A Subjective Approach to Routing in P2P and Ad Hoc Networks

Andrew Twigg

Computer Laboratory, Cambridge University, UK
andrew.twigg@cl.cam.ac.uk

Abstract. This paper presents a subjective approach to routing in peer-to-peer and ad hoc networks. The main difference between our approach and traditional routing models is the use of a trust model to mediate the risk inherent in routing decisions. Rather than blindly exchanging routing table entries, nodes 'discount' recommendations from other nodes using a distributed trust computation which allows them to avoid malicious, faulty and unreliable nodes and links in routing decisions. Adding the risk model allows energy-efficient routing decisions to be made in a wireless network, and we show how our model can be optimized for different network behaviours, including wireless networks. The model is described in the context of the DSR [1] routing algorithm, although it is equally-applicable to others, including peer-to-peer routing substrates.

1 Introduction

The game of chinese whispers is often played by a group of children (the *nodes*) sitting in a circle (the *route*). One such node (the *source*) whispers something (the *packet*) to the next node on the route, which is supposed to forward it to the next node. This continues until the packet reaches the node at the end of the route (the *destination*). At school, the source and destination are often the same, heightening the dramatic effect when the packet has been corrupted[1].

An ad hoc network has no fixed infrastructure, and the lack of implicitly trusted routers means that each node becomes part of the routing fabric. Hence the network is collaborative, since successful operation relies on nodes correctly forwarding packets which may have no direct benefit for themselves. Routing in such a network is similar to a game of chinese whispers where an entity (mobile node) must rely on other intermediate nodes to correctly forward the packet to the destination, yet in general, the intermediate nodes may have no prior contact with the sender.

In this paper, we describe how a source node can choose routes to minimize the disruption to the packets it wishes to send, by *avoiding* certain nodes. The approach is rather different to traditional routing methods. We develop a trust

[1] ScoutBase UK [www.scoutbase.org.uk/activity/games/pages/whispers.htm] report that 'once we started with "we have a new car" and ended with "someone ate a brand new car".'

model that can be used to reason explicitly about routing choices and the nodes and links which constitute them. Rather than reasoning objectively, we operate in the domain of subjective logic [2] which permits uncertainty in probabilities, so nodes form *opinions* rather than storing observations. Our main contribution is that, rather than blindly exchanging routing table entries, nodes can 'discount' these recommendations from other nodes using a distributed trust computation that takes into account others' opinions about the recommender, hence avoiding both malicious and unreliable nodes and links.

The second contribution of this work is in using the trust model to mediate the risks associated with routing decisions. In a wireless ad hoc network, this allows nodes to make informed routing decisions by trading energy requirements against the reliability of a route. The work is based on a trust-based extension to the dynamic source routing algorithm presented in [3], and the work on establishing trust in peer-to-peer systems in [4, 5]. Secure routing protocols for similar problems are presented in [6–8], though we want to consider the less 'traditional' secure notion of trust, as in [9].

The remainder of this paper is organised as follows. Section 2 briefly describes our network model, Section 3 develops the trust model and Section 4 outlines how it can be integrated with a risk model. A major part of the paper (Section 3.3) develops an inference procedure, which attempts to identify the unreliable nodes and links, given that we can only observe aggregate properties of routes.

2 Network Model

We model a *network* $\mathcal{N}$ as a set of *nodes* and a set of *links* between pairs of nodes, where packets can be sent bidirectionally across links. We model network behaviour by initially assuming that packets can be dropped at nodes only (irrespective of their source or destination). However, we show in Section 3.5 how to adapt our model to apply under the assumption that packets are instead dropped at links, which may be more useful when operating in a wireless network.

In this paper we only consider source routing, such as the dynamic source routing algorithm (DSR) [1] which is designed for mobile ad-hoc networks and forms the basis for the trust work in [3]. To send a packet from r_1 to r_n requires that the source r_1 compute the entire route $r = \langle r_1, r_2, \ldots, r_n \rangle$ and embed it into the packet. In DSR [1], each node maintains a route cache of recently used[2] routes, indexed by destination. The cache is maintained in response to changing topologies by sending 'route request' packets or 'snooping' on others' route request packets. Other types of protocol include 'hop-by-hop' (where the route is not established at source), which we hope to consider later. A good survey of peer-to-peer and ad hoc routing techniques is presented in [10].

