

# Densification Strategies for Anytime Motion Planning over Large Dense Roadmaps

Shushman Choudhury, Oren Salzman, Sanjiban Choudhury and Siddhartha S. Srinivasa<sup>1</sup>

**Abstract**—We consider the problem of computing shortest paths in a dense motion-planning roadmap  $\mathcal{G}$ . We assume that  $n$ , the number of vertices of  $\mathcal{G}$ , is very large. Thus, using any path-planning algorithm that directly searches  $\mathcal{G}$ , running in  $O(V \log V + E) \approx O(n^2)$  time, becomes unacceptably expensive. We are therefore interested in anytime search to obtain successively shorter feasible paths and converge to the shortest path in  $\mathcal{G}$ . Our key insight is to provide existing path-planning algorithms with a sequence of increasingly dense subgraphs of  $\mathcal{G}$ . We study the space of all ( $r$ -disk) subgraphs of  $\mathcal{G}$ . We then formulate and present two densification strategies for traversing this space which exhibit complementary properties with respect to problem difficulty. This inspires a third, hybrid strategy which has favourable properties regardless of problem difficulty. This general approach is then demonstrated and analyzed using the specific case where a low-dispersion deterministic sequence is used to generate the samples used for  $\mathcal{G}$ . Finally we empirically evaluate the performance of our strategies for random scenarios in  $\mathbb{R}^2$  and  $\mathbb{R}^4$  and on manipulation planning problems for a 7 DOF robot arm, and validate our analysis.

## I. INTRODUCTION

Let  $\mathcal{G}$  be a motion-planning roadmap with  $n$  vertices embedded in some configuration space (C-space). We consider the problem of finding a shortest path between two vertices of  $\mathcal{G}$ . Specifically, we are interested in settings, prevalent in motion planning, where testing if an edge of the graph is collision free or not is computationally expensive. We call such graphs Explicit graphs with Expensive Edge-Evaluation or  $E^4$ -graphs. Moreover, we are interested in the case where  $n$  is very large, and where the roadmap is dense, i.e.  $|E| = O(n^2)$ . This makes any path-finding algorithm that directly searches  $\mathcal{G}$ , subsequently performing  $O(n^2)$  edge evaluations, impractical. We wish to obtain an approximation of the shortest path quickly and refine it as time permits. We refer to this problem as *anytime planning on large  $E^4$ -graphs*.

Our problem is motivated by previous work (Sec. II) on sampling-based motion-planning algorithms that construct a *fixed* roadmap as part of a preprocessing stage [1, 16, 24]. These methods are used to efficiently approximate the structure of the C-space. When the size of the roadmap is large, even finding a solution, let alone an optimal one, becomes a non-trivial problem requiring specifically-tailored search algorithms [24]. Our roadmap formulation departs from the PRM setting which chooses a connectivity radius  $O(\log n)$  that achieves asymptotic optimality. We are interested in

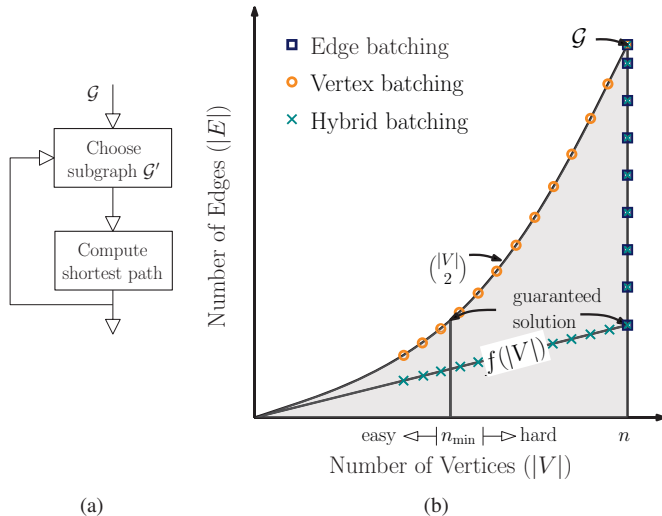


Fig. 1: (a) Our meta algorithm leverages existing path-planning algorithms and provides them with a sequence of subgraphs. (b) To do so we consider densification strategies for traversing the space of  $r$ -disk subgraphs of the roadmap  $\mathcal{G}$ . The  $x$ -axis and the  $y$ -axis represent the number of vertices and the number of edges (induced by  $r$ ) of the subgraph, respectively. A particular subgraph is defined by a point in this space. *Edge batching* searches over all samples and adds edges according to an increasing radius of connectivity. *Vertex batching* searches over complete subgraphs induced by progressively larger subsets of vertices. *Hybrid batching* uses the minimal connection radius  $f(|V|)$  to ensure connectivity until it reaches  $|V| = n$  and then proceeds like edge batching. The harder a problem, i.e. the lower the clearance between obstacles, the more vertices are needed by vertex and hybrid batching ( $n_{\min}$ ) to get their first feasible solution.

dense, nearly-complete roadmaps that capture as much C-space connectivity information as possible, and probably have one or more paths that are strictly shorter than the optimal path for the standard PRM.

Our key insight for solving the anytime planning problem in large  $E^4$ -graphs is to provide existing path-planning algorithms with a sequence of increasingly dense subgraphs of  $\mathcal{G}$ , using some *densification strategy*. At each iteration, we run a shortest-path algorithm on the current subgraph to obtain an increasingly tighter approximation of the true shortest path. This favours using incremental search techniques that reuse information between calls. We present a number of such strategies, and we address the question:

How does the densification strategy affect the time at which the first solution is found, and the quality of the solutions obtained?

We focus on  $r$ -disk subgraphs of  $\mathcal{G}$ , i.e. graphs defined by a specific set of vertices where every two vertices are connected if their mutual distance is at most  $r$ . This induces a space of subgraphs (Fig. 1) defined by the number of vertices and the connection radius (which, in turn, defines

\*Work by Sh. C., O. S. and S. S. was (partially) funded by the National Science Foundation IIS (#1409003), Toyota Motor Engineering & Manufacturing (TEMA), and the Office of Naval Research.

