

UNIVERSITY OF OSLO  
Department of Informatics

My master thesis title

My Name  
Oslo University College  
Note special front page

December 19, 2006



---

# My master thesis title

My Name  
Oslo University College  
Note special front page

December 19, 2006

## **Abstract**

*You should write an abstract that summarizes the problem you are studying and the main results you discover. Here is an example based on a paper that Mark Burgess and Tore Jonassen wrote.* We consider the implications of feedback on policy-based management of networked services, paying special attention to the case where monitoring of Service Level Agreements is used to alter policy dynamically. Using a simple model, we show that there are policies that are unstable under service provision, i.e. provide no reliable service levels (QoS), and that there are policies that can yield stable behaviour; hence we conclude that policy rule combinatorics can damage quality of service goals.

# Contents

|          |                                      |           |
|----------|--------------------------------------|-----------|
| <b>1</b> | <b>Introduction</b>                  | <b>2</b>  |
| <b>2</b> | <b>Methodology</b>                   | <b>3</b>  |
| 2.1      | Here we go . . . . .                 | 3         |
| 2.2      | Two reactive policy models . . . . . | 4         |
| 2.2.1    | Model 1: observe and react . . . . . | 6         |
| 2.2.2    | Model 2: causal chaining . . . . .   | 7         |
| 2.3      | Analysis of the models . . . . .     | 8         |
| 2.4      | Diagrams . . . . .                   | 11        |
| <b>3</b> | <b>Results</b>                       | <b>12</b> |
| 3.1      | Data . . . . .                       | 12        |
| 3.2      | Theory . . . . .                     | 12        |
| 3.3      | Describing the results . . . . .     | 13        |
| <b>4</b> | <b>Conclusions and Discussion</b>    | <b>16</b> |

# Chapter 1

## Introduction

*Your thesis needs a proper introduction to motivate the project. You should refer to previous literature and set the scene for an intelligent reader. You have written reports like this before in the lab course and in the research papers course. Here is an example:*

Service provision is normally thought of as the delivery of fixed-rate, policy-controlled transactions between willing parties. A contract of service, called a Service Level Agreement (SLA) documents the promise that the service provider makes to the client. The SLA is a form of policy for the server to uphold[1, 2]; what distinguishes it from other policy rules is that it is not deterministically enforceable, since the parties have no direct control over one another.

Bla bla bla.....

## Chapter 2

# Methodology

*You should explain your method most carefully. Do not assume that the reader has followed all the courses that you have. Do not write the report for one of the staff at this college, who might know what courses you have taken. Your methodology needs to summarize the main techniques and theory that you are using to solve the problem, so that an intelligent expert could read it.*

*If you cite words or results from the literature or from books, or (God forbid) the Internet, you MUST place it in quotes and provide a reference to it. If you copy text without a reference, you could be failed for cheating. For example:*

### 2.1 Here we go

To study the complexity that can emerge from the interaction of different policies, we begin by considering the simplest case of just two interacting peers. This scenario is already sufficiently complex to fill the present paper.

From the theory of dynamical systems, it is known that the combination of several dynamical variables using AND (multiplication) and OR (addition) in feedback loops can lead to ‘chaotic’ or non-linear behaviour. Such combinatoric ideas have been discussed to provide more expressive policies for service level agreements[4, 5, 6, 7, 8]; for example, in ref. [4], the authors discuss ‘aggregated Quality of Service parameters’ of the form:

$$S \text{ depends\_on } (SS_1 \text{ AND } SS_2) \quad (2.1)$$

as well as more general rules like:

$$QoS_A(S) = f(QoS_B(SS_1) \text{ AND } QoS_B(SS_2)). \quad (2.2)$$

They are basing service quality promises on a function of measured values. However, there are problems with combining dynamical variables non-linearly on finite precision systems. It is known that such systems can spiral out of control, or simply ebb to zero for certain ranges of parameters (for a review, see [15]). It is therefore important to understand how such behaviour emerges in dynamically regulated policies and what can be done to secure stable operation.

