

Efficient and Reconfigurable Optimal Planning in Large-Scale Systems Using Hierarchical Finite State Machines

Elis Stefansson¹ and Karl H. Johansson¹

Abstract—In this paper, we consider a planning problem for a large-scale system modelled as a hierarchical finite state machine (HFSM) and develop a control algorithm for computing optimal plans between any two states. The control algorithm consists of two steps: a preprocessing step computing optimal exit costs for each machine in the HFSM, with time complexity scaling linearly with the number of machines, and a query step that rapidly computes optimal plans, truncating irrelevant parts of the HFSM using the optimal exit costs, with time complexity scaling near-linearly with the depth of the HFSM. The control algorithm is reconfigurable in the sense that a change in the HFSM is efficiently handled, updating only needed parts in the preprocessing step to account for the change, with time complexity linear in the depth of the HFSM. We validate our algorithm on a robotic application, comparing it with Dijkstra’s algorithm and Contraction Hierarchies. Our algorithm outperforms both.

I. INTRODUCTION

A. Motivation

In today’s smart societies, with increased integrability and connectivity, there is an increasing need for efficient and optimal control of large-scale systems. Here, control algorithms should not only compute optimal plans for such systems efficiently but also be reconfigurable, i.e., be able to handle changes in parts of the system with ease.

A common approach is to consider modular systems, i.e., systems that can be decomposed into independent entities (modules). For such systems, the idea is to construct control algorithms that can separate the computation into these modules and then combine the results from the modules to obtain an optimal plan for the whole system. Moreover, a change in a modular system typically affects only some of the modules, enabling control algorithms to be reconfigurable by only updating the corresponding affected parts in the algorithm.

In this paper, we consider large-scale systems formalised by hierarchical finite state machines (HFSMs) [8]. An HFSM is a machine composed of several finite state machines (FSMs) nested into a hierarchical structure. This structure naturally makes an HFSM modular. The key research question is how to exploit the modularity in the HFSM to construct optimal control algorithms that are efficient and reconfigurable.

¹School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Sweden. Email: {elisst, kallej}@kth.se. The authors are also affiliated with Digital Futures.

This work was partially funded by the Swedish Foundation for Strategic Research, the Swedish Research Council, and the Knut and Alice Wallenberg foundation.




Fig. 1: Robot application example modelled as an HFSM with three layers.

B. An Illustrative Example

Consider an example where a robot is moving between lab houses, schematically depicted in Fig. 1. In this environment, the robot is assigned to scan a tube, located in one of the houses, with minimal operational cost. Therefore, the robot must compute a plan that with minimal cost makes it reach the tube and scan it.

Here, a key insight is that the environment can be naturally modelled as an HFSM with three layers corresponding to the layers in Fig. 1, where Layer 1 prescribes how the robot can move between the houses, Layer 2 how the robot moves inside a house, and Layer 3 how the robot can operate the lab desk and scan tubes at a given location inside a house (all formally modelled by FSMs). The first key question is then *how to construct a control algorithm that computes an optimal plan efficiently, by exploiting the hierarchical structure of the HFSM*.

Suppose now that one of the houses changes, for example, some locations in House 2 are blocked due to maintenance work, as depicted by grey areas in Fig. 2. This changes the




Fig. 2: Change in House 2 in the HiMM given by Fig. 1. Grey areas depict blocked locations.

HFSM and may thus affect the optimal plan computed by the robot (e.g., if the computed optimal plan was to go to the bottom-right location in House 2, then the robot must take a detour now to avoid the blocked locations). However, only a part of the HFSM has been changed and hence, the control algorithm of the robot should not need to reset all its computations, rather, it should only need to update the relevant parts. That is, the control algorithm should be reconfigurable. The second key question is then *how to construct a control algorithm that is reconfigurable given changes in the HFSM*.

C. Contribution

The main contribution of this paper is to formalise optimal planning in systems modelled as HFSMs, introduce changes in such systems, and provide an efficient and reconfigurable control algorithm for computing optimal plans. More precisely, our contribution is three-fold:

First, we formalise the HFSM systems extending the setup of [2] to machines with outputs, i.e., Mealy machines (MMs) [11], calling the resulting HFSM for a Hierarchical Mealy Machine (HiMM), following the formulation in [16]. We then formalise changes in HiMMs, called modifications, which can change any HiMM into any other HiMM (by a series of modifications), and specify what we mean for an algorithm to be reconfigurable with respect to modifications.

Second, we propose a control algorithm that is both efficient and reconfigurable, extending the algorithm in [16] to be reconfigurable. The algorithm is divided into two steps: A preprocessing step where optimal exit costs are computed for all MMs in the given HiMM Z (done only once for a fixed HiMM), with time complexity $O(N)$, where N is the number of MMs in Z ; and a query step, rapidly computing an optimal plan between any two states in Z using the optimal exit costs, with time complexity $O(\text{depth}(Z) \log(\text{depth}(Z)))$ for computing the next optimal input (compared to Dijkstra's algorithm [5], [6] with time complexity more than exponential in $\text{depth}(Z)$ in general) and is therefore regarded as efficient. Here, $\text{depth}(Z)$ is the number of layers in the HiMM Z . The algorithm is reconfigurable in the sense that each modification of Z only needs an update of the preprocessing step in time $O(\text{depth}(Z))$, instead of recomputing the whole preprocessing step. In all these time complexities, we have, for brevity, assumed bounds on the number of states and

inputs in an MM of Z , see Sections III and IV for the general expressions without this assumption.

Third, we validate our control algorithm on the robot application given in the motivation, comparing it with Dijkstra's algorithm [5] and Contraction Hierarchies [7]. We show that our control algorithm outperforms both.

