SYSTEM ORGANIZATIONS FOR SPEECH UNDERSTANDING:
Implications of Network and Multiprocessor Computer Architectures for AI

by

L.D. Erman, R.D. Fennell, V.R. Lesser, and D.R. Reddy

Computer Science Department*
Carnegie-Mellon University
Pittsburgh, Pa. 15213

ABSTRACT

This paper considers various factors affecting system organization for speech understanding research. The structure of the Hearsay system based on a set of cooperating, independent processes using the hypothesize-and-test paradigm is presented. Design considerations for the effective use of multiprocessor and network architectures in speech understanding systems are presented: control of processes, interprocess communication and data sharing, resource allocation, and debugging are discussed.

Keywords: **speech recognition, speech understanding, system organization, networks, multiprocessors, parallel processing, real-time systems, hardware for AI, software for AI.**

INTRODUCTION

System organizations for speech understanding systems must address many problems: effective use of multiple sources of knowledge, anticipation and goal-direction in the analysis of the incoming utterance, real-time response, continuous monitoring of input device(s), errorful nature of the recognition process, exponential increase of processing requirements with the increase of desired accuracy, and so on. A particular model of speech perception (Reddy et al., 1973) which attempts to solve the above problems involves the use of cooperating independent processes using a hypothesize-and-test paradigm. This paper examines the effect of the problem constraints and the model on system organizations, presents the structure of a system currently operational on a PDP-10 computer, and discusses the implications of multiprocessor and network architectures.

Unlike many other problems in artificial intelligence, speech understanding systems are characterized by the availability of diverse sources of knowledge, e.g., acoustic-phonetic rules, phonological rules, articulatory models of speech production, vocabulary and syntactic constraints, semantics of the task domain, user models, and so on. A major problem, then, is to develop paradigms which can make use of all the available sources of knowledge in the problem solution. At the same time, absence of one or more sources of knowledge should not cripple the system. Suppose each source of knowledge is represented within the system as a process. In order to remove or add sources of knowledge, each process must be independent, i.e., it must not require the presence of other processes in the system. But at the same time each process must cooperate with the other processes, i.e., it must be able to effectively use the information gathered by them about the incoming utterance. Thus, a major design step is to establish what information is to be shared among processes and how this information is to be communicated so as to maintain the independence of individual processes while still allowing for necessary process cooperation.

Knowledge available in the acoustic signal represents only one part of the total knowledge that is brought to bear in understanding a conversation. A good example of this is when one is interrupted by an appropriate response from the listener to a question that is as yet incomplete. In general, a human listener can tolerate a great deal of sloppiness and variability in speech because his knowledge base permits him to eliminate most of the possibilities even as he hears the first few words of the utterance (if not before!). We feel that this notion of anticipation, prediction, and hypothesis generation is essential for machine perception systems as well. In general, we expect every source of knowledge to be able to generate hypotheses in a given context, or verify hypotheses generated by others using different representations of knowledge, if necessary. The implication is that knowledge processes be organized within the system so as to reduce the problem of recognition and understanding to one of prediction and verification.

In tasks such as chess and theorem-proving, the human has sufficient trouble himself so as to make reasonably crude computer programs of interest. But, because humans seem to perform effortlessly (and with only modest error) in speech (and visual) perception tasks, similar performance is expected from machines, i.e., one expects an immediate response and will not tolerate any errors. To equal human performance, a speech understanding system must be able to understand trivial

* This research was supported in part by the Advanced Research Projects Agency of the Department of Defense under contract no. F44620-70-C-0107 and monitored by the Air Force Office of Scientific Research.

questions as soon as they are uttered. This implies that various processes within the system should be allowed to operate as soon as there is sufficient incoming data, without waiting for the completion of the whole utterance. If the processes within the system are independent and unaware of the existence of each other, then the system must provide facilities for activation, termination, and resource allocation for each of the processes. Further, if a process can be deactivated before it reaches a natural termination point, provision must be made to preserve the state of the process until it is reactivated. Also, it is necessary to provide interlocks on the data that are shared among many processes.

