

Topological belief space planning for active SLAM with pairwise Gaussian potentials and performance guarantees

Andrej Kitanov  and Vadim Indelman

The International Journal of
Robotics Research
2024, Vol. 43(1) 69–97
© The Author(s) 2023



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649231204898
journals.sagepub.com/home/ijr



Abstract

Determining a globally optimal solution of belief space planning (BSP) in high-dimensional state spaces directly is computationally expensive, as it involves belief propagation and objective function evaluation for each candidate action. However, many problems of interest, such as active SLAM, exhibit structure that can be harnessed to expedite planning. Also, in order to choose an optimal action, an exact value of the objective function is not required as long as the actions can be sorted in the same way. In this paper we leverage these two key aspects and present the topological belief space planning (t -bsp) concept that uses topological signatures to perform this ranking for information-theoretic cost functions, considering only topologies of factor graphs that correspond to future posterior beliefs. In particular, we propose a highly efficient topological signature based on the von Neumann graph entropy that is a function of graph node degrees and supports an incremental update. We analyze it in the context of active pose SLAM and derive error bounds between the proposed topological signature and the original information-theoretic cost function. These bounds are then used to provide performance guarantees for t -bsp with respect to a given solver of the original information-theoretic BSP problem. Realistic and synthetic simulations demonstrate drastic speed-up of the proposed method with respect to the state-of-the-art methods while retaining the ability to select a near-optimal solution. A proof of concept of t -bsp is given in a small-scale real-world active SLAM experiment.

Keywords

Topology, stochastic control, belief space planning, POMDP, information entropy, graph entropy, SLAM

Received 10 November 2019; Revised 1 August 2023; Accepted 28 August 2023

Senior Editor: David Hsu

1. Introduction

Decision-making under uncertainty in partially observable domains is inescapable for autonomous systems operating in the real world. An autonomous agent (robot) needs to make predictions about its own state and state of the world given often partial and uncertain data and decide its own actions accordingly. In principle, it has to perform inference about unknown quantities and evaluate the task-related objective on available actions. These operations can be quite expensive, especially in high-dimensional state spaces and with information-theoretic cost functions, such as differential entropy. Online decision-making under uncertainty in such a setting is essential in numerous problems in robotics, such as single and multi-robot active SLAM, sensor deployment, search and rescue, and autonomous environmental monitoring.

However, many problems of interest, for example, active SLAM, exhibit structure. In such problems, only subsets of variables interact directly with each other. It means there exist conditional independencies between variables and it is possible to simplify the model of their joint probability

distribution. One way to represent certain dependencies or causality among variables of a stochastic system is via a graphical model of their joint probability distribution (Koller and Friedman 2009), for example, a factor graph and a Markov Random Field.

A graphical model approach involves expressing the joint probability density/mass function (*belief*) as a product of local functions of its variables. A structure of this model, described by a topology induced from the connectivity of the variables and factors they belong to, reveals probabilistic and algorithmic features of the model (e.g., independence, learning complexity) often indicating how to efficiently perform marginalization and conditioning operations on that model

Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa, Israel

Corresponding author:

Andrej Kitanov, Department of Aerospace Engineering, Technion - Israel Institute of Technology, Technion City, Haifa 32000, Israel.
Email: akitanov@gmail.com

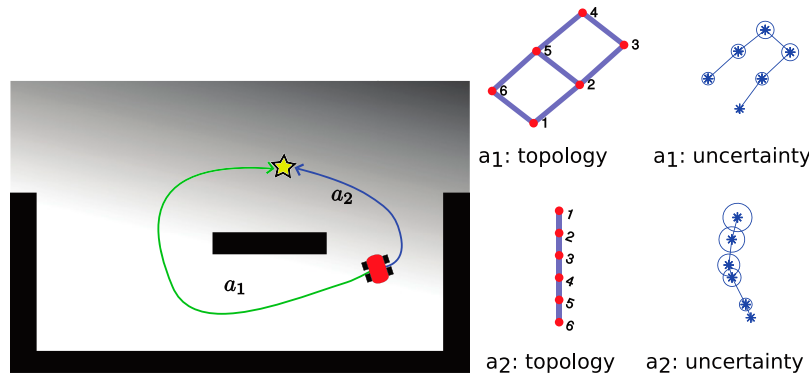


Figure 1. Conceptual illustration of topological belief space planning. Given that the objective is uncertainty minimization of the trajectory, an optimal action can be often deduced solely from posterior’s topology (*middle column*), that is, without predicting the values of future measurements and performing belief propagation (*right column*), but only by predicting that the trajectory a_1 will have loop closures due to re-observing previously mapped area, while a_2 won’t as it leads to an unknown (gray shaded) area (*left column*).

using only local computations (Lauritzen and Spiegelhalter 1988).

Graphical models have been much investigated in a design of efficient inference and learning algorithms (Sucar 2015), especially in the SLAM community (Folkesson and Christensen 2004; Frese 2006; Kaess et al. 2008, 2012; Thrun and Leonard 2008; Paskin 2003; Dellaert and Kaess 2017). However, a study of structural properties of probability distributions and their use in decision-making under uncertainty is relatively new (Kitanov and Indelman 2018; Chen et al. 2018, 2019; Khosoussi et al. 2019).

Moreover, sometimes it is possible to make decisions without explicitly performing the computationally expensive inference of future beliefs and cost function calculation by looking at some other features of the system highly related to the objective. Using this strategy, we can focus on important features first and leave the details for later if necessary, or when resources become available. The difficulty within this approach lies in identifying such features, that, on top of being discriminative, are less expensive to compute than solving the original optimization problem explicitly, and in understanding how close the solution based on them is to optimal.

In this work, we identify and study such features relevant to information-theoretic decision-making under uncertainty induced from *structural* information of the corresponding posterior probability distributions. Our approach is applicable to high-dimensional Gaussian belief space planning (BSP) problems that exhibit the specific structure of pairwise potentials, that is, each factor involves only two variables. It therefore particularly supports variants of the active SLAM problem (see, e.g., a recent survey (Placed et al. 2023)), that exhibit such a structure.

To get the intuition behind this idea, see Figure 1. Different actions generally lead to factor graphs with different topologies and more complex graphs tend to produce more accurate estimates, with smaller entropy. The idea is then to develop a computationally lightweight topological signature, that is, a function of the graph topology, that is

highly correlated with the original information-theoretic cost, and to utilize that signature for decision-making. Such a concept is therefore expected to be especially beneficial in decision-making problems with large graphs, such as in active SLAM.

1.1. Related work

BSP can be seen as a joint control and estimation problem in which an agent (robot) has to find an optimal control according to a specific task-related objective, which itself has to be estimated while accounting for different sources of uncertainty, for example, due to stochastic sensing, motion or environment. Some interesting instantiations of this problem are active SLAM (e.g., Atanasov et al., 2015; Du et al., 2011; Huang et al., 2005; Indelman et al., 2015; Kim and Eustice, 2014; Kopitkov and Indelman, 2017, 2019; Regev and Indelman, 2017; Stachniss et al., 2004; Valencia et al., 2012), active perception (Bajcsy 1988), sensor deployment and measurement selection (Joshi and Boyd 2009; Bian et al. 2006; Hovland and McCarragher 1997), graph reduction, and pruning (Kretzschmar and Stachniss 2012; Carlevaris-Bianco et al. 2014).

The general BSP problem with Markovian models can be naturally represented by a Partially Observable Markov Decision Process (POMDP). Determining globally optimal solutions to the POMDP problem is computationally intractable for most but the simplest real-world problems due to both the curse of dimensionality and the curse of history (see, e.g., Pineau et al. 2006). Computational complexity remains an issue even with simplifying assumptions on the probability distribution of the states, finite planning horizons, discrete states, actions or observations (Papadimitriou and Tsitsiklis 1987). As a result, a large subset of prior work has focused on approximately solving the POMDP problem to provide better scalability. Some examples include using sampling-based motion planners (e.g., Kavraki et al. 1996; Prentice and Roy 2009; Agha-Mohammadi et al. 2014; Davidson and

Hutchinson 2009), local optimization methods for continuous state spaces (e.g., Indelman et al. 2015; Platt et al. 2010; Van Den Berg et al. 2012), or point-based value iteration (e.g., Pineau et al. 2006; Porta et al. 2006; Kurniawati et al. 2008).

The paradigm proposed herein is related to computationally efficient information-theoretic BSP in the context of active pose SLAM. We formulate the optimization problem in a topological space where it is more easily solved, yet its solution is close to a solution of the original BSP optimization problem. Unlike existing topological methods in robot motion planning that consider a topology of the configuration space of the robot to generate or topologically classify paths (Kim et al. 2013; Bhattacharya et al. 2015) or to create and use topological maps of the environment (Ranganathan and Dellaert 2011; Choset and Nagatani 2001; Angeli et al. 2009; Paul and Newman 2010; Blöchliger et al. 2018), topology in our work refers to the structure of the belief, which encodes dependency relationships among states, and is used in planning.

An information-theoretic cost in an objective function of BSP is some measure of system uncertainty (information), typically (conditional) entropy or mutual information. Evaluating it requires determining the expected posterior belief of a candidate action. For Gaussian distributions the corresponding computations usually involve calculating a determinant of a posteriori covariance (information) matrix whose complexity is $O(n^3)$ in the general case, where n is the state dimension. Moreover, these calculations need to be performed for each candidate action. In Indelman et al. (2015) this challenge was addressed by resorting to the information form and utilizing sparsity; however, calculations still involve expensive access to marginal probability distributions. The rAMDLM approach (Kopitkov and Indelman 2017, 2019) performs a one-time calculation of marginal covariance of variables involved in the candidate actions, and then applies an augmented matrix determinant lemma (AMDLM) to efficiently evaluate the information-theoretic cost for each candidate action. Nevertheless, that approach still requires recovery of appropriate marginal covariances, the complexity of which depends on state-dimensionality and system sparsity.

Our work is motivated by the results of Khosoussi et al. (2019) in the context of measurement selection and pose-graph pruning problems in SLAM that characterizes the impact of the graph topology of SLAM (described by weighted tree-connectivity) on estimate reliability. In Chen et al. (2021), the authors extend these results to a 3D SLAM with relative pose measurements and make a comparison between the T-optimality and the D-optimality metrics.

1.2. Contributions

The main purpose of this work is to set theoretical foundations of the *topological BSP* ($t\text{-bsp}$) concept introduced in our preliminary conference publication (Kitanov and Indelman 2018), extending the main results from (Khosoussi et al. 2019) to BSP. In this paper we further

develop the potential of $t\text{-bsp}$ as an approximate method to efficiently solve BSP in high-dimensional state spaces in problems that exhibit certain structure and develop performance guarantees, as discussed below. Specifically, we consider the structure of pairwise potentials, which is often the case in active SLAM problems. The main contributions of this paper are as follows.

First, we propose new topological $t\text{-bsp}$ signatures: one is based on the number of spanning trees of the topological posterior factor graph representation, and the other on its von Neumann graph entropy (Mowshowitz and Dehmer 2012; Passerini and Severini 2009; Petz 2001). Both signatures have a strong correlation with the information-theoretic cost, but are computationally easier to calculate.

Second, we give a definition of an action consistent objective to formalize the equivalence of two criteria (functions) in terms of action ranking which extends previously defined action consistent state approximations proposed in (Elimelech and Indelman 2017b, 2017a, 2017c), building upon (Indelman, 2015, 2016). Definition 1 defines an approximation error in decision-making between any two functions (possibly on different domains) besides of the usual looking at the values of the same function for two different arguments (states or beliefs). This enables us to quantify an error between different cost functions and establish optimality guarantees. Specifically, we express the error of $t\text{-bsp}$ via bounds on the information-theoretic cost (Corollary 1), and derive the bounds analytically for an active 2D pose SLAM, considering the joint differential entropy as the cost function. These bounds are a function of topological terms, measurement noise density and an a-priori state estimate. A special emphasis is given to efficiency in calculation of error bounds to make our method appropriate for online use. Using these bounds, we are able to provide online performance guarantees of $t\text{-bsp}$, with any of the topological signatures. We emphasize these performance guarantees are with respect to a given solver (i.e., estimator) of the original information-theoretic BSP problem (6), and not with respect to the corresponding theoretical BSP problem (1). In this work we consider the solver to utilize maximum likelihood observations (Platt et al. 2010), although the bounds are applicable also to additional solvers.

Third, we propose a new method to calculate the number of spanning trees of a graph efficiently via a convergent power series. This is a result of reformulating Kirchoff’s matrix tree theorem via eigenvalues of the normalized Laplacian matrix and graph node degrees (Theorem 1). The power series reveals the connection between our proposed topological signatures and the information-theoretic cost in the limit.

Fourth, we develop an incremental variant of $t\text{-bsp}$ based on the von Neumann graph entropy that, under certain conditions, which are often met in operating conditions (e.g., SLAM), achieves near constant-time complexity.

We extensively evaluate our batch and incremental $t\text{-bsp}$ algorithms in realistic Gazebo/ROS simulations of active 2D

pose SLAM and in optimization of synthetic pose graphs. Additionally, we provide a proof of concept of τ -bsp in a small-scale real-world active SLAM experiment.

2. Notations and problem formulation

Decision-making under uncertainty can be formulated as a solution to BSP or stochastic control problem where the optimal non-myopic control action \mathcal{U}^* at planning time k is found with respect to the objective function J related to the design task as

$$\mathcal{U}^* = \operatorname{argmin}_{\mathcal{U}} J(\mathcal{U}), \text{ where for } \mathcal{U} = \mathcal{U}_{k:k+L-1},$$

$$J(b[X_k], \mathcal{U}) = \mathbb{E}_{\mathcal{Z}_{k+1:k+L}} \left\{ \sum_{l=0}^{L-1} c_l(b[X_{k+l}], \mathcal{U}_{k+l}) + c_L(b[X_{k+L}]) \right\}. \quad (1)$$

The BSP problem with Markovian models is an instance of POMDP. The state of the system X is not directly observable by the controller, but through a set of stochastic measurements \mathcal{Z} from which the future posterior beliefs $b[X_{k+l}]$ must be inferred upon optimization. The expectation in (1) is taken with respect to future (unknown) observations $\mathcal{Z}_{k+1:k+L}$. The optimal solution of BSP provides a control strategy \mathcal{U}^* for L look-ahead steps, but L can generally vary among control actions or planning sessions. We use $J(\mathcal{U}) = J(b[X_k], \mathcal{U})$ and $\mathcal{U}^* = \mathcal{U}_{k:k+L-1}^*$ to denote the objective and the optimal control, respectively, in a given planning session (at time k). The objective function in its general form reflects the design task through immediate cost functions c_l , which depend on the belief evolution $b[X_{k+l}]$ (to be defined) and a control action \mathcal{U}_{k+l} applied at time $k+l$, and through a final cost c_L . For example, the cost functions can be chosen to minimize trajectory uncertainty, time or energy required to reach a goal, state uncertainty of variables of interest at some specific time instant etc. In information-theoretic BSP, one is interested in state uncertainty minimization which can be expressed through some information-theoretic cost $c(\cdot)$ (see, e.g., Carrillo et al. 2012). This type of cost function is usually computationally the most expensive to optimize in many BSP problems in robotics.

The given notation abstracts the details of the controller design for the purpose of generality of the formulation. The basic problem we aim to solve is evaluating the objective function J , which is belief dependent (POMDP setting), for a finite sequence of controls (L look-ahead steps). This can be used in any POMDP solver no matter whether the control is closed-loop or open-loop. In this paper we implement both open-loop control (action sequences) and a Model Predictive Control (MPC) type planner and assume that a low-level controller can stabilize the system to follow the nominal path between two planning sessions using feedback based on the expected value of the state. This is similar to an LQG approach and is justified for systems that are either

linear or locally well approximated by their linearization (Van Den Berg et al. 2011).

Solving the optimization problem (1) explicitly would require belief propagation for a given control and evaluating an information-theoretic objective which we want to avoid and use a topological approach instead. Let us first look at how the belief evolves over time by separating controls and observations to those obtained by planning time k and to future controls and observations after L look-ahead steps.

2.1. Belief propagation

Let $\mathbb{P}(X_k|\mathcal{H}_k)$ represent the posterior probability density function (pdf) at planning time k over states of interest X_k of the robot. In the pose SLAM framework states of interest are the robot's current and past poses, that is, $X_k = \{x_0, x_1, \dots, x_k\}$. History $\mathcal{H}_k \doteq \{\mathcal{Z}_{1:k}, \mathcal{U}_{0:k-1}\}$ contains all observations $\mathcal{Z}_{1:k}$ and controls $\mathcal{U}_{0:k-1}$ by time k . Consider conventional state transition and observation models

$$x_{i+1} = f(x_i, u_i) + w_i, \quad z_{i,j} = g(x_i, x_j) + v_{i,j}, \quad (2)$$

with zero-mean Gaussian process and measurement noise $w_i \sim N(0, \Omega_w^{-1})$ and $v_{i,j} \sim N(0, \Omega_{vij}^{-1})$, and with known information matrices Ω_w and Ω_{vij} . The belief $\mathbb{P}(X_k|\mathcal{H}_k)$ can be written as

$$\mathbb{P}(X_k|\mathcal{H}_k) \propto \mathbb{P}(x_0) \prod_{i=1}^k \mathbb{P}(x_i|x_{i-1}, u_{i-1}) \mathbb{P}(z_i|X_i). \quad (3)$$

Let the future sampled states of the robot along one of its candidate paths \mathcal{P} generated at planning time k be $\{x_{k+1}, \dots, x_{k+L}\}$. The future posterior belief $b[X_{k+L}] = \mathbb{P}(X_{k+L}|\mathcal{Z}_{1:k+L}, \mathcal{U}_{0:k+L-1})$ that would be obtained by following the path \mathcal{P} , can be written in terms of the belief at planning time $\mathbb{P}(X_k|\mathcal{H}_k)$ and the corresponding state transition and observation models as

$$b[X_{k+L}] = \mathbb{P}(X_k|\mathcal{H}_k) \prod_{l=k+1}^{k+L(\mathcal{P})} \mathbb{P}(x_l|x_{l-1}, u_{l-1}) \mathbb{P}(z_l|X_l), \quad (4)$$

where $\mathcal{U}_{k:k+L-1}$ and $\mathcal{Z}_{k+1:k+L}$ represent controls and (unknown) observations, respectively, to be acquired by following the path \mathcal{P} .

