

Article

Neutralizing Schedulability Asymmetry Through Period Decomposition: Configured Grant Scheduling for Periodic Reservations in NR Sidelink Mode 1

Hyungjoon Shin  and Hyogon Kim * 

Department of Computer Science and Engineering, Korea University, Seoul 02841, Republic of Korea; kilk119@korea.ac.kr

* Correspondence: hyogon@korea.ac.kr

Abstract

Periodic resource reservation underpins vehicular sidelink traffic, yet its compatibility is inherently asymmetric. In New Radio (NR) sidelink Mode 1, two configured grant (CG) streams collide if and only if the greatest common divisor (GCD) of their periods divides their offset difference, so coprime periods cannot be separated by any offset on a single subchannel. Admission opportunity is thus governed by the arithmetic of the period set, not by load alone—a disparity we call schedulability asymmetry. We propose Divisor-Aligned Period Decomposition (DAPD), which fixes a common divisor d and reduces coexistence to a closed-form GCD test: periods that are multiples of d are admitted directly, and incompatible ones are reshaped into a d -aligned substitute within the delay budget rather than rejected. This levels the schedulability asymmetry at its source for any catalog period whose delay budget admits a d -compatible substitute: once every request is represented on a common divisor, no period is structurally disadvantaged, so such a period can be admitted regardless of its arithmetic relationship to the others. Under a collision-only scheduler model that isolates deterministic reservation collisions from physical-layer effects, every admitted CG is collision-free by construction. An extension, DAPD+, salvages the resulting slack with dynamic grants that never overlap an admitted CG. In simulation over 1000 seeds across vehicle densities up to subchannel saturation, DAPD+ attains the highest overall packet delivery ratio among all evaluated schedulers, including an oracle offset optimizer and a recent multi-configuration CG scheduler, while holding admitted-CG delivery at one and keeping far more per-vehicle delivery ratios above target, where the reactive baseline's per-vehicle reliability collapses with density. DAPD+ thus serves the largest volume of periodic traffic at high reliability without sacrificing aggregate throughput, equalizing admission opportunity across heterogeneous periods.

Keywords: NR sidelink; configured grant; Mode 1 scheduling; V2X; collision-free admission; schedulability asymmetry; period decomposition; dynamic grant salvage; intelligent transportation systems



Received: 10 June 2026

Revised: 8 July 2026

Accepted: 10 July 2026

Published: 13 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The compatibility of periodic reservations in a shared slotted channel is inherently asymmetric. Although each reservation has a temporally symmetric occupancy pattern that repeats with its period, different periods do not interact symmetrically: reservations with favorable common divisors can often be separated by offset selection, whereas coprime

or coprime-rich period combinations cannot be separated by offset selection on the same subchannel and therefore become structurally difficult to admit without period substitution or additional resource dimensions. As a result, admission opportunity is not determined only by load, but also by the arithmetic asymmetry of the period set, an effect that is acute in vehicular sidelink pools where heterogeneous service periods drawn from a fixed catalog must share the same resources. We refer to this disparity as schedulability asymmetry: throughout this paper, schedulability denotes the gNB scheduler's ability to admit a configured grant as a collision-free periodic reservation, so two requests with the same load but different periods can have unequal schedulability purely because of their arithmetic relationship to the already-admitted reservations. The approach taken here addresses this asymmetry directly by exploiting modular compatibility among periodic reservations: instead of treating all periods as equally packable or relying on offset search, every request is represented on a common divisor d , so that requests whose native period is d -aligned are admitted as is and the rest are reshaped onto d -aligned substitutes rather than turned away. Because all admitted reservations then share the same divisor, no period is structurally disadvantaged and any catalog period whose delay budget admits a d -compatible substitute becomes admissible regardless of its arithmetic relationship to the others; this is what levels the schedulability asymmetry, within the 3GPP periodicity catalog and the declared delay budgets, and the resources left idle by the divisor alignment are reused to serve dynamic traffic.

New Radio (NR) sidelink Mode 1 must support both periodic traffic that benefits from configured grants (CGs) and aperiodic or event-driven traffic that is more naturally served by dynamic grants (DGs) [1–4]. In Mode 1 the next-generation NodeB (gNB) configures sidelink resource pools and CGs through Radio Resource Control (RRC) signaling and issues DGs through scheduling grants, so the scheduler has the centralized visibility needed to verify coexistence among CG configurations before resources are committed. The sidelink CG periodicity field in TS 38.331 supports millisecond-scale periods, including integer values below 100 ms and coarser 100 ms multiples up to 1000 ms [4], but two catalog-valid periods can be mutually incompatible on the same subchannel: a CG with period P and offset o reserves $o + kP$, and two such sequences collide if and only if $\gcd(P_1, P_2)$ divides $o_2 - o_1$. Coprime periods therefore collide for every offset choice, and these collisions are deterministic rather than sensing- or contention-driven. Offset search alone cannot resolve them; heterogeneous CG admission is fundamentally an arithmetic compatibility problem.

This paper proposes Divisor-Aligned Period Decomposition (DAPD) for compatibility-aware CG admission. DAPD selects a common divisor d and represents every admitted CG as a d -aligned reservation structure, that is, one whose period is a multiple of d , which induces a residue-class geometry on the slot timeline: CGs in different residue classes modulo d are automatically disjoint, and CGs in the same class are tested by a closed-form GCD condition on their normalized multipliers. A user equipment (UE) request, characterized by its native traffic period p and packet delay budget $D_{\max}$, is admitted with $q = p$ when $d \mid p$, and otherwise, rather than being rejected, is mapped to one of two catalog-valid d -compatible substitutes (a single-CG candidate or a halved-pattern candidate consisting of two interleaved subsequences) whose opportunity spacing satisfies $D_{\max}$. Under the collision-only scheduler model, every admitted CG reservation is collision-free at the reservation level and achieves unit post-admission PDR. This reservation-level reliability is orthogonal to physical (PHY)-layer reliability provided by link adaptation and hybrid automatic repeat request (HARQ), which the present model deliberately does not address.

The proposed principle admits a minimal realization, GCD-only, that omits period substitution and rejects every d -incompatible request, exposing the admission test itself

as the source of the reservation-level guarantee. DAPD extends it with period substitution that broadens coverage to d -incompatible requests, a benefit concentrated in the admission-light regime, and DAPD+ adds scheduler-level dynamic grant (DG) salvage that serves pending packets of pre-admission or waiting UEs and native DG packets on resources not occupied by an admitted CG transmission, without ever preempting one. DAPD+ targets operation where broad service coverage and reservation-level reliability are both desired; its DG-salvage component mainly improves low-to-medium density operation, while high-density performance remains capacity-limited. The problem is most pronounced in dense, CG-heavy Mode 1 pools where the offered load per subchannel is high enough for deterministic period incompatibilities to dominate offset repair. When the per-subchannel load is light, the same arithmetic constraint is less frequently binding. Two PDR notions follow: the overall packet PDR includes pre-admission packets and can fall below one, while the admitted-CG PDR equals one by construction under the collision-only model. The contributions are (i) a closed-form admission test, derived from the GCD coexistence condition and the residue-class structure induced by d , that prevents deterministic reservation collisions among heterogeneous CGs at admission time; (ii) a Period Alignment Unit (PAU) and Admission Control Unit (ACU) decomposition that maps native traffic descriptors to catalog-valid substitutes and verifies coexistence in time linear in the local residue-class occupancy; (iii) DAPD+, which preserves the admitted-CG guarantee while salvaging pre-admission and waiting-phase losses through DG service on unoccupied resources; and (iv) a simulation study over 1000 seeds showing that DAPD+ attains the highest overall PDR across the studied range while sustaining admitted-CG PDR at one and keeping the largest share of per-UE delivery ratios above target even when unadmitted UEs are counted, evaluated against MC, an offline oracle offset optimizer (OPT-offset), and One-shot as non-blocking native-period baselines, the recent multi-configuration CG scheduler C-FAST [5] ported to sidelink Mode 1, GCD-only as the minimal realization, DAPD, and DAPD+. To the best of our knowledge, this is the first work to reduce heterogeneous-period CG coexistence in Mode 1 to a closed-form, order-independent arithmetic admission test rather than an offset search or a reactive repair, the first to identify and avert the coprime monopoly hazard in which a single admitted coprime-period reservation deterministically blocks every subsequent request on the subchannel, and the first to treat dynamic grants as a scheduler-level consumer of the structured slack that compatibility-aware CG admission produces.

2. Related Work

Configured-grant and semi-persistent scheduling (SPS) have been studied for ultra-reliable low-latency communication (URLLC), deterministic wireless traffic, and 5G time-sensitive networking (TSN) integration, including contention-free CG placement and multi-traffic resource optimization [5–8]. A recurring technique in this line is to adjust the timing of periodic reservations: periodicity-shifting methods perturb the reservation period or phase so that transmissions of different flows interleave without overlap [9,10]. These works exploit deterministic timing and structured reservations, but they target cellular uplink, where the base station is the receiver and orthogonal uplink resources are scheduled under a different interference and feedback structure. Two structural differences separate them from the present problem. First, in sidelink Mode 1 the gNB configures UE-to-UE transmissions over a shared resource pool, so recurring reservation overlaps among peers, rather than uplink resource contention at a single receiver, are the direct obstacle. Second, and more fundamentally, offset or periodicity adjustment cannot resolve the incompatibility addressed here: by Corollary 1, two coprime-period reservations collide at every offset, so no timing shift admits them together. This is why DAPD acts at the period level, mapping

incompatible native periods onto a divisor-aligned catalog through period substitution, rather than searching for a feasible offset that, for coprime periods, does not exist.

Centralized C-V2X and NR-V2X resource allocation has examined spectrum re-partitioning [11], preference-aware allocation [12], Mode 1 dynamic resource adaptation [13], spatial and frequency reuse [14,15], and broader NR-V2X resource-allocation mechanisms [16–18]. These centralized schemes optimize how resources are partitioned or reused across UEs, but they do not model the deterministic arithmetic coupling between heterogeneous reservation periods on a shared subchannel, which persists regardless of how the spatial or spectral dimension is allocated. Mode 2 work has instead treated sensing errors, hidden terminals, persistent collisions, and autonomous reselection [19,20]; these collision sources are contention- or sensing-driven and are mitigated probabilistically, whereas the incompatibility considered here is deterministic and independent of sensing quality, contention outcome, or initial offset search, and therefore admits a closed-form admission test rather than a statistical mitigation. To the best of our knowledge, existing schemes also do not treat DG as a scheduler-level consumer of the structured idle opportunities produced by compatibility-aware periodic CG admission. DAPD+ does, treating both unreserved resources and idle CG opportunities (slots reserved by admitted CGs but unused at a given epoch because the owner has no packet due) as DG-available resources, and using them to salvage pending packets of pre-admission or waiting UEs and native DG packets, without weakening the admitted-CG guarantee.

