capstract:2611.00045v1 [math.OC] 20 Nov 2026

Optimal checkpoint placement in sequential leave-taking

Mira Varen
Department of Applied Mathematics, Sedgewick Vale Institute

November 20, 2026

Abstract

A dinner guest who is drawn back into conversation may have to repeat farewells already made. We study the timing of coat operations that, with the host's acknowledgment, preserve progress along a prescribed farewell route. Contact durations and restart probabilities can differ by guest, and the coat operation can itself reopen conversation. We show that a fixed partition of the route is optimal among history-dependent policies and give an exact linear-time algorithm for choosing its checkpoints. For homogeneous routes with reliable acknowledgment we derive the fee below which a separate coat operation is optimal after each contact short of the door. In 120 pairs of synthetic 24-contact routes, with equal guest parameters within pairs, optimized placement reduces expected departure time relative to the best fixed stride by a mean of 7.9% when high-risk contacts are dispersed and 13.9% when they are consecutive. The free terminal passage can favor leaving a longer sequence of farewells until after the last coat operation.

Keywords: checkpoint scheduling, renewal, dynamic programming, leave-taking

1 Introduction

A guest at a dinner party can have completed seven farewells and still be required to make eight. The eighth person asks whether the guest has seen the new kitchen. By the time the guest has seen it, found their glass again, and returned to the hall, several of the first seven have grounds to believe that the departure was called off. The guest starts another circuit. A person who had put on a coat and been seen doing so by the host can sometimes resume the circuit where it stopped.

Schegloff and Sacks describe the work speakers do to reach a point at which a conversation can close, including the chance to introduce a new topic during a possible pre-closing [3]. Here the conversation is one operation in a sequence. A new topic can impose a cost on people who had already said goodbye. We represent the host's acknowledgment of a visible coat donning as a checkpoint: the guest starts again from the last acknowledgment when a farewell is interrupted. The coat may come off during the renewed conversation. The acknowledgment persists.

The timing problem resembles checkpoint placement in computing [2, 4]. In those models a checkpoint saves computations already performed against the possibility of a later interruption. Our work units are ordered and heterogeneous. Interruption occurs at a contact, and the time spent on the contact that caused it includes the time taken to discuss whatever came up. The checkpoint can itself cause an interruption. For instance, seeing the coat can remind the host that the guest has left something in the kitchen.

We show that under independent restarts the best policy can be specified before the first farewell, even when the guest is allowed to revise the plan after every attempt. A shortest-path reformulation gives a linear-time algorithm and an exact break-even condition for a single extra coat donning. The calculation also distinguishes two dinner parties with the same guests and the same individual habits. Sitting the guests who tend to reopen conversation together changes both the expected time to leave and the value of giving up regular coat intervals.

2 Model

Number the $n \geq 1$ guests in the order in which they are encountered along a prescribed route. The host witnesses coat operations and need not be on the route. The guest starts at position 0 and passes through the door immediately on completing contact n. The terminal passage takes no time and cannot fail. We measure time from the start of the first farewell.

A farewell attempt at position h succeeds with probability $p_h \in (0, 1]$ and has mean duration $a_h > 0$. On success the guest advances to the next position, with the option to make a checkpoint first. On failure the current exchange runs its course and the guest starts the current block again. A checkpoint attempt after a successful contact costs $\kappa \geq 0$ in expected time. With probability $r \in (0, 1]$ the host acknowledges the coat and the guest banks all farewells so far. With probability $1 - r$ the coat operation reopens conversation and the guest returns to the last banked position. Durations include any time spent returning to that position. A new checkpoint can be attempted only after a new farewell has succeeded. Time spent taking the coat off to stage a repeat donning is included in κ.

We take the outcomes and durations of distinct operations to be independent conditional on their types and positions. Duration can depend on the outcome of the same operation. In particular, a failed farewell can take substantially longer than a successful one, provided a_h is the unconditional mean at that contact. The parameters are known. The objective is minimum expected departure time, with every person on the route to be farewelled. A wave from the door reaches a different service level and is outside this objective.

