

Stabilizing Intelligent Route Control: Randomized Path Monitoring, Randomized Path Switching or History-aware Path Switching?

Alexandre Fonte^{1,2}, José Martins¹, Marília Curado¹, and Edmundo Monteiro¹

¹ University of Coimbra, CISUC/DEI, Coimbra, Portugal

² Polytechnic Institute of Castelo Branco, Castelo Branco, Portugal
{afonte,marilia,edmundo}@dei.uc.pt jemart@student.dei.uc.pt

Abstract. Multihoming Intelligent Route Control (IRC) plays a significant role in improving the performance of Internet accesses. However, in a competitive environment, IRC systems may introduce persistent route oscillations, causing significant performance degradation. In this study, we investigate three design alternatives to cope with this issue: Randomized Path Monitoring, Randomized Path Switching and History-aware Path Switching. The simulation results show that Randomized Path Monitoring is an effective alternative to Randomized Path switching when the sampling frequency is conservative. The results also indicate that the use of a sophisticated IRC algorithm, such as history-aware path switching, may not bring any benefit in terms of stability.

1 Introduction

Multihoming is a well-known technique to improve performance and reliability of Internet accesses. It consists on increasing the Internet connectivity, by leasing multiple broadband lines (e.g., Business DSL) from two or three ISPs (Internet Service Providers). The use of multihoming allows stub ASs (Autonomous Systems) to experience a potential performance improvement of at least 40% [1]. In this context, Intelligent Route Controllers (IRC) are, thus, being increasingly used by multi-homed stub ASs, as they provide a holistic way to solve their traffic challenges by shifting traffic between ISPs [2, 3].

Unfortunately, the use of IRC has a major weakness, that is, oscillations can take place due to factors such as its intrinsic selfish nature, self-load effects, and synchronization between the probes [4, 5]. Specifically, the last two situations were studied in [4]. As a solution, the authors proposed a set of IRC techniques that might reduce IRC oscillations by adding randomness in the path switching process. In this paper, we answer the question *if adding randomness into the sampling process or using path history to assist IRC decisions are effective alternatives to refrain IRC oscillations?*

The rest of the paper is structured as follows. Section 2 gives a brief overview of the IRC algorithms explored in this study. Next, Section 3 presents the results of the evaluation of these algorithms. Finally, the major conclusions of the work performed are summarized in Section 4.

2 A brief overview of IRC algorithms

Dynamic path switching is the key technique used by most available IRC systems to get better end-to-end performance [2–4, 6]. With this approach, the IRC selects, in every routing cycle, the next-hop ISP to forward packets that has the smallest value of the chosen performance metric. Obviously, IRCs need to previously probe all the candidate paths by sending, for instance, ICMP and TCP probes at a given frequency f_i . In addition, the path selection performed by IRCs must take into account end-to-end performance bounds, so that the perceived application quality does not suffer degradation. For instance, the ITU-Ts G.114 recommendation suggests an one-way delay bound of 150ms to maintain high quality voice.

The work described in this paper seeks to answer the question about which is the best approach to cope with the oscillations associated with IRCs. To achieve this objective, three classes of IRC algorithms are investigated, as follows:

DLV (Deterministic Last Value) - DLV is the basic IRC routing algorithm. The IRC selects as the best path, the path that has the smallest value of the metric $M(t_i)$, computed in the last time slot t_i . In other words, in DLV, the path shifts are deterministic.

FSP (Fixed Switching Probability) - FSP adds randomness to the path switching process. In this case, the IRC picks the best path with a given switching probability $P \in [0, 1]$.

LpEMA (Low pass Exponential Moving Average) - LpEMA introduces a certain degree of path history in the IRC route decisions. To compute the actual metric estimate e_i , the IRC combines the previous metric estimate e_{i-1} and the actual metric $M(t_i)$ using an adaptive Exponential Moving Average (see Eq.(1)) [7]. Then, the IRC picks the best path, the path that has the smallest metric estimate e_i .

$$\begin{cases} e_i = (1 - \alpha_i)e_{i-1} + \alpha_i M(t_i) \\ \alpha_i = \alpha_{max} \frac{1}{1 + \frac{|m_i|}{m_{norm}}} \end{cases} \quad (1)$$

where α_i is an adaptive exponential weight, which is calculated using the classical formula for low pass filter, m_i is the gradient between two metric samples (i.e., $\frac{M(t_i) - M(t_{i-1})}{t_i - t_{i-1}}$), and m_{norm} is the normative gradient calculated over a given time window (e.g., 10 times the interval $t_i - t_{i-1}$).