To avoid unnecessary complications, we attempt to answer a basic question: whether or not a reactive policy, based on mutual observation of several data points, can lead to stable behaviour, and whether the criterion of stability places any restrictions on the allowed policy used by the nodes.

Our basic viewpoint is that .....

In this paper, we base our models on this assumption and analyse their stability.

## 2.2 Two reactive policy models

We shall compare two similar reactive policy models using the language of dynamical systems; the first is a model that is well-known from game theoretical literature, the second is a modification of the first. We represent the feedback loop of negotiations between the two parties as the iteration of reactive algebraic rules over time.

Our model system is comprised of two parties who promise one another a service level for a fixed time interval  $\Delta t$ . Each party fixes its policy towards the other for this given time interval; then, after this, they re-evaluate their promise to one another for the next period using the data they have collected in previous rounds.

We use the method of difference equations to model the iterative development of policy, or SLA, between the two systems. Using this method, we represent the state and development of each peer's policy by a pair of coupled equations that relate observations of the past and present to promises about the future. Suppose we measure a value  $q(t)$ , at time  $t$ , characterizing the actual delivery of service at that time, we would then make our promise for the next period based on a function  $P(\cdot)$  of this value:

$$q(t+1) = P^{(1)}(q(t); \alpha). \quad (2.3)$$

In this example, we have a simple first order Markov chain: each new promise is based only on the results of the last iteration. In general, we can define a second order rule by:

$$q(t+1) = P^{(2)}(q(t), q(t-1), \dots; \alpha_1, \alpha_2, \dots), \quad (2.4)$$

to allow policy to track more of the historical chain of the interaction.

Note that the values  $q(t)$  are real valued functions, as is  $f$ , since service levels and promises are *averaged* quantities that are subject to fluctuations. Over an average time interval, small variations occur that allow us to treat the actual values as pseudo-real numbers, or finite precision floating point variables.

Let each host characterize its measured service level in terms of a parameter:

$$q = \frac{s}{S}, \quad (2.5)$$

where  $s$  is the actual service level provided over the time interval  $\Delta t$  and  $S$  is some constant scale that is unimportant to the discussion. For example,  $s$

## 2.2. TWO REACTIVE POLICY MODELS

---




Figure 2.1: The support for the tit-for-two-tat model. The behaviour at  $t + 1$  is based on two observations at  $t$  and  $t - 1$ .

might be the promise to deliver E-mail in less than  $s$  seconds, or the promise to deliver traffic at a rate of  $s$  bytes per second. By introducing an arbitrary scale of measurement  $S$ , we can represent measured values in a dimensionless form. (In the second model we shall call the analogous quantity  $Q$  to distinguish it from the first model.)

The service delays, due to environmental interference, and the errors made during observation affect each host at random. They are random variables, so we write these as a stochastic perturbation:

$$s = \langle s \rangle + \delta s, \quad (2.6)$$

in order to properly capture the dynamics of a realistic interaction. Here, the average service level in a given round is:

$$\langle s \rangle(t) = \frac{1}{\Delta t} \int_t^{t+\Delta t} s(x) dx \quad (2.7)$$

and  $\delta s$  is a random variable (e.g. Gaussian noise). Thus:

$$q(t) = \langle q \rangle(t) + \delta q. \quad (2.8)$$

Our promised service levels must be associated with the *average* or expected service levels, since we cannot promise more than average behaviour[3]. The revised service level agreements of the two hosts at round  $t$  is a specification of the expected service level up to  $t + 1$ .

