Chapter 7
Anti-fragile Cloud-Based Telecom Systems

While Netflix has demonstrated how to apply the five design and operational principles to develop and maintain anti-fragile software applications in the cloud, it is less clear whether the cloud facilitates the creation of anti-fragile telecom systems, because nobody has built such a system. Since we need to understand what makes a system fragile to downtime before we can make it anti-fragile, this chapter initially studies the properties of Norwegian telecom infrastructures resulting in fragility to unplanned downtime.

We first introduce three general concepts causing fragility to downtime. Next, we use the concepts to describe examples of Norwegian telecom systems' past fragility to downtime. Then, we create toy models to determine indicators of fragility to future downtime at different levels of the systems. While the models cannot predict extreme global behaviors leading to downtime in the systems, they have enough explanatory power to clarify existing vulnerabilities.

Armed with an understanding of properties that make telecom infrastructures fragile to downtime, we consider how to build and maintain anti-fragile telecom systems with much of the functionality, but not all, implemented in the cloud. We first discuss how the four design principles of modularity, weak links, redundancy, and diversity make telecom infrastructures more robust to downtime, before discussing how the fail fast operational principle makes the infrastructures anti-fragile to downtime.

As the complexity of a system increases, unintended dependencies occur and new levels and patterns emerge, changing the global behavior in unpredictable ways [2]. In particular, new dependencies can create positive feedback loops, making extreme global behavior more likely. Hence, whenever possible, we should remove fragilizing dependencies between modules rather than add new structures and functionality to combat fragility to downtime. The chapter's last part pays special attention to the removal of strong dependencies.

Since a large effort, far beyond the scope of this book, is needed to ensure highly anti-fragile infrastructures to all types of negative impacts, this chapter only points the way toward anti-fragility to downtime. The question of how to use cloud computing in

K.J. Hole, *Anti-fragile ICT Systems*, Simula SpringerBriefs on Computing 1,
DOI 10.1007/978-3-319-30070-2_7

telecom systems is an active research area [63, 64]. Because there is no commercially available cloud-based telecom solution, the current chapter is more speculative than Chap. 5 analyzing Netflix's media streaming solution.

7.1 Anti-principles Causing Fragility to Downtime

To guide our search for fragility, we exploit so-called *anti-principles* describing how not to design systems. While the existence of black swans make it impossible to precisely quantify a complex system's degree of anti-fragility to a class of impacts, we can easily detect when a system is fragile to a particular class using anti-principles [7, 10, 34]. The following three anti-principles outline how to create fragility to downtime. The author discusses these anti-principles in an earlier paper [7]. The current versions are slightly modified to emphasize downtime in telecom infrastructures.

Uniqueness A system is unique when its key services are not provided by another system. A unique infrastructure with strong dependencies between modules and little redundancy and diversity—the extreme case being a traditional monoculture [28, 29]—is particularly fragile to downtime because local failures spread easily and its many users cannot switch to an alternative infrastructure during an outage.

Connectedness A system is connected when its normal operation depends on the normal operation of another system. If an infrastructure is connected to another infrastructure, then the large overall complexity of the infrastructures causes fragility to downtime.

Closed A system is closed when stakeholders do not share technical and legal information. If only a small group of experts have deep knowledge of an infrastructure, they have a tendency to develop similar mental models for how the infrastructure works during discussions. This propensity toward groupthink is especially strong when most group members belong to the same organizational culture. A uniform group cut off from external expertise with different perspectives overlooks possible rare events causing downtime.

Similar to the design and operational principles discussed in Chap. 4, the reader may recognize some of the anti-principles as anti-patterns described in the literature on software design. Here, we use the term *anti-principle* rather than *anti-pattern* to emphasize that these general concepts are also valid outside the area of software design.

7.2 Past Fragility to Downtime

We apply the anti-principles to study past downtime incidents in Norwegian telecom systems. The major stakeholders did not predict the coincidences that resulted in the outages. Instead, the outages were analyzed by the stakeholders after the fact.

