

THE CMU SPHINX-4 SPEECH RECOGNITION SYSTEM

Paul Lamere¹, Philip Kwok¹, Evandro Gouvêa², Bhiksha Raj³,
Rita Singh², William Walker¹, Manfred Warmuth⁴, Peter Wolf³

1. Sun Microsystems Laboratories, USA
2. Carnegie Mellon University, USA
3. Mitsubishi Electric Research Labs, USA
4. University of California, Santa Cruz, USA

ABSTRACT

The Sphinx-4 speech recognition system is the latest addition to Carnegie Mellon University's repository of Sphinx speech recognition systems. It has been jointly designed by Carnegie Mellon University, Sun Microsystems Laboratories and Mitsubishi Electric Research Laboratories. It is differently designed from the earlier Sphinx systems in terms of modularity, flexibility and algorithmic aspects. It uses newer search strategies, is universal in its acceptance of various kinds of grammars and language models, types of acoustic models and feature streams. Algorithmic innovations included in the system design enable it to incorporate multiple information sources in an elegant manner. The system is entirely developed on the Java™ platform and is highly portable, flexible, and easier to use with multithreading. This paper describes the salient features of the Sphinx-4 decoder and includes preliminary performance measures relating to speed and accuracy.

1. INTRODUCTION

The Sphinx-4 speech recognition system has been jointly developed by Carnegie Mellon University, Sun Microsystems Laboratories, and Mitsubishi Electric Research Laboratories (MERL). It has been built entirely in the Java™ programming language. In this paper we describe the significant features of the Sphinx-4 decoder. It is highly modular and flexible, supporting all types of HMM-based acoustic models, all standard types of language models, and multiple search strategies. Algorithmic innovations in the system enable the concurrent use of multiple information streams. The Sphinx-4 system is an open source project. The code has been publicly available at SourceForge™ since its inception. The design, results, and team meeting notes are also publicly available.

Sphinx-4 has been developed in the Java programming language, which has inherent advantages from the stand-point of code maintenance. The Java platform is highly portable: once compiled, the bytecode can be used on any system that supports the Java platform. Its garbage collection feature simplifies memory management greatly, freeing the programmer from any memory leakage concerns. The Java programming language provides a standard way of writing multithreaded applications to easily take advantage of multi-processor machines. Also, the Javadoc™ tool automatically extracts information from comments in the code and creates html files that provide documentation about the software interface.

The Sphinx-4 architecture has been designed for modularity. Any module in the system can be smoothly exchanged for another without requiring any modification of the other modules. One can,

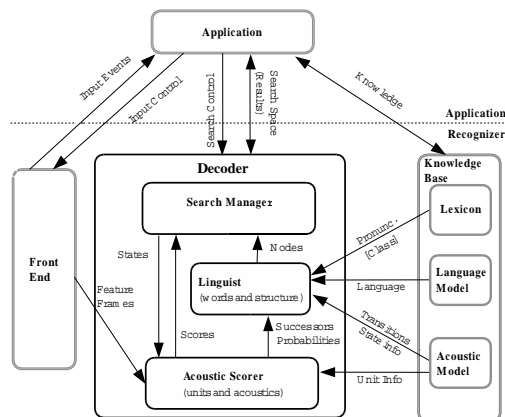


Figure 1. Sphinx-4 system architecture. The main blocks are the front end, the decoder, and the knowledge base. The decoder comprises the linguist, the search manager, and the acoustic scorer. The communication between these modules, as well as communication with an application, are depicted.

for instance, change the language model from a statistical N-gram language model to a context free grammar (CFG) or a stochastic CFG by modifying only one component of the system, namely the *linguist* (Section 2.2.2). Similarly, it is possible to run the system using continuous, semi-continuous or discrete state output distributions by appropriate modification of the *acoustic scorer* (Section 2.2.3). The system permits the use of any level of context in the definition of the basic sound units. Information from multiple information streams can be incorporated and combined at any level, *i.e.*, state, phoneme, word or grammar. The search module can also switch between depth-first and breadth-first search strategies. One by-product of the system's modular design is that it becomes easy to implement it in hardware.