### 2.2.1 Model 1: observe and react

Consider the first of our models (see fig 2.1); the two SLAs are:

$$\begin{aligned} \langle q_1 \rangle(t+1) &= \alpha_2 q_2(t) + \beta_2 q_2(t) q_2(t-1) \\ \langle q_2 \rangle(t+1) &= \alpha_1 q_1(t) + \beta_1 q_1(t) q_1(t-1) \end{aligned} \quad (2.9)$$

With the random perturbation expansion, this becomes:

$$\begin{aligned} q_1(t) &= \langle q_1 \rangle(t) + \delta q_1 \\ q_2(t) &= \langle q_2 \rangle(t) + \delta q_2, \end{aligned} \quad (2.10)$$

## 2.2. TWO REACTIVE POLICY MODELS




Figure 2.2: The support for the conscionable tit-for-tat model (model 2). The behaviour at  $t + 1$  is based on observations at  $t$  of the other player and at  $t - 1$  for one-self.

giving us two coupled equations for the time-development of the service level promises, and hence the average development of policy for the two peers.

The service level agreement has two terms; these can be given a game-theoretical interpretation:

$$\langle q_1 \rangle(t+1) = \underbrace{\alpha_2 q_2(t)}_{\text{tit-for-tat}} + \underbrace{\beta_2 q_2(t)q_2(t-1)}_{\text{tit-for-2tat}} \quad (2.11)$$

The tit-for-tat term says that we should try to reciprocate the service level that our peer gave us on the previous round. The quadratic term is like  $q_1(t)$  **AND**  $q_1(t-1)$  says that we should add on a goodwill term depending on whether the peer has been faithful for the previous two rounds, i.e. we are rewarding its generosity.

The behaviour of this policy is coarsely summarized in the following table for peer 1:

|                                  |      |      |     |      |
|----------------------------------|------|------|-----|------|
| SetPolicy <sub>1</sub> ( $t+1$ ) | High | ok   | Low | ok   |
| Measure <sub>2</sub> ( $t$ )     | High | High | Low | Low  |
| Measure <sub>2</sub> ( $t-1$ )   | High | Low  | Low | High |

Peer 1 makes a promise for the period up to  $t + 1$  based on measurements of peer 2 at  $t$  and  $t - 1$ . If peer 2 has faithfully given good service for two rounds, we reward with full service, if not Peer 1 retaliates by promising less for the next period. Conversely, peer 1 punishes bad behaviour.

We could develop this pattern to include other types of term, e.g. we also punish poor-behaviour (small  $q$ ) with a term of the form:  $(1 - q)_t(1 - q)_{t-1}$  i.e. **NOT**  $q(t)$  **AND** **NOT**  $q(t-1)$ . When  $q$  is large  $1 - q$  is small, thus it represents the complement of  $q$ .

### 2.2.2 Model 2: causal chaining

As we shall show shortly, our first policy model has a number of problems, so we consider a modification of this that has a 'conscience' as well as a reactive tit-for-tat behaviour. We define our second model in terms of policy promises

### 2.3. ANALYSIS OF THE MODELS

---

$Q$  for the two SLAs (see fig 2.2):

$$\begin{aligned}\langle Q_1 \rangle(t+1) &= \alpha_2 Q_2(t) + \beta_2 Q_2(t)(1 - Q_1(t-1)) \\ \langle Q_2 \rangle(t+1) &= \alpha_1 Q_1(t) + \beta_1 Q_1(t)(1 - Q_2(t-1))\end{aligned}\tag{2.12}$$

The interpretation of the terms is as follows:

$$\langle Q_1 \rangle(t+1) = \underbrace{\alpha_2 Q_2(t)}_{\text{tit-for-tat}} + \underbrace{\beta_2 Q_2(t)(1 - Q_1(t-1))}_{\text{conscience}}\tag{2.13}$$

