An economic perspective of message-dropping attacks in peer-to-peer overlays

Kevin W Hamlen*[1] and William Hamlen[2]

[1]Computer Science Department, The University of Texas at Dallas, USA
[2]Department of Finance and Managerial Economics, School of Management, State University of New York at Buffalo, USA

Email: Kevin W Hamlen*- hamlen@utdallas.edu; William Hamlen - mgthamle@buffalo.edu;

*Corresponding author

Abstract

Peer-to-peer networks have many advantageous security properties, including decentralization, natural load-balancing, and data replication. However, one disadvantage of decentralization is its exclusion of any central authority who can detect and evict malicious peers from the network. It is therefore relatively easy to sustain distributed denial-of-service attacks against these networks; malicious peers simply join the network and fail to forward messages.

This article shows that peer-to-peer message-dropping attacks can be understood in terms of a well-established category of economic theory: the theory of the *second best*. In particular, peers who wish to continue service during an attack seek a second best solution to a utility optimization problem. This insight reveals useful connections between economic literature on the second best and computer science literature on peer-to-peer security. To illustrate, we derive and test an economics-inspired modification to the Chord peer-to-peer routing protocol that improves network reliability during message-dropping attacks. Under simulation, networks using the modified protocol achieve a 50% increase in message deliveries for certain realistic attack scenarios.

1 Introduction

Peer-to-peer networks are an increasingly popular vehicle for highly fault-tolerant, light-weight, and low-cost distributed computing across heterogeneous hardware. Cloud computing [1], digital music exchange [2], digital libraries [3], secure data management systems [4], and bioinformatics databases [5] are just a few of the venues where this technology is being used today. Unlike traditional networks, peer-to-peer networks lack any centralized server; every agent acts as both server and client. This provides a natural resistance to many attacks, since adversaries must compromise a large number of peers instead of just a few central servers to corrupt data integrity or disrupt its availability.

Unfortunately, decentralization can leave peer-to-peer networks vulnerable to a different sort of denial-of-service attack wherein malicious agents join the network and misroute or drop messages, thereby disrupting communication. Constraining this *message-dropping* behavior is difficult since victims typically learn only that their messages weren't delivered, not who was at fault. Even when malicious behavior is localized to a particular peer, there is no central authority who can evict the peer from the network, leaving it free to perpetuate the attack. A small number of malicious agents can amplify their message-dropping with a *Sybil* *attack* [6], in which each joins the network many times under various false identities to occupy a greater percentage of the overlay space. Since dropping messages is computationally inexpensive for the attacker, a relatively small number of attackers can significantly disrupt overlay traffic.

We gain insight into this problem by connecting it to economic theories of the *second best* [7] (cf., [8]). Second best economic solutions are required when one or more units contravene the first best solution, and there is no way to remove the misbehaving units from the system. The second best solution incorporates (takes as given) the misbehavior of the deviate units to obtain a second best optimum. In the case of peer-to-peer networks, we find that application of second best theory to the secure routing problem yields recommendations that advise non-malicious peers on how to route messages optimally given observed misbehavior of their neighbors. By taking the network topology as given, such recommendations can be implemented atop any given topology to achieve better performance during a message-dropping attack.

Adopting the *piecemeal approach* of Davis and Whinston [9] further allows each peer to optimize its own behavior with only local knowledge, and without requiring cooperation from other peers. Thus, the recommendations can be implemented fully automatically by individual peer clients without a central authority that has global knowledge, and without requiring other agents to adopt the new protocol.

We show how these theoretical insights can be used in practice by deriving a simple modification to the Chord peer-to-peer networking protocol [10] that improves overlay performance during message-dropping attacks. When faced with message-dropping agents, peers approximate a second best alternative to optimal overlay routes. Simulation of our modified protocol shows that message delivery rates improve by over 50% with a corresponding increase in overall system utility under realistic attack scenarios. The modified protocol is simple to implement since it does not require any change to the network topology and uses only information that is already available in a standard Chord network.

We begin by summarizing related work in §2. Section 3 reviews the economic theory of the second best and argues its relevance to peer-to-peer systems. The individual and social welfare perspectives are each elaborated in §4, yielding a general framework for expressing peer-to-peer message-dropping attacks as second best utility optimization problems. The framework is general enough to include many different definitions of utility, including those that incorporate both reliability and risk. Section 5 applies this framework to Chord, showing its effectiveness in resisting both non-coordinated and coordinated, distributed, message-dropping attacks. Finally, §6 concludes with a summary and suggestions for future work.

2 Related Work

Over the past decade there has been an explosion of research devoted to peer-to-peer security (cf., [11]). Relevant issues include robust search [12], pollution prevention (i.e., inhibiting the spread of unwanted objects), secure data storage, and data confidentiality. Most work on these assumes a secure routing framework that facilitates reliable, robust communication between peers. For example, reputation-based trust managers such as EigenTrust [13], Credence [14], and Penny [15], aggregate local reputation information whose exchange requires a secure routing framework.

Secure routing is divisible into three sub-problems [16]: secure identifier assignment, routing table maintenance, and message forwarding. Secure identifiers prevent attackers from misrepresenting or abusing their overlay positions to intercept more than their share of traffic. Secure routing tables maintain peers' connections to appropriate sets of neighbor peers. Finally, secure message forwarding delivers messages using the secure identifiers and routing tables. This last sub-problem is the focus of our work.

Secure message forwarding has been studied from an economic perspective in the context of *selfish routing* [17]. Selfish routers enjoy the message-forwarding services provided by fellow peers, but fail to forward the messages of others. Unlike message-droppers, selfish routers desire services from the network. This has led to incentive-based solutions such as pricing networks [18], negotiating contracts for mutual message delivery [19], rewarding message delivery with increased reputation [20], and rewarding message delivery with increased quality of service [21]. In contrast, message-droppers only desire to disrupt service, making incentivization inapplicable.

Economists have also studied peer-to-peer networks in the context of *free-riding*. Free-riding peers obtain shared resources from the network but fail to share their own resources. Krishnan et al. [22] observe that the shared resources resemble *public goods* [23]—an insight that has generated a growing body of work devoted to incentivizing free-riding peers to share [24]. However, unlike public goods, the connectivity that is threatened during a message-dropping attack is not equally available to all peers; it is influenced by the position of the attackers relative to each peer. This leads us to a different economic model for message-dropping.

Message forwarding in peer-to-peer networks can be divided into protocols for structured and unstructured networks. Structured networks route queries using a multi-hop protocol that passes the message from peer to peer. Examples include CAN [25], Pastry [26], Chord [10], and Tapestry [27]. Unstructured peer-to-peer networks like Gnutella broadcast queries using a multicast protocol that floods the query to all peers within a given radius. The unstructured approach tends to be more robust against message-dropping because of its high redundancy, but it does not always scale well [28]. Although our work is potentially applicable to unstructured topologies, we focus on structured ones since message-dropping tends to be a more significant threat in those contexts.

A large body of prior work improves the robustness of structured networks by augmenting their topologies with redundant routing paths (e.g., [29–31]). Route redundancy increases the probability that at least one message replica reaches its destination. Such topologies can be improved further via adaptive techniques that adjust the topology dynamically in response to observed failures (e.g., [32, 33]) or based on the interaction history of peers (e.g., [34, 35]). In contrast to these approaches, our work adapts routing flow rates without modifying the topology. Our work therefore complements the above by optimizing routing behavior once a (possibly dynamic) topology is chosen.

Adaptive techniques, including ours, require a means of detecting message-dropping behavior. Section 5 adopts a sampling approach in which non-malicious peers test for malicious behavior by periodically sending probe messages. Such sampling has been effectively used in unstructured networks to detect and identify message-droppers [36]. Introduction graphs [37] and flow rate histories [38] can further help to identify and isolate malicious peers, but sampling has the advantage of being easy to implement atop networks that do not collect this extra information.

3 Peer-to-peer and the Second Best

3.1 Overview of the Second Best

The theory of the second best has had a significant impact on economics throughout the past half-century, being applied to such subjects as health care, antitrust, and international trade [8]. In the case of peer-to-peer overlays, the original design of the overlay topology can be thought of as a first best solution that optimizes message delivery when all peers are non-malicious. Adapting individual peer routing behavior to the presence of malicious peers can therefore be viewed as a second best solution to this optimization problem.