2.2. Assumptions

The Topological BSP concept is general as an idea of introducing a new cost function, a function of the belief's topological features, to approximate actions ranking with respect to information-theoretic cost. However, identifying relevant topological representations and signatures in general BSP is an open problem. The approach we develop aims at BSP in problems that exhibit certain structure, such as active SLAM. In this work, our definition of topological space applies to all beliefs which can be factorized to a product of pairwise factors associated with a set of independent measurements. We develop this concept for a special class of Gaussian BSP, targeting

efficient decision-making in high-dimensional state spaces, with a joint differential entropy cost function and propose two topological signatures for the case of block isotropic measurement noise (5). For clarity here we only consider the final state cost term $\mathbb{E}[c_L(b[X_{k+L}])]$, that is, the joint entropy at the end of the planning horizon. The immediate information-theoretic cost functions can be treated in a similar way.

In summary, the assumptions used in this paper to develop τ -bsp are:

- (A1) the belief is modeled by a multivariate Gaussian distribution and can be factorized to a product of pairwise factors,
- (A2) the measurement noise covariance can be written in the form

$$\Sigma = \text{diag}(\sigma_1^2 I_m, \sigma_2^2 I_m, \dots, \sigma_\kappa^2 I_m), \quad (5)$$

for some κ blocks and m measurements with variances $\sigma_1^2, \dots, \sigma_\kappa^2$,

- (A3) taking maximum likelihood observations (Platt et al. 2010), and
- (A4) the information-theoretic cost is the joint differential entropy at the end of the planning horizon. Under (A3), the objective function becomes

$$J(\mathcal{U}) = N/2 \ln(2\pi e) - 1/2 \ln|I(X_{k+L})|, \quad (6)$$

where $I(X_{k+L})$ denotes the estimated Fisher information matrix of the robot's belief $b[X_{k+L}]$, and N represents the dimension of the state X_{k+L} .

Note that τ -bsp signatures designed to be action consistent under assumption (A4), will also be consistent with respect to an information-gain reward since all actions have the same prior entropy. Specifically, in this paper we demonstrate our approach in an active 2D pose SLAM problem in which all relative positions and orientations are normally distributed with variances σ_p^2 and σ_θ^2 , respectively. Measurement noise covariance Σ then can be written as a diagonal matrix $\Sigma = \text{diag}(\sigma_p^2 I_{2m}, \sigma_\theta^2 I_m)$, with m being the number of relative pose measurements.

Another application of τ -bsp are linear systems with relative measurements $z_{ij} = x_i - x_j + v_{ij}$ such as: active pose SLAM in which the robot's orientation is known (e.g., from a compass), sensor network localization, clock synchronization and motion consensus (Barooah and Hespanha 2007) in the context of measurement selection problem etc. For those types of problems under the assumptions (A1)-(A4), τ -bsp is always action consistent according to the Definition 1 since the determinant of the posterior information matrix is given by $|I(X_{k+L})| = \prod_{p=1}^{\kappa} \sigma_p^{-2N} t^\kappa$, where t is the number of spanning trees of a topological graph associated to a posterior factor graph, so negative $J(\mathcal{U})$ equals the topological signature based on the number of spanning trees.

An extension to a feature-based SLAM and map entropy is possible under the proposed framework as long as the

measurements can be expressed in the form of pairwise relations between states. This is quite common since the robot's landmark measurements are often given relative to the robot's pose, for example, as range or bearing measurements. The Fisher information matrix of the joint state of feature-based SLAM under the ML framework with isotropic measurement noise is derived in Chen et al. (2018) in terms of its graph topology and our results apply to it.

3. General approach in topological belief space planning

The main idea of topological BSP is ranking candidate actions based on structural information (topology) of their associated beliefs. Under the τ -bsp concept, the original problem (1) is reformulated in a topological space, where it can be solved more easily, and then its solution is related to the solution of the original problem. The importance of the topology and a quality of the solution based on it will clearly be dependent on the optimization objective, the influence of non-topological information (noise distribution, geometric realization) and diversity of candidate actions. As will be shown in this work for the information-theoretic BSP, the belief topology is of great importance and can be used to improve efficiency of decision-making.

3.1. Topological representation of beliefs

The structure of the probability density function (pdf) is determined from conditional independences present in it. In this paper, we use a factor graph probabilistic graphical model (Frey et al. 1997) to represent them explicitly. A factor graph defines two types of nodes: the nodes over state variables X , and the factor nodes $F = \{f_i\}$ such that $\mathbb{P}(X) = \prod_i f_i$. Variable nodes and factor nodes are independent sets in a bipartite graph. Edges go between factors and variables that those factors depend on. One of the reasons factor graphs are important is the existence of efficient inference algorithms that operate on them, for example, the sum-product message passing algorithm (Kschischang et al. 2001), which allow recursive computation of marginal distributions conditioned on observed data for both Maximum Likelihood (ML) and Maximum-A-Posteriori (MAP) estimation frameworks.

As will be shown later on, factor graphs are also a suitable graphical representation in information-theoretic BSP because of their connection to the objective function under the ML estimation framework for Gaussian systems.

Under assumption (A1), F contains only pairwise factors. The topological representation of the posterior belief is defined as a topological graph G associated to a posterior factor graph FG of a candidate action (e.g., see Figure 2), which is a graph isomorphic to the posterior Markov Random Field (MRF), that is, $G \simeq \text{MRF}(b[X_{k+L}])$. Different candidate paths \mathcal{P} typically yield different factor graphs. A graph G herein is understood to be simple (without

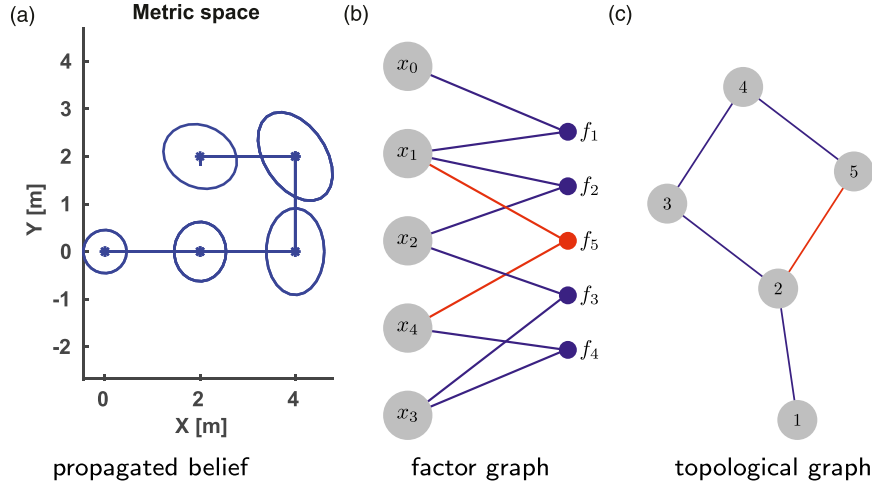


Figure 2. Each candidate action corresponds to posterior belief $b[X_{k+L}]$, that is, trajectory uncertainty (a) which can be represented with a factor graph (b) and assigned a topology. In the case of active pose SLAM, topology can be defined with a simple undirected graph $G = (V, E)$ such that graph nodes V represent robot’s poses and edges E pose constraints between them (c). Topological BSP aims to determine a graph invariant topological signature $s: G \rightarrow \mathbb{R}$ which is highly correlated with the information-theoretic cost and maintains action consistent decision-making.

self-loops or parallel edges), connected, and undirected. Each graph G corresponds to a topological space called the geometric realization (Archdeacon 1996). In the standard sense of topology, such a graph is a simplicial 1-complex (Gross and Tucker 1987). Therefore, a node in the graph, a point in its geometric realization and the corresponding point when embedding the graph in a topological space are often used interchangeably.

A topological signature $s: G \rightarrow \mathbb{R}$ is some graph invariant, that is, a function that assigns to G a number calculated from its structure according to a well-defined procedure, whose value is independent of how the graph is drawn or how its nodes are numbered.

3.2. Decision-making via t -bsp

Generally, we aim to find a topological representation G of the posterior belief corresponding to a control action \mathcal{U} and a function $s: G \rightarrow \mathbb{R}$, which we call topological signature, such that decision-making using the function s does not (or minimally) change the optimal action according to a BSP objective $J(\mathcal{U})$. In other words, suppose we can express the problem of minimizing the objective $J(\mathcal{U})$ and its solution $\mathcal{U}^\star = \arg\min_{\mathcal{U}} J(\mathcal{U})$ on the given set/interval of actions \mathcal{U} as maximizing the topological signature s on the same actions set where each action induces some topology, that is,

$$\hat{\mathcal{U}} = \arg\max_{\mathcal{U}} s(\mathcal{U}). \quad (7)$$

We say that decision-making is *action consistent* if and only if $\mathcal{U}^\star = \hat{\mathcal{U}}$. Moreover, if this is true for any subset of actions, a topological signature s preserves action order too. This is a much stronger condition than action consistency,

and as we will see, it will not always be required when we are looking only for the best action, which greatly simplifies design of a topological signature so that we can focus more on computational complexity.

Figure 3 shows all possible situations regarding the relation between the topological signature and objective function in t -bsp. In Figure 3(a), a signature is action consistent on the whole domain of control actions, and therefore the solution of t -bsp is the same as BSP. However, although ordering of actions is not kept in the whole domain in the case shown in Figure 3(b), the solution is still optimal. Figure 3(c) shows a case where an action ordering given by a signature influences the optimal action.

Clearly, topological signatures should be designed to exhibit strong correlation with the objective, but in the general case, the obtained best action $\hat{\mathcal{U}}$ may be somewhat different than the optimal action \mathcal{U}^\star from (1), leading to some error in the quality of solution.

We adopt the definition of action consistent state approximations proposed in Elimelech and Indelman (2017b) and modify it to support t -bsp and action consistent objective approximations in the following way.

Definition 1. Denote by $\hat{\mathcal{U}} = \arg\max_{\mathcal{U}} s(\mathcal{U})$ the optimal solution of t -bsp and by $\mathcal{U}^\star = \arg\min_{\mathcal{U}} J(\mathcal{U})$ the optimal solution of BSP problem (1). Let \mathcal{F} be a family of monotone decreasing real functions and $f: \mathbb{R} \rightarrow \mathbb{R} \in \mathcal{F}$ such that $f[s(\hat{\mathcal{U}})] = J(\hat{\mathcal{U}})$ and $f = \arg\min_{g \in \mathcal{F}} \gamma_g$, where the functional $\gamma_g = \max_{\mathcal{U}} |J(\mathcal{U}) - g[s(\mathcal{U})]|$. Then, the error of t -bsp is

$$\epsilon(J, s) \doteq \left| J(\mathcal{U}^\star) - f[s(\hat{\mathcal{U}})] \right|, \quad (8)$$

and γ_f the action ordering error of the signature s .




Figure 3. Different topological signatures and their influence on decision-making. The graphs show topological signature and negative objective as functions of control action.

In particular, $\epsilon(J, s) = 0$ corresponds to τ -bsp being action consistent, that is, $\hat{u} = u^*$, and when also f exists for which $\gamma_f = 0$, simplified representation preserves action order too and we call it action ordering (ranking) consistent. Essentially, f is a mapping from the signature image to the objective function image and is required for defining action consistency of objective functions.

Yet, as in Elimelech and Indelman (2017b), calculating $\epsilon(J, s)$ is essentially equivalent to solving the original problem. Therefore, a key aspect will be to provide online performance guarantees by developing bounds on $\epsilon(J, s)$ that can be evaluated online. One can then resort to topological BSP to drastically reduce computational cost while carefully monitoring a conservative estimate on the sacrifice in performance, that would be provided by the bound on $\epsilon(J, s)$. Another perspective of using this bound is to guarantee global optimality of the anytime τ -bsp algorithm we proposed in Kitanov and Indelman (2018). In that approach, actions were ranked by a topological signature and the objective function was evaluated sequentially from best to worst. Having bounds on τ -bsp error would provide a stopping condition for action consistent τ -bsp, similar to the action elimination scheme proposed in Elimelech and Indelman (2017b) and its application to belief sparsification. The optimal solution is then obtained by evaluating only a subset of candidate actions (for sufficiently tight bounds, empty set), while discarding the rest. Doing so will generally reduce the number of variables for which the marginal covariance needs to be recovered.

4. Information-theoretic topological BSP

We propose two topological signatures for solving information-theoretic BSP with uncertainty of all states considered. Both signatures are based on the spectrum of a graph (normalized) Laplacian associated to the posterior factor graph: a normalized number of spanning trees of a graph, $s_{ST}(G)$, motivated by the SLAM topological metric (Khosoussi et al. 2015) and adjusted for BSP, and a von Neumann graph entropy, $s_{VN}(G)$.

A Laplacian matrix of a graph $G = (V, E)$ is by definition $L(G) = D(G) - A(G)$, where $A(G)$ is its $|V| \times |V|$ adjacency matrix with elements

$$A(i, j) = \begin{cases} 1, & \text{if } (i, j) \in E \\ 0, & \text{otherwise} \end{cases}$$

and $D(G)$ its node degree matrix defined as a diagonal matrix with graph node degrees on its main diagonal, that is, $D(i, i) = d(i) = \sum_{j \in V} A(i, j)$. Therefore,

$$L(i, j) \triangleq \begin{cases} d(i), & \text{if } i = j \\ -1, & \text{if } i \neq j \text{ and } (i, j) \in E \\ 0, & \text{otherwise} \end{cases}$$

A normalized Laplacian matrix is defined as $\hat{L}(G) = D^{-1/2} L D^{-1/2}$ which can be also written in the form

$$\hat{L}(i, j) = \begin{cases} 1, & \text{if } i = j, d(i) \neq 0 \\ \frac{-1}{\sqrt{d(i)d(j)}}, & (i, j) \in E \\ 0, & \text{otherwise} \end{cases}$$

Both L and \hat{L} are symmetric and positive semi-definite so all their eigenvalues are real and non-negative. Hence, they can be ordered. For a given connected graph G with $n = |V|$ nodes, let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of L , and $\hat{\lambda}_1 \leq \hat{\lambda}_2 \leq \dots \leq \hat{\lambda}_n$, the eigenvalues of \hat{L} . The smallest eigenvalues $\lambda_1 = \hat{\lambda}_1 = 0$. Also, $\hat{\lambda}_n \leq 2$ and, since G has no isolated nodes, $\sum_{i=1}^n \hat{\lambda}_i = n$ (for proof see Chung 1997).

4.1. Normalized number of spanning trees

The number of spanning trees of a graph $t(G)$ can be determined using Kirchhoff's Matrix tree theorem (Biggs 1993, Theorem 6.3) which involves calculating any cofactor of a graph Laplacian, or equivalently the product of its non-zero eigenvalues, that is, $t(G) = |\tilde{L}| = 1/n \prod_{i=2}^n \lambda_i$. Here, \tilde{L} denotes L with the r -th row and column removed for some r . Typically, this number gets quite large for factor graphs in BSP so we use its logarithm in a topological signature instead to avoid arithmetic overflow in computer implementations, $\tau(G) = \ln t(G)$. As has been shown in Khosoussi et al. (2015), there exists a strong positive correlation between D-optimality criterion and $\tau(G)$ in general SLAM ML optimization while in some problems with linear measurement models, D-optimality criterion is solely characterized by $\tau(G)$. In Kitanov and Indelman (2018), we have proposed a normalization of

$\tau(G)$ based on its relation to the optimization objective in BSP which gives this BSP topological signature

$$s_{ST}(G) = \frac{\kappa}{2} \ln t(G) + \frac{n-1}{2} [\ln |\Omega_{vij}| - \ln(2\pi e)^\kappa], \quad (9)$$

normalized in such a way to account for different path lengths n and dimension of the robot's pose κ (e.g., $\kappa = 3$ in 2D active pose SLAM) that maintains a high correlation with the information-theoretic cost (6).

Determining s_{ST} is an operation with cubic complexity in the number of nodes for general dense graphs. Because SLAM graphs typically have a small number of edges due to a limited sensor range and field of view, their graph Laplacian is sparse and state-of-the-art algorithms use this fact to efficiently calculate $\tau(G)$ by performing a sparse matrix factorization. For example, let R be the lower triangular Cholesky factor of a reduced Laplacian $\tilde{L} \in \mathbb{R}^{n-1 \times n-1}$ of graph G , that is, $\tilde{L} = RR^T$, then

$$\tau(G) = \ln \det(\tilde{L}) = 2 \sum_{i=1}^{n-1} \ln R(i, i). \quad (10)$$

This approach, to be efficient, requires finding a good fill-reducing permutation P_L of \tilde{L} such that the Cholesky factor of $P_L^T \tilde{L} P_L$ is sparser than the Cholesky factor of \tilde{L} . The problem of finding the best ordering is an NP-complete problem (Yannakakis 1981) and is thus intractable, so heuristic methods are used instead. Overall, the computational time for s_{ST} depends on the state dimension and system sparsity.

4.2. Von Neumann graph entropy

The second topological signature we propose for information-theoretic BSP is the von Neumann entropy of G , $s_{VN}(G)$, and its simplification $\hat{s}_{VN}(G)$ by node degrees d of G .

The von Neumann entropy $s_{VN}(G)$ of graph G is the Shannon entropy associated with its normalized Laplacian's eigenvalues $\{\hat{\lambda}_i\}_{i=1}^n$ and was introduced in Passerini and Severini (2009). In this work we use the following definition

$$s_{VN}(G) = - \sum_{i=1}^n \hat{\lambda}_i / 2 \ln(\hat{\lambda}_i / 2). \quad (11)$$

Using Han's quadratic approximation (Han et al. 2012) and given that $\sum_{i=1}^n \hat{\lambda}_i^k = \text{Tr}[\hat{L}^k]$, it can be simplified to

$$\begin{aligned} s_{VN}(G) &\approx \frac{n \ln 2}{2} - \frac{1}{2} \sum_{i=1}^n \hat{\lambda}_i (\hat{\lambda}_i - 1) \\ &= \frac{n \ln 2}{2} - \frac{1}{2} \left(\text{Tr}[\hat{L}^2] - n \right). \end{aligned} \quad (12)$$

It can be easily seen that $\text{Tr}[\hat{L}^2] = n + \sum_{(i,j) \in E} \frac{1}{d(i)d(j)}$ from which the final expression for our approximated von

Neumann graph entropy follows which we use as a topological signature in BSP

$$s_{VN}(G) \approx \hat{s}_{VN}(G) = n/2 \ln 2 - \frac{1}{2} \sum_{(i,j) \in E} \frac{1}{d(i)d(j)}. \quad (13)$$

Notice that its computation depends only on graph node degrees and, in the general case, has quadratic complexity in the number of nodes, $O(n^2)$ but in the case of BSP, where \hat{L} is sparse, it depends only on the small number of non-zero elements of $A(G)$, that is, the number of edges $|E|$ in the graph G . Also, the expression (13) can be computed incrementally, as new edges (measurements) are added to the factor graph as the robot explores the environment or replans its path. This effectively improves computational complexity of the τ -bsp algorithm to nearly $O(1)$ since the number of measurements between two planning sessions is limited.

