



AFRL-AFOSR-JP-TR-2024-0025

Development of tele-operated quadrupedal robotic platform for disaster response

Kim, Sangbae
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVE
CAMBRIDGE, MA, 02139
USA

12/12/2023
Final Technical Report

<p>DISTRIBUTION A: Distribution approved for public release.</p>

Air Force Research Laboratory
Air Force Office of Scientific Research
Asian Office of Aerospace Research and Development
Unit 45002, APO AP 96338-5002

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE 20231212		2. REPORT TYPE Final		3. DATES COVERED	
				START DATE 20170425	END DATE 20200424
4. TITLE AND SUBTITLE Development of tele-operated quadrupedal robotic platform for disaster response					
5a. CONTRACT NUMBER		5b. GRANT NUMBER FA2386-17-1-4661		5c. PROGRAM ELEMENT NUMBER 61102F	
5d. PROJECT NUMBER		5e. TASK NUMBER		5f. WORK UNIT NUMBER	
6. AUTHOR(S) Sangbae Kim					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MASSACHUSETTS INSTITUTE OF TECHNOLOGY 77 MASSACHUSETTS AVE CAMBRIDGE, MA 02139 USA					8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD UNIT 45002 APO AP 96338-5002				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR IOA	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-JP-TR-2024-0025
12. DISTRIBUTION/AVAILABILITY STATEMENT A Distribution Unlimited: PB Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The RPC was implemented successfully on the MIT Cheetah 3 robot platform. The optimization is solved using the freely available IPOPT NLP solver [18] interfaced in C++. It runs on a separate thread, but on the same ADL embedded Quad core PC with a 2nd Gen Core i7 CPU as the main control loop as described in [4]. As mentioned throughout the report, the RPC runs asynchronously to the main control loop. Therefore while the control loop runs at a fixed 1kHz, the RPC solution time varies based on computing power available.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:				17. LIMITATION OF ABSTRACT	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR		18. NUMBER OF PAGES 23
19a. NAME OF RESPONSIBLE PERSON JERMONT CHEN					19b. PHONE NUMBER (Include area code) 315-227-7003

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

AFOSR Progress Report

Award Number: FA2386-17-1-4661

Project Title: Development of tele-operated quadrupedal robotic platform for disaster response

Report Type: Project Report

PI: Sangbae Kim

Awardee: Massachusetts Institute of Technology

Program Officer Name: Jermont Chen

Program Officer Email: jermont.chen.1@us.af.mil

1 Accomplishments

1.1 What are the major goals of the project?

Research on legged robots has seen a recent growth in interest due to advances that have demonstrated highly dynamic locomotion and ability to traverse difficult and unstructured terrains. This is in part because of developments resulting in more sophisticated controllers capable of dealing with increasingly difficult situations, as well as improved computing and actuation methods capable of robustly executing these controllers in the real world. Optimization-based controllers for legged robots have been successful in balancing, locomotion, and disturbance rejection. However, they are still lacking when attempting to solve more complex problems due to convergence and timing difficulties.

Legged robots currently use a variety of optimization-based locomotion controllers. The Atlas humanoid robot uses optimization throughout its control architecture [15]. HyQ at IIT uses a simple template model in a whole-body




Figure 1: **MIT Cheetah 3 Locomotion.** The RPC optimization finds ground reaction forces and footstep locations in real-time for realizing dynamic legged locomotion.

controller for instantaneous optimization of ground reaction forces [11] with heuristic planning to modify the desired inputs and deal with the aspects of locomotion that do not directly contribute to force controlled balance such as swing leg reflex [12]. ANYmal from ETH Zurich has implemented a hierarchical whole-body controller to execute tasks in order of priority for instantaneous force control of dynamic gaits [7, 2].

Instantaneous whole-body controllers have the limitation of only being aware of the current state rather than having knowledge about the gait and upcoming states meaning that they cannot easily control gaits with flight or longer periods of underactuation. For this reason, a heavy focus has been placed on predictive controllers, with Model Predictive Control (MPC) being the main framework of choice. MPC has been shown for many years to be theoretically successful in simulation [19, 9, 1]. However, robust hardware implementations are only recently beginning to show good performance. A whole-body MPC implementation was shown on the HRP-2 humanoid platform [14]. In this work, the authors state present many of their solutions to address the

difficulties in whole-body MPC, namely the computational cost, undesirable local minima, and discontinuities in control.

MPC is commonly used as a higher level planner rather than as the primary stabilizing controller. Many of these techniques solve the optimization for a large number of footsteps and then later execute the plan using another controller to execute and adapt the plan online. Work by Winkler et al. generates footholds and CoM trajectories for a certain look ahead period and then has a lower level controller track the motion in practice [20, 21]. These fast techniques for simultaneous foothold and CoM planning have been shown on hardware, but are geared towards planning for execution rather than real-time reactive adaptation.

Real-time nonlinear MPC techniques are gaining popularity with the improvements in computing capabilities. Quadcopters have shown success in fast implementations of NMPC [16, 13]. In legged robots, the LOLA robot uses real-time NMPC to adapt its footstep placements online under disturbances [22]. Farshidian et al. demonstrate their FastSLQ-MPC to generate a real-time whole-body nonlinear MPC to create trotting gaits [10]. A real-time whole-body NMPC implementation is shown on hardware in [17] where the robot is able to stand and do slow trotting. However, they note that computation and real-world limitations must still be overcome to take full advantage of the nonlinear optimization’s capabilities.

