On Monitoring the Health and Remaining Useful Life of Vehicle Suspension Systems

Guy E. Gallasch and Nicholas Brealey

Land Logistics, Land Division, Defence Science and Technology Group,
PO Box 1500, Edinburgh, South Australia, 5111, Australia

Abstract

Vehicle suspension plays a key role in the safety of land vehicles and their occupants. Shock absorbers are a key component of suspension systems, and failed shock absorbers can result in substantial collateral damage to suspension and steering systems if not corrected. Conventional shock absorbers dissipate vibration energy as heat, thus external temperature may provide an indication of shock absorber health. Other approaches for monitoring shock absorber health include measuring the relative motion of the attached sprung and unspring masses, or the use of embedded sensors for internal fluid temperature and pressure. We conjecture that the life of a shock absorber is a function of both the cumulative energy dissipated over its lifetime, and the intensity of individual shock events. This paper presents the results of a small study investigating the feasibility of monitoring the health of shock absorbers by retrofitting low-cost accelerometers, with the potential to support both diagnostics and prognostics. We do this through a seeded fault testing regime on a scale model test-rig. Ultimately, a programme of data collection from operational vehicles is proposed as a means of instantiating accurate life consumption models and validating our conjecture.

Keywords: Health and Usage Monitoring Systems, Vehicle Suspension Systems, Shock Absorbers, Remaining Useful Life, Prototyping, Proof of Concept

Introduction

Vehicle suspension systems are critical for vehicle handling and braking, and hence for occupant safety (and comfort). A key component of vehicle suspension is the shock absorber: a device for damping vibration. Our interest in vehicle shock absorbers stems from a general interest in vehicle Health and Usage Monitoring Systems (HUMS) and vehicle maintenance. Specifically, our interest in monitoring the health of shock absorbers has been motivated by:

- Vehicle safety: shock absorbers play a critical role in maintaining tyre contact with the road surface, hence failed or failing shock absorbers pose a significant safety risk [1].
- Anecdotal experiences from military operations in the Middle East: the use of hand-held contactless infrared (IR) thermometers to diagnose faulty shock absorbers so that they can be changed out before they induce cascading failure of other suspension and steering systems that are more expensive and time-consuming to replace.
- Conversations on vehicle HUMS priorities with Defence's Capability and Sustainment Group (CASG), including monitoring shock absorber health.
- Our group's ongoing analysis of historical maintenance records from Army's Military Integrated Logistic Information System (MILIS) suggests that suspension failure features prominently across multiple vehicle types and fleets.

Given the above, this paper documents our preliminary investigation into mechanisms for monitoring the health of vehicle suspension in general, but shock absorbers in particular. Our ultimate aim is to demonstrate how it might be possible to provide both a diagnostic capability (low cost, simple, effective detection of failing or failed shock absorbers) and a prognostic capability (the ability to pre-emptively replace shock absorbers before they fail, but in a way that minimises wasted remaining useful life, either through direct or inferred measurements of shock absorber health). Further, we wish this to be achieved with a minimum of additional instrumentation or retrofitting of new parts to existing/legacy vehicles, and a minimum of ongoing manual effort.

To this end, following a survey of monitoring techniques, we instigated a proof-of-concept study to explore the feasibility of such solutions. Intentionally designed to be inexpensive and to provide a quick turn-around, the aim of our study was not to produce final products suitable for field trials, nor was it to undertake a rigorous analytic investigation of the behaviour of suspension systems (as has been done many times before, e.g. [1]). Rather, it was to determine whether further (more expansive and hence expensive) investigation is justified.

Survey of Techniques for shock absorber HUMS

To begin, a brief literature survey was conducted into techniques for monitoring the health, usage and failure of shock absorbers. We consider a shock absorber to have failed when it no longer provides its intended damping function. This can be detected through reduced damping response, reduced operating temperature (as intimated through anecdotal operational experience), or through more direct measurements such as internal fluid volume and pressure. A (non-definitive) summary of key identified techniques from the literature and our own brainstorming activity, and an assessment of the pros and cons of each, is provided in Table 1. Note that not all of the cited material relates to detection of failure, e.g. [2] relates to assessing the potential for energy harvesting.