4.2.1. Incremental mode. Let $G_k = (V_k, E_k)$ be a topological graph of a prior belief $b[X_k]$, a prior graph at planning time k for short, and $G_{k+L} = (V_{k+L}, E_{k+L})$ the posterior topological graph corresponding to $b[X_{k+L}]$ after obtaining future predicted observations associated to a control action \mathcal{U} , a posterior graph for short. Their corresponding graph signatures, $\hat{s}_{VN}(G_k)$ and $\hat{s}_{VN}(G_{k+L})$, are

$$\hat{s}_{VN}(G_k) = (k+1)/2 \ln 2 - 1/2 \overbrace{\sum_{(i,j) \in E_k} 1/[d_k(i)d_k(j)]}^{q_k} \quad (14)$$

$$\begin{aligned} \hat{s}_{VN}(G_{k+L}) &= (k+L+1)/2 \ln 2 \\ &\quad - 1/2 \overbrace{\sum_{(i,j) \in E_{k+L}} 1/[d_{k+L}(i)d_{k+L}(j)]}^{q_{k+L}} \end{aligned} \quad (15)$$

By subtracting equation (14) from (15) we obtain the recursive update rule for the posterior graph signature

$$\hat{s}_{VN}(G_{k+L}) = \hat{s}_{VN}(G_k) + \frac{L}{2} \ln 2 - \frac{1}{2} \Delta q, \text{ where} \quad (16)$$

$$\Delta q = q_{k+L} - q_k. \quad (17)$$

Now, the key observation is that in order to calculate Δq we do not need to iterate over all edges of the corresponding graphs, but only over those whose node degrees changed or are associated to the new states added to the prior graph, that is,

$$\begin{aligned} \Delta q &= \sum_{(i,j) \in E_\delta} \left(\frac{1}{d_{k+L}(i)d_{k+L}(j)} - \frac{1}{d_k(i)d_k(j)} \right) \\ &\quad + \sum_{(i,j) \in \Delta E} \frac{1}{d_{k+L}(i)d_{k+L}(j)}. \end{aligned} \quad (18)$$

In equation (18), $\Delta E = E_{k+L} - E_k$ represents newly added edges to the prior graph G_k due to the control action \mathcal{U} yielding posterior graph G_{k+L} , V_I nodes incident to ΔE




Figure 4. Incremental t -bsp: Consider the planning session starting at time $k=4$ when the robot’s current pose is x_4 , and a control action that generates a candidate path from which states $\{x_5, x_6\}$ (marked with red circles) are sampled and added to the prior graph together with new edges

$\Delta E = \{(x_4, x_5), (x_5, x_6), (x_5, x_3)\}$ marked with blue dashed lines as a result of predicted relative pose measurements. In that case, involved states are $V_I = \{x_3, x_4\}$ (marked with blue circles), and affected edges in the prior graph $E_\delta = \{(x_2, x_3), (x_3, x_4), (x_2, x_4)\}$.

without new nodes and E_δ edges from E_{k+L} incident with V_I that belong also to E_k . An example illustrating the incremental t -bsp concept is given in Figure 4.

The complexity of updating \hat{s}_{tN} graph signature per candidate action is $O(|\Delta E| + |MB(V_I)|)$, where $MB(V_I)$ denotes the Markov blanket in G_k of the variables involved and is usually bounded owing to limits on sensor range. Since the number of actions per planning session is also bounded, for sparse graphs this means that incremental t -bsp algorithm has effective complexity $O(1)$.

4.3. Maximum likelihood estimation

In this work, we investigate error bounds of t -bsp on the active 2D pose SLAM problem. We first examine results on passive SLAM reliability regarding its graph structure described by a topological metric $\tau(G) \doteq \ln(\iota(G))$ proposed in Khosoussi et al. (2015) and then, in the next section, we show how these results can be extended to a BSP problem to provide online performance guarantees of t -bsp. In that section, it will also be clear how we chose the topological signature s_{ST} .

In the context of Maximum Likelihood Estimation (MLE) in pose SLAM, the optimal set of poses X_k for which the belief (3) is maximized can be obtained by fixing one of the poses, for example, x_0 , and treating the rest as unknown (or with an uninformative prior) while minimizing the sum of weighted squared errors between predicted and measured relative poses,¹ that is,

$$\hat{X}_k = \operatorname{argmax}_{X_k} \mathbb{P}(\Delta_k | X_k) \equiv \operatorname{argmin}_{X_k} \|\Delta_k - h(X_k)\|_{\Sigma^{-1}}^2. \quad (19)$$

In this formulation, a measurement Δ_k represents a vector of m stacked relative pose measurements $z_{i,j}^r \in \text{SE}(2)$, $r = 1, 2, \dots, m$ from motion and observation model (2) at time k with $m = |E|$, the number of edges in the topological graph of the belief (3). Relative pose measurements in pose SLAM resulting from state transitions can be obtained by the motion composition $z_{i+1,i}(x_i, x_{i+1}) = \ominus x_i \oplus x_{i+1} = \ominus x_i \oplus$

$f(x_i, u_i, w_i)$. In this work, we assume independent relative pose measurements with additive noises

$$\Delta_k = h(X_k) + v_k, v_k \sim N(0, \Sigma^{-1}). \quad (20)$$

For simplicity, we also assume a 2D pose SLAM setting in which all relative positions and orientations between poses x_i and x_j have equal variance, σ_p^2 and σ_θ^2 , respectively, that is, $\Omega_{vij} = \text{diag}(\sigma_p^{-2}, \sigma_p^{-2}, \sigma_\theta^{-2})$. Measurement noise covariance Σ in that case can be written as a diagonal matrix $\Sigma = \text{diag}(\sigma_p^2 I_{2m}, \sigma_\theta^2 I_m)$ by reordering elements of Δ_k .

The information matrix $I(X_k)$ of the MLE is $I(X_k) = H^T \Sigma^{-1} H$ (Sorenson 1980), where $H = \partial h / \partial X_k$ is a measurement Jacobian. $I(X_k)$ evaluated at the true value of X_k is known as Fisher Information Matrix (FIM) and its inverse the Cramér-Rao lower bound (CRLB). Commonly, FIM is approximated with $I(\hat{X}_k)$. In Khosoussi et al. (2015) bounds of the determinant of $I(X_k)$ are expressed in terms of pose SLAM graph topology, geometry and noise as

$$\begin{aligned} 3\tau(G) + \ln \det I_o(X_k) &\leq \ln \det I(X_k) \\ &\leq 2\tau(G) + \ln \det(\tilde{L} + \Psi I) + \ln \det I_o(X_k), \end{aligned} \quad (21)$$

where \tilde{L} is a reduced Laplacian of a graph G obtained by removing an arbitrary r -th row and r -th column of the graph Laplacian L , $I_o(X_k)$ estimated information matrix based on the odometry measurements $\Delta_k = \{z_{i+1,i}, i < k\}$, and $\Psi \doteq \xi^2 \text{dist}_{\max}^2$ where $\xi = \sigma_\theta / \sigma_p$, and $\text{dist}_{\max}^2 = \max_{i \in V} \sum_{(i,j) \in E} \|x_i - x_j\|^2$. Notice that generally Ψ depends on the noise variances, geometry, and topology of the SLAM graph. In Khosoussi et al. (2015) these bounds are not used except to prove the limiting case, $\Psi \rightarrow 0$. Therefore, for small values of Ψ a good approximation of $\ln \det I(X_k)$ is its bound that depends on $\tau(G)$, that is, lower and upper bounds become tight. Even when it is not negligible, if there is only one path realization to consider as in passive SLAM, information gain of adding relative pose measurements with a constant noise distribution to a SLAM odometry graph is solely characterized by graph G topology, that is, $\Psi = \Psi(G)$. Similar logic applies to graph pruning and the measurement selection problem (Khosoussi et al. 2019) where all graphs are only subgraphs of the original graph with the same embedding in metric space. To demonstrate why we cannot use the same metric $\tau(G)$ in BSP problems nor guarantee optimality using the bounds on $I(X_k)$, consider the example given in Figure 5.

So, in BSP we need to consider *different* path realizations X_{k+L} and therefore $\Psi = \Psi(\mathcal{U}) = \Psi(X_{k+L}(\mathcal{U}), G(\mathcal{U}))$, in contrast to graph pruning and measurement selection. Notice that we do not need to propagate belief in planning for determining Ψ under the ML observations assumption, that is, $\mathbb{E}[X_{k+L}] = [\hat{X}_k \ x_{k+1} \dots x_{k+L}]$ where future sampled poses from the path corresponding to action \mathcal{U} are added to the prior state estimate \hat{X}_k which is the same for all actions.




Figure 5. BSP vs. the measurement selection problem: In measurement selection problems, the robot considers a single path and its factor graph (e.g., marked with blue color) at planning time k and which subset of measurements (marked with blue dashed lines) to take. Pose samples X_{k+L} are fixed and minimizing entropy (6) by a measurements subset (action) is the same as maximizing $\ln |I(X_{k+L})| - \ln |I_o(X_{k+L})|$. Its bounds as can be seen from (21) depend only on the topology and ζ , that is, all actions share the same pose variables X_{k+L} which can be considered constant in optimization. However, in BSP where a robot needs to compare different paths \mathcal{U} corresponding to different factor graphs (e.g., red and green), this is not true anymore. To determine the best action according to (6), one has to account additionally for different path geometries and path lengths and larger variety of topologies.

4.4. Entropy bounds for active SLAM

In BSP, we have to consider multiple path realizations from different controls, with greater variety in both topology of factor graphs and other non-topological factors that influence estimation accuracy, for example, non-fixed geometry and different path lengths. In Kitanov and Indelman (2018), we developed a topological signature s_{ST} (9) for t -bsp. Here, we show its derivation and provide global optimality guarantees of t -bsp for any signature via estimated entropy bounds that are functions of s_{ST} for active SLAM and related problems as described in the Section 2.2.

For a control action \mathcal{U} corresponding to robot's poses $X_{k+L} = \{x_1, \dots, x_{k+L}\}$ each of dimension κ , the entropy of the future posterior belief $b[X_{k+L}]$ can be written using equation (6) as

$$J(\mathcal{U}) = (k+L)\kappa/2 \ln(2\pi e) - 1/2 \ln |I(X_{k+L})| \quad (22)$$

Notice that the number of graph nodes $|V| = n = k+L+1$. Using Lemma 2 and Lemma 3 from Khosoussi et al. (2015), for posterior belief $\mathbb{P}(X_{k+L} | \mathcal{U}_{0:k+L-1})$ corresponding to an odometry factor graph whose topological graph G^o is a tree, it follows

$$\text{Indet} I_o(X_{k+L}) = \underbrace{\tau(G^o)}_0 + \ln \left(\sigma_p^{-4} \sigma_\theta^{-2} \right)^{k+L} = (k+L) \text{Indet}(\Omega_{v_{ij}}) \quad (23)$$

Inequalities (21) and equations (22) and (23) give the entropy bounds $\mathcal{LB}[J(\mathcal{U})] \leq J(\mathcal{U}) \leq \mathcal{UB}[J(\mathcal{U})]$, where

$$\begin{aligned} \mathcal{UB}[J(\mathcal{U})] &= \frac{(n-1)\kappa}{2} \ln(2\pi e) - \frac{1}{2} [3\tau(G) + (n-1) \ln |\Omega_{v_{ij}}|] \\ &= -\frac{3}{2} \tau(G) - \frac{n-1}{2} [\ln |\Omega_{v_{ij}}| - \ln(2\pi e)^\kappa] = -s_{ST}(\mathcal{U}) \end{aligned} \quad (24)$$

Similarly, for the lower bound we get

$$\mathcal{LB}[J(\mathcal{U})] = -s_{ST}(\mathcal{U}) - 1/2 \ln |\tilde{L} + \Psi(\mathcal{U}) I_{n-1}| + 1/2 \tau(G) \quad (25)$$

$$= -s_{ST}(\mathcal{U}) - 1/2 \left[\ln |\tilde{L} + \Psi(\mathcal{U}) I_{n-1}| - \ln |\tilde{L}| \right] \quad (26)$$

$$= -s_{ST}(\mathcal{U}) - 1/2 \ln |I_{n-1} + \tilde{L}^{-1} \Psi(\mathcal{U})|. \quad (27)$$

In equation (26) we used the Matrix tree theorem that states $t(G) = |\tilde{L}|$ (recall $\tau(G) = \ln t(G)$), and in equation (27) Shur's determinant lemma. From equation (27) we see that when $\Psi \rightarrow 0$, the entropy goes to the BSP topological signature and solely depends on the belief's topology and path length. Otherwise, we have to account for the graph embedding in the metric space as well, as it appears in the scalar function Ψ of the second term in the above equations. Now, given all candidate control actions, the following corollary provides the bounds on the t -bsp error as defined by Definition 1.

Corollary 1. For the error $\epsilon(J, s)$ of t -bsp, we can write

$$\epsilon(J, s) \leq \Delta J_{max},$$

where $\Delta J_{max} = \mathcal{UB}[J(\hat{\mathcal{U}})] - \min_{\mathcal{U}} \mathcal{LB}[J(\mathcal{U})]$, and $\hat{\mathcal{U}} = \text{argmax}_{\mathcal{U}} s(\mathcal{U})$.

Proof. See Appendix B.

4.4.1. Efficient calculation of entropy bounds. The upper bound of the entropy is already determined by the topological signature s_{ST} as can be seen from equation (24). Moreover, $\mathcal{UB}[J(\hat{\mathcal{U}})] = -s_{ST}(\hat{\mathcal{U}})$ is the minimum upper bound of all actions since we select action in t -bsp according to equation (7). However, calculating the lower bound requires an additional cost due to the second term in equations (25)-(27). Calculating (25)-(26) requires evaluating the determinant of a sparse matrix $M = \tilde{L} + \Psi I_{n-1} \in \mathbb{R}^{n-1 \times n-1}$, whereas calculating (27) is even more complex because it requires inverting \tilde{L} and finding a determinant of a dense matrix.

Another idea is to find some fast method for limiting the determinant of M from above to replace the second term in equation (25) that will not introduce a big difference in tightness of the lower bound.

4.4.2. Exact lower bound. A direct approach performs from scratch some sparse matrix factorization of M , for example, Cholesky factorization $M = RR^T$ where R is the lower triangular matrix, from which then it is easy to calculate its determinant. However, this approach, to be efficient, still requires finding a good fill-reducing permutation P_M of M , but we notice that some calculations from the topological signature can be re-used. In particular, since M differs from \tilde{L} only in diagonal elements, they both have the same sparsity pattern. Therefore, if $P_{\tilde{L}}$ is the

best fill-reducing permutation of \tilde{L} (already found for determining $|\tilde{L}|$), it can be re-used for calculation of the lower bound, that is, $P_M = P_{\tilde{L}}$.

4.4.3. Hadamard bound. Since \tilde{L} is a reduced graph Laplacian, it is symmetric positive definite (SPD), and because also $\Psi > 0$, the matrix M is SPD. For large values of Ψ , the matrix M becomes diagonally dominant, in fact strongly so, and Hadamard inequality gives a good approximation of its determinant, that is,

$$|\tilde{L} + \Psi I_{n-1}| \leq \prod_{i=2}^n [d(i) + \Psi]. \quad (28)$$

Calculation of the right side of inequality (28) requires only multiplication of node degrees with added value of Ψ . Applying (28) to equation (25), we can get a somewhat more conservative but faster to compute lower bound

$$\mathcal{LB}_H[J(\mathcal{U})] = -s_{ST}(\mathcal{U}) + \frac{1}{2} \left(\tau(\mathcal{U}) - \sum_{i=2}^n \ln[d_{\mathcal{U}}(i) + \Psi(\mathcal{U})] \right). \quad (29)$$

Algorithm 1. t-bsp with performance guarantees

Input: set of factor graphs FG corresponding to control actions \mathcal{U}

Output: approximate solution to the BSP $\hat{\mathcal{U}}$ and its performance guarantees or actions subset A

Parameters: graph signature
 $s \in \{s_{ST}, s_{VN}, \hat{s}_{VN}\}$, uncertainty margin γ (optional)

```

1  $S = \emptyset$ 
2 represent each  $FG$  with a topological graph  $G(\mathcal{U})$ 
3 foreach  $\mathcal{U}$  do
4   evaluate the topological signature
    $s(\mathcal{U}) \doteq s[G(\mathcal{U})]$ 
5    $S \leftarrow S \cup \{\mathcal{U}, s(\mathcal{U})\}$ 
6 select  $\hat{\mathcal{U}} = \arg \max_{\mathcal{U}} s(\mathcal{U})$ 
7  $\mathcal{UB}[J(\hat{\mathcal{U}})] = -s_{ST}(\hat{\mathcal{U}})$ 
8 if uncertainty margin  $\gamma$  is given then
9   if  $\mathcal{UB}[J(\hat{\mathcal{U}})] < \gamma$  then
10    return  $\hat{\mathcal{U}}$ 
11  else
12    return  $A = \text{sort}(S)$  with respect to
    topological signature
13 else
14    $\mathcal{LB}_{\min} = \min_{\mathcal{U}} \mathcal{LB}_H[J(\mathcal{U})]$  using eq. (29)
15    $\Delta J_{\max} = \mathcal{UB}[J(\hat{\mathcal{U}})] - \mathcal{LB}_{\min}$ 
16    $A = \{\mathcal{U} \in S : \mathcal{LB}[J(\mathcal{U})] < \mathcal{UB}[J(\hat{\mathcal{U}})]\}$ 
17 return  $\hat{\mathcal{U}}$ , t-bsp error bound  $\Delta J_{\max}$ ,  $A$ 

```

4.5. t-bsp algorithm

Algorithm 1 summarizes our proposed method and highlights its possible uses in BSP regarding desired performance specifications. Performance guarantees can be either in the form of selected solution's entropy upper bound, that is, guarantees on the accuracy, or in bounding the error of t-bsp with respect to the optimal solution.

The first form can be used when the maximum admissible path uncertainty is known at planning time, for example, for obstacle avoidance. In that case, one can get an answer if a t-bsp solution satisfies the specification by ranking actions using the very efficient topological signature \hat{s}_{VN} and calculating only its entropy's upper bound (Alg. 1, line 7). If global optimality guarantees are required, the lower entropy bound needs to be calculated for all actions. Although in equation (29) it is given as a function of s_{ST} , in the next section we will see how it can be expressed via \hat{s}_{VN} . In the first usage, if uncertainty specification is not met by t-bsp one can still use its ranked actions set A in anytime algorithm as we proposed in Kitanov and Indelman (2018). Also, if an optimal solution needs to be found, we can use t-bsp to eliminate suboptimal actions.

