STEINWURF

Rely: The Next Generation FEC

The most flexible, dynamic, scalable, and efficient FEC software solution to propel your networks and device performance into the future

About Steinwurf and Rely

Steinwurf is the creator of Rely, the next generation FEC solution that maximises data transfer efficiency, whilst providing optimal latency, bandwidth, and reliability performance.

By utilizing the power of the patented Random Linear Network Coding (RLNC) technology, Rely has redefined the trade-off paradigm between latency, reliability, and the cost of delivering this performance, allowing clients to improve QoS, whilst driving down operational costs, bandwidth usage and carbon footprints.

Good for service providers, good for customers and good for the planet, Rely is designed to deliver industry leading reliability and ultra-low latency communications over dynamic, mobile, and wireless networks.

About this Report

This report demonstrates how Steinwurf's Rely outperforms ARQ, especially in terms of latency, and why performance and sustainability conscious SD-WAN providers and their customers would benefit from implementing Rely today.

There are many SD-WAN and VPN service providers who have either forgone the inclusion of FEC in their product, not even considered the importance of FEC to their product and customers, or realise that FEC is necessary but still depend on ARQ while they search for the right FEC to integrate in their systems. The benchmarks in this report demonstrate that:

- Rely and ARQ are both good choices when reliability is required
- For modern applications which require interaction and therefore low latency operation, Rely is able to provide reliability at a very low latency.
- Rely is able to provide this low latency and reliable connectivity at a relatively low cost of repair.

Not all FECs are Created Equal

The power of Rely is rooted in Steinwurf's patented RLNC technology, which makes it more flexible and significantly more powerful than any legacy FEC code. The RLNC technology has been actively developed across over 12 universities and research institutes, including MIT and Caltech.

Rely allows enterprises to fully optimise all key parameters and deliver best-in-industry applications that are good, fast, and cheap, unlike legacy FECs - such as Reed Solomon codes developed in the 1950s – which always require a trade-off between 1 or 2 of the key parameters below, or ARQ, which is often bandwidth optimal, but this efficiency of one parameter can come at a serious cost to latency.

A positive externality of the increased efficiency of using Steinwurf's FEC is improved utilization of networks and network devices, meaning you can serve more customers, faster, using less bandwidth and with less load on equipment, translating to a greener and cheaper solution.

FEC KPIs:

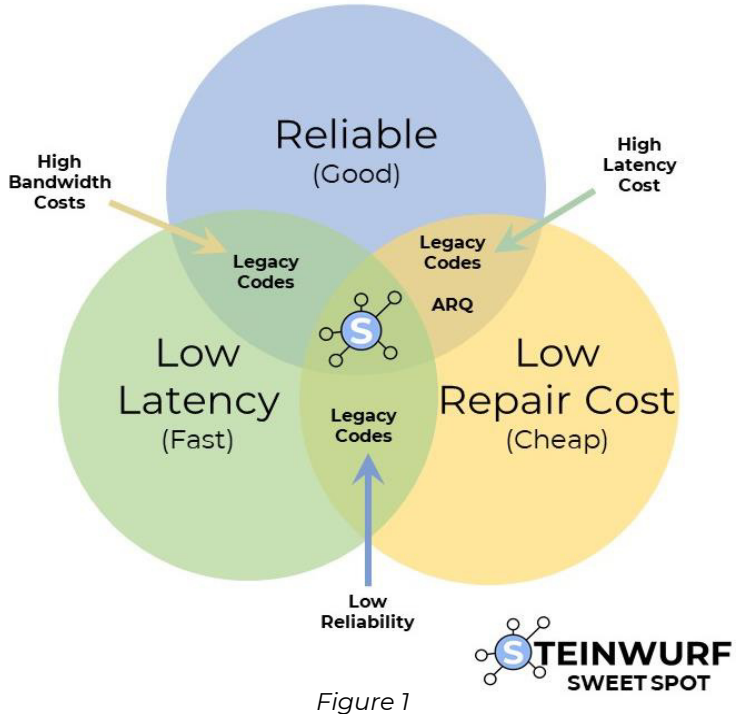

Figure 1

- **Latency**: the end-to-end delay added by the error correction mechanism.
- **Reliability**: the ability to seamlessly recover lost packets and minimize the probability of unrecovered packet loss and the resulting impairments. Good reliability leaves a network with little to no residual packet loss.
- **Repair Cost**: the amount of extra overhead/bandwidth consumed for the error correction mechanism to be able to recover and repair for lost packets.

Historic methods for ensuring reliability

To better understand and put into context Rely's performance versus a network utilising no FEC and falling back on ARQ to achieve reliability, it's useful to understand the differences in how each network operates. Historically, retransmissions were used to ensure packets arrive at the destination.

A key parameter to understand is the network loss rate – the average amount of packet loss which a network is susceptible to at the relevant time. Real world networks rarely experience fixed amount of network loss, more often it will change over time depending on various environmental and network factors.

