



Contracts: A Unified Lens on Congestion Control Robustness, Fairness, Congestion, and Generality

Anup Agarwal   

Carnegie Mellon University, Pittsburgh, PA, USA

Venkat Arun  

University of Texas at Austin, TX, USA

Srinivasan Seshan  

Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

Congestion control algorithms (CCAs) operate in partially observable environments. They cannot directly observe link capacities or competing flows. To share network resources fairly, CCAs (implicitly) communicate fair shares through observable signals. For instance, Reno encodes the fair share as $\propto 1/\sqrt{\text{loss rate}}$. We call such communication mechanisms as *contracts*. We find that the choice of contract fixes key steady-state performance metrics, including (1) robustness to errors in congestion signals, (2) fairness, (3) amount of congestion (e.g., delay, loss), and (4) generality (e.g., range of supported link rates). This leads to fundamental tradeoffs between these metrics. Further, we show that many contracts lead to starvation (extreme unfairness), and must be avoided. Hence, contracts are a powerful way to analyze tradeoffs and avoid pitfalls in CCA design and analysis. We empirically validate our findings and discuss their implications on CCA design and network measurement.

2012 ACM Subject Classification Networks → Transport protocols; Networks → Network control algorithms

Keywords and phrases Transport Protocols, Congestion Control, Fairness

Digital Object Identifier 10.4230/OASICS.NINeS.2026.8

Funding This work was supported in part by NSF grants CNS-2403026, CNS-2212102, CNS-2212390, CNS-2403026, and CCF-2422130.

Acknowledgements We are grateful to the anonymous reviewers of IMC 2025, CoNEXT 2025, and NINeS 2026, and our shepherd Radhika Mittal for feedback that helped improve our paper.

1 Introduction

Congestion control algorithms (CCAs) play a critical role in ensuring fair and efficient bandwidth allocation. Despite this, reasoning about their fairness has been ad hoc. Early CCAs were designed to avoid congestion collapse, and their fairness properties were analyzed retrospectively [16, 61, 43, 35]. With the rise of interactive applications, newer CCAs prioritized local performance metrics like latency and convergence time, where traditional CCAs performed poorly. However, this “local” perspective often led to undesirable global metrics like fairness. For instance, a version of TIMELY [49] has been shown to have infinite fixed points, causing unfair rate allocations ([72], § 2).

The result has been that a CCA that meets key efficiency (generality), latency (congestion), fairness, and robustness objectives remains elusive. Not only are real-world networks complicated and hard to predict, but empirical and theoretical evidence suggests that all objectives are not simultaneously achievable [8, 2, 72, 62, 54, 44]. Any systematization that reveals and navigates these tradeoffs can therefore guide better CCA design.



© Anup Agarwal, Venkat Arun, and Srinivasan Seshan;
licensed under Creative Commons License CC-BY 4.0

1st New Ideas in Networked Systems (NINeS 2026).

Editors: Katerina Argyraki and Aurojit Panda; Article No. 8; pp. 8:1–8:30

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We contribute to this effort by systematizing the performance consequences of the mechanisms that CCAs use to coordinate fairness. CCAs make rate decisions under uncertainty, which arises from two sources. First, CCAs are *decentralized* (i.e., flows cannot directly communicate with each other and thus, are unaware of the others’ states). Second, they operate under *partial observability* (i.e., flows lack direct visibility into network state such as topology, link capacities, queue occupancies, or the number of competing flows).

To make optimal rate decisions, flows must infer uncertain state by relying on narrow observations of their own transmissions and acknowledgments (ACKs).¹ While probes can help estimate latent parameters such as link capacity and propagation delay [14], there is no way for a flow to unilaterally infer how many flows it may be competing with (§ 2). To work around this, CCAs (implicitly) coordinate with each other by encoding fair shares into observable signals. For instance, Reno uses loss rate to coordinate fair shares, where each endpoint transmits at a rate $\propto 1/\sqrt{\text{loss rate}}$ [47]. We call such communication mechanisms *contracts*. To our knowledge, all existing fair CCAs have a contract in their design, either implicitly (e.g., Reno [47], Vegas [45]), or explicitly (e.g., TFRC [21], Swift [36], Poseidon [64]).

CCA contracts differ in the signals they use and their shape (e.g., steeper vs gradual). For example, delay-based CCAs encode fair share as: “ $1/\text{delay}$ ” [12, 67, 9], “ $1/\text{delay}^2$ ” [36], or “ $e^{-\text{delay}}$ ” [8].² Contracts effectively parameterize the space of CCAs without fully specifying a CCA. For instance, the entire family of TCP-friendly CCAs [31, 30, 21, 10] “follow” Reno’s contract (i.e. send at rates $\propto 1/\sqrt{\text{loss rate}}$), but they differ in other aspects, such as stability and convergence time properties.

Tradeoffs. This variety in contracts raises the question: “What are the tradeoffs between different contracts?” In exploring this question, we discovered a surprising result. The contract – a design aspect that may not have been chosen explicitly – fully determines four key congestion control metrics: (1) *robustness* to noise in congestion signals, (2) *fairness* in a multi-bottleneck network, (3) amount of *congestion* (e.g., delay, loss), and (4) *generality* (e.g., the range of link rates the CCA supports). Robustness and fairness are better with gradual contracts functions (e.g., Vegas [12, 45] “ $1/\text{delay}$ ”), while congestion and generality are better with steeper contracts (e.g., Swift [36] “ $1/\text{delay}^2$ ”) (§ 4). Thus, we cannot have the best in all four metrics.

Additionally, for a fixed contract, tolerating more congestion allows supporting a larger bandwidth range. In that sense congestion and generality are also at odds. Contracts also determine link utilization and total throughput. Our analysis subsumes these under fairness definitions (§ 4), where throughput and fairness are known to be at odds [50].

Similar tradeoffs have been explored before. Arun et al. [8] discuss a special case where full generality precludes simultaneous robustness and bounded variation in congestion. We find that their proposed workarounds cause unfairness due to other reasons (§ 4, § 5). Zhu et al. [72] show that delay-based CCAs cannot simultaneously ensure fixed delays and fairness. That is, if a delay-based CCA maps the same delay value to multiple rates, it does not have a contract, and cannot be fair. The NUM (Network Utility Maximization) literature [35, 61, 43, 45] studies a subset of these tradeoffs for individual CCAs. To our knowledge, we analyze a wider range of tradeoffs and also generalize them to many CCAs.

¹ Measurements like packet round-trip times (RTTs) and losses can be inferred from these observations.

² Throughout the paper, we use delay to mean a queuing delay estimate (e.g. $\text{RTT} - \text{min RTT}$).

Mistakes. Contracts also go beyond prior work by enabling us to identify previously unknown scenarios and (avoidable) design mistakes that cause severe performance degradation or suboptimal performance in a wide range of CCAs (§ 2, § 5). We find that CCAs whose contracts have **extreme shapes** (e.g., logarithmic, exponential), **shifts** (e.g., $\text{rate} = 1/(\text{delay} - c)$ vs. $\text{rate} = 1/\text{delay}$), **clamps** (e.g., $\text{delay} = \max(c, 1/\text{rate})$), or **intercepts** can starve flows. We observe this in BBR [14], ICC [33], and Astraea [42]. Further, CCAs with an explicit contract (e.g., Swift [36]) do not need AIMD (Additive Increase, Multiplicative Decrease) updates to reach fairness. Instead, MIMD (Multiplicative Increase, Multiplicative Decrease) updates allow exponentially fast convergence to both fairness and efficiency (§ 6). Finally, we show that having fixed end-to-end thresholds, e.g., AIMD on delay [59, 11, 38], causes starvation on topologies with multiple bottlenecks.

Our methodology applies to not only analytically designed CCAs, but also black-box implementations, such as Tao (Remy) [60, 68], and learned CCAs, such as PCC [18] (online-learning) and Astraea [42] (reinforcement-learning or RL).

Blueprints. Using contracts, we build blueprints for CCA design and analysis. These blueprints take a “contracts-first” approach, making the tradeoffs and mistakes intuitive and actionable (§ 2). First, the tradeoffs are obvious from the contract choice (after a few to no algebraic steps), forcing the designer to deliberately choose between uncomfortable tradeoffs and picking Pareto-optimal points. Second, contracts align with how our community designs CCAs: a variety of CCAs start with a target rate equation and build dynamics around it, e.g., Swift [36], TFRC [21].

In contrast, other methodologies often start from a utility function: local utilities (e.g., “ $\log(\text{tput}) - \log(\text{delay})$ ” in Copa [9]), global utilities (e.g., NUM/Remy [68]), or reward functions in RL [42]. These hide tradeoffs giving an illusion of unilaterally optimizing latency or throughput when these are intertwined with other metrics like fairness, robustness, and generality (§ 4). Consequently, many designs pick sub-optimal tradeoffs as inadvertent artifacts of other design decisions (§ 3, § 4, § 5).

Use cases. We envision that future work will use our blueprints for the design and analysis of CCAs enabling deliberate selection of Pareto-optimal tradeoffs and avoiding performance mistakes. We demonstrate these uses by (1) analyzing and exposing performance issues in a variety of CCAs (§ 3.1, § 5, § 6), and (2) designing a few canonical contract-based CCAs (§ 4, § 6, § 7). In another paper [3], we also use contracts to design the first CCA that provably avoids starvation on networks with delay jitter showing that the starvation result by Arun et al. [8] is not as pessimistic as it seemed.

Contributions. We build contracts to formalize the mechanisms CCAs use to coordinate fairness (§ 2, § 3). We discover and analytically quantify the tradeoffs (§ 4) and avoidable design mistakes (§ 2, § 5) exposed by contracts. We build blueprints of CCA design and analysis (§ 2) to help avoid the mistakes and navigate the tradeoffs. We validate our findings using simulation and emulation, finding a near-perfect match between analysis and measurements (§ 7). We end by discussing potential ways to work around the tradeoffs, limitations of contracts, and how we may expand their uses (§ 8, § 9). For instance, much of our current focus is on steady-state performance, and we hope to extend reasoning to short-flows and inter-CCA fairness. Further, if the network supports mechanisms like fair queueing, or if fairness is not desired, then contracts are not needed, and the tradeoffs do not apply.

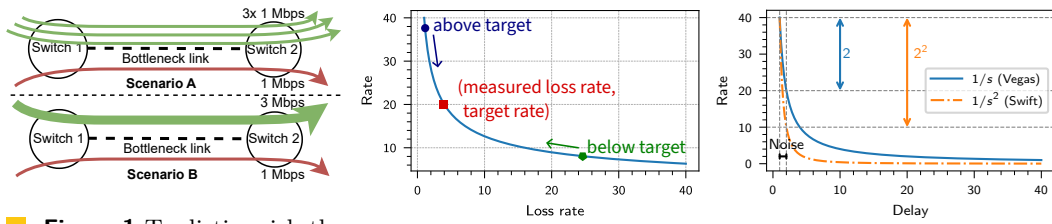


Figure 1 To distinguish the scenarios, the red (lower) flow must coordinate with green (upper) flows. **Figure 2** Reno uses loss rate to encode fair shares. Flows update rates towards the target. **Figure 3** A steeper contract implies worse robustness to noise.

2 Motivation (Why contracts?)

Since CCAs do not know how many flows they are competing against, they need some form of agreement between flows to determine their fair share. For instance, consider the scenarios in Figure 1. If the green (upper) flow in scenario B exactly emulates the cumulative effect of the three green flows in scenario A, then the red flow cannot tell the difference. Its observations (the time series of packets transmissions and acknowledgements) are identical. Yet, its fair share is different.

Contracts help flows to disambiguate between such scenarios. Since flows cannot directly communicate with each other, all flows “agree” to follow an encoding of fair share into globally observed signals. For example, Reno uses: “ $\text{rate} \propto 1/\sqrt{\text{loss rate}}$ ” [47, 21]. While this is not explicit in the design, the emergent behavior is that all flows react to losses such that they effectively measure the average packet loss rate and calculate a “target rate” using this formula. Then they increase or decrease their actual sending rate to move towards this target (Figure 2). Returning to scenario B in Figure 1, both the red and green flows measure the same loss rate, and hence calculate the same target rate, but they have different sending rates. Thus, at least one of them will change their rate until they reach a steady-state where everyone’s sending rates are equal.

Mathematically, the contract forces a *unique* equilibrium. Consider a fluid model execution of a CCA on a dumbbell topology, with n flows (f_1 , to f_n), flow f_i having an RTprop (round-trip propagation delay) of R_i seconds and a link capacity of C packets/second. The fluid model equations yield $n + 1$ independent variables with 1 independent equation: $\sum_i \text{rate}_i = C$. Where, $\text{rate}_i = \text{cwnd}_i/\text{RTT}_i$ and $\text{RTT}_i = \text{delay} + R_i$. The independent variables are cwnd_i and delay . The CCA through its contract (e.g., $\text{rate}_i = \text{cwnd}_i/\text{RTT}_i = 1/\text{delay}$) yields n additional equations to ensure a unique solution to this system.

To our knowledge, all existing CCAs use a similar method of coordination. They just use different signals and contract shapes. E.g., DCTCP, DCQCN, and MPRDMA use average ECN marking rate [4, 71, 46, 72]; Poseidon uses the maximum per-hop delay [64] collected using INT (In-band Network Telemetry); AIMD on delay uses bytes (or time) between high delay (§ 5). Even CCAs generated through machine-learning (e.g., Astraea [42]) implicitly learn to use contracts (§ 3.1).

Due to the coupling between rate and congestion, a CCA’s contract choice induces a variety of tradeoffs, e.g., halving the link capacity quadruples the steady-state loss rate for Reno (§ 4).

