

New Perspectives on Flexibility in Simple Temporal Planning

Amy Huang,* Liam Lloyd,* Mohamed Omar, James C. Boerkoel Jr.

Human Experience & Agent Teamwork Lab (heatlab.org)

Harvey Mudd College

Claremont, California 91711

{ahuang, wlloyd, omar, boerkoel} @hmc.edu

Abstract

Flexibility is generally agreed to be a desirable feature of a Simple Temporal Network (STN). However, exactly what flexibility attempts to measure has varied, making it difficult to objectively evaluate flexibility metrics. Further, past metrics tend to lose information or exhibit other undesirable properties when aggregating the flexibility measures of individual events across an entire STN. Our work is driven by the realization that the solution space of an STN is a convex polyhedron whose geometric properties convey useful information about the STN. These geometric inspirations lead to measures of an STN solution space and also motivate a set of desiderata for general flexibility metrics. We also put forth two new geometrically-inspired flexibility metrics that have some theoretical advantages over existing metrics. Finally, we comprehensively evaluate both new and existing flexibility metrics against our proposed desiderata.

Introduction

Due to communication lags of up to 24 minutes, it is desirable for Mars rovers to operate autonomously with as little communication from Earth as possible. Thus scientists would prefer to send a Mars rover a set of scheduling possibilities that maximizes *flexibility*—the room within a schedule for the rover to reschedule without violating the overall time constraints in the case of unexpected opportunities or disruptions. Additionally, given this set of constraints, the rover can decide on times to act that maximize the flexibility of its remaining schedule.

Current metrics for flexibility consider the intervals of events independently of one another, either by ignoring dependencies altogether, or by redefining the problem so that each event is assigned an independent time interval. These approaches lose important information about the original problem when calculating flexibility, particularly when the flexibility of a network of constraints is determined by summing individual event flexibilities. Rather, how flexibility is distributed throughout a network (i.e., where it is needed most) plays a large role in maximizing an agent’s autonomy to schedule or respond to unexpected events in practice.

*Primary authors listed alphabetically but contributed equally. Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we observe that the solution spaces of STNs are convex polyhedra and analyze the geometric properties of these polyhedra in order to better understand the nature of flexibility. Using these insights, we define a set of desiderata for flexibility and introduce two new flexibility metrics. The first of these is the ratio of volume to surface area, which approximates the ratio of the number of solutions to the number of possible schedule failure points. The second of these, which is inspired by finding the radius of the largest inscribed sphere in an STN’s solution space, can be computed in low-order polynomial time. We evaluate both new and pre-existing metrics against our desiderata, and see that our new metrics satisfy more of them.

Background

Simple Temporal Networks

A *simple temporal network* (STN) is a pair $S = \langle T, C \rangle$ where $t_0, t_1, \dots, t_n \in T$ represent distinct temporal events, and each constraint in C is a binary constraint on T of the form $t_j - t_i \leq c_{ij}$, for some real number $c_{ij} \in \mathbb{R}$. (Dechter, Meiri, and Pearl 1991). The zero timepoint, t_0 , is defined to be fixed at time 0, and all other events occur relative to t_0 . When both $t_j - t_i \leq c_{ji}$ and $t_i - t_j \leq c_{ij}$ are specified, we use the notation $t_j - t_i \in [-c_{ji}, c_{ij}]$.

An STN is commonly represented as a directed constraint graph, where each event in T is represented by a vertex, and each constraint C_{ij} is represented by a directed edge from t_i to t_j with label c_{ij} . Two vertices that do not share a constraint are connected by an edge with label ∞ . A *schedule* is an assignment of values to each event such that all the constraints are satisfied. An STN is *consistent* if it has at least one schedule.

The *minimal form* of an STN is the one with the tightest set of constraints. In other words, in the minimal constraint graph of an STN $S = \langle T, C \rangle$, the edge between two events t_i and t_j is the shortest path between t_i and t_j in the original constraint graph. The minimal form can be computed using a shortest path algorithm, such as Floyd-Warshall. Since the minimal form corresponds to a fully connected constraint graph, it can also be represented as a distance matrix D , where $D[i, j]$ is the shortest path from t_i to t_j in the constraint graph.

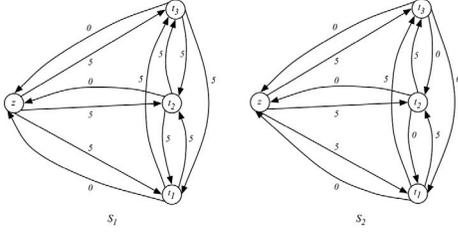


Figure 1: S_1 depicts an STN where t_1, t_2, t_3 can occur concurrently on the interval $[0, 5]$, while S_2 sequentializes these events (Wilson et al. 2014).

Existing Flexibility Metrics

A flexibility metric is defined as a map from STNs to the real numbers, capturing some notion of how much room there is to maneuver within the constraints of a given STN. STNs are traditionally used to guide scheduling decisions during execution, rather than prescriptively dictate a single, fully-determined schedule of events. Thus, STNs with more flexibility are considered desirable because they both help agents avoid failure in the event that something does not go as anticipated and also provide agents more autonomy to proactively respond to their environment. In this spirit, throughout this paper we will discuss the flexibility of temporal networks of constraints rather than of a specific schedule. However, flexibility metrics can be useful in guiding the specific scheduling decisions of agents, which motivates future research directions. Next, we present three existing flexibility metrics and discuss their relative strengths.

The Naïve Metric Perhaps the most intuitive flexibility metric simply sums the ranges of the constraint intervals about each event t , in the minimal form of an STN. We denote this by $flex_N(S)$, and so for an STN S ,

$$flex_N(S) = \sum_{t \in T} (lst(t) - est(t)),$$

where $est(t)$ and $lst(t)$ are the earliest and latest possible starting times of t . Wilson et al. (2014) discuss the limitations of this metric in detail. We will borrow an example from them, as it will be useful for our later purposes as well.