<sup>1</sup> The Robotics Institute, Carnegie Mellon University {shushmac, osalzman, sanjibac, ss5}@andrew.cmu.edu

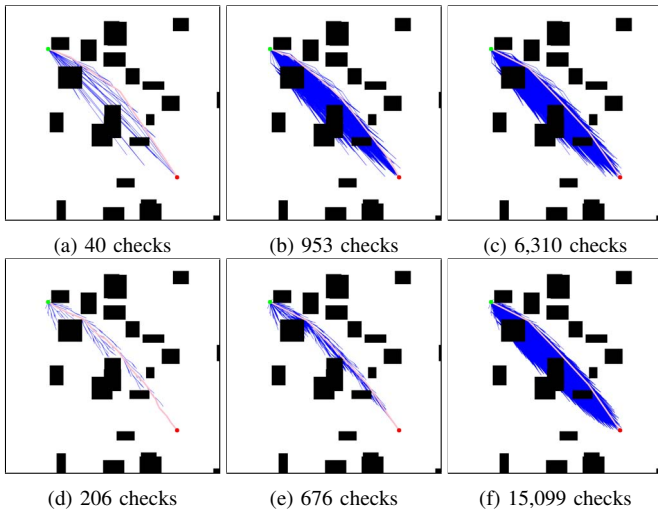


Fig. 2: A visualization of an easy  $\mathbb{R}^2$  problem, with higher average obstacle gaps. The initial and final paths are shown, as well as an intermediate path. Vertex batching (a-c) finds the first solution quickly with a sparse set of initial samples. Additional heuristics hereafter help it converge to the optimum with fewer edge evaluations than edge batching (d-f). Note that the coverage of collision checks only appears similar at the end due to resolution limits for visualization.

the number of edges). We observe two natural ways to traverse this space. The first is to define each subgraph over the entire set of vertices and incrementally add *batches of edges* by increasing  $r$  (vertical line at  $|V| = n$  in Fig. 1). Alternatively, we can incrementally add *batches of vertices* and, at each iteration, consider the complete graph ( $r = r_{\max}$ ) defined over the current set of vertices (parabolic arc  $|E| = O(|V|^2)$  in Fig. 1). We call these variants *edge batching* and *vertex batching*, respectively. Vertex batching and edge batching seem to be better suited for easy and hard problems, respectively, as visualized and explained in Fig. 2 and Fig. 3. This analysis motivates our *hybrid batching* strategy, which is more robust to problem difficulty.

Our main contribution is the formulation and analysis of various densification strategies to traverse the space of subgraphs of  $\mathcal{G}$  (Sec. IV). We analyse the specific case where the vertices of  $\mathcal{G}$  are obtained from a low-dispersion deterministic sequence (Sec. V). Specifically, we describe the structure of the space of subgraphs and demonstrate the tradeoff between effort and bounded suboptimality for our densification strategies. Furthermore, we explain how this tradeoff varies with problem difficulty, which is measured in terms of the clearance of the shortest path in  $\mathcal{G}$ .

We discuss implementation decisions and parameters that allow us to efficiently use our strategies in dense  $E^4$  graphs in Sec. VI. We then empirically validate our analysis on several random scenarios in  $\mathbb{R}^2$  and  $\mathbb{R}^4$  and on manipulation planning problems for a 7 DOF robot arm (Sec. VII). Finally, we discuss directions of future work (Sec. VIII).

## II. RELATED WORK

### A. Sampling-based motion planning

Sampling-based planning approaches build a graph, or a roadmap, in the C-space, where vertices are configurations and edges are local paths connecting configurations. A path




Fig. 3: A visualization of a hard  $\mathbb{R}^2$  problem with  $10\times$  more obstacles and lower average obstacle gaps. The same set of samples  $\mathcal{S}$  is used as in Fig. 2. Because the problem is more difficult, both vertex and edge batching require more edge evaluations for finding feasible solutions and converging to the optimum. In particular, vertex batching (a-c) requires multiple iterations to find its first solution, while edge batching (d-f) still does so on its first search, albeit with more collision checks than for the easy problem.

is then found by traversing this roadmap while checking if the vertices and edges are collision free. Initial algorithms such as PRM [16] and RRT [19] were concerned with finding a *feasible* solution. However, in recent years, there has been growing interest in finding high-quality solutions. Karaman and Frazzoli [14] introduced variants of the PRM and RRT algorithms, called PRM\* and RRT\*, respectively and proved that, asymptotically, the solution obtained by these algorithms converges to the optimal solution. However, the running times of these algorithms are often significantly higher than their non-optimal counterparts. Thus, subsequent algorithms have been suggested to increase the rate of convergence to high-quality solutions. They use different approaches such as lazy computation [1, 13, 22], informed sampling [8], pruning vertices [9], relaxing optimality [23], exploiting local information [3] and lifelong planning together with heuristics [4]. In this work we employ several such techniques in order to speed up the convergence rate of our algorithms.

### B. Finite-time properties of sampling-based algorithms

Extensive analysis has been done on *asymptotic* properties of sampling-based algorithms, i.e. properties such as connectivity and optimality when the number of samples tends to infinity [14, 15].

We are interested in bounding the quality of a solution obtained using a *fixed* roadmap for a finite number of samples. When the samples are generated from a *deterministic* sequence, Janson et. al. [12, Thm2] give a closed-form solution bounding the quality of the solution of a PRM whose roadmap is an  $r$ -disk graph. The bound is a function of  $r$ , the number of vertices  $n$  and the *dispersion* of the set of points used. (See Sec. III for an exact definition of dispersion and for the exact bound given by Janson et. al.).

Dobson et. al. [7] provide similar bounds when randomly sampled i.i.d points are used. Specifically, they consider a

PRM whose roadmap is an  $r$ -disk graph for a *specific* radius  $r = c \cdot (\log n/n)^{1/d}$  where  $n$  is the number of points,  $d$  is the dimension and  $c$  is some constant. They then give a bound on the probability that the quality of the solution will be larger than a given threshold.

