Price of Censorship in Multiple Concurrent Proposers under Sybils

Abhimanyu Nag · 3 Sep 2026

https://functor.network/user/3197/entry/1962

Special Thanks to Aditya Saraf from Cornell Tech for initial conversations that led to this post. Additional thanks to Manvir Schneider, Cardano Foundation, Sarisht Wadhwa, UMichigan and Orfeas Thyfronitis Litos, Common Prefix for their reviews and points.

A nice new paper called Price of Censorship: Censorship Resistance and Throughput under Rational Concurrent Proposers caught my attention a couple of weeks ago and I have been thinking about it ever since. Saraf et al. derived the cost to effectively censor transactions in Multiple Proposer Protocols and bounds under different transaction fee mechanisms (TFMs). Given my current interest in Sybils, I thought that it would be a good idea to look into how the bounds change in the existence of Sybil nodes and how it changes the conclusion of the paper.

First, Some Background

Taking things literally from the paper, assume that there are n concurrent proposers and a total block capacity of B. Each proposer constructs a subblock containing at most

$$\ell := \frac{B}{n}$$

transactions, where the paper assumes that n divides B (I think the notation was something like $n|B$).

The shared mempool contains m transactions with tips

$$t_1 \geq t_2 \geq \cdots \geq t_m.$$

The final block is the union of the n proposed subblocks. So we see that a transaction may appear in several proposals but it is executed only once in the final block.

For each transaction i, let

$$p_i \in [0,1]$$

denote the probability that an individual proposer includes it. In the symmetric equilibrium studied by the paper, every proposer uses the same marginal inclusion probabilities

$$\mathbf{p} = (p_1, \ldots, p_m),$$

subject to

$$\sum_{i=1}^{m} p_i = \ell.$$

Before we move further, the paper model assumes that proposers observe the same mempool, act simultaneously, randomize independently across proposers, maximize their individual expected revenue, get outside user tips not in their control and most importantly, do not collude.

Transaction Fee Mechanisms (TFMs)

We will look at the different TFMs that the censorship cost bounds were constructed under. Let k_i (would have been nicer to denote as p_i but I guess its in use) be the number of proposers that include transaction i.

Duplication-penalizing TFM

If exactly one proposer includes transaction i, that proposer receives the entire tip t_i. If more than one proposer includes it, no proposer receives the tip and the transaction pays no auction fee.

The aggregate proposer revenue from transaction i is

$$R_i^{\mathrm{DP}}(k_i) = \begin{cases} t_i, & k_i = 1, \\ 0, & k_i = 0 \text{ or } k_i \geq 2. \end{cases}$$

Even-split TFM

If $k_i \geq 1$ proposers include transaction i, each includer receives

$$\frac{t_i}{k_i}.$$

Therefore, aggregate proposer revenue is

$$R_i^{\mathrm{ES}}(k_i) = \begin{cases} t_i, & k_i \geq 1, \\ 0, & k_i = 0. \end{cases}$$

and finally

Collective TFM

The tips of all included transactions are pooled and divided equally among the n proposers, irrespective of which proposer included which transaction.

Aggregate revenue from transaction i is again

$$R_i^{\mathrm{COL}}(k_i) = \begin{cases} t_i, & k_i \geq 1, \\ 0, & k_i = 0. \end{cases}$$

SIDEWAYS: We notice that the collective and even-split mechanisms have the same aggregate revenue from a fixed union of transactions, although they vary in distribution of that revenue among proposers. I suspect that the welfare of proposers in the two TFMs are also different, with one being more unfair than the other.

SIDEWAYS: Also the mechanism may charge a base fee β which is burned rather than being paid to proposers.

Core Results

Recall initially, n proposers, each selecting ℓ transactions, with total capacity $B = n\ell$. Now if each proposer includes transaction i with probability p_i, its final inclusion probability is

$$1 - (1 - p_i)^n.$$

Normalized throughput is therefore

$$\mathrm{Th}(\mathbf{p}) = \frac{\sum_i \left[1 - (1 - p_i)^n\right]}{\min\{B, m\}}.$$

Duplication lowers throughput because repeated inclusions consume capacity without adding new transactions. I suspect this would benefit from further analysis that falls under the umbrella of spam transactions and the many works in that area.

Symmetric proposer equilibrium

The expected marginal reward from including transaction i is

$$u_i^{\mathrm{DP}}(p_i) = t_i(1-p_i)^{n-1},$$

$$u_i^{\mathrm{ES}}(p_i) = t_i\frac{1-(1-p_i)^n}{np_i},$$

and

$$u_i^{\mathrm{COL}}(p_i) = \frac{t_i}{n}(1-p_i)^{n-1}.$$

Because these rewards decrease with p_i, equilibrium equalizes marginal rewards. The paper proves that there exists a unique symmetric marginal probability vector

$$p_i^* = \widehat{f}_i(\lambda^*), \qquad \sum_i p_i^* = \ell,$$

where λ^* is the common equilibrium marginal reward and $\widehat{f}_i$ is the clipped inverse of u_i.

