Defence Research and Development Canada
Recherche et développement pour la défense Canada

Unmanned Ground Vehicle Electronic Hardware Architecture

A Flexible and Scalable Architecture for Developing Unmanned Ground Vehicles

G. Broten and S. Monckton
Defence R&D Canada – Suffield

Technical Memorandum
DRDC Suffield TM 2004-122
June 2004

Canada

Unmanned Ground Vehicle Electronic Hardware Architecture

A Flexible and Scalable Architecture for Developing Unmanned Ground Vehicles

G. Broten
S. Monckton

Defence R&D Canada – Suffield

Defence R&D Canada – Suffield

Technical Memorandum

DRDC Suffield TM 2004-122

June 2004

Author

G. Broten, S. Monckton

Approved by

D. Hanna
Head/Tactical Vehicle System Section

Approved for release by

R.W. Bide
Chairman/Document Review Panel

Abstract

The Autonomous Land Systems (ALS) initiative has been tasked with researching and developing innovative autonomous vehicles that will assist the Canadian Forces in performing their duties in the 21st century. This research will continue for many years, on many different types of unmanned ground vehicles and thus requires a flexible and scalable robot electronic hardware architecture. This architecture must be applicable to vehicles ranging from small indoor research platforms to large outdoor military vehicles such as the LAVIII. This report describes the hardware architecture, known as the Robot Nervous System, which will serve as a backbone for current and future Autonomous Land Systems development. The architecture is distributed in nature and thus strongly supports the scalable requirements of the R&D program. Presently the architecture uses three scales of processors, each well suited for their specified tasks. The processors are linked via appropriate network topologies that enable each processor type to be leveraged to its maximum capabilities. This hardware architecture is currently being implemented in a variety of autonomous vehicles that are under development at DRDC Suffield.

Résumé

L'initiative des systèmes terrestres autonomes (STA) a été chargée de conduire la recherche et la mise au point de véhicules autonomes novateurs qui appuieront les Forces canadiennes pour effectuer leurs tâches, au 21è siècle. Cette recherche continuera au cours de nombreuses années, sur plusieurs différents types de véhicules terrestres sans pilote et exige par conséquent une architecture de matériel électronique de robotique qui soit adaptable et échelonnable. Cette architecture doit être applicable aux véhicules allant des petites plateformes de recherche d'intérieur aux gros véhicules militaires d'extérieur tels que le LAVIII. Ce rapport décrit l'architecture du matériel, connue sous le nom du Système nerveux du robot qui servira de squelette aux mises au point actuelles et futures des Systèmes terrestres autonomes. L'architecture est de nature distribuée et par conséquent soutient fortement les besoins d'échelonnement des programmes de R & D. L'architecture utilise actuellement trois niveaux de processeurs, chacun bien adapté à sa tâche spécifique. Les processeurs sont liés via une topologie de réseaux appropriée permettant que chaque type de processeurs fonctionne au maximum de ses capacités. Cette architecture du matériel est actuellement implémentée dans une variété de véhicules autonomes qui sont en cours de mise au point à RDDC Suffield.

This page intentionally left blank.

Executive summary

Background: Defence R&D Canada – Suffield, on behalf of Canadian Department of National Defence, is developing new and innovative land vehicles. An important aspect of this research is the investigation and development of autonomous intelligent systems to augment and/or replace existing tele-operation capabilities. This research is being conducted under the auspices of the Tactical Vehicles Systems Section as part of their Autonomous Land Systems project and is expected to define the future capabilities of autonomous unmanned ground vehicles (UGV's).

Principle Results: Future autonomous vehicles may range in size from small indoor vehicles to large outdoor platforms the size of the current military LAVIII. It is obvious that a large vehicle is capable of carrying a large electronic payload while the payload for a small vehicle is much more limited. This wide range of vehicles sizes requires a flexible and scalable electronics hardware architecture. This scalability and flexibility is required to allow researchers to focus on the development algorithms for robotic autonomy as opposed to writing or porting software to different electronics hardware.

This report defines an electronic architecture that is termed the Robot Nervous System. It describes a topology for processors, their interconnections and the functionality implemented at each layer, which results is a flexible electronic hardware architecture that is portable across varying types of unmanned ground vehicles. This architecture defines the four functional layers with each layer of the architecture implementing a specific types of functionality. The Robot Nervous System also defines the relationship between software implemented at the functional layers and the three types of processor that host the software.

The Robot Nervous System, by using a distributed computing approach that assigns the most suitable processor(s) to implement the required robot functionality, delivers the scalability and flexibility required by the ALS program.

G. Broten, S. Monckton. 2004. Unmanned Ground Vehicle Electronic Hardware Architecture. DRDC Suffield TM 2004-122. Defence R&D Canada – Suffield.

Sommaire

Contexte : R & D pour la défense Canada – Suffield est en cours de mettre au point des véhicules terrestres nouveaux et novateurs, au nom du ministère de la Défense nationale du Canada. Un aspect important de cette recherche est l'étude et la mise au point de systèmes intelligents autonomes visant à augmenter et / ou à remplacer les capacités existantes de téléopération. Cette recherche est conduite sous les auspices de la Section des Systèmes de véhicules tactiques dans le cadre du projet des Systèmes de véhicules tactiques et vise à définir les capacités futures des véhicules au sol autonomes sans pilote.

Les résultats de principe : Les véhicules autonomes futurs peuvent varier en taille allant de petits véhicules d'intérieur aux grandes plateformes de la taille du LAVIII militaire actuel. Il est évident qu'un gros véhicule est capable de transporter une charge utile électronique importante alors que la charge utile d'un petit véhicule est beaucoup plus limitée. Cet échelonnage et cette adaptabilité est requise pour permettre aux chercheurs de se concentrer sur la mise au point des algorithmes visant à l'autonomie robotique en opposition à la création de logiciels pour chaque matériel électronique différent.

Ce rapport définit une architecture électronique appelée le Système nerveux du robot. Il décrit une topologie pour les processeurs, leurs interconnections et la fonctionnalité implémentée à chaque niveau, ce qui résulte en une architecture de matériel électronique adaptable qui est transférable à plusieurs types de véhicules au sol sans pilote. Cette architecture définit les quatre niveaux fonctionnels et chaque niveau implémente des types spécifiques de fonctionnalité. Le Système nerveux du robot définit aussi la relation entre le logiciel implémenté et les niveaux fonctionnels ainsi que les trois types de processeurs hôtes du logiciel.

