



(72) NGUYEN, John N., US


(72) PHILLIPS, Michael S., US

(71) APPLIED LANGUAGE TECHNOLOGIES, US

(51) Int.Cl.⁶ G10L 9/02, G06F 17/30

(30) 1995/05/26 (08/451,448) US

(54) **APPAREIL ET PROCÉDE PERMETTANT UNE ADAPTATION
DYNAMIQUE D'UN SYSTÈME DE RECONNAISSANCE
VOCALE A VOCABULAIRE TRES ETENDU, TENANT
COMPTE DE CONTRAINTES IMPOSEES PAR UNE BASE DE
DONNEES DE CE SYSTÈME**
(54) **METHOD AND APPARATUS FOR DYNAMIC ADAPTATION OF
A LARGE VOCABULARY SPEECH RECOGNITION SYSTEM
AND FOR USE OF CONSTRAINTS FROM A DATABASE IN A
LARGE VOCABULARY SPEECH RECOGNITION SYSTEM**



(57) Le vocabulaire d'un système de reconnaissance vocale possédant un vocabulaire très étendu est organisé

(57) The vocabulary of a large-vocabulary speech recognition system is structured to effectuate the rapid





(21) (A1) **2,220,004**
(86) 1996/05/09
(87) 1996/11/28

pour pouvoir effectuer une édition rapide et efficace des mots dans un réseau lexical enregistrant le vocabulaire. Le réseau lexical est structuré pour comporter des noeuds de contraintes phonétiques qui organisent l'information phonétique entre les mots dans le réseau et des noeuds de classes de mots qui organisent l'information syntaxique/sémantique dans le réseau. Des fragments du réseau, correspondant à des prononciations différentes de phonèmes et étiquetés pour spécifier les interconnexions permises, sont passés en revue au préalable pour faciliter la génération de la prononciation de mots nouveaux et donc faciliter une addition rapide de mots au vocabulaire, même pendant l'opération de reconnaissance vocale. Différents modèles de langages et différents vocabulaires pour différentes portions d'un discours sont invoqués d'une manière qui dépend, en partie, de l'historique du système.

and efficient edition of words to a lexical network storing the vocabulary. The lexical network is structured to include Phonetic Constraint Nodes which organize the inter-word phonetic information in the network, and Word Class Nodes which organize the syntactic/semantic information in the network. Network fragments, corresponding to phoneme pronunciations and labeled to specify permitted interconnections, are precompiled to facilitate the rapid generation of pronunciation for new words and thereby enhance the rapid addition of words to the vocabulary even during speech recognition. Different language models and different vocabularies for different portions of a discourse are invoked dependent in part on the discourse history.

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G10L 5/06, 3/00	A3	(11) International Publication Number: WO 96/37881 (43) International Publication Date: 28 November 1996 (28.11.96)
(21) International Application Number: PCT/US96/06634 (22) International Filing Date: 9 May 1996 (09.05.96) (30) Priority Data: 08/451,448 26 May 1995 (26.05.95) US (71) Applicant: APPLIED LANGUAGE TECHNOLOGIES [US/US]; 215 First Street, Cambridge, MA 02142 (US). (72) Inventors: PHILLIPS, Michael, S.; 653 Concord Avenue, Belmont, MA 02178 (US). NGUYEN, John, N.; 55 Hill Road #504, Belmont, MA 02178 (US). (74) Agents: O'DONNELL, Martin, J. et al.; Cesari and McKenna, 30 Rows Wharf, Boston, MA 02110 (US).	(81) Designated States: AM, AU, BB, BG, BR, BY, CA, CN, CZ, EE, GE, HU, IS, JP, KG, KP, KR, KZ, LK, LR, LT, LV, MD, MG, MN, MX, NO, NZ, PL, RO, RU, SG, SI, SK, TJ, TM, TT, UA, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments. (88) Date of publication of the international search report: 16 January 1997 (16.01.97)	

(54) Title: METHOD AND APPARATUS FOR DYNAMIC ADAPTATION OF A LARGE VOCABULARY SPEECH RECOGNITION SYSTEM AND FOR USE OF CONSTRAINTS FROM A DATABASE IN A LARGE VOCABULARY SPEECH RECOGNITION SYSTEM

(57) Abstract

The vocabulary of a large-vocabulary speech recognition system is structured to effectuate the rapid and efficient edition of words to a lexical network storing the vocabulary. The lexical network is structured to include Phonetic Constraint Nodes which organize the inter-word phonetic information in the network, and Word Class Nodes which organize the syntactic/semantic information in the network. Network fragments, corresponding to phoneme pronunciations and labeled to specify permitted interconnections, are precompiled to facilitate the rapid generation of pronunciation for new words and thereby enhance the rapid addition of words to the vocabulary even during speech recognition. Different language models and different vocabularies for different portions of a discourse are invoked dependent in part on the discourse history.



5 **METHOD AND APPARATUS FOR DYNAMIC ADAPTATION OF A LARGE
VOCABULARY SPEECH RECOGNITION SYSTEM AND FOR USE OF
CONSTRAINTS FROM A DATABASE IN A LARGE VOCABULARY SPEECH
RECOGNITION SYSTEM**

BACKGROUND OF THE INVENTION

A. Field of the Invention

10 The invention relates to speech recognition systems and, more particularly, to large-vocabulary speech recognition systems. The system described herein is suitable for use in systems providing interactive natural language discourse.

B. The Prior Art

15 Speech recognition systems convert spoken language to a form that is tractable by a computer. The resultant data string may be used to control a physical system, may be output by the computer in textual form, or may be used in other ways.

 An increasingly popular use of speech recognition systems is to automate transactions requiring interactive exchanges. An example of a system with limited interaction is a
20 telephone directory response system in which the user supplies information of a restricted nature such as the name and address of a telephone subscriber and receives in return the telephone number of that subscriber. An example of a substantially more complex such system is a catalogue sales system in which the user supplies information specific to himself or herself (e.g., name, address, telephone number, special identification number, credit card
25 number, etc.) as well as further information (e.g., nature of item desired, size, color, etc.) and the system in return provides information to the user concerning the desired transaction (e.g., price, availability, shipping date, etc.).