5. Von Neumann entropy based t-bsp with performance guarantees

In this section we develop a mechanism to provide global optimality guarantees as functions of the von Neumann entropy topological signature \hat{s}_{VN} (13). The key idea is to make a connection of \hat{s}_{VN} to the normalized number of spanning trees topological signature s_{ST} based on which we have already established t-bsp error bounds (see Figure 6). At the same time, we want to preserve computational efficiency. Theorem 1 is central for this, and has also much wider applicability than in the context of BSP as a reformulation of the Matrix tree theorem.

5.1. Taylor series for $\tau(G)$

First, we provide a method for calculating the number of spanning trees of a graph via a convergent Taylor series expansion as a function of only the graph's node degrees, which avoids matrix factorization, the current state-of-the-art.

Recall that $\tau(G) = \sum_{i=2}^n \ln \lambda_i - \ln n$ and representing it with a Taylor power series is possible only for all $\lambda_i \in \mathfrak{S}$ inside the interval of convergence of the logarithm function, $\mathfrak{S} = (0, 2]$ around point $\lambda_{i0} = 1$. However, the spectrum of the graph Laplacian can be outside this interval. We know that the largest Laplacian eigenvalue is bounded by the number of non-isolated vertices, that is, $\lambda_n \leq n$ (Mohar 1991, Theorem 2.2). On the other hand, the normalized graph Laplacian's spectrum without zero lies inside \mathfrak{S} for every graph G (Wilson and Zhu 2008) and therefore the von

Neumann graph entropy is well-defined via its Taylor series. This motivates us to look for an alternative representation of $\tau(G)$ that involves eigenvalues of \widehat{L} . The following theorem gives the required relation.

Theorem 1. *Let $G = (V, E)$ be a connected, undirected graph with n nodes, and let $t(G)$ be the number of spanning trees of G . Then,*

$$t(G) = \frac{\prod_{i=1}^n d_i}{\sum_{i=1}^n d_i} \prod_{i=2}^n \widehat{\lambda}_i \quad (30)$$

where $\widehat{\lambda}_2, \widehat{\lambda}_3, \dots, \widehat{\lambda}_n$ are non-zero eigenvalues of the normalized Laplacian matrix \widehat{L} of G , and d_1, d_2, \dots, d_n node degrees of G .

Proof. See Appendix C.

We will often use the logarithm of $t(G)$ in the following study

$$\tau(G) \doteq \ln t(G) = \ln \frac{\prod_{i=1}^n d_i}{\sum_{i=1}^n d_i} + \sum_{i=2}^n \ln \widehat{\lambda}_i. \quad (31)$$

Let us denote by $T_{x_0}(x)$ the Taylor series expansion of the function $f(x) = \ln x$ centered around the point x_0 . Then, on the interval of convergence \mathfrak{S} , $f(x) = T_{x_0}(x) = \lim_{r \rightarrow \infty} T_{x_0}^r(x)$ where the r -th partial sum of the Taylor series $T_{x_0}^r(x)$ is given by

$$T_{x_0}^r(x) = \ln x_0 + \sum_{i=1}^r \frac{(-1)^{i-1}}{ix_0^i} (x - x_0)^i.$$

Considering function f defined over the non-zero normalized graph Laplacian's eigenvalues $\widehat{\lambda}_i$ and the linearization point $x_0 = 1$, we obtain the following r -th order Taylor series approximation of $\ln \widehat{\lambda}_i$

$$\begin{aligned} \ln \widehat{\lambda}_i &\approx T_1^r(\widehat{\lambda}_i) = \\ &\sum_{k=1}^r \frac{(-1)^{k-1}}{k} (\widehat{\lambda}_i - 1)^k, \forall \widehat{\lambda}_i \in \text{spectrum}(\widehat{L}(G)) \setminus \{0\}. \end{aligned} \quad (32)$$

Calculating the topological signature s_{ST} directly using the Matrix tree theorem, using Theorem 1 or via a Taylor series approximation using equation (32) would require determining the graph spectrum of the associated graph matrix, which is not desirable, from a computational perspective, in high-dimensional state spaces. Similarly, applying equation (10) requires matrix factorization, whose time complexity is only by a constant scaling factor smaller than the time complexity of calculating the information-theoretic cost.

In this paper, we show how this can be avoided by developing the power series of $\tau(G)$ and expressing it with sums of traces of powers of normalized graph Laplacian which can be more easily calculated because they involve multiplication of sparse matrices. Importantly, for low order Taylor series approximations, they can be given in closed-form as simple functions of graph node degrees.

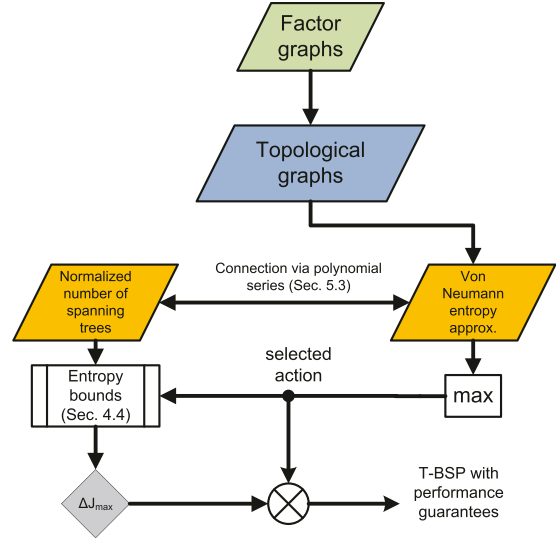


Figure 6. Schematic overview: Topological belief space planning based on approx. von Neumann entropy with performance guarantees.

Denote the k -th term in the sum given by equation (32) by

$$s_k(\widehat{\lambda}_i) \doteq \frac{(-1)^{k-1}}{k} (\widehat{\lambda}_i - 1)^k.$$

After expanding a power using Binomial theorem we can write

$$s_k(\widehat{\lambda}_i) = \sum_{l=0}^k \underbrace{\frac{(-1)^{k+l-1}}{k} \binom{k}{l}}_{a_{kl}} \widehat{\lambda}_i^{k-l}, \text{ so} \quad (33)$$

$$T_1^r(\widehat{\lambda}_i) = \sum_{k=1}^r \sum_{l=0}^k a_{kl} \widehat{\lambda}_i^{k-l}. \quad (34)$$

Observe that in equation (34), a_{kl} is defined by the horizontal underbrace in equation (33).

By convention $\widehat{\lambda}_1 = 0$ and given that $\sum_{i=2}^n \widehat{\lambda}_i^{k-l} = \sum_{i=1}^n \widehat{\lambda}_i^{k-l} = \text{Tr}[\widehat{L}^{k-l}] \Rightarrow$

$$\begin{aligned} \sum_{i=2}^n T_1^r(\widehat{\lambda}_i) &= \sum_{i=1}^n \sum_{k=1}^r \sum_{l=0}^k a_{kl} \widehat{\lambda}_i^{k-l} = \sum_{k=1}^r \sum_{l=0}^k a_{kl} \sum_{i=1}^n \widehat{\lambda}_i^{k-l} \\ &= \sum_{k=1}^r \sum_{l=0}^k a_{kl} \text{Tr}[\widehat{L}^{k-l}]. \end{aligned} \quad (35)$$

For the Taylor polynomial $T_1^r(\widehat{\lambda}_i)$, $\exists \zeta = \zeta(\widehat{\lambda}_i)$ such that a remainder term $R_1^r(\widehat{\lambda}_i) = f(\widehat{\lambda}_i) - T_1^r(\widehat{\lambda}_i)$ is

$$R_1^r(\widehat{\lambda}_i) = \frac{f^{(r+1)}(\zeta)}{(r+1)!} (\widehat{\lambda}_i - 1)^{r+1} = \frac{(-1)^r}{r+1} \left(\frac{\widehat{\lambda}_i - 1}{\zeta} \right)^{r+1}, \quad (36)$$

where for $\widehat{\lambda}_i < 1, \zeta \in [\widehat{\lambda}_i, 1]$ and for $\widehat{\lambda}_i \geq 1, \zeta \in [1, \widehat{\lambda}_i]$.

Finally, using Theorem 1 and equation (35) we obtain

$$\widehat{\tau}(G, r) = \ln \frac{\prod d_i}{\sum d_i} + \sum_{k=1}^r \sum_{l=0}^k a_{kl} \text{Tr}[\widehat{L}^{k-l}]. \quad (37)$$

$\widehat{\tau}(G, r)$ is the approximation of $\tau(G)$ by the r -th order Taylor polynomial and can be used to calculate a topological signature of $t\text{-bsp}$ with reduced complexity compared to calculating it by matrix factorization as in (10), by an operation which involves sparse matrix multiplications instead.

5.2. Performance guarantees based on $\widehat{\tau}$

The next proposition reveals a property of the remainder term which can be used to provide performance guarantees of $t\text{-bsp}$ efficiently, from the approximated number of spanning trees of the topological graph by its Taylor polynomial, without the need to know the remainder of all, but only the selected control action.

Proposition 1. *If the logarithm of the number of spanning trees of the graph G is approximated by the r -th order Taylor polynomial $\widehat{\tau}(G, r) = \ln(\prod d_i / \sum d_i) +$*

$\sum_{k=1}^r \sum_{l=0}^k a_{kl} \text{Tr}[\widehat{L}^{k-l}]$, its approximation error $\varepsilon(G, r) = \tau(G) - \widehat{\tau}(G, r)$ has the following properties:

- i. $\varepsilon(G, r) = \sum_{i=2}^n R_1^r(\widehat{\lambda}_i)$, where $R_1^r(\widehat{\lambda}_i)$ is given by equation (36),
- ii. $\lim_{r \rightarrow \infty} \varepsilon(G, r) = 0$,
- iii. $\varepsilon(G, r) \leq 0$ for odd polynomial orders r .

Proof. See Appendix D.

Figure 7 illustrates the properties of the remainder stated by the Proposition 1.

The question that naturally arises is whether we can use the approximated logarithm of the number of spanning trees to provide performance guarantees if we choose such an order r that makes $\widehat{\tau}(G, r) \geq \tau(G)$. From equation (29) it is easy to see that the lower bound of the entropy calculated from $\widehat{\tau}(G, r)$ will remain the lower bound, but the upper bound might be compromised depending on the approximation error $\varepsilon(G, r)$. So, ΔJ_{\max} needs to account for that change as well. Let's denote by $\widehat{\Delta J}_{\max}$ the $t\text{-bsp}$ error bound calculated by using $\widehat{\tau}(G, r)$ instead of the true $\tau(G)$. Then,

$$\Delta J_{\max} = \widehat{\Delta J}_{\max} - \kappa/2 \varepsilon(\widehat{U}, r), \quad (38)$$

where $\varepsilon(\widehat{U}, r) = \tau(\widehat{U}) - \widehat{\tau}(\widehat{U}, r)$.

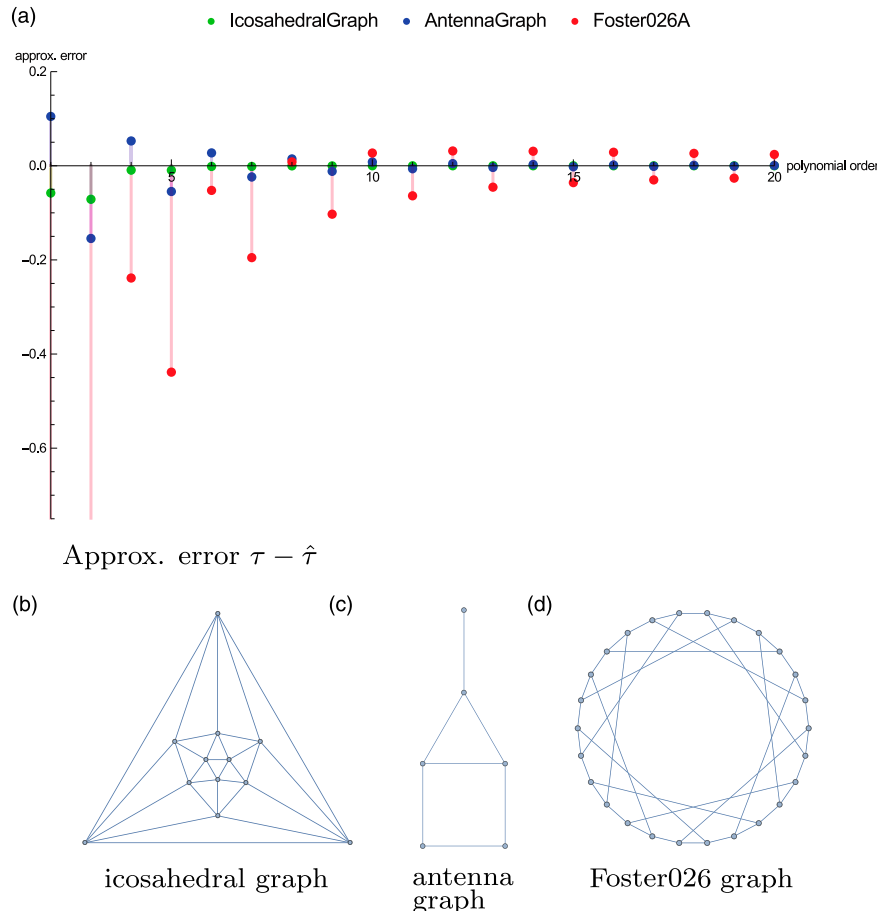


Figure 7. Taylor series approximation error $\tau - \widehat{\tau}$ of the logarithm of the number of spanning trees for three different graphs: icosahedral, antenna, and Foster026. For odd polynomial orders, this error is either zero or negative.

5.3. Relation between topological signatures

By establishing a relation between s_{ST} and \widehat{s}_{VN} it would be possible to use benefits from both worlds, speed of s_{VN} with optimality guarantees for s_{ST} .

A relation between two proposed information-theoretic BSP topological signatures s_{ST} and $\widehat{s}_{VN}(G)$ can be derived in the limit $r \rightarrow \infty$ using Theorem 1 and Taylor series expansion of the logarithms of the normalized graph Laplacian's eigenvalues, which appear in both of them (see Appendix E):

$$s_{ST}(G) = \lim_{r \rightarrow \infty} s_{ST}(\widehat{\tau}(G, r)), \text{ where} \quad (39)$$

$$\begin{aligned} s_{ST}(\widehat{\tau}(G, r)) &= \frac{\kappa}{4} r(r-1) \widehat{s}_{VN}(G) + \\ &\frac{\kappa}{2} \left(\sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] + \ln \frac{\prod d_i}{\sum d_i} \right) + \\ &\left[\frac{\kappa}{2} \left(r - \frac{r(r-1)(\ln 2 + 1)}{4} - H(r) \right) + \frac{1}{2} \ln \frac{|\Omega_{vij}|}{(2\pi e)^\kappa} \right] n + \\ &\frac{\kappa}{2} H(r) - \frac{1}{2} \ln \frac{|\Omega_{vij}|}{(2\pi e)^\kappa}, \text{ and} \end{aligned} \quad (40)$$

$H(r) = \sum_{k=1}^r 1/k$ is a harmonic number of r . In practice, we can often get very good approximation for finite small orders r .

For $r=3$, using equation (61) we find a simple closed-form expression given by equation (41), that does not involve traces of matrix powers, and in combination with Proposition 1 gives a way to formulate τ -bsp in terms of the approximated von Neumann entropy and lower bound the information-theoretic cost for all actions \mathcal{U} .

For the logarithm of the number of spanning trees of a graph G corresponding to a control action \mathcal{U} approximated by the third order polynomial we can write

$$\begin{aligned} \widehat{\tau}(G, 3) &= \ln \frac{\prod d_i}{\sum d_i} + \frac{7}{3} \left(\widehat{s}_{VN}(G) - \frac{n \ln 2}{2} \right) + \\ &\frac{1}{3} \sum_{(i,k) \in E} \widehat{l}_{ik} \left(\sum_j \widehat{l}_{ij} \widehat{l}_{jk} \right) + \frac{11}{6}. \end{aligned} \quad (41)$$

In equation (41) $\widehat{L} = (\widehat{l}_{ij})$ which means that the right-hand side is purely a function of nodes' degrees and a number of graph nodes, both of which are easy to maintain for graph operations of adding or deleting an edge encountered in BSP or graph pruning problems.

Substituting (41) into (9), we obtain the third order relation between the topological signatures

$$\begin{aligned} \widehat{s}_{ST}(G, 3) &= \frac{7\kappa}{6} \widehat{s}_{VN}(G) + \\ &\frac{\kappa}{2} \left[\ln \frac{\prod d_i}{\sum d_i} + \frac{1}{3} \sum_{(i,k) \in E} \widehat{l}_{ik} \left(\sum_j \widehat{l}_{ij} \widehat{l}_{jk} \right) + \frac{11}{6} \right] + \\ &\frac{n-1}{2} \left[\ln |\Omega_{ij}| - \frac{7\kappa \ln 2}{6} - \ln(2\pi e)^\kappa \right]. \end{aligned} \quad (42)$$

Here, we used the shorter notation $\widehat{s}_{ST}(G, r) \doteq s_{ST}(\widehat{\tau}(G, r))$. Since r is odd, after determining the lower bounds of the entropy $\mathcal{LB}(\mathcal{U})$ from $\widehat{\tau}(G, 3)$ for all actions, in order to provide performance guarantees for τ -bsp, we have to find an upper entropy bound for *only* the selected action $\widehat{\mathcal{U}} = \text{argmax}_{\mathcal{U}} \widehat{s}_{VN}(\mathcal{U})$, for example, by finding or bounding the Taylor remainder term $\varepsilon(\widehat{\mathcal{U}}, 3)$ according to equation (38). As we can see, the relation (42) enables finding $\widehat{s}_{ST}(G, 3)$ without doing expensive matrix factorizations which means that performance guarantees for τ -bsp can be established in a very efficient way.

Equation (42) reveals also how to improve normalization of $\widehat{s}_{VN}(G)$ for different candidate path lengths n which corresponds to its scaled third term, that is,

$$\widehat{s}'_{VN}(G) = \widehat{s}_{VN}(G) + \frac{\frac{n-1}{2} \left[\ln \frac{|\Omega_{vij}|}{(2\pi e)^\kappa} - \frac{7\kappa \ln 2}{6} \right]}{\left(\frac{7\kappa}{6} \right)}. \quad (43)$$

In that way, an action ranking of \widehat{s}'_{VN} follows the action ranking of $\widehat{s}_{ST}(G, 3)$ given that the node degree distribution corresponding to the second term in equation (42) does not change much and is already captured by $\widehat{s}_{VN}(G)$.

