

# On Reciprocal Altruism and Its Application to QoS

Nupur Kothari, Vartika Bhandari, and Dheeraj Sanghi

Department of Computer Science & Engineering,  
Indian Institute of Technology Kanpur,  
Kanpur, India  
{nupur.kothari, vartika.bhandari}@alumni.cse.iitk.ac.in,  
dheeraj@cse.iitk.ac.in

**Abstract.** With the growing usage of the Internet for resource-intensive applications e.g., streaming multimedia, best-effort service has ceased to be adequate. The result has been a spurt of proposals for guaranteed quality-of-service (QoS). Since network resources can be but limited, supply very often falls short of demand, thereby leading to competition to secure available resources. The factors that come into play in such situations, are network mechanisms, as well as complex economic behavior. A framework for providing cost-effective QoS needs to address both these aspects. Selfishness as a guiding motivation for action is widely observed in nature and has also been applied to QoS in the form of approaches based on competitive game theory. However, as has been observed in the natural world, long-term selfish motives may also give rise to apparently altruistic actions. This notion is well-captured in the phenomenon of Reciprocal Altruism, and has been modeled in game theory as the Iterated Prisoner's Dilemma. We propose Reciprocal Altruism as the guiding principle for a QoS framework that allows for cooperation between otherwise competing flows, leading to long-term benefit for all. We also present simulation results to validate the notion that cooperation can lead to better end-user experience.

## 1 Introduction

With the growing usage of the Internet for resource-intensive applications e.g., streaming multimedia, best-effort service has ceased to be adequate. An emphasis on guaranteed Quality-of-Service (QoS) has emerged, propelled by the increasing volumes of traffic with stringent QoS requirements. The existing Internet infrastructure is not equipped to handle such services. Hence various frameworks e.g., IntServ [1], DiffServ [2] etc. have been proposed to provide QoS. While these address the mechanisms required to provide QoS guarantees, the policies for resource allocations made, are beyond their purview. By the basic tenets of economics, network users need to pay for any guarantees sought. With limited availability of network resources, intense competition can emerge to secure the above guarantees. When the peak-rate of data transmission is

significantly different from the average rate, it becomes difficult to maintain consistent quality without the risk of over-allocation and consequent over-expenditure. This may also lead to under-utilization of resources.

Real time applications like streaming video, VoIP, online gaming, etc., often exhibit variable traffic rates, with the peak rate significantly higher than the average rate. This leads to time-varying bandwidth requirements. Variable bandwidth requirements may also be imposed by the use of protocols like Split-and-Merge [3]. The perceived quality of the transmitted content is of great importance in such applications. While it is not an easy task to map network QoS to user QoS, yet it has been observed that perceived quality is affected by such factors as Packet Loss Ratio (PLR) and the loss burst lengths (i.e. number of consecutive packets lost). The relationship between perceptual quality and PLR has been studied in numerous works e.g. [4] and [5]. Voice traffic exhibits extreme loss sensitivity with a PLR of over 2% causing appreciable voice degradation. For video streams, the loss of different kinds of packets have different end-effects, depending on the encoding. As an example, in MPEG video, certain frames are critical and as long as they are not lost, it may be possible to achieve acceptable quality even with PLRs of the order of 20%. These observations can form the basis of a notion of *survival* for multimedia flows over discrete time intervals corresponding to the transmission of a Group of Pictures (GOP). In order to ensure good quality, survival rates should be high over the flow duration, even with fluctuating data rates. At the same time, network utilization and economic concerns need to be taken into account.

Attempts have been made to partially address these issues by devising dynamic provisioning and pricing schemes to allow for adaptive QoS negotiation. Some techniques for enforcing conformance of encoder output to the pre-decided traffic envelope have also been proposed [6]. An alternative approach is to retain a simple provisioning and pricing structure with reasonably long-term service agreements, and accommodate transient bursts within this framework, so that flows are able to survive such burst phases. We propose such an approach wherein flows can go in for moderate resource allocations and cooperate with each other to tide over transient bursts and maintain consistent end-quality. Our proposal envisages a cooperative game theory based approach, as opposed to earlier proposals for application of competitive game theory, as in [7], [8] and [9], amongst others.

