

A Game Theoretic Model for Network Upgrade Decisions

John Musacchio, Jean Walrand, and Shuang Wu

Abstract—We develop a game theoretic model for studying when interconnected Internet Service Providers (ISP)s will decide to upgrade their networks. We are in particular interested in network upgrades to better accommodate real-time applications like voice over IP, but our game model is applicable to other types of upgrade. Achieving the full benefit from network upgrades requires that enough ISPs upgrade so that the users' traffic encounters only upgraded ISPs as it traverses the Internet. This situation gives rise to two important effects. One effect is that providers are reluctant to be first to invest in an upgrade when the full benefit of the upgrade cannot be achieved until their peers do the same. The second effect is that when a provider upgrades, the quality of the Internet as a whole increases, increasing demand, and thereby increasing revenue for all ISPs – including those who have not invested in the upgrade. This “free-rider” effect gives providers a temptation to not upgrade and instead enjoy the benefits of their peers upgrading. Our game model takes into account both effects, and finds sufficient conditions for it to be a Nash equilibrium the ISPs involved in the game to decide to upgrade. There are a large number of other equilibria for this game, so we also find stronger conditions under which the Nash equilibrium in which all providers upgrade is the unique Nash equilibrium. Another important feature of network upgrade costs is that they are declining as technology improves. We extend our original game model for upgrades to take into account declining upgrade costs and study what effects this has on the game's outcome.

I. MOTIVATION AND BACKGROUND

Today's Internet is so successful in large part due to the simplicity of its service model. However, this simplicity is also the root one of its biggest failures: the failure to offer differentiated service quality to real time applications. Indeed, the research community has long recognized the utility of designing a network with differentiated classes of services in order to better serve applications that have different requirements for delay, throughput, and other quality of service metrics. For example when traffic from delay sensitive applications like voice over IP share a congested network with a large amount of TCP traffic coming from applications like web browsing, the effect could be that the quality of service seen by the voice users becomes unacceptably bad. A network with the same link capacities, but with mechanisms in place to protect the voice traffic from the congestion of the web traffic, could greatly increase the welfare of the voice users at a modest cost in welfare to voice users. Thus differentiated services have the potential to

increase the social welfare of the network, without increasing its capacity. Work on technical architectures for differentiated services includes [1], [2].

Researchers have also recognized that in order to keep the higher classes of service from becoming just as congested as the lower classes of service, a network with differentiated classes of service needs to charge higher prices for its higher service classes. Work on pricing of differentiated services includes [3], [4], [5], [6], [7]. Another important effect of differentiated pricing is that it can increase the revenue of the provider. Equivalently, differentiated pricing allows the provider to capture more of the increase in social welfare that occurs as a result of upgrading to a differentiated service model. As an upgrade to a differentiated service model would require investments in equipment and software, this increase in provider revenue is essential in giving the providers a sufficient incentive to make the upgrade.

While there has been a large amount of research activity in recent years addressing differentiated services and prices, recent trends have suddenly made these ideas of vital commercial importance. Voice over IP has been available for some time, but it can be argued that voice over IP has finally reached a “tipping point” and is now well on its way to becoming adopted by a significant fraction of the public [8]. Interest in IP television has also grown sharply in the past year. The recent explosion in use of the Internet for real-time or high bit rate applications has made the problem of addressing the Internet's service model shortcomings all the more urgent. Indeed the issue is now receiving attention by the large commercial providers, the popular press, and even the U.S. Congress [9], [10].

We argue that the principal obstacle to deployment of more advanced service models is that the network is owned by different providers, and that the full benefit of upgrading the network can only come when enough providers have upgraded so that when users make a connection across the network, all providers along the way are upgraded. There are two perverse effects these facts give rise to. One is that a provider does not want to upgrade to a differentiated services model if he cannot be sure that the neighboring providers will upgrade, and without his neighbors participating in the upgrade the effect of the upgrade will be reduced. The second effect is a “free-rider” effect. A provider that does not bother to upgrade still earns some benefit from his neighbors upgrading, because the upgrades of other providers improve the network as a whole, which should increase demand in general. For such a provider, it is not enough that the absolute benefit from upgrading exceed the costs; it is necessary that the *marginal* benefit of upgrading over free-riding exceed the

Research supported in part by NSF Grant ANI-0331659.

J. Musacchio and S. Wu are with the Technology and Information Management Program, University of California, Santa Cruz, Santa Cruz, CA 95064, USA {johnm, swu}@soe.ucsc.com

J. Warand is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720, USA wlr@eecs.berkeley.edu

upgrade costs. We propose a simple game-theoretic model in this paper to more fully understand these effects.

Our work has some relation to models found in the network formation literature. See [11] for a survey. In most network formation models, agents weigh the cost of joining over the benefit of joining, and typically the benefit of joining grows with network size. These models do not usually include free-riding effects where agents enjoy a partial benefit from not making the investment to join. Another closely related area of work is agency theory. In agency theory, like in our model, agents depend upon each other to put in effort in order to achieve a mutual reward, and there is a potential for agents to free-ride [12]. Typically in agency theory there is a principal to organize the interactions between the agents, and the models are usually single period. In contrast, our model assumes that there is no trusted principal to coordinate the actions of providers, and also we are interested in effects that happen over multiple periods of time.

II. GENERAL PROBLEM FORMULATION

As we have said, our goal is to develop a model for a network consisting of several interconnected Internet Service Providers (ISPs) to help understand when these providers will want to upgrade their networks. The model we formulate could be applicable to any type of upgrade, but the type of upgrade we are most interested in is an architectural upgrade of the network to better support applications with heterogeneous quality of service (QoS) needs. An example of such an architectural upgrade would be to have a priority queue for traffic from real time applications such as voice so that congestion from the web traffic does not delay voice packets. We presume that higher QoS, in particular for voice traffic, would increase the demand for the service and allow providers to earn more revenue. The upgrade might also allow a provider to charge different prices to voice and web traffic, which might further enhance the upgrading provider's revenue. However, no one provider may want to unilaterally undertake a costly upgrade to their network architecture, if the provider can not be sure that the other ISPs that they interconnect with will upgrade as well. The simplest instance of this situation is network consisting of just two interconnected providers. This is the model we investigate first.

