

Using Repeated Games to Design Incentive-Based Routing Systems

Mike Afergan
MIT CSAIL

Email: afergan@alum.mit.edu

Abstract—Incorporating pricing information in routing systems has been explored in various contexts and fashions. In this paper we examine certain fundamental properties important in the design of such a routing protocol. The importance of these properties is derived from the underlying economic factors governing the behavior of the autonomous players.

We view the exchange of pricing information at an interconnect as a *repeated game* between the relevant players. For example, multiple ISPs competing for the business of a CDN. With this model, we see that various protocol parameters—such as protocol period, minimum bid size, and unit of measure—have a significant and important impact on the equilibrium outcome. We show how these parameters can be used to address the problem of the repeated dynamic and further that these conclusions are robust to a variety of practical assumptions including asynchronous play and heterogeneous networks. These often surprising results enable protocol designers to appreciate and leverage these seemingly benign parameters, a result that has direct practical importance.

I. INTRODUCTION

Internet routing is a dramatic example of the introduction of economic concerns into an already rich design space. Traditional design concerns include the impact of system parameters on such objectives as convergence, robustness, efficiency, and performance. Today, economic considerations play a chief role in the routing of traffic in the Internet. Each Autonomous System (AS) is an independent profit-maximizing firm, competing to generate profit by routing bits on its network. Currently, this market plays out on two very different timeframes. On a multi-month or year timeframe, networks and customers negotiate economic contracts. Then, on a timescale of seconds, routers direct traffic given a configuration which encodes these business relationships.

There are several reasons to couple these two processes more tightly. User-directed routing, such as overlays or multi-homing, shifts the balance of power, creating tension between the users and the ASes. This may be best resolved by making the incentives more explicit in the routing system [1]. Another reason is that prices make explicit what is implicit in current routing protocols and router configurations. Today's inexact tools create significant manual overhead and increase the chance for error [2] [3] [4]. By contrast, routing policy based on the explicit incentives of ISPs may be much simpler to implement. Further, as user-directed routing technologies emerge and users (e.g., companies or Content Delivery Networks) use these tools to financial ends [5], these customers may demand such service from their ISPs. This paper does not




Fig. 1. Three Networks Offering Connectivity for a Set of Routes to a Single Route Selector

defend the concept of incentives in routing but instead asks the question: “**How should one design a protocol to convey pricing information for routes?**”

To reason about the behavior of ASes under a dynamic pricing protocol, we employ the analytical framework of *repeated games*. Most distributed computations on the Internet are repeated, often among long-lived players. Repeated games are therefore an appropriate model for participants’ self-interested behavior in these systems. In repeated games, the threat (or promise) of future behavior can impact current actions, and therefore the equilibrium of one-shot and repeated games can differ significantly. This makes repeated game analysis an appropriate and practical analytical tool for a distributed protocol that will be implemented by autonomous entities. By using repeated games, this paper contrasts with most prior work in this space, namely the celebrated Feigenbaum, Papadimitriou, Sami, and Shenker (FPSS) [6] analysis.

The relevant setting for our analysis is a particular interchange between a customer and a set of networks. The customer could be an enterprise, a Content Delivery Network (CDN), or an access ISP. The customer connects to multiple networks with each network providing connectivity to the same set of destinations, as in Fig. 1, and competing for the business.¹ Based on the amount of traffic served, the customer then directly pays the provider to which it is connected.

This model is similar other applied work (e.g., [5] [7]), but contrasts with some prior theoretical work on routing that considers general network topologies. There are several reasons for this. In practice, the customer does not pay every ISP in the route, but only the ISP at the first-hop.

¹For this reason, we use the terms ‘network’, ‘AS’, and ‘player’ interchangeably in this paper.

The relevant competition therefore is between those providers directly adjacent to the customer. Our model, like [5] and [7], maintains this bilateral business relationship. Further, if incentive-based routing is deployed, it will most likely be applied at particularly useful interchanges (e.g., for CDNs who have a dramatic ability to shift traffic). They will not necessarily be used throughout the Internet graph. Of course, if dynamic pricing is eventually deployed throughout the path, the analysis of this paper also serves as an important first step. These points are discussed in more detail in [1].

In our model, the repeated dynamics of routing cause certain protocol parameters to achieve significant importance. It is well-known that in general repeated games, there exist parameters that significantly impact the equilibrium outcome. However, routing is special due both to the particulars of the problem and that it transpires via a fixed network protocol. Therefore, the contributions of this paper include not only a repeated model for routing—but also formal analysis of the particular parameters relevant to routing. We summarize those results here in the form of practical statements:

- 1) A longer protocol period (a slower protocol) can lead to a lower price.
- 2) Using a more granular format (e.g., megabits instead of megabytes) can lead to a higher price.
- 3) A wider price field in the protocol can lead to a lower price.

Given this sensitivity, we also show how protocol designers, to the degree desired, can bound prices and their sensitivity to repeated game effects. *These conclusions have clear, direct, and previously unrecognized practical significance for protocol designers.*

In this paper, we present these results by first examining a simple model to gain intuition and then exploring more involved and practical models to demonstrate how the results generalize. In Sections II and III we present an overview of repeated games and their impact on routing. We then present a simplified model of repeated routing in Section IV which we analyze to derive the above conclusions and provide clean intuition. We then significantly generalize the model. In Section VI-A, we examine a larger class of strategies whose only constraint is that the punishment be at most proportional to the deviation. This is a very significant generalization, perhaps even more general than required. In Section VI, we then consider additional relaxations and generalizations to our model, such as asynchronous play, confluent flows and multiple destinations, before ending with a discussion of the results.

II. THE NOTION OF REPEATED GAMES

Routing is inherently a repeated process – information is advertised between networks, routing decisions are made, and bits are routed. New information (perhaps dependent on actions in the prior period(s)) is then transmitted, and the process begins again. It is vital that any model of routing – and particularly any model of routing that includes incentives – appropriately captures this property. Fortunately, repeated

games are a well understood aspect of game theory. In this section, we present a targeted overview of the relevant concepts. For a more thorough summary of the topic, consult [8] or [9].

A. A Simple Model and Key Tools

A repeated game is the repeated play of a particular *stage-game*. Here, we consider the prisoners’ dilemma, a canonical example which maps nicely to our problem. Two players simultaneously choose to either cooperate (C) or defect (D), with stage-game payoffs given by Table I.

TABLE I
GAME PAYOFFS FOR THE PRISONERS’ DILEMMA

	C	D
C	(1,1)	(-1,2)
D	(2,-1)	(0,0)

In the one-shot game, both players will play D. Regardless of what the other does, it is always in the best interest for a particular player to defect. The only Nash Equilibrium (NE) of this game is therefore (D, D). This also holds when the number of rounds is finite and known, as reverse induction shows that each stage-game is equivalent to the one-shot game.

When the number of rounds is infinite or unknown, other equilibrium outcomes are possible. For example, consider the following strategy for Player 2:

- 1) Play C
- 2) If P1 ever plays D then play D forever.

If the game is infinite and players do not discount future periods, (C,C) is as an equilibrium outcome. The threat of punishment causes the selfish players to not defect.

This repeated equilibria can occur even when the game is finite and/or players are impatient. In these cases, we introduce a discount factor δ ($0 \leq \delta \leq 1$). This can capture future period discounting and/or the probability of the game ending at each period. In such a case, the players will cooperate only if δ is sufficiently large relative to the parameters of the game and the particular strategy being played.

