nodes. The adversary selects a set of system nodes to target and attack.

An *attack* is initiated and executed by the adversary in which it targets and destroys one or more system nodes. When the attack is initiated and before it is executed, the attack is referred to as a *threat*. The set of all threatened system nodes is the *current threat situation*. An attack against a system node always results in the destruction of the system node. This assumption limits the number of uncertainties in the simulation; it also represents the worst-case scenario. The assumption may also be interpreted as the assumption that each system node has equal chance of surviving an attack, (% of surviving a attack) x (% of system nodes attacked) = (% surviving system nodes). The assumption of an always-successful attack is a simplification where the result of the equation is used.

Data preservation ratio is the ratio of surviving data items after all attacks, i.e. number of surviving data items divided by the total number of data items. The data preservation ratio is the base unit for comparison between the different simulations. The data preservation ratio is a measurement of the provided service.

The *degree of replication* describes how large percentage of the existing system nodes contains a data item replica of a specific data item.

Time of commitment models which of the adversary or defender has the ability to adapt to the opponents actions. The difference in the time of commitment depends on the capabilities of the actors. For example, when the adversary issues a physical attack to destroy a data-storage a data replication algorithm may be able to decide to which system node to replicate the data to later than the adversary must decide where to aim the physical attack. Attacks' originating from different arenas require different defense measurements to survive.

The *degree of destruction* is the ratio of destroyed system nodes and is a factor with high impact on the replication requirement. For example, under a single failure assumption[1] only one replica is needed. The degree of destruction is in these simulations studied as an input parameter to the scenario.

In the simulations, the *situation awareness accuracy* models the situation understanding as a probability to know an attribute of the environment, for example if a system node is threatened. This is of interest when compared to the resulting data preservation ratio. A difficult design decision is the interpretation of accurate prediction from the situational awareness component.

The replication strategy used by the replication manager is to minimize the effect of the threats by replacing threatened nodes in the replication configuration with unthreatened nodes.

The simulations are performed in a number of trials to limit the risk of a misleading result in the case of unlikely replication configuration that performs exceptionally good or bad. The mean of the trial output values constitutes the simulation result. The simulations investigate the relations between five parameters the replication degree, the time of commitment, the degree of destruction, and the accuracy of the situational awareness of both the replication managers and the adversary. These input parameters and their effect on the resulting data preservation ratio are investigated. A uniform random distribution, (Ross, 1997), is used to generate random numbers in the simulations.

The simulation scenario consists of an adversary and a defending system. The defending system controls the replication managers. The replication managers generate a replication configuration in which the data are replicated to a specific set of system nodes. The adversary selects a set of system nodes to target and destroy. The adversary uses a situation awareness component that provides situational information to decide which system nodes to select. The situation awareness component provides the adversary with the system nodes the data items are replicated to, and the defending system with the set of system nodes targeted by the adversary.

The adversary and defending system are allowed a specific number of actions, where an action consists of selecting system nodes to replicate the data to, or target with the attack. Both systems are constrained further by the number of system nodes they select. The adversary is constrained by the number of nodes it is able to attack and the replication managers by the number of replication targets for a data item.

[1] A singe failure assumption is an assumption that only one system node is halted, or failed, at any time.

The simulations investigate different scenarios. A scenario consists of the number of system nodes, replication degree, destruction degree, accuracy of the situational awareness, and the sequence of actions. Figure 2 shows a schematic scenario.

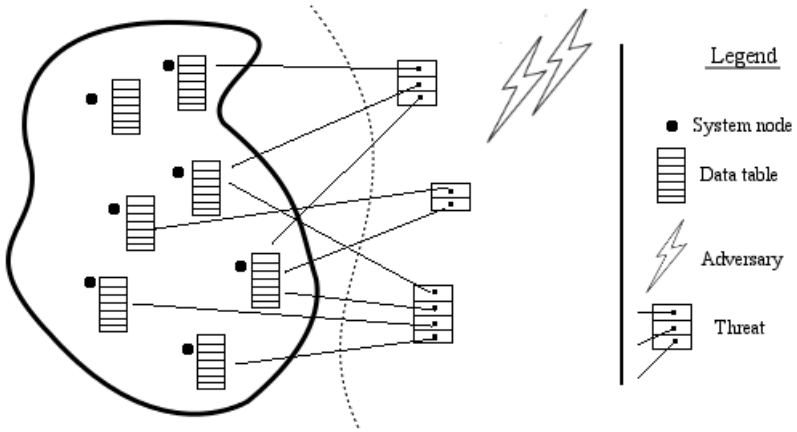

Figure 2. An example scenario.

This scenario model shows an area with seven system nodes with tables of data items and a adversary, represented by the lightning, carrying out three attacks that threaten three, two, and four specific system nodes, respectively. The lines indicate which system nodes are targeted. In the simulation, the defender and the attacker adapts to each other's actions to gain advantage. In addition to being successful, all attacks occur at the same time to reduce the number of scenarios. Since the replication managers are not able to change the replication configuration when the attacks are executed this will not effect the result. Figure 3, Figure 4, and Figure 5 shows examples of the simulation sequence.

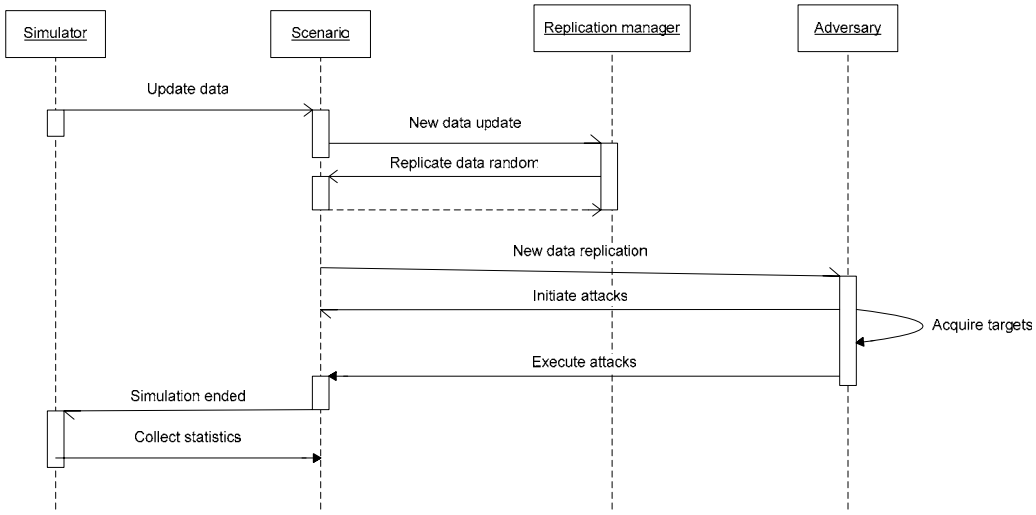

Figure 3. Attack on a randomly configured replication scheme.

The sequence diagram in Figure 3 describes a sequence of actions performed by the replication manager and the adversary. First, the scenario is generated according the parameters. Then data are added to random system nodes. The replication manager on the system nodes then replicates the data items to randomly selected system nodes. The adversary detects the replication configuration. Based on this replication configuration the adversary initiates and executes the attacks. Figure 4 shows the timeline of a scenario where replication manager detects the initiation of the attacks.

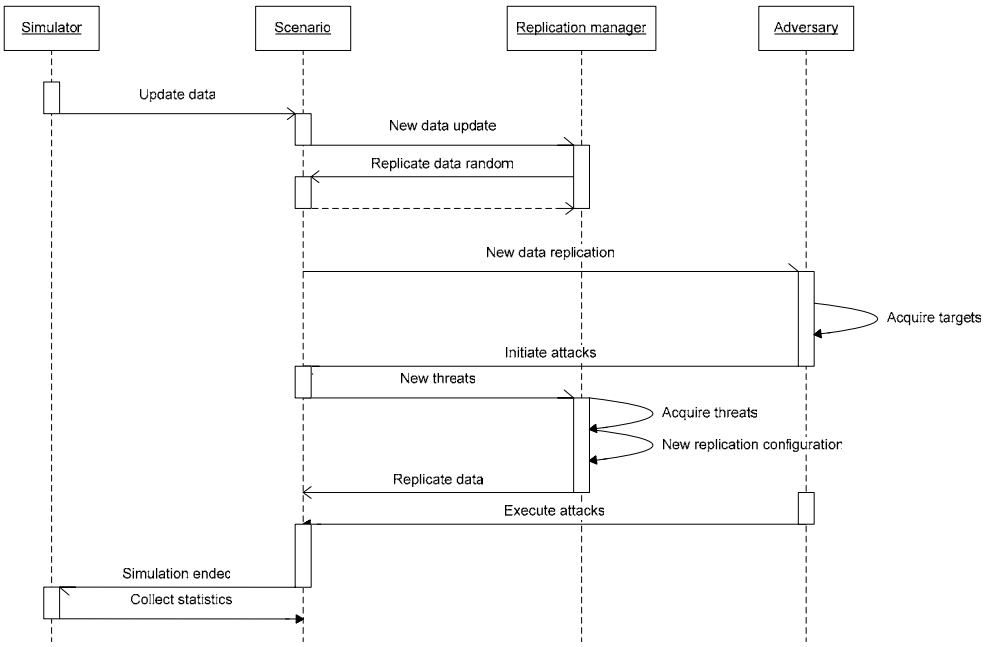

Figure 4. Attack after replication reconfiguration.

