Exact Viability and Gas Allocation in a Capacity-Limited Venting Two-Gas Atmosphere

Phillip Kingston*
Member of Technical Staff
Galactic Bioware

2 August 2026

Abstract

We solve a scalar state-constrained allocation problem with two finite resources and pressure-coupled actuator bounds, motivated by a reduced venting two-gas atmosphere. The atmosphere is isothermal and perfectly mixed; its total inventory is held constant by exact pressure regulation, and oxygen fraction is confined to a prescribed band. For constant known loads, nonbinding actuators, and a prescribed initial composition, a scalar comparison result gives the exact minimum cumulative draw of each species. Conservation and convexity determine the complete cumulative oxygen-allocation interval, whose intersection with the reserve constraints yields exact deterministic open-loop formulas for finite-horizon viability, maximum duration, viable initial compositions, and optimal initialization. Free-initial duration, reserve regimes, and reserve shadow prices follow as corollaries. For finite actuator capacities satisfying the global pointwise-authority conditions, the capacity limits induce mandatory oxygen and diluent flows. The same scalar lemma gives exact shifted allocation endpoints, an explicit viability kernel within this authority-admissible regime, and a piecewise closed-form maximum-duration formula. A residual-flexibility identity links instantaneous actuator authority to the asymptotic width growth of the cumulative-allocation set. Discretized linear programs provide deterministic same-model checks of authority cases, viability boundaries, initialization, and reduction to the full allocation range when both capacities are nonbinding. The resulting formulas provide an exact analytic benchmark for the stated scalar finite-resource allocation model.

1 Introduction and relation to prior work

We study a scalar constrained allocation problem of the form

$$\dot{y} = k(u - y),$$

in which the state must remain in an interval, the input is capacity-limited, and two cumulative resource budgets are coupled by a conservation law. A reduced two-gas atmosphere supplies a concrete and useful specialization. Historical cabin systems added oxygen and nitrogen to regulate

*Email: `phillip.kingston@galacticbioware.com`; ORCID: 0009-0005-2256-3079.

total pressure and species partial pressures [1]; later studies modeled two-gas regulation under leakage and metabolic demand [2], controlled oxygen partial pressure in recirculating breathing apparatus [3], used allowable partial-pressure ranges as operating buffers [4], and planned finite gas inventories and emergency reserves jointly [5]. A recent semi-closed-circuit respiratory-autonomy architecture with outward venting, pure-oxygen replenishment, and finite consumables provides further application context [6]. That architecture is materially richer than the abstraction below; it supplies background motivation only, and none of the present formulas transfers to it as a physiological, hardware, or certification claim. Together these systems motivate the question:

> Given a measured initial composition, finite species reserves, and finite supply capacities, which cumulative allocations are reachable, is the state viable for a required horizon, and what is the maximum feasible duration?

Present admissibility does not answer this question. A state may satisfy every instantaneous limit while lacking enough species inventory or actuator authority to survive the remaining horizon. Conversely, motion through the admissible composition interval can change the cumulative split between external oxygen and diluent. If the fixed internal inventory is N and oxygen fraction is x, the largest internal oxygen-inventory swing is

$$B = N(\overline{x} - \underline{x}). \tag{1}$$

This is an *allocation buffer*: it redistributes external species use but does not create total gas.

Viability and controlled invariance provide the general state-constraint framework [7, 8, 9]; finite-fuel control treats consumables as depleting states [10]; and state-constrained optimal control supplies classical boundary-arc tools [11]. Positive linear systems and positive-state reachability provide one close comparison-system context [12, 13]. Krastanov gives a necessary and sufficient condition for small-time controllability with respect to a cone under control and phase constraints [14], while Yang and Liu use convexification and convex optimization to compute reachable sets for linear systems with a nonconvex control constraint and additional convex control and state constraints [15]. These results address broader constrained-system questions but do not provide the cumulative species-allocation endpoints or coupled reserve formulas derived here. In recent work on networked positive systems with joint state and input capacities, Schurig et al. construct an explicit suboptimal admissible controller and compute performance bounds through a convex program [16]; the scalar pressure-coupled model considered here instead permits exact closed-form formulas. Related allocation structures occur in hybrid power-split control [17] and maximum-endurance problems [18, 19]. These sources establish the methods and neighboring problem classes. Within the literature examined, we did not identify a closed-form characterization that simultaneously combines a prescribed initial composition, a pressure-coupled cumulative allocation interval, two separate finite reserve constraints, and authority-admissible finite supply capacities in the scalar model considered here.

The contribution is an exact closed-form synthesis for the reduced model defined below. From a prescribed initial composition, scalar comparison arguments give the minimum cumulative draw of each species; conservation and convexity then yield the complete cumulative oxygen-allocation interval. Intersecting this interval with the two reserve constraints gives explicit formulas for finite-horizon viability, maximum duration, viable initial compositions, and optimal initialization. For finite capacities satisfying the global pointwise-authority conditions, mandatory flows shift the allocation endpoints and give a piecewise capacity-aware duration law. The resulting residual-flexibility identity relates instantaneous actuator authority to the asymptotic growth rate of the allocation interval. The comparison argument, convexity step, exponential boundary trajectories, and tangent conditions are standard; the model-specific closed-form synthesis is the claimed contribution. Numerical linear programs are same-model consistency checks, not physical validation. The model omits carbon dioxide, water vapor, thermal and spatial dynamics, sensor delay, valve switching, and physiological thresholds.

2 Model, conservation, and actuator authority

2.1 Pressure-regulated two-gas reduction

Assumption 2.1 (Reduced atmosphere)**.** *The control volume contains oxygen O and an inert diluent D. The gases are ideal, isothermal, and perfectly mixed. A fast pressure regulator holds the total internal molar inventory $N > 0$ constant. Well-mixed gas vents at rate $q \geq 0$, an internal process removes oxygen at rate $m \geq 0$, and external oxygen and diluent enter at nonnegative rates u_O and u_D from reserves r_O and r_D. The oxygen fraction must satisfy*

$$0 < \underline{x} < \overline{x} < 1.$$

Supply capacities are $0 \leq u_O \leq \overline{u}_O$ and $0 \leq u_D \leq \overline{u}_D$. Sections 3–6 assume $q > 0$. The pointwise-authority result in Section 2.3 remains valid when $q = 0$, but the cumulative-allocation and duration formulas are not analysed in that degenerate case.

Let n_O and n_D be internal moles and set $x = n_O/N$. Perfect-mixing balances give

$$\dot{n}_O = u_O - m - qx, \qquad\qquad \dot{n}_D = u_D - q(1-x). \tag{2}$$

Since $N = n_O + n_D$ is constant, exact pressure regulation imposes

$$u_O + u_D = m + q =: S. \tag{3}$$

The composition and reserve dynamics are therefore

$$N\dot{x} = u_O - (m + qx), \qquad \dot{r}_O = -u_O, \qquad \dot{r}_D = -u_D, \tag{4}$$

with $r_O(0) = r_{O,0}$ and $r_D(0) = r_{D,0}$. The instantaneous safe set is

$$\mathcal{S} = \{(x, r_O, r_D) : \underline{x} \leq x \leq \overline{x},\ r_O, r_D \geq 0\}.$$

Conventions and terminology. The pair (m, q) is the known constant load pair. A finite horizon $H \geq 0$ denotes the prescribed feasibility interval $[0, H]$; a starred quantity $H^\star$ (with subscripts or superscripts as needed) denotes the corresponding maximum feasible duration. For $z \in \mathbb{R}$, write $(z)_+ = \max\{z, 0\}$, and for an interval $[a, b]$ write

$$\mathrm{proj}_{[a,b]}(z) = \min\{b, \max\{a, z\}\}.$$

In this paper the actuators are called *nonbinding* when $\overline{u}_O, \overline{u}_D \geq S$, so exact pressure regulation permits the full split $u_O \in [0, S]$. Dimensionally, N, r_O, r_D and the cumulative draws C_O, C_D defined below are gas amounts; m, q, u_O, u_D are gas-amount rates; H is time; and x is dimensionless.

