Taylor & Francis
Taylor & Francis Group

The College Mathematics Journal

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ucmj20

Haste Makes Waste: An Optimization Problem

William Q. Erickson

To cite this article: William Q. Erickson (2022) Haste Makes Waste: An Optimization Problem, The College Mathematics Journal, 53:2, 122-133, DOI: 10.1080/07468342.2021.2022955

To link to this article: https://doi.org/10.1080/07468342.2021.2022955

Published online: 03 Feb 2022.

Haste Makes Waste: An Optimization Problem

William Q. Erickson

William Q. Erickson (wqe@uwm.edu, MR ID 1470622) is a doctoral student and teaching assistant at the University of Wisconsin–Milwaukee, with a research interest in representation theory and combinatorics. He taught high-school calculus for eight years, and especially loves to hear from former students thriving in their STEM professions. He also enjoys performing on the piano for anyone who will listen, and designing board games with his brothers.

The earlier you leave home, the earlier you arrive to work—but sometimes at the cost of sitting in the car, wasting fuel, stuck in traffic with everyone else who had the same idea. Suppose instead that the goal is simply to minimize your travel time in the morning; now it becomes an interesting problem to determine your optimal departure time.

Throughout this problem, we imagine that the commute to work is a stretch of freeway along which the traffic speed depends only on time (not position). We have two goals: (1) to describe the candidates for the optimal departure time geometrically, using the graph of the traffic speed function, and (2) to determine an algorithm to find this optimal time using software. We divide our treatment into three main sections. In the first section, we solve the simplified case in which traffic speed is always positive; the second section makes adjustments to allow for traffic stopping momentarily; and the third section explores a speed function that is zero on an interval (in other words, a traffic jam).

To animate our problem, we drop in on an arbitrary home, inhabited by our two heroes Ivan and Olga. The hyper-punctual Ivan leaves for work before Olga does each day, and, despite all his quirks, has apparently retained enough of Olga's affection so that she still enjoys his company in the mornings, and thus disapproves of his leaving earlier than necessary. Just this weekend, she pointed out that Ivan's usual departure time of 6:00 a.m. (a time of day when *no one* should be doing *anything*, in her opinion) must often land him in the middle of rush-hour traffic; might not a later departure *decrease* his time on the road? Ivan, realizing that Olga is correct as usual, rises on Monday morning a new man, intent on keeping his travel time to a minimum.

Simplest case: speed always positive

Ivan grabs a pen and a napkin and begins with the simplest case, in which the traffic speed on his route is always positive. Let $v(t)$ be this speed measured in miles per hour, with t in hours. Assume that v is differentiable, and for now, assume also that $v > 0$, so that traffic is always moving at least slightly forward. Then

$$s(t) := \int_0^t v(x)\,dx$$

doi.org/10.1080/07468342.2021.2022955
MSC: 26A06

is an antiderivative of $v(t)$, giving Ivan's position in miles from home if he were to leave at $t = 0$. Finally, fix D to be the distance (in miles) of Ivan's route to work. Then a given departure time t determines Ivan's travel time $h(t)$ such that

$$\int_t^{t+h(t)} v(x)\, dx = D. \tag{1}$$

This $h(t)$ is uniquely determined by t since $v > 0$. Solving for $h(t)$ from (1), we find

$$\begin{aligned} s\big(t + h(t)\big) - s(t) &= D \\ s\big(t + h(t)\big) &= s(t) + D \qquad (2) \\ t + h(t) &= s^{-1}\big(s(t) + D\big) \\ h(t) &= s^{-1}\big(s(t) + D\big) - t. \qquad (3) \end{aligned}$$

We therefore have travel time h as a function of departure time t. (In this simplified case, we know that s^{-1} exists, since $s' = v > 0$ and hence s is strictly increasing.) This travel-time function h will play the starring role in our problem, since it is the function we are trying to minimize.

Ivan, after a few minutes and a few more napkins, discovers that we can visualize the values of h more easily than the notation suggests. He starts by sketching three graphs on the same coordinate plane (see Figure 1): first the graph of $v(t)$, then the graph of $s(t)$ (which must be strictly increasing), and finally the graph of $s(t)$ shifted upwards by D miles. Then he chooses an arbitrary departure time a. Starting from the point $(a, 0)$, he draws a vertical line segment all the way up to the shifted graph of s and leaves his pen right there, at height $s(a) + D$. Now since we want to evaluate s^{-1} at this value, Ivan next draws a horizontal line segment until he intersects the graph of s, then draws another vertical line segment downwards until it intersects the t-axis, at the t-coordinate $s^{-1}\big(s(a) + D\big)$. After connecting this last point to his original point at $(a, 0)$, he has just drawn a rectangle, whose width is $s^{-1}\big(s(a) + D\big) - a$. In other words, the width of this rectangle is precisely the travel time $h(a)$. Moreover, after shading the region enclosed by this rectangle and the graph of v, Ivan sees that this region has area D, since by definition $\int_a^{a+h(a)} v(t)\, dt = D$. Ivan calls this region R_a for future reference.