Let $u_1(\cdot)$ represent the utility to player 1 of a given ordered pair of plays. Given our example and sample strategy, mutual cooperation is an equilibrium only if:

$$\sum_{t=0}^{\infty} \delta^t u_1(C, C) \geq u_1(D, C) + \sum_{t=1}^{\infty} \delta^t u_1(D, D) \quad (1)$$

which yields:

$$u_1(C, C) \geq (1 - \delta) u_1(D, C) + \delta u_1(D, D) \quad (2)$$

Substituting, the values from Table I, we must have $\delta \geq \frac{1}{2}$ in this example.

B. Repeated Games Under General Conditions

While the previous example focused on a narrow scenario, so-called ‘‘Folk Theorem’’ repeated equilibria results have been shown to be applicable to a variety of contexts. They have also been shown to be robust to a variety of assumptions – in theory and practice. These include such weakening assumptions as imperfect information [10] [11], players of different horizons [12], and even anonymous random matching [13]. In an applied work with parallels to our research, Dellarocas applies repeated game techniques to the design of reputation systems. [14].

Even more important, we often see the effects of repeated interactions and signaling in practice. This is discussed in the case of auctions by Klemperer in [15] [?]. We also see elevated prices in oligopolies. Later in the paper we will discuss how the ‘‘Price Match Guarantees’’ offered by many retailers that we are familiar with also relate to repeated games.

C. Key Terms and Notation

Here we present some standard terms and notation that we will be using in the rest of the paper:

- *Subgame*: A subgame is the subset of an original game beginning at a particular point where all (relevant) history of play is common knowledge and continuing to the end of the original game.
- *Profit Function*: $\pi(p_i, p_{-i})$: We denote the per stage-game payoff using the function $\pi(\cdot)$. The function is defined by the structure of the game. The parameters of $\pi(\cdot)$ are p_i , the play of player i , and p_{-i} , a vector representing the play of all other players. When play of the other players is symmetric, we can write p_{-i} as a single number without loss of generality.
- *Weakly dominant*: A strategy is weakly dominant if it always does at least as well as any other strategy, regardless of the strategy selected by the other players.
- *Strategyproof*: A one-shot mechanism is strategyproof if truth-telling about one’s private information is weakly dominant. In the repeated game, we define a mechanism to be strategyproof if the strategy function that always plays truthfully is weakly dominant.

III. THE CHALLENGE OF REPEATED ROUTING

The routing game, as depicted in Fig. 1, presents the same phenomenon as the repeated prisoners’ dilemma example. Under reasonable assumptions, firms can maintain an artificially higher price if their strategies include appropriately crafted threats to punish deviators.

It is important to contrast this repeated context with prior work on routing, namely the celebrated work of Feigenbaum, Papadimitriou, Sami, and Shenker (FPSS) [6]. They demonstrate that it is possible to implement the well-known Vickrey-Clarke-Groves (VCG) mechanism efficiently with a




Fig. 2. Repeated FPSS Model is Not Strategyproof in the Repeated Game

protocol that resembles BGP.² Here a strategyproof mechanism is desirable since networks do not need to expend effort strategizing about prices and since the selected routes will be more stable.

Our work builds upon their results by considering their model in the repeated setting. Thus we summarize it here:

- A set of nodes N , with $n = |N|$, representing the ASes
- A constant per packet cost c_i for each node i
- A traffic matrix $T_{i,j}$ which is exogenous and fixed (i.e., inelastic demand)
- Each AS has infinite capacity

The VCG mechanism, and thus the FPSS implementation, obtains its strategyproof property through a carefully selected payment to each node. Each node, i , on a Least Cost Path (LCP) between a source-sink pair (s, t) is paid c_i plus the difference between the cost of the LCP and the cost of the LCP if i did not exist. For example, in Fig. 2, node A is on the LCP from s to t_1 . For traffic from s to t_1 A is paid:

$$p_A = LCP_{(c_A=\infty)} - LCP + c_A = (10 + 1) - (1 + 1) + 1 = 10$$

Similarly, B is paid 10 for each traffic unit from s to t_2 .

However, it is well known that VCG mechanism is not strategyproof in the repeated game. If A and B both bid 20 until the other defects, each will be paid:

$$(20 + 10) - (20 + 1) + 20 = 29$$

We can easily show that it is possible for this to be an equilibrium for sufficiently patient players. More formally, there exists a $\bar{\delta}$ such that for all $\delta \geq \bar{\delta}$ this strategy can exist in equilibrium.³ This demonstrates that although Internet routing is a repeated setting, the VCG mechanism (and thus the FPSS implementation), is not strategyproof in the repeated routing game.

Said differently, it is known that the VCG mechanism is susceptible to collusion. But in the one shot game, without explicit agreements, such cooperation is not possible. In the repeated game, the subsequent periods provide the players

²FPSS were not the first to consider the VCG mechanism for routing [16] [17]. However, one contribution of the FPSS work is the framing of the problem with the nodes as strategic agents. This maps to the problem of AS competition and motivates us to consider the repeated game.

³For completeness, in the one-shot game, bidding 20 is not an equilibrium strategy. The other player can bid 11 and get all the traffic on both routes for a price of 29 – yielding a higher profit.

with a means of obtaining a higher price without any explicit collusion, side-payments, or constructs of any sort.⁴ While we consider a somewhat different model, this intuition about the VCG still holds. This is very troubling since routing is clearly a repeated game, not a one-shot game.

IV. A MODEL OF REPEATED ROUTING

The observation the mechanism is not strategyproof in the repeated game is worrisome for several reasons. First, to the extent that the VCG/FPSS prices are fair or desirable, we have no way of ensuring that they will occur. Second, we do not initially have any understanding of what the outcome will now be. Third, we do not understand how design decisions will impact the outcome.

To address these questions, we analyze a model of the repeated game. First, we present the formal model. We then analyze the model to determine the equilibrium outcome. Finally, we analyze this outcome to describe the impact of the fundamental design parameters on this outcome. We do not seek to impose a particular outcome (e.g., minimize price) on the system, since as we discuss in Section VII, it is unclear there is a universally correct and acceptable answer. Instead we focus on understanding the impact of these parameters and design decisions.

A. Key Intuition and Analytical Approach

Before delving into the model and analysis, we first present the high level intuition which underlies the results in the remainder of the paper. Consider a small number of firms competing for an amount of traffic. At any point in time, each firm faces a key strategic decision. One strategy is to attempt to be the low price provider and receive all the traffic. Another strategy is to offer a higher price to the market, somehow splitting the traffic with the other providers—but potentially garnering more profits due to the increased price.

The firm’s willingness to take this second, cooperative, strategy is a function of several factors. One is the *granularity of the action space* – which in this example is amount the firm needs to deviate to obtain all of the traffic. For example, if the price is \$100 and the firm can obtain all the traffic at \$99.99, it will be more likely to deviate than if it were constrained to integer prices (e.g., \$99). Another factor is the *discount factor* which manifests itself in several ways, notably in the length of the game. If the firm feels that the game will end soon, it will be more likely to decrease price to get the extra profit. On the other hand, if it feels the game will last longer, it may not want to perturb its competitors.

A key insight of this paper is that in practice these factors are directly determined by parameters of the protocol. In particular, the width of the pricing field and the representation used determines the granularity of the action space. Further, the protocol period determines the number of periods the game will be played.

⁴Certainly, with additional such collusive constructs, other equilibria are possible. We ignore those for the purposes of this paper.

Because this intuition is fundamental to the competitive dynamics of the situation, the results obtained are robust to a wide-range of practical and important assumptions. These include, but are not limited to:

- Heterogeneous networks
- Asynchronous protocols
- Multiple destinations in a network
- Confluent (BGP-like) routing
- Multi-hop networks
- A wide class of reasonable strategies for the firms

In the interest of clarity, however, we do not present a general model that contains all of these properties. Instead, we first start with a simple model that captures the essence of the game, provides for lucid analysis, and demonstrates the key intuition. After, in Section VI, we return to these assumptions and formally prove the same set of results for models that incorporate these more sophisticated assumptions.

