

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 August 2007 (30.08.2007)

PCT

(10) International Publication Number
WO 2007/098321 A1

(51) International Patent Classification:
G10L 15/18 (2006.01) *G10L 15/14* (2006.01)

(74) Agents: ISAACSON, Thomas, M. et al.; AT & T Corp.,
Room 2A-207, One AT & T Way, Bedminster, NJ 07921
(US).

(21) International Application Number:
PCT/US2007/061911

(22) International Filing Date: 9 February 2007 (09.02.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/276,151 16 February 2006 (16.02.2006) US

(71) Applicant (for all designated States except US): AT & T
CORP. [US/US]; 32 Avenue Of the Americas, New York,
NY 10013-2412 (US).

(72) Inventors; and


(75) Inventors/Applicants (for US only): BLEWETT,
Charles, Douglas [US/US]; 15 Niles Avenue, Madison,
NJ 07940 (US). BOCCHIERI, Enrico, Luigi [US/US]; 1
Wexford Road, Chatham, NJ 07928 (US).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS,
JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS,
LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY,
MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS,
RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: A SYSTEM AND METHOD FOR PROVIDING LARGE VOCABULARY SPEECH PROCESSING BASED ON
FIXED-POINT ARITHMETIC



(57) Abstract: Disclosed herein is a system, method and computer-readable medium storing instructions for controlling a computing device according to the method. As an example embodiment, the method uses a speech recognition decoder that operates or uses fixed-point arithmetic. The exemplary method comprises representing arc costs associated with at least one finite-state transducer (FST) in fixed-point, representing parameters associated with a hidden Markov model (HMM) in fixed-point and processing speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters. The method may also include computing at the decoder sentence hypothesis probabilities with fixed-point arithmetic as type Q-2e numbers.

WO 2007/098321 A1



Declaration under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

Published:

— *with international search report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

PATENT APPLICATION

**A SYSTEM AND METHOD FOR PROVIDING LARGE VOCABULARY
SPEECH PROCESSING BASED ON FIXED-POINT ARITHMETIC**

Inventors:

Charles Douglas Blewett

Enrico Bocchieri

Filing Date:

Prepared By:

The Law Office of Thomas M. Isaacson
10415 Southern Maryland Blvd.
Dunkirk, MD 20754
(410) 286-9405

A SYSTEM AND METHOD FOR PROVIDING LARGE VOCABULARY SPEECH PROCESSING BASED ON FIXED-POINT ARITHMETIC

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The present invention relates to a system and method of providing large vocabulary speech processing and more specifically to providing speech processing such as automatic speech recognition based on fixed-point arithmetic.

2. Introduction

[0002] Large-vocabulary continuous-speech recognition (LVCSR) finds wide use in consumer, military and industrial applications using embedded platforms, such as PDA's, telephone handsets, network appliances, and wearable computers. Often a speech recognition module is part of an overall spoken dialog system that includes various modules to receive speech from a user, recognize the speech (via a speech recognition module), understand the meaning of the speech or the intent of the user (via a spoken language understanding module), formulate responsive text (via a dialog management module) and generate responsive speech (via a text-to-speech module). These and variations of these modules are known in the art for carrying out a natural language spoken dialog with a person. Some systems may not utilize all of these modules but only utilize one or two, such as just providing speech recognition to convert speech to text.

[0003] An example application is in Short Message Service (SMS) that has an expected global volume in excess of 1,000 billion messages in 2005. LVCSR on embedded platforms presents a unique set of challenges. In particular, to lower hardware cost and power consumption, for longer battery life and miniaturization, the CPU's on small portable devices do not have floating-point arithmetic units. However their computational power is constantly increasing, which motivates the study of methods of enabling speech recognition on smaller devices. Traditionally, larger computers and

servers have hardware floating point units either in the CPU or in separate floating point processor chips. Accordingly, what is needed in the art is an improvement that enables highly compute intensive algorithms such as those utilized in speech processing to be able to function on smaller devices, such as portable computing devices, that do not have the computing power or size to perform floating point operations.

SUMMARY OF THE INVENTION

[0004] Additional features and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth herein.