 Recognition of natural, unconstrained speech is very difficult. The difficulty is increased when there is environmental background noise or a noisy channel (e.g., a telephone
30 line). Computer speech recognition systems typically require the task to be simplified in various ways. For example, they may require the speech to be noise-free (e.g., by using a good microphone), they may require the speaker to pause between words, or they may limit

the vocabulary to a small number of words. Even in large-vocabulary systems, the vocabulary is typically defined in advance. The ability to add words to the vocabulary dynamically (i.e., during a discourse) is typically limited, or even nonexistent, due to the significant computing capabilities required to accomplish this task on a real-time basis. The difficulty of real-time speech recognition is dramatically compounded in very large-vocabulary applications (e.g., tens of thousands of words or more).

One example of an interactive speech recognition system under current development is the SUMMIT speech recognition system being developed at M.I.T. This system is described in Zue, V., Seneff, S., Polifroni, J., Phillips, M., Pao, C., Goddeau, D., Glass, J., and Brill, E. "The MIT ATIS System: December 1993 Progress Report," *Proc. ARPA Human Language Technology Workshop*, Princeton, NJ, March 1994, among other papers. Unlike most other systems which are frame-based systems, (the unit of the frame typically being a 10ms portion of speech), the SUMMIT speech recognition system is a segment-based system, the segment typically being a speech sound or phone.

In the SUMMIT system, the acoustic signal representing a speaker's utterances is first converted into an electrical signal for signal processing. The signal processing may include filtering to enhance subsequent recognizability of the signal, remove unwanted noise, etc. The signal is converted to a spectral representation, then divided into segments corresponding to hypothesized boundaries of individual speech sounds (segments). The network of hypothesized segments is then passed to a phonetic classifier whose purpose is to seek to associate each segment with a known "phone" or speech sound identity. Because of uncertainties in the recognition process, each segment is typically associated with a list of several phones, with probabilities associated with each phone. Both the segmentation and the classification are performed in accordance with acoustic models for the possible speech sounds.

The end product of the phonetic classifier is a "lattice" of phones, each phone having a probability associated therewith. The actual words spoken at the input to the recognizer should form a path through this lattice. Because of the uncertainties of the process, there are usually on the order of millions of possible paths to be considered, each of different overall

probability. A major task of the speech recognizer is to associate the segments along paths in the phoneme lattice with words in the recognizer vocabulary to thereby find the best path.

In prior art systems, such as the SUMMIT system, the vocabulary or lexical representation is a "network" that encodes all possible words that the recognizer can identify, all possible pronunciations of these words, and all possible connections between these words. This vocabulary is usually defined in advance, that is, prior to attempting to recognize a given utterance, and is usually fixed during the recognition process. Thus, if a word not already in the system's vocabulary is spoken during a recognition session, the word will not successfully be recognized.

The structure of current lexical representation networks does not readily lend itself to rapid updating when large vocabularies are involved, even when done on an "off-line" basis, that is, in the absence of speech input. In particular, in prior art lexical representations of the type exemplified by the SUMMIT recognition system, the lexical network is formed as a number of separate pronunciation networks for each word in the vocabulary, together with links establishing the possible connections between words. The links are placed based on phonetic rules. In order to add a word to the network, all words presently in the vocabulary must be checked in order to establish phonetic compatibility between the respective nodes before the links are established. This is a computationally intensive problem whose difficulty increases as the size of the vocabulary increases. Thus, the word addition problem is a significant issue in phonetically-based speech recognition systems.

In present speech recognition systems, a precomputed language model is employed during the search through the lexical network to favor sequences of words which are likely to occur in spoken language. The language model can provide the constraint to make a large vocabulary task tractable. This language model is generally precomputed based on the predefined vocabulary, and thus is generally inappropriate for use after adding words to the vocabulary.

SUMMARY OF THE INVENTION

A. Objects of the Invention

Accordingly, it is an object of the invention to provide an improved speech recognition system.

A further object of the invention is to provide a speech recognition system which facilitates the rapid addition of words to the vocabulary of the system.

Still a further object of the invention is to provide an improved speech recognition system which facilitates vocabulary addition during the speech recognition process without
5 appreciably slowing the speech recognition process or disallowing use of a language model.

Yet another object of the invention is to provide a speech recognition system which is particularly suited to active vocabularies on the order of thousands of words and greater and total vocabularies of millions of words and greater..

Still a further object of the invention is to provide a speech recognition system which
10 can use constraints from large databases without appreciably slowing the speech recognition process.

BRIEF DESCRIPTION OF THE INVENTION

In accordance with the present invention, the lexical network containing the
15 vocabulary that the system is capable of recognizing includes a number of constructs (defined herein as "word class" nodes, "phonetic constraint" nodes, and "connection" nodes) in addition to the word begin and end nodes commonly found in speech recognition systems. (A node is a connection point within the lexical network. Nodes may be joined by arcs to form paths through the network. Some of the arcs between nodes specify speech segments,
20 i.e., phones.) These constructs effectively precompile and organize both phonetic and syntactic/semantic information and store it in a readily accessible form in the recognition network. This enables the rapid and efficient addition of words to the vocabulary, even in a large vocabulary system (i.e., thousands of active words) and even on a real-time basis, i.e., during interaction with the user. The present invention preserves the ability to comply with
25 phonetic constraints between words and use a language model in searching the network to thereby enhance recognition accuracy. Thus, a large vocabulary interactive system (i.e., one in which inputs by the speaker elicit responses from the system which in turn elicits further input from the speaker) such as a catalogue sales system can be constructed. The effective vocabulary can be very large, (i.e., millions of words) without requiring a correspondingly

large active (random access) memory because not all the words in it need be “active” (that is, connected into the lexical recognition network) at once.

In accordance with the present invention, the vocabulary is categorized into three classes. The most frequently used words are precompiled into the lexical network; typically, there will be several hundred of such words, connected into the lexical network with their phonetically permissible variations. Words of lesser frequency are stored as phonemic baseforms. A baseform represents an idealized pronunciation of a word, without the variations which in fact occur from one speaker to another and in varying context. The present invention may incorporate several hundred thousand of such baseforms, from which a word network may rapidly be constructed in accordance with the present invention. The least frequently used words are stored as spellings. New words are entered into the system as spellings (e.g., from an electronic database which is updated periodically). To make either one of the least frequently used words or a completely new word active, the system first creates a phonemic baseform from the spelling. It then generates a pronunciation network from the phonemic baseforms in the manner taught by the present invention.

