

A Metabolic Approach to Protocol Resilience

Christian Tschudin
University of Basel
CS Departement

and

Lidia Yamamoto
Hitachi Europe
Sophia Antipolis Laboratory

Technical Report CS-2004-003

University of Basel

July 2004

Abstract

The goal of this research is to apply a chemical execution model to the implementation of communication protocols. One guiding question will be the creation of robust execution circuits which can distribute over a network and which continue their service despite parts of the implementation being knocked out. Like packets that can be lost (which can be recovered by the appropriate protocols) we envisage an environment where parts of a protocol's execution can be lost. The remaining implementation elements should continue to operate and be able to recover by themselves for restoring full services again. Based on the "Fraglets" approach, we will investigate the dynamic decomposition and recombination of protocol logic inspired by metabolic pathways in cells.

Keywords: resilient communication software, autonomic communication, bio-inspired networking, sensor networks, active networking.



1 Introduction

Autonomic Communication [1] is a long-term research initiative aimed at the study of the self-organization of network elements, toward their autonomous behavior and automated evolvability. Autonomic networks must be self-managing, which includes self-monitoring and self-healing, among other self-* properties. Several areas are concerned, including security, trust, stability, resilience, control, programmability, behavior composition, and context awareness.

Two important and complementary goals of autonomic communication are resilience and self-healing capacity: resilience against internal failures and misbehavior, and self-healing ability to recover from such abnormal conditions.

Resilience and self-healing ability are essential properties of a self-organizing network, where functional and coherent protocol structures must emerge out of basic protocol submodules. The system must be able to detect and replace misbehaving software at run time, while continuing to provide the service, although perhaps less efficiently during the transitory repair phase.

In most commonly encountered current computer systems there is always a risk of full service disruption due to buggy or malicious code. The underlying software systems are usually not robust to misbehaving code, and are unable to autonomously resume their normal behavior after such misbehaving code has been installed. This fragility stems from the implicit assumption that all code should be well-behaving, predictable and correct. This assumption is unrealistic, as it can be observed daily in the form of disruptive software bugs, viruses, worms, and attacks of various kinds.

The current methodology for the design of protocols is not better. It relies on the strong assumption of full reliability of the executing components. The fragility of this approach is easy to demonstrate: remove the processor, remove the software module, remove a procedure, or even a single instruction. All these will with high probability lead to the protocol failing to provide the intended communication service.

We concentrate on the last aspect, i.e. the impact of removing a single instruction, or small code fragment, and formulate a first goal: a protocol implementation should be robust enough such that perturbing any single instruction does not lead to wrong protocol behavior. Note that the definition of a single instruction or code fragment depends on the instruction set used, and will be explained in Section 4.

The second goal is to be able to detect and correct wrong protocol behavior, provided that the amount of error is below a threshold. Such self-healing ability is essential to maintain the first goal (resilience) during a potentially long period of operation. Otherwise errors could accumulate, leading sooner or later to service disruption.

This is analogous to the desired properties of forward error correcting codes where a single bit flip still permits to recover and correct a message, or a reliable transfer protocol where message retransmission can be requested, which in our case would correspond to rectifying an incomplete protocol software.

In other words: Like the basically unreliable transmission of messages where messages can be lost, reordered etc., we assume that some subset of the execution paths of the protocol software are executed in an unreliable way. We then ask whether communication software can be written such that it is able to recover itself in such circumstances.

If this can be achieved, it means that altering a single instruction will not do any harm to the protocol in question. A consequence of this is that it becomes in theory possible to disperse the code of this resilient protocol, such that each atomic instruction would be carried out by a different

processor. Register values could be shipped in packets between the nodes. Since the protocol is robust against the loss of a single instruction, it is now robust against the crash of any of the processors involved. In practice such partitioning would not occur on a single instruction basis but at the level of modules or code compartments. In this case we would like that any compartment crash be tolerable, while keeping instruction-level robustness inside the compartment. Robustness can be examined at different levels.

