

2006

## Voice Activated Command and Control with Speech Recognition over WiFi

Brian Nolan

*Technological University Dublin*, [brian.nolan@tudublin.ie](mailto:brian.nolan@tudublin.ie)

Tony Ayres

*Technological University Dublin*

Follow this and additional works at: <https://arrow.tudublin.ie/itbinfoart>



Part of the [Engineering Commons](#)

---

### Recommended Citation

T. Ayres & B. Nolan (2006) Voice Activated Command and Control with Speech Recognition over WiFi, *Science of Computer Programming* 59 (2006) 109–126 <http://dx.doi.org/10.1016/j.scico.2005.07.007>

This Article is brought to you for free and open access by the Computational Functional Linguistics at ARROW@TU Dublin. It has been accepted for inclusion in Articles by an authorized administrator of ARROW@TU Dublin. For more information, please contact [arrow.admin@tudublin.ie](mailto:arrow.admin@tudublin.ie), [aisling.coyne@tudublin.ie](mailto:aisling.coyne@tudublin.ie), [vera.kilshaw@tudublin.ie](mailto:vera.kilshaw@tudublin.ie).



# Voice activated command and control with speech recognition over WiFi

Tony Ayres\*, Brian Nolan

*Institute of Technology Blanchardstown, Ireland*

Received 4 October 2004; received in revised form 22 February 2005; accepted 21 March 2005  
Available online 15 August 2005

---

## Abstract

This paper presents work conducted to date on the development of a voice activated command and control framework specifically for the control of remote devices in a ubiquitous computing environment. The prototype device is a Java controlled Lego Mindstorm robot. The research considers three different scenario configurations. A recognition grammar for command and control of the robot has been created and implemented in Java, in part in the recognition engine and in part on the robot. The physical topology involves Java at each node endpoint, that is, at the handheld PC (iPaq), the PC workstation, the Linux server and onboard the robot (including its Java based Lejos OS). Network communications is primarily WLAN with an element of IR where the robot is concerned. The speech recognition software used includes Sphinx4, Microsoft SAPI and the Java Speech API. We compare these speech technologies and present their benefits in the context of this research. For each given scenario we present and discuss the implementation challenges encountered and their corresponding solutions, including future plans to create additional grammars to extend the framework's range of devices.

© 2005 Elsevier B.V. All rights reserved.

*Keywords:* Speech recognition; Ubiquitous computing; Java; Command and control; Wireless networks

---

---

\* Corresponding author.

*E-mail addresses:* [tony.ayres@itb.ie](mailto:tony.ayres@itb.ie) (T. Ayres), [brian.nolan@itb.ie](mailto:brian.nolan@itb.ie) (B. Nolan).

## 1. Introduction

Research conducted by Gartner Dataquest in 2001 [3], predicted that by 2005 there would be 780 million mobile phones and PDAs being shipped worldwide. Given these statistics, it is reasonable to assume that both mobile phones and Personal Digital Assistants play an important part in many people's daily lives. The ever growing power of these mobile devices, coupled with the provision of wireless and Bluetooth networking capabilities, opens the door for a host of new applications to be developed and for old application paradigms to be applied on new frontiers. With the evolution of the Java platform and language, Java has become suitable for developing speech applications which can command and control mobile devices. The Java Speech API, specifically the Cloudgarden implementation, provides speech recognition and synthesis tools for Java applications. The Java Speech Markup Language (JSML), which is a subset of XML [17], and the Java Speech Grammar Format (JSGF) [16] offer simple but powerful development structures for speech synthesis development and recognition grammars respectively.

This research project is concerned with building a framework with applications for command and control of a remote device by voice activation with speech recognition from a local control station over a wireless network in three different scenario configurations. In each scenario Java plays a critical role. A key goal is to fit the framework within the pervasive computing paradigm, that is to leverage portable devices such as laptop computers, smart phones and personal digital assistants as control clients. Voice control of devices ranging through robotics, software systems, home appliances and in-vehicle systems (radio, mobile phones etc.) are possible applications resulting from this research.

The first of these scenarios or configurations is based on a PC workstation under the Windows operating system connected via a wireless network to a PC-based server. The server issues commands to a remote device. The second scenario involves developing a distributed speech recognition engine between an iPaq Pocket PC and a PC based server which will issue the commands to the remote device. The third scenario involves a mobile device, specifically an iPaq Pocket PC, that will connect over a wireless network to a PC-based server. The server will, again, issue commands to the remote device.

For the purposes of this research project the remote device is a Java controlled Lego MindStorms robot that moves and can undertake certain actions under instructions relayed to it over a wireless interface. The robot could be replaced with practically any computing or electronic device which has a Java Virtual Machine installed. As part of this research it is our intention to develop a distributed speech recognition framework based on the Java programming language.

## 2. Technology review

### 2.1. *Applications of speech technology*

The applications for speech recognition can be grouped into three distinct categories; these are command and control, dictation and authentication.

Command and Control applications are concerned with providing the user of these systems the means to control items within their environment with voice commands

appropriate to the domain. The appliance of command and control technology may manifest itself in the control of user interface menus in personal computing desktop applications or the control of large scale mechanical or electronic and computing devices.

Dictation applications allow the user to speak to the system and have it generate a transcript of what has been said. This is particularly useful in legal or medical areas where information is recorded in real time and making written notes would be too slow. Specialized dictation grammars exist for application domains such as this. Furthermore dictation and command/control functionality can be combined in word processing applications such as Microsoft Word.