2.1 Isn't this obvious? Why this paper?

Contracts share mathematical foundations with NUM (§ 3.2) and some of our results may seem like simple extensions of existing results. Despite this, many CCAs (including recent ones) repeatedly commit avoidable mistakes resulting in poor performance. We detail some of these mistakes below with others in § 5. To our knowledge, outside of mistake #1, other mistakes (including § 5) have not been documented before. Contracts helped us discover these issues, and make it easier to systematically avoid them in future designs.

These mistakes stem from several anti-patterns, including, treating CCA design as a mere reaction to congestion (without analytically understanding their consequences) or blindly using AIMD in hopes of ensuring fairness. The most prominent anti-pattern is myopically optimizing latency without reasoning about fairness. This manifests in various ways, including, setting constant delay targets in hand-designed CCAs (§ 5) or in reward functions of RL-based designs (§ 3, § 5). However, latency is intertwined with other metrics and cannot be optimized unilaterally (§ 4, § 5). Note that this is different from the tradeoff between latency and throughput, which stems from variations in link capacity as opposed to the coordination mechanisms used by CCAs.

Mistake #1: Not having a contract. Many CCAs do not have a contract, i.e., there is no unique mapping between steady-state rate and congestion signals to force a unique equilibrium. Consequently, these CCAs can exhibit arbitrary amounts of unfairness. For instance, Theorem 4 of [72] shows that a version of the TIMELY CCA [49] admits infinitely many solutions to the steady-state equations described above. Each solution corresponds to a different allocation of flow rates (provided they sum to link capacity), with a potentially arbitrarily large ratio of flow rates (arbitrary unfairness).

Mistake #2: Picking extreme contracts. Recent proposals like ICC [33], Astraea [42], and the exponential CCA from [8] optimize for a narrow subset of performance metrics, yielding contracts that are good for some metrics but extremely bad for others. All three CCAs cause extreme unfairness in multi-bottleneck topologies (§ 4). ICC and Astraea are also extremely susceptible to network jitter (§ 4).

Mistake #3: Conflating AIMD with contracts, and picking sub-optimal dynamics. Swift [36] uses AIMD to update cwnd despite having an explicit contract. AIMD is not necessary for fairness and unnecessarily increases convergence time. With explicit contracts, MIMD updates can reach fairness exponentially fast while maintaining a stable control loop (§ 6). For instance, in Figure 2, all flows measure the same loss rate and compute the same target rate. Flows can update their current rate to move towards the target using any increment choice (additive or multiplicative). Fairness is ensured despite MIMD updates because flows stop changing their rate *if and only if* all flow rates are equal to the target rate (and hence to each other).

Likewise, PowerTCP [1] claims that RTT-gradient based CCAs (“current-based” in [1]) are more reactive and precise than RTT/delay/loss/ECN-based CCAs (“voltage-based” in [1]). We argue that reactivity is orthogonal to the choice of congestion signal. We can have exponentially fast convergence to both fairness and efficiency even when using “voltage-based” control (§ 6).

Finally, works like Poseidon [64] artificially distinguish AIMD control and “target scaling” (contracts) even though they are mathematically equivalent. AIMD *implicitly* creates a contract (or scaling/mapping between rate and congestion), while “target scaling” is an explicit contract.

2.2 Avoiding mistakes using contracts

While there is no exhaustive list of mistakes, by following a contract-first blueprint, many common mistakes can be avoided.

Design blueprint. To design a new CCA, first, (D1) pick a contract. In § 3, we define contracts as a function. Picking a contract involves deciding its (D1.1) input (e.g., delay, loss, ECN, etc.) and output (e.g., rate, cwnd, fraction of link, etc.), (D1.2) shape (e.g., linear, square-root, exponential, etc.), and (D1.3) parameters (e.g., scale, shift, clamps, etc.). § 4, § 5 and § 8 give guidance on how these choices affect steady-state performance. Then, (D2) implement dynamics to follow the contract. § 6 gives guidance on how the dynamics impact convergence time/stability along with other design considerations. Note that unless congestion control is solved, our list of considerations is necessarily incomplete.

Analysis blueprint. For analyzing an existing CCA, one should: (A1) compute its contract (§ 3.1), (A2) see where the contract lies in the tradeoff space (§ 4), (A3) identify any obvious issues due to shifts/clamps/intercepts in the contract (§ 5), and (A4) compare dynamics with those in § 6.

Figure 1 shows why some form of agreement between flows is necessary to achieve fairness with end-to-end CCAs. While we do not formally prove this, we believe this agreement can always be represented as a contract function, and consequently, the tradeoffs induced by contracts are fundamental. The only way to work around the tradeoffs is to change the input/output in the contract to decouple performance metrics from the contract (§ 8). We believe that other efforts for improving steady-state performance are futile and will likely lead to reinventing a CCA already covered in the design/tradeoff space in § 4.

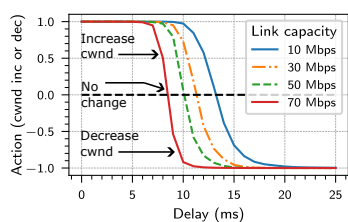
3 Contracts

► **Definition 3.1.** The contract of a CCA is a function of the form `average sending rate = func(aggregate statistic)` that “describes” the CCA’s steady-state behavior (e.g., at time infinity) when competing with itself on a dumbbell topology. The aggregate statistic (e.g., delay, loss rate, etc.) is derived from the CCA’s observations: the time series of sending and acknowledgment sequence numbers, along with any explicit signals (e.g., ECN [57], INT [64]). As a shorthand, we use “ $r = \text{func}(s)$ ”, or “ $r = f(s)$ ”. ◻

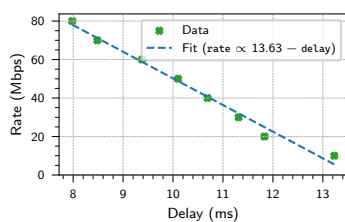
Here, “describe” depends on the form of steady-state. Most CCAs exhibit a fixed-point or a limit-cycle equilibrium. In the fixed-point case (e.g., Vegas [45]), the set of fixed-points is same as the set of input/output pairs of `func`. If this set forms a relation rather than a function, the CCA is unfair and does not have a contract (e.g., TIMELY [49]). In the limit-cycle case (e.g., sawtooth behavior in Reno [31], DCTCP [4]), we convert the cycle into a fixed-point by considering the aggregate behavior (e.g., average sending rate) over the cycle. For instance, DCTCP’s contract is “`avg cwnd = 1/(avg ECN marking rate)2`”, derived in [4, 5].

In general, it is difficult to rigorously define contracts given CCA behaviors can be complex. We discuss possible extensions in § 9. Our current definition captures all the CCAs we analyzed; for each CCA, we can either compute its contract or show that it is unfair and does not have a contract.

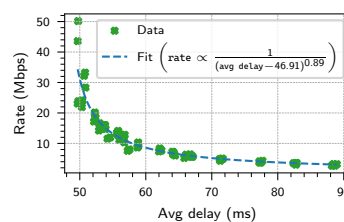
In § 2, we introduced the contract for a CCA as its encoding of fair shares into observable signals, which enforces a unique equilibrium in the fluid model equations. This “encoding” and the “steady-state behavior” are two equivalent views of a contract: the CCA follows the



■ **Figure 4** Astraea’s state to action map (from [42]). Points on the dashed black line are the fixed-points.



■ **Figure 5** Curve fitting analytical fixed-points (from Figure 4) to compute Astraea’s contract.



■ **Figure 6** Curve fitting empirical fixed-points to compute BBR’s contract. BBR is cwnd-limited with multiple flows [65].

encoding in steady-state revealing the encoding in its steady-state behavior. We adopt the steady-state view in defining contracts as it provides a *constructive* definition, allowing one to derive the contract directly from a given CCA.

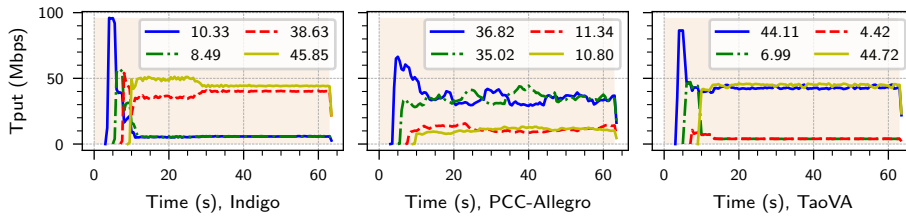
3.1 Computing contracts

We compute the contract of a CCA by analyzing its steady-state behavior. We run it analytically or empirically on a variety of dumbbell configurations (with different capacities, RTprops, buffer sizes, and flow counts) and collect the set of steady-state observations. Then, we group this set based on an aggregate statistic, such that observations with the same statistic value have the same rate. These equivalence classes define the CCA’s contract function. When the same steady-state observations result in different rates, the CCA is unfair and does not have a contract.

Often, this reduces to program analysis of the CCA’s code (see Vegas’s example below) or fluid model analysis (see [72] for DCQCN [71]). One rarely has to compute contracts oneself. The literature has already computed contracts for most CCAs; we cite these in the “`func(s)`” column of Table 3. This allows us to focus on the performance impact of contract choice, instead of computing contracts. For completeness, we show below examples of analytical and empirical contract derivations.

Note, even for CCAs where fluid modeling is hard, their equilibrium behavior is well understood. For example, BBR [14] encodes fair shares in delay when cwnd-limited [8], and in “growth in delivery rate” when rate-limited [23]. In some cases the aggregate statistic is multi-dimensional and/or the contract is a compound function. For instance, Copa emulates Reno on detecting competing loss-based flows [9]. When the loss rate is zero, Copa follows a delay-based contract versus Reno’s contract otherwise. For simplicity, we focus on homogeneous settings where flows use the same CCA and run in a single mode, e.g., we disable mode-switching in Copa. We expect the different modes to follow the tradeoffs according to the contract followed in the mode.

Analytical contract derivation (Vegas, taken from [61]). To update its cwnd, Vegas compares $\text{diff} = \text{expected_throughput} - \text{actual_throughput} = \text{cwnd}/\text{RTprop} - \text{cwnd}/\text{RTT}$ to α_{rate} , where α_{rate} is a configurable parameter. Vegas increases cwnd when diff is larger than α_{rate} , and decreases otherwise. Fixed-point steady-state occurs when Vegas has no incentive to change its cwnd, i.e., $\text{diff} = \alpha_{\text{rate}}$. This equation gives us the contract. We substitute $\text{cwnd} = \text{rate} \cdot \text{RTT}$, and $\text{RTT} = \text{RTprop} + \text{delay}$, and simplify, to get Vegas’s contract: $\text{rate} = \alpha_{\text{rate}} \cdot \text{RTprop}/\text{delay}$.



■ **Figure 7** Indigo [70], Fillp/FillpSheep (TACK [41]), PCC-Allegro/Vivace/Experimental [18], and TaoVA (Remy) [60, 68] are unfair and have no contract. We only show three for brevity. The lines show the time series of throughput and the legend shows the average throughput of the 4 flows.

■ **Table 1** Contracts computed for CCAs in Pantheon. We omit constant scaling factors for brevity.

BBR	$(\text{avg delay} - 46.91)^{-0.89}$	Vegas	$\text{avg delay}^{-1.38}$	Copa	$\text{p50 delay}^{-1.08}$
LEDBAT	$102 - \text{p50 delay}$	Sprout	$134 - \text{avg delay}$	Cubic	$\text{loss rate}^{-0.65}$

Analytical contract derivation (Astraea, adapted from [42]). We can compute contracts even for black-box CCAs without running them. We show this for the RL-based CCA Astraea [42], which implicitly learns a contract. Figure 4 shows its feedback (state) to action mapping. We obtain this by querying its neural network’s action for different feature vectors (states). For a given capacity and delay combination (state), action > 1 increases cwnd, and action < 1 decreases cwnd. Astraea is at a fixed-point when it has no incentive to change cwnd (action = 0). For each link capacity (or fair share), Astraea maintains a unique delay, given by the X-coordinates of points with action = 0. We fit a curve ($\text{rate} = a(\text{delay} + b)^c + d$) to these points to get its contract (Figure 5).

Empirical contract derivation (CCAs in Pantheon). To demonstrate that (1) contracts are easy to compute, and (2) fair CCAs have a contract, we empirically derive contracts for all CCAs in Pantheon [70] using an automated procedure. We use Pantheon to run each CCA for 60 seconds on a dumbbell topology with link capacities of 24, 48, and 96 Mbps, 40 ms RTprop, 4 BDP buffer, and vary the flow count from 2 to 8 (except Cubic, which uses 1 BDP buffer). These yield samples for throughput and three aggregate statistics: p50 delay (ms), avg delay (ms), and loss rate. We compute contracts by fitting a curve to these points (Figure 6) and pick the statistic “s” and fit that minimizes mean squared error.

The configurations we picked are not special. One should pick the dumbbell configurations that best align with the target deployment of the CCA that one is analyzing. Most CCAs produce relatively continuous steady-state behaviors and interpolating/extrapolating steady-state behavior from a few configurations is often enough for computing contracts.

Of the 17 CCAs in Pantheon, we faced deprecation/dependency issues for QUIC Cubic and Verus. We run the remainder 15. 2 CCAs (SCReAM and WebRTC) get ≈ 0 utilization (consistent with the Pantheon report [69]). 7 CCAs (Indigo, Fillp/FillpSheep, PCC-Allegro/Vivace/Experimental, TaoVA) are unfair (Figure 7) and hence, do not have a contract. Unfair means different flow rates despite similar statistics/signals, indicating the absence of a unique signal-to-rate mapping (contract). Table 1 shows the contracts derived for remainder 6 CCAs.

