

Disturbance Detection, Identification, and Recovery by Gait Transition in Legged Robots

Aaron M. Johnson, G. Clark Haynes, D. E. Koditschek

Electrical and Systems Engineering
University of Pennsylvania
200 S. 33rd St, Philadelphia, PA 19104
{aaronjoh, gchaynes, kod}@seas.upenn.edu

Abstract—We present a framework for detecting, identifying, and recovering within stride from faults and other leg contact disturbances encountered by a walking hexapedal robot. Detection is achieved by means of a software contact-event sensor with no additional sensing hardware beyond the commercial actuators’ standard shaft encoders. A simple finite state machine identifies disturbances as due either to an expected ground contact, a missing ground contact indicating leg fault, or an unexpected “wall” contact. Recovery proceeds as necessary by means of a recently developed topological gait transition coordinator. We demonstrate the efficacy of this system by presenting preliminary data arising from two reactive behaviors — wall avoidance and leg-break recovery. We believe that extensions of this framework will enable reactive behaviors allowing the robot to function with guarded autonomy under widely varying terrain and self-health conditions.

I. INTRODUCTION

This paper describes a software contact-event sensor designed to trigger a legged gait recovery transition on EduBot [1] (see Fig. 1), a relative of the RHex hexapedal robot [2]. There are two principal contributions: first, we adapt the traditional control theoretic framework of deterministic dynamical fault detection and recovery [3] to identify the need for a transition; second, we apply topologically informed gait control policies to achieve a smooth transition to desired gait timings that produce stable locomotion. In doing so, we take a small but important, novel step toward developing an operational framework for guarded autonomous legged locomotion in general terrain.

The most compelling case for legged locomotion arises from the promise of robust adaptation and graceful degradation of mobility performance in mechanically complex and highly varied environments and under conditions of changing or compromised self-health. To date most of the locomotion literature has addressed operation in the extremes: either fixed, consistent terrain wherein a specific gait can be tuned over repeated trials and accumulating experience [2, 4–8] or wildly varied footing conditions [9–12] wherein it is not at all clear that the notion of a gait is even appropriate. We are aware of only two investigations of adaptive legged locomotion in the presumably far more common middle-ground setting of challenging but “modestly” varying terrain: on-line deterministic gait parameter feedback [13]; and tuned robustness against stochastic perturbation of an open loop stride-map as an alternative to deterministic gaits [14]. In




Fig. 1: EduBot, a RHex-style hexapedal robot containing only a single actuator per leg and minimal sensor configurations.

contrast, although the promise of redundancy against individual joint or limb failure ought to be one of the major advantages of legged mobility, with few exceptions [15–17] there is little legged robotics literature on gait adaptation in the face of compromised self-health.

Inevitably, the question of how to respond to various alterations in the condition of the environment and state of self-repair hinges upon the issue of what sort of sensing is available. One way to detect the changed circumstances that may require an altered locomotion strategy is to instrument the legs with contact, force, strain, or other sensors that measure directly what is happening to them or their environment. However, instrumenting a leg may not be easy and will always have a cost both in terms of money and design constraints. In this paper, we further pursue the long-standing theme of sensor-minimal robotics [18, 19] applied to reactive locomotion in [20] and continued in [13]. Specifically, we introduce and study empirically an algorithm that can acquire the relevant information using an estimator driven only by a motor mounted encoder that would typically be included in any actuator package.

The paper is organized as follows. This introduction concludes with a more detailed look at the motivation for and prior literature most relevant to our software contact detector. Section II details its algorithmic constituents, briefly

describes the problem of legged gait transitions, and lays out the manner in which the contact detector is integrated into an existing transition planner to trigger the desired adaptation. Section III documents the operation of this leg-contact triggered gait transition mechanism in achieving two illustrative behaviors: wall avoidance and leg fault recovery. The paper concludes with a brief summary and assessment of next steps in Section IV.

A. Disturbance Detection

At its core, this paper documents the value of adding into a leg contact detector an internal state model patterned on the decades-long tradition of industrial on-line fault detection [3] (translated more recently into the setting of robot execution monitoring [21] and hybrid systems diagnosis [22]), suggesting the architecture depicted in Figure 2. As we will detail in Section II, the measured output (leg position) is compared with estimates generated by an independent dynamical observer to form a “residual” (error signal) containing clues about how the physical plant’s behavior departs from modeled expectations to be processed by downstream diagnostics. These estimates could also have been generated via a dynamic bayes network (as in [23]), or a particle filter method (as in [24]), or other estimation technique. However the targeted application domain presents very starkly and characteristically distinctive dynamics that seem well captured by the simple, deterministic models and well classified by the modest, deterministic finite state automaton we introduce. Very likely, in settings requiring the classification of many different terrains the more complex stochastic methods will justify their significantly greater calibration effort (e.g., selection of priors) and lengthier transients. Such an inquiry lies considerably beyond the scope of the present study.