In this work, the controller is designed for the MIT Cheetah 3 robot platform pictured in Figure 1. Two separate balance controllers have been implemented on the robot. The first being a QP based balance controller that uses instantaneous ground reaction forces at the contact points to regulate the CoM states [4]. The second being a linear, convex MPC implementation which is able to stabilize various dynamic gaits as described in [6]. However, neither implementation optimizes footstep locations and assumes the designed footsteps are adequate. Previous work showed promising results with a similar quadruped model using nonlinear optimization to pick both footsteps and ground reaction forces in simulation [5]. This research presents the Regularized Predictive Control (RPC) hardware implementation which is an extension of that work that relies heavily on heuristic regularization rather than a high-fidelity dynamics model.

1.2 Contribution

The contribution of this work lies with the culmination of various important research efforts into a functional hardware implementation. As pointed out, the system integration of a nonlinear optimization-based controller is non-trivial and as such it is not yet widely used as the main controller in legged robots. It is the intention of this work to present some of the more critical details that led to the successful implementation of the RPC framework for real-time control of a dynamic quadruped. The hope is that other NMPC-style methods can benefit and overcome the computation limitations that currently plague real-time nonlinear optimization for control. The results show the nonlinear predictive controller is able to perform as adequately or better than the previous controllers on the robot during the experiments. We believe that this is a big step in showing that nonlinear predictive controllers are approaching a point where they can be used robustly to control dynamic legged robots.

We describe the nonlinear RPC optimization framework with the cost function and constraints as written in the theory. Then present a survey of the implementation details that allow the theoretical controller to be implemented on the hardware in real-time. Results are shown for various experimental situations. Finally, we will discuss brief conclusions of the work.

2 Major Activities

2.1 Nonlinear RPC Optimization Framework

The RPC framework described here is an extension of previous work by the authors in [5]. That work showed that adding heuristic regularization directly into the optimization allowed the controller to exploit the known dynamics or the system resulting in faster computation time and better conditioned cost spaces. However, that version was presented purely in simulation where time could be completely stopped while the optimization searched for a solution. This is not realistic on the actual hardware and so while the theoretical controller and optimization framework has not been modified significantly, the details of how it is actually used are drastically different.

The optimization uses massless leg and linear CoM dynamics assumptions to simplify some of the more nasty nonlinearities and make the dynamics calculation and its associated gradients simple to solve. This allows us to use

the dynamic model as a rough feasibility constraint rather than a strict dominating condition. Finally the optimization framework with its cost function and all explicitly written constraints is posed as

$$\begin{aligned}
& \min_{\mathbf{x}, \mathbf{u}} \quad \sum_{k=0}^{N-1} \underbrace{(\tilde{\mathbf{x}}_k^T \mathbf{Q}_k \tilde{\mathbf{x}}_k)}_{\text{Dynamics}} + \underbrace{\tilde{\mathbf{u}}_k^T \mathbf{R}_k \tilde{\mathbf{u}}_k}_{\text{Regularization}} + \underbrace{\tilde{\mathbf{x}}_N^T \mathbf{Q}_N \tilde{\mathbf{x}}_N}_{\text{Terminal State}} \\
& \text{subject to} \quad \text{Simplified Discrete Dynamics} \\
& \quad \mathbf{x}_{k+1} - (\mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k h(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k) = 0 \\
& \quad \text{Kinematic Leg Limits} \\
& \quad \|\mathbf{r}_{ref,i} - \mathbf{r}_{h,i}\| - z_0 \tan(\frac{\pi}{6}) \leq 0 \\
& \quad \text{Foot placed on ground} \\
& \quad \hat{z}_{g,i}(p_{x,i}, p_{y,i}) - p_{z,i} = 0 \\
& \quad \text{Foot Stationary During Stance} \\
& \quad \mathbf{s}_{\phi,k+1} \mathbf{s}_{\phi,k} (\mathbf{p}_{k+1} - \mathbf{p}_k) = 0 \\
& \quad \text{Positive Ground Force Normal} \\
& \quad -\hat{\mathbf{f}}_i \cdot \nabla \hat{\mathbb{G}}(\hat{\mathbf{p}}_f) \leq 0 \\
& \quad \text{Lateral Force Friction Pyramids} \\
& \quad -\mu \mathbf{f}_{z,i} \leq \mathbf{f}_{x,i} \leq \mu \mathbf{f}_{z,i} \\
& \quad -\mu \mathbf{f}_{z,i} \leq \mathbf{f}_{y,i} \leq \mu \mathbf{f}_{z,i}
\end{aligned}$$

The framework enforces the necessary constraints for physical realizability characteristic of a point foot legged system as depicted in Figure 2. For example the forces are constrained so that the feet cannot pull on the ground. However, we make the assumption that the feet will not slip on the ground and write it as a constraint even though it is possible for the feet to slip. We attempt to avoid this by adding the friction constraint to avoid large lateral forces. Although many of the terrain conditions are explicitly written in the constraints, it is worth noting that the hardware experiments in this shown are conducted through blind locomotion, meaning that the robot has no external perception of the environment. Therefore assumptions about the terrain are made for the optimization constraints.