D. Related Work

HFSMs are typically used to model the control law of an agent [8], [12], [14], prescribing how the agent reacts to external inputs (e.g., "low battery"), where FSMs in the HFSMs commonly corresponds to subtasks (e.g., "charge battery"). Here, we instead use an HFSM formalism to model a control system where the agent can choose the inputs, for example, the robot can choose to go left in House 2 in the robot application in the motivation. The aim is to steer the system (e.g., the robot) to a desirable state with minimal cost. In discrete-event systems [3], another HFSM formalism, called state tree structures, has been developed to compute safe executions [9], [18]. We use a different formalism and focus instead on optimal planning minimising cost.

There is a wide range of algorithms for optimal path planning on graphs such as Dijkstra's algorithm [5], see [1] for a survey. Here, a common approach is to use a preprocessing step to achieve faster planning at run-time [1], similar to our preprocessing step. In particular, hierarchical algorithms [4], [7], [10], [13] using a preprocessing step are the ones most reminiscent of our approach. However, these approaches are typically based on heuristics which in the worst case perform no better than Dijkstra's algorithm. Furthermore, to use those methods, one also needs to flatten the HFSM to an equivalent FSM. This hides the hierarchical structure making these algorithms less suitable for reconfiguration. In this paper, we instead exploit the hierarchical structure to give formal performance guarantees and make the algorithm reconfigurable to changes. The paper [17] considers a variant of an HFSM formalism and seeks an execution with the minimal length between two configurations. In this paper, we have non-negative costs, where minimal length is a special case by setting all costs to unit costs.

In [2], HFSMs without outputs are formalised, and used to decompose an FSM into an HFSM. However, planning was not considered. In this paper, we extend their formalism to HFSMs with outputs and consider optimal planning.

Finally, this paper builds on the work [16]. We use the same HFSM formalism as in [16] and extend it to formally define changes in such a system, called modifications. We then extend the control algorithm in [16] to be reconfigurable, able to handle modifications efficiently.

E. Outline

The outline of the remaining paper is as follows. Section II formalises HiMMs, including modifications of HiMMs, and formulates the problem statement. Section III presents the efficient control algorithm, and Section IV extends it to be reconfigurable. Section V then conducts numerical evaluations. Finally, Section VI concludes the paper with a discussion and

future directions. An extended version of this paper can be found at [15] that contains all the proofs in the Appendix.

F. Notation

Let $f : A \rightarrow B$ denote a partial function from set A to set B , and let $f(a) = \emptyset$ mean that f is not defined for $a \in A$. For a tree T , we use the notation $(X \xrightarrow{v} Y) \in T$ meaning that there is an arc from X to Y labelled v in T , and denote the depth of T by $\text{depth}(T)$, that is, the maximum number of arcs one can traverse in T starting from the root. Finally, $|A|$ is the cardinality of the set A , and $\mathbb{R}^+ = \{x \geq 0 : x \in \mathbb{R}\}$.

II. PROBLEM FORMULATION

A. Hierarchical Mealy Machines

In this section, we formalise the notion of an HiMM, also found in [16], extending the formalism in [2] to machines with outputs. We start with the notion of an MM.

Definition 1 (Mealy Machine): An MM is a tuple $M = (Q, \Sigma, \Lambda, \delta, \gamma, s)$, where Q is the finite set of states, Σ the finite set of inputs, Λ the finite set of outputs, $\delta : Q \times \Sigma \rightarrow Q$ the transition function, $\gamma : Q \times \Sigma \rightarrow \Lambda$ the output function, and $s \in Q$ the start state. We sometimes use the notation $Q(M)$, $\Sigma(M)$, $\Lambda(M)$, δ_M , γ_M and $s(M)$ to stress that e.g., $Q(M)$ is the set of states of M .

The intuition is as follows. The MM $M = (Q, \Sigma, \Lambda, \delta, \gamma, s)$ starts in $s \in Q$, and transits to the next state $q' = \delta(q, x)$ and outputs $\gamma(q, x)$ given current state $q \in Q$ and input $x \in \Sigma$, or stops if $\delta(q, x) = \emptyset$. Repeating this process, we obtain a trajectory of M :

Definition 2: A sequence $(q_i, x_i)_{i=1}^N \in (Q \times \Sigma)^N$ is a trajectory of an MM $M = (Q, \Sigma, \Lambda, \delta, \gamma, s)$ if $q_{i+1} = \delta(q_i, x_i) \in Q$ for $i \in \{1, \dots, N-1\}$.

In our setting, we assume that an agent (e.g., a robot) can choose the inputs. Hence, we also define a plan to be a sequence $u = (x_i)_{i=1}^N \in \Sigma^N$ of inputs, and $z = (q_i, x_i)_{i=1}^N$ to be the induced trajectory of u starting from q_1 if z is a trajectory. We next provide the definition of an HiMM.

Definition 3 (Hierarchical Mealy Machine): An HiMM is a tuple $Z = (X, T)$ consisting of a set X of MM (the machines in Z) with a common input set Σ and output set Λ , and a tree T with nodes given by X (how the MM are composed in Z). More precisely, each node $M \in X$ in T has $|Q(M)|$ outgoing arcs $\{M \xrightarrow{q} N_q\}_{q \in Q(M)}$ where either $N_q \in X$ (meaning that state q in M corresponds to the MM N_q one layer below in the hierarchy of Z) or $N_q = \emptyset$ (meaning that state q is simply just a state having no further refinement). For brevity, we call $Q_Z := \cup_{X_i \in X} Q(X_i)$ the nodes of Z , and $S_Z := Q_Z \cap \{q : N_q = \emptyset\}$ the states of Z . We assume that all state sets are disjoint, that is, $Q(X_i) \cap Q(X_j) = \emptyset$ for all $X_i, X_j \in X$ such that $X_i \neq X_j$. Furthermore, we have corresponding notions for the start state, transition function, and output function:

- (i) The start function: The start state is a function $\text{start} : X \rightarrow S_Z$ given by

$$\text{start}(X_i) = \begin{cases} \text{start}(X_j), & (X_i \xrightarrow{s(X_i)} X_j) \in T \\ s(X_i), & \text{otherwise.} \end{cases}$$




Fig. 3: Example of an HiMM depicted in a tree-form.