Speech technology can also be used for authentication purposes as part of a security system. The signal analysis algorithms employed as part of a speech recognition front end generate a feature vector which can be matched to a pre recorded sample of a users voice. Given that each person has a unique voice print this can be used for authentication purposes.

## *2.2. Types of speech recognition systems*

Speech recognition systems can be classified according to whether they are speaker dependent or speaker independent [1,12]. Speaker dependent speech recognition engines require the user to train a profile of their voice for the engine to use when performing recognition. This process typically involves reading sample passages of text to the engine. Conversely, speaker independent systems do not require the user to train them before achieving high recognition accuracy. In this instance a pre-recorded corpus of words is compared to the input speech vectors to generate recognition result.

In general terms speaker dependent systems achieve greater recognition accuracy given that the engine will be customized for a specific user's voice, however speaker independent systems can achieve comparable accuracy levels where the grammar is constrained to a specific application domain and, as such, are ideally suited for applications which may have a large number of different users and where training a speaker dependent engine is not feasible.

## *2.3. Process of speech recognition*

The components of speech recognition systems include a speech corpus (database), a frontend processing system and a speech decoding unit. The frontend is responsible for analyzing the speech input and extracting feature vectors which will be used in the decoding process (Fig. 1). The speech corpus, in the case of speaker independent recognition engines will contain acoustic information for all the words and phonemes which the corpus contains. The speech decoding process compares input features generated by the frontend with those in the corpus, the result is a probability score, representing the engine's confidence in the accuracy of any match found.

Modern speech recognition systems use stochastic techniques to model and decode speech signal data. Hidden Markov Models (HMM) [7,5] have become the most successful statistical method for speech recognition. The HMM phase of speech recognition comes after an initial analysis and feature extraction process on the incoming speech signal. The

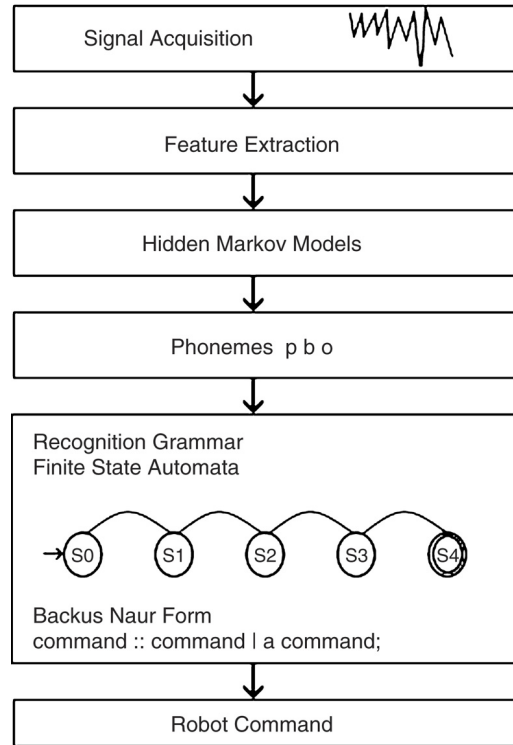


Fig. 1. Speech recognition process.

feature extraction process generates feature vectors which are used as the input to the HMM.

A Hidden Markov Model is specified by the states  $Q$ , the set of transition probabilities  $A$ , defined start and end states and a set of observation likelihoods  $B$ . A Hidden Markov Model formally differs from a Markov model by adding two other requirements. Firstly it has a set of observation symbols which is not drawn from the same alphabet as the state  $Q$ . Secondly the observation likelihood function  $B$  is not limited to the values 1 and 0, in the HMM the probability can take any value between 0 and 1.

The parameters needed to define a HMM are as follows:

- A set of states
- Transition probabilities
- Observation likelihoods
- Initial distribution
- Accepting states.

In order to extract a suitable output for speech recognition we must parse the representation which the Markov model contains.

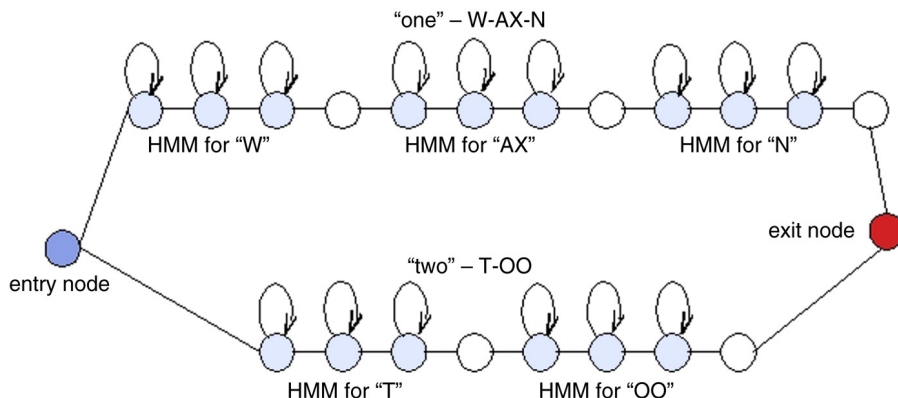


Fig. 2. Hidden Markov Model.

Fig. 2 shows the HMM graph for the words ‘one’ and ‘two’ according to the Sphinx4 speech decoder. HMMs are generated for each phoneme that constitutes a word. Each HMM has a transition from to various nodes in the graph and to itself. The Viterbi [8] algorithm is used to find the best path through the graph based on the highest score of each transition.