III. TWO PROVIDERS

We first present the case where there are only two interconnected ISPs connected, and then we generalize the model to n ISPs. The game is a multi-period game, and rewards earned n periods in the future have a present value that is discounted by a factor of δ^n . In each period, each provider chooses between two actions – “Upgrade” and “Not-Upgrade”. The first time either provider chooses “Upgrade,” the upgrading ISP incurs a one-time charge of U to pay for the upgrade. When both providers have not upgraded we assume that each provider earns a payoff of 0 per period. In reality a network that is not upgraded still earns revenues, but without

	Upgrade	Not Upgrade
Upgrade	R, R	a, f
Not Upgrade	f, a	$0, 0$

TABLE I

THE PER-PERIOD REVENUES OF EACH PROVIDER IN THE TWO-PROVIDER NETWORK GAME.

loss of generality, we will “normalize” all other payoffs by subtracting the appropriate amount.

When one provider has upgraded, and the other has not, the upgraded provider earns a and the provider who did not upgrade earns f . Here, a is a mnemonic for early “adopter,” while f is a mnemonic for “free-rider” benefit. We assume that the early adopter ISP will earn enhanced revenue in an upgraded network, because the QoS improvement will increase the demand of voice users – both increasing the number of users that want to use the service, as well as increasing the price they are willing to pay. However, some of the benefits of this increased demand may also accrue to the free-rider ISP, because the ISPs are interconnected, and increased traffic on one ISP increases traffic on the other. The benefit to the free-rider ISP may not be quite as much as the benefit to the early adopter ISP, because the free-rider ISP using the old-style architecture may not have the ability to charge differentiated prices. We therefore assume $a > f > 0$.

When both providers in the network choose to upgrade, the improvement in QoS should be more than if only a single provider upgrades. Therefore, the subsequent increase in demand for voice traffic should also be greater. Thus if both providers have upgraded, they both earn a payoff of R , where $R > a > f$. The per-period revenues of the game are summarized by Table I.

Theorem 1 which we state and prove below demonstrates that there are many possible equilibria of a two-provider network game. In Theorem 1, and in the rest of this paper, we are interested in pairs of strategies, called strategy profiles, that are Subgame Perfect Nash Equilibria (SPE). In multi-stage games, a player's strategy is a plan of what action to take in any period, given what the player has observed happen in the past. Roughly, a SPE is a strategy profile for which each player's strategy is a best response to the other, and furthermore, a player cannot do better by deviating from his strategy as the game evolves over time. We will also be interested in when strategies are dominant strategies – meaning that they are a best response no matter what action(s) the other player takes. An SPE that consists of dominant strategies is more robust than an ordinary SPE in the following sense. We should expect rational players to actually play a SPE in dominant strategies, whereas rational players may not play an ordinary SPE if they don't have the correct expectation of what their opponents might do. See [13] for a more detailed discussion of these concepts

As one should expect, both providers upgrade in equilibrium, only if the present value of the “marginal” revenue between upgrading and free-riding, which is $\frac{R-f}{1-\delta}$, exceeds

the cost U of upgrading. However, the conditions for which upgrading is each provider's only strategy are stronger. An interesting case occurs when the early adopter benefit for upgrading exceeds the upgrade cost, but the marginal revenue between upgrading and free-riding does not. This corresponds to the condition $\frac{R-f}{1-\delta} < U < \frac{a}{1-\delta}$. In this case, it cannot be a SPE for both providers to upgrade, as one provider can improve his payoff by free-riding. Similarly, it cannot be a SPE for both providers to not upgrade, because a provider can improve his payoff by unilaterally upgrading. In this case, the only pure SPE is for one provider to upgrade and the other to not upgrade. One might be uncomfortable with such a solution because the provider's strategy spaces and payoff functions are symmetric, and yet there is an asymmetric outcome in equilibrium. However, we can find a symmetric SPE if we turn to mixed strategies, strategies where a provider chooses his action each period from his set of possible actions according to a probability distribution which may depend on the actions that provider has observed up until that point in time. Theorem 1 shows that there exists a symmetric mixed equilibrium where each provider chooses to upgrade with a probability α in each period, but sticks to free-riding if the other provider is seen to upgrade first.

Theorem 1: The two-provider network can have the following SPE depending on the payoff parameters as indicated below:

Upgrade Immediately: Both providers choose to upgrade in the current period. This is a SPE if $U \leq \frac{R-f}{1-\delta}$. Furthermore, "Upgrade Immediately" is an SPE in dominant strategies if

$$U < \min\left(\frac{R-f}{1-\delta}, \frac{a}{1-\delta}\right).$$

No First Upgrade: Each provider is unwilling to upgrade until the other upgrades, and consequently no provider upgrades. If one provider were to deviate by upgrading, the other provider would upgrade in the subsequent period. "No First Upgrade" is a SPE if $U \geq a - R + \frac{R}{1-\delta}$ and $U \leq \frac{R-f}{1-\delta}$.

Delayed Upgrade: Each provider waits until a period n to upgrade, where n is arbitrary. However, if one provider were to deviate by upgrading, the other provider would upgrade in the subsequent period. This is a SPE if

$$\frac{a}{1-\delta^n} + \frac{(\delta - \delta^n)R}{(1-\delta)(1-\delta^n)} \leq U \leq \frac{R-f}{1-\delta}$$

Never Upgrade: No provider ever upgrades. Even if one provider were to deviate by upgrading, the other provider would still not upgrade. "Never Upgrade" is the only SPE if $U > \frac{a}{1-\delta}$ and $U > \frac{R-f}{1-\delta}$.

Asymmetric Free-ride: One provider upgrades in slot 1, while the other provider never upgrades. "Asymmetric Free-ride" is a SPE if $\frac{R-f}{1-\delta} \leq U \leq \frac{a}{1-\delta}$.

Mixed Free-ride: In each subgame that begins with no provider upgraded, each provider upgrades with probability α . In a subgame that begins with one provider upgraded, the other provider remains not-upgraded. "Mixed Free-ride" is a SPE if $\frac{R-f}{1-\delta} \leq U \leq \frac{a}{1-\delta}$. (The value of α for which "Mixed Free-ride" is a SPE depends on the game data in a way we will make clear in the proof.)