B. The Repeated Incentive Routing Game (RIRG)

Our model is based on the FPSS model, presented above. We extend their model to capture both the repeated nature of routing and the properties of routing protocols. We also introduce some simplifying restrictions to make the model more tractable for the initial analysis and appropriate for the particular problem in which we are interested. In Section VI we relax and address many of these assumptions.

The game analyzed in this section is depicted in Fig. 1. This could be competition among ASes for the traffic of a CDN or perhaps a third AS. As discussed in the Introduction and in [1], this analysis examines a particular interchange instead of a general graph. This is similar to the analysis of [5] and [7].

We note that Shakkotai and Srikant [7] also consider a repeated model of routing, but in a very different context. A key difference is that we examine protocol design where pricing and competition occur electronically, whereas they examine competition in the abstract (presumably through paper contracts). Other important differences include the fact that we consider a much larger space of potential ISP strategies.

We now define the particulars of the game:

Repeated Incentive Routing Game Model

- There is only one source and one destination.
- Each of the N networks connects directly to both the source and the destination with exactly one link, as depicted in Fig. 1.
- Each network has an identical, constant per-packet cost c for transiting the network⁵, identical quality, and infinite capacity.
- Bids are represented as fields in packets and thus are discrete. The maximum granularity of the representation, equivalent to the minimum change in a bid, is represented by b . For simplicity, we assume that c is a multiple of b .

⁵We note that this maps cleanly to average-based billing, a common billing technique in practice. A richer discussion of volume-based versus rate-based models and their prevalence in industry is beyond the scope of this paper.

- Each AS is perfectly-patient with respect to the time value of money.
- The low cost bid in each period is common knowledge. Specifically, before the next time an AS advertises a price, it knows the lowest price bid in the prior round.
- The game is finite. The game length is represented as an exponential random variable with mean D . $E(D)$ is known to all players but the actual value is unknown to the source or any of the other ASes. The duration corresponds to the expected period of time for which the other factors will be stable.⁶

Play of the Game

- 1) The game proceeds in a series of rounds, each of length d , a constant of common knowledge. For simplicity, we assume that D is a multiple of d . Thus, we can relate d and D as:

$$d = D(1 - \delta) \quad (3)$$

where $(1 - \delta)$ is a constant representing the per-period probability of the game coming to an end.

- 2) At the start of each round, each of the N players advertises its per-packet price simultaneously.
- 3) For the entire period, traffic is routed over the provider with the lowest price. In the event of a tie, traffic is split evenly among the providers with the lowest price.
- 4) Each provider is paid for the number of packets that transit its network. The rate paid for each packet is the price it advertised at the beginning of the round (first price auction).⁷

C. Equilibrium Notion and Strategy Space

It is important to refine the space of strategies and equilibria, since in a repeated game, the set can be quite large. Our significant refinement is to consider only strategies which are subgame perfect. Subgame perfection, defined below, means that in every stage game, all players must play a strategy that is optimal, given the remainder of the game. It thus takes the entire discounted stream of payments into account, precluding myopic strategies but permitting long-term thinking. While not vital, for clarity, we also restrict ourselves to symmetric equilibria, where all players play the same strategy, and pure strategies. This allows us to speak of a single strategy being played.

Subgame Perfection: A strategy α is subgame perfect if i) α is a Nash equilibrium for the entire game and ii) α is a Nash equilibrium for each subgame.

In a repeated, simultaneous-move game such as ours, the set of Subgame Perfect Equilibria (SPEs) can still be quite large. One class of strategies are “trigger price strategies” [11]. In this context, players offer some desirable price, p^* , so long as all other players do. If a player deviates, offering some $p' <$

p^* , the other players *punish* the deviating player by playing some $\hat{p} < p^*$. In general, trigger price strategies allow for the players to return to p^* after some period of time. The intuition of these strategies is that in equilibrium the threat of punishment can maintain a higher price.

While there is empirical studies supporting the existence of such behavior, such a severe and coordinated practice may seem implausible in many contexts. For example, in the bandwidth market, we have not observed such wild swings. Instead, as costs decrease and competitive pressure has increased, we have seen prices move down rather smoothly and steadily. Such a phenomenon may be better modeled by a *price matching* strategy where players play the lowest price observed in the prior period.

The key difference between these two classes of strategies is how we perceive the reaction to a deviation. To the extent that it is a *punishment*, trigger-price strategies are appropriate. To the extent that it is simply a *protective reaction or learning mechanism*, price matching seems more appropriate. Price matching may even be too severe, as more appropriate strategy may be to price match for a certain period before returning to p^* . However, in all strategies, deviation of a player leads to decreased profit for some number of future periods.

We can generalize this space of strategies. In particular, in a parameterized space, price-matching is a mild punishment for an infinite amount of time. By parameterizing the punishment time and severity, we can consider a larger class of strategies. In this paper, motivated by space and clarity, we first discuss and analyze price matching strategies. However, in Section VI-A, we show that our results hold for a much larger class of strategies, namely all strategies where the punishment is no greater than a constant multiple of the deviation. (This is important as it also permits strategies that return to p^* .)

We are now ready to formally define the price-matching strategy:

Price Matching (PM) Strategy

- S1) At $t = 0$, offer some price p_i^0 . For $t > 0$:
- S2) $p_i = \max(c, \min_j(p_j^{t-1}))$

where p_j^t is the price offered by player j in period t .

While the second step of the strategy is clear, it is not immediately obvious how a player should select the initial p_i^0 . From our definition of subgame perfection, we have that the player cannot benefit from deviating. From price matching, we have that in equilibrium $p^t = p^0 \forall t$. (Thus, we drop the superscript notation and simply write p .) This means that if price matching at p is an SPE then:

$$\sum_{t=0}^{\infty} \delta^t \pi_i(p, p) \geq \pi_i(p - b, p) + \sum_{t=1}^{\infty} \delta^t \pi_i(p - b, p - b) \quad (4)$$

for a given (δ, b) . Just as Eqn (1) simplifies to (2), we can simplify Eqn (4) to:

$$\pi_i(p, p) \geq (1 - \delta)\pi_i(p - b, p) + \delta\pi_i(p - b, p - b) \quad (5)$$

Informally, this condition says that we will accept p only if the payoff to playing p forever is greater than the payoff from

⁶The property of having stability for a sufficiently-long, finite, and unknown period of time corresponds very well to the true nature of Internet interconnects.

⁷Payment mechanisms, enforcement, etc are outside the scope of this game.

TABLE II
SUMMARY OF KEY TERMS

Term	Meaning
N	Number of firms competing for the traffic
b	Minimum bid change size
p	Price
p^*	Profit maximizing price
$\pi(\cdot)$	Per-firm profit function
δ	Per-period chance of the game ending
d	Period of the protocol
D	Expected stability of network topology
T	Total amount of traffic

deviating once and suffering the consequences. The $\pi_i(p, p)$ term represents the payoffs of playing p ; $(1 - \delta)\pi_i(p - b, p)$ is the weighted payoff to deviating by some amount b ; and $\delta\pi_i(p - b, p - b)$ captures the payoffs in the future. Note that it is strictly dominant to make the smallest possible deviation from p . Thus, we use b , the size of the minimum bid change, as the magnitude of the deviation without loss of generality.

Of all the values of p that satisfy Eqn (5), we consider the profit-maximizing value, which we define to be p^* . Therefore,

$$p^* = \max_p \text{ s.t. } \pi_i(p, p) \geq (1 - \delta)\pi_i(p - b, p) + \delta\pi_i(p - b, p - b) \quad (6)$$

(We solve for p^* explicitly for our game by expanding $\pi(\cdot)$ in the following section.)