Example 1. We define two STNs S_1 and S_2 , whose distance graphs are shown in Figure 1, such that $S_1 = \langle T_1, C_1 \rangle$ and $S_2 = \langle T_2, C_2 \rangle$.

$$\begin{aligned} T_1 = T_2 &= \{t_0, t_1, t_2, t_3\}, \\ C_1 &= \{0 \leq t_i - t_0 \leq 5 \mid i = 1, 2, 3\} \cup \\ &\quad \{-5 \leq t_i - t_j \leq 5 \mid 1 \leq i \leq j \leq 3\}, \\ C_2 &= \{0 \leq t_i - t_0 \leq 5 \mid i = 1, 2, 3\} \cup \\ &\quad \{0 \leq t_i - t_j \leq 5 \mid 1 \leq j \leq i \leq 3\}. \end{aligned}$$

The three events in T_2 must occur sequentially, while the events in T_1 can occur independently of one another.

As Wilson et al. (2014) illustrate, $flex_N$ produces the counter-intuitive result that both these STNs have the same

flexibility, since $flex_N(S_1) = 15 = flex_N(S_2)$. We would expect S_1 to be more flexible than S_2 , since every schedule that satisfies S_2 also satisfies S_1 , but S_1 admits many schedules that S_2 does not—such as the schedule that assigns $t_1 = 4, t_2 = 3$, and $t_3 = 2$. We define two classes of STNs that generalize those introduced in Example 1, which we will refer to throughout this paper.

Definition 1. Let \mathcal{C} (concurrent) and \mathcal{S} (sequential) be two classes of STNs, which are generalizations of the STNs S_1 and S_2 . Members of these classes will be denoted \mathcal{C}_n and \mathcal{S}_n , where n is the number of events in the STN. The events in \mathcal{C}_n can happen at any time within an interval $[a, b]$, while the events in \mathcal{S}_n must occur sequentially on an interval $[a, b]$.

The Hunsberger Metric A second metric that avoids some of the limitations of the naïve metric is the Hunsberger metric. Hunsberger’s metric (Hunsberger 2002), which originally measured rigidity but can be adapted to measure flexibility (Wilson et al. 2014), adds to the naïve metric the tightest constraints on the difference between each pair of events. Thus if D_S is the distance matrix representing the minimal form of an STN S , this metric $flex_H(S)$ is given by

$$flex_H(S) = flex_N(S) + \sum_{i=1}^n \sum_{j>i}^n (D_S[i, j] + D_S[j, i]).$$

Since this takes into account constraints between pairs of edges it produces different results for S_1 and S_2 . $flex_H(S_1) = 45$, while $flex_H(S_2) = 30$, so S_1 is more flexible, as expected.

However, Wilson et al. (2014) point out a limitation of this metric. For sequential STNs \mathcal{S}_k (see Definition 1),

$$flex_H(\mathcal{S}_k) = 5 \times \frac{k^2 - k}{2}.$$

From this we can see that as k goes to infinity so does $flex_H$. This is counter-intuitive; we would expect that cramming more sequential events into a fixed interval would decrease flexibility and that this limit would approach 0 rather than infinity.

The Wilson Metric In an effort to capture dependencies between events, Wilson et al. (2014) devised a new flexibility metric, which we will label $flex_W$, that computes flexibility using the interval schedules of an STN. An interval schedule is a valid assignment of an independent time interval to each event, such that scheduling an event inside its interval will not further restrict when any of the other events can be scheduled. More precisely, it is a decoupling of an STN into a set of intervals $[t_i^-, t_i^+]$ for each event t_i , such that for all $t_i, t_j \in T$ and all $v_i \in [t_i^-, t_i^+]$ and all $v_j \in [t_j^-, t_j^+]$, the assignment σ that maps $\sigma(t_i) = v_i$ and $\sigma(t_j) = v_j$ is a valid schedule for the STN. The Wilson metric is the sum $\sum_{v_i \in T} (t_i^+ - t_i^-)$ of the interval sizes of the interval schedule that maximizes this sum.

Thus, $flex_W(S_2) = 5$, since the maximal sum of non-overlapping intervals about each of the events in S_2 must be the size of the total interval on which all events must occur.

This result is independent of the number of events in S_2 , so as the number k of events approaches infinity flexibility remains 5. While this is much more intuitive than flexibility going to infinity as it does in $flex_H$, it still fails to capture the reality that as we add more events there will be less room for error. We see another peculiarity of this metric by comparing S_2 to a new STN $S_3 = \langle T_3, C_3 \rangle$:

$$\begin{aligned} T_3 &= \{t_0, t_1, t_2, t_3\}, \\ C_3 &= \{0 \leq t_1 - t_0 \leq 3, 3 \leq t_2 - t_0 \leq 4, \\ &\quad 4 \leq t_3 - t_0 \leq 5\}. \end{aligned}$$

Observe that $flex_W(S_2) = 5 = flex_W(S_3)$. Yet any schedule satisfying S_3 also satisfies S_2 , and S_2 admits additional schedules, such as one which assigns $t_1 = 1, t_2 = 2, t_3 = 3$. This seems counter-intuitive; we would expect that if the solutions to an STN are a proper subset of the solutions to another, the second should be more flexible. Hence all of these existing metrics exhibit potentially undesirable or un-intuitive properties. Next, we consider geometric interpretations of STNs as a way to better interpret existing flexibility metrics while also inspiring new ones.

Geometric Interpretations of STNs

Visualizing an STN as a geometric object can give us new insights into its properties. The set of constraints that defines an STN naturally lends itself to representation as a polyhedron. Our hope is that this in turn will lead to both more concrete definitions of STN flexibility and more intuitive flexibility metrics.