Economic theory generally begins with economic subunits such as consumers, households, or firms that attempt to optimize their own objective functions subject to various constraints. These economic subunits work within a larger *administrative perspective* that is defined by the set of rules governing the system and the degree of decentralization. This leads to a larger *fundamental problem* first formulated by Nobel laureate Paul Samuelson [39], who observed that every society acts as if it is attempting to maximize a (known or unknown) *social welfare function* subject to constraints.[a] The welfare function is so named because it incorporates the goals of all economic subunits. For any viable system, the optimal conditions derived from the fundamental problem must also be consistent with the optimal behavior of the economic subunits.

Following Samuelson and the later literature [7], we write this problem in its most general form as:

$$\text{maximize } F(x_1, \ldots, x_n) \text{ subject to } G(x_1, \ldots, x_n) = 0 \tag{1}$$

The objective function is F, the binding constraint is G, and the choice variables are x_i with $i \in 1..n$. In a decentralized, consumer-oriented economy these choice variables are selected by independent economic subunits (e.g., consumers). The objective function is related to the maximization of the consumers' objective functions (*utility functions*) subject to constraints reflecting consumer income limitations. In a peer-to-peer network the choice variables are selected by individual peers and the constraints are those imposed by the routing protocol and the overlay topology.

The first best solution to the problem formalized by Equation 1 has the following necessary conditions for an optimal solution [7]:

$$F_i + \lambda G_i = 0 \quad \forall i \in 1..n \tag{2}$$

where F_i and G_i are partial first derivatives and λ is a Lagrange multiplier[b] associated with the constraint in Equation 1. To obtain the first best solution, individual subunits solve their respective constrained optimization problems, and these must be consistent with the necessary conditions given by Equation 2.

Samuelson was the first to recognize that when some economic subunits fail to adhere to Equation 2, the ramifications extend beyond that deviate sector:

> *"First, what is the best procedure if for some reason a number of optimum conditions are not realized? What shall we do about the remaining ones which are in our power? Shall we argue that 'two wrongs do not make a right' and attempt to satisfy those we can? Or is it possible that failure of a number of conditions necessitates modifying the rest? Clearly the latter alternative is the correct one."* [39, p. 252]

In other words, if the first best solution is unobtainable because one sector behaves suboptimally, then the second best solution does not necessarily imply that the remaining sectors should continue to satisfy their first best optimization conditions; many subunits may need to change their optimal behaviors.

Lipsey and Lancaster [7] formalized the second best solution with an additional constraint containing the first-order behavior of the deviating sector. We use the concise presentation of Henderson and Quandt [40, p. 316], who append the following constraints to the problem given by Equation 1:

$$F_j = h_j G_j, \; h_j \neq \lambda, \; h_j \neq 0 \tag{3}$$

The second best optimal condition that replaces Equation 2 is[c]

$$F_i + \lambda G_i + \mu(F_{ji} - h_j G_{ji}) = 0 \quad \forall i \in 1..n \tag{4}$$

where F_{ji} and G_{ji} are the partial cross-derivatives between subunits i and j and μ is the new Lagrange variable associated with constraint 3. Thus, unless the cross-derivatives are zero, each sector $i \in 1..n$ has optimal behavior that deviates from Equation 2.

Negishi [41] (cf., [42]) further found that under standard economic assumptions of a competitive, decentralized economy (viz., concave functions and at least one interior solution) it is as if society seeks to maximize a specific social welfare function expressible as a weighted sum of the consumer utility functions $U_i(x_i)$. The final form of the society's objective function is therefore

$$F(x_1, \ldots, x_n) = \sum_{i=1}^{n} \alpha_i U_i(x_i) \tag{5}$$

The coefficients α_i are the final weights given by the implicit social welfare to the utility of the economic subunits.[d]

These weights have a specific economic interpretation: Each consumer or household's weight is the reciprocal of its marginal utility of income. Thus, if a wealthy family places relatively low value on the last dollar earned, it receives a relatively high weight in the objective function of the society. In essence, the system tends to weight the more successful households (in terms of income).

We find a similar result in the case of peer-to-peer systems. During a message-dropping attack, non-malicious peers can be defined as those that derive utility from message deliveries. This leads to a social welfare (administrative objective) function that weights peers by the reciprocal of their relative reliability—i.e., their message delivery success rates. When all weights are equal (a standard assumption when applying utility theory to information systems problems [43]), this corresponds to the special case where all peers in the network are equally reliable.

Davis and Whinston [9] used Negishi's result to provide a piecemeal approach to the second best problem. They concluded that the difficulties of implementing a complex set of second best conditions is often overestimated since, based on Equation 5, many of the second derivatives are zero in a decentralized system and can therefore be dropped. In the context of a peer-to-peer network, this implies that it is possible to approach the second best solution without the need to share global information amongst peers. Instead, each peer adapts its own behavior individually, based on its own knowledge.

Based on the economic development above, there are two separate aspects of the problem: the mathematical method proposed by Lipsey and Lancaster, and the welfare and utility considerations discussed by Samuelson and Negishi. Section 3.2 examines the former aspect and §4.2 the latter.

3.2 Example Application of the Second Best to a Peer-to-peer System

Application of the second best approach to resisting message-dropping attacks in peer-to-peer overlays can be illustrated by a simple example. We begin by considering a single peer p of degree k in a network of size n (with $k \ll n$). Peer p periodically receives messages for other peers, each of which it must forward to one of its k neighbors. We assume a flat identifier space for this example, and that peer p may forward each message to any one of its k neighbors (though for any given message, certain neighbors are better positioned than others to deliver it).

The first best solution to this problem is the one addressed when designing the overlay topology, which assigns p an optimal set of k neighbors given an assumed distribution of message destinations seen by p. Specifically, we seek the k neighbors that minimize the absolute distance between each message's final destination and the nearest of the k neighbors to that destination:

$$\underset{N_1,\ldots,N_k}{\text{minimize}} \sum_{i=1}^{k} \int_{(N_i+N_{i-1})/2}^{(N_i+N_{i+1})/2} |x - N_i|\, D_p(x)\, dx \tag{6}$$

where N_i $\forall i \in 1..k$ are the desired neighbor identifiers in ascending order, $N_0 = 0$ and $N_{k+1} = n$ are the limits of the identifier space, and D_p is the probability density of message destinations seen by p. For example, if D_p is a uniform distribution, the optimal first order conditions (derived by setting the derivative of Equation 6 with respect to N_i to zero) are:

$$N_i = \tfrac{1}{2}(N_{i-1} + N_{i+1}) \quad \forall i \in 1..k \tag{7}$$

That is, a uniformly distributed set of random messages is delivered most effectively when the k neighbors of p have identifiers that are evenly spaced along the interval $[0, n)$.

The above implicitly assumes that peers behave optimally, forwarding each message to the neighbor closest to its destination. A second best solution is needed when some peers behave suboptimally and there is no way (short of centralizing the system) to force optimal behavior. One option in this case is to implement a new first best solution, but this requires global information about the overlay topology, which isn't typically available to peers once the network has been deployed and malicious behavior becomes evident. The second best solution takes the topology as given and re-solves the optimization problem to obtain a new recommended optimal behavior for peer p given the suboptimal behavior of its neighbor(s).