In DSR, a 'route error' packet is sent to the source (or a timeout occurs) when the packet could not be delivered, due to an intermediate or destination node being unreachable, or as the result of an intermediate node dropping the

[2] The cache replacement algorithm is not fixed, although LRU is often used.

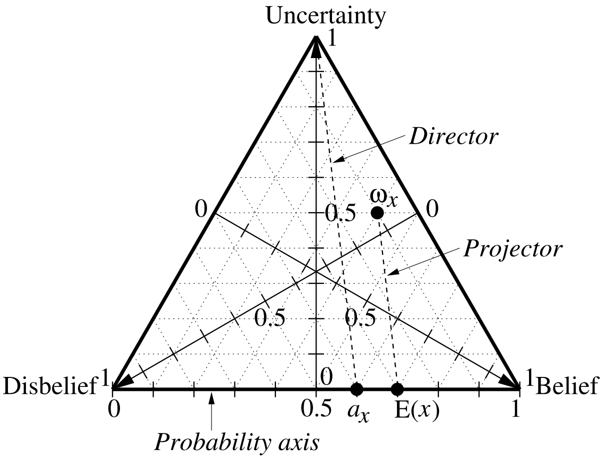

Fig. 1. The opinion space Ω showing the opinion $\omega_x = (0.4, 0.1, 0.5)$ as an example. The 'weight' of the opinion is $a_x = 0.6$, which is used to determine the expectated value $\mathrm{E}(x)$ since the *projector* lies parallel to the *director*

packet (or even faking the route error packet). Such packets form the basis of the observations upon which our trust model is built.

3 Trust Model

"Trust is a precious commodity; easily damaged but difficult to mend."

3.1 Trust Values

Before developing the trust model, we present a few brief definitions. *Trust values* are elements of a complete lattice $(T, \leq)$, P is the set of *principals* and the *trust space* $\mathcal{T}$ is a partial function $\mathcal{T} : (P \rightharpoonup T)$. Initially, let P be the set of nodes in the network, hence $\mathcal{T}(u) \in T$ is our trust in node u. In Section 3.5 we show a use for other definitions of P.

How does one assign meaning to T? A starting point is to follow [3] and take trust values to be arbitrary scalars $\in [0, 1]$ where $0, 1$ represent complete distrust and trust, respectively. Using arbitrary values (in the sense that they do not represent any measurable quantity) presents problems in both understanding and analysis. Asking 'what does it mean for a node to have trustworthiness 0.5?' become subjects of philosophical debate, and asking 'how should I update the trustworthiness of a node given it performed action a in context Γ with outcome o' are open to subjective interpretation.

Subjective Logic. We take elements of T to be *opinions* in subjective logic [2]. An *opinion* $\omega = (b, d, u)$ is an element in the Barycentric *opinion space* Ω

shown in Figure 1, where $b, d, u \in [0,1]$ represent belief, disbelief and uncertainty, respectively, and $b+d+u=1$. The horizontal line representing $u=0$ is the *probability axis* and represents situations without uncertainty (*dogmatic* opinions), equivalent to a traditional probability model. Uncertainty is caused by lack of evidence to support either belief or disbelief, and so the opinion space provides a way of continuously distinguishing between notions of 'unknown', $(0,0,1)$ and 'least trust', $(0,1,0)$. Observations are made in the *evidence space* Φ and are transformed into the opinion space using a simple bijection, derived in [2]. Let $\varphi_r^v = (x,y)$ be an *observation* by principal v about route r where x, y are the numbers of successful and unsuccessful packet transmissions, respectively. Let the opinion corresponding to φ be $\omega(\varphi) = (b,d,u)$ where

$$b = x/(x+y+r) \tag{1}$$

$$d = y/(x+y+r) \tag{2}$$

$$u = r/(x+y+r) \tag{3}$$

where $r \geq 1$ is a parameter controlling the rate of loss of uncertainty, which can be used to tune the use of uncertainty in the model for the requirements of different scenarios (we often take $r=1$).

Given the above and the trust space, we say that $\mathcal{T}(v)$ is our *opinion* about the proposition 'principal v forwards packets successfully', based on observations on routes containing v. We make use of three operators on opinions, the definitions of which are reproduced in Appendix A.