Other key parameters are the one-way link latency (latency between the source and destination), whether the application being served has a latency budget or not, and the round trip time (RTT) which is twice the link latency.

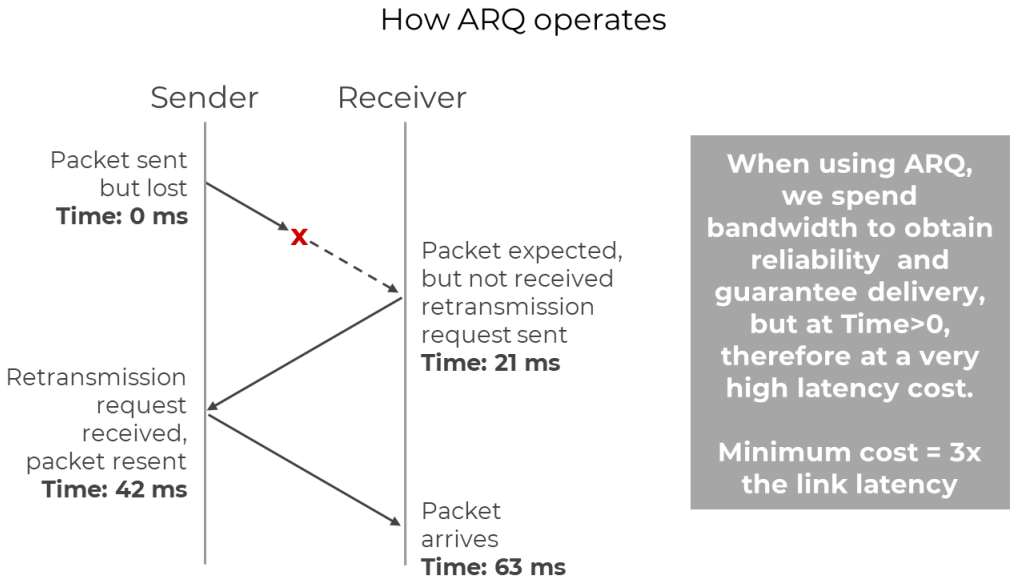

Figure 2

Figure 2 shows a system with a one-way latency of 21ms. ARQ is not pre-emptive in nature but reactive. When packets are believed to be lost, the target application will automatically send a request back to the sender to retransmit a packet and keep doing so until all missing packets in a sequence are received. This has the advantage of being very bandwidth efficient compared to FEC in general as retransmissions should only occur when packets are believed to be lost, Whereas repair traffic may be configured to be constantly generated, even when no packet loss may occur over a time period.

The new state of the art for reliability

A system using Rely can be configured and adapt to utilising any repair rate, and this can be changed over time to adapt to any variance in the network loss rate, whether network losses drop to 0% or increase to 5% or more, Rely can adapt to ensure maximum efficiency. This repair rate is repair traffic sent by Rely in advance of packet loss which can be used to repair transmissions where data packets are lost. The key is to always have the same amount or slightly more repair traffic than network losses to be able to fix any reliability issues caused by packet loss without the need for delay generating retransmissions.

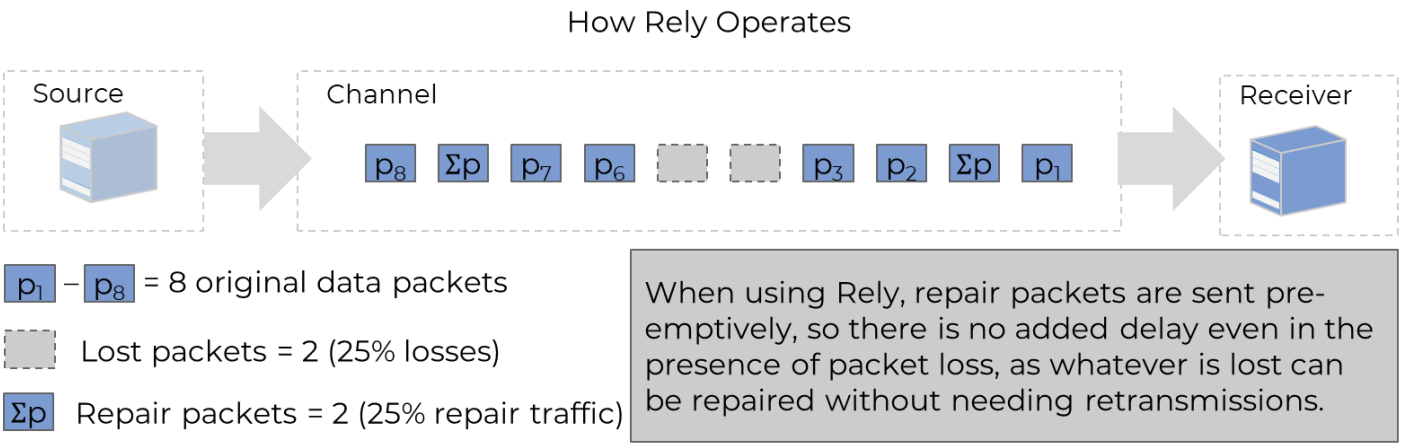

Figure 3

While a major downside of using ARQ is that when packet losses occur the application must wait for at least 3x the link latency to receive a packet which has been lost, Rely operates without delay, making it much more suitable for modern day applications such as streaming live voice and video for conference calls and VoIP, or any other live or interactive application which cannot tolerate delay without sacrificing the user experience.