5.3.1. Keeping action consistency. In general, given two different candidate actions \mathcal{U}_1 and \mathcal{U}_2 , when can we say that ordering them by one topological signature is the same as ordering by another, that is, $s_{ST}(\mathcal{U}_1) < s_{ST}(\mathcal{U}_2) \Leftrightarrow \widehat{s}_{VN}(\mathcal{U}_1) < \widehat{s}_{VN}(\mathcal{U}_2)$? For any r and $G = G(\mathcal{U})$

$$\widehat{s}_{ST}(G, r) = \alpha_r \widehat{s}_{VN}(G) + \beta_r(D(G)) + \gamma_r(n),$$

where α_r , β_r , and γ_r are some functions that depend on the polynomial order r used for approximating $\tau(G)$ that follow from equation (40). For a given r , α_r is a constant, β_r is a function of graph's node degrees, and γ_r an affine function of the number of its nodes.

The difference of topological signatures \widehat{s}_{ST} of order r between graphs G_1 and G_2 with $n_1 = |V_1|$ and $n_2 = |V_2|$ can be written

$$\begin{aligned} \Delta \widehat{s}_{ST}(G_1, G_2, r) &= \widehat{s}_{ST}(G_1, r) - \widehat{s}_{ST}(G_2, r) \\ &= \alpha_r \Delta \widehat{s}_{VN}(G_1, G_2) + \Delta \beta_r(D(G_1), D(G_2)) + \Delta \gamma_r(n_1, n_2). \end{aligned} \quad (44)$$

Then, for any two graphs G_1 and G_2 from the control actions set, $\Delta \widehat{s}_{ST}(G_1, G_2, r) < 0$ if and only if

$$\Delta \widehat{s}_{\text{VN}}(G_1, G_2) < - \frac{\Delta \beta_r(D(G_1), D(G_2)) + \Delta \gamma_r(n_1, n_2)}{\alpha_r}. \quad (45)$$

This is another way of formulating the relation between topological signatures in terms of action consistency between them. If the normalization is applied to \widehat{s}_{VN} using (43), then $\Delta \gamma_r(n_1, n_2) = 0$.

6. Results

We evaluated our approach in Gazebo simulations of a single-robot active pose SLAM and in optimization of synthetic pose graphs. $t\text{-bsp}$ was compared to the standard BSP with the ML inference problem solved using GTSAM library (Dellaert 2012). We studied empirically different topological signatures and their correspondences to information-theoretic cost (6) together with a time required for decision-making.

In active pose SLAM simulations, the robot performs LIDAR-based 2D SLAM assuming perfect data association. A probabilistic roadmap (PRM) (Kavraki et al. 1996) is used to discretize the environment and generate the roadmap while k -diverse shortest path algorithm (Voss et al. 2015) is used to generate topologically diverse candidate paths over it. The action generation process uses a ground truth map to validate actions and avoid possible re-planning stages, since they do not contribute to the goal of BSP analysis, which are a result of incremental map building. Local path corrections due to localization errors or environment changes are handled by the low-level controller and obstacle avoidance algorithm. Of course, in the deployment of active SLAM, an inferred map will be used in generation of candidate paths using the same principles.

Topological properties of candidate actions are determined by the configuration of obstacles in the environment and by the parameters of the PRM algorithm (e.g., number of samples) and k -diverse shortest path algorithm (e.g., k , path length, paths' distance) which control the number and

diversity of candidate actions with respect to their geometrical properties and homology class. Topological graphs corresponding to these candidate actions depend on the robot's trajectory and sensor's characteristics and pre-processing (SLAM front-end). In particular, we assume that a link between certain trajectory samples will exist whenever the Euclidean distance between them is below a threshold δ and relative orientation below ϑ , which model the range and the field of view with high reliability of ICP method used for matching LIDAR scans. In all simulations when predicting future measurements, we assume that the robot can follow the candidate path perfectly and that trajectory pose samples with LIDAR scans are taken after each robot's relative displacement of 1 m and at the desired predefined goal poses.

6.1. Small-scale realistic simulation

First we validated our theoretical results in a small-scale scenario with the main goal to demonstrate properties of $t\text{-bsp}$ in elementary planning modes (exploration and exploitation) under different probabilistic models of relative pose measurements.

The robot was performing two planning sessions (S1 and S2) with both exploration (in S1) and exploitation trajectories (in S2) considered to show influence of different candidate path topology and geometry in $t\text{-bsp}$. Figure 8 shows the considered scenario in Gazebo and the generated candidate paths for the robot in two planning sessions S1 (Figure 8a) and S2 (Figure 8b). To demonstrate the effect of noise level on $t\text{-bsp}$, we simulated three cases. All relative pose measurements had standard deviation of orientation error σ_θ either 0.01, 0.035, or 0.085 rad while we kept the same standard deviation of the position error $\sigma_p = 0.1$ m. This corresponds to three different values of $\xi \in \{0.1, 0.35, 0.85\}$ for each planning session. The parameter δ was set to 2 m and ϑ to 2π radians, that is, without the constraint on relative orientation in topology prediction model.

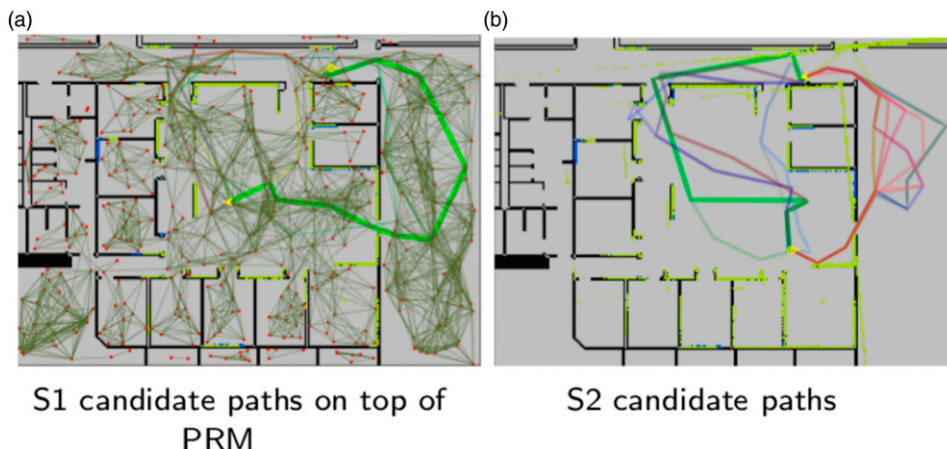


Figure 8. Gazebo scenario with robot's start pose marked with \blacktriangle and goal pose with \odot in each planning session.

During planning session S1, the robot is performing mainly exploration as the environment was completely unknown in the beginning and its first goal is set far in the unknown area so its future posterior beliefs have topologies that resemble tree graphs, with loop closing edges between poses nearby in time. In planning session S2, the robot was instructed to return to the previously mapped area causing larger diversity among candidate actions. Topologies of the most and least complex graph of each planning session are shown in Figure 9 together with their corresponding posterior belief. In both planning sessions t -bsp based on all proposed signatures was action consistent, that is, correctly identified the best action which can be seen from Figure 11g-11l. A maximum of a topological signature is indeed a minimum of the joint entropy. Notice that it does not imply necessarily that the pose uncertainty will be the lowest at the goal or that the shortest trajectory will be selected as is evident from S1's solution (Figure 9(a)-9(b)) since currently the entropy/uncertainty of the entire system is considered. This opens one other relevant research direction where uncertainty of a subset of state variables might be considered and its relation to topological information. Among exploitation trajectories, t -bsp prefers the one with larger loop closings leading to highest information gain (see Figure 9(c)-9(d)).

Our results in a realistic simulation show the promise of the proposed approach for efficiently solving BSP. All topological signatures of the posterior factor graphs are strongly correlated with the information-theoretic cost (Figure 11(g)-11(l)); yet, t -bsp based on an exact von

Neumann entropy (s_{VN}) and the number of spanning trees (s_{ST}) of a graph outperforms standard BSP by an order of magnitude in terms of time complexity as shown in Figure 10. This is due to operation in a topological (less dimensional) space instead of a metric state space. The dimension of a graph Laplacian matrix in this problem is $\kappa = 3$ times smaller than the dimension of an information matrix because every 2D pose in the state vector corresponds to one node in the graph. The relative speed of s_{VN} and s_{ST} is similar as both of them require determining a graph spectrum of the associated Laplacian or its cofactor. On the other hand, the advantage of using \hat{s}_{VN} over s_{VN} and s_{ST} is that it is much faster, $O(|E|)$, with possible incremental calculation that depends only on the node degrees which we investigate in the large-scale simulation in Subsection 6.4. It effectively enables real-time performance of t -bsp given that typically at each measurement update a limited number of factors is added to the graph. Also visible from Figure 10 is that t -bsp error bounds calculation adds a small, especially in the case with Hadamard bound, additional cost.

While t -bsp was action consistent in all sessions for all proposed topological signatures ($\epsilon(J, s) = 0$), action ordering consistency ($\gamma_f = 0$) was kept only for s_{ST} in all cases, while only in the first planning session for s_{VN} and \hat{s}_{VN} . $\gamma_f \neq 0$ can happen when topologies among actions are very similar and then other factors, for example, noise or geometry, determine the solution. However, γ_f was still small. As we already stated, the determined bounds of the entropy can always be used to provide a globally optimal solution but at the cost of evaluating the objective function

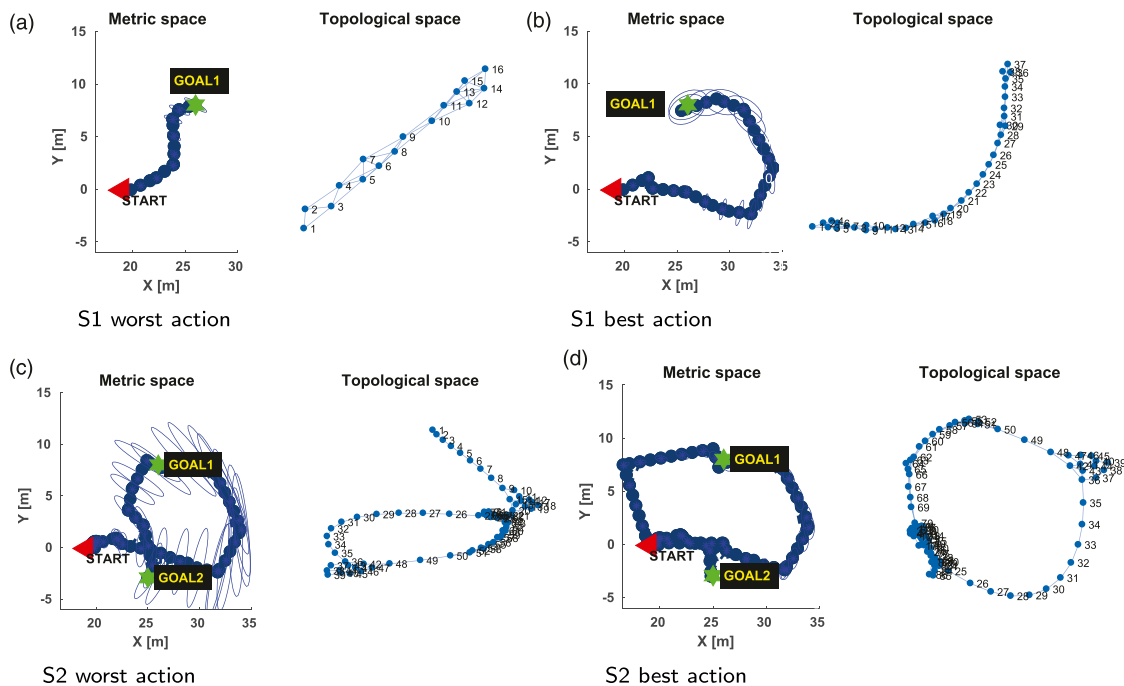


Figure 9. Trajectory uncertainty after optimization and topologies of the worst/best actions calculated by topological BSP.

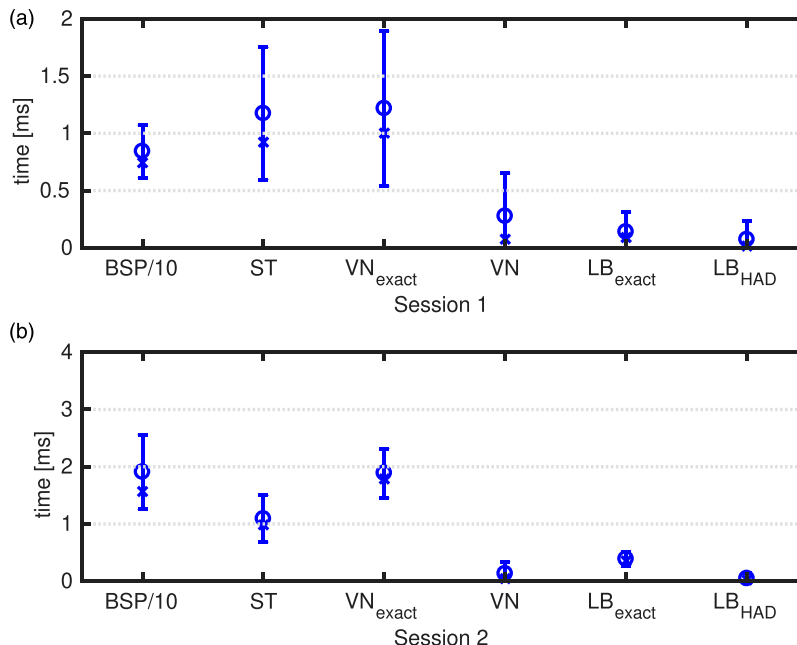


Figure 10. Computation time per candidate action in each planning session (circle represents the mean, “x” the median and line interval one σ confidence region) for standard BSP, τ -bsp by signatures based on the number of spanning trees (ST), the von Neumann entropy (VN exact and approx.) and for bounds determination, exact (equation (25)) and Hadamard (equation (29)). Notice that computation time of the standard BSP approach has been scaled by 10.

of actions inside them, that is, the ones whose lower bound is below an upper bound of the selected action. From Figures 11(a)–11(f) we can see that a negative topological signature s_{ST} is a very good approximation of the estimated posterior entropy even with large orientation noise, that is, large ζ values. On the other hand, its determined lower bound is very sensitive to it. For larger ζ values, bounds are less tight and therefore our performance guarantees are more conservative. In practice, however, more important is that the upper bound is close to the real entropy since we make decisions based on it. We can also notice that for higher ζ , Hadamard approximation of the lower bound is getting better, meaning that we practically have very small computational cost in determining the bounds accurately.

6.2. Real-world experiment

The experiment was performed on a real robot using a fixed measurement noise model in inference and planning, considering a similar scenario as in the simulation from Section 6.1. The results are presented in Extension 1 as a proof of concept of our method.

6.3. Synthetic simulation

Apart from the SLAM setting, next we consider general pose graphs randomly created and optimized by the Maximum Likelihood Estimator (MLE) and their relation to

the topological signatures developed in this paper. We show that good action consistency is still kept even when comparing significantly varying pose graphs. We generated pose graphs with varying number of nodes n and edges m of which $n - 1$ correspond to odometry and the rest to $LC = m - n + 1$ loop closing factors. We choose $LC = \Delta\bar{d} n/2$, causing average node degree of a graph increasing in steps of $\Delta\bar{d} \in \{0.1, 0.5, 0.75, 1, 2\}$ from a given odometry graph of size n (whose average node degree is approx. 2). For each pair (n, m) we create a group of 10 random topologies, calculate their topological signatures s_{ST} and \hat{s}_{VN} , and measure the average time for their calculation, t_{ST} and t_{VN} , respectively. The results of this simulation are given in Figure 12. We can see that correlation with the posterior entropy of FIM of MLE is still high for s_{ST} on the whole range of topologies while \hat{s}_{VN} follows the overall trend and is high for graphs with the same n , but breaks the action consistency with jumps of n (Figure 12(b), green curve). This suggests that a normalization of \hat{s}_{VN} should be applied according to its relation to the objective according to the Definition 1. As explained in Section 5.3, after applying correction term given by the equation (43), we can see that actions consistency improves significantly (Figure 12(b), black curve).

Although less accurate, the plot in Figure 12(a) demonstrates why \hat{s}_{VN} might still be interesting to use for fast decision-making in some situations. Clearly, the gain in speed is significant in high-dimensional BSP problems.

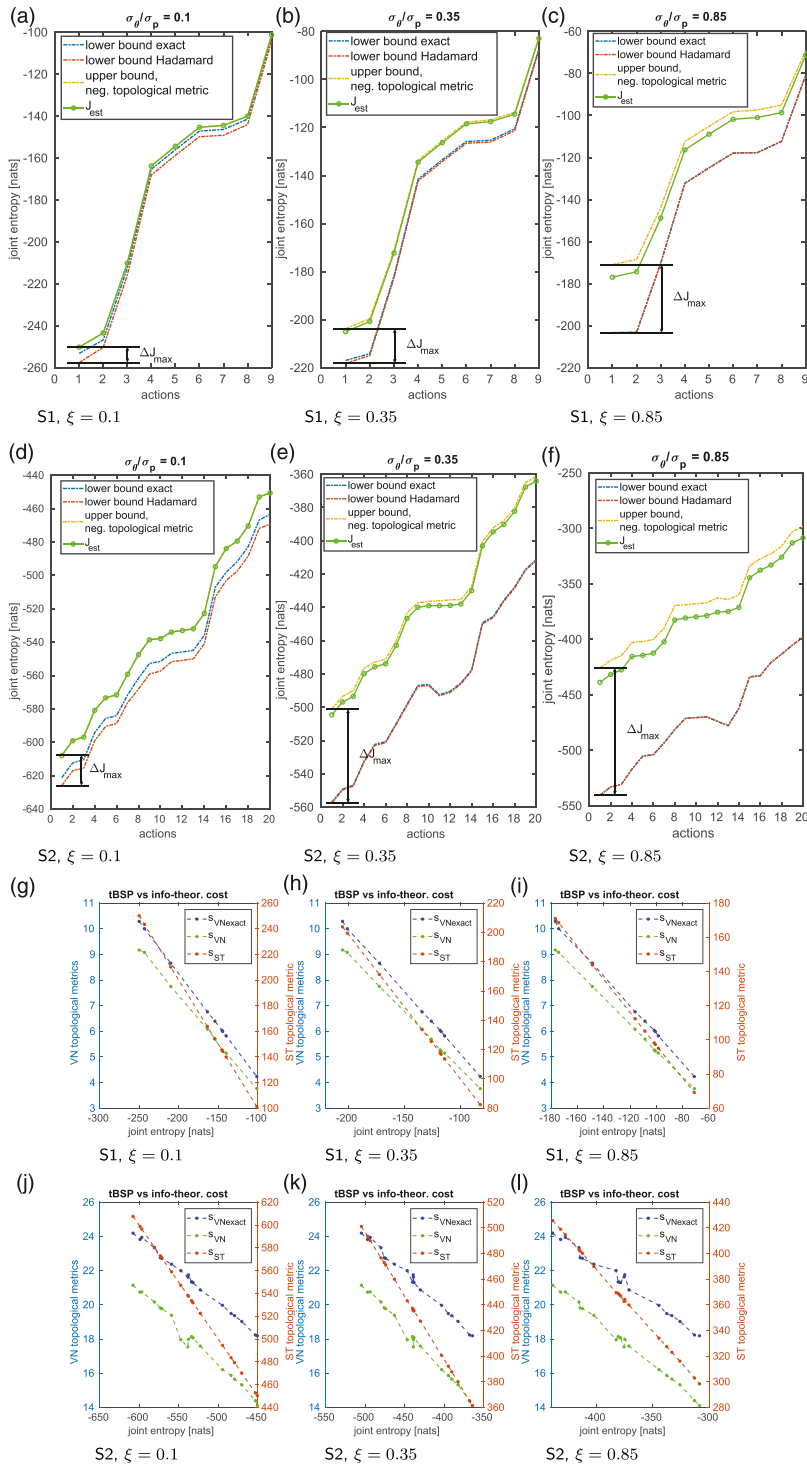


Figure 11. Estimated entropy of posterior belief and its bounds determined by t -bsp (*top*) and correlation of t -bsp and standard BSP (*bottom*) for three different values of ξ in planning sessions S1 and S2. Notice different y-axis on the correlation plots since only action ordering consistency is important in BSP.