3 Renewal costs and optimal policies

Let $C(i, j)$ be the time in expectation to complete the contacts after position i through j with no intermediate checkpoint, restarting from i on failure. Define the prefix quantities $P_0 = 1$, $S_0 = 0$ and

$$P_j = \prod_{h=1}^{j} p_h, \qquad S_j = \sum_{h=1}^{j} a_h P_{h-1}.$$

Lemma 1 (Block cost)**.** *For $0 \leq i < j \leq n$,*

$$C(i, j) = \sum_{h=i+1}^{j} \frac{a_h}{p_h \cdots p_j} = \frac{S_j - S_i}{P_j}. \tag{1}$$

The block cost with a required checkpoint at $j < n$ is $(C(i, j) + \kappa)/r$.

Proof. On one pass, contact h is attempted with probability P_{h-1}/P_i. Its duration conditional on reaching it has mean a_h. The mean pass duration is therefore $(S_j - S_i)/P_i$, and the probability of

finishing the pass is P_j/P_i. Write the mean total time as the mean first pass duration plus the probability of failure times the mean total time. Solving gives Eq. (1). To add a checkpoint, first complete the block, then attempt the checkpoint. The mean cost is $C(i,j)+\kappa$ plus, with probability $1-r$, the mean cost of beginning the same procedure again. ∎

The block cost is sensitive to the order of guests. If a farewell at the end of a block fails, the retry budget includes every earlier farewell in that block. Permuting contacts preserves the probability that a pass succeeds and changes the expected time spent reaching a failed contact.

Theorem 2 (Optimality of a fixed partition)**.** *An optimal policy among all history-dependent policies is given by a partition $0=b_0<b_1<\cdots<b_m=n$. It makes a checkpoint after each $b_1,\dots,b_{m-1}$ and exits after b_m. Its value is the minimum over partitions of*

$$C(b_{m-1},n)+\frac{1}{r}\sum_{v=1}^{m-1}\bigl(C(b_{v-1},b_v)+\kappa\bigr). \tag{2}$$

Proof. Use state (i,k), where i is the last banked position and k the last successfully completed position in the current block. A failure sends the state to (i,i). From (i,k) with $i<k<n$, the guest can attempt a checkpoint or the next farewell. At $k=i$ only the next farewell is allowed. At $k=n$ the route ends. The state and the selected operation determine the next-state law and expected time.

This is a finite stochastic shortest-path problem [1]. From any state, a run of $2n$ successful operations suffices for absorption under any policy, since each checkpoint needs a new farewell. The conditional probability of absorption in each such window is at least $(\min\{r,p_1,\dots,p_n\})^{2n}>0$. This gives a uniform geometric bound on the operation count and, using the finite maximum of the one-operation mean durations, a uniform bound on expected residual cost. The finite-horizon value functions converge uniformly to the infinite-horizon value function. Selecting a minimizing operation in each state of the limiting Bellman equation gives a deterministic stationary policy that attains this value.

Fix such a policy. From (i,i), consider its decisions on a pass in which all operations succeed. Its first checkpoint is at a definite position j, or the policy goes directly to n. After any failure before banking j, it returns to (i,i) and uses the same decisions. On banking j it repeats this construction from (j,j). This yields a fixed partition, whose cost follows from Lemma 1. Every such partition is feasible. ∎

3.1 The marginal value of a coat operation

Consider one block $(i,j]$ that ends at the door. Inserting a checkpoint after k splits it into $(i,k]$ and $(k,j]$. From Eq. (1),

$$C(i,j)=C(i,k)\frac{P_k}{P_j}+C(k,j).$$

The change is beneficial exactly when

$$\frac{P_j}{P_k}<\frac{r}{1+\kappa/C(i,k)}. \tag{3}$$

The left side is the probability of completing the remaining farewells in one pass. On the right, a long stretch of farewells already completed raises the value of preserving them. If the coat operation is less reliable than the entire remainder of the route, no nonnegative checkpoint fee makes this particular split strictly beneficial. This condition compares one split with no split. Downstream checkpoints can change the comparison.