## 2 Characteristics of Video Traffic

There are generally three types of frames in encoded video. In the case of MPEG-4 encoding, they are referred to as the Intra Frame (I-Frame), Predictive Frame (P-Frame), and Bidirectional Frame (B-Frame). An I-Frame is the most basic frame used in video compression and stores all the data required to display the frame. A P-Frame is smaller in size as it builds upon the previous I or P-Frame

and only stores the difference. A B-Frame is built upon two frames, one I or P before it and one I or P after it. Therefore the loss of a P-Frame or an I-Frame would affect all P and B-Frames dependent on it. In fact a notion of Group of Pictures (GOP) exists wherein a GOP is a block of frames comprising one or more I Frames followed by P or B-Frames. It is largely self-sufficient in that inter-GOP dependence is either non-existent or restricted to a single B-Frame. Hence, losses in one GOP do not really affect the quality of other GOPs. These observations about video traffic open up the possibility of looking at QoS as a per-GOP phenomenon.

### 3 Reciprocal Altruism

Altruism is defined as social behavior that benefits an unrelated individual while being detrimental to the individual displaying altruism. Motivations for altruistic behavior may be varied. However, it has been observed that often apparently altruistic behavior is actually motivated by self-interest. This is termed as *Reciprocal Altruism* [10]. This kind of behavior is characterized by the acceptance of short-term costs/losses in the expectation of a long-term benefit based on reciprocation of the *altruistic* gesture by the current beneficiary.

*Reciprocal Altruism* has been observed in numerous organisms e.g. vampire bats (*Desmodus rotundus*) [11]. Vampire bats have a very high metabolic rate and hence starve to death if unable to find food for two days running. As such, it has been observed that a bat that was unable to find food solicits the same from a roost-mate and is often helped based on the expectation that if the benefactor ever fails to find food, the current beneficiary would be willing to help in return.

*Reciprocal Altruism* can sustain itself as a long-term behavioral model only if the gain to beneficiary is much more than the cost to the benefactor. This does happen in the case of vampire bats, since bats with large *margin-to-death* donate food to a starving bat who is close to death. *Reciprocal Altruism* has been modeled in Game Theory as a non-zero sum game, where it is possible to come up with a situation where all parties stand to gain.

#### 3.1 Iterated Prisoner's Dilemma

The Iterated Prisoner's Dilemma (IPD) is a classic example of a non-zero sum game. The two-person Prisoner's Dilemma (PD) is characterized by two prisoners being interrogated in isolation about a joint crime. Each prisoner has the choice of either confessing (defection) or maintaining silence (cooperation). The payoff matrix is as shown in Figure 1. The payoff here denotes the years of imprisonment they will receive. As can be seen, if both cooperate, they can get away with a minimal sentence. However, since each has no means of knowing what the other will do, the most obvious choice would be to defect, and this is also the Nash

		B	
		Cooperate	Defect
A	Cooperate	1, 1	20, 0
	Defect	0, 20	10, 10

**Fig. 1.** Payoff Matrix for Prisoner’s Dilemma showing the possible jail terms (A’s term, B’s term) prisoners A and B may get (cooperate *implies* be silent, defect *implies* testify)

equilibrium for this game. The IPD extends the PD into a multi-move game. In this case, the prisoners can derive from previous history to decide their next move. The Alternating IPD is a variant in which the participants do not make their next move simultaneously but in turn.

### 3.2 Reciprocal Altruism as an Alternating IPD

Reciprocal Altruism has been modeled in game theory as an Alternating IPD. Various strategies for determining the next move in the Alternating IPD have been studied. Axelrod and Hamilton [12] ran a tournament whereby they attempted to determine the best strategy, which in their case was simple Tit-For-Tat (TFT) with an initial cooperating move. Another mechanism for determining an optimal strategy is an evolutionary game, in which individual strategies evolve till only the most evolutionarily stable one(s) remain. When the environment is noisy, a misunderstanding may arise between individuals asking for/giving help. In such a case, a strategy like TFT fails to achieve cooperation. Hence introduction of an element of clemency is required.

## 4 Reciprocal Altruism as a QoS Paradigm

The great intra-GOP dependence between frames implies that the post-reconstruction quality of received frames belonging to a GOP is inextricably interlinked with each other. The inter-GOP independence indicates that packet loss in one GOP has no effect on the end-quality of other GOPs. This suggests the treatment of a GOP as an atomic unit having either acceptable or unacceptable quality from the end-user view point. Borrowing terminology from the biological world, we therefore introduce a notion of *survival* wherein a flow either survives or dies for a particular GOP. *Survival* corresponds to having acceptable end-quality, whereas *death* corresponds to the contrary. In this section we formalize this notion and propose a paradigm for Reciprocal Altruism amongst flows that would lead to higher survival rates for the participant flows.