The phonetic constraint nodes (referred to hereinafter as PC nodes or PCNs) organize the inter-word phonetic information in the network. A PC node is a tuple, PC ($x, y, z, \dots$) where the x, y, z are constraints on words that are, or can be, connected to the particular node. For example, x may specify the end phone of a word; y the beginning phone of a word with which it may be connected in accordance with defined phonetic constraints; and z a level of stress required on the following syllable. While tuples of any desired order (the order being the number of constraints specified for the particular PCN) may be used, the invention is most simply described by tuples of order two, e.g., PCN (x, y). Thus, PCN (null, n) may specify a PCN to which a word with a “null” ending (e.g., the dropped “ne” in the word “phone” is connected and which in turn will ultimately connect to words beginning with the phoneme / n /.

Word Class Nodes (referred to hereinafter as WC nodes or WCNs) organize the syntactic/semantic information in the lexical network and further facilitate adding words to the system vocabulary. Examples of word class nodes are parts-of-speech (e.g., noun, pronoun, verb) or semantic classes (e.g., “last name”, “street name”, or “zip code”). Both

the words that form the base vocabulary of the speech recognition system (and therefore are resident in the lexical network to define the vocabulary that the system can recognize), as well as those that are to be added to this vocabulary are associated with predefined word classes.

5 Words are incorporated into the lexical network by connecting their begin and end nodes to WC nodes. The WCNs divide the set of words satisfying a particular PCN constraint into word classes. There may be a general set of these word classes, e.g., nouns, pronouns, verbs, "last name", "street name", "zip code", etc. available for connection to the various PCNs. On connecting a specific instance of a set member (e.g., "noun") to a PCN, it
10 is differentiated by associating a further, more specific characteristic to it, e.g., "noun ending in /n/", "noun ending in "null", etc. Each specific instance of a WCN connects to only one particular PCN. So, for example, there may be a "noun" WCN connected to the (null, n) PCN which is separate from a "noun" WCN connected to the (vowel, n) PCN. To qualify for connection to a given WC node, a word must not only be of the same word class as the
15 WC node to which it is to be connected, but must also be characterized by the same phonetic constraint as the PC node to which the WC node is connected, e.g., noun ending in "null".

The PC nodes are interconnected through word connection nodes (hereinafter referred to as CONN nodes) which define the allowable path between the end node of a word and the begin node of a following word. Effectively, CONN nodes serve as concentrators,
20 that is, they link those PC nodes which terminate a word with those PC nodes which begin a succeeding word which may follow the preceding word under the phonetic constraints that are applicable. These constraints are effectively embedded in the WC nodes, the PC nodes, the CONN nodes, and their interconnections.

In order to add a word to the lexical network, it is necessary first to create a
25 pronunciation network for that word. A given word will typically be subject to a number of different pronunciations, due in part to the phonetic context in which they appear (e.g., the end phoneme /n/ in the word "phone" may be dropped when the following word begins with an /n/ (e.g., "number"), and in part to other factors such as speaker dialect, etc. Variations in pronunciation which are due to differing phonetic context are commonly modeled by standard
30 rules which define, for each phoneme, the ways in which it may be pronounced, depending on

the surrounding context. In the present invention, network fragments corresponding to the operation of these rules on each phoneme are precompiled into binary form and stored in the system, indexed by phoneme. The precompiled network fragments include labels specifying allowed connections to other fragments and associations with PCNs. These labels are of two types: the first refers to the phoneme indexes of other pronunciation networks; the second refers to specific branches within the pronunciation networks which are allowed to connect to the first pronunciation network. Pronunciation networks for phonemes precompiled according to this method allow the rapid generation of pronunciations for new words to thereby facilitate word addition dynamically, i.e., during the speech recognition process itself.

10 In adding a word to the lexical network, the word is associated with a phonemic baseform and a word class. Its pronunciation network is generated by choosing the network fragment associated with phonemes in the phonemic baseform of the word and then interconnecting the fragments according to the constraints at their end nodes. The ensuing structure is a pronunciation network typically having a multiplicity of word begin and word end nodes to allow for variation in the words which precede and follow. The resultant
15 pronunciation network is linked to the word class nodes in the manner described above.

In the present invention, the words are organized by word class, and each added word is required to be associated with a predefined word class in order to allow use of a language model based on word classes during the search of the lexical network, even with added
20 words. Predefined words are not required to belong to a word class; they may be treated individually. The language model comprises functions which define the increment to the score of a path on leaving a particular word class node or word end node or on arriving at a particular word class node or word begin node. A function may depend on both the source node and the destination node.

25 In accordance with the present invention, constraints from electronic databases are used to make the language vocabulary task tractable. The discourse history of speech frequently can also provide useful information as to the likely identification of words yet to be uttered. In the present invention, the discourse history is used in conjunction with a database to invoke different language models and different vocabularies for different portions of the
30 discourse. In many applications, the system will first pose a question to the user with words

drawn from a small-vocabulary domain. The user's response to the question is then used to narrow the vocabulary that needs to be searched for a subsequent discourse involving a large domain vocabulary. As an example, in a catalogue sales system, the system may need to determine the user's address. The system will first ask: "What is your zip code?", and then
5 use the response to fill in the "street name" word class from street names found in the database that have the same zip code as that of the stated address. The street names so determined are quickly added to the system vocabulary, and street names previously in the vocabulary are removed to provide the requisite room in active memory for the street names to be added and to reduce the size of the network to be searched. The language model, i.e.,
10 the probabilities assigned to the various street names so selected, may be established a priori or may be based on other data within the system database, e.g., the number of households on each street, or a combination of these and other information items. Similarly, the system may ask for the caller's phone number first, then use the response and an electronic phonebook database to add to the vocabulary the names and addresses corresponding to the
15 hypothesized phone numbers.

The extensive use of large electronic databases while interacting with the user necessitates an efficient database search strategy, so that the recognition process is not slowed appreciably. In accordance with the present invention, hash tables are employed to index the records in the database, and only that information which is needed for the task at
20 hand is stored with the hash tables.

BRIEF DESCRIPTION OF THE DRAWINGS

For a fuller understanding of the nature and objects of the invention, reference should
25 be had to the following detailed description of the invention, taken in connection with the accompanying drawings, in which:

Fig. 1 is a block and line diagram of a conventional form of speech recognition system;

Fig. 2 is a diagram showing a portion of a lexical network as currently used in speech
30 recognition systems;

Fig. 3 is a diagram showing the word connection portion of a lexical network in accordance with the present invention;

Fig. 4 illustrates the creation of the pronunciation network for a word in accordance with the present invention;

5 Fig. 5 is a flow diagram showing the manner in which words are added to the system vocabulary; and

Fig. 6 is a flow diagram of a lexical network search in accordance with the present invention.