3 Evaluation of IRC algorithms

This section presents the performance evaluation of the DLV, FSP and LpEMA Intelligent Route Control algorithms. The evaluation was performed on a J-Sim simulation model for IRCs³. Two sampling processes were used to compare the three algorithms, namely, a periodic sampling process (Periodic sampling) and a pseudo-Poisson sampling process (p-Poisson sampling), with N_i samples uniformly distributed over a slot of time t , afterwards referred as the window t .

³ <http://www.j-sim.org/>

3.1 Simulation Setup and IRC parameterization

The network topology was built using the BRIT⁴ topology generator and it contains 100 ASs with a ratio of ASs to inter-domain links of 1:3. During the tests, 300 IRCs sources send homogeneous traffic aggregates to remote prefixes. Each traffic aggregate is composed of a fixed number of multiplexed Pareto flows (i.e., VoIP flows) with Poisson arrivals.

The RTT bound for DLV, FSP and LpEMA is set to 300ms. In this study, the results obtained with DLV are used as a reference for comparison with FSP and LpEMA. The performance of FSP is studied under two different values of the switching probability P (0.1 and 0.5). On the other hand, the LpEMA configuration relies on α_{max} , typically ranging from 0.5 to 5. In this study, we pick out a middle value for α_{max} (2.5).

To avoid the self-load effect [4], we combined latency and spare bandwidth into a single metric, i.e., $M = \alpha_1 \cdot latency_t + \alpha_2 \cdot \frac{1}{abw_t}$, where $latency_t$ is the median of the measured RTTs in a window t of 30s, and abw_t is the estimated spare bandwidth in the peering link during t . To facilitate the tuning of $\alpha_i, i = 1, 2$, we adopted the framework in [6]. Finally, DLV, FSP and LpEMA use a route threshold (R_{th}) set to 10, as common stability mechanism [6].

3.2 Comparison of IRC algorithms

Figures 1 to 6 illustrate two empirical Complementary Cumulative Distribution Functions (CCDF) of the number of path shifts performed by the three IRC schemes for two bin sizes, used to count the path shifts along the simulation. If the probability of a number of Path Shifts (PS), in a bin, is greater than or equal to x is high (i.e., $P(\text{PS in a bin} \geq x)$ is high), it means route oscillations are highly present in every bin. It must be pointed out that a line starting at a value smaller than 1, means that a fraction of the bins do not have path shifts.

Three main conclusions can be drawn from the analysis of these results. First, the comparison between the DLV and FSP shows that IRC oscillations can be reduced by adding a certain degree of randomness in the route control decisions. In fact, these results show similar trends of a previous study [4]. However, when IRCs use the RTT performance metric rather than the loss rate metric, DLV and FSP present a similar number of path shifts for path switching probabilities below 0.5. This behavior might be explained by a weaker aggressiveness of the RTT metric when compared to the loss rate metric used in [4].

Second, the results also show that the oscillations can be reduced by adding a certain degree of randomness in the sampling process. In fact, figures 3 and 4, illustrate that, for lower frequencies of sampling, DLV combined with p-Poisson sampling performs better than when combined with periodic sampling. In particular, when $f = 0.2Hz$, DLV combined with p-Poisson sampling has a smaller probability (20% smaller) to reach the same number of path shifts in a bin. However, when the sampling frequency is higher, both DLVs perform similarly.

⁴ <http://www.cs.bu.edu/brite/>




Fig. 1. CCDF of PS, DLV vs FSP (P=0.1), Periodic, $f = 0.2Hz$.




Fig. 2. CCDF of PS, DLV vs FSP (P=0.5), Periodic, $f = 0.2Hz$.