4 Exact optimization

Let D_j be the least expected time to bank a checkpoint at $j < n$, with $D_0 = 0$. Let F_n be the least time to reach the door. Equation (2) gives

$$\begin{aligned} D_j &= \min_{0 \le i < j} \bigl\{ D_i + (C(i,j) + \kappa)/r \bigr\}, \quad 1 \le j < n, \\ F_n &= \min_{0 \le i < n} \bigl\{ D_i + C(i,n) \bigr\}. \end{aligned} \tag{4}$$

The indices are the vertices of a directed acyclic graph with one arc for each eligible pair. A direct scan of incoming arcs costs $O(n^2)$ operations. The ratio form in Eq. (1) permits a smaller computation.

Theorem 3 (Linear-time evaluation)**.** *The value F_n and an optimal partition can be computed in $O(n)$ arithmetic operations and $O(n)$ space.*

Proof. For $j < n$ put $x_j = 1/(r P_j)$. Substitution in Eq. (4) gives

$$D_j = \kappa/r + S_j x_j + \min_{i<j} \{ D_i - S_i x_j \}. \tag{5}$$

The lines have strictly decreasing slopes $-S_i$ in insertion order. The query arguments x_j are nondecreasing. Keep the candidate lines in a deque. While the second line is no higher than the first at a query, remove the first line. It cannot become preferable at a later query because its slope is greater. The remaining first line attains the minimum. On inserting a new line, discard the current last line whenever its intersection with its predecessor is at or to the right of its intersection with the incoming line. Such a line has no interval of strict optimality. This maintains increasing intersection arguments among the remaining lines. An insertion or removal is charged once to the line involved.

Evaluate the terminal minimum by a single scan of all the D_i values. Its argument $1/P_n$ may lie to the left of the preceding query when $r < 1$, so the discarded lines cannot be ignored at the door. Keep the minimizing index for every evaluated state, and trace indices back from n to recover the partition. ∎

The complexity in Theorem 3 is an arithmetic count. For long routes the raw products P_j may underflow in floating point, and consecutive weighted prefix sums can become indistinguishable. The direct scan can instead compute arc costs using $C(i,j) = (C(i,j-1) + a_j)/p_j$ with $C(i,i) = 0$. The numerical study below uses short routes.

5 Homogeneous routes

For $p_h = p < 1$ and $a_h = a$, a block of ℓ contacts has cost

$$H_\ell = \frac{a(p^{-\ell} - 1)}{1 - p}, \qquad H_{\ell+1} - H_\ell = a p^{-(\ell+1)}.$$

For reliable checkpoints and a fixed number of blocks, the optimal lengths differ by at most one. Transferring a contact from a block that is at least two contacts longer than another to the shorter block strictly reduces cost. With $r < 1$, the terminal block receives a discount: its time is unscaled. For an intermediate block of length $\ell > 1$ and a terminal block of length $f > 1$, the two one-contact transfers imply

$$\left| f - \ell - \frac{\log(1/r)}{\log(1/p)} \right| \le 1 \tag{6}$$

at an optimum with that number of blocks. To obtain the lower bound, transferring one contact to the terminal block costs $a p^{-(f+1)}$ and saves $a p^{-\ell}/r$. The reverse transfer gives the upper bound. The free exit can therefore justify a longer last batch even on a route whose guests are otherwise interchangeable.

Proposition 4 (Long-route coat interval)**.** *Define $K_0 = 0$ and $K_\ell = \ell(H_{\ell+1} - H_\ell) - H_\ell$ for $\ell \ge 1$. The limiting minimum expected time per guest is*

$$\lim_{n \to \infty} \frac{F_n}{n} = \min_{\ell \ge 1} \frac{H_\ell + \kappa}{r\ell}. \tag{7}$$

An interval of ℓ attains the minimum if $K_{\ell-1} \le \kappa \le K_\ell$. In particular, the minimizing interval has length one when $\kappa < a(1-p)/p^2$.

