

Hierarchical Multi-modal Navigation for Experience-informed Quadrupedal Rebar Grid Traversal

Max Asselmeier¹, Eohan George², Patricio A. Vela³ and Ye Zhao¹

Abstract—In this work, we adapt a hierarchical multi-modal contact planning paradigm to the domain of quadrupedal navigation. In order to adapt this framework to the task of legged locomotion, we incorporate domain-specific modules including contact sequences and reachability analysis into the higher-level foothold sequence planning. We additionally integrate kinodynamically-aware TO-based footstep controllers into the lower-level foothold transition generation layer in an effort to tightly integrate planning and control. We also embed this multi-modal contact planner into a broader navigation framework to both provide an additional contact planning heuristic in the form of a guiding torso path and also demonstrate an avenue for future integration of perception into the loop. We validate the efficiency of our multi-modal navigation framework through a case study involving the task of rebar grid traversal, where the robot is required to locomote exclusively along rebar poles to reach a specified goal region while avoiding obstacles, and show that the hierarchical integration of planning and control leads to more informed and successful planning.

I. INTRODUCTION

The task of legged locomotion elicits a hybridized planning space with a discrete set of combinations between robot end effectors and environmental artifacts that can support footsteps along with continuous footstep positions along such artifacts. Methods that opt to perform joint contact and footstep position planning often struggle to do so in a tractable manner because of this hybrid planning space. The need for kinodynamically-feasible solutions to this planning problem further compounds the computational efforts required.

An alternative way to resolve contact planning is through a hierarchical approach in which discrete contact sequences are generated at the higher level and then continuous whole body trajectories that abide by the generated contact sequence are synthesized at the lower level. The separation of the planning space into discrete and continuous components greatly simplifies the overall planning problem.

In this work, we take inspiration from the domain of multi-modal motion planning (MMMP) which often leverages such a hierarchical decomposition. Namely, we employ the structure of a mode transition graph for high-level contact sequence planning. We additionally integrate a modified version of the experience heuristic [1] used in MMMP




Fig. 1: Illustration of a quadrupedal robot performing rebar traversal in a construction environment [2].

which weights the mode transition graph based on the results of offline planning trials. At the lower level, we opt to deploy model-based trajectory optimization (TO) to combine the high-level contact sequence planning with low-level TO-based footstep controllers. In summary, our main contributions are:

- 1) Adaptation of experience-based MMMP paradigms to the domain of quadrupedal contact planning
- 2) Integration of high-level contact sequence planning with lower-level TO in order to naturally embed footstep controllers and tightly integrate kinodynamically-aware optimal costs into mode transition graph weights
- 3) Deployment of this multi-modal contact planner within a global navigation framework to expedite planning and provide avenues for deeper perception integration
- 4) Experimental validation of the proposed framework on the task of rebar grid traversal

II. PRELIMINARIES

In this section, we provide background information about the planning problem we aim to solve and elaborate on important adapted concepts from the ALEF framework.

A. Environment specifications

While the proposed navigation framework is environment-agnostic, we ground many of the concepts introduced in this work in the task of quadrupedal rebar grid traversal. Therefore, it is crucial to elucidate the specifics of this planning environment before delving into subsequent sections.

Rebar grids are comprised of a collection of bars denoted as a set B . A bar $b \in B$ is characterized by an initial point

¹M. Asselmeier, and Y. Zhao are with the School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30308, USA. mass@gatech.edu

²E. George is with SkyMul Inc, Atlanta, GA 30339, USA.

³P.A. Vela is with the School of Electrical and Computer Engineering and the Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30308, USA.

*This work is supported through a sponsored research grant from SkyMul Inc. and a Georgia Tech IRIM/IPaT Aware Home Seed Grant

$p_0 = (x_0, y_0, z_0)$ in the rebar grid frame \mathcal{B} , a length l , and an orientation θ with respect to the rebar grid frame in the $x - y$ plane of the grid. These grids are comprised of N_h horizontal bars and N_v vertical bars. Additionally, grids incorporate an assortment of obstacles represented by O positioned throughout their surface. The poses and shapes of obstacles are assumed to be known a priori. Various grids with distinct features are shown in Figure 2.