Our automated procedure is for illustration purposes and is not full-proof. These contracts are for the range of networks we experiment with. The shifts in delays (e.g., 46.91 in BBR and 102 in LEDBAT) depend on the RTprops or buffer sizes, e.g., for BBR the shift is equal to RTprop. One would require more experiments to deduce this empirically. Similarly, the

unfair CCAs may be fair and exhibit contracts on a narrower range of scenarios. For instance, we find that PCC-Vivace and PCC-Experimental exhibit the contracts “`avg delay`^{-0.4}” and “`p50 delay`^{-0.46}” respectively, if we only consider data from the 96 Mbps link. Finally, the contracts may be compound functions that take different shapes on different ranges of networks and use other statistics than the three we considered. For instance, [53] shows Reno’s contract under different operating regimes: loss detections are dominated by timeouts vs duplicate ACKs.

3.2 Scope

We only consider strictly decreasing contracts, with closed intervals as their domain and range, ensuring they are continuous and invertible. All contracts that we are aware of are decreasing: the statistic typically measures congestion; an increasing contract suggests increasing the rate with increasing congestion. This further increases congestion, creating a positive feedback loop.

Rate- vs cwnd-based contracts. A contract can be written in terms of rate or cwnd, and using `rate = cwnd/RTT` to convert between the two forms. Independent of the form, the CCA can be implemented using either rate or cwnd (§ 6). Without loss of generality, we consider rate-based contracts.

Exploration over time. CCAs often use control loops that periodically explore the network to estimate latent parameters, and the effects of this exploration are reflected in the aggregate statistic used in their contract. For example, Copa deliberately oscillates its sending rate to alternately drain and build queues: draining enables estimation of the round-trip propagation delay (RTprop, also called baseRTT or minRTT), while building queues enables estimation of the “standingRTT”. Copa then computes the aggregate statistic as $s = \text{delay} = \text{standingRTT} - \text{RTprop}$. Similarly, BBR launches bandwidth probes every 8 RTprops and RTT probes every 10 seconds. When BBR is rate-limited, the aggregate statistic captures the growth in ACK rate during a bandwidth probe; when cwnd-limited, the RTT-probes estimate the RTprop and consequently queueing delay, analogous to Copa.

In general, a CCA’s aggregate statistic can be complex, combining multiple measurements and the dynamics of the CCA’s exploration, rather than reflecting a single congestion signal at a specific time instant. We do not exhaustively explore all the nuances in how these statistics are computed, though these can also significantly affect performance. Nevertheless, we hope that by viewing CCAs through the lens of contracts designers can reason about and navigate these complexities.

Non-globally observed signals and RTT or RTprop bias. CCAs often rely on signals that are not globally observed, i.e., different flows may see different values for the signal. For example, flows can experience different RTTs or RTprops. These non-global signals can introduce biases in the sending rates. Contracts explicitly capture these biases. For instance, Reno allocates more rate to flows with lower RTTs, following the contract $\text{rate} \propto 1/(\text{RTT} \sqrt{\text{loss_rate}})$. Similarly, Vegas favors flows with higher RTprop, as expressed by the contract $\text{rate} = \alpha_{\text{rate}} \cdot \text{RTprop}/\text{delay}$.

Although such biases have attracted significant attention in the community, they can often be removed by adjusting the contract. For example, many Vegas implementations eliminate the bias by setting $\alpha_{\text{rate}} = \alpha_{\text{pkts}}/\text{RTprop}$ [29]. The same can be done for Reno by removing the RTT term from its contract and implementing the contract using TFRC [21].

Because differences due to non-globally observed signals can be easily removed, and many “good” CCAs already do this, we omit these factors in the bulk of our analysis. Nevertheless, CCA designers and analysts should remain aware of these effects and their potential impact on performance.

Contracts in NUM. The concept of contracts can also be described using NUM [35, 61, 43, 44]. A contract maps the aggregate statistic (congestion measure or price) to rate, same as the demand function (target rate) for a given price [35]. The inverse of a contract represents the link price (e.g., target delay) for a given load (e.g., link’s ingress rate) [35]. Consequently, the utility function that the CCA optimizes is (derived in [43]): $U(r) = \int \text{func}^{-1}(r) dr$ (1). Here, U is the utility derived from a rate of r . Eq. 1 only holds for statistics that add up over links (e.g., delay). For other statistics (e.g., max per-hop delay), the utility is different [64].

4 Metrics and tradeoffs

Metrics. We list the metrics that CCAs typically optimize, and study how contracts affect them.

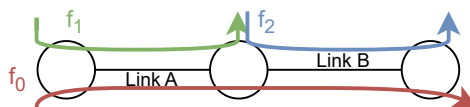
1. Link utilization, flow throughput
2. Amount of congestion (e.g., delay, loss)
3. Stability, and convergence time/reactivity
4. Fairness (on general topologies)
5. Robustness to noise in congestion signals
6. Generality (range of link rates, flow counts)

Of these metrics, fairness notions (see below) subsume Pareto-optimality; that is, the bottleneck links are fully utilized. They also subsume the tradeoff between total throughput and fairness (see below). Stability and convergence time are concerned with transient behavior as opposed to steady-state equilibrium. While contracts may affect them, we can often independently optimize them by choosing how fast a CCA’s sending rate moves toward the target rate (§ 6). We are left with the following four metrics: robustness, fairness, congestion, and generality.

Approach. The choice of contract determines these four performance metrics. We show this by expressing each metric explicitly as a function of the contract `func`. Conversely, specifying a desired value for one metric constrains the contract, which in turn constrains the other metrics (i.e., a tradeoff). We derive these tradeoffs quantitatively by characterizing the feasible values of one metric under a desired constraint on another metric. We keep our metric definitions unit-less to minimize dependence on network parameters (e.g., link capacity, RTprop).

We use running examples of Vegas [12, 45] ($r = 1/s$) and Swift [36] ($r = 1/s^2$), where s is delay. We omit constant factors that ensure that the units are consistent. E.g., the unit of 1 in Vegas’s and Swift’s contract is bytes, and byte · seconds respectively, so that the function value has the unit bytes/second. Table 3 shows where existing CCAs fall in the tradeoff space.

Generality (e.g., range of link rates supported). Practical constraints limit congestion signals to a finite range, $s \in [S_{\min}, S_{\max}]$. Small delays and losses are difficult to measure due to noise, while large delays and losses increase application perceived latency. Since the contract maps congestion signals to rates, this range confines the CCA’s operating range



■ **Figure 8** Parking lot topology.

■ **Table 2** Rates under different fairness notions on the parking lot topology (Figure 8). Both links have capacity $C = 1$. r_i shows the rate of flow f_i . From top to bottom, fairness (rate equality) improves but total throughput (or rate) decreases.

Fairness notion	α	Global utility	r_0	$r_1 = r_2$	Total rate, $\sum_i r_i$	Example CCA
Max throughput	0	$\sum_{i=0}^k r_i$	0	1	2	
Proportional	1	$\prod_{i=0}^k r_i$	1/3	2/3	5/3	Vegas
Min potential delay	2	$\sum_{i=0}^k -1/r_i$	$\sqrt{2} - 1$	$2 - \sqrt{2}$	$3 - \sqrt{2}$	Reno
Max-min	∞	$\min_i r_i$	1/2	1/2	3/2	Poseidon

of link rates to $[\text{func}(S_{\max}), \text{func}(S_{\min})]$, where $\text{func}(S_{\min}) > \text{func}(S_{\max})$ because func is decreasing. CCAs must select the constants in func to suit the networks they operate on. For instance, if delay spans a $200\times$ range, e.g., $\text{delay} \in [0.5 \text{ ms}, 100 \text{ ms}]$, Vegas can support only a $200\times$ range of sending rates (e.g. 1 Mbps to 200 Mbps), whereas Swift can support a $200^2\times$ range (e.g. 1 Mbps to 40 Gbps). In general, steeper contracts enable a broader range of bandwidths for the same domain of the congestion signal.

Note, our community does not generally say that a given CCA does not support a particular link rate. Here, by saying that a CCA cannot support a link rate, we mean that when the CCA is running on that on that link rate, some performance metric (e.g., high utilization, fairness, low delay, etc.) breaks down. For example, at low link rates, multiple competing Astraea flows can exhibit starvation (Figure 11). Conversely, at high link rates, all “delay-convergent” CCAs exhibit flow starvation [8].

Robustness to noise. We study the impact of “ δs ” error in the statistic “ s ” on the CCA’s rate. We quantify the impact by looking at the (worst-case) ratio of rates without and with the noise: Error factor = $\max_s \frac{\text{func}(s)}{\text{func}(s+\delta s)}$ (2). A higher error factor means that the CCA is more sensitive to noise and is less robust. From Eq. 2, the error factors for Vegas ($r = 1/s$) and Swift ($r = 1/s^2$) are “ $1 + \delta s/S_{\min}$ ”, and “ $(1 + \delta s/S_{\min})^2$ ” respectively. If δs error perturbs Vegas by $2\times$, then the same error perturbs Swift by $4\times$. Steeper contracts yield higher (worse) error factors (Figure 3).

The magnitude of noise “ δs ” that a CCA must tolerate depends on the chosen statistic. For example, delay estimates may exhibit tens of milliseconds of noise [24], whereas loss-rate estimates may fluctuate by a few percent [37]. CCAs often attempt to mitigate such noise by filtering or smoothing raw congestion signals (e.g., using min, max, or averages over raw delay samples). In this setting, δs represents the residual noise in the statistic *after* filtering, rather than noise in the raw signal itself.

Even small residual noise can significantly distort rate allocations when combined with certain contract shapes. For instance, Figure 15 shows that flows can starve under BBR and Copa despite the use of filtering mechanisms (BBR takes the max over ACK rates; Copa takes the min over delays). While the contract’s shape and input/output choice play a central role in determining robustness, they may not be sufficient in isolation. CCA designers should holistically reason about how noise in raw signals propagates through the aggregate statistic and the contract function to ultimately affect rate allocation.

■ **Table 3** Existing CCAs and the tradeoff space. We organize into 3 sections: good contracts (blue, top), corner points (green, middle), and bad contracts (pink, bottom). For Reno, s is loss rate. For DCTCP, s is ECN marking rate. For Poseidon s is max per-hop delay. For AIMD on delay, s is bytes between high delay. For all other CCAs s is delay. We assume loss and ECN marking rates are small enough to approximate `fold` as \sum . For Poseidon, `fold` is max and for AIMD on delay, `fold` is harmonic sum (§ 5). `fold` is \sum for all other CCAs. In ICC and Astraea, unfairness and robustness error are worst when $s \approx S_0$.

CCAs	$\text{func}(s) = \mathbf{f}(s)$	Robustness error for δs noise	Unfairness with k hops	Congestion for n flows	Range of bandwidths
		$\max \frac{\mathbf{f}(s)}{\mathbf{f}(s+\delta s)}$	$\max \frac{\mathbf{f}(s_k)}{\mathbf{f}(\text{fold}(s_k, \dots))}$	$\max \frac{\mathbf{f}^{-1}(C/N)}{\mathbf{f}^{-1}(C)}$	$\frac{\mathbf{f}(S_{\min})}{\mathbf{f}(S_{\max})}$
		Lower is better	Lower is better	Lower is better	Higher is better
		Want gradual \mathbf{f}	Want gradual \mathbf{f}	Want steeper \mathbf{f}	Want steeper \mathbf{f}
FAST, Vegas, Copa	$\frac{1}{s}$ [67, 43, 8]	$1 + \frac{\delta s}{S_{\min}}$	k	n	$\frac{S_{\max}}{S_{\min}}$
Swift, DCTCP	$\frac{1}{s^2}$ [36, 4, 5]	$\left(1 + \frac{\delta s}{S_{\min}}\right)^2$	k^2	\sqrt{n}	$\left(\frac{S_{\max}}{S_{\min}}\right)^2$
Reno	$\frac{1}{\sqrt{s}}$ [31, 47, 43]	$\sqrt{1 + \frac{\delta s}{S_{\min}}}$	\sqrt{k}	n^2	$\sqrt{\frac{S_{\max}}{S_{\min}}}$
Poseidon	e^{-s} [64]	$e^{\delta s}$	1	$\log n$	$e^{S_{\max} - S_{\min}}$
α -fair (§ 4)	$\frac{1}{\sqrt[s]{s}}$ (Eq. 1)	$\sqrt[\alpha]{1 + \frac{\delta s}{S_{\min}}}$	$\sqrt[\alpha]{k}$	n^α	$\sqrt[\alpha]{\frac{S_{\max}}{S_{\min}}}$
Exp.	e^{-s/S_0} [8]	$e^{\delta s/S_0}$	$e^{k \cdot S_{\max}/S_0}$	$\log n$	$e^{\frac{S_{\max} - S_{\min}}{S_0}}$
ICC	$\log\left(\frac{S_0}{s}\right)$ [33]	∞	∞	$\left(\frac{S_0}{S_{\min}}\right)^{\frac{n-1}{n}}$	$\frac{\log S_0 - \log S_{\min}}{\log S_0 - \log S_{\max}}$
Astraea	$C_0(1 - \frac{s}{S_0})$ (§ 3)	∞	∞	$\frac{nC_0 - C_{\max}}{n(C_0 - C_{\max})}$	$\frac{S_0 - S_{\min}}{S_0 - S_{\max}}$
AIMD on delay	$\frac{1}{\sqrt{s}}$ (§ 5)	$\sqrt{1 + \frac{\delta s}{S_{\min}}}$	∞	n^2	$\sqrt{\frac{S_{\max}}{S_{\min}}}$

Fairness. For general network topologies, there exist multiple notions of fairness (Table 2). These are typically parameterized using α -fair utility functions [50] given by: $U_\alpha(r) = \frac{r^{1-\alpha}}{1-\alpha}$, where U represents the utility a flow derives from a rate of r , and $\alpha > 0$. The rate allocation maximizes the sum of utilities across all flows ($\sum_i U_\alpha(r_i)$) (global utility). Larger values of α indicate greater fairness at the cost of total throughput [62]. Typically, $\alpha \geq 1$ is desired.