SIDEWAYS: *As Orfeas points out, the notation implies that the proposer will choose ℓ transactions on average but that will often not be the case.*

For duplication-penalizing and collective,

$$f_i^{\mathrm{DP}}(\lambda) = 1 - \left(\frac{\lambda}{t_i}\right)^{1/(n-1)},$$

$$f_i^{\mathrm{COL}}(\lambda) = 1 - \left(\frac{n\lambda}{t_i}\right)^{1/(n-1)}.$$

The even-split inverse is computed numerically. The equilibrium itself can be found by binary search over λ^*.

Censorship Cost Theorem

Suppose an attacker wants transaction c censored with probability at least s. Independence implies

$$(1-p_c)^n = s,$$

so the target equilibrium inclusion probability must be

$$p_c^* = 1 - s^{1/n}.$$

Remove c and solve the equilibrium for the remaining transactions using residual capacity $\ell - p_c^*$. Let λ_{-c} denote the resulting marginal reward.

The paper's cut-and-paste theorem gives the minimum unconditional bribe per omitting proposer:

$$\boxed{b_c^* = [u_c(p_c^*) - \lambda_{-c}]_+ .}$$

Since each proposer omits c with probability $s^{1/n}$, expected total expenditure is

$$\boxed{C_c(s) = n s^{1/n} b_c^*.}$$

For the three TFMs, this becomes

$$b_{c,\mathrm{DP}}^* = \left[t_c s^{(n-1)/n} - \lambda_{-c}^{\mathrm{DP}} \right]_+ ,$$

$$b_{c,\mathrm{ES}}^* = \left[\frac{t_c(1-s)}{n(1-s^{1/n})} - \lambda_{-c}^{\mathrm{ES}} \right]_+ ,$$

and

$$b_{c,\mathrm{COL}}^* = \left[\frac{t_c}{n} s^{(n-1)/n} - \lambda_{-c}^{\mathrm{COL}} \right]_+ .$$

For complete censorship, $s = 1$, so

$$b_{c,\mathrm{DP}}^* = \left[t_c - \lambda_{-c}^{\mathrm{DP}} \right]_+ ,$$

$$b_{c,\mathrm{ES}}^* = \left[t_c - \lambda_{-c}^{\mathrm{ES}} \right]_+ ,$$

and

$$b_{c,\mathrm{COL}}^* = \left[\frac{t_c}{n} - \lambda_{-c}^{\mathrm{COL}} \right]_+ .$$

Economic Censorship Resistance

The paper defines economic censorship resistance as the attacker's expected expenditure relative to the transaction's payment:

$$\boxed{\mathrm{eCR}_c(s) = \frac{C_c(s)}{P_c}.}$$

For even-split and collective,

$$P_c^{\mathrm{ES}} = P_c^{\mathrm{COL}} = \beta + t_c.$$

For duplication-penalizing,

$$P_c^{\mathrm{DP}} = \beta + t_c \frac{n p_c (1-p_c)^{n-1}}{1-(1-p_c)^n}.$$

The paper's simulations find that eCR grows approximately linearly with the number of proposers. Duplication initially reduces throughput but throughput eventually plateaus while censorship resistance continues to increase.

Greater congestion and fee burning reduce eCR, while the duplication-penalizing TFM generally achieves the highest throughput and censorship resistance among the three mechanisms studied.

This is where the paper's model ends and our analysis begins. I suspect that the linear scaling is not going to hold under Sybils but let us look at it more formally.

Censorship under Sybils and Proposer Collusion

The original model contains n independent proposers. Let us look at the case where these proposers collude and thus are partitioned among q disjoint groups:

$$\Pi = \{C_1, \ldots, C_q\}, \qquad \bigcup_{g=1}^{q} C_g = [n].$$

The three benchmark cases are

$$q = n \quad (\text{no collusion}),$$

$$1 < q < n \quad (\text{partial collusion}),$$

and

$$q = 1 \quad (\text{full collusion}).$$

For coalition C_g, we will denote the utility from deviating from censorship and including target transaction c as

$$d_g(c) = \left[V_g^{\mathrm{in}}(c) - V_g^{\mathrm{out}}(c) \right]_+,$$

where $V_g^{\mathrm{out}}(c)$ is its optimal payoff while omitting c and $V_g^{\mathrm{in}}(c)$ is its optimal payoff when it may include c holding all other coalitions to censorship.