For example, suppose that peer p discovers that of the a_m messages it forwarded to peer $m \in 1..k$ during some sampling period, only a_m^* of them were ultimately delivered to their final destinations. In the first best solution given by Equation 7, peer p forwards an average of w/k of its messages to each of its neighbors, where w is the total number of messages; thus, when $a_m^* < w/k$, neighbor m is behaving suboptimally (possibly due to the suboptimal behavior of its neighbors). The second best solution with respect to peer m's actions requires appending the following constraint to Equation 6, with new Lagrange multiplier β:

$$\beta(a_m^* - w/k) \geq 0 \tag{8}$$

When not all messages that peer p forwards to peer m are ultimately delivered (i.e., $a_m < w/k$), we obtain a different optimal identifier value for the neighbor on each side of neighbor m (obtained by summing Equations 6 and 8, setting the derivative to zero, and solving for the unknowns):

$$N^*_{m-1} = \tfrac{1}{2}(N_{m-2} + N_m) + \tfrac{1}{2}\beta \tag{9}$$

$$N^*_{m+1} = \tfrac{1}{2}(N_{m+2} + N_m) - \tfrac{1}{2}\beta \tag{10}$$

Even though peer p cannot change the identifiers of its neighbors, it should forward its messages as if neighbor $m-1$ had identifier N^*_{m-1} and neighbor $m+1$ had identifier N^*_{m+1}. As peer m's reliability decreases, β increases and N^*_{m-1} and N^*_{m+1} approach N_m. Peer p therefore forwards fewer messages to neighbor m since fewer destinations are closer to N_m than to N^*_{m-1} or N^*_{m+1}. In the limiting case where $a^*_m = 0$ (i.e., neighbor m is completely unreliable), peer p only forwards to m those messages whose final destinations are m itself.

We next generalize this approach to a larger class of topologies and message-dropping attacks, and we examine the second best approach from both the perspective of individual peers and that of the system as a whole.

4 Attack Resistance As Utility Optimization

Economic formulations of the second best typically have two aspects: an individual perspective in which individuals in the society seek to maximize their own utility subject to individual constraints, and an administrative perspective in which the society as a whole acts as if it is seeking to maximize an objective function subject to administrative constraints. In this section we develop the each of these aspects as they relate to message-dropping attacks in peer-to-peer overlays. The analysis of the individual perspective yields a piecemeal approach [9] to resisting message-dropping attacks, wherein each peer individually adjusts its optimal behavior to account for the suboptimal behavior of its immediate neighbors. The analysis of the administrative perspective yields a measure of the network's success in resisting message-dropping attacks, providing a means to evaluate defense effectiveness. We initially consider only peer reliability; risk is added in §4.3.

4.1 The Individual Perspective

The optimization problem presented in §3.2 can be cast in the more general framework of a utility maximization problem. Following the economic literature, we assume that all decision-makers seek to maximize their own utility when making choices. The utility $U_i(P_i)$ of a non-malicious peer $i \in 1..n$ is a function of its message-delivering reliability $P_i \in [0, 1]$. Following the standard assumptions on utility functions, we assume that utility increases with reliability but at a decreasing rate:

$$\frac{dU_i}{dP_i} > 0 \quad \text{and} \quad \frac{d^2U_i}{dP_i^2} < 0 \qquad \forall i \in 1..n \tag{11}$$

The first derivative of utility is the *marginal utility* and the second assumption above is the *law of diminishing marginal utility.*

Since each peer in the overlay forwards messages through its neighbors, its reliability P_i can be expressed in terms of the reliabilities of its $k \leq n$ neighbors:

$$P_i = \sum_{j=1}^{k} w_{ij} R_{ij}(w_{ij}) \quad \text{and} \quad \sum_{j=1}^{k} w_{ij} = 1 \quad \forall i \in 1..n \tag{12}$$

where $w_{ij} \in [0, 1]$ is the relative share of its messages that peer i forwards to neighbor j, and $R_{ij}(w_{ij})$ is the reliability that peer i estimates for neighbor j. R_{ij} is a function of w_{ij} because the observed reliability of neighbor j typically varies with the share of messages it receives from i. As j receives a larger share, a greater portion of their destinations are farther from j, making those messages harder for j to deliver. Thus, $\partial R_{ij}/\partial w_{ij} \leq 0$.

The optimization problem confronting each peer i is that of maximizing utility subject to the constraint above. Written in the Lagrange (Kuhn-Tucker [44]) format it is:

$$\underset{w_{i1},\ldots,w_{ik}}{\text{maximize}} \; U_i(P_i) + \phi_i \left(1 - \sum_{j=1}^{k} w_{ij}\right) + \lambda_i \left(P_i - \sum_{j=1}^{k} w_{ij} R_{ij}(w_{ij})\right) \tag{13}$$

where ϕ_i and λ_i are the Lagrange multipliers.

For example, adopting reliability function $R_{ij}(w_{ij}) = d_{ij} w_{ij}^{-c}$ (generalizing the example in §3.2 in which $d_{ij} = 1$ and $c = 1$), using the natural logarithmic utility function $U_i = \log$, and solving the resulting Kuhn-Tucker conditions (see the appendix), we find that if peer i deems neighbor m to be a fraction $\varepsilon \in [0, 1]$ less reliable than neighbor j (i.e., $R_{im}(w) = \varepsilon R_{ij}(w)\ \forall w \in [0, 1]$), then peer i's optimal relative use of m compared to j is

$$w_{im} = \sqrt[c]{\varepsilon}\, w_{ij} \tag{14}$$

Hence, we advise peer i to use neighbor m less than it uses neighbor j by a factor of $\sqrt[c]{\varepsilon}$.

4.2 The Administrative Perspective

In this section we derive a measure of the social welfare (i.e., overall system utility) of a peer-to-peer system using the results from the previous sections. This provides a general measure of the performance of the system during a message-dropping attack in terms of peer utilities. In particular, networks that attain higher social welfare can be characterized as more robust against message-dropping attacks. Section 4.3 illustrates the generality of this metric by showing how it can incorporate risk as well as reliability, and §5 uses this measure to evaluate the performance of our method when implemented in an actual peer-to-peer network.

The peer-to-peer system acts as though it has an administrator who seeks to maximize some vector $(U_1(\hat{P}_1), \ldots, U_n(\hat{P}_n))$ of the utility functions of all peers in the system [41, 42], where $\hat{P}_i = P_i / \sum_{j=1}^{n} P_j$ denotes the *relative reliability* of peer i. Such an administrator might not actually exist, but the system behaves as if it does. Recall that this equates to maximizing the weighted sum in Equation 5. The administrator's optimization problem can therefore be written as:

$$\underset{\hat{P}_1,\ldots,\hat{P}_n}{\text{maximize}} \; \sum_{i=1}^{n} \alpha_i U_i(\hat{P}_i) \;\; \text{subject to} \;\; \sum_{i=1}^{n} \hat{P}_i = 1 \tag{15}$$

Thus, the system behaves as if an administrator determines the relative reliabilities of all members, knowing that the selection must ultimately account for the various utility functions of the members as well as their respective optimization behaviors.

A peer-to-peer network's success in resisting a message-dropping attack can therefore be measured by computing the objective function in Equation 15. Combining the optimality result of the individual behavior described by Equation 13 with the administrative optimal behavior described by Equation 15, we find (see the appendix) that consistency between the individual behavior and the administrator's optimal solution requires that $\alpha_i = (\delta/\phi_i)\hat{P}_i$, where δ is the Lagrange multiplier associated with the constraint in Equation 15 and $\phi_i = (\partial U_i / \partial \hat{P}_i)\hat{P}_i$. This implies that weight α_i is directly related (up to common factor δ) to peer i's relative reliability and inversely related to its marginal utility. Using $U_i = \log$ yields $\phi_i = 1$ for all $i \in 1..n$. Since constant factor δ has no effect on optimization problem 15, this simplifies to $\alpha_i = \hat{P}_i$.

Thus, a peer-to-peer network's success in resisting a message-dropping attack is measurable via the following social welfare function:

$$\sum_{i=1}^{n} \hat{P}_i \log(\hat{P}_i) \tag{16}$$

where $\hat{P}_i$ is the fraction of the messages forwarded by peer i that were ultimately delivered to their destinations. This can be interpreted as a familiar result from information theory. It is the negation of the *Shannon entropy* of the peer-to-peer system, and the administrative problem therefore reduces to the problem of minimizing the entropy subject to the constraints. When there are no additional constraints, the optimal solution is obviously one in which all peers are equally reliable—i.e., $\hat{P}_i = \hat{P}_j \; \forall i, j \in 1..n$. In the case where there are additional constraints (e.g., some peers are malicious and therefore have constrained reliabilities), the optimal solution is non-trivial, as we see in §5.