- (ii) Hierarchical transition function: Let $q \in Q(X_j)$ with $X_j \in X$ and $v = \delta_{X_j}(q, x)$. The hierarchical transition function $\psi : Q_Z \times \Sigma \rightarrow S_Z$ is given by

$$\psi(q, x) = \begin{cases} \text{start}(Y), & v \neq \emptyset, (X_j \xrightarrow{v} Y) \in T, Y \in X \\ v, & v \neq \emptyset, \text{otherwise} \\ \psi(w, x), & v = \emptyset, (W \xrightarrow{w} X_j) \in T, W \in X \\ \emptyset, & v = \emptyset, \text{otherwise.} \end{cases}$$

- (iii) Hierarchical output function: Let $q \in Q(X_j)$ with $X_j \in X$ and $v = \delta_{X_j}(q, x)$. The hierarchical output function $\chi : Q_Z \times \Sigma \rightarrow \Lambda$ is given by

$$\chi(q, x) = \begin{cases} \gamma_{X_j}(q, x), & v \neq \emptyset \\ \chi(w, x), & v = \emptyset, (W \xrightarrow{w} X_j) \in T, W \in X \\ \emptyset, & v = \emptyset, \text{otherwise.} \end{cases}$$

An HiMM $Z = (X, T)$ works analogously as an MM: The HiMM $Z = (X, T)$ starts in state $\text{start}(M_0)$ with M_0 being the root MM of T , and for each state $q \in S_Z$ and input $x \in \Sigma$ the machine transits to $q' = \psi(q, x) \in S_Z$ with output $\chi(q, x) \in \Lambda$, or stops if $\psi(q, x) = \emptyset$. A trajectory of an HiMM is defined analogously to a trajectory of an MM.

Remark 1 (Intuition): Definition 3 is technical, but the intuition is simple. To see it, depict an HiMM $Z = (X, T)$ as a tree similar to T with some extra details, as in Fig. 3. Here, $X = \{A, B, C, D\}$ are the MM of Z with dependency given by lines, e.g., B is a child of A in T . States of Z are depicted as circles. We also depict the transitions of each MM given by arrows, where the label of the arrow specifies the input needed to make the transition. We omit writing the outputs explicitly. For example, in MM A , one goes from (the state corresponding to) B to (the state corresponding to) C with input y . Also, small arrows indicate start states (e.g., $s(D) = 4$). We now provide the intuition of Definition 3 with an example. Assume we are at state 5 and apply input x . If there were a transition from 5 with x (as from state 6), Z would just follow that transition. However, since this is not the case, Z instead goes upwards iteratively in the tree T until it finds a node that supports this input (or stops if it does not find one), in this case, C in the MM A has a transition with x (to B). Thus, Z moves to B . After this, Z follows the start states downwards until it reaches a state of Z , in this case, it goes to D then 4. Here, the transition is complete and thus $\psi(5, x) = 4$. Furthermore, the output $\chi(5, x)$ corresponds to the output where the transition step

occurred, in our case going from C with input x . This is to support modularity neatly where we care about that we got from C with input x (e.g., x can encode that we did a task in C), and not how we managed to exit C (e.g., how we did the task in C). Other setups are considered as future work.

1) *Optimal Planning*: In this paper, the output $\chi(q, x)$ for an HiMM Z will exclusively model the cost for applying input x from state q , therefore, we henceforth assume $\Lambda \subset \mathbb{R}^+$. We also consider the cost of a trajectory $z = (q_i, x_i)_{i=1}^N$, called the cumulative cost, defined by $C(z) := \sum_{i=1}^N \chi(q_i, x_i)$ if all $\chi(q_i, x_i) \neq \emptyset$, and $C(z) = +\infty$ otherwise.¹ In this setting, we want to find a plan taking Z from a state s_{init} to a state s_{goal} with minimal cumulative cost. In detail, let $U(s_{\text{init}}, s_{\text{goal}})$ be the set of all plans taking Z from s_{init} to s_{goal} , i.e., $u \in U(s_{\text{init}}, s_{\text{goal}})$ has an induced trajectory $z = (q_i, x_i)_{i=1}^N$ such that $q_1 = s_{\text{init}}$ and $q_{N+1} := \psi(q_N, x_N) = s_{\text{goal}}$. Then we want to find an optimal plan $\hat{u} \in \arg \min_{u \in U(s_{\text{init}}, s_{\text{goal}})} C(z)$ (where z is the induced trajectory of u). We call \hat{u} an optimal plan to the planning objective $(Z, s_{\text{init}}, s_{\text{goal}})$ and the corresponding trajectory for an optimal trajectory. In Section III, we present a control algorithm that efficiently computes optimal plans.

B. Modifications

In this section, we formalise changes in an HiMM, called modifications. More precisely, we introduce four kinds of modifications: state addition, state subtraction, arc modification and composition, specified by Definitions 4 to 7. These modifications are schematically depicted in Fig. 4.

Definition 4 (State Addition): Let $Z = (X, T)$ be an HiMM and $M \in X$ be an MM of Z . We say that an HiMM $Z_{\text{add}} = (X_{\text{add}}, T_{\text{add}})$ is added to M in Z forming a new HiMM Z' with tree T' equal to T except that a new arc $M \xrightarrow{q} N$ has been added from M to the root MM N of T_{add} (connecting Z_{add} to M) labelled with some (arbitrary but distinct) state $q \notin Q(M)$, and M has updated its state set to $Q \cup \{q\}$. Adding a state to M in Z is identical except that the arc $M \xrightarrow{q} N$ is instead a terminal arc $M \xrightarrow{q} \emptyset$.