Figure 4 shows a scenario where the replication manager detects the initiated attacks and reconfigures the replication to replicate the data to unthreatened system nodes. The attacks are then executed based on the previously intercepted replication configuration. Figure 5 shows the timeline of a scenario where the adversary detects the changed replication reconfiguration has.

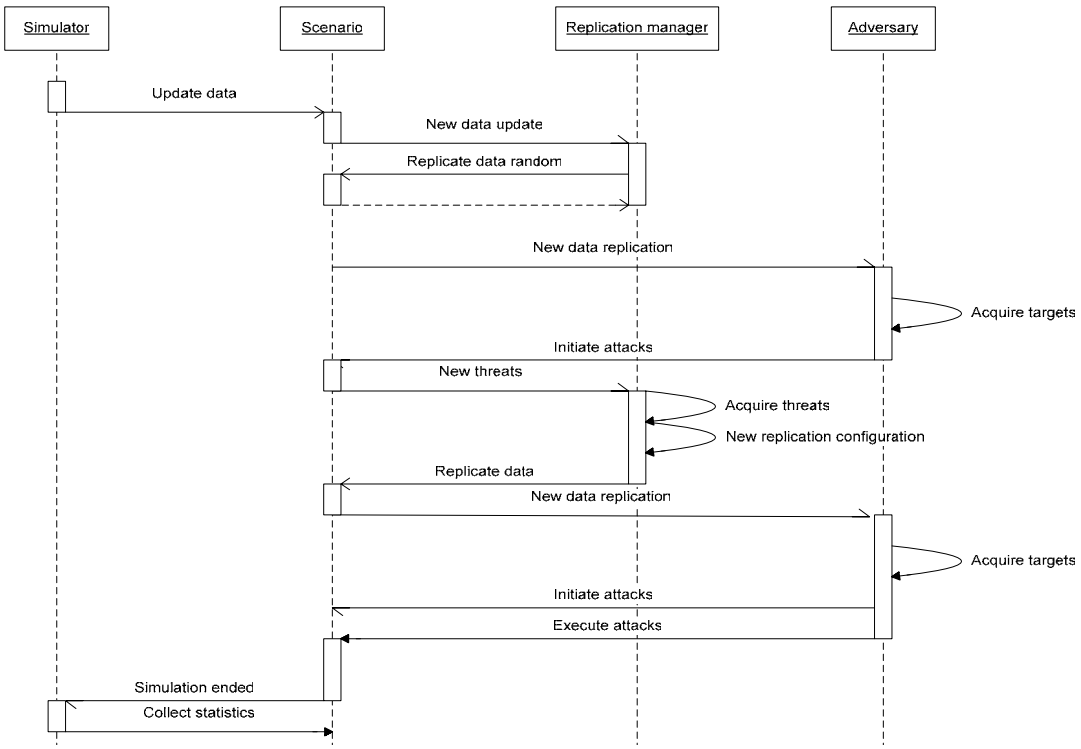

Figure 5. Reinitiating and executing the attack after the replication were reconfigured

Figure 5 shows a scenario where the adversary reinitiates the attack as a reaction to the reconfigured replication before executing the attacks. The sequences shown in Figure 4 may be interpreted as if the replication manager is faster than the adversary and is able to modify the replication configuration before the attack. The sequences

shown in Figure 5 may be interpreted as if the adversary is faster and able to modify the attacks.

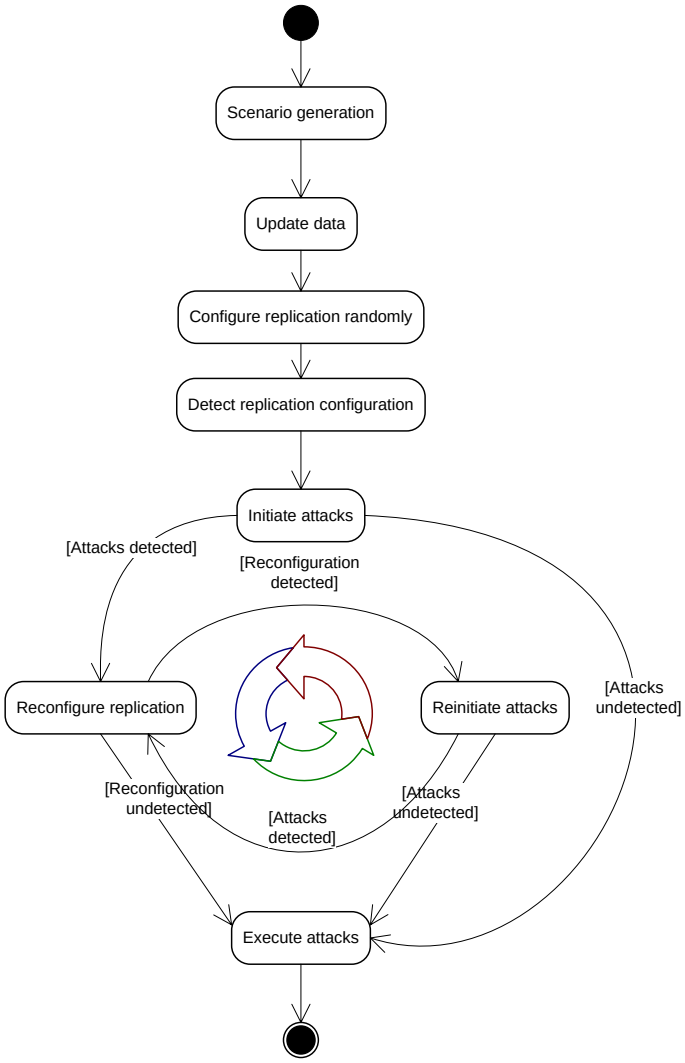

Figure 6. A model of the simulation control flow.

Figure 6 shows the control flow of a simulation in an activity diagram.. It is assumed that if an actor detects an action, the actor has time to adapt to the action. First, a scenario is generated according to the specified simulation parameters. The data items are then updated on random system nodes. All system nodes have approximately equal probability of a data update. The replication manager on each selected system node then replicates the data to random system nodes. The adversary detects the replication configuration and initiates the attacks. The scenario ends if the scenario parameters dictate that the defending system does not detect the attack, as shown in Figure 3. If the defending system detects the attack, the replication managers at the threatened system nodes copy all data items to unthreatened system nodes. This results in the adversary executing the attacks on system nodes based on the previous random replication configuration, as is the case in Figure 4. Figure 5 shows a case when the adversary detects the replication reconfiguration and reinitiates the attacks to target the set of system nodes currently in the replication configuration. The filled arrows in Figure 6 illustrate the potential detection-reconfiguration loop.

3.1.2 Simulator design

The design of the simulator enables experimentation with the parameters to investigate. The simulator supports investigation of the relationships between the parameters and does not include real-time simulations. Figure 7 shows a conceptual model of the simulator where the arrows denote knowledge of the other entity. As an example, the scenario generator knows of the scenario.

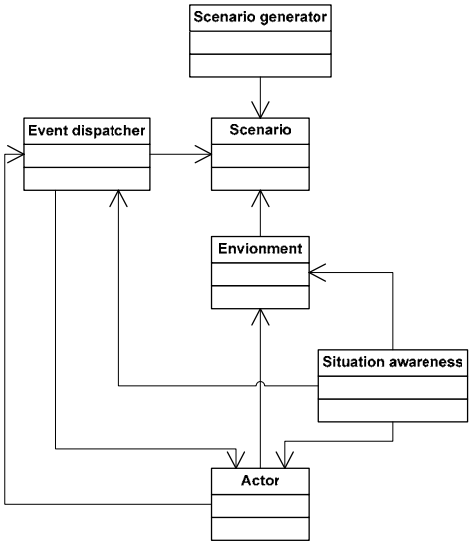

Figure 7. Model describing the relations between the simulation components

A scenario is generated based on the current parameter configuration. A scenario includes information on number of system nodes, replication degree, destruction degree, accuracy of the situational awareness for both actors, and the time of commitment. From this generated scenario, the environment is initialized.

The environment represents the static world in the simulation. The environmental representation contains information not readily available to the system actors if they were real systems in an unbounded environment. This information is included in the simulation environment and provided to both the adversary and the defending system to avoid that system visibility and details of the network configuration affected the simulation results. The simulations could be extended to include a partitioning of the environment where the adversary and the defending system do not necessarily see the same partitions. This would extend what is included in the situational awareness concept in the simulation and shift the simulation focus to what information the situational awareness component is able to acquire, and from the focus in this thesis, how a system benefits from situational awareness.

