

Lazy Validation of Experience Graphs

Victor Hwang*, Mike Phillips*, Siddhartha Srinivasa*, Maxim Likhachev*

Abstract—Many robot applications involve lifelong planning in relatively static environments e.g. assembling objects or sorting mail in an office building. In these types of scenarios, the robot performs many tasks over a long period of time. Thus, the time required for computing a motion plan becomes a significant concern, prompting the need for a fast and efficient motion planner. Since these environments remain similar in between planning requests, planning from scratch is wasteful. We propose using Experience Graphs (E-Graphs) to accelerate the planning process by reusing parts of previously computed paths to solve new motion planning queries more efficiently. This work describes a method to improve planning times with E-Graphs given changes in the environment by lazily evaluating the validity of past experiences during the planning process. We show the improvements with our method in a single-arm manipulation domain with simulations on the PR2 robot.

I. INTRODUCTION

In situations where a robot is doing lifelong planning, the efficiency of a motion planner becomes very important. However, in tasks such as assembling objects or performing sorting and pick-and-place tasks, motion planners still suffer from poor planning times depending on the complexity of the environment. In a robot manipulation environment, it is simple to find standard scenarios where planning takes a non-trivial amount of time, despite the fact that these environments are relatively structured.

In this work, we leverage learning from prior experiences to tackle difficult planning scenarios where a robot’s movement in a structured environment is encumbered by unpredictable clutter. Figure 1 shows a mailroom environment, where the general task of sorting mail is a structured one. However, this environment typically has clutter appearing inside cubby holes or on shelves. By learning from experiences, we expect a mail-sorting robot to accumulate enough experience over time so that planning to constrained spaces (cubbyholes, boxes, etc) becomes very fast. At the same time, we expect this framework to be robust against the fact that portions of its prior experiences are constantly becoming invalid due to clutter.

Our work modifies an existing Experience Graph planner [15] to provide fast planning times in this kind of challenging environment. The Experience Graphs (E-Graphs) framework acts as a collection of previously planned paths and augments A*-like searches to guide a search to reuse these past experiences. In addition, the planner accumulates experiences throughout its lifetime by feeding back solutions. This results in faster planning times while maintaining completeness guarantees and a bound on the sub-optimality of the solution with respect to the optimal path.

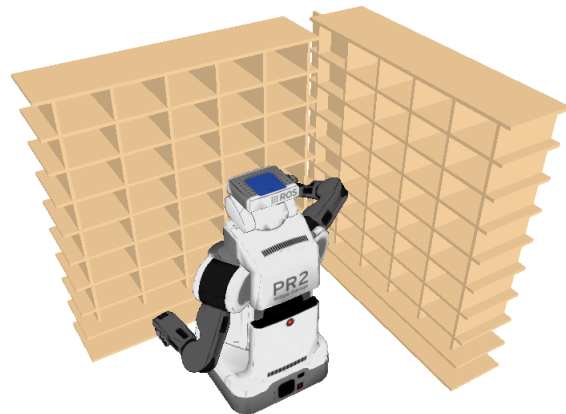


Fig. 1: PR2 in a simulated mailroom environment.

While [15] provides an efficient planning framework for a completely static environment with a large number of prior experiences, planning times suffer once the possibility of unpredictable clutter is introduced. The E-Graph planner is only able to plan with a completely valid set of experiences; even a negligible change in the environment forces the planner to validate its entire set of experiences to ensure a valid plan. Considering the robot may be working tirelessly for long periods of time, the E-Graph can grow substantially and introduce more and more overhead into every single planning request.

In this paper, we address one facet of the problem that affects any data-driven algorithm: *does it scale?* We show efficient experience validation by introducing a general method that lazily validates prior experiences that the planner deems potentially relevant to the planning task. This method leads to performance improvements in changing, cluttered environments. Specifically, we show speedups using a simulated environment where the PR2 conducts single arm, 7D manipulation tasks in a constrained mailroom environment.