A further line of work applies optimization, graph-based formulations, and machine learning to V2X resource allocation. Deep reinforcement learning (DRL) has been used to mitigate collisions of the Mode 4 semi-persistent scheduling process by selecting the resource reservation interval and transmit power with NOMA/SIC-assisted reception, jointly optimizing information freshness and energy consumption [21], and broader surveys catalog optimization- and learning-driven NR-V2X allocation mechanisms [18]. The DRL line is, moreover, not directly comparable as a baseline here: it operates on Mode 4 autonomous sensing-based selection, controlling the resource reservation interval and transmit power per UE under freshness and energy objectives, whereas this paper concerns centralized Mode 1 admission of heterogeneous-period CGs under a deterministic collision predicate, so the control variables, the objective, and the system mode all differ. More fundamentally, whether two reservations collide is a closed-form number-theoretic predicate of their periods and offsets, so there is no unknown stochastic dynamics for a learning agent to estimate; the placement problem such an agent would learn is exactly the offset-selection problem whose offline optimum OPT-offset already realizes. These methods occupy a different layer of the problem than the one addressed here: they search or learn a placement within the feasible set of offsets and resources, but the feasible set itself is fixed by the arithmetic of the period set. By Corollary 1, no placement policy, whether exhaustive, graph-colored, or learned, can admit two coprime-period reservations collision-free on the same subchannel while the native periods are kept, so the reservation-level performance of any placement-layer method is bounded by the same deterministic collision structure regardless of how sophisticated the search is. This is why the evaluation of Section 5 represents this entire class by an offline oracle, OPT-offset, which jointly optimizes every reservation's offset with full arrival and departure knowledge and thus serves as a near-upper bound of the class, rather than by individual heuristics. DAPD instead operates on the feasible set itself, substituting incompatible native periods with catalog-valid divisor-aligned structures, which no placement-layer method can emulate while keeping native periods. Closest to our setting is the C-FAST configured grant scheduler of Larrañaga-Zumeta et al. [5], developed for 5G-TSN integration on the cellular uplink: a centralized scheduler makes heterogeneous-period CG flows coexist by giving each flow multiple CG configurations

with per-packet activation vectors computed over the hyperperiod, steering packets that would collide onto pre-planted avoidance configurations. It thus attains reservation-level protection within the native-period feasible set by multiplying configurations up to the standard's per-UE cap, at the cost of hyperperiod-scale conflict enumeration at admission and a growing configuration footprint, whereas DAPD attains the same protection by substituting the period itself onto a divisor-aligned catalog, with a closed-form order-independent test, at most two configurations per flow, and structured slack that a salvage layer can reuse; Section 5 ports C-FAST to sidelink Mode 1 and compares the two directly. Table 1 summarizes the approach families by mechanism and capability.

Table 1. Prior approach families versus DAPD by mechanism and capability.

Approach Family	Representative Works	Mechanism	Separates Coprime Periods?	Reservation-Level Guarantee
Offset selection, min-collision	[9,10]; MC, OPT-offset	offset or phase adjustment at fixed period	No (Corollary 1)	No
Reactive repair	[22,23]; One-shot	per-collision deflection to idle resources	No (repairs, does not prevent)	No
Multi-CG avoidance	[5]; C-FAST	per-packet activation vectors over the hyperperiod	Yes (multiplexing up to per-UE CG cap)	Yes (while configurations suffice)
Centralized partitioning, reuse	[11,12,14,15]	spectrum or spatial dimension	Orthogonal to period arithmetic	No
Learning-based allocation	[18,21]	learned placement policy	Bounded by same feasible set	No
DAPD, DAPD+ (this work)		period substitution onto common divisor	Yes (by construction)	Yes ($PDR_{\text{adm}} = 1$)

Research gap. What is missing in the literature is a mechanism that makes heterogeneous catalog periods coexist on a shared subchannel when their arithmetic forbids it. Offset-selection and periodicity-shifting schemes are insufficient because, by Corollary 1, coprime-period reservations collide at every offset, so no timing adjustment at fixed periods can admit them together. Centralized partitioning and reuse schemes are insufficient because they allocate spatial or spectral dimensions and leave the time-domain arithmetic coupling between periods untouched on every subchannel they create. Mode 2 sensing enhancements and learning-based allocation are insufficient because they mitigate contention-driven or stochastic collisions, whereas the incompatibility considered here is deterministic, recurs at every least common multiple of the conflicting periods (Corollary 2), and is independent of sensing quality or learned policy. Multi-configuration avoidance in the style of C-FAST [5] does make incompatible periods coexist, but by enumerating conflicts over the hyperperiod and spending up to the per-UE configuration cap on each flow; it leaves the period arithmetic in place and pays for the incompatibility in configurations and computation rather than removing it. No existing scheme, to our knowledge, modifies the period structure itself under catalog and delay budget constraints so that admission becomes a closed-form feasibility test; that is the gap this paper fills.

3. System Model

3.1. Notation Summary

For ease of reference, Table 2 collects the symbols used throughout the paper. Symbols are introduced again at their first contextual use, but the table is intended as a single lookup for the reader.

Table 2. Summary of notation used in the paper.

Symbol	Meaning
Resource pool and CG catalog	
T_s	slot duration (1 ms)
P, o	CG reservation period and offset
P_{3GPP}	3GPP sidelink CG periodicity catalog
d	common divisor selected by DAPD
C	d -compatible catalog, $C = \{q \in P_{3GPP} : d \mid q\}$
Request and substitute	
p	native traffic period of a request
$D_{\max}$	one-way packet delay budget
q	catalog-valid period chosen by the PAU
m	number of subsequences (1: single, 2: halved)
b	normalized multiplier, $b = q/d$
ϕ	normalized phase index, $\phi \in \{0, \dots, b-1\}$
r	residue class index, $r \in \{0, \dots, d-1\}$
n_r	subsequences occupying residue class r
Provisioning and DG salvage	
$\eta(p, q, m)$	provisioning ratio, mp/q (Equation (4))
$\bar{\eta}$	average provisioning ratio over traffic mix
N_{op}	OP slots over a hyper-period (Equation (8))
$Q(t)$	DG queue at slot t (native plus waiting CG)
$R(t)$	DG-available slots at slot t
Metrics	
PDR	overall packet PDR
PDR_{adm}	admitted-CG PDR, $N_{\text{del}}^{\text{adm}} / N_{\text{tx}}^{\text{adm}}$
η_{reuse}	slack reuse efficiency, $N_{\text{op} \rightarrow \text{DG}} / N_{\text{op}}$
PDR_{salv}	DG-salvage packet PDR
Traffic and load	
ρ	UE density (UEs/km)
L	road segment length
T_{sjn}	mean UE sojourn time on the segment
γ	CG fraction of offered load

3.2. Resource Pool, CG Requests, and Mixed Traffic

We consider a single-subchannel NR sidelink resource pool with slot duration $T_s = 1$ ms throughout this paper. The arithmetic structure that drives DAPD is a per-subchannel time-domain property: two CG reservations on the same subchannel collide if and only if the GCD of their periods divides their offset difference, and this condition is independent of any other subchannel. We develop the admission rule and the bulk of the evaluation on a single subchannel, which isolates the time-domain arithmetic incompatibility that is the focus of this paper; the multi-subchannel composition of the same per-subchannel test is then evaluated separately in Section 5.8. In a multi-subchannel pool, the same GCD test applies independently per subchannel, while cross-subchannel coupling such as half-duplex constraints requires an additional placement layer that is outside the scope of this paper. On a single subchannel, by contrast, the half-duplex constraint is vacuous under a collision-free admission rule: every slot carries at most one transmission, so a UE transmitting in slot t forgoes no concurrent reception in that slot, and half-duplex becomes binding only when multiple subchannels carry simultaneous transmissions, the case revisited in Section 5.8. A CG with period P and offset o reserves $R(P, o) = \{o + kP : k \in \mathbb{Z}_{\geq 0}\}$. Let $P_{3GPP} = \{1, 2, \dots, 100, 200, \dots, 1000\}$ ms be the 3GPP sidelink CG periodicity catalog, as configured by the periodicity parameter of the sidelink

configured grant in TS 38.331 [4], and let DAPD pick a common divisor $d > 1$ with the d -compatible catalog

$$C = \{q \in P_{3GPP} : d \mid q\}. \quad (1)$$

A CG request contains arrival time t_0 , native traffic period p , and one-way packet delay budget $D_{\max}$. The traffic descriptors used by the admission rule are available to the gNB through standard procedures. The native period, timing offset, and message size of a periodic sidelink flow are reported through the sidelink UE assistance information procedure of TS 38.331: the UE includes in its UEAssistanceInformation message the sl-UE-AssistanceInformationNR field, a list of SL-TrafficPatternInfoentries each carrying the traffic periodicity, timing offset, message size, and the associated sidelink QoS flow identity [4]. The packet delay budget is part of the PC5 quality-of-service (QoS) profile established at QoS flow setup [24,25]. The overall flow is the standard Mode 1 procedure: a UE entering gNB coverage requests sidelink resources through SidelinkUEInformationNR with per-destination QoS information, supplies its traffic pattern through the assistance report above, and receives its CG configuration within an RRCReconfiguration message [1,4]. The signaling cost is therefore one RRC assistance report per flow at setup, followed by one CG configuration (RRC for Type 1, RRC plus one downlink control information activation for Type 2) [1]; no per-packet control signaling is involved, and DAPD adds no signaling beyond this standard Mode 1 CG configuration because the alignment and admission logic runs entirely at the gNB on already-reported descriptors. The report is event-triggered: it is refreshed only when the traffic pattern changes, so for periodic V2X services whose generation patterns are stable over the sojourn timescale, the update frequency is far below the scheduling timescale. A mismatch between the declared and actual pattern degrades only the efficiency of the affected reservation, whose unused opportunities become reclaimable slack for DG salvage, not the collision-freedom of other admitted CGs, which depends only on the configured periods and offsets.

The mixed CG/DG regime is parameterized by the CG fraction of offered load

$$\gamma = \frac{L_{CG}}{L_{CG} + L_{DG}}, \quad (2)$$

where L_{CG} and L_{DG} are the offered packet rates from periodic CG flows and aperiodic native DG packets, respectively. The simulator holds the total offered load constant while sweeping γ from 0 (pure native DG) to 1 (pure periodic CG). Native DG packets have deadline $D_{\max}^{DG}$ and enter a DG queue, which also receives waiting-phase CG packets generated during the pre-admission interval of CG UEs that have not yet been admitted; the DG queue is therefore $Q(t) = Q^{\text{nat}}(t) \cup Q^{\text{wait}}(t)$. This regime captures the operating point in which divisor-induced over-provisioning matters: CG traffic creates structured collision-free reservations, and the DG queue draws on resources that are not occupied by admitted CG transmissions at the current epoch, comprising unreserved resources and idle CG opportunities at slots whose reserving owner has no packet to send.