The situation awareness component provides the actors with information. In the simulations the replication managers and the adversary receives different information. The replication managers receive a list of threatened system nodes. Each threatened system node has a probability of the situation awareness accuracy of being included in the list and known to the replication managers. This corresponds to accurately deducing a threatened system node from the situational data, e.g., intelligence reports, sensor data, etc. The adversary has a probability of its situation awareness accuracy to receive knowledge of all system nodes a data item is replicated to, for all data items. This corresponds to the adversary intercepting the replication table of the data item.

The event dispatcher controls the simulation. The event dispatcher is initialized from the scenario. The situational awareness component uses information from the environment and event dispatcher to analyze the situation. The actor entity is either the adversary or the defending system. The same component provides both systems with information and the systems are shown in the model as one conceptual entity. Note that the adversary and defender may have different probabilities of acquiring accurate information. The events from the event generator affect the actors. If the replication configuration changes, or set of targeted system nodes, information regarding the change propagates through the event generator. In Table 1, there are four event types.

‘Data event’	A ‘data event’ signals that data have changed on a system node. The data are either a data item or a data item replica to be stored at the system node
‘Threat event’	A ‘threat event’ signals that the attacker has changed its set of threatened system nodes.
‘Initiation event’	This event is generated in the beginning of the simulation before any of the actors has acted.
‘Simulation has ended event’	A ‘simulation has ended event’ signals that all the allowed actions are executed. At this point, all attacks are carried out and the statistics of the simulation are collected**.**

Table 1. Event types used in the simulation.

3.1.2.1 The rule set of the adversary

The purpose of the adversary is to destroy all data items. The adversary tries to identify all system nodes to where each data item and its replicas are stored and destroy them. The following event-condition-action (ECA) rules describe the algorithm executed by the adversary. The algorithms is described as ECA rules since they react to external events and the ECA notation is a simple and clear notation to describe event-based algorithms.

ON (‘initiation event’) IF () DO (acquire targets)
ON (‘data event’) IF (replication configuration has changed) DO (acquire new targets)

Table 2. ECA rules describing the actions of the adversary.

When an ‘initiation event’ occurs, no system nodes are threatened. To identify targets the adversary acquires a list of suitable targets from the situation awareness component. When a ‘data event’ occurs, the replication configuration has changed and the adversary needs to identify to which new system nodes each data item are

replicated to and try to target the new system nodes. The adversary uses the situational awareness component to identify these new system nodes.

3.1.2.2 The rule set of the replication manager

The aim of the replication managers is to preserve all data items stored in the distributed database. Each replication manager does this by including unthreatened system nodes in the replication configuration for all data items stored at the system node. The replication manager acquires a list of all threatened system nodes the situational awareness component has detected and avoids replicating data items to those nodes. The following event-condition-action (ECA) rules describe the algorithm executed by the replication managers.

ON (‘initiation event’) IF () DO (create a replication configuration)
ON (‘threat event’) IF (new system nodes are targeted) DO (update the replication configuration)

Table 3. ECA rules describing the actions of the replication manager

When an ‘initiation event’ occurs, the replication manager assigns randomly chosen system nodes as replication targets. To identify which system nodes are targeted, the replication manager acquires a list of all targets from the situation awareness component and avoids these when creating the replication configuration. If there are too few unthreatened nodes to reach the required replication degree, the replication manager will include threatened nodes. When a ‘threat event’ occurs, new system nodes are threatened and system nodes have changed and the replication manager needs to update the replication configuration. The replication manager at each system node checks if the system node it is located on is threatened. If the system node is threatened, the replication manager tries to save all the data stored at the system node. To save each data item, the replication manager finds an unthreatened system node that does not contain a replica of the data item. If such a system node is located the replication manager replicates the data item to that system node, otherwise a random system node that does not contain the data item is selected.

3.1.2.3 Assumption model

During the design of the simulator some simplifying assumptions are made. This section enumerates and explains the assumptions.

1) All operations on the data occur before the simulation begins. This assumption limits the number of cases to investigate and gives the adversary a complete view of the situation.
2) Each actor performs all its actions one at a time. This reduces the number of cases to investigate.
3) Attacks
   a) An attack against a system node by the adversary is always successful.
   b) When an attack is executed, the targeted system node is destroyed and permanently removed. The system nodes are not repaired.

   c) When a system node is destroyed all data stored on that node is lost.
   d) All attacks occur simultaneously.
4) The adversary and defending system are allowed a specific number of actions
   a) The defending system and the adversary have the same view of the world.
5) The systems are provided with the situational information with no regard to its distribution. In a real application the situational information needs to be distributed to the system nodes and may not be available when needed. The simulation is not concerned with these problems.
6) Situation awareness
   a) In the simulations, the situation awareness is simulated as provided by a system or a component. The method by witch the awareness is acquired may include automated systems and human operators. The inner working and methods of this system is not considered in the simulations.

Assumption 1 represents the worst case for the defender since it gives the adversary the opportunity to know all potential targets. Assumption 2 is made to avoid complications from interleaving actions. Assumption 3.a is an assumption that merges two parameters; number of targets and probability of a successful attack, see section 3.1.1. Assumption 3.b limits the opportunity for the defender and represents the worst case. Assumption 3.d simplifies the simulation and does not effect the result because of assumption 3.b and assumption 2 since a system node is not repaired. Even if the attacks occur over time, the same system nodes are destroyed and the result is the same. Assumption 4 is the end condition of the simulation. It represents the ability to adapt to the opponent. Assumption 5 and 5.a removes the effects of situational information distribution from the simulation.

3.2 Simulation goal

The goal of the investigation is to investigate the effect of situational-awareness through simulation. The comparison is made based on the ratio of preserved data items between an adaptive replication configuration strategy and a static replication. Different environmental conditions are investigated.

This investigation only aims to study the effects of these factors. In a real application numerous other problems concerning the ratio of preserved data items, number of replication targets, data consistency, real-time data access requirements, and scalability arises (Ramamritham, 1993).

3.2.1 Simulation scenarios

Four simulation scenarios are investigated. The design of the scenarios aims to investigate different aspects of the relations between the parameters presented in section 3.2 and the hypothesis presented in section 2.6.3.1 . Each scenario is designed to investigate a set of parameter configurations. The scenarios aim to prove the over all hypothesis. First, a general description is given followed by a definition of the experiment setup. Finally, the scenario specific assumptions are presented. The scenarios investigate three sequences of actions. In the first sequence, the replication managers are not notified of the attack and the adversary executes the attacks undisturbed. In the second sequence, the replication manager detects the initiated attacks and changes the replication configuration to protect the data. In the third sequence, the adversary reinitiates the attacks based on the adapted reconfigured replication.

3.2.1.1 Scenario one, situation-aware replication manager

This scenario shows the hypothesis in the simplest case, an unaware adversary. The scenario investigates a situation-aware replication manager compared to an unaware replication manager, i.e. an adversary with no situational-awareness. This is a baseline scenario and shows the basic behavior of the simulator. The adversary selects system nodes at random in this simulation. Figure 4 describes the sequence of actions executed in this scenario. In this scenario, the adversary does not know to which system nodes the data items are replicated. The adversary targets system nodes at random, each system node has equal risk of being selected. Then the replication managers are notified of the new threats, depending on the accuracy of the situational-awareness. A replication manager that is notified of a threat against its system node tries to locate new system nodes to store the data replicas. This leads to a partially changed replication configuration. This scenario is extended with a situation-aware adversary in scenario three. The parameter setup tested in this simulation is

Parameter	Values
Accuracy of adversary situational awareness	0%, i.e. random predictions
Accuracy of replication manager awareness	0%, 25%, 50%, 75%, 100%
Degree of destruction	20%, 50%, 80%, 90%

Table 4. Parameter setup in scenario one.

This scenario assumes that the adversary has no situational awareness. The adversary attacks system nodes without any knowledge of the replication configuration. When the situational awareness component used by the adversary is unable to give accurate predictions there is a random chance that system nodes is targeted and destroyed. When the adversary selects system nodes at random, they are selected independently. This simulation investigates if situational awareness improves data preservation ratio under this independent attack assumption.

3.2.1.2 Scenario two, situation-aware adversary

This scenario investigates the data preservation ratio of a random replication configuration when attacked by a situation-aware adversary. Figure 3 describe the sequence of actions executed in this scenario and illustrates a case where the replication manager does not detect the approaching attack and does not reconfigure the replication. In this scenario, the adversary has a probability of knowing to which system nodes the data items are replicated. The adversary tries to locate and attack the system nodes data items are replicated to. The replication managers are not notified of the new threats and do not change the replication configuration. The parameter setup tested in this simulation is:

Parameter	Values
Accuracy of adversary situational awareness	0%, 25%, 50%, 100%
Accuracy of replication manager awareness	0%
Degree of destruction	20%, 80%

Table 5. Parameter setup in scenario two.

3.2.1.3 Scenario three, situation-aware replication manager vs. situation-aware adversary

This scenario investigates how the level of accuracy of the adversary's situation awareness affects the performance of the replication manager. This scenario is extension of the first scenario with the same sequence but a larger number of simulated parameter configurations. Figure 4 describes the sequence of actions executed in this scenario and illustrates a case where the replication manager detects the attacks upon the randomly configured replication and changes the replication configuration to increase the number of surviving data items.