### C. Efficient path-planning algorithms

We are interested in path-planning algorithms that attempt to reduce the amount of computationally expensive edge expansions performed in a search. This is typically done using heuristics such as for A\* [11], for Iterative Deepening A\* [18] and for Lazy Weighted A\* [4]. Some of these algorithms, such as Lifelong Planning A\* [17] allow recomputing the shortest path in an efficient manner when the graph undergoes changes. *Anytime* variants of A\* such as Anytime Repairing A\* [20] and Anytime Nonparametric A\* [27] efficiently run a succession of A\* searches, each with an inflated heuristic. This potentially obtains a fast approximation and refines its quality as time permits. However, there is no formal guarantee that these approaches will decrease search time and they may still search all edges of a given graph [28]. For a unifying formalism of such algorithms relevant to  $E^d$  graphs and additional references, see [6].

## III. NOTATION, PROBLEM FORMULATION AND MATHEMATICAL BACKGROUND

We provide standard notation and define our problem concretely. We then provide necessary mathematical background about the *dispersion* of a set of points.

### A. Notation and problem formulation

Let  $\mathcal{X}$  denote a  $d$ -dimensional C-space,  $\mathcal{X}_{\text{free}}$  the collision-free portion of  $\mathcal{X}$ ,  $\mathcal{X}_{\text{obs}} = \mathcal{X} \setminus \mathcal{X}_{\text{free}}$  its complement and let  $\rho : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  be some distance metric. For simplicity, we assume that  $\mathcal{X} = [0, 1]^d$  and that  $\rho$  is the Euclidean norm. Let  $\mathcal{S} = \{s_1, \dots, s_n\}$  be some sequence of points where  $s_\ell \in \mathcal{X}$  and denote by  $\mathcal{S}(\ell)$  the first  $\ell$  elements of  $\mathcal{S}$ . We define the  $r$ -disk graph  $\mathcal{G}(\ell, r) = (V_\ell, E_{\ell, r})$  where  $V_\ell = \mathcal{S}(\ell)$ ,  $E_{\ell, r} = \{(u, v) \mid u, v \in V_\ell \text{ and } \rho(u, v) \leq r\}$  and each edge  $(u, v)$  has a weight  $w(u, v) = \rho(u, v)$ . See [14, 25] for various properties of such graphs in the context of motion planning. Finally, set  $\mathcal{G} = \mathcal{G}(n, \sqrt{d})$ , namely, the complete<sup>1</sup> graph defined over  $\mathcal{S}$ .

For ease of analysis we assume that the roadmap is complete, but our densification strategies and analysis can be extended to *dense* roadmaps that are not complete. Furthermore, our definition assumes that  $\mathcal{G}$  is embedded in the C-space. Thus, we will use the terms *vertices* and *configurations* as well as *edges* and *paths* in C-space interchangeably.

A query  $\mathcal{Q}$  is a scenario with start and target configurations. Let the start and target configurations be  $s_1$  and  $s_2$ , respectively. The obstacles induce a mapping  $\mathcal{M} : \mathcal{X} \rightarrow \{\mathcal{X}_{\text{free}}, \mathcal{X}_{\text{obs}}\}$  called a *collision detector* which checks if a configuration or edge is collision-free or not. Typically, edges are checked by densely sampling along the edge, hence the

<sup>1</sup>Using a radius of  $\sqrt{d}$  ensures that every two points will be connected due to the assumption that  $\mathcal{X} = [0, 1]^d$  and that  $\rho$  is Euclidean.

term *expensive edge evaluation*. A feasible path is denoted by  $\gamma : [0, 1] \rightarrow \mathcal{X}_{\text{free}}$  where  $\gamma[0] = s_1$  and  $\gamma[1] = s_2$ . Slightly abusing this notation, set  $\gamma(\mathcal{G}(\ell, r))$  to be the shortest collision-free path from  $s_1$  to  $s_2$  that can be computed in  $\mathcal{G}(\ell, r)$ , its clearance as  $\delta(\mathcal{G}(\ell, r))$  and denote by  $\gamma^* = \gamma(\mathcal{G})$  and  $\delta^* = \delta(\mathcal{G})$  the shortest path and its clearance that can be computed in  $\mathcal{G}$ , respectively. Note that a path has clearance  $\delta$  if every point on the path is at a distance of at least  $\delta$  away from every obstacle.

Our problem calls for finding a sequence of increasingly shorter feasible paths  $\gamma_0, \gamma_1 \dots$  in  $\mathcal{G}$ , converging to  $\gamma^*$ . We assume that  $n = |\mathcal{S}|$  is sufficiently large, and the roadmap covers the space well enough so that for any reasonable set of obstacles, there are multiple feasible paths to be obtained between start and goal. Therefore, we do not consider a case where the entire roadmap is invalidated by obstacles. The large value of  $n$  makes any path-finding algorithm that directly searches  $\mathcal{G}$ , thereby performing  $O(n^2)$  calls to the collision-detector, too time-consuming to be practical.

### B. Dispersion

The *dispersion*  $D_n(\mathcal{S})$  of a sequence  $\mathcal{S}$  is defined as  $D_n(\mathcal{S}) = \sup_{x \in \mathcal{X}} \min_{s \in \mathcal{S}} \rho(x, s)$ . Intuitively, it can be thought of as the radius of the largest empty ball (by some metric) that can be drawn around any point in the space  $\mathcal{X}$  without intersecting any point of  $\mathcal{S}$ . A lower dispersion implies a better *coverage* of the space by the points in  $\mathcal{S}$ . When  $\mathcal{X}$  is the  $d$ -dimensional Euclidean space and  $\rho$  is the Euclidean distance, deterministic sequences with dispersion of order  $O(n^{-1/d})$  exist. A simple example is a set of points lying on grid or a lattice.

Other low-dispersion deterministic sequences exist which also have low discrepancy, i.e. they appear to be random for many purposes. One such example is the *Halton sequence* [10]. We will use them extensively for our analysis because they have been studied in the context of deterministic motion planning [2, 12]. For Halton sequences, tight bounds on dispersion exist. Specifically,  $D_n(\mathcal{S}) \leq p_d \cdot n^{-1/d}$  where  $p_d \approx d \log d$  is the  $d^{\text{th}}$  prime number. Subsequently in this paper, we will use  $D_n$  (and not  $D_n(\mathcal{S})$ ) to denote the dispersion of the first  $n$  points of  $\mathcal{S}$ .