3.3. Scope and Metrics

The reliability guarantee in this paper is reservation-level: the analysis isolates deterministic CG-to-CG reservation collisions and does not model PHY decoding errors, HARQ failures, control-channel loss, or fading. Unit admitted-CG PDR therefore means an admitted reservation is collision-free under the scheduler model, not that PHY decoding succeeds in a fading channel. A CG opportunity is reclaimable for DG service only when its owner has no packet due, in which case the slot joins the DG-available pool; reclaim is cross-UE, and DAPD+ never preempts an admitted CG transmission. The overall packet PDR is

$$\text{PDR} = \frac{N_{\text{del}}^{\text{CG}} + N_{\text{del}}^{\text{DG}}}{N_{\text{gen}}^{\text{CG}} + N_{\text{gen}}^{\text{DG}}}, \quad (3)$$

which includes packets generated before CG admission and can be below one. The admitted-CG PDR is $\text{PDR}_{\text{adm}} = N_{\text{del}}^{\text{adm}} / N_{\text{tx}}^{\text{adm}}$, equal to one by construction for every scheduler that enforces the GCD admission test (GCD-only, DAPD, DAPD+, and their DG-salvage variants) under the collision-only model, since an admitted reservation that passes the test is collision-free for the remainder of its sojourn. This guarantee is orthogonal to PHY reliability: if a link delivers an admitted transmission with probability $1 - \text{PER}_{\text{PHY}}$, the effective post-admission delivery is, to first order, $\text{PDR}_{\text{adm}} \cdot (1 - \text{PER}_{\text{PHY}}) = 1 - \text{PER}_{\text{PHY}}$; the separation is approximate under HARQ and correlated fading, but the point is that this paper prevents the deterministic collision factor and leaves the PHY factor to link-level mechanisms. The two loss sources also differ in kind, not only in layer: by Corollary 2, a collision between incompatible reservations recurs deterministically at every least common multiple of the conflicting periods for the lifetime of both reservations, so it cannot be averaged away by HARQ or link adaptation, whereas PHY losses are stochastic and are precisely what HARQ targets; preventing the deterministic factor is therefore a prerequisite for, not a substitute for, PHY reliability mechanisms, and it also spares the pool the recurring retransmission load that deterministic collisions would otherwise inject into the pool. All delivery ratios reported in this paper are accordingly reservation-level results that upper-bound end-to-end delivery. When a native period p is mapped to a d -compatible substitute with reservation period q served by m subsequences, the substitute can reserve more opportunities than the native rate requires. The provisioning ratio

$$\eta(p, q, m) = \frac{mp}{q} \quad (4)$$

captures this overhead: $\eta = 1$ when the substitute is native-compatible, $\eta > 1$ when $q < p$ or when a halved pattern with $m = 2$ is used. The average provisioning ratio $\bar{\eta}$ is the expectation of $\eta(p, q, m)$ over the native traffic mix and is used in Section 3.4 to compare divisor choices. Slots reserved by substitutes with $\eta > 1$ but not used by their owners are termed over-provisioned (OP) slots; their count over a hyper-period is derived in (8). The slack reuse efficiency $\eta_{\text{reuse}} = N_{\text{op} \rightarrow \text{DG}} / N_{\text{op}}$ counts only OP slots used to salvage native DG packets, so η_{reuse} isolates the cross-UE consumption of idle CG slack by native DG traffic from other DG-available resources. To answer the question of whether DG salvage actually rescues the UEs it targets, we further define a salvage-target metric. Among DAPD+ packets, the subset eligible for DG salvage consists of all pre-admission CG packets generated by a UE before its CG admission, and the DG-salvage packet PDR is

$$\text{PDR}_{\text{salv}} = \frac{N_{\text{DG,del}}^{\text{preCG}}}{N_{\text{DG,elig}}^{\text{preCG}}}, \quad (5)$$

with the corresponding packet error rate $1 - \text{PDR}_{\text{salv}}$, where the eligible count denotes pre-admission CG packets whose owning UE is waiting for CG admission, and the delivered count denotes those that the DG salvage rule in Algorithm 3 to be explained later actually delivers within the packet delay budget. This metric is independent of how many UEs eventually receive CG admission and isolates the rescue rate of the salvage mechanism itself.

3.4. Divisor Selection

The common divisor d jointly controls catalog preservation, substitute feasibility, residue-class separation, and over-provisioning, so it is selected offline from the expected service mix rather than treated as a universal constant. Two criteria apply. First, divisors

of 100 preserve the 100 ms-aligned coarse catalog $\{100, 200, \dots, 1000\}$ ms, so long-period services naturally aligned with that catalog remain admissible without period substitution, whereas a non-divisor d forces these long periods onto d -compatible substitutes and thereby incurs unnecessary over-provisioning. Second, the halved-pattern substitute requires $d \leq 2D_{\max}$ to support tight-deadline traffic: the smallest d -compatible catalog period is d itself and the effective opportunity spacing of the halved pattern is $q/2$, so the deadline condition $q/2 \leq D_{\max}$ yields $d \leq 2D_{\max}$, and a divisor outside this range cannot place urgent services within their deadlines. Figure 1 visualizes the two criteria for the traffic mix in Table 4, with uniform service weights. Figure 1a sets $D_{\max} = p$ for every service and exposes the feasibility boundary of the second criterion: because of the urgent service (with a period and deadline of 10 ms), every $d > 20$ and the odd values 11 to 19 leave at least one service inadmissible under the PAU (triangles), so infeasible divisors are screened out immediately; within the feasible region, divisors and non-divisors alike sit near $\bar{\eta} \approx 1$ (1.10 at $d = 20$), and some non-divisors even dip slightly lower, because generous deadlines let them find close substitutes with little over-provisioning. Figure 1b tightens every deadline to $D_{\max} = 10$ ms and exposes the first criterion: divisors of 100 admit the long-period services natively and meet the deadline through offset placement, keeping $\bar{\eta}$ low (1.77 at $d = 20$), whereas non-divisors must substitute the 100 to 1000 ms periods onto tiny d -compatible periods to meet the deadline, and their over-provisioning explodes to $\bar{\eta}$ as high as 33; divisors that cannot meet the deadline at all are infeasible and omitted from the panel. The selection then follows: among divisors of 100 satisfying both criteria, $d = 20$ ms is the largest practical value, meeting $d \leq 2D_{\max} = 20$ ms with equality for the urgent service; being the largest feasible divisor, it leaves the most residue classes for offset placement while keeping the average provisioning ratio low under the halved pattern, whereas $d = 25$ and $d = 50$ preserve the catalog but violate the deadline bound and are excluded, exactly the failure that Section 5.9 later measures. Accordingly, we use $d = 20$ ms throughout the evaluation. The divisor is thus a deployment-time design output, not a universal constant: a deployment with a different expected service distribution reruns the same selection procedure and may obtain a different d , and the sensitivity of the results to this choice is examined in Section 5.9. Online adaptation of d during operation is not treated, because all admitted CGs share the divisor and changing it would require RRC reconfiguration of the active reservation population.




Figure 1. Divisor selection based on average provisioning ratio $\bar{\eta}$. (a) Normal traffic, $D_{\max} = p$. (b) Urgent traffic, $D_{\max} = 10$ ms.

4. Divisor-Aligned CG/DG Scheduling

4.1. Period Compatibility and Residue Classes

Theorem 1 (Periodic collision condition). *Two periodic reservation sequences with periods P_1, P_2 and offsets o_1, o_2 collide if and only if $\gcd(P_1, P_2) \mid (o_2 - o_1)$, and are collision-free if and only if $(o_2 - o_1) \bmod \gcd(P_1, P_2) \neq 0$.*

Proof. A collision requires $o_1 + kP_1 = o_2 + mP_2$, i.e., $kP_1 - mP_2 = o_2 - o_1$, which has an integer solution if and only if $\gcd(P_1, P_2) \mid (o_2 - o_1)$. A nonnegative solution follows by shifting along the standard solution family. $\square$

Corollary 1 (Coprime-period incompatibility). *If $\gcd(P_1, P_2) = 1$, the two sequences collide for every offset pair.*

Corollary 2 (Recurring collisions). *If an incompatible CG pair (i, j) collides, that is $g_{ij} = \gcd(P_i, P_j) \mid (o_j - o_i)$, then the collisions recur periodically with period $L_{ij} = \text{lcm}(P_i, P_j)$, so the pair collides once in every L_{ij} slots indefinitely rather than at an isolated instant.*

Theorem 1 shows that collisions are governed entirely by an arithmetic divisibility condition on the offset difference, and its coprime corollary (Corollary 1) shows that offset selection alone cannot remove heterogeneous-period collisions when the periods are coprime. Corollary 2 further shows that an admitted incompatible pair contributes a periodically recurring collision stream rather than a one-time event, so the loss keeps accumulating at every period L_{ij} for as long as both reservations remain active. DAPD therefore represents every admitted CG as a d -aligned reservation structure, a multiple of a common divisor d , reshaping rather than rejecting requests whose native period is not, which induces a residue-class geometry on the slot timeline. The following lemma states that d -compatible CGs in different residue classes modulo d are automatically disjoint, and the subsequent theorem reduces same-class coexistence to a GCD test on the normalized period multipliers, so that admission control becomes a closed-form arithmetic test rather than an exhaustive offset search.

Lemma 1 (Residue-class invariance). *If $d \mid P$, then a CG with period P and offset o occupies only one residue class modulo d : $o + kP \equiv o \pmod{d}$ for all $k \geq 0$. Two d -compatible CGs whose offsets lie in different residue classes modulo d are therefore collision-free.*

Theorem 2 (d -space coexistence). *Let $P_i = b_i d$ for $i \in \{1, 2\}$ and suppose $o_1 \equiv o_2 \pmod{d}$, so that the normalized offset difference $\Delta_d = (o_2 - o_1)/d$ is an integer. Two d -compatible CGs sharing the same residue class are collision-free if and only if $\Delta_d \bmod \gcd(b_1, b_2) \neq 0$.*

Proof. $\gcd(b_1 d, b_2 d) = d \gcd(b_1, b_2)$ and $o_2 - o_1 = d \Delta_d$ reduce Theorem 1 to $\gcd(b_1, b_2) \nmid \Delta_d$. $\square$

Corollary 3 (Identical-multiplier capacity). *Within a single residue class, at most b CGs with the same multiplier b can coexist without collision, and the bound is tight.*

Corollary 4 (Residue-class monopoly). *A CG with $q = d$ has multiplier $b = 1$ and cannot share its residue class with any other d -compatible CG.*

4.2. Period Alignment and Admission Control

The above arithmetic structure determines the scheduler design. Choosing d partitions the slot timeline into d residue classes, across which coexistence is automatic by Lemma 1,

and within which coexistence reduces to a GCD test on the normalized period multipliers by Theorem 2; Corollary 3 bounds the residue-class capacity, and Corollary 4 identifies the special case where the smallest d -compatible period $q = d$ monopolizes a residue class. This motivates the DAPD packing rule, which preserves empty residue classes for $b = 1$ requests whenever possible and packs $b > 1$ requests into already occupied but compatible classes when doing so preserves future schedulability. DAPD then separates admission into a Period Alignment Unit (PAU) and an Admission Control Unit (ACU). The PAU maps the native pair $(p, D_{\max})$ to one or more catalog-valid d -compatible candidates: if $d \mid p$ it keeps $q = p$, and otherwise it constructs two substitutes. A single-CG substitute uses opportunity spacing $q_1^* = \max\{q \in C : q \leq D_{\max}\}$ and corresponds to a single RRC-configured CG of period q_1^* . A halved-pattern substitute uses two interleaved subsequences of common period q_2^* separated by $q_2^*/2$ slots, with effective spacing $q_2^*/2$, where

$$q_1^* = \max\{q \in C : q \leq D_{\max}\}, \quad q_2^* = \max\{q \in C : q \leq 2D_{\max}\}. \quad (6)$$

The halved pattern does not request a non-catalog period $q_2^*/2$: it is implemented as two RRC-configured CGs sharing the catalog-valid period q_2^* , with offsets o and $o + q_2^*/2$ chosen by the ACU, so each subsequence is a standard 3GPP CG configuration; TS 38.331 supports multiple simultaneous sidelink CG configurations per UE through the configured grant configuration list [4], so the halved pattern requires no extension to the CG semantics. Conformance with the CG configuration semantics and the QoS constraints then holds by construction: every period configured by the PAU is drawn from $C \subset P_{3GPP}$, so DAPD never requests a periodicity value outside the standardized set, and the PAU selects the substitute so that the opportunity spacing never exceeds the packet delay budget of the PC5 QoS profile ($q_1^* \leq D_{\max}$ for the single-CG candidate and effective spacing $q_2^*/2 \leq D_{\max}$ for the halved pattern), so substitution can only tighten, never violate, the declared delay bound; the resource cost of that tightening is exactly what the provisioning ratio η of (4) quantifies and what DAPD+ reclaims. The offset separation $q_2^*/2$ is an integer at $d = 20$ ms. Since the smallest d -compatible catalog period is d , the single-CG candidate is unavailable when $D_{\max} < d$, and no d -compatible substitute exists when $2D_{\max} < d$; in that case the request is deferred with interim DG service rather than force-fitted, so no admitted configuration ever violates its declared QoS. The PAU depends only on the traffic descriptor and d and can be implemented as a constant-size lookup table indexed by $(p, D_{\max}, d)$.