Definition 2.2 (Admissible deterministic control)**.** *For the prescribed interval $[0, H]$, an admissible nonbinding control is a measurable function $u_O : [0, H] \to [0, S]$, with $u_D = S - u_O$, whose unique absolutely continuous trajectory satisfies $x(t) \in [\underline{x}, \overline{x}]$ for all t. It is reserve-feasible from $(r_{O,0}, r_{D,0})$*

if $C_O(H) \le r_{O,0}$ and $C_D(H) \le r_{D,0}$. Since both cumulative draws are nondecreasing, terminal reserve feasibility is equivalent to nonnegative reserves throughout the horizon. A state is viable over $[0, H]$ if such a reserve-feasible open-loop control exists. Here open-loop *means that the control is selected as a measurable function of time for the known constant load pair and initial state, rather than as a feedback law depending on subsequent observations. Accordingly, the viability sets below are existential deterministic finite-horizon sets for the stated model; they are not disturbance-robust or partial-information viability kernels.*

For later use define

$$\begin{aligned} a_L &= m + q\underline{x}, & a_U &= m + q\overline{x}, \\ b_U &= q(1-\overline{x}), & b_L &= q(1-\underline{x}), \end{aligned} \tag{5}$$

so that $a_L + b_L = a_U + b_U = S$.

2.2 Conservation and the allocation buffer

For a horizon H, write

$$C_O(H) = \int_0^H u_O(t)\,\mathrm{d}t, \qquad C_D(H) = \int_0^H u_D(t)\,\mathrm{d}t.$$

Integrating (2) yields

$$C_O(H) = N(x(H) - x(0)) + mH + q\int_0^H x(t)\,\mathrm{d}t, \tag{6}$$

$$C_D(H) = N(x(0) - x(H)) + q\int_0^H (1 - x(t))\,\mathrm{d}t. \tag{7}$$

Adding gives the invariant

$$C_O(H) + C_D(H) = SH. \tag{8}$$

Composition can therefore change only the species split of the external supply. Equation (1) bounds the internal inventory available for that reallocation, while (8) fixes the total external draw.

2.3 Pointwise actuator authority

Controlled invariance here concerns the composition interval alone with reserve depletion ignored, whereas finite-horizon viability also imposes the two reserve budgets over $[0, H]$. By *pointwise actuator authority* we mean that, at each composition boundary, a pressure-compatible instantaneous input exists for which the composition vector field points inward or tangent to the admissible interval.

Exact pressure regulation restricts oxygen flow to

$$u_O \in [\alpha, S - \beta], \qquad \alpha = (S - \overline{u}_D)_+, \qquad \beta = (S - \overline{u}_O)_+. \tag{9}$$

Thus a diluent capacity limit imposes a mandatory oxygen flow α, and an oxygen capacity limit imposes a mandatory diluent flow β.

Theorem 2.3 (Necessary and sufficient pointwise authority)**.** *Ignoring reserve depletion, the composition interval $[\underline{x}, \overline{x}]$ is controlled invariant under* (4) *and exact pressure regulation if and only if*

$$\alpha + \beta \le S, \qquad \alpha \le a_U, \qquad \beta \le b_L. \tag{10}$$

For nonnegative capacities these conditions are equivalently

$$\overline{u}_O + \overline{u}_D \ge S, \qquad \overline{u}_O \ge a_L, \qquad \overline{u}_D \ge b_U. \tag{11}$$

Proof. The state-independent admissible input interval in (9) is nonempty exactly when $\alpha+\beta \le S$; if this condition fails, no pressure-compatible control exists. If $\beta > b_L$, then at $x = \underline{x}$ every admissible input satisfies $u_O \le S - \beta < a_L$, so $N\dot{x} = u_O - a_L < 0$ and the trajectory immediately leaves the band. Similarly, if $\alpha > a_U$, then at $x = \overline{x}$ every admissible input satisfies $u_O \ge \alpha > a_U$, so $N\dot{x} = u_O - a_U > 0$. These contrapositives establish necessity of all three conditions in (10).

Conversely, suppose (10) holds. At $x = \underline{x}$, the inequality $\beta \le b_L$ is equivalent to $S - \beta \ge a_L$, so an admissible input with $\dot{x} \ge 0$ exists. At $x = \overline{x}$, the inequality $\alpha \le a_U$ supplies an admissible input with $\dot{x} \le 0$. Because the control interval is independent of x and the interior imposes no tangent restriction, the one-dimensional Nagumo tangent condition [9] makes these boundary conditions sufficient for controlled invariance of the whole interval. Expanding the positive parts gives (11). For $q > 0$, the joint-capacity inequality is genuinely separate, because

$$a_L + b_U = S - q(\overline{x} - \underline{x}) < S.$$

When $q = 0$, $a_L + b_U = S$, so the joint-capacity inequality follows from the two boundary-capacity inequalities. $\square$

Remark 2.4 (Oxygen-only boundary trajectory)**.** *If $u_D \equiv 0$, pressure regulation forces $u_O = S$ and*

$$x(t) = 1 - (1 - x_0)e^{-qt/N}.$$

This standard perfect-mixing law [20] is the diluent-free boundary trajectory that generates the upper-composition extremal construction used below.

3 Exact cumulative allocation from a prescribed state

Throughout Sections 3–6, $m \ge 0$ and $q > 0$ are constant. The initial composition $x(0) = x_0 \in [\underline{x}, \overline{x}]$ is prescribed. We first package the standard scalar comparison mechanism used on both species so that it can be invoked without repeating the same argument. Its coast-then-hold structure is a state-constrained first-order analogue of classical minimum-fuel switching laws [21]. The lemma packages the scalar calculation used for both species.

Lemma 3.1 (Constrained scalar minimum-input law)**.** *Let $k > 0$, $0 < L < U$, $y(0) = y_0 \in [L, U]$, and*

$$\dot{y} = k(u - y), \qquad u(t) \in [c, \overline{u}] \tag{12}$$

for measurable inputs, where $0 \le c \le \min\{U, \overline{u}\}$ and $\overline{u} \ge L$. Among trajectories satisfying $y(t) \in [L, U]$ on $[0, H]$, define

$$\tau_c(y_0) = \begin{cases} k^{-1} \log \dfrac{y_0 - c}{L - c}, & c < L, \\ 0, & c \ge L. \end{cases} \tag{13}$$

Then the exact minimum cumulative input is

$$\boxed{\min \int_0^H u(t)\,\mathrm{d}t = cH + (L - c)_+ \left(H - \tau_c(y_0)\right)_+ .} \tag{14}$$

If $c < L$, an attaining input uses $u = c$ until $y = L$ and then $u = L$; if $c \ge L$, the constant input $u = c$ attains the minimum.