Proof-Of-Concept Study

Based on the literature survey, our assessment of the balance of pros and cons, and our stated aims, we identified the use of accelerometers and load cells as a technique worth investigating. Specifically, they are smaller and hence easier to retrofit and less susceptible to physical damage than linear potentiometers, while avoiding the potential for environmental distortion inherent in temperature measurements. Further, they do not require the replacement of existing (still functional) shock absorbers, hence avoid wasting remaining useful life. Finally, once fitted, they can monitor condition automatically with minimal human intervention.

Our choice of test shock absorber was dictated by organisational restraints (timeliness in construction and Work Health and Safety considerations) and a desire to minimise overall expense. We chose the smallest authentic shock absorber we could obtain readily, designed for a 1/5 scale remote controlled off-road vehicle [3]. A small shock absorber, we reasoned, would allow us to easily experiment with a light weight test-rig, without the need for large electric, hydraulic, or pneumatic actuation and associated organisational overhead.

We then designed and built two test-rigs, each using a different mechanism to actuate the shock absorber. The first (Fig. 1 (a)) used a large solenoid as the actuator, with the idea that a solenoid would allow some variability in actuation frequency and amplitude, and hence be able to simulate a wider range of road surface profiles than a fixed cam or crank.

Table 1: Summary of brainstorming and survey of approaches for monitoring shock absorber health

Method	Pros	Cons
No HUMS (as-is)	• System operates 'as is' with no remedial cost or effort	• Potential for significant collateral damage in event of a failure
Manual physical testing		• Costly (time, effort) to remove and replace
Temperature:		General for temperature: • No scope for analysis of health of springs • Susceptible to variation induced by environment (e.g. puddles)
• Hand-held contactless thermometer	• Simple, cheap, reasonably effective	• 'Handraulic' (manual effort) • No digital record - hard to trend over time
• External (attached) thermometer	• Can produce a digital record, hence trends over time	• Susceptible to external damage • Requires replacement of existing (still functional) shock absorbers
• Internal or embedded thermometer	• Can produce a digital record, hence trends over time	• Requires replacement of existing (still functional) shock absorbers
• Temperature and internal pressure [4]	• Can produce a digital record, hence trends over time	• Requires replacement of existing (still functional) shock absorbers
Dynamic response:		
• Dynamic displacement measurement via attached measurement device (e.g. linear potentiometer [2])	• Direct measurement of displacement, velocity and acceleration can be derived • Avoids the issue of drift present with accelerometers	• Potential physical incompatibility - may be difficult to retrofit to existing suspension systems (springs etc. in the way) • Susceptible to physical damage
• External accelerometer attached to the top of the shock absorber mount [5][1]	• Simpler, less sensors to be retrofitted and from which data is collected	• May not be possible in real-time • Relies on 'normal' datasets known a priori
• External accelerometers attached "top and bottom"	• Smaller, easier to retrofit, and less susceptible to physical damage • Does not require replacement of existing shock absorbers	• Drift becomes a factor that must be mitigated when deriving velocity and displacement
• Internal embedded accelerometers	• Less susceptible to physical damage	• Requires replacement of existing (still functional) shock absorbers
• Load cells	• Allows direct measurement of force exerted on damper, rather than deriving this from velocity and damping coefficient	• Velocity and displacement calculations require knowledge of mass of sprung and unsprung parts

[1] Note that [5] is based on simulated data only, and focuses on spring faults rather than shock absorbers.

However, the use of a solenoid was problematic: to gain the desired pull strength and throw length required the largest readily available solenoid that we could obtain, which in turn required a large electric current to drive, resulting in a very short duty cycle. The solenoid was later replaced with a hand-actuated crankshaft (not pictured).

Hence, the second test-rig (Fig. 1 (b)) utilised a more conventional electric motor and cam. The cam provided regular, fixed amplitude actuation of the shock absorber, although the frequency could be varied by changing the speed of the electric motor. The gearing on the motor restricted us to relatively low speeds.