[0005] The present invention provides a fixed-point decoding approach to speech processing modules. Previous work on fixed-point decoding concern either small-vocabulary continuous-speech tasks or large-vocabulary tasks with deterministic grammars. Instead, the present inventors present fixed-point algorithms for LVCSR, including word N-gram language models, frame-synchronous beam-search Viterbi decoder and HMM likelihood computation. The concepts disclosed herein provide floating-point systems in LVCSR experiments on various tasks, with different feature extraction front-ends. The fixed-point parameters do not require task-dependent or feature-dependent calibrations. Experiments are run on the Darpa Switchboard task and on fluently spoken telephone speech from an AT&T customer care application, with up to fifty words/sentence. Even for these long utterances, the accumulation of log-likelihoods scores during fixed-point decoding is not problematic. An example target

computing device for the present invention is a 32-bit integer CPU's (e.g. StrongARM), but the approach may be suitable for 16-bit CPU's with 32-bit accumulators as well or other varieties of CPUs that can utilize the features of the invention. However, any computing device is contemplated for use with this invention.

[0006] The increasing computational power of embedded CPU's motivates the implementation of highly accurate large-vocabulary continuous-speech (LVCSR) algorithms in fixed-point arithmetic, to achieve the same performance on the device as on the server. Disclosed herein are example algorithms for the fixed-point implementation of the frame-synchronous beam-search Viterbi decoder, HMM likelihood computation, and language models (including word N-grams), that yield the same accuracy as floating-point recognizer in LVCSR experiments on the DARPA Switchboard task and on an AT&T proprietary task. Experiments are presented on the DARPA Resource Management task, performed in an embedded system based on the StrongARM-1100 206 MHz CPU.

[0007] The invention relates to a system, method and computer-readable medium storing instructions for controlling a computing device according to the method. As an example embodiment, the method uses a speech recognition decoder that operates or uses fixed-point arithmetic. The exemplary method comprises representing arc costs associated with at least one finite-state transducer (FST) in fixed-point, representing parameters associated with a hidden Markov model (HMM) in fixed-point and processing speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters. The method may also include computing at the decoder sentence hypothesis probabilities with fixed-point arithmetic as type Q-2e numbers.

[0008] In another aspect of the invention, the method relates to performing speech processing by generating a fixed-point speech recognition transducer by converting a

floating point speech recognition transducer via quantization into a fixed-point format, receiving input speech to the fixed-point speech recognition transducer and generating a best hypothesis of the received input speech using the fixed-point speech recognition transducer. The idea is to represent the "cost" of the speech recognition transducers as fixed-point numbers in a Q-2e format, where "e" is a specified parameter, and to accumulate the decoder log-likelihoods as Q-2e numbers.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0010] FIG. 1 illustrates an example system embodiment of the invention;

[0011] FIG. 2 illustrates the automatic speech recognition (ASR) system in floating-point;

[0012] FIG. 3 illustrates the ASR system in fixed-point;

[0013] FIG. 4 illustrates the L (lexicon) FST;

[0014] FIG. 5 illustrates the G (grammar or language model) FST;

[0015] FIG. 6 illustrates a method embodiment of the invention; and

[0016] FIG. 7 illustrates an accuracy versus recognition time on a CPU.

DETAILED DESCRIPTION OF THE INVENTION

[0017] Various embodiments of the invention are discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without parting from the spirit and scope of the invention. Embodiments may include a system (portable computing device, laptop, computer server, computer cluster or grid, etc.)

[0018] Fig. 1 illustrates a block diagram of an exemplary processing device 100 which may be used to implement systems and methods consistent with the principles of the invention. Processing device 100 may include a bus 110, a processor 120, a memory 130, a read only memory (ROM) 140, a storage device 150, an input device 160, an output device 170, and a communication interface 180. Bus 110 may permit communication among the components of processing device 100.

[0019] Processor 120 may include at least one conventional processor or microprocessor that interprets and executes instructions. Memory 130 may be a random access memory (RAM) or another type of dynamic storage device that stores information and instructions for execution by processor 120. Memory 130 may also store temporary variables or other intermediate information used during execution of instructions by processor 120. ROM 140 may include a conventional ROM device or another type of static storage device that stores static information and instructions for processor 120. Storage device 150 may include any type of media, such as, for example, magnetic or optical recording media and its corresponding drive. In some implementations consistent with the principles of the invention, storage device 150 may store and retrieve data according to a database management system.