Fig. 2: (a) Normal rebar grid, (b) skewed grid, (c) variably spaced grid, (d) grid with obstacles on surface.

B. Contact manifolds

Many MMMP frameworks that reason about contact exploit the manifold nature of contact constraints. In this work, we maintain the terminology used in the Augmented Leafs with Experience on Foliations (ALEF) framework [1]. We briefly cover the terminology in this section and refer readers to [1] for more details. In legged locomotion, a contact *mode* ξ is a unique set of foothold positions along a set of steppable objects. From this mode, an $(n - n_\xi)$ -dimensional mode manifold \mathcal{M}^ξ embedded in the configuration space $\mathcal{Q} \in \mathbb{R}^n$ arises. As foothold positions vary along the steppable objects, different contact manifolds arise.

All contact modes corresponding to the same group of steppable objects can be binned into a *mode family* Ξ . From each mode family Ξ_i arises a foliation \mathcal{F}_{Ξ_i} defined by a n_χ -dimensional *transverse manifold* X and a set of non-overlapping $(n - n_\chi)$ -dimensional *leaf manifolds* $\mathcal{L}_\chi \forall \chi \in X$. The elements $\chi \in X$ are called *coparameters*, and a coparameter χ uniquely parameterizes a mode ξ . A mode ξ can be viewed as the tuple of a mode family Ξ and a coparameter χ , and the union of leaf manifolds along the set of coparameters $\bigcup_{\chi \in X} \mathcal{L}_\chi$ recovers the foliation \mathcal{F}_{Ξ_i} .

A leaf or mode manifold $\mathcal{L}_\chi = \mathcal{M}^\xi$ can be implicitly defined through a constraint function $F^\xi : \mathbb{R}^n \rightarrow \mathbb{R}^{n_\xi}$ where a robot configuration \mathbf{q} lies on the mode manifold if $F^\xi(\mathbf{q}) = \mathbf{0}$. For rebar traversal, we choose a mode to represent three stance feet in contact with three rebar segments. Therefore, an example constraint function for a contact mode is

$$F^\xi(\mathbf{q}) := [F_1^\xi(\mathbf{q}), F_2^\xi(\mathbf{q}), F_3^\xi(\mathbf{q})]^T, \quad (1)$$

where for a foot j in contact with bar $b = \langle p_0, l, \theta \rangle$, then

$$F_j^\xi(\mathbf{q}) := \begin{bmatrix} \text{FK}_x^j(\mathbf{q}) - (x_0 + \chi_j \cdot l \cdot \cos(\theta)) \\ \text{FK}_y^j(\mathbf{q}) - (y_0 + \chi_j \cdot l \cdot \sin(\theta)) \\ \text{FK}_z^j(\mathbf{q}) - z_0 \end{bmatrix} = \mathbf{0}, \quad (2)$$

where $\text{FK}^j(\mathbf{q})$ gives the position of foot j given \mathbf{q} via forward kinematics (FK).

III. LOW-LEVEL MULTI-MODAL PLANNING

Given the prior information, the planning problem we are trying to solve is as follows. Assume a quadrupedal robot with a configuration space $\mathcal{Q} \subset \mathbb{R}^{n_q}$. We seek to find a collision-free path $\mathbf{q}(s)$ with $s \in [0, 1]$ from a start configuration $\mathbf{q}(0) = \mathbf{q}_{\text{start}}$ to a goal configuration $\mathbf{q}(1) = \mathbf{q}_{\text{goal}}$. Contact must strictly be made with the set of bars B , and collisions with obstacles O should be avoided.

A. Mode Transition Graph Construction

We employ a mode transition graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ in which the mode families comprise the set of vertices \mathcal{V} , and the edges \mathcal{E} are formed between mode families for which kinematically feasible transitions exist. We utilize locomotion-specific domain knowledge to implicitly define what transitions are feasible within the graph.

First, a user-defined contact sequence informs the graph on what order in which the four feet will step. This constraint is embedded into the graph by adding a transition from a vertex v_i to vertex v_j only if the currently planning leg at v_i and the currently planning leg at v_j occur sequentially in the user-defined contact sequence.