4.3 Risk

In coordinated, distributed, message-dropping attacks, malicious peers vary their behavior over time, dropping some but not all messages they receive in an effort to evade detection. Malicious peers may even coordinate their behavior changes so as to keep each individual peer's reliability relatively high while keeping overall availability of network services low. Peer reliability alone is not an adequate measure of malicious behavior during such an attack; one must also consider variance or *risk*.

Following the work of von Neumann and Morgenstern [45], there has been strong agreement that individuals tend to maximize *expected utility* in the presence of risk. The negative exponential utility function is most commonly used to examine effects of both mean and variance in such contexts:

$$U_i(P_i) = -\exp(-a_i P_i) \quad \forall i \in 1..n \tag{17}$$

where parameter $a_i > 0$ is a measure of the *risk aversion* of peer i. (Equation 17 satisfies the assumptions about utility given by Equation 11, and approximates the natural logarithmic utility function used in §4 [46].[e]) When random variable P_i ($i \in 1..n$) has an approximately normal distribution with mean $E[P_i]$ and variance $Var[P_i]$, the expected utility is

$$E[U_i(P_i)] = -\exp\left(-a_i E[P_i] + \tfrac{1}{2} a_i^2 \, Var[P_i]\right) \tag{18}$$

With risk, the administrative perspective differs from §4.2. There, utility U_i is a function of only one argument P_i and increases monotonically with P_i; but with risk, expected utility is a function of two arguments $E[P_i]$ and $Var[P_i]$ and is concave. Assuming there is at least one interior, feasible solution, we can again write the administrative objective function as a weighted sum of the expected utilities of the peers. Substituting these into Equation 5, the administrative optimization problem becomes

$$\text{maximize} \sum_{i=1}^{n} \alpha_i E[U_i(\hat{P}_i)] \tag{19}$$

$$\text{subject to} \sum_{i=1}^{n} E[\hat{P}_i] = 1 \text{ and } \sum_{i,j=1}^{n} Cov[\hat{P}_i, \hat{P}_j] = 0 \tag{20}$$

where $\hat{P}_i$ is the relative reliability of peer i as defined in §4.2 and $Cov[X, Y]$ denotes the covariance of random variables X and Y.

The first best solution is obviously one in which all peers are invariably equally reliable. Since zero variation implies $E[\hat{P}_i] = \hat{P}_i$, this reduces to the same optimization problem as derived in §4.2; we therefore conclude that $\alpha_i = \hat{P}_i$ as before. When the variance is non-zero for some peers, we seek a second best solution. In that case the form of the social welfare function stays the same but is evaluated at the second best solution. Thus, $\alpha_i = \hat{P}_i$ in that case as well, and we conclude that social welfare can be measured by weighting each peer's expected utility by its relative reliability. The optimal conditions for Equations 19–20 are derived in the appendix.

We next consider how individual peers make optimal decisions in the presence of risk in a way that is consistent with the administrative perspective above. Using Equation 12, the expected reliability and variance of each peer $i \in 1..n$ is

$$E[P_i] = \sum_{j=1}^{k} w_{ij} E[P_j] \tag{21}$$

$$Var[P_i] = \sum_{j=1}^{k} w_{ij}^2 \, Var[P_j] + 2 \sum_{j=1}^{k} \sum_{h=j+1}^{k} w_{ij} w_{ih} \, Cov[P_j, P_h] \tag{22}$$

where $w_{ij} \in [0, 1]$ is again the relative share of messages that peer i forwards to neighbor j. Substituting these into the Kuhn-Tucker conditions for Equations 19–20 (see the appendix) yields the following system of linear equations that must be solved to find peer i's optimal relative use w_{ij} of each of its neighbors $j \in 1..k$:

$$\frac{1}{a_i}(E[P_b] - E[P_j]) = \sum_{h=1}^{k} w_{ih} (Cov[P_h, P_b] - Cov[P_h, P_j]) \tag{23}$$

Since the shares w_{ij} are all relative, we choose some arbitrary *benchmark neighbor* $b \in 1..k$ in terms of which peer i computes the other optimal shares.

One approach to implementing the above in actual peer-to-peer client software is with a linear constraint solver. The system of linear equations given by 23 can be expressed as a matrix computation of the form

$$(\mathbf{C} - \mathbf{D})\mathbf{w} = \mathbf{E} \tag{24}$$

where $\mathbf{C}_{hj} = Cov[P_h, P_j]$ is the covariance matrix; $\mathbf{D}_{hj} = Cov[P_h, P_b]$ if $h \neq b$ and $\mathbf{D}_{bj} = 0$ otherwise; $\mathbf{E}_j = (E[P_b] - E[P_j])/a_i$ if $j \neq b$ and $\mathbf{E}_b = 1$ otherwise; and $\mathbf{w}$ is the unknown length-k vector of relative shares for which the system must be solved. The problem can be further simplified if we assume that under normal conditions most of the covariance terms for any individual peer are likely to vanish. We can therefore approximate the above solution by setting them to zero, which simplifies to the following formula for computing the optimal relative use of neighbor $j \in 1..k$ by peer $i \in 1..n$:

$$w_{ij} = \frac{E[P_j] - E[P_b]}{a_i \, Var[P_j]} + \frac{Var[P_b]}{Var[P_j]} w_{ib} \tag{25}$$

Peers that use Equation 25 to guide their relative usage of their neighbors tend to maximize expected reliability and minimize risk as they forward messages. This can have some interesting ramifications for peer behavior. For example, depending on their risk aversion a_i, they may sometimes forward messages through less reliable peers to avoid a more reliable but much riskier one. Risk-averse peers also tend to diversify their message-forwarding behavior similar to an investor's diversification of a portfolio. This can result in a better outcome when resending a dropped message since there is a higher chance that the message will not take the same route to its destination even when the overlay topology remains static.

5 Implementation

To put our approach into practice, we implemented it within a Chord network [10]. We begin with a review of Chord's overlay structure and routing protocol in §5.1. Section 5.2 then formulates the Chord protocol as a utility optimization problem using the second best. Finally, §5.3 describes our experimental methodology and results.

5.1 The Chord Protocol

Chord [10] is a structured peer-to-peer protocol with a ring-shaped overlay. Each peer's ring position is defined by an integer *identifier*. Identifiers are derived via secure hash functions so that attackers cannot easily choose their positions. Each peer is directly connected to $k = \lfloor \log_2 n \rfloor$ neighbors, where n is the size of the identifier space. For example, in a Chord network that can accommodate 2^{160} peers, each peer has 160 neighbors.

The neighbor set of peer i is densest near i and thins farther away. Specifically, the jth neighbor of peer i is the peer whose identifier is closest to (but no less than) $(id_i + 2^{j-1}) \bmod n$ $(\forall j \in 1..k)$. Thus, peer i's first neighbor is its successor in the ring, each subsequent neighbor is approximately twice as far from i as its previous neighbor, and peer i's last neighbor is approximately halfway around the ring. To send a message to peer h, peer i forwards it to the neighbor whose identifier is closest to but no greater than h's identifier (modulo n). When all peers adhere to this protocol, messages are delivered to their final destinations in at most $O(\log_2 n)$ hops because each hop at least halves the distance from the message's current position to its destination. Without malicious peers, the topology is naturally load-balancing in that a uniform distribution of message sources and destinations tends to solicit equal relative use of each peer's k neighbors.

During a message-dropping attack, however, malicious peers drop the messages they receive instead of forwarding them. Since Chord is deterministic, a single malicious peer on the route from i to h can thereby prevent i from sending any messages to h until the topology changes (e.g., due to churn). With multiple attackers, the identifier assignment process tends to distribute attackers approximately uniformly across the identifier space. As a result, attackers can intercept a significant portion of the overlay traffic. For example, even when malicious peers comprise only 10% of the network they can intercept about 40% of the messages on average [47].

5.2 Applying the Second Best Approach to Chord

Peers can forward messages via different neighbors than the ones prescribed by the Chord protocol at the expense of longer message delivery paths. This flexibility allows peers to potentially improve message delivery rates in the presence of malicious peers via a second best routing strategy. Specifically, a peer can potentially forward each message to any neighbor between itself and the message's intended destination, not just the closest one to the destination.