Definition 5 (State Subtraction): Let $Z = (X, T)$ be an HiMM and $M \in X$ be an MM of Z . We say that we subtract a state $q \in Q(M)$ and form a new HiMM Z' with tree T' equal to T except that the q -arc from M is removed, and M is the modified MM with state q removed.²

Definition 6 (Arc Modification): Let $Z = (X, T)$ be an HiMM and $M = (Q, \Sigma, \Lambda, \delta, \gamma, s) \in X$ be an MM of Z . Let $\delta' : Q \times \Sigma \rightarrow Q$, $\gamma' : Q \times \Sigma \rightarrow \Lambda$ and $s' \in Q$ be the new transition function, output function and start state, and form $M' = (Q, \Sigma, \Lambda, \delta', \gamma', s')$. Replace M with M' in Z with new HiMM Z' as a result, called an arc modification.

Definition 7 (Composition): Let M be an MM with states $Q = \{q_1, \dots, q_{|Q|}\}$ and $Z_{\text{seq}} = \{Z_1, \dots, Z_n\}$ be a set of HiMMs such that $n \leq |Q|$. The composition of Z_{seq} with

¹In fact, only $\chi(q_i, x_i)$ may be empty, since z is a trajectory. Also, we put $C(z) = +\infty$ in the latter case since we do not want Z to stop.

²For simplicity, we forbid the start state to be removed. However, using an arc modification, one can change the start state and then remove the previous start state.




Fig. 4: The four modifications of an HiMM depicted. The dashed rectangle specifies the MM M subject to change.

respect M is the HiMM Z' with tree T' having root M and arcs $M \xrightarrow{q_i} R_i$ for $1 \leq i \leq n$, where R_i is the root MM of T_i (connecting $Z_i = (X_i, T_i)$ to M), and $M \xrightarrow{q_i} \emptyset$ for $i > n$.

Remark 2 (Intuition): The composition replaces the first n states in the MM M with corresponding HiMMs from Z_{seq} , while the remaining $|Q| - n$ states in M are left unchanged.

These four modifications are general enough to account for any change in the sense that any HiMM $Z = (X, T)$ can be changed to any other HiMM $Z' = (X', T')$ using these modifications. To see it, first add all the children of the root of Z' (i.e., the root MM of T') to the root of Z (via repeated state addition), then do an arc modification so that the added part to the root of Z is identical to the root of Z' (including start state), and finally remove the original children of the root of Z (using state subtraction). Then Z is identical to Z' . Of course, this is a brute-force change, replacing the whole

Algorithm 1 Static Hierarchical Planning

Input: HiMM $Z = (X, T)$ and states $s_{\text{init}}, s_{\text{goal}}$.

Output: Optimal plan u to $(Z, s_{\text{init}}, s_{\text{goal}})$.

- 1: **Optimal Exit Computer:**
 - 2: $(c_x^M, z_x^M)_{x \in \Sigma}^{M \in X} \leftarrow \text{Compute_optimal_exits}(Z)$
 - 3: **Optimal Planner:**
 - 4: $z \leftarrow \text{Reduce_and_solve}(Z, s_{\text{init}}, s_{\text{goal}}, (c_x^M, z_x^M)_{x \in \Sigma}^{M \in X})$
 - 5: $u \leftarrow \text{Expand}(z, (z_x^M)_{x \in \Sigma}^{M \in X}, Z)$
-

tree of Z with the whole tree of Z' , and there may be other ways that use much fewer modifications to change Z into Z' . In particular, the true benefit of using modifications is instead when we have small changes in the system, such as removing just a subtree of Z or changing the transitions in one of the MMs, captured by using just a few modifications.

1) *Reconfigurability:* We say that a control algorithm is reconfigurable if a change in the HiMM given by any of the four modifications can be handled with low time complexity (where the change is assumed to be known to the algorithm). We propose such a control algorithm in Section IV.

C. Problem Statement

We now formalise the problem statement. In this paper, we want to find a control algorithm that, for a given HiMM Z :

- 1) Efficiently computes an optimal plan u to any planning objective $(Z, s_{\text{init}}, s_{\text{goal}})$, where efficiency means that the computation has low time complexity.
- 2) Is reconfigurable in the sense of Section II-B.1.

We present in Section III a control algorithm from [16] that fulfils the first criteria and extend this control algorithm in Section IV to fulfil the second criteria.

III. EFFICIENT HIERARCHICAL PLANNING

In this section, we present the control algorithm from [16] for efficiently computing an optimal plan for a given and static (i.e., none-changing) HiMM Z . Then, in Section IV, we extend this algorithm to the reconfigurable case.

We start with an overview of the algorithm, summarised by Algorithm 1. Given an HiMM $Z = (X, T)$, the algorithm consists of two steps. In the first step, the Optimal Exit Computer preprocesses the HiMM Z , computing optimal exit costs $(c_x^M)_{x \in \Sigma}$ for each MM $M \in X$, and corresponding trajectories $(z_x^M)_{x \in \Sigma}$ (line 2). This step needs to be done only once for the given Z . In the second step, the Online Planner considers any two states s_{init} and s_{goal} , and computes an optimal plan to $(Z, s_{\text{init}}, s_{\text{goal}})$ by first reducing irrelevant parts of the tree of Z (intuitively, parts of the tree of Z not on the path between s_{init} and s_{goal}) using the result from the Optimal Exit Computer, then obtains an optimal trajectory z for the reduced HiMM (line 4), and finally expands z to get an optimal plan u to the original HiMM Z (line 5).

We provide the details of the Optimal Exit Computer in Section III-A since it is crucial for the reconfigurable case in Section IV. However, the Online Planner is identical in the reconfigurable case. We therefore only highlight the efficiency of it in Section III-B, refer to [16] for details.