We define certain notions that shall be used throughout this paper.

**Definition 1. (Renewal Time)** *Renewal Time( $T_r$ ) is the time span over which the effect of loss of crucial packets lingers. After this time has elapsed, the flow can renew itself to its original quality.*

**Definition 2. (Instantaneous PLR)** *Instantaneous PLR ( $PLR_i$ ) is defined as the ratio of the number of packets lost ( $lp_i$ ) to the total number of packets received ( $tp_i$ ) over a time span ( $t_i$ ) of duration  $T_r$ .*

**Definition 3. (Survival)** *As long as the Instantaneous PLR remains below a certain threshold ( $PLR_S$ ), the perceived end-quality is acceptable over that time span ( $t_i$ ). This condition is termed as survival during the time span  $t_i$ .*

**Definition 4. (Survival Rate)** *Survival Rate ( $\sigma$ ) for a flow (over the consecutive time intervals  $t_1, \dots, t_N$  of length  $T_r$  may be defined as*

$$\sigma = \sum_{i=1}^N S_i / N \text{ where } S_i = \begin{cases} 0, & \text{if } PLR_i \geq PLR_S; \\ 1, & \text{otherwise;} \end{cases}$$

The basic premise of our proposal is that *consistent end-quality may be quantified by the survival rate of a flow over its entire duration*, and each flow would seek to maximize this rate. However a sudden burst of traffic can lead to high losses ( $PLR_i$ ) and resultant failure to survive (flow death) in  $t_i$ . The core concept behind this notion of survival is that the perceptual quality of the reconstructed video at the receiver's end should conform to some minimum level. Given a particular GOP, its reconstruction becomes extremely difficult either if an I Frame packet is lost or if the PLR becomes rather high. In order to avoid this, a flow may seek help from other flows that have low PLR. These flows may decide to temporarily lend their resources at some (but not critical) expense to themselves, in expectation of future reciprocation. At the same time, we recognize a concept of *Posthumous Donation (PD)* whereby a flow that has already incurred heavy packet losses may decide against sending further packets for the current GOP (i.e. the current  $T_r$ ) and instead grant assistance to another flow in need. Such action would be motivated by the fact that the current GOP is already beyond reconstruction. Hence it is advisable to garner good-will for the future. Notions similar to PD have been proposed earlier to avoid sending useless packets [13].

## 5 Mechanisms for Cooperation

It is possible to incorporate *Reciprocal Altruism* into QoS scheduling policies via two broad categories of approaches.

A possible mechanism for cooperation is to deploy intelligent agents that act on behalf of each flow. These agents would monitor flow and network state. Based on decision parameters specified by the flow and the current state, these would make decisions regarding obtaining/providing help. There are numerous issues

related to inter-flow communication etc. that need to be studied to implement such a mechanism. This has not been looked at in this paper.

A more simple mechanism comprises the assumption of the task of enforcing cooperation by a centralized authority (say the ISP). Flows may specify their survival thresholds and other requirements. The centralized authority (which would have to be trusted) thereafter decides on behalf of each flow about the need to seek help as well as the feasibility of admitting a request for help. Such a mechanism would have the advantage of having lower overheads as all the functionality would be built into the service discipline deployed by the service-provider. This paper presents results for a preliminary implementation of such a centralized scheme, that assumes complete trust and willingness-to-help.

## 6 Validation by Simulation

We present a preliminary validation of our proposal by simulating simplistic schemes for centralized imposition of reciprocal altruism. We analyze the effect of the same on flow performance.

### 6.1 Simulation Model

We consider a scenario wherein flows contend for bandwidth on a single link. We first consider a situation where there are only two flows. The topology corresponding to this scenario is depicted in Fig. 2. Links  $l_1$  and  $l_2$  are of 3.5 Mbps capacity each. The bottleneck link is  $l_3$ . Flows  $f_1$  and  $f_2$  each have a certain reserved bandwidth allocation  $a_1$  and  $a_2$  respectively over  $l_3$ . It is assumed in all simulations that  $a_1 + a_2 = capacity(l_3)$ . We assume lossless links. The traffic corresponds to MPEG-4 traces of various movies [14] as listed in Table 1, sent over UDP/IP. A constant packet size of 200 bytes (data+headers) is considered. We assume the use of the TOS field of the IP header to mark the type of frame (I, P or B) to which each packet belongs.