II. RELATED WORK

Motion planning has seen considerable interest recently. Most of the approaches were initially focused on treating each new motion planning request as a fresh request for planning. There was little, if any, reuse of information from experience gained while carrying out a series of motion plans. However, recent work has seen more reuse of previous information. In [12], Lien et. al. constructed roadmaps for obstacles, stored them in a database and reused them during

* Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

motion planning while Bruce et. al. [4] biased an RRT search towards waypoints remembered from previous motion planning attempts. Related work can also be found in [18], [8].

Other approaches to exploiting experience have included the use of trajectory libraries. These libraries were used to adapt policies for control of underactuated systems and high-dimensional systems in [16]. In [1], entirely new trajectories were generated by combining nearby trajectories. A different application of such techniques can also be found in [13] while transfer of policies across tasks was presented in [17]. The reuse of environment information, coupled with information about previous motion plans, was presented in [9]; machine learning methods were coupled with paths stored in a database to generate new motion plans based on the environment and types of obstacles. Jiang et. al. [10] present an approach to use a database of older motion plans to draw a bi-directional RRT search towards a path stored in the database that is most similar to the new motion plan request. Recent work [2] attempts to repair previous plans from a database using randomized planners. As mentioned in [15], the use of a database of motion plans is also core to Experience Graphs.

Search-based planning with Experience Graphs offers several advantages over these approaches. Experience Graphs attempt to use all information from previous experiences instead of attempting to find the nearest or closest motion plan from a database. They are thus capable of reusing many parts of the previous experiences when possible. This method provides guarantees on completeness and solution quality, which the other methods lack. It also builds the graph using paths from prior tasks in contrast to approaches like Probabilistic Roadmaps [11] which rely on sampling the whole space. At the same time, there are parallels between the Lazy Probabilistic Roadmap [3] and the work we present in that both methods only do expensive validation once a plan has been generated.

III. EXPERIENCE GRAPHS

This section provides a brief overview of the E-Graphs framework. For a more in-depth analysis, see [15].

A. Terminology

We first introduce a brief set of used terminology.

- $G(V^G, E^G)$ is a graph representation of the original motion planning problem. V^G is a set of vertices, and E^G is a set of edges connecting vertices.
- $G^E(V^E, E^E)$ is a subgraph of G that represents the E-Graph.
- h^G and h^E represent the original graph heuristic function and the E-Graph heuristic function, respectively.
- $c(s, s')$ describes the scalar weight of an edge between two states s and s' .
- A *successor* describes a child state $s_{successor}$ which can be directly reached from its parent state s_{parent} .
- An *expansion* of a state involves computing all valid successors for a given parent state.

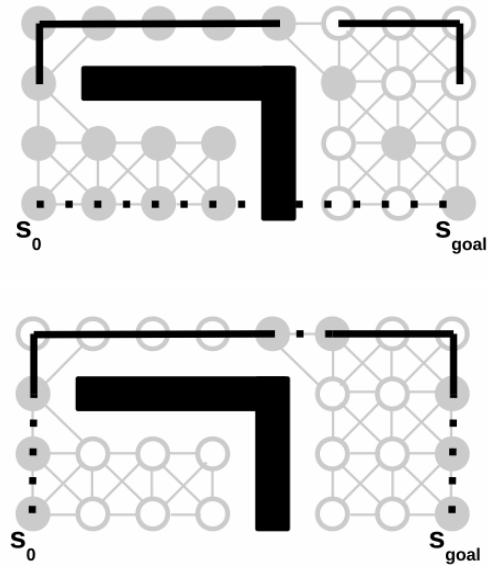


Fig. 2: The above two images show the differences between heuristics when varying ϵ^E . When ϵ^E is close to one, the heuristic effectively becomes h^G (in this example, Euclidean distance), as seen in the top image. The dashed line indicates that it is a heuristic edge. As ϵ^E increases, the heuristic guides along previous experiences, as shown by the solid lines in the bottom figure.

B. Algorithm Description

E-Graphs augment A* search to reduce search computation by reusing relevant pieces of prior experience. It is composed of two aspects - the Experience Graph that maintains all the prior experience, and a heuristic that drives a weighted A* search towards these prior experiences during the planning process.