Fig. 3. CCDF of PS, DLV, Periodic, $f = 0.2Hz$ and $f = 1Hz$.




Fig. 4. CCDF of PS, DLV, p-Poisson, $f = 0.2Hz$ and $f = 1Hz$.

This arises from similar overlappings of the IRCs measurement windows, which are also significant for p-Poisson when the sampling frequency is high.

Third, the results in figures 5 and 6 indicate the use of a sophisticated IRC algorithm, such as LpEMA, may not bring any benefit in terms of the reduction of oscillations. LpEMA can be viewed as a cascaded DLV and a *low-pass filter* block, and so in theory it should reduce the number of path shifts. However, its ideal tuning depends on the particular stability pattern of the network [7].

Figures 7 to 12 illustrate the CCDF distribution of the traffic latency. The main conclusions of the analysis of these results are described next. First, the results show that there is a correlation between the frequency of sampling and the observed RTTs. More specifically, as the sampling frequency increases, the probability of the traffic to experiment higher RTT also increases. This is even more noticeable for RTT smaller than 70ms. However, in case of Periodic sampling, there is a clear turn-over point, and a long tail over 70ms, over which higher frequency leads to smaller probabilities to get higher RTTs. This characteristic is common to all IRC algorithms. Second, the results in figures 7 and 8 confirm that when DLV is combined with p-Poisson sampling, the traffic experiments a




Fig. 5. CCDF of PS, DLV vs LpEMA, Periodic, $f = 0.2Hz$.




Fig. 6. CCDF of PS, DLV vs LpEMA, p-Poisson, $f = 0.2Hz$.




Fig. 7. CCDF of RTTs, DLV, Periodic.




Fig. 8. CCDF of RTTs, DLV, p-Poisson.

potential performance benefit of 20%, while a smaller number of path shifts is needed. However, when the sampling frequency is high, the improvement is negligible. Third, when changing the FSP probability P (from $P = 0.5$ to $P = 0.1$) to reduce oscillations its effects over the RTT distribution are negligible.

4 Conclusions

In this paper, the evaluation of the three main classes of Intelligent Route Control algorithms was performed in order to find out which is the best alternative to cope with the oscillations associated with this type of schemes. The results showed that the addition of a randomness component to the route control process drastically reduces the number of path shifts needed to meet the traffic challenges. However, the addition of a randomness component to the path monitoring mechanism is an effective alternative to this solution. But, the decision to adopt a randomized path monitoring solution depends on the time scale used by the IRC system to probe all candidate paths. In particular, short time scales might hurt its effectiveness due to the overlapping of the IRC measurement win-




Fig. 9. CCDF of RTTs, FSP, $P = 0.1$, Periodic.




Fig. 10. CCDF of RTTs, FSP, $P = 0.5$, Periodic.




Fig. 11. CCDF of RTTs, LpEMA, $\alpha_{max} = 2.5$, Periodic.




Fig. 12. CCDF of RTTs, LpEMA, $\alpha_{max} = 2.5$, p-Poisson.

dows. Finally, the use of sophisticated IRC algorithms, such as history-aware path switching, is questionable, since these mechanisms require additional tuning, according to the particular stability pattern of the network.

References

1. A. Akella et al., A Measurement-Based Analysis of Multihoming. In the Proc. of ACM SIGCOMM 2003, Karlsruhe, Germany, 2003.
2. Cisco Systems, Inc., Optimized Edge Routing.
3. Internap Networks Inc., Flow Control Platform.
4. R.Gao et al., Avoiding Oscillations due to Intelligent Route Control Systems. In the Proc. of IEEE INFOCOM 2006, Barcelona, Spain, April 2006.
5. L. Qiu et al., On Selfish Routing in Internet-Like Environments. In Proc. of ACM SIGCOMM 2003, Karlsruhe, German, August 2003.
6. M. Yannuzzi et al., On the Advantages of Cooperative and Social Smart Route. In Proc. of ICCCN 2006, Arlington, Virginia, USA, October 2006.
7. L. Burgstahler and M. Neubauer, New Modications of the Exponential Moving Average Algorithm for Bandwidth Estimation, In Proc. of the 15th ITC Specialist Seminar, July 2002.