

On the Stability of Adaptive Service Level Agreements

Kyrre Begnum, Mark Burgess, Tore M. Jonassen and Siri Fagernes
Oslo University College

Abstract— We consider some implications of non-linear feedback, due to policy combinatorics, on policy-based management of networked services. We pay special attention to the case where the monitoring of certain aspects of Service Level Agreements is used to alter future policy dynamically, according to a control feedback scheme. Using two simple models, we show that non-linear policies are generally unstable to service provision, i.e. provide no reliable service levels (QoS). Hence we conclude that automated control by policy-rule combinatorics can damage quality of service goals.

Keywords— **Service level agreements, adaptive policy, feedback, dynamical systems.**

I. INTRODUCTION

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.

Any service provider's policy is subject to environmental uncertainties and, for that reason, the promise can never be more than an expectation or hopeful prediction of average service levels[3]. To verify that a promise has been kept, the client must monitor the service level. Service Level Agreements are often complex. They combine several measured values into composite metrics that are then compared to a template[4], [5], [6], [7], [8]. This has important implications for the stability of policy relationships that we explore in this paper.

In more complex networked collaborations, hosts (nodes) often provide services for one another mutually, in clusters. For example, network service providers undergo peering agreements to route and carry traffic for one another without explicit payment. These are bilateral agreements[9], [10]. If one provider carries a greater load than the other for a significant time, it might consider the agreement to be inequitable.

There is a *voluntary cooperation* implicit in a peering arrangement. Each party in the agreement is risking its resources for the sake of the other, and undertakes that risk for the potential benefits it will receive in return. Since peering parties often have no control over one another's behaviour, other than by the threat of reprisal, it is natural for cooperating peers to monitor one another's behaviour to verify their compliance with the Service Level Agreement.

The concept of Quality of Service attaches to fixed rate service provision. It describes an umbrella of techniques

which attempts to address and reduce the uncertainties in data rates so that service guarantees can be made. Of course, absolute guarantees are impossible in the short term, but average guarantees over longer times can be arranged as long as each agent has some form of control over its resources. The dilemma in voluntary agreements is that cooperating parties do not control each other's resources. Consequently, all they can do to maintain fairness is to adapt their policy towards one another based on monitored behaviour. This is a form of feedback regulation.

Feedback regulation of policy is an idea that has received more attention recently, due to interest in autonomous system operation; it has been discussed in a variety of circumstances[11], [12], [13], [14] for different service types ranging from web services to configuration management. Feedback regulation occurs regardless of whether it is carried out autonomously by machines or by humans, whenever there is a dynamical interplay between monitoring and SLA.

In a dynamical environment, such as in ad hoc networks and pervasive computing settings, resources are both more limited and more transitory than in fixed infrastructure provision. A service provider might be nothing more than a simple hand-held device with a varying list of customers, meaning that there is little room for over-provision and capacity reservation. In such a dynamical environment one might be tempted to give up on the idea of quality of service altogether, since there are many uncertainties. However, devices can monitor, remember and react to failures to live up to SLAs, using adaptive policy between collaborating parties. The question then becomes: how far can we go to secure fair and stable promises between collaborating parties[4], [15]?

In this paper we consider one aspect of adaptive service level agreements, namely adaptive behaviour based on proportional feedback metrics of peer compliance; i.e. we study the implications of basing service levels on past performance of peers with respect to their promised agreements. To do this, we employ a toy model which abstracts the essential elements of the interactions for proportional metrics. Further, we allow adjustments to agreed values based on composite measurements, using logical AND and OR combinations. Such methods have been suggested in a related setting in ref. [4]. The idea of voluntary cooperation thus enters, as understood from cooperative game theory, and we use well-established techniques from the study of dynamical systems to derive our conclusions.

II. VOLUNTARY COOPERATION AND POLICY FEEDBACK

To study the complexity that can emerge from the interaction of different policies, we begin by considering a simple model based on two interacting peers. This scenario is already sufficiently complex to fill the present paper and it is a stepping stone to more complex models. If two hosts cannot coordinate their policies, it is highly unlikely that increasing the number will help.