Mixed Upgrade: In each subgame that begins with no provider having yet upgraded, each provider upgrades with probability α . In a subgame that begins with one provider upgraded, the not-upgraded provider upgrades. A sufficient condition for "Mixed Upgrade" to be a SPE is

$$a + \delta \frac{R}{1-\delta} \leq U \leq \frac{R-f}{1-\delta}.$$

(The value of α for which "Mixed Upgrade" is a SPE depends on the game data.)

Proof: We examine each equilibrium type introduced in the Theorem statement.

Upgrade Immediately:

Suppose both players follow the "Upgrade Immediately" strategy profile. If one player (call him player A) were to deviate from the profile by not upgrading, the deviating player would earn a payoff of $f + \delta f + \delta^2 f + \dots = \frac{f}{1-\delta}$. In contrast, if player A followed the "Upgrade Immediately" profile, he would earn a payoff of $\frac{R}{1-\delta} - U$. Player A would earn a payoff in-between these values by upgrading at some point in the future. Thus if $U \leq \frac{R-f}{1-\delta}$, player A's best response is to follow the "Upgrade Immediately" profile. Once upgraded, both players remain in the upgraded configuration because once the one-time upgrade fee has been paid, the revenues of an upgraded network dominate those of a not-upgraded network. Thus "Upgrade immediately" is a SPE if $U \leq \frac{R-f}{1-\delta}$.

The preceding argument shows that if $U \leq \frac{R-f}{1-\delta}$, that "Upgrade immediately" is a SPE. Now we will show that under a stronger condition that "Upgrade immediately" is the unique SPE by iterated strict dominance. Consider a subgame after one player has upgraded and the other has not. Then the player who has not upgraded (call him Player B without loss of generality) could choose not to upgrade now or ever, and earn a payoff of $\frac{f}{1-\delta}$. Player B could also choose to upgrade now and earn $\frac{R}{1-\delta} - U$. Upgrading at some point in the future produces a payoff in-between these two values. If $U < \frac{R-f}{1-\delta}$, then Player B's dominant strategy is to upgrade immediately.

Now consider a subgame in which no player has upgraded up until now. (Note that the original game is such a subgame.) If $U < \frac{R-f}{1-\delta}$, a player that upgrades now (call him Player A without loss of generality) knows that the other player (Player B) will have a dominant strategy of upgrading in the next turn. However, Player A cannot be certain what player B will do in the current period. Supposing player B does not upgrade in the current period, Player A's payoffs

under his two options are

$$\begin{cases} a + \delta \frac{R}{1-\delta} - U & \text{A upgrades now} \\ [0, \delta \frac{R}{1-\delta} - \delta U] & \text{A does not upgrade now.} \end{cases}$$

The second payoff is specified as a range, because A cannot know what B will do in the future if A does not upgrade now. The first option surely has a higher payoff if $U < \frac{a}{1-\delta}$. Supposing player B does upgrade in the current period, Player A's expected payoffs under his two options are

$$\begin{cases} \frac{R}{1-\delta} - U & \text{A upgrades now} \\ f + \delta \frac{R}{1-\delta} - \delta U & \text{A does not upgrade now.} \end{cases}$$

The second payoff takes into account that A would upgrade in the next period if B were in fact to upgrade now. The first of these options has a higher payoff if $U < \frac{R-f}{1-\delta}$. Thus A can be assured that upgrading now is the best action, regardless of what B will do this period, if U satisfies

$$U < \min\left(\frac{R-f}{1-\delta}, \frac{a}{1-\delta}\right).$$

Delayed Upgrade:

Suppose that $U \leq \frac{R-f}{1-\delta}$ and a player, which we call player A follows the Delayed Upgrade strategy, with a plan to upgrade in period n unless the opposing player (player B) upgrades first. In that case A would upgrade just after B. Player B has the following three options for her counter strategy. The first option is to upgrade in the beginning period, which would yield her a payoff of $(a - U) + \frac{\delta R}{1-\delta}$. This is because player A would be induced to upgrade in the subsequent period. The second option is to upgrade at time n which would yield a payoff of $\delta^n \left[\frac{R}{1-\delta} - U \right]$. Finally, the third option is to never upgrade which would yield a payoff of $\delta^n \frac{f}{1-\delta}$. There are other strategies, like waiting to upgrade until some arbitrary time before or after time n , but all such strategies have a payoff that lies between the payoffs of either the first and second options, or the second and third options we enumerated. Therefore to find conditions for player B to follow the Delayed Upgrade strategy at time n , we must check that upgrading in period n is at least as desirable as upgrading in period 0 (option 1), or never upgrading (option 3). These conditions reduce to

$$\frac{a}{1-\delta^n} + \frac{(\delta - \delta^n)R}{(1-\delta)(1-\delta^n)} \leq U \leq \frac{R-f}{1-\delta}.$$

No First Upgrade:

Suppose that a player, which we call player A, follows the "No First Upgrade" strategy profile as described in the theorem statement. We will evaluate the best response of player B. If player B upgrades, player A will upgrade in the subsequent period. Consequently, player B's payoff function has the form

$$\begin{cases} a - R + \frac{R}{1-\delta} - U & \text{B upgrades now} \\ 0 & \text{B does not upgrade.} \end{cases}$$

Note the form of first expression is due to the fact that $a + \delta R + \delta^2 R + \delta^3 R + \dots = a - R + \frac{R}{1-\delta}$. Thus, if $U \geq a - R +$

$\frac{R}{1-\delta}$, player B's best response would be to not to upgrade. Now suppose that player A were to deviate from the "No First Upgrade" profile by upgrading. We verify that player B would prefer to upgrade in this scenario, which is the action stipulated by the "No First Upgrade" profile. Player B would earn $\frac{f}{1-\delta}$ by not upgrading, and would earn $\frac{R}{1-\delta} - U$ by upgrading. Thus, player B would upgrade $U \leq \frac{R-f}{1-\delta}$. We conclude that "No First Upgrade" is a SPE if

$$a - R + \frac{R}{1-\delta} \leq U \leq \frac{R-f}{1-\delta}.$$

Never Upgrade:

Suppose that $U > \frac{a}{1-\delta}$ and $U > \frac{R-f}{1-\delta}$. A player, which we call player A, has no incentive to unilaterally upgrade because his net payoff for doing so would be negative. If the opposing player were to upgrade, player B would still not want to upgrade because the payoff from free-riding exceeds the payoff from upgrading. Thus, we have verified that under the conditions $U > \frac{a}{1-\delta}$ and $U > \frac{R-f}{1-\delta}$, never upgrading is the dominant strategy. Thus, under these conditions, "Never Upgrade" is the unique SPE by strict dominance.