A. Optimal Exit Computer

In this section, we provide the details of the Optimal Exit Computer computing the optimal exit costs. To this end, we first need to define what we mean by optimal exit costs:

Definition 8 (Optimal Exit Cost): Consider HiMM $Z = (X, T)$ and MM $M \in X$. We say that a state $q \in S_Z$ is contained in M if q_i is a descendant of M in T (e.g., 4 is descendant of B in Fig. 3), and an (M, x) -exit trajectory is a trajectory $z = (q_i, x_i)_{i=1}^N$ such that all q_i are contained in M , while $\psi(q_N, x_N)$ is not, and $x_N = x$ (i.e., exits M with x). To an (M, x) -exit trajectory z , let the (M, x) -exit cost be given by $\sum_{i=1}^{N-1} \chi(q_i, x_i)$ (we exclude the cost of (q_N, x_N) since the transition is happening outside the subtree with root M), and let c_x^M be the minimum over all (M, x) -exit costs. We call c_x^M an optimal exit cost.

We now provide the details of the Optimal Exit Computer, computing the optimal exit costs $(c_x^M)_{x \in \Sigma}$ for each MM $M \in X$ recursively over T , as follows. Let $M \in X$ be given with states $Q(M) = \{q_1, \dots, q_l\}$. For $q \in Q(M)$, define c_x^q to be zero if $q \in S_Z$ is a state of Z , or $c_x^q = c_x^{M_q}$ otherwise, where M_q is the MM corresponding to q . Intuitively, this is the extra cost one needs to exit q with x in Z , where, by recursion, we may assume that each c_x^q is known. To compute $(c_x^M)_{x \in \Sigma}$, form the augmented MM \hat{M} given by:

$$\hat{M} := (\{q_1, \dots, q_l\} \cup \{E_x\}_{x \in \Sigma}, \Sigma(M), \Lambda(M), \hat{\delta}, \hat{\gamma}, s(M)).$$

Intuitively, \hat{M} is the same as M except that whenever we are in a state q_i of M that does not support an input $x \in \Sigma$ (i.e., $\delta_M(q_i, x) = \emptyset$), then \hat{M} transits to E_x instead (hence the extra states $\{E_x\}_{x \in \Sigma}$ in \hat{M}). More precisely, the transition function $\hat{\delta}$ is given by

$$\hat{\delta}(q_i, y) := \begin{cases} \delta_M(q_i, y), & \delta_M(q_i, y) \neq \emptyset \\ E_y, & \text{otherwise,} \end{cases}$$

and $\hat{\delta}(E_x, y)$ is immaterial. Furthermore, the cost is

$$\hat{\gamma}(q_i, y) := \begin{cases} c_y^{q_i} + \gamma_M(q_i, y), & \delta_M(q_i, y) \neq \emptyset \\ c_y^{q_i}, & \text{otherwise,} \end{cases}$$

and $\hat{\gamma}(E_x, y)$ is immaterial. With this, we can search in \hat{M} from $s(M)$ using Dijkstra's algorithm to find the minimal cumulative costs to all $(E_x)_{x \in \Sigma}$. Then, c_x^M is equal to the minimal cumulative cost to E_x [16, Proposition 1]. In Dijkstra's algorithm, we also get the corresponding optimal trajectories $(z_x^M)_{x \in \Sigma}$ to $(c_x^M)_{x \in \Sigma}$ for free. The algorithm is summarised by Algorithm 4 (also given in [16] as Algorithm 2) removing lines 3-5 and 16, with time complexity $O(N[b_s|\Sigma| + (b_s + |\Sigma|)\log(b_s + |\Sigma|)])$ [16, Proposition 2]. Here, b_s is the maximum number of states in an MM of Z , and N is the number of MMs of Z . In particular, assuming a bound on the number of states and inputs in an MM of Z , we get time complexity $O(N)$. For further details, see [16].

B. Optimal Planner

In this section, we highlight the time complexity of the Online Planner, see [16] for further details. More precisely, the

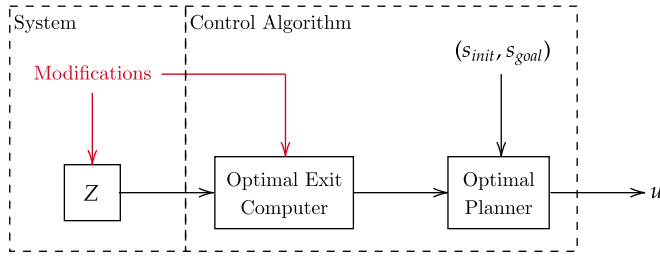


Fig. 5: The control algorithm and its interaction with the changing system Z .

time complexity of `Reduce_and_solve` (line 4 in Algorithm 1) is $O(|\Sigma|^2 \text{depth}(Z) + |\Sigma| \text{depth}(Z) \log(|\Sigma| \text{depth}(Z))) + O([b_s |\Sigma| + b_s \log(b_s)] \text{depth}(Z))$. In particular, assuming a bound on the number of inputs and states in an MM, we get time complexity $O(\text{depth}(Z) \log(\text{depth}(Z)))$. Moreover, the time complexity for `Expand` (line 5 in Algorithm 1) is $O(\text{depth}(Z)|u|)$ for obtaining the whole plan u at once (where $|u|$ is the length of u), or $O(\text{depth}(Z))$ for obtaining the next input in u (ideal for sequential executions). In particular, the total time complexity for obtaining the next optimal input in u is bounded by $O(\text{depth}(Z) \log(\text{depth}(Z)))$. This should be compared with Dijkstra's algorithm, having a time complexity lower bounded by $O(V \log(V))$ [6], where V is the number of states in the HiMM, which in general could be exponential in $\text{depth}(Z)$. The Online Planner therefore potentially saves huge computing time for large HiMMs.