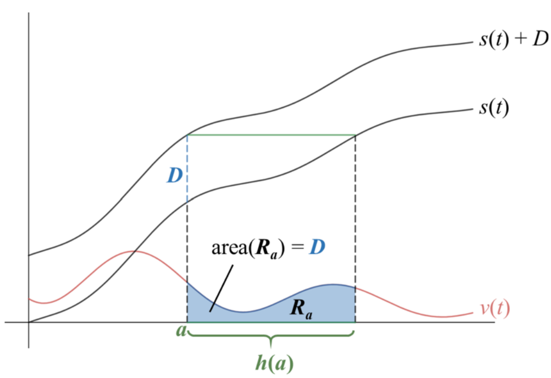

Figure 1. The width of the rectangle with left edge at $t = a$ is precisely $h(a)$. Using this construction, we have a correspondence between each departure time a and the region R_a.

This is a good time to specify the domain of h. We let $t = 0$ correspond to the earliest feasible departure time, due to sleep or breakfast or both, and we fix T to be the report time at Ivan's workplace. Then, Ivan's latest possible departure time is the unique solution ℓ to the equation $\int_{\ell}^{T} v(x)\,dx = D$, that is, $s(\ell) = s(T) - D$ or

$$\ell = s^{-1}\big(s(T) - D\big). \tag{4}$$

Hence we are trying to find the time that yields the absolute minimum value of h on the interval $[0, \ell]$. (See Figure 2.)

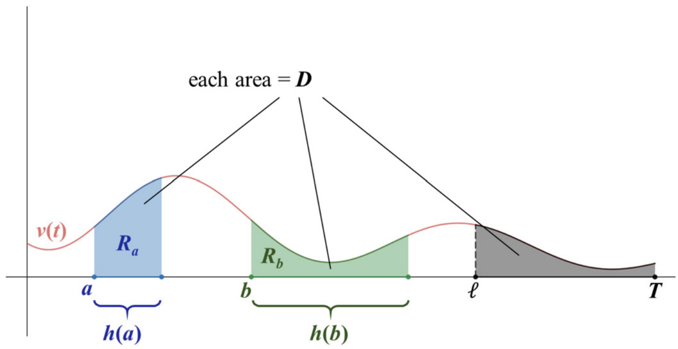

Figure 2. A visual summary of the problem; our goal is to find the departure time within the interval $[0, \ell]$ that minimizes the travel time function h. In other words, we want to find the time whose corresponding region has the least possible width. In the example above, clearly a is a better departure time than b since the width of the blue region R_a is less than the width of the green region R_b.

Since one of our goals is to characterize visually the candidates for the best departure time, this picture in Figure 2 will be crucial. To summarize, each possible departure time t corresponds to a shaded region R_t under the graph of v, with area D and width $h(t)$. The left edge of this shaded region shows the departure time t, and the right edge shows arrival time $t + h(t)$.

Characterizing the local minima geometrically. While Ivan ponders these shaded regions and makes coffee for Olga, he asks himself whether anything can be known *a priori* about the solution to this problem. One possibility, of course, is that the absolute minimum of h occurs at one of the two endpoints of its domain; as a special example, the minimum occurs at time 0 on a day when v is a decreasing function, and at time ℓ on a day when v is an increasing function.

The other possibility is that the absolute minimum of h occurs at a critical point of h, meaning that it is also a *local* minimum. To find critical points, Ivan first needs to compute the derivative $h'(t)$, and then determine the t-values at which it equals 0 or does not exist. Rather than differentiate the explicit but unpleasant formula (3), he glances a few lines above, at his equation (2). (This will just allow him to avoid the s^{-1}. He didn't like differentiating inverse functions in Calculus I, and he doesn't like it now.) Turning the napkin over, he differentiates implicitly with respect to t:

$$s\big(t + h(t)\big) = s(t) + D \qquad \text{from (2)}$$

$$s'\big(t + h(t)\big) \cdot \frac{d}{dt}\big(t + h(t)\big) = s'(t) \qquad \text{by the chain rule}$$

$$v\big(t + h(t)\big) \cdot \big(1 + h'(t)\big) = v(t)$$

$$h'(t) = \frac{v(t)}{v\big(t + h(t)\big)} - 1. \tag{5}$$

Equation (5) shows that h' exists and is continuous everywhere, due to our earlier assumption that v is not only differentiable but also positive everywhere. Therefore, the only critical points for h are those at which $h' = 0$.