Unknown principals The opinion space helps overcome the problem of unknown principals being associated with 'least trust', by assigning them the opinion $\omega = (0,0,1)$. As a metric to order opinions, Jøsang [2] suggests computing the expected probability value by projecting the opinion onto the probability axis, parallel to the director (see Figure 1). If the director cuts the axis at ρ ($\rho = 0.6$ in the Figure), then $E(\{b,d,u\}) = (b+u)/(b+d+u/\rho)$.

The meaning of this is to assign a 'newcomer' principal with $\mathcal{T}(v) = (0,0,1)$ the expectation ρ, whilst maintaining independence between the notions of 'unknown' and 'untrusted'. One could also take ρ as the expectation of some distribution representing all the current principals' expectations.

3.2 Building the Trust Model

To make a good estimate of the trust space $\mathcal{T}$, one needs a good inference procedure and a good set of observations (and, implicitly, a good way of gathering them).

The Need for Inference. In our model, observations can be made only on routes. Rather than take P to be the set of routes, we make use of an inference procedure whose job is to make estimations about *principals* given observations on *routes*. By using knowledge about principals on other routes, opinions about new unused routes can be formed.

An *inference procedure* Θ maps observations on routes to opinions about principals, i.e. $\Theta : \Phi \rightarrow (P \rightarrow \Omega)$ where Φ is the evidence space and Ω is the opinion space. Separating the inference procedure from the trust computation allows each node to have its own procedure, for example some nodes could become authorities on recommendations because they have the power to perform stronger inference, whereas nodes with less power can perform weaker inference. An inference procedure based on a least mean-squares error approach is developed in Section 3.3.

Gathering observations. A useful strategy is to supplement a node's direct observations with observations from other nodes, known as *recommendations*. Recommendations are made by piggypacking some observations onto the routing table entries (RTEs) returned in response to a route request packet. A typical set of routing table entries returned from a node v would be as in Figure 2.

$$\left\{ \left(\text{route A}, \varphi^{v}_{\text{route A}}\right), \left(\text{route B}, \varphi^{v}_{\text{route B}}\right), \left(\text{route C}, \varphi^{v}_{\text{route C}}\right) \right\}$$

Fig. 2. Making recommendations: route observations are piggybacked onto the routing table entries (RTEs) returned from a node v, in response to a route request packet

The advantage of using recommendations is that it removes the need to store and explore many routes, and reduces the chance that little can be inferred about a route's trustworthiness. However, we must deal with the fact that principals may lie (making bad inferences is not a problem since only direct observations are communicated).

The gathering of recommendations via the route request mechanism can be used in conjunction with a strategy for separately gathering recommendations by sending dummy packets (though one needs to consider the extra network traffic incurred). Some strategies will provide a higher amount of information than others, but this is difficult to quantify as it depends on the inference procedure too. One approach might be to gather recommendations along routes which currently have a high uncertainty, or when a node moves into a new region of space. A promising approach is to make *direct observations* of increasing length along a route (e.g. r_1 then r_1, r_2 then r_1, r_2, r_3 etc.) by sending out dummy packets.

Malicious and Colluding Principals. There are two types of threat from malicious and colluding principals. Malicious principals can collude to attempt to make each other look more trustworthy. Alternatively, principals can try to make other principals appear malicious by spreading bad recommendations. In the trust computation, we assume that principals are rational in the game-theoretic sense, i.e. that a principal makes true recommendations if they forward packets (and are trustworthy), and attempts to hide their misbehaviour by making false recommendations if they do not forward packets (and are untrustworthy). Hence $\mathcal{T}(v)$ is an appropriate discounting for the credibility of v's alleged observations.

The trust computation ensures that recommendations from colluding nodes are heavily discounted (effectively ignored) unless one of the following holds:

1. Principals we trust make good recommendations on a colluding principal;
2. We make good direct observations on a colluding principal.

Under the game-theoretic assumption, these both fail unless we also collude. In practice, a principal in a clique of colluding principals can build up one's trust by forwarding packets correctly, then use this to falsely make good observations about the others in the clique (and violating our assumption). Whether or not the assumption will be valid in reality remains a point of discussion, particularly with regard to faulty (misconfigured) principals.

One can avoid making the assumption by using separate trust spaces for 'participation' (forwarding packets) and 'recommendation' (the ability to make accurate recommendations) as in [**?**]. Taking this approach complicates the trust computation, but can easily be integrated into the model presented here.