The applications of resilient and self-healing protocols are numerous: they would enable safe automated installation of new protocols, protocol upgrade, run-time customization of protocols to adapt to different network situations, distributed protocol implementations in sensor networks, spray computers [2], support for ambient intelligence and other networks of small devices, the dynamic placing of middle-box elements such as proxies and caches, and so on.

As a first approach to the problem we take inspiration from metabolic pathways in cells. These chemical processes are highly interlocked and surprisingly robust. This is of major interest to the pharmaceutical industry that is faced with the problem of identifying the multiple change points in a metabolic pathway in order to alter a cell's production levels (e.g. reproduction of a virus, cancer cell, etc.), where a single inhibition point is in general hard to find.

We build upon previous work on fraglets [3] as it permits to easily demonstrate a simple example of a robust piece of software. The fraglet model comprises a unified code/data format and execution engine inspired by metabolic networks in molecular biology.

The rest of this paper is organized as follows: Section 2 discusses related techniques including self-testing software, fault tolerant systems, resilience in today's protocols, and so on. Section 3 describes the fraglet model and its current instruction set. Section 4 gives a few initial hints on how resilient protocols can be constructed such that the loss of a fraglet has no impact on the service but leaves some traces behind which can be used to trigger a self-healing process. Section 5 presents our current conclusions and ideas to motivate a new branch of autonomic communication dedicated to resilient and self-healing code for autonomic network protocols.

2 Related work

2.1 Protocol robustness today

Resilience is a key requirement for any network. At the hardware level it is common to have redundant links, and redundant parts in core-network routers. At the software level, most current protocols and network services incorporate some form of robustness. We give a few examples below.

Robustness to link or node failure is generally achieved by rerouting traffic to alternative paths. In OSPF (Open Shortest Path First), node or link failures are detected via link advertisement messages, and new routes are recomputed accordingly. In BGP (Border Gateway Protocol), as well as in MPLS (Multi-Protocol Label Switching) route restoration is achieved via backup paths, such that service can be preserved during failure of the main path.

At the transport layer, TCP recovers from packet loss, duplicate packets, and congested paths, via an integrated retransmission and congestion control mechanism. Other adaptive transmission schemes are able to recover from similar disturbances by using mechanisms adapted to the nature of the traffic they transport, such as real-time streaming, or loss-tolerant voice/video.

The DNS (Domain Name Server) resolution service is a key pillar of the Internet. Resilience is a paramount concern, and is achieved through redundant servers.

Resilient Overlay Networks (RON) [4] are application-layer overlays on top of the current Internet. They seek to improve end-to-end reliability and performance by dynamically avoiding overloaded or faulty paths. RON nodes monitor path quality in order to detect and select higher quality paths. This makes it possible for the end systems to self-organize into more robust topologies that could not otherwise be offered by the standard Internet routing mechanisms.

These classic robustness protocols focus on external events such as node or link failure, link errors or congestion and are not robust to failure of the protocol implementation itself, except if the failure of an implementation on a node corresponds to the failing of the full node.

2.2 Fault tolerance in distributed systems

The complexity of distributed systems and their dependence upon the underlying hardware and network infrastructure expose them to several possible faults. Fault tolerance must therefore be an inherent part of their design, so that they can keep delivering the intended services with acceptable levels of performance and safety, even in the presence of failures.

There are multiple techniques for fault tolerance in distributed systems [5, 6, 7, 8], but most of them use variants or combinations of two main building blocks: state persistence and redundancy.

State persistence can be achieved via checkpointing and/or logging. A checkpoint is a snapshot of the process state at a given execution point. A log contains incremental state changes. Both can be used to resume execution after a failure, at the latest stage for which state has been consistently saved. This can diminish the impact of the failure.

Redundancy can take the form of multiple copies of the same process running on different machines, or multiple versions of a software component which implement the same functionality in different ways. For example, if failure of a version is detected, one can substitute the bad version with another one and try again.