In the next section we dive into some metrics by looking at some benchmarks for two scenarios, one for an application with no latency budget, and one for an application with a strict latency budget, where packets received after a certain amount of time are no longer useful and essentially lost.

Rely provides great reliability without added latency

Since both ARQ and Rely provide great reliability, there's no need to deep dive much in to figure 4 below

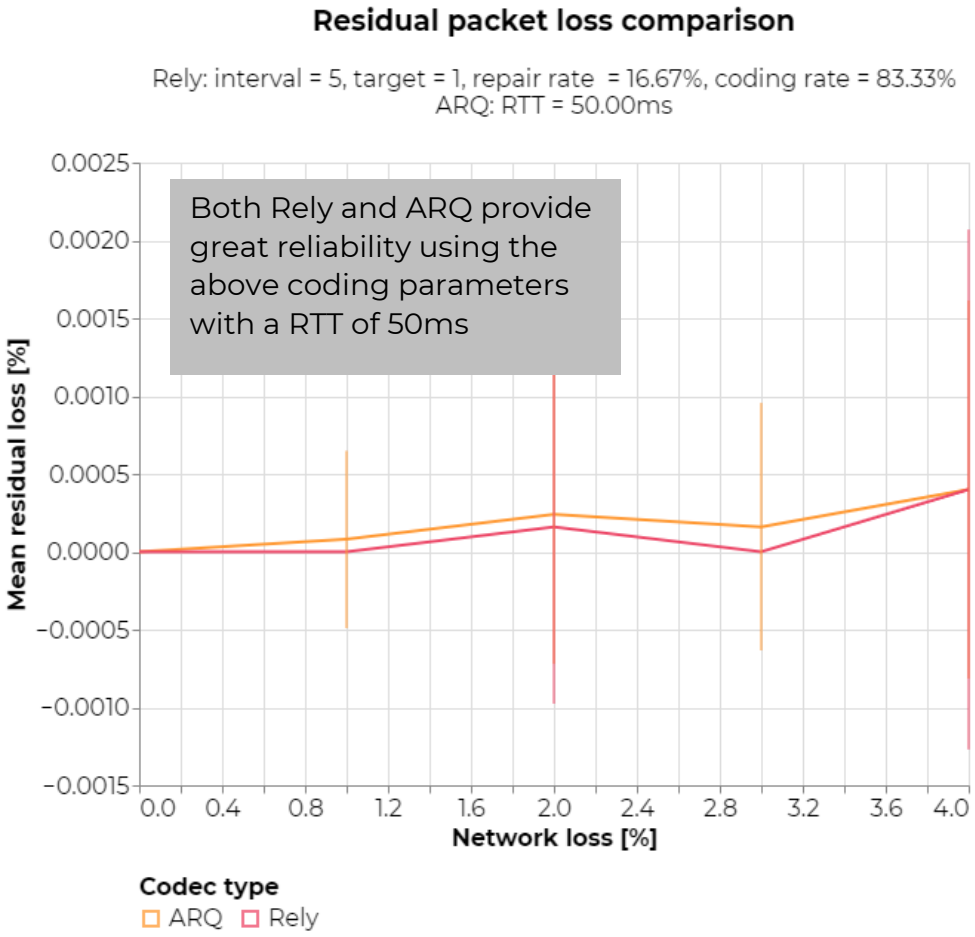

Figure 4

So for applications where time is not a factor and a user can wait to receive information, both Rely and ARQ can both make sure packets reach the destination. However, even for applications without a strict latency budget, that's not the whole story.