Le Système nerveux du robot permet l'échelonnage et l'adaptation requise par le programme STA en utilisant une méthode de traitement distribué qui désigne le ou les processeurs les plus adaptés à implémenter la fonctionnalité requise du robot.

G. Broten, S. Monckton. 2004. Unmanned Ground Vehicle Electronic Hardware Architecture. DRDC Suffield TM 2004-122. R & D pour la défense Canada – Suffield.

Table of contents

List of figures

List of tables

This page intentionally left blank.

1. Introduction

Continuing a tradition of innovative development, the Tactical Vehicle Systems Section of Defence R&D Canada-Suffield has launched the Autonomous Land Systems project to define the future capabilities of autonomous unmanned ground vehicles (UGV's) for the Department of National Defence.

Dudek[2] has defined unmanned ground vehicles[1] as either fully or semi-autonomous systems. A fully autonomous robot operates without full-time external human control. Though a semi-autonomous robot requires a full-time human operator, certain decisions can be performed without human intervention. This latter group is further composed of tele-robotic and tele-operated systems. A tele-robotic system possesses software that interprets the operator commands, while a tele-operated robot is directly and completely controlled by the operator.

1.1 Target Vehicles

The ALS project anticipates that future UGV's will range from small indoor platforms to larger outdoor vehicles comparable to the Canadian Forces LAVIII.

Figure 1: *Vehicle Platforms: LAVIII, Hummer, Outdoor Testbed, Co-operative Robot Group, Extreme Testbed*

Figure 1 illustrates the potential UGV platforms and payload capacities that range from hundreds of kilograms and cubic meters of space to a payload of a few kilograms with only a minimal amount of space available. This wide range of platform sizes requires a flexible electronic hardware architecture that can be scaled to suit the size of each individual platform. This scalability requirement affects both the hardware's physical size, weight and power requirements.

1.2 UGV Capabilities

The electronic hardware architecture must accommodate both current and future UGV capabilities. A short list of the major capabilities are:

- Learning

[1]The terms unmanned ground vehicles, autonomous ground vehicles, robots and mobile robot are used synonymously throughout this report.

- Perception
- World Representations
- Planning
- Navigation
- Behaviors and Skills
- Communication

Given this list, computing resources constitute a key component of the the electronics hardware architecture. While most of these features require pure computing resources, some elements will require the ability to communicate with and control external devices.

1.2.1 Sensor and Equipment Compliment

A fully equipped UGV will carry a variety of sensors, processors, communications devices, motors and power sources. The sensors allow the UGV to sense the external environment and a modern UGV may include:

- Cameras
- Laser ranging devices such as the SICK scanner
- Rate Gyro
- Accelerometers
- Inclinometers
- Global Positions System
- Optical Encoders
- Magnetometers
- Infrared or mechanical proximity sensors
- Electric Motors
- Hydraulics
- Internal Combustion Engines

The electronics hardware architecture must interface with and control a wide variety of sensors and other equipment. Some of these devices will be intelligent, with auto configuration, calibration, diagnostics, and high level Ethernet services, while others will function at a very low level, exploiting analog I/O or simple serial protocols for control.

2. Background

Fictional robots have been staple characters of science fiction novels since the 1940's[2] but it has only been since the late 1960's[3], with the advent of modern digital computers, that recognizable mobile robots have been realized. The ALS program builds upon a now significant body of background research in the definition of its electronic hardware architecture. Specifically, the DRDC architecture has been influenced by research conducted at Carnegie Mellon University[1], MIT[4], and JPL[5]. The UGV tele-operation research that has been performed at DRDC Suffield has also played an important role in shaping the ALS program[6]. Finally, research outside the robotics community has also shaped this electronic hardware architecture. In particular, distributed control systems that is exemplified by distributed control architecture that is implemented at the Canadian Light Source particle accelerator[7].

3. Robot Architectural Models

From early optimism of the 1960's, through disappointment of the 1970's and '80's to the pragmatic expectations of machine intelligence at the turn of the century, the history of robot architectures closely follows the history of Artificial Intelligence. Though the earliest work in mobile robotics used analog networks [8, 9] , early research in AI assumed that the root of intelligent behavior lay in cognitive ability and problem solving (e.g SHRDLU[10] and STRIPS[11]). This approach embedded the 'Sense-Model-Plan-Act' (SMPA) model, commonly called the deliberative approach to intelligence. Manipulation robots eventually adopted much of this linear approach to intelligence, using familiar programming methods to exploit machine vision, CAD world modeling, and precision control. Mobile robot applications quickly found weaknesses in this purely linear implementation, encountering computing and sensing limits that challenged the feasibility of a purely symbolic SMPA approach (notably Brooks [12, 4]). By the mid 1980's, innovative alternatives to symbolic SMPA appeared, decomposing and, at times, simplifying the general SMPA problem into multiple, smaller subproblems and exploiting multiple computing resources and simplified sensing over a range of time scales (from micro-controllers to workstations) effectively generating multiple small control threads. Indeed, MIT's subsumption architecture[4], Georgia Tech's Motor Schema[13] among others used multiple parallel Sense-Act or 'reactive' subsystems using relatively fine grained networks of simple processors and little or no state retention. In the most extreme case digital computing has even been replaced by simple analog networks reminiscent the early analog robots[14] .

Modern robotic systems have evolved to include both deliberative and reactive elements, exploiting networked resources operating at a variety of time scales and with a wide range of capabilities. Rather than binding robot hardware to a particular

[2] Issac Asimov published the influential novel Runaround, in 1942, where he introduced the now famous three laws of robotics.