The simulator used is ns-2 [15] into which we have built our own scheduling module. We assume a simple scheduling policy loosely modeled on Round Robin scheduling, wherein each flow gets assigned a quantum equal to its legitimate bandwidth share over a certain time interval. The quantum decreases with each packet sent. At the end of the interval, the remaining quantum is flushed, and quantum is re-assigned for the next interval. Bandwidth unused by a flow may be used by the other if it has a packet to send. The time interval for quantum assignment coincides with the *renewal time*  $T_r$  defined in Section 4. The value of  $T_r$  is set to 0.45s which is the approximate display time of one GOP (IPBBPBBPBB) in the MPEG-4 traces used. However, this is only an approximate measure employed for preliminary validation and needs to be replaced by a mechanism that conforms exactly to GOP boundaries.

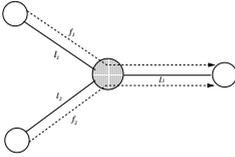


Fig. 2. Simulation Topology 1

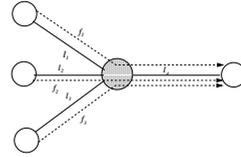


Fig. 3. Simulation Topology 2

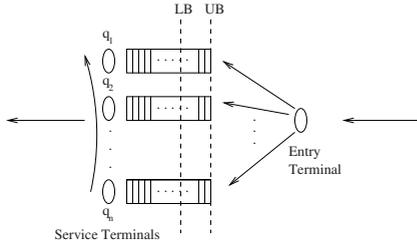
Table 1. Traffic Characteristics of MPEG-4 Traces Used

Flow	Content	Encoding	Mean BR	Peak/Mean
$f_1$	<i>Star Wars IV</i> Video	MPEG-4	1.9e+05	6.81
$f_2$	<i>Jurassic Park</i> Video	MPEG-4	7.7e+05	4.37
$f_2$	<i>The Firm</i> Video	MPEG-4	2.9e+05	6.96

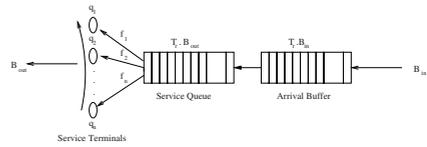
We have built in *Reciprocal Altruism* atop this strategy by tracking intra- $T_r$ -PLRs and estimating the need for seeking help, as well as the feasibility of granting it. The architecture of the scheme is depicted in Fig. 4. The scheduler maintains per-flow queues with two occupancy limits defined, viz.  $LB$  and  $UB$ .  $LB$  is the lower bound and as long as the queue length remains below it, all incoming packets are enqueued. As soon as the length exceeds  $LB$ , an arrival burst is assumed. If the length is equal to  $UB$ , the packet is immediately dropped. However, when the queue length lies between  $LB$  and  $UB$ , if PLR is within acceptable limits, the packet may be dropped, unless it is an I frame packet. In that case, the flow attempts to accommodate it by dropping a lower priority packet (P/B Frame) from the queue. If the PLR is already hovering at dangerous levels, the flow may seek help from other flows in the form of a promise for an extra transmission turn(s). If such a promise is forthcoming from any other flow, the packet is retained, in the belief that the backlog shall soon return to normal levels (i.e.  $< LB$ ). It is to be noted that a promise is not binding, in that the current PLR levels of the benefactor flow are checked prior to actually passing on its turn(s). The turn is passed on only if PLR is still within acceptable limits, thereby avoiding situations in which an earlier mis-estimation might lead to critical losses for the benefactor flow. The value of  $PLR_S$  is set to 0.2. This corresponds to a figure often used as a loose upper bound on acceptable video losses. At  $PLR > 0.2$ , the end-quality generally becomes unacceptable even with error-correction etc. However, it is to be noted that the value of  $PLR_S$  may be different for different flows depending on how stringent their quality requirements are, and may vary significantly from the value used in this paper. The values of  $UB$  and  $LB$  used for the simulations presented here are 40 and 100 respectively. The pseudo-code for the entry terminal and service terminal procedures are shown in the next page.

We have also looked at a buffer-based scheme for implementing *Reciprocal Altruism*, as a reference for comparison with the estimation-based strategy. This

scheme utilizes a buffer to obtain a lookahead of one  $T_r$  and makes advance decisions on exchange of help for this upcoming interval.