The minimum payments needed to sustain complete censorship are therefore

$$\boxed{C_{\Pi}^{\mathrm{op}}(c) = \sum_{g=1}^{q} d_g(c).}$$

This gives the new computation

$$\boxed{\mathrm{eCR}_{\Pi}^{\mathrm{op}}(c) = \frac{\sum_g d_g(c)}{P_c},}$$

where P_c is the transaction's payment conditional on inclusion.

Notice that this definition does not automatically increase when one proposer creates more keys/identities.

No collusion (Benchmark)

When all n proposers act independently, the paper's original bounds apply. To censor c with probability s, its equilibrium inclusion probability must satisfy

$$p_c^* = 1 - s^{1/n}.$$

The resulting expected bribe expenditures are

$$C_{\mathrm{NC}}^{\mathrm{DP}}(c,s) = n s^{1/n} \left[t_c s^{(n-1)/n} - \lambda_{-c}^{\mathrm{DP}} \right]_+,$$

$$C_{\mathrm{NC}}^{\mathrm{ES}}(c,s) = n s^{1/n} \left[\frac{t_c (1-s)}{n(1-s^{1/n})} - \lambda_{-c}^{\mathrm{ES}} \right]_+,$$

and

$$C_{\mathrm{NC}}^{\mathrm{COL}}(c,s) = n s^{1/n} \left[\frac{t_c}{n} s^{(n-1)/n} - \lambda_{-c}^{\mathrm{COL}} \right]_+.$$

For complete censorship, $s = 1$:

$$C_{\mathrm{NC}}^{\mathrm{DP}}(c,1) = n[t_c - \lambda_{-c}^{\mathrm{DP}}]_+,$$

$$C_{\mathrm{NC}}^{\mathrm{ES}}(c,1) = n[t_c - \lambda_{-c}^{\mathrm{ES}}]_+,$$

and

$$C_{\mathrm{NC}}^{\mathrm{COL}}(c,1) = [t_c - n\lambda_{-c}^{\mathrm{COL}}]_+.$$

The approximately linear dependence on n relies on the proposers being independent economic actors.

Partial collusion

Let $\rho_g^T(c)$ denote C_g's opportunity cost of making room for c under mechanism T.

For duplication-penalizing and even-split, a coalition can include c exactly once and collect the full tip. Hence

$$d_g^{\mathrm{DP}}(c) = [t_c - \rho_g^{\mathrm{DP}}(c)]_+,$$

$$d_g^{\mathrm{ES}}(c) = [t_c - \rho_g^{\mathrm{ES}}(c)]_+.$$

The corresponding censorship costs are

$$C_\Pi^{\mathrm{DP}}(c) = \sum_g [t_c - \rho_g^{\mathrm{DP}}(c)]_+,$$

and

$$C_\Pi^{\mathrm{ES}}(c) = \sum_g [t_c - \rho_g^{\mathrm{ES}}(c)]_+.$$

Under the collective TFM, suppose coalition C_g controls k_g identities and therefore receives fee share

$$\alpha_g = \frac{k_g}{n}.$$

Its utility from causing c to enter the common pool is only $\alpha_g t_c$, giving

$$d_g^{\mathrm{COL}}(c) = [\alpha_g t_c - \rho_g^{\mathrm{COL}}(c)]_+,$$

and

$$C_\Pi^{\mathrm{COL}}(c) = \sum_g [\alpha_g t_c - \rho_g^{\mathrm{COL}}(c)]_+.$$

So this is probably the most interesting case out of the three. It mostly hinges on ρ_g. I imagine that ρ_g will be sensitive to coupled utilities of other honest proposers in the setting.

SIDEWAYS: How can we figure out the best way to derive ρ_g? Maybe it's a coupled utility function with other (honest) agents in the system? Maybe it can be approximated as a bound of utilities that can be derived from different quantities of q? Open work here.

Full collusion

Now suppose all proposers form a full sybil collusion with total capacity B.

For all three TFMs, including a transaction exactly once produces aggregate proposer revenue t_i, while duplication cannot increase aggregate revenue and consumes capacity. The coalition therefore selects the B highest tip transactions exactly once.