Janson et. al. bound the length of the shortest path computed over an  $r$ -disk roadmap constructed using a low-dispersion deterministic sequence [12, Thm2]. Specifically, given start and target vertices, consider all paths  $\Gamma$  connecting them which have  $\delta$ -clearance for some  $\delta$ . Set  $\delta_{\text{max}}$  to be the maximal clearance over all such  $\delta$ . If  $\delta_{\text{max}} > 0$ , then for all  $0 < \delta \leq \delta_{\text{max}}$  set  $c^*(\delta)$  to be the cost of the shortest path in  $\Gamma$  with  $\delta$ -clearance. Let  $c(\ell, r)$  be the length of the path returned by a shortest-path algorithm on  $\mathcal{G}(\ell, r)$  with  $\mathcal{S}(\ell)$  having dispersion  $D_\ell$ . For  $2D_\ell < r < \delta$ , we have that

$$c(\ell, r) \leq \left(1 + \frac{2D_\ell}{r - 2D_\ell}\right) \cdot c^*(\delta). \quad (1)$$

Notably, for  $n$  random i.i.d. points, the lower bound on the dispersion is  $O((\log n/n)^{1/d})$  [21] which is strictly larger than for deterministic samples.




Fig. 4: Regions of interest for the space of subgraphs. *Vertex Starvation* happens in the region with too few vertices to ensure a solution, even for a fully connected subgraph. *Edge Starvation* happens in the region where the radius  $r$  is too low to guarantee connectivity.

For domains other than the unit hypercube, the insights from the analysis will generally hold. However, the dispersion bounds may become far more complicated depending on the domain, and the distance metric would need to be scaled accordingly. This may result in the quantitative bounds being difficult to deduce analytically.

#### IV. APPROACH

We now discuss our general approach of searching over the space of all ( $r$ -disk) subgraphs of  $\mathcal{G}$ . We start by characterizing the boundaries and different regions of this space. Subsequently, we introduce two densification strategies—edge batching and vertex batching. As we will see, these two are complementary in nature, which motivates our third strategy, which we call hybrid batching.

##### A. The space of subgraphs

To perform an anytime search over  $\mathcal{G}$ , we iteratively search a sequence of graphs  $\mathcal{G}_0(n_0, r_0) \subseteq \mathcal{G}_1(n_1, r_1) \subseteq \dots \subseteq \mathcal{G}_m(n_m, r_m) = \mathcal{G}$ . If no feasible path exists in the subgraph, we move on to the next subgraph in the sequence, which is more likely to have a feasible path.

We use an incremental path-planning algorithm that allows us to efficiently recompute shortest paths. Our problem setting of increasingly dense subgraphs is particularly amenable to such algorithms. However, any alternative shortest-path algorithm may be used. We emphasize again that we focus on the meta-algorithm of choosing which subgraphs to search. Further details on the implementation of these approaches are provided in Sec. VI.

Fig. 4 depicts the set of possible graphs  $\mathcal{G}(\ell, r)$  for all choices of  $0 < \ell \leq n$  and  $0 < r \leq \sqrt{d}$ . Specifically, the graph depicts  $|E_{\ell, r}|$  as a function of  $|V_{\ell}|$ . We discuss Fig. 4 in detail to motivate our approach for solving the problem of anytime planning in large  $E^d$ -graphs and the specific sequence of subgraphs we use. First, consider the

curves that define the boundary of all possible graphs: The vertical line  $|V| = n$  corresponds to subgraphs defined over the entire set of vertices, where batches of edges are added as  $r$  increases. The parabolic arc  $|E| = |V| \cdot (|V| - 1)/2$ , corresponds to complete subgraphs defined over increasingly larger sets of vertices.

Recall that we wish to approximate the shortest path  $\gamma^*$  which has some minimal clearance  $\delta^*$ . Given a specific graph, to ensure that a path that approximates  $\gamma^*$  is found, two conditions should be met: (i) The graph includes some minimal number  $n_{\min}$  of vertices. The exact value of  $n_{\min}$  will be a function of the dispersion  $D_{n_{\min}}$  of the sequence  $\mathcal{S}$  and the clearance  $\delta^*$ . (ii) A minimal connection radius  $r_0$  is used to ensure that the graph is connected. Its value will depend on the sequence  $\mathcal{S}$  (and not on  $\delta^*$ ).

Requirement (i) induces a vertical line at  $|V| = n_{\min}$ . Any point to the left of this line corresponds to a graph with too few vertices to prove any guarantee that a solution will be found. We call this the *vertex-starvation* region. Requirement (ii) induces a curve  $f(|V|)$  such that any point below this curve corresponds to a graph which may be disconnected. We call this the *edge-starvation* region. The exact form of the curve depends on the sequence  $\mathcal{S}$  that is used. The specific value of  $n_{\min}$  and the form of  $f(|V|)$  when Halton sequences are used is provided in Sec. V.

Any point outside the starvation regions represents a graph  $\mathcal{G}(\ell, r)$  such that the length of  $\gamma(\mathcal{G}(\ell, r))$  may be bounded. For a discussion on specific bounds, see Sec. II-B. For a visualization of the different regions, see Fig. 4.

##### B. Edge and vertex batching

Our goal is to search increasingly dense subgraphs of  $\mathcal{G}$ . This corresponds to a sequence of points on the space of subgraphs (Fig. 4) that ends at the upper right corner of the space. Two natural strategies emerge from this. We defer the discussion on the choice of parameters used for each strategy to Sec. VI.

1) *Edge batching*: All subgraphs include the complete set of vertices  $\mathcal{S}$  and the edges are incrementally added via an increasing connection radius. Specifically,  $\forall i \ n_i = n$  and  $r_i = \eta_e r_{i-1}$  where  $\eta_e > 1$  and  $r_0$  is some small initial radius. Here, we choose  $r_0 = O(f(n))$ , where  $f$  is the *edge-starvation* boundary curve defined previously. Using Fig. 4, this induces a sequence of points along the vertical line at  $|V| = n$  starting from  $|E| = O(n^2 r_0^d)$  and ending at  $|E| = O(n^2)$ .