However, this flexibility must be exercised in moderation to avoid unacceptably long routing paths, since forwarding messages in very small hops greatly increases the worst-case path length bound given in §5.1. For example, if each peer forwards messages to its nearest neighbor, the worst-case path length is $O(n)$, which is clearly unreasonable when $n \approx 2^{160}$. More generally, when peers forward messages to their rth-closest neighbors, the worst-case path length increases by a factor of $(r - \log_2(2^r - 1))^{-1}$. Hence, forwarding to the 2nd-closest neighbor multiplies the worst-case path length by a factor of about 2.4, and forwarding to the 3rd-closest multiplies it by a factor of over 5.

To keep the worst-case path length reasonable, we therefore modify the Chord protocol to allow (non-malicious) peers to forward each message only to the closest neighbor or 2nd-closest neighbor to the message's intended final destination. In our experiments we found that allowing peers to forward to other neighbors is seldom useful, since during a message-dropping attack greatly increased path lengths almost always include at least one malicious peer.

Given this restriction, the reliability P_i of any peer $i \in 1..n$ can be expressed in terms of the reliabilities of its neighbors as follows:

$$P_i = s_1 P_1 + \sum_{j=2}^{k} s_j \big(w_{ij} P_j + d(1 - w_{ij}) P_{j-1} \big)$$

where $s_j \in [0, 1]$ is the fraction of the messages seen by peer i whose destinations are closest to neighbor j, $w_{ij} \in [0, 1]$ is the share of those s_j messages that peer i chooses to forward to neighbor j (instead of to neighbor $j - 1$), and $d \in [0, 1]$ models a distance penalty for forwarding messages to the 2nd-closest neighbor instead of to the closest one. (In our implementation we used $d = 0.8$, but other values in the interval $[0.1, 0.9]$ performed similarly.) The expected value and variance of P_i are

$$E[P_i] = s_k w_{ik} E[P_k] + \sum_{j=1}^{k-1} \big(s_{j+1} d(1 - w_{i\,j+1}) + s_j w_{ij}\big) E[P_j] \tag{26}$$

$$\mathit{Var}[P_i] = s_k^2 w_{ik}^2\, \mathit{Var}[P_k] + \sum_{j=1}^{k-1} \big(s_{j+1} d(1 - w_{i\,j+1}) + s_j w_{ij}\big)^2\, \mathit{Var}[P_j] \tag{27}$$

respectively, for all $i \in 1..n$.

Individual utility is modeled by Equation 17 with the addition of a unit constant to force non-negative utilities [48]; hence, $U_i(P_i) = 1 - \exp(-a_i P_i)$. Solving the resulting optimization problem given by Equation 23 (and zeroing the covariance terms as in §4.3), yields a system of linear equations of the form $\mathbf{A} = \mathbf{0}$, where $\mathbf{A}$ is the $n \times k$ matrix defined by

$$\begin{aligned}\mathbf{A}_{ij} = -s_{ij}(E[P_{ij}] - dE[P_{i\,j-1}]) + as_{ij}\big[\big(s_{ij}w_{ij} + s_{i\,j+1}d(1{-}w_{i\,j+1})\big)\, \mathit{Var}[P_{ij}] - \\ d\big(s_{i\,j-1}w_{i\,j-1} + s_{ij}d(1{-}w_{ij})\big)\, \mathit{Var}[P_{i\,j-1}]\big]\end{aligned} \tag{28}$$

for all $j \in 2..k-1$, and for $j \in \{1, k\}$ we have $\mathbf{A}_{i1} = 0$ and

$$\mathbf{A}_{ik} = -s_{ik}(E[P_{ik}] - dE[P_{i\,k-1}]) + as_{ik}\big[s_{ik}w_{ik}\, \mathit{Var}(P_{ik}) - d\big(s_{i\,k-1}w_{i\,k-1} + s_{ik}d(1{-}w_{ik})\big)\, \mathit{Var}[P_{i\,k-1}]\big] \tag{29}$$

Fully solving the above system subject to constraint $w_{ij} \in [0, 1]$ requires a mixed integer programming algorithm. However, the implementation can be significantly simplified by approximating the solution iteratively. We used a quasi-Newtonian approximation obtained by computing the Hessian matrix of Equations 28–29, zeroing the cross-derivatives, and solving for $\mathbf{w}$. This yields the following rule for updating share w_{ij}:

$$\Delta w_{ij} = \frac{-(\mathbf{A}_{ij})}{(\mathbf{A}_{ij})^2 + a_i s_{ij}^2(\mathit{Var}[P_{ij}] + d^2\, \mathit{Var}[P_{i\,j-1}])}\psi \qquad \forall j \in 2..k \tag{30}$$

$$w_{i1} = 1 \tag{31}$$

where $\mathbf{A}$ is defined by Equations 28–29 and ψ controls the rate of convergence. (In our implementation we used $\psi = 1$.) Equation 31 reflects the inflexibility of traffic forwarded to a peer's first neighbor (since no neighbors fall between a peer and its first neighbor).

In summary, non-malicious peers in our modified system continuously adjust their relative usage of neighbors in small increments, based on the most current information available concerning neighbor reliability and riskiness. These adjustments are made so as to maximize reliability and minimize risk. That is, each peer optimizes its own expected utility subject to the constraints imposed by the routing protocol.

5.3 Experimental Results

To test our solution, we simulated a Chord network in which non-malicious peers maximize expected utility by adapting their relative use of their neighbors according to Equation 30. Malicious peers drop some or all messages they are asked to forward. To assess the network's success in resisting the attack, we computed the social welfare (Equation 16) that it attained over each simulation. We also measured the total percentage of messages that were successfully delivered. Each simulation involved sending a total of one million randomly generated messages through the overlay, and simulation results were averaged over 50 trials each.

We assume that senders learn whether their messages were ultimately delivered, but not who dropped undelivered messages. This is consistent with networks in which delivered messages solicit unforgeable, direct responses from recipients. For example, object lookups in Chord solicit a direct response that does not use

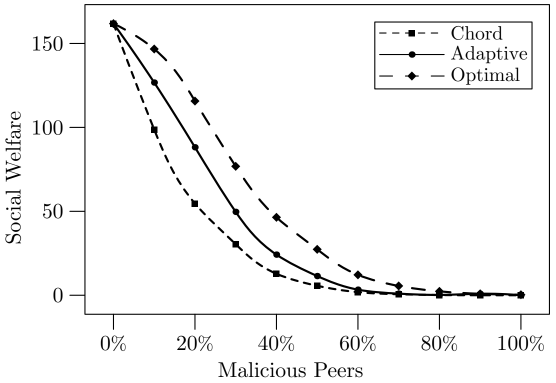

Figure 1: Social welfare attained during message-dropping attacks

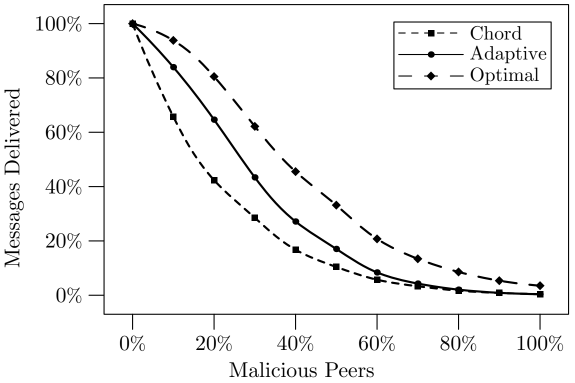

Figure 2: Message delivery success rates during message-dropping attacks

the overlay, and that can be authenticated via cryptographic message signing. This information allows each peer i to estimate a running mean $E[P_{ij}]$ and running variance[f] $Var[P_{ij}]$ for each neighbor $j \in 1..k$. Each peer also tracks its own relative usage s_{ij} of each neighbor.