10 DESCRIPTION OF THE PREFERRED EMBODIMENT OF THE INVENTION

In Fig. 1, a speech recognition system 10 comprises a signal processor 12 which converts the acoustic input from a speaker into a form suitable for computer processing, e.g., a sequence of vectors, each representing the frequency spectrum of a portion of the speech. The resultant signals are then analyzed in a segmenter 14 to hypothesize boundaries between the sub-constituents (termed "segments") of the input. These boundaries delineate
15 hypothesized "phones" corresponding to phonemes. (A phoneme is a symbol that designates a basic sound of the language, e.g., the sound of the letter "n"; it is enclosed between slashes: /n/. A phone is a symbol representing the pronunciation of a phoneme; it is enclosed between brackets, e.g., [n]. The ARPABET system of notation will be used in the following
20 to represent phones and phonemes). Several begin- and end-times may be hypothesized for each phone, resulting in overlapping hypothesized regions. The segmentation is performed in accordance with acoustic models 16 available to the system.

The output of the segmenter is applied to a phonetic classifier 18 which generates a phonetic lattice representation of the acoustic input. This lattice describes the various phones
25 corresponding to the hypothesized segments, and the probabilities associated with each. Paths through this lattice, each of which represents a hypothesis of a sequence of phones corresponding to the acoustic input, are compared with possible paths in a corresponding recognition network 20 in a search stage 22. Language models 24 guide the search for the best match between the phonetic paths generated by classifier 18 and the paths traced

through recognition network 20. The resultant is an output 26 representing a recognized communication, e.g., a sentence.

Fig. 2 illustrates a portion of a typical lexical recognition network commonly used in phonetic-based speech recognition systems, and shows how words have commonly been connected in a lexical network prior to the present invention. This figure illustrates typical permitted combinations involving the words "phone", "number", and "and" as an aid in explaining the present invention. Each phone of the respective words is represented as an arc (e.g., arc 30) between two nodes (e.g., node 32 and node 34). Phones are connected sequentially to form a word. For example, the word "and" is formed of the phones /ae/ (arc 36), /n/ (arc 38), and /d/ (arc 40). A particular phoneme may have alternative pronunciations in a word. For example, the final phoneme in the word "number" may be /er/ (arc 42) or /ah/ (arc 44). A phoneme may sometimes be omitted. For example, the final /n/ (arc 50) in the word "phone" is sometimes omitted in speech; this is indicated by alternatively terminating the word in the "null" phone (arc 52).

The permitted connections between words is shown by links connecting the end nodes of one word with the begin nodes of other words. For example, link 60 connects the end node 54 of the word "phone" to the begin node 62 of the word "number", while link 64 connects the end node 54 of the word "phone" to the begin node 66 of the word "and".

It will be noticed that the null end node 56 of the word "phone" has a connection to the begin node 62 of the word "number" via link 68, but has no connection to the begin node 66 of the word "and". This indicates that a pronunciation which drops the final phoneme (/n/) in "phone" when pronouncing the successive words "pho(ne)" and "and" is not permitted, but such a pronunciation would be permitted in pronouncing the string "pho(ne) number".

It should be understood that Fig. 2 is shown in an expanded form for purposes of clarity, that is, the words are repeated in order to show the possible connections between them more clearly. In an actual implementation, the words need not be repeated. Further, it will be noted that a string such as "phone phone" is shown as a possible sequence in Fig. 2, since this is phonetically permissible, although syntactically and semantically unlikely.

In order to add words to a vocabulary structured in the manner shown in Fig. 2, it is necessary to test the compatibility of the begin and end nodes of the words against the begin and end nodes of all other words in order to determine the allowed connections between words. This requires substantial computational processing when the size of the vocabulary is large, and would be unacceptably slow for the user when "on-line" expansion of the vocabulary is sought, that is, when it is sought to add words to the vocabulary during a dialogue with the system. This has been a major impediment to the construction of large vocabulary interactive speech recognition systems.

Fig. 3 shows the structure of the word connection portion of a lexical representation network in accordance with the present invention which solves this problem. In Fig. 3, nodes 80, 82 and 84 are connected to receive the end phoneme of a particular pronunciation of the each of the words "sewn", "phone", and "crayon", respectively. Node 86 is connected to receive the end phoneme of the word "can". The nodes 80-86 are referred to as "end nodes", since they end the word connected to them. Node 80 is connected via an arc 90 to a Word Class node 92; word class nodes organize word pronunciations according to phonetic as well as syntactic and/or semantic information. For example, the class of word pronunciations connected to node 92 might be "past participles ending in /n/ in which the /n/ is not pronounced". Similarly, node 82 is connected via a link 94 to a Word Class node 96 to which node 84 is also connected via a link 98. Word Class node 96 may comprise, for example, the class "nouns ending in /n/ in which the /n/ is not pronounced."

Word Class nodes 92 and 96, in turn, are connected to a phonetic constraint node 100 via arcs 102, 104, respectively. Phonetic constraint nodes embody phonetic constraints. For example, phonetic constraint node 100 may comprise the constraint pair (null, n), indicating that nodes connected into it from the left are characterized by "null" final phoneme, while those to which it connects to the right are characterized by an initial /n/. In most instances word end nodes are connected to phonetic constraint nodes through word class nodes. However, some words in the lexical network may in effect form their own word class. Such words are directly connected to a phonetic constraint node. Such is the case with the word "can", whose end node 86 is directly connected to a phonetic constraint node 106 via a link 108.