The framework we choose for examining the problem is that of iterative dynamics[16]. Recently several authors have introduced classical control theory into the field of network management[17]. This embodies a range of techniques for linear feedback modelling that parallel iterative dynamics to a certain degree. Our motivation with using iterative dynamics is three-fold: first, control theory is at its best when feedback is linear, so its efficacy is unclear in this situation; second, there is already a large body of general knowledge about iterative dynamics which we can simply use by mapping our situation onto that framework; finally, control theory is fundamentally an imposed regulation of state—it is not about voluntary cooperation. For that reason its *raison d'être* is at odds with the idea of service level agreements and we choose a passive form of analysis which is more in the manner of decision theory than control theory.

Why would one consider non-linear feedback? The combination of several dynamical variables using AND (multiplication) and OR (addition) leads to at least quadratic order in control variables. 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:

$$S \text{ depends_on } (SS_1 \text{ AND } SS_2) \quad (1)$$

as well as more general rules like:

$$QoS_A(S) = f(QoS_B(SS_1) \text{ AND } QoS_B(SS_2)). \quad (2)$$

They base 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 [16]). It is therefore important to understand how such behaviour emerges in dynamically regulated policies and what can be done to secure stable operation[18].

We use a continuum model, since time-average metrics are continua over long enough time-scales[3], and define a valuation of a service quality observation as follows:

$$QoS \rightarrow \begin{cases} = 0 & \text{no service given} \\ \propto q(t) & \text{performance is } q(t) \text{ in } [t, t+dt]. \end{cases}$$

Based on such a valuation, we observe that combinatoric valuation must be at least quadratic in q :

Lemma 1: Let q_A and q_B be two proportional observations of service quality. Any valuation q_{AB} , which represents $q(A \text{ AND } B)$ must be quadratic in q_A .

The proof of this is straightforward.

Proof: For q_{AB} to be a valuation it must be zero when q_A and q_B are zero, and proportional to both of these. In other words, we must have

$$q_{AB} \propto q_A q_B f(q_A, q_B), \quad (3)$$

where $f(\cdot)$ is some function of q_A, q_B . However, in order for q_{AB} to be strictly proportional to the service level, we require $f = \text{constant}$. ■

This conclusion mirrors the fact that, in order to combine independent probabilities, one uses the formulae[19]:

$$\begin{aligned} \text{AND } (p_i) &= \prod_i p_i \\ \text{OR } (p_i) &= 1 - \prod_i (1 - p_i) \end{aligned} \quad (4)$$

An inclusive OR requires a low degree of non-linearity, but here also a small degree of non-linearity is present, depending on the policy. In this paper, we do not combine probabilities but rather measured ‘expectation’ values which are compiled from statistical histories of a system. These expectations combine like probabilities due to the stochastic nature of the values and the theorem above. The statistical independence can be assumed, in spite of the iterative process, because the variables are combined at different times and hence they belong to different measures.

Note that it does not matter exactly what type of measurement is used in the SLA: whether it is an average response time or a percentile of recent statistics. All we require is that the measure be a proportional measure. If the measure is not proportional, the non-linearity does not disappear but rather it becomes more complicated.

To avoid unnecessary complications then, we take an idealized view of the SLAs and abstract the problem to generalized rules at quadratic order. We then attempt to answer a single basic question: whether or not a reactive policy, $q \rightarrow f(q \text{ AND/OR } q')$, based on observation of two data items, can lead to stable behaviour, and whether the criterion of stability places any restrictions on the allowed policy used by the nodes. This is then the simplest non-trivial model we can investigate in the spirit of the problem.

III. STABILITY

Our basic viewpoint is that stability to small perturbations in $q(t)$ is an important criterion for any policy that adapts dynamically, as any realistic system model must include noise from environmental conditions which will be testing this stability. Stability can mean many things[20].

By stability we mean the following: let $q_1(t)$ and $q_2(t)$ be two observable quality levels between parties 1 and 2, and let I be an iterative mapping from one set of levels to the next:

$$I : \langle q_1(t), q_2(t) \rangle \rightarrow \langle q'_1(t'), q'_2(t') \rangle \quad (5)$$

This mapping is said to be stable if it has at least one non-trivial *attractive fixed point* $\langle p_1, p_2 \rangle$, such that

$$I : \begin{cases} \langle p_1, p_2 \rangle \rightarrow \langle p_1, p_2 \rangle \\ \langle p_1 + \delta q_1, p_2 + \delta q_2 \rangle \rightarrow \langle p_1, p_2 \rangle \end{cases} \quad (6)$$

Indeed, we posit that stability is more important than maintaining a constant maximum envelope for service. If one does not have a basic stability over long times, the services will not be usable. Environment means that there is a constant fluctuation around the deterministically calculated values, so we must investigate whether such fluctuations could grow out of control. Hence, we are interested in policies that are stable to small fluctuations.