Parameter	Values
Accuracy of adversary situational awareness	0%, 25%, 50%, 75%, 100%
Accuracy of replication manager awareness	0%, 25%, 50%, 75%, 100%
Degree of destruction	80%

Table 6. Parameter setup in scenario three.

3.2.1.4 Scenario four, situation-aware replication manager vs. faster situation-aware adversary

This scenario investigates the effect on the data preservation ratio when the adversary adapts its attack based on the replication configuration created by a situation-aware replication manager. Figure 5 describe the sequence of actions executed in this scenario. It illustrates a case where the replication managers are notified of the threats and change the replication configuration to increase the number of surviving data items. In this case, the adversary detects the reconfiguration and is able to change the targets to reflect the new replication configuration.

Parameter	Values
Accuracy of adversary situational awareness	0%, 25%, 50%, 75%, 100%
Accuracy of replication manager awareness	0%, 25%, 50%, 75%, 100%
Degree of destruction	80%

Table 7. Parameter setup in scenario four.

4 Results

This chapter presents the results of this investigation. The simulation results are presented with a description of the simulation scenarios. The initial simulations showed the importance of the accuracy of situation awareness of both the adversary and the replication managers. The amount of knowledge provided to the replication manager has a large effect on the simulation results. The focus of the simulation is the investigation of the effects of the situation awareness accuracy of systems.

4.1 Simulation results

The simulations are executed with 100 system nodes and 10 data items. During each simulation, 10 trials are performed; a trial is a simulation execution with a specific set of simulation parameters. The mean value of these trials is presented in the diagrams as the simulation result. The error bars in the diagrams indicate the estimated standard deviation at each observation point. The results from each scenario are presented together with a discussion of the investigated hypotheses.

4.1.1 Scenario one, situation-aware replication manager 3.2.1.1

This scenario investigates the effect of a situation-aware replication manager in environments with different degrees of destruction compared to an unaware replication manager. Section 3.2.1.1 describes this scenario further and Figure 4 shows the sequence of action in the simulation.

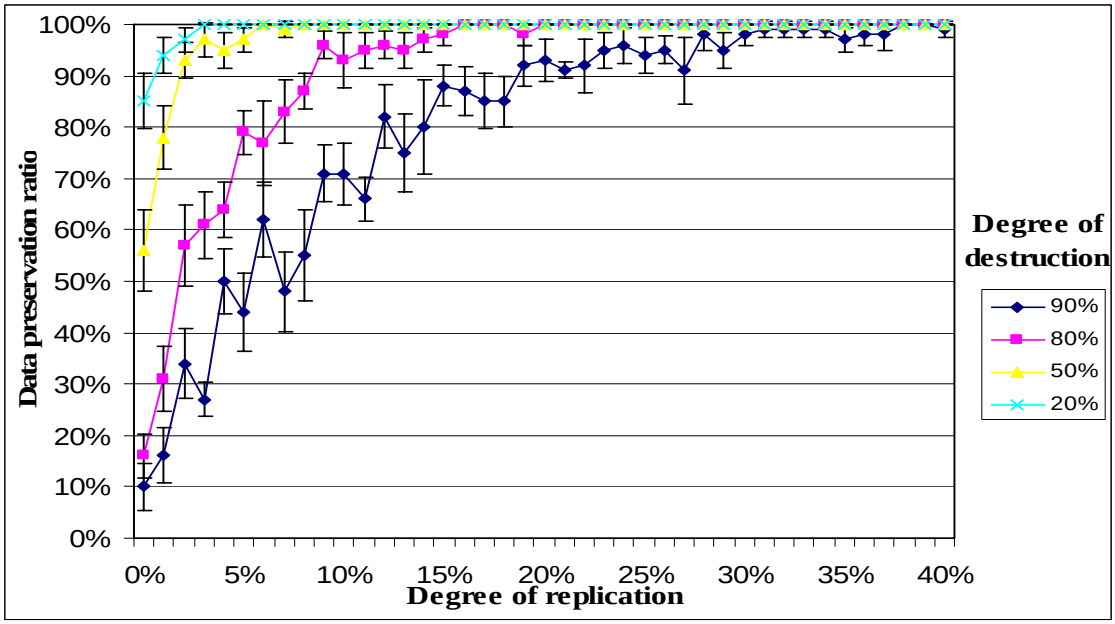

Figure 8. Data preservation ratio of an unaware replication manager and an unaware adversary.

Figure 8 shows the data preservation of unaware, random replication configurations at different levels of destruction. When the degree of destruction is increased, the data preservation ratio is lowered as expected. The result of this simulation is considered a baseline and a verification of the simulation model. Data survival is achieved by a random replication strategy. This indicates that a random replication configuration may be sufficient to protect the data when it is assumed that a majority of the system survives.

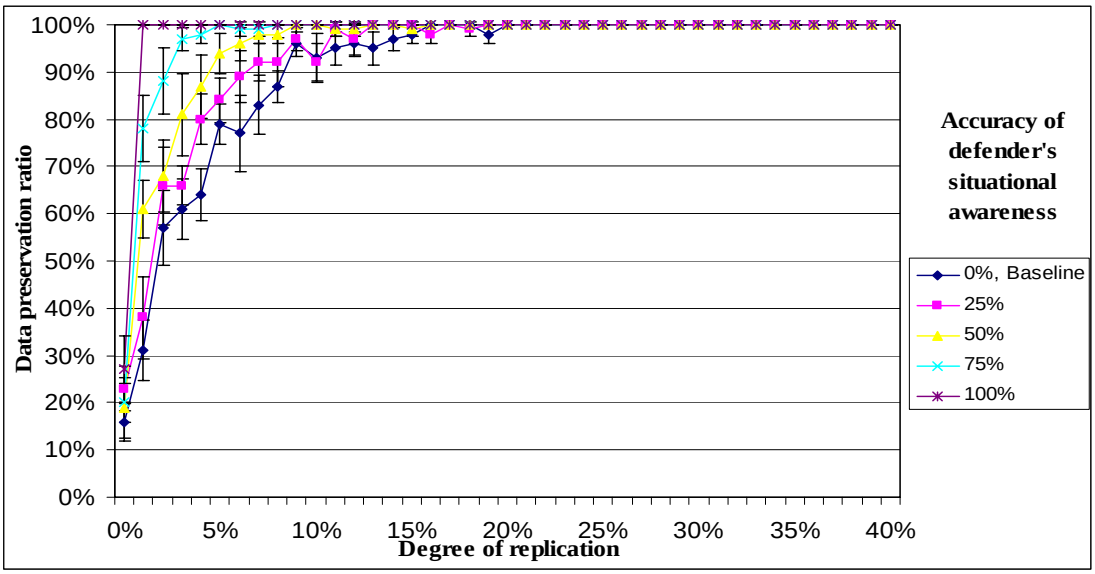

Figure 9. Data preservation ratio at 80% system destruction and an unaware adversary.

Figure 9 shows the performance of a situation-aware replication manager at different levels of accuracies when facing an unaware adversary. It shows that the number of replication targets required to achieve 100% data preservation ratio is reduced as the accuracy increases but at 80% of destruction and random attacks, the difference is not large.

The conclusion drawn from the first scenario is that the importance of using an adaptive replication manager depends on how destructive system environment is to be expected.

4.1.2 Scenario two, situation-aware adversary 3.2.1.2

This scenario investigates the data preservation ratio of a random replication scheme when faced with a situation-aware adversary. Section 3.2.1.2 describes this scenario and Figure 3 shows the sequence of actions in the simulation.

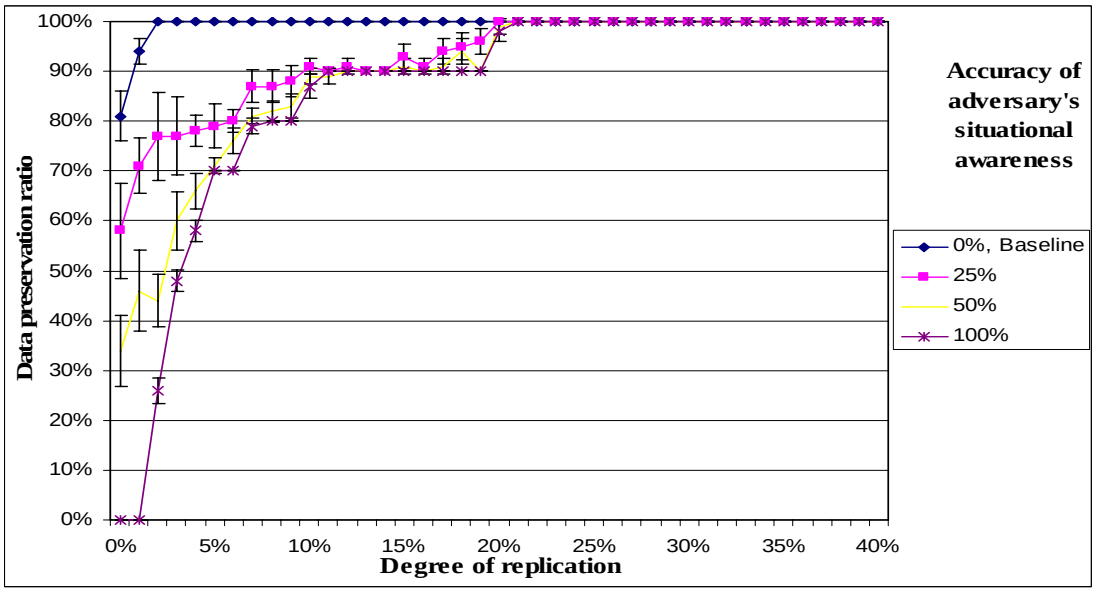

Figure 10. Data preservation ratio at 20% system destruction and an unaware replication manager.