**Fig. 4.** Estimation Based Scheme



**Fig. 5.** Buffer Based Scheme

We have also performed basic simulations for three-flow interactions by simulating the topology depicted in Fig. 3, wherein flows  $f_1$ ,  $f_2$  and  $f_3$  contend for bandwidth on bottleneck link  $l_4$ . Links  $l_1$ ,  $l_2$  and  $l_3$  are of 3.5 Mbps capacity each. Allocations  $a_1$ ,  $a_2$  and  $a_3$  are all equal and  $a_1 + a_2 + a_3 = capacity(l_4)$ .

### 6.2 Results and Analysis

We present here the results obtained from simulation. Figs. 6 and 7 depict the survival rates for all participant flows (each having equal allocation) with variation in bottleneck link bandwidth for the topologies in Figs. 2 and 3 respectively. All flows correspond to runs of the MPEG-4 trace of *Star Wars IV* with varyingly staggered start times. All three schemes, viz. non-altruistic, altruistic-estimation-based and altruistic-buffer-based are simulated, and all flows are allocated equal shares of the link bandwidth. It may be seen that the estimation-based scheme consistently tends to perform well in terms of having higher survival rates. It significantly out-performs the other two schemes at low link bandwidths. However at higher link bandwidths, the buffer-based scheme performs rather well, especially in the three-flow case.

Since the above results point to similar trends for 2-flow as well as 3-flow interactions, it seems that the results for 2-flow interaction are indicative of general trends in n-flow interaction. Hence further simulations have only been carried out for the two-flow case. Besides, only the non-altruistic scheme and the estimation-based scheme have been considered.

Figs. 8 and 9 show the distribution of bytes received and bytes sent per interval of duration  $T_r$  for Flow 0 over a total duration of approximately 3500s for MPEG-4 traces of *Star Wars IV* staggered by 90.0s for the non-altruistic and estimation-based schemes respectively. The bottleneck link bandwidth is 0.6 Mbps and both flows have allocations of 0.3 Mbps each (that also corresponds closely to their mean rates). As may be clearly seen, in the former case, the

```

handle_arrived_packet(packet p){
    f= flow_id(p);
    plr=pkts_drpd_sofar_thisTr/(pkts_drpd_enqueue_packet(p,flowqueue(f));
    sofar_thisTr+pkts_sent_sofar_thisTr);
    if((bytes_drpd_sofar_thisTr>0.1*
    quantum_allocated) and
    (plr>death_thresh(f))
    &&(frametype(p)!=I))
        posthumous_donation(f)=yes;
    if(length(flowqueue(f))>UB){
        drop_packet(p);
        return;
    }
    if(length(flowqueue(f))>LB){
        if(plr>=0.7*death_thresh(f))
            alert=1;
        if(frametype(p)==I){
            \\Giving priority to I
            sacrificial_lamb=look_for_some_
            PorB_packet(flowqueue(f));
            if(sacrificial_lamb_not_found)
                alert=1;
            else if(alert==0)
                drop_packet(sacrificial_lamb);
        }
        else if(alert==0){
            \\If no alert then drop
            drop_packet(p);
            return;
        }
    }
    if(posthumous_donation(f)==yes)
        \\if dead no alert
        alert=0;
    if(alert){
        seek_help(f);
        if(!help_obtained){
            drop_packet(p);
            return;
        }
    }
    seek_help(flow f){
        for each flow g!=f{
            if(posthumous_donation(g)==yes){
                help_obtained=1;
                return;
            }
            plr=pkts_drpd_sofar_thisTr/
            (pkts_drpd_sofar_thisTr
            +pkts_sent_sofar_thisTr);
            if(((plr/death_thresh(g))*
            (length(flowqueue(g)/LB)<0.3)and
            (length(flowqueue(g))<0.8*LB)and
            (plr<0.8*death_thresh)and
            (not_already_giving_help)){
                help_obtained=1;
                return;
            }
        }
        service_flowqueue(flow f){
            if(giving_help(f))
                service(beneficiary(f));
            else
                if(!empty(flowqueue(f))
                service(f);
                turn=next_flow;
            }
        }
        revolution(){
            service_flowqueue(turn);
        }
    }
}

```

number of bytes sent per interval of duration  $T_r$  remains close to the allocation. In the latter case, it may be seen that the flows are able to send at rates higher than their allocation a greater number of times, consequent to their cooperation. For non-altruistic scheduling, the number of time-intervals in which bytes sent were greater than allocation was 2800, whereas for altruistic scheduling it was 3097, i.e. an increase of around 10%.