STNs as Convex Polyhedra

A convex polyhedron can be thought of as a generalization of a convex polygon to general dimensions, but is not necessarily bounded. It is defined as a region in n dimensions bounded by a finite number of hyperplanes. It is usually denoted either by a set of inequalities or by a list of its vertices and extreme rays. Note that an STN is represented by an unbounded polyhedron if it has one or more event that is unconstrained; otherwise, it has a bounded polyhedron.

An STN with n events can be represented as a polyhedron in n -dimensional space, where each axis represents the value of an event. Then, each constraint $t_j - t_i \leq c_{ij}$ defines a half-space. The intersection of these half-spaces is bounded by the finite list of hyperplanes $t_j - t_i = c_{ij}$, which exactly defines the space of possible solutions, forming a convex polyhedron. Figure 2 illustrates \mathcal{S}_2 (see Definition 1) as a polyhedron where constraints are represented by dashed lines and the solution space is highlighted.

We can also look at the polyhedra of interval schedules of STNs. Because each event in an interval schedule has a domain independent of other events, its polyhedron is a hyperrectangle with edges parallel to the axes. Therefore its volume is the product of the lengths of each interval. Wilson et. al. use their flexibility metric to find an interval schedule for an STN that maximizes the sum of the size of intervals. If we were to visualize this interval schedule as a geometric object, it would form a hyperrectangle, inscribed within the polyhedron of the original STN. The polyhedron view

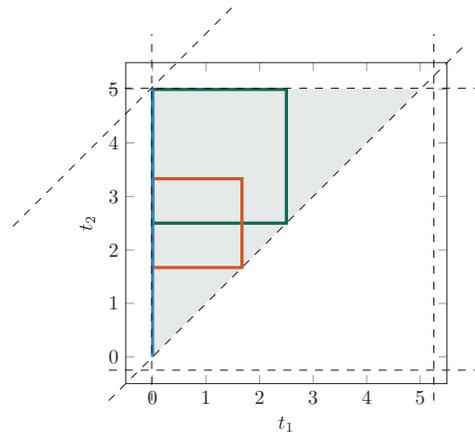


Figure 2: A graphical representation of \mathcal{S}_2 . The dashed lines are its constraints and the shaded triangle is its solution space. The squares and line are the solution spaces of interval schedules. Notice that the square schedule ($t_1 \in [0, 2.5], t_2 \in [2.5, 5]$) and the line schedule ($t_1 \in [0, 0], t_2 \in [0, 5]$) both sum to 5, even though the square has a larger 2-dimensional volume (i.e., area).

suggests maximizing the volume of the polyhedron for an interval schedule as an alternative to maximizing the sum of its intervals.

Measuring the Size of an STN Solution Space

When considering STNs geometrically, any valid schedule for an STN represents a point in n -dimensional space that falls inside the polyhedron defined by the STN. Therefore, the relative volume of a polyhedron captures the size of an STN's solution space, where relative volume is the volume the polyhedron takes in the affine space spanned by its vertices. This affine space can be understood as the largest dimension in which the polyhedron is full-dimensional, so if we had an STN with three events and its solution space happened to be a 2-dimensional polyhedron, the affine space would be the two dimensional space spanned by that polyhedron. To see why we want relative volume rather than volume, consider an STN specification that contains a fixed event (or fixed edge). The corresponding polyhedron will subsequently be flat in one direction, and consequently have 0 volume. Thus, when measuring the solution space, what is actually desired is relative volume. It is less clear (i.e., application and context dependent) whether to consider volume or relative volume when considering volumes for the purpose of flexibility metrics, as we will discuss further later on. Finally, note that each integer lattice point inside a polyhedron is a valid schedule assignment. Hence, the number of integer lattice points in a polyhedron describes the total number of valid integer schedules of an STN. However, the number of integer lattice points is well approximated asymptotically by the relative volume of the polyhedron.

It is of interest that the relative volume of an STN's polyhedron quantifies the number of schedules satisfying it, since

previous work has proposed this quantity as a definition of flexibility (Policella et al. 2004). However, computing the volume of a general convex polyhedron is known to be #P-hard (Dyer and Frieze 1988). Despite the fact that STNs have special structure (for instance, all coefficients can only be 1 or -1), one or more of the parallel boundaries may be superfluous due to other constraints (see Figure 2), and thus we cannot exploit fast methods for computing volumes of the special class of polytopes called parallelotopes (i.e., n -dimensional parallelepipeds) (Gover and Krikorian 2010). Thus, for STNs of more than about ten dimensions, finding the exact volume is infeasible (Ge and Ma 2015).

The fastest known approximation for volume uses a hit-and-run Monte Carlo method and runs in $O^*(mn^3)$, where m is the number of constraints and n is the number of variables (Ge and Ma 2015). Computing the number of lattice points contained in a polytope is also somewhat more feasible than computing exact volume if the polytope has rational-coordinate vertices. This computation can be done in polynomial time if dimension is held constant, but is generally exponential in the number of variables and becomes intractable above 30 dimensions (De Loera et al. 2004).

While measuring the number of solutions to an STN is useful in some application domains, volume does not capture an STN’s flexibility as well as it initially seems like it might. This is because it does not take into account how those solutions are distributed. To illustrate the importance of the distribution of an STN’s solutions, consider two STNs with two events whose constraints define rectangles. Call the first S_1 , and let its constraints define a 10×10 rectangle, and call the second S_2 and let its constraints define a 50×2 rectangle. Suppose then that one of the events in these STNs is delayed by 1 time unit, decreasing flexibility along that axis by one. For S_1 , the volume of the solution set after this delay will be 90, no matter which event is delayed, while for S_2 it will be either 98 or 50 depending on which axis is delayed. Supposing that either event is equally likely to be delayed we can average these to get an expected volume of 74. Thus we see that S_1 is less damaged by a delay than S_2 because its *shape* is more advantageous, and should therefore be considered more flexible. Thus while volume both measures a useful property (size of the solution space) and will also motivate our new flexibility metrics, volume as a flexibility metric is not without flaws.