Figure 10 shows a situation where the replication manager does not achieve 100% data preservation by random replication until the replication degree exceeds the degree of destruction. The baseline results show higher data preservation ratio than the simulations where the adversary is situation aware. Figure 10 indicates a reduced importance of the situational awareness accuracy of the adversary as the replication degree increases. Figure 11 shows the same simulation configuration but in a more destructive environment, i.e. an 80% degree of destruction.

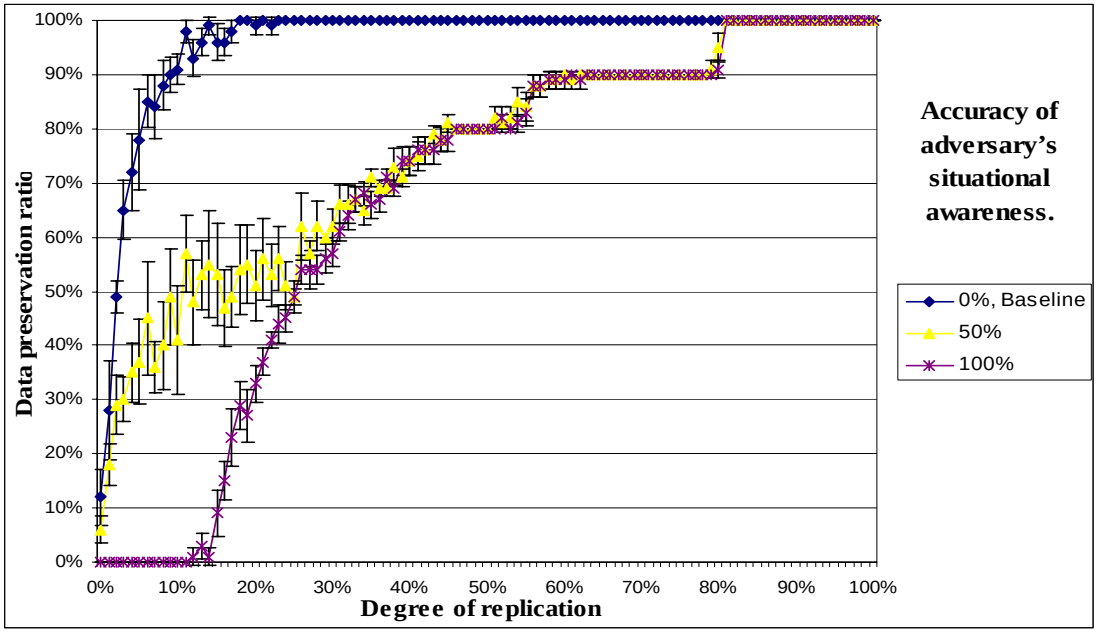

Figure 11. Data preservation ratio of an unaware replication manager at 80% of system destruction.

In Figure 11, the performance of the replication manager is worse than in Figure 10, i.e. a lower data preservation ratio at the equal degree of replication. When the degree of destruction has increased to 80%, the data preservation ratio is lower than 100% even when degree of replication is as large as 70%. The baseline results show higher data preservation ratio than the simulations where the adversary is situation aware. Figure 11 also indicates that the importance of the accuracy of the situational awareness is reduced as the replication degree increases.

From this scenario two conclusions are drawn. First, if the adversary uses knowledge of the defending system to issue attacks, a defending random replication configuration needs to cover more than 80% of the available system nodes to survive. Second, the efficiency in terms of data preservation ratio per destruction degree of the adversary depends to a larger extent on the accuracy of the adversary's situational awareness when the data are replicated to fewer system nodes. However, when the data replication degree approaches the destruction degree, the difference in efficiency between levels of situation awareness accuracy disappears. This may be an effect of the simulation taking place in a bounded environment. The differences may increase in the case of unbounded environments. When the adversary is situation-aware, the defending system needs to be situation-aware to avoid massive replication. As in the previous scenario, the situational awareness may be used at design or configuration time. To be able to configure the replication manager accurately, knowledge of the adversary's capability is needed. A random replication configuration achieves a lower data preservation ratio than the situation aware replication managers in 4.1.1. Figure

10 and Figure 11 indicates that a random replication configuration is required to have larger replication coverage than the assumed destruction degree to provide a data preservation ratio of 100%.

4.1.3 Scenario three, situation-aware replication manager vs. situation-aware adversary

This scenario investigates how the level of respective accuracy of the replication manager's and the adversary's situation awareness affects the data survival ratio. Figure 12 shows the performance of a replication manager with a 25% accurate situational awareness. Section 3.2.1.3 describes this scenario further and Figure 4 shows the sequence of action in the simulation.

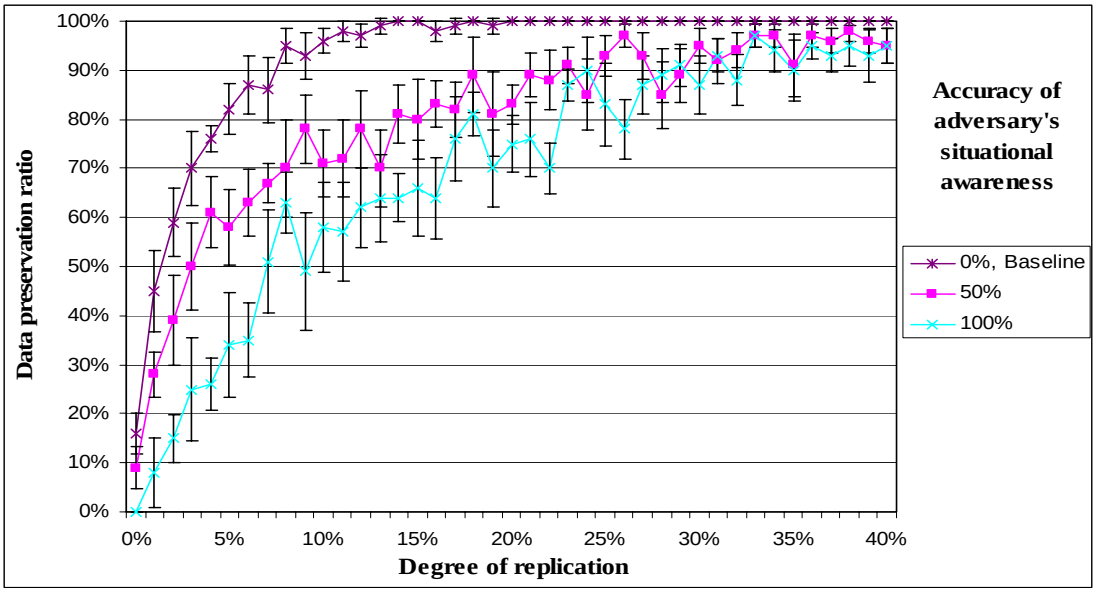

Figure 12. Data preservation ratio when replication manager has 25% situation awareness accuracy, 80% of system nodes are destroyed and a situation aware adversary attacks them.

The data preservation ratio of the in the case of situation-aware replication managers is inversely proportional to the accuracy of the adversary. In other words, an adversary with a more accurate situation awareness is more accurate in its attacks and eliminates more data items even when the adversary uses the same amount of resources (number of destroyed system nodes).

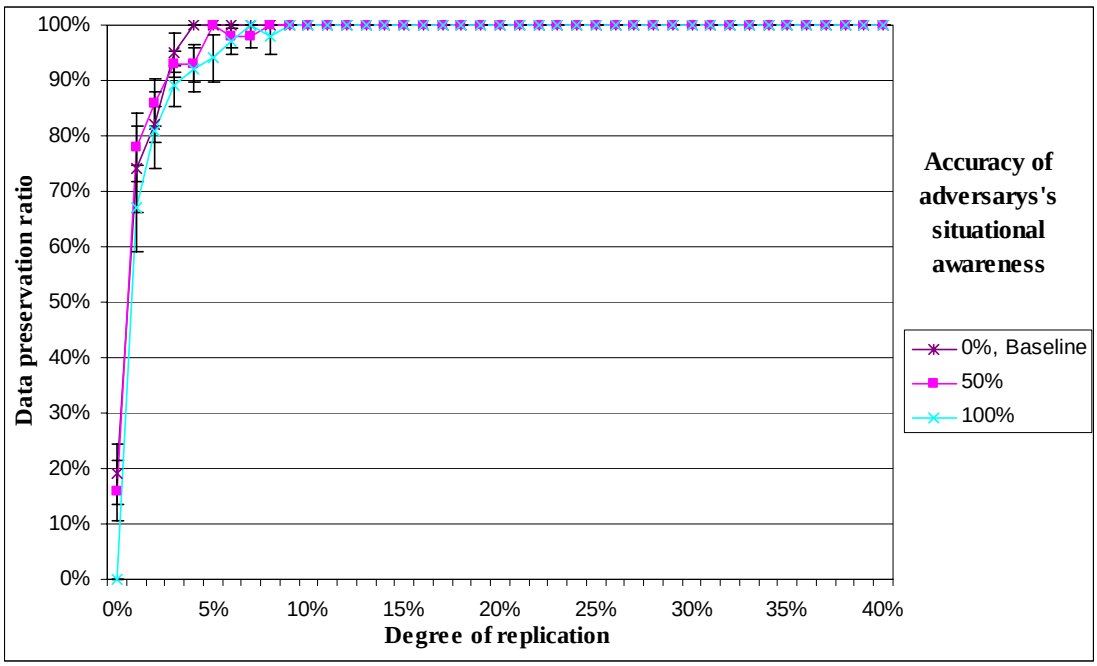

Figure 13. Data preservation ratio when replication manager has 75% situation awareness accuracy, 80% of system nodes are destroyed and a situation aware adversary attacks them.