Not all contracts correspond to an α -fair utility. As a proxy, to compare the fairness of contracts, we study their behavior in the parking lot topology (Figure 8). This topology exposes differences between fairness notions (Table 2), and is common in data centers (e.g., when inter- and intra-rack flows compete). In Figure 8, each link has a capacity of C bytes/second and flows experience congestion signals from multiple congested links. The *long* flow (f_0) observes signals from two hops, while the *short* flows (f_1 and f_2) observe signals from only one hop. More generally, for k hops, we consider flows $\langle f_0, f_1, \dots, f_k \rangle$, where f_0 observes signals from all k hops, and other flows see signals from a single hop.

We quantify unfairness as the (worst-case) throughput ratio of f_k to f_0 . We derive unfairness using three steady-state equations. First, at each hop, the sum of the incoming rates is equal to the link capacity: $\forall i \geq 1, r_0 + r_i = C$ (3). Second, each flow's rate (r_i) and statistic (s_i) follow the contract $\forall i \geq 0, r_i = \text{func}(s_i)$ (4). Third, f_0 sees an accumulation of signals from all the hops: $s_0 = \text{fold}(s_1, s_2, \dots, s_k)$ (5). Here, `fold` describes the statistic f_0 sees when other flows see s_1, s_2, \dots, s_k . For delay, $s_0 = \sum_{i=1}^k s_i$, as delays add up over hops. For loss rate and ECN marking rate, $s_0 = 1 - \prod_{i=1}^k (1 - s_i)$, as survival probability gets multiplied over the hops. Note, when the absolute value of loss rates (or ECN marking rates) is small (e.g., 10^{-2}), then `fold` can be approximated as a sum. For max per-hop delay (e.g., Poseidon [64]), $s_0 = \max(s_1, s_2, \dots, s_k)$.

From Eq. 3, we get $r_1 = r_2 = \dots r_k$. Given, func is invertible, we get $s_1 = s_2 = \dots = s_k$. Then, the worst-case throughput ratio is given by: $\max_C \frac{r_k}{r_0} = \frac{\text{func}(s_k)}{\text{func}(s_0)} = \max_C \frac{\text{func}(s_k)}{\text{func}(\text{fold}(s_k, s_k, \dots, k \text{ times}))}$ (6). Here, s_k solves $\text{func}(s_k) + \text{func}(\text{fold}(s_k, \dots, k \text{ times})) = C$ (from Eq. 3).

Substituting the contracts for Vegas and Swift, and fold as \sum , we get the throughput ratios of k and k^2 respectively. Steeper contracts imply worse fairness.

Congestion growth. We study how congestion grows with (decreasing) fair share or (increasing) flow count. Consider a dumbbell topology with capacity C and n flows. For each flow to get its fair share of $r_i = C/n$, the statistic needs to satisfy: $\forall i, r_i = C/n = \text{func}(s_i)$. This implies that $\text{func}^{-1}(C/n) = s_n = s_i$. We define worst-case growth (i.e., (s for n flows)/(s for 1 flow)) as: $\text{growth}(n) = \max_C \frac{\text{func}^{-1}(C/n)}{\text{func}^{-1}(C)}$ (7). For Vegas and Swift, $\text{growth}(n)$ is n and \sqrt{n} respectively. Steeper contracts imply slower growth.

Tradeoff summary. Gradual contracts give better robustness and fairness, while steeper contracts give better congestion and generality. In Appendix A, we mathematically show these tradeoffs by considering each pair of contending metrics. E.g., if we want error factor $\leq \epsilon_r$ for δs noise, then $\text{growth}(n)$ is $\Omega\left(\delta s \frac{\log(n)}{\log(\epsilon_r)}\right)$. These bounds are tight, i.e., we show CCAs that meet these bounds (see below).

The tradeoffs involving fairness depend on fold . We consider $\text{fold} \in \{\sum, \max, \min\}$. The tradeoff exists for $\text{fold} = \sum$, e.g., delay, loss rate, and ECN marking rate (assuming loss and ECN marking rates are small enough). There is no tradeoff for \max and \min , e.g., max per-hop delay, as this allows independently ensuring max-min fairness. So, for CCAs like Poseidon [64], the only tradeoff is between robustness versus congestion and generality.

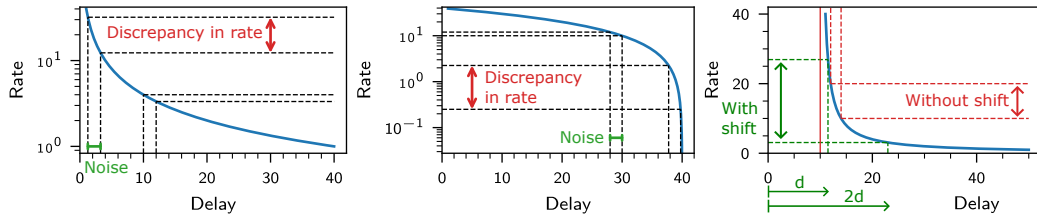
There are exactly two corner points in this tradeoff space. If we fix desired robustness, e.g., we want to tolerate δs noise, then the Exponential CCA [8] gives the best generality and congestion growth, but it has poor fairness (Exp. in Table 3). If we fix desired fairness, e.g., proportional fairness ($\alpha = 1$), then Vegas gives the best generality and congestion growth (for $\text{fold} = \sum$). In general, if we want α -fairness, then the contract “ $\text{func}(s) = \frac{1}{\sqrt[s]{s}}$ ” gives the best generality and congestion growth (for $\text{fold} = \sum$). We mathematically show that these are corner points in Appendix A.

4.1 Guidance on picking contract (D1)

(D1.1) Input/output. Unless one chooses a workaround to the tradeoffs (§ 8), the choice of contract input is one of delay, loss rate, or ECN marking rate. This choice depends on buffer sizes, availability of ECN, and constraints on dynamics (§ 6). Similarly, outside of the workaround of coordinating fraction of link use (§ 8), the output of the contract is rate.

All the workarounds in § 8 either require in-network support (e.g., max per-hop-delay signal) or fundamental congestion control research. For instance, [8] shows that to ensure robustness and full generality (arbitrarily large link rates), CCAs must deliberately vary rate and create delay variations. This goes against our desire to have a stable application-level rate. In fact, RL-based CCAs [42] and Remy [68] update cwnd based on moving averages of historical signals. CCAs that are deterministic responses to such histories are incapable of creating deliberate rate/delay variations.

(D1.2) Shape. For the above input/output choices, the tradeoffs exist. One should decide which of the two corner points in the tradeoff space is preferred. The tradeoff point fixes the contract shape (e.g., $\text{rate} = 1/s$) and the asymptotics in the tradeoffs. Next, we want



■ **Figure 9 (Left)** Vegas’s contract ($\text{rate} = 1/\text{delay}$). The same noise creates larger discrepancies with increasing link capacities. **(Middle)** Astraea’s contract ($\text{rate} = S_0 - \text{delay}$) has an X-intercept. Discrepancies increase as delay approaches the intercept. **(Right)** Shifting the contract changes the fixed-point delays and rates on the parking lot topology. Solid red line shows Y-Axis without any shift in contract. The dotted lines show the rates of two flows at steady-state (red shows without shift and green shows with shift). Their delays are $2\times$ apart (i.e., d and $2d$). We omit units (e.g., Mbps or ms) here as the shape/shift/intercept matter not the scale.

to tune the contract for deployment, i.e., picking parameters that control shift or scaling (e.g., a, b, c in $\text{rate} = a/(s - b) + c$). The scale (e.g., a) affects the constant factors in the tradeoffs, and shifts/clamps (e.g., b, c) should be avoided.

(D1.3) Shift/clamps. Shifting the contract right (positive b) changes the fairness properties (§ 5). Shifting left (negative b) restricts the range of link capacities by introducing a Y-intercept. Shifting down (negative c) introduces an X-intercept which severely degrades robustness and fairness (§ 5). Shifting up (positive c) only makes sense if there is a known lower bound on the fair share of a flow (e.g., fair share ≥ 512 Kbps). Clamping (e.g., $s = \min(a, 1/\text{rate})$) has similar consequences.

(D1.3) Scale. We discuss the scaling choice for the two corner contracts. We use C_0 and S_0 respectively to represent scaling of rate and s , e.g., $\text{rate} = C_0 e^{-s/S_0}$. One can set these parameters to meet a desired value for one of the four metrics. The tradeoffs decide the other metrics.

The general form of the exponential contract is $\text{rate} = C_0 e^{-s/S_0}$. C_0 is the maximum rate at which the CCA can ever send. If the link capacity is higher, the CCA would under-utilize the link. S_0 controls the robustness error (throughput ratio) for δs error in $s = C_{\max} e^{-s/S_0} / C_{\max} e^{-(s+\delta s)/S_0} = e^{\delta s/S_0}$. Larger S_0 gives better robustness but worse congestion. These parameters also decide the range of delays (S_{\min}, S_{\max}) produced by the CCA for a deployment’s range of fair shares (C_{\min}, C_{\max}).

The general form of the α -fair contract is $\text{rate} = C_0 S_0^\alpha / s^\alpha$, or $\text{rate} = C_0 S_0 / s$ for $\alpha = 1$. The contract passes through $\langle S_0, C_0 \rangle$: S_0 is the delay maintained when the fair share is C_0 . The deployment’s range of fair shares (C_{\min}, C_{\max}) implicitly decides the range of congestion (S_{\min}, S_{\max}), and robustness to noise. For small capacities ($C_{\min} \rightarrow 0$), S_{\max} is ∞ . For large capacities ($C_{\max} \rightarrow \infty$), S_{\min} is 0 and the error factor is ∞ (as $\delta s / S_{\min} \rightarrow \infty$). When C_{\max} is known (e.g., a data center), and one desires to maintain a minimum delay of S_{\min} (to ensure utilization despite variations [42], or bound error factor by “ $1 + \delta s / S_{\min}$ ”), an example contract choice is $\text{rate} = C_{\max} S_{\min} / \text{delay}$.

The exponential contract ensures a finite error factor at the cost of not supporting arbitrarily large link capacities. The α -fair contract supports arbitrarily large capacities but cannot bound the error factor. This is the same as the fundamental tradeoff in [8].

5 Learnings from contracts

Avoid extreme shape and intercepts. Astraea (linear contract) and ICC (logarithmic contract) exhibit poor robustness and fairness due to extreme shapes and X-intercepts. The issue is that all the low rates (e.g., 0.1 to 1 Mbps) map to delays near the X-intercept (e.g., 40 ms) (Figure 9 middle). Consequently, small delay jitter creates large discrepancies in inferring fair shares. Similarly, on parking lot (Figure 8), if the hop delays are close to the intercept delay, the long flow (f_0) observes their sum, which exceeds the intercept delay. Attempting to reduce this “excessive delay”, f_0 reduces its rate to zero and starves.

Figure 11 demonstrates these empirically. Ambient emulation noise causes Astraea to starve flows when the fair share is low (4 flows on a dumbbell topology with 10 Mbps capacity, 30 ms RTprop, and infinite buffer). ICC starves the long flow (f_0) on a parking lot topology with 3 flows (2 hops) and the same network parameters. Importantly, contract analysis reveals these performance issues without running the CCAs or understanding their internals.

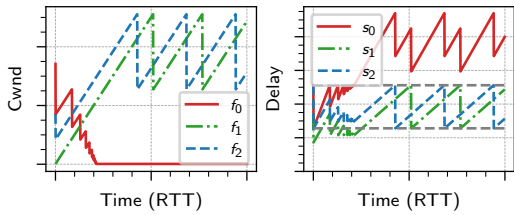
Note, unlike [8], the starvation here is not fundamental and is easily avoided by removing the X-intercept (e.g., $\text{rate} = 1/\text{delay}$). This increases the delay at low fair shares, however, this is unavoidable as transmission delays anyway grow as $1/C$ with decreasing capacity. In contrast, the issue in [8] is fundamental and is caused by the asymptote on the Y-axis, where all the high fair shares (e.g., > 100 Mbps) map to near-zero delays (Figure 9 left). This asymptote is hard to remove while supporting arbitrarily large rates with a monotonically decreasing contract.

Avoid shifting contracts. CCAs like Swift [36] define target delay as: $\text{target_delay} = b + 1/\text{rate}$. In such formulations, it may seem convenient to shift the contract (set a positive b) to maintain a minimum delay for high utilization. However, shifting changes the steady-state fixed-point causing undesired unfairness (Figure 9 right). For instance, in the parking lot topology with 2 hops, the short and long flows see delays of d and $2d$ respectively. These delays are such that the corresponding contract rates add up to the link capacity. Shifting the contract rightward changes these delays in a way that increases discrepancies in the rates. We show this empirically for BBR in § 7. To avoid this issue, one should tune the scaling (§ 4) to maintain a minimum delay instead.

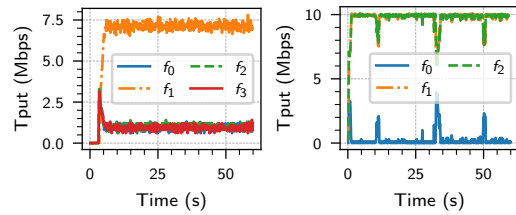
While we described shifting in the context of delay-based contracts, its impact more generally depends on the aggregate statistic and how it accumulates over hops (i.e., the `fold` function). The effect of such shifts can be analyzed using the parking lot steady-state equations (Eq. 3, Eq. 4, Eq. 5). For example, with Poseidon’s max per-hop delay statistic, shifting the contract does not introduce unfairness: in the parking lot topology, both short and long flows observe the same max per-hop delay even after shifting.