The Experience Graph G^E maintains a collection of previously planned paths as a set of edges E^E and vertices V^E . While weighted A* still searches the original graph G , a new heuristic function $h^E(s)$ is introduced that encourages the search to explore G^E rather than G :

$$h^E(s_0) = \min_{\pi} \sum_{i=0}^{N-1} \min\{\epsilon^E h^G(s_i, s_{i+1}), c^E(s_i, s_{i+1})\} \quad (1)$$

In Equation 1, $h^E(s)$ is defined as computing the cost of the shortest path π where π is a path $\langle s_0 \dots s_{N-1} \rangle$ and $s_{N-1} = s_{goal}$. This path is constructed using two kinds of edges. The first edge cost, $\epsilon^E h^G(s_i, s_{i+1})$, represents the original search heuristic $h^G(s, s')$ and describes the underestimating cost between states s and s' for all states in G . The second term represents the edge cost $c^E(s, s')$ between states on G^E . The scalar ϵ^E (which is ≥ 1) is a parameter to inflate the cost of the first term to favor the second. This effectively biases the minimum path to avoid using edges not on the E-Graph and instead, favor using edges from previous experiences.

The qualitative effect of ϵ^E can be seen in Figure 2. This image shows the search from s_0 to s_{goal} and the required




Fig. 3: An example of shortcut successors (arrows) being used to avoid expanding unnecessary states on an E-Graph component.

state expansions to find the solution. In both images, the grey circles represent states expanded during the search and the solid black line represents a E-Graph edges of a prior experience. We use Euclidean distance as the original heuristic h^G , which is represented as the dashed line. When the ϵ^E is close to one, h^E draws the search directly towards the goal with a heuristic edge (dashed line) and avoids using any E-Graph edge. However, this heuristic guidance isn't particularly helpful, since the search is led straight into an obstacle (black rectangles). We see that many states end up being expanded (shown in gray). As ϵ^E increases (bottom image), h^E instead guides the search towards the prior experience (black lines), which effectively dodges the obstacle and completes the search with fewer states expanded.

In the standard weighted A* search, we can use an admissible heuristic that is inflated by some scalar factor ϵ^G . This allows us to find a bounded, sub-optimal path that is within some ϵ^G multiple of the optimal solution cost. In [15], it was shown that the E-Graph heuristic is ϵ^E consistent. Therefore, when using this heuristic with weighted A* search (whose heuristic is normally inflated by an ϵ^G factor), it is guaranteed that the solution cost will be no worse than $\epsilon^G \cdot \epsilon^E$ times the optimal solution cost, rather than just ϵ^G .

C. Experience Graph Shortcuts

One important element to the E-Graph algorithm is the use of shortcuts. E-Graph shortcuts are a mechanism for further reducing the number of expansions by introducing a new type of successor $s_{shortcut}$ to the search. Given a connected component C where $C \subseteq V^E$

$$s_{shortcut} = \arg \min_{s \in C} h^G(s) \quad (2)$$

When expanding any state $s \in C$, $s_{shortcut}$ becomes an additional valid successor of s , allowing the search to avoid expanding intermediate states in C .

In Figure 3, we see two components that are not directly connected to the goal. For each component, we compute a shortcut state as per Equation 2, generating the two labeled nodes. The arrows show an example of expanding one node on the component, which has the shortcut node as a successor. This leads to the planner avoiding expanding states

along a particular E-Graph component by jumping right to the end.

The utility of this successor can be seen more clearly when similar start/goal pairs are repeatedly solved by the planner (assuming a valid E-Graph). The first instance of planning for a given pair will result in many expansions. However, subsequent expansions will only require expanding a few additional node besides the start, and will immediately end up at the goal state, regardless of how many E-Graph edges lie in between the start and goal.

IV. LAZY VALIDATION OF E-GRAPHS

This section describes the general method for lazy validation of E-Graphs, along with two possible design choices for its implementation (on-the-fly validation and post-validation). These two design options offer trade-offs between speed and path quality and can be exercised when implementing a planner for a specific robotic domain. We present the post-validation version first to introduce the general idea of lazy validation, and then present the subtleties behind post-validation and on-the-fly validation.