Since knowing the impact of an incident influences how the incident is assessed, it is necessary to be careful when stating the results of a study. Humans have a tendency to concoct explanations for events after they have occurred, making them seem less surprising and more predictable than they really were. This hindsight bias misleads stakeholders into simplifying the causes of an accident, highlighting a single element as the cause and potentially overlooking multiple contributing factors [6, 9, 17, 18, 36].

First, we consider the anti-principle of *uniqueness*. Mobile phone networks are essentially unique infrastructures because the users of one network cannot generally connect to another network when their network experiences problems. If an entire network goes down, millions of mobile phones become useless as communication devices. The largest mobile phone network in Norway went down for about 11 h on June 10, 2011 [65, 66]. The restart of a central node with upgraded software initiated a signal storm that exceeded the network's signalling capacity. The outage affected nearly 3 million customers, or approximately 60 % of all Norwegians. According to top management, the incident was not supposed to happen because earlier restarts of the same node had not caused any problems. The management's surprise and the rare, highly negative outcome of a common operation qualify the signal storm as a black swan incident, at least to top management.

On June 17, 2011, parts of the same mobile phone network went down again due to a new signal storm [67, 68]. According to the Norwegian Post and Telecommunications Authority (NPTA), both signal storms were caused by insufficiently understood dependencies between central nodes in the network combined with insufficient capacity to handle the increasing signal traffic from the many new smartphones. The difficulty in pinpointing the causes of the extreme behavior became evident when, more than a month later, careful technical analyses of the events finally revealed that the signal storms were primarily due to a programming error and not insufficient signal capacity to serve new smartphones [69]. However, the network owner also discussed the need to change the system design and to increase the signal traffic capacity.

The network owner's difficulty in determining the causes of the downtime illustrates that the mobile phone network is indeed a complex adaptive system prone to surprising global behavior. The NPTA publicly stated that the owner needed to improve the network's risk management. While better risk management can assess and mitigate more incidents, perhaps providing longer periods of stable network operation, large-impact incidents will still occur because the network has too many dynamic interactions for humans to reliably foresee rare and extreme behavior. The fundamental problem is not bad risk management but that the four design principles, especially the principles of modularity and weak links, were not fully adhered to when the system was created. Hence, a major outage affecting many customers was bound to happen sooner or later due to the system's uniqueness.

Second, we consider the anti-principle of *connectedness*. All mobile phone networks in Norway are connected to the national power grid. The normal operation of each network depends on a nearly continuous supply of electricity. In late December 2011, the networks went down in a large area of Norway when a storm with hurricane-force winds damaged many power lines, leaving more than 700 base stations belonging to the different networks without electrical power [70]. While the base stations had backup batteries, most lasted a maximum of only four hours. Because landline phone and fixed Internet access were also disrupted in the same area, most people were without communication capabilities following the storm. Although both the power and phone companies worked hard to repair the extensive damage, it took more than a week to restore services to all customers. The problem in 2011 was that the telecom systems were too dependent on the power grid, that is, the telecom system and the power grid were strongly connected systems.

Third, we consider the anti-principle of being *closed*. The telecom networks' strong dependence on the nearly continuous delivery of electrical power came as a surprise to leading Norwegian politicians. Their initial response was to severely criticize the mobile phone companies. According to the NPTA's director general, the people of Norway had come to depend more on the mobile phone systems than the agency had realized before the storm. The fact that both leading politicians and the NPTA were surprised indicates that the consequences of the telecom systems' strong dependencies on the power grid were not fully understood. It is reasonable to suspect that this surprise was due to insufficient information sharing between the network owners and the NPTA, obscuring the fragility to downtime.

The Norwegian Directorate for Civil Protection reports that Norway's largest network owner does not provide major stakeholders, including the Directorate itself, with enough information about changes in the telecom infrastructure [71]. It seems that the owner makes major changes to its infrastructure without informing important stakeholders such as the Norwegian Public Roads Administration, the National Air Navigation service, or the police. Therefore, it has been hard for these institutions to determine the level of exposure they face by using the telecom infrastructure. More publicly available information is needed to discuss and understand the real dangers associated with the use of the telecom infrastructures.