While regularization is often used in optimal control, this method attempts




Figure 2: **RPC Optimization.** Physically realistic constraints are enforced for the simplified dynamics, footsteps locations, and ground reaction forces.

to regularize to a time-varying, nonlinear function composed of heuristics designed to closely approximate the optimal solution in various operating situations. When implemented on the hardware, the choice and design of the heuristic regularization became much more important than the actual dynamics model being optimized over. Solutions perform well towards the beginning of the prediction horizon, but breakdown towards the later parts of a long horizon as the dynamics are not highly accurate to the system. For this reason, it is desirable to constantly replan as fast as possible.

2.2 Implementation Details

The theory behind the controller has been shown in simulation, but robotics requires engineering to actually implement. While in Section 2.1 the overall block diagram was presented with a single block representing the theoretically derived RPC, in practice it is not so simple to just run the optimization assuming that it will be solved instantaneously. Figure 4 extends on the theoretical derivation and shows the algorithms and modifications designed

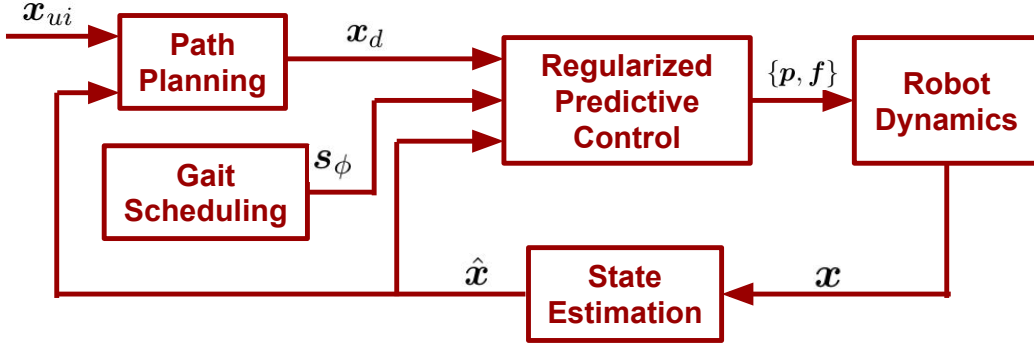


Figure 3: **Simplified Control Framework.** Simplified block diagram showing the general flow of information throughout the control framework. In theory, the RPC block is simply the optimization with no modifications needed.

to aide the optimization overcome the real-time implementation challenges. This section describes the blocks in the diagram and their function in the control framework resulting in successful locomotion on the hardware platform.

A lot of the success comes from managing the tradeoff between accuracy and computation time, as well as preprocessing the NLP inputs by exploiting prior knowledge of the system. The MIT Cheetah 3 robot is designed to be robust to inaccuracies in ground reaction forces and has been shown to be adequately modeled with massless legs which affords the ability to use the simple dynamics model in the optimization. Preprocessing is a one time calculation per solution that has powerful effects on the conditioning of the problem’s initial guesses and cost landscape.

2.2.1 Direct Transcription

Various optimization methods are available for solving optimal control problems. The choice of which has large effects on the solution accuracy and speed with each having their own benefits and drawbacks. For the controller we decided to use direct transcription which uses piecewise constant input and a piecewise linear state trajectories with both the states and the inputs used as decision variables in the optimization [8].

By posing the optimization as a direct transcription we lose some of the accuracy we may gain using a more sophisticated method, but also reduce the computation times. The loss in accuracy is also not critical since the

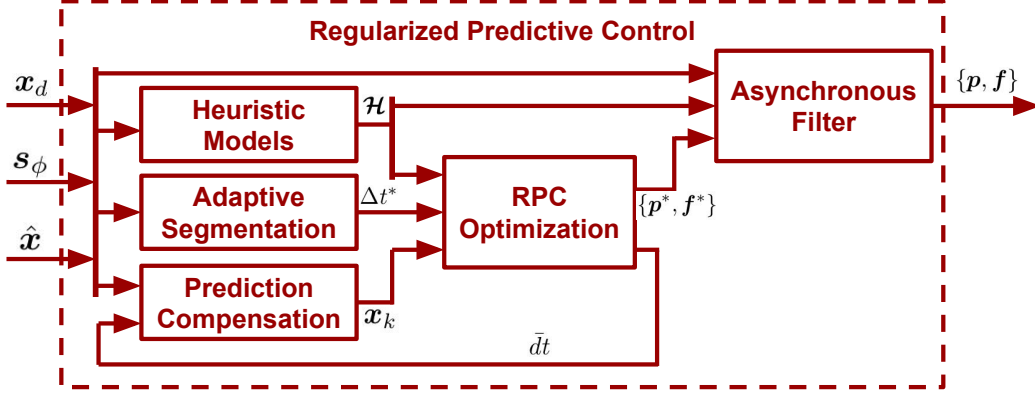


Figure 4: **Modified RPC Implementation.** The theoretical RPC optimization is aided by various other algorithms to overcome some of the limitations and difficulties that come with real world implementation.

model we are using is an approximation rather than high-fidelity. As such, tolerances and convergence conditions can be greatly relaxed. It is not imperative to solve for a high degree of optimality on a model that is already a large approximation. We found that while the predicted robot states further along the prediction horizon rarely match with the actual states at that future time, the robot states generally trended towards the predicted directions. However, since the controller is constantly replanning, new inputs are modified to account for the discrepancies arising from the coarseness and simplicity of the optimization model dynamics.