A. Lazy Post-Validation

Lazy post-validation is built on top of the original E-Graph planner described in [15]; it uses the same heuristic computation to focus the search towards prior experiences. However, rather than maintaining E-Graph feasibility and validity as an invariant, we loosely assume that the E-Graph is valid until the E-Graph planner returns a complete solution. Only then do we validate the vertices and edges in the solution path that lie on the E-Graph. During this validation, we also know that the only E-Graph edges that could be invalid are those used in the E-Graph shortcut - any other type of successor would be validated during the search.

Algorithm 1

lazyPlan(s_{start}, s_{goal}, G^E)

- 1: while not valid π do:
 - 2: initializeEGraph(s_{goal})
 - 3: $\pi = \text{computePath}(s_{start}, s_{goal})$
 - 4: $G^E = G^E \cup \pi$
 - 5: $(V_{invalid}, E_{invalid}) = \text{returnInvalidElements}(\pi)$
 - 6: $G^E = G^E \setminus (V_{invalid}, E_{invalid})$
 - 7: if G^E is empty:
 - 8: return π
-

Algorithm 2

initializeEGraph(s_{goal})

- 1: precomputeShortcut
 - 2: compute heuristic according to Equation 1
-

Algorithm 1 shows a high level overview of the logic for lazy post-validation. Given a pre-built E-Graph G^E , we run the standard procedure for E-Graph planning: initializing relevant E-Graph mechanisms (Algorithm 2), which pre-computes the shortcuts and the E-Graph heuristic. Once we compute a path π , we feed it back into the E-Graph in

order to further accumulate experiences for later use. In *returnInvalidElements*, we validate elements of G^E related to π and return those that are invalid. In line 7, we remove these from G^E , and rerun the planning cycle. Given that the E-Graph contains a sufficient amount of experience and the environment has not change dramatically, the number of total replans remains low.

The goal of *returnInvalidElements* goes beyond discovering points on the returned path that are invalid. Because the while loop continues to run until a valid path is returned, it is advantageous to discover as many invalid edges and vertices in G^E as possible in a single run of *returnInvalidElements*, especially those that are not explicitly part of π . In the ideal case, one could run a pre-computation to determine vertex connectivity in terms of the validation procedure., i.e. return all $V_{invalid}$ given that a particular vertex is invalid. However, this is likely computationally expensive, so faster approximations can be used.

In our implementation of lazy E-Graphs for robot arm planning, we check that the 3D location of the end effector for $V_{invalid}^E$ is not an obstacle. If it is, we look up all other V^E that result in the same end effector position, and invalidate those as well.

An example run of the post-validation method can be seen in Figure 4. Two replans are required before a completely valid path is found in the fourth image.

B. On-the-Fly Validation

In the previous method, when a path was found, it needed to be checked for validity. The only parts of a path that could be invalid were the shortcut successors, which use the unverified E-Graph edges. In on-the-fly validation, we check the edges in a shortcut successor $s_{shortcut}$ (as described in III-C) during the search as soon as the shortcut is generated. Therefore, when a path is found, it is guaranteed to be valid. There is no need to check for validity, and the loop in algorithm 1 is no longer needed.

C. Theoretical Properties

The general idea of lazy validation still maintains theoretical guarantees described in [15]. Specifically, there is an upper limit on the solution cost with respect to the optimal cost.

Theorem 1: For a finite graph G , and finite graph G^E , the planner terminates and finds a path in G if one exists.

Theorem 2: For a finite graph G , and finite graph G^E , the planner terminates and the solution returned is no worse than $\epsilon \cdot \epsilon^E$ times the optimal solution cost in G .

D. Post-Validation vs. On-the-Fly Validation

We now explore the subtle differences between these two implementations. As presented above, the post-validation algorithm validates the E-Graph paths only after a path has been successfully found. However, in certain domains, there can be substantial overhead in completely restarting the search (recomputation of heuristics, reinitializing data structures, etc.), especially if the environment changes drastically, resulting in many invalid E-Graph edges and many replans.