Proof. If $c < L$, set $\widehat{y} = y - c$ and $\widehat{u} = u - c \geq 0$. Then $\dot{\widehat{y}} = k(\widehat{u} - \widehat{y})$ and the state floor is $\widehat{y} \geq L - c > 0$. Since $\widehat{u} \geq 0$, comparison gives

$$\widehat{y}(t) \geq \widehat{y}(0)e^{-kt}.$$

The exponential comparison curve reaches $L - c$ at $\tau_c(y_0)$; the controlled trajectory need not make contact there.

If $H \leq \tau_c(y_0)$, then

$$\int_0^H \widehat{u}(t)\,\mathrm{d}t \geq 0 = (L - c)\,(H - \tau_c(y_0))_+\,.$$

If $H > \tau_c(y_0)$, use the exponential comparison on $[0, \tau_c(y_0)]$, the state floor $\widehat{y} \geq L - c$ on $[\tau_c(y_0), H]$, and $\widehat{y}(H) \geq L - c$ in

$$\int_0^H \widehat{u}(t)\,\mathrm{d}t = \int_0^H \widehat{y}(t)\,\mathrm{d}t + k^{-1}(\widehat{y}(H) - \widehat{y}(0)).$$

The boundary terms cancel and give

$$\int_0^H \widehat{u}(t)\,\mathrm{d}t \geq (L - c)(H - \tau_c(y_0)).$$

Combining the two cases yields

$$\int_0^H \widehat{u}(t)\,\mathrm{d}t \geq (L - c)\,(H - \tau_c(y_0))_+\,.$$

The coast-to-boundary-then-hold input attains equality; $\overline{u} \geq L$ makes the hold input admissible. Adding the mandatory input cH proves (14). If $c \geq L$, pointwise input feasibility gives $\int_0^H u \geq cH$, while the constant input $u = c$ has an equilibrium in $[L, U]$ because $c \leq U$ and therefore remains in the interval. □

For the nonbinding two-gas problem, $0 \leq u_O, u_D \leq S$. Set

$$k = \frac{q}{N}, \qquad \tau_O(x_0) = \frac{1}{k}\log\frac{m + qx_0}{a_L}, \qquad \tau_D(x_0) = \frac{1}{k}\log\frac{1 - x_0}{1 - \overline{x}}. \tag{15}$$

Here $\tau_O(x_0)$ is the oxygen-free descent time from x_0 to $\underline{x}$, and $\tau_D(x_0)$ is the diluent-free ascent time from x_0 to $\overline{x}$.

Proposition 3.2 (Exact minimum species draws)**.** *Every admissible trajectory from x_0 over $[0, H]$ satisfies*

$$C_O(H) \geq a_L\,(H - \tau_O(x_0))_+\,, \tag{16}$$
$$C_D(H) \geq b_U\,(H - \tau_D(x_0))_+\,. \tag{17}$$

Both bounds are attained for every $H \geq 0$.

Proof. For oxygen set $y = m + qx$. Then $y \in [a_L, a_U]$ and $\dot{y} = k(u_O - y)$. Since $0 < a_L < a_U < S$, the hypotheses of Lemma 3.1 hold for $(L, U, c, \overline{u}) = (a_L, a_U, 0, S)$, giving (16). The attaining control uses $u_O = 0$ until the lower boundary is reached and then $u_O = a_L$. For diluent set $w = q(1 - x)$. Then $w \in [b_U, b_L]$ and $\dot{w} = k(u_D - w)$; since $0 < b_U < b_L < S$, the same hypotheses hold for $(L, U, c, \overline{u}) = (b_U, b_L, 0, S)$, giving (17) and the symmetric ascent-and-hold control. □

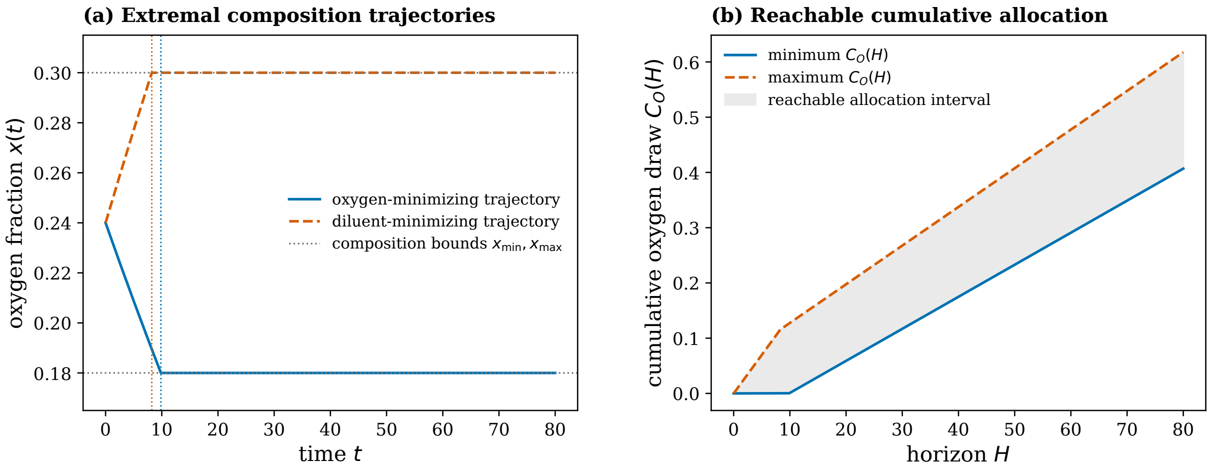

Figure 1: Fixed-state endpoint geometry for $N = 1$, $\underline{x} = 0.18$, $\overline{x} = 0.30$, $m = 0.004$, $q = 0.01$, and $x_0 = 0.24$. (a) The oxygen-minimizing and diluent-minimizing trajectories start from the same measured state, reach the appropriate composition boundary at $\tau_O = 9.8440$ and $\tau_D = 8.2238$, and then hold it. (b) Their cumulative draws are the exact lower and upper endpoints of the reachable oxygen-allocation interval. Color is reinforced by line style for grayscale reproduction.

Theorem 3.3 (Fixed-state reachable cumulative allocation)**.** *The set of cumulative oxygen draws attainable from the prescribed initial composition x_0 over $[0, H]$ is exactly*

$$\boxed{\mathcal{C}_O(H \mid x_0) = \left[a_L \left(H - \tau_O(x_0)\right)_+ ,\ SH - b_U \left(H - \tau_D(x_0)\right)_+\right].} \tag{18}$$

Proof. Proposition 3.2 gives the lower endpoint. By (8), maximizing C_O is equivalent to minimizing C_D, which gives the upper endpoint. Both endpoint trajectories start at the same x_0. For any $\lambda \in [0, 1]$, their pointwise convex combination preserves the initial state, affine dynamics, pressure equality, convex input constraints, and composition band; its cumulative oxygen draw is the corresponding convex combination. Hence every point between the endpoints is attainable, and Proposition 3.2 excludes every point outside them. □

Corollary 3.4 (Free-initial allocation)**.** *If the initial composition may be selected freely in the band, define*

$$t_O = \frac{1}{k} \log \frac{a_U}{a_L}, \qquad t_D = \frac{1}{k} \log \frac{1 - \underline{x}}{1 - \overline{x}}. \tag{19}$$

Then the attainable cumulative oxygen set is

$$[a_L(H - t_O)_+,\ SH - b_U(H - t_D)_+]\,. \tag{20}$$

Proof. The lower endpoint is minimized at $x_0 = \overline{x}$ because τ_O is increasing, and the upper endpoint is maximized at $x_0 = \underline{x}$ because τ_D is decreasing. These endpoint trajectories begin at different boundary compositions, which is permitted because the initial composition is a decision variable in this corollary. Their pointwise convex combination begins at the corresponding convex combination of $\underline{x}$ and $\overline{x}$, remains in the band, and fills the interval. □

4 Viability, duration, and initialization

Let the initial reserves be $r_{O,0}, r_{D,0} \geq 0$. By conservation, the two reserve inequalities impose the algebraic interval

$$I_R(H) = [SH - r_{D,0},\ r_{O,0}]. \tag{21}$$

The physically attainable values are obtained only after intersecting $I_R(H)$ with the reachable allocation set, which is contained in $[0, SH]$.