In contrast, heretofore, we have relied upon memoryless contact detectors, for example, examining directly the discrepancy between commanded and actual motor shaft output [13, 20] and the difficulty in getting these schemes to function robustly serves as a strong motivation for the present work. We document in Section II the comparative benefit of this internal model approach to diagnostics relative to the memoryless alternatives. Such difficulties had previously motivated our group to undertake the significant effort of instrumenting a direct physical contact sensor [25], but this is a particularly challenging exercise on the continually circulating legs of RHex-style machines¹.

With non-recirculating legs the more modest cost and complexity of leg contact hardware can be justified by the documented benefits — e.g., in climbing unknown vertical substrates [5] or highly irregular level ground surfaces [26] (albeit, note these authors described the physical touchdown

¹Bringing the sensor signals across the unconstrained rotating legs necessitates some non-contact communications channel. We reported on a wireless scheme in [25] but this has proven very challenging to maintain in robust operating form. We have also experimented with infra-red and even slip-ring communications bridges between body and legs. All of these hardware approaches can be coaxed into functional operating form, but — as long argued in the “sensor-minimal robotics” literature [18, 19] — each incurs its own additional fragilities and operational complexity.




Fig. 2: Overview of fault detection, identification, and recovery architecture for a single state system. See [3]

sensors as not “adequate”). Even in these settings, running a model based observer will provide a good reference for accurately determining what is an expected disturbance and what is an unexpected disturbance. We believe that the broader virtues proposed in the industrial fault detection literature [3] for “analytical redundancy” will make the state-based software contact sensor we explore here a useful adjunct to such hardware solutions, but these considerations also lie beyond the scope of the present paper.

B. Disturbance Identification and Recovery

In concert with the overall framework of [3] (echoed in [21]), our detector’s residual signal is passed through a decision logic block for purposes of disturbance identification. In this paper the decision block takes the form of a hand-designed and hand-coded finite state machine depicted in Fig. 6. In the longer term, as the range of possible environments broadens and the diversity of potential fault sources increases, we suspect that automated methods of [27] will be required for the reliable and robust generation of such decision blocks, and, as mentioned above, it seems likely that a stochastic formulation may be required [23, 24].

The problem of fault recovery represents a vast, important domain in its own right that is still relatively unexplored in robotics. Bongard et al [16] compare the sensor-motor signatures of their robot with a physics simulation based upon generated self-models, for purpose of detecting the design configuration of the robot mechanism. While similar in purpose to our methods, we are focused upon models in which the implicit physics are simple, rather than making use of accurate full-body simulation, a difficult undertaking for dynamic legged robots. Other work on loss of limb [15] and reduced limb functionality [17] shows interesting gait strategies, but differs from our results as we make use of no sensory information other than actuators, and employ a robot with a minimal number of actuators for locomotion.

II. ALGORITHM

A key factor underlying the success of our observer-based sensor is that the dynamical properties of an EduBot leg in flight are extremely simple to model: it is essentially a one degree of freedom proportional-derivative reference tracking loop, decoupled from all the other degrees of freedom. In contrast, it is well understood that modeling contact is hard: characterizing a leg’s interaction with complex substrates lies

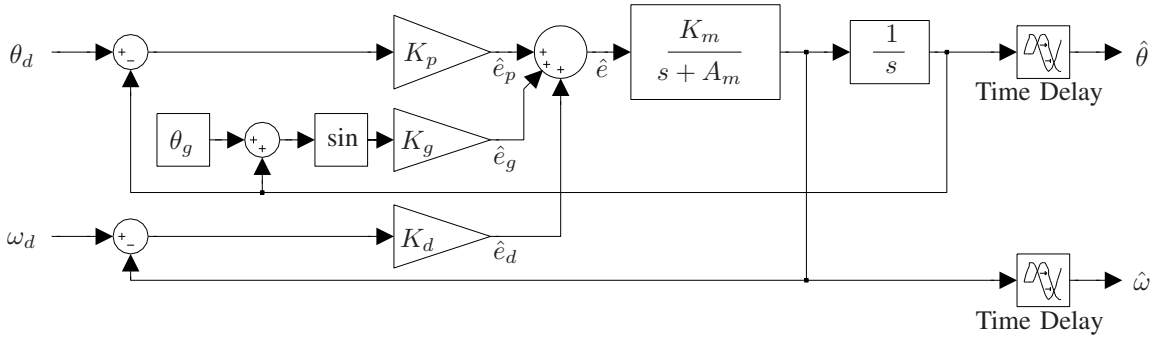


Fig. 3: Block diagram of leg observer. The estimated angle $\hat{\theta}$ and angular velocity $\hat{\omega}$ are generated from the same reference angle θ_d and velocity ω_d as the actual leg using a simple model of the motor and controller.