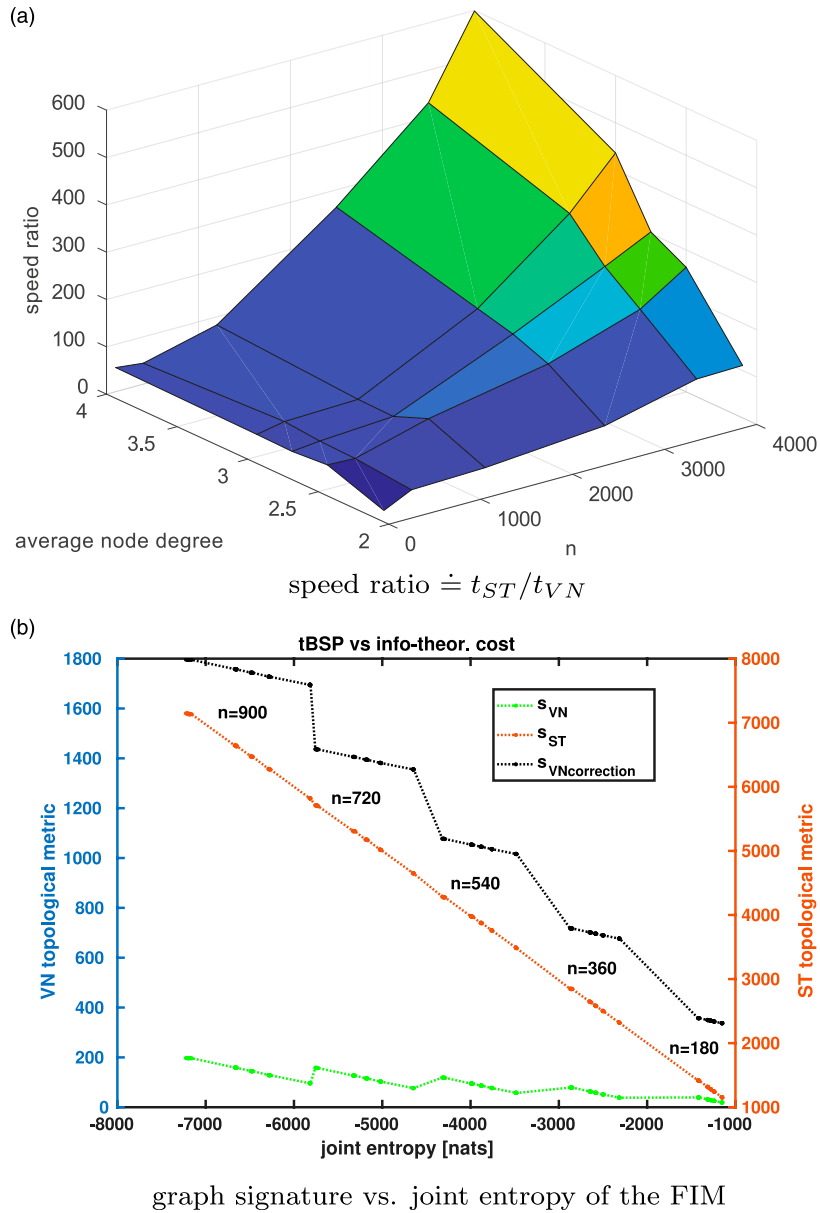


Figure 12. Performance on simulated random graphs.




Figure 13. Gazebo simulation environment with Pioneer-3AT ground robot marked with white bounding box. Grid cell size is 1 m x 1 m.

6.4. Large-scale realistic simulation

A large-scale scenario tests t -bsp performance in a complex environment depicted in Figure 13 with larger number of planning sessions. Its main goal is to investigate the proposed approach in the settings closer to active SLAM application domains. It serves also as a test bed for comparing the batch and incremental BSP algorithms.

For the second purpose, we design two experiments which mainly differ in the planning horizon length L considered:

- (a) a goal-to-goal planning, and
- (b) Model Predictive Control (MPC) style planning.




Figure 14. Run-time comparison of batch BSP algorithms. Standard BSP is based on Levenberg-Marquardt optimization and explicit calculation of objective function (6) by doing Cholesky factorization of information matrix. t -bsp based on s_{ST} uses the same method of Cholesky factorization to calculate the number of spanning trees of a topological graph.

The experiment a) was primarily designed to test the execution time of the batch BSP algorithms whereas the experiment b) for their incremental counterparts. Since action consistency does not depend on the particular operation mode, it is enough to show it only once. In the first experiment, the robot is instructed to visit ten predefined goals in a given order and to plan with a horizon sufficient to connect consecutive goals. The second considers the same set of goals with the difference that re-planning is done also on several intermediate waypoints on the selected route between goals choosing the next waypoint as a sub-goal. We call this MPC mode, because the robot considers changing its immediate control action on a way. In this simulation, however, for ensuring similar conditions in both experiments, the robot does not take any detours from the original route. This mode would make sense in a dynamic environment or if the inferred map is used by the action generator method due to partial map availability and possibility that a better global path exists.

Consequentially, in goal-to-goal planning, the robot executes 10 planning sessions and in each, many measurements are added during the motion and prediction steps. Contrary to that, MPC mode has more planning sessions with fewer new measurements acquired between them. In both, the topology prediction parameter δ was set to 2 m and ϑ to 1 rad. Standard deviations of relative orientation and position measurement noise were kept constant at $\sigma_p = 0.1$ m and $\sigma_\theta = 0.01$ rad.

6.4.1. Batch BSP algorithms. The standard batch BSP inference was solved by Levenberg-Marquardt algorithm from GTSAM library with the information-theoretic cost (6) calculated exactly using Cholesky factorization of the information matrix of the posterior belief required for determining the determinant. The same Cholesky method was used to factorize the reduced Laplacian matrix of the topological posterior graph needed for t -bsp based on the number of spanning trees topological signature s_{ST} . Also, in every planning session, the approximated von Neumann graph

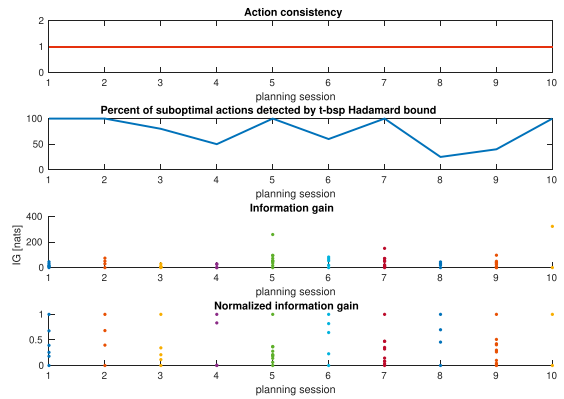


Figure 15. Action consistency of t -bsp based on approx. von Neumann entropy \hat{s}_{VN} and optimality guarantees in relation to information gain of candidate actions (each dot of the same color represents one candidate action in a single planning session). Optimal action is selected by t -bsp in all planning sessions, but its optimality can be guaranteed in sessions 1, 2, 5, 7, and 10 where one action is clearly distinctive in information gain.

entropy \hat{s}_{VN} was calculated from scratch using the expression (13) and then normalized using (43).

As expected, the run-time of the t -bsp based on s_{ST} is by a constant multiplicative factor lower than the standard BSP (Figure 14) since the sparsity pattern of both information matrix and reduced Laplacian matrix is the same, whereas the latter's dimension is κ times reduced, κ being the dimension of the robot's pose. On the other hand, the proposed t -bsp based on \hat{s}_{VN} shows almost linear-time complexity on average planning time per candidate action and achieves by three orders of magnitude higher speed by the last planning session.

Importantly, an optimal action was selected in all planning sessions by the proposed t -bsp based on \hat{s}_{VN} as can be seen from Figure 15. Its only limitation is that optimality cannot always be guaranteed using topological entropy bounds. Information gain of candidate actions, in those planning sessions where it is the case, is indeed very

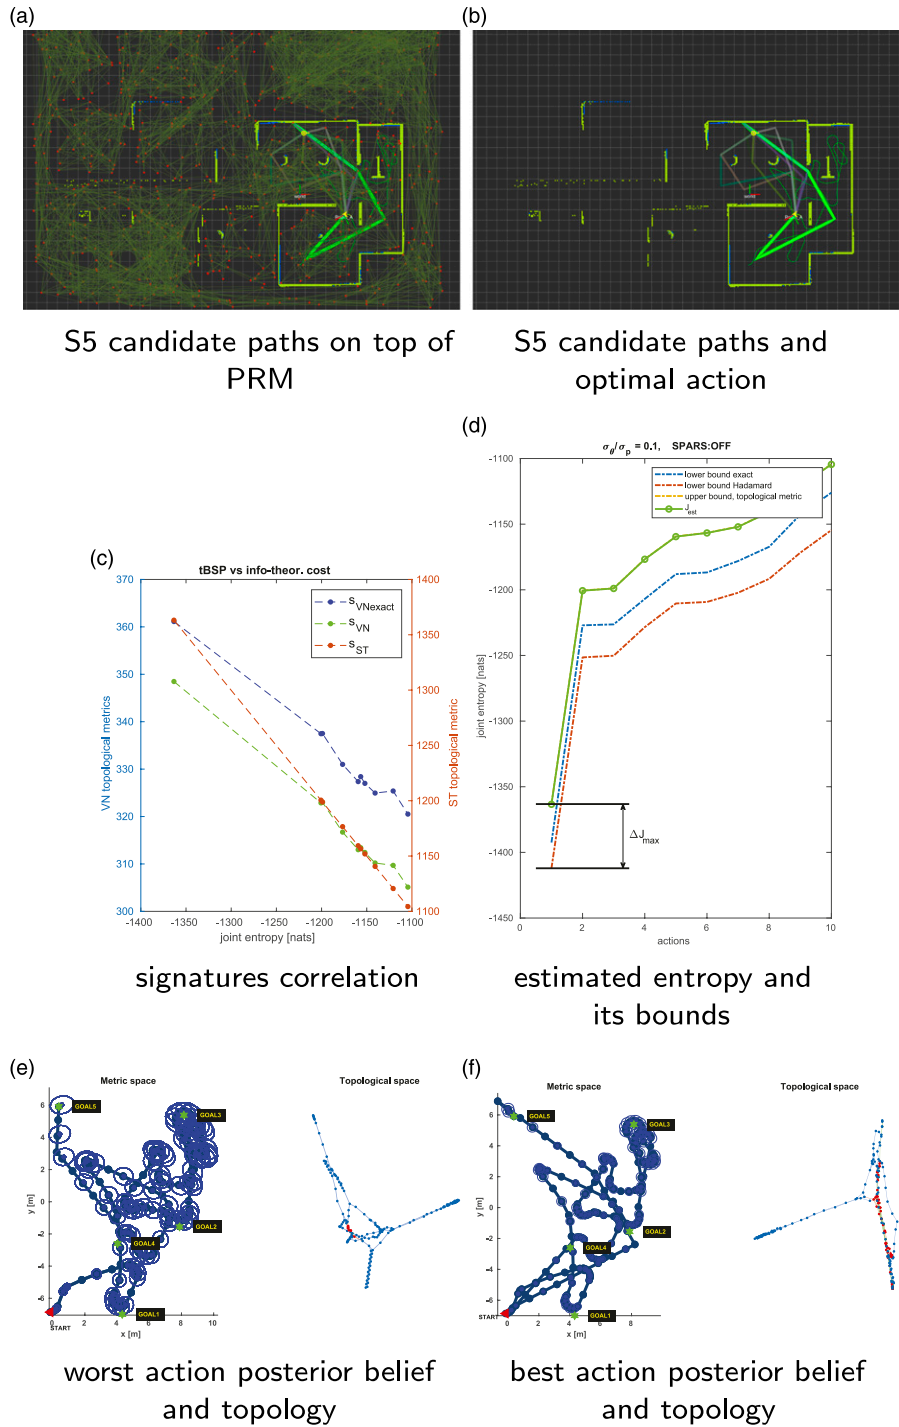


Figure 16. Planning session 5.

similar and then non-topological factors determine the solution.

Let us have a closer look into two extreme cases: one in which the optimal action can be guaranteed and the other with a small percentage of suboptimal actions detected by t -bsp. The representative of the first case is the planning session 5 (S5) with relevant plots shown in Figure 16, and of the second, planning session 8 (S8) analyzed in Figure 17. The whole simulation and

analysis of the rest of the planning sessions can be seen in Extension 2.

S5’s candidate paths together with the inferred map and the past robot’s trajectory are shown in Figures 16(a)–16(b). It can be seen that there exists a single candidate path (marked with light green color) which leads to the previously mapped area and from the information-theoretic perspective is optimal as it significantly reduces uncertainty of the robot’s trajectory as shown in Figure 16(f). For

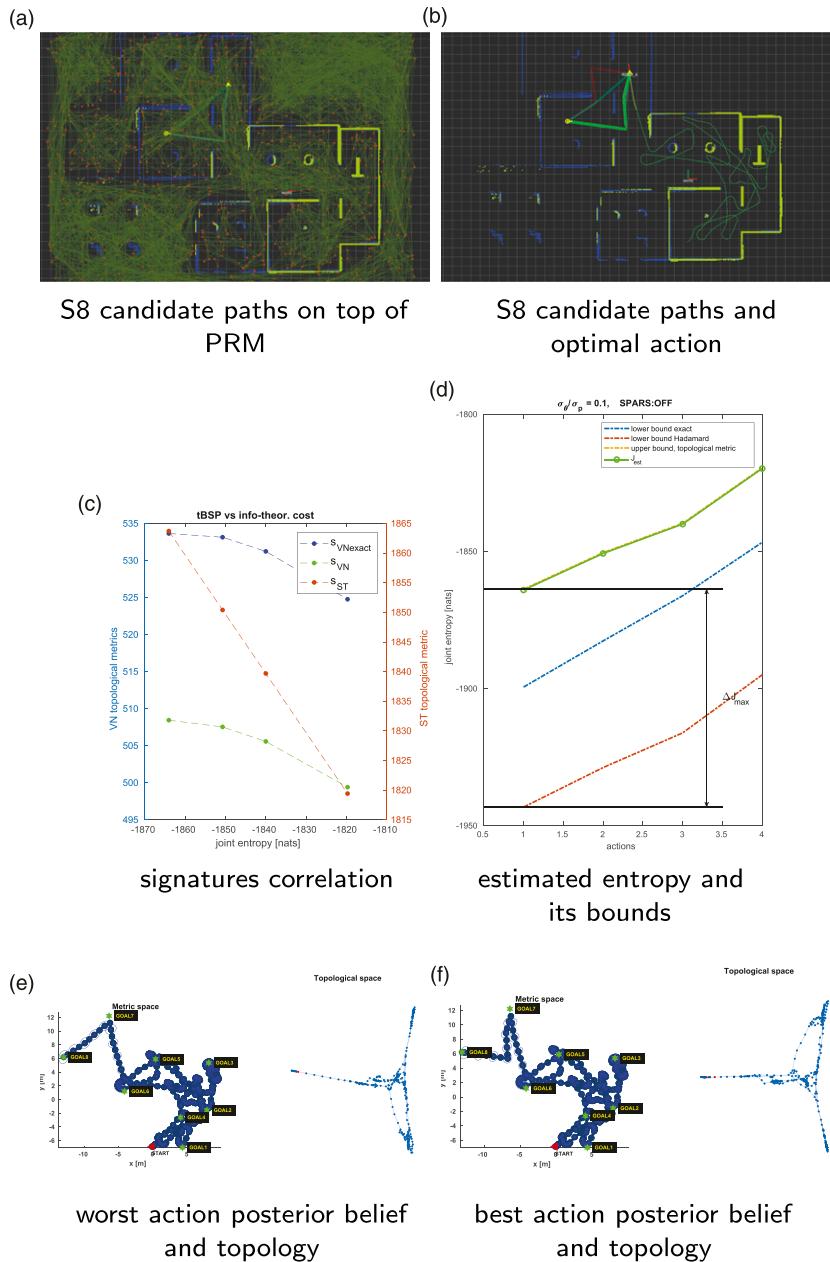


Figure 17. Planning session 8.

illustration, we also show the uncertainty of the worst action in that session in Figure 16(e) where ellipses represent the marginal pose’s uncertainty in both plots. In this case, topology solely is enough to disqualify other actions (Figure 16(d)).

S8’s candidate paths together with the inferred map and the past robot’s trajectory are shown in Figures 17(a)–17(b). First we notice that in that session there was a smaller number of candidate paths (only 4) which belong to two homology classes, either going left or right around the wall obstacle. Three of them are quite similar in the geometrical shape and all of them passing through already mapped areas. Therefore, it makes sense that the expected accuracy (path entropy) obtained by following either path will roughly be the same as can be seen in

Figures 17(e)–17(f) for the worst and best actions from that set. Even in that unfavorable situation for τ -bsp, it still managed to find an optimal action (Figure 17(c)) where maximum of \hat{s}_{VN} corresponds to the minimum of the joint entropy, but not to guarantee its optimality. Two other actions need to be inspected further for that, which is visible from Figure 17(d). We can use an upper topological bound which gives an uncertainty margin for deciding that. We argue that in most practical applications this difference in information gain will be insignificant and τ -bsp is good for coarse-level planning especially, when a loss of precision up to the model approximation gap is higher than a difference between actions’ objective values.