To show that PM is SPE, we will use the one-deviation principle which states that a strategy is a SPE if and only if it is not possible to profitably deviate in exactly one stage-game. This allows us to consider simple one-stage deviations as opposed to more complicated multi-stage deviations. We state the principle, whose proof can be found in [9], below:

Theorem 1 (One Stage Deviation Principle): In an infinite horizon multi-stage game with observed actions where the payoffs are a discounted sum of per-period payoffs and the per-period payoffs are uniformly bounded; strategy profile α is subgame perfect if and only if it satisfies the condition that no player i can gain by deviating from α in a single stage and conforming to α thereafter.

Lemma 1: PM is a SPE.

Proof: First we note that the RIRG game satisfies the technical conditions of the principle and the fact that the game is finite ensures that the discount factor, $\delta < 1$. Therefore, we can apply the theorem and consider only one-stage deviations. We look at each stage of the specified strategy:

- S1) By construction, assuming that other players offer p^* , it is optimal to offer p^* . By definition of p^* , bidding a lower value decreases the discounted stream of profits. A higher price leads to no profits in this period and no prospect of higher profits in the future.
- S2) Again by construction, there is no benefit to decreasing price. Likewise, increasing price given that others are playing PM does not help.

Since we have examined all one-stage deviations, we have that PM is a SPE. ■

D. Analysis of the Model

With a model and equilibrium notion, we can now examine the equilibrium conditions. The first step is to derive an explicit expression for p^* in terms of the parameters of the game.⁸

Theorem 2: In the RIRG, the unique equilibrium price when all players play Price Matching is given by:

$$p^* = \frac{b(\delta bN - \delta b - N)}{1 - N + \delta N - \delta} + c$$

⁸As discussed, price is discrete. However, for notational simplicity, we analyze the continuous variable p such that the market price is $\lfloor \frac{p}{b} \rfloor b$.

Proof: Since the firms seek to maximize profit, we consider the profit-maximizing price matching strategy, which bids p^* as given by Eqn (6). This means we have:

$$\pi_i(p, p) = (1 - \delta)\pi_i(p - b, p) + \delta\pi_i(p - b, p) \quad (7)$$

where p is the price advertised.

Define $m = p - c$ for notational simplicity, and T as the total amount of traffic. We now expand π_i based on the definition of the game:

$$\pi_i(p_i, p_{-i}) = \begin{cases} \left(\frac{T}{N}\right) m, & p_i = p_j, \forall i \neq j \\ T * m, & p_i < p_j \forall j \neq i \\ 0, & \text{otherwise} \end{cases}$$

This yields:

$$\left(\frac{T}{N}\right) m = (1 - \delta)(m - b)T + \delta \left(\frac{T}{N}\right) (m - b) \quad (8)$$

Solving, we have:

$$m = \frac{b(\delta N - \delta - N)}{1 - N + \delta N - \delta} \quad (9)$$

or

$$p = \frac{b(\delta N - \delta - N)}{1 - N + \delta N - \delta} + c \quad (10)$$

Since all players are homogeneous and since we consider only symmetric equilibria, this is thus the unique equilibrium. ■

V. UNDERSTANDING THE RESULT

Given an expression for the equilibrium price, we turn to the practical questions that we seek to understand.

A. Protocol Period

We examine the model parameter tied to the period, δ , holding the other factors (including D) constant. Intuitively, it may seem that the period of the game should have no impact on prices. Alternatively, a shorter period—corresponding to a faster protocol—would perhaps help to keep the market more competitive. This is not necessarily the case.

Lemma 2: The protocol period and the market price are positively correlated – or $\frac{\partial p}{\partial \delta} > 0$.

Proof deferred to appendix.




Fig. 3. Price as a function of δ for $N = 2$, $c = 0$, and $b \in \{0.01, 0.05, 0.1\}$. Margin increases with δ and is very sensitive to δ when δ is large.

Recall now that $d = D(1 - \delta) \rightarrow \frac{\partial d}{\partial \delta} < 0$. This coupled with Lemma 2 yields

$$\frac{\partial p}{\partial d} < 0 \quad (11)$$

In other words, *as the protocol period increases, the price decreases – a surprising and initially counterintuitive result!*

Careful consideration provides us with the rationale behind this conclusion. When a player deviates, it enjoys a one-time increased payoff at the expense of diminishing the future stream of payoffs. Consequently, the longer the period is before a competitor can match the price, the bigger the benefit to deviating. Furthermore, a longer period means fewer expected future periods. As a result, as we increase the protocol period, we increase the propensity for a player to lower its price. It is well-known in the repeated game theory of oligopolies that fewer periods can increase price. But it is interesting to realize that the protocol period, typically analyzed in the context of information flow and convergence, in practice also defines the number of rounds and thus significantly impacts the equilibrium.

We depict this relationship between p and δ by graphing Eqn (10) in Fig. 3. As can be observed, p is strictly increasing in δ but converges readily to $c + \frac{Nb}{N-1}$ as $\delta \rightarrow 0$.⁹

Although this phenomenon may seem counter-intuitive at first, most consumers are familiar with it in the form of “Price Match Guarantees” offered by many major retailers. [18] [19] While the policies vary, the notion is that Firm A will match any competitor’s advertised price that is lower than A’s price. While there are other factors at play in these markets, this practice can be abstracted in the notion of a protocol period. Instead of waiting some period to match the competitor’s price (e.g., in the next week’s circular, in the following day, etc.) a price match guarantee effectively brings the period to zero. Once a firm lowers its price, the other firm effectively matches price immediately. Thus, one result of these policies is to dissuade competitors from lowering price, since, it can be

⁹For $N > 2$, $c + \frac{Nb}{N-1}$ yields $c + b$ when discretized to a multiple of b . In the one-shot game bidding c or $c + b$ are both Nash Equilibria. We return to this subject in Section VI-C.

TABLE III
IMPACT OF PROTOCOL PARAMETERS ON PRICE

Variable	Impact on Price
N : Number of players	Decreases
b : Minimum bid size	Increases
d : Period of the protocol	Decreases
D : Stability period for the topology	Increases

TABLE IV
THE IMPACT OF A \$1 PRICE CHANGE WITH MEGABYTE AND MEGABIT REPRESENTATION FORMATS

Format	Traffic	Price	Revenue	New Price	New Revenue
Mbits	1000	\$100	\$100,000	\$99	\$99,000
MBytes	125	\$800	\$100,000	\$799	\$99,875

argued, it will not provide that competitor with any additional revenue.

B. Additional Parameters

We now consider the other relevant parameters in similar fashion. For each parameter found in the expression for the equilibrium price, we present the main result and some intuition to provide better understanding. The results are summarized in Table III.

1) *Minimum Bid Size, b* : Similar to the analysis of period, we can show that $\frac{\partial p}{\partial b} > 0$ —as we increase the minimum bid size, the equilibrium price increases. This again comes from the firm’s decision which weighs the one-time benefit of deviating versus the longer-term cost. The less a firm is able to decrease p and still get all the market, the more profit it garners in the short-term and the less punishment it suffers in the long term. Therefore, it is more likely to deviate.

It is important to understand that this is more than just a matter of precision. Fig. 3 plots equilibrium price versus δ for $N = 2$, $b \in \{0.01, 0.05, 0.1\}$. One can note not only that p changes significantly but moreover that the change in p is qualitatively greater than the change in b .