IV. RECONFIGURABLE HIERARCHICAL PLANNING

In this section, we extend Algorithm 1 to be reconfigurable. To first see the need for this, note that the computational bottleneck of Algorithm 1 typically comes from the Optimal Exit Computer, with time complexity $O(N)$, whereas the Optimal Planner computes the next optimal input in time $O(\text{depth}(Z) \log(\text{depth}(Z)))$. Therefore, the Optimal Planner is typically much faster than the Optimal Exit Computer. This is not a problem for a non-changing HiMM Z , since the optimal exit costs need to be computed only once, which can then be used repeatedly by the Optimal Planner to rapidly compute optimal plans. However, if Z changes, then one would need to re-compute all the optimal exit costs every time Z is changed, which is inefficient, especially if the changes are small.

To address this issue, we extend Algorithm 1 so that the Optimal Exit Computer updates only the needed parts of its computation whenever Z changes. The resulting control algorithm and its interaction with the (changing) HiMM Z are depicted in Fig. 5. Here, a change in Z is given by a sequence of modifications. The control algorithm gets informed of the change and the modifications. By specifying the modifications, only some of the MMs of Z need to be recomputed, marked by the Optimal Exit Computer. After all modifications have been reported, the Optimal Exit Computer recomputes the optimal exit costs, but only for the marked MMs. This potentially saves huge computation time.

Algorithm 2 Reconfigurable Hierarchical Planning

Input: HiMM Z , **modifications** m_{seq} and states s_{init}, s_{goal} .
Output: Optimal plan u to (Z, s_{init}, s_{goal}) .

```

1: Optimal Exit Computer:
2: for  $m$  in  $m_{seq}$  do
3:    $\text{Modify}(Z, m)$   $\triangleright$  This line is executed in System
4:    $\text{Mark}(Z, m)$ 
5: end for
6:  $(c_x^M, z_x^M)_{x \in \Sigma}^{M \in X} \leftarrow \text{Compute\_optimal\_exits}(Z)$ 
7: Optimal Planner:
8:  $z \leftarrow \text{Reduce\_and\_solve}(Z, s_{init}, s_{goal}, (c_x^M, z_x^M)_{x \in \Sigma}^{M \in X})$ 
9:  $u \leftarrow \text{Expand}(z, (z_x^M)_{x \in \Sigma}^{M \in X}, Z)$ 

```

In more detail, the control algorithm is summarised by Algorithm 2, where changes with respect to Algorithm 1 are in red. Here, m_{seq} is the sequence of modifications, read sequentially (line 2-4): For m in m_{seq} , Z is first changed with respect to m (line 3)³, then, the Optimal Exit Computer marks corresponding MMs (line 4), given by Algorithm 3. After m_{seq} has been read, the optimal exit costs for the marked MMs are updated (line 6), given by Algorithm 4 (red parts accounts for the marking). We stress that we only need to execute the Optimal Exit Computer whenever Z changes. The Optimal Planner is identical to the one in Algorithm 1 and can be used repeatedly as long as changes are up-to-date.

A. Reconfigurability

We now show that Algorithm 2 is reconfigurable. We first explain the marking procedure in detail. Then, we prove the correctness of the Optimal Exit Computer (Proposition 1) and show that the time complexity is low for marking and updating the optimal exit costs for each of the four modifications (Proposition 2), thus, Algorithm 2 is reconfigurable.

1) *Marking:* The marking of $Z = (X, T)$ given a modification m is done by Algorithm 3. In brief, the procedure is to mark the MM considered by m and all ancestors to it. In more detail, consider the case when m is an arc modification modifying an MM M in Z . With M modified, the optimal exit costs of M might have changed. We therefore mark M (to recompute its optimal exit costs) in Z (line 1). Also, modifying M could affect the optimal exit costs of the MMs above M in the hierarchy of Z , i.e., the ancestors of M . We therefore mark all the ancestors of M (line 3). However, note that all other MMs in Z have unaffected optimal exit costs, and we are therefore done. The case for state subtraction and state addition are analogous, and in the composition case, we only need to mark M since M has no ancestors.

Convention 1: When initialising an HiMM Z , all MMs are by default marked since no optimal exit costs have been computed. This enables us to use Algorithm 4 when initialising Z and not only when updating Z (after modifications). Also, whenever we modify an HiMM we always mark it accordingly (as above). Thus, we do not need to mark Z_{add}

³This line is executed by the system and not the control algorithm, see Fig. 5. However, the algorithm tracks the change, so we keep it for clarity.

Algorithm 3 Mark

Input: Modification m and (modified) HiMM Z **Output:** Appropriate marking of Z

```

1: mark_itself( $M, Z$ )
2: if  $m$  not composition then
3:   mark_ancestors( $M, Z$ )
4: end if

```

in state addition or $\{Z_1, \dots, Z_n\}$ in composition since their optimal exit costs do not change by the modification, and are already correctly marked (note that some of their optimal exit costs might have already been computed by Algorithm 4). See the proof of Proposition 1 for details.

2) *Computing optimal exit costs:* After marking, the optimal exit costs for the marked MMs are computed using Algorithm 4, where the computation of the optimal exit costs for a given MM is as in Section III-A. The correctness of this procedure follows from the following proposition:

Proposition 1: Let Z be an HiMM, modified by the sequence of modifications m_{seq} (possibly empty, e.g., if Z is only initialised) and marked using Algorithm 3. Then Algorithm 4 returns the correct optimal exit costs.