Theorem 4.1 (Deterministic viability kernel and maximum duration)**.** *Under the assumptions of Section 3, a state $(x_0, r_{O,0}, r_{D,0}) \in \mathcal{S}$ is viable over $[0, H]$ if and only if*

$$a_L\,(H - \tau_O(x_0))_+ \leq r_{O,0}, \tag{22}$$

$$b_U\,(H - \tau_D(x_0))_+ \leq r_{D,0}, \tag{23}$$

$$SH \leq r_{O,0} + r_{D,0}. \tag{24}$$

Consequently the exact deterministic kernel is

$$\mathcal{V}_H^{\det} = \{(x_0, r_{O,0}, r_{D,0}) \in \mathcal{S} : (22)\text{–}(24)\}\,, \tag{25}$$

and the maximum feasible duration is

$$\boxed{H^\star(x_0, r_{O,0}, r_{D,0}) = \min\left\{\frac{r_{O,0} + r_{D,0}}{S},\ \tau_O(x_0) + \frac{r_{O,0}}{a_L},\ \tau_D(x_0) + \frac{r_{D,0}}{b_U}\right\}.} \tag{26}$$

The optimum is attained.

Proof. The dynamically reachable interval is (18), while reserve feasibility is (21). The latter is nonempty exactly under (24). Conditional on both intervals being nonempty, their intersection is nonempty exactly when the reachable lower endpoint does not exceed $r_{O,0}$ and $SH - r_{D,0}$ does not exceed the reachable upper endpoint. These are (22) and (23). Each inequality is nondecreasing in H and is equivalent to the corresponding upper bound in (26). Attainment follows from Theorem 3.3: a reserve-compatible cumulative allocation in the interval has an admissible realizing trajectory. $\square$

The three terms in (26) correspond to total-gas exhaustion, oxygen shortage after the exact oxygen-free descent available from x_0, and diluent shortage after the symmetric ascent. Unlike the free-initial convention, these credits are determined by the measured initial state.

Corollary 4.2 (Viable initial compositions)**.** *For fixed $H, r_{O,0}, r_{D,0}$ define*

$$x_-(H, r_{O,0}) = \frac{a_L \exp\Big(k\,(H - r_{O,0}/a_L)_+\Big) - m}{q}, \tag{27}$$

$$x_+(H, r_{D,0}) = 1 - (1 - \overline{x}) \exp\Big(k\,(H - r_{D,0}/b_U)_+\Big)\,. \tag{28}$$

If $SH \leq r_{O,0} + r_{D,0}$, the set of viable initial compositions is exactly

$$\boxed{\mathcal{X}_H(r_{O,0}, r_{D,0}) = [x_-(H, r_{O,0}),\ x_+(H, r_{D,0})]} \tag{29}$$

when $x_- \leq x_+$, and is empty otherwise. If the total-gas inequality fails, it is empty.

Proof. The oxygen condition in Theorem 4.1 is equivalent to $\tau_O(x_0) \geq (H - r_{O,0}/a_L)_+$, which inverts to $x_0 \geq x_-$. The diluent condition similarly inverts to $x_0 \leq x_+$. At zero positive part, the two formulas reduce to $\underline{x}$ and $\overline{x}$, respectively. $\square$

This design problem treats x_0 as freely selectable at the beginning of the accounting horizon and does not charge the gas or time required to establish it. If preparation draws from the same reserves, its cost must be included separately. The result below assumes nonbinding actuators; capacity-aware optimal initialization is not derived here.

Proposition 4.3 (Optimal initialization)**.** *Define*

$$F_O(x) = \tau_O(x) + \frac{r_{O,0}}{a_L}, \qquad F_D(x) = \tau_D(x) + \frac{r_{D,0}}{b_U},$$
$$\Lambda = \frac{a_L}{1-\overline{x}} \exp\left[k\left(\frac{r_{D,0}}{b_U} - \frac{r_{O,0}}{a_L}\right)\right], \qquad x_{\text{bal}}^{\text{tr}} = \operatorname{proj}_{[\underline{x},\overline{x}]}\left(\frac{\Lambda - m}{q+\Lambda}\right).$$

Let

$$T = \frac{r_{O,0} + r_{D,0}}{S}, \qquad G^\star = \max_{x\in[\underline{x},\overline{x}]} \min\{F_O(x), F_D(x)\}.$$

If $T \geq G^\star$, $x_{\text{bal}}^{\text{tr}}$ maximizes $H^\star(x, r_{O,0}, r_{D,0})$ over the composition band. If $T < G^\star$, every $x \in \mathcal{X}_T(r_{O,0}, r_{D,0})$ is optimal.

Proof. Differentiating gives

$$F_O'(x) = \frac{N}{m+qx} > 0, \qquad F_D'(x) = -\frac{N}{q(1-x)} < 0.$$

Thus $\min\{F_O, F_D\}$ increases until the two functions cross and decreases afterward. If no crossing lies in the band, the maximum is at the appropriate boundary; this is exactly projection of the formal crossing. The equation $F_O(x) = F_D(x)$ is equivalent to

$$\frac{m+qx}{1-x} = \frac{a_L}{1-\overline{x}} \exp\left[k\left(\frac{r_{D,0}}{b_U} - \frac{r_{O,0}}{a_L}\right)\right] = \Lambda,$$

which gives $x = (\Lambda - m)/(q+\Lambda)$. Hence $x_{\text{bal}}^{\text{tr}}$ attains $G^\star$. If $T \geq G^\star$, truncation by the total-gas term does not alter that optimizer. If $T < G^\star$, the optimal value is T, and the optimal compositions are exactly those satisfying $F_O(x) \geq T$ and $F_D(x) \geq T$, namely the viable interval $\mathcal{X}_T(r_{O,0}, r_{D,0})$ from Corollary 4.2. $\square$

Remark 4.4 (Consistency of fixed-state and free-initial optimization)**.** *The fixed-state optimization in Proposition 4.3 and the free-initial formula below might appear to compare* $\max_x \min\{F_O, F_D\}$ *with* $\min\{\max_x F_O, \max_x F_D\}$. *At an interior crossing $F_O(x) = F_D(x) = v$, however,*

$$T - v = \frac{(a_L + b_U - S)v - a_L\tau_O(x) - b_U\tau_D(x)}{S} \leq 0,$$

because $a_L + b_U - S = q(\underline{x} - \overline{x}) < 0$. Thus the total-gas term is already active at the only interior point where a strict max–min gap could arise. If the crossing lies above the band, then $F_O(\overline{x}) \leq F_D(\overline{x})$ and monotonicity gives

$$\max_x \min\{F_O(x), F_D(x)\} = F_O(\overline{x}) = \min\Big\{\max_x F_O(x), \max_x F_D(x)\Big\};$$

if it lies below the band, the symmetric identity holds with $F_D(\underline{x})$. Hence the boundary-projection cases have equality directly, while the total-gas term removes the possible interior discrepancy. This reconciles the two formulations.