This has several implications for system organization. The system must monitor the input device continuously to determine whether speech is present; this requires non-trivial processing. If the system is unable to process the incoming data, automatic buffering must be provided. If the system is to run on a time-sharing system, provision must be made to ensure that no data is lost because the program is swapped out for a period of time. If the speech understanding system is to consist of a set of cooperating independent processes, it is further necessary that they be able to be interrupted at unprepprogrammed points -- if the microphone monitoring program is not activated in time to process the incoming utterance, it could lead to irrevocable loss of data. These considerations lead to two additional requirements that are not commonly available on existing time-sharing systems, viz., process-generated interrupts of other processes and user servicing of interrupts.

One of the characteristics of speech understanding systems is the presence of error at every level of analysis. To control such errors and permit recycling with improved definitions of the situation, one uses techniques such as feedforward, feedback, and probabalistic backtracking. If such facilities do not exist within the system, they have to be programmed explicitly.

Speech, by its nature, appears to be computer intensive. A substantially unrestricted system capable of reliably understanding connected speech of many speakers using a large vocabulary is likely to require systems of the order of a proposed AI machine (Bell, Freeman, et al., 1971a), i.e., processing power of 10 to 100 million instructions per second and memory of 100 to 1000 million bits.* To obtain such processing power, it appears necessary to consider multiprocessor architectures. Decomposition of speech processing systems to effectively use distributed processing power requires careful consideration even with primitive systems. Our model of cooperating independent processes, each representing a source of knowledge, leads to a natural decomposition of the algorithms for such machine architectures.

THE CURRENT Hearsay SYSTEM

In this section we briefly describe the Hearsay speech understanding system as it now exists at C-MU. (More detailed descriptions of the system are given in Reddy et al., 1973,1973a (this volume); Erman, 1973; and Neely, 1973.) We shall stress those aspects of its organization which are responsive to the constraints and model outlined above. This system represents a first attempt to solve those problems; thus, some of the constraints are only partially or poorly met, while others are satisfied in a more constricted way than necessary. We shall point out these limitations as they are described; later sections on closely-coupled and loosely-coupled processor network architectures describe possible corrections and improvements of the system.

* Smaller and substantially cheaper systems can be built to perform useful but restricted speech understanding tasks.

The Hearsay system is implemented as a small number of parallel coroutines (see figure). Each coroutine (module) is realized as a separate job in the PDP-10 time-sharing system; thus the time-sharing monitor is the primary scheduler for the modules. In general, the modules may achieve a high degree of (pseudo-) parallel activity (through the use of shared memory and a flexible inter-process message system*), but, in practice, we limit the parallelism to a very modest amount. This limitation is imposed for two reasons: first, since the PDP-10 is a uniprocessor system, there is nothing to be gained (in the time domain) by increasing the parallelism; and, second, the greater the amount of parallelism, the more difficult it is to control and debug the programs within a time-sharing system that is not designed for cooperating processes (jobs).

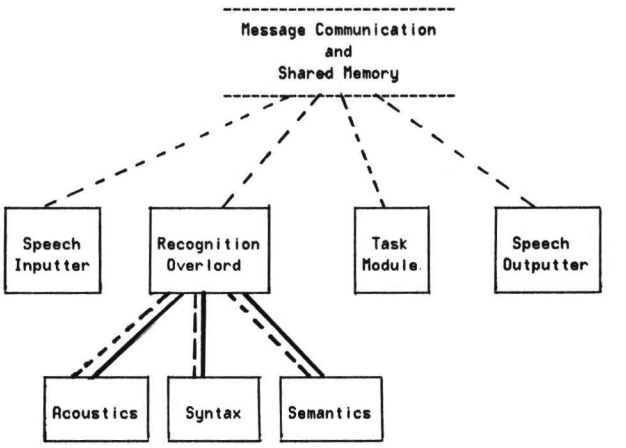

Decomposition of processes in the current Hearsay system.

The model of recognition specifies that there be separate processes, each representing a different domain of knowledge. We have chosen three major domains of knowledge: acoustic-phonetics, syntax, and semantics:

1. The acoustic-phonetic domain, which we refer to as just **acoustics**, deals with the sounds of the language and how they relate to the speech signal produced by the speaker. This domain of knowledge has traditionally been the only one used in most previous attempts at speech recognition.
2. The **syntax** domain deals with the ordering of words in the utterance according to the grammar of the input language.
3. The **semantic** domain considers the meaning of the utterances of the language, in the context of the task that is specified for the speech understanding system.