Non-malicious peers begin the simulation with $w_{ij} = 1$ and $s_{ij} = 1$ for all $i \in 1..n$ and $j \in 1..k$. That is, each peer initially behaves as in a traditional Chord network, forwarding each message to the closest neighbor and using all neighbors approximately equally. At regular intervals, non-malicious peers modify w_{ij} according to Equation 30. (If w_{ij} rises above 1 or descends below 0, it is truncated down to 1 or up to 0, respectively.) In our simulation, peers recomputed w_{ij} after every 1000 messages they sent. We used a convergence rate of $\psi = 1$, a distance penalty of $d = 0.8$, and a risk aversion of $a_i = 1$ to strike a roughly even balance between reliability maximization and risk minimization.

Figure 1 shows the overall utilities (i.e., social welfare) attained by a traditional Chord network, an adaptive Chord network that uses our utility optimization procedure, and an optimal Chord network in which peers have global knowledge of the overlay topology and know the identifiers of all malicious peers. For each attack, malicious peers dropped all messages they received (other than the messages intended for themselves) and were roughly uniformly distributed throughout the overlay.

The curve for the optimal Chord network was computed by exhaustively deciding for each possible source-destination pair whether there exists a route through the overlay that does not include any malicious peers (subject to the constraint that non-malicious peers must not route messages farther away than their 2nd-closest neighbor to the message's intended destination). We simulated networks with up to 10K peers, but the number of peers did not influence any of our results (except that computing the optimal curve for very large networks was not feasible). The curves shown in Figure 1 are for a network with 256 peers.

As Figure 1 illustrates, once malicious peers comprise over half the network, our adaptive approach is unable to provide substantial improvement; however, we see significant gains in the more typical scenarios where malicious peers comprise 10%–30% of the network. Under those conditions the adaptive approach resulted in about a 60% increase in social welfare compared to a traditional Chord network—about 55% of the gain that was possible even with global information. Figure 2 verifies that this increase in social welfare translates to a corresponding increase in message deliveries. Message delivery success rates were increased by about 52% when malicious peers comprised 10%–30% of the network—about 58% of what was possible with global information.

We next considered a more sophisticated message-dropping attack in which the attackers vary their behavior over time in an effort to avoid detection. Malicious peers coordinate these behavior changes to keep each malicious peer's observed reliability relatively high while keeping overall network connectivity low. In our simulation, attackers chose their reliabilities from a normal distribution of mean 0.3 and variance

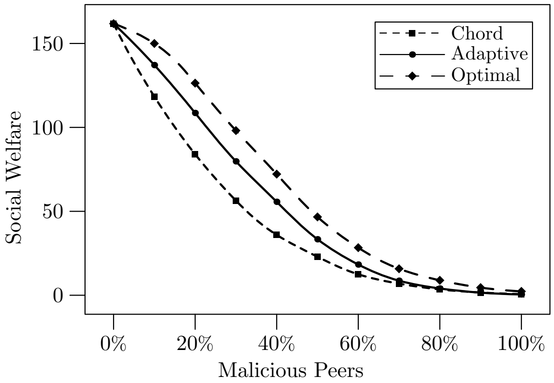

Figure 3: Social welfare attained during coordinated, distributed attacks

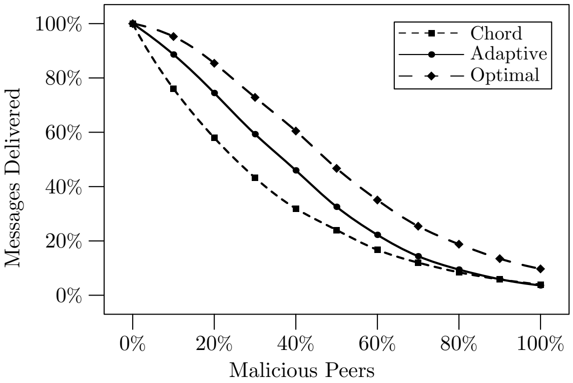

Figure 4: Message delivery success rates during coordinated, distributed attacks

0.12. A coordinated, distributed, denial-of-service attack of this kind can be quite effective against defense mechanisms that rely on average reliability as the sole indicator of maliciousness. Our protocol's inclusion of risk as a secondary indicator was therefore important for resisting this attack.

Figure 3 shows that our adaptive approach continues to provide effective recommendations to non-malicious peers during the attack. The three curves are closer than in the non-coordinated attack since malicious peers in this simulation deliver at least some messages. Nevertheless, the adaptive network is still able to achieve a 30% increase in social welfare (58% of what was possible with global information) when attackers comprise 10%–30% of the network. Likewise, the message delivery rates reported in Figure 4 show a 28% increase (60% of what was possible with global information). Most importantly, a point-by-point comparison reveals that introducing attacker behavior variations only *increased* social welfare and message deliveries compared to when they dropped all messages. Thus, behavior variations were ineffective when attacking Chord networks equipped with our adaptive routing protocol.

During testing, our utility optimization strategy demonstrated little sensitivity to parameter changes and implementation details. For example, different convergence rates ψ, different distance penalties d, different approximation methods for Equations 28–29, and different refresh rates for recomputing shares w_{ij} resulted in little or no change to the results reported here (except when parameters were set to extreme values). This seems to indicate that our method is does not require much manual tuning to perform well.

6 Conclusion

The theory of the second best has played a significant role over the past several decades in solving numerous important problems in economics. In this article we have shown that it also applies to the problem of resisting message-dropping attacks in peer-to-peer overlay networks. If one views the design of the underlying overlay topology as an optimization problem, the second best solution yields recommendations on how to make optimal use of that existing topology in the presence of malicious peers who drop messages.

This has implications both for individual peers and for the peer-to-peer system as a whole. For individual peers, the second best solution provides peer-specific recommendations on how to forward messages so as to maximize each peer's individual reliability and minimize its risk. For the system as a whole, it maximizes the overall objective of the system given the misbehavior of the attackers. We found that in this context the overall objective can be expressed as a weighted sum of the utility functions of the individual peers, where the weights are the relative reliabilities of the peers. When individual utility functions are standard Bernoulli logarithmic functions, this equates to minimizing the Shannon entropy of the peer-to-peer system.

As a practical application of our work, we solved the above optimization problem for the Chord peer-to-peer network protocol [10] and implemented it in a simulator. Non-malicious peers in our modified network forward messages according to the recommendations prescribed by the second best solution. Rather than compute the second best solution directly, each peer approximates it iteratively using an efficient quasi-Newtonian algorithm. We simulated simple message-dropping attacks in which attackers drop all messages, as well as coordinated, distributed message-dropping attacks in which attackers vary their behavior to avoid detection. The modified protocol achieves a 50% increase in message deliveries and a 60% increase in social welfare when malicious peers comprise about 20% of the network. Behavior variations are not effective as a means of disguising the attack; they only result in higher message delivery rates and higher social welfare in networks equipped with our adaptive protocol.

In future work we intend to apply our approach to other distributed computing paradigms, such as clouds. Tapestry [27] networks are more densely connected than Chord networks, incorporating extra routing links for improved fault tolerance. These extra links could provide more opportunities for second best optimization. CAN [25] poses interesting mathematical challenges for our method since it uses a multidimensional identifier space.

In addition, much prior work on adaptive overlay routing has focused on adapting the overlay topology in response to observed peer behavior and performance. Since our second best optimization approach takes the topology as given, it could be implemented atop one of these adaptive topologies. Future work should investigate the interaction between these two approaches.

Finally, we plan to investigate second best optimization approaches to protecting these networks from other forms of attacks, such as message misrouting, message integrity and confidentiality violations, and reputation mismanagement. These are all significant current-day threats to large, distributed data management systems, and would likely benefit from second best optimization.

