Military Diagnostics, Prognostics,

And Logistics Whitepaper -

A Way Forward

Document Summary

The industry discussion of comprehensive diagnostic, prognostics, and the ensuing logistics support has been observed and studied by Pi Shurlok. This discussion has shifted to broad brush comprehensive communications networks, logistics, and massive database management of fleet prognostic data manned by analysts looking for trends. Presented here is a way forward to achieve the lower level application of the diagnostics, prognostics and logistics problem. The main focus is on individual vehicle or even subsystem and how to achieve the benefits of advanced diagnostics and prognostics without the cost overhead of a more expansive system. The way forward is to build on COTS (Commercial Off The Shelf) intellectual property from the automotive domain and apply this to a military environment. A phased implementation that will scale and grow with current vehicles, in the form of retro-fit kits, is proposed and also an approach to new vehicles as they are developed. The proposal is to use well proven techniques from the COTS automotive domain and combine these with innovative new ideas to address the diagnostic, prognostic, and logistic requirements of military users that will deliver a robust, achievable, and common solution.

Authors: David Stamm, David Price

Contents

1. Introduction and Scope

This document describes the commonly expressed desires of the US military and suppliers regarding ground vehicle diagnostics, prognostics, and logistics. These desires are contrasted against a concept of Commercial Off the Shelf (COTS) Intellectual Property (IP), now termed COTS-IP. This whitepaper is based on the view of a military electrical system supplier with experience in the automotive industry.

It has been observed that some military suppliers do not leverage existing technologies to provide a robust low-cost solution. These bespoke designs run into issues of infrastructure, standardization, and high development costs. Through observation, Pi Shurlok noted similarities between the military needs for diagnostics and prognostics and the existing solutions being used in the commercial vehicle industry.

COTS-IP seeks to harvest the most suitable technologies from the consumer automotive domain and apply them with suitable systems engineering for military applications. The rationality for this approach is intended to improve the time to field of new technologies and to reduce costs.

The COTS-IP approach was applied to a technology leveraged, distributed system such as ground vehicle diagnostics, prognostics, and logistics because it has the largest Venn diagram intersection with the commercial vehicle market. Many of the issues the military is facing have been solved in part by a variety of vehicle technologies. In some select cases only minor robustness improvements and repackaging are required to field these technologies. Pi Shurlok has significant experience from the commercial world and has also gained experience over recent years with the application of systems within these types of architectures on tactical wheeled vehicles.

This paper specifically addresses the issues for implementing a comprehensive diagnostic and prognostic system on a military vehicle. The scope of this whitepaper is limited to the transmission of the data from the vehicle. No explicit direction is provided as to what should be done with the data once transmitted off the vehicle.

2. Terminology and Definitions

Most critically to this discussion is the agreement of a standard set of definitions. It is a common error that between industries terms are often mixed and misused. Thus to facilitate a common understanding the following definitions are offered.

2.1 Diagnostics

<u>Dictionary Definition</u>

"serving to identify or characterize; being a precise indication"

<u>Industry Definition</u>

A system that estimates the <u>current</u> status of other systems.

In the transportation industry, the concept of a diagnostic is not unfamiliar. Diagnostics have existed in several different forms over the decades. In the simplest form, it is the red oil pressure lamp on the dashboard. In modern vehicles these are software algorithms that automate many of the same historically manual procedures a technician may perform to determine the root cause of a failure.

2.1.1 Components of a Diagnostic

For the purposes of this paper diagnostics will be treated as software algorithms that actively monitor the health of another system and provide status to the user. These diagnostics are matched with pre-defined thresholds above or below which an alarm will be triggered. In this sense there are two distinct aspects to the diagnostic, the monitor and the alarm.

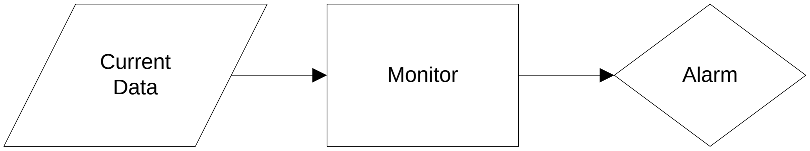

Figure 1 - Diagnostic Components

2.1.1.1 Diagnostic Monitors

The sophistication of the monitor can vary widely depending on the complexity of the system being diagnosed, or the level of inference made into the system. In the most direct scenario there might be a sensor directly measuring the item you wish to diagnose, such as oil pressure. The monitor in this case can merely detect if the pressure is below a predefined threshold.