These processes, according to the model, are to be independent and removable; therefore the functioning (and very existence) of each must not be necessary or crucial to the others. On the other hand, the model also requires that the processes cooperate and that the recognition should run efficiently and with

* The facilities provided for inter-job control and communication are similar to those developed for the Stanford Hand-Eye system (Feldman and Sproull, 1971).

good error recovery; these dictates imply that there be a great deal of interaction among the processes. Thus we seem to have opposing requirements for the system. These opposing requirements led to the design of the following structure:

> Each process interfaces externally in a uniform way that is identical across processes; no process knows what or how many other recognition processes exist.
>
> A mediator, ROVER (Recognition OVERlord), handles the interface to each of the processes and thus serves as the linkage connecting the processes; the processes are called ROVER's "sons."

The interface is implemented as a global data structure which is maintained by ROVER. Each of ROVER's sons puts information into this data structure in a uniform way. Each may access information submitted by its brothers, but in a manner which leaves the source of that information anonymous. This mechanism is analogous to a bulletin board on which messages can be left by several people and for which there is a monitor who accepts the message and arranges them in appropriate places on the board for others to react.

This anonymous interface structure is appropriate only if the global data structure can be designed in such a way as to allow the processes to communicate meaningfully; i.e. there must be a common language which allows them to transmit the kind of information they need to help each other to work on the problem. We resolve this problem by using the **word** as the basic unit of discourse among the processes.

The basic element of the global data structure is the **word hypothesis** which represents an assertion that a particular word (of the input language lexicon) occurs in a specified position in the spoken input. A **sentence hypothesis** is an ordered linear sequence of word hypotheses; it represents an assertion that the words occur in the sentence in the order that the word hypotheses appear in the sentence hypothesis. In addition, the unique "word" **FILLER** may appear as a word hypothesis; this is a placeholder and represents the assertion that zero or more as yet unspecified words occur in this position in the spoken sentence. In general, there may be any number of sentence hypotheses existing at any one time.

The interactions among the source-of-knowledge processes are carried out using the hypothesize-and-test paradigm prescribed by the model. In general, any process may make a set of hypotheses about the utterance; all the processes (including the hypothesizer) may then verify (i.e. reject, accept, or re-order) these hypotheses. In particular, hypothesization occurs when a recognition process (Acoustics, Syntax, or Semantics) chooses a FILLER word from a sentence hypothesis and associates with it one or more **option words**, each of which it asserts is a candidate to replace all or part of the FILLER. Verification consists of each process examining the option words and rating them in the context of the rest of the sentence hypothesis.

Several restrictions have been placed on the implementation of this general scheme. First, at any time only one part of the shared, global data structure (i.e., one sentence hypothesis) is accessible to the processes for hypothesization and verification. Second, the processes go through the hypothesization and verification stages (and several other subsidiary stages) in a synchronized and non-interruptable manner. Finally, only one process is allowed to hypothesize at any one time. Again, these restrictions were imposed both because parallelism on a uniprocessor does not accomplish any throughput increase and because the available programming and operating systems make a more general implementation difficult to specify, debug, and instrument. These restrictions are mitigated somewhat by carefully adjusting the time grain of the processing so that each non-interruptable phase is not "excessively large."

Each sentence hypothesis has a confidence rating associated with it which is an estimate of how well it describes the spoken utterance. This rating is calculated by ROVER, based on information supplied by the recognition processes. Errors in processing become evident when the overall rating given to a sentence hypothesis begins to drop; at that point, attention is focused on some other sentence hypothesis with a higher rating. This switching of focus is the mechanism that provides the error recovery and backtracking that is necessary in any speech understanding system.

CLOSELY-COUPLED PROCESSOR SYSTEM ORGANIZATIONS

As discussed in the introduction, in order to do real-time speech understanding a substantial amount of computing power is required. Recent trends in technology indicate that this computing power can be economically obtained through a closely-coupled network of "simple" processors, where these processors can be interconnected to communicate in a variety of ways (e.g., directly with each other through a highly multiplexed switch connected to a large shared memory (Bell et al., 1971), or through a regular or irregular network of busses (Bell et al., 1973)). However, the major problem with this network approach to generating computing power is finding algorithms which have the appropriate control and data structures for exploiting the parallelism available in the network. The model for a speech understanding system as previously discussed, which is decomposed into a set of independent processes cooperating through a hypothesize-and-test paradigm, represents a natural structure for exploiting this network parallelism.