A related but distinct mechanism arises with active queue management (AQM), such as RED-based ECN marking [20, 57]. Here, it is possible to *effectively* shift queue buildup (and hence delay) without shifting either the contract or the aggregate statistic (i.e., ECN marking probabilities). Specifically, under fixed network parameters and topology, increasing the RED thresholds K_{\min} and K_{\max} raises the steady-state queuing while leaving ECN marking probabilities – and resulting rate allocations – unchanged. This behavior follows from jointly solving the parking lot steady-state equations and the ECN marking probability equation.

Avoid fixed thresholds (e.g., delay targets) at end-hosts. We illustrate how AIMD on delay (i.e., MD when delay crosses a fixed threshold and AI otherwise) causes starvation. AIMD on delay appears in 1RMA [59] and the CCAs (SMaRTT [11] and STrack [38]) proposed for standardization in the Ultra Ethernet Consortium [63]. It was also proposed by [8] to work around their impossibility result.



■ **Figure 10** AIMD on delay starves the long flow (f_0) on parking lot topology.



■ **Figure 11 (Left)** Astraea starves under ambient noise. **(Right)** ICC starves on parking lot.

On a dumbbell topology, AIMD on delay works fine. It induces a Reno-like contract: “ $r = 1/\sqrt{s}$ ”, where s is the high-delay probability (inverse of bytes between high-delay events) instead of loss probability. However, it creates starvation with multiple bottlenecks. All hosts attempt to maintain delay around the same (fixed) target value, but when flows observe delays from different combinations of hops, they cannot simultaneously maintain the same delay. Figure 10 shows this on a 2-hop parking lot. The short flows (f_1 and f_2) oscillate delay between the delay threshold and half the threshold (shown in gray). The long flow (f_0) observes their sum, which always exceeds the threshold. Consequently, f_0 never increments cwnd and starves.

We can resolve this by either (1) not using a fixed threshold, e.g., Swift scales the (target delay) threshold with rate, which transforms the contract to use delay instead of high-delay probability, allowing different flows to maintain different delays, or (2) moving the threshold to links instead of end-hosts, e.g., packet drops (Reno) and ECN marks (DCTCP) occur when delay crosses a threshold at the links.

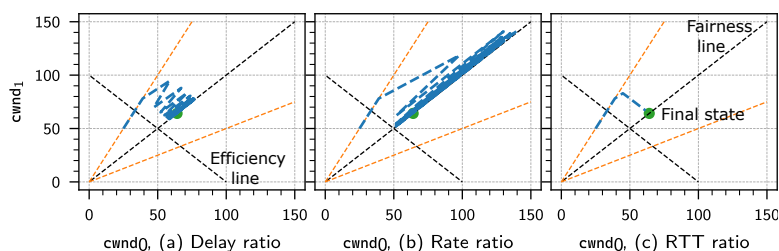
This issue may arise for any CCA that uses fixed end-to-end thresholds to decide when to increase or decrease their rate, e.g., BBRv3 uses `BBRLossThresh = 2%` [15].

Minor changes in CCA change the contract and consequently steady-state performance. MPRDMA [46], DCTCP [4], and Reno [31] all perform AIMD on binary feedback. In fact, [28] uses MPRDMA as an approximation of DCTCP. However, the minor differences in their design result in different contracts: Reno: “ $r = 1/\sqrt{s}$ ”, DCTCP: “ $r = 1/s^2$ ”, MPRDMA: “ $r = 1/s$ ”. Similarly, we found an algebraic mistake in the Linux implementation of TCP Vegas that changes its contract from “ $r = 1/\text{delay}$ ” to “ $r = 1/\text{delay}^2$ ” when `RTprop` is small. We have confirmed this with the maintainers. This bug has existed for 17 years due to a refactoring commit [40]. Such bugs may be caught immediately if CCA implementations explicitly delineate contracts.

6 Canonical CCA dynamics (D2)

We discuss how to best implement a CCA (cwnd or rate updates) to follow a contract. We can implement updates using either rate or cwnd, regardless of whether the contract is based on rate or cwnd. We discuss cwnd-based updates to follow a rate-based contract. Other combinations are similar.

While the contract fixes steady-state performance, the updates determine dynamics (stability and convergence time). Existing CCAs correspond to various ways of implementing cwnd updates to follow a contract. Vegas uses AIAD, Swift uses AIMD, Posiedon uses MIMD (using rate ratio), FAST uses (a different kind of) MIMD (using RTT ratio), and Copa uses AIAD with increasing gains when cwnd updates occur in the same direction.



■ **Figure 12** MIMD using RTT ratio is more stable than rate or delay ratio. Dynamics stay above the efficiency line showing no throughput loss, contrary to PowerTCP’s claim on voltage-based CCAs [1].

We argue that TCP FAST’s method is the best way to implement a contract. AIMD/AIAD are sub-optimal in convergence time. Copa uses convex – instead of concave (e.g., ETC [26]) – changes to cwnd, which creates overshoots and instability. Ideally, we want to dampen changes to cwnd when it is already close to convergence [26]. Poseidon’s MIMD creates instability depending on network parameters (see below). FAST’s update is stable and converges exponentially fast to both efficiency and fairness. We illustrate (Figure 12) the issue with Poseidon’s MIMD and how FAST’s MIMD fixes the issue in the context of delay, and later discuss other congestion signals.

MIMD using rate or delay ratio (C.f. Poseidon). We can interpret a contract (e.g., $\text{rate} = 1/\text{delay}^2$), in two ways: (1) *target rate* given *current delay* ($\text{target_rate} = 1/\text{current_delay}^2$), or (2) “*target delay* as a function of *current rate*” ($\text{target_delay} = 1/\sqrt{\text{current_rate}}$). Where, $\text{current_rate} = \text{cwnd}/\text{RTT}$, and $\text{current_delay} = \text{RTT} - \min \text{RTT}$. These yield two natural MIMD updates to implement a contract:

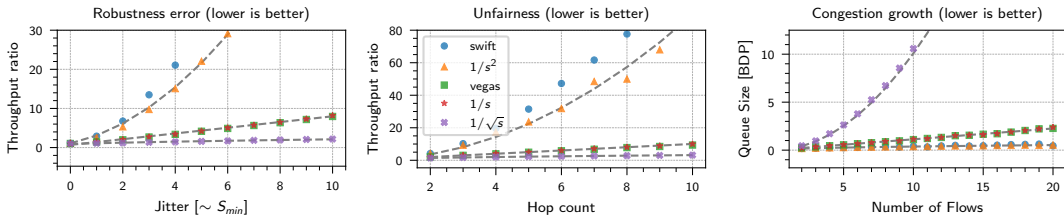
$$\text{target_cwnd} \leftarrow \text{cwnd} \frac{\text{target_rate}}{\text{current_rate}} \text{ or } \text{cwnd} \frac{\text{target_delay}}{\text{current_delay}}$$

Where cwnd moves towards the target “ $\text{next_cwnd} \leftarrow (1 - \alpha) \cdot \text{cwnd} + \alpha \cdot \text{target_cwnd}$ ”, with optional clamps bounding the change “ $\text{cwnd} \leftarrow \text{clamp}(\text{next_cwnd}, \beta_{\min} \text{cwnd}, \beta_{\max} \text{cwnd})$ ”. Outside of the contract “ $\text{rate} = 1/\text{delay}$ ”, rate and delay ratios are different yielding different dynamics. For each α , clamp choice, both updates are only stable for specific network parameters (e.g., capacity, flow count).

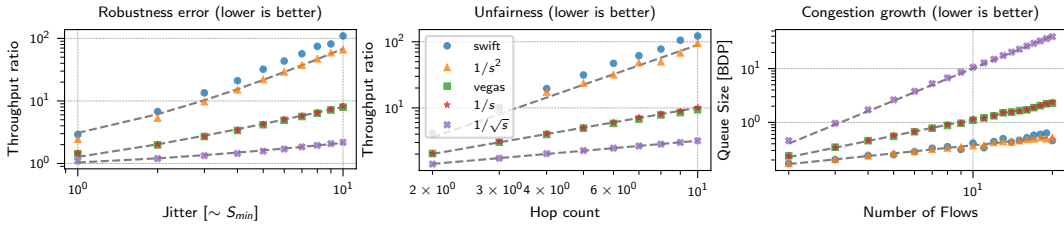
MIMD using RTT ratio (C.f. FAST). To ensure stability, we want to consider how rate and delay affect or are affected by changes in cwnd. The following update achieves stability without requiring any averaging (α), clamps, or assumptions on network parameters:

$$\text{cwnd} \leftarrow \text{cwnd} \frac{\text{target_RTT}}{\text{current_RTT}} = \text{cwnd} \frac{\min \text{RTT} + \text{target_delay}}{\min \text{RTT} + \text{current_delay}}$$

Intuitively, cwnds map to “packets in queue + packets in pipe”. We only want to move the “packets in queue” part from the current delay to the target delay. This cwnd update isolates the term responsible for queuing delay and only scales that term using the delay ratio. We can also interpret the update as: $\text{cwnd} \leftarrow \text{rate} * \text{target_RTT} = \text{cwnd}/\text{RTT} * (\min \text{RTT} + \text{target_delay})$.



■ **Figure 13** Contracts-based performance estimates match packet-level simulation. The markers show empirical data and gray lines show performance estimated by contracts. Figure 14 shows log-log scale. Note: markers may be hard to see because they overlap.



■ **Figure 14** Figure 13 on log-log scale for visual clarity.

Other congestion signals. We discussed optimal dynamics for delay-based contracts. For other signals, e.g., ECN or loss, the best cwnd update may differ. The current and target RTT in the update depends on the relation between RTT and aggregate statistic. For instance, for RED-based ECN, given `target_ECN_rate`, the `target_RTT` is “ $\min RTT + (1/C) \cdot (K_{min} + (K_{max} - K_{min}) * \text{target_ECN_rate})$ ”. We derive this by inverting the mapping from queue size (and hence queueing delay) to ECN marking probability. This works directly when ground truth link capacity “ C ” is known (e.g., a data center), otherwise either the capacity needs to be estimated or one needs to use delay or rate ratio with a value of α tuned for the range of network parameters one wants to support.

TFRC [21] explores cwnd update for loss-based contracts. The challenge is that while a constant cwnd creates a constant delay, it may not create a constant loss rate (e.g., when the loss rate is less than one loss per window). Thus, we may need cwnd variations to maintain a persistent loss rate even when the target and current loss rates are equal.

Other design considerations. A complete the CCA needs other decisions including: (O1) how to aggregate multiple statistic samples, (O2) how long to measure the samples, (O3) time between cwnd updates, and (O4) how to compute any other estimates (e.g., bandwidth or RTprop estimate). For instance, Copa computes standing RTT by taking the “minimum” (O1) over RTT samples in the last “half `srtt`” (standard smoothed RTT) (O2). It updates cwnd “every ACK” (O3) and uses minimum RTT over last 10 seconds to estimate RTprop (O4). Alternatively, one can estimate RTprop using BBR’s RTT probes. Implementations may also require other features like using rate instead of cwnd when the BDP is less than 1 packet [36]. Such details are orthogonal to contracts.

7 Empirical validation

We empirically validate the trends in metrics and tradeoffs predicted by contracts. For visual clarity in plots, we only show a handful of contracts/CCAs. This section complements the empirical results in § 3, § 4, and § 5; where we already showed performance issues and contracts for a large set of CCAs including Sprout, PCC, Indigo, ICC, Astraea, AIMD on delay (1RMA, SMaRTT, STrack, etc.).

Methodology. We use packet-level simulation and emulation. Simulations allow controlling noise and isolating one source of tradeoff at a time where emulation always has ambient noise, e.g., due to OS scheduling jitter. We use `htsim` [28] for simulation (also used in MPTCP [55], NDP [27], EQDS [52]), and `mahimahi` [51], `Pantheon` [70], and `mininet` [17] for emulation.

In simulation only, for two reasons, we give oracular knowledge of RTprop to all CCAs. First, CCAs like Swift target data center deployments where RTprop may be known. For consistency, we provide RTprop to all CCAs. Second, CCAs like Vegas do not explicitly drain queues resulting in misestimating RTprop and poor performance. We remove this source of poor performance as this is not fundamental unlike the tradeoffs imposed by contracts. For instance, Copa and BBR explicitly drain queues to estimate RTprop accurately (at least in the absence of noise).

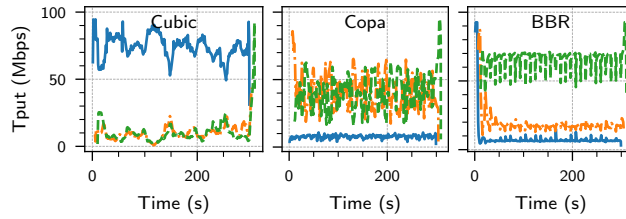
Note that we are validating negative results (i.e., tradeoffs). Simplifications only make the validation stronger. If tradeoffs exist with oracular knowledge of RTprop, then performance is only worse without it, e.g., under-estimation causes under-utilization, and over-estimation increases congestion.