The mean latency figures for Rely using 16.67% repair and ARQ with a RTT of 100ms are shown in figure 5

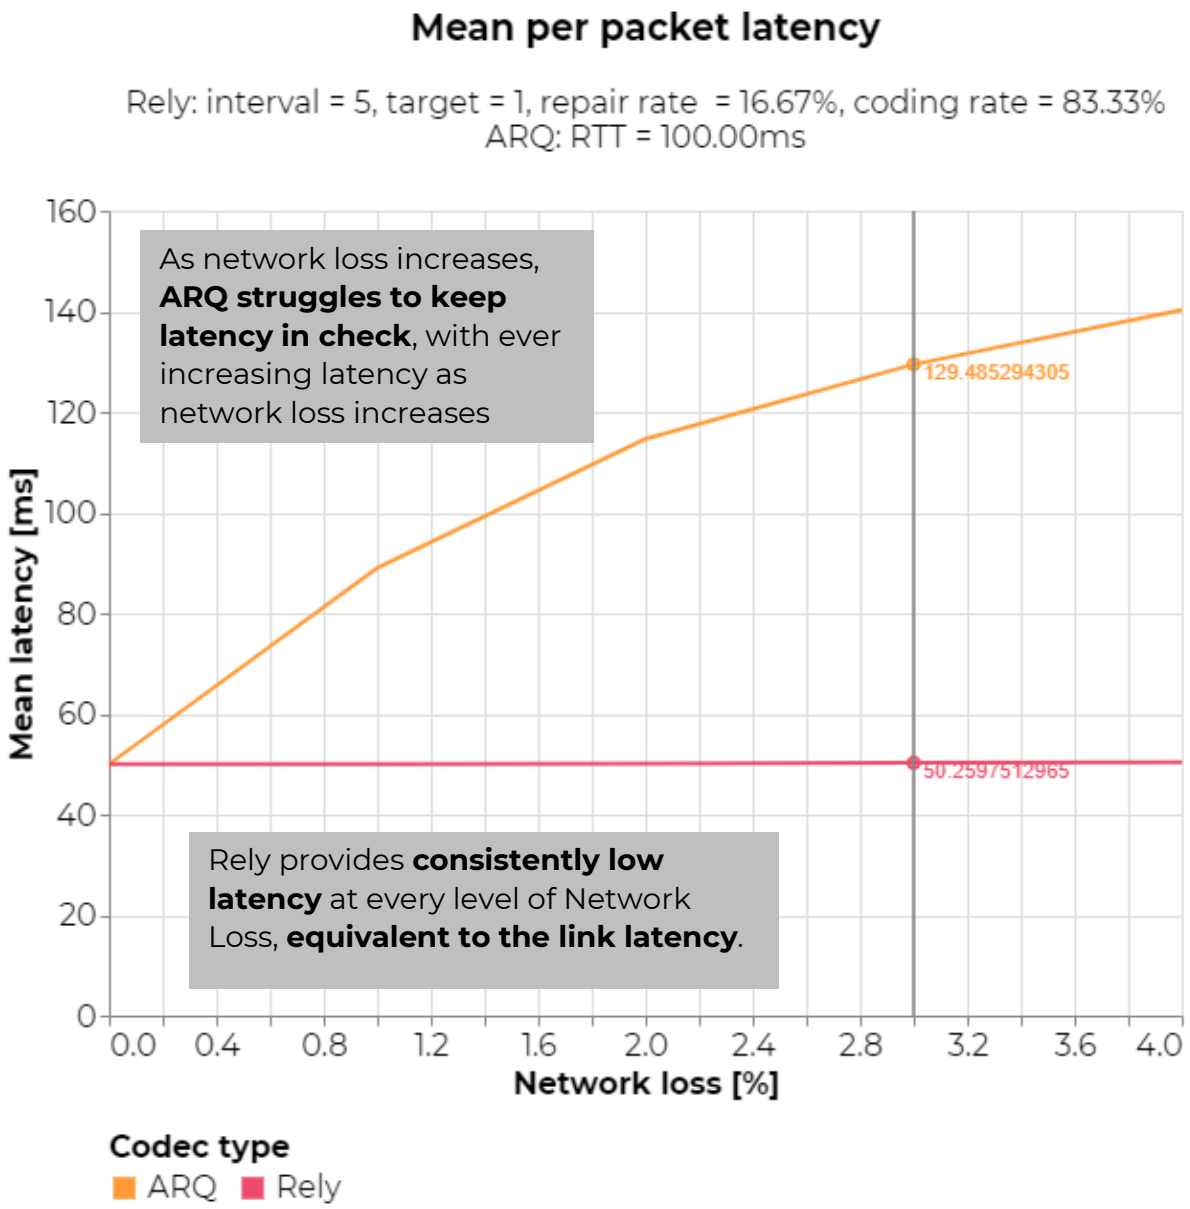

Figure 5

Figure 5 reveals the latency disparity between using ARQ, and next gen erasure coding. ARQ was designed and implemented before a time when interactive applications were both ubiquitous and central to the successful operation of so many cloud services such as audio/video conference; VoIP; SD-WAN; VPN; gaming; remote support and remote operation of machinery, and so many more applications.

Figure 5 only shows mean latency figures, but looking at the per packet latency data for the 100,000 simulations conducted at each level of network loss, reveals more disparity.