STN Flexibility Desiderata

While flexibility has been discussed with some regularity as a desirable quality in temporal planning, (Do and Kambhampati (2003), Davenport, Gefflot, and Beck (2014), Say, Cire, and Beck (2016)), it has not yet been concretely characterized. Flexibility has been called an “aggregate measure of slack” (Lund et al. 2017), or defined as the number of possible schedules in an STN (Policella et al. 2004), and it is often only defined in terms of the metric a particular author uses to calculate it (Wilson et al. (2014), Hunsberger (2002)). However, there is currently no clear consensus as to which, if any, of these metrics is the best one. Different applications of flexibility suggest optimizing for different qualities, further complicating efforts to put forth a single, concise definition.

Indeed, this has thwarted our own attempts to define flexibility rigorously, particularly in a way that is consistent with past work. Thus, instead, we propose a set of key properties that characterize critical features of a flexibility metric.

Since flexibility is generally agreed to be at least partly a measure of the ability of an STN’s satisfying schedules to resist scheduling perturbations, flexibility metrics are forced to make assumptions about the likelihood and nature of such perturbations. It is not clear which assumptions are ideal, since STNs contain no information about such uncertainty. For the purposes of this paper, we assume that all events in an STN are equally vulnerable to perturbations of equal magnitude. We characterize the following key properties as desirable under this assumption, but recognize that applications where different assumptions hold may characterize and optimize for these properties differently.

Wilson et al. (2014) make a start at defining ideal qualities for flexibility metrics by identifying desirable limiting behavior of flexibility metrics on the concurrent (\mathcal{C}) and sequential (\mathcal{S}) classes of STNs introduced in Definition 1. Wilson et. al. point out that a flexibility metric *flex* should have the property:

$$\lim_{n \rightarrow \infty} \frac{flex(\mathcal{C}_n)}{flex(\mathcal{S}_n)} = \infty.$$

This captures that events happening concurrently on a interval are less constrained than events happening sequentially on the same interval, and that this difference increases as the number of events approaches infinity.

We suggest a pair of stronger qualities inspired by this property, which we believe capture further intuition about concurrent and sequential STNs. We follow these with two geometrically inspired desiderata.

Desideratum 1 (Simplicity). *If S is an STN, and we add an event to create a new STN S' , even if the new event is independent from all existing events, then*

$$flex(S') \leq flex(S).$$

Additionally, a metric exhibits strong simplicity when $flex(S') = flex(S)$ if and only if the new event is both independent of the rest of the network and has a domain of ∞ . Otherwise, we say it exhibits weak simplicity.

The flexibility of STN should not increase just because the number of events increases. Essentially, the more events there are, the higher the chance that one of them will lead to a scheduling failure, arguing for *simple* networks. Adding an unconstrained event that is independent of the existing network and cannot fail should not affect flexibility, and certainly should not cause the flexibility to go to infinity.

To illustrate why this is desirable, consider an example where five rovers must each observe an event for 10 minutes between 3:00 and 3:30, and their observation times must overlap. Then, we add another rover that needs to communicate independently with a satellite between 5:00 and 5:30. Intuitively, this new schedule should be less flexible than the original schedule, as the constraints on the original two rovers remain the same while a new constrained event is added. With more constrained events, there are more ways for a schedule to fail, rendering the schedule less flexible.

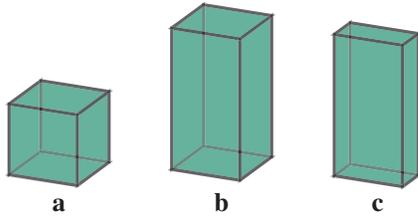


Figure 3: If a flexibility metric measures *sphericity*, it will consider **a** more flexible than **c**. If it measures *containment*, it will consider **b** more flexible than **a** and **c**.

Desideratum 2 (Density). For STNs $\mathcal{S}_n \in \mathcal{S}$ composed of a set of n sequential events that occur within a fixed, finite interval $[a, b]$, a flexibility metric $flex$ should have the property

$$\lim_{n \rightarrow \infty} flex(\mathcal{S}_n) = 0.$$

Density reflects the idea that as more sequential events are crammed into a single interval, there is less space for each of them. For infinite events, each (or all but finitely many, since infinite events could occur simultaneously to give nonzero size intervals to finitely many events) will have only an exact timepoint when they can happen to satisfy the STN.

Here, we'll look at an example where a rover must collect five samples sequentially between 5:00 and 6:00. If we increase the number of samples to 50, the rover will have a smaller interval of time in which to collect each sample, and there will be less room for error. Thus, the more sequential events there are, the smaller the flexibility.

Desideratum 3 (Sphericity). A flexibility metric is spherical if and only if for two STNs with the same number of events and whose solution spaces are the same size (in other words, their polyhedra have the same dimension and volume), the STN with the larger inscribed sphere is considered to be more flexible.

The first of our geometrically inspired desiderata comes from the intuition that not all flexibility is equal. That is, it seems more meaningful to give an extra unit of flexibility to an event which had 0 than to an event which had 100. Drawing on the geometric interpretation of an STN, we call this property *sphericity*. Thus, as illustrated in Figure 3, a spherical metric would assign more flexibility to the STN represented by polyhedron **a** than to **c**, despite both containing the same volume/number of schedules.

To illustrate sphericity, consider an example where one rover must observe an event between 3:00 and 3:01, while another rover must observe an event between 3:00 and 9:00. A schedule with a polyhedron of equivalent volume would be one in which both rovers observe events between 3:00 and 4:00. In the second example, neither rover is as likely to miss their event as is the rover in the first example with a one-minute time interval. Therefore, the second example should be considered more flexible.

