

An HMM-based reputation model

Ehab Elsalamouny, Vladimiro Sassone

► **To cite this version:**

Ehab Elsalamouny, Vladimiro Sassone. An HMM-based reputation model. *Advances in Security of Information and Communication Networks*, Sep 2013, Cairo, Egypt. Springer Berlin Heidelberg, 381, pp.111-121, 2013, <10.1007/978-3-642-40597-6_9>. <hal-00831401>

HAL Id: hal-00831401

<https://hal.inria.fr/hal-00831401>

Submitted on 6 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An HMM-based reputation model

Ehab ElSalamouny^{1,2} and Vladimiro Sassone³

¹ INRIA, France

² Faculty of Computer and Information Science, Suez Canal University, Egypt

³ ECS, University of Southampton, UK

Abstract. In modern global networks, principals usually have incomplete information about each other. Therefore trust and reputation frameworks have been recently adopted to maximise the security level by basing decision making on estimated trust values for network peers. Existing models for trust and reputation have ignored dynamic behaviours, or introduced ad hoc solutions. In this paper, we introduce the HMM-based reputation model for network principals, where the dynamic behaviour of each one is represented by a hidden Markov model (HMM). We describe the elements of this novel reputation model. In particular we detail the representation of reputation reports. We also describe a mixing scheme that efficiently approximates the behaviour of a trustee given multiple reports about it from different sources.

1 Introduction

In modern global networks, the principals have incomplete information about each other. Therefore it is a challenging problem for a principal to take decisions regarding interactions with others. One approach that has recently adopted is to base these decisions upon a level of ‘trust’ associated with each network peer. We consider specifically the *probabilistic trust*, where the trust in a peer (trustee) te is expressed as a probability distribution over the potential outcomes of an interaction with te .

Many systems (e.g. [13, 4, 18, 5]) have adopted the so-called *Beta model* [11], where the behaviour of a trustee te is modelled by a fixed rating θ approximating the probability that an interaction with te yields ‘success’. This probability is learnt from past interactions with the trustee te . One limitation of the Beta-based systems is that they assume a fixed probabilistic behaviour for each principal; that is for each principal, there exists a fixed probability distribution over possible outcomes of its interactions. This assumption is indeed not realistic in practice. For example the behaviour of a principal can be significantly different when it is corrupted by an attacker.

As a step forward to handling dynamic behaviours, several papers, e.g. [11, 4, 19], adopted the ‘decay’ principle which was first introduced in [11]. It aims at capturing the recent behaviour of the trustee by letting older observations ‘decay’, so as to give higher weight to recent interactions over older ones. However, we showed in [7] that the decay principle is only useful when the trustee’s behaviour is highly *stable*, i.e., when the probability distribution over the observables is unlikely to change.

For coping with this limitation, we introduced in [8] the foundations of a novel HMM-based trust model to evaluate trust in trustees exhibiting dynamic behaviours. A

trustee is characterised by a set of (behavioural) states, each associated with a probability distribution over observable outcomes of interactions. It proceeds by performing (unobservable) probabilistic state transitions, which in turn determine changes in the statistical properties of the (observable) outcomes of interactions. These assumptions are met exactly by representing the trustee with a finite-state hidden Markov model (HMM) [17], called the ‘real’ model. In particular, the trustee’s state transitions are hidden, and trusters observe only the outcomes of their interactions with it.

For evaluating the trust, past observations are used to learn the trustee’s behaviour, and then predict the outcomes of future interactions. The key information for that is the *real* predictive distribution, i.e., the probability of each potential outcome in the next interaction between a truster and a trustee te , given the outcomes of past interactions. Yet, since the real model λ for te is unknown, the truster can only estimate the real predictive distribution. In our approach this is done using the *Baum-Welch* algorithm [17] that yields an ‘approximation’ η of λ , and then use η to evaluate the so-called *estimated* predictive distribution, which ultimately defines the truster’s trust in te .

In many cases, the sequence of direct observations available to the truster is not sufficiently long to learn the behaviour of the trustee with a satisfactory accuracy. To handle this shortage of information, the *reputation* information is used in the learning process. This information is seen as *reports* - about the trustee - given by other principals called *reputation sources*. These reports enrich the truster’s knowledge about the trustee and therefore enhance the approximation of λ . This also implies a better estimate of the predictive distribution. Thus the notion of reputation raises two main questions:

- Which representation is appropriate for a reputation report ?
- How are reports, from different sources, utilised to enhance estimations ?

Clearly, the answers depend on the assumptions made about the real model of the trustee’s behaviour. For example, in systems based on the Beta reputation model and its extensions (e.g., [11, 10, 14]), a reputation report consists of the count of each outcome experienced by its source in its interactions with the trustee. Multiple reputation reports from different sources are therefore mixed by adding up the corresponding counts of outcomes. This is not so easy in our case, where we must take into account that observations seen by different sources could correspond to different (hidden) states of the trustee, and can not therefore be summed together. In fact, this is a major technical challenge we face, one which demands a new approach.

In this paper we introduce a reputation model that matches the ‘dynamic’ nature of the trustee’s behaviour. This model answers the above two questions when the behaviour is represented by an HMM, and therefore completes the basic HMM-based trust model of [8] with a reputation handling mechanism. We also point to experimental evidence, through simulations, in support of our model. To the best of our knowledge, this is the first trust-and-reputation model for multi-state, dynamic systems. Thus, it provides the first complete answer to the research challenge launched in [14].