Generally in practice, the minimum bid size is not an explicit parameter but rather it implicitly manifests itself in two protocol parameters. The first is the width of the pricing field. In most any protocol this width is likely to be fixed. Here we see that *increasing the width of the pricing field can decrease the price in the system*. Another means by which the minimum bid size manifests itself is via the unit of measure. Given a fixed granularity on prices, it makes a difference if we represent quantities in megabits or megabytes. For example, consider a system in which prices are set at whole dollar increments. Using *megabytes* as opposed to *megabits* provides for larger price values and thus more granularity in the prices, holding all other parameters constant. In Table IV we see that a \$1 decrease when using megabits causes a 1% decrease in revenue whereas a \$1 decrease when using megabytes causes a 0.125% decrease in revenue. Per the logic outlined above,

we see that *using megabytes instead of megabits can lead to a lower price.*

2) *Stability Period, D:* The stability period is likely not under the control of the protocol per se, but it is still useful to understand its impact on prices, holding other parameters (including d) constant. We have that $\frac{\partial p}{\partial \delta} > 0$. Since $d = D(1 - \delta) \rightarrow (1 - \frac{d}{D}) = \delta$ we have that $\frac{\partial \delta}{\partial D} > 0$. Thus, $\frac{\partial p}{\partial D} > 0$. This should come as no surprise given the prior two examples. As we increase the expected duration of the game, the relative importance of the stream of future payoffs increases. Thus, a player is less willing to deviate.

3) *Number of Players, N:* While the number of players is generally assumed to be constant, it is useful to note that similar to the other variables, we can show that $\frac{\partial p}{\partial N} < 0$. This conclusion is perhaps the most likely to be obvious *a priori*. As the number of firms increases, the profit is split among more players. Thus, as the number of firms increases, so too does the benefit from a one-stage deviation—and thus the propensity to deviate. This corresponds with the basic intuition that with more firms we approach perfect competition.

C. Discussion

Because the rest of the paper consists of various relaxations and further analysis of the results presented above, we pause here to make a few observations:

- *There are several protocol parameters which – unexpectedly – may significantly impact the equilibrium price.* These include the protocol period, the width of the pricing field and the unit of measure. Unlike some properties that one might readily be able to identify and reason about (such as the number of players); *a priori* it is unclear that these parameters have any affect. Further, it is unclear which way they push the equilibrium price. These often counter-intuitive results are therefore quite revealing.
- *The conclusions about these parameters are directly applicable to system design.* Understanding the impact of these parameters is a useful result. What is perhaps most important, however, is that *a priori* a protocol designer may not have even considered these parameters as relevant at all! Therefore, merely understanding that they are relevant, let alone understanding how they impact the equilibrium, is an important conclusion.
- *We have **not** shown that these repeated outcomes will always occur, but still believe consideration of the parameters is important.* We have shown that it is possible to obtain increased prices in repeated equilibria, but have not shown that this result is robust to all variations to our model. In Section VI, we will consider various relaxations to our model and show that similar results can be obtained. However, in a general setting, it is possible to construct degenerate scenarios in which such increased prices are not possible. Therefore, our argument is not that increased prices will always result if the parameters are not considered. Instead, we argue that in general (and in changing) environments, such outcomes *may* result, and in some cases will *likely* result. In practice, it will be

the rare case when one is certain that such outcomes will not occur. Since a good protocol should be applicable to a wide range of circumstances, we therefore believe that the protocol designer should and must take these parameters into account.

VI. IMPORTANT GENERALIZATIONS

In this section we show that the key intuition and spirit of the results from the simple Repeated Incentive Routing Game (RIRG) hold in more general networks, mechanisms, and/or assumptions. In all cases examined, while the model is more general, the key intuition (presented in Section IV-A) from the simpler example still holds. Consequently, the core results—the impact of the granularity of the action space and its manifestation in the protocol parameters—also carry over. The underlying reason for this is that in all of the games firms face the same decision: attempt to be the lowest-priced provider and take the whole market, or split the market at a higher price with multiple firms.

In this section we present four key generalizations. First we significantly expand the strategy space, allowing the ASes significantly greater flexibility. We then consider asynchronous play (vital for routing protocols), a set of FPSS-like assumptions including confluent flows, multiple destinations, and a second-price auction, and heterogeneous costs. In the first three cases, where the proof is particularly insightful, we present a full formal proof of our results. In the case of heterogeneous costs, where the extension is relatively simple, we present only a discussion of the modification.

A. Generalizing the Strategy Space

There are several reasons why the strategy considered thus far, price-matching, may be too restrictive. Perhaps most important, it assumes that prices never return to the original p^* . Further, the punishment phase is limited at $p - b$. For example, we could relax both of these assumptions and consider a set of strategies which punish at $p - kb$ for T periods before returning to p^* .

In this section, we generalize our results to a larger set of strategies, which we call *porportional-punishment strategies*. The punishment of these strategies is porportional to the deviation, and we permit prices to return to some higher price, including p^* . This is not the only sufficient condition but one that we believe is the most general and more important most indicitive of strategies that would be used in practice.¹⁰ For this set of strategies, we seek to understand the highest possible market price, \bar{p} , given a fixed (N, δ, b) tuple. We derive a bound and show that this bound is tight. Using this bound we can show that our conclusions regarding the parameters still hold. Further, we can show that it is possible for the protocol designer, if she indeed desired, to bound prices using the parameters to $p^I + \epsilon$ where p^I is the price in the one-shot, first-price auction.

¹⁰If anything, we believe that our model here is in fact too general, but that is not a problem in this context.

To begin our analysis, we introduce some new notation: Let α_h^t be the price proscribed by strategy α in period t given history h . Let p_i^t be the bid of player i in period t . Finally, let $D(\sigma, h, s) = \{t > s \mid \exists p_i^t < \alpha_{h^{t-1}}^t\}$. These are the periods where there has been a deviation.

We now define Δ which represents the set of all one-stage deviations for a strategy σ .

Definition: For all history pairs $(h, \hat{h}, \sigma) \in \Delta$

- 1) σ is a SPE strategy.
- 2) $h^t = \hat{h}^t \forall t < t_0$
- 3) In $t_0, \exists p' = p_i^{t_0} < \alpha_{\hat{h}^{t_0-1}}^{t_0}$
- 4) $D(\sigma, \hat{h}, t_0) = D(\sigma, h^t, t_0)$

Informally, this says that the two histories are the same until t_0 (2), there is a deviation at t_0 (3), and that there are no further deviations (4). We now define the set of *Porportional Punishment Strategies* (PP) as follows:

Definition: $\sigma \in PP_k$ iff σ is a symmetric SPE and $\forall (h, \hat{h}, \sigma) \in \Delta, \forall t > t_0, \forall p'$:

$$\sigma_h^t - \sigma_{\hat{h}}^t \leq k(\sigma_{\hat{h}}^{t_0} - p')$$

This captures our definition of porportional punishments. We now turn to analyzing the equilibrium conditions. To do so, we introduce two new terms: p^α and \bar{p}_k :

Definition: p^σ is the highest price obtained by σ in any period, that is:

$$p^\sigma = \max_{\forall t, h} \sigma_h^t$$

Definition: We define \bar{p}_k to be the maximum price for all strategies in PP_k . More formally,

$$\bar{p}_k = \max_{\alpha \in PP_k} p^\alpha$$

Theorem 3: If $\sigma \in PP_k$, then $p^\sigma \leq \frac{b(\delta N - N - \delta k)}{(\delta - 1)(N - 1)}$. Further this bound is tight, that is, $\exists \bar{\alpha}$ such that $p^{\bar{\alpha}} = \bar{p} = \frac{b(\delta N - N - \delta k)}{(\delta - 1)(N - 1)} + c$.

Given this bound of \bar{p}_k , we can also analyze the bound just as we did for the bound we derived for price matching. Simply by taking the partial derivatives, we see that we get the same qualitative results and resulting intuition.