Asymmetric Free-ride:

Suppose $\frac{R-f}{1-\delta} \leq U \leq \frac{a}{1-\delta}$. Under this condition a provider would prefer to unilaterally upgrade over having no one upgrading, but a provider would prefer even more strongly to free-ride by not upgrading and having the other provider upgrade. Thus both providers simultaneously upgrading is not a SPE, and both providers not upgrading is also not a SPE. Suppose a provider, which we call provider A, upgrades in period 1. Then provider B's payoff for not upgrading would be, $\frac{f}{1-\delta}$ and his payoff for upgrading would be $\frac{R}{1-\delta} - U$. As we have assumed $U \geq \frac{R-f}{1-\delta}$, B's best response is not to upgrade. This does not change if we consider subgames starting after period 1.

Suppose now B's strategy is not to upgrade. Then A's payoff for upgrading in period 1 is $\frac{a}{1-\delta} - U$, his payoff for never upgrading is 0, and his payoff for upgrading later is in between these two values. As we have assumed $U \leq \frac{a}{1-\delta}$, A's best response is to upgrade in period 1.

Thus we have shown that A and B's strategies are best responses to each other, and that "Asymmetric Free-ride" is a SPE.

Mixed Free-ride:

Again we suppose that $U \geq \frac{R-f}{1-\delta}$ and $U \leq \frac{a}{1-\delta}$. Unlike the "Asymmetric Free-ride" strategy profile, we will try to find a symmetric strategy profile that is a SPE. To do so, we will make use of mixed strategies. Suppose that in each subgame that begins with both players not-upgraded, player B upgrades with probability α . In other words, player B follows the "Mixed Free-ride" strategy described in the theorem statement. The expected reward to player A for upgrading is

$$\alpha \frac{R}{1-\delta} + (1-\alpha) \frac{a}{1-\delta} - U$$

because there is a chance α of player B upgrading simultaneously. We define J to be the expected reward for player A

not upgrading. The expected reward J satisfies the equation

$$J = \alpha \frac{f}{1-\delta} + (1-\alpha)\delta J.$$

This is because there is chance α of player B upgrading this turn giving a free-rider benefit to A, and with probability $1-\alpha$ we enter a subgame with a state identical to that of the current subgame. This equation algebraically reduces to

$$J = \frac{\alpha}{1-\delta+\alpha\delta} \cdot \frac{f}{1-\delta}.$$

The reward for player B upgrading minus the reward for player B not upgrading has the expression:

$$\frac{\alpha(R-a)+a}{1-\delta} - U - \frac{\alpha}{1-\delta+\alpha\delta} \cdot \frac{f}{1-\delta}. \quad (1)$$

When α is 1, expression (1) has value $\frac{R-f}{1-\delta} - U$, which is non-positive under the assumptions we have made for this case. When α is 0, expression (1) has value $\frac{a}{1-\delta} - U$, which we know is non-negative in this case. Because expression (1) is continuous in α , there exists an α^* , with $0 < \alpha^* < 1$ where expression (1) has value 0. (Note that the value of α^* can be solved in closed form, but the expression is quite complicated so we omit it here. For our purposes it is more important to know that such a value exists and is between 0 and 1.) Thus when player A upgrades with probability α^* in each subgame for which no player has upgraded yet, player B is indifferent between upgrading and not upgrading if no one has upgraded yet. A best response, though not unique best response, is for player B to upgrade with probability α^* in each subgame in which no one has upgraded yet. Because the payoff functions are symmetric, when both players play this strategy, no player has an incentive to deviate, and hence this is a SPE.

Mixed Upgrade:

Suppose that player B follows the ‘‘Mixed Upgrade’’ profile described in the theorem statement. Also suppose that $\frac{R-f}{1-\delta} \geq U$, so that if one player upgrades, the other player’s dominant strategy is to upgrade as well. We evaluate the best response of player B. The expected reward for player B upgrading is

$$\alpha R + (1-\alpha)a + \delta \frac{R}{1-\delta} - U.$$

We define J to be the expected reward for player B not upgrading. The reward J satisfies the following equation

$$J = (1-\alpha)\delta J + \alpha \left(f + \delta \frac{R}{1-\delta} - \delta U \right)$$

because if A were to upgrade, B would have a dominant strategy of upgrading in the next period, otherwise if A does not we enter a new subgame in the same state as the current subgame. The above equation reduces to

$$J = \left(\frac{\alpha}{1-\delta+\alpha\delta} \right) \left(f + \delta \frac{R}{1-\delta} - \delta U \right).$$

For player B to be neutral between the options of upgrading or not, the payoffs from the two options must be the same.

Equating the two payoffs, and isolating U algebraically, we find

$$U = \frac{1-\delta+\alpha\delta}{1-\delta} \left(\alpha R + (1-\alpha)a + \frac{\delta R}{1-\delta} \right) - \frac{\alpha}{1-\delta} \left(f + \frac{\delta R}{1-\delta} \right). \quad (2)$$

When α is set to 1, the right side of (2) has a value of $\frac{R-f}{1-\delta}$. When α is set to 0, the right side of the equation has a value of $a + \delta \frac{R}{1-\delta}$. Because the right side of (2) is continuous in α , for every U satisfying

$$a + \delta \frac{R}{1-\delta} \leq U \leq \frac{R-f}{1-\delta}$$

there exists an $\alpha \in [0, 1]$ for which player B is neutral between upgrading and not-upgrading. Therefore, a best-response, though not unique best-response, would be for player B to follow the ‘‘Mixed Upgrade’’ profile with the α satisfying (2). ■

IV. N PROVIDERS

We now consider the case in which there are $N > 2$ interconnected service providers. As we did when we analyzed the two player case in the proof of Theorem 1, we subtract the payoff each player earns in a period in which no player has upgraded from all payoffs. So without loss of generality, we may assume that the per-period payoff when no player has upgraded is 0. When there are more than two providers, the benefits to both free-riding ISPs and early-adopting ISPs will depend on the number other ISPs that have upgraded so far. Therefore, we define $f(j)$ to be the reward earned by a free-rider ISP if j ISPs have upgraded. Similarly, we define $a(j)$ to be the reward earned by an early adopter ISP when j ISPs have upgraded. To simplify notation, we adopt the convention that $a(N) = R$ and $f(0) = 0$.