Structure of the paper. The next section concisely introduces hidden Markov models. In Section 3 we recall the basic model of HMM-based trust, whilst in Section 4 we detail the elements of our HMM-based reputation model. Section 5 sketches an experimental analysis of our reputation model against some of its predecessors. Finally, we

conclude our results in Section 6. For reason of space we have omitted the proofs from the body of the paper. The interested reader can find them in the appendix.

2 Hidden Markov Models (HMMs)

A *Hidden Markov Model (HMM)* [1] is a probabilistic model essentially based on a notion of system state. Underlying any HMM there is a Markov chain modelling (probabilistically) system's transitions between a set of states. Each state in this chain is associated with a particular probability distribution over the set of possible outcomes (observables). The output of an HMM is a sequence of outcomes where each outcome is sampled according to the probability distribution of the underlying state. In the following, we denote the state of the HMM and the outcome at time t by q_t and o_t respectively.

Definition 1 (hidden Markov model). A (discrete) *hidden Markov model (HMM)* is a tuple $\lambda = (S, V, \pi, \mathbf{A}, \mathbf{B})$ where $S = \{1, 2, \dots, N\}$ is a finite set of *states*; $V = \{z_1, z_2, \dots, z_K\}$ is a finite set of possible *observables*; π is a distribution on S , the *initial distribution*; $\mathbf{A} : S \times S \rightarrow [0, 1]$ is the *state transition matrix*, with $A_{ij} = P(q_{t+1} = j \mid q_t = i)$ and $\sum_{j \in S} A_{ij} = 1$; and $\mathbf{B} : S \times V \rightarrow [0, 1]$ is the *emission matrix*, with $B_i(z_k) = P(o_t = z_k \mid q_t = i)$, $\sum_{z_k \in V} B_i(z_k) = 1$.

As an example, Figure 1 shows a two-state HMM with the observation set $\{\mathbf{s}, \mathbf{f}\}$, where transitions $\mathbf{1} \mapsto \mathbf{2}$ and $\mathbf{2} \mapsto \mathbf{1}$ have probabilities 0.1, 0.12 respectively. The probabilities of self-transitions $\mathbf{1} \mapsto \mathbf{1}$ and $\mathbf{2} \mapsto \mathbf{2}$ are therefore 0.9, 0.88 respectively. The other parameters π, \mathbf{B} are shown.




Fig. 1. Example Hidden Markov Model.

HMMs provide the computational trust community with several obvious advantages: they are widely used in scientific applications, and come equipped with efficient algorithms for computing the probabilities of events and for parameter estimation [17]. It is worth noticing that HMMs are a generalisation of the Beta model [11] to a multi-state model. In fact, representing the behaviour of a trustee te by an HMM λ provides, for each state j of te , a distribution \mathbf{B}_j over candidate observables V .

According to Definition 1 of HMM, the probability of a sequence of outcomes $h = o_1 o_2 \dots o_T$ given an HMM λ is given by the following equation.

$$P(h \mid \lambda) = \sum_{q_1, \dots, q_T \in S} \pi(q_1) \cdot B_{q_1}(o_1) \cdot A_{q_1 q_2} \cdot B_{q_2}(o_2) \cdots A_{q_{T-1} q_T} \cdot B_{q_T}(o_T). \quad (1)$$

The above probability is evaluated efficiently by the *forward-backward* algorithm ([17]), which evaluates the above probability inductively on the time t where $1 \leq t \leq T$.

In our work we assume that utilised HMMs are *ergodic*. This corresponds to demanding that the Markov chain underlying an HMM is irreducible and aperiodic (c.f. [9, 15, 3]). An HMM is *irreducible* if each state is reachable (with non-zero probability) from every other one. It is *aperiodic* if it has at least one aperiodic state. A state i is aperiodic if it does not recur with a cyclic period, that is if the greatest common divisor of times $t > 0$ such that $P(q_{1+t} = i \mid q_1 = i) > 0$ is 1. To guarantee aperiodicity it is sufficient that one state has a self-transition with non-zero probability. Such conditions are not overly restrictive for practical systems and typical applications.

2.1 Baum-Welch algorithm

Given a fixed set S of states and a fixed set V of observables, the Baum-Welch (BW) algorithm [2, 17] iteratively finds an HMM η which maximises the probability $P(h \mid \eta)$ of a given sequence h . This iterative algorithm starts with an initial HMM η' having parameters π' , A' , and B' . Then at each iteration, the a priori HMM η' is refined to obtain a posteriori HMM η with parameters π , A , and B . These parameters are evaluated by mathematical equations which we here summarise only informally by means of Equations (2), (3), and (4), respectively.

$$\pi_i = \text{the probability of being in state } i \text{ at time } (t = 1). \quad (2)$$

$$A_{ij} = \frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}. \quad (3)$$

$$B_i(z_k) = \frac{\text{expected number of times in state } i \text{ and observing symbol } z_k}{\text{expected number of times in state } i}. \quad (4)$$

In the above equations, the expected values are computed given the sequence h of observations, and the probability distributions defined by the a priori model η' . The resulting a posteriori model becomes the a priori one for the next iteration. The algorithm stops when the a priori and a posteriori models have the same parameters. More details about this algorithm can be found in [17]. One limitation of this algorithm is that it only converges to a local maxima for the probability function rather than the global one.