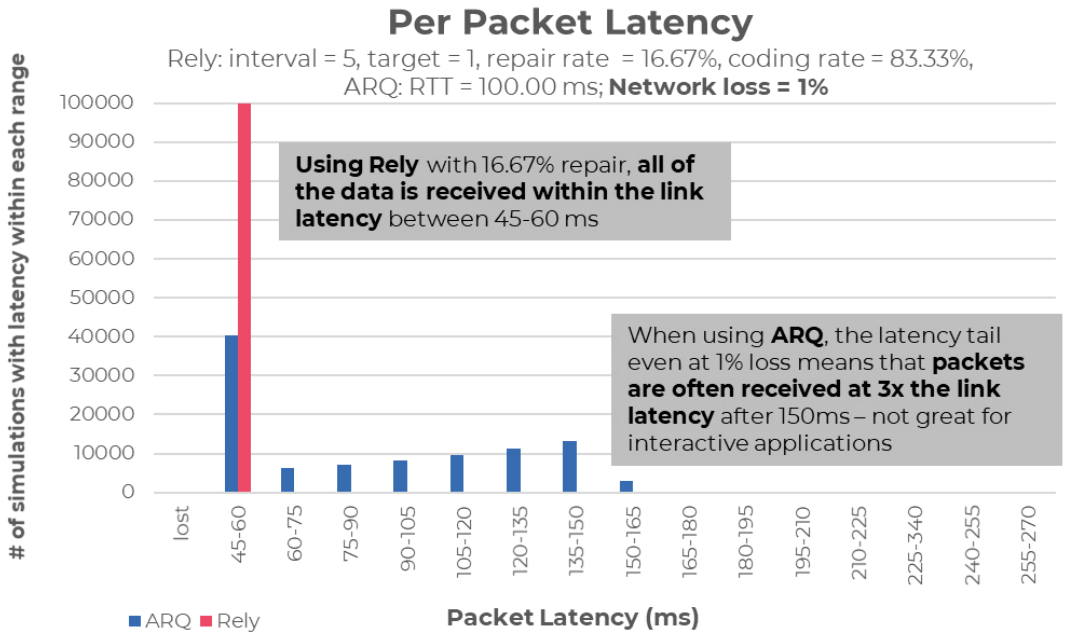

Figure 6

Figures 6 & 7 show how dependable Rely is at delivering packets reliably within the link latency, whereas ARQ introduces increasingly larger latency to cope at higher network loss. This would completely break the user experience for real time interactive applications, so any service provider still relying on ARQ to deliver such services will be at a serious disadvantage

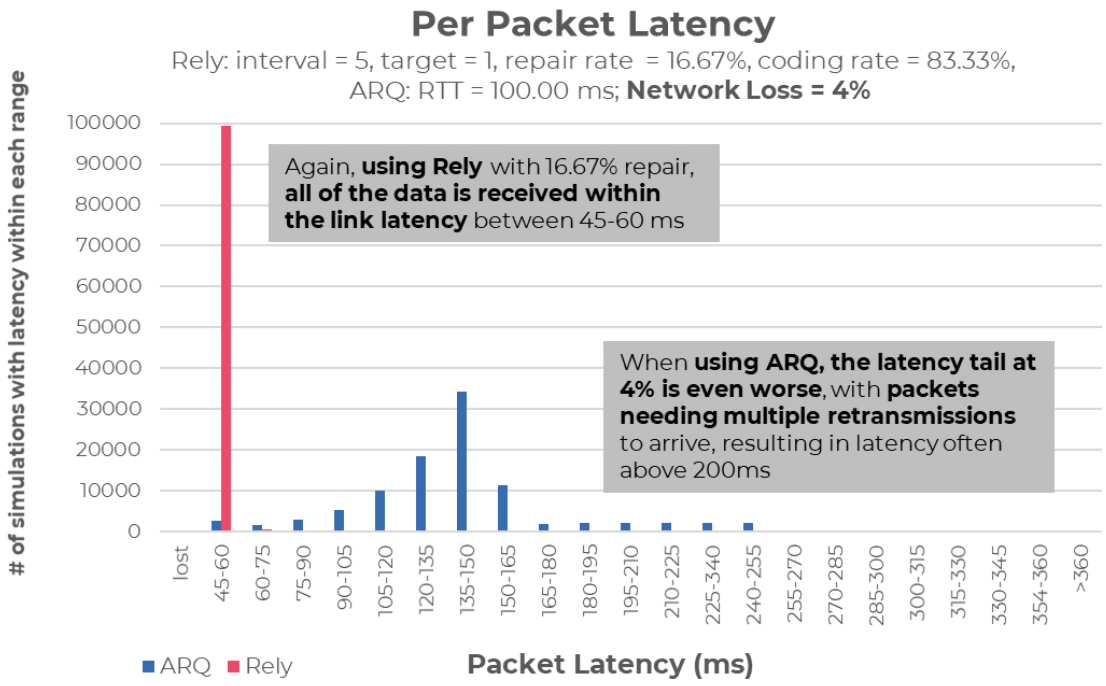

Figure 7

Rely provides great reliability without added latency, and can be used to guarantee performance within certain latency bounds.