We now state and prove Theorem 2 which verifies, as one might expect, that it is a SPE for all providers to upgrade if the marginal benefit of the last provider to upgrade is larger than the upgrade cost. We also show that in order for all providers upgrading to be the only SPE by iterated strict dominance, it must be that the marginal benefit of each provider upgrading, the quantity $\frac{a(j+1)-f(j)}{1-\delta}$ for each j to be precise, must be greater than U .

The following is a short explanation of iterated strict dominance. A strictly dominated strategy of a player is a strategy for which there exists another strategy that gives that same player a higher payoff no matter what the competing players do [13]. Iterated strict dominance means that for purposes of analysis we eliminate dominated strategies, and then furthermore we eliminate strategies that are dominated on this reduced strategy space. Each time we reduce the strategy space, we may find more strategies that are dominated, which is why it is called ‘‘iterated’’ strict dominance. If by this process we are left with one strategy profile, we say that we have a SPE by iterated strict dominance. Such an SPE is relatively robust in that rational players that know that if

their opponents are rational they should be expected to play such an SPE.

Theorem 2: If $U \leq \frac{R-f(N-1)}{1-\delta}$, then every provider upgrading in the first period is a SPE. If the stronger condition that

$$U < \min_{j=0,\dots,N-1} \left[\frac{a(j+1) - f(j)}{1-\delta} \right]$$

is true, then the only SPE is for every provider to upgrade by iterated deletion of dominated strategies.

Proof: Suppose that $U \leq \frac{R-f(N-1)}{1-\delta}$ and that all providers follow the strategy of upgrading in the beginning period. Any one provider can earn a payoff of $\frac{R}{1-\delta} - U$ for sticking to the strategy, a payoff of $\frac{f(N-1)}{1-\delta}$ for unilaterally deciding to never, upgrade, and a payoff in-between these values by upgrading at some point in the future. Under the condition on U that we supposed, sticking to the strategy of upgrading immediately is the best response. This shows that upgrading immediately is a SPE. We turn to verifying the stronger conditions under which it is a SPE by iterated strict dominance.

We use backwards induction – studying what happens if we reach a subgame where all but one providers have upgraded, and then studying what happens in a subgame where all but two have upgraded, and so on.

Suppose that $U < \frac{R-f(N-1)}{1-\delta}$ and that we begin a subgame with all but one provider upgraded. The best response for the one not-upgraded provider is to upgrade immediately.

We now proceed by induction. Suppose that we have shown that if

$$U < \min_{j=N-i,\dots,N-1} \left[\frac{a(j+1) - f(j)}{1-\delta} \right] \quad (3)$$

and that if we enter a subgame in which all but i providers have upgraded, that the remaining i providers should upgrade by iterated strict dominance. Note that we have just verified this for the $i = 1$ case.

Now suppose that we enter a subgame in which all but $i+1 \in \{1, \dots, N\}$ providers have upgraded, and also suppose (3) is true. Consider one of the providers, which we call provider A. If provider A upgrades, the other providers will be induced to upgrade in the following period. However, provider A cannot be sure what the other providers will do in the current period. Supposing that $k \in \{0, \dots, i\}$ other providers upgrade in the current period, A's payoff function has the form

$$\begin{cases} a(N-i+k) + \delta \frac{R}{1-\delta} - U & \text{A upgrades now} \\ f(N-i+k-1) + \delta \frac{R}{1-\delta} - \delta U & \text{A waits} \end{cases}$$

where the payoff for the ‘‘A waits’’ is an upper bound on the payoff A could achieve by not upgrading in the current period. By comparing the above payoffs for each k , we find that Provider A's payoff for upgrading is better than the payoff for not upgrading, independent of k , if

$$U < \min_{j=N-i-1,\dots,N-1} \left[\frac{a(j+1) - f(j)}{1-\delta} \right].$$

By induction, we conclude that the only SPE, by iterated deletion of dominated strategies, is for all providers to upgrade if

$$U < \min_{j=0,\dots,N-1} \left[\frac{a(j+1) - f(j)}{1-\delta} \right]. \quad \blacksquare$$

V. DECLINING UPGRADE COSTS

In our previous models, we assumed that the cost to upgrade a provider's network would be the same in any period. However as technology improves, the costs of upgrading the network should decline over time. Declining upgrade costs put a pressure on providers to postpone their upgrades until a time when it is less expensive. In this section we investigate how this effect influences the upgrade game.

We switch to a continuous time formulation, as this turns out to make the analysis easier and more intuitive. The revenue functions of the providers in this game are as follows. When both providers have upgraded, both players earn revenue of R per unit time. When one provider is upgraded, and the other is ‘‘free-riding’’, the upgraded provider earns a per unit time and the free-rider earns f per unit time. Without loss of generality when no provider is upgraded, both players earn no revenue. The present value of revenues earned t time units in the future is discounted by the factor $e^{-\delta t}$. For example if both providers decided to upgrade at time 0, the present value of their revenues would be $\int_0^\infty e^{-\delta t} R = \frac{R}{\delta}$. A provider that upgrades his network at time t' incurs a cost of $e^{-\gamma t'} U$. The factor γ includes both the effects of declining upgrade costs as well as discounting, so $\gamma > \delta$.