Phonetic constraint nodes 100 and 110 feed into a word connection node 120 via links 122, 124, respectively. Similarly, node 106 feeds into connection node 126 via an arc 128. The pattern outward from the connection nodes is the reverse of the pattern inward to the nodes, that is, connection node 120 expands to phonetic constraint nodes 130, 132, 134 via links 136, 138 and 140, respectively. Similarly, connection node 126 expands to phonetic constraint node 142 via a link 144, as well as to nodes 143, 145 via links 147, 149, respectively. Phonetic constraint node 130 is connected to a Word Class node 150 via a link 152 and then is connected to the begin node 154 of a qualified word pronunciation (e.g., a typical pronunciation of the word "number") via a link 156. The word "qualified" indicates that the begin node to which a connection node leads satisfies the phonetic constraints associated with the end nodes feeding into that connection node. In Fig. 3, for example, the connection node 120 is linked to words normally ending in "n" in which the final "n" may be dropped if the following word begins with "n". Thus, the final "n" in the word "phone" may be dropped when the word "phone" is followed by the word "number". In similar fashion, phonetic constraint node 132 is connected via arc 160 to Word Class node 162 and then via arc 164 to word begin node 166. Node 132 is also connected via arc 168 to Word Class node 170 and then via arc 172 to word begin node 174. Finally, phonetic constraint node 142 is directly connected to word begin node 176 via arc 178 without connection through an intermediary Word Class node.

The structure of the lexical networks shown in Fig. 3 greatly facilitates the dynamic addition of words to the lexical network, that is, the addition of such words during the progress of a discourse and concurrent with speech recognition. Specifically, permissible word classes and phonetic constraints are defined in advance and the connections between the Word Class nodes, the phonetic constraint nodes and the connection nodes are precompiled and stored in the lexical network of the system prior to speech recognition. The end nodes and begin nodes of a defined "base" vocabulary are connected to the Word Class nodes in the manner illustrated in Fig. 3 in cases where such words belong to the defined Word Classes, and are connected directly to phonetic constraint nodes otherwise (e.g., when the word is of sufficient frequency to, in effect, form its own Word Class).

Words which are not precompiled in the baseform lexical network are added to the network dynamically by first forming the pronunciation network of the word (described in more detail in connection with Fig. 4) and then connecting the end nodes and the begin nodes of the words to the appropriate Word Class nodes in accordance with the word-class and phonetic constraints embodied in the respective nodes. For example, a pronunciation of a noun ending in "n" in which the final "n" may be dropped if the following word begins with "n" would be connected to a noun Word Class node which itself is connected to a phonetic constraint node having the constraint (null, n) which specifies the constraint of a word terminating in a "null" sound which must be followed by a word beginning with "n".

Each word end and word begin node contains the index of the PCN it connects to (through a WCN). The present invention performs a simple lookup (thereby eliminating the need for run-time computations) to establish phonetic compatibility between the end nodes and begin nodes of the words being added to the system vocabulary and thus enables the addition of the words in "real time", that is, concurrent with interaction with the user of the system. This in turn can dramatically reduce the size of the vocabulary that must be precompiled and stored in active (random access) memory, and thus enables expansion of the effective range of the vocabulary for a given system, since the needed vocabulary can be formed "on the fly" from word spellings or from phonemic baseforms requiring rather limited storage capacity in comparison to that required for a vocabulary which is completely precompiled and integrated into the lexical network.

Fig. 4A shows the manner in which pronunciation networks are generated from phonemic baseforms and stored in accordance with the present invention. In particular, the word "phone" has a phonemic base form /f ow n/. A phonemic baseform is the "ideal" pronunciation of the word, i.e., without the slurring that occurs in rapid or casual speech.

The network fragment for each phoneme of the word is created by applying rules which embody phonetic constraints that depend on the context of the phoneme, that is, on the surrounding phonemes. As illustrated in Fig. 4A, the phoneme /f/ may be pronounced as [f] or as [f pause] if it is followed by a vowel; is pronounced as [f] or is omitted (null pronunciation) if it is followed by /f/; and is pronounced as [f] if it is followed by any phoneme except another /f/ or a vowel. Similarly, the phoneme /ow/ is pronounced as [ow]

or as [ow w] whenever followed by a vowel, and is pronounced as [ow] in all other cases. Finally, the phoneme /n/ is pronounced as [n] or is omitted (null pronunciation) when followed by /n/, and is pronounced as [n] in all other cases. The phonetic fragments corresponding to the individual phonemes of the word “phone”, as well as the resultant word
5 formed from connecting the fragments together in accordance with the phonetic constraints embedded in them, is shown in Fig. 4A.

In accordance with the present invention, the phonetic fragments are precompiled (that is, in advance of speech recognition) and stored in binary form. They are indexed by phoneme for rapid retrieval. Given a phoneme’s baseform, the appropriate phonetic
10 fragments are strung together to form words. It will be noted from Figure 4A that the words thus formed may have multiple connections; these may occur at the beginning of the word, at the end, or both, and are referred to as “frayed edges”. Each node at a frayed edge corresponds to a word begin or word end node of a particular pronunciation of the word. Each word begin or word end node connects to exactly one PCN (through a WCN in the
15 case of dynamically added words).

The phonetic constraints according to which the fragments are linked to each other and to PCNs (through WCNs) are expressed in terms of the phoneme indexes. In some cases, further restrictions may need to be imposed on the connections of the word into the lexical network. For example, the phoneme /ax/ (“schwa” in ARPABET notation) may at
20 some times be pronounced as the phone [ax] and at other times as the phone [null]. Similarly, the phoneme /l/ may at times be pronounced as the phone [l] and at other times as the phone [el] (syllabic /l/). (An example is the differing pronunciations of the word “vowel” as [v aw ax l] or alternatively as [v aw el].) The conditions under which the respective pronunciations are permissible are shown in Fig. 4B. For the standard English pronunciation
25 of /ax l/, the combination [ax][l] is permissible, as is the combination [null][el], but the combinations [ax][el] and [null][l] are not. In accordance with the present invention, the permissible combinations may be formed, and the impermissible combinations prevented, by “tagging” the “ends” of a given fragment and restricting the connection of a fragment end to those nodes having not the permitted phonetic constraint but also the same tag as the
30 fragment end to which they are to be connected.