With direct transcription the dynamics constraints are only dependent on the current timestep and the next timestep. This implicitly constraints the dynamics to be enforced throughout the prediction rather than explicitly adhering to the dynamics from the first to the last timestep. Similarly, even though footsteps may extend over several timesteps, by only constraining adjacent stance positions the constraint does not need to be explicitly written over all the relevant timesteps. This removes a lot of redundancies and better conditions the problem to be a series of small linear approximations linked together to approximate the overall nonlinear function. Only each small step needs to be constrained to be dynamically feasible rather than the full, highly nonlinear function. This results in more accurate dynamics between timesteps as the linear dynamics hold more fidelity to the real nonlinear dynamics over smaller time intervals.

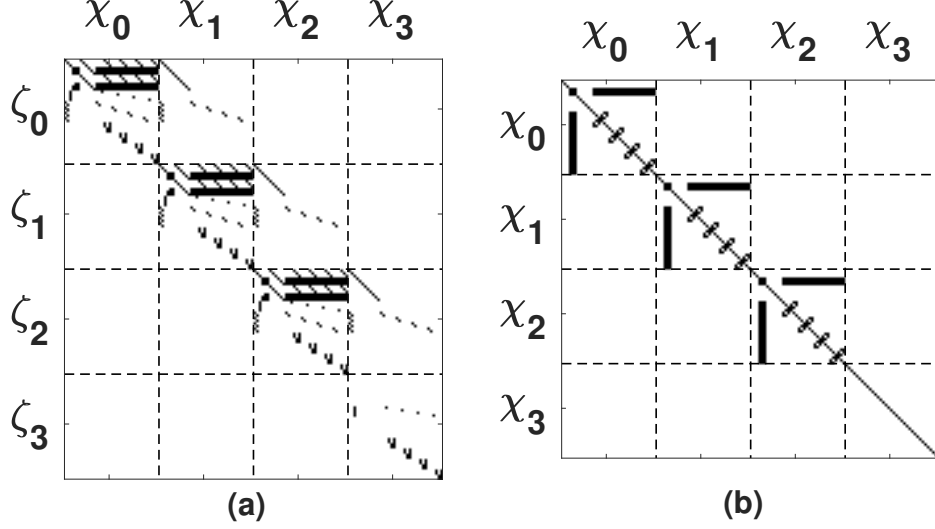


Figure 5: **Sparsity Patterns.** Both the constraint Jacobian (a) and Lagrangian Hessian (b) feature a highly sparse, repeated structure throughout the matrix which contributes to the flexibility and speed of the algorithm.

Consequently, Figure 5 shows the sparsity matrices for the Constraint Jacobian (5a) and the Hessian of the Lagrangian (5b) where the dark spots indicate the indices in the matrices where the entries are possible to be non-zero. χ signifies the decision variables in the optimization, ζ is the set of constraints, while the subscripts denote the prediction timestep number. In this example $N = 4$, but in practice it can be chosen to be anything. By design, the matrices are both sparse and although the size of the Constraint Jacobian and Lagrangian Hessian grow proportional to N^2 , the non-zero entries of each grows linearly as, $\mathcal{O}(N)$.

All parts of the equations and constraints that relate to contact points are defined by the contact state boolean variable, $\mathbf{s} = \{0 = \text{swing}, 1 = \text{contact}\}$. This provides a simple way to automatically adapt the active decision variables at each contact mode without having to explicitly write a different optimization for each combination of foot contacts. This propagates through the prediction horizon and activates or deactivates different contact modes in the gradients and Hessians according to the gait schedule regardless of the order or duration of each.

2.2.2 Adaptive Timing Segmentation

Since Legged locomotion is defined by the discrete changes between contact and non-contact modes of each of the legs, the optimization timestep segmentations must reflect these modes. As the controller is run continuously rather than at a scheduled time, the time step vector to be optimized over cannot simply be prescribed. As such, an algorithm must be developed to adaptively segment the time vector regardless of the scheduled gait pattern. The algorithm takes a nominal discretization timestep sequence and attempts to calculate the closest set of timestep segmentations while respecting the gait schedule. Discrete scheduled contact switches must signify a new division as the robot enters a new control mode. Since this often happens irregularly, the algorithm chooses the number of timesteps within the same contact mode and spaces them out evenly. Algorithm 1 presented on the next page outlines this process. This allows for an incredible amount of freedom in choosing gait patterns and frequencies. Instead of constraining the predicted timestep vector to be divided evenly, any gait schedule can be optimized over without needing to make any change in the controller.

Algorithm 1: Adaptive Timestep Segmentation

```

1 initialization;
2 nominal timestep:  $\Delta t_{nominal}$ ;
3 scheduled contact to swing phase:  $\phi_{c \rightarrow \bar{c}} = \frac{T_p - T_{\bar{c}}}{T_p}$ ;
4 for  $k = 0$  to  $NUM\_PREDICTIONS - 1$  do
5     number of feet in swing:  $N_{\bar{c}} = 0$ ;
6      $\Delta t_{max} = \infty$ 
7     for  $f = 0$  to  $NUM\_FEET - 1$  do
8         if  $\phi_{f,k} < \phi_{c \rightarrow \bar{c}}$  then
9             foot in contact:  $s_{f,k} = 1$ ;
10            time until foot switches to swing:  $\Delta t_{switch} = T_p(\phi_{c \rightarrow \bar{c}} - \phi_{f,k})$ ;
11        else
12            foot out of contact:  $s_{f,k} = 0$ ;
13            time until foot switches to contact:  $\Delta t_{switch} = T_p(1 - \phi_{f,k})$ ;
14             $N_{\bar{c}}++$ ;
15        if  $\Delta t_{switch} < \Delta t_{max}$  then
16             $\Delta t_{max} = \Delta t_{switch}$ ;
17    choose the current segment duration:  $\Delta t_k^*$ ;
18    if  $NUM\_FEET = N_{\bar{c}}$  then
19         $\Delta t_k^* = \Delta t_{max}$ ;
20         $flight\_phase = 1$ ;
21    else if  $\Delta t_{max} < \Delta t_{nominal}$  then
22         $\Delta t_k^* = \Delta t_{max}$ ;
23         $flight\_phase = 0$ ;
24    else
25        nominally divide time till switch evenly:  $\Delta t_k^* = \frac{\Delta t_{max}}{round(\frac{\Delta t_{max}}{\Delta t_{nominal}})}$ ;
26         $flight\_phase = 0$ ;
27     $\phi_{f,k+1} = \text{mod}(\phi_{f,k} + \frac{\Delta t_k^*}{T_p}, 1)$ ;