The first term is a short term tit-for-tat retaliation, as before. The latter term mitigates the retaliation by checking whether or not it is justified based on prior provocation (it has the form  $Q_2(t)$  **AND NOT**  $Q_1(t-1)$ ). In other words, it checks what one peer does to provoke the other before jumping to conclusions about reprisals. The behaviour of this policy is summarized roughly by the following table:

|                                  |      |      |      |     |
|----------------------------------|------|------|------|-----|
| SetPolicy <sub>1</sub> ( $t+1$ ) | Low  | Good | ok   | ok? |
| Measure <sub>2</sub> ( $t$ )     | Low  | Good | Good | ok? |
| Measure <sub>1</sub> ( $t-1$ )   | High | Bad  | Good | ok? |

If peer 1 gives good service and peer 2 gives poor, then peer 1 returns poor service. If peer 1 realizes it gives poor service, then peer 2 gives good service, peer 1 returns good service. If both peers gave good service, the service level is slightly reduced but not by much, so that the peers back off a little so as not to get too carried away! In the final column, it seems reasonable that an okay result from both peers would lead to an okay promise in return, so we expect that there will be a balance or fixed point at an 'ok' value.

To test these claims, we move to a more formal approach that can be analyzed impartially.

### 2.3 Analysis of the models

The simplified and reduced models can be written the following form, for the special symmetrical case  $\alpha_1 = \alpha_2 \rightarrow \alpha$  and  $\beta_1 = \beta_2 \rightarrow \beta$ . Note that, for the second model, the parameters are adjusted  $\alpha - \beta \rightarrow \alpha$  to write the expressions in this form. Model 1 becomes:

$$\begin{aligned}q_1(t+1) &= \alpha q_2(t) + \beta q_2(t)q_2(t-1) \\ q_2(t+1) &= \alpha q_1(t) + \beta q_1(t)q_1(t-1)\end{aligned}\tag{2.14}$$

Model 2 becomes:

$$\begin{aligned}Q_1(t+1) &= \alpha Q_2(t) + \beta Q_2(t)Q_1(t-1) \\ Q_2(t+1) &= \alpha Q_1(t) + \beta Q_1(t)Q_2(t-1)\end{aligned}\tag{2.15}$$

### 2.3. ANALYSIS OF THE MODELS

---

We would like to know: do these mappings have fixed points, i.e. points at which the behaviour falls into a dynamical equilibrium of constant service levels? Does both stable and chaotic behaviour emerge for some values of the parameters? How should we choose the parameters in order to get stable behaviour?

The procedure for analyzing dynamical systems of this type is somewhat technical and will be presented elsewhere. We outline only the main steps and the results. Given that the dynamical system is relatively simple, one can proceed algebraically, thereby eliminating the dangers of working with non-linear numerical simulations and arriving at an erroneous conclusion. (Note however, that any policy engine basing itself on finite precision calculations is subject to problems of instability in practice; this inherent problem will not affect our conclusions here, however.)

The models have six degrees of freedom ( $q_1, q_2$ , at three different times), thus with two constraints, the models are four dimensional. To see this more clearly, we rewrite the iterated policies. For Model 1 let  $x_t = q_1(t)$ ,  $y_t = q_2(t)$ ,  $z_t = q_1(t-1) = x_{t-1}$  and  $w_t = q_2(t-1) = y_{t-1}$ . Then  $z_{t+1} = x_t$  and  $w_{t+1} = y_t$ . Substitution of the variables  $x, y, z, w$  into the equations for  $S_1$  gives

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t w_t \\ y_{t+1} &= \alpha x_t + \beta x_t z_t \end{aligned} \quad (2.16)$$

By adding the equations for  $z$  and  $w$  we obtain a two parameter first order system  $(\alpha, \beta)$ . where

$$F_\zeta(x, y, z, w) = (\alpha y + \beta y w, \alpha x + \beta x z, x, y).$$

The induced dynamical system is given by

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t w_t \\ y_{t+1} &= \alpha x_t + \beta x_t z_t \\ z_{t+1} &= x_t \\ w_{t+1} &= y_t \end{aligned} \quad (2.17)$$

This expression is called the first order representation of Model 1. The mapping  $F_\zeta$  is called the *map representation* of the model. The corresponding mapping in the symmetric case will be denoted by  $F_\eta$  where  $\eta = (\alpha, \beta)$ .