This is illustrated more clearly in Figure 4B which shows the phonetic fragments for the phoneme /ax/ when it is followed by any phoneme except the phoneme /l/ as well as when it is followed by /l/. Figure 4B also shows the phonetic fragment for the phoneme /l/, both when it is preceded by the phoneme /ax/ and otherwise. The fragment corresponding to the phoneme /ax/ followed by the phoneme /l/ has tags "a" and "b" associated with the alternate pronunciations of phones [ax] and [null], respectively. Similarly, the fragment corresponding to the phoneme /l/ preceded by the phoneme /ax/ has tags "a" and "b" associated with the alternate pronunciations [l] and [el], respectively. When connecting the fragments together in the sequence /ax/ /l/, a node tagged as "a" in the phoneme /ax/ can connect only to a node similarly tagged as "a" in the phoneme /l/, and the same is true for the nodes tagged as "b" in these phonemes. It will be understood that this structure is used between words, as well as within words, that is, this structure is also used in connecting words into the lexical network where appropriate. Thus, the connection into the network may be done on the basis of matching the phonetic constraints (based on phoneme indexes) at the begin and end nodes of a word to the phonetic constraints of a word class or phonetic constraint node, while at the same time preserving the further constraints that apply to the particular connection. This facilitates rapid addition of words to the lexical vocabulary.

During speech recognition, the system searches through the lexical network to find matches between the sounds emitted by the speaker and words stored in the system. As illustrated in Fig. 5, if a word candidate is found to be stored in the system at the time of recognition (step 190), the system proceeds with the recognition process (step 192). If it is not found, the system checks to determine whether a baseform of the word is stored in the system (step 194). If it is, the system forms the word from the baseform as described in connection with Figure 4, connects it into the lexical network as described in connection with Figure 3, and proceeds with the recognition process. If a baseform of the word being sought is not stored in the system, the baseform is first created in accordance with defined text-to-phoneme rules (step 198).

Turning now to Fig. 6, a flow diagram of the search process through the lexical network is shown. The form of the search overall is of the type described by Viterbi, but expanded to use the capabilities of the present invention. The goal of the search is to find the

most probable path through the lexical network at each stage in the network (i.e., at each node in the network). In achieving this, the cumulative probability of occurrence of the linguistic condition associated with a given node is recorded and stored for subsequent use.

As an example, refer again to Fig. 3. The end nodes 80, 82, 84, and 88 have "scores" associated with them of -75, -50, -100, and -80, respectively. These numbers are proportional to the log of the probabilities at the respective nodes of the most probable paths into those nodes. The task is to propagate those probabilities through the lexical network to the "begin" nodes 154, 166, etc. This is accomplished as follows:

The first step (Fig. 6, step 200) is to maximize the scores from the various end nodes into the left-side word class nodes. This is accomplished by choosing the maximum score, at each word class node, from the scores of the end nodes connected to the particular word class node. Thus, in Fig. 3, node 96 is connected to both node 82 (with a score of -50) and node 84 (with a score of -100). Accordingly, the higher of these scores (-50) is propagated to node 96. Similarly, node 92 is shown as being connected to only a single end node, node 80, which has a score of -75. Accordingly, the score -75 is propagated to node 92.

Next (Fig. 6, step 202), the scores into the left side phonetic constraint nodes are maximized. In the algorithm for calculating the scores, the general functions $f_l()$, $f_r()$, $g()$, and $h()$ are used. The functions take as arguments a node (actually, the word or word class associated with the node) or a pair of nodes. By specifying these functions, different language models can be implemented. The score into a left-side phonetic constraint pair comprises the left side entering score s_l plus a function $f_l(u)$, where s_l is the end node score in the case where the source node (that is, the node preceding the phonetic constraint node) is an end node and s_l is the score of the word class node when the source node is a word class node. Similarly, $f_l(u)$ is a function defined on the end node u in the case where the source node is an end node, and defined on the word class node u in the case where the source node is a word class node. Typically, $f_l(u)$ is a probability, i.e., it defines the likelihood of occurrence of the condition associated with a given node. For example, it may define the likelihood of a particular word class within the system vocabulary. The maximum value of the score so computed for each left-side phonetic constraint node is then determined and assigned to the corresponding phonetic constraint node. For example, node 100 is connected

to node 92, which has a score of -75, and also to node 96, which has a score of -50. The score assigned to node 100 is then determined by selecting the maximum of $-75 + f_i(92)$ (the total score of node 92) and $-50 + f_i(96)$ (the total score of node 96). For purposes of simplicity only, $f_i()$ will be chosen to be zero in this example, and the score -50 is assigned to PCN 100.

The maximum phonetic constraint node scores into the connection nodes are then determined (Fig. 6, step 204). Thus, in Fig. 3, node 100 has a score of -50 and this can be seen to be larger than the score from any other node entering into connection node 120. Accordingly, node 120 is assigned the score -50. The scores for the other connection nodes are computed similarly.

The next step (Fig. 6, step 206) is to propagate the connection node scores to the right side phonetic constraint nodes. Thus, the score -50 for node 120 is propagated to nodes 130, 132, and 134 which are connected to node 120. The scores from the other connection nodes in the network are similarly propagated to the phonetic constraint nodes to which they are respectively connected.

The scores into the right side word class nodes are then maximized (Fig. 6, step 208) in those cases in which the phonetic constraint nodes are connected to the begin nodes via word class nodes; in cases where the phoneme constraint pairs are connected directly to the begin nodes without first traversing intermediate word class nodes, the score into the begin nodes are maximized. For example, in Fig. 3, phonetic constraint node 130 is connected to begin node 154 via word class node 150. Conversely, phonetic constraint node 134 is connected directly to node 120 without traversing a word class node.

The score which is being maximized can be expressed as $s_r + f_r(u)$ where s_r is the score of the right side phonetic constraint node which is connected to either the word class node or the begin node, and $f_r(u)$ is a function that specifies a score increment (typically a log probability) associated with the node u (either word class node or begin node) to which the phonetic constraint node in question is connected. For example, the function $f_r(u)$ may be defined as $\log(C(u)/N)$, where $C(u)$ is the word count in a representative text of the word or word class associated with the particular node u and N is the total number of words in the representative text.

For purposes of illustration, assume that word class node 150 comprises the class of singular nouns beginning with the letter "n" and that in a sample text of one million words, ten thousand such nouns were found. The function $f_r(u) = f_r(150)$ would then be evaluated as $f_r(150) = \log(10,000 / 1,000,000) = -2$. Accordingly, the score at this node is given as $s_r + f_r(u) = -50 + (-2) = -52$. This score is shown at node 150 in Fig. 3 but "crossed out" for reasons which will shortly be explained.