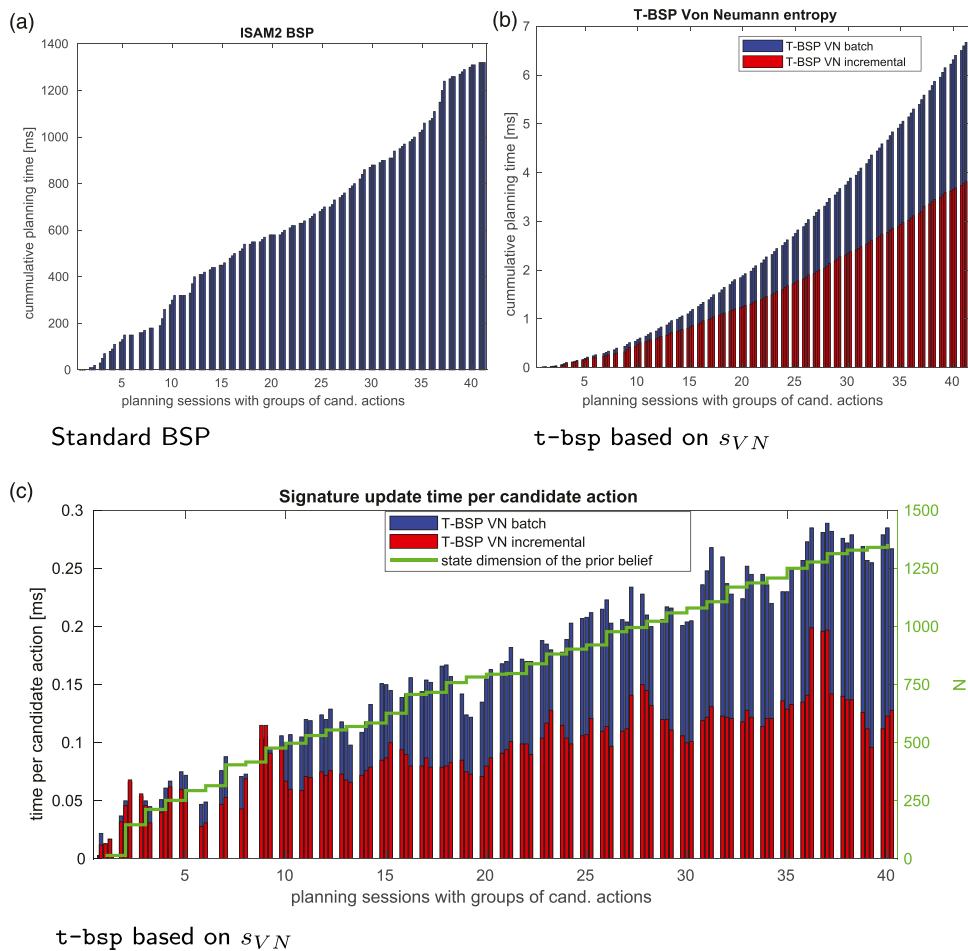


Figure 18. Run-time comparison of incremental BSP algorithms in MPC mode. Standard BSP is based on ISAM2 algorithm for belief propagation and information matrix factorization (Figure 18(a)). Topological BSP is based on approximated von Neumann entropy \hat{s}_{VN} signature (Figure 18(b)). Topological signature update per candidate action demonstrates almost constant-time complexity of incremental t -bsp and near linear-time of its batch equivalent (Figure 18(c)). The plot in Figure 18(b) shows cumulative planning time, that is, summation of planning times over the sessions and candidate actions from Figure 18(c).

6.4.2. Incremental BSP algorithms. In the experiment b) we compare running times of incremental BSP algorithms. Standard incremental BSP was realized by employing state-of-the-art incremental smoothing and mapping algorithm ISAM2 from GTSAM library. It updates Cholesky factorization of an information matrix incrementally using operations on a graphical model called Bayes tree (Kaess et al. 2010). For the purpose of this analysis, we measure the run-time of only updating the Bayes tree of posterior beliefs corresponding to the candidate actions, since after obtaining this form calculating the information-theoretic cost is a simple operation of calculating a determinant of a triangular matrix. Topological BSP is based on the approximated von Neumann entropy \hat{s}_{VN} signature calculated incrementally using the formula (16). The cumulative planning time across all planning sessions in MPC mode is shown in Figure 18(a) for standard BSP and in Figure 18(b) for the t -bsp based on \hat{s}_{VN} . Again, t -bsp outperforms standard BSP (by two orders of magnitude). For reference, we also show non-incremental t -bsp based on \hat{s}_{VN} which is also much faster than ISAM2 BSP. The graph in Figure 18c presents

planning time per candidate action with the change of state dimension over time. As the theoretical analysis suggests, time complexity is near $O(1)$ for the incremental t -bsp and near $O(m)$ for its batch equivalent, where m is the number of edges in the topological graph, which in this case increases almost linearly with the state dimension.

7. Conclusions

This work provides theoretical foundations for information-theoretic topological belief space planning (t -bsp), a novel paradigm for highly efficient decision-making under uncertainty. The main idea consists of transforming the original optimization problem to a topological space where a solution is easier to find and then establish the error bounds with respect to the optimal solution. Doing so, we avoid belief propagation and explicitly calculating an objective function which offers significant speed-up to decision-making. Very importantly, even when using such approximations, an optimal BSP solution can still be obtained in cases when two optimization problems are action

consistent, a concept introduced in this work for objective function approximations. In the general case, one can resort to t -bsp to drastically reduce computational cost while carefully monitoring a conservative estimate on the sacrifice in performance, that would be provided by the error bounds, or to find globally optimal solution by evaluating generally a much smaller number of candidate actions.

Topological signatures are defined on topologies of factor graphs that correspond to future posterior beliefs associated to candidate actions and their design depends on the objective function. In this work, we consider information-theoretic objective functions and pairwise factors whose topologies can be represented with simple graphs. Two topological signatures are proposed: normalized number of spanning trees of a graph and approximation of its von Neumann entropy by a function of node degrees. The first shows better action consistency, but the special focus in this paper is given to the second signature as it enables real-time performance of BSP while still offering good action consistency. It can be used in either batch or incremental mode. For both signatures, we derived error bounds of t -bsp that depend on topological parameters, measurement noise, and prior state estimate under maximum likelihood estimation framework and the bounds can be calculated with a small additional computational cost.

We presented an evaluation of our approach in realistic active SLAM simulations and in random pose-graph optimizations. Results show t -bsp was optimal in all planning sessions, but optimality could not always had been guaranteed depending on the measurement noise characteristics and diversity of candidate actions (topological and geometrical). However, in those sessions candidate paths carry very similar information gain and which action will be chosen is often irrelevant in practice, but when it is important, t -bsp can be used as a pre-processing tool to reduce the action search space. Also, decision-making is done by the upper entropy bound which is much more resilient to noise than the lower bound needed for the guarantees. Time performance improved drastically compared to state-of-the-art. Incremental t -bsp based on von Neumann entropy exhibits $O(1)$ asymptotic complexity in active SLAM MPC-style planning. Also, update of the topological signature is very intuitive and its complexity is independent on which variables loop closure affects.

This work opens several research directions which might be interesting in the context of simplifications approaches for BSP. One is to investigate if a wider class of objective functions to which t -bsp is applicable exists, that is, where the solution is strongly characterized by the belief topology. Also, we might consider relaxing the assumptions on Gaussian beliefs, isotropic measurements noise or pairwise relations in measurement model used in this work. A first work in that direction is (Shienman et al. 2021) where uncertainty minimization of a predefined subset (focused set) of variables and uncorrelated differently weighted measurements are addressed, building upon the t -bsp approach for the unfocused case introduced here. Improving

tightness of the error bounds would be of great interest for guarantees in fine-level planning, when control actions are topologically similar. Finally, topology prediction models could be further improved to increase prediction accuracy, for example, by taking into account probabilities of data associations. Overall, the presented approach contains both fundamental and practical contributions to belief space planning and we expect it will motivate further research and find diverse applications.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was partially supported by the Israel Science Foundation (grant No. 371/20).

ORCID iD

Andrej Kitanov  <https://orcid.org/0000-0002-4183-2896>

Supplemental Material

Supplemental material for this article is available online.

Note

1. Since the state x_0 is considered deterministic, here we estimate the rest of the variables $X_k = \{x_1, x_2, \dots, x_k\}$ and $I(X_k)$ denotes their joint information matrix from now on.

References

- Agha-Mohammadi AA, Chakravorty S and Amato NM (2014) FIRM: sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *Intl J of Robotics Research* 33(2): 268–304.
- Angeli A, Doncieux S, Meyer J, et al. (2009) Visual topological SLAM and global localization. In IEEE Intl. Conference On Robotics and Automation (ICRA), Kobe, Japan, 2009. pp. 4300–4305.
- Archdeacon D (1996) Topological graph theory: a survey. In *Congressus Numerantium*. Utilitas Mathematica Pub. Incorporated, Volume 115, pp. 5–54.
- Atanov N, Ny JL, Daniilidis K, et al. (2015) Decentralized active information acquisition: theory and application to multi-robot SLAM. In IEEE Intl Conf On Robotics and Automation (ICRA), Seattle, WA, USA, 2015, pp. 4775–4782.
- Bajcsy R (1988) Active perception. *Proceedings of the IEEE* 76(8): 996–1005.
- Barooh P and Hespanha JP (2007) Estimation on graphs from relative measurements. *IEEE Control Systems Magazine* 27(4): 57–74.

- Bhattacharya S, Ghrist R and Kumar V (2015) Persistent homology for path planning in uncertain environments. *IEEE Trans Robotics* 31(3): 578–590.
- Bian F, Kempe D and Govindan R (2006) Utility based sensor selection. In Proceedings of the 5th Intl Conference On Information Processing in Sensor Networks, Nashville, TN, USA, 2006, pp. 11–18.
- Biggs N (1993) *Algebraic Graph Theory*. 2nd edition. Cambridge University Press.
- Blöchliger F, Fehr M, Dymczyk M, et al. (2018) Topomap: topological mapping and navigation based on visual SLAM maps. In IEEE Intl Conf On Robotics and Automation (ICRA), Brisbane, QLD, Australia, 2018, pp. 3818–3825.
- Brouwer AE and Haemers WH (2011) *Spectra of Graphs*. Springer.
- Carlevaris-Bianco N, Kaess M and Eustice RM (2014) Generic node removal for factor-graph SLAM. *IEEE Trans Robotics* 30(6): 1371–1385.
- Carrillo H, Reid I and Castellanos J (2012) On the comparison of uncertainty criteria for active SLAM. In IEEE Intl. Conf. On Robotics and Automation (ICRA), Saint Paul, MN, USA, 2012, pp. 2080–2087.
- Chen Y, Huang S, Fitch R, et al. (2018) Efficient active SLAM based on submap joining, graph topology and convex optimization. In IEEE Intl. Conf. On Robotics and Automation (ICRA), Brisbane, QLD, Australia, 2018, pp. 5159–5166.
- Chen Y, Huang S, Fitch R, et al. (2019) On-line 3D active pose-graph SLAM based on key poses using graph topology and sub-maps. In IEEE Intl Conf On Robotics and Automation (ICRA), Montreal, QC, Canada, 2019, pp. 169–175.
- Chen Y, Huang S, Zhao L, et al. (2021) Cramér-rao bounds and optimal design metrics for pose-graph SLAM. *IEEE Transactions on Robotics* 37: 627–641. DOI: [10.1109/TRO.2020.3001718](https://doi.org/10.1109/TRO.2020.3001718).
- Choset H and Nagatani K (2001) Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Trans Robot Automat* 17(2): 125–137.
- Chung FR (1997) *Spectral Graph Theory*. American Mathematical Soc.
- Davidson JC and Hutchinson SA (2009) A sampling hyperbelief optimization technique for stochastic systems. In *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*. Springer, pp. 217–231.
- Dellaert F (2012) Factor graphs and GTSAM: a hands-on introduction. In *Technical Report GT-RIM-CP&R-2012-002*. Georgia Institute of Technology.
- Dellaert F and Kaess M (2017) Factor graphs for robot perception. *Foundations and Trends in Robotics* 6(1–2): 1–139.
- Du J, Carlone L, Ng MK, et al. (2011) A comparative study on active SLAM and autonomous exploration with particle filters. In Proc. of the IEEE/ASME Int. Conference on Advanced Intelligent Mechatronics, Budapest, Hungary, 03–07 July 2011.
- Elimelech K and Indelman V (2017a) Consistent sparsification for efficient decision making under uncertainty in high dimensional state spaces. In IEEE Intl. Conference On Robotics and Automation (ICRA), Singapore, 29 May 2017 - 03 June 2017, pp. 3786–3791.
- Elimelech K and Indelman V (2017b) Fast action elimination for efficient decision making and belief space planning using bounded approximations. In Proc Of the Intl Symp of Robotics Research (ISRR). Puerto Varas, Chile, 2017.
- Elimelech K and Indelman V (2017c) Scalable sparsification for efficient decision making under uncertainty in high dimensional state spaces. In IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS), Vancouver, BC, 24–28 September 2017, pp. 5668–5673.
- Folkesson J and Christensen H (2004) Graphical SLAM – a self-correcting map. In IEEE Intl. Conf. on Robotics and Automation ICRA), New Orleans, LA, 26 April 2004 - 01 May 2004, pp. 383–390.
- Frese U (2006) Treemap: an $O(\log n)$ algorithm for indoor simultaneous localization and mapping. *Autonomous Robots* 21(2): 103–122.
- Frey B, Kschischang F, Loeliger HA, et al. (1997) *Factor graphs and algorithms*. In Proc 35th Allerton Conf Communications, Control, and Computing, University of Illinois, 1997, pp. 666–680
- Gross JL and Tucker TW (1987) *Topological Graph Theory*. New York, NY: Wiley.
- Han L, Escolano F, Hancock ER, et al. (2012) Graph characterizations from von Neumann entropy. *Pattern Recognition Letters* 33(15): 1958–1967.
- Horn RA and Johnson CR (2013) *Matrix Analysis*. 2nd edition. Cambridge University Press.
- Hovland GE and McCarragher BJ (1997) Dynamic sensor selection for robotic systems. *IEEE Intl. Conf. on Robotics and Automation (ICRA)* Leuven, Belgium, May 16–20, 1998.
- Huang S, Kwok N, Dissanayake G, et al. (2005) Multi-step look-ahead trajectory planning in SLAM: possibility and necessity. In IEEE Intl. Conf. On Robotics and Automation (ICRA), Barcelona, Spain, 2005, pp. 1091–1096.
- Indelman V (2015) Towards information-theoretic decision making in a conservative information space. In *American Control Conference*, Chicago, IL, USA, 2015, pp. 2420–2426.
- Indelman V (2016) No correlations involved: decision making under uncertainty in a conservative sparse information space. *IEEE Robotics and Automation Letters* 1(1): 407–414.
- Indelman V, Carlone L and Dellaert F (2015) Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl J of Robotics Research* 34(7): 849–882.
- Joshi S and Boyd S (2009) Sensor selection via convex optimization. *IEEE Transactions on Signal Processing* 57(2): 451–462.
- Kaess M, Ranganathan A and Dellaert F (2008) iSAM: incremental smoothing and mapping. *IEEE Trans Robotics* 24(6): 1365–1378.
- Kaess M, Ila V, Roberts R, et al. (2010) The Bayes tree: an algorithmic foundation for probabilistic robot mapping. In *Intl Workshop on the Algorithmic Foundations of Robotics*. Springer.

- Kaess M, Johannsson H, Roberts R, et al. (2012) iSAM2: incremental smoothing and mapping using the Bayes tree. *Intl J of Robotics Research* 31(2): 217–236.
- Kavraki L, Svestka P, Latombe JC, et al. (1996) Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot Automat* 12(4): 566–580.
- Khosoussi K, Huang S and Dissanayake G (2015) Good, bad and ugly graphs for SLAM. In *Robotics: Science and Systems (RSS). Workshop on the Problem of Mobile Sensors*. Springer.
- Khosoussi K, Giamou M, Sukhatme GS, et al. (2019) Reliable graph topologies for SLAM. *Intl J of Robotics Research* 38(2–3): 260–298.
- Kim A and Eustice RM (2014) Active visual SLAM for robotic area coverage: theory and experiment. *Intl J of Robotics Research* 34(4–5): 457–475.
- Kim S, Bhattacharya S, Ghrist R, et al. (2013) Topological exploration of unknown and partially known environments. In *IEEE/RSJ Intl Conf On Intelligent Robots and Systems (IROS)*, Tokyo, Japan, 03-07 November 2013, pp. 3851–3858.
- Kitanov A and Indelman V (2018) Topological multi-robot belief space planning in unknown environments, In *IEEE Intl Conf On Robotics and Automation (ICRA)*, Brisbane, QLD, 21-25 May 2018, pp. 5726–5732.
- Koller D and Friedman N (2009) *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press.
- Kopitkov D and Indelman V (2017) No belief propagation required: belief space planning in high-dimensional state spaces via factor graphs, matrix determinant lemma and reuse of calculation. *Intl J of Robotics Research* 36(10): 1088–1130.
- Kopitkov D and Indelman V (2019) General purpose incremental covariance update and efficient belief space planning via factor-graph propagation action tree. *Intl J of Robotics Research* 38(14): 1644–1673.
- Kretschmar H and Stachniss C (2012) Information-theoretic compression of pose graphs for laser-based SLAM. *Intl J of Robotics Research* 31(11): 1219–1230.
- Kschischang F, Frey B and Loeliger HA (2001) Factor graphs and the sum-product algorithm. *IEEE Trans Inform Theory* 47(2): 498–519.
- Kurniawati H, Hsu D and Lee WS (2008) SARSOP: efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems (RSS)*. Singapore, 2008.
- Lauritzen SL and Spiegelhalter DJ (1988) Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B (Methodological)* 50(2): 157–224.
- Mohar B (1991) The Laplacian spectrum of graphs. In *Graph Theory, Combinatorics, and Applications*. Wiley, pp. 871–898.
- Mowshowitz A and Dehmer M (2012) Entropy and the complexity of graphs revisited. *Entropy* 14(3): 559–570.
- Papadimitriou C and Tsitsiklis J (1987) The complexity of Markov decision processes. *Mathematics of operations research* 12(3): 441–450.
- Paskin M (2003) Thin junction tree filters for simultaneous localization and mapping. In *Intl Joint Conf On AI (IJCAI)*. San Francisco, CA, USA, August 2003, pp. 1157–1164.
- Passerini F and Severini S (2009) Quantifying complexity in networks: the von Neumann entropy. *International Journal of Agent Technologies and Systems (IJATS)* 1(4): 58–67.
- Paul R and Newman P (2010) Fab-map 3d: topological mapping with spatial and visual appearance. In *IEEE Intl Conf On Robotics and Automation (ICRA)*, Anchorage, AK, USA, 2010, pp. 2649–2656.
- Petz D (2001) Entropy, von Neumann and the von Neumann entropy. In *John Von Neumann and the Foundations of Quantum Physics*. Springer, pp. 83–96.
- Pineau J, Gordon GJ and Thrun S (2006) Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research* 27: 335–380.
- Placed JA, Strader J, Carrillo H, et al. (2023) A survey on active simultaneous localization and mapping: state of the art and new frontiers. *IEEE Trans Robotics* 39(3): 1686–1705.
- Platt R, Tedrake R, Kaelbling L, et al. (2010) Belief space planning assuming maximum likelihood observations. In *Robotics: Science and Systems (RSS)*. Zaragoza, Spain: MIT, pp. 587–593.
- Porta JM, Vlassis N, Spaan MT, et al. (2006) Point-based value iteration for continuous POMDPs. *Journal of Machine Learning Research* 7: 2329–2367.
- Prentice S and Roy N (2009) The belief roadmap: efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research* 28(11–12): 1448–1465.
- Ranganathan A and Dellaert F (2011) Online probabilistic topological mapping. *Intl J of Robotics Research* 30(6): 755–771.
- Regev T and Indelman V (2017) Decentralized multi-robot belief space planning in unknown environments via efficient re-evaluation of impacted paths. In *Autonomous Robots Special Issue on Online Decision Making in Multi-Robot Coordination*. Springer.
- Shienman M, Kitanov A and Indelman V (2021) FT-BSP: focused topological belief space planning. *IEEE Robotics and Automation Letters* 6(3): 4744–4751.
- Sorenson HW (1980) *Parameter Estimation: Principles and Problems*. M. Dekker, Volume 9.
- Stachniss C, Haehnel D and Burgard W (2004) Exploration with active loop-closing for FastSLAM. In *IEEE/RSJ Intl Conf on Intelligent Robots and Systems IROS*, Sendai, Japan, 28 September 2004 - 02 October 2004, pp. 1505–1510.
- Sucar LE (2015) *Probabilistic Graphical Models Principles and Applications*. London: Springer.
- Thrun S and Leonard JJ (2008) Simultaneous localization and mapping. In *Springer Handbook of Robotics*. Springer Verlag, pp. 871–889.
- Valencia R, Miró JV, Dissanayake G, et al. (2012) Active pose SLAM. In: *IEEE/RSJ Intl Conf On Intelligent Robots and Systems (IROS)*. Vilamoura-Algarve, Portugal, 2012, pp. 1885-1891.
- Van Den Berg J, Abbeel P and Goldberg K (2011) LQG-MP: optimized path planning for robots with motion uncertainty and imperfect state information. *The International Journal of Robotics Research* 30(7): 895–913.