Trust computation. Here we describe the local trust computation carried out at each node. Let Θ be the inference procedure used. The trust space $\mathcal{T}$ is described by an iterative fixpoint computation over the set of observations $\{\varphi^w\}$ (both our direct observations and recommendations from other nodes), as below:

$$\mathcal{T}(v) = \bigoplus_{\varphi^w} \{\mathcal{T}(w) \otimes \Theta(\varphi^w)(v)\} \tag{4}$$

where $\oplus$ is the *Bayesian consensus* operator (to combine opinions), and $\otimes$ is the *discounting* operator, defined in Appendix A.

A more descriptive view of Equation (4) is that we are trying to determine the trust value (our opinion) of principal v. Given the set of direct observations and recommendations, we apply our inference procedure to obtain opinions on principals. To account for false recommendations, these opinions are discounted by our opinion of their observer (under the assumption made earlier), and the resulting opinions combined using the consensus operator. But v's trust value affects the weighting of v's observations, and so on. The solution is analogous to eigen problems in sparse graphs (such as the WWW), so techniques similar to PageRank [11] can be employed to solve it. A similar computation and its approximate solution is described in [4]. Note that node v's direct observations are discounted by v's opinion of itself.

A consequence of having to *infer* trust values from observations, in a distributed context, is that each node may estimate the trust values in a different way, based on the same observations. It would be interesting to see if this leads to greater subjective variation of trust throughout the system (as in real life).

3.3 Inference: Spreading the Blame

The goal of this section is to develop an inference procedure to estimate principals' *prior* behaviours from *a priori* observations on routes. Assume one makes the observation $\varphi_r = (x, y)$. What can be concluded? Given no prior information about principals, the fairest is to assume each principal behaved equally badly. Perhaps an obvious solution would be to distribute the observation by uniformly shifting the uncertainty component in the opinion of each principal on the route. But the number of observations on a principal that a route observation provides decreases exponentially from the source, since each principal drops some *proportion* of its packets along the way and so fewer packets reach nodes further along the route. This section presents a theory for spreading the 'blame' or 'praise' among nodes, and uses a least-mean squares error approach.

Our inference consists of two parts. We first estimate the most likely *packet-level* behaviour of principals then use these to form new opinions on principals. An opinion about a route r can then be formed as the conjunction of the opinions of nodes in r. Consider the observation $\varphi_r = (x, y)$ along the route $r = \langle r_0, r_1, \ldots, r_n, r_{n+1}\rangle$ (where r_0 is the source and r_{n+1} the destination). Let the total number of observations, $T = x + y$. We 'explain' the observation by assuming each principal r_i successfully transmitted a *proportion* α_i of its packets, for the duration of the observation.

First, consider the behaviour of node r_1. Since T packets enter r_1 and α_1 are dropped, the estimated observation is $\varphi_{r_1} = (\alpha_1 T, (1-\alpha_1)T)$. For principal r_2 the situation is similar, since $(\alpha_1 T)$ packets are expected to leave r_1 and r_2 is expected to successfully forward a proportion α_2 of them. Hence we have the expected observations for each principal

$$\begin{aligned}
\varphi_{r_1} &= (\alpha_1 T, (1-\alpha_1)T) \\
\varphi_{r_2} &= (\alpha_1\alpha_2 T, (1-\alpha_2)\alpha_1 T) \\
&\vdots \\
\varphi_{r_i} &= \left(\prod_{j=1}^{i} \alpha_j T, (1-\alpha_i)\prod_{j=1}^{i-1} \alpha_j T\right)
\end{aligned}$$

One can check that this does in fact explain the observation, since $(\prod_{j=1}^{n} \alpha_j T) = x$ packets leave r_n, as observed. Transforming these from the evidence space to the opinion space gives the opinions

$$\begin{aligned}
\omega(\varphi_{r_1}) &= \left(\frac{\alpha_1 T}{T+1}, \frac{(1-\alpha_1)T}{T+1}, \frac{1}{T+1}\right) \\
\omega(\varphi_{r_2}) &= \left(\frac{\alpha_1\alpha_2 T}{\alpha_1 T+1}, \frac{(1-\alpha_2)\alpha_1 T}{\alpha_1 T+1}, \frac{1}{\alpha_1 T+1}\right) \\
&\vdots \\
\omega(\varphi_{r_i}) &= \left(\frac{\prod_{j=1}^{i} \alpha_j T}{\prod_{j=1}^{i-1} \alpha_j T+1}, \frac{(1-\alpha_i)\prod_{j=1}^{i-1} \alpha_j T}{\prod_{j=1}^{i-1} \alpha_j T+1}, \frac{1}{\prod_{j=1}^{i-1} \alpha_j T+1}\right)
\end{aligned}$$