3) *Time complexity:* Finally, we address the time complexity. The time complexity of Algorithm 3 is $O(\text{depth}(Z))$ since there are at most $\text{depth}(Z)$ MMs to mark (the MM itself and all its ancestors). The time complexity of Algorithm 4 varies with the modifications that have been done. In the worst case, all optimal exit costs need to be computed with time complexity $O(N[b_s|\Sigma| + (b_s + |\Sigma|) \log(b_s + |\Sigma|)])$ [16, Proposition 2]. However, for just one modification, we get:

Proposition 2: Let HiMM Z be given and consider the Optimal Exit Computer with only one modification $m_{\text{seq}} = m$. Assume Z have all optimal exit costs computed before the modification, and that the same holds for Z_{add} if m is a state addition, or $\{Z_1, \dots, Z_n\}$ if m is a composition. Then, the time complexity for executing the Optimal Exit Computer if m is a state addition, state subtraction, or arc modification is $O(\text{depth}(Z)[b_s|\Sigma| + (b_s + |\Sigma|) \log(b_s + |\Sigma|)])$, and $O([b_s|\Sigma| + (b_s + |\Sigma|) \log(b_s + |\Sigma|)])$ if m is a composition. In particular, assuming a bound on the number of states and inputs of the MMs in Z , the time complexity for any modification is $O(\text{depth}(Z))$. Thus, Algorithm 2 is reconfigurable.

V. NUMERICAL EVALUATIONS

In this section, we present numerical evaluations, validating the proposed control algorithm given by Algorithm 2. We consider the robot application from the motivation and compare our method with Dijkstra's algorithm and Contraction Hierarchies, see [15] for implementation details.

A. Case studies

We first formalise the robot application from the motivation, also given in [16]. The HiMM Z is constructed using three MMs, corresponding to the layers in Fig. 1.

The first MM (for Layer 1) has 10 states corresponding to the houses. The houses are ordered in a line where the robot

Algorithm 4 Compute_optimal_exits

Input: HiMM $Z = (X, T)$ with markings.**Output:** Computed $(c_x^M, z_x^M)_{x \in \Sigma}$ for each MM M of Z

```

1: Optimal_exit( $M_0$ )      ▷ Run from root MM  $M_0$  in  $T$ 
2: Optimal_exit( $M$ ):      ▷ Recursive help function
3: if  $M$  is not marked then
4:   return  $(c_x^M, z_x^M)_{x \in \Sigma}$  ▷ Have already been computed
5: else
6:   for each state  $q$  in  $Q(M)$  do
7:     if  $q \in S_Z$  then
8:        $(c_x^q)_{x \in \Sigma} \leftarrow 0_{|\Sigma|}$ 
9:     else
10:      Let  $M_q$  be the MM corresponding to  $q$ .
11:       $(c_x^q, z_x^q)_{x \in \Sigma} \leftarrow \text{Optimal\_exit}(M_q)$ 
12:    end if
13:  end for
14:  Construct  $\hat{M}$ 
15:   $(c_x^M, z_x^M)_{x \in \Sigma} \leftarrow \text{Dijkstra}(s(M), \{E_x\}_{x \in \Sigma}, \hat{M})$ 
16:  Unmark  $M$ 
17:  return  $(c_x^M, z_x^M)_{x \in \Sigma}$ 
18: end if

```

can move to a neighbouring house as seen by the arrows in Fig. 1, at a cost of 100. Start state is $s(M_1) = 1$.

The second MM (for Layer 2), M_2 , corresponds to the locations in a house with $10 \cdot 10 + 1 = 101$ states in a grid-like formation where the robot can move to any neighbouring grid-point, at a cost of 1, as in Fig. 1, and start state $s(M_2)$ marked S, interpreted as the entrance to the house. From $s(M_2)$, the robot can also move left/right. However, these two inputs make the robot exit M_2 and instead, in Z , move it to the corresponding neighbouring house in M_1 .

The third MM (for Layer 3), M_3 , corresponds to the lab desk at a location, with $10 \cdot 9 + 1 = 91$ states. Here the robot starts at $s(M_3)$, marked S in Fig. 1. From $s(M_3)$, the robot can choose to go up, down, left or right, which quits M_3 and instead moves to the corresponding location in M_2 ; It can also start to steer a robot arm over a 3×3 test tube rack analogous to the location-grid in M_2 , where the robot arm (i, j) ($1 \leq i, j \leq 3$) is initially over tube (1,1). The robot can also scan a tube with the arm. The scanned tube (i, j) is then remembered and no further tubes can be scanned. Finally, the robot can go back to $s(M_3)$ when the arm is again over tube (1,1). All costs in M_3 are set to 0.5, except scanning which costs 10. The hierarchy of M_1 , M_2 and M_3 yields Z with $10 \cdot 101 \cdot 91 = 91910$ states.

The case studies we consider are as follows. In Study 1, we consider Z just described. In Study 2, we change Z by adding a house, House 11, being a neighbour house to House 10 (see Fig. 1), formally done by a state addition (adding House 11), followed by an arc modification (connecting House 10 and 11). All MMs of House 11 are marked, i.e., their optimal exit costs have not been computed yet. In Study 3, we also change Z (from Study 1), but now by blocking some locations in House 2 as in Fig. 2, formally removing

TABLE I: Computing times for the case studies.

| | Study 1 | Study 2 | Study 3 |
|--------------------------------|----------|----------|-----------|
| Compute optimal plan | 0.0254 s | 0.0271 s | 0.0228 s |
| Compute optimal exit costs | 2.23 s | 0.208 s | 0.00234 s |
| Compute all optimal exit costs | 2.23 s | 2.24 s | 1.96 s |
| Dijkstra's algorithm | 3.73 s | 4.01 s | 1.37 s |
| CH – preprocessing | 1196 s | 1320 s | 1175 s |
| CH – compute optimal plan | 0.0166 s | 0.0187 s | 0.0155 s |

locations using state subtraction. In all studies, the robot starts in the bottom-right corner of House 1 with an arm over tube (2, 2), with the goal to scan tube (2, 2) in the bottom-right corner of House 10, 11 and 2 in Study 1, 2 and 3 respectively.