The free-initial result below treats the initial composition as a cost-free design variable at the start of the accounting horizon; any gas or time required to establish that composition must be charged separately.

Corollary 4.5 (Free-initial duration and reserve geometry)**.** *Optimizing over the initial composition gives*

$$H^\star_{\mathrm{free}}(r_{O,0}, r_{D,0}) = \min\left\{\frac{r_{O,0}+r_{D,0}}{S},\ t_O + \frac{r_{O,0}}{a_L},\ t_D + \frac{r_{D,0}}{b_U}\right\}. \tag{30}$$

For fixed total reserve $R > 0$ and oxygen share $\rho = r_{O,0}/R$, the total-gas term is active exactly for

$$\max\left\{0, \frac{a_L}{S} - \frac{a_L t_O}{R}\right\} \le \rho \le \min\left\{1, 1 - \frac{b_U}{S} + \frac{b_U t_D}{R}\right\}. \tag{31}$$

Moreover $H^\star_{\mathrm{free}}$ is continuous, concave, nondecreasing, and piecewise affine in the two reserves. On a uniquely active branch its reserve gradient is (S^{-1}, S^{-1}), $(a_L^{-1}, 0)$, or $(0, b_U^{-1})$. These gradients are the shadow prices of the two reserves: the marginal endurance gained per additional unit of stored oxygen or diluent in the active regime.

Proof. Apply the interval-intersection argument of Theorem 4.1 to the free-initial interval in Corollary 3.4. Corollary 3.4 realizes every intermediate allocation using the initial composition generated by the same convex combination, while Remark 4.4 gives the corresponding reconciliation with the fixed-state max–min formulation. The result is (30). Substituting $r_{O,0} = \rho R$ and $r_{D,0} = (1-\rho)R$ and requiring the total term not to exceed either species term gives (31). The regularity and gradients follow because (30) is the pointwise minimum of three affine nondecreasing functions of the reserve vector. □

The composition credits obey

$$t_O \le \frac{B}{a_L}, \qquad t_D \le \frac{B}{b_U}, \tag{32}$$

with strict inequalities for a nonzero band. Indeed,

$$t_O = \frac{N}{q}\log\left(1 + \frac{q(\overline{x} - \underline{x})}{a_L}\right) < \frac{B}{a_L},$$

and the diluent inequality is symmetric, by $\log(1+z) < z$ for $z > 0$. The linear full-band inventory-swing estimates B/a_L and B/b_U are loose for the same reason: the extremal transition consumes none of the scarce species, but its state-dependent exponential discharge yields a smaller time credit than dividing the full inventory swing by the boundary hold rate.

Remark 4.6 (Reserve perturbations)**.** *Let*

$$L_r = \max\{S^{-1}, a_L^{-1}, b_U^{-1}\}.$$

Since $H^\star_{\mathrm{free}}$ is the pointwise minimum of affine branches whose gradients have ℓ_∞ norm at most L_r,

$$|H^\star_{\mathrm{free}}(r) - H^\star_{\mathrm{free}}(\widetilde{r})| \le L_r \|r - \widetilde{r}\|_1.$$

Thus an ℓ_1 reserve-measurement error of at most ε changes the predicted free-initial maximum duration by at most $L_r\varepsilon$.

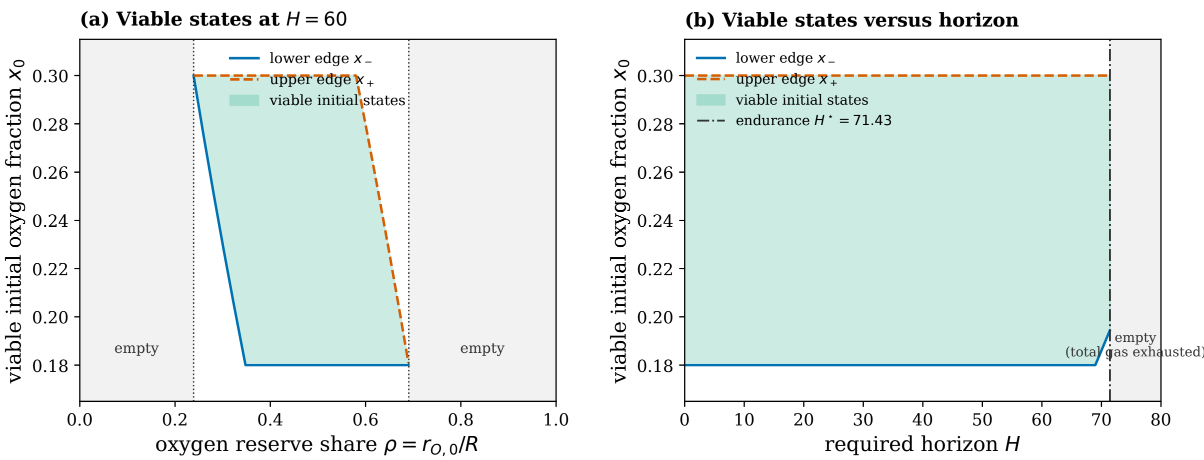

Figure 2: Closed-form viable-initial-composition geometry for the common model parameters in Table 1 and total reserve $R = 1$. (a) At $H = 60$, viable initial states exist exactly for $0.23893 \le \rho = r_{O,0}/R \le 0.69076$; the gray outer regions are empty because $x_- > x_+$. (b) For $(r_{O,0}, r_{D,0}) = (0.40, 0.60)$, total gas binds at $H^\star = 71.43$. The viable interval is still nonempty at $H^\star$ and is empty for $H > H^\star$; the boundary curves are therefore masked beyond the exact maximum-duration line.

5 Finite capacities under global pointwise authority

We now retain the prescribed initial state and constant loads and impose finite capacities $\overline{u}_O, \overline{u}_D$. Throughout this section we restrict attention to capacities satisfying the global pointwise-authority conditions of Theorem 2.3, so that the entire composition band is controlled invariant when reserve depletion is ignored. Admissibility in Definition 2.2 is therefore understood with the pressure-compatible input interval $u_O \in [\alpha, S - \beta]$ from (9), where α and β are the mandatory oxygen and diluent flows. Define

$$\tau_O^\alpha(x_0) = \begin{cases} \dfrac{1}{k} \log \dfrac{m + qx_0 - \alpha}{a_L - \alpha}, & \alpha < a_L, \\ 0, & \alpha \ge a_L, \end{cases} \tag{33}$$

$$\tau_D^\beta(x_0) = \begin{cases} \dfrac{1}{k} \log \dfrac{q(1 - x_0) - \beta}{b_U - \beta}, & \beta < b_U, \\ 0, & \beta \ge b_U. \end{cases} \tag{34}$$

When a mandatory flow lies below the corresponding boundary-hold rate, these are the times required to reach that boundary under the mandatory flow. If it equals or exceeds the hold rate, the mandatory input has an equilibrium inside the composition band and no boundary switch is required.