#### 2.4. Selecting a speech engine

The Java Speech API [15] is used to provide a platform independent speech recognition and synthesis interface for Java applications. Sun Microsystems supply a reference standard for the Java Speech API, but an implementation is not provided. The Java Speech API implements no speech processing functionality of its own but allows Java applications to plug into functionality available on the host operating system.

Recent advances in signal processing algorithms coupled with the development of HMM based decoding techniques, have led to the development of many highly accurate speech recognition engines. The majority of these speech engines are commercial products which include text to speech capabilities in tandem with their recognition functionality. These products include Microsoft Speech API version 5 [11] (SAPI5), IBM Via Voice [6] and Dragon Naturally Speaking [13]. One defining characteristic of all these engines is that they are speaker dependent.

In the open source domain the Sphinx project is the only suitably developed candidate for our purposes. The Sphinx project is concerned with developing HMM based speaker independent recognition systems. The project has 3 engines available as source code downloads, namely Sphinx2, Sphinx3 and Sphinx4.

Sphinx2 is a real time speech decoder and its features include continuous speech decoding, a single best or several alternative recognitions, and support for bigram, trigram, or finite-state grammar language models. Sphinx3 is a state of the art speech decoder written in C. While it has a slower decoding speed than Sphinx 2, it provides more accurate recognition. Initially it was developed to perform batch speech decoding from audio files but is now capable of live decoding.

Sphinx4 is the latest speech decoder to be released by the project. Initially it started as a port of Sphinx3 to the Java programming language; however the engine evolved to become more flexible than Sphinx3. A defining characteristic of Sphinx4 is the configuration of the engine which is achieved through an XML configuration file. Each Java object in the Sphinx4 system can be instantiated through this file. This helps to keep the application code clear of the Sphinx4 code which makes for easier debugging. Sphinx4 also includes an implementation of the Java Speech API. Its object oriented structure and easy XML configuration make it an ideal choice for conducting research.

This research project defines a number of application scenarios. A different speech engine is implemented in each of the scenarios. Scenario 1 uses Microsoft SAPI5 with the Cloudgarden JSAPI [9] implementation; Scenario 2 uses the Sphinx4 speech recognition engine which includes its own JSAPI implementation while Scenario 3, which is still in development, will use a lightweight Java/C++ based speech decoder based on the Sphinx engine.

We tested the performance of the Java Speech API with Sphinx4 and SAPI5 in the context of developing this framework for command and control. The test provided an insight into the characteristics of speaker dependent and speaker independent recognition engines and presented an opportunity to compare the current state of the art of both approaches to speech recognition.

### 3. Performance analysis

#### 3.1. Performance analysis of the Java speech API

Under Windows XP, the set-up time for a JSAPI recognizer under both speech engines is in a similar range, although SAPI5 is marginally faster. Most notable is the length of time Sphinx takes to process the first command; after this initial command has been processed the time drops back to just over 3 s for each subsequent command. The SAPI5 configuration is extremely quick for all commands; although SAPI could be prone to a high number of errors, this occurred under JDK1.4 where the average error rate was 1.5 (accuracy of 63%). The graph shown in Fig. 2 highlights the speed of the SAPI engine in processing the commands.

Fig. 3 also shows that Java 1.5 is faster than 1.4; this is most noticeable with the Sphinx4 test. While the set-up times are almost identical, the recognition time drops to around 2 s in comparison to 3 s with 1.4. The SAPI5 test also ran faster with JDK1.5, although the difference is marginal when compared to the performance gain Sphinx4 achieves.

Sphinx encountered a trough in recognition accuracy with the Pentium IV/1.4 configuration, with at least one error occurring in the majority of tests, thus yielding an average of 1.4 errors which translates to an accuracy of 65%. In addition, the recognizer set-up time was much longer than those on the Windows XP test machine. With JDK 1.5 Sphinx4 yielded a recognition accuracy of 93%, with only 0.2 errors.

One noticeable difference is the recognition times under JDK 1.5; we encountered numerous spikes where decoding of commands took between 8 and 30 s, this occurred on three occasions while issuing the 'Backward' command. As a result the average command time for JDK1.5 on the Windows 2000 machine is slower than JDK1.4. In Fig. 4 the graph

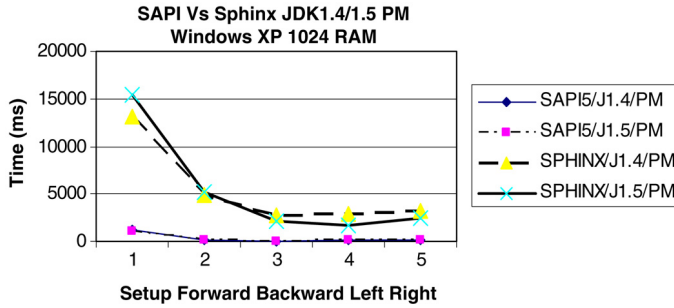


Fig. 3. SAPI vs Sphinx Windows XP.