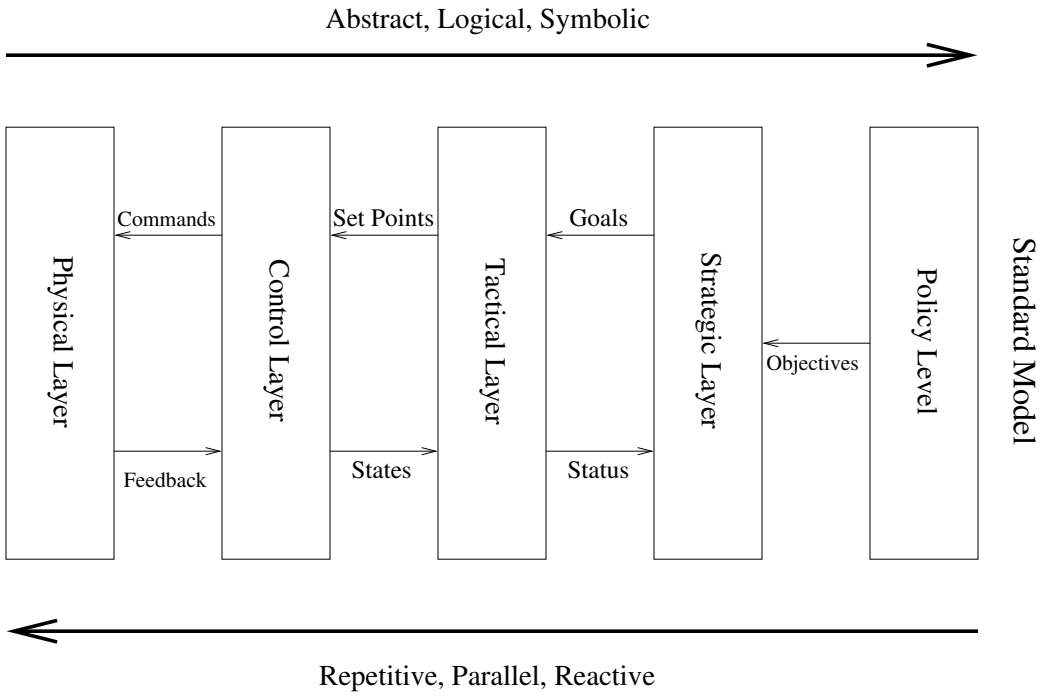

Figure 2: *Standard Model [1]*

'intelligence model', modern robot architectures now emphasize software middleware over philosophy. In effect, battles such as reactivity vs deliberation have disappeared, replaced by more technical issues centred on maintainability, flexibility, and robustness[15]. As a consequence a number of open source and commercial middleware packages used by a number of different groups – including DRDC– are in circulation. For example: ANCAEUS[6], JAUS, ACE, IPC, and RTC share similar abilities to dynamically link multiple modules within a computing network. At DRDC Suffield these packages are referred to as modularity toolkits they enhance the modularity and portability of software by insulating it from the details of how data is transfered between different software modules.

3.1 Standard Model

Regardless of architecture, all mobile robots must cope with the environment and, therefore, share some basic structural features. A standard robot model has been proposed that is suitable for all classes of robots[1] and this model is illustrated in Figure 2. This hierarchical model provides insight into autonomous robot structure that features abstract, symbolic and logical processing at high levels and repetitive, parallel and reactive computations at lower levels.

Further, the higher layers exhibit greater computational needs, less deterministic algorithms and longer reaction times. Each of the layers in the hierarchy implements a

set of functions that must work together in unison to allow the robot to achieve autonomy. The following sections describe, at a high level, the basic functionality implemented by each layer in the hierarchy.

3.1.1 Policy Layer

The Policy Layer encapsulates the robot's *raison d'etre*, providing the mission objectives whether defined by humans, deliberately encoded, or emerging as a result of robot intelligence. It is important to note that while these high level goals are abstract and, in general, do not significantly change over time[3], they *can* be modified. For example: during co-operative operations, group interactions between the robots may redefine the goal of a given machine to suit group objectives[4]. In military applications this layer might encode the "rules of engagement" prescribing a given robot's strategy during encounters with opposition forces.

3.1.2 Strategic Layer

The Strategic Layer is the layer that implements the intelligence of the robot. It corresponds to the deliberative, logical and goal-generating component of the robot. It deals with large encompassing issues confronting the robot such as:

- The generation of a world representation.
- Planning of the movement of the robot.
- Navigation through the world, in the attempt to achieve a goal.
- Perception of the environment, in terms of interpreting what the robot "sees".
- The understanding of its environment.
- Situational awareness.
- Multi-robot coordination and cooperation.
- Interpretation of status reports from the lower layers.

The Strategic Layer must process the information for all of these components and derive a plan of action to be taken within the confines of the directives of the Policy Layer.

[3] An example high level goal is for the robot to stay alive, which is abstract and is not subject to compromises.

[4] An example of this type of goal change could be when a robot, upon attaining a desirable location, acts as a communications relay between a group of co-operating robots.

Issues confronting the Strategic Layer tend to be computationally complex, often requiring a search through a large number of possible alternatives. Many solutions involve artificial intelligence and learning systems, where the robot learns the appropriate actions via investigation and feedback.

The Strategic Layer operates on complex data over a relatively long term. Given the computational complexity, optimal solutions over long and uncertain periods often consume considerable time and resources.

3.1.3 Tactical Layer

Conceptually this layer is the overall supervisor of the robot. It maintains maps and generates short term plans for the control layer at a faster rate than the Strategic Layer. From a military perspective, this layer implements tactical behaviors such as seeking cover when under fire.

This layer views the directives from the Strategic Layer as recommendations, overriding them when necessary. It hosts the navigator component that, if possible, ensures the robot follows the route proposed by the Strategic Layer, using a local map to plan and execute route deviations.

The Tactical Layer runs on a tactical time scale which is shorter and therefore the Tactical layer, more responsive than the Strategic Layer, swiftly responds to both Control Layer feedback and any new information provided by the robot's perception system.

3.1.4 Control Layer

The Control Layer deals with the mundane, but essential, services that make the robot move and react to the stimuli from external sensors. It deals with generating signals to drive motors and reading the feedback from devices such as optical encoders. Given the Control Layer's continuous interaction with external devices, this layer must run under real-time constraints. This layer's limited, but essential, autonomy, consists mainly of emergency responses such as stopping or cutting power to motors or actuators when it is unable to follow commands. An example of such a situation could be when the instrumented bumper detects the presence of an obstacle blocking the forward progress of the robot.

3.1.5 Physical Layer

The Physical Layer corresponds to processes allowing the Control Layer to communicate with external devices.