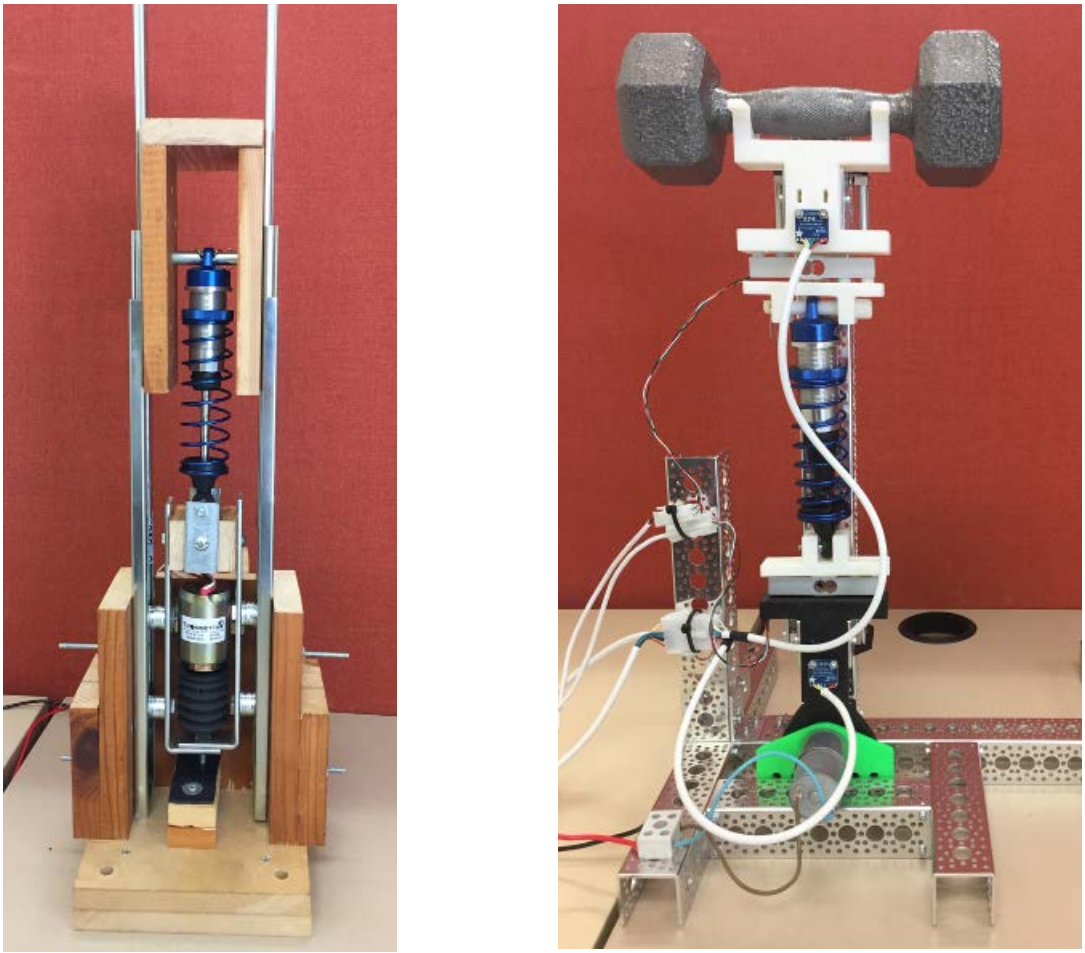

(a) (b)

Fig. 1: 1/5 scale shock absorber test-rigs

Our test-rigs incorporated load cells and accelerometers, with one pair at each of the top and bottom mounting points of the shock absorber. In this fashion we could measure the accelerations and actual forces acting on each end of the shock absorber. To collect this data we used an Arduino-based logger that connected to the four sensors. In keeping with our desired intent, the load cells and accelerometers were inexpensive 'hobbyist-grade' models and thus with limited responsiveness and accuracy.

When operating the test-rigs, we record accelerometer readings for a period of a few seconds before we begin actuation, akin to the period after the driver enters the vehicle but before driving commences, to calibrate each sensor's reading of gravity prior to cumulative numerical integration to get velocity. From this point, we retrospectively compute and subtract a moving average of velocity over a fixed window, centrally located around each

computed point, to alleviate drift. We can do this because: a) as suggested in [6] we know that the accelerometer displacement is highly physically constrained, i.e. displacement cannot fall outside of known physical bounds, and b) that the accelerations experienced by each sensor are not 'lost in the drift'. A window size of around 3-5 wavelengths of the primary oscillation worked well as a trade-off between removing the general trend of drift and attenuating the oscillation itself.