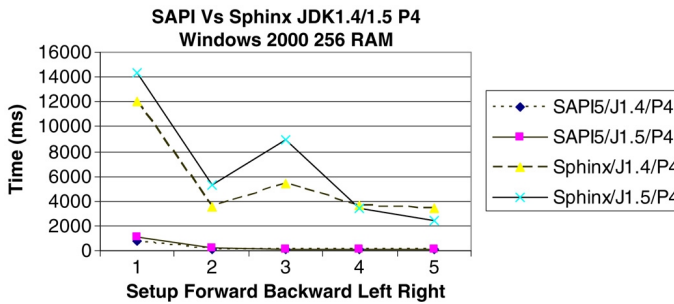


Fig. 4. SAPI vs Sphinx on Windows 2000.

shows the sharp peaks and troughs in the Sphinx performance under JDK1.5. The SAPI5 results on this machine were in line with the previous test; recognition time averages below 500 ms and the set-up time averages around 1 s. As with the Sphinx tests on this machine, the SAPI performance under JDK1.5 was not faster than JDK1.4.2.

Fig. 5 presents a bar graph indicating the recognition accuracy of the test configuration. SAPI5 under Windows 2000 under JDK1.5 was the most accurate with a score of 98%, Sphinx4 on the same system and the same JDK received the second highest score with 95%. SAPI also holds the lowest score of 63% on Windows XP with JDK1.4. The average accuracy for both engines was 84%. Both engines had one trough of poor performance where the accuracy dropped considerably and as a result both engines received the same average accuracy percentage. If we ignore the trough values, then SAPI5 achieves the highest average score of 91% with Sphinx scoring 89%.

We tested the Sphinx4 engine under Mandrake Linux 10.1 running the release version of JDK 1.5. This allows us to compare the performance of the Java runtime environment on a second operating environment for the purposes of speech recognition.

Furthermore, Sphinx4 is one of the few large scale speech recognition engines available for the Linux operating system. In comparison to the Windows based tests the Linux test (Fig. 6) achieved the greatest accuracy with a score of 97%. Again the initial engine set-up time is longer than SAPI5/Windows, however it does perform better than Sphinx4 on Windows in both recognition processing time and accuracy.



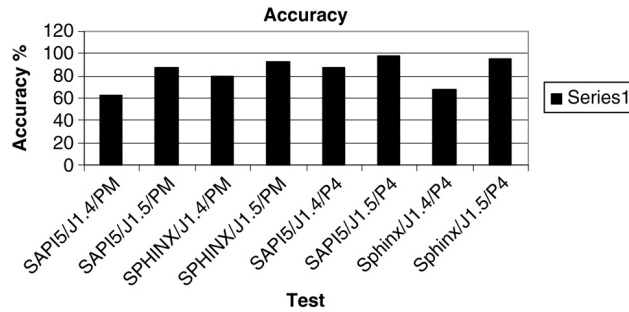


Fig. 5. Accuracy percentage.

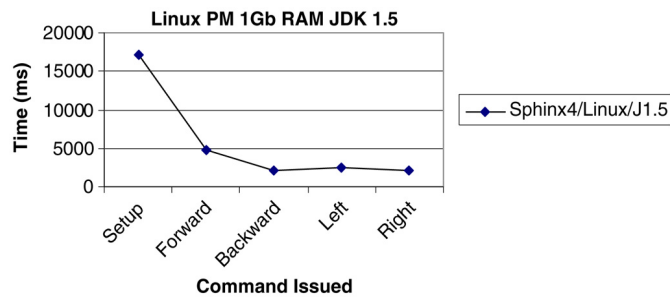


Fig. 6. Linux performance results.

The results of these tests show the relative strengths of both speech engines and performance-wise there is little to separate them. There are, however, significant functional differences between them. SAPI5, as a speaker dependent engine using a speech profile to decode speech input, has the ability to recognize any speech utterance from the user who trained the profile. Sphinx4 is limited to recognizing only those words which are contained within the language model. This limitation in Sphinx4 makes it less suitable for dictation applications, but its high accuracy and open source code make it excellent for conducting research and building command and control applications. SAPI5 is an excellent engine for both dictation and command and control, but it is less adaptable for research purposes as it a closed system.

## 4. Application topology

### 4.1. Scenario one

To date we have been developing the overall system architecture and evaluating speech technologies which will meet our requirements. Fig. 7 illustrates the architecture of scenario one. In scenario one we have developed a speech recognition application in Java using the Cloudgarden implementation of the Java Speech API. Cloudgarden provides JSAPI functionality in conjunction with a SAPI 4/5 compliant speech engine on the host

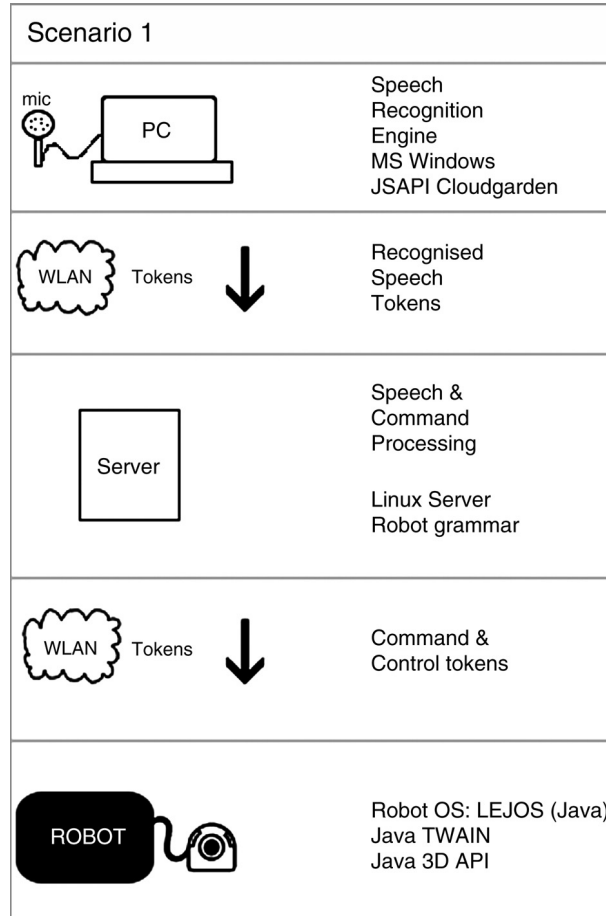


Fig. 7. Scenario one.

operating system. Supported SAPI speech engines include Microsoft SAPI, IBM Via Voice and Dragon Naturally Speaking.