In the following sections we describe the overall architecture and salient design aspects of the system, focusing on newer algorithmic improvements that have been included in it.

2. SYSTEM ARCHITECTURE

Figure 1 shows the overall architecture of the Sphinx-4 decoder. The speech signal is parameterized at the front-end module, which communicates the derived features to the decoding block. The decoding block has three components: the search manager, the linguist, and the acoustic scorer. These work in tandem to perform the decoding. In the following sections we describe each of these components in greater detail.

2.1 Front end

Figure 2 shows a detailed representation of the front-end module. The module consists of several communicating blocks, each with an input and an output. Each block has its input linked to the output of its predecessor. When a block is ready for more data, it reads data from the predecessor, and interprets it to find out if the incoming information is speech data or a control signal. The control signal might indicate beginning or end of speech - important for the decoder - or might indicate data dropped or some other problem. If the incoming data is speech, it is processed and the output is buffered, waiting for the successor block to request it.

One of the features of this design is that the output of any of the blocks can be tapped. Similarly, the actual input to the system need not be at the first block, but can be at any of the intermediate blocks. The current implementation permits us to run the system using not only speech signals, but also spectra, cepstra, etc. In addition, any of the blocks can be replaced. Thus, other kinds of features such as auditory representations can also be plugged in. Additional blocks can also be introduced between any two blocks, to permit noise cancellation or compensation on the signal, on its spectrum or on the outputs of any of the intermediate blocks. Features computed using independent information sources, such as visual features, can be directly fed into the decoder, either in parallel with the features from the speech signal, or bypassing the latter altogether.




Figure 2. Sphinx-4 front end. The current implementation is compatible with state of the art standards. Each block is independent of the others, and can be easily replaced as needed.

The system is capable of running continuously from a stream of input speech. It is also capable of three other modes of operation: fully endpointed, where the system performs explicit endpointing, determining both beginning and ending endpoints of a speech segment automatically, click to talk, where the user indicates the beginning of a speech segment but the system determines when the speech ends automatically, and push to talk, where the user indicates both the beginning and the end of a speech segment.

Currently, endpoint detection is performed by a very simple algorithm that compares the energy level to three threshold levels. Two of these are used to determine start of speech, and one for end of speech. The endpointer detects the start and/or end of speech from the incoming audio, as required, and discards non-speech segments, sending only speech segments to the decoder. The decoder thus does not waste time processing non-speech segments.

2.2 Decoder

The decoder block consists of three modules: search manager, linguist, and acoustic scorer. These are described below.

2.2.1. Search Manager

The primary function of the search manager is to construct and search a tree of possibilities for the best hypothesis. The construction of the search tree is done based on information obtained from the linguist. In addition, the search manager communicates with the acoustic scorer to obtain acoustic scores for incoming data.

The search manager makes use of a token tree. Token trees have been utilized in other speech recognition systems [1]. The token tree consists of a set of tokens that contain information about the search and provides a complete history of all active paths in the search. Each token contains the overall acoustic and language scores of the path at a given point, a *SentenceHMM* reference (see Section 2.2.2), an input feature frame identification, and a reference to the previous token, thus allowing backtracing. The *SentenceHMM* reference allows the search manager to fully categorize a token to its senone, context-dependent phonetic unit, pronunciation, word, and grammar state.

The search algorithm maintains a set of active tokens in an *active list*. The active list represents the tips of the active search branches. During the search, each input feature frame is scored against the acoustic models associated with each token in the active list, and low scoring branches are pruned. The search manager updates the active list based on the successor *SentenceHMM* states of the tokens remaining after pruning.

Sphinx-4 provides a beam pruner that constrains the scores to a configurable minimum relative to the best score, while also keeping the total number of active tokens to a configurable maximum. New implementations can easily be created that provide alternative methods of storing and pruning the active list. The garbage collector automatically reclaims unused tokens, thus simplifying the implementation of the pruner by merely allowing a token to be removed from the set of active tokens. The active list is available as part of the final recognition result. Applications can use the active list to inspect the highest scoring paths. This can be used, for instance, to construct an N-Best list.