This procedure can be seen as shifting the unknown component of the opinions, but relative to the expected number of observations made on a node. An interesting result is that one has to make around $(1/\prod_{j=1}^{n}\alpha_j)$ times as many observations to achieve an equal reduction in the unknown of the nth node, compared to the unknown in the 1st node. This implies that the learning rate increases with the α's, i.e. our opinions will be less uncertain when observing well-behaved routes. A similar argument can be made for shorter routes, since an observation is spread less thinly over the nodes on the route. Together, these suggest a strategy that a node should use to gather observations in order to maximise its 'information gain'.

Estimating the behaviours: picking the α_i's. So far we have not discussed how the α_i's are chosen. To *correctly* explain the observation, we must satisfy the constraint $\prod_{i=1}^{n}\alpha_i = (x/T)$. A good start would be to assume no prior knowledge of the current principals' behaviours and hence that all principals behaved equally, that is $\alpha_1 = \alpha_2 = \cdots = (x/T)^{1/n}$.

Let us now take into account the current opinions about principals. For example, if we currently have high belief in r_1 and high disbelief in r_2 and make (or receive) an observation with high disbelief, the 'best' explanation is that which most closely resembles our current opinions: that r_1 behaved well and r_2 behaved badly, rather than penalising them both equally. We now try to make this approach more formal.

Let $\beta_i \in [0,1]$ be the expected value of the current opinion of principal r_i. We want to approximate (in the least squares error sense) the β_i's with a set of α_i's (whilst obeying the constraint). Our approach is to renormalize the opinions such that they satisfy the constraint, which effectively performs a *ratio scaling* on the β's. We conjecture that by minimizing the total square change in our current opinions (β's), the least mean square error is also minimized.

The argument is as follows. Our current β's don't satisfy the constraint (if they do, then $\alpha_i = \beta_i$). Hence we can write

$$\prod_{i=1}^{n}\beta_i \neq (x/T) \tag{5}$$

Taking logs gives

$$\log\prod_{i=1}^{n}\beta_i \neq \log(x/T) \tag{6}$$

Now consider the actual value of the LHS of Equation 6; some unique $d \neq \log(x/T)$. Therefore

$$\log\prod_{i=1}^{n}\beta_i = d \neq \log(x/T) \tag{7}$$

Let $\kappa = \log(x/T)/\log\prod_{i=1}^{n}\beta_i$. Multiplying Equation (7) by κ, simplifying then removing the logs gives

$$\kappa \log \prod_{i=1}^{n} \beta_i = \log(x/T) \tag{8}$$

$$\prod_{i=1}^{n} \beta_i^{\kappa} = (x/T) \tag{9}$$

and hence $\alpha_i = \beta_i^{\kappa}$ where $\kappa = \log(x/T)/\log\prod_{i=1}^{n}\beta_i$.

One can easily check this is a generalisation of the special case described earlier. In fact, the case appears whenever our prior opinions of each principal are equal, not just when there is no knowledge. Setting $\beta_1 = \beta_2 = \cdots = \beta$:

$$\begin{aligned} \alpha_i &= \beta^{\log(x/T)/\log\beta^n} && (10)\\ &= \beta^{\frac{1}{n}\log_\beta(x/T)} && (11)\\ &= \left(\beta^{\log_\beta(x/T)}\right)^{1/n} && (12)\\ &= (x/T)^{1/n} && (13) \end{aligned}$$

Finally, our inference procedure is as follows. Given an observation φ_r, it computes opinions for each principal (node) on the route r (except the source and destination). All other principals are assigned the 'unknown' opinion $\omega = (0, 0, 1)$.

Related approaches. Using arbitrary trust values naturally leads to arbitrary inference procedures. Consider the approach taken in [3]; if the packet was lost, the source node has no way of determining which node in r caused the fault, hence the best it can do is to decrease the trust value of each node on the route. Even ignoring this uniform distribution of blame and the fact that more observations are made on nodes closer to the source, no 'correct' update procedure is obvious. In practice, the values are updated 'exponentially', i.e. for a successful transmission $t = t + (1 - t)/20$ and for an unsuccessful transmission $t = t - t/20$ for each node on the route. It is difficult to make anything other than an empirical case for this procedure, and it stems from using arbitrary trust values.