3 Hmm-based trust model

The *HMM-based trust model* [8] is based on the assumption that the behaviour of the trustee te is dynamic. This ‘unknown’ behaviour is modelled by an HMM λ , called the ‘real’ model of te . Since λ is generally unknown, each truster tr approximates it by a finite-states HMM η , which we call the ‘approximate’ model of te .

In this approximation (learning) process, the truster tr uses past outcomes of its direct interactions with te as follows. Given a sequence h of outcomes of direct interactions between tr and te , the truster tr applies the BW-algorithm [2, 17] to h . As described earlier, this algorithm yields an HMM that maximises the probability of h , and therefore defines the required approximate model η of te . We remark here that the size of η is fixed by the truster, and represents the approximation level of the trustee’s behaviour.

Using the in-hand approximate model η for te , the truster tr can estimate a probability distribution over possible outcomes of its next interaction with te . This distribution, is called an ‘estimated’ predictive distribution of te (from tr ’s point of view), and represents the trust of tr in te . Actually, this distribution is meant to be an ‘estimate’ for the ‘real’ predictive distribution which is determined by the real model λ of the trustee. The quality of this estimation is quantified by the difference between the ‘real’ and ‘estimated’ predictive distributions as follows.

Estimation error. The Kullback-Leibler (KL) divergence [12, 6] is an appropriate measure for the difference between distributions. It has an information-theoretic flavour, and is technically understood as a measure for the lost information when a probability distribution is approximated by another. In our case, we write the real and estimated predictive distributions as $P(\cdot)$ and $\mathcal{H}(\cdot)$ respectively. The KL-divergence from $P(\cdot)$ to $\mathcal{H}(\cdot)$ is called the ‘estimation error’ and is defined as follows.

$$D_{KL}(P(\cdot) \parallel \mathcal{H}(\cdot)) = \sum_{z \in V} P(z) \log \left(\frac{P(z)}{\mathcal{H}(z)} \right). \quad (5)$$

Note that the above estimation error is specific to a pair of distributions corresponding to a particular sequence of outcomes. Thus, we evaluate the quality of our trust model as the *average* over all possible sequences. This average is the expected value of above divergence, and therefore is referred to as the ‘expected estimation error’.

4 Hmm-reputation model

Now we describe our proposed reputation model, which enhances the trust evaluation using supplementary feedback reports about the trustee. This model consists of two main components: a formalism of reputation reports exchanged between the network peers; and a mixing scheme which uses multiple reputation reports about a trustee te to evaluate the trust in te . As this trust is an estimated predictive distribution, our goal is to design those components such that the expected estimation error, described in Section 3 is minimised.

We start by linking the expected estimation error (for a trustee te) to the sequences of outcomes observed by all reputation sources. Let h be a sequence of outcomes resulting from past interactions between a trustee te and a single reputation source. Let also λ be the unknown real model of te . It is shown in [8] that with any approximate model η , the expected estimation error converges (as the length T of h grows) to the following limit which we refer to as the *asymptotic estimation error*.

$$Error(\lambda, \mathcal{H}_\eta) = C(\lambda) - H(\lambda, \eta), \quad (6)$$

where $C(\lambda) = \lim_{T \rightarrow \infty} \mathbf{E}[\log P(o_{T+1} | q_T, \lambda)]$, $H(\lambda, \eta) = \lim_{T \rightarrow \infty} \mathbf{E}[\log P(o_{T+1} | h, \eta)]$. In the language of information theory, $-C(\lambda)$ is the expected *entropy* of the real predictive distribution determined by λ (the real model); and $-H(\lambda, \eta)$ is the expected *cross-entropy* between the real and estimated predictive distributions where the latter is

determined by both λ and η . By the asymptotic properties of ergodic HMMs [1, Theorem 3.2], the log probability of any T -length observation sequence h , generated by λ , is related to $H(\lambda, \eta)$ as follows.

$$(1/T) \log P(h | \eta) \xrightarrow{a.s.} H(\lambda, \eta), \quad (7)$$

that is, the left-hand term converges *almost surely* (with probability 1) to $H(\lambda, \eta)$ as $T \rightarrow \infty$. Now consider the interactions between a trustee te and a set of reputation sources $\mathcal{M} = \{1, 2, \dots, M\}$ until a certain time instant. For every reputation source $u \in \mathcal{M}$, let h^u be the sequence of outcomes observed by u . Let also T_u be the length of h^u . By Eq. (7), the average of the quantities $(1/T_u) \log P(h^u | \eta)$ over the elements of \mathcal{M} approximates $H(\lambda, \eta)$. Thus, minimising the asymptotic estimation error, expressed by (6), amounts to choosing the approximate model η that maximises such an average, i.e. choosing the approximate model η that satisfies the following equation.

$$\eta = \operatorname{argmax}_{\mathcal{R}_n} \mathcal{G}(h^1, h^2, \dots, h^M | \mathcal{R}_n), \quad (8)$$

where $\mathcal{G}(h^1, h^2, \dots, h^M | \mathcal{R}_n)$ is an objective function whose value depends on an n -state HMM \mathcal{R}_n . This objective function is defined as

$$\mathcal{G}(h^1, h^2, \dots, h^M | \mathcal{R}_n) = \sum_{u \in \mathcal{M}} (1/T_u) \log P(h^u | \mathcal{R}_n). \quad (9)$$

Maximising the above objective function requires full access to all the sequences h^u . However, it is not practical for principals to exchange their entire observed sequences because each one of these sequences gets longer over time. Therefore we alternatively maximise a tight lower bound for this function. For doing so, we assume that each reputation source u uses its own observation sequence h^u to learn an ‘a priori’ approximate HMM η^u for the trustee. In terms of the a priori HMMs and other variables, the following lemma provides the required lower bound for the objective function.