Fig. 10 depicts flow behavior trends for a wide range of bottleneck link bandwidths, as well as varying proportion of allocation of the same to the flows. The figure essentially illustrates the survival rate against the individual allocations of both flows. Though in a real world situation, flows may not actually go in for such a widely varying allocation spectrum, it is useful to study the same.

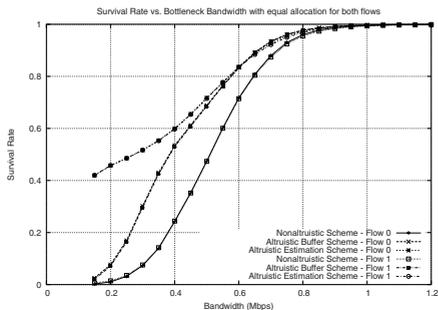


Fig. 6. Survival Rates for 2 Flow Scenario

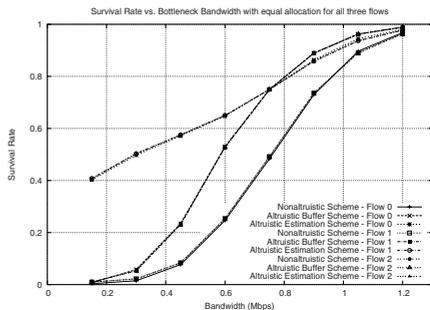


Fig. 7. Survival Rates for 3 Flow Scenario

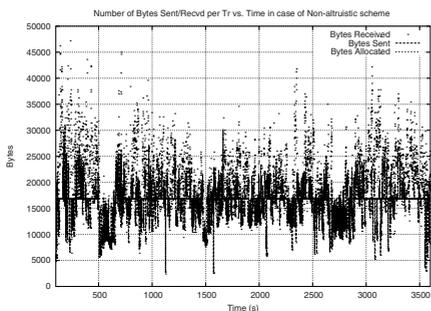


Fig. 8. Distribution for Non-Altruistic Scheme

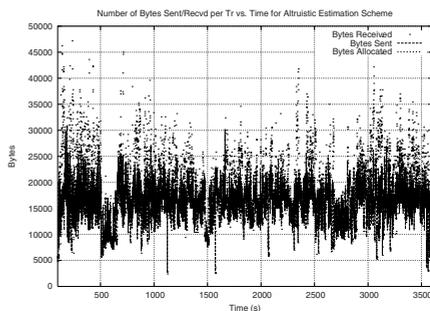
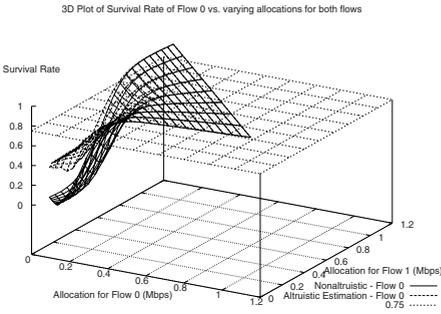
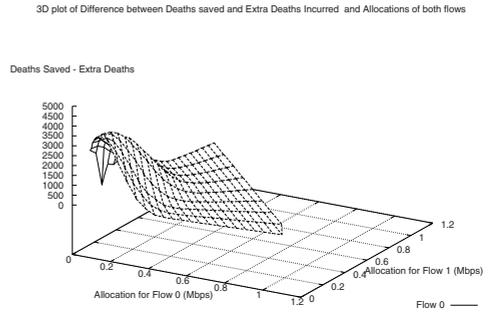


Fig. 9. Distribution for Altruistic Estimation Scheme

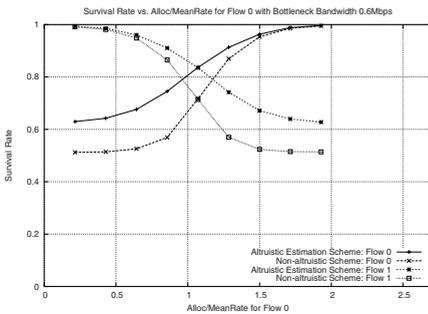
Fig. 11 depicts the net number of deaths saved for Flow 0 over a wide allocation spectrum for Flows 0 and 1 (varying bottleneck link bandwidth divided between the two flows in varying proportions). Once again it is indicated that the net improvement is greater when allocations are lower. Fig. 12 depicts the survival rates of both flows versus the ratio of allocation to mean rate of one flow, for bottleneck link bandwidth of 0.6 Mbps respectively. This bandwidth corresponds to the situation in which the total available capacity is just enough to allow traffic at the mean rates of both flows to pass through. The intent is to determine the degree of incentive a flow might have to go in for a much lower or higher allocation than its mean rate. At a bottleneck capacity of 0.6 Mbps it is seen that, at an allocation lower than the mean rate, there is significant increase in survival rate due to altruism. At allocations much higher than the mean rate, the performance gain gradually diminishes. At the mean rate, we find that there is significant improvement in performance and the new survival rate lies in a very desirable range of above 80%. So it seems that there is sufficient motivation for everyone to seek allocations close to the mean rate, as then the post-altruism survival rates are fairly reasonable. Fig. 13 depicts the percentage of flow deaths saved on using the estimation-based altruistic scheme