2) *Vertex batching*: In this variant, all subgraphs are complete graphs defined over increasing subsets of the complete set of vertices  $\mathcal{S}$ . Specifically  $\forall i \ r_i = r_{\max} = \sqrt{d}$ ,  $n_i = \eta_v n_{i-1}$  where  $\eta_v > 1$  and the base term  $n_0$  is some small number of vertices. Because we have no priors about the obstacle density or distribution, the chosen  $n_0$  is a constant and does not vary due to  $n$  or due to the volume of  $\mathcal{X}_{\text{obs}}$ . Using Fig. 4, this induces a sequence of points along the parabolic arc  $|E| = |V| \cdot (|V| - 1)/2$  starting from  $|V| = n_0$  and ending at  $|V| = n$ . The vertices are chosen in




Fig. 5: Visualization of the work required by our densification strategies as a function of the problem’s hardness. Here work is measured as the number of edges evaluated. This is visualized using the gradient shading where light gray (resp. dark grey) depicts a small (resp. large) amount of work. Assuming  $n > n_{\min}$ , the amount of work required by edge batching remains the same regardless of problem difficulty. For vertex batching the amount of work required depends on the hardness of the problem. Here we visualize an easy and a hard problem using  $n_{\min}$  (easy) and  $n_{\min}$  (hard), respectively.

the same order with which they are generated by  $\mathcal{S}$ . So,  $\mathcal{G}_0$  has the first  $n_0$  samples of  $\mathcal{S}$ , and so on.

Intuitively, the relative performance of these densification strategies depends on problem hardness. We use the clearance of the shortest path,  $\delta^*$ , to represent the hardness of the problem. This, in turn, defines  $n_{\min}$  which bounds the vertex-starvation region. Specifically we say that a problem is easy (resp. hard) when  $\delta^* \approx \sqrt{d}$  (resp.  $\delta^* \approx \Omega(D_n(\mathcal{S}))$ ). For easy problems, with larger gaps between obstacles, vertex batching can find a solution quickly with fewer samples and long edges, thereby restricting the work done for future searches. In contrast, assuming that  $n > n_{\min}$ , edge batching will find a solution on the first iteration but the time to do so may be far greater than for vertex batching because the number of samples is so large. For hard problems vertex batching may require multiple iterations until the number of samples it uses is large enough and it is out of the vertex-starvation region. Each of these searches would exhaust the fully connected subgraph before terminating. This cumulative effort is expected to exceed that required by edge batching for the same problem, which is expected to find a feasible albeit sub-optimal path on the first search. A visual depiction of this intuition is given in Fig. 5.

### C. Hybrid batching

Vertex and edge batching exhibit complementary properties for problems with varying difficulty. Yet, when a query  $\mathcal{Q}$  is given, the hardness of the problem is not known a-priori. In this section we propose a hybrid approach that exhibits favourable properties, regardless of the hardness of the problem.

This hybrid batching strategy commences by searching over a graph  $\mathcal{G}(n_0, r_0)$  where  $n_0$  is the same as for vertex

batching and  $r_0 = O(f(n_0))$ . As long as  $n_i < n$ , the next batch has  $n_{i+1} = \eta_v n_i$  and  $r_i = O(f(n_i))$ . When  $n_i = n$  (and  $r_i = O(f(n))$ ), all subsequent batches are similar to edge batching, i.e.,  $r_{i+1} = \eta_e r_i$  (and  $n_{i+1} = n$ ).

This can be visualized on the space of subgraphs as sampling along the curve  $f(|V|)$  from  $|V| = n_0$  until  $f(|V|)$  intersects  $|V| = n$  and then sampling along the vertical line  $|V| = n$ . See Fig. 1 and Fig. 5 for a mental picture. As we will see in our experiments, hybrid batching typically performs comparably (in terms of path quality) to vertex batching on easy problems and to edge batching on hard problems.

## V. ANALYSIS FOR HALTON SEQUENCES

In this section we consider the space of subgraphs and the densification strategies that we introduced in Sec. IV for the specific case that  $\mathcal{S}$  is a Halton sequence. We start by describing the boundaries of the starvation regions. We then continue by simulating the bound on the quality of the solution obtained as a function of the work done for each of our strategies.

### A. Starvation-region bounds

To bound the vertex starvation region we wish to find  $n_{\min}$  after which bounded sub-optimality can be guaranteed to find the first solution. Note that  $\delta^*$  is the clearance of the shortest path  $\gamma^*$  in  $\mathcal{G}$  connecting  $s_1$  and  $s_2$ , that  $p_d$  denotes the  $d^{\text{th}}$  prime and  $D_n \leq p_d/n^{1/d}$  for Halton sequences. For Eq. (1) to hold we require that  $2D_{n_{\min}} < \delta^*$ . Thus,

$$2D_{n_{\min}} < \delta^* \Rightarrow 2 \frac{p_d}{n_{\min}^{1/d}} < \delta^* \Rightarrow n_{\min} > \left( \frac{2p_d}{\delta^*} \right)^d.$$

Indeed, one can see that as the problem becomes harder (namely,  $\delta^*$  decreases),  $n_{\min}$  and the entire vertex-starvation region grows.

We now show that for Halton sequences, the edge-starvation region has a linear boundary, i.e.  $f(|V|) = O(|V|)$ . Using Eq. (1) we have that the minimal radius  $r_{\min}(|V|)$  required for a graph with  $|V|$  vertices is

$$r_{\min}(|V|) > 2D_{|V|} \Rightarrow r_{\min}(|V|) > \frac{2p_d}{(|V|)^{1/d}}.$$

For any  $r$ -disk graph  $\mathcal{G}(\ell, r)$ , the number of edges is  $|E_{\ell, r}| = O(\ell^2 \cdot r^d)$ . In our case,

$$f(|V|) = O(|V|^2 \cdot r_{\min}^d(|V|)) = O(|V|).$$

### B. Effort-to-quality ratio

We now compare our densification strategies in terms of their worst-case anytime performance. Specifically, we plot the cumulative amount of work as subgraphs are searched, measured by the maximum number of edges that may be evaluated, as a function of the bound on the quality of the solution that may be obtained using Eq. (1). We fix a specific setting (namely  $d$  and  $n$ ) and simulate the work done and the suboptimality using the necessary formulae. This is done for an easy and a hard problem. See Fig. 6.