In order to compute Nash Equilibrium strategies of this game, we will first need to define two ‘‘critical’’ times t_f^* and t_a^* . Roughly, time t_f^* is the time after which a provider would prefer upgrading over free-riding. This notion will be made precise when we proceed to the proofs of this section. The value of t_f^* is given by

$$\begin{aligned} t_f^* &= \arg \max_{t \in \mathbb{R}^+} \left[e^{-\delta t} \frac{R-f}{\delta} - U e^{-\gamma t} \right] \\ &= \left[\frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{R-f} \right) \right]^+ \end{aligned} \quad (4)$$

This last expression for t_f^* is because the first derivative of the expression being maximized is $(R-f)e^{-\delta t} - U\gamma e^{-\gamma t}$, which has a unique zero at $t = \frac{1}{\gamma-\delta} \log \left(\frac{\gamma U}{R-f} \right)$. If the zero of the first derivative occurs at a $t < 0$, then the expression's maximum on \mathbb{R}^+ occurs at $t = 0$ because in this case the value of the first derivative can be shown to be negative for all nonnegative t .

The time t_a^* is the time after which a provider would prefer to be the first to upgrade (be an early ‘‘adopter’’). This notion will become more precise by the subsequent proofs. We define its value by

$$\begin{aligned} t_a^* &= \arg \max_{t \in \mathbb{R}^+} \left[e^{-\delta t} \frac{a}{\delta} - U e^{-\gamma t} \right] \\ &= \left[\frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{a} \right) \right]^+ \end{aligned}$$

The last expression for t_a^* can be derived in the same way as was done for expression (4).

Theorem 3: Suppose $t_a^* \geq t_f^*$. Then, the only SPE is one in which both providers upgrade at time t_f^* whether or not the other player is observed to upgrade earlier.

We state and prove two lemmas we will need to prove Theorem 3.

Lemma 1: Suppose provider A upgrades at time $t' \in [0, t_f^*]$. Consider a subgame starting at time $\tau \in (t', t_f^*]$ in which provider B has not yet upgraded. Provider B's best response is to upgrade at time t_f^* .

Proof: Provider B's best response in the subgame is to upgrade at time t satisfying

$$\begin{aligned} t &= \arg \max_{t \geq \tau} \left[e^{-\delta t} \frac{R}{\delta} - (e^{-\delta \tau} - e^{-\delta t}) \frac{f}{\delta} - Ue^{-\gamma t} \right] \\ &= \arg \max_{t \geq \tau} \left[e^{-\delta t} \frac{R-f}{\delta} - Ue^{-\gamma t} \right] \\ &= t_f^*. \end{aligned}$$

■

Lemma 2: Suppose provider A upgrades at time $t' \in [t_f^*, \infty)$. Consider a subgame that starts at time $\tau = t'$ and in which provider B has not yet upgraded. Provider B's dominant strategy in the subgame is to upgrade immediately.

Proof: Provider B's response in the subgame is to upgrade at time t satisfying

$$\begin{aligned} t &= \arg \max_{t \geq t'} \left[e^{-\delta t} \frac{R}{\delta} - (e^{-\delta t'} - e^{-\delta t}) \frac{f}{\delta} - Ue^{-\gamma t} \right] \\ &= \arg \max_{t \geq t'} \left[e^{-\delta t} \frac{R-f}{\delta} - Ue^{-\gamma t} \right]. \end{aligned} \quad (5)$$

The expression being maximized in (5) has a first derivative of $(R-f)e^{-\delta t} - U\gamma e^{-\gamma t}$, which can be shown to be negative for values of $t > t_f^*$. Therefore provider B's response is to upgrade at time t' . ■

We can now prove Theorem 3.

Proof: [Proof of Theorem 3] Suppose provider A's strategy is to upgrade at time t_f whether or not B is observed to upgrade earlier. Then provider B chooses his best-response upgrade time by comparing the best upgrade times on the intervals $[0, t_f^*]$ and $[t_f^*, \infty)$. These times are given by the expressions:

$$\begin{aligned} &\arg \max_{t \leq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} + (e^{-\delta t} - e^{-\delta t_f^*}) \frac{a}{\delta} - Ue^{-\gamma t} \right] \\ &= \arg \max \left[e^{-\delta t} \frac{a}{\delta} - Ue^{-\gamma t} \right] = \min(t_a^*, t_f^*) = t_f^*, \end{aligned}$$

and

$$\begin{aligned} &\arg \max_{t \geq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} - Ue^{-\gamma t} \right] \\ &= \min \left[t_f^*, \frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{R} \right) \right] = t_f^*. \end{aligned}$$

Hence, B's best response is to upgrade at time t_f^* . By Lemma 1, if B observes A upgrade before time t_f^* , his best response in the ensuing subgame is to wait until time t_f^* to

upgrade. Because our hypothesized SPE is symmetric, A's best response if B upgrades at time t_f^* is also to upgrade at time t_f^* . This shows that it is a SPE for both providers to upgrade at time t_f whether or not the other player is observed to upgrade earlier.

We now turn our attention to showing that this is the only SPE. We show by contradiction that in SPE, the first provider to upgrade must upgrade at time t_f^* . We consider two cases: that the first provider to upgrade does so before time t_f^* , and that the first provider to upgrade does so after time t_f^* .

Suppose there is a SPE in which the first provider to upgrade, which we refer to as provider A, upgrades sometime strictly before time t_f^* . By Lemma 1, provider B's best response is to upgrade at time t_f^* . As a result, provider A would prefer to deviate from his strategy of upgrading before t_f^* because his best response would be to upgrade at time

$$\begin{aligned} &\arg \max_{t \leq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} + (e^{-\delta t} - e^{-\delta t_f^*}) \frac{a}{\delta} - Ue^{-\gamma t} \right] \\ &= \arg \max_{t \leq t_f^*} \left[e^{-\delta t} \frac{a}{\delta} - Ue^{-\gamma t} \right] = \min(t_a^*, t_f^*) = t_f^*. \end{aligned}$$

This is a contradiction with our original assumption, and thus there can be no SPE in which the first provider to upgrade upgrades before t_f^* .

Now suppose there is a SPE in which the first provider to upgrade, which we will again refer to as provider A, upgrades sometime t' strictly after t_f^* . By Lemma 2, Provider B will be induced to upgrade immediately after provider A upgrades. Consequently, player A would prefer to deviate from his strategy of upgrading strictly after t_f^* because his best response would be to upgrade at time

$$\begin{aligned} &\arg \max_{t \geq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} - Ue^{-\gamma t} \right] \\ &= \min \left[t_f^*, \frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{R} \right) \right] = t_f^*. \end{aligned}$$

Thus there can be no SPE in which the first provider to upgrade upgrades after t_f^* .