There exist three major areas for exploitation of parallelism in the structure of this speech understanding system: preprocessing, hypothesization and verification, and the processing specific to each source of knowledge. The **preprocessing task** involves the repetition of a sequence of simple transformations on the acoustic data, e.g., detection of the beginning and end of speech, amplitude normalization, a simple phoneme-like labeling, smoothing, etc. This sequence of transformations can be structured as a pipeline computation in which each transformation is a stage in the pipe. Thus, through this pipeline decomposition of the preprocessing task, a limited amount (i.e., 4) of parallel activity is generated.

The **hypothesize-and-test paradigm** for sequencing the activity of the different sources of knowledge can also be structured so as to exhibit parallelism, but the amount of parallelism is potentially much greater. This parallel activity is generated by the simultaneous processing of multiple sentence hypotheses and the simultaneous hypothesization and verification by all sources of knowledge. The simultaneous processing of multiple sentence hypotheses, rather than processing just the currently most likely candidate, can conceptually introduce unnecessary work. But in practice, because of the errorful nature of the processing, there may be a considerable amount of necessary backtracking to find the best matching sentence hypothesis. It is appropriate to quote a conjecture of Minsky and Papert (1969, Section 12.7.6) on this point:

> [While for the exact match problem] relatively small factors of redundancy in memory size yield very large increases in speed, . . . [for the best match problem] . . . for large data sets with long word lengths there are no practical alternatives to large searches that inspect large parts of the memory.

Thus, the parallel activity generated by simultaneous processing of more than one sentence hypothesis can result in a

proportional speed-up of the recognition process.* Correspondingly, simultaneous hypothesization and verification by all sources of knowledge also results in a proportional speed-up of the recognition process because each source of knowledge is independent and is designed so that its knowledge contirbution is additive.

Finally, the **verification algorithm** of each source of knowledge can be decomposed into a set of parallel processes in two ways: The first kind of decomposition is based on the fact that verifications are performed on a set of option words rather than a single word at a time. Thus, for each source of knowledge there can be multiple instantiations of its verification process, each operating on a different option word. The second kind of decomposition involves the parallelizing of the verification algorithms themselves; thus, each instantiation of a verification process may itself be composed of a set of parallel processes. However, this set of instantiations may not be totally independent because the rating produced by the verification process may be dependent on the particular set of option words to be verified and also on the local data base which is common to all the instantiations. For example, the acoustic verification process is a hierarchical series of progressively more sophisticated tests. The first few levels of testing look only at the context of a single option word, while the more sophisticated tests compare one option word against another. Thus, only at the first few levels of tests can the acoustic verification algorithm be parallelized in a straightforward manner.

The parallelism generated by parallelizing the hypothesize-and-test control structure and the verification processes are multiplicative in their parallel activity (i.e. performing in parallel the updating of **n** sentence hypothesis where each hypothesis invokes **m** verification processes and each verification process operates on **o** option words leads to a potential parallelism of **n*m*o**). This parallelism, together with the pipeline parallelism of the preprocessing, leads to what appears to be a large amount of potential parallelism to be exploited by a closely-coupled network. However, it is still not clear just how much potential parallel activity exists over the entire recognition system; nor is it known how much of this potential will be dissipated because of software and hardware overhead.

In order to answer these questions, a parallel decomposition of the Hearsay speech understanding system is now being implemented on C.mmp, a closely-coupled network of PDP-11's which communicate through a large shared memory (Bell et al., 1971). The C.mmp hardware configuration can contain up to 16 PDP-11's; the highly multiplexed switch that connects processors to memory permits up to 16 simultaneous memory references if these references are not to the same memory module. Thus, if processors are referencing different memory modules, then each processor can run at full speed. In addition, C.mmp can be configured for a specific application (e.g., speech) by replacing a processor by a special purpose hardware device which directly accesses memory (e.g., a signal processor).