Theorem 4: In the RIRG, if players play strategies in PP, the value of \bar{p} varies with the parameters (δ, b, N) in the same manner as the optimal price matching price. That is: $\frac{\partial \bar{p}}{\partial \delta} < 0$, $\frac{\partial \bar{p}}{\partial b} > 0$, $\frac{\partial \bar{p}}{\partial D} > 0$, and $\frac{\partial \bar{p}}{\partial N} < 0$.

Finally, we can consider the implementation design question of how to limit prices.¹¹ Here we obtain the following result:

Theorem 5: For an instance of the RIRG with N players playing strategies in PP, $\forall \epsilon > 0$ there exists a tuple (δ, b) such that $p_m < c + \epsilon$ where p_m is the price realized in the market.

¹¹Of course, as we discuss in Section VII, limiting prices is not necessarily a desirable goal.

B. Asynchronous Play

The assumption of synchronized play in the RIRG clearly does not match the reality of Internet routing. Moreover even if synchronization were desirable, it would be a hard property to achieve. While synchronous play is the normal model for repeated games, a limited amount of recent work has explored asynchronous models of repeated games, albeit in other contexts [20]. While the analysis of the previous section relied on this assumption, the key intuition of the problem (presented in Section IV-A) does not. Therefore, we are able to obtain essentially the same results in the asynchronous case, which we present below.

While the analysis below suggests that the asynchronous play does not change the game, that is not correct. Indeed, the asynchronous play has a significant impact on the set equilibrium strategies that can be played. For example, the grim strategy of setting price equal to cost in response to a defection is no longer a SPE strategy.

Lemma 3: If $\sigma \in PP_k$, then $p^\sigma \leq p^\sigma \leq \frac{kb(-N\phi + N\phi\delta - \delta^N)}{1 - N\phi + N\phi\delta - \delta^N}$ in the asynchronous game. Further this bound is tight, that is, $\exists \bar{\alpha}$ such that $p^{\bar{\alpha}} = \bar{p} = p^\sigma \leq \frac{kb(-N\phi + N\phi\delta - \delta^N)}{1 - N\phi + N\phi\delta - \delta^N}$

Proof: We know that in equilibrium:

$$\sum_{t=t_0}^{\infty} \delta^t p^\sigma \geq (p^\sigma - b)N + \sum_{t=t_0+1}^{\infty} \delta^t \beta_{(i,t)}(p^\sigma - b, t_0, \sigma_{-i})$$

where $\beta(\cdot)$ specifies the continuation payoff to player i for a deviation at t_0 with other players playing σ_{-i} . Since $\sigma \in PP_k$, let us consider the most severe punishment. This yields:

$$\sum_{t=t_0}^{\infty} \delta^t p^\sigma \geq \sum_{i=t_0}^{N-1} \frac{(p^\sigma - kb)N}{(i+1)} \delta^i + \delta^N \sum_{t=t_0+N}^{\infty} \delta^i (p^\sigma - kb) \quad (12)$$

For notational simplicity, we define:

$$\phi = \sum_{i=0}^{N-1} \frac{\delta^i}{(i+1)}$$

This is the sum of the first N terms of the Harmonic series with discounting.) We can restate Eqn.(12) as:

$$p^\sigma \geq (1 - \delta)\phi(p^\sigma - kb)N + \delta^N(p^\sigma - kb) \quad (13)$$

Solving for p^σ yields:

$$p^\sigma \leq \frac{kb(-N\phi + N\phi\delta - \delta^N)}{1 - N\phi + N\phi\delta - \delta^N} \quad (14)$$

From this Theorem, we can easily derive our two conclusions:

Theorem 6: In the **asynchronous** RIRG, if players play strategies in PP, the value of \bar{p} varies with the parameters (δ, b, N) in the same manner as the optimal price matching price. That is: $\frac{\partial \bar{p}}{\partial \delta} < 0$, $\frac{\partial \bar{p}}{\partial b} > 0$, $\frac{\partial \bar{p}}{\partial D} > 0$, and $\frac{\partial \bar{p}}{\partial N} < 0$.

We then can show:

Theorem 7: For an instance of the **asynchronous** RIRG with N players playing strategies in PP, $\forall \epsilon > 0$ there exists a tuple

(δ, b) such that $p_m < c + \epsilon$ where p_m is the price realized in the market.

C. FPSS-Like Assumptions

The RIRG assumption of splittable flows and a single destination maps to a variety of practical contexts. For example, many ISPs offer a single price for all Internet routes, making “the Internet” the single destination. Further, while flows over BGP are confluent, routing technologies such as multihoming and overlays enable customers to split traffic among the providers, by time or destination. This split can be done in time or by selecting some granularity smaller than the destination advertised by the ISP. Nonetheless, it is insightful to relax both of these assumptions.¹²

We formally define the new game below. In summary there are three key differences as compared to the RIRG:

- 1) Flows are confluent
- 2) There are two destinations
- 3) A second-price auction sets the allocation and prices

These three relaxations map directly to the FPSS model and our counter-example from Section III. Consider, for example, the limited topology depicted in Fig. 4 with two players. As in FPSS, each provider advertises a single bid for its network. Despite the changes from the RIRG, the players face a similar decision: bid low to obtain more traffic or concede one link to the other player in a repeated equilibrium. We define this game formally as a game among N players below. Each player P_i will be the low cost provider to some destination t_i .

We define the game formally as:

Model: N -Player Repeated VCG Routing Game

- There is only one source and N destinations, with $\frac{T}{N}$ units of flow from s to each of t_1, \dots, t_N .
- There are N networks in the game (P_1, \dots, P_N) , each connected to the source. These N networks have identical cost c and all paths have equal quality.
- Between the P_i s and the t_i s are other networks providing connectivity. Each network has a fixed price, c_h or c_l where $c_h > c_l$. Each P_i is connected to t_i via a network with cost c_l and connected to $t_j, j \neq i$, via a network with cost c_h .
- Each AS is infinitely patient with respect to the time value of money.
- All bids are common knowledge as in the prior game.
- δ models the finite but unknown duration

Play of the Game

- 1) The game proceeds in a series of rounds, each of length d , a constant that is common knowledge.
- 2) At the start of each round, each of the players advertises a *single bid* simultaneously. This value represents a (perhaps truthful) per-packet cost.
- 3) For the entire period, for each destination, traffic is routed over the provider with the lowest bid. In the event

¹²Note that it is natural to relax these two assumptions together. Relaxing one without the other produces a game that is either a generalization of the RIRG or game equivalent to one-shot Bertrand competition.




Fig. 4. The N -Player Repeated VCG Routing Game with $N = 2$. With $c_H > c_L$, A is on the LCP to t_1 whereas B is on the LCP to t_2 .

of a tie, traffic is sent to the lexicographic first network. Thus, all flows are confluent.

- 4) Each provider is paid for the number of packets that transit its network. The price per packet is set by the (second-price) VCG mechanism.

We define the critical price here in a similar fashion to the prior game:

$$p^* = \max_p \text{ such that } \pi_i(p, p) \geq (1-\delta)\pi_i(p-x, p) + \delta\pi_i(p-x, p-x)$$

for a given tuple (δ, π) and any $x \geq b$. However, note that the profit function, $\pi(\cdot)$, and selection of p^* is more subtle than before:

- The single bid requirement forces the minimum profitable deviation to be larger than the minimum bid size. For example, if two players are at a given p and P1 decreases its price by some $\epsilon < c_h - c_l$, there will be no benefit to this deviation.
- The deterministic tie-breaking causes an asymmetry. A player later in the lexicographic ordering must exhibit a (slightly) larger price decrease to gain the additional traffic.
- The nature of the second price auction is that if $(p-x)$ is the lowest bid, then $\pi_i(p-x, p) = Tp$ not $T(p-x)$ as in the first-price auction.