3.4 Selecting routes

Recall $\mathcal{T}(v)$ is our opinion about the proposition 'principal v forwards packets successfully'. We can form an opinion about a route $r = \langle r_0, r_1, r_2, \ldots, r_n, r_{n+1}\rangle$ by combining opinions about principals using the *conjunction* operator, $\wedge$ [2]:

$$\omega_r = \mathcal{T}(r_1) \wedge \mathcal{T}(r_2) \wedge \cdots \wedge \mathcal{T}(r_n) \tag{14}$$

The remaining problem is ordering opinions, to choose a route from those in the route cache. This is done by ordering routes according to their expected

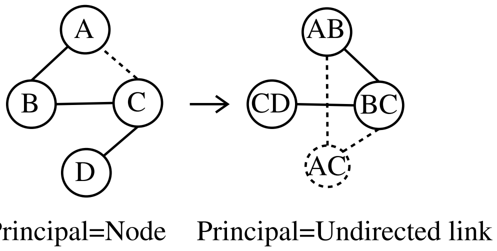

Fig. 3. Optimising the model for wireless networks: to model the behaviour where packets are dropped on links rather than at nodes, the principal space is transformed. In the network on the left, principals represent nodes. The network on the right is equivalent in that principals are now the links from the original network. Applying the trust model to the new 'network' allows it to reason about link, rather than node, behaviours. The dashed edges show the effect of adding the physical link $A \rightarrow C$

probability value $\mathrm{E}(\omega_r)$, as described in Section 3.1. To avoid overloading nodes with slightly higher trust values, the routes should be chosen with probability proportional to their relative trust expected probability values. This provides a form of trust-based load balancing, and gives previously untrusted nodes a chance to redeem themselves.

3.5 A better model for wireless networks

So far the set of principals has been the set of nodes (the *network principal space*), and we have assumed the *weak network model* where packets are dropped only at nodes. Yet often it is the links rather than nodes which are the cause of network disruption, and this is particularly true for wireless networks - contrast this with the Internet, where the large distances between hops and a reliable medium mean that malicious or faulty nodes are the major concern.

We now show how the same routing model can be used in the *undirected-link network model*, where packets can be dropped at individual links rather than nodes (this corresponds to a stronger 'adversary'). This assumption can be modelled by taking P to be the set of links as in Figure 3. The source routing algorithm can be extended to include this transformation as follows. The request 'route from A to C' provides choices $\langle A, B, C\rangle$ and $\langle A, C\rangle$ from the route cache, as before. These are transformed into routes $\langle AB, BC\rangle$ and $\langle AC\rangle$ respectively in the model principal space (where a principal is an undirected link), about which opinions can now be formed as before. Observations are gathered in a similar way; an observation on a route $\langle B, C, D\rangle$ transforms to $\langle BC, CD\rangle$ in the model principal space and is handled as before.

We may consider an even stronger adversary which can block packets at links based on their direction (the *directed-link network model*), hence the model principal space induces a directed graph. More devious adversaries can block packets on links, based on the route taken *so far* of the packet. Such a network model can be simulated in the [un]directed-link network model where the model principal space is the set of possible routes to a node in the network, from all sources.

4 Risk

> *"There is no need to trust anyone unless there is risk involved."*

Routing decisions based on trust need to be mediated by the risk of doing so. For example, less trust is needed to send a packet of low importance than to send a message whose safe arrival is critical. Previous works such as [3, 4, **?**] consider trust in complete isolation, and in this section we show how trust can mediate the risk-inherent action of sending packets.

Nodes in a wireless network often have limited battery power, so energy conservation is an important topic when considering routing decisions. In a sense, the *risk* to a node of sending a packet along a particular route is proportional to the energy needed to make the first hop, and the probability of that energy being 'wasted' (if the packet is not successfully sent). Since our trust values represent a meaningful quantity, *i.e.* the probability of a successful transmission, then the inverse represents the expected number of retransmissions, assuming retries are to the same node (although more complex schemes can be handled). We say that trust *mediates* the energy risks in wireless routing, as described in Figure 4.