Looking closer at (5), we observe that $h'(t)$ is positive whenever $v(t) > v\big(t + h(t)\big)$. Visually, this occurs whenever the left edge of the shaded region R_t is taller than the right edge. Intuitively, this makes sense: h' positive means that at the present moment, waiting to depart is causing Ivan's travel time to *increase*, which is what we would expect if we know that the traffic is going to become worse. Likewise, $h'(t)$ is clearly negative whenever $v(t) < v\big(t + h(t)\big)$.

This leaves the most important case: $h'(t) = 0$ if and only if $v(t) = v\big(t + h(t)\big)$. In visual terms, this tells us the following:

Lemma 1. *The travel-time function $h(t)$ has a critical point at t_0 if and only if the shaded region R_{t_0} has left and right edges of equal height.*

Ivan's next goal, after remembering too late to place a mug under the coffeemaker, is to distinguish the local minima from the rest of the critical points. Intuitively, a local minimum for h occurs at the "best" departure time t_0 within some open interval: in the moments just before t_0, waiting to depart should *decrease* travel time, whereas just after t_0, waiting to depart should *increase* travel time. But from our observations above, this is equivalent to the condition that at t_0, the sign of $v(t) - v\big(t + h(t)\big)$ changes from negative to positive. This suggests that as t ranges over a sufficiently small interval containing t_0, the left edge of the region R_t is growing faster than the right edge, meaning that $v'(t_0) > v'\big(t_0 + h(t_0)\big)$.

Hoping to verify this intuition, Ivan now computes $h''(t)$, since by the second derivative test, a critical point t_0 is a local minimum if $h''(t_0)$ is positive. Out of napkins, he grabs a coffee filter and calculates using the quotient rule on (5):

$$h''(t) = \frac{v\big(t + h(t)\big) \cdot v'(t) - v(t) \cdot v'\big(t + h(t)\big) \cdot \big(1 + h'(t)\big)}{\big[v\big(t + h(t)\big)\big]^2}.$$

But we are assuming that we already have a critical point t_0, and so Ivan evaluates the above expression at this t_0 to see what happens: first of all, $h'(t_0) = 0$, and moreover, we know from Lemma 1 that $v(t_0) = v\big(t_0 + h(t_0)\big)$. Using these substitutions, Ivan rewrites and reduces:

$$h''(t_0) = \frac{v'(t_0) - v'\big(t_0 + h(t_0)\big)}{v(t_0)}.$$

Since $v(t_0)$ is positive, Ivan concludes that if t_0 is a critical point, then $h''(t_0)$ is positive whenever $v'(t_0) > v'\big(t_0 + h(t_0)\big)$, exactly as he had predicted. In other words, a critical point t_0 is a local minimum if v' is greater at the left edge of R_{t_0} than it is at the right edge. So we need only compare the acceleration of the traffic at Ivan's departure time and his arrival time; or more visually, we need only compare the slopes of the tangent lines to v at the two upper vertices of R_{t_0}.

(A word about the unlikely case where these two slopes are equal, i.e., where $h''(t_0) = 0$, since then the second derivative test is inconclusive. We certainly could

use the *higher-order derivative test*, as explained in [2], to determine whether or not the critical point is a local minimum; in fact, a quick inductive calculation shows that at a critical point t_0, the sign of the higher-order derivative $h^{(n+1)}(t_0)$ is just the sign of the difference $v^{(n)}(t_0) - v^{(n)}\big(t_0 + h(t_0)\big)$. But since our goal here is a simple *visual* characterization from the graph of v, we leave this as a special case instead of including it in the result below.)

Just as Ivan hears Olga stirring from bed, he flips over the nearest coaster and jots down his result:

Theorem 1. *Let t_0 be some time in the interval $(0, \ell)$, with corresponding region R_{t_0}. (Exclude the special case mentioned above where $v'(t_0) = v'\big(t_0 + h(t_0)\big)$, although even then we could just replace condition 2 below with the higher-order derivative test.) Then t_0 is a local minimum for h—and thus a candidate for the optimal departure time—if and only if:*

1. *The left and right edges of R_{t_0} are of equal height; and*
2. *The tangent line to v at the upper-left vertex of R_{t_0} has greater slope than that at the upper-right vertex.*

An algorithm for the solution. A groggy Olga stumbles into the kitchen, only to have a mug of mediocre coffee and a defaced coaster thrown before her. Ivan, inordinately proud of his result, shows her everything he now knows about the candidates for optimal departure time. Olga listens patiently, stifles a tired yawn, wonders whether the term "theorem" might be a bit grandiose, and then asks the obvious question: So what time will he have to leave today?