Theorem 5.1 (Capacity-constrained cumulative allocation)**.** *Define*

$$G_O(H \mid x_0) = \alpha H + (a_L - \alpha)_+ \left(H - \tau_O^\alpha(x_0)\right)_+, \tag{35}$$

$$G_D(H \mid x_0) = \beta H + (b_U - \beta)_+ \left(H - \tau_D^\beta(x_0)\right)_+. \tag{36}$$

Under the authority conditions of Theorem 2.3, the exact capacity-constrained cumulative oxygen-allocation set is

$$\boxed{\mathcal{C}_O^{\mathrm{cap}}(H \mid x_0) = [G_O(H \mid x_0),\ SH - G_D(H \mid x_0)]\,.} \tag{37}$$

Proof. For oxygen, $y = m + qx$ satisfies $\dot{y} = k(u_O - y)$, $y \in [a_L, a_U]$, and $u_O \in [\alpha, S - \beta]$. The authority conditions give

$$\alpha \le a_U, \qquad \alpha \le S - \beta \iff \alpha + \beta \le S, \qquad S - \beta \ge a_L \iff \beta \le b_L,$$

so Lemma 3.1 applies with $(L, U, c, \overline{u}) = (a_L, a_U, \alpha, S - \beta)$ and yields (35). If $\alpha < a_L$, the attaining input uses $u_O = \alpha$ until the lower boundary and then $u_O = a_L$; if $\alpha \ge a_L$, it uses the constant input $u_O = \alpha$.

For diluent, $w = q(1 - x)$ satisfies $\dot{w} = k(u_D - w)$, $w \in [b_U, b_L]$, and $u_D \in [\beta, S - \alpha]$. Authority gives

$$\beta \le b_L, \qquad \beta \le S - \alpha \iff \alpha + \beta \le S, \qquad S - \alpha \ge b_U \iff \alpha \le a_U,$$

so Lemma 3.1 gives (36). Conservation gives the upper oxygen endpoint. Both endpoint trajectories start at x_0 and use the same convex capacity interval, so pointwise convex combinations fill the interval. □

Corollary 5.2 (Deterministic viability kernel under global pointwise authority)**.** *Under the authority conditions of Theorem 2.3, the exact capacity-limited kernel is*

$$\boxed{\mathcal{V}_{H,\mathrm{cap}}^{\mathrm{det}} = \left\{ (x_0, r_{O,0}, r_{D,0}) \in \mathcal{S} : \begin{array}{l} G_O(H \mid x_0) \le r_{O,0}, \\ G_D(H \mid x_0) \le r_{D,0}, \\ SH \le r_{O,0} + r_{D,0} \end{array} \right\}.} \tag{38}$$

Proof. Intersect the reachable interval in Theorem 5.1 with the reserve interval (21), exactly as in Theorem 4.1. □

Scope boundary: transient viability outside the global-authority regime. The conditions of Theorem 2.3 are necessary and sufficient for controlled invariance of the entire composition band when reserve depletion is ignored; they are not necessary for nonzero finite-horizon viability from a particular initial state. For example, let $\overline{u}_O = 0$, $\overline{u}_D \ge S$, and $x_0 = \overline{x}$. Pressure regulation then forces $u_O = 0$ and $u_D = S$. If $r_{D,0} \ge SH$, the resulting trajectory is reserve-feasible and remains in the composition band for $0 \le H \le \tau_O(\overline{x})$, although $\overline{u}_O \ge a_L$ fails. The formulas in this section therefore characterize the globally authority-admissible capacity regime and do not cover all transiently viable capacity–state pairs.

To invert the two species budgets, define for $d > 0$, $c, \tau, r \ge 0$

$$g_{c,d,\tau}(H) = cH + (d - c)_+ (H - \tau)_+, \tag{39}$$

and its generalized inverse $\Gamma(r; c, d, \tau) = \sup\{H : g_{c,d,\tau}(H) \le r\}$. Since $d > 0$, g has a strictly positive affine tail, so $\Gamma(r; c, d, \tau)$ is finite for every $r \ge 0$.

Lemma 5.3 (Closed-form reserve inverse)**.** *The generalized inverse of* (39) *is*

$$\Gamma(r; c, d, \tau) = \begin{cases} r/c, & c \ge d > 0, \\ \tau + r/d, & c = 0 < d, \\ r/c, & 0 < c < d,\ r \le c\tau, \\ \tau + (r - c\tau)/d, & 0 < c < d,\ r > c\tau. \end{cases} \tag{40}$$

Proof. If $c \geq d$, the positive-part term vanishes and $g = cH$. If $c < d$, then $g = cH$ on $[0, \tau]$ and $g = c\tau + d(H - \tau)$ afterward. Inverting those two affine pieces gives (40); both branches agree at $r = c\tau$. $\square$

Theorem 5.4 (Exact duration under global pointwise authority)**.** *Under the authority conditions of Theorem 2.3, the maximum feasible duration with finite capacities is*

$$\boxed{H^{\star}_{\text{cap}}(x_0, r_{O,0}, r_{D,0}) = \min\left\{\frac{r_{O,0} + r_{D,0}}{S},\ \Gamma(r_{O,0}; \alpha, a_L, \tau_O^{\alpha}(x_0)),\ \Gamma(r_{D,0}; \beta, b_U, \tau_D^{\beta}(x_0))\right\}.} \tag{41}$$

The formula is continuous and piecewise affine in the reserves and reduces to (26) *when* $\alpha = \beta = 0$.

Proof. By Theorem 5.1, horizon H is feasible exactly when $G_O(H \mid x_0) \leq r_{O,0}$, $G_D(H \mid x_0) \leq r_{D,0}$, and $SH \leq r_{O,0} + r_{D,0}$. Lemma 5.3 inverts the first two inequalities. The interval theorem supplies attainment. $\square$

When $0 < \alpha < a_L$, the oxygen-specific reserve horizon has slope $1/\alpha$ before the reserve threshold $\alpha\tau_O^{\alpha}$ and slope $1/a_L$ afterward. The first phase pays only the mandatory flow; after the lower boundary is reached, the higher hold rate applies. The diluent side is symmetric.

Corollary 5.5 (Authority–allocation identity)**.** *Define the residual allocation-flexibility rate*

$$\delta_{\text{eff}} = S - \max\{\alpha, a_L\} - \max\{\beta, b_U\}. \tag{42}$$

Under the authority conditions, $\delta_{\text{eff}} \geq 0$ *and*

$$\delta_{\text{eff}} = q(\overline{x} - \underline{x}) - (\alpha - a_L)_+ - (\beta - b_U)_+. \tag{43}$$

For

$$H \geq H_{\text{aff}}(x_0) := \max\{\tau_O^{\alpha}(x_0), \tau_D^{\beta}(x_0)\}, \tag{44}$$

the endpoint laws are affine and

$$\operatorname{diam} \mathcal{C}_O^{\text{cap}}(H \mid x_0) = \delta_{\text{eff}} H + (a_L - \alpha)_+ \tau_O^{\alpha}(x_0) + (b_U - \beta)_+ \tau_D^{\beta}(x_0), \tag{45}$$

and therefore

$$\lim_{H\to\infty} \frac{\operatorname{diam} \mathcal{C}_O^{\text{cap}}(H \mid x_0)}{H} = \delta_{\text{eff}}. \tag{46}$$