Studies in cooperative behaviour show that, from a payoff viewpoint, predictable relationships must be built up over time, and only persist if customers are regular patrons of a service provider[21]. Customer trust is based on two ingredients: that customers and providers form relatively long lasting relationships, and that one punishes misdeeds by some kind of tit-for-tat response[22].

The policy of tit-for-tat implies its commonplace meaning: to simply do against your opponents what they did to you on the previous encounter. The strength in this strategy lies in that it punishes bad behaviour, while it rewards cooperation. A modified version of tit-for-tat, is the strategy tit-for-two-tats, which offers a certain amount of leniency by always being cooperative unless the opponent act selfishly twice in a row. This strategy has the advantage of not punishing occasional bad behaviour as bad as systematic bad behaviour. These strategies have been made famous by the ubiquitous Prisoner's Dilemma model of cooperation[22].

IV. TWO REACTIVE POLICY MODELS

Our model is an abstraction, based on SLA policies which combine observations of previous compliance with SLA conditions, thus its conclusions are implementation independent. In this framework, we can identify interesting cases which play the roles of game-theoretical strategies for cooperative behaviour[21].

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 system is comprised of two parties, which promise one another a service level for a fixed time interval Δt . 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 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 $P(\cdot)$, with some policy parameter α , of this value:

$$q(t+1) = P^{(1)}(q(t); \alpha). \quad (7)$$

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:

$$q(t+1) = P^{(2)}(q(t), q(t-1), \dots; \alpha_1, \alpha_2, \dots), \quad (8)$$

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:

$$q = \frac{s}{S}, \quad (9)$$

where s is the actual service level metric, provided over the time interval Δt and S is some constant scale that is unimportant to the discussion. For example, s 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:

$$s = \langle s \rangle + \delta s, \quad (10)$$

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

$$\langle s \rangle(t) = \frac{1}{\Delta t} \int_t^{t+\Delta t} s(x) dx \quad (11)$$

and δs is a random variable (e.g. Gaussian noise). Thus:

$$q(t) = \langle q \rangle(t) + \delta q. \quad (12)$$

Our promised service levels must be associated with the *average* or expected service level valuations, since we cannot promise more than average behaviour[3]. Note that this does not mean that only average performance can be used in an SLA. Other proportional measures, such as percentiles, can also be used, but one still seeks *average stability* of these so this does not alter the argument significantly.




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

A. Model 1: observe and react

Consider the first of our models (see fig 1); the two SLAs are:

$$\begin{aligned} \langle q_1 \rangle(t+1) &= \alpha_2 q_2(t) + \beta_2 q_2(t)q_2(t-1) \\ \langle q_2 \rangle(t+1) &= \alpha_1 q_1(t) + \beta_1 q_1(t)q_1(t-1) \end{aligned} \quad (13)$$

With the random perturbation expansion, this becomes:

$$\begin{aligned} q_1(t) &= \langle q_1 \rangle(t) + \delta q_1 \\ q_2(t) &= \langle q_2 \rangle(t) + \delta q_2, \end{aligned} \quad (14)$$

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:

$$\langle q_1 \rangle(t+1) = \underbrace{\alpha_2 q_2(t)}_{\text{tit-for-tat}} + \underbrace{\beta_2 q_2(t)q_2(t-1)}_{\text{tit-for-2tat}} \quad (15)$$

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 heuristic behaviour of this policy is coarsely summarized in the following table for peer 1, explained below:

SetPolicy ₁ ($t+1$)	High	Medium	Medium	Low
Measure ₂ (t)	High	High	Low	Low
Measure ₂ ($t-1$)	High	Low	High	Low

This table illustrates, in heuristic terms, the magnitude of payoff (or service quality) peer 1 chooses to deliver on the next round to peer 2, given the level it received in the previous two rounds. Three levels only (low, medium and high) are used; these are relative measures, since the absolute values depend on the policy constants α, β . This gives one a feel for the policy that is used by the agent when it measures characteristic service levels from its peer.

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, peer 1 rewards it with a high level of service, if not it retaliates by promising less, by




Fig. 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 oneself.

degrees, 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** $q(t-1)$ (NOR). When q is large $1-q$ is small, thus it represents the complement of q .