A typical way to implement a fault-tolerant service is by replicating the servers at several independent locations and coordinating the updates so that at least one of the servers is available. When redundant processes (either via simple replication of code or via multiple versions) are used to achieve a result, a voting scheme is usually employed to decide on the output that must be actually produced.

Reconfiguration can be performed to recover from faulty processors by replicating the process to an operational processor, or to replace a malfunctioning version with a correct one. A model to dynamically reconfigure software in distributed systems is presented in [9]. It analyzes the dependencies among the different processes in the system to determine the impact of a given reconfiguration operation. The model has been used to build a fault-tolerant environment based on ARMOR processes (Adaptive Reconfigurable Mobile Objects of Reliability). However the model itself does not advise on which kind of reconfiguration should be performed under which circumstances. It leaves this task to system administrators or possibly other software components with the ability to make such decisions. Moreover, if the checkpointing or replication logic becomes corrupted, fault tolerance will not be achievable anymore. We call this an extrinsic approach where fault tolerance logic is incorporated into a system as an add on. Our interest is in an intrinsic solution where the robustness is part of the (protocol) software itself.

At this place we also mention “core wars”, where two programs share the same core memory and compete against each other by attacking the other program’s instructions. Robustness (and self-healing) are necessary attributes for letting a program survive in such a hostile environment.

2.3 Self-testing and self-correcting software

The usual methodologies for software testing and debugging rely on running the program under a controlled environment using a subset of possible inputs. The corresponding outputs must be known beforehand, and in most cases bugs arise after the software is deployed, due to untested combinations of input data or events.

Self-testing and self-correcting programs [10, 11] have been proposed to improve the reliability of software systems. Given a program P , a self-testing program for P is another, simpler program designed to make calls to P on a number of inputs, and to check whether the corresponding outputs are correct. A self-corrector for the same program P attempts to return the correct output for each input value, even in the presence of abnormal behavior of P , provided that the probability of P producing the wrong output is small enough.

While these techniques can in principle be useful for software debugging and testing, and to improve software reliability at run time, in practice especially the self-correcting functions can be hard to design, since they are very specific for each program. Another shortcoming of self-correcting functions is that they do not correct the program itself, but just attempt to correct wrong outputs.

In [12] the authors extend the notion of self-testing/correcting functions to distributed processing or protocols. Their approach has the same advantages and limitations of [10, 11], but applied to protocols in distributed systems, which makes the testing and correcting algorithms more complex.

Again, self-testing or correction is an extrinsic approach that works on a given program: it does not apply to self-testing and correcting its own operation.

2.4 Fault Tolerant Quantum Computing

A key aspect of quantum computing is that quantum state (e.g. in a qubit) should not interact with its environment in an uncontrolled way. Unlike classical analog computers, where state can be measured, restored and copied, it is not possible to copy quantum state or to remove inevitable noise from quantum operations by some threshold scheme. However, it was shown that with an appropriate encoding of a qubit as a codeword over several qubits, it is possible to implement error *detection* capabilities (a) with quantum operations, (b) such that the detection circuitry can also be subject to potential errors and (c) that errors can be *corrected* – again with quantum operations – permitting arbitrarily long quantum computations [13].

This corresponds exactly to the approach followed in this paper, where we want a protocol’s implementation to be robust in an intrinsic way such that errors can occur in the “core” protocol implementation as well as in the detection and correction part. The main difference is that we do not want to (re-)model communication software as a quantum system and that we aim directly at highlevel description and execution of protocols. That is, we do not want to (re-)implement basic logic gates and construct computation networks out of them as quantum computing theory does. Another important difference is that we are interested in software and arbitrary interconnections between code pieces, instead of fixed hardware circuits. Finally, in a first step we are interested in a programming environment where communications and computations can be executed without intermediate encoding step.