The ACU places a candidate against the active reservation set $A(t)$ by scanning feasible offsets in the deadline window of the request. Given a request that arrives at slot t_0 with delay budget $D_{\max}$, the deadline window is the slot interval $[t_0, t_0 + D_{\max}]$ within which the first CG transmission must occur for the packet to be useful. The ACU scans candidate first-transmission offsets o over this window. For each candidate subsequence (o, bd) and each active subsequence $(o', b'd)$ in the same residue class, it tests

$$((o' - o)/d) \bmod \gcd(b, b') \neq 0. \quad (7)$$

The candidate is admitted only if all of its subsequences pass against all conflicting actives. If no candidate can be placed, the request is deferred: it receives interim one-shot DG service while waiting and is periodically retried for CG admission, so that once a compatible placement becomes available it is configured as a collision-free CG and leaves the DG queue. Figure 2 summarizes the scheduler, and Algorithms 1 and 2 give pseudocode. With $|K| \leq 2$ candidates per request, W feasible offsets in the deadline window, and n_r subsequences already occupying the target residue class, the admission test has complexity $O(|K|Wn_r)$, which is strictly smaller than scanning all active CGs because d partitions the timeline into d residue classes and only the local occupancy matters.

This is the operational advantage of the common-divisor restriction over an unrestricted pairwise-GCD admission rule that admits any native period as long as it passes a pairwise feasibility check against every active CG: the unrestricted rule must scan the entire active set per arrival ($O(|A(t)|)$) and, because native catalog periods need share no common structure, it fragments the timeline into placements with no residue-class locality, so a feasible placement for one request can render a later request infeasible in a way that is expensive to predict. The restriction is not merely a complexity optimization: admitting native periods directly is unsafe when the catalog contains periods that are coprime to the rest. By Corollary 1, a CG with a prime period such as 17 ms collides at every offset with any CG whose period is not a multiple of 17, and the 3GPP catalog contains no other multiple of 17; once such a reservation is admitted it blocks every subsequent request on that subchannel, so admission order alone can monopolize the resource. Period substitution averts this hazard at admission time by mapping every native period onto the divisor-aligned catalog before placement, so no coprime native period ever enters the active set. Fixing d converts coexistence into the residue-class-local test of Theorem 2 and bounds the admitted population through the capacity of Corollary 3, trading a small amount of catalog flexibility (resolved by period substitution) for a closed-form, locally checkable, and order-independent admission rule.

This residue-class locality yields a measurable admission-time advantage over admitting native periods directly. Table 3 reports the number of pairwise GCD comparisons per admission attempt, measured by an admission-only simulator that shares the arrival process, sojourn distribution, and traffic mix used in the evaluation of Section 5; one comparison is one GCD feasibility test between a candidate subsequence and one active subsequence. The unrestricted native-period rule tests each candidate against the entire active set, so its per-attempt cost scales with the active population $O(|A(t)|)$, whereas DAPD tests only the occupants of the target residue class and scales with the per-class occupancy $O(|K|Wn_r)$. The measured locality advantage is roughly $8\times$ to $17\times$ across the density range. This cost gap compounds the qualitative argument above: admitting native periods directly is not only more expensive but unsafe, because a single coprime period such as 17 ms, once admitted, blocks every later non-multiple of 17 on the subchannel, so the unrestricted rule is both order-dependent and vulnerable to resource monopoly. DAPD's substitution excludes coprime periods at admission time, making the rule cheap, residue-class-local, and order-independent.




Figure 2. DAPD scheduler architecture: admission pipeline, DG queue, retry path, and outputs.

Table 3. Admission-check cost: comparisons per CG admission attempt for the unrestricted native-period rule and for DAPD.

ρ (UEs/km)	Unrestricted Native	DAPD	Ratio
6	15.8 ± 2.0	0.9 ± 0.2	$17.3\times$
12	27.4 ± 2.2	3.1 ± 0.7	$8.8\times$
18	62.2 ± 29.0	6.9 ± 1.4	$9.0\times$
24	74.7 ± 8.0	9.0 ± 0.7	$8.3\times$
30	120.8 ± 27.5	12.2 ± 0.9	$9.9\times$

Algorithm 1 DAPD candidate generation (PAU).**Require:** Native period p , deadline $D_{\max}$, divisor d , catalog P_{3GPP}

```

1:  $C \leftarrow \{q \in P_{3GPP} : d \mid q\}$ ,  $K \leftarrow \emptyset$ 
2: if  $d \mid p$  then
3:   Add native candidate  $q = p$  to  $K$ 
4: else
5:    $q_1^* \leftarrow \max\{q \in C : q \leq D_{\max}\}$  if nonempty
6:    $q_2^* \leftarrow \max\{q \in C : q \leq 2D_{\max}\}$  if nonempty
7:   if  $q_1^*$  exists then
8:     Add single-CG candidate  $q_1^*$  to  $K$ 
9:   end if
10:  if  $q_2^*$  exists and  $q_2^*/2 \in \mathbb{Z}$  then
11:    Add halved-pattern candidate  $(q_2^*, q_2^*/2)$  to  $K$ 
12:  end if
13: end if
14: return  $K$  sorted by reservation overhead

```

Algorithm 2 DAPD admission test (ACU).**Require:** Candidate set K , active set $A(t)$, deadline window

```

1: for candidate  $K \in K$  do
2:   for feasible offset  $o$  in the deadline window do
3:     Decompose  $K$  into its subsequences  $S$  (one for single-CG, two for halved-pattern);
      $ok \leftarrow \text{true}$ 
4:     Let  $A(t)' \leftarrow A(t)$  extended with the other subsequence(s) of  $K$  {halved-pattern
     subsequences are also tested against each other}
5:     for candidate subsequence  $(o, bd) \in S$  do
6:       for subsequence  $(o', b'd) \in A(t)'$  in the same residue class do
7:         if  $((o' - o)/d) \bmod \gcd(b, b') = 0$  then
8:            $ok \leftarrow \text{false}$ 
9:         end if
10:      end for
11:    end for
12:    if  $ok$  then
13:      admit  $K$ , update  $A(t)$ ,
14:      return success
15:    end if
16:  end for
17: end for
18: return defer

```

4.3. Over-Provisioning and DAPD+ Slot-Level Scheduling

Using the provisioning ratio $\eta(p, q, m)$ of (4), the number of over-provisioned opportunities created by a non- d -compatible native period mapped to a shorter d -compatible substitute can be derived in closed form. Over a hyper-period $L = \text{lcm}(p, q)$, the substitute

reserves mL/q opportunities while the native packet-generation process produces only L/p packets, so the number of OP slots is

$$N_{\text{op}}(p, q, m) = m \frac{L}{q} - \frac{L}{p}, \quad (8)$$

independent of the phase between the packet-generation process and the reservation stream. A native-compatible single CG has $\eta = 1$ and $N_{\text{op}} = 0$; a substitute with $q < p$ or a halved pattern with $m = 2$ has $\eta > 1$, and DAPD+ records the resulting reclaimable slots at admission time.

Proposition 1 (Collision-free invariant). *If every newly admitted CG passes (7), with a halved pattern treated as two interleaved subsequences, the active set $A(t)$ remains pairwise collision-free for all future repetitions.*

Proof. By induction over admissions. The active set is initially empty; at each admission, every new subsequence is checked against every active subsequence in the same residue class, and, for a halved pattern, the two new subsequences are also checked against each other; different residue classes are disjoint by Lemma 1, so Theorem 2 guarantees that no new collision is introduced. Departures only remove reservations. $\square$

DAPD+ runs the slot-level rule in Algorithm 3. The DG salvage mechanism is scheduler-level and cross-UE: the DG-candidate packet pool consists of pending packets of pre-admission or waiting UEs together with native DG packets, not of admitted UEs reusing their own unused CG opportunities. The DG-available resource pool at slot t is the set of slots not currently occupied by an admitted CG transmission, comprising (i) unreserved slots that no admitted CG holds and (ii) reclaimable over-provisioned slots held by an admitted CG whose owner has no packet due at slot t . Admitted CG transmissions have strict priority on every slot they occupy, so DAPD+ never overlaps an admitted CG transmission and the collision-free invariant of Proposition 1 is preserved with admitted-CG PDR equal to one at the reservation level under the collision-only scheduler model.

Algorithm 3 DAPD+ slot-level scheduling (single subchannel).

