# A Hybrid Approach to Mitigating Jamming

Jamming attacks against the lightning network can be characterized as:

1. Fast jamming: a continuous flow of fast-resolving, failing payments that occupy a target node's resources. These payments are difficult to identify, as they mimic honest payments.
2. Slow Jamming: payments that are held for the maximum allowable time (before causing a force close) to occupy a target node's resources. These payments are more easily identifiable as an attack, though honest nodes will occasionally suffer slow-resolving payments as well.

Both types of jamming attacks - and combinations of the two - could be addressed by an idealized solution that charges for the time a HTLC is held, as targeted nodes are fairly compensated for their resources. This mechanism is unavailable in trustless, decentralized networks such as lightning [1], so a hybrid approach to target each type of attack is suggested.

Fast jamming - Unconditional Fees:

- Attaching a small fee to failed payments is trivial to an end user, but costly to an attacker sending continuous streams of payments.
- Since fast jamming attacks rely on a large number of rapidly resolving payments, a failure fee of 1-2% of success case fees can compensate nodes targeted by attacks[2].
- These payments cannot appropriately compensate for slow jamming attacks, as per-payment pricing that prohibits slow jams would affect usability for honest users.

Local Reputation to address slow jamming:

- If allowed full access to liquidity and slots, an attacker can instantly trivially slow jam a lightning channel by filling it with seemingly harmless HTLCs then holding them until their expiry.
- A reputation scheme based on forwarding fees that rewards honest actions (fast resolving, successful HTLCs) and punishes undesirable actions (slow resolving HTLCs) can be used to limit access to resources to nodes that demonstrate honest behavior over time, and revoke access from bad behaving peers.

Local reputation tracking can ensure that nodes are only granted full access to resources (and thus full ability to orchestrate a jamming attack) once they have built up a history of good behavior, and revoke access if they begin to misbehave. However, any reputation metric can inherently be gamed to fall *just below* the threshold that is defined as "good behavior", so upfront fees are proposed in combination with local reputation to fill this gap.

In [3] we outline a proposal for local reputation to mitigate slow jamming attacks. A sketch of a simplified upfront fees scheme is available in [4], although specification is ongoing at the time of writing.