In the nonbinding case this limit is $q(\overline{x} - \underline{x})$*, exactly the gap between the sum of the two boundary-specific capacity minima* $a_L + b_U$ *and the joint capacity* S *required by pressure regulation.*

Proof. For $H \geq H_{\text{aff}}(x_0)$, Eqs. (35)–(36) have slopes $\max\{\alpha, a_L\}$ and $\max\{\beta, b_U\}$. Subtracting them from the total slope S gives (42). Since $S - a_L - b_U = q(\overline{x} - \underline{x})$, expanding the maxima gives (43). To prove nonnegativity directly, if only the oxygen positive part is active, authority gives $\alpha - a_L \leq a_U - a_L = q(\overline{x} - \underline{x})$; the diluent-only case is symmetric. If both are active, $\alpha + \beta \leq S$ gives

$$(\alpha - a_L) + (\beta - b_U) \leq S - (a_L + b_U) = q(\overline{x} - \underline{x}).$$

The case with neither active is immediate. The constant terms in the endpoint laws yield (45); division by H proves (46). $\square$

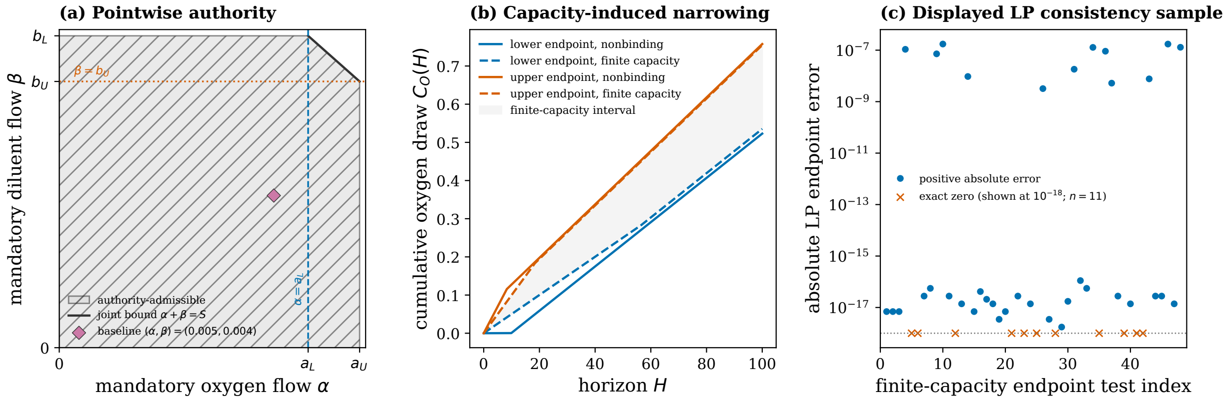

Figure 3: Finite-capacity structure for the scenario in Table 1. (a) The hatched polygon is the pointwise authority set; $\alpha + \beta = S$ is its joint boundary, whereas $\alpha = a_L$ and $\beta = b_U$ are internal hold-policy thresholds. The diamond marks $(\alpha, \beta) = (0.005, 0.004)$. (b) Finite capacities narrow the reachable cumulative-allocation interval relative to nonbinding actuators; color identifies the endpoint and line style identifies the capacity regime. (c) A displayed subset of 24 seeded finite-capacity cases gives 48 LP endpoint comparisons. Eleven exact-zero discrepancies are shown at the declared plotting floor 10^{-18}; the largest displayed error is 1.73×10^{-7}.

Interpretive limitation. The identity in Corollary 5.5 depends on the scalar two-species conservation law $C_O + C_D = SH$ and the affine endpoint slopes in (35)–(36). It should not be read as a general identity for multispecies, time-varying, or dynamically pressure-regulated systems.

Remark 5.6 (Edge cases)**.** *The formulas permit $m = 0$ and zero reserve of either species. The cases $\alpha = a_L$ and $\beta = b_U$ are covered by the $c \geq d$ branch of Lemma 5.3. If $\delta_{\text{eff}} = 0$, long-run allocation flexibility has zero width growth although a finite intercept may remain. The cumulative-allocation and duration analysis in Sections 3–6 assumes $q > 0$. When $q = 0$, the vent-driven diluent requirement and logarithmic transition-time formulas degenerate, so those results require a separate direct analysis.*

6 Numerical verification

All numerical examples are dimensionless. The portable verification program holds u_O constant on each step of length Δt and integrates (4) exactly:

$$x_{j+1} = e^{-q\Delta t/N} x_j + \frac{1 - e^{-q\Delta t/N}}{q}(u_{O,j} - m). \tag{47}$$

The linear program enforces every node in the composition band, the pressure-compatible capacity interval, and both cumulative reserve budgets. Under constant input the scalar trajectory moves monotonically toward its equilibrium, so node constraints preserve the band within each step. The program is an independently coded consistency check of the same reduced model, not physical validation.

Table 1: Numerical scenarios used in Figs. 1–3. The common model parameters apply to all panels; Figs. 1–2 use nonbinding actuators, while Fig. 3 uses the stated finite capacities.

Parameter	Value	Parameter	Value
Common model parameters			
N	1	x_0	0.24
$\underline{x}, \overline{x}$	$0.18, 0.30$	m, q	$0.004, 0.010$
S	0.014	a_L, a_U	$0.0058, 0.0070$
b_U, b_L	$0.0070, 0.0082$	τ_O, τ_D	$9.8440, 8.2238$
Nonbinding actuators: Figs. 1–2			
$\overline{u}_O, \overline{u}_D$	$0.014, 0.014$	α, β	$0, 0$
Finite capacities: Fig. 3			
$\overline{u}_O, \overline{u}_D$	$0.010, 0.009$	α, β	$0.005, 0.004$
$\tau_O^\alpha, \tau_D^\beta$	$55.9616, 18.2322$	δ_{eff}	0.0012

Table 2: Deterministic verification summary at $\Delta t = 0.1$. Values are worst observed discrepancies in the stated finite test suites, not probabilistic error bounds. Endpoint entries are absolute cumulative-flow errors.

Test family	Checks	Worst observed result
Fixed-state allocation endpoints	60 cases	5.63×10^{-7} maximum absolute error
Finite-capacity allocation endpoints	60 cases	2.65×10^{-7} maximum absolute error
Capacity-duration brackets	60 cases	$0.0974 < \Delta t$ theorem-to-grid gap
Viable-composition interval	139 points	0 classification mismatches
Optimal initialization and max–min consistency	100 cases	0 discrepancy to a 4001-point grid
Structured authority and reduction cases	7 cases	7.15×10^{-8} maximum absolute error
Capacity-limited kernel	50 cases	0 feasibility mismatches
Reduction to nonbinding formulas	100 cases	4.26×10^{-14} maximum formula discrepancy

The archived verification artifact [22] provides the code, numerical results, figures, and reproduction instructions.