Appendix

We here sketch derivations of the solutions of the three main optimization problems presented throughout the paper. We begin with the problem of individual utility optimization presented in §4.1. The necessary (Kuhn-Tucker) conditions for the problem given by Equation 13 are

$$\frac{\partial U_i}{\partial P_i} + \lambda_i \geq 0 \qquad P_i\left(\frac{\partial U_i}{\partial P_i} + \lambda_i\right) = 0 \tag{32}$$

$$\phi_i - \lambda_i d_{ij} w_{ij}^{-c}(1-c) \geq 0 \qquad w_{ij}(\phi_i - \lambda_i d_{ij} w_{ij}^{-c}(1-c)) = 0 \qquad \forall j \in 1..k \tag{33}$$

$$1 - \sum_{j=1}^{k} w_{ij} \geq 0 \qquad \phi_i\left(1 - \sum_{j=1}^{k} w_{ij}\right) = 0 \tag{34}$$

$$P_i - \sum_{j=1}^{k} d_{ij} w_{ij}^{(1-c)} \geq 0 \qquad \lambda_i\left(P_i - \sum_{j=1}^{k} d_{ij} w_{ij}^{(1-c)}\right) = 0 \tag{35}$$

for all $i \in 1..n$, where ϕ_i is the Lagrange multiplier associated with the constraint.

Equation 33 implies that $\phi_i > 0$ for any positive share $w_{ij} > 0$ with $d_{ij} > 0$, and that for any two positive shares $w_{ij}, w_{im} > 0$ and $d_{ij}, d_{im} > 0$ with $c > 0$ the ratio of shares is

$$\frac{w_{im}}{w_{ij}} = \left(\frac{d_{im}}{d_{ij}}\right)^{1/c}$$

The assumption in §4.1 that peer m is a factor ε less reliable than peer j implies that $d_{im}{=}\varepsilon d_{ij}$. This yields the result given by Equation 14.

We next consider the administrative optimization problem presented in §4.2. The optimal (Kuhn-Tucker) conditions associated with Equation 15 are[g]

$$\alpha_i \frac{\partial U_i}{\partial \hat{P}_i} - \delta \leq 0 \qquad \hat{P}_i \left(\alpha_i \frac{\partial U_i}{\partial \hat{P}_i} - \delta \right) = 0 \qquad \forall i \in 1..n \tag{36}$$

$$1 - \sum_{i=1}^{n} \hat{P}_i \geq 0 \qquad \delta \left(1 - \sum_{i=1}^{n} \hat{P}_i \right) = 0 \tag{37}$$

where δ is the Lagrange multiplier associated with the constraint in Equation 15. By combining the optimality result of the individual behavior given by Equations 32–35 with the optimal behavior from the administrative perspective given by Equations 36–37, we see that whenever $\hat{P}_i > 0$, consistency between the individual behavior and the administrator's optimal solution requires that

$$\alpha_i = \frac{\delta}{\phi_i} \hat{P}_i \quad \text{and} \quad \phi_i = \frac{\partial U_i}{\partial \hat{P}_i} \hat{P}_i$$

Section 4.2 describes how this result leads directly to the conclusion that $\alpha_i = \hat{P}_i$.

Finally, we consider the optimization problem given in §4.3, which introduces risk. The necessary (Kuhn-Tucker) conditions for the administrative optimization problem given by Equations 19–20 are

$$\alpha_i a_i E[U_i(\hat{P}_i)] + \delta \geq 0 \qquad E[\hat{P}_i](\alpha_i a_i E[U_i(\hat{P}_i)] + \delta) = 0 \tag{38}$$

$$\tfrac{1}{2}\alpha_i a_i^2 E[U_i(\hat{P}_i)]\, Var[\hat{P}_i] - \gamma \leq 0 \qquad \left(\tfrac{1}{2}\alpha_i a_i^2 E[U_i(\hat{P}_i)]\, Var[\hat{P}_i] - \gamma \right) Var[\hat{P}_i] = 0 \tag{39}$$

for all $i \in 1..n$, and

$$\delta \left(1 - \sum_{i=1}^{n} E[\hat{P}_i] \right) = 0 \tag{40}$$

where δ and γ are the Lagrange multipliers. Recall that $\alpha_i = \hat{P}_i$ (see §4.3); therefore the individual optimization problem in Equations 21–22 contributes the following additional (Kuhn-Tucker) conditions:

$$w_{ij} \left(\frac{\partial E[U_i]}{\partial w_{ij}} - \phi_i \right) = 0 \tag{41}$$

$$1 - \sum_{j=1}^{k} w_{ij} \geq 0 \qquad \phi_i \left(1 - \sum_{j=1}^{k} w_{ij} \right) = 0 \tag{42}$$

Equation 41 can be reexpressed as $w_{ij}(E[U_i]z_{ij} - \phi_i)$, where z_{ij} is defined by

$$z_{ij} = -a_i E[P_j] + a_i^2 \sum_{h=1}^{k} w_{ih}\, Cov[P_j, P_j]$$

In the above, ϕ_i is the Lagrange multiplier associated with the constraint requiring that each peer i's relative usage of its neighbors sums to 1. In most situations the constraint would be binding and $\phi_i > 0$. This complicates the solution for any single w_{ij}. We can, however, examine the ratio $E[U_i]z_{ij} = \phi_i$ to $E[U_i]z_{ib} = \phi_i$ for any two shares $w_{ij}, w_{ib} > 0$. The benchmark neighbor b is arbitrarily chosen by peer i as the standard by which its other neighbors are assessed. From the conditions above we see that $z_{ij} = z_{ib}$, leading to the system of linear equations given in Equation 23.

Competing Interests

The authors declare that they have no competing interests.

Authors' Contributions

KWH conceived of the research, implemented the Chord simulator, carried out the experiments, and analyzed the results. WH provided economic interpretations and solved the optimization problems. Both authors drafted and approved the final manuscript.

Acknowledgement

This work includes material supported by the National Science Foundation under award NSF-0959096. Any opinions or findings expressed are those of the authors and not of the NSF.

Notes

[a] This "as if" approach is attributed to Nobel laureate Milton Friedman [49], who recognized that participants need not know or understand a model for it to adequately explain an outcome. Kenneth Arrow [50], another Nobel laureate, later proved that there are limits on what any welfare function can achieve. [b] Samuelson [39] and others used the Lagrange method, which requires binding constraints. Subsequent work used the Kuhn-Tucker conditions [44], which allow non-binding constraints. [c] If $\lambda = h_j$ in Equation 4, the Lagrange multiplier μ is zero, re-attaining the first best solution. Only when $\lambda \neq h_j$ does the added constraint change the solution. [d] Hamlen and Hamlen [51] show that Negishi's solution, while important in defining the social welfare function, can be evaluated only at the equilibrium solution and therefore does not aid in obtaining the solution. [e] This follows from isoelasticity of natural logarithm [46]. [f] Neither running mean nor running variance require maintaining any message history. We computed running variance via $Var[P_{ij}] = E[P_{ij}^2] - E[P_{ij}]^2$. [g] Partial derivatives ∂ are used here to remind us that the administrator's objective function contains all P_i $\forall i \in 1..n$.

References

1. Marozzo F, Talia D, Trunfio P: **A Peer-to-peer Framework for Supporting MapReduce Applications in Dynamic Cloud Environments**. In *Cloud Computing: Principles, Systems and Applications*. Edited by Antonopoulos N, Gillam L, Springer 2010.
2. Alexander PJ: **Peer-to-Peer File Sharing: The Case of the Music Recording Industry**. *Review of Industrial Organization* 2002, **20**:151–161.
3. Amrou A, Maly K, Zubair M: **Freelib: Peer-to-peer-based Digital Libraries**. In *Proceedings of the 20th International Conference on Advanced Information Networking and Applications (AINA)*, Vienna, Austria 2006:9–14.
4. Hamlen KW, Thuraisingham B: **Secure Peer-to-peer Networks for Trusted Collaboration**. In *Proceedings of the 3rd International Conference on Collaborative Computing: Networking, Applications and Worksharing*, White Plains, New York 2007:58–63.
5. Chou PH, Ortega RB, Borriello G: **The Chinook Hardware/Software Co-synthesis System**. In *Proceedings of the 8th International Symposium on System Synthesis (ISSS)*, Cannes, France 1995:22–27.
6. Douceur JR, Donath JS: **The Sybil Attack**. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, Cambridge, MA 2002:251–260.
7. Lipsey RG, Lancaster K: **The General Theory of Second Best**. *The Review of Economic Studies* 1956, **24**:11–32.
8. Hamlen W: **Gleanings into the Second Best Debate**. *Business Modelling: Multidisciplinary Approaches—Economics, Operational, and Information Systems Perspectives (In Honor of Andrew B. Whinston)* 2002.
9. Davis OA, Whinston AB: **Piecemeal Policy of the Second Best**. *The Review of Economic Studies* 1967, **34**(3):323–331.
10. Stoica I, Morris R, Karger D, Kaashoek MF, Balakrishnan H: **Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications**. *IEEE/ACM Transactions on Networking* 2003, **11**:17–32.
11. Wallach DS: **A Survey of Peer-to-Peer Security Issues**. In *Proceedings of the International Symposium on Software Security—Theories and Systems, Mext-NSF-JSPS (ISSS)*, Tokyo, Japan 2002:42–57.
12. Risson J, Moors T: **Survey of Research Towards Robust Peer-to-peer Networks: Search Methods**. *Computer Networks* 2006, **50**(17):3485–3521.
13. Kamvar SD, Schlosser MT, Garcia-molina H: **The EigenTrust Algorithm for Reputation Management in P2P Networks**. In *Proceedings of the 12th International World Wide Web Conference (WWW)*, Budapest, Hungary 2003:640–651.