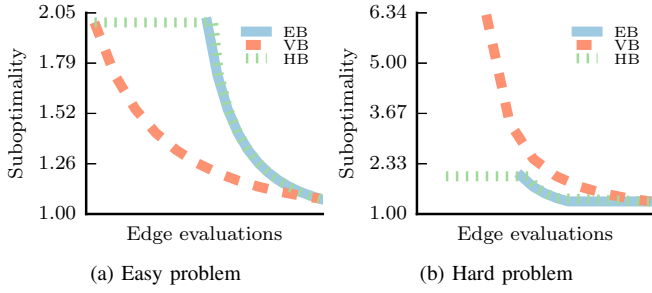


Fig. 6: A simulation of the work-suboptimality tradeoff for vertex, edge and hybrid batching. Here we chose  $n = 10^6$  and  $d = 4$ . The easy and hard problems have  $\delta^* = \sqrt{d}/2$  and  $\delta^* = 5D_n$ , respectively. The plot is produced by sampling points along the curves  $|V| = n$  and  $|E| = |V| \cdot (|V| - 1)/2$  and using the respective values in Eq. (1). Note that x-axis is in log-scale.

Indeed, this simulation coincides with our discussion on properties of both batching strategies with respect to the problem difficulty. Vertex batching outperforms edge batching on easy problems and vice versa. Hybrid batching lies somewhere in between the two approaches with the specifics depending on problem difficulty.

## VI. IMPLEMENTATION

### A. Search Parameters

We choose the parameters for each densification strategy such that the number of batches is  $O(\log_2 n)$ .

1) *Edge Batching*: We set  $\eta_e = 2^{1/d}$ . Recall that for  $r$ -disk graphs, the average degree of vertices is  $n \cdot r_i^d$ , therefore this value (and hence the number of edges) is doubled after each iteration. We set  $r_0 = 3 \cdot n^{-1/d}$ .

2) *Vertex Batching*: We set the initial number of vertices  $n_0$  to be 100, irrespective of the roadmap size and problem setting, and set  $\eta_v = 2$ . After each batch we double the number of vertices.

3) *Hybrid Batching*: The parameters are derived from those used for vertex and edge batching. We begin with  $n_0 = 100$ , and after each batch we increase the vertices by a factor of  $\eta_v = 2$ . For these searches, i.e. in the region where  $n_i < n$ , we use  $r_i = 3 \cdot n^{-1/d}$ . This ensures the same radius at  $n$  as for edge batching. Subsequently, we increase the radius as  $r_i = \eta_e r_{i-1}$ , where  $\eta_e = 2^{1/d}$ .

### B. Optimizations

Our analysis and intuition is agnostic to any specific algorithms or implementations. However, for these densification strategies to be useful in practice, we employ certain optimizations.

1) *Search Technique*: Each subgraph is searched using Lazy A\* [4] with incremental rewiring as in LPA\* [17]. For details, see the search algorithm used for a single batch of BIT\* [9]. This lazy variant of A\* has been shown to outperform other path-planning techniques for motion-planning search problems with expensive edge evaluations [6].

2) *Caching Collision Checks*: Each time the collision-detector  $\mathcal{M}$  is called for an edge, we store the ID of the edge along with the result using a hashing data structure. Subsequent calls for that specific edge are simply lookups

in the hashing data structure which incur negligible running time. Thus,  $\mathcal{M}$  is called for each edge at most once.

3) *Sample Pruning and Rejection*: For anytime algorithms, once an initial solution is obtained, subsequent searches should be focused on the subset of states that could potentially improve the solution. When the space  $\mathcal{X}$  is Euclidean, this, so-called ‘‘informed subset’’, can be described by a prolate hyperspheroid [8]. For our densification strategies, we prune away all existing vertices (for all batching), and reject the newer vertices (for vertex and hybrid batching), that fall outside the informed subset.

Successive prunings due to intermediate solutions significantly reduces the average-case complexity of future searches [9], despite the extra time required to do so, which is accounted for in our benchmarking. Note that for Vertex and Hybrid Batching, which begin with only a few samples, samples in successive batches that are outside the current ellipse can just be rejected. This is cheaper than pruning, which is required for Edge Batching. Across all test cases, we noticed poorer performance when pruning was omitted.

In the presence of obstacles, the extent to which the complexity is reduced due to pruning is difficult to obtain analytically. As shown in Theorem 1, however, in the assumption of free space, we can derive results for Edge Batching. This motivates using this heuristic.

**Theorem 1** *Running edge batching in an obstacle-free  $d$ -dimensional Euclidean space over a roadmap constructed using a deterministic low-dispersion sequence with  $r_0 > 2D_n$  and  $r_{i+1} = 2^{1/d}r_i$ , while using sample pruning and rejection makes the worst-case complexity of the total search, measured in edge evaluations,  $O(n^{1+1/d})$ .*

*Proof*: Let  $c_{\text{best}}^i$  denote the cost of the solution obtained after  $i$  iterations by our edge batching algorithm, and  $c_{\text{min}} = \rho(s_1, s_2) \leq \sqrt{d}$  denote the cost of the optimal solution. Using Eq. (1),

$$c_{\text{best}}^i \leq (1 + \varepsilon_i) c_{\text{min}}, \quad (2)$$

where  $\varepsilon_i = \frac{2D_n}{r_i - 2D_n}$ . Using the parameters for edge batching,

$$\varepsilon_{i+1} = \frac{2D_n}{r_{i+1} - 2D_n} = \frac{2D_n}{2^{1/d}r_i - 2D_n} \leq \frac{\varepsilon_i}{2^{1/d}}. \quad (3)$$

Let  $i_{\text{max}}$  be the maximum number of iterations and recall that we have  $i_{\text{max}} = O(\log_2 n)$ .