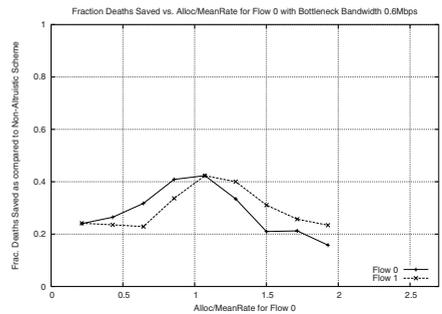
**Fig. 10.** Variation in Survival Rate with Allocations



**Fig. 11.** Variation in Net Deaths Saved with Allocations



**Fig. 12.** Survival Rate vs. Allocation/Mean for Flow 0(0.6 Mbps Link)



**Fig. 13.** Fraction of Deaths Saved vs. Allocation/Mean for Flow 0(0.6 Mbps Link)

for bottleneck link bandwidth of 0.6 Mbps. It may be seen that the maximum percentage is obtained close to the mean rate allocation. Fig. 14 shows survival rates for a situation in which one flow cheats i.e. though it has a true  $PLR_S$  of 0.2 (on which actual survival depends), it advertises a different  $PLR_S$  (on which help decisions are made). The figure clearly depicts that the cheating flow gains no advantage in terms of survival rate by advertising a higher or lower  $PLR_S$ . However, if one looks at the average PLR obtained by this flow in time intervals where it survived (Fig. 15), one finds that its PLR significantly reduces at lower advertised death thresholds. Thus the flow can gain in terms of better end-quality during periods of survival. It is therefore advisable to employ game strategies in real-world situations where complete trust may not be assumed.

Fig. 16 depicts survival rates versus allocation for one flow, in the case when the two flows correspond to two different MPEG-4 traces viz., *Jurassic Park* and *The Firm*. We obtain trends similar to Fig. 12 which show that the results obtained hold even for interacting flows with different traffic characteristics. Fig. 17 depicts survival rates versus allocation for one flow, keeping the allocation

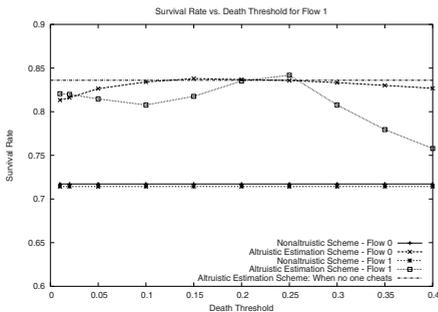


Fig. 14. Survival Rates with one cheating flow

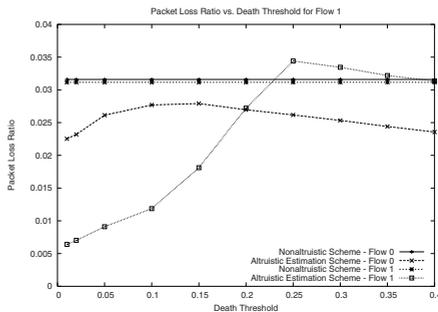


Fig. 15. Avg. PLR over all Periods of Survival with one cheating flow

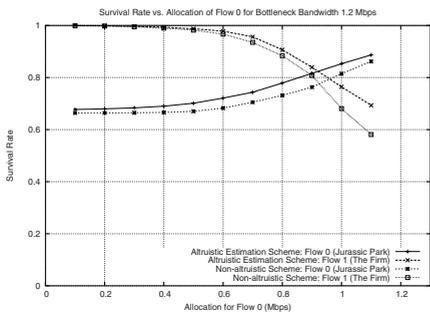


Fig. 16. Survival Rate vs. Allocation/Mean for Flow 0 (1.2 Mbps Link)