Lemma 1. *Let $\mathcal{M} = \{1, 2, \dots, M\}$ be a set of reputation sources. For all $u \in \mathcal{M}$, let h^u and $\eta^u = (S, V, \pi^u, \mathbf{A}^u, \mathbf{B}^u)$ be, respectively, the sequence of outcomes observed by u , and the corresponding a priori HMM. For any sequence q of states, let $P(q | h^u, \eta^u)$ denote the probability of q given h^u and η^u . Thus, it holds for every a posteriori HMM $\eta = (S, V, \pi, \mathbf{A}, \mathbf{B})$ that*

$$\mathcal{G}(h^1, h^2, \dots, h^M | \eta) \geq \sum_{u \in \mathcal{M}} (1/T_u) Q(\eta^u, h^u, \eta) + \sum_{u \in \mathcal{M}} (1/T_u) \mathcal{R}(\eta^u, h^u),$$

$$\text{where} \quad Q(\eta^u, h^u, \eta) = \sum_q P(q | h^u, \eta^u) \log P(h^u, q | \eta), \quad (10)$$

$$\mathcal{R}(\eta^u, h^u) = - \sum_q P(q | h^u, \eta^u) \log P(q | h^u, \eta^u); \quad (11)$$

and the equality holds when $\eta = \eta^1 = \eta^2 = \dots = \eta^M$.

Lemma 1 provides a lower bound for $\mathcal{G}(h^1, h^2, \dots, h^M | \eta)$ given any a posteriori model η . Note that this bound is tight in the sense that it is equal to the objective function when the a priori models η^u are all equal to the a posteriori model η . Thus, we set our objective in the following to compute the *optimal* a posteriori model η^* which we define as the one maximising the above lower bound. That is,

$$\eta^* = \operatorname{argmax}_{\eta} \left(\sum_{u \in \mathcal{M}} (1/T_u) \mathcal{Q}(\eta^u, h^u, \eta) + \sum_{u \in \mathcal{M}} (1/T_u) \mathcal{R}(\eta^u, h^u) \right). \quad (12)$$

We will show that a truster wanting to compute η^* needs to collect only certain statistics about every observation sequence h^u , rather than the entire sequence. For fixed sets of states S and observables V , let the sequence of outcomes observed by the reputation source u be $h^u = o_1^u o_2^u \dots o_{T_u}^u$, where o_t^u is the outcome at time t . Similarly, let $q^u = q_1^u q_2^u \dots q_{T_u}^u$ be the (hidden) sequence of states underlying the observed sequence h^u . The optimal a posteriori HMM η^* in (12) is computed by the following theorem.

Theorem 1. *Given a set \mathcal{M} of reputation sources, the parameters of the optimal a posteriori HMM $\eta^* = (S, V, \pi^*, A^*, B^*)$, are given by the following equations.*

$$\pi_i^* = \frac{\sum_{u \in \mathcal{M}} \frac{1}{T_u} P(q_1^u = i | h^u, \eta^u)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u}}, \quad (13)$$

$$A_{ij}^* = \frac{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=2}^{T_u} P(q_{t-1}^u = i, q_t^u = j | h^u, \eta^u)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=2}^{T_u} P(q_{t-1}^u = i | h^u, \eta^u)}, \quad (14)$$

$$B_i^*(z_k) = \frac{\sum_{u \in \mathcal{M}} \sum_{t=1, o_t^u = z_k}^{T_u} P(q_t^u = i | h^u, \eta^u)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=1}^{T_u} P(q_t^u = i | h^u, \eta^u)}. \quad (15)$$

In the context of the BW-algorithm [17, 16], $P(q_t = i | h, \eta)$, the probability of visiting state i at time t given an observation sequence h and an HMM η is denoted by the variable $\gamma_t(i)$. Also $P(q_{t-1} = i, q_t = j | h, \eta)$, the probability of visiting states i and j at times $t-1$ and t respectively is denoted by the variable $\xi_{t-1}(i, j)$. In the same manner, we use the variables $\gamma_t^u(i)$ and $\xi_{t-1}^u(i, j)$ to denote the probabilities $P(q_t^u = i | h^u, \eta^u)$ and $P(q_{t-1}^u = i, q_t^u = j | h^u, \eta^u)$, respectively. Using these variables, (13-15) can be written in shorter forms as follows.

$$\pi_i^* = \frac{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \gamma_1^u(i)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u}}, \quad A_{ij}^* = \frac{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=2}^{T_u} \xi_{t-1}^u(i, j)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=2}^{T_u} \gamma_{t-1}^u(i)}, \quad B_i^*(z_k) = \frac{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=1, o_t^u = z_k}^{T_u} \gamma_t^u(i)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u} \sum_{t=1}^{T_u} \gamma_t^u(i)}.$$