The client application accepts speech input via microphone and performs speech recognition using the JSAPI. A successful output from the speech recognition process is one or more string tokens. The application delivers the recognized speech strings to a Java based server application over TCP/IP via a wireless network. The server application is running on a Linux operating system. The server processes the recognized speech strings in order to determine if they match terms in the robot control grammar. If a match is found, that particular command is issued to the robot over a wireless network.

The robot will eventually be fitted with a camera and we are investigating the possibility of using Java API's including the Java Twain and the Java 3D API to capture and process images returned to the client application from the robot.

Table 1  
iPaq Pocket PC hardware configuration

HP iPaq Pocket PC 5550	
Processor	Intel XScale 400 MHz
Memory	128 MB RAM
ROM	48 MB ROM
Networking	802.11 b WLAN and Bluetooth
Operating system	Pocket PC 2003

In the next section we describe scenario two which involves a distributed speech recognition engine between an iPaq Pocket PC and a PC based server which will issue the commands to the remote device.

#### 4.2. Scenario two

Scenario two (Fig. 8) offers an alternative approach. The PC workstation is replaced by a PDA with wireless networking capabilities, in this case an HP iPaq 5550 running Pocket PC 2003. The application on the PDA is concerned with capturing a speech input signal and sending it through the wireless interface to the server. The server hosts a processing engine for the robot grammar, as in scenario one. However it also has an additional layer of functionality, namely, the speech recognition engine. The PDA provides mobility but has limited processing power.

The configuration of the iPaq Pocket PC is shown in Table 1.

In order to minimise the load on the PDA the speech recognition engine is distributed between the server and PDA. The PDA has the Jeode runtime environment installed on it. Jeode is a Personal Java compatible runtime environment. Personal Java [14] is fully compatible with JDK 1.1, however JDK 1.1 does not have sufficient libraries to perform signal capture or speech recognition. To circumvent this obstacle, the Java Native Interface will be used to plug into the sound capture functionality of the PDA hardware from which a digitized speech signal is returned. This digitized speech is sent to the server application where the remainder of the speech recognition process takes place. Processing of the speech signal is the most processor intensive activity in the recognition process. This function is therefore assigned to the server in this scenario and thus memory and processing power on the mobile device become of less importance. This approach presents greater complexity in acquiring the signal, sending it over the network and reading it back into the speech processing engine.

#### 4.3. Scenario three

In scenario three (Fig. 9), we port the speech recognition engine from the PC workstation in scenario one to the PDA. Given the resources of the PDA, a simple copy of the speech engine from scenario one is not feasible.

The Cloudgarden JSAPI implementation requires a SAPI compliant engine and Java Development Kit 1.3 or better and as these are not available on the PDA an alternative is needed. The open source Sphinx 2 can be used to provide speech recognition functionality

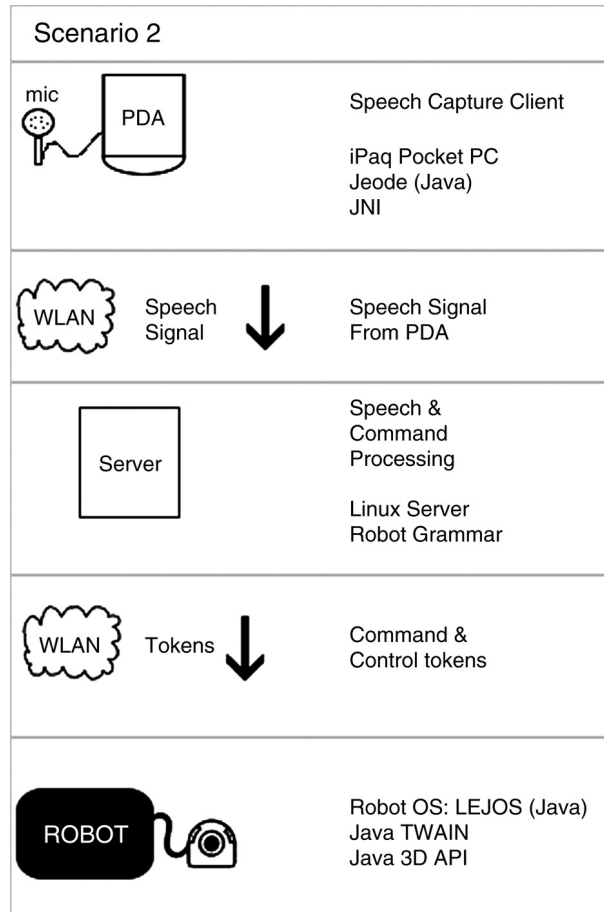


Fig. 8. Scenario two.

on the device. Sphinx2 is written in C, therefore the Java client application will use the Java Native Interface to plug into the methods which are provided by Sphinx. The recognized speech will be sent over a WLAN link (TCP/IP) to the server application.