at the cutting edge of contemporary applied physics research [28], and, even on simple substrates, modeling the complex Lagrangian mechanics characterizing a robot’s joints while contacting a surface remains challenging. Thus, our problem formulation establishes a leg’s swing phase as its nominal operational state to be contrasted with “disturbances” caused by either ground or obstacle collision, which the subsequent fault logic can then readily classify as either expected, unexpected, or missing contacts.

A. Disturbance Detection

The leg flight model in RHex-style machines takes the particularly simple form of a single, decoupled, servo motor with estimated state (shaft position, $\hat{\theta}$, and velocity, $\hat{\omega}$) driven by a linear estimated error, \hat{e} ,

$$\hat{\omega} = \frac{K_m}{s + A_m} \hat{e} \quad (1)$$

$$\hat{\theta} = \frac{1}{s} \hat{\omega} \quad (2)$$

arising from a proportional-derivative (PD) tracker,²

$$\hat{e} = \hat{e}_p + \hat{e}_d \quad (3)$$

$$\hat{e}_p = K_p \cdot (\theta_d - \hat{\theta}) \quad (4)$$

$$\hat{e}_d = K_d \cdot (\omega_d - \hat{\omega}) \quad (5)$$

excited by the desired reference position, θ_d , and velocity, ω_d , signals issued from the “Buehler clock” that defines a RHex gait [2]. In addition, as our legs operate in the vertical plane, they are affected by gravity which we model (instead of trying to cancel actively) by adjoining a third known “reference error” term, \hat{e}_g , to the tracking controller’s input. Equation 3 is thus replaced by:

$$\hat{e} = \hat{e}_p + \hat{e}_d + \hat{e}_g \quad (6)$$

$$\hat{e}_g = K_g \cdot \sin(\hat{\theta} + \theta_g) \quad (7)$$

²Note that since θ_d and $\hat{\theta}$ lie on the circle (\mathbb{S}^1), their difference in Equation 4 is taken to be in the range $(-\pi, \pi]$, and computed by a standard modulus function. Also note that the physical interpretation of the parameters is standard and not essential to the paper’s central contribution. We discuss how to calibrate these parameters in the following paragraphs.

where K_g is the magnitude of the effect and θ_g is the angular offset. Finally, there is a time delay T_{del} at the end to synchronize the observer with the physical plant³. Figure 3 depicts the model just described.

The unknown parameters (K_m , A_m , K_g , θ_g , K_d , and T_{del}) are calibrated via the Nelder-Mead algorithm [29] using a hand-tuned starting simplex. The final parameter, K_p , is taken to be the value used by the higher level controller⁴. Furthermore, T_{del} is taken to be constant across all legs. For all experiments listed in this paper, these parameters were trained on a dataset collected when the robot was allowed to spin all legs freely and with a standard alternating tripod gait. The speed of the gait was ramped up over time from approximately 0.6 to 2.5 strides per second.

The outputs of this observer, $\hat{\theta}$ and $\hat{\omega}$, are compared with the actual achieved angle θ and angular velocity ω , as reported by the motor-mounted encoder, to form the observer residual vector⁵:

$$\begin{bmatrix} r_\theta \\ r_\omega \end{bmatrix} = \begin{bmatrix} \hat{\theta} - \theta \\ \hat{\omega} - \omega \end{bmatrix} \quad (8)$$

To test the accuracy of the observer, we collected a second dataset at a moderate speed of one stride per second for 15 seconds. This yielded a median position residual of $r_\theta = 0.0271$ radians (1.59 degrees) and a median velocity residual of $r_\omega = 0.4267$ rad/s (4.075 rpm) over all six legs. The maximum residuals were $r_\theta = 0.1038$ radians (5.95 degrees) and $r_\omega = 3.8857$ (37.11 rpm). A section of this raw data from the first leg is shown in Figure 4 for the robot in the air and Figure 5 for the robot making ground contact. In both figures parts (a) and (c) we plot r_θ and r_ω , respectively. For comparison, errors calculated from the position and velocity tracking ($\theta_d - \theta$ and $\omega_d - \omega$) are plotted in parts (b) and (d), respectively. The green shaded portion indicates the nominal stance phase of the gait.