Layer	Robot Nervous System	Standard Robot Model
1	High Level Intelligence	Policy and Strategic
2	Intermediate Intelligence	Tactical
3	Control	Control
4	Physical	Physical

Table 1: *Robot Nervous System vs. Standard Robot Model*

3.2 Robot Nervous System

Section 3.1 describes the Standard Model of a robot which was DRDC Suffield's starting point for formulating a robot electronic hardware architecture. This electronic hardware architecture has been termed the Robot Nervous System (RNS) and it implements a hierarchy that is largely based upon the Standard Robot Model. Table 1 shows the correspondence between the RNS and the Standard Robot Model.

As can be seen from the table, the RNS combines the Policy and Strategic Layers of the standard robot model into a single layer called the High Level Intelligence Layer. But otherwise, the two topologies are the same.

The term Robot Nervous System highlights the similarities between the defined architecture and the central nervous system found in humans. While biological correspondence is not an endorsement of an engineered solution, nature has evolved many clever solutions to challenging problems. These solutions can be an excellent source of motivation, inspiration and guidance for the development of future devices[16].

These similarities between the RNS and the central nervous system are shown in Figure 3. The central nervous system is comprised of four major regions: the spinal cord; the brain stem; the cerebellum and the cerebrum[17]. These regions closely correspond to the layers defined by the RNS.

The RNS defines the software's functionality at each layer and suggests the prerequisite computing hardware. This approach assigns appropriate hardware, in terms of electronics and processing capabilities, to each layer in the hierarchy.

The most significant functions performed by each component of the central nervous system and the equivalent functionality of the RNS are listed in Table 2.

4. Implementation

Thus far this document has described the RNS in high level terms. Figure 4 shows a fictitious implementation of the RNS in terms of sensor inputs, intelligence implemented, outputs from each layer, and communications between layers.

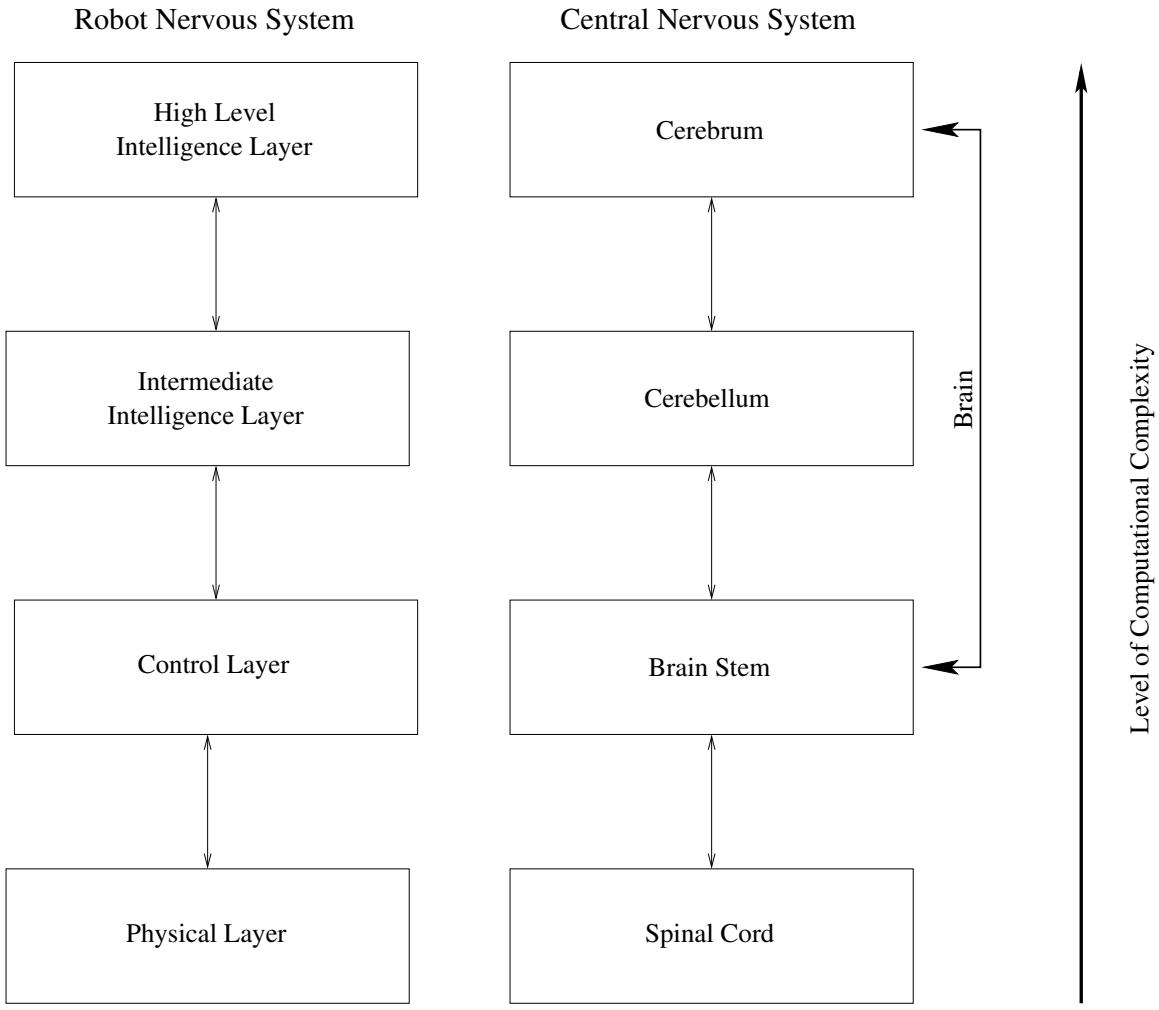

Figure 3: *Robot Nervous System and Central Nervous System Comparison*

Central Nervous System		Robot Nervous System	
Component	Function	Layer	Function
Cerebrum	●The conscious functions of the nervous system. ●Deliberative responses with active choices.	High Level Intelligence Layer	●Mission objectives. ●Deliberative, logical, goal generating.
Cerebellum	●The coordinator in efferent voluntary muscle systems. ●Conditioned responses (default choice) that can be overridden.	Intermediate Intelligence Layer	●Robot supervisor. ●Prioritizes, rationalizes and translates high level directives into set points.
Brain Stem	●Connecting link between cerebral cortex, spinal cord and cerebellum. ●Center of integration for visceral functions such as control of heart and respiratory rates. ●Integration center for various motor reflexes. ●Reactive reflex response (no choice).	Control Layer	●Synthesizes and monitors low level signals that generate robot motion. ●Monitors sensors to determine robot status and pose. ●Reactive responses.
Spinal Cord	●Connecting link between brain and body. ●Low level center for integration of motor activity.	Physical Layer	●The communications network that links the various devices together.