The load cells, due to their low sampling rate (approx. 10Hz) did not contribute substantially to the overall results obtained during our experiments.

Results and Discussion

Despite the use of inexpensive 'hobbyist-grade' accelerometers, we were able to detect changes in the dynamic response of the shock absorbers as the fluid level within the shock absorber was progressively reduced, hence reducing its ability to provide a damping effect. The graphs in Fig. 2 show representative time-slices of simple measures of (drift-mitigated) dynamic response. Fig. 2 (a) shows the velocities of the lower (in red) and upper (in blue) mount points of the shock absorber. Because the input oscillation at the lower mount point is between 1.2 and 1.5 Hz, well below the natural undamped frequency of our sprung mass (approx. 6Hz), we see the input oscillation at the lower mount point being amplified (rather than reduced) at the upper mount point. The amplification increases as fluid levels (and hence damping ability) decreases. Fig. 2 (b) shows relative displacement between the two mount points of the shock absorber (noting that the vertical axes have the same scale, but different ranges, due to the compounding of accelerometer drift when integrating a second time to get displacement, and our drift mitigation technique not being perfect). Again, the change in dynamic response (in this case the increasing amplitude of the 'ripple' in displacement) is evident as the fluid level decreases.

In both cases, there is little change in observed response between the full shock absorber and the case when 3 ml and 6 ml of fluid had been removed. The change becomes noticeable when 9 ml, or just over half, of the fluid is removed, and becomes increasingly prominent as fluid levels move toward zero. Visual inspection confirmed that the piston within the shock absorber begins to project beyond the fluid during its travel when around 9 ml of fluid is removed, hence corroborating the observed results from accelerometer data.

Although we've investigated only one specific frequency of oscillation, we suggest that 'norms' for dynamic response can be established for specific vehicles (an ideal application for machine learning) and deviations from expected responses could be readily detected. Implementing such a mechanism is beyond the scope of this study. Further, although the shock absorber we used was a simple oil-filled monotube design, it encapsulates the common operating principles of all shock absorbers, and hence we believe this approach can be extended to shock absorbers with more elaborate designs.

Neither test-rig was sufficiently robust for us to achieve one of our original goals; to actuate the suspension vigorously enough and for long enough to see significant changes in shock absorber temperature. Hence, calculating energy dissipated by the shock absorber from empirical data was also problematic due to a) the changing damping coefficient as fluid was removed, and b) the small amplitude of compression and extension of the shock absorber at the low frequency of oscillation supported by the test-rig (approx. 1.5 Hz).