This section looks at performance of Rely and ARQ when delivering a latency sensitive service such as live video streaming. Rely has the unique ability of being able to add more repair for more important traffic in a transmission and utilise less repair for less important traffic. We call this content aware coding, and for video streaming it is applied by protecting the relatively more important i-frames in a video stream with a higher repair rate, as they are more important for the video QoE.

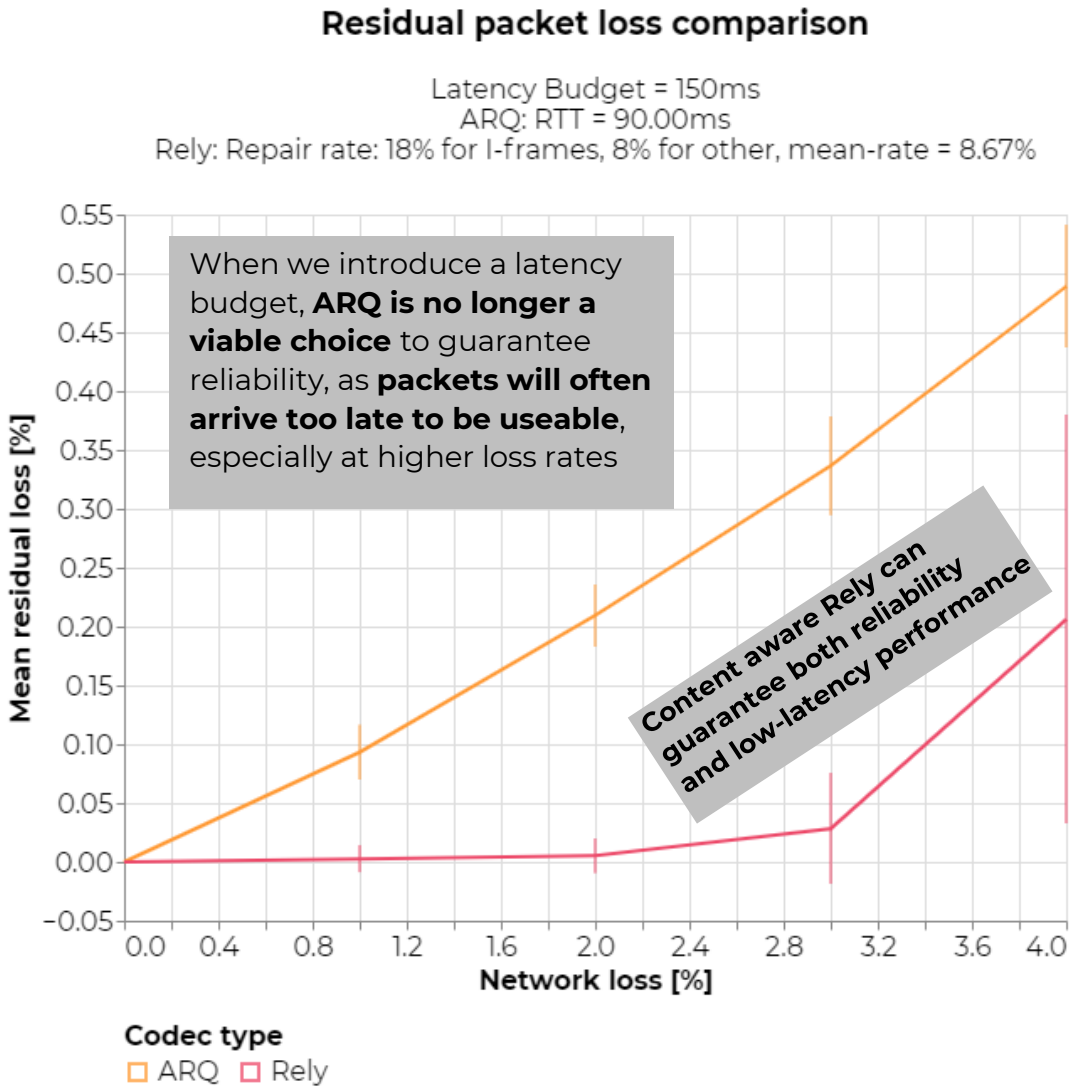

Figure 8

Figure 8 shows that when a latency budget is introduced, ARQ is no longer fit for purpose, as when packets need 3x the link latency or more to arrive, they are no longer useful, and the QoE for the service breaks down. In fact Rely is also able to dynamically adjust its repair rate to increase when network loss rates increase, so the curve for rely can be flattened further very easily, but that is beyond the scope of these simulations.