Proof. Write g_* for the minimum on the right of Eq. (7). It is attained since H_ℓ/ℓ diverges. Repeating a minimizing interval up to a bounded terminal block gives an upper bound $n g_* + O(1)$. Each intermediate block costs at least g_* times its length. For a terminal block of any length f, the difference $H_f - g_* f$ is bounded below, again by exponential growth of H_f. This gives the matching lower bound.

The difference between the per-guest costs at $\ell + 1$ and ℓ has the sign of $K_\ell - \kappa$. The K_ℓ increase strictly because H_ℓ is strictly convex in ℓ. Finally, $K_1 = H_2 - 2H_1 = a(1-p)/p^2$. ∎

For $a = 30$ s and $p = 0.60$, this first cut-off is 33.3 s. A reliable 30 s coat operation makes it optimal to put on the coat after every farewell except the last, for any n. Indeed, for $r = 1$, splitting a block of length $\ell \ge 2$ into blocks of lengths $\ell - 1$ and one saves at least K_1 before the checkpoint fee. Taking longer to remove the coat is charged at the same rate as taking longer to put it on.

6 Computational experiments

6.1 Eight-contact instance

We first use a route with attempt times $(24, 40, 26, 30, 35, 20, 37, 28)$ s and success probabilities $(.96, .98, .97, .55, .97, .94, .96, .57)$. The coat operation costs 30 s. The fixed-stride comparison selects the best stride from 1 through n for each instance, with checkpoints at all positive multiples below n. The exit-only rule makes no intermediate checkpoint.

We evaluated the three rules in Table 1 using Eq. (2) and by simulating 50,000 independent departure runs per row. Each attempted operation takes its specified mean duration in the simu-

lator. The intervals are normal Monte Carlo intervals for simulation error, computed with the sample standard deviation of departure times. The analytic times do not depend on these simulations.

Table 1: Exact costs and simulated mean times on the eight-contact route. Times and Monte Carlo half-widths are in seconds.

r	Rule	Coat after	Exact	Simulated (95% CI)
1.00	Exit only	None	681.95	$\mathbf{679.40 \pm 4.56}$
1.00	Fixed stride	3, 6	391.41	$\mathbf{390.85 \pm 0.83}$
1.00	Optimal	3, 7	367.31	$\mathbf{367.61 \pm 0.63}$
0.70	Exit only	None	681.95	$\mathbf{687.91 \pm 4.61}$
0.70	Fixed stride	3, 6	509.12	$\mathbf{508.82 \pm 1.60}$
0.70	Optimal	3, 7	503.67	$\mathbf{503.54 \pm 1.74}$

The optimal rule moves the second checkpoint one guest closer to the door. The two guests most likely to reopen conversation then both start a block. With reliable acknowledgment the move saves 24 s in expectation at the same coat count. Five of the six simulation intervals include the exact mean. The second exit-only simulation misses high, though its exact cost is unchanged because it never uses a coat operation.

6.2 Paired synthetic routes

We generated 120 sets of 24 contacts, each with six high-risk and 18 low-risk guests. Each attempt time is independently uniform on $[15, 45]$ s. Success probabilities are independent uniforms on $[0.48, 0.68]$ for the high-risk guests and $[0.94, 0.995]$ otherwise. For each set we made two routes. In the dispersed route the high-risk guests occur every fourth position, with a uniformly drawn starting position from 1 to 4. In the clustered route they occupy six consecutive positions, with start drawn uniformly from 1 to 19. Within each type the guest order is the same in both routes. The pairing holds the sum of the attempt-time means and the full-pass probability constant.

All policies in Table 2 were scored by their exact expected times on the same routes. The every-guest rule checkpoints after every contact except the last. Counts refer to distinct checkpoint positions specified by a rule. Under unreliable acknowledgment a guest may attempt the same coat operation more than once.

Table 2: Mean model departure time (min) and mean number of distinct coat positions at $\kappa = 30$ s and $r = 1$. Each layout has 120 routes.