³The EduBot’s distributed control architecture and bus structure incurs a time delay from motherboard (where θ_d and ω_d are generated) through the network to the local hip controller and then back up to the motherboard where the residuals are calculated. Since our system is time invariant, we can combine these delays into an overall delay of twice the average one way network transport time. Our observer outputs are thus held in a buffer for a total of T_{del} .

⁴Due to the implementation of the derivative feedback in our controller, K_d had to be calculated

⁵As with Equation 4, the value $\hat{\theta} - \theta$ is taken to lie in the range $(-\pi, \pi]$.

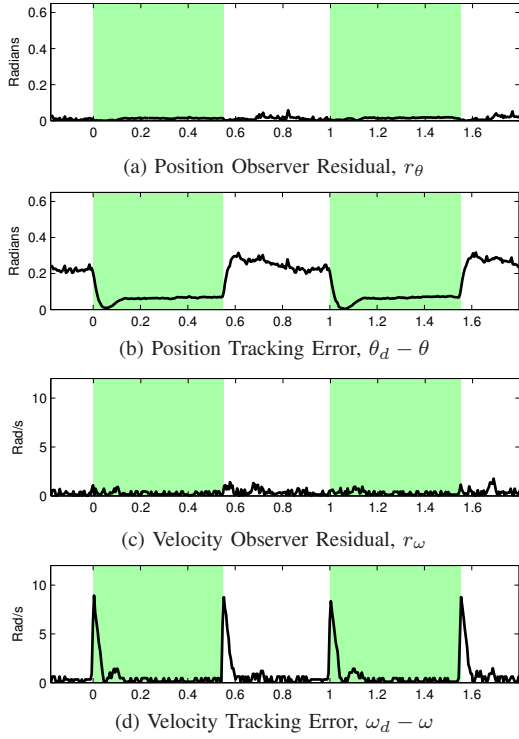


Fig. 4: Observer residual contrasted with controller tracking error in both position and velocity under conditions of free leg swing (no “disturbance” from any ground contact). The abscissas display time in seconds. Green shading indicates the expected stance phase of the gait.

These plots suggest the significantly greater utility of observer residuals relative to mere controller tracking errors in assessing a leg’s relationship to the ground. Whereas the simple tracking errors exhibit sizable and varying excursions even when the leg has no load, the observer residuals account for the predictable causes of such variation, and only exhibit excursions when contact conditions change. More specifically, during normal operation, due to the nature of the proportional-derivative controller, velocity tracking cannot account for abrupt changes in reference velocity (which the motor cannot perfectly follow), and position tracking must lag as a function of the commanded and actual speeds. These structural features of the PD error signals are particularly onerous because they are strongest just at the moments of the putative ground interaction of true interest. The change in gait phase to slow the leg down for stance by definition should happen around the same time as the touchdown event we are trying to detect. In contrast, these expected dynamical variations in the normal tracking error are accounted for in the observer, as is evidenced by the low level of error for both r_θ and r_ω in Figure 4a and 4c. While both estimated states, $\hat{\theta}$ and $\hat{\omega}$, provide useful information, we have found that using only θ is sufficient for disturbance identification.

B. Disturbance Identification

Given an informative disturbance signal, r_θ , we introduce a simple output logic stage to classify the conditions of

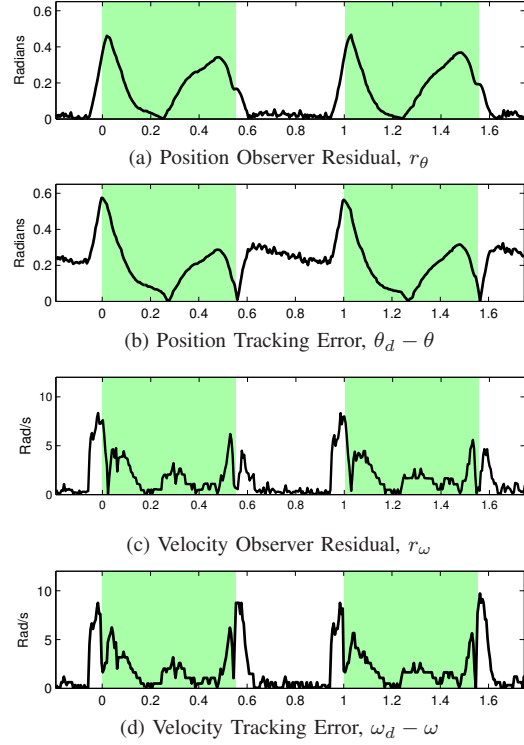


Fig. 5: Observer residual contrasted with controller tracking error in both position and velocity under conditions of cyclic ground contact (one of the disturbances of interest). The abscissas display time in seconds. Green shading indicates the expected stance phase of the gait.