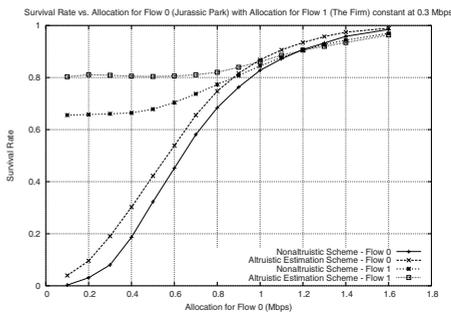


Fig. 17. Survival Rate vs. Allocation for Flow 0 (0.3 Mbps to Flow 1)

for the other flow constant. The flow with a constant allocation of 0.3 Mbps corresponds to the MPEG-4 trace for *The Firm* while the other corresponds to the MPEG-4 trace for *Jurassic Park*. We once again find that though the flow for *Jurassic Park* does obtain some performance improvement due to altruism at allocations lower than its mean rate, yet they are not significant enough to be an incentive for deliberate under-allocation. Besides, we find that the survival statistics of the other flow (which has gone in for a fair allocation) do not deteriorate due to altruism. Rather, it also sees a performance improvement.

## 7 Conclusions

This paper presents *Reciprocal Altruism* as a paradigm for providing QoS to multimedia flows. It also introduces a notion of per-GOP survival. A preliminary investigation into the feasibility of the same has been undertaken and a simplistic centralized mechanism has been simulated. The results obtained so far are indication of the potential of this paradigm. They serve to validate the concept of *Reciprocal Altruism* as a QoS Paradigm. The simplistic estimation-based scheme

described here vindicates the possibility of coming up with lightweight on-the-fly estimation strategies that have no foreknowledge of traffic characteristics. However, there is need to come up with a comprehensive architecture of a service discipline for centralized imposition of reciprocal altruism. Such a discipline would need to handle multi-flow interactions and allow for various game strategies (e.g. TFT, Pavlov etc.) instead of assuming complete trust. One would also need to look at efficiency issues in the implementation, as well as the interplay of multimedia flows with best-effort traffic. The possibility of an agent based approach also merits investigation. Such an approach would be particularly useful to large organizations which could then define their own custom policy, based on economic concerns. Another major issue is that of looking at multiple points of contention along flow-paths in an integrated manner, and basing decisions thereof. A comprehensive investigation into these issues needs to be undertaken.

## References

1. Braden, R., Clark, D., Shenker, S.: Integrated services in the internet architecture: an overview. Internet RFC 1633 (1994)
2. Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W.: An architecture for differentiated services. Internet RFC 2475 (1998)
3. Liao, W., Li, V.O.K.: The split and merge protocol for interactive video-on-demand. *IEEE Multimedia Magazine* **Oct.-Dec.** (1997) 51–62
4. Verscheure, O., Frossard, P., Hamdi, M.: User-oriented qos analysis in mpeg-2 video delivery. *Journal of Real-Time Imaging (special issue on Real-Time Digital Video over Multimedia Networks)* **5** (1999) 305–314
5. Feamster, N., Balakrishnan, H.: Packet loss recovery for streaming video. In: *Proceedings of 12th International Packet Video Workshop*. (2002)
6. Verscheure, O., Frossard, P., Boudec, J.Y.L.: Joint smoothing and source rate selection for guaranteed service networks. In: *Proceedings of IEEE Infocom (2)*. (2001) 613–620
7. Shenker, S.: Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking* **3** (1995) 819–831
8. Park, K., Sitharam, M., Chen, S.: Quality of service provision in noncooperative networks with diverse user requirements. *Decision Support Systems* **28** (2000) 101–122
9. Garg, R., Kamra, A., Khurana, V.: A game-theoretic approach towards congestion control in communication networks. *ACM SIGCOMM Computer Communications Review* **32** (2002)
10. Trivers, R.: The evolution of reciprocal altruism. *Quarterly Review of Biology* **46** (1972) 35–57
11. Wilkinson, G.S.: Food sharing in vampire bats. *Scientific American* **262** (1990) 76–82
12. Axelrod, R., Hamilton, W.D.: The evolution of cooperation. *Science* **211** (1981) 1390–1396
13. Wu, J., Hassan, M.: The issue of useless packet transmission for multimedia over the internet. *Computer Communications* (2003)
14. Fitzek, F.H.P., Reisslein, M.: Mpeg-4 and h.263 video traces for network performance evaluation. *IEEE Network* **15** (2001) 40–54
15. NS: ns notes and documentation. (<http://www.isi.edu/nsnam>)