Desideratum 4 (Containment). Consider two arbitrary STNs S and S' that have the same number of events and

where the polyhedron of S is a proper subset of the polyhedron of S' . A flexibility metric captures strong containment if S' is considered more flexible than S , and it captures weak containment if S' is at least as flexible as S .

Containment captures the idea that an STN with more solutions in its solution space than another should have a greater flexibility, given that the two STNs have the same number of events. This desideratum has no bearing on two STNs that have different numbers of events; it is capturing a notion of size, and volumes of different dimensional polyhedra are not meaningfully comparable. As illustrated in Figure 3, any solution of STN **a** is also a solution of STN **b**, and thus **b** would be assigned more flexibility.

To illustrate this, we return to the first example, where five rovers must each observe an event for 10 minutes between 3:00 and 4:00, and their observation times must overlap. If we instead the observation must occur between 2:00 and 5:00, this definition suggests that flexibility should increase, as any solution to the first problem remains a solution to the second problem, and the rovers in the second problem have more options.

Geometrically Inspired Flexibility Metrics

Having determined a set of desirable properties for a flexibility metric that characterize the distribution of flexibility, it seems natural to seek metrics that satisfy as many as possible. We propose two new metrics, inspired by the geometric interpretation of STNs.

It should be noted that both of these metrics involve making a decision about whether to use volume or relative volume in cases where an STN's polyhedron is not full-dimensional. If an STN's polyhedron is not full-dimensional, this means some event is fixed to a specific value, or two events are fixed relative to each other. If we can ensure that events fixed to a time happen at that time and that events fixed relative to each other happen the correct distance apart, it makes sense to use relative volumes. If not, then it makes sense to use non-relative volume, which measures 0 for STNs with fixed (pairs of) events. We, however, will use the term "volume" throughout for simplicity.

Ratio of Volume to Surface Area

For an STN S , where $V(S)$ is the (relative) volume of the STN's polyhedron and $SA(S)$ is its (relative) surface area, we define the volume to surface area metric of S to be

$$flex(S) = \frac{V(S)}{SA(S)}$$

This ratio captures the proportion of valid schedules that are close to (roughly within a unit of) the polyhedron's boundary, and thus are most at risk of becoming invalidated when perturbed. We will show in the next section that volume to surface area ratio satisfies two of our desiderata, strong simplicity and density. Interestingly, it can be shown that volume alone also satisfies two desiderata, density and strong containment. However, we choose to propose the ratio of volume to surface area as a metric rather than volume alone because it captures the distribution of solutions in addition to the size of the solution space.

The Sphere Metric

In considering geometrically inspired flexibility metrics, the sphere inscribed in the polyhedron described by an STN is a natural starting point, since its center maximizes the minimum distance of any point within the polyhedron to its nearest boundary. An STN whose polyhedron has a large inscribed sphere will have many schedules that are not close to any facet; that is, it has many schedules that can withstand some degree of perturbation. Further, spheres exhibit desirable behavior as their dimensionality increases. The volume of a sphere with fixed radius approaches 0 as its dimension approaches infinity. Further, the n^{th} root of the volume monotonically decreases as dimension increases, which we have identified as desirable. Thus we define the *sphere flexibility* of an STN as the n^{th} root of the volume of the largest inscribed sphere of its polyhedron.

Though sphere flexibility, like other metrics, measures the flexibility of an STN rather than a schedule, it does provide some intuitive guidance toward the “most flexible” schedule. This would be the schedule that occupies the center of sphere inscribed in the STN’s polyhedron, since this is the schedule that could withstand the greatest perturbation in any direction without violating any constraints.

Computing the Radius of the Inscribed Sphere The radius of the largest inscribed sphere corresponding to an STN $S = \langle T, C \rangle$ can be computed using a simple linear program (Murty 2009):

$$\begin{aligned} &\text{maximize: } r \\ &\text{subject to: } t_j - t_i + \sqrt{2}r \leq c_{ij} \quad \forall t \in T; i, j \neq 0 \\ &\quad t_j - t_i + r \leq c_{ij} \quad \forall t \in T \text{ if } i = 0 \text{ or } j = 0 \end{aligned}$$

Note that the resulting values of t_1, t_2, \dots, t_n define the Chebyshev center of the polyhedron, a (not necessarily unique) point that is at least a distance of r away from every boundary.

Once the radius of the inscribed sphere is determined, which is achievable in the runtime of the chosen linear program solver, the (n -dimensional) volume can be computed with the equation $V_n(r) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)} r^n$. Taking the n^{th} root of the result gives us the sphere flexibility. Since both the gamma function and the n^{th} root can be computed in low order polynomial time, the whole process remains computationally efficient.

Evaluation

We will here evaluate our new flexibility metrics as well as pre-existing metrics against each of our proposed desiderata. Figure 4 summarizes our comparisons. While our evaluation is analytically focused, we also evaluated metrics empirically to see if their expected performance behaved qualitatively differently than our theoretical evaluation would suggest. However, we omitted empirical results in the interest of space, since they led to no evidence that would cause us to deviate from our theoretical findings.

The Naïve Metric

The naïve metric only satisfies *weak containment*.

D1 If we add an independent event t to an STN S , the naïve flexibility will increase by exactly t ’s domain. This is because t is independent, so the domains of the other events will not change. Thus the naïve metric increases as independent events are added to an STN, violating *simplicity*.

D2 Since in a sequential STN \mathcal{S}_n all events can happen between a and b , each event has a range of possible values of size $b - a$. Hence the naïve flexibility of \mathcal{S}_n is $n(b - a)$, which approaches infinity as $n \rightarrow \infty$. Thus the naïve metric does not exhibit *density*.