Ivan, having somewhat forgotten that the goal is actually to find the solution and not just to visualize the candidates, bounds out of the room and reappears in a second with Olga's laptop. Olga blinks a few times, dims the screen brightness, and fires up Mathematica.

Example 1. Ivan (on principle) will not leave home before 6 a.m., and must be at work by 8 a.m. His workplace is located just off the freeway 30 miles away, so $D = 30$, and his incredible live traffic app informs him that today's freeway speed will be modeled by $v(t) = 50 + 36t^2e^{-2t}$ where t is given in hours since 6 a.m. (so $T = 2$). Olga can now define the functions s and h, and use (4) to calculate ℓ:

```
d = 30; T = 2;
v[t_] := 50 + 36*t^2*Exp[-2*t];
s[t_] := Integrate[v[x], {x, 0, t}];
h[t_] := InverseFunction[s][s[t] + d] - t;
l = InverseFunction[s][s[T] - d]
1.438
```

From here, Olga plots the function h on the interval $[0, 1.438]$, obtaining Figure 3(a), from which she sees that neither endpoint is the absolute minimum; instead, that minimum occurs near $t = 0.8$, and so Olga uses `FindMinimum` to locate the local minimum `{0.548, {t -> 0.751}}`. Therefore, Ivan's shortest possible travel time is 0.548 hours or 33 minutes, which results from his leaving at $t_0 = 0.751$ or 6:45 a.m.

Olga begins to close the laptop, but Ivan is dying to see his theory confirmed in practice, and so he plots the function v along with the region $R_{0.751}$ corresponding to Olga's solution (see Figure 3(b)). He is not disappointed. The solution does indeed satisfy both conditions of his theorem: the left and right edges of $R_{0.751}$ have equal

height (54.521), and furthermore, the tangent line at the upper-left vertex clearly has a greater slope than that at the upper-right vertex.

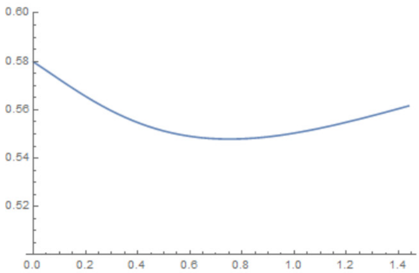

(a) The graph of the travel-time function $h(t)$.

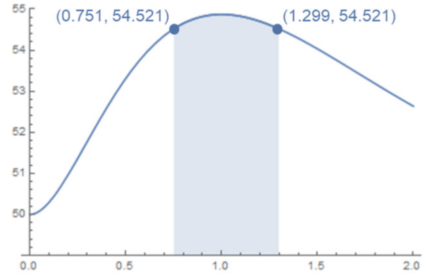

(b) The graph of the traffic speed $v(t)$, along with the shaded region $R_{0.751}$, clearly satisfying both conditions in Theorem 1.

Figure 3. The results from Example 1.

Next case: speed momentarily zero

Olga, by now mostly awake and sipping Ivan's latest attempt at coffee, has been scratching her head ever since hearing of his $v > 0$ assumption in the beginning. Glancing out the kitchen window at the motionless gridlock on the freeway, she raises one cranky eyebrow in a manner which, after years of marriage, Ivan knows to translate as "Your traffic problem doesn't account for *traffic jams?*" Ivan admits that this is a rather crippling flaw, but still he insists on keeping v strictly greater than 0, for the two reasons he has already encountered:

1. If v is allowed to take the value 0, then s may not be strictly increasing anymore, and so we can no longer guarantee the existence of s^{-1}. Since h is defined in terms of s^{-1}, this means that h may not exist everywhere.
2. Furthermore, we have seen that $h'(t) = \frac{v(t)}{v(t+h(t))} - 1$, and so $h'(t_0)$ does not exist whenever $v(t_0 + h(t_0)) = 0$. Hence, h would have critical points besides those at which $h' = 0$, which would invalidate Ivan's clean little theorem.

Ivan may believe in beauty over realism, but Olga, who shares no such scruples, reminds him of what someone once said Einstein once said about elegance and tailors; and, never one to complain about a problem without offering a fix, she suggests that they first consider functions v which assume the value 0 at discrete times, rather than on an interval. Although scarcely more realistic than before, still, this will relieve Ivan's mind of his first objection entirely: after all, as long as there is no interval on which $v = 0$, then s is still *strictly* increasing. (The quick 6 a.m. proof from Olga: for $a < b$, the integral $\int_a^b v(t)\, dt$ is positive since v is continuous, nonnegative, and not identically 0 on the interval (a, b). Therefore, $s(a) < s(b)$. QED and more coffee, please.) From the practical perspective, then, Olga's code will still produce a solution since s^{-1} still exists; it is only Ivan's description of local minima which may be at risk.