- Van Den Berg J, Patil S and Alterovitz R (2012) Motion planning under uncertainty using iterative local optimization in belief space. *Intl J of Robotics Research* 31(11): 1263–1278.
- Voss C, Moll M and Kavraki LE (2015) A heuristic approach to finding diverse short paths. In IEEE Intl Conf On Robotics and Automation (ICRA), Seattle, WA, USA, 2015, pp. 4173–4179.
- Wilson RC and Zhu P (2008) A study of graph spectra for comparing graphs and trees. *Pattern Recognition* 41(9): 2833–2841.
- Yannakakis M (1981) Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods* 2(1): 77–79.

Appendix

Appendix A: Index to Multimedia Extensions

1. Video Active SLAM small-scale real-world experiment ([Supplemental Video 1](#))
2. Video Active SLAM large-scale simulation ([Supplemental Video 2](#))

Appendix B: Corollary 1

Proof. Rewrite ΔJ_{\max} as

$$\Delta J_{\max} = \left\{ J(\hat{u}) + \underbrace{UB[J(\hat{u})]}_{\delta(\hat{u}) \geq 0} - J(\hat{u}) \right\} + \left\{ -J(u^*) + \underbrace{J(u^*) - \min_u \mathcal{L}B[J(u)]}_{\delta(u^*) \geq 0} \right\}$$

Therefore, $\delta = \delta(\hat{u}) + \delta(u^*) \geq 0$ and $\Delta J_{\max} = J(\hat{u}) - J(u^*) + \delta \geq 0$ because u^* is the minimum of the entropy. Also, $UB[J(\hat{u})] = -s_{ST}(\hat{u})$ is the minimum upper bound of all actions if we select action in $t\text{-bsp}$ according to equation (7). We know that the optimal action must have entropy below or equal to this value, otherwise the selected action would be better which is a contradiction by itself. Then, $\Delta J_{\max} = J(\hat{u}) - J(u^*) + \delta = |J(\hat{u}) - J(u^*)| + \delta \Rightarrow |J(\hat{u}) - J(u^*)| = \varepsilon(J, s) = \Delta J_{\max} - \delta \leq \Delta J_{\max}$. \square

Appendix C: Theorem 1

Proof. Let $D = \text{diag}(d_1, d_2, \dots, d_n)$, and $S = D^{-1/2}$. By definition, the normalized graph Laplacian $\hat{L} \doteq D^{-1/2} L D^{-1/2} = S L S^T$, since $S = S^T$.

\hat{L} is a real symmetric matrix with dimension n which implies all its eigenvalues are real and \hat{L} has n linearly independent eigenvectors, i.e. can be diagonalized by an orthogonal matrix according to the spectral theorem (Horn and Johnson 2013, theorems 2.3.1, 4.1.5). Let the spectral decomposition of the normalized graph Laplacian be $\hat{L} = \hat{Q} \hat{\Lambda} \hat{Q}^T$ where $\hat{\Lambda} = \text{diag}(\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_n)$ represents a diagonal matrix with eigenvalues of \hat{L} on the main diagonal and \hat{Q} an orthogonal matrix ($\hat{Q}^{-1} = \hat{Q}^T$) whose columns are unit eigenvectors of \hat{L} associated to them. Even if there are degenerate (repeated) eigenvalues, it is always possible to find an orthogonal basis of \mathbb{R}^n consisting of n eigenvectors of \hat{L} e.g. by the Gram-Schmidt orthonormalization process (Horn and Johnson 2013, §0.6.4). Analogously, the spectral decomposition of the graph Laplacian is given by $L = Q \Lambda Q^{-1}$ with $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$ where $\lambda_i, i = 1, 2, \dots, n$, are the eigenvalues of L and columns of Q are their corresponding orthogonal unit eigenvectors. \hat{Q}, Q and S are invertible, thus we can write $\hat{\Lambda} = \underbrace{\hat{Q}^T}_{M} S Q \underbrace{\Lambda Q^T S^T}_{M^T} \hat{Q}$.

We can perform partitioning $M = \begin{pmatrix} a & b \\ c & D \end{pmatrix}$, $M^T = \begin{pmatrix} a & c \\ b & D^T \end{pmatrix}$, $\Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \check{\Lambda} \end{pmatrix}$ and $\hat{\Lambda} = \begin{pmatrix} \hat{\lambda}_1 & 0 \\ 0 & \check{\Lambda} \end{pmatrix}$, where $\lambda_1 = \hat{\lambda}_1 = 0$. The multiplicity of 0 as a Laplacian eigenvalue of an undirected graph equals the number of connected components of the graph (Brouwer and Haemers 2011). Both graph Laplacians are positive semidefinite, and therefore for connected graphs, have *exactly one* zero eigenvalue making $\check{\Lambda} > 0$,

$$\check{\Lambda} = \bar{D} \check{\Lambda} \bar{D}^T. \quad (46)$$

Taking the determinant of equation (46), we get

$$\prod_{i=2}^n \hat{\lambda}_i = \det(\bar{D}) \prod_{i=2}^n \lambda_i \det(\bar{D}^T). \quad (47)$$

Since $\det(\bar{D}) = \det(\bar{D}^T)$ and $t(G) = 1/n \prod_{i=2}^n \lambda_i$ (matrix tree theorem) \Rightarrow

$$t(G) = \frac{\prod_{i=2}^n \hat{\lambda}_i}{n \det(\bar{D})^2} \quad (48)$$

Let $S = \text{diag}(s)$, $\hat{Q} = (\hat{q}^1 \dots \hat{q}^n)$ and $Q = (q^1 \dots q^n)$ with $s \in \mathbb{R}^n$ being the main diagonal of S , and \hat{q}^i, q^i , the i -th column vector of \hat{Q} and Q , respectively. Then, after multiplication

$$M = \begin{pmatrix} (s \circ \hat{q}^1)^T q^1 & (s \circ \hat{q}^1)^T q^2 & \dots & (s \circ \hat{q}^1)^T q^n \\ (s \circ \hat{q}^2)^T q^1 & (s \circ \hat{q}^2)^T q^2 & \dots & (s \circ \hat{q}^2)^T q^n \\ \vdots & \vdots & \ddots & \vdots \\ (s \circ \hat{q}^n)^T q^1 & (s \circ \hat{q}^n)^T q^2 & \dots & (s \circ \hat{q}^n)^T q^n \end{pmatrix}$$

Now we can read partitions of M : $a = (s \circ \hat{q}^1)^T q^1 = (S\hat{q}^1)^T q^1$ and $b = \left((s \circ \hat{q}^1)^T q^2 \dots (s \circ \hat{q}^1)^T q^n \right)$. Here the operator ‘ \circ ’ represents element-wise multiplication of its vector operands.

Graph Laplacian always has the eigenpair $\lambda_1 = 0$, $q^1 = 1/\sqrt{n}(1 \dots 1)^T = 1/\sqrt{n}\mathbf{1}$ (sum of its rows/columns is zero). Also, if v is an eigenvector of \hat{L} , after rescaling by S it becomes a generalized eigenvector u of L , i.e.

$$\begin{aligned} \hat{L}v &= \hat{\lambda}v \\ D^{-1/2}LD^{-1/2}v &= \hat{\lambda}v \\ \underbrace{LD^{-1/2}v}_u &= \hat{\lambda}D^{1/2}v \\ Lu &= \hat{\lambda}Du \end{aligned}$$

$Lq^1 = \hat{L}\hat{q}^1 = 0 \Rightarrow \hat{q}^1 = D^{1/2}\mathbf{1} / \sqrt{\sum_{i=1}^n d_i}$, the unit eigenvector associated to $\hat{\lambda}_1 = 0$. Now we can determine

$$a = (S\hat{q}^1)^T q^1 = \left(D^{-1/2} \frac{D^{1/2}\mathbf{1}}{\sqrt{\sum_{i=1}^n d_i}} \right)^T \frac{\mathbf{1}}{\sqrt{n}} = \dots \quad (49)$$

$$= \frac{\mathbf{1}^T \mathbf{1}}{\sqrt{n \sum_{i=1}^n d_i}} = \sqrt{\frac{n}{\sum_{i=1}^n d_i}}. \quad (50)$$

Also, because $S\hat{q}^1 = \mathbf{1}/\sum_{i=1}^n d_i$ it follows that

$$b = \frac{1}{\sum_{i=1}^n d_i} \left(\underbrace{\mathbf{1}^T q^2}_0 \dots \underbrace{\mathbf{1}^T q^n}_0 \right) = 0$$

Elements of the vector part of b are scalar products of the last $n - 1$ eigenvectors of L with the first eigenvector which all equal to zero since eigenvectors are orthogonal. \bar{D} is non-singular as can be seen from equation (47) so its Schur complement exists. Because of the orthogonality of Q and $\hat{Q} \Rightarrow |\det(Q)| = |\det(\hat{Q})| = 1$ and after applying the matrix determinant lemma we get

$$\begin{aligned} |\det(M)| &= |\det(S)| = |\det(\bar{D})| |\det(a - b\bar{D}^{-1}c)| \Rightarrow \\ |\det(\bar{D})| &= \sqrt{\bar{d}} \prod_{i=1}^n s(i), \text{ where } \bar{d} = \frac{\sum_{i=1}^n d_i}{n} \end{aligned} \quad (51)$$

Finally, substituting (51) into (48), the theorem statement follows

$$t(G) = \frac{\prod_{i=2}^n \hat{\lambda}_i}{n|\bar{D}|^2} = \frac{\prod_{i=1}^n d_i}{\sum_{i=1}^n d_i} \prod_{i=2}^n \hat{\lambda}_i \quad \square \quad (52)$$

Appendix D: Proposition 1

Proof. Using Theorem 1 and equation (35), we can write $\tau(G) = \ln(\prod d_i / \sum d_i) + \sum_{i=2}^n \ln \hat{\lambda}_i = \ln(\prod d_i / \sum d_i) + \sum_{i=2}^n [R_1^r(\hat{\lambda}_i) + R_1^r(\hat{\lambda}_i)]$. Now we can see that the error $\varepsilon(G, r) = \sum_{i=2}^n R_1^r(\hat{\lambda}_i)$. Since $\hat{\lambda}_i \in (0, 2]$ for every $i > 1$ of a connected graph G and this interval is an interval of convergence of the Taylor series of the logarithm function, it follows $\lim_{r \rightarrow \infty} \varepsilon(G, r) = 0$.

Consider the remainder term $R_1^r(\hat{\lambda}_i)$ given by the equation (36) for $r = 2k - 1$ and $k \in \mathbb{N}$.

For $\hat{\lambda}_i \geq 1 \Rightarrow R_1^r(\hat{\lambda}_i) \leq 0$.

Also, for $\hat{\lambda}_i < 1 \Rightarrow R_1^r(\hat{\lambda}_i) = (-1)^r (-1)^{r+1} (1 - \hat{\lambda}_i)^{r+1} / ((r+1)\xi^{r+1}) < 0$.

Therefore, $\forall \hat{\lambda}_i$ and $r = 2k - 1 \Rightarrow \varepsilon(G, r) \leq 0 \Rightarrow \hat{\tau}(G, r) \geq \tau(G)$. \square

Appendix E: Relation between topological signatures

Recall that $\hat{\tau}(G, r) = \ln(\prod d_i / \sum d_i) + \sum_{k=1}^r \sum_{l=0}^k a_{kl} \text{Tr}[\hat{L}^{k-l}]$ with

$$a_{kl} = \frac{(-1)^{k+l-1}}{k} \binom{k}{l}$$

$$\text{Tr}[\hat{L}^0] = n - 1$$

$$\text{Tr}[\hat{L}] = n.$$

Also,

$$\hat{s}_{\text{VN}}(G) = \frac{n \ln 2}{2} - \frac{1}{2} (\text{Tr}[\hat{L}^2] - n). \quad (53)$$

Grouping all the coefficients next to $\text{Tr}[\hat{L}^2]$, $\text{Tr}[\hat{L}]$ and $\text{Tr}[\hat{L}^0]$, i.e. for all $k - l = 2, 1, 0 \Rightarrow$

$$\begin{aligned} \hat{\tau}(G, r) &= \ln \frac{\prod d_i}{\sum d_i} + \sum_{k=1}^r \sum_{l=0}^k a_{kl} \text{Tr}[\hat{L}^{k-l}] \\ &= \ln \frac{\prod d_i}{\sum d_i} + \sum_{k=2}^r a_{k, k-2} \text{Tr}[\hat{L}^2] \\ &\quad + \sum_{k=1}^r a_{k, k-1} \text{Tr}[\hat{L}] + \sum_{k=1}^r a_{k, k} \text{Tr}[\hat{L}^0] + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\hat{L}^{k-l}] \end{aligned} \quad (54)$$

$$\begin{aligned}
&= \ln \frac{\prod d_i}{\sum d_i} + \sum_{k=2}^r \frac{(-1)^{2k-3}}{k} \binom{k}{k-2} \text{Tr}[\widehat{L}^2] \\
&+ \sum_{k=1}^r \frac{(-1)^{2k-2}}{k} \binom{k}{k-1} \text{Tr}[\widehat{L}] \quad (55)
\end{aligned}$$

$$\begin{aligned}
&+ \sum_{k=1}^r \frac{(-1)^{2k-1}}{k} \binom{k}{k} \text{Tr}[\widehat{L}^0] + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] \\
&= \ln \frac{\prod d_i}{\sum d_i} + \text{Tr}[\widehat{L}^2] \sum_{k=2}^r \frac{(-1)^{2k-3}}{k} \frac{k(k-1)}{2} + \text{Tr}[\widehat{L}] \sum_{k=1}^r \frac{1}{k} k \\
&+ \text{Tr}[\widehat{L}^0] \sum_{k=1}^r \frac{(-1)^{2k-1}}{k} + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] \quad (56)
\end{aligned}$$

$$\begin{aligned}
&= \ln \frac{\prod d_i}{\sum d_i} - \frac{\text{Tr}[\widehat{L}^2]}{2} \frac{r(r-1)}{2} + nr - (n-1)H(r) \\
&+ \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] \quad (57)
\end{aligned}$$

$$\begin{aligned}
&= \ln \frac{\prod d_i}{\sum d_i} - \frac{\text{Tr}[\widehat{L}^2]}{2} \frac{r(r-1)}{2} + \left[\frac{r(r-1)}{4} n - \frac{r(r-1)}{4} n \right] \\
&+ nr - (n-1)H(r) + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] \quad (58)
\end{aligned}$$

$$\begin{aligned}
&= \ln \frac{\prod d_i}{\sum d_i} - \frac{r(r-1)}{2 \cdot 2} [\text{Tr}[\widehat{L}^2] - n] - \frac{r(r-1)}{4} n + nr \\
&- (n-1)H(r) + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}]. \quad (59)
\end{aligned}$$

Putting (53) into (59) \Rightarrow

$$\begin{aligned}
\widehat{\tau}(G, r) &= \ln \frac{\prod d_i}{\sum d_i} + \frac{r(r-1)}{2} \left(\widehat{s}_{\text{VN}}(G) - \frac{n \ln 2}{2} \right) \\
&- \frac{r(r-1)}{4} n + nr - (n-1)H(r) + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] \quad (60)
\end{aligned}$$

where $H(r) = \sum_{k=1}^r 1/k$ is a harmonic number of r .

After rearranging,

$$\begin{aligned}
\widehat{\tau}(G, r) &= \frac{r(r-1)}{2} \widehat{s}_{\text{VN}}(G) + \sum_{k=3}^r \sum_{l=0}^{k-3} a_{kl} \text{Tr}[\widehat{L}^{k-l}] + \ln \frac{\prod d_i}{\sum d_i} \\
&+ \left[r - \frac{r(r-1)(\ln 2 + 1)}{4} - H(r) \right] n + H(r) \quad (61)
\end{aligned}$$

and substituting it into (9) we obtain power series relation between topological signatures given by equation (40).