Phonetic constraint node 142 is directly linked to begin node 176, and does not pass through any intermediate word class node. Accordingly, the score of begin node 176 is evaluated via the function $s_r + f_r(u)$. As shown in Fig. 3, the value of s_r with respect to node 176 is the value of node 142, namely, -25. Note that, in Fig. 3, node 176 is shown as the beginning of the word "the". Assume that, in a text of one million words, the word "the" occurs ten thousand times. The function $f_r(176)$ is then evaluated as $f_r(176) = \log \left[\frac{10,000}{1,000,000} \right] = -2$ and the score at node 176 is then $-25 - 2 = -27$.

Once the maximum scores from all PCNs into the begin nodes or WCNs are evaluated, the system then determines whether any bigram statistics for the word pairs in question exist in the language model. In accordance with the present invention, the maximum score for each bigram pair (f, t) is determined as $s_r + g(f, t)$, where s_r is the score of a "from" node (i.e., a word class node in cases where an end node is connected through a word class node, and an end node otherwise) on the left side of a connection node (see Fig. 3) and $g(f, t)$ is a function defined jointly on the "from" node (i.e., a word class node in cases where an end node is connected through a word class node, and an end node otherwise) and on a "to" node (i.e., a word class node in cases where a begin node is connected through a word class node, and a begin node otherwise) on the right side of a connection node that connects the respective "from" and "to" nodes.

As a specific example to illustrate this, consider the relationship between "from" node 96 (e.g., the word class of "nouns ending in "n" but pronounced with the n omitted") and "to" node 150 (the word class of "singular nouns starting with "n"). $g(f, t)$ may be defined as $\log [C(f, t) / C(f)]$, where $C(f, t)$ is the count of the word class pair f, t and $C(f)$ is the count of word class f in some representative text. Assume that it has been determined that $g(96, 150)$

= $\log [C(96, 150) / C(96)] = -1$. The value of s_f at node 96 is -50, as shown. Thus, the function $s_f + g(f, t) = -50 + (-1) = -51$. Since this is greater than the prior value, -52, associated with node 150, node 150 takes on the new value -51. Assuming that all other paths to node 150 through CONN node 120 lead to lower scores, -51 is the final value for node 150.

In cases where the begin nodes are reached through right-side word class nodes, as opposed to directly from right-side phonetic constraint nodes, a final adjustment to the score is performed (Fig. 6, step 212). This adjustment comprises determination of the function $s_w + h(b, w)$, where s_w is the score of the word class node that has so far been determined and $h(b, w)$ is a function which depends on the source word class node and the destination word begin node. In the present example, $h(b, w)$ may be defined as the frequency of occurrence of the word at the begin node divided by the frequency of occurrence of the word class node in question. For example, for word class node 150 and begin node 154, the previously computed score, s_w , is -51. The function $h(b, w) = h(154, 150) = \log [C(154, 150) / C(150)]$ where $C(154, 150)$ is the count of occurrences of word 154 when it belonged to word class 150. If there are 1,000 instances of the word "number" used as a noun in a given text, and 10,000 instances of the word class "singular noun starting with 'n'", the function $h(154, 150)$ has a value of -1. The final score at begin node 154 is then $-51 + (-1) = -52$. Similar calculations occur for the other nodes.

Having determined the maximum scores at the begin nodes, the paths which defined those scores are selected as the most likely connection between the words terminating at the end nodes and the following words beginning at the begin nodes. The search process is then continued in similar manner until the complete path through the lexical network is determined.

The ability to add words dynamically (i.e., concurrent with speech recognition) to the system vocabulary enables the system to use to advantage various constraints which are inherent in a domain-specific database, and thereby greatly facilitates recognition of speech input. For example, consider the case of an interactive sales-order system to which the present system may be applied. In such a database, the vocabulary includes, at least in part, records of associated fields containing elements of information pertinent to the user. During

discourse with the user, the system prompts the user for various spoken inputs and, on recognizing the inputs, provides information or further prompts to the user in order to process an order. Because of the large number of elements (e.g., "customer names") that may need to be included in the vocabulary to be recognized, it may be impractical to have all potential elements of a field "active" (i.e., available for recognition) at one time. This problem is addressed in accordance with the present invention by using a related field in the record to limit the words which must be added to the active lexical network in order to perform the requisite spoken word recognition.

For example, as shown in Fig. 7, a record 300 may include a user's zip code, 302; his or her address, 304; the user's last name, 306; and the user's first name, 308. The fields 302, 304, 306, and 308 are related to or "associated with" each other in that they pertain to the same user. The zip code consists of a vocabulary of ten words (the ten digits). Furthermore, the number of users associated with a zip code, or several hypothesized zip codes is smaller than the total number of users. In accordance with the present invention, the user is first prompted for his or her zip code. The speech recognition system then determines the zip codes that are most likely to correspond to the actual zip code spoken by the user and adds to the lexical recognition network those last names corresponding to the zip codes so determined. This dramatically limits the number of names that must be added to the recognition network at a given time and therefore facilitates recognition of the last name. This capability is therefore expected to be widely useful.

CLAIMS

- 1 1. In a speech recognition system, means for dynamically adding words to an existing
2 recognition vocabulary defined by a lexical network, said means comprising:
- 3 A. means defining a plurality of phonetic constraint nodes characterized by
4 phonetic constraint tuples (x, y, z, ...) of order two or greater, where x, y, z, ..
5 are constraints on the word to be added to the vocabulary;
- 6 B. means defining a plurality of word class nodes, each characterized by a word
7 class and a phonetic constraint;
- 8 C. means defining a plurality of connection nodes, each characterized by a
9 phonetic constraint;
- 10 D. means for establishing links between said word class nodes and said phonetic
11 constraint nodes, and between said phonetic constraint nodes and said
12 connection nodes, for a base vocabulary, prior to the time of dynamic word
13 addition to said lexical network, in accordance with a predefined congruence
14 between the phonetic constraints of said word class nodes, said phonetic
15 constraint nodes and said connection nodes; and
- 16 E. means for adding to said network words defined by at least beginning and
17 ending phoneme nodes and characterized by at least one phonetic constraint,
18 said means comprising means for establishing links between said beginning
19 and ending phoneme nodes and said word class nodes in accordance with a
20 predefined congruence between the phonetic constraints of said words and
21 those of said word class nodes.
- 1 2. A speech recognition system according to claim 1 in which said phonetic constraint
2 tuples define at least a first phonetic constraint on a word and at least a second phonetic
3 constraint on an adjacent word in a path through said lexical network.
- 1 3. A speech recognition system according to claim 1 in which said phonetic constraints
2 comprise phonemes chosen from the class comprising the first and the last phoneme of said
3 word and said adjacent word, respectively.