2.5 Resilience and self-healing properties of biological systems

Emergent behavior in biological systems leads to self-regulatory feedback systems that are robust to external perturbation and to failure of its constituent parts [14]. These systems tend to be highly decentralized, the emergent behavior resulting from simple interactions among autonomous agents that make decisions based solely on local views. These simple agents are often anonymous and non-specific, leading to intrinsic fault tolerance and self-healing properties, as other agents can easily take over the roles of failed or missing ones. Several examples can be cited, such as evolutionary selection with survival of the fittest, flock of birds, colonies of social insects, etc. In this section we concentrate on the biochemical processes occurring in cells, which inspired the Fraglets paradigm.

Within a living cell or microorganism there are several chemical processes responsible for maintaining the various cellular functions. These processes can be represented by graphs (networks) where the nodes are the chemical compounds and the links are the reactions that transform these compounds. Biochemical pathways describe the sequence of chemical reactions inside a biochemical network. Among the numerous cellular biochemical networks we can distinguish metabolic networks, responsible for the cell's energy cycle.

These cellular processes are known to be highly robust against mutations, due to several redundancy and diversity mechanisms, such as [15]: the presence of multiple genes with similar functions; interactions among genes with unrelated functions, such as in the case of recessive mutations; the existence of alternative routes in large metabolic networks; the scale-free nature of metabolic networks [16], which are dominated by a few highly connected hubs, while the vast majority of the nodes have a small number of connections, making them inherently robust to random errors.

3 The Fraglet Paradigm

The Fraglet paradigm [3] has been proposed as part of our search for feasible ways to achieve automated synthesis of protocol implementations. It is based on a chemical model where "molecules" interact with each-other or undergo some internal transformation. Formally, it is an instance of Gamma systems, which were proposed by Banatre and Metayer [17, 18]. Adopting a chemical model (with Fraglets as the only objects) has the benefit of being able to handle code and data at the same level and to integrate code deployment into protocols in a natural way.

A fraglet is a string of symbols $[s_1 : s_2 \dots : s_n]$ which represents data and/or protocol logic. It represents, so to speak, a fragment of a distributed computation. Fraglets may reside inside a node's fraglet store or may be carried in packets, where successive fraglet symbols are analogous to successive header fields in today's regular data packets. Upon arrival, a fraglet packet is injected into the local fraglet store or context.

The fraglet processing engine continuously executes tag matching operations on the fraglets in the store, in order to determine the actions that should be applied to them. Fraglet operations, except for the transmission, have the property that they can be carried out in constant time.

Formally, the store is a multiset: several instances of the same fraglet may be simultaneously present. This is indicated by a suffix counter value as in $[data : item]3$ (meaning that 3 copies of fraglet $[data : item]$ are stored in this context).

The fraglet instruction set currently contains two types of actions: transformation of a single fraglet, and reaction between two fraglets. Table 1 shows some transformation rules defined so

far. Table 2 shows the reaction rules.

Table 1: Transformation rules

<i>Op</i>	<i>Input</i>	<i>Output</i>
dup	[<i>dup</i> : <i>t</i> : <i>u</i> : <i>tail</i>]	[<i>t</i> : <i>u</i> : <i>u</i> : <i>tail</i>]
exch	[<i>exch</i> : <i>t</i> : <i>u</i> : <i>v</i> : <i>tail</i>]	[<i>t</i> : <i>v</i> : <i>u</i> : <i>tail</i>]
new	[<i>new</i> : <i>t</i> : <i>tail</i>]	[<i>t</i> : n_{i+1} : <i>tail</i>]
split	[<i>split</i> : <i>t</i> : ... : * : <i>tail</i>]	[<i>t</i> : ...], [<i>tail</i>]
send	A [<i>send</i> : B : <i>tail</i>]	B [<i>tail</i>] (unreliably)

The semantics of the transformation rules are:

- *dup*: Duplicates the symbol at the third position (*u*); the second field (*t*) becomes the new fraglet’s head symbol.
- *exch*: Swaps the two symbols at the third and fourth position (*u* and *v* respectively).
- *new*: Creates a new symbol n_{i+1} which is unique in this context.
- *split*: Breaks the fraglet into two parts at the first marker position (*).
- *send*: Sends the fraglet unreliably to the destination context specified by the second symbol (*B*). The subscript prefix $N[\dots]$ denotes the context where the fraglet is stored.