Note that the fact that vertices and edges are pruned away, does not change the bound provided in Eq. (2). To compute the actual number of edges considered at the  $i$ th iteration, we bound the volume of the prolate hyperspheroid  $\mathcal{X}_{c_{\text{best}}^i}$  in  $\mathbb{R}^d$  (see [8]) by,

$$\mu(\mathcal{X}_{c_{\text{best}}^i}) = \frac{c_{\text{best}}^i \left( (c_{\text{best}}^i)^2 - c_{\text{min}}^2 \right)^{\frac{d-1}{2}} \xi_d}{2^d}, \quad (4)$$

where  $\xi_d$  is the volume of an  $\mathbb{R}^d$  unit-ball. Using Eq. (2) in Eq. (4),

$$\mu(\mathcal{X}_{c_{\text{best}}^i}) \leq \varepsilon_i^{\frac{d-1}{2}} (1 + \varepsilon_i) (2 + \varepsilon_i)^{\frac{d-1}{2}} \Gamma_d, \quad (5)$$




Fig. 7: Experimental results in random unit hypercube scenarios for  $\square$  vertex batching  $\square$  edge batching, and  $\square$  hybrid batching. The  $y$ -axis is the ratio between the length of the path produced by the algorithm and length of  $\gamma^*$  (the shortest path on  $\mathcal{G}$ ) for that problem. The naive strategy of searching the complete graph requires the following times to find a solution - (a) 44s, (b) 200s, (c) 12s and (d) 56s. In each case this is significantly more than the time for any other strategy to reach the optimum. Figure best viewed in color.

where  $\Gamma_d = \xi_d \cdot (c_{\min}/2)^d$  is a constant. Using Eq. (3) we can bound the volume of the ellipse used at the  $i$ 'th iteration, where  $i \geq 1$ ,

$$\begin{aligned} \mu(\mathcal{X}_{c_{\text{best}}^i}) &\leq \varepsilon_i^{\frac{d-1}{2}} (1 + \varepsilon_0) (2 + \varepsilon_0)^{\frac{d-1}{2}} \Gamma_d \\ &\leq \eta^{-\frac{i(d-1)}{2}} \varepsilon_0^{\frac{d-1}{2}} (1 + \varepsilon_0) (2 + \varepsilon_0)^{\frac{d-1}{2}} \Gamma_d \quad (6) \\ &\leq 2^{-\frac{i(d-1)}{2d}} \mu(\mathcal{X}_{c_{\text{best}}^0}) \end{aligned}$$

Furthermore, we choose  $r_0$  such that  $\mu(\mathcal{X}_{c_{\text{best}}^0}) \leq \mu(\mathcal{X})$ . Now, the number of vertices in  $\mathcal{X}_{c_{\text{best}}^i}$  can be bounded by,

$$n_{i+1} = \frac{\mu(\mathcal{X}_{c_{\text{best}}^i})}{\mu(\mathcal{X})} n \leq 2^{-\frac{i(d-1)}{2d}} n. \quad (7)$$

Recall that we measure the amount of work done by the search at iteration  $i$  using  $|E_i|$ , the number of edges considered. Thus,

$$|E_i| = O(n_i^2 r_i^d) = O\left(n^2 2^{-\frac{i(d-1)}{d}} \left(r_0 2^{\frac{i}{d}}\right)^d\right) = O\left(n 2^{\frac{i}{d}}\right) \quad (8)$$

Finally, the total work done by the search over all iterations is

$$O\left(\sum_{i=0}^{\log_2 n} n 2^{\frac{i}{d}}\right) = O\left(n \sum_{i=0}^{\log_2 n} 2^{i/d}\right) = O\left(n^{1+\frac{1}{d}}\right). \quad (9)$$

## VII. EXPERIMENTS

Our implementations of the various strategies are based on the publicly available OMPL [5] implementation of BIT\* [9]. Other than the specific parameters and optimizations mentioned earlier, we use the default parameters of BIT\*. Notably, we use the Euclidean distance heuristic, an approximately sorted queue, and limit graph pruning to changes in path length greater than 1%.

### A. Random scenarios

The different batching strategies are compared to each other on problems in  $\mathbb{R}^d$  for  $d = 2, 4$ . The domain is the unit hypercube  $[0, 1]^d$  while the obstacles are randomly generated axis-aligned  $d$ -dimensional hyper-rectangles. All problems have a start configuration of  $[0.25, 0.25, \dots]$  and a goal configuration of  $[0.75, 0.75, \dots]$ . We used the first

$n = 10^4$  and  $n = 10^5$  points of the Halton sequence for the  $\mathbb{R}^2$  and  $\mathbb{R}^4$  problems, respectively.

Two parameters of the obstacles are varied to approximate the notion of problem hardness described earlier – the number of obstacles and the fraction of  $\mathcal{X}$  which is in  $\mathcal{X}_{\text{obs}}$ , which we denote by  $\zeta_{\text{obs}}$ . Specifically, in  $\mathbb{R}^2$ , we have easy problems with 100 obstacles and  $\zeta_{\text{obs}} = 0.33$ , and hard problems with 1000 obstacles and  $\zeta_{\text{obs}} = 0.75$ . In  $\mathbb{R}^4$  we maintain the same values for  $\zeta_{\text{obs}}$ , but use 500 and 3000 obstacles for easy and hard problems, respectively. For each problem setting ( $\mathbb{R}^2/\mathbb{R}^4$ ; easy/hard) we generate 30 different random scenarios and evaluate each strategy with the same set of samples on each of them. Each random scenario has a different set of solutions, so we show a representative result for each problem setting in Fig. 7.

The results align well with our intuition about the relative performance of the densification strategies on easy and hard problems. Notice that the naive strategy of searching  $\mathcal{G}$  with A\* directly requires considerably more time to report the optimum solution than any other strategy. We mention the numbers in the accompanying caption of Fig. 7 but avoid plotting them so as not to stretch the figures. Note the reasonable performance of hybrid batching across problems and difficulty levels.