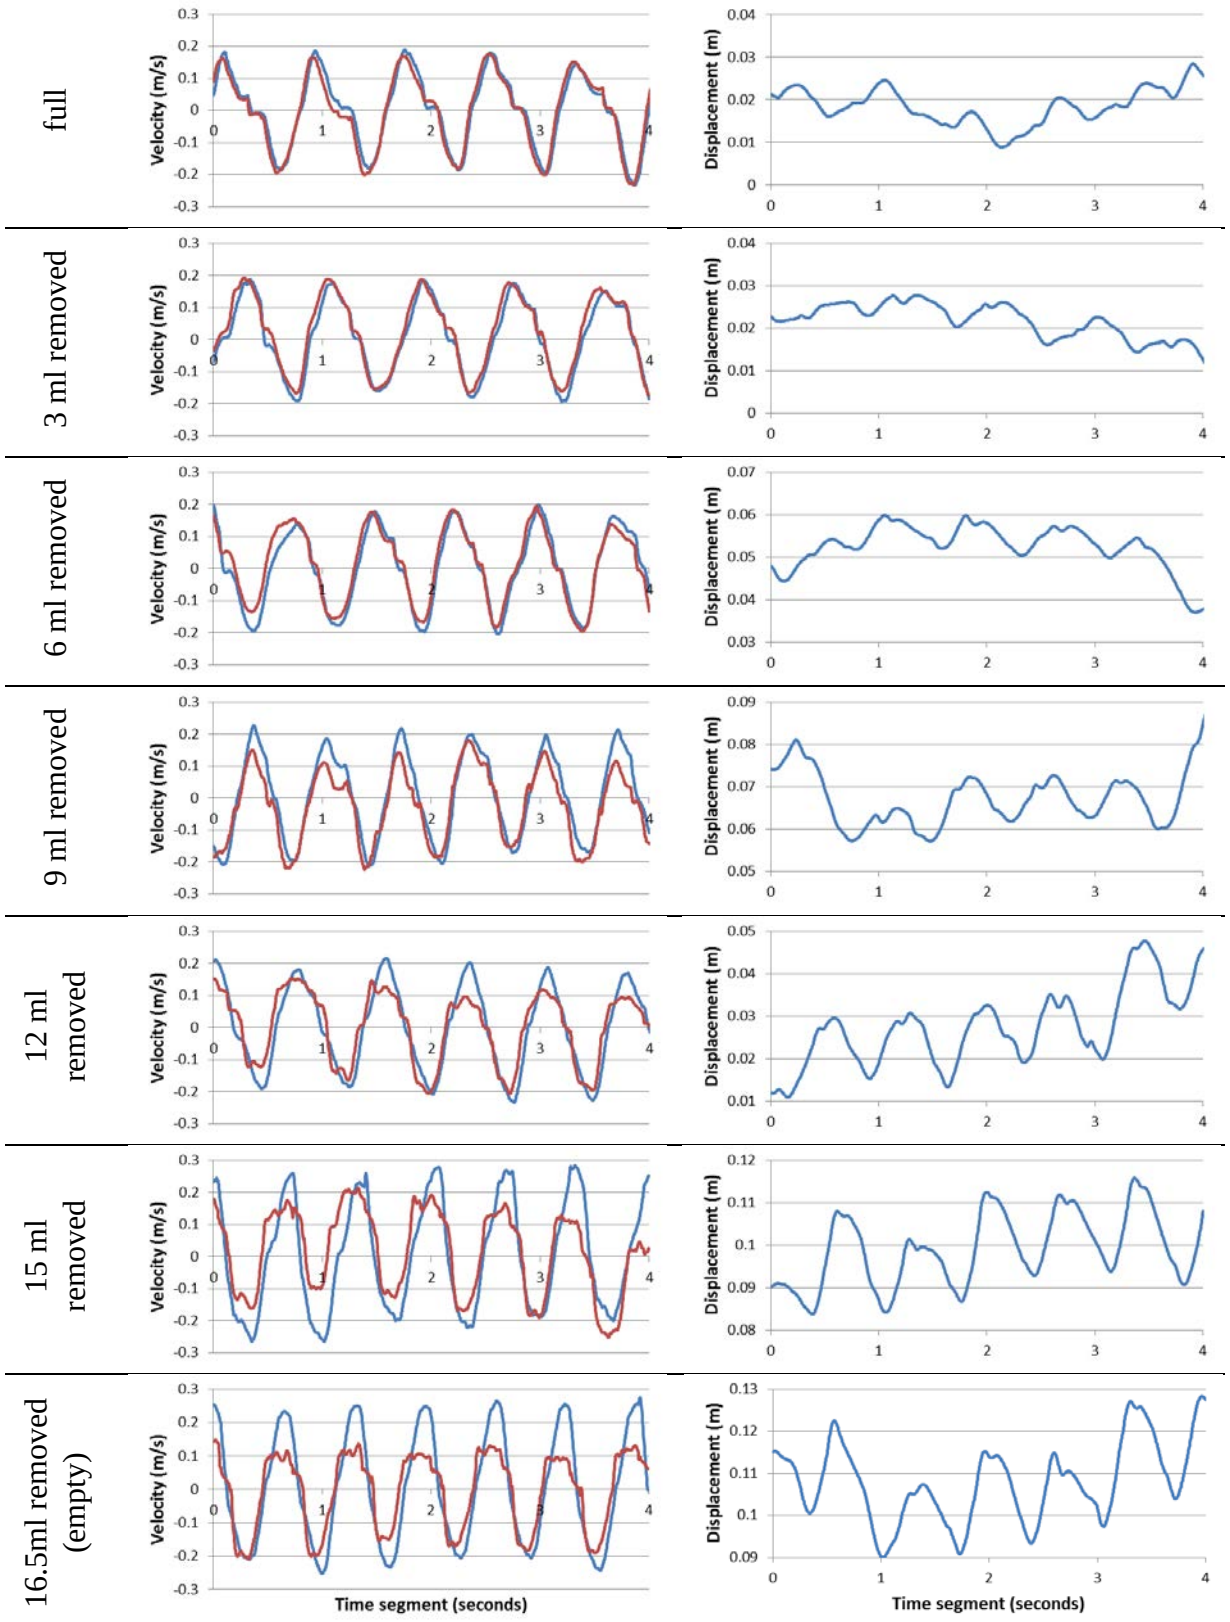

(a) (b)

Fig. 2: Shock absorber dynamic response as internal fluid is progressively removed.

Conclusions

Our goal was to quickly and inexpensively test a proof-of-concept where we used externally fitted sensors to detect shock absorber failure. We have shown that this is indeed possible. However, our study was not entirely successful. Through our desire to save time and expense in the design and fabrication of a test-rig, we did sacrifice accuracy, precision, and robustness. The balance between these objectives is something we would revisit. Due to the relative fragility of our test-rigs we were unable to stress the shock absorber as much as we planned. Further, while our approach shows promise for failure detection, we have done little work on developing a prognostic capability.

We believe it is possible to extend our analysis to determine the health of springs. Using the position information from the sensors we are able to determine the cumulative energy stored and released by the spring. Although we have made no explicit investigation into this matter, our data appears to support this approach, at least as an approximation. Ideally, we would like to extend our investigation to include data collection from in-service vehicles (through suitably ruggedized hardware) where we could observe examples of suspension failure and thus develop life consumption models and validate (or not) our conjecture.

Finally, having satisfied our goal (at least to some extent), we must acknowledge that as continuing advancements in suspension systems (e.g. semi-active magneto-rhetorical dampers [7]) make their way into military vehicles, and the ongoing trend for digitisation in vehicles more generally, embedded sensors in shock absorbers will soon become not