Table 2: Reaction rules

<i>Op</i>	<i>Input</i>	<i>Output</i>
match (merge)	[<i>match</i> : <i>s</i> : <i>tail</i> ₁], [<i>s</i> : <i>tail</i> ₂]	[<i>tail</i> ₁ : <i>tail</i> ₂]
matchp (persist)	[<i>matchp</i> : <i>s</i> : <i>tail</i> ₁], [<i>s</i> : <i>tail</i> ₂]	[<i>tail</i> ₁ : <i>tail</i> ₂], [<i>matchP</i> : <i>s</i> : <i>tail</i> ₁]

We introduced two simple reaction rules, listed in Table 2. The “merge” instruction (*match*) concatenates two fraglets with matching tags. The “persist” variant (*matchP*) moreover keeps a copy of the initial [*matchP* : ...] fraglet in the store, thus acts like a catalyst.

It is important to emphasize that our research in defining the instruction set is still in progress, thus the current state should be interpreted as a snapshot of an evolving work. In spite of this apparent limitation, in [3] two protocol implementations using fraglets have been shown: a very simple confirmed delivery protocol and a more complex flow control protocol with send credit and packet reordering.

4 Resilient Protocols

In this section we demonstrate a simple computation task in a communication context whose implementation is robust against partial erasure of its code base. We develop a “signaling frequency

doubler” program using Fraglets that is robust at the instruction level and we discuss the resulting code and insights gained from this exercise.

4.1 Robustness

Today we have a methodology of designing protocols that makes rather strong assumptions on the reliability of the executing components. The single execution environments are presumed to be stable: Reliable protocols that execute on these components are supposed to handle “only” the unreliable aspects of networking, like broken links, lost packets or transmission errors. That is, robustness applies to the harsh (communication) environment in which computers operate, not to the execution support itself.

Here, we redefine protocol resilience, or protocol robustness, as the ability to survive instruction failures: Even when portions of the protocol’s implementation are lost (or duplicated, or changed), the protocol is still able to provide the intended service, albeit with some loss of efficiency. Ideally, a robust protocol implementation will not only be able to detect but also to recover from code losses (self-healing).

Our definition of “protocol robustness” is linked to an implementation, and therefore to the instruction set it is based on. In this paper we use the Fraglet system as our instruction set and consider an example where any fraglet rule (which is a fraglet by itself) can be removed without changing the outcome of the computation.

4.2 A Robust Message Doubler with Fraglets

The simple task we want to solve is the doubling of a signal stream: for each message x we want two messages z to leave a node or to be available for further processing inside that node. For instance, the messages x could be the ticks from a Geiger counter sensor or any other source which encodes information as a frequency.

Rewriting the signal x (represented as a Fraglet $[x]$) to the new name can be done by the single (non-robust) fraglet rule

$$[matchp : x : z]$$

where the $[x]$ will be replaced by $[z]$. The doubling is done by a slightly larger rule, namely

$$[matchp : x : split : z : * : z]$$

which says that fraglet $[x]$ is replaced by a $[split...]$ fraglet. This new fraglet will break apart in the next processing step and produce two $[z]$ Fraglets, which is the desired effect.

Unfortunately, this program is not robust: Removing the rule will effectively erase the program. A trivial way to obtain robustness in the Fraglet is to double the rules:

$$\begin{aligned} &[matchp : x : split : z : * : z] \\ &[matchp : x : split : z : * : z] \end{aligned}$$

Here, any one rule instance can be removed, leading to the other rule to be invoked twice as much as before. This seems to satisfy the resilience condition: However, it is not possible to determine, from the program’s output, whether a fraglet (rule) has been lost or not. After the first loss, the program is no longer resilient, although it continues to run and produce results, until the single remaining copy is also lost. In order to achieve long-term resilience, the loss of the first Fraglet must be signaled, permitting to trigger a self-healing process.