The seven structured cases comprise one nonbinding reduction case, the four hold-threshold equalities $\alpha = a_L$, $\alpha = a_U$, $\beta = b_U$, and $\beta = b_L$, one joint-boundary case $\alpha + \beta = S$ (which has zero residual width), and one additional zero-width case. Viable-composition tests checked 120 analytic interior points and 19 points immediately outside the viable interval; capacity-kernel tests covered 50 cases. Initialization tests compared the closed form with a 4001-point composition grid and directly asserted equality with the free-initial optimum. For the representative grid-invariance case, the maximum absolute endpoint discrepancies at $\Delta t = 0.2, 0.1, 0.05$ were each 2.78×10^{-17}. In that case $H = 12 < \tau_O^\alpha = 77.29$ and $\beta > b_U$, so neither endpoint policy switches within the horizon and both constant endpoint controls are exactly representable on every grid. This is therefore a no-switch grid-invariance check, not a discretization-convergence result at a switching time. Across the broader endpoint suites, discrepancies may nevertheless be much smaller than Δt because the one-step dynamics are integrated exactly and a constant control within the switching cell can approximate an interior switch. Maximum-duration estimates remain bracketed on the time grid.

7 Interpretation, limitations, and conclusion

Scalar comparison arguments identify the minimum cumulative draw of each species from the prescribed initial composition. Conservation and convexity convert those extrema into the complete reachable allocation interval, whose intersection with the reserve constraints yields the finite-horizon viable set and maximum duration. Within the globally authority-admissible capacity regime, mandatory flows shift the scalar comparison problems while preserving the interval structure.

Operationally, the composition band acts as an allocation buffer. It can delay a species-specific shortage, determine which initial compositions are viable, and create a range of cumulative species splits, but it cannot exceed the hard total-gas lifetime. Finite actuator limits narrow that range in two ways: an opposite-actuator limit imposes a mandatory flow before a boundary is reached, and sufficiently strong mandatory flow can eliminate the boundary-discharge phase entirely. The residual-flexibility identity in Corollary 5.5 connects these cumulative effects to instantaneous hardware authority.

The model assumes perfect mixing, constant temperature, exact pressure regulation, constant loads, and two gas species. It omits spatial transport, carbon dioxide, water vapor, contaminants, regulator dynamics, sensing and control delay, hybrid valves, and physiological constraints. Its composition bounds are symbolic. Application to human endurance or hardware certification would require an application-specific model and validation campaign.

Within these limits, the paper gives an exact open-loop characterization of reachable cumulative allocation and finite-horizon reserve viability, together with optimal initialization for nonbinding actuators and capacity-aware maximum-duration formulas under the global pointwise-authority conditions. The free-initial, nonbinding formula is recovered as a corollary rather than imposed as an operating assumption.

Code and data availability

The verification code, numerical summaries, generated figures, environment specification, reproduction instructions, and checksum manifest are archived in the verification artifact [22].

References

[1] J. K. Jackson, *Design development and test: Two-gas atmosphere control subsystem*, NASA-CR-134190, NASA NTRS 19740008730, 1974. NASA.

[2] W. Hofacker and P. Rosenthal, "Simulation of the N_2/O_2 Pressure Control Function of Life Support Systems with Respect to Gas Losses due to Structural Leak and Crew Metabolism," SAE Technical Paper 932135, 1993. doi:10.4271/932135.

[3] F. Garofalo, S. Manfredi, and S. Santini, "Modelling and control of oxygen partial pressure in an underwater breathing apparatus with gas recycle," in *Proceedings of the 2003 European Control Conference*, pp. 868–873, 2003. doi:10.23919/ECC.2003.7085067.

[4] R. N. Schaezler and A. J. Cook, "Report on ISS O_2 Production, Gas Supply and Partial Pressure Management," International Conference on Environmental Systems, NASA NTRS 20150010425, 2015. NASA.

[5] R. A. Sturtz, R. W. Koelsch, and M. J. Sargusingh, "Optimizing Consumable Gas Tank Usage while Maintaining Emergency Reserves Between Two Pressure Control Systems on Gateway," International Conference on Environmental Systems, NASA NTRS 20230015200, 2024.

[6] P. Kingston and N. Johnston, "Physicochemical-Neural Fusion for Semi-Closed-Circuit Respiratory Autonomy in Extreme Environments," *arXiv*, 2026. doi:10.48550/arXiv.2603.26697.

[7] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability Theory: New Directions*. Springer, 2011.

[8] P. Saint-Pierre, "Approximation of the viability kernel," *Applied Mathematics and Optimization*, vol. 29, pp. 187–209, 1994. doi:10.1007/BF01204182.

[9] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999. doi:10.1016/S0005-1098(99)00113-2.

[10] D. B. Rokhlin and G. Mironenko, "Regular finite fuel stochastic control problems with exit time," *Mathematical Methods of Operations Research*, vol. 84, no. 1, pp. 105–127, 2016. doi:10.1007/s00186-016-0536-2.

[11] R. F. Hartl, S. P. Sethi, and R. G. Vickson, "A Survey of the Maximum Principles for Optimal Control Problems with State Constraints," *SIAM Review*, vol. 37, no. 2, pp. 181–218, 1995. doi:10.1137/1037043.

[12] L. Farina and S. Rinaldi, *Positive Linear Systems: Theory and Applications*. Wiley, 2000. doi:10.1002/9781118033029.

[13] C. Guiver, R. Rebarber, and S. Townley, "Positive state reachability for positive linear control systems in continuous time," *Journal of the Franklin Institute*, vol. 363, no. 1, art. 108244, 2026. doi:10.1016/j.jfranklin.2025.108244.

[14] M. I. Krastanov, "On the constrained small-time controllability of linear systems," *Automatica*, vol. 44, no. 9, pp. 2370–2374, 2008. doi:10.1016/j.automatica.2008.01.007.

[15] R. Yang and X. Liu, "Reachable set computation of linear systems with nonconvex constraints via convex optimization," *Automatica*, vol. 146, art. 110632, 2022. doi:10.1016/j.automatica.2022.110632.

[16] R. Schurig, D. Ohlin, A. Rantzer, E. Tegling, and R. Findeisen, "Model Predictive Control for Constrained Linear Positive Systems on Graphs," *IEEE Control Systems Letters*, vol. 10, pp. 481–486, 2026. doi:10.1109/LCSYS.2026.3698436.

[17] T. A. C. van Keulen, J. Gillot, B. de Jager, and M. Steinbuch, "Solution for state constrained optimal control problems applied to power split control for hybrid vehicles," *Automatica*, vol. 50, no. 1, pp. 187–192, 2014. doi:10.1016/j.automatica.2013.09.039.

[18] G. Sachs and T. Christodoulou, "Endurance increase by cyclic control," *Journal of Guidance, Control, and Dynamics*, vol. 9, no. 1, pp. 58–63, 1986. doi:10.2514/3.20067.

[19] S. Li and L. Rodrigues, “Maximum Endurance of Hybrid-Electric Aircraft in Cruise,” in *Proceedings of the 2025 European Control Conference*, pp. 1322–1327, 2025. doi:10.23919/ECC65951.2025.11187156.

[20] P. V. Danckwerts, “Continuous flow systems: Distribution of residence times,” *Chemical Engineering Science*, vol. 2, no. 1, pp. 1–13, 1953. doi:10.1016/0009-2509(53)80001-1.

[21] M. Athans, “Minimum-Fuel Feedback Control Systems: Second-Order Case,” *IEEE Transactions on Applications and Industry*, vol. 82, no. 65, pp. 8–17, 1963. doi:10.1109/TAI.1963.5407855.

[22] P. Kingston, *Verification Artifact for Exact Viability and Gas Allocation in a Capacity-Limited Venting Two-Gas Atmosphere*, Zenodo, 2026. doi:10.5281/zenodo.21761429.