We can also consider aggregate properties, including retransmissions and so on. Assuming that packets will be delivered after an infinite number of retransmissions, consider the action 'send p to v within time t'. Our outcomes are likely to be $\{p$ delivered late, p delivered successfully$\}$ (the outcome 'p is not delivered' is not possible because of the retransmission assumption). Our policy will be to carry out the sub-action 'send p to node v' until we observe the outcome 'p successfully sent'. Assuming a geometric distribution, i.e. independent retransmissions[3], the expected time before we observe this is 'latency'$/(1 - \mathrm{Pr}($'p delivered successfully'$))$. This can be used to parameterise a distribution over the time, and calculate $\mathrm{Pr}(\text{time} < t)$ from which the expected cost of the original action can be found, for a given route. This kind of reasoning permits one to justify *why* a particular route should be chosen.

5 Conclusion

This paper has presented a subjective approach to routing in peer-to-peer and ad hoc networks. The main difference between our approach and traditional routing models is the ability to reason subjectively about trust in the network through the use of opinions. Rather than blindly exchanging routing table entries, nodes can 'discount' these recommendations from other nodes using a distributed trust computation that takes into account others' opinions about the recommender.

[3] This is not quite true; the route may be switched if the trust values change enough during transmissions.

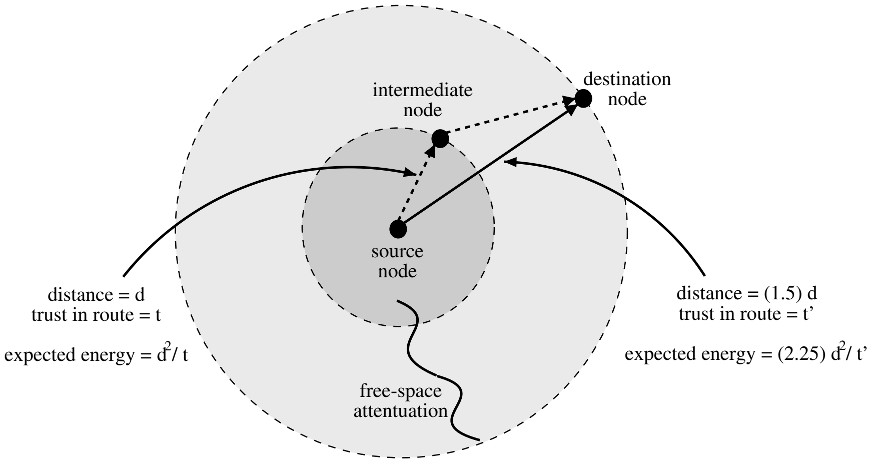

Fig. 4. Using trust to mediate the energy risks in wireless routing. Assuming Friss free-space attentuation, the energy needed for wireless transmission over distance d is proportional to d^2. The source node has two possible routes - the dashed route involving the intermediate node and the direct route. The dashed route requires less energy for the first hop, yet the direct route may have a higher probability of succeeding. Given the trust values t and t' of the routes and that the direct route is 50% further, which route should the source pick to minimize the expected energy required? Equating the two energy equations, we find the direct route is most energy-efficient if $t' \geq 2.25 \cdot t$, *i.e.* over twice as trustworthy

Hence malicious and faulty participants can be avoided. Finally, we showed how the model can be optimised for various uses, including wireless networks.

Source-routing algorithms such as DSR can easily be augmented with the model, and we hope to apply similar ideas to hop-by-hop algorithms such as Pastry [12].

Acknowledgements. The author wishes to thank Dave Gamble, David Ingram, Brian Shand and members of the Cambridge Computer Laboratory for interesting discussions and insightful comments. This work has been inspired by the EU-funded SECURE project (IST-2001-32486), part of the EU Global Computing initiative. As a result, the author thanks BRICS at Århus, Denmark for their initial work in developing a formal trust model [13]. The author is supported by BTexact Technologies and EPSRC.

References

1. Johnson, D., Maltz, D., Broch, J.: DSR: A dynamic source routing protocol for multihop wireless ad hoc networks. In: David B. Johnson, David A. Maltz, and Josh Broch. DSR The Dynamic Source Routing Protocol for Multihop Wireless

Ad Hoc Networks. In Ad Hoc Networking, edited by Charles E. Perkins, chapter 5, pages 139–172. Addison-Wesley, 2001. (2001)