Search through the token tree and the sentence HMM is performed in two ways: *depth-first* or *breadth-first*. Depth-first search is similar to conventional stack decoding. The most promising tokens are expanded in time sequentially. Thus, paths from the root of the token tree to currently active tokens can be of varying lengths. In breadth-first search, on the other hand, all active tokens are expanded synchronously, making the paths from the root of the tree to the currently active tokens equally long.

In Sphinx-4, breadth-first search is performed using the standard Viterbi algorithm as well as a new algorithm called *Bush-derby*. During search, competing units (phoneme, word, grammar etc.) are each represented by a directed acyclic graph (dag). Each dag has a source and a sink, as illustrated by the examples in Figure 3. Figure 3a shows two-node dags for two competing phonemes AX and AXR. The dags associated with the competing word units CAT and RAT in Figure 3b are more complicated.

In Viterbi, each competing unit is scored by the probability of the (single) best path through it and the unit with maximum best-path score wins. For example, if the probabilities on the edges were (0.9,0.02,0.01) and (0.2,0.7,0.6) for the phonemes AX and AXR, respectively, then the scores would be 0.9 and 0.7 and AX would win. If we used the sum of the probabilities of all paths rather than

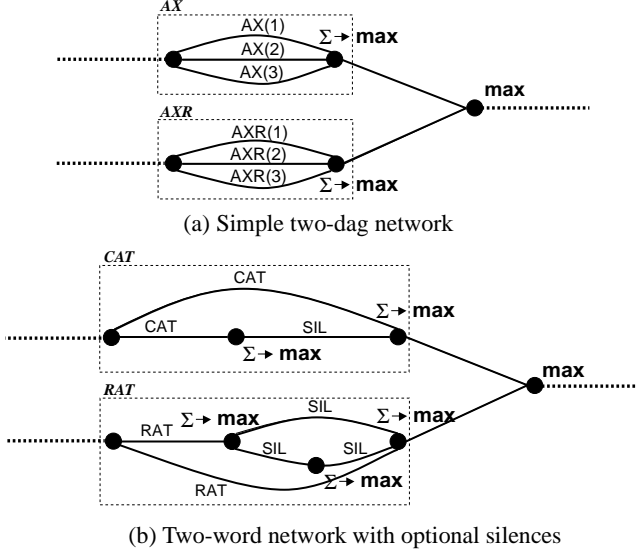


Figure 3. Bush-derby algorithm: Rather than considering only the highest scoring path for each unit (surrounded by dotted boxes), this algorithm combines all paths of the unit into a single score.

the maximum, then the phoneme AXR would come out higher.

Bush-derby computes an η -score for each dag (where $\eta \geq 1$ is a parameter) that interpolates between the sum and the maximum. The essential difference to Viterbi is that all paths (*i.e.*, the entire “bush” of paths) contribute to the score of a unit. The sink nodes of competing units are combined at a *max* node. So among competing units, the one with the highest score still wins.

The η -scores are computed as follows. Source nodes have score one and the η -score of any other node i is computed recursively from the scores of the predecessors j in the dag and the probabilities of the edges into node i :

$$\text{score}(i) = \left(\sum_j (\text{score}(j) \text{prob}(\text{edge}(j, i)))^\eta \right)^{\frac{1}{\eta}} \quad (1)$$

This function is computed at all the $\Sigma \rightarrow \mathbf{max}$ nodes of Figure 3. By expanding Equation (1) one can see that the η -score of the sink node of a unit U can also be expressed as a sum over all source to sink paths π in a unit U :

$$\text{score}(U) = \left(\sum_{\pi \in U} \text{prob}(\pi)^\eta \right)^{\frac{1}{\eta}} \quad (2)$$

This is the η norm of a vector where the probabilities of the paths in U are the components. Note that when $\eta = 1$, then the score associated with a unit is simply the total probability of all paths. As η increases the most likely paths contribute the most to the score, and as $\eta \rightarrow \infty$, the η -score reduces to the probability of the best path (*i.e.*, Viterbi). Note that the overall computation needed for computing the new scores via Equation (1) is of the same order as Viterbi. The essential difference between the Viterbi and Bush-derby algorithms is that the former returns the sequence of units that represent the global best path through the graph, whereas the latter returns the sequence of units that win the local max competitions between the fancier η -scores of the units.