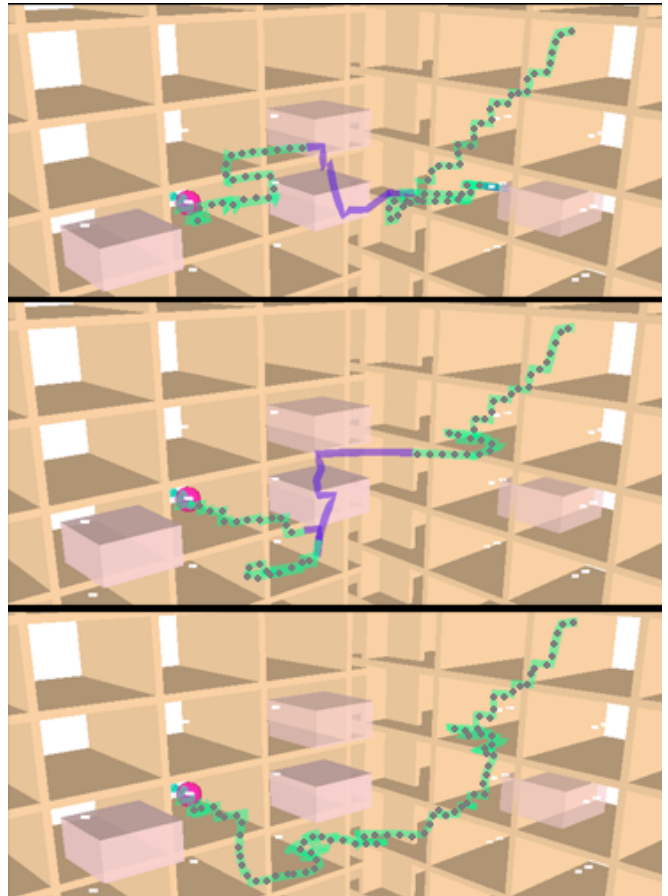


Fig. 4: Three iterations of the post-validation method. While the plans are 7-dimensional, the pictures show their corresponding end-effector trajectories. The dark regions in the top two images show the invalid regions of the end effector path, while the dotted lines are the valid regions of the path. The last image shows a final (without shortcutting) path of the end effector.

One option to avoid this overhead is to validate the E-Graph edges during the search process. Specifically, this would occur whenever an E-Graph shortcut successor is generated. If the shortcut is found to be invalid, the successor is thrown out, and the search proceeds in the standard fashion. This avoids the overhead problem associated with the post-validation method because the search does not have to restart when this occurs. However, the major caveat is that the E-Graph structure changes in the middle of the search. In the case where the E-Graph heuristic is pre-computed (which is standard practice for heuristics used in robot manipulation and other high dimensional domains), there is now a potential problem for the heuristic to become extremely misinformed.

This is illustrated in detail in Figure 5. We begin with a two paths leading to the goal. When we introduce an obstacle that obstructs one of the prior experiences, a segment of the experience is taken out. At this point, we hope that the method will use the second path to the goal to quickly find a solution. Because the post-validation method reruns the




Fig. 5: A scenario where on-the-fly validation will fail due to a misinformed heuristic, whereas post-validation will succeed. Before any invalidation occurs, the heuristic guides the search along the bottom path. When both methods discover that the experience passes through the obstacle, the corresponding edges are invalidated. Post-validation will recompute the heuristic and guide the search towards the top segment of experience. On-the-fly, however, will not recompute the heuristic, and get stuck trying to bypass the obstacle.

heuristic computation every time the heuristic is invalidated, the search will lead to the second path. The on-the-fly method’s heuristic, however, becomes uninformed because the heuristic is still drawn to follow the E-Graph as if the invalid segment still connected to the goal.

The correct choice between these two methods is very application dependent - if the overhead of recomputing the heuristic is negligible compared to overall planning time, then post-validation remains the better choice because the heuristic will be better informed and result in higher path quality for the solution. However, if the average planning time is low enough that the overhead becomes significant, then on-the-fly validation can be used to improve planning times at the expense of success rate due to pathologically bad cases. In our experiments, these bad cases almost never occur.