Table 2: *Robot Nervous System and Central Nervous System Functions*

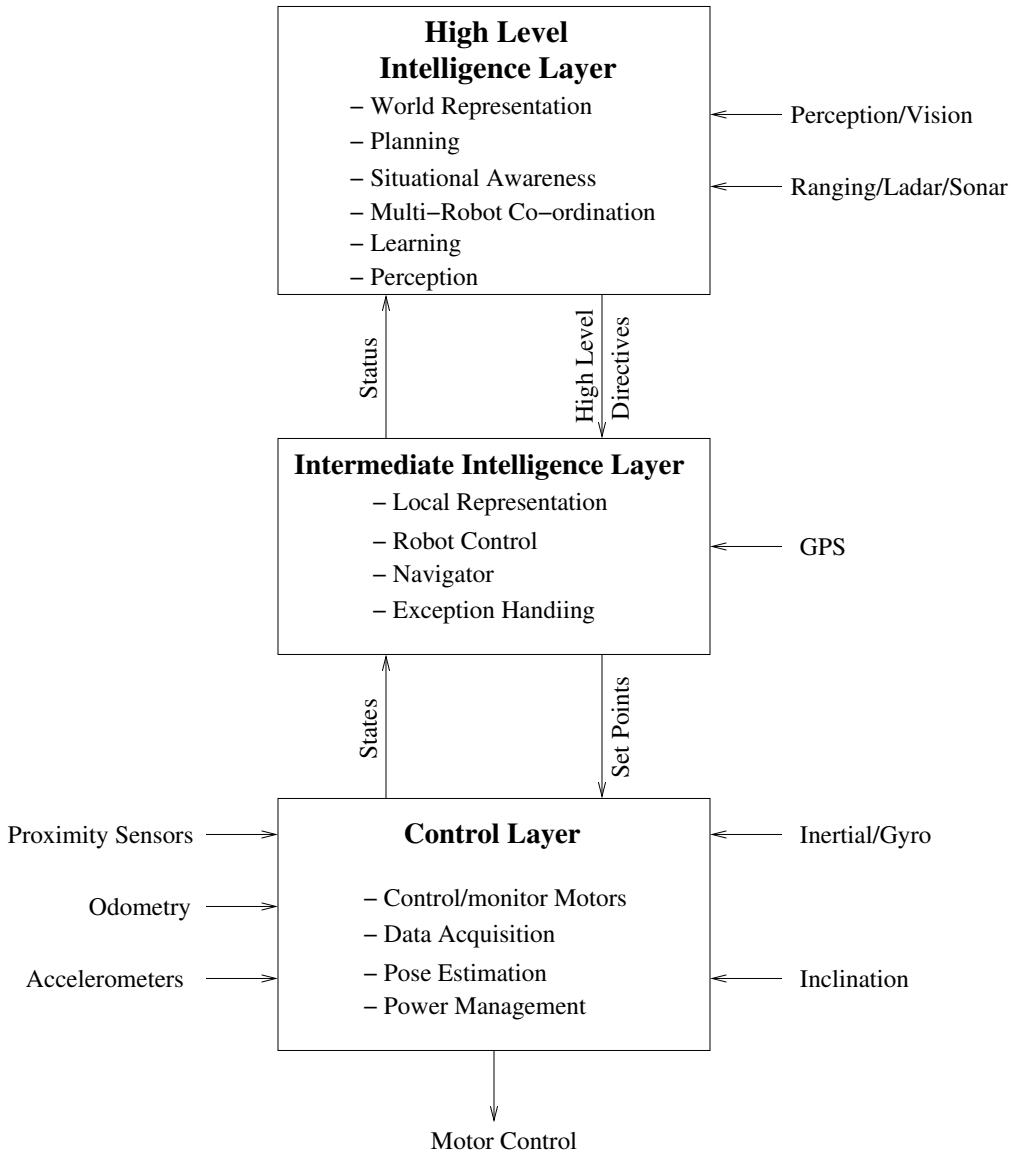

Figure 4: *Robot Nervous System Implementation*

The requirements of each layer, shown in Figure 4, are detailed in the following sections.

4.1 High Level Intelligence Layer

This layer implements the algorithms that allow the robot to be aware of its environment and enables the robot to make intelligent decisions about how to interact with and to traverse its environment. It is at this layer where external direction, such as requests from a battlefield commander, are accepted and integrated into the decision making process.

The capabilities implemented within the High Level Intelligence Layer are, in general, very computationally demanding. These tasks require hardware that can supply the required computing power. Such hardware doesn't have to interact with external devices but must be computationally powerful. It could be a general purpose processor, a Digital Signal Processing (DSP) chip, or a specially programmed FPGA. To supply this computational power, multiple processors may be necessary and thus the hardware at this layer must be amenable to multi-processor implementations.

Real-time responses are not required at this layer and thus a general purpose multi-tasking operating system would be suitable for use.

4.2 Intermediate Intelligence Layer

The Intermediate Intelligence Layer implements the supervisory functions that ensure the functional operation of the robot. It receives directives from the High Level Intelligence Layer, converts these high level directives to commands that drive the Control Layer.

The processor used at the Intermediate Intelligence Layer should feature moderate processing power and have relatively small in physical size. The processor must also have a low power consumption and support operation from a battery source. The operating system at this level must be responsive but hard real-time capabilities are not required. Soft real-time capabilities at this layer would suffice.

4.3 Control Layer

The Control Layer handles interactions with the physical world. It interfaces with a host of external devices. Some of these external device may include:

- Proximity sensors (e.g. instrumented bumpers).
- Odometry devices (e.g. encoders connected to drive shafts).
- Accelerometers (e.g. to measure vibrations to quantify trafficability parameters).

- Inertial measurement units (e.g. to monitor heading).
- Inclinometers (e.g. determine vehicle stability).
- Numerous motors of various types.

At this level the processor should possess abundant capabilities to both interface with and to control peripheral devices. Capabilities such as analog to digital conversion (A/D), pulse width modulation (PWM), input/output (I/O), time processing are a mandatory requirement. Given that the processor at this level must deal with external hardware in the real world, a real time operating system is also a necessity, with hard real-time response times in the sub milli-second range.