Rule	Dispersed (min)	Coat positions	Clustered (min)	Coat positions
Exit only	169.44	0.00	257.04	0.00
Every guest	26.07	23.00	26.07	23.00
Fixed stride	19.83	6.78	22.01	10.23
Optimal	18.21	5.59	18.95	6.08

The mean of the within-route percentage savings over the best fixed stride is 7.9% for dispersed guests and 13.9% for clustered guests. Regular spacing copes with the cluster by shortening the coat interval across the whole route. The optimized policy needs about four fewer coat positions than the best fixed stride in the clustered case.

The left panel of Figure 1 shows the saving as the fee varies. At zero cost the stride-one rule is optimal. The right panel shows how many checkpoint positions the optimal rule uses. The clustered layout uses fewer positions than the dispersed one once the fee reaches 45 s on the tested grid.

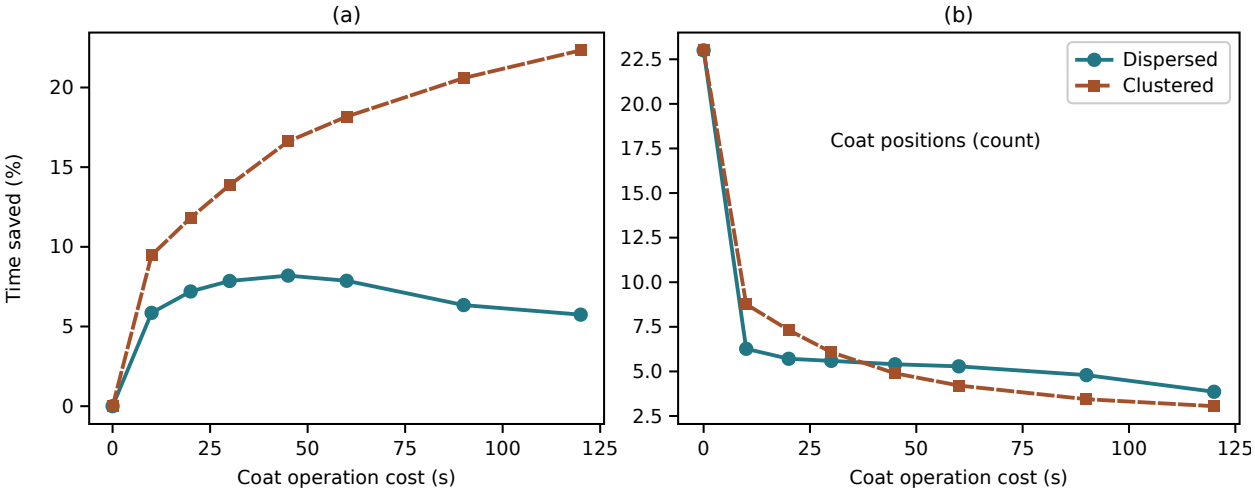


Figure 1: Sensitivity to the cost of a coat operation ($r = 1$). Left: mean of per-route percentage savings over the best fixed stride. Right: mean number of checkpoint positions in the optimal partition. Each point uses the same 120 paired sets of guests.

Reducing r raises the value of reserving some farewells for the free exit. In Figure 2, the clustered layout starts out slower under both rules. Under optimal placement it becomes faster than dispersal at $r = 0.70$ and 0.50 on the tested grid. The difference at $r = 0.70$ is 0.40 min. The reversal pertains to these paired instances. It does not require a reversal in the guests' individual probabilities, which are the same in both layouts.