The HYDRA software operating system (Wulf, 1972), which is associated with C.mmp, provides an appropriate kernel set of facilities for implementing the parallel version of the speech system. These facilities permit control of real-time devices, convenient building of a tree of processes, message queues and shared data base communication among processes, user-defined scheduling strategies, arbitrary interruption of running processes, and dynamic creation of new processes. Building up from this base, a debugging system will be constructed which, in addition to the normal features, will permit the recording of all communication among processes, the tracing of all process activity, and the monitoring of global variables (including a recording of which processes have modified them). These additional capabilities are crucial for isolating errors and understanding the dynamic behavior patterns of the parallel system.

The major software problem to be investigated in this parallel implementation of the Hearsay system is how to efficiently map virtual parallelism (process activity) into actual parallelism (processor activity). This mapping problem in turn centers on three design issues, each of which relates to how processes interact:

1. the design of the interlock structure for a shared data base,
2. the choice of the smallest computational grain at which the system exhibits parallel activity, and
3. the techniques for scheduling a large number of closely-coupled processes.

The first design issue is important because in a closely-coupled process structure many processes may attempt to access a shared data base at the same time. In a uniprocessor system, the sequentialization of access to this shared data base does not significantly affect performance because there is only one process running at a time. In a multiprocessor system, however, if the interlock structure for a shared data base is not properly designed so as to permit as many non-interfering accesses as possible, then access to the shared data base becomes a significant bottleneck in the system's performance (McCredie, 1972).

The second issue relates to how closely-coupled processes can interact. If the grain of decomposition is such that the overhead involved in process communication is significant in relation to the amount of computation done by the process, then the added virtual parallelism achieved by a finer decomposition can decrease, rather than increase, the performance of the system. Thus, understanding the relationship between the grain of decomposition and the overhead of communication is an important design parameter.

The third issue relates to a phenomenon called the "control working set" (Lesser, 1972). This phenomenon predicts that the execution of a closely-coupled process structure on a multiprocessor may result in a significant amount of supervisory overhead caused by a large number of process context switches. The reason for this high number of process context switches is analogous to the reason for "thrashing" within a data working set (Denning, 1968). For example, in a uniprocessor system if two parallel processes closely interact with each other, then each time one process is waiting for a communication from the other it would have to be context switched so as to allow the other process to execute. If these two processes communicate often then there would be a large number of context switches. However, if there were two processors, each containing one of the processes, then there would be no process switching.

The implications of this phenomenon on constructing process structures are the following:

1. Processes should be formed into clusters where communication among cluster members is closely-coupled whereas communication among clusters is loosely-coupled. This process structuring paradigm has also been been suggested as a model for the operation of complex human and natural systems (Simon, 1962).
2. The size of a process cluster cannot be chosen independent of the particular hardware configuration that will be used to execute it. For example, a cluster size of 8 may be appropriate

* Simulation studies are currently being carried out on evaluating this speed-up factor. These studies are based on data generated from the current version of the Hearsay system.

for a hardware system containing 16 processors while being inappropriate for a system containing 6 processors.

3. The scheduler of a multiprocessor system should use a strategy that schedules process clusters rather than single processes. (This is analogous to the advantage of preloading the data working set rather than dynamically constructing the working set at each context swap.)
4. The use of process structures to implement inherently sequential, though complex, control structures (e.g., coroutines, etc.) may lead to inefficient scheduling of process structures on a multiprocessor system (i.e., the scheduling strategy should be able to easily differentiate those processes that can go on in parallel from those that are sequentialized).

NETWORK ORGANIZATIONS

The multiprocessor type organization described earlier implies a closely-coupled set of processes on a set of closely-coupled processors cooperating to accomplish the common goal of utterance recognition. The key idea in such a system is that both the processes and processors are closely-coupled -- that is, the cost of communication between processes or processors is relatively cheap with respect to the amount of computation to be done by any individual process. Indeed, in the multiprocess system described earlier, much interprocess communication and data sharing may be achieved by actually having shared physical address spaces. However, such a system usually also implies a certain homogeneity or physical proximity of the processors and memory.