Require: Slot t , admitted-CG set $A(t)$, DG queue $Q(t)$ { $Q(t)$ holds packets that have arrived by slot t and are not yet expired, from pre-admission/waiting UEs and native DG sources}

```

1: if some admitted CG in  $A(t)$  has a packet due on slot  $t$  then
2:   Schedule that admitted CG on slot  $t$  {strict priority}
3: else
4:   Slot  $t$  is DG-available (either unreserved or an idle CG opportunity)
5:   Remove from  $Q(t)$  any packet whose deadline is earlier than  $t$ 
6:   if  $Q(t) \neq \emptyset$  then
7:      $x^* \leftarrow$  earliest-deadline-first (EDF) packet in  $Q(t)$ 
8:     Schedule  $x^*$  by DG on slot  $t$ 
9:     if  $x^*$  is native DG and slot  $t$  is an idle CG opportunity then
10:      Increment  $N_{\text{op} \rightarrow \text{DG}}$ 
11:     end if
12:   Remove  $x^*$  from  $Q(t)$ 
13:   end if
14: end if

```

5. Performance Evaluation

We simulate one gNB serving a 2 km road segment. UEs arrive as a Poisson process and traverse the segment at a speed drawn from a truncated normal distribution with

a mean of 100 km/h and a standard deviation of 10 km/h, clipped to [70, 130] km/h, giving a mean sojourn time of $\bar{T}_{\text{soj}} = 72$ s. Density varies from 3 to 30 UEs/km; beyond this range the offered load saturates the single-subchannel pool and the regime becomes one of congestion control rather than admission, which is outside the scope of this paper. In the CG-dominant setting each UE requests one CG reservation, with native period and packet delay budget (PDB) drawn jointly from Table 4; the mix is a representative synthetic combination of advanced V2X services and basic safety messaging, consistent with 3GPP service requirements [3], with the 100 ms period dominant at 45%. It is not a measurement trace, but it is standard-derived in its components: every row of Table 4 is anchored to a documented 3GPP or industry service definition, from the advanced V2X sets of TS 22.186 to the CAM/BSM generation rules and 5GAA evaluation traffic listed in the last column. A validation against a measured trace is, moreover, not yet possible: to the best of our knowledge, no public trace of a heterogeneous NR sidelink service mix is available, and C-V2X has not yet seen a large-scale deployment, so the empirical traffic distribution of mixed advanced services is not yet known to the community; deployed V2X traffic remains confined to single-service periodic safety messaging (CAM/BSM), corresponding to the 100 ms row that already dominates the mix, and standard-derived synthetic mixes are therefore the common evaluation practice. The mix intentionally includes non-divisor periods such as 17 and 33 ms to stress heterogeneous catalog compatibility, so it doubles as an arithmetic stress test for the admission rule. Under this mix 78% of periods are d -compatible at $d = 20$ ms and 22% require a PAU substitute; the GCD guarantee holds for any mix, and robustness to the divisor choice, which also varies the compatible/coprime composition of the offered traffic, is examined in Section 5.9. We use $T_{\text{sim}} = 300,000$ ms ($\approx 4\bar{T}_{\text{soj}}$) and discard the first $T_{\text{warmup}} = 72,000$ ms as warmup. We compare seven schedulers. Min-Collision (MC), OPT-offset, and One-shot (OS) are non-blocking baselines that keep native periods: MC picks the offset with fewest overlaps against already-placed neighbors; OPT-offset is an offline oracle that, knowing every arrival and departure in advance, jointly optimizes all UEs' offsets within their deadline windows by coordinate descent on the total pairwise collision cost, initialized at the MC solution, so it improves monotonically on MC and converges to a local optimum that serves as a near-upper bound of the offset-selection family, heuristic, graph-based, or learned; and One-shot deflects a collided packet once, deterministically, to the nearest idle resource within its deadline, a deterministic variant of the stochastic one-shot transmission prescribed by SAE J3161/1 [22,23]. The three represent the offset-search and reactive-repair family, which by Corollary 1 cannot eliminate coprime collisions, so its collapse isolates the structural limit the GCD rule overcomes. C-FAST ports the configured grant scheduler of Larrañaga-Zumeta et al. [5], to our knowledge the closest prior work, from the cellular uplink to sidelink Mode 1: the gNB enumerates all collisions over the hyperperiod at admission and pre-assigns each flow up to several CGs with per-packet activation vectors, so that packets that would collide are statically steered to avoidance configurations; the number of configurations per UE is capped at eight to respect the sidelink limit ($\text{maxNrofCG-SL-r16} = 8$, TS 38.331), against twelve on the uplink, and a request that cannot be accommodated within eight configurations is rejected. GCD-only applies the test of (7) but omits substitution, rejecting d -incompatible requests. DAPD adds PAU substitution, and DAPD+ adds the slot-level DG mechanism of Algorithm 3. DAPD+ is the complete evaluated variant; GCD-only and DAPD are progressive realizations that isolate each layer. Table 5 lists default parameters.

Table 4. V2X periodic-traffic distribution: native period, sampling probability, PDB, and representative application.

Period (ms)	Probability (%)	PDB (ms)	Representative Application/Traffic Model
10	3	10	Cooperative collision avoidance; high-density platooning (advanced V2X Set 2, 100 Hz)
17	3	17	Public-safety MCVideo over sidelink relay (1080p, 60 fps)
20	3	20	Platooning cooperative driving and information sharing (50 msg/s)
25	5	25	Normal-density platooning (advanced V2X Set 1, 40 Hz)
33	7	25	Platooning cooperative driving (30 msg/s)
50	4	50	Periodic group-start and cooperative maneuver coordination
100	45	100	CAM/BSM 10 Hz baseline; CPM/VAM minimum generation interval; automated-driving info sharing (10 msg/s)
200	10	100	CAM baseline traffic in 5GAA NR-V2X highway/urban evaluation
500	5	500	Platooning status reporting (2 msg/s)
600	5	100	BSM congestion-control maximum inter-transmission interval (vMaxITT)
1000	10	100	CAM maximum generation interval (1 Hz lower bound)

We evaluate the schedulers under the following metrics. The PDR, the fraction of generated packets delivered within their deadline windows, is measured at three scopes:

- OverallPDR: Taken over all generated packets, including those a UE produces before obtaining a CG reservation, so it falls below one even when admission is perfect.
- Admitted-CG PDR_{adm} : Restricted to packets generated after a UE is admitted as a CG. This is the reservation-level invariant, and every GCD-safe scheduler (GCD-only, DAPD, DAPD+) holds it at one by construction. In all delivery ratios, packets generated by a UE that exits the scenario before its next transmission opportunity are excluded from the accounting, since no scheduler could have served them.
- Per-UE: For each UE, admitted or not, the PDR over its entire packet stream from arrival to departure. This includes the pre-admission packets generated while a UE waits for a reservation and the full stream of a UE that is never admitted, whose only delivered packets are those rescued by DG salvage. Counting unadmitted UEs (which enter at a per-UE PDR near zero) avoids survivorship bias. It is reported as a survival curve $\mathbb{P}\{\text{per-UE PDR} \geq \theta\}$ against a reliability target θ (e.g., 0.95 or 0.99 for stringent applications), which reveals whether reliability is shared across the population or concentrated in a fortunate subset. This scope differs from the admitted-CG PDR above, which counts only post-admission packets of admitted UEs; a scheduler can hold $\text{PDR}_{\text{adm}} = 1$ yet fall well below target here if it loses pre-admission packets or leaves UEs unadmitted.

The remaining metrics characterize timing and admission behavior:

- Admitted-CG jitter: The within-UE standard deviation of the per-packet queueing delay (generation epoch to serving slot), averaged over admitted UEs and taken over delivered packets only, since an undelivered packet has no defined delay. It captures the timing variation introduced by period substitution in DAPD and by reactive deflection in One-shot relative to a fixed native offset.
- Admission rate: The long-run fraction of native CG requests eventually placed as a collision-free CG, reported both at the first attempt (initial) and at any point during a UE's sojourn (eventual, counting requests admitted later via retry once a compatible slot frees up).
- Admission delay: The time from a request's arrival to its successful CG admission, reported as a per-density CDF across UEs to expose the latency cost of the strict admission rule.

- DG-salvage PDR PDR_{salv} (Equation (5)): The fraction of pre-admission CG packets that DG salvage delivers within the delay budget, a diagnostic for the DAPD+ salvage layer reported in the ablation of Section 5.7.

Together these decompose performance into packet-level delivery, reservation-level reliability, per-UE fairness, timing variation, and admission behavior. We additionally report an ablation of the DAPD+ mechanisms in Section 5.7.

Table 5. Default simulation parameters.

Parameter	Value
Common divisor d	20 ms
Slot duration	1 ms
Road length	2 km
Speed (km/h)	truncated $\mathcal{N}(100, 10)$ on $[70, 130]$
UE density ρ	3 to 30 UEs/km
Native period and PDB	Table 4
Simulation time T_{sim}	3×10^5 ms ($\approx 4\bar{T}_{soj}$)
Warmup time T_{warmup}	72,000 ms ($=\bar{T}_{soj}$)

5.1. Overall Packet PDR

We evaluate overall packet PDR in two scenarios: a CG-dominant scenario with density swept from 3 to 30 UEs/km, and a mixed CG/DG scenario at $\rho = 15$ UEs/km where a native one-shot DG stream is added so that $\gamma = \text{CG load} / (\text{CG load} + \text{DG load})$ sweeps from 0.2 to 1.0. The upper end of the density range is set by resource saturation: as reported in Section 5.2, the total channel occupancy of DAPD+ reaches about 0.86 at $\rho = 27$ and 0.94 at $\rho = 30$, so beyond this range the single subchannel is essentially full and the operating regime becomes one of congestion control rather than admission, which is outside the scope of this paper. Figure 3 reports the two scenarios; curves carry 95% confidence intervals over 1000 seeds, mostly smaller than the markers. In the CG-dominant case (Figure 3a), DAPD+ is the upper envelope across the entire sweep: it stays at 1.00 through $\rho = 12$, at 0.9997 at $\rho = 15$, and degrades most gracefully thereafter (0.994 at $\rho = 21$, 0.96 at $\rho = 27$, 0.92 at $\rho = 30$). One-shot matches it only at light load (at or above 0.999 for $\rho \leq 12$) and then falls away as idle deflection slots grow scarce, to 0.96 at $\rho = 21$ and 0.77 at $\rho = 30$. C-FAST tracks between the two through moderate load (0.99 at $\rho = 12$, 0.97 at $\rho = 15$) but drops to 0.73 at $\rho = 30$: its admitted reservations remain protected (Section 5.2), but it has no salvage path, so pre-admission packets and the streams of rejected UEs are lost, and its multi-CG avoidance consumes pool capacity that DAPD+ instead returns to delivery. Min-Collision is the lower envelope, falling from 0.97 at $\rho = 3$ to 0.57 at $\rho = 30$ as deterministic recurrence (Corollary 2) defeats naive offset selection. OPT-offset, the offline oracle of the offset family, improves on MC by at most 0.05 anywhere in the sweep (0.62 vs. 0.57 at $\rho = 30$) despite jointly optimizing every offset with full knowledge of arrivals and departures: by Corollary 1 no offset placement exists for the coprime pairs in the mix, so the shortfall it retains measures the feasible set rather than the search quality, and the near-coincidence of the oracle with the greedy heuristic is itself the empirical signature of an empty feasible set. DAPD tracks below DAPD+ (from 1.00 at $\rho = 3$ to 0.62 at $\rho = 30$); the widening gap between them, about 0.26 to 0.30 for $\rho \geq 21$, is the contribution of DG salvage, which grows as the waiting population that salvage serves expands. Under the heaviest loads in range, DAPD+ preserves admitted-CG reliability (Section 5.2) but its overall PDR still declines, because once the admitted population saturates the single-subchannel pool the offered load exceeds capacity and the residual loss is pre-admission and unadmitted traffic. In the mixed scenario (Figure 3b),

the CG fraction γ splits the range into two regimes. For $\gamma \leq 0.4$ the native DG traffic dominates and all six schedulers nearly coincide (0.39 to 0.41 at $\gamma = 0.2$ and 0.77 to 0.81 at $\gamma = 0.4$): single-shot packets with a 20 ms deadline are served alike by every scheduler, so the CG discipline carries little aggregate weight, though the substitution-based group already sits visibly above the offset family. Beyond $\gamma = 0.6$ the schedulers separate into the same order as the density sweep: DAPD+ leads throughout and reaches 0.9998 at $\gamma = 1.0$, One-shot follows at 0.996, and C-FAST plateaus near 0.98. Min-Collision, OPT-offset, and DAPD are non-monotonic: they peak at $\gamma = 0.6$ (0.88 to 0.91) and decline to 0.81 to 0.85 at $\gamma = 1.0$, because a growing CG fraction means more deterministic collisions for the offset family and more pre-admission losses for salvage-free DAPD, while the idle DG capacity that delivered their remaining traffic shrinks away. That Min-Collision and DAPD end close together on this aggregate metric is informative rather than coincidental: their losses arise from different mechanisms of comparable magnitude, recurring collisions for the former and pre-admission drops for the latter, so the aggregate PDR cannot separate a scheduler with no reservation guarantee from one whose admitted reservations are collision-free; that separation appears only in the admitted-CG PDR of Section 5.2 and the per-UE reliability of Section 5.3. The mixed scenario thus shows that the salvage layer, not the admission rule alone, is what lets DAPD+ dominate even the strongest reactive baseline whenever spare DG capacity exists.