2.2.2. Linguist

The linguist translates linguistic constraints provided to the system into an internal data structure called the *grammar* which is usable by the search manager. Linguistic constraints are typically provided in the form of context free grammars, N-gram language models, finite state machines etc. The grammar is a directed graph, where each node represents a set of words that may be spoken at a particular time. The nodes are connected by arcs which have associated language and acoustic probabilities that are used to predict the likelihood of transiting from one node to another. Sphinx-4 provides several grammar loaders that load various external grammar formats and generate the internal grammar structure. The pluggable nature of Sphinx-4 allows new grammar loaders to be easily added to the system.

The grammar is further compiled into a *SentenceHMM*, which is a directed state graph where each state in the graph represents a unit of speech. Grammar nodes are decomposed into a series of word states, one for each word represented by the node. Words states are further decomposed into pronunciations states, based on pronunciations extracted from a dictionary maintained by the linguist. Each pronunciation state is then decomposed into a series of unit states, where units may represent phonemes, diphones, etc. and can be specific to contexts of arbitrary length. Each unit is then further decomposed to its sequence of HMM states. The SentenceHMM thus comprises all of these states. States are connected by arcs that have language, acoustic and insertion probabilities associated with them.

Although the contents of a SentenceHMM are well defined by the linguist, there are a number of strategies that can be used in constructing the SentenceHMM that affect the search. By altering the topology of the SentenceHMM, the memory footprint, perplexity, speed and recognition accuracy can be affected. The pluggable nature of Sphinx-4 allows different SentenceHMM compilation methods to be used without changing other aspects of the search.

With large grammars, the SentenceHMM can grow to be quite large. Sphinx-4 provides a mechanism that allows subsets of the SentenceHMM to be dynamically constructed as needed and discarded when no longer needed. We expect this to allow Sphinx-4 to support very large grammars such as is typically required for general dictation recognition.

2.2.3. Acoustic Scorer

The task of the acoustic scorer is to compute state output probability or density values for the various states, for any given input vector. The acoustic scorer provides these scores on demand to the search module. In order to compute these scores, the scorer must communicate with the front-end module to obtain the features for which the scores must be computed.

The scorer retains all information pertaining to the state output densities. Thus, the search module is ignorant of whether the scoring is done with continuous, semi-continuous or discrete HMMs. Any heuristic algorithms incorporated into the scoring procedure for speeding it up can be performed locally within the search module. Where such heuristics can benefit from additional information derived from the search module, however, this information can be obtained on demand. The modular nature of the design permits us to include conventional Gaussian scoring procedures, as well as several new algorithms based on game-theoretic considerations.

3. EXPERIMENTAL EVALUATION

In this section we report results on some benchmark tests of increasing complexity. As a comparator, we report performance on the CMU Sphinx-3 [2] speech recognition system on the same tasks. For all experiments acoustic models were trained with the training module of the CMU Sphinx-3 system. The following benchmark tests have been conducted so far:

Isolated digits: This was conducted on the TI46 isolated digits database [3]. The acoustic models were trained on speech from the adult speakers in the TIDigits database [4].

Connected Digits: This test was conducted on a subset of the TIDigits database containing adult speakers. The acoustic models used were the same as those for isolated digits.

Small Vocabulary: This was conducted on CMU’s Census database (AN4) [5]. The entire database (including components conventionally marked as “training” and “test”) was used for recognition. Acoustic models were trained using the Wall Street Journal (WSJ) database [6]. In this benchmark we ran several tests. These included tests on control words (words) and spelling (spelling) components of AN4 as well as on the entire database (full). Performances with a flat unigram LM as well as a trigram LM were evaluated.