only the norm, but will be essential for correct operation of the suspension system. Moreover, such sensors will not be limited to shock absorbers, but will be embedded on systems throughout the vehicle. A multitude of sensors will enable techniques, such as data fusion, to provide a better picture of the vehicle health, and to provide enhanced diagnostic and prognostic capabilities.

References

1. Wong, J. Y., *Theory of Ground Vehicles*, 4th edition, John Wiley and Sons, Inc., Hoboken, New Jersey, 2008.
2. Gill, C., Knight, C. and McGarry, S., “Measured Power Dissipation of Shock Absorbers on Light and Heavy Commercial Vehicles”, *SAE International Journal of Commercial Vehicles*, Vol. 7, No. 2, October 2014, pp. 718-725.
3. Horizon Hobby, “5IVE-T 1/5-Scale 4WD Off-Road Racing Truck”, accessed 11 January 2017, via http://www.losi.com/Products/Default.aspx?ProdID=LOSB0019BLKBD
4. Ventura, P. J. C., Ferreira, C. D. H., Neves, C. F. C. S., Morais, R. M. P., Valente, A. L. G. and Reis, M. J. C. S., “An embedded system to assess the automotive shock absorber condition under vehicle operation”, IEEE SENSORS conference, Leece, Italy, 2008, pp. 1210-1213.
5. Yin, S. and Huang, Z., “Performance Monitoring for Vehicle Suspension System via Fuzzy Positivistic C-Means Clustering Based on Accelerometer Measurements”, *IEEE/ASME Transactions on Mechatronics*, Vol. 20, No. 5, October 2015.
6. Woodman, O. J., “An Introduction to Inertial Navigation”, University of Cambridge Computer Laboratory Technical Report, UCAM-CL-TR-696, August 2007, 37 pages.
7. Gordanlnejad, F., Fuchs, A., Dogrour, U., Kanakas, S., Liu, Y., Hu, B. and Evrensel, C., “A New Generation of Magneto-Rheological Fluid Dampers”, Final Progress Report, Department of Mechanical Engineering, University of Nevada, Reno, August 2004.