Figure 3. Overall packet PDR versus load. (a) CG-dominant scenario, density sweep. (b) Mixed scenario, CG-fraction sweep.

5.2. Admitted-CG PDR

The overall PDR mixes pre-admission waiting losses with admitted-CG losses. To isolate the admission-rule guarantee from the waiting-phase effect, Figure 4 reports the admitted-CG PDR PDR_{adm} across the CG-dominant density sweep; this metric is the reservation-level invariant under the collision-only scheduler model. DAPD and DAPD+ maintain $\text{PDR}_{\text{adm}} = 1.000$ at every density tested (the two coincide on this metric and are plotted as one curve), validating the admission-rule invariant of Proposition 1 in the packet-level simulator: once a CG reservation passes the divisor-aligned GCD admission test, it is collision-free for the remainder of its sojourn. C-FAST also protects admitted reservations, holding PDR_{adm} at 1.000 across the sweep as well: pre-planted activation-vector configurations are an effective avoidance mechanism, and their cost appears not here but in the overall PDR of Figure 3a, the configuration footprint of up to eight CGs per UE, the hyperperiod-scale conflict enumeration at admission, and the growing jitter of Section 5.4. The remaining schedulers fall below one because none enforces a coexistence test, but they fall differently. One-shot, which deflects a collided admitted-CG packet once within its deadline, recovers much of the loss at low-to-medium density: its admitted-CG

PDR stays at or above 0.999 for $\rho \leq 12$ and declines only as idle deflection slots grow scarce, to 0.96 at $\rho = 21$, 0.85 at $\rho = 27$, and 0.77 at $\rho = 30$. Min-Collision, with no repair, falls monotonically from 0.97 at $\rho = 3$ to 0.57 at $\rho = 30$. OPT-offset barely improves on it, reaching only 0.62 at $\rho = 30$, a gain of at most 0.05 anywhere in the sweep: even the jointly optimized oracle placement over the entire offset space cannot hold the admitted-CG PDR at one, because for coprime-period pairs no collision-free placement exists to find (Corollary 1), which empirically confirms that the limitation of this family is the feasible set defined by the period arithmetic, not the quality of the search. The difference between the GCD-enforcing schedulers and the offset family on this metric is thus one of guarantee, not degree: DAPD and DAPD+ hold PDR_{adm} at one because admission is collision-free by construction, C-FAST attains the same protection by spending configurations rather than by reshaping the period arithmetic, whereas One-shot only patches collisions reactively and Min-Collision and OPT-offset do not address them at all, so the offset family falls as the pool fills. The same simulation also measures the total channel occupancy of DAPD+ (admitted-CG transmissions plus DG-salvage transmissions, with collided slots counted as occupied), which serves as the resource-efficiency accounting of the evaluation: at fixed offered load it ties the delivered PDR directly to spectral efficiency on the single subchannel, and it is read jointly with the over-provisioning cost η of (4) and the slack reuse efficiency η_{reuse} . Although we omit the plot for space, the measured occupancy rises from 0.11 at $\rho = 3$ to 0.51 at $\rho = 15$, 0.70 at $\rho = 21$, 0.86 at $\rho = 27$, and 0.94 at $\rho = 30$, confirming that the studied range ends just as the subchannel saturates.




Figure 4. Admitted-CG (reservation-level) PDR PDR_{adm} across the CG-dominant density sweep.

5.3. Per-UE Reliability

Aggregate PDR does not reveal whether losses concentrate on a few UEs or spread across the population. Figure 5 reports per-UE reliability at thresholds 0.95 and 0.99 across the density sweep. To avoid survivorship bias, the per-UE population includes every UE, not only those that obtain a CG: a UE that departs without ever being admitted contributes its full arrival-to-departure packet count to the denominator, with only its DG-salvaged packets (if any) counted as delivered, so an unadmitted UE that sends nothing correctly enters the distribution at a per-UE PDR of zero. Under this accounting, DAPD+ keeps the largest share of UEs above the 0.95 target at every density, staying at 1.00 through $\rho = 12$ and declining only gently to 0.94 at $\rho = 27$ and 0.90 at $\rho = 30$ (0.87 at the 0.99 threshold), because its salvage path rescues the pre-admission packets of both late-admitted and never-admitted UEs. C-FAST is the strongest baseline on this metric: it stays within a few points of DAPD+ throughout (0.98 at $\rho = 15$, 0.93 at $\rho = 21$, 0.88 at $\rho = 30$ for $\theta = 0.95$), since its admitted UEs are fully protected and its rejections are few; at the strictest corner ($\theta = 0.99$, $\rho = 30$) it even edges ahead by 0.007 (0.872 vs. 0.865), at the limit of the confidence intervals, while DAPD+ retains the advantage at $\theta = 0.95$ and everywhere below $\rho = 30$ and pairs its per-UE robustness with a far higher overall PDR (Figure 3a).

DAPD trails (0.93 at $\rho = 15$ falling to 0.82 at $\rho = 30$): without salvage, the pre-admission packets of deferred UEs and the entire stream of never-admitted UEs are lost. One-shot holds up to moderate density (0.99 at $\rho = 15$) but collapses beyond it, to 0.89 at $\rho = 21$ and 0.41 at $\rho = 30$ (0.28 at $\theta = 0.99$), because reactive deflection cannot break the recurrence of Corollary 2 and a growing fraction of UEs is effectively starved. Min-Collision and OPT-offset give almost no per-UE protection: both fall from 0.79 at $\rho = 3$ to below 0.06 by $\rho = 21$ and near 0.01 at $\rho = 30$, and the oracle is indistinguishable from the greedy heuristic on this metric, the per-UE restatement of Corollary 1. The per-UE view reinforces the aggregate result while exposing what the aggregate hides: once UEs that never obtain service are counted, only reservation-level protection, by GCD admission or by C-FAST's configuration spending, keeps reliability spread across the population, and the salvage layer then determines how much of the remaining stream is delivered.




Figure 5. Per-UE reliability versus density. (a) Threshold $\theta = 0.95$. (b) Threshold $\theta = 0.99$.

5.4. Admitted-CG Jitter

The reservation-level guarantee is collision-free but not jitter-free: because DAPD maps a native period p to a d -compatible substitute $q^* \neq p$, the spacing between a packet's generation epoch and its serving CG slot varies, whereas a scheduler that keeps the native period and never relocates a packet serves every packet at a fixed offset. Figure 6 reports the admitted-CG jitter, the within-UE standard deviation of this queueing delay averaged over UEs. Min-Collision and OPT-offset have identically zero jitter: both keep the native cadence and never relocate a packet, so every admitted-CG packet is served at its fixed offset. One-shot also keeps the native period, but its reactive deflection relocates each collided packet to a later idle slot, and this relocation injects timing variation that grows with density as collisions become frequent: its jitter rises from 0.2 ms at $\rho = 3$ to 1.5 ms at $\rho = 15$, 3.9 ms at $\rho = 21$, and 9.0 ms at $\rho = 30$, overtaking DAPD+ beyond $\rho \approx 11$. DAPD and DAPD+ have identical jitter, since DG salvage does not alter admitted-CG timing, and it stays flat near 0.72 to 0.78 ms across the entire range, well over an order of magnitude below the tightest delay budget of Table 4. C-FAST starts lowest (0.09 ms at $\rho = 3$) but grows steadily to 0.99 ms at $\rho = 30$, crossing DAPD+ near $\rho = 21$: each avoidance configuration serves its packets at a different offset, so as load pushes more packets onto secondary configurations the within-UE timing spread widens. The substitution-induced jitter of DAPD is therefore bounded and density-insensitive, the configuration-multiplexing jitter of C-FAST is bounded but rising, and the deflection-induced jitter of One-shot is unbounded in density; the divisor-aligned structure trades a small fixed timing offset for cheap, order-independent admission, while reactive repair pays an increasing timing penalty exactly where load is heaviest.




Figure 6. Admitted-CG jitter across the density sweep; DAPD and DAPD+ coincide.

5.5. Admission Rate

Figure 7 characterizes admission for the blocking schedulers, with MC shown as the non-blocking reference at a rate of one (one-shot and OPT-offset are likewise non-blocking and omitted); DAPD+ shares DAPD's admission rule exactly. Figure 7a plots the simulated eventual admission rate, the long-run fraction of requests configured as a collision-free CG at any point during their sojourn, with 95% confidence intervals. DAPD admits nearly all requests at the lightest densities (eventual rate 1.00 at $\rho \leq 6$) and declines gracefully to 0.95 at $\rho = 30$ as the pool fills. GCD-only, which rejects d -incompatible requests outright, follows a different admission policy and sits well below DAPD throughout, flat near 0.78: the incompatible 22% of requests never enter its admitted population, whereas DAPD substitutes them onto d -compatible structures and admits them, raising the admitted volume; the flat curve matches the 78% natively compatible fraction of the mix. C-FAST admits nearly everything as well (0.9998 at $\rho = 3$, 0.985 at $\rho = 24$, 0.964 at $\rho = 30$), rejecting only flows that cannot be accommodated within eight configurations; at the top of the range it admits slightly more than DAPD (0.964 vs. 0.948), but each admission carries a multi-CG footprint, and Figure 3a shows that admitted volume does not translate into delivered packets without a salvage path. The gap between DAPD and GCD-only is therefore the direct contribution of period substitution to admission coverage. Figure 7b decomposes DAPD's eventual rate into the first-attempt (initial) component and the retry contribution: the initial component dominates throughout, declining from 1.00 at $\rho = 3$ to 0.81 at $\rho = 30$, while the retry contribution grows steadily to 0.14 at $\rho = 30$, so at the densest load about one admitted reservation in seven is secured on a retry after a compatible slot frees up.




Figure 7. Simulated admission rate versus UE density. (a) Eventual admission rate. (b) Initial vs. retry breakdown for DAPD.