14. Walsh K, Sirer EG: **Experience with an Object Reputation System for Peer-to-Peer Filesharing**. In *Proceedings of the 3rd Symposium on Networked System Design and Implementation (NSDI)*, San Jose, California 2006.

15. Tsybulnik N, Hamlen KW, Thuraisingham B: **Centralized Security Labels in Decentralized P2P Networks**. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, Miami, Florida 2006:315–324.

16. Castro M, Druschel P, Ganesh A, Rowstron A, Wallach DS: **Secure Routing for Structured Peer-to-peer Overlay Networks**. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI)*, Boston, Massachusetts 2002:299–314.

17. Roughgarden T, Tardos É: **How Bad is Selfish Routing?** *Journal of the ACM* 2002, **49**(2):236–259.

18. Cole R, Dodis Y, Roughgarden T: **Pricing Network Edges for Heterogeneous Selfish Users**. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, San Diego, California 2003:521–530.

19. Feldman M, Chuang J, Stoica I, Shenker S: **Hidden-action in Multi-hop Routing**. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*, Vancouver, British Columbia 2005:117–126.

20. Blanc A, Liu YK, Vahdat A: **Designing Incentives for Peer-to-Peer Routing**. In *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Miami, Florida 2005:374–385.

21. Ngan TW, Wallach DS, Druschel P: **Incentives-Compatible Peer-to-Peer Multicast**. In *Proceedings of the 2nd Workshop on the Economics of Peer-to-peer Systems (IPTPS)*, Berkeley, California 2003:149–159.

22. Krishnan R, Smith M, Telang R: **The Economics of Peer-to-peer Networks**. *Journal of Information Technology Theory and Application (JITTA)* 2003, **5**(3):31–44.

23. Hardin G: **The Tragedy of the Commons**. *Science* 1968, **162**:1243–1248.

24. Feldman M, Chuang J: **Overcoming Free-Riding Behavior in Peer-to-Peer Systems**. *ACM SIGecom Exchanges* 2005, **5**(4):41–50.

25. Ratnasamy S, Francis P, Handley M, Karp R, Schenker S: **A Scalable, Content-addressable Network**. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, San Diego, California 2001:161–172.

26. Rowstron A, Druschel P: **Pastry: Scalable, Decentralized Object Location and Routing for Large-scale Peer-to-peer Systems**. In *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, Heidelberg, Germany 2001:329–350.

27. Zhao BY, Huang L, Stribling J, Rhea SC, Joseph AD, Kubiatowicz JD: **Tapestry: A Resilient Global-scale Overlay for Service Deployment**. *IEEE Journal on Selected Areas in Communications (JSAC)* 2004, **22**:41–53.

28. Lv Q, Cao P, Cohen E, Li K, Shenker S: **Search and Replication in Unstructured Peer-to-peer Networks**. In *Proceedings of the 16th International Conference on Supercomputing*, New York 2002:84–95.

29. Artigas MS, López PG, Skarmeta AG: **A Novel Methodology for Constructing Secure Multipath Overlays**. *IEEE Internet Computing* 2005, **9**(6):50–57.

30. Fiat A, Saia J, Young M: **Making Chord Robust to Byzantine Attacks**. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, Mallorca, Spain 2005:803–814.

31. Naor M, Wieder U: **A Simple Fault Tolerant Distributed Hash Table**. In *Proceedings of the 2nd International Workshop on Peer-to-peer Systems (IPTPS)*, Berkeley, California 2003:88–97.

32. Ren S, Guo L, Jiang S, Zhang X: **SAT-Match: A Self-adaptive Topology Matching Method to Achieve Low Lookup Latency in Structured P2P Overlay Networks**. In *Proceedings of the 18th International Symposium on Parallel and Distributed Processing (IPDPS)*, Santa Fe, New Mexico 2004:83–91.

33. Hong F, Li M, Yu J, Wang Y: **PChord: Improvement on Chord to Achieve Better Routing Efficiency by Exploiting Proximity**. In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW)*, Columbus, Ohio 2005:806–811.

34. Condie T, Kamvar AD, Garcia-molina H: **Adaptive Peer-to-Peer Topologies**. In *Proceedings of the 4th International Conference on Peer-to-Peer Computing (P2P)*, Zurich, Switzerland 2004:53–62.

35. Gatani L, Re GL, Gaglio S: **An Adaptive Routing Protocol for Ad Hoc Peer-to-peer Networks**. In *Proceedings of the 6th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Taormina, Italy 2005:44–50.

36. Xie L, Zhu S: **Message Dropping Attacks in Overlay Networks: Attack Detection and Attacker Identification**. *ACM Transactions on Information and Systems Security (TISSEC)* 2008, **11**(3).

37. Danezis G, Lesniewski-Laas C, Kaashoek MF, Anderson R: **Sybil-resistant DHT Routing**. In *Proceedings of the 10th European Symposium on Research in Computer Security (ESORICS)*, Milan, Italy 2005:305–318.

38. Stavrou A, Locasto ME, Keromytis AD: **Pushback for Overlay Networks: Protecting against Malicious Insiders**. In *Proceedings of the 6th International Conference on Applied Cryptography and Network Security (ACNS)*, New York, New York 2008:39–54.

39. Samuelson PA: *Foundations of Economic Analysis*. Harvard University Press 1947.

40. Henderson JM, Quandt RE: *Microeconomic Theory: A Mathematical Approach*. Economics Handbook Series, McGraw Hill, 3rd edition 1980.

41. Negishi T: **Welfare Economics and Existence of an Equilibrium for a Competitive Economy**. *Metroeconomica* 1960, **12**:92–97.

42. Takayama A: *Mathematical Economics*, Cambridge University Press. 2nd edition 1985 :117.

43. Varian H: **Economic Mechanism Design for Computerized Agents**. In *Proceedings of the 1st Usenix Workshop on Electronic Commerce*, New York 1995:13–21.

44. Kuhn HW, Tucker AW: **Nonlinear Programming**. In *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability* 1951:481–492.

45. von Neumann J, Morgenstern O: *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition 1944.

46. Merton RC: **Lifetime Portfolio Selection Under Uncertainty: The Continuous-time Case**. *Review of Economics and Statistics* 1969, **51**:247–257.

47. Cooke J: **A Comparison of the Effectiveness of Message Dropping Attacks Against Structured Peer-to-peer Networks**. *Senior undergraduate honors thesis, University of Texas at Dallas, Richardson, TX* 2008.

48. Panzar JC, Sibley DS: **Public Utility Pricing Under Risk: The Case of Self-Rationing**. *American Economic Review* 1978, **68**(5):888–895.

49. Friedman M: *Essays in Positive Economics*. University of Chicago Press 1953.

50. Arrow KJ: **A Difficulty in the Concept of Social Welfare**. *Journal of Political Economy* 1950, **58**(4):328–346.

51. Hamlen W, Hamlen KW: **A Closed System of Production Possibility and Social Welfare**. *Computers in Higher Education Economics Review* 2006, **18**:15–18.