As with scenario one the server processes the recognized speech strings in order to determine if they match terms in the robot control grammar. If a match is found, the command is issued to the robot over a wireless network.

#### 4.4. Lego Mindstorm and Lejos

The prototype Lego Mindstorm robot is shown in [Photo 1](#). The Lego RCX Microcontroller (yellow box (white box in print version) in [Photo 1](#)) is the main computing device. It contains 32 kB of RAM and inputs for sensors and motors. Communication with the robot is achieved through the infrared (IR) tower which is attached to a personal computer running Windows.

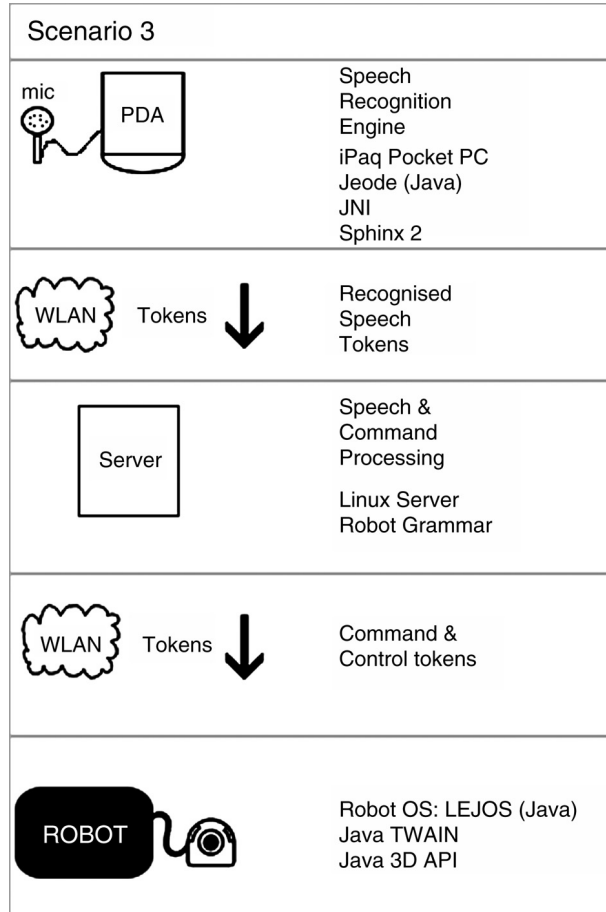


Fig. 9. Scenario three.




Photo 1. Lego Mindstorm Robot.

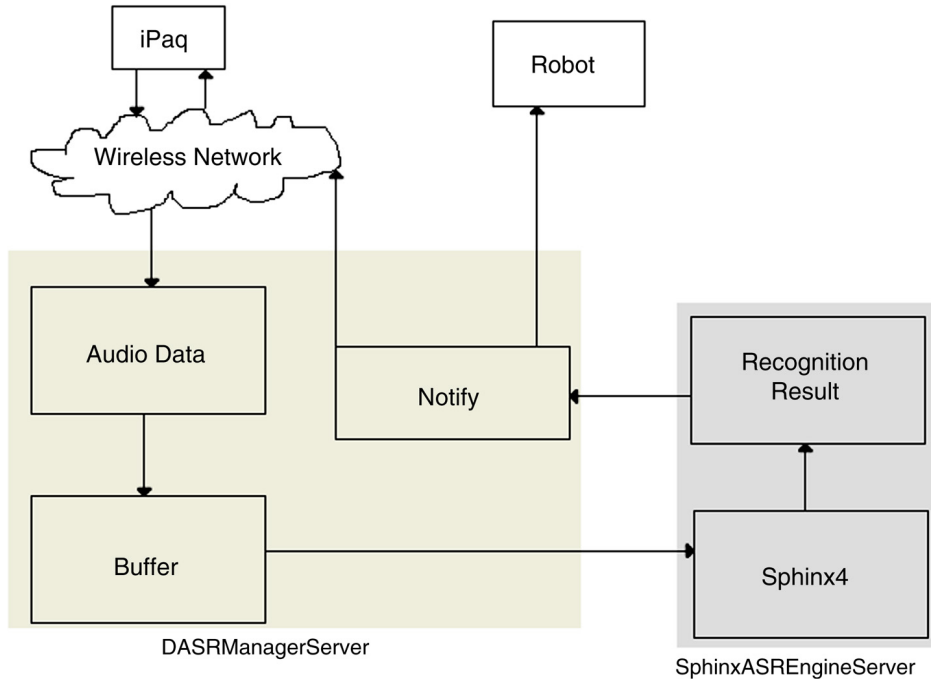


Fig. 10. Distributed speech recognition under scenario two.

We have replaced the pre-installed Lego operating system with Lejos [10]. Lejos is an operating environment for the Lego RCX which allows application programs to be written in Java. This allows for greater granular control of the robot than is available through use of the Lego operating system. With Lejos installed on the RCX, 14 kB of RAM remains for both storing and running programs.

## 5. Design

### 5.1. Distributed speech recognition under scenario 2

Scenario two requires a distributed speech recognition application topology. The architecture of this design is shown in Fig. 10.