Despite these differences, we are able to obtain a similar result for the protocol period, namely that $\frac{\partial p}{\partial \delta} > 0$.

Theorem 8: In the N -Player Repeated VCG Routing Game, $\frac{\partial p}{\partial \delta} > 0$.

Proof: Deferred to Appendix ■

While the impact of δ is the same in the 1st and 2nd price mechanisms, the impact of b is different. One difference is that unlike the first price auction, a deviation must be larger than the minimum bid size to have impact. That is, the minimum bid size is not relevant in determining equilibrium price (beyond rounding). This can be viewed as a positive or negative. In the repeated first price mechanism, as $\delta \rightarrow 0$, $p^I \rightarrow b \frac{N}{N+1} + c$ (Theorem 5). However, in the repeated second price mechanism, as $\delta \rightarrow 0$, $p^{II} \rightarrow y(N-1) + c$ (Lemma 4). Both prices correspond to the maximal values of the set of undominated strategies in the one-shot game. Thus, from an implementation perspective, while it is possible to force $p^I \approx c$ independent of topology, p^{II} may be significantly bound away from c even with the slightest possibility of repetition. These

elevated prices and lack of control can be viewed as additional weaknesses of the VCG mechanism in the repeated game.¹³

D. Heterogeneous Costs

Above we saw that cooperation is possible in spite of – and in fact facilitated by – heterogeneous cost structures. Examining heterogeneous costs provides better insight into the strengths and weaknesses of our result.

With heterogeneous costs non-trivial repeated equilibria can and still *may* exist. Let the function $p^*(c)$ evaluate to the p^* in a game with homogeneous costs of c . Consider a two player game where $c_1 < c_2$. Here, P1 has the choice of i) selecting a repeated equilibrium in which the market is split or ii) pricing below c_2 and taking the whole market. This corresponds to bidding $p^*(c_1)$ or $c_2 - b$ respectively. For example, if we have $c_1 = 1$, $c_2 = 1.1$, and $p^*(1) = 3$, then we would expect the equilibrium price to be 3, even though P1 could undercut and price at say 1.09.¹⁴ This logic can be generalized to derive an equilibrium price in the case of heterogeneous costs. Given this price, the results from prior sections – the impact of δ and b in the first price auction and the impact of δ in the second price auction – still hold. Further, in practice, heterogeneous costs can aid in allowing firms to signal and cede one market to a competitor in exchange for another market (e.g., domestic vs international).

This example also underlines a lesson for protocol designers. With the assumption of heterogeneous costs, it is possible to construct examples where the repeated outcome is the same as the static outcome and the protocol period and field width are not of great importance. However, there still exists a large class of instances where the repeated case will be the relevant one. Because one can rarely be sure about the parameter range over which a designed protocol will be run, consideration of results presented in this paper are therefore important.

VII. DISCUSSION AND FUTURE WORK

The key conclusion of our work is that ***basic properties of the underlying protocol can have a significant impact on the equilibrium price.*** Robust to multiple assumptions and conditions, this leads us to several interesting conclusions and grounds for future work:

- 1) *Tools for Protocol Designers:* We have endowed the protocol designer with a set of new tools, perhaps previously hidden. For example, in a simple first-price setting we can achieve lower prices through a longer period (increase δ), a wider price field in the protocol (smaller b), and/or a less granular bandwidth using (smaller effective b). Holding other concerns aside, this means a consumer who has control over the protocol and seeks to limit price may find these useful.
- 2) *New Unavoidable Questions for Mechanism and Market Selection:* These tools are a double-edged sword as they

¹³Note that this is in addition to the reasons outlined in [21] which relate to information revelation. In this game, all information is common knowledge.

¹⁴Note that in this case $3 < p^*(c_2)$. So the profit to P2 is less than if both had cost c_2 .

raise questions of what the designer *should* do. For example, the interests consumers and suppliers may be at odds with each other. This could induce protocol alterations or issues of market and mechanism selection. This is a classic “tussle” [22].

Further troubling is that these parameters are *unavoidable*. Removing period restrictions creates a period implicitly defined by the players’ reaction time. Likewise, in any networked protocol, there is a maximum level of granularity. This poses interesting questions regarding the possibility of flexible and/or self-adapting protocols and frameworks.

- 3) *A Bridge Between the One-Shot and Repeated Models:* The parameters yield an understanding of the relationship between the one-shot and repeated routing games. The FPSS result and the counter-example presented from Section III are not different games, but rather two ends of a spectrum of games.
- 4) *The Importance of Repeated Games:* Another key insight from our work is the importance of repeated games. Almost all networking applications have an element of repetition. In this work we have seen that consideration of the repeated game is vital as the outcomes can be qualitatively different from the one-shot game. We have also seen that the repeated model can prove to be a useful analytical tool with conclusions of practical importance. Our current research explores repeated games in the context of additional practical networking problems.

Finally there are several assumptions in our models that can be relaxed. In this paper we do not examine the case of finite capacity, elastic demand, nor full networks of active players. We believe that the existence of non-trivial repeated equilibria and the relationships presented in this paper are robust to these relaxations. Nonetheless, there is room here for considerable future work. Furthermore, the correct equilibrium notions for repeated games is an open question; and it would be interesting to understand the interaction between model assumptions (e.g., asynchronous play) and the set of SPE strategies.

VIII. SUMMARY

In this paper we develop a model of incentive (or price) based routing that captures the notion of repetition, which is a vital aspect of practical applications. We see that the FPSS result does not directly hold here since it is not strategyproof in the repeated game. For a simple general model we are able to show that while prices can increase in general settings, their value is tied closely to certain, seemingly benign, properties of the underlying protocol. As such, we see that the protocol designer has greater control on the market than otherwise realized. We also show that these conclusions hold in more general settings, such as a 2nd price auction, the case of multiple destinations, and the case of heterogeneous costs. Taken together, these results present an interesting and novel relationship between routing protocol design and economic considerations of practical importance.

I would particularly like to thank Rahul Sami who provided invaluable support and advice for this research and this paper. I would also like to thank (in alphabetical order) Steve Bauer, Arthur Berger, Rob Beverly, Dave Clark, Peyman Faratin, Drew Fudenberg, Bobby Kleinberg, and John Wroclawski for their very valuable feedback at various stages of this research effort. Last, I'd like to thank the attendees at the poster sessions of PODC and SIGCOMM for their valuable feedback.