```

Figure 6 shows an example of how the algorithm modifies the nominal, evenly spaced prediction timestep vector for several gaits. The algorithm always has a segmentation when a contact state is scheduled to be changed from contact to swing or swing to contact while attempting to best segment longer periods evenly. The ambling gait has quick contact switches with 8 over the full gait

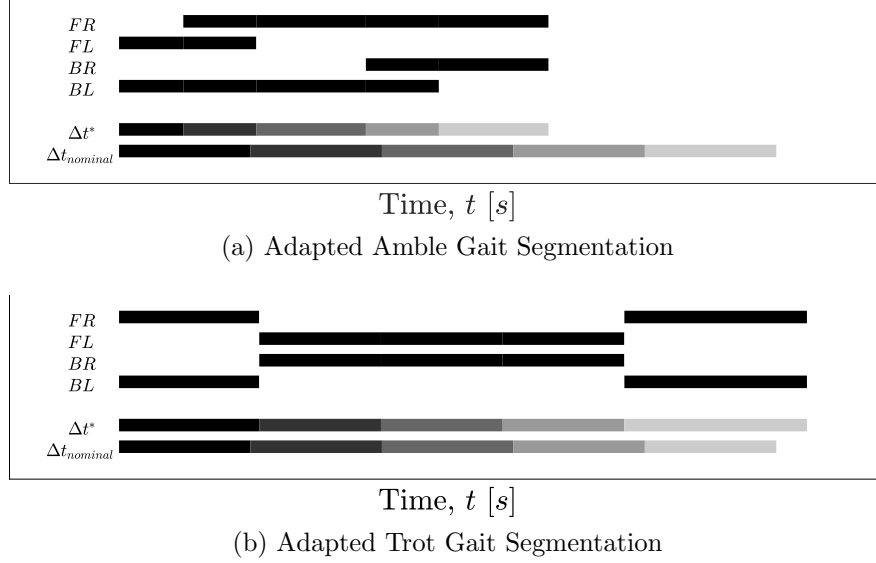


Figure 6: **Adaptive Timestep Segmentation.** Top four bars represent the upcoming gait scheduled contact sequences for each of the feet. The bottom bars are the nominal timestep segmentation and the adapted segmentation based on contact mode switches where each segment gets lighter further along the prediction horizon.

cycle and therefore the algorithm shortens the overall prediction horizon. The trotting gait has 2 contact switches over the nominal horizon, since diagonal pairs of legs switch simultaneously and can be captured more easily by the nominal prediction horizon. For those more complex gaits, it is easy to simply increase the number of prediction discretization in the optimization. This algorithm is $\mathcal{O}(N)$, so it is straight-forward and computationally inexpensive to add checks for various boolean flags that get used in the calculations. For example, if a foot is determined to be in contact, we can check if it was previously in contact or not, which signifies a touchdown event. Keeping track of touchdown events lets the optimization know if it should use the heuristic reference footstep location or the footstep location from the previous iteration during initialization. The no foot slip constraints will be enforced either way so it is not necessary to explicitly flag touchdown events, but initializing and regularizing with already correct footstep logic decreases the number of iterations and complexity needed to converge.

2.2.3 Prediction Delay Compensation

Since the optimization is running at a different frequency on a separate thread from the main control loop, there are often timing delays and discrepancies between the current state of the robot with its associated scheduled stance legs and the optimal solution being returned by the optimization. Namely, that the prediction solve time computation is slower than the $1kHz$ main control loop and may return desired forces for legs that have since switched into a different scheduled contact state.

The state used as the initial condition for the optimization, \mathbf{x}_k is not actually the current state of the robot but rather an expected predicted future state solved for by the linear dynamics approximation

$$\mathbf{x}_k = \mathbf{A}(\bar{dt})\hat{\mathbf{x}}(t) + \mathbf{B}_k(\bar{dt})h(\hat{\mathbf{x}}(t), \mathbf{u}(t)) + \mathbf{d}_k(\bar{dt}) \quad (1)$$

where $\hat{\mathbf{x}}(t)$ is the actual current estimated robot state and $\mathbf{u}(t)$ is the current input from the last solution. The dynamics are integrated forward with \bar{dt} being the filtered average optimization solve time.

This starts the optimization at the approximate state that the robot can be expected to be in when the optimization finishes. This is more appropriate to calculate inputs for than the instantaneous state when starting the prediction. If the instantaneous state is used, by the time the solution is returned, the calculated inputs are optimized for a state that the robot is no longer in. By predicting forward to a future state as the initial state to optimize over, it is more likely that the solution is optimized for a state close to the robot's current state.