Simulation CCAs. We implement and test Swift [36] and Vegas [12]. For reference, we also show 3 canonical (§ 6) contract implementations: $1/\sqrt{s}$, $1/s$, and $1/s^2$, where $s = \text{delay} = \text{RTT} - \text{RTprop}$. These avoid RTT-bias unlike vanilla Vegas/Swift, and use MIMD instead of AIAD/AIMD. In the canonical CCAs, we update `cwnd` every 2 RTTs and aggregate delay as the minimum over delay samples since the last `cwnd` update. To isolate the impact of contract shape, we also tune the scale parameters symmetrically. All CCAs have the same $S_{\min} = 1.2 \mu s = 0.1 \text{ RTprop}$ and $C_{\max} = 100 \text{ Gbps} = \text{link capacity}$. S_{\max} is then decided by C_{\min} and the contract shape.

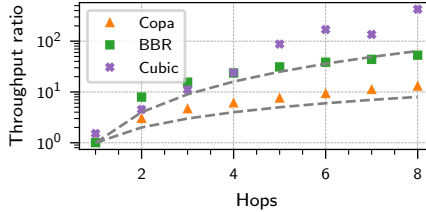
Simulation scenarios. We set link capacity = 100 Gbps, RTprop = 12 μs , packet size = 4 KB. These are default parameters in `htsim` for data-center deployments. In § 4 we defined the metrics in to be unitless and the tradeoffs we showed exist for all choices of network parameters (link capacity, RTprop, etc.). Consequently, the specific parameter values are of little importance and we could have used any other values as well. We set the buffer size to be infinite to remove any effects of packet losses, since we use delay-based CCAs. To measure robustness error, we use a dumbbell topology with 2 flows where one of them witnesses noise. We inject controlled error by adding a hop that persistently delays packets by “ δs ” μs , and vary δs . We do not include this in the RTprop provided to the CCAs. We inject noise this way to show trends. In emulation, we show the impact of realistic noise. To measure unfairness and congestion growth, we instantiate parking lot (with varying hops) and dumbbell (with varying flow count) topologies respectively. For the CCAs we test, generality is just the inverse of congestion growth (Table 3), so we do not show generality.

Simulation results. Empirical performance matches that estimated by contracts (Figure 13). Swift’s RTT-bias causes slightly worse fairness and robustness than the equivalent “ $1/s^2$ ” canonical CCA which removes the bias. Vanilla Vegas has RTprop-bias (instead of RTT-bias). Since the RTprop is the same for all flows, the steady-state performance of Vegas is the same as the canonical “ $1/s$ ” CCA.

Emulation CCAs, scenarios, and results. We run Cubic [25], BBR [14] (Linux kernel v5.15.0) and Copa [9, 7]. The empirical contract derivations (§ 3) already showed generality and congestion growth, e.g., increase in delay or loss rate with decreasing fair share (increasing



■ **Figure 15 Robustness error.** Cubic starves the orange/green flows that do not witness jitter. Copa starves the blue flow that witnesses jitter. BBR starves the blue flow with the smallest RTprop.



■ **Figure 16 Unfairness.** Grey lines show $y = x$ and $y = x^2$. Copa matches the unfairness predicted by its contract. BBR and Cubic are worse due to shift and RTT-bias.

flow count). In Figure 15 and Figure 16, we show robustness and fairness. We run flows for 5 mins on dumbbell and parking lot topologies with capacity = 100 Mbps, buffer = 1 BDP, and describe RTprop and flow count below. Emulation does not scale to data-center link speeds (unlike htsim). Here, our parameter choices align with Internet deployments. Again, the tradeoffs are independent of the absolute parameter values and we get qualitatively similar results with other values.

Note that noise in the raw signal may not create an equivalent amount of error in the statistic used by the CCA. Hence, we do not see a persistent trend in robustness error with varying noise. To validate that robustness is an issue, we show that the CCAs incur large throughput ratios (starvation) with small delay jitter. We inject jitter in two ways: (J1) slightly different RTprops (3 flows with RTprop of 10, 20, and 30 ms), and (J2) ACK-aggregation (3 flows with RTprop of 32 ms but 1 flow additionally witnesses 32 ms of ACK aggregation). We emulate ACK-aggregation in the same way as Pantheon [70]. In Figure 15, Cubic and Copa show starvation with J2 and BBR shows starvation with J1. Note, BBR’s unfairness in J1 is different from RTT-unfairness in traditional CCAs [31, 25]. For BBR, a small difference in RTprops leads to large unfairness that increases with the link rate [8].

Figure 16 shows unfairness on parking lot with 5 ms RTprop. Copa matches the trend estimated by the contract. With BBR, the shape (derivative) of the contract is same as Copa. However the shift and RTT-bias in BBR’s contract causes worse unfairness. For Cubic, the throughput ratio should be at least $\text{hops}^{4/3}$ (contract is $\text{rate} = \text{loss rate}^{-0.75}$ [39]). Reality is worse due to RTT-bias.

8 Working around the tradeoffs

As discussed in § 2, we believe the only way to work around the tradeoffs is to pick the input/output of the contract in a way that decouples physical quantities (e.g., rate or delay) from the contract. We show this for the four metrics.

Note that compound contract functions that take different shapes on different link rates or switch the shape on the fly do not alleviate the tradeoffs. In the worst-case, all the scenarios may occur simultaneously, e.g., multiple flows per hop on a parking lot topology with noise.

- **Fairness.** As mentioned in § 4, statistics that accumulate using max or min, like max per-hop delay, decouple multi-bottleneck fairness, trivially ensuring max-min fairness. However, such accumulation often relies on in-network support [64].
- **Congestion and robustness.** The congestion growth metric describes growth in the statistic and not congestion. Decoupling statistic and congestion allows independently bounding congestion. For instance, [72] shows use of a PI controller to obtain different ECN marking probabilities for the same queue buildup (i.e., different statistic for the same congestion). Likewise, explicit communication in packet headers (using enough precision) may eliminate noise to meet robustness [34].
- **Generality.** The domain of the statistic limits generality. Existing CCAs encode fair shares using a “unary” encoding. As proposed in [8], we can improve encoding efficiency using a “binary” encoding that communicates fair shares over time – similar to deriving multi-bit feedback from single ECN bit [24]. Another work around is coordinate the “fraction of link use” (i.e., a quantity between 0 and 1) instead of “absolute fair shares” (i.e., an arbitrarily large number). This reduces the range of output values that a contract needs to support. BBR’s rate-limited mode does this [23], but BBR often operates in cwnd-limited mode [65], without fully leveraging this workaround.

9 Limitations and future work

Improving expressivity. In defining contracts, we faced a tradeoff between expressivity (breadth of statements we can make) and tractability (mathematically backing the statements). We erred on the side of tractability to mathematically derive tradeoffs in Appendix A. For instance, our current definition makes assumptions about the network and other flows. As a result, we are unable to reason about inter-CCA fairness. Due to similar reasons, we also found it hard to prove/disprove that contracts are necessary or sufficient for fairness. Below, we describe the challenges, benefits, and possible approaches to improve expressivity.

Challenges. Two CCAs may achieve fairness even when they have different contracts. Consider Copa, it employs a delay-based contract when competing with itself, but it switches to emulating Reno when it detects Reno flows. Even CCAs like Vegas, that employ a delay-based contract all the time, may compete fairly with Reno depending on the network conditions [43]. For instance, RED-based packet drops [20] create a mapping between queuing delay and loss rates. If the loss-rate-based and delay-based fair shares match, then Vegas and Reno may compete fairly.

Benefits. Improving expressivity can further guide CCA design. For instance, to be TCP-friendly, BBRv3 [15] leaves “headroom” for loss-based CCAs. If contracts are necessary for fairness, then it is better to explicitly follow Reno’s contract on detecting competing Reno flows than leaving headroom which may or may not cause BBRv3 to follow Reno’s contract.

Approach. We hope to extend reasoning using the formal methods literature that also uses contract-like abstractions to reason about distributed algorithms. For instance, [6] defines contracts as a *set* of traces described using ω -regular grammars, while we defined contracts as a *function*.

Extending contracts and blueprints to cover more scenarios, metrics, and CCA design nuances. Our current blueprints do not capture all scenarios that CCAs must handle, e.g., short flows, incast, micro-bursts [58, 32], frequent flow arrivals and departures, or bursty traffic. In such settings, flows may never reach steady state; nevertheless, we still desire fairness guarantees for long-lived flows. Similarly, we do not model all metrics that may be

relevant to applications, such as flow or co-flow completion times. Supporting these scenarios and metrics often requires complementary mechanisms beyond contracts, including slow start, receiver- or credit-based control, and flow scheduling.

Moreover, while we showed that contracts determine key aspects of CCA performance, they are not all-encompassing. Other design nuances can also significantly impact performance. For instance, CCAs periodically explore latent network parameters and smooth or filter noise from raw congestion signals. Although these choices influence the observed aggregate statistics and thus appear in the resulting contract, the contract itself does not directly tell us how such mechanisms should be designed. Nevertheless, we believe that the lens of contracts can help designers reason about these choices indirectly, by clarifying their impact on steady-state behavior and tradeoffs.

In the future, we envision extending our blueprints into a living repository that incorporates these evolving scenarios, metrics, and design nuances, providing a comprehensive view of CCA behavior. Outside CCA design and analysis, we believe a contracts-first approach would also improve reverse-engineering and classification of CCAs in the wild [22, 19, 66, 48].

Network measurement. Measurement of workloads and network environments can guide which metrics to prioritize in the tradeoff space. For instance, assessing fairness requires understanding how often flows experience multi-hop congestion and the typical number of flows or hops involved. Recent work [13] suggests that contention may be rare on the Internet, implying fairness may be less critical in some contexts. For robustness, it would be useful to quantify the size and frequency of non-congestive delays and losses. On the workload front, we want to understand which network-level metrics correlate with application-level metrics. For instance, [56] shows that unfairness is better for AI collectives.

10 Conclusion

We showed that contracts determine key performance metrics, resulting in tradeoffs. We identify pitfalls to avoid when designing CCAs. We hope that with our work, contracts will be a conscientious design choice rather than an afterthought. Contracts should be a direct consequence of desired steady-state performance, and rate updates should be a consequence of desired reactivity/convergence time.

References

- 1 Vamsi Addanki, Oliver Michel, and Stefan Schmid. PowerTCP: Pushing the performance limits of datacenter networks. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 51–70, Renton, WA, April 2022. USENIX Association. URL: <https://www.usenix.org/conference/nsdi22/presentation/addanki>.
- 2 Anup Agarwal, Venkat Arun, Devdeep Ray, Ruben Martins, and Srinivasan Seshan. Towards provably performant congestion control. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 951–978, Santa Clara, CA, April 2024. USENIX Association. URL: <https://www.usenix.org/conference/nsdi24/presentation/agarwal-anup>.
- 3 Anup Agarwal, Venkat Arun, and Srinivasan Seshan. Frcc: Towards provably fair and robust congestion control. In *23rd USENIX Symposium on Networked Systems Design and Implementation (NSDI 26)*, Renton, WA, May 2026. USENIX Association.
- 4 Mohammad Alizadeh, Albert Greenberg, David A. Maltz, Jitendra Padhye, Parveen Patel, Balaji Prabhakar, Sudipta Sengupta, and Murari Sridharan. Data center tcp (dctcp). In *Proceedings of the ACM SIGCOMM 2010 Conference, SIGCOMM '10*, pages 63–74, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1851182.1851192.

- 5 Mohammad Alizadeh, Adel Javanmard, and Balaji Prabhakar. Analysis of dctcp: stability, convergence, and fairness. In *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '11, pages 73–84, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993744.1993753.
- 6 Ashwani Anand, Anne-Kathrin Schmuck, and Satya Prakash Nayak. Contract-based distributed logical controller synthesis. In *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control*, HSCC '24, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3641513.3650123.
- 7 Venkat Arun. venkatarun95/genericcc. [Online; accessed 2025-04-17]. URL: <https://github.com/venkatarun95/genericCC>.
- 8 Venkat Arun, Mohammad Alizadeh, and Hari Balakrishnan. Starvation in end-to-end congestion control. In *Proceedings of the 2022 ACM SIGCOMM 2022 Conference*, SIGCOMM '22, Amsterdam, Netherlands, 2022. Association for Computing Machinery.
- 9 Venkat Arun and Hari Balakrishnan. Copa: Practical Delay-Based congestion control for the internet. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 329–342, Renton, WA, April 2018. USENIX Association. URL: <https://www.usenix.org/conference/nsdi18/presentation/arun>.
- 10 D. Bansal and H. Balakrishnan. Binomial congestion control algorithms. In *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213)*, volume 2, pages 631–640 vol.2, 2001. doi:10.1109/INFCOM.2001.916251.
- 11 Tommaso Bonato, Abdul Kabbani, Daniele De Sensi, Rong Pan, Yanfang Le, Costin Raiciu, Mark Handley, Timo Schneider, Nils Blach, Ahmad Ghalayini, Daniel Alves, Michael Papamichael, Adrian Caulfield, and Torsten Hoeffler. Fastflow: Flexible adaptive congestion control for high-performance datacenters, 2024. arXiv:2404.01630.
- 12 L.S. Brakmo and L.L. Peterson. Tcp vegas: end to end congestion avoidance on a global internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–1480, 1995. doi:10.1109/49.464716.
- 13 Lloyd Brown, Yash Kothari, Akshay Narayan, Arvind Krishnamurthy, Aurojit Panda, Justine Sherry, and Scott Shenker. How i learned to stop worrying about cca contention. In *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*, HotNets '23, pages 229–237, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3626111.3628204.
- 14 Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. Bbr: Congestion-based congestion control. *ACM Queue*, 14, September-October:20–53, 2016. doi:10.1145/3012426.3022184.
- 15 Neal Cardwell, Ian Swett, and Joseph Beshay. BBR Congestion Control. Internet-Draft draft-ietf-ccwg-bbr-01, Internet Engineering Task Force, Work in Progress. URL: <https://datatracker.ietf.org/doc/draft-ietf-ccwg-bbr/01/>.
- 16 Dah-Ming Chiu and Raj Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, 17(1):1–14, 1989. doi:10.1016/0169-7552(89)90019-6.
- 17 Mininet Project Contributors. Mininet: An instant virtual network on your laptop (or other pc) - mininet. [Online; accessed 2025-04-17]. URL: <https://mininet.org/>.
- 18 Mo Dong, Qingxi Li, Doron Zarchy, P Brighten Godfrey, and Michael Schapira. Pcc: Re-architecting congestion control for consistent high performance. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 395–408, 2015.
- 19 Margarida Ferreira, Ranysha Ware, Yash Kothari, Inês Lynce, Ruben Martins, Akshay Narayan, and Justine Sherry. Reverse-engineering congestion control algorithm behavior. In *Proceedings of the 2024 ACM on Internet Measurement Conference*, IMC '24, pages 401–414, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3646547.3688443.