But that is a risk Ivan is not willing to take. To determine whether a salvage is required, we must deal with the following issue: $h'(t_0)$ does not exist when $v(t_0 + h(t_0)) = 0$, and hence such a t_0 is a critical point for h which might not obey the equal-edge-heights property in Lemma 1. Visually, this occurs whenever the right edge of

region R_{t_0} has height 0, which means Ivan arrives at work just as traffic momentarily stops moving. Ivan and Olga suddenly notice from (5) that as long as $v(t_0)$ is not also 0, then $\lim_{t \to t_0} h'(t) = \infty$. The two of them exchange a puzzled glance; what does this mean? But they are reassured by a moment's thought at the infinitesimal level, where the area of R_{t_0} is considered as a right Riemann sum of arbitrarily thin rectangles:

Delaying Ivan's departure time t_0 by an arbitrarily small increment results in "losing" the leftmost rectangle in R_{t_0}; now, in order to maintain the constant area D in this new region, we need to compensate by extending the right edge of R_{t_0} farther to the right. But since the rightmost rectangle in R_{t_0} has height 0, extending this edge to the right adds 0 area to the region: effectively, we would have to move that edge infinitely far to the right in order to compensate for the lost area on the left.

The result: although $v(t_0 + h(t_0)) = 0$ does imply that t_0 is a critical point for h, nevertheless—at least as long as $v(t_0)$ is not also 0—this critical point cannot possibly be a local minimum for h, since travel time at that moment is increasing at an infinite rate. Hence, this situation poses no threat to Ivan's theorem.

Nor is there any danger in the reverse scenario where $v(t_0) = 0$ and $v\big(t_0 + h(t_0)\big) \neq 0$. In this case, clearly $h'(t_0) = -1$, so t_0 is not a critical point. (If Olga and Ivan were aware that there is a reader, then they would encourage that reader to verify that -1 makes sense by imagining R_{t_0} as a left Riemann sum, as in the thought experiment above.)

Finally, what about the unlikely possibility that traffic speed is 0 at *both* the departure time t_0 *and* the arrival time $t_0 + h(t_0)$? In this case, our hypotheses imply that v has a local minimum at each of those two times, and since v is differentiable, this means that $v'(t_0) = v'(t_0 + h(t_0)) = 0$. But this is precisely the case excluded from Ivan's theorem; moreover, as the persistent reader may enjoy verifying via L'Hospital's Rule and some limit manipulations, it turns out that such a t_0 is a local minimum for h under precisely the same higher-order derivative conditions mentioned in the first section. Therefore, Ivan's theorem still applies.

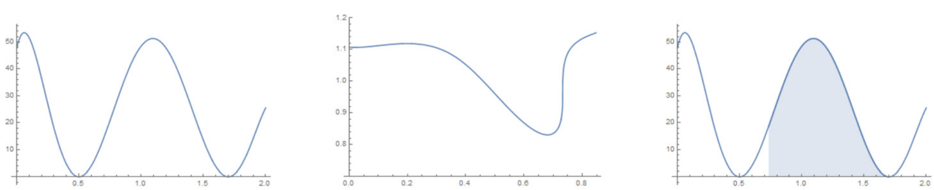

(a) Traffic speed $v(t)$, with zeros at 0.5 and 1.7.

(b) The travel-time function $h(t)$. Note the vertical tangent line at $t = 0.733$.

(c) The region $R_{0.733}$ with right edge having height 0, corresponding to the fact that $h'(0.733) = \infty$.

Figure 4. Visualizing Example 2.

Example 2. The details are the same as in Example 1, except this time the morning traffic speed model is

$$v(t) = (t - 0.5)^2(t - 1.7)^2(66 + 630t - 300t^2), \quad 0 \leq t \leq 2.$$

(See Figure 4(a).) There are two times at which traffic momentarily halts, namely $t = 0.5$ (or 6:30 a.m.) and $t = 1.7$ (or 7:42 a.m.). Olga's code from the last example

determines $\ell = 0.847$ and generates a plot for the travel-time function h (see Figure 4(b)). This time, however, she and Ivan notice what seems to be a vertical tangent line near $t = 0.75$. But this is no surprise: they already predicted that $h'(t) = \infty$ whenever $v(t + h(t)) = 0$. In this case, they are convinced that the zero of v at 1.7 is the cause of this vertical tangent, but just to be sure, they have the computer solve $t + h(t) = 1.7$ for t, which indeed yields $t = 0.733$. That is, the right edge of $R_{0.733}$ has height 0, as shown in Figure 4(c).