Consider now the task of integrating the knowledge of many different research groups in various widespread geographical locations, each with its own computing facilities and each with its own areas of specialization. In an attempt to avoid unnecessary duplications of effort, one would desire a scheme whereby each group could develop pieces of a total recognition system (which pieces might represent new sources of knowledge, such as a new and improved vowel classification algorithm) using local computing resources (i.e., using an arbitrary machine configuration and program structure). Those pieces of the system would then be incorporated into a distributed "total recognition system" by appropriate (hopefully minimal) linkage and protocol conventions and their contributions to the entire system evaluated. The geographical constraints suggest the use of a computer network facility as a means by which one might assemble this total recognition system. We are currently undertaking the task of designing and implementing such a system for use on the ARPA network of computing facilities (Roberts and Wessler, 1970). The usefulness of such a network organization for a speech understanding system lies in its potential ability to combine and evaluate the various algorithms and sources of knowledge of a wide variety of research groups. In particular, the objective of the network organization is to create a research tool rather than to produce a highly efficient recognition system.

As an example, suppose a group wishes to add a new source of knowledge (a new vowel classification algorithm, for instance) to the network system. This knowledge-source is provided in the form of a process (or a set of processes) running on a local computer connected to the ARPA network. System integration is then achieved by adding linking instructions to the process (perhaps interactively) for notifying a centralized controlling process of the set of pre-conditions (e.g., conditions relating to the incoming speech wave or the current state of the recognition) that must be met in order to activate this process (Adams, 1968), as well as the required inputs and created outputs (and their formats). The central controller is then responsible for activating the new knowledge source at appropriate times, supplying the requested inputs, and updating a global data base to reflect the results of the activated process. Knowledge source processes may communicate with one another via a message service facility provided by the central controller. The marked increase of indirection with respect to communication and data sharing as compared with a closely-coupled multiprocessor approach is a result of the goal to serve a wide geographic region of users and to allow cooperation between essentially autonomous knowledge sources.

The problems that occur in this network concept are of a nature different from that of those occurring in the multiprocessor structure described previously. The many sources of knowledge are no longer necessarily closely-coupled. In fact, we might term such a network organization to be "loosely-coupled" in the sense that process communication and data base sharing must be achieved by some form of message switching scheme since the system is now operating on an indefinite number of (nonhomogeneous) computers. In particular, there is no longer the ability for all processes to share data and communicate by sharing physical address spaces. The problems of data base sharing and shipping now abound: one would like not to have multiple copies of a given data structure due to updating synchronization problems, but the message switching involved in maintaining and updating a single, centralized data structure may be overwhelmingly inefficient.

It is intended that, besides serving as a research tool for testing various recognition algorithms and combinations thereof, such a network organization will become an interesting experiment in its own right. There remains much investigation to be conducted regarding the tradeoffs involved in passing and sharing data through channels having low communication rates, as well as investigating the means of coordination of many autonomous knowledge sources. Points of interest for systems design also exist in creating the appropriate interfaces between any given group's knowledge source process and the central controlling process. Specification for data base requirements and formats (for both input and output) and specifications for determining the pre-conditions upon which a process should be activated must be easily specified for each new process to be added. In particular, the new process should not need to know the details of the global data structures it may need to access -- the linkage interface should take care of such details (Parnas, 1971,1971a).

Issues of user control over the entire system and the human interface in general are considered vital, demanding much investigation for any system organization which intends to run as a set of parallel cooperating (whether closely- or loosely-coupled) processes. The user must have the ultimate control over halting the entire recognition system or some subset of processes involved therein and interrogating (and perhaps altering) the instantaneous state of any given process. Protocols for debugging and controlling any knowledge source process should be provided via the interface linkage setup. Systems allowing the amount of user control that might be desired are not easily achievable given the current state of the art, primarily due to a general lack of experience in multiprocess environments (however, see Swinehart, 1973). Given a well-defined problem environment such as the speech understanding task, which lends itself readily to a multiple-process decomposition, investigation into the realms of multiprocess debugging and control might now be given more definite aims. Indeed, the problems involved in controlling a set of independent parallel processes that are cooperating to solve a single problem reach beyond the issues raised in the development of present multiprogramming systems (e.g., monitoring and controlling the interactions involving shared data structures and process intercommunications demand that new debugging systems and strategies be formulated).

SUMMARY

The main focus of this paper has been to illustrate the issues of system organization that arise when one attempts to build a general speech understanding system which can equal human performance. In practice, however, one can finesse a large number of these issues by working with pre-recorded data and relaxing other requirements, such as real-time response. However, unless the system is organized with the eventual goals firmly in mind, one is likely to end up with dead-end systems, necessitating a complete reformulation of the problem solution. The complexity of the hardware and software problems raised by real-time requirements explains why there are very few systems which can accept or attempt recognition of live connected speech.