The process begins with the client application installed on the iPaq. Presently the client is written in C# but will be converted to Java in due course. The application records and streams speech audio data from the client over the iPaq's wireless network interface to the server application. The server application consists of two Java classes namely *DASRManagerServer* and *SphinxASREngineServer*. *DASRManagerServer* is responsible for receiving client connections and subsequently audio data sent from the client. The class buffers the audio data before sending it to *SphinxASREngineServer* for speech recognition processing. The recognition result is sent back to the Manager class and the notification

```

grammar robots;
public <forward> = forward | forward <distance>;
public <backward> = backward | reverse;
public <turn_left> = (turn left <distance>) | (left <distance>);
public <turn_right> = (turn right <distance>) | (right <distance>);
public <distance> = (five {five}) | (fifty {fifty}) | (ninety {ninety});

public interface RobotGrammarInterface {
    public int moveForward();
    public int moveBackward();
    public int moveForward(int distance);
    public int moveBackward(int distance);
    public int turnLeft();
    public int turnRight();
    public double turnLeft(int degrees);
    public double turnRight(int degrees);
    public void stop();
}

```

Code fragment 1. JSGF grammar and Java interface.

process begins. If a command has been detected in the speech, this is sent to the robot (or other end device) and a notification of the command is returned to the client.

## 5.2. Grammar design

The commands appropriate to the target device are defined in a JSGF grammar file; this is used by the speech recognition engine to define what words the application can recognize. These commands are also built as methods into a Java interface. The interface is used by classes which wish to implement the real world functionality of a particular device.

The prototype device for this is a Lego Mindstorm robot and we have devised a grammar which maps its functionality to software. The robot's range of movements include forward, backward, left turn and right turn. We can enhance this further by specifying a distance or a number of degrees for the robot to move or turn. The grammar is robust enough to cater for this. [Code fragment 1](#) shows the JSGF grammar and the Java interface which defines the methods in code.

The implementation of the methods will vary depending on the application of the robot device, but the range of motions will be the same for the any application involving this robot.

## 6. Implementation challenges

### 6.1. Ipaq audio capture

A suitable Java runtime environment must be installed on the ipaq in order to run Java applications. The environment used in this research is Jeode, which provides JDK 1.1.8 compatibility and support for Java applets. The sand box security model implemented in Java prevents code from access to underlying system hardware. This has many security

advantages but creates problems for application developers wishing to use low level device functionality, in our case access to the sound card. More recent versions of Java (1.3 and better) include the Java Sound API [4] which provides access to this functionality on more powerful desktop computers.

In order to gain access and record audio from the sound card the Java Media Framework (JMF) [4] was downloaded and a customized Java only version was built specifically for the device. The ability to customize the JMF is a relatively unknown and, at this time, poorly documented feature of the software. Following numerous unsuccessful attempts to record audio through the JMF, it was ascertained that the sound functionality was not available to the JMF on the iPaq.

A pure Java solution to this problem does not currently exist. Under Scenario 2 the client is developed using C# and the Microsoft .NET Compact Framework [18] which includes low level access to the iPaq audio capabilities. The C# code is similar to Java and the runtime characteristics are identical i.e. the code is interpreted, not compiled, as is the case with Java.

Scenario 3, which is currently being developed, will use C++ code to capture the audio and populate a Java object with the data, and this will be returned to a Java program using the Java Native Interface.

## 6.2. Robot movement

The robot's movement is controlled by timers which activate the motors for a specified time before stopping. In order to move the robot accurately, the timings of the motor movements must be synchronized properly. The robot moves on two caterpillar tracks each controlled by a separate motor, therefore in order to move forward both motors must be engaged in the forward direction, the same is true for reversing. Turning is more complicated; in order to turn left, the left motor must be stopped while the right motor continues to move forward and vice versa when dealing with right turns. To allow for a greater degree of control, a distance to turn can also be specified, to execute such a command the time to pause the motor must be calculated. The equation to calculate this is show [Equation 1](#).

**Formula**  

$$(\text{Degrees\_To\_Turn} / 360) * \text{Robot 360 Time}$$

**Code Implementation**

```
public double turnRight(int degrees) {
    return (degrees / 360) * FULL_CIRCLE_TIME;
}
```

Equation 1. Robot turn movement formula.

The time it takes the robot to complete a 360 degree turn was measured and found to be 10000 ms (10 s), this is set as a constant value in the code.



Table 2  
Robot communication protocol functions

Binary value	Byte value	Function
1000110	70	Begin turn right procedure
1000111	71	Begin turn left procedure
1001000	72	Begin forward procedure
1001001	73	Begin reverse procedure

### 6.3. Robot communication protocol

The Java control program installed on the Lego Mindstorm robot acts as server which listens for incoming data packets from its control PC. The control program simply uses the built-in data port functionality to do this. In order to send commands successfully to the robot a protocol for communication needed to be designed and implemented.

Specific byte values within the RCX computer are allocated to perform certain tasks (Table 2) related to motor and sensor control, therefore these values needed to be avoided to ensure interoperability between the RCX and the Java control program. The byte range of values from 70 to 79 were found to be free for external use.

In the case of each function an additional value can follow the function value. This value indicates a distance or degree parameter for the function. If this value is not present the execution continues with the default values. This allows the protocol to comply with the application grammar.

## 7. Conclusions and future work

The physical architecture and network topology has been built for scenario one. We have implemented scenario one using both Cloudgarden/SAPI5 and Sphinx4/JSAPI. In addition, we have successfully installed the Lejos Java based operating system on the Lego Mindstorm robot and written Java code to control a simple robot. Finally a control grammar for the robot was designed and mapped to the corresponding methods in the Lejos code. The successful implementation of scenario one indicates that a framework for the command and control of remote devices is both feasible and practical when considerable computing power is available.