In an indirect scenario there may not be a sensor directly measuring the item being diagnosed, such as when trying to calculate the life remaining in a suspension shock absorber. In this scenario multiple sensors, relating data to a mathematical model of the system, are required to determine if the damping in the shock absorber is within desired specifications or has degraded. This type of prognostic might use accelerometers in the sprung and un-sprung masses in the vehicle. The mathematical model would look at these two acceleration signals and determine if the damping has changed or degraded. The amount of degradation can be quantified and related back to percentage of life remaining.

Diagnostics can take variable amounts of time to execute. Diagnostics are not required to produce instant results, although certain diagnostics are capable of doing so. Certain diagnostics, that analyze trends, may take a large sample of data before it is capable of making a determination. These diagnostics can take minutes or hours or even days to execute fully.

2.1.1.2 Diagnostic Alarms

The alarm is how the monitor output is related to the end user. The alarm is effectively a filter of information. The alarm continuously monitors the diagnostic for the end-user and only alerts the end-user when the monitor has exceeded some pre-defined threshold. This is a simple yet effective method to allow the end-user to focus on other tasks, and only focus on a system when an alarm becomes active.

A significant area of concern is the cost of false-pass and false-fail determinations in the diagnostics alarms. The commercial vehicle industry has struggled with this in the past. In the commercial vehicle industry a false-pass often results in a federal non-compliance for regulations, where a false-fail will often result in increased warranty and customer annoyance. Translated to the military domain a false-pass has serious impacts in placing people in situation where their mobility is compromised. A false-fail has similar impact to the commercial industry in increasing costs and annoying the end user.

2.1.1.3 Diagnostic Outputs

Once an alarm has been set, there are a number of other pieces of information that can be logged to assist in further investigation of the issue.

- Diagnostic Not Run Flag – current ignition cycle
- Diagnostic Fail Counter – current ignition cycle
- Diagnostic Pass Counter – current ignition cycle
- Diagnostic Fail Counter – total failures since reset
- In-Use Rate – ratio of number of completions to number of enablements
- Freeze Frame Data – snapshot of other sensor data at the moment a fault occurs

This information can be stored away for use by maintenance personnel. It is not expected that the vehicle operator will require this information in the course of normal operation. This information is commonly stored in civilian passenger cars and provides vital information after a fault to help the maintenance personnel better diagnose and repair the root cause.

2.1.1.4 Diagnostic Summary

Diagnostics provide an assessment of a <u>current</u> state of a system. Several key observations should be made based on this statement.

- Diagnostics are statistically more robust than prognostics, however diagnostics are still estimates.

- Diagnostics are intended to communicate information that is of near-term importance.
- Diagnostic histories can be very useful to add depth to the data.
- Diagnostic notification may vary based on role (operator, maintenance, logistics, etc).

2.2 Prognostics

Dictionary Definition

“predictive of something in the future”

Industry Definition

A system that estimates the future status of other systems.

The civilian transportation industry is rapidly adopting the concept of prognostics. The enhanced level of information provided by prognostics offers military mission planners, logisticians, and the crew increased awareness regarding the health of the equipment under their control. This increase in information allows for a litany of benefits, many have been discussed broadly in other publications and presentations. A few benefits are listed here for reference

- Condition based maintenance – fix only what requires repair
- Reduced logistics tail for forward operating forces – less manpower, resources, and hardware required to maintain vehicles.
- Increased vehicle availability (MTBR)

2.2.1 Components of a Prognostic

The components of a prognostic are not particularly different than a diagnostic. The workhorse of the prognostic is the monitor. The goal of the monitor changes in the prognostic case in that the monitor is observing the current system to predict future behavior. Figure 2 - Prognostic Components - Fielded shows the rudimentary nature of a fielded prognostic and how it is similar to a diagnostic.