Therefore we need an implementation such that the side effect of a loss is not harmful to its intended core functionality, but permits to identify which fraglet rule became unavailable, operating as a signal that can be used as input for a self-healing process. The next version of the doubling program, also shown in Figure 1, is:

```
[matchp : x : split : s1 : * : s1]
[matchp : x : split : s2 : * : s2]
[matchp : s1 : split : z : * : a]
[matchp : s1 : split : z : * : b]
[matchp : s2 : split : z : * : c]
[matchp : s2 : split : z : * : d]
```

In this case, we have *two different* rules that transform an input fraglet x into a *split* fraglet. This means that with a 50% chance, one of these two rules will be picked. Depending on which rule was picked, either two $[s1 : \dots]$ or two $[s2 : \dots]$ fraglets will be produced. At this level we have again two rules which transform an $[s1 : \dots]$ into two different $[split : \dots]$ fraglets. In fact, it is at this level that the doubling of the original messages is happening: half of the $[x]$ fraglets will become $[s1 : \dots]$ fraglets, but because of the two levels and branches of $[split : \dots]$ (resulting in a quadrupling), we have an overall doubling of the incoming flux of $[x]$ fraglets.

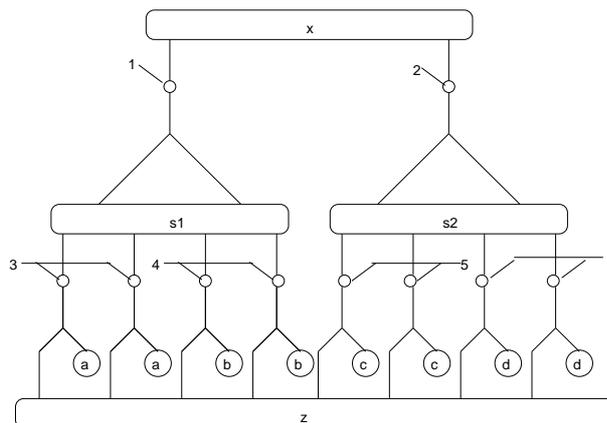


Figure 1: The processing of Fraglets for the robust doubler program.

One can observe that removing any one of these 6 fraglet rules will not change the net outcome. If the first rule is removed, the second rule will react with 100% of the influx: for each $[x]$ fraglet there will be two $[s2 : \dots]$ fraglets which each produce a $[z]$ fraglet as desired. Similarly, removing any of the other rules will just lead to its “homologue” rule to take over all intermediate fraglets instead of processing only 50% of them.

4.3 Discussion

The main “trick” of our demo program is the doubling of rules which compete against each other. This so called “soup”-aspect of chemical execution model makes it much easier to write robust code which continues operation despite some losses: any sequential execution model where the control flow must pass through a single instruction is potentially not amendable to robustness,

because there is no fallback execution path.

4.3.1 Self-Healing:

An important aspect of our program is that it produces a constant stream of a , b , c and d fraglets. These can be used to monitor the health of the program and to trigger repair mechanisms. Adding two simple rules like

$$\begin{aligned} &[matchp : a : match : c] \\ &[matchp : b : match : d] \end{aligned}$$

will immediately remove these signal streams in case everything works well. As soon as one of the Fraglet rules is removed, this balance is disturbed: By looking at the relative weights of $[a]$, $[match : a]$ etc. debris, one can infer for almost all cases which rule was removed. We note here that despite all these detection activities, the program continues to produce twice as many $[z]$ as $[x]$ Fraglets.