The distributed speech recognition and command and control model described in scenario two is almost complete. Upon completion, it will be clear how the framework fits within the ubiquitous computing paradigm. The initial tests undertaken during development indicate that the distributed model of speech recognition will be successful in the context of this application. Scenario two will also be developed further to incorporate a Java client to augment the current C# client implementation.

The framework can be easily extended and adapted to suit a variety of devices or applications. We propose to conduct further development of the framework utilizing additional software and hardware end devices. Devices presently under consideration include a web camera and programmable TINI board [2] (Fig. 11). Voice activated remote control of these devices is in line with the functional abstraction available within the

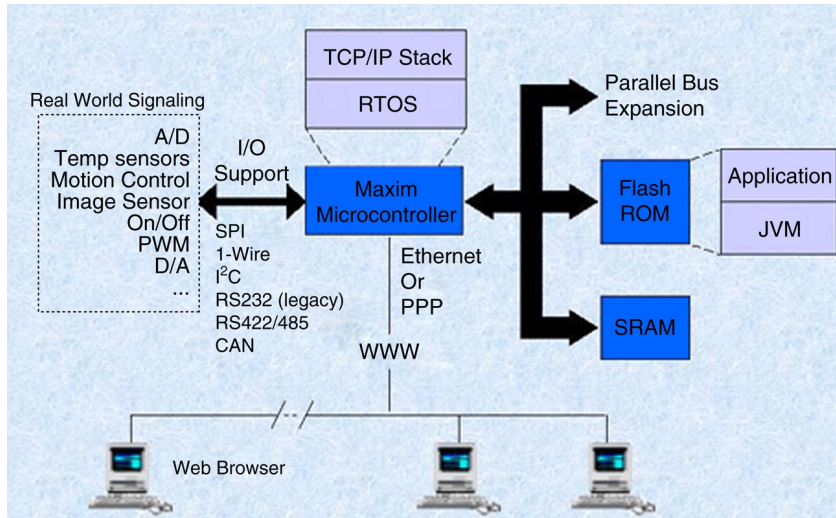


Fig. 11. TINI board schematic (from: <http://www.maxim-ic.com/products/tini/networkedmicros.cfm>).

framework to enable plug n' play control of network devices containing a Java virtual machine. All each device will need is the availability of a suitable command and control grammar of commands and parameters. Incorporating additional devices will demonstrate the flexibility of the command and control framework and extensibility of the grammar mechanism.

For future research we intend to move, in scenario three, the speech recognition even closer to the ubiquitous computing paradigm. Scenario three will involve attempting to place a minimal functionality speech recognition engine onto the iPaq. We are not certain we can do this. However, a usable and easily programmable speech recognition engine which can run efficiently on a mobile device with limited resources will enable many interesting possibilities for the development of mobile command and control applications. An important constraint is the limited resources including memory and processing power available on the iPaq and whether this will be sufficient to do speech recognition with a reasonable response time. Should we be able to achieve this successfully then this level of distribution of the speech recognition will further enhance our speech recognition framework.

## References

- [1] C. Becchetti, L.P. Ricotti, *Speech Recognition: Theory and C++ Implementation*, John Wiley & Sons, 1999.
- [2] Dallas Semi-Conductors, TINI Board. Online at <http://www.maxim-ic.com/TINI/platform.cfm>.
- [3] Gartner Group, Research on PDA and mobile phone sales by 2005. Online at <http://www.info-edge.com/samples/EM-2058sam.pdf>.
- [4] R. Gordon, S. Talley, *Essential JMF: Java media framework*, Prentice Hall PTR, 1999.
- [5] X. Huang, A. Acero, H.-W. Hon, *Spoken Language Processing: A Guide to Theory, Algorithm and System Development*, Pearson Education, 2001.
- [6] IBM, Via Voice. Online at <http://www-306.ibm.com/software/voice/viavoice/>.

- [7] F. Jelinek, *Statistical Methods for Speech Recognition (Language, Speech and Communication)*, The MIT Press, 1998.
- [8] D. Jurafsky, J.H. Martin, *Speech and Language Processing An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, Prentice Hall, New Jersey, 2000.
- [9] J. Kinnersley, *Cloudgarden Java Speech Api Implementation*. Online at <http://www.cloudgarden.com>.
- [10] Lejos, Online at <http://lejos.sourceforge.net>.
- [11] Microsoft Corporation, *Microsoft Speech and SAPI 5*. Online at <http://www.microsoft.com/speech/>.
- [12] L. Rabiner, B.-H. Juang, *Fundamentals of Speech Recognition*, Prentice Hall PTR, 1993.
- [13] ScanSoft, *Dragon Naturally Speaking*. Online at <http://www.scansoft.com/naturallyspeaking/>.
- [14] Sun Microsystems Ltd, *Java 2 Micro Edition: Personal Java*. Online at <http://java.sun.com/products/cdc/index.jsp>.
- [15] Sun Microsystems Ltd, *Java Speech API*. Online at <http://java.sun.com/products/java-media/speech/>.
- [16] Sun Microsystems Ltd, *Java Speech Grammar Format Specification*. Online at <http://java.sun.com/products/java-media/speech/forDevelopers/JSGF/>.
- [17] Sun Microsystems Ltd, *Java Speech Markup Language Specification*. Online at <http://java.sun.com/products/java-media/speech/forDevelopers/JSML/>.
- [18] A. Wigley, M. Sutton, R. MacLeod, R. Burbidge, S. Wheelwright, *Microsoft .Net Compact Framework (Core Reference)*, Microsoft Press, Redmond, 2003.