V. EXPERIMENTAL RESULTS

We tested our approach in a simulated mailroom environment, where the PR2 robot is sorting objects between cubby holes, as seen in Figure 1. We use an existing benchmarking framework to run all experiments [6]. All the tasks involve manipulation with a single 7DOF arm of the PR2. The goals are specified as the position and orientation of the end effector.

We first build up a usable E-Graph by running 200 trials of random start goal pairs between the 18 cubby holes in the environment. These 200 trials are run with a very low bound ($\epsilon = 1.5$) to ensure high quality experiences between locations. This results in an E-Graph of near 7000 vertices. Figure 6 shows a visualization of the E-Graph - each vertex represents the 3D location of the end effector of a particular V^G and the edges represent end effector motions.

For testing, we run 50 trials where each trial introduces six randomly placed obstacles inside the cubbies. We plan using three different E-Graph based methods and the OMPL




Fig. 6: An E-Graph with 7000 vertices. Each vertex represents the end effector location of a state, and edges represent the 3D movement of the end effector.

	Post	Full	On-the-fly	RRTConnect
Median CC/request	5280	54235	8304	191248
Average CC/request	14222	54245	11574	221606

TABLE I: Collision check count comparison between three E-Graph methods and RRTConnect. All E-Graph methods are initialized with an E-Graph of 7000 vertices.

[7] implementation of the RRT. We do not report any results for Probabilistic Roadmaps or RRT* since they were both unable to successfully plan in this scenario. For the E-Graph based methods, we plan with a higher ϵ and ϵ^E , allowing for more reuse of experiences. All paths undergo shortcutting in a post-processing phase. After each trial, the E-Graph is reloaded with the 200 trial training set. All statistics are computed on trials where all the planners succeed.

We compare the two lazy validation (post and on-the-fly) methods with the naive solution: full validation. In full validation, we validate every vertex and edge in G^E once a planning request has been received. This computation is factored into the planning time. In all algorithms, the validation of any E-Graph edge is equivalent to running a full collision check of the robot arm against the environment.

The first experimental result we present is statistics on the number of collision checks that each planner issued. Because collision checking generally takes up a large portion of planning times, this offers an implementation-independent view of performance (compared to planning times, which depend on the particular collision checking implementation). Table I shows both the median and average number of collision checks for all planning methods. We see that the post and on-the-fly methods have lower average collision checks overall. We expect full validation to have a high average, because it’s directly proportional to the number of E-Graph vertices. RRTConnect has a difficult time in this environment, with an order of magnitude more collision checks due to the narrow passages from the cubbies.

	Post	Full	On-the-fly	RRTConnect
Fast CC	1.45	1.24	.97	8.99
Medium CC	2.44	6.83	1.76	42.3
Slow CC	7.16	29.9	5.09	98.4
Success Rate	76%	80%	73%	84%

TABLE II: Average planning time in seconds using different speeds of collision checkers. The “fast” collision checker takes 2.0×10^{-5} sec/collision check. The “medium” collision checker takes 1.1×10^{-4} sec/collision check. The “slow” collision checker takes 4.9×10^{-4} sec/collision check. As a reference, the SBPL approximate collision checker is the “fast” checker. The FCL collision checker is slightly slower than our “medium” collision checker at 2.4×10^{-4} sec/collision check.

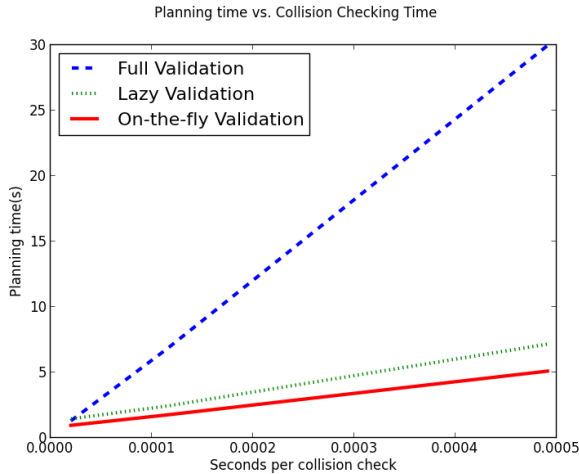


Fig. 7: A comparison between the post-validation and full validation method over different speeds of collision checking. These results were generated by artificially slowing the SBPL collision checker performance to various speeds.