From these formulae we notice that the computation of η^* does not require full knowledge about the observation sequences h^u , where $u \in \mathcal{M}$, but only some statistical functions of these sequences computed ‘locally’ in individual reputation sources.

$$\begin{array}{ll}
\bar{\gamma}_1^u = [\bar{\gamma}_1^u(1) \bar{\gamma}_1^u(2) \dots \bar{\gamma}_1^u(N)], & \text{where } \bar{\gamma}_1^u(i) = \frac{1}{T_u} \gamma_1^u(i) \\
\bar{\gamma}_{T_u}^u = [\bar{\gamma}_{T_u}^u(1) \bar{\gamma}_{T_u}^u(2) \dots \bar{\gamma}_{T_u}^u(N)], & \text{where } \bar{\gamma}_{T_u}^u(i) = \frac{1}{T_u} \gamma_{T_u}^u(i) \\
\bar{\gamma}^u = [\bar{\gamma}^u(1) \bar{\gamma}^u(2) \dots \bar{\gamma}^u(N)], & \text{where } \bar{\gamma}^u(i) = \frac{1}{T_u} \sum_{t=1}^{T_u-1} \gamma_t^u(i) \\
\bar{\xi}^u = \begin{bmatrix} \bar{\xi}^u(1,1) & \bar{\xi}^u(1,2) & \dots & \bar{\xi}^u(1,N) \\ \bar{\xi}^u(2,1) & \bar{\xi}^u(2,2) & \dots & \bar{\xi}^u(2,N) \\ \vdots & \dots & \ddots & \vdots \\ \bar{\xi}^u(N,1) & \bar{\xi}^u(N,2) & \dots & \bar{\xi}^u(N,N) \end{bmatrix}, & \text{where } \bar{\xi}^u(i,j) = \frac{1}{T_u} \sum_{t=2}^{T_u} \xi_{t-1}^u(i,j) \\
\bar{\omega}^u = \begin{bmatrix} \bar{\omega}^u(1,1) & \bar{\omega}^u(1,2) & \dots & \bar{\omega}^u(1,K) \\ \bar{\omega}^u(2,1) & \bar{\omega}^u(2,2) & \dots & \bar{\omega}^u(2,K) \\ \vdots & \dots & \ddots & \vdots \\ \bar{\omega}^u(N,1) & \bar{\omega}^u(N,2) & \dots & \bar{\omega}^u(N,K) \end{bmatrix}, & \text{where } \bar{\omega}^u(i,k) = \frac{1}{T_u} \sum_{t=1, o_t^u=z_k}^{T_u} \gamma_t^u(i)
\end{array}$$

Fig. 2. The elements of an HMM-based reputation report

It is essential to ensure that the a priori HMMs η^u have the same set S of states (as required by Lemma 1 and Theorem 1). Therefore we define a parameter $\bar{\eta}$ for the reputation protocol to be an initial N -state HMM. The parameter $\bar{\eta}$ is shared by all principals and is regarded as the ‘default’ trustee’s behaviour which is refined through the learning process to η^u according to the sequence of outcomes h^u seen by u .

Thus, we describe the HMM based reputation model as follows. Using $\bar{\eta}$ as an initial HMM, each reputation source u applies the BW-algorithm to its own observations h^u about the trustee te . This learning process, performed by u , yields an HMM η^u approximating the behaviour of te (from u ’s point of view) and also the variables $\gamma_t^u(i)$, $\xi_t^u(i,j)$ for $1 \leq t \leq T_u$, and all $i, j \in \{1, 2, \dots, N\}$. In terms of these variables, every reputation source u constructs its reputation report about te as the tuple $(T_u, \bar{\gamma}_1^u, \bar{\gamma}_{T_u}^u, \bar{\gamma}^u, \bar{\xi}^u, \bar{\omega}^u)$. While T_u is clearly the length of h^u , each other element in this tuple (report) is basically a matrix defined in Figure 2.

Now we describe our reputation mixing scheme. Suppose that multiple reports about te , from a set \mathcal{M} of reputation sources, are available to a truster tr . Note that these reports include the one constructed by tr itself about te . Using the elements of these reports, tr computes the optimal a posteriori HMM η^* of te as follows

$$\pi_i^* = \frac{\sum_{u \in \mathcal{M}} \bar{\gamma}_1^u(i)}{\sum_{u \in \mathcal{M}} \frac{1}{T_u}}, \quad A_{ij}^* = \frac{\sum_{u \in \mathcal{M}} \bar{\xi}^u(i,j)}{\sum_{u \in \mathcal{M}} \bar{\gamma}^u(i)}, \quad B_i^*(z_k) = \frac{\sum_{u \in \mathcal{M}} \bar{\omega}^u(i,k)}{\sum_{u \in \mathcal{M}} (\bar{\gamma}^u(i) + \bar{\gamma}_{T_u}^u(i))}.$$

Using η^* and the past observations h^{tr} (seen by tr), the trust of tr in te is evaluated as an estimated predictive distribution, i.e. the probability - given h^{tr} and η^* - of every possible outcome $z_k \in V$ for a new interaction with te . This distribution is expressed as

$$P(z_k | h^{tr}, \eta^*) = P(z_k, h^{tr} | \eta^*) / P(h^{tr} | \eta^*) \quad \forall z_k \in V.$$

In the above equation, the probabilities on the right side are defined by Eq. 1, and efficiently evaluated by the *forward-backward* algorithm ([17]).