B. 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 Q for the two SLAs (see fig 2):

$$\begin{aligned} \langle Q_1 \rangle(t+1) &= \alpha_2 Q_2(t) + \beta_2 Q_2(t)(1 - Q_1(t-1)) \\ \langle Q_2 \rangle(t+1) &= \alpha_1 Q_1(t) + \beta_1 Q_1(t)(1 - Q_2(t-1)) \end{aligned} \quad (16)$$

The interpretation of the terms is as follows:

$$\langle Q_1 \rangle(t+1) = \underbrace{\alpha_2 Q_2(t)}_{\text{tit-for-tat}} + \underbrace{\beta_2 Q_2(t)(1 - Q_1(t-1))}_{\text{conscience}} \quad (17)$$

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 ₁ ($t+1$)	Low	Good	Medium	Medium?
Measure ₂ (t)	Low	Good	Good	Medium?
Measure ₁ ($t-1$)	High	Bad	Good	Medium?

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.

V. 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:

$$\begin{aligned} q_1(t+1) &= \alpha q_2(t) + \beta q_2(t)q_2(t-1) \\ q_2(t+1) &= \alpha q_1(t) + \beta q_1(t)q_1(t-1) \end{aligned} \quad (18)$$

Model 2 becomes:

$$\begin{aligned} Q_1(t+1) &= \alpha Q_2(t) + \beta Q_2(t)Q_1(t-1) \\ Q_2(t+1) &= \alpha Q_1(t) + \beta Q_1(t)Q_2(t-1) \end{aligned} \quad (19)$$

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 omitted; it is not of general interest beyond this investigation so we refer readers to the literature. 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

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t w_t \\ y_{t+1} &= \alpha x_t + \beta x_t z_t \end{aligned} \quad (20)$$

By adding the equations for z and w we obtain a two parameter first order system (α, β) . 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

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t w_t \\ y_{t+1} &= \alpha x_t + \beta x_t z_t \\ z_{t+1} &= x_t \\ w_{t+1} &= y_t \end{aligned} \quad (21)$$

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

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

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t z_t \\ y_{t+1} &= \alpha x_t + \beta x_t w_t \end{aligned} \quad (22)$$

and hence a map given by

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

The induced dynamical system is then given by

$$\begin{aligned} x_{t+1} &= \alpha y_t + \beta y_t z_t \\ y_{t+1} &= \alpha x_t + \beta x_t w_t \\ z_{t+1} &= x_t \\ w_{t+1} &= y_t \end{aligned} \quad (23)$$

This expression is called the first order representation of Model 2, and G_ζ 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

$$\begin{aligned} F_\zeta(x, y, z, w) &= (x, y, z, w) \\ G_\zeta(x, y, z, w) &= (x, y, z, w) \end{aligned} \quad (24)$$

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

$$\begin{aligned} x &= \alpha y + \beta y w \\ y &= \alpha x + \beta x z \\ z &= x \\ w &= y \end{aligned} \quad (25)$$

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

$$\begin{aligned} x &= \alpha y + \beta y^2 \\ y &= \alpha x + \beta x^2 \end{aligned} \quad (26)$$

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.

VI. RESULTS

The properties of the *fixed points* of a discrete dynamical system are key to the overall behaviour of voluntarily co-operative systems. This section provides a brief technical

summary of results, so that experts in dynamical systems will understand the basis for our conclusions. We suggest that non-experts skim this section only.

The behaviour of the fixed points of discrete deterministic dynamical systems given by Model 1 and Model 2 can be analyzed by simple techniques found in any standard textbook on dynamical systems, e.g. see ref. [23].

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:

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

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 α , the first order tit-for-tat strategy which drives this; β , 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 4×4 Jacobian matrices DF_η 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 β and depends only on α , thus it is the tit-for-tat component of policy that leads to stability here, not the **AND**-ed combinatorics.

The fixed points p_3 and p_4 exist only for the set of parameters $\alpha \in (-\infty, -1] \cup [3, \infty)$. We have given the graph of the norm of these eigenvalues in figure 3 below. For discussing service levels, we are only interested in positive values, so again we ignore negative roots.

