REAL-TIME SPEECH RECOGNITION SYSTEM

Mitchel Weintraub

SRI International
Speech Research and Technology Program
Menlo Park, CA, 94025

PROJECT GOALS

SRI and U.C. Berkeley are developing hardware for a real-time implementation of spoken language systems (SLS). Our goal is to develop fast speech recognition algorithms and supporting hardware capable of recognizing continuous speech from a bigram or trigram based 10,000 word vocabulary or a 1,000 to 5,000 word SLS system.

RECENT RESULTS

The special-purpose system achieves its high computation rate by using special-purpose memories and data paths, and is made up of the following several components:

- A special-purpose HMM-board with eight newly designed integrated circuits that does the HMM inner-loop processing to implement the word-recognition algorithms.
- An output-distribution board made of off-the-shelf components for computing HMM discrete-density state-output probabilities.
- A multi-processor TMS32030 board for computing the statistical language processing. This board has a custom high-speed interface to the HMM-board.
- A general-purpose CPU board to perform system control.
- A DSP board with A/D convertor for computing the feature extraction.
- A Sun workstation for computing the spoken language system database retrieval and human machine interface.

SRI and U.C. Berkeley's recent accomplishments on this project include:

- Completed the construction of a working hardware prototype. This prototype has been demonstrated running the Resource Management (RM) task as

well as the Airline Travel Information System (ATIS) task.

- Began intensive use of the hardware for a real-time Airline Travel Information System (ATIS) task.
- Completed the design and construction of a second generation multiprocessor TMS32030 grammar processing board. Testing is currently in progress.
- Revised and corrected errors in several of the custom VLSI chips that are used for the HMM word-recognition processor.

PLANS FOR THE COMING YEAR

- Complete the construction and testing of the second generation multiple-processor TMS32030 board with a high I/O bandwidth to interface with the special-purpose HMM-board.
- Implement multiple types of grammars using this hardware.
- Collect data about man-machine speech interactions using the real-time hardware.
- Integrate the real-time recognizer into our research to shorten the development cycle for new systems.
- Evaluate the current architecture to determine the computational and algorithmic bottlenecks.
- Deliver a hardware prototype to DARPA.