Using the same substitutions, with the obvious modification  $q_i = Q_i$ , in the Model 2 we obtain

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t z_t \\ y_{t+1} &= \alpha x_t + \beta x_t w_t \end{aligned} \quad (2.18)$$

and hence a map given by

$$G_\zeta(x, y, z, w) = (\alpha y + \beta y z, \alpha x + \beta x w, x, y).$$

## 2.4. DIAGRAMS

---

The induced dynamical system is then given by

$$\begin{aligned}x_{t+1} &= \alpha y_t + \beta y_t z_t \\y_{t+1} &= \alpha x_t + \beta x_t w_t \\z_{t+1} &= x_t \\w_{t+1} &= y_t\end{aligned}\tag{2.19}$$

This expression is called the first order representation of Model 2, and  $G_\zeta$  is its map representation.

We now have the basic ‘equations of motion’ of the models, and would like to determine whether they have any fixed points of stable operation, i.e. can we find values for the service provision that satisfy

$$\begin{aligned}F_\zeta(x, y, z, w) &= (x, y, z, w) \\G_\zeta(x, y, z, w) &= (x, y, z, w)\end{aligned}\tag{2.20}$$

that is, for the former, the following set of four equations in four unknowns:

$$\begin{aligned}x &= \alpha y + \beta y w \\y &= \alpha x + \beta x z \\z &= x \\w &= y\end{aligned}\tag{2.21}$$

It is easily seen that this gives the two coupled quadratic equations

$$\begin{aligned}x &= \alpha y + \beta y^2 \\y &= \alpha x + \beta x^2\end{aligned}\tag{2.22}$$

using that  $z = x$  and  $w = y$ . Clearly  $x = y = 0$  will solve this equation, and hence  $(x, y, z, w) = (0, 0, 0, 0)$  is a fixed point of  $S_1$  for all parameter values. We will denote this trivial fixed point by  $p_0$ . This corresponds to no service provided by either system (which is clearly a stable situation, if not an interesting one). Non-trivial solutions are determined by equating and reducing the expressions above, and looking for the polynomial roots of the systems of equations. We shall not show these details.

## 2.4 Diagrams

*To explain your model you should include helpful diagrams and explain any calculations that you are going to use. Do not assume that the reader will think like you do. Make sure you explain your method and give references to textbooks or papers that explain the method you have used.*

Refer to your figures in text, e.g. see fig. 2.3.




Figure 2.3: Every diagram of figure should have a caption explaining what it is, and highlighting any special features that the reader should notice about it. Notice that this figure has four parts (a), (b), (c) and (d).

## Chapter 3

# Results

*In presenting your results you should show data as tables and as graphs. Do not forget to label every graph and write a caption explaining what it is and what it tells us. Each figure should stand for itself, without needing too much explanation in the text.*

### 3.1 Data

*Do not forget to illustrate and estimate your uncertainties. The worst thing you can do is to hide the uncertainties in your work. This looks unprofessional.*

Tables of data should include estimates of the experimental errors (see table 3.1). Note that the values you claim should never be cited to greater accuracy than your basic uncertainty, i.e. you should never write  $\langle q \rangle = 0.677239972 \pm 0.1$  (this is silly).

| $t$ | Value $q(t)$ | $\Delta q$ |
|-----|--------------|------------|
| 0.1 | 14.44        | 0.01       |
| 0.2 | 13.65        | "          |
| 0.3 | 12.22        | "          |

Table 3.1: Every table and figure should have a caption telling the reader what they are looking at and drawing attention to anything special about these data.