Model 2 has only two fixed points, the trivial fixed point $p_1 = (0, 0)$ and the non-trivial fixed point

$$p_2 = \left(\frac{1 - \alpha_1\alpha_2}{\beta_1 + \alpha_1\beta_2}, \frac{1 - \alpha_1\alpha_2}{\beta_2 + \alpha_2\beta_1} \right)$$

provided that the denominators in the expressions above are different from zero. The same notation as in Model 1




Fig. 3. The norm of the eigenvalues of the fixed point p_3 as a function of α whenever it exist. As seen from the graph, it is not an attractive fixed point. α is along the horizontal axis. The stability criterion 1 is also plotted.

applies here too, as $z = x$ and $w = y$. We note that this has the same basic behaviour as Model 1. The stable service level is now determined by the geometric mean of the first order policies $\alpha_1\alpha_2$, and is regulated by the second order term β . The tit-for-tat terms α_i do not have to be so small now, since it is their product that drives the agreed service levels. We can again assume that $\alpha_1 = \alpha_2, \beta_1 = \beta_2$ for simplicity, though since the agents are autonomous, there is no reason to demand this.

Again, the topological type of the fixed points is given by the norm of the eigenvalues of the Jacobian DG_ζ evaluated at the points p_1 and p_2 . A simple computation shows that the eigenvalues at the trivial fixed point $p_1 = (0, 0)$ are given by 0 (multiplicity 2) and $\pm\alpha$. At the fixed point p_2 we easily compute them in terms of α_i and $\beta_i, i = 1, 2$, but the expressions are, again, rather large. For $\alpha = \alpha_1 = \alpha_2$ and $\beta = \beta_1 = \beta_2$ the eigenvalues are given by:

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

The dynamics of Model 1 and Model 2 may now be simulated using the `NestList` function in *Mathematica* applied to the mappings F_η and G_ζ . We cannot represent all of the dimensions without projecting them, but we shall attempt an illustration.

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 Model 2, the behaviour is more sophisticated and the behaviour is shown in fig. 3.¹ For the deterministic (non-noisy) case, this model exhibits attractive, periodic orbits,

¹Several types of ‘phase’ or bifurcations may be observed, and the system has stable periodic orbits of high periods. In the symmetric case, we observe a double Hopf bifurcation resulting in an invariant two-dimensional torus in $\mathbf{R}^4$. The projection of an attracting orbit on this torus to $\mathbf{R}^2$ is shown in figure 3 for randomly chosen initial values in the unit square.




Fig. 4. An attracting orbit on an invariant manifold, projected onto the two first coordinates for a symmetric version of Model 2. The initial point is chosen at random. This shows a region of stable behaviours in which service levels will oscillate within a bounded region.




Fig. 5. The effect of noise of figure 4. We have used a random initial value in the unit 4-cube. The form of the figure is quite sensitive to noise, which is a powerful illustration of our contention that, unless an interactive agreement is stable to small perturbations, it will be useless in practice, since it might as well correspond to no agreement.

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 $\alpha > 1, \beta < 0$ 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 6, 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 4, the corresponding projection of the random attractor is




Fig. 6. The orbit of stochastic version of the map in figure 4, where white noise has been added to each of the two non-linear components of the symmetric version of Model 2.

given in figure 5. 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. [24].

VII. SIMULATION

To test the validity of our idealized model predictions, we simulate the two models with a numerical iteration, for a few test values of the parameter sets. We add some real-world assumptions to the mathematical models, namely that the service level falls within fixed bounds: it can never fall below zero and it can never exceed a maximum ceiling.

In addition, if we see that the service level is falling below a minimum threshold, or exceeding a maximum threshold, we can try to regulate the rates by adjusting them accordingly, up or down, to try to prevent them from running away. This is what an administrator would do in a genuine exchange. To simulate this, we add a simplistic rule: if value exceeds threshold, subtract 0.1 from q .

The results of the simulations are shown in figures 7 and 8. Both graphs consist of two pairs of two lines; in both cases, the lines belonging to a given model are so close together that they are almost indistinguishable. The steepest gradients belong to model 1: an essentially square wave oscillation over time. Slower gradients in model 2 show more easily regulatable behaviour. The values we choose are as follows, for both models:

- Model 1 has a fixed point for: $\alpha_1 = \alpha_2 = 0.5$ and $\beta_1 = \beta_2 = 1$.
- Model 1 has runaway growth for: $\alpha_1 = \alpha_2 = 0.5$ and $\beta_1 = 1.001, \beta_2 = 1.00001$.
- Model 1 has runaway decay for: $\alpha_1 = \alpha_2 = 0.4$ and $\beta_1 = 1.001, \beta_2 = 1.00001$.