The design of control loops – which for example regenerate a lost $[matchp : \dots]$ rule is not trivial and, as our first experiments show, will probably lead to solutions where perfect robustness is not achievable anymore because the result stream can show temporary distortions. Nevertheless, the resulting program performs in an elastic way without fully disrupting the processing. Potentially, one should be able to produce programs which regain correct status after the healing process and given that the error rate is not too high.

4.3.2 Resource Control:

In a fully self-healing process we have to resolve tricky resource control issues: unlike the simple demo example above, which reacts on the *loss* of fraglet rules, one also has to anticipate an *excess* of rules. This leads to a completely dynamic program where the $[matchp : \dots]$ rules are replaced by $[match : \dots]$ rules. These later fraglets have to be constantly regenerated, thus leading to a program that is constantly rewriting itself according to the sensed performance of itself. Metaphorically, we need transcription signals which control the gene expression.

4.3.3 Distribution

As mentioned in the introduction, our program can be put in a communication context. By changing the intermediate fraglet types such that they undergo a “transmission”-transformation, one can have different parts of the process to occur on different nodes. Assuming that we have four nodes n , m , o and p , we would have the rules:

$$n[\text{matchp} : x : \text{send} : m : \text{split} : s1 : * : s1]$$

$$n[\text{matchp} : x : \text{send} : o : \text{split} : s2 : * : s2]$$

$$m[\text{matchp} : s1 : \text{split} : \text{send} : p : z : * : \text{send} : p : a]$$

$$m[\text{matchp} : s1 : \text{split} : \text{send} : p : z : * : \text{send} : p : b]$$

$$o[\text{matchp} : s2 : \text{split} : \text{send} : p : z : * : \text{send} : p : c]$$

$$o[\text{matchp} : s2 : \text{split} : \text{send} : p : z : * : \text{send} : p : d]$$

$$p[\text{matchp} : a : \text{match} : c]$$

$$p[\text{matchp} : b : \text{match} : d]$$

From node n we rewrite $[x]$ fraglets into fraglets that transfer themselves to nodes m and o and split there. On these nodes, the doubling is performed and the resulting stream is redirected to node p . Note that here we start to blur the distinction between traditional robustness protocols, as the crash of a node m (resulting in removing two rules) would be observable by an imbalance of the monitoring signals and could trigger the necessary repair actions.

5 Conclusions and Outlook

In this paper we have demonstrated a simple program that continues to perform its task despite the removal of any of its instructions. This “intrinsic” robustness is different from the usual extrinsic fault-tolerance approaches as it weaves self-monitoring and self-healing into the proper processing. Unlike quantum computing, where intrinsic robustness has also been used to make computations robust against errors in the hardware execution environment, our aim is to produce resilient communication *software*.

Redundancy through a multiset model and parallelism are built-in concepts in the Fraglet model, on which we based our example program. They can be exploited to achieve intrinsic fault tolerance both in a local context, as shown, but potentially also on a network wide basis by using the seamless moving of fraglets between nodes without data and code distinction.

Transforming existing protocols into a robust fraglet implementation is probably not a straight forward process. Instead, we have chosen a special case of rate encoding similar to the chemical process inside a cell where control loops are realized through the flux of chemical molecules. Another venue, similar to the path taken by the quantum computing community, could be the transformation of programs into resilient forms and to execute them in an encoded form.

The list of potential research topics in this area is immense. Resilience should be extended to a more general scope, beyond single instruction failure. The current fraglet model does not incorporate resource control features. It is possible for a fraglet to generate an infinite number of other fraglets and/or send them to other nodes. Maybe one solution would be to impose rules similar to conservation of matter (or symbols, in this case), to limit the amount of new fragments that can be generated. Economically-inspired schemes could also be incorporated, such that the services provided by fraglet-based protocols only remain active when there are clients able to pay for them. Or maybe an energy abstraction could be used where chemical reactions need energy to take effect, and the available energy is finite.