The discussed incidents show that the Norwegian telecom systems were fragile to downtime in 2011 due to the anti-principles of *uniqueness*, *connectedness*, and, most likely, being *closed*.

7.3 Indicators of Fragility to Future Downtime

We now turn our attention to properties of telecom infrastructures that indicate fragility to future outages. First, consider the building blocks of the single generic telecom infrastructure in Fig. 7.1 [72]. While the model is quite coarse, it is adequate for our purpose. The model contains one *transport network* and multiple *access networks* [71]. The transport network is the backbone of the telecom infrastructure and moves data over long distances. The access networks give users access to the infrastructure. Some of the access networks consist of one or a few base stations, serving wireless terminals in the vicinity, while others consist of local broadband networks, connecting, for example, homes and offices.

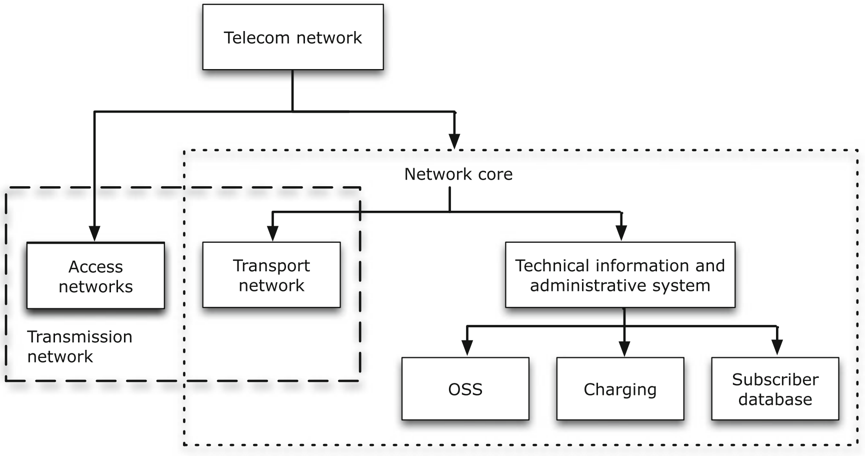

Fig. 7.1 Hierarchy of networks and systems in a generic telecom infrastructure

Entities owning and operating a transport network and/or access networks are called operators. These operators each have a technical information and administrative system consisting of three smaller systems. Together with the transport network, the three systems constitute the network core. The operations support system (OSS) in Fig. 7.1 configures and provisions the core network nodes. Factors impacting the configuration are the number of subscribers, peak hour call rates, the nature of the services, and geographical preferences. The OSS system also collects network statistics, monitors alarms, and logs various actions of network nodes. The subscriber database contains information on all customers and the charging system calculates the costs chargeable to the customers. Because the subsystems in the network core are needed to set up and take down all user communications, we arrive at the following conclusion:

- A unique technical information and administrative system in the network core indicates fragility to future downtime.

From Sect. 2.7, a system $\mathcal{X}$ depends on a system $\mathcal{Y}$ if a failure in $\mathcal{Y}$ negatively affects the functionality of $\mathcal{X}$. The details of the dependencies between different telecom infrastructures are generally unknown to analysts without close ties to operators, but the main dependencies between access and transport operators are usually known. Figure 7.2 depicts publicly known dependencies between Norwegian operators early in 2014 [72]. The arrows show the direction of these high-level dependencies.