4.4 Hardware Architecture

While the Standard Robot Model concerns itself with only the functional hierarchy of a robot, the RNS defines an electronic hardware architecture. This hardware architecture provides the appropriate electronics that allow the robot to be constructed in a flexible and scalable manner. The electronic hardware architecture of the RNS provides this flexibility by isolating the software from the physical details of each UGV platform. This hardware architecture is distributed by design, thus allowing additional capabilities to be added as required. Figure 5 illustrates a potential implementation of this hardware architecture.

Distributed processing is a key aspect of this architecture. It allows the computationally intensive High Level Intelligence Layer to be distributed across multiple processors. The intelligence implemented at the Intermediate Layer should require a distributed approach, but if multiple processors are required, the architecture does not preclude it. Significantly, the Control Layer, often controlling multiple physically remote devices, may also be distributed across multiple processors. The hardware associated with each of the layers is described in more detail in the following sections.

4.4.1 Computational Element

This element is comprised of hardware that is tailored for computationally intensive algorithms. The algorithms implemented at the High Level Intelligence Layer can be: world representations, path planning, and learning are either iterative in nature, or process large volumes of external data. They all consume large amounts of processing power and tend to be scalable. Hence if more processing power is available it can be consumed by simple configuration changes such as resolutions or depth of searches. The hardware that is well suited for this type of computation burden includes:

- The latest release of Intel general purpose processors[5]. While the Intel processors are fast and simplify the implementation of software they are

[5]Currently (2004) the Pentium 4 series at speeds of over 3 GHz.

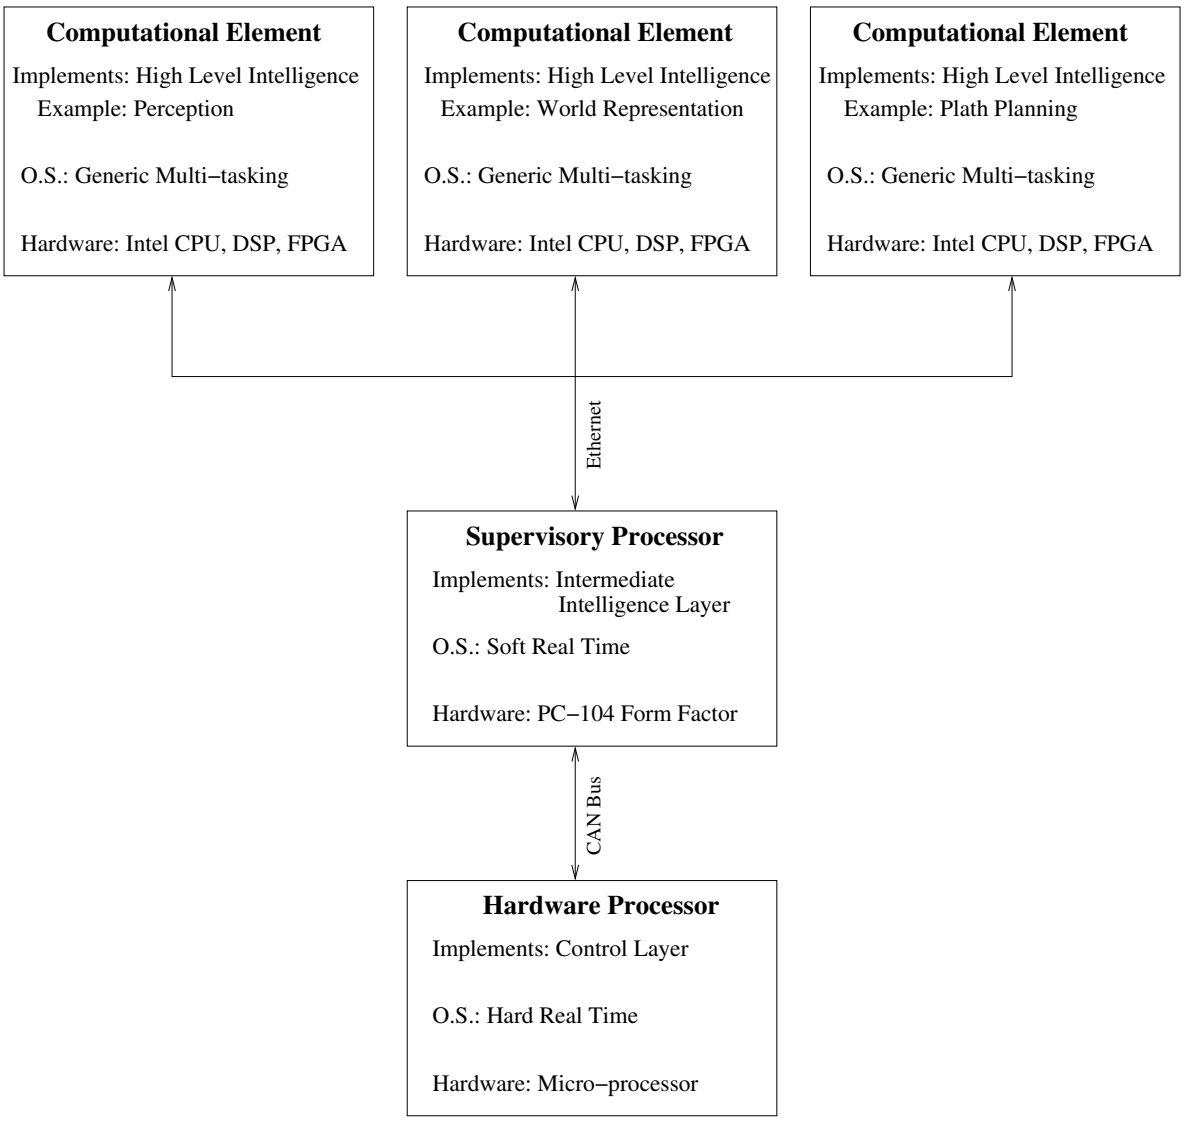

Figure 5: *Robot Nervous System Hardware Architecture*

power hungry and dissipate approximately 100W of heat. When a complete system is considered an Intel Pentium 4 based system can draw in the realm of 300W power. While this is not an issue for some UGV platforms, for others this large power draw is an issue, especially if multiple processors are used. On the plus side, hardware and software that supports multiple processors on a single motherboard are readily available and at a very attractive price level. One final limitation of the Intel architecture is the large physical size of implementations based on this processor series.