In this example, there are actually three local extrema for h, not just one. Clearly the optimal departure time is the local minimum near $t = 0.7$, and the Mathematica output `{0.831, {t -> 0.681}}` translates to a minimal travel time of 50 minutes if Ivan leaves home at 6:41 a.m. (See Figure 5(a).) But there is another, barely perceptible local minimum for h near the vertical axis of the graph, at $t = 0.020$, and so Ivan thinks it is worth noticing that $R_{0.020}$ is another incarnation of his Theorem 1 (see Figure 5(b)). Finally, the perceptive reader will notice a local *maximum* on the graph of h, at 0.200, and sure enough, as a critical point, its corresponding region $R_{0.200}$ has edges of equal height—but now we have the *opposite* of criterion 2 from Ivan's theorem, since we are dealing with a maximum instead of a minimum (see Figure 5(c)).

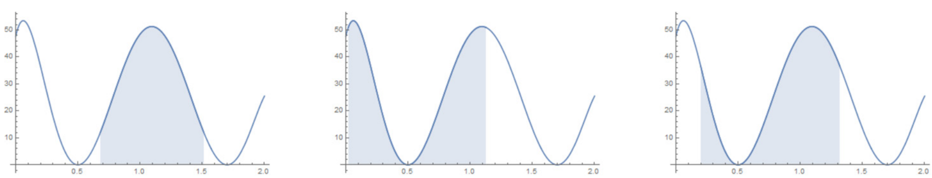

(a) The region $R_{0.681}$ corresponding to the optimal departure time. Note that both conditions in Ivan's theorem are satisfied.

(b) The region $R_{0.020}$ corresponding to the other local minimum for h: another instance of Ivan's theorem.

(c) The region $R_{0.200}$, corresponding to the local maximum for h, satisfies condition 1 but not condition 2 of the theorem.

Figure 5. The three local extrema for h in Example 2, viewed on the graph of v.

Final case: traffic jams

Having allowed the speed function v to assume the value 0 at discrete times, we are now ready to consider speed functions that are 0 on an interval, i.e., traffic jams. Ivan's first objection from the previous section is now a real obstacle: if $v = 0$ on some interval $[b, c]$, then s is constant on this interval, and so the inverse function s^{-1} is no longer defined. Since the function h we are trying to minimize is defined in terms of s^{-1}, this means that neither Olga's code nor Ivan's theorem will apply anymore.

Ivan sighs; that theorem has become like a child to him. But as Olga and a third cup of coffee point out, all is not lost: the notion of "arrival time" is still well-defined, with or without traffic jams. After all, given a certain distance from home, there must still be a unique time at which Ivan's car *first* attains that distance, regardless of whether he spends more time stuck at that same position. It seems to Olga that we may be able to replace s^{-1}, which assigned the unique time input to each position output, with a different sort of function that assigns the *earliest* possible time input to each position output. Ivan, eyes aglow with admiration, and suddenly reminded of what first drew

him to Olga all those years ago, is already scrawling across the nearest paper plate, translating her idea as

$$s^{-}(y) := \inf\{t \mid s(t) \geq y\}. \tag{6}$$

The advantage is that in our problem, this s^{-} always defines a function, whereas the usual inverse s^{-1} sometimes does not. Ivan is apt to start calling s^{-} the "Olga inverse," although Olga (blushing slightly) is fairly certain that she is merely adapting the notion of a *generalized inverse* which she once encountered while reading through [**1**].

While Olga tinkers with her code to make this change from s^{-1} to s^{-}, Ivan tries to imagine the graph of this s^{-}. Now, if there is no interval on which $v = 0$, then s^{-} is the same as s^{-1}, meaning that its graph is just the reflection of the graph of s across the line passing through the origin at 45°. Ivan therefore supposes that there *is* some interval $[b, c]$ on which $v = 0$, meaning that s is constant on $[b, c]$. It is this horizontal segment in the graph of s which requires special care when reflecting to obtain s^{-}. Ivan knows that this entire segment must be collapsed to a single point on the graph of s^{-}, since s^{-} assigns only the *least* corresponding input to each output of s; specifically, when the graph of s is reflected, the horizontal segment over the interval $[b, c]$ is mapped to the single point $(s(b), b)$. As a result, s^{-} has a jump discontinuity at $s(b)$, caused by the traffic jam on $[b, c]$, and the graph jumps by precisely $c - b$ hours (the duration of the traffic jam). Moreover, despite this discontinuity, the one-sided derivatives of s^{-} at $s(b)$ are both ∞, since $s'(b) = s'(c) = 0$. (See Figure 6, as well as [**1**] for more details.)

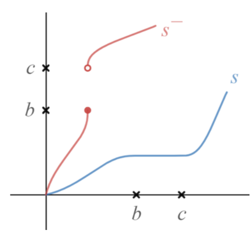

Figure 6. The graph of s in blue, and its generalized inverse s^{-} in red.