Two systems are *interdependent* when each is dependent on the other. There are interdependencies between the three transport operators in Fig. 7.2. Dependent and interdependent infrastructures allow cascading failures to pass infrastructure boundaries [6]. Figure 7.2 illustrates that it is important to prevent a failure in a single infrastructure from spreading to other infrastructures. Furthermore, since most of the dependency paths in the figure end up in the transport network of the largest operator

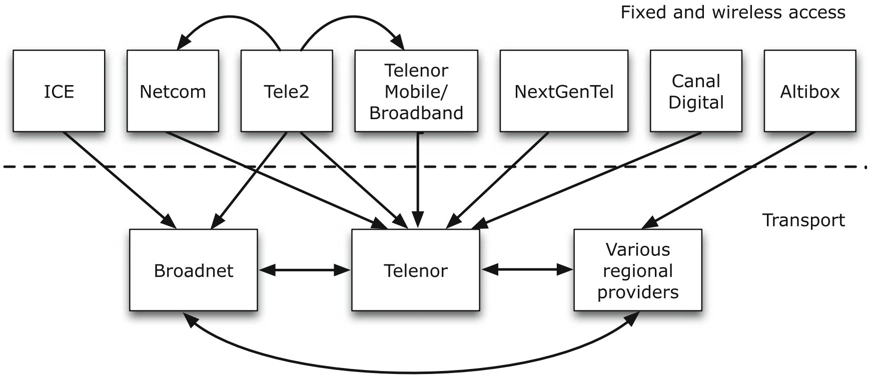
Fixed and wireless access
ICE
Netcom
Tele2
Telenor Mobile/ Broadband
NextGenTel
Canal Digital
Altibox
Transport
Broadnet
Telenor
Various regional providers

Fig. 7.2 Access and transport operators in Norway. The *arrows* show the dependencies between the operators. There are interdependencies between the transport operators

(Telenor), fragility to downtime in this transport network is especially serious. If the transport network has an outage, then the outage will spread to nearly all other infrastructures. We therefore note the following:

- A unique transport network connecting multiple operators signals fragility to downtime.

To illustrate the fragility of transport networks, we consider the transport network of Norway's largest operator. While the network's full topology is not publicly known, it includes "self-healing" rings. Figure 7.3 shows a particularly simple ring with two transmission paths between network nodes. If there is a break in one line, the other line may still be available, providing the second is not in close proximity to the first and also damaged. For best protection against failure, different physical routes are used for the two lines. All data are transmitted on the working or active line, while the protection line is on standby. When the active line fails, the two network nodes affected immediately switch to the protection line.

Even a self-healing ring fails, however. On May 23, 2011, both transmission lines of a ring in the largest operator's transmission network failed [73]. Due to roadwork, one line was temporarily moved aboveground by installing a temporary cable. An excavator broke this cable by accident. About seven minutes later, a falling tree cut the other line somewhere else in the country, causing an outage affecting mobile phone customers in large parts of Norway for about three and a half hours. Air traffic

Fig. 7.3 Conceptual self-healing ring

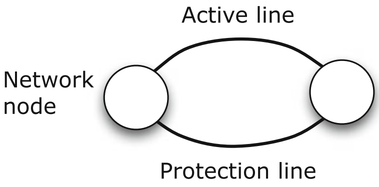

was also affected because a regional flight control center lost its phone connections. The incident shows that having only two cables transporting the major part of the telecom traffic between different geographical parts of a country is unsafe, especially when both cables are aboveground.

7.4 Robust Access Networks

Using the above understanding of fragility to downtime in telecom infrastructures, we now consider how to make future infrastructures robust to downtime. The expected seamless integration of Wi-Fi and mobile network technologies and the emerging Internet of Things will lead to a massive increase in the number of mobile and stationary devices connecting wirelessly to telecom networks. Many believe that machine-to-machine communications supporting smart grids, smart homes and cities, and electronic health will be particularly important. All the new devices and new data-hungry services will lead to a huge increase in wireless data traffic. Examples of devices are notebooks, mobile phones, tablets, televisions, kitchen appliances, smartwatches, 3D glasses, drones, robots, sensors, and actuators, while examples of services involve high-definition video available anywhere, continuous real-time interactions between individuals, and medical sensors monitoring people's health.