Looking at the mean latency figures in figure 9 for the live video streaming scenario with a latency budget again shows a similar story to before, with ARQ adding increasing amounts of latency to ensure delivery, and Rely ensuring a consistent latency target can be met

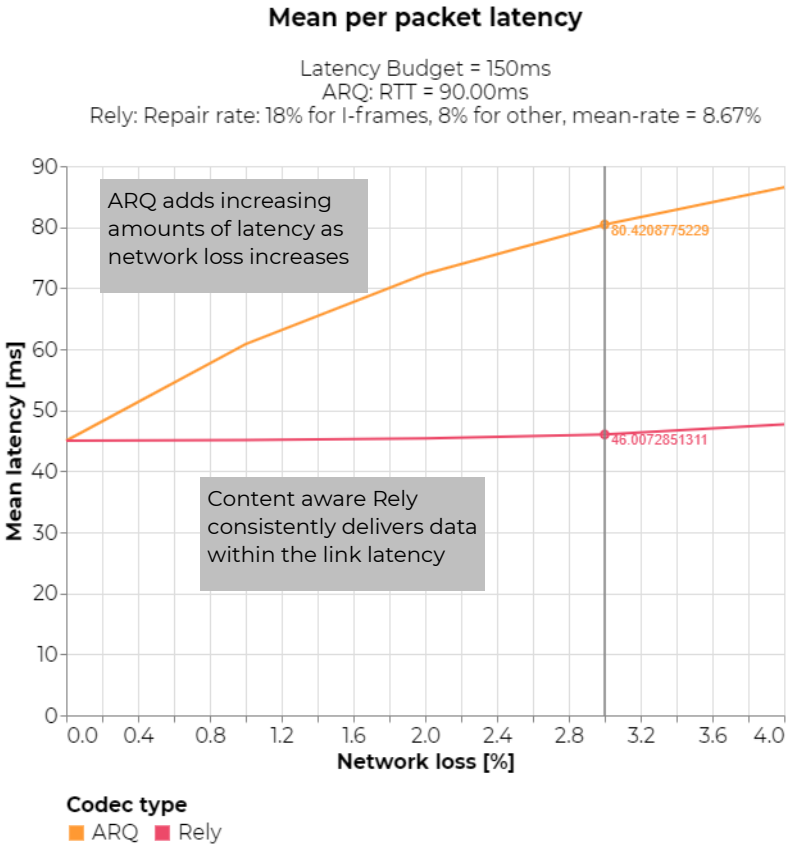

Figure 9

Figure 10 shows the full latency picture from 250,000 simulations run at 1% packet loss comparing content aware Rely with ARQ.

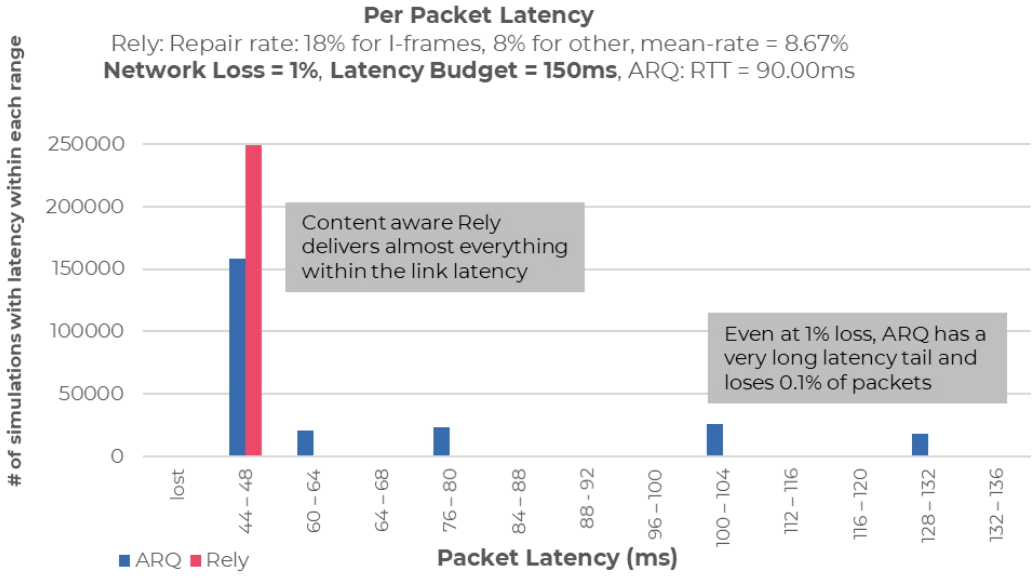

Figure 10

Figure 11 covers the same per packet latency view, but for 4% network loss.