Ivan is eager to run an example, so he looks over and asks Olga what exactly is taking her so much longer than usual; Olga growls and replies, almost sweetly, that she had the code written minutes ago, but waiting for the computer to perform a `MinValue` calculation for every single input is bound to take a very long time, and perhaps her idea was far better in theory than in practice. Ivan, who considers this one of the greatest compliments an idea can receive, smiles affectionately; Olga glares at him and snaps the computer shut. Even one brief interval where $v = 0$, she admits, disrupts their method more than she had expected; if only they could just ignore the traffic jam altogether . . .

Ivan nods slowly; then his eyes light up and he asks Olga to repeat what she just said. Olga protests—it was a foolish dream—they cannot just *ignore* traffic jams when they occur, since they very clearly *do* affect travel time. But Ivan, who has an endless supply of faith in his better half, jumps up and insists on trying. Soon both of them, looking around frantically for unused paper products and finding none, exchange a shrug and take to the nearest wall, which is overdue for a paint job anyway:

Ivan first draws the graph of v with a traffic jam on $[b, c]$. (See Figure 7(a).) Now, he says, suppose they take the domain of v and temporarily collapse the interval $[b, c]$

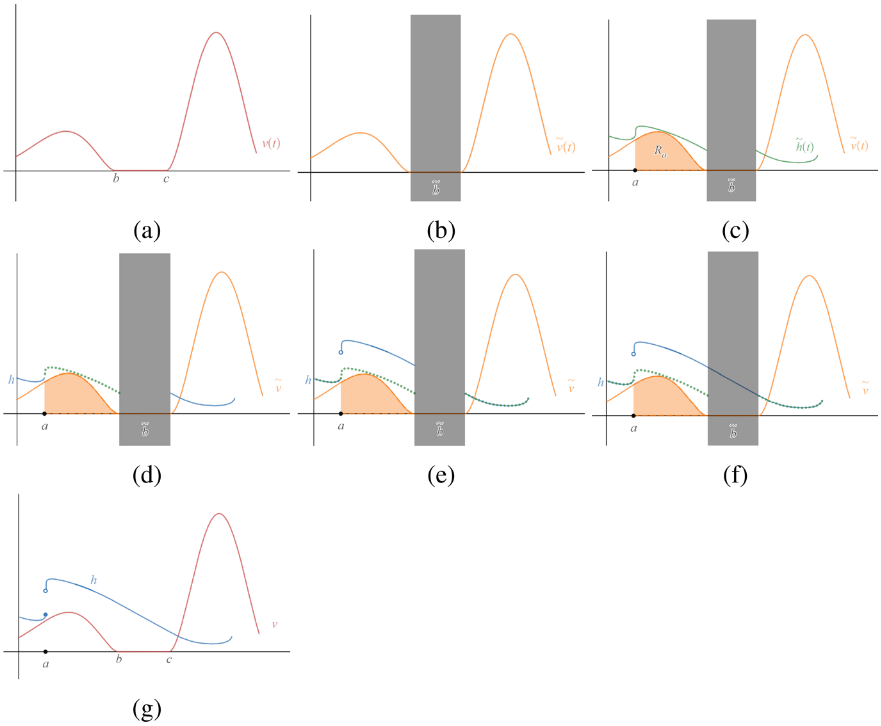

Figure 7. Given a speed function v involving a traffic jam, we can construct the graph of its travel-time function h by the steps illustrated above.

into a single point $\tilde{b}$. Even though the gap between b and c is still there on the wall, we imagine it is gone, so that an ant crawling from the left toward b would pass $\tilde{b}$ and immediately see values greater than c. Giving the name $\tilde{v}$ to this new speed function with the collapsed domain, we notice that $\tilde{v}$ agrees with v everywhere (meaning $\tilde{v}(\tilde{b}) = 0$ still). Effectively, $\tilde{v}$ is simply v with the traffic jam removed, so that $\tilde{v}$ assumes the value 0 only at the single point $\tilde{b}$. (See Figure 7(b), with the gray strip deleted from the domain.) Olga sees what is happening—they have already dealt with this kind of v in the previous section.

Back in familiar territory with this $\tilde{v}$, our heroes know they can define its corresponding travel-time function $\tilde{h}$ just as before. (See the green graph in Figure 7(c).) Now, their whole goal here is to describe the travel-time function for v, which they will still denote as h. Immediately, they are happy to observe that

$$h(t) = \tilde{h}(t) \text{ for } t > c, \tag{7}$$

since after the traffic jam is over, v agrees with $\tilde{v}$ and h depends solely on the future behavior of v. (See Figure 7(d).)