Second, we incorporate a kinematic reachability approximation that captures what ground support surfaces are reachable by the robot's legs. Reachable areas or volumes can be approximated through polygons [3] ellipses [4], and the method that we employ, superquadrics [5]. To obtain a parameterized estimate of the reachable area, we fix the robot torso pose at a nominal height above the ground and randomly sample joint configurations for the legs. Samples are then projected into ground contact and those that satisfy a distance threshold $\epsilon = 1$ cm are kept. Then, superquadric parameters are tuned to encapsulate the samples. The set of points $S \subset \mathbb{R}^2$ within the superquadric centered at (x_0, y_0) is

$$S = \{(x, y) \in \mathbb{R}^2 \mid \left| \frac{x - x_0}{A} \right|^a + \left| \frac{y - y_0}{B} \right|^b \leq 1\}, \quad (3)$$

where scalars A, B and a, b control the dimensions and curvature in the x - and y -dimensions respectively. Visualizations of the tuned superquadrics along with the projected samples can be seen in Figure 3.

B. Mode Transition Graph Search

We formulate the task of finding a foothold sequence as a graph search problem over the mode transition graph. Similar to the ALEF framework, we discretize along the transverse manifold of each mode family to generate "slices" of the foliations that correspond to intervals of foothold positioning. Transitions are then added between all slices of the source and destination mode families in which edges exist in the graph. This discretization allows for the search to provide a candidate *lead*: a sequence of mode families and coparameter values that define a foothold sequence from the start to goal.

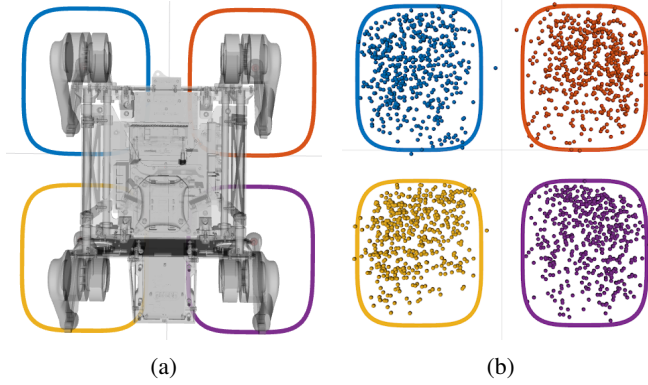


Fig. 3: (a) Superquadrics centered at robot feet, (b) superquadrics with projected samples. Some outlying samples are ignored to prioritize feasibility over reachability.

1) *Edge cost*: For a transition between source mode $\xi_{\text{src}} = \langle \Xi_{\text{src}}, \chi_{\text{src}} \rangle$ and destination mode $\xi_{\text{dst}} = \langle \Xi_{\text{dst}}, \chi_{\text{dst}} \rangle$, the graph edge $e = (\xi_{\text{src}}, \xi_{\text{dst}})$ is assigned the weight

$$\Delta c(\xi_{\text{src}}, \xi_{\text{dst}}) = w_{\mathcal{D}} \cdot \mathcal{D}^{\Xi_{\text{src}}, \Xi_{\text{dst}}}(\chi_{\text{src}}, \chi_{\text{dst}}) + w_d \cdot d_{\text{CoM}}(\xi_{\text{src}}, \xi_{\text{dst}}) + w_{\tau} \cdot d_{\tau}(\xi_{\text{src}}, \xi_{\text{dst}}), \quad (4)$$

where the distribution $\mathcal{D}^{\Xi_{\text{src}}, \Xi_{\text{dst}}}(\chi_{\text{src}}, \chi_{\text{dst}})$ captures the difficulty of transitioning from ξ_{src} to ξ_{dst} , $d_{\text{CoM}}(\xi_{\text{src}}, \xi_{\text{dst}})$ is the Euclidean distance between nominal CoM positions for ξ_{src} and ξ_{dst} , and $d_{\tau}(\xi_{\text{src}}, \xi_{\text{dst}})$ is the deviation of modes ξ_{src} and ξ_{dst} from a suggested torso path. Weights $w_{\mathcal{D}}$, w_d , and w_{τ} are all scalar positive values. We employ an A* search algorithm to obtain the suggested torso path, but any global navigation planner that provides a sequence of torso poses can be integrated.