interest with respect to an intuitively developed partition of the signal space as follows. The circle, \mathbb{S}^1 , of leg phase angles is partitioned into four intervals labeled “ground” (G) — leg angles that the Buehler clock associates with ground contact by commanding lower ω_d — and “air” (A) — leg angles that the clock associates with free flight by commanding higher ω_d — together with two intermediating phase angles labeled “takeoff” (T) — an interval over which the transition from low to high ω_d is expected to occur — and “landing” (L) — an interval over which the transition from high to low ω_d is expected to occur. Similarly, the circle, \mathbb{S}^1 , of residual position angle errors is partitioned into three intervals labeled “high” (H) — large residual values that experience suggests should be expected only in conjunction with stance — “low” (L) — small residuals associated with typical free flight conditions — and “medium” (M) — a pair of disconnected intervals that separate the “low” and “high” intervals. We use these symbols to trigger the transition of a simple hand-designed FSM with four normal states — stance, possible takeoff, flight, or possible landing. The FSM includes two additional error states — unexpected disturbance and missing ground. An unexpected disturbance occurs when r_θ increases but the leg is not in a phase of the gait where it could hit the ground. Missing ground is when r_θ does not increase but the leg is in a phase of the gait where it should have contacted the ground. There is also a

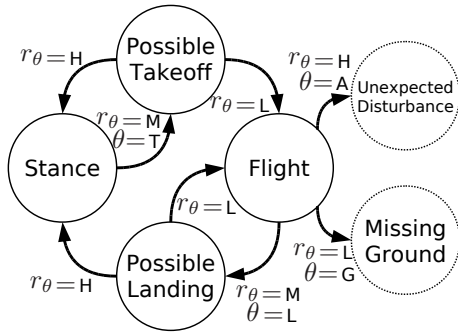


Fig. 6: State transition diagram for fault identification

minimum lingering time in each state of the FSM to avoid quick transitions due to noise spikes. The state transition diagram is depicted in Figure 6.

The possible takeoff and possible landing states were added to improve accuracy over a wider range of gait speeds. For instance the leg may be in a state with $\theta \in T$ and $r_\theta \in M$ but it should not be construed as having taken off unless at some point in the near future r_θ continues to decrease. In contrast, if r_θ goes back up then we should treat the leg as if it is still on the ground. The possible landing state allows our ground detector to trigger at the medium error level but not declare a confirmed ground contact unless r_θ continues to rise up to the high level.

As mentioned in I-A, the ad hoc construction of this diagnostic state machine could surely be improved by recourse to the more formal methods of supervisory control [27]. However, for this application, we believe the details of its cell structure and transition logic are less important than the broader design insight that a sound decision about the nature of the current disturbance must be based on the available information: θ , ω , r_θ , r_ω , and their evolution in time.

C. Disturbance Recovery

Reacting to and recovering from disturbances or damage to limbs during locomotion should be a strength of multi-legged platforms, given the intuitive understanding that a multitude of legs confers redundancy. Re-coordination of the limbs of a within-stride walking machine while attempting to guarantee gait stability, however, is not entirely straightforward. Particularly with RHex-style machines such as EduBot, equipped with a single actuator per leg, the coordination of phase during active locomotion while avoiding a fall can be challenging. In planning and executing the gait recovery mechanism that we report here, we build upon prior work in topological gait classification, analysis, and control [30]. For this application, we signal a gait recovery transition for the scenario where the robot accidentally breaks a leg, resulting in an asymmetric five-legged gait for a hexapedal machine.