When the situation awareness accuracy of the replication managers increases to 75%, the data preservation ratio increases greatly. See Figure 12 and Figure 13. However, the relations between the baseline simulation and the simulations with a situational awareness with accuracy of 50% and 100% are similar to the results presented in Figure 12. Figure 14 shows the data preservations ratio when an adversary attacks with a situation awareness accuracy of 75% the replication managers.

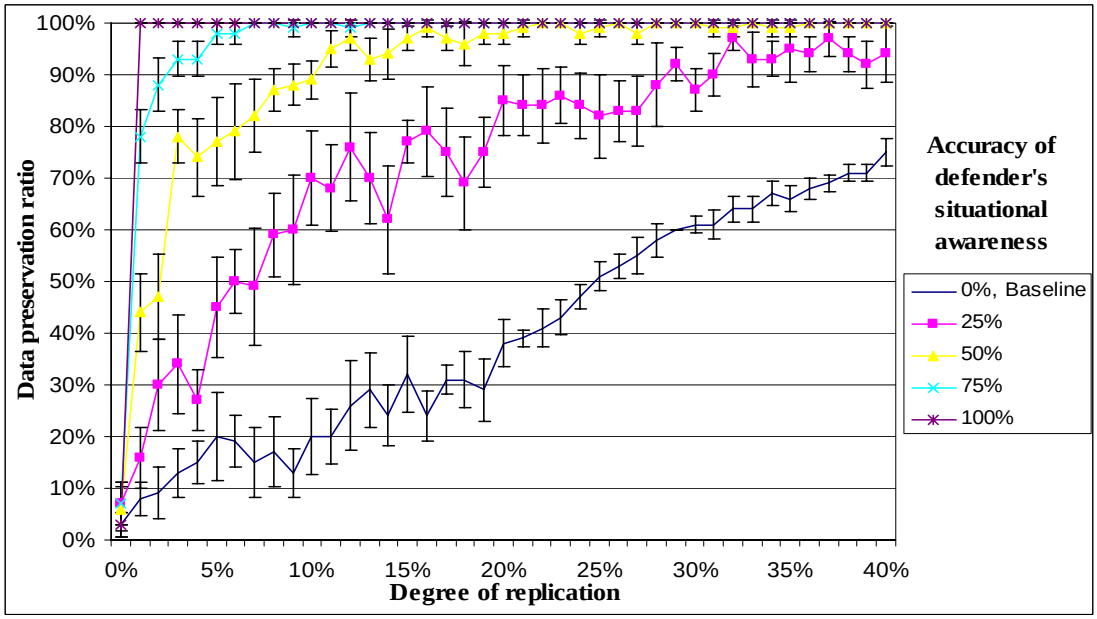

Figure 14. Data preservation ratio when the adversary has a situation awareness accuracy of 75% and destroys 80% of the system nodes.

When comparing Figure 14 with the baseline from Figure 11, unaware replication managers attacked by an unaware adversary, the results indicate a similar data

preservation ratio when the replication managers and the adversary have similar situation awareness accuracy. The rapid reduction of the data preservation ratio in Figure 14 when the situation awareness accuracy of the adversary increases. The defending system needs to achieve at least as good accuracy of the situational awareness as the opponent, or configure the replication to cover a large part of the available system nodes to achieve a high data preservation ratio. The baseline simulation results in not as good data preservation as the simulations where the adversary is situation-aware and the data preservation is increased with each higher level of situational awareness accuracy of the replication manager.

The conclusion from this scenario is that in a highly destructive environment it is not enough being situational aware to avoid massive replication. If the defending system has access to more accurate situational knowledge the replication configuration does not need to cover a large percentage of the system nodes. Large in this context depends on the required data preservation ratio. If it is 100% it still might be sufficient with a degree of replication of 15%. A situation-aware adversary reduces the data preservation ratio and a situational aware replication manager improves the data preservation ratio also when facing a situational aware adversary.

4.1.4 Scenario four, situation-aware replication manager vs. faster situation-aware adversary

This scenario investigates how the data preservation is effected when the adversary adapts its attack based on a replication reconfiguration created by a situation-aware replication manager. Section 3.2.1.4 describes this scenario further and Figure 5 shows the sequence of action in the simulation. The figures in this section are compared to the figures from scenario three. Figure 15 shows the data preservation when a 25% accurate adversary has the ability to act after the defending system.

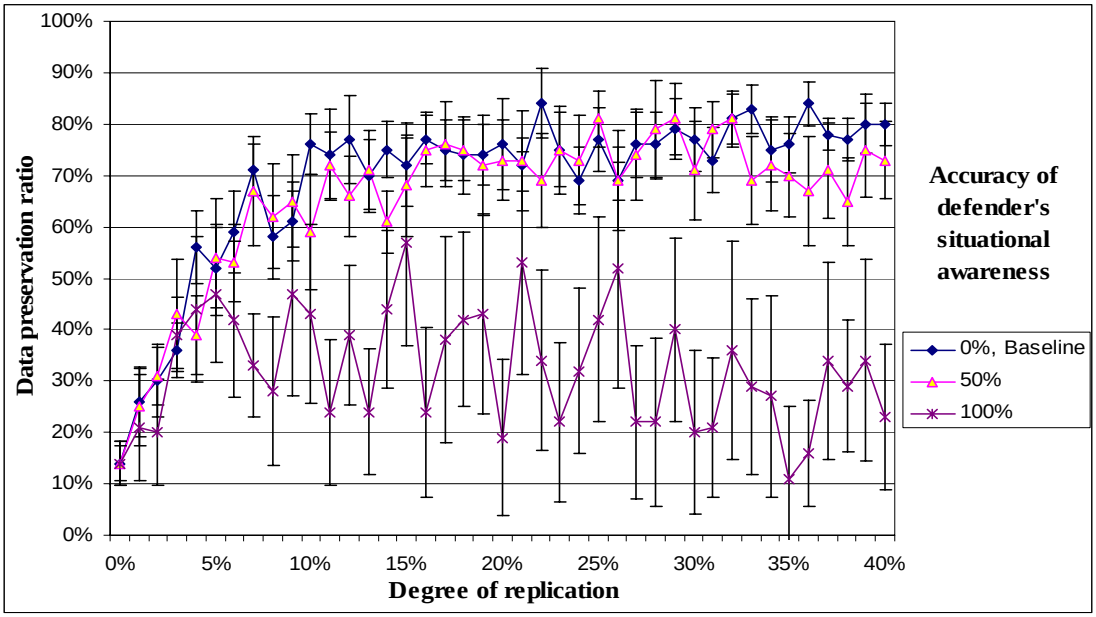

Figure 15. Data preservation ratio when the adversary have a situation awareness accuracy of 25% and destroys 80% of the system nodes.

Figure 15 shows a large variance and that a replication manager with a complete picture of the situation has the lowest preservation ratio. Figure 16 shows the same situation but with an adversary with 75% accurate situational predictions. When

Figure 15 is compared to Figure 12, it is evident that when the adversary adapts to a situational aware replication configuration the data preservation ratio is reduced. An interesting aspect of this simulation is the inverse relationship between the accuracy of the replication managers' situational awareness and the data preservation ratio shown in Figure 15.