Table 1 shows the results of the benchmark tests. Both word error rate (WER) and recognition speed in times real time are shown in the table. All experiments were run on a Sun Blade™ 1000 workstation with dual 750 MHz UltraSPARC® III processors

Table 1: Benchmark tests comparing Sphinx-3 (S3) and Sphinx-4 (S4)

| Task Name | WER(%) | | Speed (xRT) | |
|----------------------|--------|------|-------------|------|
| | S3 | S4 | S3 | S4 |
| Isolated Digits | 0.7 | 0.2 | 0.79 | 0.26 |
| Connected Digits | 1.1 | 0.6 | 0.79 | 0.22 |
| Small vocabulary | | | | |
| -Words | 9.8 | 8.1 | 6.48 | 1.21 |
| -spelling | 30.4 | 20.6 | 6.53 | 1.52 |
| -full | 25.8 | 25.4 | 6.57 | 3.6 |
| Trigram small vocab. | | | | |
| -words | 3.4 | 3.9 | 6.49 | 1.29 |
| -spelling | 22.3 | 16.8 | 6.52 | 1.82 |
| -full | 15.0 | 18.3 | 6.57 | 3.56 |

The results in Table 1 show that the Sphinx-4 is clearly comparable to the Sphinx-3 in recognition performance and significantly superior in terms of speed, on the tasks evaluated. Note that the Sphinx-4 performance has not been optimized on the trigram tasks at the time of this submission. We have also not completed evaluations on medium and large vocabulary tasks. The optimized trigram tasks and the medium and large vocabulary evaluations will be completed by the time the paper is presented. They will also be reported on SourceForge as and when they are completed.

4. SUMMARY

In this paper we have described the salient features of the architecture of CMU’s latest Sphinx-4 speech recognition system. We have described the advantages of the system’s modular design, its flexibility in the usage of various kinds of acoustic and language representations, and the features due to the Java platform that it has been developed on. The recognizer incorporates a completely new algorithm in its search, the Bush-derby algorithm, in addition to the conventional Viterbi search algorithm. The Bush-derby algorithm enables the system to include all paths through a unit into the score for the unit. It also facilitates the use of multiple sources of information, such as audio and visual features, or parallel feature streams, in a natural way. In addition, the η parameter associated with the Bush-derby algorithm provides an elegant way of emphasizing or de-emphasizing the relative contribution of high-likelihood features during the search. Benchmark tests show that the Sphinx-4 is comparable to Sphinx-3 in recognition performance, and is much faster. Although all benchmark tests have not yet been performed, these will be reported on SourceForge as they are completed.

ACKNOWLEDGEMENTS

The authors thank Prof. Richard Stern at CMU, Robert Sproull at Sun Microsystems, and Joe Marks at MERL, for making this team possible. We also thank SourceForge.net for their contribution to the open source community. The authors also thank Sun Microsystems Laboratories and the current management for their continued support and collaborative research funds. Manfred Warmuth was sponsored by the National Science Foundation under Grant No. NSF CCR 9821087. Rita Singh was sponsored by the Space and Naval Warfare Systems Center, San Diego, under Grant No. N66001-99-1-8905. The content of this publication does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

REFERENCES

- [1] S.J. Young, N.H. Russel, and J.H.S. Russel (1989). “Token passing: A simple conceptual model for connected speech recognition systems,” *Technical Report*, Cambridge University Engineering Dept.
- [2] P. Placeway, S. Chen, M. Eskenazi, U. Jain, V. Parikh, B. Raj, M. Ravishankar, R. Rosenfeld, K. Seymore, M. Siegler, R. Stern, and E. Thayer (1997). “The 1996 hub-4 sphinx-3 system,” in *Proc. DARPA Speech Recognition Workshop*.
- [3] G.R. Doddington and T.B. Schalk (1981). “Speech recognition: Turning theory to practice,” *IEEE Spectrum*, vol 18, no. 9, pp 26-32.
- [4] R.G. Leonard and G. R. Doddington (1984). “A database for speaker independent digit recognition,” in *Proc. Intl. Conf. on Acoustics Speech and Signal Processing*, vol. 3, p. 42.11.
- [5] A. Acero, *Acoustical and environmental robustness in automatic speech recognition*, Kluwer Academic Publishers, 1993.
- [6] D. Paul and J. Baker (1992). “The design of the wall street journal-based csr corpus,” in *Proceedings of ARPA Speech and Natural Language Processing Workshop*. pp 357-362.