Our gait classification formalism [30] adapts Young Tabloids [31], representations of the partitions of a set of unique integers, $1 \dots n$, to represent the unique gait timings of multi-legged systems. For an individual tabloid, elements within a row correspond to legs in phase, while consecutive rows dictate the out of phase cyclic recirculation order of

legs. An example is the alternating tripod, a commonplace gait for hexapedal systems in which two sets of three legs are grouped into individual tripods of support. Represented as a Young Tabloid, this gait is:

$$\begin{array}{|c|c|c|} \hline 1 & 3 & 5 \\ \hline 2 & 4 & 6 \\ \hline \end{array} \quad (9)$$

Our previous work in this problem domain associates Young Tabloids with gait cycles on a high dimensional torus — the phase space of multi-legged gaits — and uses a simple control policy to place global attractors at desired gait cycles [30]. While that prior work deals with the use of tabloids in correspondence with an algebra over which to plan gait transitions, we are now faced with the scenario of changing dimensionality, by which losing a leg forces us to consider gaits on a 5-dimensional torus rather than six.

A total of 1082 unique gait orderings exist for a hexapedal robot [30], of which the alternating tripod is both uniquely fast and stable. Upon loss of a leg, however, we have a five-legged system, for which only 150 unique timings exist, and the now “5-legged” tripod gait no longer maintains stability. Depending upon which leg has been lost, we transition to a crawl gait that maximizes static stability for the walking machine⁶. In the case of loss of the robot’s sixth leg, the back right actuator, we transition to a crawl gait with leg recirculation ordering,

$$\begin{array}{|c|} \hline 3 \\ \hline 1 \\ \hline 5 \\ \hline 2 \\ \hline 4 \\ \hline \end{array} \quad (10)$$

The transition is executed when the FSM indicates missing ground by a given leg. As EduBot and RHex only contain a single actuator per leg, we speed up and slow down legs in recirculation to achieve the phase relationships of our new desired gait [30]. Our control policies change the gait immediately, and converge fully to the desired crawl gait within a few strides, as witnessed by an example of this transition shown in the following section.

III. REACTIVE BEHAVIORS

A. Wall Avoidance

The first behavior that we have implemented using this software contact detector and disturbance classifier is a simple wall avoidance algorithm. Instead of whiskers or antennae [32], the robot must touch the wall with its leg and, upon unexpected disturbance (see Section II-B), back up to turn away. While not necessarily an efficient solution to any sort of maze problem, this simple, useful behavior illustrates the reliability of our contact detector for disturbance detection and identification.

In this case disturbance recovery is quite simple. Once the robot knows there is an obstacle in front of it, it must immediately move backwards, turn, and continue on its

⁶Our current recovery strategy uses specifically chosen crawl gaits. Future versions of this work will take into account the stability margins of all 150 potential gaits in order to automatically select suitable gaits.




Fig. 7: Overhead plot, in meters, of the center of mass of an EduBot running a wall avoidance behavior inside a closed rectangular region, using the described method as the only sensing strategy. The twelve contact points are labeled.

way. For simplicity of solution, here the robot always turns right. An overhead plot of this behavior within an enclosed rectangle seen in Fig. 7, recorded using a Vicon motion capture system [33]. A second experiment with an available exit is shown in the attached video.

B. Leg Fault Recovery


To test the effectiveness of our detection of and recovery from disturbances, we now exhibit behavioral strategies on the EduBot machine in which the robot must quickly adapt to dramatically changed locomotion capabilities.

The robot was rigged with purposefully weakened legs designed to fracture and fall off after only a few steps. Without disturbance detection, the now five-legged machine naïvely continues to make use of the alternating tripod gait, but does so with dramatically reduced stability. As shown in Fig. 8a, the machine impacts the ground (chattering in vertical acceleration) during each stride due to loss of stability, greatly decreasing its locomotive quality.


Upon execution of a transition to a five-legged crawl gait triggered by a missing ground event using strategies outlined in Sections II-B and II-C, the robot is still handicapped with just five legs, seen in Fig. 8b at steady-state behavior, but maintains stability and does not impact the ground. A plot of a transition is shown in Fig. 9, in which a gait transition smoothly switches gaits upon noticing a locomotive failure of the sixth leg. An example of both the reactive and non-reactive cases is included in the attached video.

IV. CONCLUSION

We have introduced and documented empirically the performance of a software contact-event driven disturbance identification and recovery system, based upon the established fault detection and isolation methods of industrial control [3], operating on a hexapedal walking machine. Initial results demonstrate that our ground contact estimates successfully cue appropriate behavioral transitions, including



(a) Alternating Tripod Gait with 5 Legs



(b) Five-legged Crawl Gait

Fig. 8: Gait comparison using inertial measurements of a robot walking with a missing sixth leg. For the bottom plot of each, black regions indicate stance. On five legs, the tripod gait loses stability and impacts the ground, seen by strong impacts in the vertical acceleration. The crawl gaits retains stability, reduces roll, and induces only small pitching moments.

effective reaction to the sudden and unexpected loss of a limb during locomotion followed by smooth, safe transition to a new, more stable gait.