5 Experimental evaluation

To evaluate our reputation model experimentally, we simulate an HMM λ representing a trustee te . We consider two partners: tr and rs . Each interaction between te and a partner is simulated by allowing λ to make a state transition, and produce an observation in one partner. Over time we allow the reputation source rs to serve a reputation report about te to the truster tr , which combines it with its local trust information to produce a new estimated predictive distribution for te . We then evaluate the expected estimation error, defined in Section 3, using Monte-Carlo approach (see, e.g., [3]).

Let λ be a 4-state HMM representing a ‘stable’ behaviour, where the probability of making self-transition is high (0.9), and other transitions are equally likely. Let the observation alphabet $V = \{1, 2\}$, where the emission matrix is

$$B_\lambda = \begin{bmatrix} 1.0 & 0.0 \\ 0.7 & 0.3 \\ 0.3 & 0.7 \\ 0.0 & 1.0 \end{bmatrix} \quad (16)$$

At each interaction with te , we assume the tr and rs are equally likely (with probability 0.2) to be te ’s partner, while it remains a probability 0.6 that any other principal is the partner. Figure 3, shows the impact of using the HMM-based reputation model on the expected estimation error. The graph on the left compares the HMM trust model *with* and *without* reports from reputation sources. The higher curve shows the estimation error resulting from using only tr ’s observations (1 report). The lower curve shows the error when tr uses also the reputation report collected from rs (2 reports). The improvement resulting from mixing in the reputation report is indicated by the vertical gap between the curves. Observe that this becomes less significant as the total number of interactions grows. This is because in the case of long sequences, the observations made individually by tr tend to be sufficient to learning the trustee’s behaviour with accuracy.

With the same simulation framework, the right-hand side of Figure 3 compares the expected estimation errors of our model against the Beta reputation model. Observe that for a relatively low number of total interactions T , the beta model outperforms the HMM reputation model by exhibiting a lower expected estimation error. In this case, the combined length of the observation sequences is not sufficient to capture the ‘dynamicity’ of te ’s behaviour. As a result, the learning algorithm working on such input produces a low-quality approximate behaviour HMM η . Hence the large estimation error compared to using Beta reputation reports. However, when the number of interactions grows and the sequence become long enough to compute a good estimate, the




Fig. 3. The expected estimation error using the HMM-based reputation model

HMM reputation model exhibits a substantially lower error than the Beta model. The Beta model falls short here: since it ignores te 's dynamic behaviour and only learns an 'average' probability distribution over possible outcomes, it can not keep the expected estimation level low. It is also apparent from Figure 3, that incorporating a decay factor in the Beta model (viz., 0.9 in the example) reduces the expected estimation error. Further details about the effect of decay can be found in [7].

6 Conclusion

We proposed a model for reputation which completes our basic HMM-based trust model in [8] and yields the first trust-and-reputation model for multi-state, dynamic systems. The reputation model enhances the quality and reliability of trust judgements and their evaluation process by using feedback information about the trustee in the form of *reputation reports*. The latter are a 'digest' of a principal's interactions with the trustee, conceptually nothing but an abstraction on the simple idea of ratings given by *reputation sources* about trustees. The model provides mixing equations which can be used by the truster to combine reputation reports collected from different sources, together with its own trust information, in order to evaluate its trust in the trustee.

We used the same experimental approach previously used in [8] to evaluate and compare trust models in terms of the expected estimation error. This allows us to investigate the impact of the HMM-reputation model on trust evaluation, as well as to compare this model against its predecessors. We found that the estimation error is significantly reduced when multiple reputation reports are used in the trust evaluation process. We also discussed how the improvement due to reputation reports gets less significant as the total number of interactions with the trustee gets larger. This is because a larger number of total interactions with the trustee implies that the single sequence experienced by the truster itself tends to provide a sufficient accurate basis to learn the trustee's behaviour.

A comparison with the Beta reputation model, using the same number of reputation sources, yielded that the Beta reputation model outperforms our HMM-based reputation when the total number T of interactions is relatively small. As T gets larger, the HMM-based reputation model gradually improves in terms of the estimation error, and eventually outperforms the Beta model very significantly. This is because longer ob-

ervation sequences imply more accurate approximate models of the trustee's dynamic behaviour, information which, in contrast, the Beta model ignores altogether.