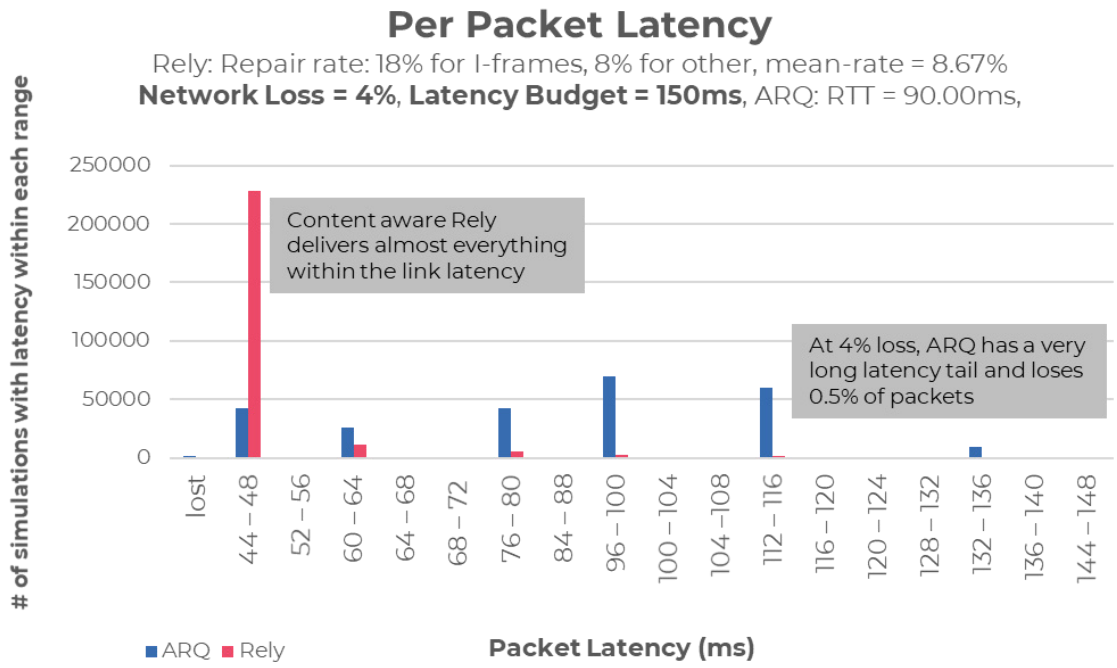

Figure 11

Figures 10 and 11 show that at 1% and 4% network loss, content aware Rely, consistently performs within a narrow latency window, making **Rely the only FEC choice for low latency applications**. When **using Rely, SD-WAN providers can provide performance guarantees** within a certain latency window, e.g., in this case, that all packets will arrive within the 150ms mark.

Interestingly, since content aware Rely prioritises delivery of video frames more important to the QoE (i-frames), even if Rely is losing a few packets in these simulations, lost frames will have little to no impact on the overall QoE. Whereas packets lost when using ARQ may often be very important to the QoE, compounding the negative effect on the user experience from the added latency when using ARQ.

At 1% loss, ARQ uses about 1% extra bandwidth to retransmit the 1% of packets lost and Rely is using 8.67% repair traffic (cost of Repair). However, with a RTT of 90ms, ARQ adds 34% extra latency – not a great trade off. In relative terms Rely is able to provide reliability and ultra low latency communication at a very low cost.

Modern customers using video/audio conferencing services, VoIP, accessing financial trading data, collaborating remotely or using any other application requiring ultra-low latency for a great QoE will look to their SD-WAN providers to make sure they have incorporated the latest technology and have best performing system. Service providers who have not yet upgraded their systems from using ARQ based reliability mechanisms are liable to lose clients who expect a better QoE for their real time interactive applications. **Steinwurf's FEC solutions are a core component** for SD-WAN providers who want to deliver on this performance promise.

Read More about Steinwurf

23/04/2021

Reduce Operating Costs When Protecting Audio Streams with Rely FEC

RLNC based FEC codes from Steinwurf can help level up performance while using less overhead and translating to lower operating costs for better quality of experience and a more reliable service.

Read More

19/03/2021

Steinwurf's next-gen FECs aren't a choice for SD-WAN, they're an imperative

SD-WAN companies need to simultaneously minimize latency and maximize reliability – as a consequence, continued use of ARQ or traditional block ECC / FEC solutions won't deliver the performance required.

Read More

27/02/2021

Overcoming packet loss – the bane of online gamers

The only viable option is to have a system that can operate flawlessly even in the presence of packet loss, and mask it from the user such that they don't experience any drop outs, lag, delay, lost inputs and connection instability. Steinwurf's next generation FEC is a must have component for a seamless gaming experience.

Read More

Get in touch to try Steinwurf's high performance Next Generation FEC codes in your SD-WAN solution for free!

Reach out today for expert guidance and details on our free evaluation and flexible pricing to suit your business model.

CONTACT

Niels Jernes Vej 10
9220 Aalborg East
Denmark

contact@steinwurf.com

+44 7932 073 224

www.steinwurf.com

Follow us