5.6. Admission Delay

The eventual admission rate alone does not reveal how long a deferred UE waits before being configured as a CG. Figure 8 reports the CDF of admission delay for DAPD+ at $\rho \in \{6, 12, 18, 24, 30\}$ UEs/km, conditioned on UEs that are eventually admitted. The CDFs are mixtures of an atom at zero delay (first-attempt admissions) and a continuous tail (retry successes). At light load the atom dominates: 0.995 of admitted UEs at $\rho = 6$ are admitted immediately, and 0.95 at $\rho = 12$. As density rises the immediate fraction declines gently to 0.90 at $\rho = 18$ and 0.85 at $\rho = 30$, while the retry tail lengthens: the 95th percentile of the delay grows from zero at $\rho = 12$ to 16 s at $\rho = 18$, 28 s at $\rho = 24$, and 37 s at $\rho = 30$, the 99th percentile reaches 57 s at $\rho = 30$ but stays within the 72 s mean sojourn, and the conditional mean among retried UEs grows from about 21 s at $\rho = 12$ to 29 s at $\rho = 30$. The trend is therefore that most UEs are still admitted on the first attempt even at the densest load, while the waiting time of the deferred minority increases steadily with density. Two effects keep this admission delay from eroding DAPD+'s per-UE reliability. First, the immediate-admit fraction stays high even at the densest load (0.85 at $\rho = 30$), so most UEs carry few or no pre-admission packets and reach a per-UE PDR of one outright. Second, for the deferred minority, DAPD+ keeps serving their packets through DG one-shot transmissions on unoccupied resources throughout the waiting interval, salvaging the pre-admission stream rather than dropping it. Together these limit the per-UE reliability decline reported in Section 5.3 even as admission delay grows with density, whereas DAPD, which lacks the salvage path, loses the pre-admission packets of late-admitted UEs and the entire stream of never-admitted UEs, and consequently trails DAPD+ on per-UE delivery. UEs that are never admitted are not represented in this CDF: they appear in Figure 7 as the gap between the admission rate and one, and the DG-salvage measurements in Section 5.7 quantify how their pre-admission packets fare. DAPD shows near-identical behavior because both schedulers use the same admission rule and the same retry mechanism; for clarity, only DAPD+ is shown.

Because the CDF tail lengthens with density, one might worry that the eventual rate is inflated by UEs admitted with almost no sojourn left; this is not the case: even at the densest load, 0.995 of eventually admitted UEs retain 10 s or more of remaining sojourn after admission and 0.97 retain 30 s or more, so admitted reservations are typically usable for a substantial fraction of the UE's residence, with DAPD+ providing interim DG service over the wait.




Figure 8. CDF of admission delay for DAPD+ across the density sweep, conditioned on UEs that are eventually admitted.

5.7. Ablation

To isolate the contribution of each mechanism in the progressive realization, we run five variants at $\rho = 15$ UEs/km and $\gamma = 0.8$: GCD-only (admission test, no substitution and no salvage), DAPD (admission test plus PAU substitution, no salvage), DAPD^{idle-only}

(substitution plus native DG service on idle CG opportunities, no pre-admission salvage), $\text{DAPD}^{\text{pre-only}}$ (substitution plus pre-admission CG salvage, no native DG service), and full DAPD^+ (substitution plus both salvage paths). All five use the verified admission rule, so the reservation-level admitted-CG PDR is exactly one in every variant under the collision-only model. Figure 9 reports overall PDR. The two layers contribute in sequence. Substitution is the first layer: it raises overall PDR from 0.33 for GCD-only to 0.68 for DAPD , an increment of 0.35, because GCD-only rejects the d -incompatible 22% of requests outright while substitution admits them as d -compatible CGs; this is also why GCD-only, although it never suffers the coprime-blocking hazard of admitting native periods, cannot serve the incompatible traffic at all and is included only as the substitution-free control. Neither of these two variants serves the aperiodic native DG stream, which at $\gamma = 0.8$ is one packet in five, so both sit well below their mixed-scenario counterparts. DG salvage is the second layer: it raises overall PDR from 0.68 for DAPD to 0.998 for full DAPD^+ , an increment of 0.31, and the intermediate variants price its two paths separately: native DG service alone ($\text{DAPD}^{\text{idle-only}}$) recovers +0.20, reaching 0.88, and pre-admission CG salvage alone ($\text{DAPD}^{\text{pre-only}}$) recovers +0.12, reaching 0.80. The two paths are almost exactly additive ($0.20 + 0.12$ against the joint 0.31) because they serve disjoint packet populations, the aperiodic native DG stream and the waiting CG stream, and at this operating point the idle-slot supply accommodates both. The decomposition is also consistent across figures: $\text{DAPD}^{\text{idle-only}}$ at 0.88 coincides with the DAPD baseline of the mixed scenario at $\gamma = 0.8$ (Figure 3b), which serves the native DG stream by scenario convention, and full DAPD^+ matches its mixed-scenario value of 0.998.

To characterize the regime in which pre-admission CG salvage is effective, we also measure its conditional success rate PDR_{salv} defined in (5), the fraction of pre-admission CG packets that DAPD^+ actually delivers via idle slots. At low-to-moderate density every such packet is rescued: the success rate is 1.00 through $\rho = 15$. As density grows the rescue rate falls as idle slots become scarce, to 0.97 at $\rho = 21$, 0.87 at $\rho = 27$, and 0.78 at $\rho = 30$, because the volume of pre-admission packets grows faster than the supply of idle slots within their delay windows. Pre-admission salvage is therefore most effective when the waiting population is small, which is precisely the regime where the strict admission rule already keeps admitted-CG PDR at one; as the pool approaches saturation the binding reliability lever is the admission rule itself, not DG salvage.




Figure 9. Ablation of substitution and salvage at $\rho = 15$ UEs/km, $\gamma = 0.8$.

5.8. Multi-Subchannel Operation

The incompatibility addressed by DAPD is a per-subchannel property (Lemma 1), so the test composes across K subchannels: each runs Algorithm 2 on its own residue-class state, and the gNB places each UE on the least-loaded subchannel whose test passes. To isolate the effect of subchannelization from a change in load, we sweep K from 1 to 8 while scaling the total density in proportion to K , so that the per-subchannel offered load is

held fixed and only the number of independent subchannels grows. We report the per-UE reliability at the stringent threshold $\theta = 0.99$ at two per-subchannel load levels: a moderate level, total $\rho = 15K$ UEs/km (Figure 10a), and a heavy level, total $\rho = 25K$ UEs/km (Figure 10b). Because the per-subchannel load is held constant, the reliability of each scheduler is essentially flat across K in both figures, which is the intended check: the per-subchannel GCD test composes without degradation as subchannels are added, so the relative ordering is set by the per-subchannel load, not by K . Only One-shot drifts mildly upward with K , as pooling more subchannels smooths the availability of idle deflection slots; every other scheduler is flat within its confidence interval. At the moderate level (Figure 10a), DAPD+ holds the fraction of UEs meeting $\theta = 0.99$ at 0.999 to 1.00 for every K ; One-shot and C-FAST sit near 0.97 to 0.99, DAPD near 0.92, and Min-Collision and OPT-offset stay near zero. At the heavy level (Figure 10b), the families separate sharply: DAPD+ still holds 0.93 to 0.95 and C-FAST 0.89, but One-shot collapses to 0.57 to 0.63 while DAPD settles near 0.83, because at a saturating per-subchannel load reactive deflection can no longer find idle slots, whereas pre-planted protection, whether by GCD admission or by C-FAST's configurations, is unaffected; Min-Collision and OPT-offset are at the floor, the per-subchannel restatement of Corollary 1. The reservation-level advantage of DAPD+ is therefore a per-subchannel property that widens as per-subchannel load grows and is preserved under constant-load subchannelization rather than being an artifact of any particular K . This is not a complete multi-subchannel evaluation: it only verifies that the per-subchannel test composes under independent subchannel choice. Half-duplex coupling, a first-order NR sidelink constraint, is not modeled and would require an outer placement layer coordinating transmit and receive across subchannels. Concretely, such a layer decides which transmissions may share a slot across subchannels. Designing it is an open problem in its own right, and we do not claim a definitive mechanism here; one example is space division, in which the same slot is reused on different subchannels only by transmitter-receiver pairs whose neighborhoods are geographically disjoint, in the spirit of the adaptive spatial-reuse scheduling of Sempere-García et al. [14], so that no UE is scheduled to receive in a slot in which it transmits; under this example the layer serializes transmissions across subchannels within a single mutual-reception neighborhood, and the multi-subchannel gain is then harvested across, rather than within, neighborhoods. The layer operates on top of, and orthogonally to, the per-subchannel GCD condition: it only decides which transmissions may be concurrent, while the period arithmetic within each subchannel is unchanged; it determines cross-subchannel performance and is left as future work.




Figure 10. Per-UE reliability ($\theta = 0.99$) vs. subchannels K at fixed per-subchannel load. (a) Moderate load, total $\rho = 15K$. (b) Heavy load, total $\rho = 25K$.

5.9. Sensitivity to Common Divisor d

The choice $d = 20$ ms used throughout the evaluation was motivated in Section 3.4 by catalog preservation and tight-deadline feasibility. To verify what the divisor choice does and does not affect, we repeat the CG-dominant experiment with $d \in \{10, 20, 25, 50\}$ ms at two operating points, $\rho = 15$ UEs/km (admission-light) and $\rho = 30$ UEs/km (near saturation). The d -compatible catalog fraction under Table 4 is 0.85, 0.78, 0.84, and 0.79 for $d = 10, 20, 25$, and 50 respectively, so the four divisors cover a ± 10 pp range around the 0.78 value of the default. Because varying d changes how much of the traffic mix is natively d -compatible, this sweep also serves as a proxy sensitivity test for the compatible/coprime composition of the offered traffic. The four divisors, however, do not all satisfy the design constraints: $d = 25$ and $d = 50$ violate the tight-deadline feasibility bound $d \leq 2D_{\max}$ for the 10 ms delay-budget flows, so under either divisor the PAU has no admissible substitute for an urgent request and must reject it outright. Figure 11 reports overall packet PDR, and the sweep is informative in both directions. Within the feasible region ($d = 10, 20$), the results are insensitive to the divisor: DAPD moves by about 0.02 (0.846 vs. 0.855 at $\rho = 15$) and DAPD+ is unchanged. Outside it, salvage-free DAPD is penalized structurally rather than statistically: its overall PDR drops from 0.855 to 0.626 at $\rho = 15$ and from 0.620 to 0.580 at $\rho = 30$, with $d = 25$ and $d = 50$ coinciding because the same urgent flows are excluded in both cases. DAPD+, by contrast, is flat to the third decimal at every divisor (0.9997 to 0.9998 at $\rho = 15$ and 0.923 to 0.924 at $\rho = 30$): whatever admission rejects, the salvage layer serves, at the delivery cost already quantified in the ablation. Admitted-CG PDR remains identically one at every d , as expected from Proposition 1. The sweep therefore does not merely show robustness; it shows that the offline selection procedure of Section 3.4 is what stands between a safe divisor and a structurally penalized one: a divisor chosen by the procedure ($d \mid 100$ and $d \leq 2D_{\max}$, which together select $d = 20$) is representative of the feasible design region, and a divisor violating the deadline bound fails in exactly the way the constraint predicts.




Figure 11. Sensitivity of overall packet PDR to common divisor d at $\rho = 15$ and $\rho = 30$ UEs/km.