The fraglet model has been designed with the intention to achieve automated protocol synthesis. An interesting research topic would be to investigate methods to synthesize fraglet code from

autonomic communication rule sets [1].

Realistic methodologies for protocol synthesis will certainly take a long time to mature. These methodologies will have to rely on robustness features and code introspection, since one cannot assume that error-free protocol implementations will be always derived, which are at the same time able to function in a highly asynchronous and dynamic environment. The next open question is then how to synthesize self-healing protocols.

References

- [1] Smirnov, M.: *Autonomic Communication: Research Agenda for a New Communication Paradigm*. White paper, Fraunhofer FOKUS (2003)
- [2] Zambonelli, F., Gleizes, M.P., Mamei, M., Tolksdorf, R.: *Spray Computers: Frontiers of Self-Organization*. In: *Proceeding of 1st International Conference on Autonomic Computing (ICAC'04)*, New York, USA (2004) 268–269
- [3] Tschudin, C.: *Fraglets - a Metabolic Execution Model for Communication Protocols*. In: *Proceeding of 2nd Annual Symposium on Autonomous Intelligent Networks and Systems (AINS)*, Menlo Park, USA (2003)
- [4] Andersen, D.G., Balakrishnan, H., Kaashoek, M.F., Morris, R.: *Resilient Overlay Networks*. In: *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP 2001)*, Banff, Canada (2001)
- [5] Jalote, P.: *Fault Tolerance in Distributed Systems*. ISBN 0133013677. Pearson Education (1994)
- [6] Dialani, V., Miles, S., Moreau, L., Roure, D.D., Luck, M.: *Transparent Fault Tolerance for Web Services based Architectures*. In: *In Proceedings of 8th International Europar Conference (EURO-PAR'02)*, Paderborn, Germany (2002) 889–898
- [7] Mullender (Ed.), S.: *Distributed Systems, Second Edition*. ACM Press, Addison-Wesley (1993)
- [8] Torres-Pomales, W.: *Software Fault Tolerance: A Tutorial*. Technical Report TM-2000-210616, NASA (2000)
- [9] Whisnant, K., Kalbarczyk, Z.T., Iyer, R.K.: *A system model for dynamically reconfigurable software*. *IBM Systems Journal* **42** (2003) 45–59
- [10] Blum, M., Luby, M., Rubinfeld, R.: *Self-Testing/Correcting with Applications to Numerical Problems*. *Journal of Computer and System Sciences* **47** (1993) 549–595
- [11] Wasserman, H., Blum, M.: *Software Reliability via Run-Time Result-Checking*. *Journal of the ACM (JACM)* **44** (1997) 826–849
- [12] Franklin, M., Garay, J., Yung, M.: *Self-Testing/Correcting Protocols*. In Jayanti, P., ed.: *7th International IS&N Conference on Intelligence in Services and Networks (ISN'00)*. Springer-Verlag LNCS 1693, Bratislava (1999) 269–283

- [13] Preskill, J.: Fault-Tolerant Quantum Computation (1997) arXiv:quant-ph/9712048.
- [14] Anthony, R.J.: Emergence: A Paradigm for Robust and Scalable Distributed Applications. In: International Conference on Autonomic Computing (ICAC-04), New York, USA (2004)
- [15] Fontana, W., Wagner, A.: Mutational Robustness, Modularity and Evolvability. Research focus area robustness, Santa Fe Institute (2002) <http://www.santafe.edu/sfi/research/focus/robustness/projects/mutationalRobustness.html>.
- [16] Jeong, H., Tombor, B., Albert, R., Oltvai, Z., Barabasi, A.L.: The large-scale organization of metabolic networks. *Nature* **407** (2000) 651–654
- [17] Banatre, J.P., Metayer, D.L.: A new computational model and its discipline of programming (1986) Technical Report RR0566, INRIA.
- [18] Banatre, J.P., Fradet, P., Radenac, Y.: Principles of chemical programming (2004) Fifth International Workshop on Rule-Based Programming (RULE'04).