- The new PowerPC 970 processor, from IBM, is expected to attain performance levels close to that achieved by the latest Pentium processors with the advantage that the PowerPC traditionally requires significantly less power than the Intel processors[6]. It is expected that this processor will run the Linux O.S. and that they will be released in symetric multi-processor (SMP) configurations for enhanced computational power. This processor is currently available in high-end Apple computers.

- Digital signal processors (DSP) could also be considered for use as a computational element. The main advantage of a DSP is that it's special instruction set and architecture delivers more processing power per watt consumed than any conventional general purpose processor. In general the fastest DSPs deliver a computational capability that is close to the processing power available from the fastest Intel products.

- Field programmable gate arrays (FPGA) or programmable array logics (PAL) could also be considered for specific specialized algorithms. They allow an algorithm to be implemented in hardware logic and thus can compute in parallel at the propagation speed through the electronic device. This may allow a potential order of magnitude in processor power improvement over the fastest general purpose processors available and consume only a fraction of the power[7]. The challenge this approach poses is the implementation of complicated algorithms in a manner that can be executed by logic level devices.

As a general rule a processor used as a computational element would run a multi-tasking operating system. DSPs and FPGAs/PALs would be an exception to this rule.

4.4.2 Supervisory Processor

The supervisory processor is the hardware that implements the Intermediate Intelligence Layer. While a given UGV may have none to multiple

[6]The 970 PowerPC at 1.8 GHz (2003) requires 42 W, while the 2.8 GHz Pentium consumes 68.4 W. This represents roughly a 40% to 50% reduction in power consumption.

[7]An FPGA or PAL also has a very small physical footprint.

computational elements[8], every UGV must have at least one supervisory processor. Since the supervisory process must be present on all UGV's, it must have a physical form factor and power requirements that are acceptable for all UGV platforms. The RNS defines a PC-104+ form factor board as the supervisory processor with a preference for the board utilizing a CPU that is power efficient. Current UGV implementations have a choice of two different PC-104+ boards: a PC-104+ board based upon the PowerPC 405GPr processor from IBM; and a board based upon the Transmeta Crusoe processor. Both of the boards feature a small form factor and low power consumption. The PowerPC 405GPr board from Embedded Planet that has an overall power requirement of 7.5 W and a clock rate of up to 400 MHz[9]. The Transmeta Crusoe board from IBT Technologies utilizes the TM5400 processor at a 533 MHz clock rate with processor requiring up to 7.3 watts for power.

- The IBM[10] embedded processors target the high performance consumer market including: high-end PDAs and internet appliances; video games and multimedia edutainment systems requiring realistic imaging, motion and sound; photo-quality scanning and printing; digital imaging; voice recognition; and data communications [18]. These applications are characterized by high performance floating-point processing requirements. The embedded PowerPC line will benefit from the forces that are driving the expansion of this new market. The embedded PowerPC market includes milliwatt sensitive applications to gigahertz applications and, because of this large range, has been designed to be a scalable architecture that allows for the easy porting of code across this range of processors. The 405GRr hosts the following features:
  - 333 MHz clock speeds with a 400 MHz implementation available.
  - On-board 100 Mbit Ethernet transceiver.
  - Two on-board USB ports (V1.1).
  - PC-104 Plus interface, allowing access to PCI based cards.
  - On-board serial peripheral interface (SPI) interface.
- The Transmeta Crusoe processors are targeted at compact designs that are power sensitive. A key aspect of the Crusoe processor is its x86 software compatibility which run all x86 binary applications. Crusoe processors are available in frequency ranges from 500 MHz to 1 GHz. The IBT PC-104+ board features all the standard components that are associated with a standard desktop computer, such as:

[8] It is conceivable that on a simple UGV that no computational elements are required and all High Level Intelligence Layer functions are implemented on the supervisory processor.

[9] The author notes that while clock speed is not a definitive measure computational capabilities it does tend to be indicative of the overall processing capabilities of the processor.

[10] IBM and Motorola share the PowerPC architecture and both offer compatible implementations.

- 10/100 Mbit Ethernet Controller.
- VGA graphics.
- 2 Communications ports, a parallel port and a floppy drive.
- Real-time clock
- Keyboard controller
- IDE hard drive interface
- 1 USB port
- Watch dog timer

The software executing on the supervisory processor requires, at a minimum, soft real-time response characteristics.

4.4.3 Hardware Processor

The hardware processor must have inherent abilities to interface with and control a wide variety of peripheral devices. It must also be extensible thus allowing it to easily incorporate additional interface capabilities. The hardware processor should also have a significant amount of processing power allowing it to implement signal conditioning as well a low level control strategies.

The micro-controller world exhibits a huge variety of implementations ranging from small 8-bit processors to 32-bit RISC implementations with on-board FPU's. The dominate supplier of micro-controllers is Motorola, who is especially strong in the 32-bit micro-controller field, with an extensive offering for both the automotive and communications industries. While Motorola is not the only vendor of micro-controllers, its market share, high end offerings and DRDC Suffield's extensive use of Motorola processors make it a prime candidate for the hardware processor. The microprocessor selected as the hardware processor is the MPC555 implemented on an Intec Automation Ltd. credit card sized board. While the current MPC555 runs at a clock speed of 40 MHz, the Motorola road map calls for the release of MPC5550 with a clock speed of 400 MHz and DSP capabilities by late 2003 [19]. It should be noted that this embedded processor line is driven by the requirements of numerous industries such as: the automotive industry, industrial controls, and avionics [20]. This large consumer base should translate into a continuing stream of innovative offerings as can be seen by the MPC5500 next generation embedded controller.

The MPC555 micro-controller from Motorola has a 32-bit PowerPC core and is targeted directly at high performance embedded control applications. It features inherent capabilities to control and interface with numerous types of external devices. The capabilities of the MPC555 are summarized as follows:

- 40MHz PowerPC Core with FPU
- 8 Channels of PWM
- 10 Dual Action Timer pins
- 32 Channels of A/D (up to 100KHz sampling rate)
- 2 Time Processing Engines
- 2 RS-232 Ports
- 2 CAN Ports
- 16 General Purpose I/O Lines
- Queued Serial Peripheral Interface (QSPI)

The capabilities listed above give the MPC555 the ability to interface with numerous devices. The flexibility of the MPC555 makes it a very suitable hardware processor. The QSPI allows the MPC555 to easily expand its capabilities allowing for the addition of extra RS-232 ports or finer resolution A/D converters.