D3 Since the naïve metric merely sums the range of possible values for each event, it will consider an STN where one event has a range of 100 and the other has a range of 0 more flexible than an STN where both events have a range of 49. Thus the naïve metric does not satisfy *sphericity*.

D4 As we saw in the background section, the naïve metric assigns the same flexibility to \mathcal{S}_k and \mathcal{C}_k , even though \mathcal{S}_k ’s solution set is a proper subset of \mathcal{C}_k ’s. Thus, the naïve metric does not satisfy *strong containment*. However, if the solution set of an STN S' contains the solution set of the STN S , $flex_N(S') \geq flex_N(S)$. This is because since every solution to S is a solution to S' , all the events in S will have at least their same range of possible values in S' . Thus the naïve metric does exhibit *weak containment*.

The Hunsberger Metric

The Hunsberger metric satisfies only *containment*, though it is notably the only metric that exhibits *strong containment*.

D1 The corresponding example for the naïve metric applies here, so this metric does not satisfy *simplicity*.

D2 As Wilson et al. (2014) shows, for a sequential STN \mathcal{S}_n , $flex_H = (b - a) \frac{n^2 - n}{2}$. As $n \rightarrow \infty$, $flex_H(\mathcal{S}_n) \rightarrow \infty$. Hence the Hunsberger metric violates *density*.

D3 For STNs with exclusively independent events, the Hunsberger metric acts like the naïve metric. Thus the counterexample showing that the naïve metric does not exhibit *sphericity* works for the Hunsberger metric as well.

D4 From a geometric perspective, the Hunsberger metric sums the Manhattan distance between all pairs of parallel boundaries of an STN’s polyhedron, as given by the minimal form constraints. It should be noted that some of these boundaries may not be facets of the polyhedron; some may simply pass through vertices. However, every facet of the polyhedron must be a boundary. Now, if the STN S ’s polyhedron is a proper subset of the STN S' ’s polyhedron, there must be space between some facet of S and the facets of S' . Hence, there is room to extend at least one of the distances between boundaries. Since S' and S are of the same dimension they have the same number of boundary pairs. The distances between those pairs will be at least as long in S' as they are in S , and at least one will be longer. Thus, the Hunsberger metric must be larger for S' , and the Hunsberger metric satisfies *strong containment*. Interestingly, it is the only metric we know of that does so.

	Naïve	Huns.	Wilson	V/SA	Sphere
Simplicity	×	×	×	✓	✓*
Density	×	×	×	✓	✓
Sphericity	×	×	×	×	✓
Containment	✓*	✓	✓*	×	✓*
Poly-time	✓	✓	✓	×	✓

*Indicates that the metric satisfies a weak version of this property.

Figure 4: A summary of flexibility metrics properties.

The Wilson Metric

The Wilson metric exhibits only *weak containment*.

D1 As in the previous two metrics, adding an independent event t to an STN will increase flexibility by the domain of t . Therefore this metric does not satisfy *simplicity*.

D2 Wilson et al. (2014) prove that for a sequential STN where all events must occur in a time interval of size k , the flexibility is k , for any number of events. Thus it does not approach zero as the number of events approaches infinity and so does not exhibit *density*. We will note, however, that if we consider the product of the interval sizes under the Wilson decoupling to be the flexibility rather than the sum we would satisfy density, since the Wilson decoupling defines a box inscribed in the STN’s polyhedron, and since the polyhedra for sequential STNs approach zero volume as $n \rightarrow \infty$ (we will see this in the evaluation of volume to surface area).

D3 The counterexample showing that the naïve metric does not satisfy *sphericity* applies to the Wilson metric as well, as long as the two events are independent.

D4 As we saw in the Background section, there exist pairs of STNs of equal Wilson flexibility where the solution set of one is a proper subset of the solution set of the other. Thus the Wilson metric does not recognize strong containment. However, if the solution set of an STN S' contains the solution set of an STN S , $flex_W(S') \geq flex_W(S)$, because the interval schedule contained in S from which $flex_W(S)$ is calculated will also be contained in S' . Thus the Wilson metric has *weak containment*.

Volume to Surface Area Ratio

The ratio of volume to surface area (V/SA) exhibits *strong simplicity* and *density*, but not *sphericity* or *containment*.

D1 Let S be an STN with n events. If you construct a new STN S' by adding an event t_{n+1} that is independent from all other events and whose domain has size $\ell > 0$, then the polyhedron defined by S' will be a prism with the polyhedron defined by S as its base. Hence, $V(S') = \ell V(S)$ and $A(S') = 2V(S) + \ell A(S)$. Thus,

$$flex(S') = \frac{V(S')}{A(S')} = \frac{\ell V(S)}{2V(S) + \ell A(S)} = \frac{flex(S)}{\frac{2}{\ell} flex(S) + 1}$$

Since ℓ is positive, $flex(S') \leq flex(S)$. Additionally,

$$\lim_{\ell \rightarrow \infty} flex(S') = flex(S),$$

which confirms that adding an independent unbounded event has no effect on flexibility. Thus, V/SA exhibits *strong simplicity*.

D2 Let \mathcal{S}_n be an STN with n events where the events occur sequentially within a time interval $[a, b]$. \mathcal{S}_n is defined by the inequalities

$$\begin{aligned} t_1 - t_0 &\geq a \\ t_2 - t_1 &\geq 0 \\ &\vdots \\ t_n - t_{n-1} &\geq 0 \\ t_n - t_0 &\leq b \end{aligned}$$

Then the volume to surface area ratio of the polyhedron defined by \mathcal{S}_n in n -dimensional space approaches 0 as n approaches ∞ :

$$\lim_{n \rightarrow \infty} flex(\mathcal{S}_n) = 0$$

as required for *density*. Due to space considerations, we present a sketch of the proof.