References

1. L. E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite-state Markov chains. *Annals of Mathematical Statistics*, 37(6):1554–1563, Dec 1966.
2. L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970.
3. P. Brémaud. *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. Springer, 1998.
4. S. Buchegger and J. Le Boudec. A Robust Reputation System for Peer-to-Peer and Mobile Ad-hoc Networks. In *P2PEcon 2004*, 2004.
5. V. Cahill, E. Gray, J. M. Seigneur, C. D. Jensen, Y. Chen, B. Shand, N. Dimmock, A. Twigg, J. Bacon, C. English, W. Wagealla, S. Terzis, P. Nixon, G. Serugendo, C. Bryce, M. Carbone, K. Krukow, and M. Nielsen. Using trust for secure collaboration in uncertain environments. *IEEE Pervasive Computing*, 2(3):52–61, 2003.
6. T. M. Cover and J. A. Thomas. *Elements of Information Theory 2nd Edition*. Wiley Series in Telecommunications and Signal Processing. Wiley-Interscience, July 2006.
7. E. ElSalamouny, K. Krukow, and V. Sassone. An analysis of the exponential decay principle in probabilistic trust models. *Theoretical Computer Science*, 410(41):4067 – 4084, 2009.
8. E. ElSalamouny, V. Sassone, and M. Nielsen. Hmm-based trust model. In P. Degano and J. Guttman, editors, *Formal Aspects in Security and Trust*, volume 5983 of *Lecture Notes in Computer Science*, pages 21–35. Springer Berlin / Heidelberg, 2010.
9. G. Grimmet and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, third edition, 2001.
10. A. Jøsang and J. Haller. Dirichlet reputation systems. In *The Second International Conference on Availability, Reliability and Security, 2007. ARES 2007.*, pages 112–119, 2007.
11. A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings from the 15th Bled Conference on Electronic Commerce, Bled*, 2002.
12. S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, March 1951.
13. L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation (for ebusinesses). In *Proceedings from 5th Annual Hawaii International Conference on System Sciences (HICSS'02)*, page 188. IEEE, 2002.
14. M. Nielsen, K. Krukow, and V. Sassone. A bayesian model for event-based trust. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 172:499–521, April 2007.
15. J. R. Norris. *Markov chains*. Cambridge University Press, 1997.
16. L. Rabiner and B. H. Juang. *Fundamentals of Speech Recognition*. Prentice Hall PTR, united states ed edition, April 1993.
17. L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, February 1989.
18. W. Teacy, J. Patel, N. Jennings, and M. Luck. Travos: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12:183–198, 2006.
19. L. Xiong and L. Liu. PeerTrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on knowledge and data engineering*, 16(7):843–857, 2004.

A Proofs

A.1 Proof of Lemma 1

For any given h^u , let q denote the random sequence of states underlying h^u . We have

$$\begin{aligned}
\log P(h^u | \eta) &= \log \left(\sum_q P(h^u, q | \eta) \right) = \log \left(\sum_q P(q | h^u, \eta^u) \frac{P(h^u, q | \eta)}{P(q | h^u, \eta^u)} \right) \\
&= \log \left(\mathbf{E}_q \left[\frac{P(q, h^u | \eta)}{P(q | h^u, \eta^u)} \mid h^u, \eta^u \right] \right) \stackrel{(1)}{\geq} \mathbf{E}_q \left[\log \left(\frac{P(q, h^u | \eta)}{P(q | h^u, \eta^u)} \right) \mid h^u, \eta^u \right] \\
&= \sum_q P(q | h^u, \eta^u) \log \frac{P(q, h^u | \eta)}{P(q | h^u, \eta^u)} \\
&= \sum_q P(q | h^u, \eta^u) \log P(q, h^u | \eta) - \sum_q P(q | h^u, \eta^u) \log P(q | h^u, \eta^u) \\
&= Q(\eta^u, h^u, \eta) + \mathcal{R}(\eta^u, h^u)
\end{aligned}$$

The inequality (1) is obtained by applying Jensen's inequality (see e.g. [6, Theorem 2.6.2]) to the \log function. The equality holds when $\eta = \eta^u$. It easily follows that.

$$\sum_u \frac{1}{T_u} \log P(h^u | \eta) \geq \sum_u \frac{1}{T_u} Q(\eta^u, h^u, \eta) + \sum_u \frac{1}{T_u} \mathcal{R}(\eta^u, h^u)$$

where the equality holds when $\eta = \eta^1 = \eta^2 = \dots = \eta^M$. Using the definition (9) of $\mathcal{G}(h^1, h^2, \dots, h^M | \eta)$ to substitute the summation in the left hand side of the above equation, the proof is complete. \square

A.2 Proof of Theorem 1

Note that the summation $\sum_u \frac{1}{T_u} \mathcal{R}(\eta^u, h^u)$ is independent of η . The maximisation in Eq. (12) amounts, therefore, to maximising $\sum_u \frac{1}{T_u} Q(\eta^u, h^u, \eta)$. Let $h^u = o_1^u, o_2^u, \dots, o_{T_u}^u$ denote the observation sequence observed by reputation source u .

Let also $q^u = q_1^u, q_2^u, \dots, q_{T_u}^u$ denote the (hidden) sequence of states underlying the observation sequence h^u . We start by expressing $\log P(h^u, q^u | \eta)$ in terms of the parameters of η as follows.

$$\log P(h^u, q^u | \eta) = \log \pi_{q_1^u} + \sum_{t=2}^{T_u} \log A_{q_{t-1}^u q_t^u} + \sum_{t=1}^{T_u} \log B_{q_t^u}(o_t^u) \quad (17)$$

where π_i denotes the probability that the initial state (q_1^u) is i . A_{ij} is the probability of transition from state i to state j . $B_i(z_k)$ is the probability of observing the outcome z_k at state i . Refer to the description of the HMM elements in Section 2 for more details about these notations.