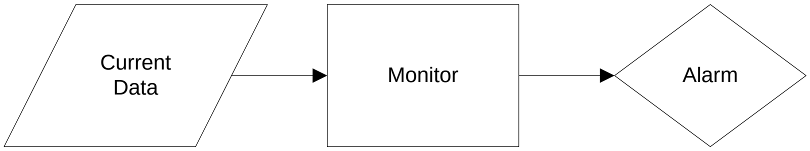

Figure 2 - Prognostic Components - Fielded

During the development of the prognostic algorithm it is common to acquire additional data to facilitate the tuning and optimization of the prognostic predictions. Figure 3 - Prognostic Components - Development shows how the block diagram changes for the Development of the prognostic. This distinction is mentioned here since the development process is the one most people associate with prognostics. The overhead in acquiring the data is often perceived to be a perpetual burden, and not an early phase in the development of a prognostic.

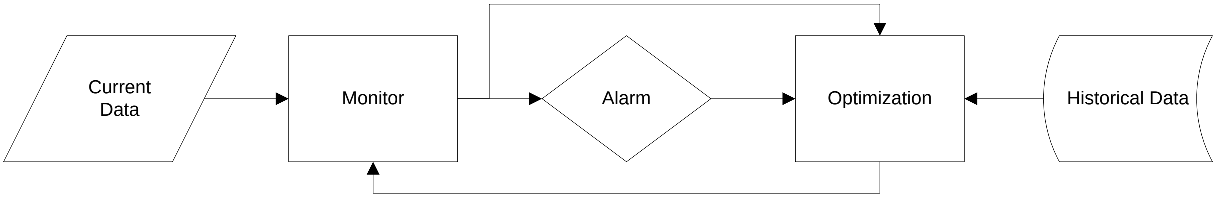

Figure 3 - Prognostic Components - Development

2.2.1.1 Prognostic Monitors

Prognostic monitors, although at a high level are very similar to diagnostic monitors, differ significantly depending on the complexity of the system. One key distinction is the temporal nature of the monitor. The traditional concept of a prognostic monitor is to predict the future performance of a system, or the remaining life in a given component.

True forward looking monitors of systems, or components, will require multiple sensor inputs and mathematical models to relate sensor data to life expectancy. The creation of these monitors will require the following items.

- Intimate knowledge of the system or component under analysis
- Statistically significant data samples relating model inputs to product/system durability
- Mathematical models relating inputs to product/system durability

As these items can be daunting to acquire, other options are available to segue into prognostics with less upfront investment. Simpler prognostic monitors can be employed to provide rudimentary estimates of accumulated usage.

Rather than investing resources to determine the life-span of an unknown component or system, it is often more economic to simply monitor the usage of the component or system. Once a statistically significant sample of data is available from these monitors, an estimate of life can be made. Often this data is more robust and reliable than laboratory or simulation based data. The regression of this raw data on-vehicle is a key facet to the future architecture discussion later in the paper.

If the component or system was tested to destruction during development validation (DV) or production validation (PV), this data can be used to jumpstart a prognostic algorithm. Often times this data will not completely represent the field use, and will need adaptation to match the field data. If this DV/PV data is not available to help set the initial prognostic life estimates, the in-use field data can be utilized.

If failures are being observed in advance of the expected lifespan two main causes can be blamed. First, the component has changed in some way and now has a flaw leading to pre-mature failures. At this point the prognostic is not useful since the data used to seed the prognostic algorithm was done on a component with different properties. Second, the prognostic algorithm is incorrect and does not reflect the operating environment of the actual vehicle. This would drive a modification of the prognostic to better accommodate the real-world data observed.

If failures are not being observed around the expected lifespan, similar conclusions can be made as in the case of pre-mature failures. In each case, a potential modification of the prognostic algorithm may be chosen to better reflect the real-world test data becoming available.

With these problems of usage duty cycle and the percent life calculation notes, in many cases simply knowing the usage of a given component or system is sufficient to meet the needs of the end-user of the prognostic. Not all systems require a high-resolution calculation of remaining life. In the future architecture an alternative to the percent life calculations is presented.