Of course, a great variety of errors may be encountered during robot locomotion, ranging from the most simple, such as unexpected obstacles in a robot's path, to more complex self-failures due to a variety of causes (potentially including, e.g., motors overheating, electrical shorts, seized gearboxes, broken legs, etc.) and it is important that a legged robot respond to and recover from each disturbance appropriately. While readily available and effective sensors exist for a variety of applications, our method relies solely upon a robot's actuators and offers robust, model-based results for both minimalism and redundancy during operation.

For future studies, we are interested in incorporating additional actuator information, such as measured motor current, into our model, and we also believe that use of r_ω , in addition to r_θ , can provide further information regarding system state. The detector will benefit from additional learning techniques, both on- and off-line, to remove hand-tuned elements as well as fit performance characteristics during operation. Whereas our current methods only estimate discrete leg contact, we




Fig. 9: Sequence of leg recirculation events for each leg during a transition from a hexapedal alternating tripod gait to a five-legged crawl, black regions corresponding to leg stance. Around $t = 30.5$ the robot loses its sixth leg and compensates via the gait transition.

are looking into estimation for the magnitudes of ground reaction forces as well [34]. In the longer term, we plan to adapt our estimation techniques to terrain classification, incorporating, for example, IMU data and frequency analysis as in [35]. This will allow the robot to adapt gaits based upon terrain, such as when suddenly encountering a sandy rather than than expected hard surface [28]. Finally, this paper has focused on robots with single actuators per leg, but we are now contemplating the introduction of such simple diagnostics on other legged robots for which additional actuator feedback may be useful for more robust behaviors.

ACKNOWLEDGMENTS

This was supported in part by the ARL/GDRS RCTA and NSF FIBR projects as well as an Intelligence Community Postdoctoral Fellowship held by the second author and funding from the DARPA STOMP project for the last author.