Substituting Expression (17) in (10), the function $\mathcal{Q}(\eta^\mu, h^\mu, \eta)$ can be written as follows.

$$\begin{aligned} \mathcal{Q}(\eta^\mu, h^\mu, \eta) &= \sum_{i=1}^N P(q_1^\mu = i | h^\mu, \eta^\mu) \log \pi_i + \\ &\sum_{i=1}^N \sum_{j=1}^N \sum_{t=2}^{T_u} P(q_{t-1}^\mu = i, q_t^\mu = j | h^\mu, \eta^\mu) \log A_{ij} + \\ &\sum_{i=1}^N \sum_{k=1}^K \sum_{t=1}^{T_u} P(q_t^\mu = i | h^\mu, \eta^\mu) \delta(o_t^\mu, z_k) \log B_i(z_k) \end{aligned} \quad (18)$$

where N is the number of states, K is the number of possible observation symbols, and the δ -function $\delta(o_t^\mu, z_k)$ is defined as:

$$\delta(o_t, z_k) = \begin{cases} 1 & \text{if } o_t = z_k \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

Now we are ready to expressing the sum $\sum_{u=1}^M \frac{1}{T_u} \mathcal{Q}(\eta^\mu, h^\mu, \eta)$ in (12) by scaling Eq. (18) by $1/T_u$, and then summing over the available reputation sources $\{1, 2, \dots, M\}$. For convenience, we write the resulting sum as follows.

$$\sum_{u=1}^M \frac{1}{T_u} \mathcal{Q}(\eta^\mu, h^\mu, \eta) = \mathcal{Q}_\pi(\boldsymbol{\pi}) + \sum_{i=1}^N \mathcal{Q}_{A_i}(\mathbf{A}_i) + \sum_{i=1}^N \mathcal{Q}_{B_i}(\mathbf{B}_i) \quad (20)$$

where $\boldsymbol{\pi} = [\pi_1, \pi_2, \dots, \pi_N]$ is the vector representing the initial state probability distribution, $\mathbf{A}_i = [A_{i1}, A_{i2}, \dots, A_{iN}]$ is the vector representing the probability distribution over state transitions from state i to other states, and $\mathbf{B}_i = [B_i(z_1), B_i(z_2), \dots, B_i(z_K)]$ is the vector representing the emission probability distribution over outcomes given state i . The functions $\mathcal{Q}_\pi(\boldsymbol{\pi})$, $\mathcal{Q}_{A_i}(\mathbf{A}_i)$, and $\mathcal{Q}_{B_i}(\mathbf{B}_i)$ in the above equation are defined as follows.

$$\mathcal{Q}_\pi(\boldsymbol{\pi}) = \sum_{i=1}^N \left(\sum_{u=1}^M \frac{1}{T_u} P(q_1^\mu = i | h^\mu, \eta^\mu) \right) \log \pi_i \quad (21)$$

$$\mathcal{Q}_{A_i}(\mathbf{A}_i) = \sum_{j=1}^N \left(\sum_{u=1}^M \frac{1}{T_u} \sum_{t=2}^{T_u} P(q_{t-1}^\mu = i, q_t^\mu = j | h^\mu, \eta^\mu) \right) \log A_{ij} \quad (22)$$

$$\mathcal{Q}_{B_i}(\mathbf{B}_i) = \sum_{k=1}^K \left(\sum_{u=1}^M \frac{1}{T_u} \sum_{t=1}^{T_u} P(q_t^\mu = i | h^\mu, \eta^\mu) \delta(o_t^\mu, z_k) \right) \log B_i(z_k) \quad (23)$$

Observe that each term in Equation (20) is a function of a probability distribution which parametrises the HMM η . These distributions ($\boldsymbol{\pi}$, \mathbf{A}_i , $\mathbf{B}_i \forall i : 1 \leq i \leq N$) are independent of each other, that is the choice of one of them does not affect the choice

of the others. Therefore the overall sum (20) is maximised by maximising each term in (20) separately. Observe furthermore that each of equations (21),(22), and (23) is in the following form.

$$F(y_1, y_2, \dots, y_V) = \sum_{v=1}^V w_v \log y_v \quad \text{where} \quad \sum_{v=1}^V y_v = 1 \quad (24)$$

Using the *Lagrange multiplier* technique for optimising a function subject to a constraint, the constrained function F defined above can be easily proved to have a global maximum at the point $(\bar{y}_1, \bar{y}_2, \dots, \bar{y}_V)$, where \bar{y}_v is given by

$$\bar{y}_v = \frac{w_v}{\sum_{v=1}^V w_v}$$

Using the above fact, the parameters of the optimal a posteriori model η^* are given as follows.

$$\pi_i^* = \frac{\sum_{u=1}^M \frac{1}{T_u} P(q_1^u = i | h^u, \eta^u)}{\sum_{u=1}^M \frac{1}{T_u}}$$

$$A_{ij}^* = \frac{\sum_{u=1}^M \frac{1}{T_u} \sum_{t=2}^{T_u} P(q_{t-1}^u = i, q_t^u = j | h^u, \eta^u)}{\sum_{u=1}^M \frac{1}{T_u} \sum_{t=2}^{T_u} P(q_{t-1}^u = i | h^u, \eta^u)}$$

$$B_i^*(z_k) = \frac{\sum_{u=1}^M \sum_{t=1}^{T_u} P(q_t^u = i | h^u, \eta^u) \delta(o_t^u, z_k)}{\sum_{u=1}^M \frac{1}{T_u} \sum_{t=1}^{T_u} P(q_t^u = i | h^u, \eta^u)}$$

□