Proof. (sketch) It can be shown that the sequential STN \mathcal{S}_n defines an n -simplex. The volume of an n -simplex in n dimensions with vertex set v_0, v_1, \dots, v_n where v_0 is at the origin is

$$\frac{\det(v_1, v_2, \dots, v_n)}{n!}.$$

If we set one of our vertices to be the origin, all other vertices will be of the form $(k, 0, \dots, 0), (k, k, 0, \dots, 0)$, and so on up to (k, \dots, k) . Here, $k = b - a$. Thus the volume of our polyhedron is

$$V(P(\mathcal{S}_n)) = \frac{k^n}{n!}.$$

The facets of a simplex are $n - 1$ dimensional simplices, and using similar reasoning we can establish that the relative surface area of $P(\mathcal{S}_n)$ is

$$A(P(\mathcal{S}_n)) \geq \frac{1}{(n-1)!} (n+1)(k^{n-1}).$$

We have a lower bound rather than equality because some facets of the simplex lie on planes not parallel to any axis, resulting in values larger than k on the diagonals of their matrices. Hence

$$\lim_{n \rightarrow \infty} \frac{V(P(\mathcal{S}_n))}{A(P(\mathcal{S}_n))} \leq \lim_{n \rightarrow \infty} \frac{k}{n^2 + n} = 0,$$

as desired. \square

D3 V/SA does not capture *sphericity*. As an example, consider two STNs $S_1 = \langle T_1, C_1 \rangle$ and $S_2 = \langle T_2, C_2 \rangle$, where

$$\begin{aligned} T_1 &= T_2 = \{t_0, t_1, t_2\}, \\ C_1 &= \{0 \leq t_1 - t_0 \leq 6, 0 \leq t_2 - t_0 \leq 2, -2 \leq t_2 - t_1 \leq 0\}, \\ C_2 &= \{0 \leq t_1 - t_0 \leq 2.5, 0 \leq t_2 - t_0 \leq 0.8\}. \end{aligned}$$

Figure 5 shows the polyhedra of S_1 and S_2 . The V/SA flexibilities of S_1 and S_2 are 0.29 and 0.30, respectively. S_1 has a larger inscribed sphere, with a radius of 0.5, but S_2 has a greater volume to surface area ratio, although it has a smaller radius of 0.4. It is noteworthy, though, that the polyhedron of a given volume with the greatest volume to surface area ratio is a sphere, so this metric does favor “sphere-like” polyhedra in the limit.

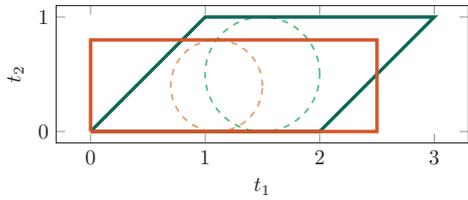


Figure 5: A graphical representation of STNs S_1 (parallelogram) and S_2 (rectangle) demonstrate that the volume to surface area ratio does not track sphericity.

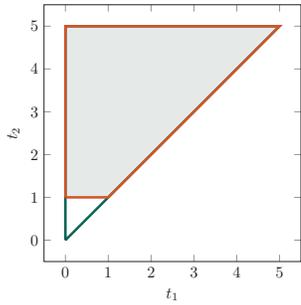


Figure 6: A graphical representation of STNs S_1 (triangle) and S_2 (shaded area) demonstrate that the volume to surface area ratio does not track containment.

D4 V/SA does not exhibit *containment*. Given an STN where three events occur concurrently in the interval $[0, 1]$, extending one event's domain to $[0, 5]$ does indeed increase the volume to surface area ratio. However, consider the following counterexample, depicted in Figure 6. Let $S_1 = \langle T_1, C_1 \rangle$ be defined:

$$T_1 = \{z, t_1, t_2\},$$

$$C_1 = \{0 \leq t_1 - z \leq 5, 0 \leq t_2 - t_1 \leq 5, 0 \leq t_2 - z \leq 5\}$$

Define S_2 by shrinking the domain of $t_1 - z$ to $[1, 5]$. It can be verified that S_2 has a greater volume to surface area ratio than S_1 , even though we decreased the interval of t_1 so that S_2 is contained in S_1 .

This seems counterintuitive, but the volume to surface area ratio captures the likelihood that a randomly selected valid schedule (as opposed a schedule that an agent gets to choose) would remain valid after being perturbed. Therefore, adding in the corner of a triangle is detrimental to the volume to surface area ratio, as any schedule in that corner has very little room for error. Unlike the inscribed sphere, which measures the room for error of the most optimally-placed schedule, the ratio of volume to surface area measures the proportion of schedules that are most vulnerable.

The Sphere Metric

The sphere metric satisfies all of our desiderata at some level, though it only weakly satisfies simplicity and containment.

D1 If an n -dimensional STN S 's polyhedron has an inscribed sphere of radius r and an independent event is added to the STN to get S' , if that event's range of possible values is smaller than r , the radius, will decrease. If instead the new event's range of possible values is greater than or equal to r , the radius of S' 's sphere will also be r . Thus the radius r' of S' is less than or equal to r . Hence

$$flex(S') \leq \sqrt[n+1]{\frac{\pi^{\frac{n+1}{2}}}{\Gamma(\frac{n+1}{2} + 1)}} (r)^{n+1} = \pi r \sqrt[n+1]{\frac{\pi^{\frac{1}{2}}}{\Gamma(\frac{n+1}{2} + 1)}}.$$

Thus we can see that since $\sqrt[n+1]{\pi^{\frac{1}{2}}} < \sqrt[n]{\pi^{\frac{1}{2}}}$ and $\sqrt[n+1]{\Gamma(\frac{n+1}{2} + 1)} > \sqrt[n]{\Gamma(\frac{n+1}{2} + 1)}$, then $flex(S') < flex(S)$ for all $n \geq 1$, satisfying *weak simplicity*. However, if an unbounded event is added to S to make S' , the radius of S' 's inscribed sphere will still be r and S' 's polyhedron will still be $n + 1$ dimensional, so the above argument that $flex(S') < flex(S)$ still holds. Therefore the sphere metric does not satisfy *strong simplicity*.