2.2.1.2 Prognostic Alarms

Prognostic alarms do not differ from diagnostic alarms discussed earlier. In essence both diagnostics and prognostic alarms are indications of a fault in the system. The distinction is that one is in the future. For the purposes of this paper no further distinctions will be made between the two alarms.

2.2.1.3 Prognostic Outputs

In addition to the prognostic alarm indication, additional (regressed) data should be provided to place the alarm in context. The prognostic output content differs from the diagnostic output. For a prognostic to be relevant there must be an objective measurement of the system tied to a temporal measurement of the system.

An example would be pump charge rate (quantitative) and pump hours (temporal). The pump charge rate can be monitored by the prognostic and when a decrease in the charge rate is detected, this can be the alarm trigger for the prognostic. The decrease in charge rate may be an early indication of a future pump failure. What places the pump charge rate data in context of usage is the temporal data. The appropriate measurement for this is the hours logged of pump operation. The addition of temporal data provides statistical context of how the data relates to other historical data. Providing the temporal data allows the data to be mined later for more statistical significance. Outliers and trends can be identified more easily. Distinctions can be made between early-life failures, mid-life failures, and long-term wear out. Although temporal data can also be added to the diagnostics, it is not as pertinent to the purpose of the diagnostic.

2.2.1.4 Prognostics Summary

Prognostics provide an assessment of a future state of a system. Several key observations should be made based on this statement.

- Prognostic results are estimates
- Prognostics results can be changed based on future vehicle fleet behavior
- Prognostics results are by definition less accurate than diagnostic results
- Prognostics intend to communicate information that is of medium to long-term importance.

3. Future System Architecture

Section 2 of this whitepaper sought to set a level ground of terminology from which we can form a meaningful discussion about future system architectures to support ground vehicle diagnostics and prognostics. Section 3 is the basis of the collective experience of Pi Shurlok on our commercial vehicle and military vehicle projects. As we have worked in both areas on similar projects we have observed similarities and areas for economies of scale to benefit both the commercial and military vehicle industries.

The architecture of future systems is discussed here since it has a critical bearing on the cost of the overall system. Poor choices in architecture can result in sharp increases in resource requirements for the downstream users. The architecture suggestions focus on a few key points.

- Live monitoring of trends, variation, and or deviations from normal instead of elaborate percent life computations
- Regression of raw data on-vehicle
- Store only regressed data
- Sub-system ECUs perform control and diagnostic/prognostic functions (where possible)

- Sub-system ECUs report to vehicle level computer for off-vehicle transmission
- Diagnostics: Transmit upon change only
- Prognostics: Transmit only when requested

One of the most important facets of the proposed architecture is the on-vehicle regression of the data for prognostics. This concept meshes with the prognostic attitude to monitoring the data for trends, variation, or deviations, rather than elaborate computations of percent life remaining. This approach is beneficial since the diagnostic and prognostic algorithms begin to look very similar. This reduces software development costs, full-life testing costs to determine useful life, and computational overhead. Additionally it provides a more objective measurement of the health of the system, rather than an abstracted percent life remaining metric.

From these key points we focus the architecture to use as much COTS-IP as possible to minimize the development costs of a system. By using methods and protocols that are already in use in other industries significant cost reductions can be realized. The COTS-IP leveraged in this architecture is:

- CAN communication
- OBD software methodology from the commercial vehicle industry
- SAE communication protocols
- Fault and alarm communication from the commercial vehicle industry

The following sections will discuss how the key points of the architecture and the COTS-IP have been integrated and applied to provide a comprehensive diagnostic and prognostic suite. Pi Shurlok has experience of implementing a number of systems and developing this type of architecture on military vehicles in recent years.