Thus the only SPE can be one in which the first provider to upgrade does so at time t_f^* . We have seen that if a provider upgrades at time t_f^* , the best response of the other provider is also to upgrade at time t_f^* . Thus, the only SPE is the one in which both providers upgrade at time t_f^* . ■

We now state a theorem that addresses the case if $t_a^* < t_f^*$.

Theorem 4: If $t_a^* < t_f^*$, it is a SPE for one provider, which we call provider A, to upgrade at time t_a^* , and for the other provider (provider B) to upgrade at time t_f^* . If provider A were to see provider B upgrade before time t_a^* , provider A would wait until time t_f^* to upgrade.

In order to prove Theorem 4, we will need to state and prove two lemmas.

Lemma 3: Suppose $t_a^* < t_f^*$. Suppose provider A's strategy is to upgrade at time t_f^* . Then provider B's best response is to upgrade at time t_a^* , unless provider A were seen to deviate by upgrading before time t_a^* . In that case, provider B's best counter response would be to upgrade at time t_f^* .




Fig. 1. The dependence on t_f^* , the time until both providers are upgraded in SPE, on the rate of decline γ of upgrade costs and on the ratio $\frac{U}{R-f}$. The discount factor δ is fixed at 0.05 for all the curves illustrated.

Proof: Obviously, provider B's best time to upgrade must occur either on the interval $[0, t_f^*]$ or on the interval $[t_f^*, \infty)$, therefore provider B's best upgrade time must take the value given by one of the two following expressions:

$$\begin{aligned} \arg \max_{t \leq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} + (e^{-\delta t} - e^{-\delta t_f^*}) \frac{a}{\delta} - U e^{-\gamma t} \right] \\ = \arg \max_{t \leq t_f^*} \left[e^{-\delta t} \frac{a}{\delta} - U e^{-\gamma t} \right] = t_a^*, \quad (6) \end{aligned}$$

or

$$\begin{aligned} \arg \max_{t \geq t_f^*} \left[e^{-\delta t} \frac{R}{\delta} - U e^{-\gamma t} \right] \\ = \min \left[t_f^*, \frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{R} \right) \right] = t_f^*. \quad (7) \end{aligned}$$

Expression (7) shows that the best time to upgrade on the interval $[t_f^*, \infty)$ occurs at time t_f^* , which also lies on the interval $[0, t_f^*]$. Therefore, the best time to upgrade must occur on the interval $[0, t_f^*]$, and expression (6) shows that this best time is t_a^* .

If A were seen to deviate by upgrading before time t_a^* , then we would enter a subgame as described in the statement of Lemma 1. By Lemma 1, B's best counter response would be to upgrade at time t_f^* . ■

Lemma 4: Suppose $t_a^* < t_f^*$. Suppose provider B's strategy is to upgrade at time t_a^* if A is not observed to upgrade by that time, otherwise if A were to upgrade before t_a^* , B

upgrades at time t_f^* . The best response for provider A is to upgrade at time t_f^* .

Proof: Obviously, provider A's best time to upgrade must occur either on the interval $[0, t_a^*)$ or on the interval $[t_a^*, \infty)$. Provider A's best time to upgrade on the interval $[0, t_a^*)$ is given by

$$\begin{aligned} \arg \sup_{t < t_a^*} \left[(e^{-\delta t} - e^{-\delta t_f^*}) \frac{a}{\delta} + e^{-\delta t} \frac{R}{\delta} - U e^{-\gamma t} \right] \\ = \arg \sup_{t < t_a^*} \left[e^{-\delta t} \frac{a}{\delta} - U e^{-\gamma t} \right] = t_a^- \end{aligned}$$

where the “-” sign in t_a^- reminds us that the value was arrived at by finding the sup on the open set $[0, t_a^*)$. In terms of the game, the payoff of upgrading at time t_a^- is the limiting payoff of upgrading just enough before time t_a^* so that provider B would see it in time to postpone his upgrade until time t_f^* . Provider A's best time to upgrade on the interval $[t_a^*, \infty)$ is given by

$$\begin{aligned} \arg \sup_{t \geq t_a^*} \left[(e^{-\delta t_a^*} - e^{-\delta t}) \frac{f}{\delta} + e^{-\delta t} \frac{R}{\delta} - U e^{-\gamma t} \right] \\ = \arg \sup_{t \geq t_a^*} \left[e^{-\delta t} \frac{R-f}{\delta} - U e^{-\gamma t} \right] = t_f^*. \quad (8) \end{aligned}$$

Thus the highest payoff occurs either at time t_f^* or t_a^- . By (8), the payoff of upgrading at time t_f^* is greater than the

payoff of upgrading at time t_a^* , which is

$$e^{-\delta t_a} \frac{R}{\delta} - U e^{\gamma t_a^*}.$$

The payoff for upgrading at time t_a^* is

$$(e^{-\delta t_a^*} - e^{-\delta t_f^*}) \frac{a}{\delta} + e^{-\delta t_f} \frac{R}{\delta} - U e^{-\gamma t_a^*},$$

and thus the difference in payoffs between upgrading at time t_a^* and time t_a^* is $(e^{-\delta t_a} - e^{-\delta t_f^*}) \frac{R-a}{\delta}$, which is positive.

Combining these observations, we conclude that provider A's payoff for upgrading at time t_f^* is better than the payoff of upgrading at either times t_a^* or t_a^* , and thus A's best response is to upgrade at time t_f^* .

Now suppose that provider B has deviated from his expected strategy, and that provider B has not upgraded by time t_f^* . If provider A upgrades, provider B will be induced to upgrade immediately after A by Lemma 2. Thus provider A's best response in this subgame is given by

$$\arg \max_{t \geq t_f^*} e^{-\delta t} \frac{R}{\delta} - U e^{-\delta t}$$

which can be shown to be equal to t_f^* . ■

We are now ready to prove Theorem 4.

Proof: [Proof of Theorem 4] Suppose Provider A's strategy is to upgrade at time t_f^* . Then by Lemma 3, B's best response strategy is the following: upgrade at time t_a^* if A is not seen to upgrade before time t_a^* . Otherwise if A is seen to upgrade before time t_a^* , B would upgrade at time t_f^* .