We start the variables q_i with initial values of 0.5.




Fig. 7. The unregulated simulation shows that the two curves from model 1 run away exponentially, while model 2 remains relatively stable. After just four iterations, model 1 diverges significantly from a stable value.

Model 1 is fundamentally unstable. For different policy parameters, it very quickly either grows to saturation or falls off to zero. If we apply floor and ceiling regulation, and manual adjustment of the values, we end up with long period oscillations (essentially a square wave) that are artificial. In other words, it oscillates between all or nothing. Clearly this model is characterized by basic instabilities. We note that model 1 has a fixed point, but even the smallest amount of noise will send it off into runaway behaviour.

Model 2 also has runaway behaviour, but the real-world ceiling is sufficient for it to regulate it, causing oscillations without administrative intervention. It thus shows more long-term *average* stability. It does not run away as readily as model 1. However, there are still wild fluctuations in the values, and the frequency of these oscillations is some ten times greater than that of model 1; what is important, however, is that these recover without intervention and the growth is slower.

It is worth noting, however, that since the two parties' dynamical policies track one another very closely for similar initial values, they gain little benefit from monitoring one another's behaviour, at such close quarters. The service level policies would be better regulated over much longer time-scales, with increased trust, for greater stability. Indeed, it is a general rule that increased trust leads to more stable cooperative behaviour [21], [22].

VIII. CONCLUSIONS

We have studied a simple algebra of one important aspect service level policies: how average expectations and service demands should be set, if they are based on observation of previous compliance. 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




Fig. 8. The regulated simulations both show oscillations over long times. Two patterns can be seen: a square wave (model 1) and a sawtooth pattern (model 2).

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. [20] 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.

REFERENCES

- [1] M. Sloman, “Policy driven management for distributed systems,” *Journal of Network and Systems Management*, vol. 2, pp. 333, 1994.
- [2] D. Verma et al., “Policy based sla management in enterprise networks,” in *Policy Workshop 2001*. 2001, Springer Verlag.
- [3] M. Burgess, “On the theory of system administration,” *Science of Computer Programming*, vol. 49, pp. 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)*, p. 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)*, p. 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)*, p. 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*, 2003, p. 215.
- [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*, 2003, p. 563.
- [9] W.B. Norton, “The art of peering: The peering playbook,” Tech. Rep., Equinix.com, 2001.
- [10] W.B. Norton, “Internet service providers and peering,” Tech. Rep., 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)*, p. 42, 2002.
- [12] M. Burgess, “Computer immunology,” *Proceedings of the Twelfth Systems Administration Conference (LISA XII) (USENIX Association: Berkeley, CA)*, p. 283, 1998.
- [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)*, vol. LNCS 2506, pp. 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*, 2003, p. 158.
- [15] Mark Burgess, “An approach to policy based on autonomy and voluntary cooperation,” *Lecture Notes on Computer Science*, vol. 3775, pp. 97–108, 2005.
- [16] Arun V. Holden, Ed., *Chaos*, Manchester University Press, 1986.
- [17] J.L. Hellerstein, Y. Diao, S. Parekh, and D.M. Tilbury, *Feedback Control of Computing Systems*, IEEE Press/Wiley Interscience, 2004.
- [18] M. Suneel, “Dynamical systems as logic gates.” <http://www.citebase.org/cgi-bin/citations?id=oai:arxiv.org:nlin/0511070>,” 2005.
- [19] A. Høyland and M. Rausand, *System Reliability Theory: Models and Statistical Methods*, J. Wiley & Sons, New York, 1994.
- [20] M. Burgess, *Analytical Network and System Administration — Managing Human-Computer Systems*, J. Wiley & Sons, Chichester, 2004.
- [21] R. Axelrod, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton Studies in Complexity, Princeton, 1997.
- [22] R. Axelrod, *The Evolution of Co-operation*, Penguin Books, 1990 (1984).
- [23] D.K. Arrowsmith and C.M. Place, *An Introduction to Dynamical Systems*, Cambridge University Press, 1990.
- [24] T.M. Jonassen, “On the concept of hyperbolicity. the yerevan and japan talks,” Tech. Rep. 21, Oslo University College Report Series, 2002, ISBN 82-579-4155-7.