Let

$$t_1 \geq t_2 \geq \cdots \geq t_m.$$

Its optimal revenue is

$$R^* = \sum_{i=1}^{B} t_i.$$

If target c belongs to the top B, censoring it replaces it with transaction $B+1$:

$$R^*_{-c} = R^* - t_c + t_{B+1}.$$

Thus, for every TFM,

$$\boxed{\Delta_c = R^* - R^*_{-c} = [t_c - t_{B+1}]_+.}$$

Consequently,

$$\boxed{C^{\mathrm{DP}}_{\mathrm{FC}}(c) = C^{\mathrm{ES}}_{\mathrm{FC}}(c) = C^{\mathrm{COL}}_{\mathrm{FC}}(c) = [t_c - t_{B+1}]_+.}$$

Since an optimally coordinated coalition includes c exactly once, its conditional payment is $P_c = \beta + t_c$. Therefore,

$$\boxed{\mathrm{eCR}_{\mathrm{FC}}(c) = \frac{[t_c - t_{B+1}]_+}{\beta + t_c}.}$$

Interestingly, this bound does not grow linearly with the number of proposers.

Under heavy congestion, t_{B+1} may be close to t_c, making censorship cheap. If an equally valuable substitute exists, then

$$t_c = t_{B+1} \quad \Longrightarrow \quad \Delta_c = 0.$$

Comparison Between Censorship Bounds

For complete censorship, the central expressions are:

TFM	No collusion	Partial collusion	Full collusion
Duplication-penalizing	$n[t_c - \lambda_{-c}^{\mathrm{DP}}]_+$	$\sum_g [t_c - \rho_g^{\mathrm{DP}}(c)]_+$	$[t_c - t_{B+1}]_+$
Even-split	$n[t_c - \lambda_{-c}^{\mathrm{ES}}]_+$	$\sum_g [t_c - \rho_g^{\mathrm{ES}}(c)]_+$	$[t_c - t_{B+1}]_+$
Collective	$[t_c - n\lambda_{-c}^{\mathrm{COL}}]_+$	$\sum_g [\alpha_g t_c - \rho_g^{\mathrm{COL}}(c)]_+$	$[t_c - t_{B+1}]_+$

We have to be careful before we assign a universal ordering under partial collusion because the values λ_{-c}^{T} and $\rho_g^T(c)$ differ across mechanisms.

I do think I can make some points though. Duplication-penalizing performs strongly without collusion but coalitions can avoid its duplication penalty through coordination. On the other hand, even-split does not destroy the user's payment under duplication but it is vulnerable to Sybils. Finally collective sharing creates weaker inclusion incentives and greater free riding.

So technically under full collusion, all three mechanisms collapse to the same aggregate censorship bound.

SIDEWAYS: It is tempting to conjecture that

$$\mathrm{eCR}_{\mathrm{NC}} \geq \mathrm{eCR}_{\Pi} \geq \mathrm{eCR}_{\mathrm{FC}}\,.$$

But as I said, I would be careful in doing so.

Addendum: All Sybil Versus Single Proposer

Aditya in our meeting mentioned that the full Sybil case simplifies down to the single proposer case. Let us try it out.

Suppose one operator controls every MCP proposer. Its aggregate capacity is

$$n\ell = B.$$

Because it coordinates all subblocks, it chooses exactly the same allocation as a single proposer with capacity B: the B highest tip transactions and each once.

Therefore,

$$\boxed{C^{\mathrm{opp}}_{\mathrm{MCP,\ full\ ownership}}(c) = C^{\mathrm{opp}}_{\mathrm{single\ proposer}}(c) = [t_c - t_{B+1}]_+.}$$

Likewise,

$$\boxed{\mathrm{eCR}^{\mathrm{op}}_{\text{MCP, full ownership}}(c) = \mathrm{eCR}_{\text{single proposer}}(c) = \frac{[t_c - t_{B+1}]_+}{t_c + \beta}.}$$

Hence proved.

SIDEWAYS: *As Manvir pointed out, Sybils and collusion should be separated out more. If the multiple proposer model is stake weighted then there is some cost associated with acquiring enough stake to be viable to be another proposer. That should be part of the calculations as well to accurately pinpoint the cost for Sybils. Also different proposers may still benefit from forming a collusion because it actually lets them avoid this cost which ultimately comes down to a bribery attack again.*

Another bound that comes from conversation with Manvir and Orfeas is as follows: If U_g^{ind} is the payoff for Sybils including the cost of creation of new proposers and U_g^{coord} is the payoff for collusion with existing players instead then $\Gamma_g = U_g^{\mathrm{coord}} - U_g^{\mathrm{ind}}$ is the surplus. This surplus can offset the revenue lost through censorship where if censoring c costs $d_g(c)$, a collusion formed specifically for the attack may require only $[d_g(c) - \Gamma_g]_+$ in additional compensation. Can be a good idea to pursue.

Interesting point by Sarisht: When is it incentive compatible for me to break the collusion? It may happen that I may make more utility by not colluding than staying as well. What are the conditions for that?

Thanks for reading. Reach out to talk more. I have been Abhimanyu Nag.