Next, assume that b is large enough so that there exists a time a such that $a + \tilde{h}(a) = \tilde{b}$. (If not, then we can replace a by an arbitrary negative number, meaning that (8) becomes vacuous and the domain in (9) becomes $0 \le t < b$.) Then immediately, we have

$$h(t) = \tilde{h}(t) \text{ for } t \le a. \tag{8}$$

(See Figure 7(d).) Moreover, by the previous section's results, $\tilde{h}'(a) = \infty$ and so the derivative of h at a, at least from the left, is also ∞. But the value of h must jump at a, since any delay in departure after time a causes Ivan's arrival to fast-forward to some time after c; more precisely, there is a jump discontinuity at which the value of h jumps by exactly $c - b$ hours, which is the duration of the traffic jam. (Ivan is thrilled almost to the point of tears to see this same jump discontinuity from s^- appearing so clearly in h.) In fact, on the entire interval (a, b), necessarily h behaves exactly as $\tilde{h}$ does, except shifted upwards by $c - b$ hours, since $\tilde{h}$ does not "see" the traffic jam ahead. Therefore, we have

$$h(t) = \tilde{h}(t) + (c - b) \text{ for } a < t < b. \tag{9}$$

(See Figure 7(e).) Finally, since $v(t) = 0$ for every $t \in [b, c]$, our results from the previous section tell us that $h' = -1$ on $[b, c]$. We know the value of $h(b)$ must be $\tilde{h}(\tilde{b}) + (c - b)$, and so we have determined h on all of $[b, c]$, and its graph is just a line segment with slope -1:

$$\begin{aligned} h(t) &= \tilde{h}(\tilde{b}) + (c - b) - (t - b) \\ &= \tilde{h}(\tilde{b}) + c - t \qquad\qquad \text{for } b \le t \le c. \end{aligned} \tag{10}$$

(See Figure 7(f).) Combining results (7) through (10), Ivan and Olga have shown that h is completely determined by $\tilde{h}$:

$$h(t) = \begin{cases} \tilde{h}(t), & 0 \le t \le a, \\ \tilde{h}(t) + c - b, & a < t < b, \\ \tilde{h}(\tilde{b}) + c - t, & b \le t \le c, \\ \tilde{h}(t), & c < t \le \ell. \end{cases} \tag{11}$$

Hence, to obtain the graph of h from that of $\tilde{h}$:

1. Take the portion of the graph of $\tilde{h}$ on the interval (a, b) and shift it upwards by $c - b$ hours.
2. Fill in the interval $[b, c]$ with the unique line segment that makes h continuous. (See Figure 7(g) for the final result.)

In this way, the case of a traffic jam actually reduces to the case in the previous section, which Ivan and Olga understand completely. Finally, although this example involved only one traffic jam, this algorithm can be iterated for each traffic jam in a given speed function v, proceeding left to right.

Conclusion

We leave Olga and Ivan to finish one of many more breakfasts together, and summarize our results. Given a time-dependent traffic speed function, along with bounds on departure and arrival time, we have determined a method using software (Example 1) to find the departure time which minimizes the time spent on the road. Along the way, we discovered a nice visual characterization of the candidates for this departure time (Theorem 1). Because traffic jams—i.e., time intervals on which the traffic speed is zero—truly do complicate our method, we have presented two workarounds in the final section: one elegant tactic employs a generalized inverse function (6), but seems

difficult to program efficiently, whereas another strategy gives an algorithm (11) to reduce the problem to an easier case, by deleting the traffic jams and compensating accordingly.

Some further interesting questions come to mind, perhaps for Ivan and Olga to attack together on a rainy day. What if the speed function v is allowed to take negative values? Of course, traffic never moves backward in reality, but it would certainly add a twist to the theoretical solution to consider negative velocity. Or what if we allow the commute distance D to vary instead of keeping it fixed? Again, although we would lose all veneer of real-world application (not much of a sacrifice in this problem), it could be worthwhile to explore how the optimal departure time t_0 depends on the value of D; for instance, very small D would cause t_0 to occur whenever v is at its maximum, while very large D would force t_0 to occur near 0. This little problem proved to be more involved than it first seemed, and it can certainly be explored further still.

Acknowledgment. The author would like to thank the referee for the many helpful suggestions in improving the presentation of the problem.

Summary. We pose the following optimization problem: given a fixed distance to be traveled in one dimension and a known time-dependent velocity function, what departure time minimizes the time spent traveling? We determine a method to solve this problem with computer assistance, and along the way we discover a nice geometric characterization of the candidates for the solution.

References

[1] Embrechts, P., Hofert, M. (2013). A note on generalized inverses. *Math. Methods Oper. Res.* 77: 423–432.
[2] Gkioulekas, E. (2014). Generalized local test for local extrema in single-variable functions. *Int. J. Math. Ed. Sci. Tech.* 45(1): 118–130.