To fulfill future communication needs in a power- and frequency-efficient manner, the deployment of multiple layers of radio coverage is most likely necessary where traditional macrocell towers provide a blanket of coverage while, under the blanket, thousands of small cells provide high data rates in areas such as malls, airports, arenas, public plazas, urban parks, and business districts [63]. Because most devices will be close to base stations, it is possible to provide high data rates while keeping the signal power low. The use of small low-power cells enables the increased reuse of frequencies across cells. Today's national telecom infrastructure with many access networks already has thousands of expensive base stations. In fact, the base stations constitute a large percentage of the total cost of current telecom networks. One attractive possibility to limit the costs of even more base stations is to move much of the stations' functionality to the cloud.

High-speed links between the base stations and the cloud are needed to satisfy the stringent delay requirements enabling radio signal processing in the cloud. Data from multiple base stations can be used to alleviate the increased multi-cell interference due to reduced cell size by dynamically adjusting the radio signaling according to channel conditions. To further limit processing delays, use of a highly distributed cloud architecture with local access network clouds is possible, as depicted in Fig. 7.4, where each access network cloud consists of a cluster of commodity and special-purpose hardware. Hence, the access networks will become much more intelligent than they are today.

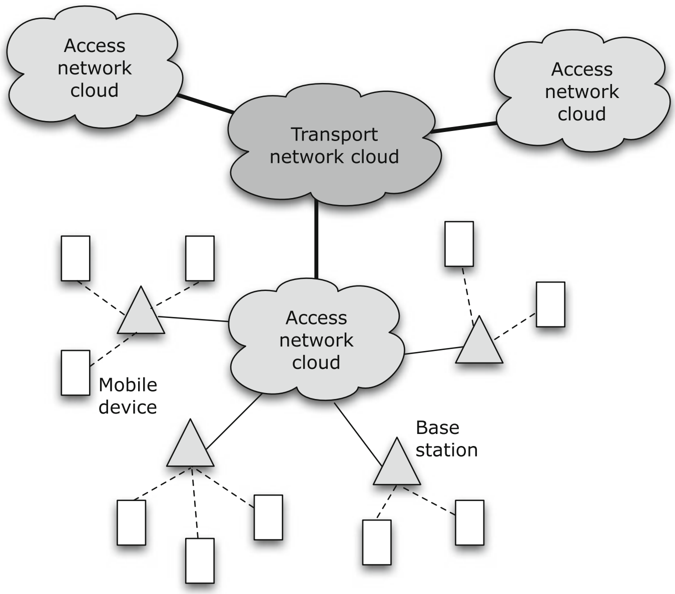

Fig. 7.4 Possible architecture of cloud-based telecom infrastructure

Flexibility in how much processing is done in the access network clouds is needed to support varying types of services and link delays. The upper layers of the base stations' radio protocol stack can most easily be moved to the cloud, such as admission/congestion control and radio resource management algorithms. If low-delay links are available, then lower-layer functions, for example, parts of the physical and medium access control layers, can also be moved to the cloud [63, 64].

By moving most of the base stations' functionality from the cell sites to access network clouds, a provider no longer needs to build enough processing capacity into every base station to handle peak traffic conditions [63, 64]. Instead, the provider can allocate processing resources to the parts of the access networks where they are most needed at any given time. For example, in the evening hours, the processing resources can be adjusted according to shifting service demands as phone users move from a city's business district and to its suburbs.

When there are low-delay links between the base stations and the clouds, each base station is reduced to a radio and an antenna array. While today's access networks are costly and time-consuming to upgrade, future generations of access networks would simply be software upgrades in the cloud. Using the four design principles of modules, weak links, redundancy, and diversity together with the implementation ideas introduced by Netflix, it is possible to develop cloud-based implementations of base station functionality that are robust to downtime. Cloud implementations will isolate local failures to break any positive feedback loops, as discussed in Chap. 5.