- 20 S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, 1993. doi:10.1109/90.251892.
- 21 Sally Floyd, Mark Handley, Jitendra Padhye, and Jörg Widmer. Equation-based congestion control for unicast applications. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '00, pages 43–56, New York, NY, USA, 2000. Association for Computing Machinery. doi:10.1145/347059.347397.
- 22 Sishuai Gong, Usama Naseer, and Theophilus A Benson. Inspector gadget: A framework for inferring tcp congestion control algorithms and protocol configurations. In *Network Traffic Measurement and Analysis Conference*, 2020.
- 23 Google. bbr/Documentation/startup/gain/analysis/bbr_drain_gain.pdf at master · google/bbr, November 2024. [Online; accessed 14. Nov. 2024]. URL: https://github.com/google/bbr/blob/master/Documentation/bbr_bandwidth_based_convergence.pdf.
- 24 Prateesh Goyal, Anup Agarwal, Ravi Netravali, Mohammad Alizadeh, and Hari Balakrishnan. ABC: A simple explicit congestion controller for wireless networks. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 353–372, Santa Clara, CA, February 2020. USENIX Association. URL: <https://www.usenix.org/conference/nsdi20/presentation/goyal>.
- 25 Sangtae Ha, Injong Rhee, and Lisong Xu. Cubic: A new tcp-friendly high-speed tcp variant. *SIGOPS Oper. Syst. Rev.*, 42(5):64–74, 2008. doi:10.1145/1400097.1400105.
- 26 Feixue Han, Qing Li, Peng Zhang, Gareth Tyson, Yong Jiang, Mingwei Xu, Yulong Lan, and ZhiCheng Li. ETC: An elastic transmission control using End-to-End available bandwidth perception. In *2024 USENIX Annual Technical Conference (USENIX ATC 24)*, pages 265–284, Santa Clara, CA, July 2024. USENIX Association. URL: <https://www.usenix.org/conference/atc24/presentation/han>.
- 27 Mark Handley, Costin Raiciu, Alexandru Agache, Andrei Voinescu, Andrew W. Moore, Gianni Antichi, and Marcin Wójcik. Re-architecting datacenter networks and stacks for low latency and high performance. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '17, pages 29–42, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3098822.3098825.
- 28 Mark Handley, Costin Raiciu, Chih-Yuan Chang, and Mihai Brodschi. csg-htsim, November 2024. [Online; accessed 18. Nov. 2024]. URL: <https://github.com/Broadcom/csg-htsim>.
- 29 Stephen Hemminger and David S. Miller. [tcp]: Add tcp vegas congestion control module. · torvalds/linux@b87d856, June 2005. [Online; accessed 2025-06-02]. URL: <https://github.com/torvalds/linux/commit/b87d8561d8667d221b728ccdc18eb95b16a687b>.
- 30 Tom Henderson, Sally Floyd, Andrei Gurtov, and Yoshifumi Nishida. RFC 6582: The NewReno Modification to TCP’s Fast Recovery Algorithm, April 2012. [Online; accessed 20. Nov. 2024]. URL: <https://datatracker.ietf.org/doc/rfc6582>.
- 31 Janey C. Hoe. Improving the start-up behavior of a congestion control scheme for tcp. In *Conference Proceedings on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '96, pages 270–280, New York, NY, USA, 1996. Association for Computing Machinery. doi:10.1145/248156.248180.
- 32 Mohammad Hosseini, Sina Darabi, Hannaneh B. Pasandi, Mohammad Nakhjiri, and Patrick Eugster. Poison comes in small packages: Application-driven reexamination of datacenter microbursts. *Proc. ACM Meas. Anal. Comput. Syst.*, 9(2), 2025. doi:10.1145/3727126.
- 33 Wanchun Jiang, Haoyang Li, Jia Wu, Kai Wang, Fengyuan Ren, and Jianxin Wang. Intropective congestion control for consistent high performance. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, pages 428–445, New York, NY, USA, 2025. Association for Computing Machinery. doi:10.1145/3689031.3696084.
- 34 Dina Katabi, Mark Handley, and Charlie Rohrs. Congestion control for high bandwidth-delay product networks. *SIGCOMM Comput. Commun. Rev.*, 32(4):89–102, August 2002. doi:10.1145/964725.633035.

- 35 Frank Kelly. Fairness and stability of end-to-end congestion control*. *European Journal of Control*, 9(2):159–176, 2003. doi:10.3166/ejc.9.159-176.
- 36 Gautam Kumar, Nandita Dukkipati, Keon Jang, Hassan M. G. Wassel, Xian Wu, Behnam Montazeri, Yaogong Wang, Kevin Springborn, Christopher Alfeld, Michael Ryan, David Wetherall, and Amin Vahdat. Swift: Delay is simple and effective for congestion control in the datacenter. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, pages 514–528, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3387514.3406591.
- 37 T.V. Lakshman and U. Madhow. The performance of tcp/ip for networks with high bandwidth-delay products and random loss. *IEEE/ACM Transactions on Networking*, 5(3):336–350, 1997. doi:10.1109/90.611099.
- 38 Yanfang Le, Rong Pan, Peter Newman, Jeremias Blendin, Abdul Kabbani, Vipin Jain, Raghava Sivaramu, and Francis Matus. Strack: A reliable multipath transport for ai/ml clusters, 2024. doi:10.48550/arXiv.2407.15266.
- 39 Rodolfo I. Ledesma Goyzueta and Yu Chen. A deterministic loss model based analysis of cubic. In *2013 International Conference on Computing, Networking and Communications (ICNC)*, pages 944–949, 2013. doi:10.1109/ICCNC.2013.6504217.
- 40 Doug Leith and David S. Miller. tcp: tcp_vegas cong avoid fix · torvalds/linux@8d3a564, December 2008. [Online; accessed 2025-06-04]. URL: <https://github.com/torvalds/linux/commit/8d3a564da34e5844aca4f991b73f8ca512246b23>.
- 41 Tong Li, Kai Zheng, Ke Xu, Rahul Arvind Jadhav, Tao Xiong, Keith Winstein, and Kun Tan. Tack: Improving wireless transport performance by taming acknowledgments. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, pages 15–30, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3387514.3405850.
- 42 Xudong Liao, Han Tian, Chaoliang Zeng, Xinchun Wan, and Kai Chen. Astraea: Towards fair and efficient learning-based congestion control. In *Proceedings of the Nineteenth European Conference on Computer Systems*, EuroSys '24, pages 99–114, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3627703.3650069.
- 43 S.H. Low. A duality model of tcp and queue management algorithms. *IEEE/ACM Transactions on Networking*, 11(4):525–536, 2003. doi:10.1109/TNET.2003.815297.
- 44 S.H. Low. *Analytical Methods for Network Congestion Control*. Synthesis Lectures on Learning, Networks, and Algorithms. Springer International Publishing, 2022. URL: <https://books.google.com/books?id=tYFyEAAAQBAJ>.
- 45 Steven H. Low, Larry L. Peterson, and Limin Wang. Understanding tcp vegas: a duality model. *J. ACM*, 49(2):207–235, 2002. doi:10.1145/506147.506152.
- 46 Yuanwei Lu, Guo Chen, Bojie Li, Kun Tan, Yongqiang Xiong, Peng Cheng, Jiansong Zhang, Enhong Chen, and Thomas Moscibroda. Multi-Path transport for RDMA in datacenters. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 357–371, Renton, WA, April 2018. USENIX Association. URL: <https://www.usenix.org/conference/nsdi18/presentation/lu>.
- 47 Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis Ott. The macroscopic behavior of the tcp congestion avoidance algorithm. *SIGCOMM Comput. Commun. Rev.*, 27(3):67–82, 1997. doi:10.1145/263932.264023.
- 48 Ayush Mishra, Lakshay Rastogi, Raj Joshi, and Ben Leong. Keeping an eye on congestion control in the wild with neby. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, pages 136–150, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672223.

- 49 Radhika Mittal, Vinh The Lam, Nandita Dukkipati, Emily Blem, Hassan Wassel, Monia Ghobadi, Amin Vahdat, Yaogong Wang, David Wetherall, and David Zats. Timely: Rtt-based congestion control for the datacenter. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM '15*, pages 537–550, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2785956.2787510.
- 50 J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000. doi:10.1109/90.879343.
- 51 Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. Mahimahi: Accurate Record-and-Replay for HTTP. In *2015 USENIX Annual Technical Conference (USENIX ATC 15)*, pages 417–429, Santa Clara, CA, July 2015. USENIX Association. URL: <https://www.usenix.org/conference/atc15/technical-session/presentation/netravali>.
- 52 Vladimir Olteanu, Haggai Eran, Dragos Dumitrescu, Adrian Popa, Cristi Baciuc, Mark Silberstein, Georgios Nikolaidis, Mark Handley, and Costin Raiciu. An edge-queued datagram service for all datacenter traffic. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 761–777, Renton, WA, April 2022. USENIX Association. URL: <https://www.usenix.org/conference/nsdi22/presentation/olteanu>.
- 53 Jitendra Padhye, Victor Firoiu, Don Towsley, and Jim Kurose. Modeling tcp throughput: A simple model and its empirical validation. In *Proceedings of the ACM SIGCOMM '98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM '98*, pages 303–314, New York, NY, USA, 1998. Association for Computing Machinery. doi:10.1145/285237.285291.
- 54 Lucian Popa, Arvind Krishnamurthy, Sylvia Ratnasamy, and Ion Stoica. Faircloud: sharing the network in cloud computing. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X*, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/2070562.2070584.
- 55 Costin Raiciu, Sebastien Barre, Christopher Pluntke, Adam Greenhalgh, Damon Wischik, and Mark Handley. Improving datacenter performance and robustness with multipath tcp. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, pages 266–277, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/2018436.2018467.
- 56 Sudarsanan Rajasekaran, Manya Ghobadi, Gautam Kumar, and Aditya Akella. Congestion control in machine learning clusters. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks, HotNets '22*, pages 235–242, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3563766.3564115.
- 57 K. K. Ramakrishnan, Sally Floyd, and David L. Black. RFC 3168: The Addition of Explicit Congestion Notification (ECN) to IP, September 2001. [Online; accessed 21. Nov. 2024]. URL: <https://datatracker.ietf.org/doc/rfc3168>.
- 58 Danfeng Shan, Fengyuan Ren, Peng Cheng, Ran Shu, and Chuanxiong Guo. Micro-burst in data centers: Observations, analysis, and mitigations. In *2018 IEEE 26th International Conference on Network Protocols (ICNP)*, pages 88–98, 2018. doi:10.1109/ICNP.2018.00019.
- 59 Arjun Singhvi, Aditya Akella, Dan Gibson, Thomas F. Wenisch, Monica Wong-Chan, Sean Clark, Milo M. K. Martin, Moray McLaren, Prashant Chandra, Rob Cauble, Hassan M. G. Wassel, Behnam Montazeri, Simon L. Sabato, Joel Scherpelz, and Amin Vahdat. Irma: Re-envisioning remote memory access for multi-tenant datacenters. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM '20*, pages 708–721, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3387514.3405897.
- 60 Anirudh Sivaraman, Keith Winstein, Pratiksha Thaker, and Hari Balakrishnan. An experimental study of the learnability of congestion control. In *Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14*, pages 479–490, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2619239.2626324.