Similarly, the phase of each foot will be different by the time the optimization returns a value and therefore may also be scheduled to be in a different contact state than it is currently in. Each leg phase can therefore be modified according to the filtered average solve time with

$$\bar{d}\phi = \frac{\bar{dt}}{T_P} \quad (2)$$

$$\phi_k = \text{mod}(\phi(t) + \bar{d}\phi, 1) \quad (3)$$

where T_P is the nominal gait period time and $\phi(t)$ is the current phase of the foot. Again, this starts the prediction with the contact mode that the robot will likely be in when a solution is returned. If a foot is nearing the end of stance, it is not a good idea to solve for a ground reaction force assuming

that the foot will be in stance when the optimization converges. If there is a mode change before the end, then whatever solution was found is no longer possible to execute. By solving with the expected future state of the robot, it is more likely that forces will be found for the appropriate contact feet.

2.2.4 Extracted Heuristic Models

In previous work it was demonstrated that the performance of the controller could improve drastically simply by adding improved heuristic regularization to the RPC. Particularly since most NLP solvers are designed for smooth continuous problems and legged locomotion is by definition a discrete hybrid problem, the regularization heuristics help guide and attract the solutions towards known operating regions. Discontinuities can prevent the solution from converging or finding local minima so the regularization warps the cost space in a favorable way to find a solution quickly.

In the majority of situations, the heuristics provide a "good enough" solution for locomotion. By simply executing the heuristics, the robot should be able to remain upright for at least a small amount of time. This gives the RPC enough time to return a new solution without falling over. The heuristics are embedded directly into the optimization through the error term in the quadratic cost functions as

$$\tilde{\mathbf{x}} = \mathcal{H}_x(\hat{\mathbf{x}}, \mathbf{x}_d) - \hat{\mathbf{x}} \quad (4)$$

$$\tilde{\mathbf{u}} = \mathcal{H}_u(\hat{\mathbf{x}}, \mathbf{x}_d) - \mathbf{u} \quad (5)$$

This regulates the states and inputs towards the designed heuristics rather than the simple user inputted commands. Desired commands act as a higher level goal for the robot, but are modified with the heuristics in order to better solve the optimization.

Designing these heuristics can be challenging, but the heuristic extraction framework discussed in [3] presents methods for determining heuristics from data analysis as well as expert design. The results in the work showed a roughly $\sim 2x$ improvement in maximum forward, lateral, and turning velocities simply by exploiting the heuristics. No changes were made to the control gains or the controller structure to incite the performance boost. For the purposes of this work, we will assume that a set of heuristics, \mathcal{H} , is pre-determined.

2.2.5 Asynchronous Solution Filtering

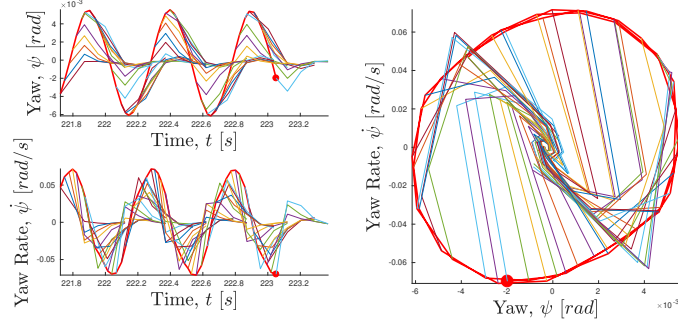
The prediction delay compensation in Section 2.2.3 starts the optimization acting on a future state and possibly contact mode based on the moving average of the optimization solution time. However, this means that if the actual gait schedule has a differing contact mode than the resulting RPC solution, there must be a way to deal with this discrepancy. For both the footstep and ground reaction forces of each leg, the nominal values are given to be the extracted heuristics discussed in Section 2.2.4. As stated, the heuristics are designed precisely for this purpose. Therefore in the event that the optimization does not return a solution or cannot converge before the contact mode changes, the robot simply executes the heuristic and is able to temporarily get away without a new optimization solution. This temporary control action is designed for idealized conditions and will not stabilize the robot indefinitely. However, it will provide a "good enough" solution to avoid a fall.

This is especially important in gaits that have very distinct contact modes, such as trotting. In a version of the trot gait, the diagonal pairs of legs are picked up and put down in an alternating pattern. Even with the prediction delay compensation, it is not guaranteed that the predicted phase and the scheduled phase signal the same contact state. Therefore, without this heuristic there would be no force command from the stance legs to keep the robot upright. When the optimized solution contact mode and the scheduled gait sequence return to a synchronized state, the results from the RPC can be confidently used again.

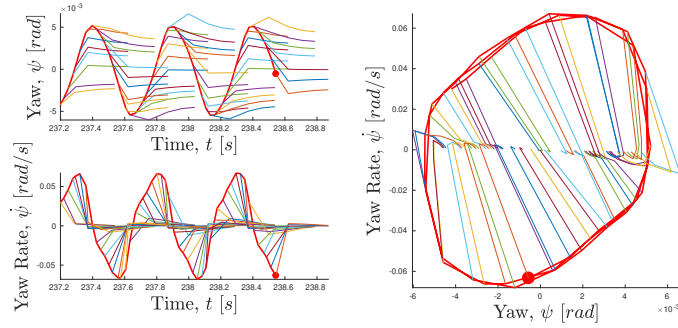
2.2.6 Gain Tuning

With most controllers, the tuning of the gains is a critical component to its success. To tune the gains, simple desired setpoints were used as we automatically swept over the individual gains for each state. Rough bounds of stability for the gains were found that allowed the robot to automatically tune itself without falling over in simulation by simply finding the gains that produced the best tracking.