7.5 Robust Network Core

There are at least three ways to make a network core (see Fig. 7.1) robust to unplanned downtime. First, telecom operators should realize the functionality of nodes in the network core in a transport network cloud, as shown in Fig. 7.4. Again, the design principles of modules, weak links, redundancy, and diversity make it possible to isolate local failures. Furthermore, cloud implementation allows operators to upgrade software without taking down and restarting nodes. This advantage is important, since it takes a long time to restart central nodes in today's systems. The cloud also increases the programmability and controllability of the network core because good development and monitoring tools exist for the cloud. Operators may use a private cloud to control its hardware layer. To reduce costs, several operators could share a community cloud infrastructure.

Second, to reduce the chance of a major outage, operators should improve the redundancy of the transmission paths in their transport networks. A risk analysis [74] from 2012 of the largest transport network in Norway recommends additional redundant paths to avoid incidents similar to the outage of May 23, 2011, described at the end of Sect. 7.3.

Third, to make it harder for local failures to spread, operators could deploy equipment from different vendors to increase the hardware and software diversity of the nodes in the network core [72]. However, the advantage of added vendor diversity must be weighed against the extra resources needed to operate and maintain a diverse system. In particular, the use of equipment from different vendors could lead to compatibility issues.

7.6 Reduced Dependency on the Power Grid

The December storm discussed in the second half of Sect. 7.2 revealed the Norwegian telecom infrastructures' strong dependence on the national power grid [70]. Strong winds damaged many power lines, causing more than 700 base stations to go down after their backup batteries were quickly depleted. To reduce the dependency on a nearly continuous supply of electricity, the NPTA required network operators to improve their backup power solutions to ensure that the 1,000 base stations covering the most critical areas of Norway have backup power for at least 72 h [70]. All other base stations were required to have at least six hours of backup power. In addition, the operators were told to prepare more resources and develop better contingency plans to enable local crews to quickly repair damaged power lines and base stations.

7.7 Reduced Dependency on One Infrastructure

In general, a unique complex system should not implement a service of critical national importance when the impact of a black or gray swan is intolerable [7]. The impact of extreme global behavior is reduced by realizing a critically important service using two different systems. It is vital that a failure in one system does not cause a failure in the other. It is not sufficient to deploy two identical systems because they obviously have common vulnerabilities; the systems must have diverse designs or implementations. Of course, simultaneously targeted attacks can still bring down both systems, but diversity is likely to make such attacks costly and difficult to successfully carry out.

After the 2011 outages discussed earlier, customers wondered why they were not switched to another network when their home network went down. The simple answer is that none of the networks had the capacity to service a huge number of additional users. The network infrastructure needed large and expensive changes to facilitate such a switch. However, it is economically viable to give a limited group of people with important responsibilities during a crisis access to several operators' networks, either by giving them phones with multiple subscriber identity module (SIM) cards or by adding functionality to switch the group members between networks.

To further reduce the impact of outages in commercial telecom networks, emergency services in Norway have their own telecom network, called the Norwegian Public Safety Network. This network covers all populated areas of Norway. While current commercial networks are based on the same technology (LTE), the emergency network is based on another technology (TETRA), especially developed for emergency communication. None of its roughly 2,000 base stations have less than eight hours of backup power and 15 % of the stations have 48 h worth (http://dinkom.no/en). The emergency network only supports low-speed data communication.

7.8 Anti-fragility to Downtime

While implementations of the four design principles make cloud-based telecom infrastructures robust to downtime, the operational principle must be implemented to make the systems anti-fragile to downtime, that is, robustness must be maintained over time by learning from small incidents.

The stakeholders of telecom infrastructures may balk at the idea of deliberately introducing failures to quickly detect vulnerabilities. Granted, it may be a bad idea to induce failures in today's infrastructures. However, if a telecom infrastructure is designed and implemented in the cloud according to the four design principles, then it should be possible to induce local failures without creating a significant danger of systemic failure causing prolonged downtime. Cloud implementation then makes it possible to quickly discover vulnerabilities in administrative systems, in nodes in the network core, and in the base station functionality.