### 3.2 Theory

*You will probably need to point to (if not explain) some theoretical background that you use to analyze the data in your project. This is different from the theory that motivates the problem itself.*

As noted above, it is possible to compute the fixed points for Model 1 exactly in the case we have outlined. We have  $z = x$  and  $w = y$  and hence we only need the two first components in the points in order to have all four components. The model can have up to four real fixed points given by:

$$\begin{aligned} p_1 &= (0, 0) \\ p_2 &= \beta^{-1}((1 - \alpha), (1 - \alpha)) \end{aligned}$$

### 3.3. DESCRIBING THE RESULTS

---

Note that, since negative service levels have no meaning for us here, we are limited to solutions in which the values are positive. This means that we must require  $0 < \alpha \ll 1$  for a high level of stable service. Notice also that it is  $\alpha$ , the first order tit-for-tat strategy which drives this;  $\beta$ , from the second order term, is only a regulating parameter that prevent the fixed point from blowing up to infinity.

If  $\alpha \in (-\infty, -1] \cup [3, \infty)$  then there are two additional fixed points given by

$$p_{3,4} = (2\beta)^{-1} (\pm\gamma(\alpha) - (\alpha + 1), \mp\gamma(\alpha) - (\alpha + 1))$$

where  $\gamma = \sqrt{(\alpha - 3)(\alpha + 1)}$ . The *topological type* of these fixed points is determined by the norm of the eigenvalues of  $4 \times 4$  Jacobian matrices  $DF_{\eta}$  evaluated at the fixed points. A simple computation shows that the eigenvalues of the Jacobian at  $p_1$  are 0 (multiplicity 2) and  $\pm\alpha$ . The eigenvalues at  $p_2$  are:

$$\begin{aligned} \lambda_{1,2} &= \frac{1 \pm \sqrt{5 - 4\alpha}}{2} \\ \lambda_{3,4} &= \frac{-1 \pm \sqrt{4\alpha - 3}}{2} \end{aligned}$$

The expressions for the eigenvalues of the Jacobians at the points  $p_3$  and  $p_4$  are large, and we will not give the formulas here; however they are easily found by a computer algebra system such as *Mathematica*. It should be noted that these eigenvalues are independent of  $\beta$  and depends only on  $\alpha$ , thus it is the tit-for-tat component of policy that leads to stability here, not the **AND**-ed combinatorics.

### 3.3 Describing the results

*Do not just present reams of data without explanation. You should explain what you believe the data tell you. For every quantitative set of measurements, there should be a qualitative description and explanation.*

We find that, for Model 1, for a large set of parameters, initial points chosen in the unit square have two different fates: either the orbit will converge to the trivial fixed point at zero; or the orbit escapes to infinity (saturates the service levels), if it starts from a sufficiently high value. In other words, either service levels grow to their maximum ceiling, or they fall to zero. The set of points that could have a stable fixed point is so narrow that even the smallest perturbation would destroy its stability.

For the deterministic (non-noisy) case, this model exhibits attractive, periodic orbits, i.e. service levels that chase each other through oscillations within a stable range of values. These have many different periods and bear a resemblance to models of population dynamics in biology or economics.

What is interesting about Model 2 is that it has a much wider region of parameters that lead to stable fixed points than does Model 1. Whereas Model 1 essentially has no regions that are stable to perturbation, Model 2 has several ‘phases’. In the first phase, only the zero fixed point  $p_0$  is stable, then as  $\alpha >$

### 3.3. DESCRIBING THE RESULTS

---




Figure 3.1: An attracting orbit projected into the plane for a symmetric version of Model 2. This shows a region of stable behaviours in which service levels will oscillate within a bounded region.




Figure 3.2: The effect of noise of figure 3.1. We have used a random initial value in the unit 4-cube.

### 3.3. DESCRIBING THE RESULTS

---




Figure 3.3: The projection of an orbit of a random system where the deterministic system exhibit an stable orbit of period 5. The random system is comuted by adding noise on each iteration.