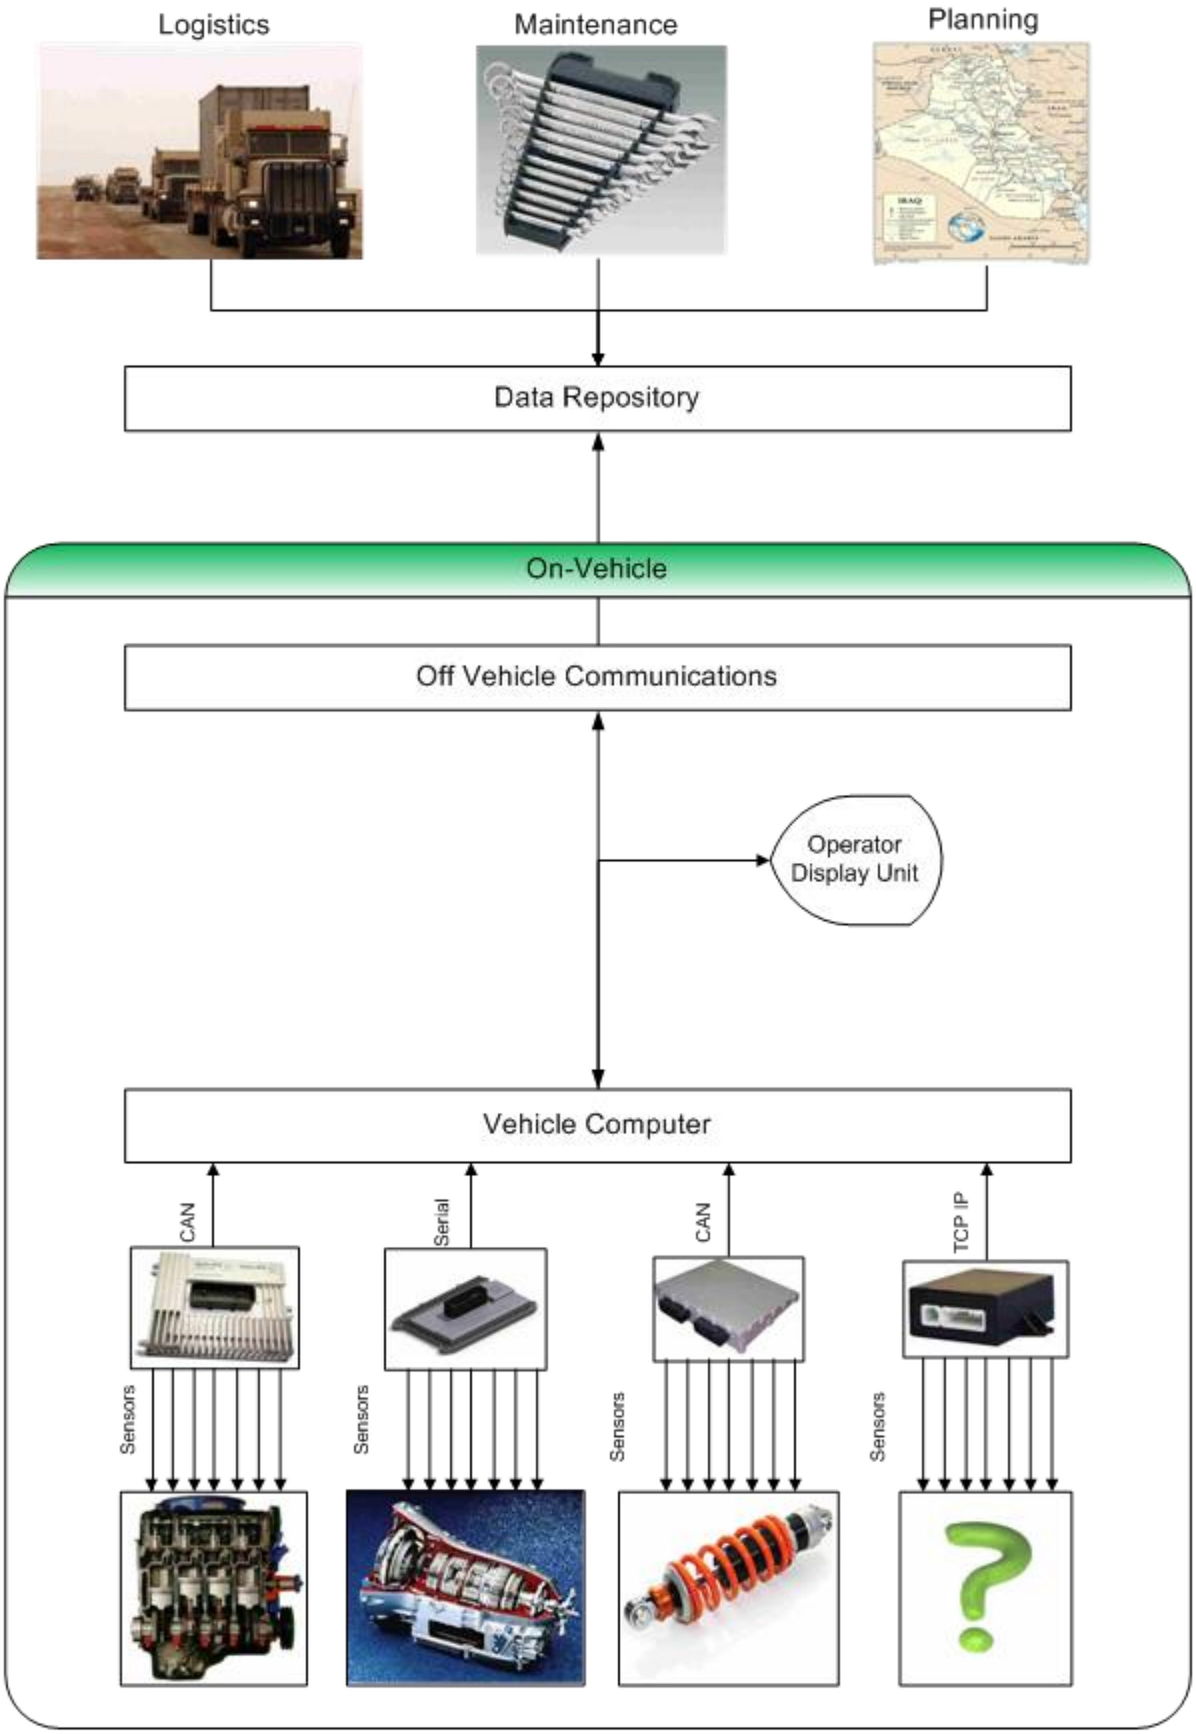

Figure 4 - Simple Block Diagram of Architecture

3.1 Positive Architecture Aspects

The architecture proposed here focuses on minimizing the off-vehicle communications and mimimizing the computational impact to a system when prognostics and diagnostics are added. To accomplish this will require a higher degree of computing power for each subsystem. The modest increase in computing power at the sub-system level will pay benefits in the areas of communications bandwidth and modular expansion.

Additional sub-systems can be added to a vehicle with only fractional impact to the communications bandwidth. Since each system is reporting only bytes of data for status and metrics, the increment in communications bandwidth is minimal. New systems can be fielded without maintenance of the diagnostic and prognostic infrastructure on the vehicle.

Storing only the regressed data reduces the memory requirements of the system (overall). Development may still require additional data acquisition.

Each subsystem reports data up a chain abstracted from off-vehicle communications or data storage. Changes can be made upstream without affecting the lower level systems.

On-board regressed data is a key enabler for real-time status monitoring and health awareness of the vehicle. Without on-board data regression the operators will not have real-time awareness of the state of the vehicle.

Monitoring trends in data, variation, and or deviations provides an objective measurement of system health and performance that an abstracted percent life calculation can not provide.

Significant cost savings can be realized since the extensive development testing to provide full-life durability data can be eliminated.

Establishing threshold values for a prognostic alarm is objective and based on system requirements for performance, not based on abstract percent life calculations.

Allowing the subsystems to have ownership of the diagnostic and prognostic allows those people most knowledgeable about the system to own the diagnostic and prognostic algorithms.

With on-board data regression the communications off-vehicle can be conducted at a lower priority. Data can be transmitted on an as-needed basis when the off-vehicle end-user requests the data. This keeps the communications bandwidth low.

Since raw data will not need to be transmitted off-vehicle, the on-board prognostic has an opportunity to provide a more accurate result as the fidelity of the data will not be encumbered by resolution or sampling time reductions due to off-vehicle communications.