Figure 7 shows the two extreme bounds of tuning the yaw behavior of the robot while trotting. While both are stable, they produce very different results. The red lines are the actual measured states of the robot and the multicolored lines are the predicted states returned from the RPC. Low yaw



(a) Low Gain, $R_\psi = 1$



(b) High Gain, $R_\psi = 500$

Figure 7: **Intuitive Gain Tuning.** Red lines indicate the actual motion of the robot and the thin multicolored lines represent the predicted states at each moment in time as the yaw rate gain is swept from $R_\psi = 1$ in (7a) to $R_\psi = 500$ in (7b).

rate gain, R_ψ , produces a more natural and smooth sinusoidal pattern when not forcing the yaw rate to go to zero as is the case with the high gain. The tracking error for the states is used as a metric for tuning, as is the predicted state error. When the predicted states also have good tracking, then the controller is better conditioned and results produce smoother, less sporadic results.

The predicted state trajectory does not line up with the actual state because the prediction model is based on simple uncoupled linear dynamics while the actual robot is a highly nonlinear hybrid system that must adhere to its natural physics regardless of how much the controller fights it. This is addressed in related work. However, for the purposes of tuning the gains, it is enough to attempt to track as closely to a simple setpoint as possible.

3 Significant Results

The RPC was implemented successfully on the MIT Cheetah 3 robot platform. The optimization is solved using the freely available IPOPT NLP solver [18] interfaced in C++. It runs on a separate thread, but on the same ADL embedded Quad core PC with a 2nd Gen Core i7 CPU as the main control loop as described in [4]. As mentioned throughout the report, the RPC runs asynchronously to the main control loop. Therefore while the control loop runs at a fixed $1kHz$, the RPC solution time varies based on computing power available.

We note that the stability of the robot improves as solve frequency increases, with good stability over $40Hz$, decent stability $20-40Hz$, and unstable under $20Hz$. Since generally the heuristic-based initial guess is designed to provide an adequate solution to at least remain stable, we can limit the maximum solve time to be $50ms$ which corresponds to the lowest stable frequency and rely on the heuristics if a solution was not found. In practice the robot generally runs at $60-80Hz$ on the robot during steady state locomotion and around $40-60Hz$ during worst-case large disturbances as the optimization must search further from the heuristic regularization for better solutions.

The xyz ground reaction forces resulting from the RPC are shown in Figure 8. The gait schedule contact sequence for the trotting gait is also depicted for the six footsteps. Each diagonal pair of legs is capable of stabilizing the robot while trotting in place. A discrete jump in the forces can be seen at the beginning of each new contact mode, signifying the switch from the short period of heuristic force calculation to the RPC solution.




Figure 8: **Ground Reaction Forces.** Optimized ground reaction forces during a trotting gait. Forces are smoothly filtered even though the RPC returns discrete constant forces at a lower frequency than the control loop.

In addition to selecting the ground reaction forces at the contact points, future footsteps are also chosen. The benefit of simultaneously optimizing for both is that footsteps are chosen in the same context as future stabilizing forces. Figure 9 shows two example predicted footstep optimization results during $1.4 \frac{m}{s}$ forward locomotion using the robot with a trotting gait. The

figure depicts a 2D projection in the XY plane of the (x, y) coordinates for the CoM, predicted CoM, contact feet, and predicted footstep locations. In Figure 9a, the front left (blue) and back right (green) feet are on the ground while the next footstep locations for the front right and back left feet are solved for. The prediction horizon is long enough such that the next footsteps for the front left and back right feet are also already being solved for before they have entered their swing phase.




Figure 9: **Predicted Footstep Locations.** Solid colored dots represent the current contact foot locations while the colored circles signify the predicted footstep locations. The solid circle inside of the predicted footsteps corresponds to the future time segment of touchdown, where the darker the circle, the closer to the current time. CoM predicted trajectory is shown in the middle with decreasingly dark dots over time.

The robot was able to closely track the desired velocities during the experiment both forwards and backwards. The body's orientation throughout is modulated close to level with minor fluctuations, never exceeding more than 5° which usually occurs during the sharp velocity changes. However, it returns flat when the robot is in steady state. Figure 10 shows the robot's ability to track commanded translational velocities. The maximum commanded velocity is $1.4 \frac{m}{s}$ during trotting. Turning rate commands were also sent to the robot as seen in Figure 11. Maximum turning rate given was $2 \frac{rad}{s}$ with a mean orientation error in roll of 0.47° and 0.48° in pitch throughout.




Figure 10: **Velocity Tracking.** The robot was given forward and backwards velocity commands in the x direction and was able to closely track the desired values while keeping its body orientation close to flat.