REFERENCES

- [1] J. D. Weingarten, D. E. Koditschek, H. Komsuoglu, and C. Massey, "Robotics as the delivery vehicle: A contextualized, social, self paced, engineering education for life-long learners," in *Robotics Science and Systems Workshop on "Research in Robots for Education"*, 2007.
- [2] U. Saranli, M. Buehler, and D. E. Koditschek, "RHEx: A Simple and Highly Mobile Hexapod Robot," *The International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.
- [3] P. M. Frank, "Fault Diagnosis in Dynamic Systems Using Analytical and Knowledge-based Redundancy—A Survey and Some New Results," *Automatica*, vol. 26, no. 3, pp. 459–474, 1990.
- [4] J. D. Weingarten, G. A. D. Lopes, M. Buehler, R. E. Groff, and D. E. Koditschek, "Automated gait adaptation for legged robots," in *Proceedings of the International Conference on Robotics and Automation*, vol. 3, April 2004, pp. 2153–2158.
- [5] M. A. Spenko, G. C. Haynes, A. Saunders, A. A. Rizzi, M. Cutkosky, R. J. Full, and D. E. Koditschek, "Biologically inspired climbing with a hexapedal robot," *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 223–242, 2008.
- [6] S. Chernova and M. Veloso, "An evolutionary approach to gait learning for four-legged robots," in *In Proceedings of IROS'04*, September 2004.
- [7] V. Zykov, J. Bongard, and H. Lipson, "Evolving dynamic gaits on a physical robot," in *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, M. Keijzer, Ed., Seattle, Washington, USA, 26 Jul. 2004.
- [8] N. Kohl and P. Stone, "Machine learning for fast quadrupedal locomotion," in *In Proceedings of the National Conference on Artificial Intelligence*, July 2004.
- [9] K. Hauser, T. Bretl, J. C. Latombe, and B. Wilcox, "Motion planning for a six-legged lunar robot," in *Workshop on the Algorithmic Foundations of Robotics*, New York, NY, 2006.
- [10] P. Vernaza, M. Likhachev, S. Bhattacharya, S. Chitta, A. Kushleyev, and D. D. Lee, "Search-based planning for a legged robot over rough terrain," in *Proceedings of the International Conference on Robotics and Automation*, 2009.
- [11] D. Pongas, M. Mistry, and S. Schaal, "A robust quadruped walking gait for traversing rough terrain," in *Proceedings of the International Conference on Robotics and Automation*, 2007.
- [12] K. Byl and R. Tedrake, "Dynamically diverse legged locomotion for rough terrain," in *Proceedings of the International Conference on Robotics and Automation*, video submission, May 2009.
- [13] J. D. Weingarten, R. E. Groff, and D. E. Koditschek, "A framework for the coordination of legged robot gaits," in *Proceedings of the IEEE Conference on Robotics, Automation and Mechatronics*, vol. 2, 2004, pp. 679–686.
- [14] K. Byl and R. Tedrake, "Metastable walking machines," *International Journal of Robotics Research*, 2009.
- [15] C. Ferrell, "Failure recognition and fault tolerance of an autonomous robot," *Adaptive Behavior*, vol. 2, no. 4, pp. 375–398, 1994.
- [16] J. Bongard, V. Zykov, and H. Lipson, "Resilient Machines Through continuous Self-Modeling," *Science*, vol. 314, pp. 1118–1121, 2006.
- [17] J.-M. Yang, "Two-phase discontinuous gaits for quadruped walking machines with a failed leg," *Robotics and Autonomous Systems*, vol. 56, no. 9, pp. 728 – 737, 2008.
- [18] M. T. Mason, "Kicking the sensing habit," *AI Magazine*, vol. 14, no. 1, pp. 58–59, 1993.
- [19] J. F. Canny and K. V. Goldberg, "A RISC approach to sensing and manipulation," *Journal of Robotic Systems*, vol. 12, no. 6, 1995.
- [20] H. Komsuoglu, D. McMordie, U. Saranli, N. Moore, M. Buehler, and D. E. Koditschek, "Proprioception based behavioral advances in a hexapod robot," in *Proceedings of International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [21] O. Pettersson, "Execution monitoring in robotics: A survey," *Robotics and Autonomous Systems*, vol. 53, pp. 73–88, 2005.
- [22] S. Mcilraith, G. Biswas, D. Clancy, and V. Gupta, "Hybrid systems diagnosis," in *In Proceedings of Hybrid Systems: Computation and Control, Lecture Notes in Computer Science*. Springer, 2000, pp. 282–295.
- [23] U. Lerner and B. Moses, "Monitoring a complex physical system using a hybrid dynamic bayes net," in *In UAI*. UAI, 2002.
- [24] X. Koutsoukos, J. Kurien, and F. Zhao, "Estimation of distributed hybrid systems using particle filtering methods," in *In Hybrid Systems: Computation and Control (HSCC 2003)*. Springer Verlag Lecture Notes on Computer Science. Springer, 2003, pp. 298–313.
- [25] P.-C. Lin, H. Komsuoglu, and D. E. Koditschek, "A leg configuration sensory system for dynamical body state estimates in a hexapod robot," in *Proceedings of International Conference on Robotics and Automation*, Taipei, Taiwan, September 2003, pp. 1391–1396.
- [26] S. Chitta, P. Vernaza, R. Geykhman, and D. D. Lee, "Proprioceptive localization for a quadrupedal robot on known terrain," in *Proceedings of the International Conference on Robotics and Automation*, 2007.
- [27] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. C. Teneketzis, "Failure diagnosis using discrete-event models," *Transactions on Control Systems Technology*, vol. 4, no. 2, March 1996.
- [28] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, "Sensitive dependence of the motion of a legged robot on granular media," *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 3029–3034, 2009. [Online]. Available: <http://www.pnas.org/content/106/9/3029.abstract>
- [29] J. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.
- [30] G. C. Haynes, F. R. Cohen, and D. E. Koditschek, "Gait transitions for quasi-static hexapedal locomotion on level ground," in *International Symposium of Robotics Research*, Lucerne, Switzerland, August 2009.
- [31] B. Sagan, *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*. New York: Springer-Verlag, 2001.
- [32] N. J. Cowan, E. J. Ma, M. Cutkosky, and R. J. Full, "A biologically inspired passive antenna for steering control of a running robot," in *In International Symposium on Robotics Research*. Springer, 2003, pp. 541–550.
- [33] Vicon Motion Systems, <http://www.vicon.com/>.
- [34] C.-S. Liu and H. Peng, "Inverse-Dynamics Based State and Disturbance Observers for Linear Time-Invariant Systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 124, pp. 375–381, 2002.
- [35] E. M. DuPont, C. A. Moore, E. G. Collins, and E. Coyle, "Frequency response method for terrain classification in autonomous ground vehicles," *Autonomous Robotics*, vol. 24, pp. 337–347, 2008.