1 4. A speech recognition system according to claim 1 in which said phonetic constraint
2 tuples define at least the last phoneme of a first word included in, or to be added to, the
3 lexical network and the first phoneme of a second word included in, or to be added to, the
4 lexical network in a path adjacent said first word.

1 5. A speech recognition system including means for rapidly augmenting a recognition
2 vocabulary, comprising:

- 3 A. means for establishing a lexical network formed of phonemic sequences
4 defining the words forming said vocabulary;
- 5 B. means forming constraint nodes defining constraints on the permitted
6 interconnection between selected phonemic sequences in said vocabulary;
- 7 C. means for establishing links between words in said vocabulary only through
8 said constraint nodes.

1 6. A speech recognition system according to claim 5 in which said constraints comprise
2 constraints chosen from the class of phonetic, syntactic, and semantic constraints.

1 7. A speech recognition system according to claim 6 in which said phonetic constraints
2 comprise segmental or suprasegmental phonetic constraints.

1 8. A speech recognition system according to claim 6 in which said constraint nodes
2 include at least

- 3 A. word class nodes for incorporating syntactic or semantic information in said
4 network;
- 5 B. phonetic constraint nodes characterized by at least one phoneme from each of
6 a pair of words and defining a permissible word-pair sequence; and
- 7 C. means for connecting to said word class nodes words having the same word
8 class, and having congruent phonetic constraints, as the nodes to which they
9 are connected.

1 9. A speech recognition system according to claim 8 which further includes connection
2 nodes for connecting congruent phonetic constraint nodes.

1 10. A speech recognition system according to claim 9 in which the connection between
2 said phonetic constraint nodes and said connection nodes are formed concurrent with the
3 formation of the base vocabulary.

1 11. A speech recognition system according to claim 10 which includes means for
2 dynamically adding words to said vocabulary by connecting said words to word class nodes
3 having the same phonetic constraint tuple and the same word class.

1 12. A speech understanding system according to claim 11 in which said words are
2 connected to said word class nodes via word end nodes and word begin nodes.

1 13. A method for rapidly generating word pronunciation networks for use in a speech
2 recognition system, comprising the steps of :

3 A. storing a set of phonemes to be recognized;

4 B. for each phoneme, storing a set of pronunciation networks that have been
5 formed in accordance with predetermined rules that define alternative
6 pronunciations of the phoneme dependent on phonemes with which it may be
7 connected;

8 C. creating a pronunciation network for a given word by selecting a permissible
9 pronunciation network for each phoneme in said word in accordance with the
10 specific phonemes in the word and linking the pronunciation networks so
11 selected.

1 14. The method of claim 13 in which said pronunciation networks are indexed in storage
2 by phoneme.

1 15. The method of claim 13 in which said pronunciation networks are stored in the form
2 of binary records that include fields defining permissible connections to other phonemes.

1 16. The method of claim 13 in which said pronunciation networks are stored in the form
2 of binary records that include fields defining permissible connections to preceding and
3 succeeding phonemes.

1 17. The method of claim 13 in which said pronunciation network is created at least in part
2 concurrent with recognition of a word not previously included in said network.

1 18. In a speech recognition system, a method of efficiently searching a database for a
2 word to be recognized, comprising:

3 A. providing a plurality of databases specific to a corresponding plurality of
4 discourse states, each said database being utilized in connection with its
5 corresponding discourse state;

6 B. restricting search in a second database in accordance with data obtained in
7 response to searching a first database.

1 19. The method of claim 18 which includes the step of establishing the likelihood of a
2 word or word sequence in at least one of said databases on the basis of data obtained from
3 another of said databases.

1 20. A method of efficiently generating a word string in a phoneme-based speech
2 recognition system, comprising the steps of:

3 A. storing word forms of the speech being recognized in the form of a lexical
4 network in which said word forms comprise nodes;

5 B. linking at least some of said word forms into word class nodes, each word
6 class node defined by a common phonetic constraint associated with the word
7 form linked to said node;

8 C. linking said word class nodes into phonetic constraint nodes in accordance
9 with at least one common phonetic constraint of said word class nodes;

10 D. defining functions to determine path scores on word class nodes and phonetic
11 constraint nodes in accordance with linguistic constraints; and

12 E. traversing alternative paths through said network to determine those paths
13 associated with the highest scores.

1 21. The method of claim 20 which includes the step of defining further probabilities on
2 said word class nodes in accordance with syntactic constraints to thereby further constrain
3 the path determination.

1 22. The method of claim 21 in which the further precede defined on each word class
2 defines the probability of occurrence of the word class in the vocabulary of the system.

1 23. The method of claim 21 which includes the step of defining still further probabilities
2 on selected ones of said word class nodes in accordance with the conditional likelihood of
3 occurrence of said nodes given the occurrence of other word class nodes in the vocabulary of
4 the system.

1 24. The method of claim 21 which includes the step of defining still further probabilities
2 on said word class nodes in accordance with the conditional likelihood of occurrence of
3 nodes to which said word class nodes are connected, given the occurrence of said word class
4 nodes.

1 25. In a speech recognition system having a vocabulary including a plurality of words
2 defining records of associated fields, each field containing information elements specific to
3 the field, a method of facilitating recognition of spoken input, comprising the steps of:



4 A. identifying at least one candidate word corresponding to a first information
5 element in response to a spoken input to the system;

6 B. adding to the vocabulary, concurrent with discourse and responsive to said at
7 least one candidate word, a plurality of words associated with a second
8 information element of said records to thereby facilitate recognition of words
9 associated with said second information element on subsequent input to the
10 system.

1 26. The method of claim 25 in which said first information element stands in a one-to-
2 many relation to said second information element.

1 27. The method of claim 26 in which said first information element comprises a zip code
2 and said second information element comprises a proper name.


1 28. The method of claim 26 in which the number of said first information elements in the
2 vocabulary is less than the number of said second information elements.

**FIG. 1**

300

FIG. 7

2/7

**FIG. 2**

SUBSTITUTE SHEET (RULE 26)

317








FIG. 3
SUBSTITUTE SHEET (RULE 26)

**FIG. 4A****SUBSTITUTE SHEET (RULE 26)**


5/7

**FIG. 4B****FIG. 4A cont.**

6/7

**FIG. 5**

7/7

**FIG. 6**