5.10. Limitations

This study has five limitations, consolidated here for clarity. First, the reliability guarantee is reservation-level: the collision-only scheduler model excludes PHY decoding errors, HARQ, control-channel loss, and fading, so the reported delivery ratios upper-bound end-to-end performance, as formalized in the scope discussion of the system model. Second, the admission analysis is per-subchannel; the multi-subchannel evaluation of Section 5.8 verifies that the per-subchannel test composes under independent subchannel choice, but half-duplex coupling across subchannels is not modeled and requires an outer placement layer that decides which transmissions may share a slot across subchannels; designing that layer is itself an open problem for which this paper claims no solution, and space-division spatial reuse, in which a slot is reused across subchannels only by geo-

graphically disjoint transmitter-receiver pairs, is one example, as elaborated in Section 5.8. Third, the traffic mix of Table 4 is a standard-derived synthetic stress mix rather than a measured trace, and its coprime-rich composition is intentionally adversarial to the admission rule. Fourth, the common divisor d is selected offline from the expected service mix; online adaptation of d to a drifting service distribution would require RRC reconfiguration of the active CG population and is not treated. Fifth, the scheduler consumes traffic descriptors reported through the sidelink UE assistance procedure and does not model reporting errors beyond the robustness argument of the system model, under which a misdeclared pattern wastes only the affected reservation's own opportunities.

6. Conclusions

Heterogeneous CG periods in NR sidelink Mode 1 can be deterministically incompatible, causing recurring reservation collisions even under centralized scheduling. DAPD prevents these at admission time by mapping requests to d -compatible catalog-valid CG structures and applying a closed-form GCD coexistence test, and DAPD+ adds cross-UE DG salvage without relaxing this rule. The critical quantitative outcomes are as follows. Across the studied density range of 3 to 30 UEs/km, DAPD and DAPD+ hold the admitted-CG packet delivery ratio at exactly one at every density under the collision-only model, whereas the reactive One-shot baseline falls to 0.77, Min-Collision to 0.57, and even the oracle OPT-offset placement to 0.62 at 30 UEs/km. DAPD+ simultaneously leads all baselines on overall packet delivery (0.92 versus 0.77 for One-shot and 0.73 for the multi-configuration C-FAST scheduler at 30 UEs/km) and keeps 0.90 of all UEs, admitted or not, at or above a 0.95 per-UE delivery target where One-shot retains only 0.41. The admission rule is also cheap: divisor alignment reduces the admission-time comparison count by a factor of 8 to 17 relative to unrestricted native-period admission, and the substitution-induced jitter stays bounded at 0.7 to 0.8 ms across all densities while the reactive baseline's deflection jitter grows to 9.0 ms. The ablation attributes +0.35 of overall delivery to period substitution and +0.31 to DG salvage at the reference operating point. The divisor sweep shows that DAPD+ is insensitive to the divisor choice, while salvage-free DAPD is structurally penalized when d violates the tight-deadline bound $d \leq 2D_{\max}$, confirming the necessity of the offline selection procedure of Section 3.4. The principle decomposes into progressive realizations: GCD-only exposes the admission test as the source of the guarantee, DAPD adds substitution that recovers the coverage GCD-only sacrifices, and DAPD+ adds transient recovery through salvage; the divisor-aligned structure further makes admission cheap and order-independent, avoiding the resource monopoly that admitting coprime native periods directly would cause, and the oracle OPT-offset baseline shows that this gap is not closable by better placement at native periods. The broader lesson is to enforce arithmetic compatibility so that admitted grants are collision-free at the reservation level by construction, broaden coverage through period substitution, then let dynamic grants recover the utilization cost of being strict. Framed in terms of symmetry, the schedulability of periodic reservations is governed not by any visual or topological symmetry but by the number-theoretic relationship among their periods; the divisor-aligned catalog levels the resulting schedulability asymmetry, under which coprime-rich period combinations are structurally penalized, by reducing every admitted period to a common modular structure. The scope boundaries of the study are consolidated in Section 5.10; half-duplex coordination and a cross-subchannel placement layer, integration with PHY reliability mechanisms such as HARQ, online divisor adaptation, and handover-prepared CG search are left as future work, as is validation against a realistic, less coprime-heavy V2X trace rather than the synthetic stress mix used here.

Author Contributions: Conceptualization, H.K.; methodology, H.S. and H.K.; software, H.S.; validation, H.S. and H.K.; formal analysis, H.K.; investigation, H.S.; writing—original draft preparation, H.S. and H.K.; writing—review and editing, H.K.; supervision, H.K.; funding acquisition, H.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Technology Innovation Program (No. 20026341, Development of ultra-thin SiP under 3.5 mm and roof-integrated 5G TCU) funded by the Ministry of Trade, Industry and Energy (MOTIE, Korea).

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- 3GPP. NR; *Medium Access Control (MAC) Protocol Specification*; Technical Specification (TS) 38.321; 3rd Generation Partnership Project: Sophia Antipolis, France, 2024.
- 3GPP. NR; *Physical Layer Procedures for Data*; Technical Specification (TS) 38.214; 3rd Generation Partnership Project: Sophia Antipolis, France, 2024.
- 3GPP. *Service Requirements for Enhanced V2X Scenarios*; Technical Specification (TS) 22.186; 3rd Generation Partnership Project: Sophia Antipolis, France, 2022.
- 3GPP. NR; *Radio Resource Control (RRC); Protocol specification*; Technical Specification (TS) 38.331; 3rd Generation Partnership Project: Sophia Antipolis, France, 2024.
- Larrañaga-Zumeta, A.; Lucas-Estañ, M.C.; Gozávez, J.; Arriola, A. 5G configured grant scheduling for seamless integration with TSN industrial networks. *Comput. Commun.* **2024**, *226*–227, 107930. <https://doi.org/10.1016/j.comcom.2024.107930>.
- Zhang, T.; Hu, X.S.; Han, S. Contention-free configured grant scheduling for 5G URLLC traffic. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*; IEEE: Piscataway, NJ, USA, 2023; pp. 1–6. <https://doi.org/10.1109/DAC56929.2023.10247842>.
- Pan, Y.; Mahfouzi, R.; Samii, S.; Eles, P.; Peng, Z. Multi-traffic resource optimization for real-time applications with 5G configured grant scheduling. *ACM Trans. Embed. Comput. Syst.* **2024**, *23*, 63:1–63:31. <https://doi.org/10.1145/3664621>.
- Riaz, S.M.; Lucas-Estañ, M.C.; Coll-Perales, B.; Gozávez, J. Predictive configured grant scheduling for deterministic wireless communications. In *2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–5.
- Jiang, N.; Aijaz, A.; Jin, Y. Recursive periodicity shifting for semi-persistent scheduling of time-sensitive communication in 5G. In *2021 IEEE Global Communications Conference (GLOBECOM)*; IEEE: Piscataway, NJ, USA, 2021; pp. 1–6. <https://doi.org/10.1109/GLOBECOM46510.2021.9685475>.
- Lu, J.; Zhao, J.; He, X.; Yin, C. Periodic shifting semi-persistent scheduling for multi-user communication in 5G-TSN systems. In *2023 International Conference on Wireless Communications and Signal Processing (WCSP)*; IEEE: Piscataway, NJ, USA, 2023; pp. 523–527. <https://doi.org/10.1109/WCSP58612.2023.10404264>.
- Sabeeh, S.; Wesołowski, K.; Sroka, P. C-V2X centralized resource allocation with spectrum re-partitioning in highway scenario. *Electronics* **2022**, *11*, 279. <https://doi.org/10.3390/electronics11020279>.
- Yuan, D.; Hu, D.; Chen, X. Resource allocation in C-V2X Mode 3 based on the exchanged preference profiles. *Electronics* **2023**, *12*, 1071. <https://doi.org/10.3390/electronics12051071>.
- Sabeeh, S.; Wesołowski, K. On adaptation of resources in New Radio Vehicle-to-Everything Mode 1 dynamic resource allocation. *Electronics* **2025**, *14*, 77. <https://doi.org/10.3390/electronics14010077>.
- Sempere-García, D.; Sepulcre, M.; Gozávez, J. LTE-V2X Mode 3 scheduling based on adaptive spatial reuse of radio resources. *Ad Hoc Netw.* **2021**, *113*, 102351. <https://doi.org/10.1016/j.adhoc.2020.102351>.
- Garcia, M.H.C.; Molina-Galan, A.; Boban, M.; Gozávez, J.; Coll-Perales, B.; Şahin, T.; Kousaridas, A. A tutorial on 5G NR V2X communications. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 1972–2026. <https://doi.org/10.1109/COMST.2021.3057017>.
- Lien, S.Y.; Deng, D.J.; Lin, C.C.; Tsai, H.L.; Chen, T.; Guo, C.; Cheng, S.M. 3GPP NR sidelink transmissions toward 5G V2X. *IEEE Access* **2020**, *8*, 35368–35382. <https://doi.org/10.1109/ACCESS.2020.2973706>.
- Ali, Z.; Lagen, S.; Giupponi, L.; Rouil, R. 3GPP NR V2X Mode 2: Overview, models and system-level evaluation. *IEEE Access* **2021**, *9*, 89554–89579. <https://doi.org/10.1109/ACCESS.2021.3090855>.
- Wael, C.B.A.; Dogheche, E.H.; Armi, N.; Subekti, A.; Dayoub, I. Leveraging 3GPP Features and Optimization Techniques for 5G NR-V2X Resource Allocation: A Survey. *IEEE Open J. Intell. Transp. Syst.* **2025**, *6*, 967–994. <https://doi.org/10.1109/OJITS.2025.3584024>.

19. Nikooroo, M.; Estrada-Jimenez, J.; Machalek, A.; Härrä, J.; Engel, T.; Turcanu, I. Mitigating collisions in sidelink NR V2X: A study on cooperative resource allocation. In *2024 22nd Mediterranean Communication and Computer Networking Conference (MedComNet)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1–4. <https://doi.org/10.1109/MedComNet62012.2024.10578242>.
20. Daw, S.; Kar, A.; Tamma, B.R. On enhancing semi-persistent scheduling in 5G NR V2X to support emergency communication services in highly congested scenarios. In *Proceedings of the 24th Int. Conference on Distributed Computing and Networking (ICDCN)*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 245–253. <https://doi.org/10.1145/3571306.3571409>.
21. Zhang, Z.; Wu, Q.; Fan, P.; Cheng, N.; Chen, W.; Letaief, K.B. DRL-based optimization for AoI and energy consumption in C-V2X enabled IoV. *IEEE Trans. Green Commun. Netw.* **2025**, *9*, 2144–2159. <https://doi.org/10.1109/TGCN.2025.3531902>.
22. *Standard J3161/1; On-Board System Requirements for LTE-V2X V2V Safety Communications*. SAE International: Warrendale, PA, USA, 2022.
23. Gholmieh, R.A.; Abbasi, H.I. C-V2X (LTE-V2X) performance enhancement through SAE J3161/1 probabilistic one-shot transmissions. In *2021 IEEE Global Communications Conference (GLOBECOM)*; IEEE: Piscataway, NJ, USA, 2021; pp. 1–6. <https://doi.org/10.1109/GLOBECOM46510.2021.9685839>.
24. 3GPP. *Architecture Enhancements for 5G System (5GS) to Support Vehicle-to-Everything (V2X) Services*; Technical Specification (TS) 23.287; 3rd Generation Partnership Project: Sophia Antipolis, France, 2024.
25. 3GPP. *System Architecture for the 5G System (5GS)*; Technical Specification (TS) 23.501, 3rd Generation Partnership Project: Sophia Antipolis, France, 2024.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.