REFERENCES

- [1] Mike Afegan and John Wroclawski. On the Benefits and Feasibility of Incentive Based Routing Infrastructure. In *PINS '04: Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pages 197–204. ACM Press, 2004.
- [2] Nick Feamster. Practical Verification Techniques for Wide-Area Routing. *SIGCOMM Computer Communication Review*, 34(1):87–92, 2004.
- [3] R Farrow. Routing Instability on the Internet. *Network Magazine*, 2002.
- [4] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. *SIGCOMM Comput. Commun. Rev.*, 32(4):3–16, 2002.
- [5] David K. Goldenberg, Lili Qiuy, Haiyong Xie, Yang Richard Yang, and Yin Zhang. Optimizing Cost and Performance for Multihoming. In *Proceedings of ACM SIGCOMM 2004*, pages 79–92. ACM Press, 2004.
- [6] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. A BGP-Based Mechanism for Lowest-Cost Routing. In *PODC '02: Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*, pages 173–182, New York, NY, July 2002.
- [7] Srinivas Shakkotai and R. Srikant. Economics of Network Pricing with Multiple ISPs. In *Proceedings of INFOCOM 2005*, 2005.
- [8] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [9] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [10] D. Fudenberg, D. Levine, and E. Maskin. The Folk Theorem in Repeated Games with Imperfect Public Information. *Econometrica* 62, pages 997–1039, 1994.
- [11] Edward J. Green and Robert H. Porter. Noncooperative Collusion under Imperfect Price Information. *Econometrica*, (1):87–100, Jan 1984.
- [12] D. Fudenberg and D. Levine. Efficiency and Observability in Games with Long-Run and Short-Run Players. *Journal of Economic Theory* 62, pages 103–135, 1994.
- [13] G. Ellison. Cooperation in the Prisoner's Dilemma with Anonymous Random Matching. In *Review of Economics Studies* 61, pages 567–588, 1993.
- [14] Chrysanthos Dellarocas. Analyzing the Economic Efficiency of eBay-like Online Reputation Reporting Mechanisms. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 171–179. ACM Press, 2001.
- [15] P. Klemperer. What Really Matters in Auction Design. *Journal of Economic Perspectives*, (1):169–189, 2002.
- [16] Noam Nisan and Amir Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35(1):166–196, 4 2001.
- [17] J. Hershberger and S. Suri. Vickrey Prices and Shortest Paths: What is an Edge Worth? In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, page 252. IEEE Computer Society, 2001.
- [18] Staples: What is Staples' price-match policy? http://www.staples.com/content/help/using/general_match.asp.
- [19] Price Match Guarantee- Sears Canada Inc. http://www.sears.ca/e/customerservice/price_match.htm.
- [20] Roger Lagunoff and Akihiko Matsui. Asynchronous Choice in Repeated Coordination Games. *Econometrica*, 65(6):1467–1478, 1997.
- [21] Michael H Rothkopf, Thomas J Teisberg, and Edward P Kahn. Why Are Vickrey Auctions Rare? *Journal of Political Economy*, 98(1):94–109, February 1990.
- [22] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. In *Proceedings of ACM SIGCOMM 2002*, Pittsburgh, PA, August 2002.

Proof of Lemma 2

The protocol period and the market price are positively correlated – or $\frac{\partial p}{\partial \delta} > 0$. *Proof:* From equation (9) we have:

$$\frac{\partial p}{\partial \delta} = \frac{b}{(-1 + \delta)(-N + N\delta - \delta + 1)} \quad (15)$$

We seek to show that this ratio is positive. Clearly, the numerator $b > 0$ and the first term of the denominator $(-1 + \delta) < 0$.

Considering the other term of the denominator, we have:

$$\begin{aligned} (-N + N\delta - \delta + 1) &= N(\delta - 1) + (1 - \delta) \\ &= (1 - \delta) - N(1 - \delta) = (1 - \delta)(1 - N) \end{aligned}$$

Since $(1 - \delta) > 0$ and $(1 - N) < 0$, we have $(-N + N\delta - \delta + 1) < 0$. Thus,

$$\frac{\partial p}{\partial \delta} = \frac{+}{(-)(-)} > 0$$

as desired. ■

Proof of Theorem 3

Proof: Since σ is a SPE, we know that the one-stage deviation property must hold for every history and time step. Therefore, we examine a given strategy at a given decision point.

For any σ cooperating in a period yields $\frac{p_\sigma^t - c}{N}$ whereas deviating yields $p_\sigma^t - c$. Therefore, the benefit of deviating is:

$$B = (p_\sigma^t - b) - \frac{p_\sigma^t}{N}$$

The cost to this one-stage deviation is:

$$C = \delta \sum_{t > t_0} \delta^t (\sigma_h^t - \sigma_h^t)$$

Since $\sigma_h^t - \sigma_h^t < kb$, we therefore have:

$$C \leq \frac{\delta kb}{(1 - \delta)}$$

The one-stage deviation property holds iff $B \leq C$ or:

$$(p_\sigma^t - b) - \frac{p_\sigma^t}{N} \leq \frac{\delta kb}{(1 - \delta)}$$

We can rewrite this in standard form as:

$$p \geq (p - b)N(1 - \delta) + \delta(p - kb)$$

which we can solve to yield:

$$p \leq \frac{b(\delta N - N - \delta k)}{(\delta - 1)(n - 1)}$$

Further, if $\sigma_h^t - \sigma_h^t = kb$, then this expression holds with equality. ■

Proof of Theorem 5 For an instance of the RIRG with N players playing strategies in PP, there exists a tuple (δ, b) such that $p_m < c + \epsilon$.

Proof: From Theorem 3 we have that the highest possible price in PP_k is given by the strategy which punishes by a factor of k in each period which is in turn bound by:

$$p_m \leq \frac{b(\delta N - N - \delta k)}{(\delta - 1)(N - 1)} + c \quad (16)$$

This bound also holds for any equilibrium even when players play different SPEs as any punishment weaker than consistently punishing by k will only decrease the market price.

Therefore if we seek to have $p_m < c + \epsilon$, we can set:

$$\epsilon > \frac{b(\delta N - N - \delta k)}{(\delta - 1)(N - 1)} \quad (17)$$

which can be readily solved.

In particular, for a fixed b we have:

$$\delta < \frac{-\epsilon N + \epsilon + bN}{-\epsilon N + \epsilon + bN - bk} \quad (18)$$

and for a fixed δ we have:

$$b > \frac{\epsilon(\delta - 1)(N - 1)}{\delta N - N - \delta k} \quad (19)$$

where $x \geq y$. Since each side is monotonic in p , we consider only the case where $x = y$ to yield:

$$(p + y - c) \geq (1 - \delta)2(p - c) + \delta(p - y - c)$$

We can solve to obtain:

$$\frac{y(1 + \delta)}{1 - \delta} + c \geq p \quad (20)$$

Thus, in equilibrium, we have:

$$p^* = \frac{y(1 + \delta)}{1 - \delta} + c \quad (21)$$

And the VCG price, p_V is given by:

$$p_V = p^* + y = \frac{y(1 + \delta)}{1 - \delta} + 2c + y \quad (22)$$

Proof of Lemma 8

Proof: We take the partial derivative from Eqn(22):

$$\frac{\partial p_V}{\partial \delta} = \frac{y}{1 - \delta} + \frac{y(1 + \delta)}{(1 - \delta)^2} > 0 \quad (23)$$

since $0 < \delta < 1$.

Lemma 4: In the N-Player Repeated VCG Routing Game, the equilibrium price is given by:

$$p^* = \frac{yN(2N + \delta N - 2\delta - 1) + \delta(Nc - y - c) - y + c}{(1 - \delta)(1 - N)}$$

Proof: Define $y = c_h - c_l$.

We can expand $\pi(\cdot)$ based on its definition:

$$\pi_1(p_1, p_2) = \begin{cases} T(p_2 - c), & p_1 \leq p_2 - y \\ \frac{T}{2}(p_2 + y - c), & p_2 + y \geq p_1 > p_2 - y \\ 0, & \text{otherwise} \end{cases}$$

where $y = c_h - c_l$. Due to the deterministic selection in a tie, $\pi_2(\cdot)$ is slightly different:

$$\pi_2(p_1, p_2) = \begin{cases} T(p_1 - c), & p_1 < p_2 - y \\ \frac{T}{2}(p_1 + y - c), & p_1 + y > p_2 > p_1 - y \\ 0, & \text{otherwise} \end{cases}$$

These values follow directly from the VCG calculation. Note that per the VCG calculation, if $\pi_i(p_i, p_{-i}) > 0$ it is dependent only on p_{-i} .

We now turn to the question of the equilibrium conditions. For the purposes of this analysis we focus only on the first player (lexicographically). This is because the second player (which must make a larger sacrifice) will deviate if and only if the first player will. We know that in equilibrium, bidding p must be better than deviating by some x , or:

$$\frac{T}{2}(p + y - c) \geq (1 - \delta)T(p - c) + \frac{T}{2}\delta(p - x - c)$$