2. Jøsang, A.: A logic for uncertain probabilities. Available at citeseer.nj.nec.com/392196.html (2001)
3. Keane, J.: Trust based dynamic source routing in mobile ad hoc networks. msc. thesis, trinity college dublin. (2002)
4. Xiong, L., Liu, L.: Building trust in decentralized peer-to-peer electronic communities. In: Fifth International Conference on Electronic Commerce Research (ICECR-5), Canada. (2002)
5. Aberer, K., Despotovic, Z.: Managing trust in a peer-2-peer information system. In: CIKM. (2001) 310–317
6. Awerbuch, B., Holmer, D., Nita-Rotaru, C., Rubens, H.: An on-demand secure routing protocol resilient to byzantine failures (2002)
7. Hu, Y., Perrig, A., Johnson, D.: Ariadne: A secure on-demand routing protocol for ad hoc networks (2002)
8. Hu, Y.C., Johnson, D.B., Perrig, A.: SEAD: Secure efficient distance vector routing in mobile wireless ad hoc networks. In: Fourth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '02). (2002) 3–13
9. Jøsang, A.: The right type of trust for distributed systems. In: C. Meadows, editor, Proc. of the 1996 New Security Paradigms Workshop. ACM, 1996. (1996)
10. Schollmeier, R., Gruber, I., Finkenzeller, M.: Routing in mobile ad hoc and peer-to-peer networks. a comparison (2002)
11. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. available as tech. Rep., computer science department, stanford university (1998)
12. Rowstron, A., Druschel, P.: Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In: IFIP/ACM International Conference on Distributed Systems Platforms (Middleware). (2001) 329–350
13. Carbone, M., Danvy, O., Damgaard, I., Krukow, K., Møller, A., Nielsen, J.B., Nielsen, M.: A model for trust (2002) EU Project SECURE IST-2001-32486, Deliverable 1.1.

A Subjective Logic Operator Definitions

Definition 1 (Bayesian Consensus). *Let $\omega^A = (b^A, d^A, u^A)$ and $\omega^B = (b^B, d^B, u^B)$ be opinions respectively held by prinicpals A and B about the same proposition. Let $\omega^{A,B} = (b^{A,B}, d^{A,B}, u^{A,B})$ be the opinion such that*

$$\begin{aligned} b^{A,B} &= (b^A u^B + b^B u^A)/\kappa \\ d^{A,B} &= (d^A u^B + d^B u^A)/\kappa \\ u^{A,B} &= (u^A u^B)/\kappa \end{aligned}$$

where $\kappa = 1 - (1 - u^A)(1 - u^B) = u^A + u^B - u^A u^B$ such that $\kappa \neq 0$ (i.e. u^A and u^B cannot both be 0, otherwise one is trying to combine dogmatic opinions which leave no room for uncertainty). Then $\omega^{A,B} = \omega^A \oplus \omega^B$ is called the Bayesian consensus between ω^A and ω^B. Furthermore, $\oplus$ is commutative and associative.

Definition 2 (Discounting). *Let* $\omega_B^A = (b_B^A, d_B^A, u_B^A)$ *be principal A's opinion about principal B, and let* $\omega_p^B = (b_p^B, d_p^B, u_p^B)$ *be B's opinion about some proposition p. Let* $\omega_p^{A:B} = (b_p^{A:B}, d_p^{A:B}, u_p^{A:B})$ *be the opinion such that*

$$\begin{aligned} b_p^{A:B} &= b_B^A b_p^B \\ d_p^{A:B} &= b_B^A d_p^B \\ u_p^{A:B} &= d_B^A + u_B^A + b_B^A u_p^B \end{aligned}$$

Then $\omega^{A:B} = \omega^A \otimes \omega^B$ *is called the discounting of* ω_p^B *by* ω_B^A*. Furthermore,* $\otimes$ *is associative but non-commutative.*

Definition 3 (Conjunction). *Let* $\omega_p = (b_p, d_p, u_p)$ *and* $\omega_q = (b_q, d_q, u_q)$ *be a principal's opinions about two distinct propositions* p, q*. Let* $\omega_{p\wedge q} = (b_{p\wedge q}, d_{p\wedge q}, u_{p\wedge q})$ *be the opinion such that*

$$\begin{aligned} b_{p\wedge q} &= b_p b_q \\ d_{p\wedge q} &= d_p + d_q - d_p d_q \\ u_{p\wedge q} &= b_p u_q + u_q b_q + u_p u_q \end{aligned}$$

Then $\omega_{p\wedge q} = \omega_p \wedge \omega_q$ *is called the conjunction of* ω_p *and* ω_q*. As expected,* $\wedge$ *commutes and associates.*