In order to get a sense of planning times, one experiment was run to demonstrate the planning times for different speeds of collision checkers. We begin with the SBPL collision checker, which operates very quickly because of its sphere-based approximate model of the robot. However, in domains where a robot is manipulating objects, an exact collision checking model may be required, such as the Flexible Collision Library [14].

To get a sense of how the E-Graph methods scale with time required for a single collision check, we artificially increase the amount of computation in the SBPL collision checker [5] to various speeds and measure the required planning time. Figure 9 shows a plot of the planning times, and Table II shows the times for several speeds of the collision checker. Overall, we see the clear relationship that, as collision checking time increases, full validation linearly increases in planning time, while post-validation and on-the-fly validation do not suffer. As a reference, FCL takes roughly 2.4×10^{-4} seconds per collision check. Interpolating on our results, FCL would spend about 12 seconds to plan

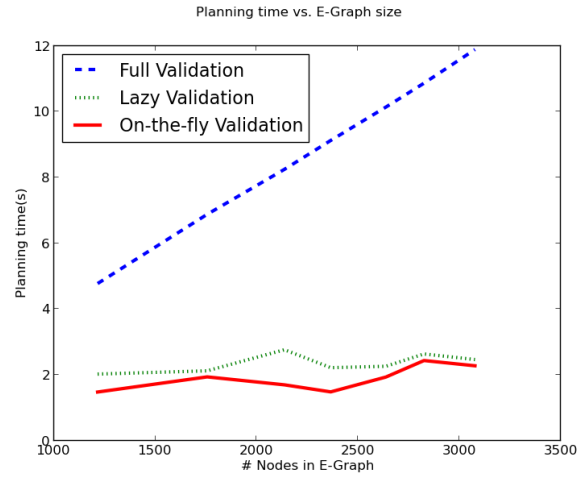


Fig. 8: A comparison between the post, on-the-fly, and full validation methods over different sizes of E-Graphs. We see that the full validation method planning time increases linearly, while post and on-the-fly validation remain relatively constant.

for full validation, while post-validation still remains close to two seconds.

We see that the on-the-fly is slightly faster in planning time than the post-validation method due to the reduced overhead associated with restarting a search and recomputing heuristics. However, in terms of path quality, on-the-fly does slightly worse, as seen in Table III. This is expected, since the quality of the heuristic with the on-the-fly method is reduced as more E-Graph edges are obstructed by obstacles.

In Figure 8, we show how the planning times for each method changes with the size of the E-Graph measure by the number of nodes it contains. We see that, without the use of lazy validation, the E-Graph planner does not scale well - planning times increase linearly with the size of the E-Graph. Post-validation and on-the-fly are on-par in performance, though vary with the randomness of the E-Graph.

A. Post-Validation vs On-the-Fly Analysis

In the following experiment, we engineer a situation where on-the-fly fails, while post-validation succeeds. Here, a small set of experiences have been added to the environment, and the robot is planning from the top half of the cabinet to the bottom half.

As explained before, the plan here fails due to structural changes in the E-Graph, which causes the heuristic to become misinformed. Specifically, we include a block that intentionally disrupts the original end effector path from the top to bottom of the block, which renders the heuristic useless, and the plan is unable to finish.

In Figure 9a, we introduce the obstacle, which disrupts the original experience. Figure 9b reflect the state of both algorithms after updating the E-Graph structure - with post-validation, this occurs once the full path has been found, and for on-the-fly, it happens during the search. Post-validation

	Post	Full	On-the-fly	RRTConnect
Average Path Quality	1.60	1.69	1.72	2.15

TABLE III: Path quality comparison between three E-Graph methods and RRTConnect. All E-Graph methods are initialized with an E-Graph of 7000 vertices.

succeeds the second time after recomputing the heuristic, which generates the path seen in Figure 9c. On-the-fly validation fails to plan because its heuristic continues to drive the search into the block, unaware that the environment has changed.