When infrastructures are able to confine the impact of local failures, operators can constantly adjust their systems to keep them within the bounds of normal operations. The adjustments involve tinkering with selected parts and processes of the systems. Not all tinkering will have the desired effect, because it is hard to foresee the consequences of changes to complex systems, especially large telecom infrastructures. To react quickly to unusual behavior, it is necessary to monitor an infrastructure. It is not enough to just monitor each part of the infrastructure. Because a system's complexity first and foremost stems from the many interactions between its parts, a global view of the system's behavior is necessary. A system must be monitored at all times, especially when it experiences problems.

Because failures will occur in complex systems no matter how many resources are used on high-quality risk analysis, reactive measurements are needed to limit the impact of surprising incidents. In practice, there is a trade-off between proactive and reactive measures to reduce downtime. A risk analysis [74] of Norwegian telecom networks using Internet protocol version 4 (IPv4) suggests that sometimes it is better to improve operations, maintenance, and the ability to quickly react to problems than to make specific parts of the incumbent network infrastructure more robust to downtime. These actions will benefit all Norwegian providers (see Fig. 7.2), while physical changes will mostly provide local or regional benefits.

7.9 Discussion and Summary

Although we do not have any general method to measure an information and communications technology (ICT) system's degree of anti-fragility to downtime, it is possible to determine when a system is fragile to outages. Here, we applied three anti-principles to determine fragilities to downtime. These anti-principles were selected because they proved useful during the investigation. While we have only applied the anti-principles to telecom systems in Norway, it is not hard to apply them to other types of ICT systems. Experience with anti-principles indicates that many systems are fragile to downtime. Many more anti-principles exist (http://sourcemaking.com). Additional work is needed to determine other anti-principles that reveal fragility in ICT systems.

We should move as much functionality of a telecom infrastructure as possible to the cloud and apply the five design and operational principles to create anti-fragility to downtime. A service-oriented architecture (SOA) with microservices is most likely a good way to achieve anti-fragility in practice. There is no need to use public cloud infrastructures. A better solution is to use private, specialized clouds with, perhaps, custom hardware for signal processing, in addition to commodity hardware. The important point is that the amount of custom hardware can be reduced significantly compared to today's telecom systems.

Due the diverse expertise and huge amount of work required, it is outside the scope of this book to determine and analyze all aspects of telecom systems leading to fragility to downtime. In particular, we mention the need to study the protocols of telecom systems. Future systems are likely to use IPv6, which has known vulnerabilities, including fragility to denial-of-service attacks denying regular customers access to telecom services. The creation of anti-fragile protocols is an interesting research topic.

What to learn from Part II

Part II has studied the fragility, robustness, and anti-fragility of Netflix's media streaming solution, the Norwegian e-government system Altinn, and Norway's telecom infrastructure. The case studies provide strong evidence that careful application of the five design and operational principles introduced in Chap. 4 can provide anti-fragility to downtime, that is, the principles lead to systems with less downtime than today's strongly connected, highly optimized systems with little ability to handle unforeseen events. Furthermore, the cloud facilitates the realization of these principles, although the principles are also believed to be valid for non-cloud systems. While the five principles are easy to understand at an abstract level, the case studies demonstrate that the challenge is to determine how to implement the principles in real systems.

Anti-fragile software solutions in the cloud should be based on SOA with microservices, preferably implemented and operated by development and operations (DevOps) teams with "skin in the game." SOA and microservices together model a software solution as a set of independently deployable and scalable services with well-defined interfaces. This architecture style supports the development and management of services by multiple, largely independent teams using different programming languages, continuous deployment, and highly redundant and scalable data storage. The use of microservices with limited functionality makes it possible to ensure the graceful degradation of an application's functionality. Each service's limited functionality facilitates the development of automated fallback responses in the case of local failures. When a local failure affects a service, other services depending on this malfunctioning service receive a standardized response.