2) *Cost-to-go*: For the A* search, we also assign a contact mode $\xi = \langle \Xi_{\text{src}}, \chi_{\text{src}} \rangle$ the search heuristic value

$$g(\xi) = w_d \cdot d(\text{CoM}(\xi), \text{CoM}(\Xi_{\text{goal}})) \quad (5)$$

to ensure an admissible search heuristic and therefore provide optimal foothold sequences with respect to edge weights.

C. Whole Body Trajectory Optimization

At the lower level of the framework, we use TO to obtain continuous paths that transition the robot from mode ξ_{src} to mode ξ_{dst} . While this change sacrifices the probabilistic completeness that the original framework provides, TO offers a way to embed the dynamics-aware footstep controllers that are prevalent in legged locomotion, enabling kinodynamically-aware planning which was not accomplished in the prior framework. The TO problem that we formulate solves over the robot state \mathbf{x} which contains the floating base pose \mathbf{q}_b , the leg joint configuration \mathbf{q}_j , and the centroidal momentum \mathbf{h} as well as the robot input \mathbf{u} which contains the contact forces \mathbf{f} and joint velocities \mathbf{v} . This formulation is given in Equation 6.

The TO formulation is shown as:

$$\begin{aligned} & \|\mathbf{x}[N] - \mathbf{x}^{\text{des}}[N]\|_{Q_f}^2 + \\ \min_{\mathbf{x}, \mathbf{u}} & \sum_{k=0}^{N-1} \left(\|\mathbf{x}[k] - \mathbf{x}^{\text{des}}[k]\|_{Q}^2 + \|\mathbf{u}[k]\|_{R}^2 \right) \\ \text{subject to} & \\ \text{(Mode)} & F^{\xi_{\text{src}}}(\mathbf{q}[k]) = \mathbf{0}, F^{\xi_{\text{dst}}}(\mathbf{q}[N]) = \mathbf{0} \quad (6a) \\ \text{(Dynamics)} & \begin{bmatrix} \dot{\mathbf{h}}[k] \\ \dot{\mathbf{q}}_b[k] \\ \dot{\mathbf{q}}_j[k] \end{bmatrix} = \begin{bmatrix} \mathcal{D}(\mathbf{q}[k], \mathbf{f}[k]) \\ \mathbf{A}_b^{-1}(\mathbf{h}[k] - \mathbf{A}_j \mathbf{v}_j[k]) \\ \mathbf{v}_j[k] \end{bmatrix} \quad (6b) \\ \text{(Friction)} & \mathbf{f}_j[k] \in \mathcal{F}_j(\mu, \mathbf{q}) \quad \forall j \in \mathcal{C}_{\text{src}} \quad (6c) \\ & \mathbf{f}_j[k] = \mathbf{0} \quad \forall j \notin \mathcal{C}_{\text{src}} \quad (6d) \\ \text{(Collision)} & g(\mathbf{q}[k]) \geq 0 \quad \forall k \in [0, N] \quad (6e) \end{aligned}$$

where \mathcal{C}_{src} represents the set of stance feet for the source mode. The dynamics constraints (6b) come from the centroidal dynamics model [6] where $\mathbf{A} \in \mathbb{R}^{6 \times n_q}$ is the centroidal momentum matrix. \mathbf{x}^{des} is generated in two steps, first, mode constraint functions (6a) are used to project a randomly sampled target configuration into contact satisfying the source and destination modes. If this step is successful, cubic splines are synthesized for the swing feet to build out the remainder of \mathbf{x}^{des} . We formulate the above TO as a Sequential Quadratic Program (SQP) and solve through the OCS2 library [7] with a time horizon of $T = 0.5$ seconds, $N = 50$ knot points, and maximal 250 iterations. If the optimal cost of an attempted mode transition is above a threshold \mathcal{J}_{max} , then the planning trial is terminated early.