### B. Manipulation problems

We also run simulated experiments on HERB [26], a mobile manipulator designed and built by the Personal Robotics Lab at Carnegie Mellon University. The planning problems are for the 7-DOF right arm, on the problem scenarios shown in Fig. 8. We use a roadmap of  $10^5$  vertices defined using a Halton sequence  $\mathcal{S}$  which was generated using the first 7 prime numbers. In addition to the batching strategies, we also evaluate the performance of BIT\* [9], using the same set of samples  $\mathcal{S}$ . BIT\* had been shown to achieve anytime performance superior to contemporary anytime algorithms. The hardness of the problems in terms of clearance is difficult to visualize in terms of the C-space of the arm, but the goal regions are considerably constrained. As our results show (Fig. 8), all densification strategies solve the difficult planning problem in reasonable time, and generally outperform the BIT\* strategy on the same set of samples.

## VIII. CONCLUSION AND FUTURE WORK

We present, analyze and implement several densification strategies for anytime planning on large  $E^d$  graphs. We




Fig. 8: We show results on 2 manipulation problems for ■ vertex batching ■ edge batching, ■ hybrid batching and ■ BIT\*. For each problem the goal configuration of the right arm is rendered translucent. Both of the problems are fairly constrained and non-trivial. The problem depicted in (c) has a large clear area in front of the starting configuration, which may allow for a long edge. This could explain the better performance of vertex batching. The naive strategy takes 25s for (b) and 44s for (d) respectively.

provide theoretical motivation for these densification techniques, and show that they outperform the naive approach significantly on difficult planning problems.

In this work we demonstrate our analysis for the case where the set of samples is generated from a low-dispersion deterministic sequence. A natural extension is to provide a similar analysis for a sequence of random i.i.d. samples. Here,  $f(|V|) = O(\log |V|)$  [14] instead of  $O(|V|)$ . When out of the starvation regions we would like to bound the quality obtained similar to the bounds provided by Eq. (1). A starting point would be to leverage recent results by Dobson et. al. [7] for Random Geometric Graphs under expectation, albeit for a *specific* radius  $r$ .

Another question we wish to pursue is alternative possibilities to traverse the subgraph space of  $\mathcal{G}$ . As depicted in Fig. 1, our densification strategies are essentially ways to traverse this space. We discuss three techniques that traverse relevant boundaries of the space. But there are innumerable trajectories that a strategy can follow to reach the optimum. It would be interesting to compare our current batching methods, both theoretically and practically, to those that go through the interior of the space.

#### REFERENCES

- [1] R. Bohlin and L. E. Kavraki. Path planning using lazy PRM. In *IEEE International Conference on Robotics and Automation*, pages 521–528, 2000.
- [2] M. S. Branicky, S. M. LaValle, K. Olson, and L. Yang. Quasi-randomized path planning. In *IEEE International Conference on Robotics and Automation*, pages 1481–1487, 2001.
- [3] S. Choudhury, J. D. Gammell, T. D. Barfoot, S. S. Srinivasa, and S. Scherer. Regionally accelerated batch informed trees (rabit\*): A framework to integrate local information into optimal path planning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4207–4214. IEEE, 2016.
- [4] B. J. Cohen, M. Phillips, and M. Likhachev. Planning single-arm manipulations with n-arm robots. In *RSS*, 2014.
- [5] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 2012.
- [6] C. M. Dellin and S. S. Srinivasa. A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors. In *International Conference on Automated Planning and Scheduling*, pages 459–467, 2016.
- [7] A. Dobson, G. V. Moustakides, and K. E. Bekris. Geometric probability results for bounding path quality in sampling-based roadmaps after finite computation. In *IEEE International Conference on Robotics and Automation*, 2015.
- [8] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot. Informed RRT\*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2997–3004, 2014.
- [9] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot. Batch informed trees (BIT\*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *IEEE International Conference on Robotics and Automation*, pages 3067–3074, 2015.
- [10] J. H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.*, 2(1):84–90, 1960.
- [11] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems, Science, and Cybernetics*, SSC-4(2):100–107, 1968.
- [12] L. Janson, B. Ichter, and M. Pavone. Deterministic sampling-based motion planning: Optimality, complexity, and performance. *CoRR*, abs/1505.00023, 2015.
- [13] L. Janson, E. Schmerling, A. Clark, and M. Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *I. J. Robotics Res.*, pages 883–921, 2015.
- [14] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *I. J. Robotics Res.*, 30(7):846–894, 2011.
- [15] L. E. Kavraki, M. N. Kolountzakis, and J. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. Robotics and Automation*, 14(1):166–171, 1998.
- [16] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. Robotics*, 12(4):566–580, 1996.
- [17] S. Koenig, M. Likhachev, and D. Furcy. Lifelong planning A\*. *Artificial Intelligence*, 155(1):93–146, 2004.
- [18] R. E. Korf. Iterative-deepening-A\*: An optimal admissible tree search. In *Joint Conference on Artificial Intelligence*, pages 1034–1036, 1985.
- [19] S. M. LaValle and J. J. Kuffner. Randomized kinodynamic planning. In *IEEE International Conference on Robotics and Automation*, pages 473–479, 1999.
- [20] M. Likhachev, G. J. Gordon, and S. Thrun. ARA\*: Anytime A\* with provable bounds on sub-optimality. In *Advances in Neural Information Processing Systems*, pages 767–774, 2003.
- [21] H. Niederreiter. *Random Number Generation and quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, 1992.
- [22] O. Salzman and D. Halperin. Asymptotically-optimal motion planning using lower bounds on cost. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 4167–4172. IEEE, 2015.
- [23] O. Salzman and D. Halperin. Asymptotically near-optimal RRT for fast, high-quality motion planning. *IEEE Trans. Robotics*, 32(3):473–483, 2016.
- [24] K. Solovey, O. Salzman, and D. Halperin. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *I. J. Robotics Res.*, 35(5):501–513, 2016.
- [25] K. Solovey, O. Salzman, and D. Halperin. New perspective on sampling-based motion planning via random geometric graphs. In *RSS*, 2016.
- [26] S. S. Srinivasa, D. Ferguson, C. J. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. J. Kuffner, and M. V. Weghe. HERB: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.
- [27] J. van den Berg, R. Shah, A. Huang, and K. Y. Goldberg. Anytime nonparametric A. In *Association for the Advancement of Artificial Intelligence*, pages 105–111, 2011.
- [28] C. M. Wilt and W. Ruml. When does weighted A\* fail? In *Symposium on Combinatorial Search*, pages 137–144, 2012.