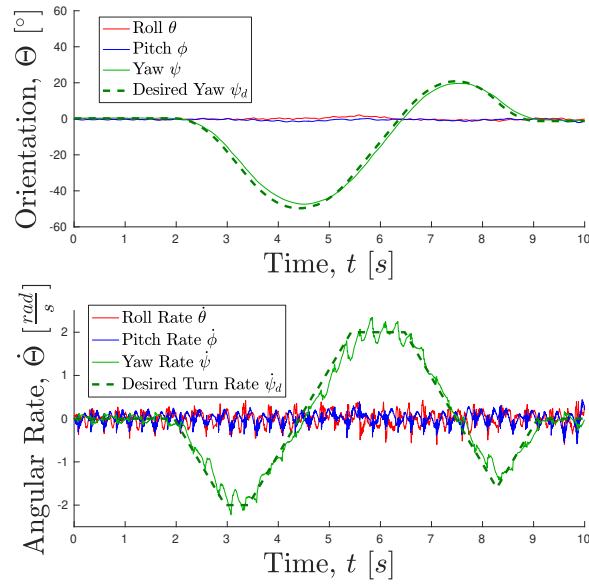


Figure 11: **Turn Rate Tracking.** The robot was given a turning rate command which it was able to track with little error over the test duration.

References

- [1] T. Apgar, P. Clary, K. Green, A. Fern, and J. W. Hurst. Fast online trajectory optimization for the bipedal robot cassie. In *Robotics: Science and Systems*, 2018.
- [2] C. D. Bellicoso, F. Jenelten, C. Gehring, and M. Hutter. Dynamic locomotion through online nonlinear motion optimization for quadrupedal robots. (99), 2018-01-17.
- [3] G. Bledt and S. Kim. Extracting legged locomotion heuristics with regularized predictive control. In *Proceedings of Robotics: Science and Systems*, Freiburg, Germany, June. 2019. Submitted.
- [4] G. Bledt, M. J. Powell, B. Katz, J. D. Carlo, P. M. Wensing, and S. Kim. MIT cheetah 3: Design and control of a robust, dynamic quadruped robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Madrid, Spain, Oct. 2018.
- [5] G. Bledt, P. M. Wensing, and S. Kim. Policy-regularized model predictive control to stabilize diverse quadrupedal gaits for the mit cheetah. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vancouver, Canada, Sept. 2017.
- [6] J. D. Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim. Dynamic locomotion in the MIT cheetah 3 through convex model-predictive control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Madrid, Spain, Oct. 2018.
- [7] C. Dario Bellicoso, F. Jenelten, P. Fankhauser, C. Gehring, J. Hwangbo, and M. Hutter. Dynamic locomotion and whole-body control for quadrupedal robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3359–3365, Sep. 2017.
- [8] M. Diehl, H. Bock, H. Diedam, and P.-B. Wieber. Fast direct multiple shooting algorithms for optimal robot control. In M. Diehl and K. Mombaur, editors, *Fast Motions in Biomechanics and Robotics*, volume 340 of *Lecture Notes in Control and Information Sciences*, pages 65–93. Springer Berlin Heidelberg, 2006.

- [9] D. Dimitrov, A. Sherikov, and P.-B. Wieber. A sparse model predictive control formulation for walking motion generation. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 2292–2299, sept. 2011.
- [10] F. Farshidian, E. Jelavic, A. Satapathy, M. Gifftthaler, and J. Buchli. Real-time motion planning of legged robots: A model predictive control approach. *CoRR*, abs/1710.04029, 2017.
- [11] M. Focchi, A. del Prete, I. Havoutis, R. Featherstone, D. G. Caldwell, and C. Semini. High-slope terrain locomotion for torque-controlled quadruped robots. *Autonomous Robots*, 41(1):259–272, Jan 2017.
- [12] M. Focchi, R. Orsolino, M. Camurri, V. Barasuol, C. Mastalli, D. G. Caldwell, and C. Semini. Heuristic planning for rough terrain locomotion in presence of external disturbances and variable perception quality. *CoRR*, abs/1805.10238, 2018.
- [13] M. Kamel, K. Alexis, M. Achtelik, and R. Siegwart. Fast nonlinear model predictive control for multicopter attitude tracking on $so(3)$. In *2015 IEEE Conference on Control Applications (CCA)*, pages 1160–1166, Sep. 2015.
- [14] J. Koenemann, A. D. Prete, Y. Tassa, E. Todorov, O. Stasse, M. Bennewitz, and N. Mansard. Whole-body model-predictive control applied to the hrp-2 humanoid. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 3346–3351, Sept 2015.
- [15] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, Mar 2016.
- [16] M. Neunert, C. de Crousaz, F. Furrer, M. Kamel, F. Farshidian, R. Siegwart, and J. Buchli. Fast nonlinear model predictive control for unified trajectory optimization and tracking. In *IEEE Int. Conf. on Rob. and Automation*, pages 1398–1404, May 2016.
- [17] M. Neunert, M. Stäuble, M. Gifftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli. Whole-body nonlinear model pre-

- dictive control through contacts for quadrupeds. *CoRR*, abs/1712.02889, 2017.
- [18] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, Mar. 2006.
 - [19] P.-B. Wieber. Trajectory free linear model predictive control for stable walking in the presence of strong perturbations. In *Humanoids*, 2006.
 - [20] A. W. Winkler, F. Farshidian, M. Neunert, D. Pardo, and J. Buchli. On-line walking motion and foothold optimization for quadruped locomotion. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5308–5313, 2017.
 - [21] A. W. Winkler, F. Farshidian, D. Pardo, M. Neunert, and J. Buchli. Fast trajectory optimization for legged robots using vertex-based ZMP constraints. *CoRR*, abs/1705.10313, 2017.
 - [22] R. Wittmann, A. Hildebrandt, D. Wahrmann, D. Rixen, and T. Buschmann. Real-time nonlinear model predictive footstep optimization for biped robots. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 711–717, Nov 2015.