D2 It was shown earlier that the volume of the polyhedron defined by a sequential STN \mathcal{S}_n has volume

$$V = \frac{k^n}{n!}.$$

We can see then that for the polyhedron $P(\mathcal{S}_n)$,

$$\lim_{n \rightarrow \infty} V = 0.$$

As volume approaches zero, the volume of the inscribed sphere must also approach zero, so the sphere metric satisfies *density*.

D3 Given two STNs of the same dimension, the one with a larger inscribed sphere radius will have an inscribed sphere with larger volume; thus this metric exhibits *sphericity*.

D4 In an STN with n independent events with domain $[0, k]$, increasing the domain of one event will not affect the dimensions of the inscribed sphere, though it will produce a polyhedron containing the original STN's polyhedron. For example, see Figure 3. This metric does, however, exhibit *weak containment*, since if a polyhedron Q contains a polyhedron P the inscribed sphere of Q must be at least as large as the inscribed sphere of P .

Discussion

In this paper, we use the geometric representation of STNs to better understand the concept of flexibility. We have defined a set of desiderata that flexibility metrics should exhibit, which will allow more objective and better understood comparisons between flexibility metrics. We also developed two new flexibility metrics, both of which satisfy more of these desiderata than previous metrics. In particular, the sphere flexibility metric can be computed in low-order polynomial time, and satisfies all of the desiderata, at least weakly.

There remains ample room for further study in the realm of STN geometry. In the future, we would like to search for a flexibility metric that strongly satisfies all our desiderata, possibly by combining features from multiple metrics. It would be particularly interesting to imbue another metric

with the Hunsberger metric’s ability to strongly satisfy containment. We would also like to investigate methods for efficiently approximating volume to surface area ratios, building on existing volume approximation methods.

Further, the geometric interpretations that we present in this paper indicate that not all feasible schedules within an STN are equally flexible to disruptions. For instance, the V/SA metric indicated that flexibility could be improved by pruning schedules that appear in ‘corners’ of STNs (where the volume to surface area is low). Similarly, the sphere metric implicitly argues that the most flexible schedules are those at the center of the largest inscribed sphere. Both of these examples point towards future work that explores using our new metrics to guide agents’ scheduling decisions in a way that improves performance.

Finally, it would also be interesting to investigate which desiderata are the most important for the empirical performance of a flexibility metric for real-world temporal planning applications. Particularly since we have yet to identify a metric that satisfies all of them, this could be useful for deciding which desiderata can be least detrimentally sacrificed. Additionally, we here assumed that all events in an STN were equally vulnerable to perturbations of equal magnitude; it would be interesting to investigate how to modify our desiderata and flexibility metrics for situations where this assumption does not hold, such as Probabilistic STNs (PSTNs). Our work here might also be beneficially applied to finding maximally controllable solutions to STNs with Uncertainty (STNUs) and maximally flexible temporal decouplings.

Acknowledgements

Funding for this work was graciously provided by the National Science Foundation under grant IIS-1651822. Thanks to the anonymous reviewers, HMC faculty, staff and fellow HEATlab members for their support and constructive feedback.

References

Davenport, A.; Gefflot, C.; and Beck, C. 2014. Slack-based techniques for robust schedules. In *Proc. of the 6th European Conference on Planning (ECP-14)*, 43–49.

De Loera, J.; Hemmecke, R.; Tauzer, J.; and Yoshida, R. 2004. Effective Lattice Counting in Rational Convex Polytopes. *Journal of Symbolic Computation* 38:1273–1302.

Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. In *Knowledge Representation*, volume 49, 61–95.

Do, M. B., and Kambhampati, S. 2003. Improving temporal flexibility of position constrained metric temporal plans. In *Proc. of 13th International Conference on Automated Planning and Scheduling (ICAPS-03)*, 42–51.

Dyer, M. E., and Frieze, A. M. 1988. On the complexity of computing the volume of a polyhedron. *SIAM Journal on Computing* 17(5):967–974.

Ge, C., and Ma, F. 2015. A fast and practical method to estimate volumes of convex polytopes. In *International Workshop on Frontiers in Algorithmics*, 52–63.

Gover, E., and Krikorian, N. 2010. Determinants and the volumes of parallelotopes and zonotopes. *Linear Algebra and its Applications* 433(1):28–40.

Hunsberger, L. 2002. Algorithms for a temporal decoupling problem in multi-agent planning. In *Proc. of 18th National Conference on Artificial Intelligence (AAAI-02)*, 468–475.

Lund, K.; Dietrich, S.; Chow, S.; and Boerkoel, J. 2017. Robust execution of probabilistic temporal plans. In *Proc. of 31st AAAI Conference on Artificial Intelligence (AAAI-17)*, 3597–3604.

Murty, K. G. 2009. Ball centers of special polytopes. Technical report.

Policella, N.; Smith, S.; Cesta, A.; and Oddi, A. 2004. Generating robust schedules through temporal flexibility. In *Proc. of 14th International Conference on Automated Planning and Scheduling (ICAPS-04)*, 209–218.

Say, B.; Cire, A. A.; and Beck, J. C. 2016. Mathematical programming models for optimizing partial-order plan flexibility. In *In Proc. of 22nd European Conference on Artificial Intelligence (ECAI-16)*, 1044–1052.

Wilson, M.; Klos, T.; Witteveen, C.; and Huisman, B. 2014. Flexibility and decoupling in simple temporal networks. *Artificial Intelligence* 214:26–44.