1,  $\beta < 0$  there is a region in which  $p_2$  is stable. Finally, there is another region in which there is a Hopf bifurcation and more complex behaviour emerges.

In figure 3.3, the deterministic version of Model 2, we see an attractive orbit of period 5. This means that service levels between the peers can jump discontinuously between five different rates as time progresses.

The original models, given in equation (7) and equation (11), use random variables to represent the noisy perturbation of the environment that can drive instabilities, in the first place. We have simulated these models by adding Gaussian noise in the defining functions for the discrete dynamical systems. In the case of the attractor shown in figure 3.1, the corresponding projection of the random attractor is given in figure 3.2. This shows us the effect of environmental noise.

The corresponding random system shows a similar behaviour, which means that service levels will jump between five different levels as the system evolves over time! This is clearly rather complex behaviour from the innocent looking **AND** construction, which is a powerful indicator that the stability of Service Level Agreements is not a trivial matter.

Further examples of using *Mathematica* in various discrete dynamical systems simulations including chaotic dynamics may be found in ref. [20].

## Chapter 4

# Conclusions and Discussion

*Conclusions begin by stating the original problem again and explaining how far you have come in answering the questions you had. If you did not have any questions, you are in big trouble!*

We have studied a simple algebra of service level policies, coupled back-to-back in a feedback loop. This models a two-party collaborative agreement, of the kind that is typified by BGP peering and other mutual service agreements like Quality of Service (QoS) levels. The two toy-models that we have used here are the simplest models that exhibit non-linearity (combinatoric logic): they are easy to analyze yet they still provide us with some interesting information. We look for dynamical *fixed points*, or *stable* behaviour since service guarantees (QoS) require stability.

In Model 1, we have simple reactive tit-for-tat term, and a bonus tit-for-two-tats mixture of strategies that rewards faithfulness and punishes persistent defection. We find that, for only a few policies  $(\vec{\alpha}, \vec{\beta})$ , is there a possibility of stable cooperative service levels. Thus, the presence of noise is likely to destabilize these except in carefully tuned scenarios.

Put another way, the robustness of Model 1 is questionable. In many cases, the service levels will simply decay to zero: as noise leads to deviations from the agreement, each host will punish the other until the services levels spiral down to zero. For a small range of parameters, the tit-for-tat strategy allows service levels to grow to a stable point, and the tit-for-two-tats strategy limits it to a sensible level (as  $\beta \rightarrow 0$ , the fixed service point goes off to infinity, so we must have it).

In Model 2, we have tit-for-tat and a mitigating conscience term that observes whether retaliation was justified or not. In this Model, the qualitative behaviour is similar, but now there is an even greater sensitivity to the policy parameters  $(\vec{\alpha}, \vec{\beta})$  and there is a wider region of stability for constant (guaranteeable) service levels.

So, if the parties start at too low a level, their SLAs will mistrustingly decay to zero. If they start at a high level, their promises will grow stronger and stronger until they reach a maximum service level (theoretically infinite, according our model).

Thus, if one is interested in autonomic stability, it is far from clear that

---

one should employ complex policies based on monitoring and reprisal. There are inherent instabilities in service level policies based on combined measurements (using **AND** and **OR** ).

The main message of this paper is that feedback regulation of SLA-like policies is a highly complex matter that must be parameterized with caution. Often we tend to assume that more constraint parameters leads to tighter and more controlled behaviour, but there is evidence here to suggest that the opposite can be true. Stability seems to emerge mainly from the simplistic first order tit-for-tat ‘dumb’ term, because it allows for ‘recovery’ of the cooperation. (Note that the second model has an implicit term of this form, since it includes an ‘**AND NOT**’ operation that brings in first order behaviour from the second order term, but that one then needs to tune the parameters of the model carefully to find the regions of stability.)