IV. PLANNING WITH EXPERIENCE

The objective of experience-based planning is to acquire a continuous function that captures the level of difficulty associated with undertaking the contact transitions within the mode transition graph. These distributions $\mathcal{D}^{\Xi_{\text{src}}, \Xi_{\text{dst}}}(\chi_{\text{src}}, \chi_{\text{dst}})$ are then used as weights in the mode transition graph search. Experience accumulation can be viewed as a form of proprioception where the robot interacts with the environment offline to “practice” planning in order to determine what foothold sequences reliably lead to successful paths to the goal. It is through this experience accumulation that the multi-modal planner is able to discover the affordances that the environment provides in terms of which objects can support collision-free footholds and facilitate reaching the goal region.

A. Optimal Cost Integration

There are two potential outcomes of a transition attempt:

- 1) Contact projection fails to generate a target configuration, suggesting an infeasible transition ($\mathcal{J}_{\text{src}, \text{dst}} = \infty$)
- 2) Contact projection generates a target configuration, triggering a TO instance ($\mathcal{J}_{\text{src}, \text{dst}} = \text{Equation 6}$)

While a contact transition between two modes might be infeasible, transitions for all source and destination coparameters between the two mode families might still be feasible. Therefore, discarding the entire transition could impede the

discovery of a viable path to the goal. To address this issue, it becomes crucial to devise an experience distribution scheme capable of handling either of these two outcomes. To accommodate for this, the costs $\mathcal{J}_{\text{src,dst}}$ are passed through a weighted hyperbolic tangent (\tanh) function

$$\delta_{\text{src,dst}} = w_1 \tanh(w_2 \cdot \mathcal{J}_{\text{src,dst}} + w_3), \quad (7)$$

where w_1 , w_2 , and w_3 are positive scalar values, to map the costs to finite positive *penalties* that can be used to populate the edge weights in the mode transition graph.

B. Experience incorporation

The smoothness of contact manifolds allows us to exploit planning results we obtain to obtain informed estimates of the difficulty of contact transitions near those attempted. This is predicated on the idea that since foliations are smooth, coparameters that are nearby on the transverse manifold parameterize modes that are similar in difficulty.

To incorporate penalties into weight distributions, we employ function regression techniques to estimate the continuous distribution of the average penalty value at different contact positions. Each time an edge in the mode transition graph is traversed, a weighted radial basis function (RBF) is added to the transition weight distribution in the form of

$$f^{\Xi_{\text{src}}, \Xi_{\text{dst}}}(\chi_{\text{src}}, \chi_{\text{dst}}) = w_e \cdot \exp\left(\frac{-d(\chi_{\text{src}}, \chi_{\text{dst}})^2}{2 \cdot \sigma^2}\right) \quad (8)$$

where

$$w_e = (\mathcal{J}_{\text{src,dst}} - \bar{\mathcal{J}}) \quad (9)$$

where $\mathcal{J}_{\text{src,dst}}$ is the cost obtained from the attempted mode transition, $\bar{\mathcal{J}}$ is the average transition cost between ξ_{src} and ξ_{dst} , $d(\chi_{\text{src}}, \chi_{\text{dst}})$ represents the distance of a coparameter to $(\chi_{\text{src}}, \chi_{\text{dst}})$, and σ represents the standard deviation of the RBF. This is equivalent to adding an RBF centered at $(\chi_{\text{src}}, \chi_{\text{dst}})$ to the distribution which is scaled by the traversal's deviation from the mode transition's average cost, meaning that as the number of traversals increases and this average value becomes more accurate, the distribution itself converges. By integrating footstep controllers and using optimal costs to weight the mode transition graph, the higher-level contact sequence planner is informed of the overall performance of the lower-level footstep controllers, yielding tight integration between planning and control.

V. EXPERIMENTAL RESULTS

We perform offline experience accumulation and investigate the resultant whole-body motion plans. The environment used in this case study along with the final reference trajectory can be seen in Figure 4. We perform an ablation study where the navigation framework is run with and without the guiding torso path, and we record computation times and overall planning results. Results are shown in Figure 5.