Conversely, we suppose that B's strategy is to upgrade at time t_a^* if A is not seen to upgrade before then, and to wait until time t_f^* if A is seen to upgrade. Then by Lemma 4, A's best strategy is to upgrade at time t_f^* . ■

In both the case of $t_f^* < t_a^*$, covered by Theorem 3, and the case of $t_a^* < t_f^*$, covered by Theorem 4, it is not until time t_f^* that both providers are upgraded. It is useful to see how t_f^* depends on γ , the rate of decline of upgrade costs, and on the ratio of upgrade costs U to the "marginal" revenue of upgrading over free-riding, $R - f$. These dependencies are illustrated in Figure 1. The discount factor δ is fixed at 0.0 in the figure.

Figure 1 shows that for small enough values of $\frac{U}{R-f}$, the effect of having upgrade costs decrease more rapidly (γ larger), is that the time until providers upgrade in SPE increases. It is perhaps somewhat paradoxical that rapidly decreasing upgrade costs, corresponding to rapidly advancing technology, would actually increase how long it takes until providers are willing to upgrade. However, the correct intuition in this case is that providers are more willing to postpone investments if the costs will be less in the future. For larger values of $\frac{U}{R-f}$, the effect is reversed, more rapidly decreasing upgrade costs hastens the time until providers are willing to upgrade. Here the intuition is that if upgrade costs are prohibitively expensive today, rapidly decreasing costs will hasten the day when the upgrade costs are not too expensive.

Finally, we can use Figure 1 to understand the effect of the network being controlled by two providers versus if it were

controlled by one provider end to end. A single provider would choose his upgrade time to maximize $e^{-\delta t} \frac{R}{\delta} - U e^{-\gamma t}$ which is given by $\left[\frac{1}{\gamma - \delta} \log \left(\frac{\gamma U}{R} \right) \right]^+$. Therefore the optimum time to upgrade for a single provider would be given by the curves of Figure 1, except with the abscissa taken to be $\frac{U}{R}$ instead of $\frac{U}{R-f}$. Therefore the effect of switching from one provider to two providers is like moving to the right on Figure 1 by an amount related to the free-rider benefit f . This has the effect of increasing the time we have to wait until the network gets upgrades, but it may also push us from a regime where improving the rate of technology development γ retards network upgrades to a regime where improvements in the rate of technology development hastens network upgrades.

VI. CONCLUSION

We have developed a simple game-theoretic model for understanding the interactions between interconnecting ISPs that are contemplating upgrading their networks. We have shown that there are a large number of equilibria that are possible and that fairly strong conditions are required for all providers upgrading to be the unique SPE by iterated strict dominance. We have also studied the effect that declining upgrade costs has on the game. When the ratio of the upgrade cost to the marginal revenue of upgrading over free-riding is sufficiently small, increasing the rate at which upgrade costs decline actually had the effect of delaying when the providers upgrade in Nash equilibrium. We also see that the amount of time until the network is upgraded is larger for a two-provider network than it would be for a single provider network.

There are several possible directions for future work, including understanding the upgrade game when improvements in one network can reduce the profitability of another network through competition. Note that in our model it is possible that the networks both compete as well as complement each other, so long as the effect of one provider upgrading has a net positive effect on the other's revenues. A more general model would allow for the net effect to be negative. Another extension would be to consider a pricing game that would ensue after networks made their upgrade decisions. In our model, we assumed that that a provider would earn a revenue that is a function of only the number of providers that have upgraded. In reality, following the upgrade decisions there may be a pricing game that can have many possible Nash Equilibria. The uncertainty about the outcome of the ensuing pricing game may effect the provider's willingness to upgrade. Note that other researchers have studied the behavior of multistage games that have a pricing stage that is preceded by a capacity investment stage. For instance, [14] considers a two-stage game where providers in competition make a capacity decision followed by a pricing decision.

REFERENCES

- [1] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," RFC 2638, IETF, July 1999.

- [2] Y. Bernet, P. Ford, R. Yavatkar, F. Baker, L. Zhang, M. Speer, R. Braden, B. Davie, J. Wroclawski, E. Felstaine, "A Framework for Integrated Services Operation over Diffserv Networks," RFC 2998, IETF, November 2000.
- [3] M. Mandjes. "Pricing Strategies and Service Differentiation," *Netnomics* 6: 5981, 2004.
- [4] L. He and J. Walrand. "Pricing and Revenue Sharing for Internet Service Providers," To appear in *Journal on Selected Areas of Communication*, 2006.
- [5] J. Shu, P. Varaiya, "Pricing Network Services," *IEEE Infocom 2003*, San Francisco, CA, April, 2003.
- [6] J. MacKie-Mason and H. Varian, "Pricing the Internet," in *Public Access to the Internet*, B. Kahin and J. Keller, Eds. Englewood Cliffs, NJ: Prentice-Hall 1995.
- [7] A. Odlyzko, "Paris Metro Pricing for the Internet," In *Proceedings of ACM Conference on Electronic Commerce*, pp 140-147, 1998.
- [8] B. Bergstein, "Number of Internet-phone consumers soars," *USA Today*, August 17, 2005.
- [9] H. Bray, "Telecoms want their products to travel on a faster Internet," *Boston Globe*, December 13, 2005.
- [10] V. Kopytoff, "Panel dumps Net neutrality / House committee drops amendment banning two-tier Internet," *San Francisco Chronicle*, April 27, 2006.
- [11] M.O. Jackson, "A Survey of Models of Network Formation: Stability and Efficiency," Chapter 1 in *Group Formation in Economics; Networks, Clubs and Coalitions*, edited by G. Demange and M. Wooders, Cambridge University Press, Cambridge U.K., 2004.
- [12] M. Babaioff, M. Feldman, and N. Nisan, "Combinatorial Agency," *ACM Conference on Electronic Commerce 2006 (EC'06)*: 18-28.
- [13] D. Fudenberg and J. Tirole, *Game Theory*, MIT Press, Cambridge, MA, 1991.
- [14] D. Acemoglu, K. Bimpikis, A. Ozdaglar, "Pricing and Capacity Competition," *Allerton Conference on Communication, Control and Computing*, Monticello, IL, 2006.