Usually the term "parallel processing" is used as if it will resolve all of one's problems. The intent here is mainly to indicate that speech understanding systems naturally decompose into a set of cooperating, independent processes. Whether one uses a single processor (as we now do) or many processors (as we propose to do), the program structure and organization tends to be similar. The main question, then, is how much computational power is available on the system to attempt real-time recognition of connected speech. The multiprocessor and network organizations provide an opportunity to study and evaluate relative merits of various computer architectures in this context.

Finally, we believe that the issues of system organization raised here are relevant to a large class of current problems in AI, e.g., vision, robotics, chess, chemistry, etc., where performance is the main criterion for acceptability and where many sources of knowledge are available. In particular, the notions of hypothesize-and-test and cooperating independent processes seem equally applicable to these areas as well.

ACKNOWLEDGMENTS

We would like to thank our colleagues C.G. Bell, R. Neely, A. Newell, D. Parnas, and W. Wulf for many interesting discussions on this subject.

BIBLIOGRAPHY

Adams, D.A. (1968), "A Computation Model with Data Flow Sequencing", Tech. Rep. CS-117 (Ph.D. Thesis), Comp. Sci. Dept., Stanford Univ.

Bell, C.G., W. Broadley, W. Wulf, A. Newell, et al. (1971), "C.mmp: The CMU Multi-mini-processor Computer", Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Bell, G., P. Freeman, et al. (1971a), "C.ai: A Computing Environment for AI Research", Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Bell, C.G., R.C. Chen, S.H. Fuller, J. Grason, S. Rege, and and D.P. Siewiorek (1973), "The Architecture and Application of Computer Modules: A Set of Components for Digital Systems Design", COMPCON 73, San Francisco, Ca.

Denning, P.J. (1968), "The Working Set Model for Program Behavior", **Comm. ACM**, 11, 5, 323-333.

Erman, L.D. (1973, in preparation), "An Environment and System for Machine Recognition of Connected Speech", (Ph.D. Thesis), Stanford Univ., to appear as Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Feldman, J.A. and R.F. Sproull (1971), "System Support for the Stanford Hand-Eye System", Second Inter. Joint Conf. on Artificial Intelligence, 183-189.

Lesser, V.R. (1972), "Dynamic Control Structures and Their Use in Emulation", (Ph.D. Thesis), Tech. Rep. CS-309, Comp. Sci. Dept., Stanford Univ.

McCredie, J.W. (1972), "Analytic Models of Time-Shared Computing Systems: New Results, Validations and Uses", (Ph.D. Thesis), Comp. Sci. Dept., Carnegie-Mellon Univ., Chapter 5.

Minsky, M. and S. Papert (1969), **Perceptrons**, MIT Press, Cambridge, Mass.

Neely, R.B. (1973), "On the Use of Syntax and Semantics in a Speech Understanding System", (Ph.D. Thesis), Stanford Univ., to appear as a Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Parnas, D.L. (1971), "Information Distribution Aspects of Design Methodology", Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Parnas, D.L. (1971a), "On the Criteria to be Used in Decomposing Systems into Modules", Tech. Rep., Comp. Sci. Dept., Carnegie-Mellon Univ.

Roberts, L.G. and B.D. Wessler (1970), "Computer Network Development to Achieve Resource Sharing", Proc. SJCC, 36.

Reddy, D.R., L.D. Erman, and R.B. Neely (1973), "A Model and System for Machine Recognition of Speech", **IEEE Trans. Audio and Electro-acoustics**, in press.

Reddy, D.R., L.D. Erman, R.D. Fennell, and R.B. Neely (1973a), "The Hearsay Speech Understanding System: An Example of the Recognition Process", Third International Joint Conference on Artificial Intelligence (this volume).

Simon, H.A. (1962), "The Architecture of Complexity", Proc. Am. Phil. Soc. 106.

Swinehart, D.C. (1973, in preparation), "A Multiple-Process Approach to Interactive Programming Systems", (Ph.D. Thesis), Comp. Sci. Dept., Stanford Univ.

Wulf, W.A. (1972), "C.mmp: A Multi-Mini-Processor", Proc. FJCC.