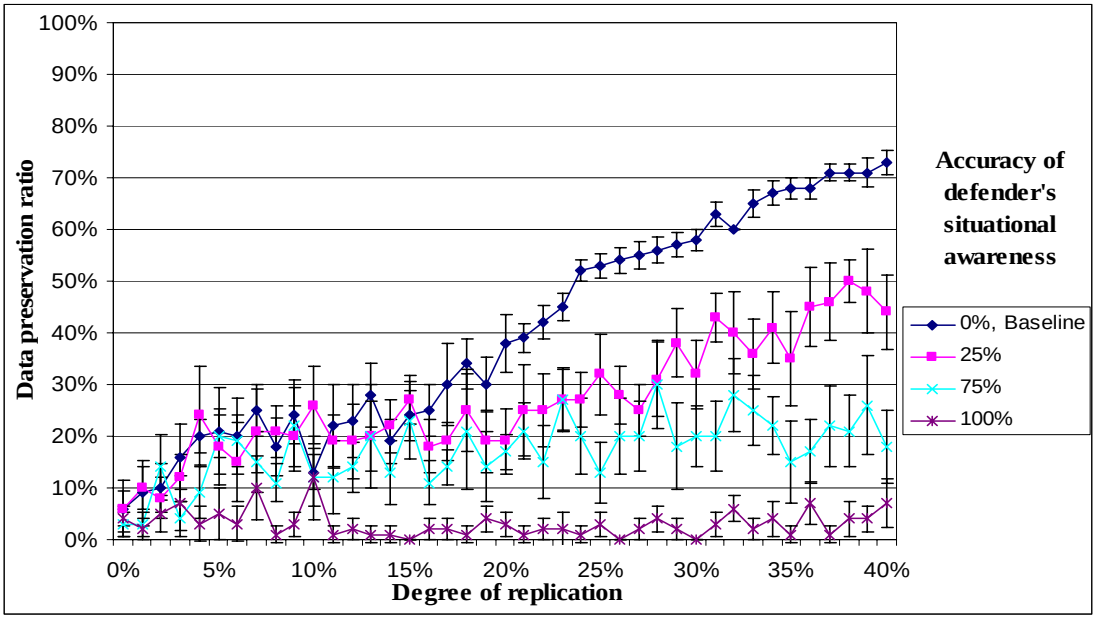

Figure 16. Data preservation rate when the adversary have a situation awareness accuracy of 75% and destroy 80% of the system nodes.

Figure 15 and Figure 16 shows two cases that only differs by the accuracy of the situational predictions provided to the adversary. In these cases, increased situation awareness results in replication configurations that achieve lower preservation ratio. The conclusion from this scenario is that there is a clear distinction between situation awareness and intelligence, to know what happens and to be able to use the information to ones benefit. The result is an effect of the replication managers storing all data item replicas on the system nodes they perceive as unthreatened, which leads to a concentration of data item replicas on fewer system nodes. When the replication managers know which system nodes are threatened, they replicate all data item replicas to the unthreatened system nodes. After this, the adversary can target those fewer system nodes and destroys all data items. A replication manager with a lower accuracy of the situational awareness creates replication configurations with a higher degree of randomly included system nodes. This results in a replication configuration that distributes the data items amongst a larger part of the system nodes and more data survive. The data preservation ratio is reduced when the adversary adapts to a non-random replication configuration.

4.1.5 Summary of simulation

The simulations have highlighted a number of characteristics of the situational awareness that are important for the system performance.

The first scenario investigates the effect of the replication managers' situational awareness when facing random but predictable threats. This scenario shows that a system that adapts to the situation results in a higher data preservation ratio and that data survival is achieved at a lower replication degree. The difference increases with

the destruction degree. The most evident implication is that even if the system is unable to adjust to the environmental changes, a situational awareness is required to decide the degree of replication coverage. The system builders must decide if the environment is stable enough to allow manual reconfiguration of the system when the environment changes.

The second scenario shows that when the adversary is situation-aware, the replication managers must replicate the data items to at least the number of destroyed system nodes in order to achieve a data preservation ratio of 100%. The scenario also shows that the replication manager achieves lower data preservation when facing a situation-aware adversary. This scenario shows that the replication managers need to know how destructive and situation aware the adversary is when they create a replication configuration that meets the data preservation requirement.

The third scenario shows that it is not the accuracy of the situation predictions that is important but the relation between the accuracy of the replication managers and the adversary. The scenario shows that a situation-aware adversary reduces the data preservation ratio and a situational aware replication manager improves the data preservation ratio.

The fourth scenario identifies cases where an increased situational awareness reduces the performance of the defending system. This scenario investigates how the data preservation is effected when the adversary modifies the ongoing attack based on the replication configuration created by the situation-aware replication managers. The scenario shows that the data preservation ratio is reduced when the adversary adapts to a non-random replication configuration. The scenario also shows that if the replication managers only include unthreatened system nodes in the replication configuration, the adversary may herd the data to a fewer number of system nodes and is able to destroy all data items and their data replicas. A system needs to use its knowledge wisely.

4.2 Discussion

4.2.1 Environmental characteristics

The simulations indicate that a number of environmental characteristics are important when designing a system configuration, for example the destruction degree and the accuracy of the adversary's situational awareness. The replication managers need to know the degree of destruction to decide the replication degree to use. These characteristics may be provided by a situational awareness component to allow the system to be adapted and optimized for the current situation as predicted. The information can be provided at design or deploy time if the environment is static.

4.2.2 Situation-awareness

The situation-awareness of the opponent and his ability to use the situation-awareness is also an important factor when deciding the number of replication targets. The simulations indicate that the ability to detect the opponent's actions and hide your own intentions is an important characteristic to know when designing and reconfiguring the system. To use a detailed situational awareness might be devastating if nothing is done to countermeasure any attempt to herd the system into a vulnerable state. Situational awareness and knowledge must be used with caution and intelligence; the situation may be misleading and the opponent deceiving. The simulations performed do give indicates some important factors and characteristics of

both the situation awareness, how the information is structured and suitable interfaces to provide the information through, and the investigated system, destruction degree and the accuracy of the adversary's situational awareness . These factors and characteristics need to be investigated further in realistic simulation environments. Situational awareness, when provided as a product to a system may be of different level of detail or abstraction. On low abstraction level, detailed knowledge of the situation guides algorithms and optimize their performance. On a high abstraction level the situational awareness might guide operational mode changes, switch to a faster random replication configuration to act faster then the opponent or replicate the data to a safe area.

4.2.3 The hypothesis

The hypothesis this thesis aims to prove as formulated in section 2.6.3.1 is "A replication strategy that uses situation awareness to decide where to replicate data, avoids loosing data with fewer replication targets than a replication strategy that selects replication targets randomly."

The first scenario shows that the effect of situational awareness is more evident when the environment is more destructive. When the system resources are limited, it is important to use the available resources in an optimal way in the current situation. In the second scenario, the adversary uses information of the replication scheme to issue the attacks. To defend the data against intelligent orchestrated attacks the defender needs a high degree of replication or use knowledge of the situation when creating the replication configuration. When the defender uses accurate situation awareness to create a replication configuration it needs fewer replication targets. The third scenario shows that the relation between the situation awareness of the adversary and the defender is important. It is important to know how the environment changes. The first three scenarios are all in agreement with the hypothesis.

The fourth scenario contradicts the hypothesis. When the attacker can predict the defenders reactions, increased situation awareness of the defender reduces the defender's data preservation ratio. Under these circumstances, the defender will suffer a lower data preservation ratio when it uses accurate situation awareness.

5 Related work

Survivability and dependability have many common or similar concepts. This relates work within both dependable and survivable systems to this thesis. Mobile, ad hock, sensor networks have similar requirements on adaptively and locality as survivability.

5.1.1 Control theoretic approach to Survivability

The background to this section is given in section 2.3.2. The control system architecture needs to be decentralized to handle the huge size of the system. Since both the controlled system and the controlling system face changes, the control system needs to be adaptive. The system to control is assumed large, and a hierarchical system structure allows for an efficient information flow in the system. The control need is different on different system levels.

On leaf nodes, for example, local information must be readily available to allow fast tactical decisions. The local information is then sent up in the hierarchy where it is aggregated to a compound view of the current state, allowing for strategic decisions. This global state together with strategic control is fed back downwards in the hierarchy to allow for adaptation of the system. Knight, Sullivan, Elder & Wang (2000) describe the architecture further. It focuses on non-local faults and distinguishes between local faults and non-local faults. Knight et al. (2000) assume that local faults are dealt with by a local fault handling or masking mechanism. A non-local fault affects multiple nodes in the network and is often difficult to mask. An unmasked fault means that the system cannot continue in normal operation.

Survivability implemented as a control system adds requirements both to the controlled, and to the controlling system. The controlled application needs to be reconfigurable and to provide a design diversity that allows for reconfiguration. This poses a challenge when legacy applications are a part of the system. Legacy applications are often essential to the infrastructure and do not provide the design flexibility to allow runtime reconfiguration (Knight et al., 2000). The control system itself is a likely target when an adversary launches an attack. An attack towards the survivability control system can have severe effects on the system as a whole since it may allow both changing the information on which the decisions are based but also to control the system reconfiguration. The demands on the reconfiguration mechanism are high. It is likely that there are a high number of fault types and the actions taken as a response to each type of fault may differ. This results in a large number of possible system transitions. The control system must complete a reconfiguration within a bounded time to support systems with real-time needs. Since a reconfiguration is a system transformation started by some specific event in the environment it is important that no vulnerabilities are introduced, neither during nor after the reconfiguration. The control theoretic solutions are based on general system properties and no simulation where the information used is varied has been found.