- 61 Rayadurgam Srikant and Tamer Başar. *The mathematics of Internet congestion control*. Springer, 2004.
- 62 Ao Tang, Jiantao Wang, and S.H. Low. Counter-intuitive throughput behaviors in networks under end-to-end control. *IEEE/ACM Transactions on Networking*, 14(2):355–368, 2006. doi:10.1109/TNET.2006.872552.
- 63 Ultra Ethernet Consortium. Working Groups - Ultra Ethernet Consortium, December 2023. [Online; accessed 18. Nov. 2024]. URL: <https://ultraethernet.org/working-groups>.
- 64 Weitao Wang, Masoud Moshref, Yuliang Li, Gautam Kumar, T. S. Eugene Ng, Neal Cardwell, and Nandita Dukkipati. Poseidon: Efficient, robust, and practical datacenter CC via deployable INT. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 255–274, Boston, MA, April 2023. USENIX Association. URL: <https://www.usenix.org/conference/nsdi23/presentation/wang-weitao>.
- 65 Ranysha Ware, Matthew K. Mukerjee, Srinivasan Seshan, and Justine Sherry. Modeling bbr’s interactions with loss-based congestion control. In *Proceedings of the Internet Measurement Conference, IMC ’19*, pages 137–143, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3355369.3355604.
- 66 Ranysha Ware, Adithya Abraham Philip, Nicholas Hungria, Yash Kothari, Justine Sherry, and Srinivasan Seshan. Ccanalyzer: An efficient and nearly-passive congestion control classifier. In *Proceedings of the ACM SIGCOMM 2024 Conference, ACM SIGCOMM ’24*, pages 181–196, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672255.
- 67 David X. Wei, Cheng Jin, Steven H. Low, and Sanjay Hegde. Fast tcp: Motivation, architecture, algorithms, performance. *IEEE/ACM Transactions on Networking*, 14(6):1246–1259, 2006. doi:10.1109/TNET.2006.886335.
- 68 Keith Winstein and Hari Balakrishnan. Tcp ex machina: Computer-generated congestion control. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM ’13*, pages 123–134, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2486001.2486020.
- 69 Francis Y. Yan, Jestin Ma, Greg D. Hill, Deepti Raghavan, Riad S. Wahby, Philip Levis, and Keith Winstein. Pantheon emulation result for scream and webrtc. 12 mbps bottleneck with 30 ms rtp and 1 bdp buffer, April 2018. [Online; accessed 2025-05-26]. URL: <https://pantheon.stanford.edu/result/2422/>.
- 70 Francis Y. Yan, Jestin Ma, Greg D. Hill, Deepti Raghavan, Riad S. Wahby, Philip Levis, and Keith Winstein. Pantheon: the training ground for internet congestion-control research. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pages 731–743, Boston, MA, July 2018. USENIX Association. URL: <https://www.usenix.org/conference/atc18/presentation/yan-francis>.
- 71 Yibo Zhu, Haggai Eran, Daniel Firestone, Chuanxiong Guo, Marina Lipshteyn, Yehonatan Liron, Jitendra Padhye, Shachar Raindel, Mohamad Haj Yahia, and Ming Zhang. Congestion control for large-scale rdma deployments. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM ’15*, pages 523–536, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2785956.2787484.
- 72 Yibo Zhu, Monia Ghobadi, Vishal Misra, and Jitendra Padhye. Ecn or delay: Lessons learnt from analysis of dcqcn and timely. In *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies, CoNEXT ’16*, pages 313–327, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2999572.2999593.

A Tradeoff derivations

We show that metrics (M1) robustness and (M2) fairness are at odds with (M3) congestion and (M4) generality. M1 & M2 require the contract to be gradual, while M3 & M4 require the contract to be steeper. We derive these tradeoffs quantitatively. We assume the contract function func has domain $[S_{\min}, S_{\max}]$ and range $[C_{\min}, C_{\max}]$, i.e., $C_{\max} = \text{func}(S_{\min})$, and $C_{\min} = \text{func}(S_{\max})$.

Our strategy to derive the tradeoffs is as follows, choice of a metric puts constraints on the contract, and this in-turn puts constraints on other metrics. So we fix one metric and see what constraints it puts on other metrics.

The tradeoffs involving fairness depend on fold . We consider $\text{fold} \in \{\sum, \max, \min\}$. The tradeoffs exists for \sum . We do not get tradeoffs for \max and \min , i.e., we can achieve max-min fairness independent of requirements on robustness/generality. We also show derivation steps for arbitrary fold , if future work wants to consider other statistics that accumulate differently across hops.

A.1 M1 vs. M3: robustness vs. congestion growth

Say we want the robustness error factor to be at most $\epsilon_r > 1$, then we compute a lower bound on $\text{growth}(n)$. We start from the definition of robustness error factor (Eq. 2), i.e., $\forall s, \frac{\text{func}(s)}{\text{func}(s+\delta s)} \leq \epsilon_r$. We evaluate it at $s = S_{\min}, s = S_{\min} + \delta s, s = S_{\min} + 2 * \delta s, \dots$, and perform the following algebraic manipulations:

$$\begin{aligned} & \frac{\text{func}(S_{\min})}{\text{func}(S_{\min} + \delta s)} \leq \epsilon_r \\ \implies C_{\max} = \text{func}(S_{\min}) & \leq \epsilon_r * \text{func}(S_{\min} + \delta s) \\ & \leq \epsilon_r * \epsilon_r * \text{func}(S_{\min} + 2 * \delta s) \\ & \leq \epsilon_r^k * \text{func}(S_{\min} + k * \delta s) \\ \implies \frac{C_{\max}}{\epsilon_r^k} & \leq \text{func}(S_{\min} + k * \delta s) \end{aligned} \tag{8}$$

$$\implies \text{func}^{-1} \left(\frac{C_{\max}}{\epsilon_r^k} \right) \geq S_{\min} + k * \delta s \tag{9}$$

$$\begin{aligned} \implies \frac{\text{func}^{-1} \left(\frac{C_{\max}}{\epsilon_r^k} \right)}{\text{func}^{-1}(C_{\max})} & \geq \frac{S_{\min} + k * \delta s}{S_{\min}} \\ \implies \text{growth}(\epsilon_r^k) & \geq \frac{\text{func}^{-1} \left(\frac{C_{\max}}{\epsilon_r^k} \right)}{\text{func}^{-1}(C_{\max})} \geq \frac{S_{\min} + k * \delta s}{S_{\min}} \end{aligned} \tag{10}$$

$$\implies \text{growth}(n) \geq \frac{\text{func}^{-1} \left(\frac{C_{\max}}{n} \right)}{\text{func}^{-1}(C_{\max})} \geq \frac{S_{\min} + \frac{\log(n)}{\log(\epsilon_r)} * \delta s}{S_{\min}} \tag{11}$$

Note, we get Eq. 9 by taking func^{-1} on both sides, the inequality flips as func and func^{-1} are monotonically decreasing. We get the left inequality in Eq. 10 from the definition of growth (Eq. 7). Eq. 11 shows that lower error factor implies higher signal growth. The inequalities become equalities when (by substituting $S_{\min} + k * \delta s$ by s in Eq. 8): $\text{func}(s) = C_{\max} \epsilon_r^{\frac{s - S_{\min}}{\delta s}}$. This is same as the exponential contract from [8].

A.2 M1 vs. M4: robustness vs generality

We compute an upper bound on $\frac{C_{\max}}{C_{\min}}$ given we want the error factor to be at most $\epsilon_r > 1$. We start from Eq. 8, and substitute $S_{\min} + k * \delta s$ by S_{\max} :

$$\begin{aligned} \frac{C_{\max}}{\epsilon_r^k} &\leq \text{func}(S_{\min} + k * \delta s) \\ \implies C_{\max} \epsilon_r^{-\frac{S_{\max} - S_{\min}}{\delta s}} &\leq \text{func}(S_{\max}) = C_{\min} \\ \implies \frac{C_{\max}}{C_{\min}} &\leq \epsilon_r^{\frac{S_{\max} - S_{\min}}{\delta s}} \end{aligned}$$

Lower the robustness error, lower is the range of bandwidths we can support. The inequality becomes equality for the exponential contract.

A.3 M2 vs. M3: fairness vs. congestion growth

We derive a lower bound on $\text{growth}(n)$ given that we want the throughput ratio in parking lot for k hops to be $\text{ratio}^*(k)$. For the parking lot ratio to be $\text{ratio}^*(k)$, we need:

$$\forall k. \quad \max_s \frac{\text{func}(s)}{\text{func}(\text{fold}(s, s, \dots, s))} = \text{ratio}^*(k)$$

Say, the maximum of LHS is achieved when the link capacity in the parking lot is $C = C^*(k)$. $\text{ratio}^*(\cdot)$ and $C^*(\cdot)$ are functions of k . For brevity, we drop the k , and refer to them as r^* and C^* respectively. The statements are true for all positive k .

Instead of directly computing a constraint on $\text{growth}(n)$, we first derive a constraint on $\text{growth}_{C^*}(n)$. This eventually constrains $\text{growth}(n)$. Where, $\text{growth}_C(n)$ is defined as $\frac{\text{func}^{-1}(C/n)}{\text{func}^{-1}(C)}$. I.e., $\text{growth}(n) = \max_C \text{growth}_C(n)$ (from Eq. 7). Note, $\text{growth}_C(1) = 1$ from this definition. We will use this later.

Consider the execution of the contract (CCA) on a parking lot topology with k hops and link capacity C^* . In steady-state, we have:

$$\begin{aligned} r_0 + r_k &= C^* \quad \text{same as Eq. 3, and,} \\ \frac{r_0}{r_k} &= r^* \quad \text{from the definition of } r^* \text{ above.} \end{aligned}$$

On solving these, we get:

$$r_0 = C^*/(r^* + 1) \quad \text{and,} \quad r_k = C^*r^*/(r^* + 1)$$

Say the aggregate statistic seen by flow f_i in the parking lot is $s_i = \text{func}^{-1}(r_i)$. Then $\text{growth}_{C^*}(C^*/r_i) = \frac{\text{func}^{-1}(C^*/(C^*/r_i))}{\text{func}^{-1}(C^*)} = \frac{s_i}{\text{func}^{-1}(C^*)}$. We define $s^* = \text{func}^{-1}(C^*)$, and substitute i by 0. On rearranging, we get:

$$s_0 = \text{func}^{-1}(C^*) \cdot \text{growth}_{C^*}(C^*/r_0) = s^* \cdot \text{growth}_{C^*}(r^* + 1)$$

$$\text{Likewise, } s_k = s^* \cdot \text{growth}_{C^*}(1 + 1/r^*)$$

In steady-state:

$$\begin{aligned} s_0 &= \text{fold}(s_k, s_k, \dots, k \text{ times}) \quad \text{same as Eq. 5,} \\ &\quad \text{on substituting } s_0 \text{ and } s_k \text{ just computed, we get} \\ \implies s^* \cdot \text{growth}_{C^*}(r^* + 1) &= \text{fold}(s^* \cdot \text{growth}_{C^*}(1 + 1/r^*), \dots, k \text{ times}) \end{aligned}$$

If multiplication distributes over accumulation, this is true for $\text{fold} \in \{\sum, \max, \min\}$, then:

$$\text{growth}_{C^*}(r^* + 1) = \text{fold}(\text{growth}_{C^*}(1 + 1/r^*), \dots, k \text{ times})$$

We consider under different choices of fold , what constraint ratio^* puts on growth_{C^*} . For fold as \max or \min . We can vacuously meet this constraint for $\text{ratio}^* = 1$ (corresponding to max-min fairness), and there is no tradeoff. For $\text{fold} = \sum$, the constraint becomes:

$$\text{growth}_{C^*}(r^* + 1) = k \cdot \text{growth}_{C^*}(1 + 1/r^*)$$

Let $n = r^* + 1 = \text{ratio}^*(k) + 1$, i.e., $k = \text{ratio}^{*-1}(n - 1)$, where ratio^{*-1} is the inverse of ratio^* then:

$$\text{growth}_{C^*}(n) = \text{ratio}^{*-1}(n - 1) \cdot \text{growth}_{C^*}(1 + 1/(n - 1))$$

Since, func is decreasing, growth_C is increasing for all C . So $\text{growth}_{C^*}(1 + 1/(n - 1)) \geq \text{growth}_{C^*}(1) = 1$. So, we get:

$$\text{growth}_{C^*}(n) \geq \text{ratio}^{*-1}(n - 1)$$

If we want better fairness, we need ratio^* to be slow growing, then ratio^{*-1} is fast-growing, and so growth_{C^*} is fast-growing, and so congestion growth $\text{growth} \geq \text{growth}_{C^*}$ needs to be fast-growing.

For proportional fairness, $\text{ratio}^*(k) = k$, and so $\text{growth}(n) = O(n)$, this is met by Vegas. For α -fairness, $\text{ratio}^*(k) = \sqrt[\alpha]{k}$, and so $\text{growth}(n) = O(n^\alpha)$, this is met by the contract $\text{func}(s) = \frac{1}{r^\alpha}$.

A.4 M2 vs. M4: fairness vs. generality

We compute an upper bound on $\frac{C_{\max}}{C_{\min}}$ given we want the throughput ratio in parking lot for k hops to be at most $\text{ratio}^*(k)$. For the throughput ratio to be at most $\text{ratio}^*(k)$, we need:

$$\forall s, k. \frac{\text{func}(s)}{\text{func}(\text{fold}(s, s, \dots, s))} \leq \text{ratio}^*(k)$$

The derivation beyond this depends on fold . For $\text{fold} \in \{\max, \min\}$, the above constraint is vacuously true and there is not tradeoff. For $\text{fold} = \sum$, we need:

$$\forall s, k. \text{func}(s) \leq \text{ratio}^*(k) \cdot \text{func}(k * s)$$

Picking $s = S_{\min}$, and $k = \frac{S_{\max}}{S_{\min}}$, we get:

$$\begin{aligned} \text{func}(S_{\min}) &\leq \text{ratio}^* \left(\frac{S_{\max}}{S_{\min}} \right) \cdot \text{func}(S_{\max}) \\ \implies \frac{C_{\max}}{C_{\min}} &\leq \text{ratio}^* \left(\frac{S_{\max}}{S_{\min}} \right) \end{aligned}$$

Better fairness implies smaller ratio $\text{ratio}^*(\cdot)$ and a smaller range of supported bandwidths.

For fairness to be at least as good as proportional fairness, we need $\text{ratio}^*(k) \leq k$, or:

$$\forall s, k. \text{func}(s) \leq \text{ratio}^*(k) \cdot \text{func}(k * s) \leq k \cdot \text{func}(k * s)$$

Substituting $k = S_{\min}/s$, we get:

$$\text{func}(s) \leq \frac{S_{\min} C_{\max}}{s}$$

Equality occurs when $\text{func}(s) = \frac{S_{\min} C_{\max}}{s}$, which is the same contract as Vegas.

Likewise, for α -fairness, $\text{ratio}^*(k) = \sqrt[\alpha]{k}$, and so $\frac{C_{\max}}{C_{\min}} \leq \sqrt[\alpha]{\frac{S_{\max}}{S_{\min}}}$. The equality occurs for the contract: $\text{func}(s) = \frac{1}{r^\alpha}$.