B. Result

The results of Study 1 to 3 are summarised in Table I. The first row shows the time it takes the Optimal Planner in Algorithm 2 to compute an optimal plan. In all three studies, the computation time is around 20-30 ms. This should be compared with Dijkstra's algorithm (row 4) with times around 1-4 s, two orders of magnitudes slower. Contraction Hierarchies (row 6) has instead a computation time around 16-19 ms, therefore slightly faster than Algorithm 2, though still the same order of magnitude. However, the slightly faster times come at the expense of a slower preprocessing step for the Contraction Hierarchies taking around 1200-1300 s to finish (row 5), compared with significantly faster computation times of Algorithm 4 (row 2): 2.23 s, 0.208 s and 0.00234 s for Study 1, 2 and 3, respectively. Here, we also see the importance of reconfigurability, explaining the time differences of Algorithm 4. More precisely, in Study 1, we initialise Z and need to compute all optimal exit costs. This takes the same time (2.23 s) as computing all optimal exit costs (row 3). However, in Study 2, we only modify Z from Study 1, hence, only some optimal exit costs need to be recomputed, in this case, all MMs of Z corresponding to House 11 (since they have not been computed yet) and the root MM of Z (since it has been modified). The result is a much faster computation time (0.208 s), compared to computing all optimal exit costs (2.24 s). In Study 3, the modification results in even fewer MMs that need to be recomputed, with resulting time 2.34 ms. This shows the benefit of reconfigurability, being several orders of magnitude faster than needing to compute all optimal exit costs (as in [16]).

VI. CONCLUSION

In this paper, we have considered a planning problem for a large-scale system modelled as an HiMM Z and developed a control algorithm for computing optimal plans between any two states. The control algorithm consists of two steps. A preprocessing step computing optimal exit costs for each MM in the Z , and a query step computing an optimal plan between any two states, using the optimal exit costs. Given bounds on the number of states and inputs of an MM in Z , the preprocessing step has time complexity $O(N)$ (where N is the number of MMs in Z), while the query step has

time complexity $O(\text{depth}(Z) \log(\text{depth}(Z)))$ for obtaining the next optimal input. The latter enables rapid optimal plan computations. The control algorithm is also reconfigurable in the sense that changes, given by modifications, can be handled with ease, where a modification in Z only needs a minor update of the preprocessing step in time $O(\text{depth}(Z))$. We validated the control algorithm on a robotic application, comparing it with Dijkstra's algorithm and Contraction Hierarchies. We noted that our control algorithm outperforms Dijkstra's algorithm, computes optimal plans in the same order of magnitude as Contraction Hierarchies but with a significantly faster preprocessing step, and handles changes in the HiMM efficiently.

Future work includes extensions to stochastic hierarchical systems and additional numerical evaluations. Another interesting direction is to modify the decomposition method in [2] to automatically decompose an MM into an equivalent HiMM and then use our planning algorithm on the HiMM to efficiently compute optimal plans.

REFERENCES

- [1] H. Bast, et al. Route planning in transportation networks. In *Algorithm engineering*, pages 19–80. Springer, 2016.
- [2] O. Biggar, M. Zamani, and I. Shames. Modular decomposition of hierarchical finite state machines. *arXiv preprint arXiv:2111.04902*, 2021.
- [3] C. G. Cassandras and S. Lafortune. *Introduction to discrete event systems*. Springer, 2008.
- [4] J. Dibbelt, B. Strasser, and D. Wagner. Customizable contraction hierarchies. *Journal of Experimental Algorithmics (JEA)*, 21:1–49, 2016.
- [5] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 1(1):269–271, dec 1959.
- [6] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. In *25th Annual Symposium on Foundations of Computer Science, 1984.*, pages 338–346, 1984.
- [7] R. Geisberger, P. Sanders, D. Schultes, and C. Vetter. Exact routing in large road networks using contraction hierarchies. *Transportation Science*, 46(3):388–404, 2012.
- [8] D. Harel. Statecharts: A visual formalism for complex systems. *Science of computer programming*, 8(3):231–274, 1987.
- [9] C. Ma and W. M. Wonham. Nonblocking supervisory control of state tree structures. *IEEE Transactions on Automatic Control*, 51(5):782–793, 2006.
- [10] J. Maue, P. Sanders, and D. Matijevic. Goal-directed shortest-path queries using precomputed cluster distances. *ACM J. Exp. Algorithmics*, 14, 2010.
- [11] G. H. Mealy. A method for synthesizing sequential circuits. *The Bell System Technical Journal*, 34(5):1045–1079, 1955.
- [12] I. Millington and J. Funge. *Artificial intelligence for games*. CRC Press, 2018.
- [13] R. H. Möhring, et al. Partitioning graphs to speedup Dijkstra's algorithm. *ACM J. Exp. Algorithmics*, 11, 2007.
- [14] P. Schillinger, S. Kohlbrecher, and O. Von Stryk. Human-robot collaborative high-level control with application to rescue robotics. In *2016 IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 2796–2802. IEEE, 2016.
- [15] E. Stefansson and K. H. Johansson. Efficient and reconfigurable optimal planning in large-scale systems using hierarchical finite state machines. *arXiv preprint arXiv:2303.16567*, 2023.
- [16] E. Stefansson and K. H. Johansson. Hierarchical finite state machines for efficient optimal planning in large-scale systems. In *2023 European Control Conference (ECC)*, pages 1–8, 2023.
- [17] O. N. Timo, et al. Reachability in hierarchical machines. In *2014 IEEE 15th Int. Conf. on Information Reuse and Integration (IEEE IRI 2014)*, pages 475–482, 2014.
- [18] X. Wang, Z. Li, and W. M. Wonham. Real-time scheduling based on nonblocking supervisory control of state-tree structures. *IEEE Transactions on Automatic Control*, 66(9):4230–4237, 2020.