VI. CONCLUSION

In this work, we have presented a general method for leveraging a large amount of prior experiences to accelerate planning in scenarios with random clutter. We augment the existing E-Graph planning framework to lazily validate experiences in the E-Graph, which allowed our method to only expend computation on validating portions of prior experience that are relevant to the particular planning episode. We show experimental results that demonstrate the speedup in a mailroom environment compared to a number of other possible methods.


In future work, we will explore obstacle modeling methods to decrease the number of overall replans required when many obstacles are added to the environment. In addition, we will look into possible offline computations to improve the accuracy of edge invalidations along with pruning or improving the quality of the E-Graph when there is spare time.

VII. ACKNOWLEDGEMENTS


We thank Google for their support of these efforts. This research was also sponsored by ONR grant N00014-09-1-1052. We also thank Ben Cohen for his invaluable input.

REFERENCES


- [1] C. Atkeson and J. Morimoto. Nonparametric representation of policies and value functions: A trajectory-based approach. In *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- [2] Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. A robot path planning framework that learns from experience. In *ICRA*, 2012.
- [3] Robert Bohlin and EE Kavraki. Path planning using lazy prm. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 1, pages 521–528. IEEE, 2000.
- [4] J. Bruce and M. Veloso. Real-time randomized path planning for robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.
- [5] Benjamin Cohen, Sachin Chitta, and Maxim Likhachev. Single- and dual-arm motion planning with heuristic search. *The International Journal of Robotics Research*, page 0278364913507983, 2013.
- [6] Benjamin Cohen, Ioan Alexandru Sucan, and Sachin Chitta. A generic infrastructure for benchmarking motion planners. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 589–595. IEEE, 2012.
- [7] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 2012. To appear.
- [8] D. Ferguson, N. Kalra, and A. T. Stenz. Replanning with rrt. In *IEEE International Conference on Robotics and Automation*, May 2006.



(a) An existing E-Graph is disrupted by an obstacle.



(b) The E-Graph structure once both methods (post-validation and on-the-fly) have been updated to reflect the new environment.



(c) The post-validation solution, which succeeds because the heuristic is updated. On-the-fly fails to plan in this situation.

Fig. 9

- [9] Nikolay Jetchev and Marc Toussaint. Trajectory prediction: Learning to map situations to robot trajectories. In *IEEE International Conference on Robotics and Automation*, 2010.
- [10] Xiaoxi Jiang and Marcelo Kallmann. Learning humanoid reaching tasks in dynamic environments. In *IEEE International Conference on Intelligent Robots and Systems*, 2007.
- [11] L. E. Kavraki, P. Svestka, J.-C Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [12] J. Lien and Y. Lu. Planning motion in environments with similar obstacles. In *Proceedings of the Robotics, Science and Systems Conference*, 2005.
- [13] C. Liu and C. G. Atkeson. Standing balance control using a trajectory library. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009.
- [14] Jia Pan, Sachin Chitta, and Dinesh Manocha. Fcl: A general purpose library for collision and proximity queries. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3859–3866. IEEE, 2012.
- [15] Mike Phillips, Benjamin Cohen, Sachin Chitta, and Maxim Likhachev. E-Graphs: Bootstrapping Planning with Experience Graphs. In *Proceedings of the Robotics, Science and Systems Conference*, 2012.
- [16] M. Stolle and C. Atkeson. Policies based on trajectory libraries. In *IEEE International Conference on Robotics and Automation*, 2006.
- [17] M. Stolle, H. Tappeiner, J. Chestnutt, and C. Atkeson. Transfer of policies based on trajectory libraries. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.
- [18] M. Zucker, J. Kuffner, and M. Branicky. Multipartite rrt* for rapid replanning in dynamic environments. In *IEEE International Conference on Robotics and Automation*, 2007.