The quadratic terms determine the boundedness of the behaviour, but also lead to complexity and multiple patterns of stability (see ref. [16] for concepts of stability). Without these terms, the non-trivial fixed points are based on ‘blind trust’, i.e. one has to start with a high level SLA and this will then grow to a maximum value and remain there, always saturated.

We show that there exist policies that have no stable fixed point that can be associated reliable service levels, other than the trivial case of no service. If two parties enter into a policy agreement of this type, they will either undergo a long fibrillation of changing behaviour, or they will spiral down to offering zero service levels.

The results we have presented here are only for a pair of interacting nodes. The extension to multiple systems joined in more complex networks is, in principle, doable though the dimension of the problem increases for each new node, as does the possibility for complex behaviour. We shall investigate this behaviour in future work.

# Bibliography

- [1] M. Sloman. Policy driven management for distributed systems. *Journal of Network and Systems Management*, 2:333, 1994.
- [2] D. Verma et al. Policy based sla management in enterprise networks. In *Policy Workshop 2001*. Springer Verlag, 2001.
- [3] M. Burgess. On the theory of system administration. *Science of Computer Programming*, 49:1, 2003.
- [4] G.B. Rodosek. Quality aspects in it service management. *IFIP/IEEE 13th International Workshop on Distributed Systems: Operations and Management (DSOM 2002)*, page 82, 2002.
- [5] D. Daly, G. Kar, and W. Sanders. Modeling of service-level agreements for composed services. *IFIP/IEEE 13th International Workshop on Distributed Systems: Operations and Management (DSOM 2002)*, page 4, 2002.
- [6] A. Sahai et al. Automated sla monitoring for web services. *IFIP/IEEE 13th International Workshop on Distributed Systems: Operations and Management (DSOM 2002)*, page 28, 2002.
- [7] P. Flegkas et al. Design and implementation of a policy-based resource management architecture. In *Proceedings of the VIII IFIP/IEEE IM conference on network management*, page 215, 2003.
- [8] M. Debusmann and A. Keller. Sla-driven management of distributed systems using the common information model. In *Proceedings of the VIII IFIP/IEEE IM conference on network management*, page 563, 2003.
- [9] W.B. Norton. The art of peering: The peering playbook. Technical report, Equinix.com, 2001.
- [10] W.B. Norton. Internet service providers and peering. Technical report, Equinix.com, 2001.
- [11] Y. Diao, J.L. Hellerstein, and S. Parekh. Optimizing quality of service using fuzzy control. *IFIP/IEEE 13th International Workshop on Distributed Systems: Operations and Management (DSOM 2002)*, page 42, 2002.
- [12] M. Burgess. Computer immunology. *Proceedings of the Twelfth Systems Administration Conference (LISA XII) (USENIX Association: Berkeley, CA)*, page 283, 1998.

## BIBLIOGRAPHY

---

- [13] M. Burgess. Two dimensional time-series for anomaly detection and regulation in adaptive systems. *IFIP/IEEE 13th International Workshop on Distributed Systems: Operations and Management (DSOM 2002)*, page 169, 2002.
- [14] Y. Diao et al. Generic on-line discovery of quantitative models for service level management. In *Proceedings of the VIII IFIP/IEEE IM conference on network management*, page 158, 2003.
- [15] Arun V. Holden, editor. *Chaos*. Manchester University Press, 1986.
- [16] M. Burgess. *Analytical Network and System Administration — Managing Human-Computer Systems*. J. Wiley & Sons, Chichester, 2004.
- [17] R. Axelrod. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton Studies in Complexity, Princeton, 1997.
- [18] R. Axelrod. *The Evolution of Co-operation*. Penguin Books, 1990 (1984).
- [19] D.K. Arrowsmith and C.M. Place. *An Introduction to Dynamical Systems*. Cambridge University Press, 1990.
- [20] T.M. Jonassen. On the concept of hyperbolicity. the yerevan and japan talks. Technical Report 21, Oslo University College Report Series, 2002. ISBN 82-579-4155-7.