5.1.2 AQuA

AQuA (Adaptive Quality of Service Availability), Cukier et al. (1998), is a middleware that allows distributed applications to specify a desired level of availability. AQuA is a process-oriented middleware that builds upon quality objects (QuO) and offers a CORBA interface that allows clients to specify a desired level of availability. AQuA is a middleware that focuses on how a specified level of services is achieved without any regard to what level is appropriate or how it is best achieved.

The goal of AQuA and the replications strategies developed are similar; to provide a desired level of service, but the developed strategies strive to take advantage of the system context to provide the service. In the simulation tests of the middleware, a fixed scenario is used and the middleware is configured for the scenario.

5.1.3 Immune system

The Immune system (Narasimhan et al., 1999) is a framework to provide survivability for existing CORBA applications by intercepting the communication between the CORBA objects and subject it to a deterministic voting scheme. To support critical survivable systems both the client and server objects are actively replicated and both input voting and output voting is employed (Narasimhan et al., 1999). The immune system is, similar to AQuA, a process-based middleware to provide a desired level of service. The immune system approaches a larger range of survivability related problems. The immune system also focuses on providing the service without regard to the current situation or to the system context.

6 Conclusions

Survivability is an important concept that together with dependability and quality of service are key issues in the systems of the future such as infrastructural systems, business applications, and everyday desktop applications. Its importance and its wide spread usage together with the complexity of those systems make middleware and frameworks for survivability imperative to system builders. A difficult task when building a survivable system is to avoid including unnecessary and harsh requirements. An example is a requirement of a never halting system. This is a requirement that in practice is impossible to guarantee without including unrealistic assumptions, never failing power for example, or using huge amount of resources trying to comply with the requirement.

The thesis presents an ontology for describing critical systems and outlines the environment and threats these systems have to encounter. The simulations show that to use the resources efficiently, situational awareness is required. Depending on the capabilities of the system and the availability of situation awareness, a system developer may use situational awareness at design and configuration, and build a system that uses the situation awareness at runtime. The situation information consists of different levels of abstractions, where the different levels provide input to decisions at different levels, from algorithmic optimizations to design diversity. It is important to use the information wisely since a purely reactive usage of the knowledge makes the system predicable and vulnerable.

6.1 Contributions

This thesis contributes with an analysis of the intersection of dependable systems, survivable systems, network based defense, and situation awareness and highlights common areas in which the different fields enriches each other. The thesis emphasizes the usability of situation-aware systems and the need of systems that is adapted at runtime to optimize the resource usage. Two interesting aspects of situational awareness that are presented are how the awareness should be used and when situation awareness might be damaging to the system that uses it.

6.2 Future work

6.2.1 Simulation

The simulations are discrete event-based and do not capture the real circumstance of two combating systems. A real-time simulation with timed OODA-loops (Wik, 2003) and system reconfiguration times that investigates the effects of the interleaving of actions in detail would be beneficial.

6.2.2 Survivability

The field of survivability needs to reach a common ground. During the course of this project, the need for a benchmarking scenario has been the most pressing. There is a need for a number of realistic scenarios describing as many facets of reality as possible in various domains. The description should at least cover the malicious activities, system requirements, and system-usage patterns and cover physical activity and pure computer system threats. Benchmarking scenarios should be developed for a number of small, well contained, domains. When these scenarios have been implemented benchmarking scenarios that cover larger domains and multiple domains

should be developed. These scenarios provide a good foundation to base survivability frameworks on and make it possible to compare and differentiate the effects of different solutions for survivability. Based on the scenarios, different framework prototypes should be developed and evaluated to test the different proposed architectural styles – hierarchical, emergent, etc. – in different domains and circumstances. During the construction of these prototypes, it is important to base the work on the usage requirements and to show where the prototypes fail, not where they succeed.

6.2.3 Situation awareness

To be able to utilize situational awareness, approach of the design and construction of information, control, and command systems may differ. To be able to construct situation-adapted systems it must be established that the increased development cost can be motivated. Constructing systems that take advantage of situational knowledge is difficult and to verify and validate those systems poses a huge challenge. To simplify construction and evaluation of these systems a simulation-based approach is suggested. A simulation of the system environment together with the situation awareness component should be performed to enable testing and experimentation with the adaptive system. To simulate situation awareness accurately an investigation on classification of the knowledge provided by a system for situation awareness is proposed. This classification is also important to build systems that use the situation awareness with modular design. In the classification, a granularity concept should be present. It needs to be investigated how the accuracy of the predictions and the level of detail of the predictions affect each other and how they affect the performance of the system utilizing the awareness to optimize performance.

6.2.4 System adaptability

Survivability can be seen as a quality of service among others. To study survivability from this perspective would set the problem in a richer context that may both provide new solutions and show earlier solutions infeasible. This also enriches the study of quality of service and provides a concrete domain of study.

6.2.5 System integration

The unbounded network domain foreseen as the domain of survivable systems is a domain in which system integration is the predominant way of constructing systems. This leads to the conclusion that survivable system integration poses a significant challenge. To integrate systems, a method to model the assumptions built into the systems would simplify verification and validating, and make it possible to make correct design- or assembly-time decisions.

References

Anderson R. H., Hearn A. C. & Hundley R. O. (1997). RAND studies of cyberspace security issues and the concept of U.S. minimum essential information infrastructure. Proceedings of the 1997 Information survivability workshop. IEEE Computer Society, San Diego, Calif. 12-13 February 1997.

Avižienis A., Laprie J., Randell B. & Landwehr C. E. (2004). Basic Concepts and Taxonomy of Dependable and Secure Computing. IEEE Transactions on Dependable Secure Computing, (1), 11-33.

Bendz, J. H. & Johannisson, P. & Jönsson, P. & Öhlund, G., (2003), The Swedish Armed Forces Enterprise Architecture, Stockholm: Swedish Defense Materiel Administration (FMV).

Cukier M., Ren J., Sabnis C., Henke D., Pistole J., Sanders W. H., Bakken D. E., Berman M. E., Karr D. A. & Schantz R. E. (1998). AQuA: An Adaptive Architecture that Provides Dependable Distributed Objects, Proc. of 17th IEEE Symposium on Reliable Distributed Systems, 1998.

Ellison R., Fisher D., Linger R., Lipson H., Longstaff T. & Mead, N. (1997). Survivable Network Systems: An Emerging Discipline (Technical Report CMU/SEI-97-TR-013). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University.

Ellison R., Linger R., Longstaff T. (1999).An approach to survivable systems. NATO IST Symposium on Protecting Information Systems in the 21st Century. Washington, DC, October 25-27, 1999.

Fisher D. A. & Lipson H. F. (1999). Emergent Algorithms – A new method for enhancing survivability in unbounded systems. Proceedings of the 23ed Hawaii international conference on System sciences (1999).

Ellison R., Linger R., Lipson H., Mead N. & Moore A. (2002).Foundations for survivable systems engineering,. *The Journal of Defense Software Engineering*, pp. 10.15, July 2002.

Hemly A., Garg S. & Nahata N. (2005). CARD: A Contract-based architecture for resource discovery in wireless ad hock networks. Mobile Networks and applications 10(1-2), 99-113, February 2005.

Knight, J. C., Strunk, E. A. & Sullivan, K. J. (2003). Towards a Rigorous Definition of Information System Survivability. DARPA Information Survivability Conference and Exposition, 1, 78-90.

Knight, J.C. & Sullivan K.J. (2000). On the Definition of Survivability, Technical Report CS-TR-33-00. University of Virginia, Department of Computer Science

Knight J. C., Sullivan K. J., Elder M. C. & Wang C. (2000).Survivability Architectures: Issues and Approaches. DARPA Information Survivability Conference and Exposition (DISCEX 2000), Hilton Head SC (January 2000)

Laprie J. C. (ed.) (1992). Dependability: Basic Concepts and Terminology. Springer-Verlag, 1992.

Law A. M. & Kelton D. W. (1991). Simulation Modeling and Analysis, 2nd edition. McGraw-Hill, Inc., New York.

Narasimhan P., Kihlstrom K. P., Moser L. E. & Melliar-Smith P. M. (1999) Providing support for survivable CORBA applications with the immune system. *in Proceedings of the 19th IEEE International Conference on Distributed Computing Systems*, (Austin, TX), pp. 507--516, May 1999.

Ramamritham K. (1993). Real-Time Databases, International Journal of Distributed and Parallel Databases, 1(2), 199-226.

Ross S. M. (1997). *Introduction to probability models, 6th ed*. Academic Press.

Sullivan K. J., Knight J. C., Du X. & Geist S. (1999). Information Survivability Control Systems. International Conference on Software Engineering, 184-192.

Tolk A. & Pullen M. J. (2003). Ideas for a Common Framework for Military M\&S and C3I Systems (2003), Euro Simulation Interoperability Workshop Stockholm, Sweden, June 2003.

Warston H. & Persson H., (2004), Ground surveillance and fusion of ground target sensor data in a network based defense, The 7th International Conference on Information Fusion June 28 to July 1, 2004 in Stockholm, Sweden, 1195-1201.

Wik M., W. (2003). What is Network-Based Defence (NBD) and the Impact on the Future Defence? KKrVAHT nr 5 2003, Royal Swedish Academy of War Sciences, Defence Materiel Administration, FMV, Stockholm.