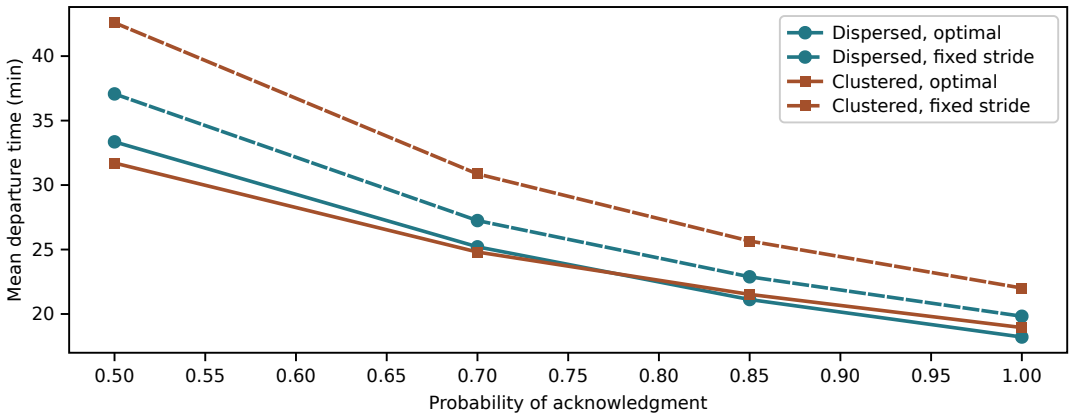


Figure 2: Mean departure time versus probability of acknowledgment, with the fee fixed at 30 s. Lines join exact means over each 120-route group.

For an independent check on the linear algorithm, we enumerated all partitions of 240 additional random routes of one to nine contacts. The minimum cost and the evaluated cost of the reconstructed partition agreed with the algorithm's reported value to a relative tolerance of 10^{-12} on every instance. The route generator, simulator and enumeration check use separate pseudorandom streams. Appendix A gives the inputs needed to regenerate the results.

7 Discussion

The long exit-only times in Table 2 come from repeatedly forfeiting earlier farewells. They do not estimate the usual length of a dinner-party departure. Actual guests may keep partial credit without an acknowledgment, shorten a farewell on a later pass, or leave after an unusually long interruption. All three behaviors reduce the loss from an unsuccessful pass. Estimating that loss would require following the same guest across attempts, including contacts an observer might otherwise record as repetitions and discard.

The independence assumption also rules out a host becoming harder to convince after witnessing several coat operations. If r falls with the number of previous attempts, a guest can learn something operationally relevant from the attempt history. The fixed-partition result would then no longer follow from the same state description. Counting coat operations and counting acknowledgments are distinct measurement tasks in that setting.

In our model, the appropriate place to put on the coat can be well inside the house. The position depends on how much of the route the guest would have to walk again after the next conversation.

A Reproduction details

The 120 paired sets use the Python Mersenne Twister with seed 81423. For each set, draw an attempt time then a success probability for each of the 18 low-risk contacts, then for each of the six high-risk contacts. Draw the two offsets in order (dispersed, clustered) using uniform integer sampling, then fill each layout in the order of the draws within each type. All uniform real draws use Python's uniform method. For Table 1, seed 41385 and take one uniform draw per attempted farewell or checkpoint in table order. A draw below the operation's success probability advances

the state. For the exhaustive check, seed 23988. Draw the route length uniformly from 1 through 9 and the attempt times uniformly from the integers 5 through 45 s. Draw the contact probabilities from $\{0.4, 0.55, 0.8, 0.98, 1\}$, the fee from $\{0, 15, 30, 120\}$ s and the reliability from $\{0.5, 0.85, 1\}$, uniformly over each set. The program uses no data from observed gatherings.

References

[1] Dimitri P. Bertsekas and John N. Tsitsiklis. An Analysis of Stochastic Shortest Path Problems. *Mathematics of Operations Research*, 16(3):580–595, 1991. doi: 10.1287/moor.16.3.580.

[2] John T. Daly. A Higher Order Estimate of the Optimum Checkpoint Interval for Restart Dumps. *Future Generation Computer Systems*, 22(3):303–312, 2006. doi: 10.1016/j.future.2004.11.016.

[3] Emanuel A. Schegloff and Harvey Sacks. Opening up Closings. *Semiotica*, 8(4):289–327, 1973. doi: 10.1515/semi.1973.8.4.289.

[4] John W. Young. A First Order Approximation to the Optimum Checkpoint Interval. *Communications of the ACM*, 17(9):530–531, 1974. doi: 10.1145/361147.361115.