For this grid, the added torso planner greatly expedites experience accumulation. The guiding torso path biases the multi-modal contact planner towards collision-free transitions which take far less time to generate in TO. Without the torso path, the mode transition graph search initially


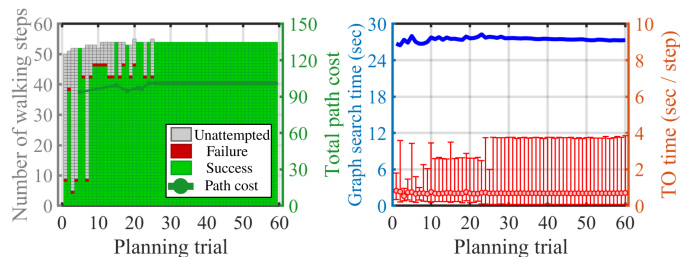



Fig. 4: Rebar grid used in case study along with reference trajectory visualization.



(a) Results for experience accumulation with guiding torso path.



(b) Results for experience accumulation without guiding torso path.

Fig. 5: Results for rebar grid traversal case study. Left plots showcase results of all attempted subproblems – success, failure, or not attempted due to early trial termination – as well as the total path costs for the trials that reached the goal. Right plots show computation times of the graph search triggered for each planning trial as well as average, minimum, and maximum TO solve times across all of the attempted subproblems within each planning trial.

tries to take the shortest path from start to goal which runs through the obstacle, leading to much higher TO times and less informative experience. One drawback of instituting the guiding torso path is that the graph search times increase significantly due to the introduction of the complicated torso path deviation term into the graph edge weight function.

VI. CONCLUSION

In this work, we adapt an efficient multi-modal contact planner to the task of quadrupedal rebar traversal. In the future, we intend on deploying this framework as an online, perception-informed navigation framework. To do so, we plan on using perception-aware navigation planners [8], [9] to generate guiding torso paths and incorporating vision-based heuristics into the mode transition graph search [10], [11]. We would like to acknowledge SkyMul [12] for providing robot and rebar grid models to use for testing.

REFERENCES

- [1] Z. Kingston and L. E. Kavraki, "Scaling Multimodal Planning: Using Experience and Informing Discrete Search," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 128–146, Feb. 2023.
- [2] Clearpath Robotics, "Clearpath Additional Simulation Worlds." [Online]. Available: https://github.com/clearpathrobotics/cpr_gazebo
- [3] L. Wellhausen and M. Hutter, "Rough Terrain Navigation for Legged Robots using Reachability Planning and Template Learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2021.
- [4] M. Geisert, T. Yates, A. Orgen, P. Fernbach, and I. Havoutis, "Contact Planning for the ANYmal Quadruped Robot using an Acyclic Reachability-Based Planner," Apr. 2019, arXiv:1904.08238 [cs]. [Online]. Available: <http://arxiv.org/abs/1904.08238>
- [5] O. Melon, R. Orsolino, D. Surovik, M. Geisert, I. Havoutis, and M. Fallon, "Receding-Horizon Perceptive Trajectory Optimization for Dynamic Legged Locomotion with Learned Initialization," Apr. 2021, arXiv:2104.09078 [cs]. [Online]. Available: <http://arxiv.org/abs/2104.09078>
- [6] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, Nov. 2014, pp. 295–302, iSSN: 2164-0580.
- [7] "OCS2: An open source library for Optimal Control of Switched Systems." [Online]. Available: <https://github.com/leggedrobotics/ocs2>
- [8] M. Asselmeier, Y. Zhao, and P. A. Vela, "Dynamic Gap: Formal Guarantees for Safe Gap-based Navigation in Dynamic Environments," Oct. 2022, arXiv:2210.05022.
- [9] S. Feng, Z. Zhou, J. Smith, M. Asselmeier, Y. Zhao, and P. A. Vela, "GPF-BG: A hierarchical vision-based planning framework for safe quadrupedal navigation," *Under Review*, 2022.
- [10] D. Driess, J.-S. Ha, and M. Toussaint, "Deep Visual Reasoning: Learning to Predict Action Sequences for Task and Motion Planning from an Initial Scene Image," in *Robotics: Science and Systems XVI*. Robotics: Science and Systems Foundation, Jul. 2020.
- [11] D. Driess, O. Oguz, J.-S. Ha, and M. Toussaint, "Deep Visual Heuristics: Learning Feasibility of Mixed-Integer Programs for Manipulation Planning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 9563–9569.
- [12] "SkyMul: Modular Robots for Construction." [Online]. Available: <https://skymul.com/>