Additionally there is less risk in making incorrect diagnostic or prognostic decisions due to data being time-shifted as a result of to off-vehicle communication issues. The data is being used live, as opposed to overlaying multiple data streams after the fact.

3.2 Negative Architecture Aspects

Any chosen architecture will have some deficiencies or areas of compromise. One downside of this architecture is that the end user has no control over prognostic algorithms for the regression of data. With a system that simply streams raw data to the end user; the prognostic algorithms can be changed easily. This should be less of a concern as the prognostics mature.

The deployment of updates or enhancements to prognostic algorithms is multiplied by the number of vehicles fielded. This issue is not much different to the upgrade of other electronic systems on ground vehicles.

The computational overhead will increase for each subsystem. Additional computational power will be required to calculate and regression of diagnostic and prognostic data. It has been observed that the code required to provide diagnostics and prognostics for a system can be equal to the code needed to control the system. This is a 100% increase in code space and throughput at a subsystem level.

3.3 Data End-Users

When the three main end-users of the data repository coordinate and share information efficiencies can be gained in the overall operation. The details of this coordination and use of the data repository is not the focus of this paper, however some brief notes regarding the potential usage and relationships are provided below.

3.3.1 Logistics

The architecture discussed here does not significantly change the interface to the logistics community as discussed in other papers on the topic. Using the architecture discussed, the logistics community can mine the data repository at a high frequency for diagnostic information. In general the logistics role here is to flag issues and prepare for a worst case scenario. The details of the repair activities and the specific failure are not of a concern here, only the logistical implications of a specific diagnostic fault or prognostic fault being set on a vehicle. The related logistics might be items such as:

- Jump-start the requisition process to get parts and services arranged for repairing vehicles.
- Provide compensation plan for loss of vehicle
- Communicate downstream to maintenance and planning regarding vehicle status.

3.3.2 Maintenance

The maintenance organization will take cues from the logistics community to prepare for the receipt of a vehicle in need of repair. The following actions could be taken by the maintenance organization by utilization of the data in the repository.

- Preparation, schedule, and repair of vehicle
- Provide status to Logistics and Planning organizations
- Provide feedback regarding robustness of diagnostic and prognostic algorithms

3.3.3 Planning

The mission planning organization can also make use of the data repository to understand the following facets of the vehicle fleet.

- Health and status of each vehicle
- Ability to hand-select the vehicles used based on their suitability and health

4. Future Work

The work performed thus far on tactical wheeled vehicles can still benefit from improvements and further integration with other systems. The architecture and the implementations of actual vehicles that have followed can benefit from more extensive integration of the various systems. The following sections discuss areas of future work that should be investigated to advance the cause of diagnostic and prognostics on military ground vehicles.

4.1 Intra-Vehicle Communications Architecture

To simplify the introduction of future systems, a common communications architecture at the physical layer should be adopted. Current vehicles use a variety of methods for communication of data within a vehicle. Communications vary from CAN, Serial, Ethernet and wireless. Adoption of a common communications architecture that all fielded vehicles must adhere will save cost and provide a basis for more competitive systems development.

By increasing the accessibility of intra-vehicle communications, many optionally equipped systems may not require as many sensors, thus reducing cost. If data can be more readily shared across a vehicle platform, redundancy can be reduced. Sharing of common parameters such as vehicle speed, ambient temperature and battery voltage can be easily facilitated with a common communications architecture.

Systems developments often flourish when open communications architectures are adopted. When one system can gain additional knowledge via another sub-system the effect can be to improve the performance, reliability, or robustness of the system in question. This fact can be backed up by the direct experiences of Pi Shurlok on a number of military projects that did and did not adopt an open communications protocol.

It is the opinion of the authors that adoption of a CAN based communications protocol between the sub-system ECU and the vehicle computer would be the best combination of communication architectures. The sub-system ECUs may require other interfaces to handle the specific needs of their system, but standardizing the diagnostic and prognostic interface up to the vehicle computer provides the gains mentioned later in the document.