The MPC555 implementation from Intec Automation Ltd. does not feature a real-time operating system. While a real-time operating system is not required, there are many benefits accrued through the use of a real-time operating system. The MPC555 supports numerous real-time operating systems, ranging from commercial implementations to open source offerings.

4.4.4 Communications

The RNS defines a distributed processing architecture. A non-centralized distribution of processors requires communications hardware to allow the various processors to communicate with each other. The RNS will use the 100 MBit Ethernet interface as the medium of communications between the Computational Elements and the Supervisory Processor. The Ethernet standard is nearly universally accepted and offers excellent performance levels including an easy migration path to a 1 Gbit interface.

The communications standard for networking the Supervisory Processor to the Hardware Processor is the CAN bus. The CAN bus is a robust multi-master standard that is commonly used by the automotive industry. It features multi-drop capabilities at data rates up to 1 Mbit/second.

5. Conclusions

This report has introduced the “Robot Nervous System” which is a distributed, flexible electronic hardware architecture that is portable across a variety of unmanned ground vehicles. The ALS program is developing UGVs that vary greatly in size and capability. To simplify development, an electronic hardware architecture is required that has the flexibility to easily support new platform implementations. The RNS, using a distributed computing architecture, assigns the most suitable processor(s) to the required robot functionality and delivers the flexibility required by the ALS project. The RNS is flexible and scalable allowing the same architecture to control a variety of UGVs from small portable mobiles robots to large autonomous military vehicles.

References

1. Kelly, A. J. (1995). Predictive Control Approach to the High-Speed Cross-Country Autonomous Navigation Problem. Ph.D. thesis. Carnegie Mellon University. Pittsburg, Pa.

2. Dudek, G. and Jenkin, M. (2000). Computational Principles of Mobile Robotics, Cambridge, UK: Cambridge University Press. p. 10.

3. Nilsson, N. (1969). A mobile automation: An application of artificial intelligence techniques. In *Proc. IJCAI*. Reprinted in Autonomous Mobile Robots: Control, Planning and Architecture. Vol.2 233-239.

4. Brooks, Rodney A. (1986). A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, *RA-2*(1), 14–23.

5. SPIE Aerosense Conference (2003). CLARAty: An Architecture for Reusable Robotic Software, SPIE Aerosense Conference. Orlando, Florida.

6. Broten, G., Erickson, D., Giesbrecht, J., Monckton, S., and Verret, S. (2003). Engineering Review of ANCAEUS/AVATAR An Enabling Technology for the Autonomous Land Systems Program?. (DRES TR-2003-167). DRDC Suffield.

7. Mathais, E. (2002). Canadian Light Source Control System Technical Specifications. (CLS Technical Report 7.439.1, Rev. 2). Canadian Light Source. p. 11.

8. Walter, W.G. (1950). An Imitation of Life. *Scientific American*, *182*(5), 42–45.

9. Walter, W.G. (1950). A Machine that Learns. *Scientific American*, *185*(2), 60–63.

10. Winograd, T. (1972). Understanding Natural Language, Academic Press, New York.

11. Nilsson, Nils J. (1984). Shakey The Robot. (Technical Report 323). AI Center, SRI International. 333 Ravenswood Ave., Menlo Park, CA 94025.

12. Brooks, R.A. (1989). A Robot that Walks: Emergent Behaviours from a Carefully Evolved Network, Artificial Intelligence at MIT, Ch. 24, pp. 28–39. The MIT Press.

13. Arkin, R.C. (1989). Motor Schema Based Mobile Robot Navigation. *International Journal of Robotics Research*, *8*(4), 92–112.

14. Hasslacher, B. and Tilden, M. (1995). Living Machines, Elsivier Publishers.

15. Gowdy, J. (2000). Emergent Architectures: A Case Study for Outdoor Mobile Robots. Ph.D. thesis. Carnegie Mellon University.

16. Bar-Cohen, Y. and Breaszeal, C. (2003). Biologically Inspired Intelligent Robots, Bellingham, WA: SPIE Press. p. 5.

17. Curtis, B., Jacobson, S., and Marcus, E. (1972). An Introduction to the Neurosciences, Philadelphia, Pa: W. B. Saunders Company.

18. Allison, A. (2002). PowerPC: High-Performance Embedded Processors of Choice. Independent Industry Analyst.

19. Motorola (2002). MPC500 Family Road Map - External. Motorola Literature.

20. Bettelhiem, R. (2002). The MPC500 Family of 32-bit Embedded Controllers from Motorola. Motorola Literature.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence R&D Canada – Suffield
PO Box 4000, Medicine Hat, AB, Canada T1A 8K6

2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable).

UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title).

Unmanned Ground Vehicle Electronic Hardware Architecture

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Broten, G. ; Monckton, S.

5. DATE OF PUBLICATION (month and year of publication of document)

June 2004

6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc).

30

6b. NO. OF REFS (total cited in document)

20

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered).

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address).

Defence R&D Canada – Suffield
PO Box 4000, Medicine Hat, AB, Canada T1A 8K6

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Specify whether project or grant).

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written).

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique.)

DRDC Suffield TM 2004-122

10b. OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution beyond the audience specified in (11) is possible, a wider announcement audience may be selected).

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

The Autonomous Land Systems (ALS) initiative has been tasked with researching and developing innovative autonomous vehicles that will assist the Canadian Forces in performing their duties in the 21st century. This research will continue for many years, on many different types of unmanned ground vehicles, and thus requires a flexible and scalable robot electronic hardware architecture. This architecture must be applicable to vehicles ranging from small indoor research platforms to large outdoor military vehicles such as the LAVIII. This report describes the hardware architecture, known as the Robot Nervous System, which will serve as a backbone for current and future Autonomous Land Systems development. The architecture is distributed in nature and thus strongly supports the scalable requirements of the R&D program. Presently the architecture uses three scales of processors, each well suited for their specified tasks. The processors are linked via appropriate network topologies that enable each processor type to be leveraged to its maximum capabilities. This hardware architecture is currently being implemented in a variety of autonomous vehicles that are under development at DRDC Suffield.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title).

Unmanned ground vehicles,
robotics,
autonomous vehicles,
electronic hardware architecture,
robot nervous system