The basic CAN architecture is simple, low cost, and robust. This architecture would not be exclusive or exotic, thus more potential suppliers could be involved. Additionally, expansion of the standards to address diagnostic and prognostic information for the military would not be a significant expansion or re-invention.

4.1.1 Communications Standard

Similar to what the SAE has done for the passenger car and heavy-duty truck markets (J1939), the military should follow suit. Development, or more accurately the enhancement, of a communications standard for intra-vehicle communications should be established. The needs of the military are slightly different than the commercial vehicle markets and thus a straight adoption of the existing standards would not be appropriate.

The benefits of a common communications protocol is well documented regarding the cost benefits. Allowing systems developers to focus efforts on technology development tasks, rather than system integration tasks, provides one of the largest cost reductions.

The development of a standard will take several years, many meetings, cost several million dollars, and be an arduous task for such a large organization such as the military. It is, however, worth the investment based on the potential returns as provided by the J1939 specification developed to the heavy-vehicle industry.

4.2 Up-Integration of Diagnostic Systems

As legacy vehicles are phased out, the new systems, designed into the replacement vehicles, need to consider the problem of up-integrating the diagnostic systems. Up-integration involves the points discussed earlier about de-centralizing the diagnostics and placing the diagnostic responsibility into the ECUs for each subsystem. External “bolt-on” HUMS (Health and Utilization Monitoring System) will be gone, and each subsystem will be responsible for its own diagnostics and prognostics. Up-integrating culminates with the systematic design of interactions between all systems at an initial design stage, as opposed to adding overlay / appliqué systems after the vehicle is designed.

On existing fielded vehicle platforms there can be overlay / appliqué systems that are installed after the vehicle was initially constructed. These systems add sensors, wiring, and a processor to monitor other systems on the vehicle and do diagnostic and prognostics. Appliqué systems should be abandoned and the diagnostic and prognostic systems permanently integrated into the vehicle architecture. Significant benefits can be realized if these systems are integrated earlier in the vehicle development process. Size, weight, and power (SWaP) can be optimized if the original vehicle manufacturer integrates the various systems, rather than field retro-fit kit installations which can often compromise the overall SWaP efficiency of the system installation.

Finally, system developers need to protect their ECUs for the additional requirements of up-integrated diagnostics and prognostics. These protections need to include microprocessor throughput, memory, and software enhancements. These ECUs will also need to be able to communicate according to the standard communications protocols as defined earlier in this whitepaper.

5. Summary

This whitepaper has discussed the experiences of Pi Shurlok in providing a comprehensive diagnostic and prognostic suite on a several military ground vehicle applications. It was observed that many military systems developers have taken paths that did not take advantage of existing (COTS) approaches to the problem of diagnostics and prognostics. Through these projects Pi Shurlok made observations of the similarities between the military needs for diagnostics and prognostics and the existing solutions being used in the commercial vehicle industry.

Military and commercial vehicle designers perform vehicle development along very similar paths. There is good cross-over in the mechanical arena, but relatively little cross-over has been observed in the electronic / control domains.

The successful deployment of diagnostics and prognostics can be performed by leveraging current technology, methods, and systems. This leveraging can provide a significant increase in awareness for the operator, logistics, and maintenance personnel. However with additional efforts to standardize some of the common interfaces and architecture across vehicles, further benefits can be achieved.

The way forward on diagnostics and prognostics appears to be a two-pronged approach. First, continue to leverage COTS-IP to develop diagnostic and prognostic systems for legacy and new vehicle programs. Secondly, start efforts to develop standards from which more integrated systems can be developed for future vehicles.

Standardization of the diagnostics and prognostics communication infrastructure will provide an effective platform to ensure fair and open competition among suppliers while reducing overall cost. Standardization of communications is as significant as standardizing whether SAE or metric fasteners will be used on a vehicle and what that choice will mean to the overall cost of development, operation, and sustainment over the course of the vehicle life.

These architecture changes at a low-level in the vehicle are intended to provide the end-users of the data with the most robust, timely, and relevant data possible to make their respective decisions. Essentially to meet the end-users needs, the problem must be solved at the individual vehicle level, and this starts with the proper vehicle architecture.