[0020] Input device 160 may include one or more conventional mechanisms that permit a user to input information to system 200, such as a keyboard, a mouse, a pen, a voice

recognition device, a microphone, a headset, etc. Output device 170 may include one or more conventional mechanisms that output information to the user, including a display, a printer, one or more speakers, a headset, or a medium, such as a memory, or a magnetic or optical disk and a corresponding disk drive. Communication interface 180 may include any transceiver-like mechanism that enables processing device 100 to communicate via a network. For example, communication interface 180 may include a modem, or an Ethernet interface for communicating via a local area network (LAN). Alternatively, communication interface 180 may include other mechanisms for communicating with other devices and/or systems via wired, wireless or optical connections.

[0021] Processing device 100 may perform such functions in response to processor 120 executing sequences of instructions contained in a computer-readable medium, such as, for example, memory 130, a magnetic disk, or an optical disk. Such instructions may be read into memory 130 from another computer-readable medium, such as storage device 150, or from a separate device via communication interface 180.

[0022] Processing device 100 may be, for example, a personal computer (PC), or any other type of processing device capable of creating and sending messages. In alternative implementations, such as, for example, a distributed processing implementation, a group of processing devices 100 may communicate with one another via a network such that various processors may perform operations pertaining to different aspects of the particular implementation.

[0023] We now turn to the details of the invention and introduce the decoding problem and a fixed-point representation usable with the present invention. Most dialog systems have a floating point implementation of ASR. In a fixed-point implementation, the program's dynamic range of arithmetic operation is smaller. Therefore, one issue is to make sure that a range of different quantities that are manipulated, for example added or

multiplied, in a quantization process will stay within the dynamic range. So wherever there is quantization, the invention preserves enough precision so essentially so that the system does not lose accuracy in the computation.

[0024] In general, a fixed-point number representation in computing is a real data type for a number that has a fixed number of digits after the decimal (binary or radix) point. Fixed-point numbers are useful for representing fractional values in native two's complement format if the executing processor has no floating point unit (FPU) or if fixed-point provides improved performance. Most low cost embedded processors do not have an FPU. The bits to the left of the radix point are magnitude bits that represent integer values, the bits to the right of the radix point represent fractional values. Each fractional bit represents an inverse power of 2. Thus the first fractional bit is $1/2$, the second is $1/4$, the third is $1/8$ and so on.

[0025] A floating-point number is a digital representation for a number in a certain subset of the rational numbers, and is often used to approximate an arbitrary real number on a computer. In particular, it represents an integer or fixed-point number (the significand or, informally, the mantissa) multiplied by a base (usually 2 in computers) to some integer power (the exponent). When the base is 2, it is the binary analogue of scientific notation (in base 10).

[0026] Fig. 2 shows the implementation of a floating-point recognizer that transcribes "input speech" into "output text", discussed for background information. The ASR frontend 202 computes from the input spoken speech (typically a sentence of seconds in duration) a sequence of feature (or observation) vectors, typically at a rate of 100 vectors/second. Therefore, each vector can be thought as representing the input speech over a duration of 10 msec. Fig. 2 shows that the sequence of observation vectors is fed into the decoder 204 that in turn outputs the corresponding text sentence. Even though a language can be thought as composed by linguistic units (phonemes), the pronunciation

of phonemes is affected by speaker dialect, accent, gender, age and emotion as well as other factors such as background noise and acoustic environment. It is also affected by the phonetic context in which the phoneme is pronounced (co-articulation or context-dependency). Phonemes are pronounced with varying durations, depending on the speech rate. Therefore, the operation of the decoder is fairly complex: the decoder cannot map directly the feature vectors (acoustic measurements of speech) into phonemes and then into words and sentences, for example using a simple table look-up scheme. Instead, the decoder 204 evaluates a multitude of possible text sentences (sentence hypotheses), and it selects the “best” hypothesis that “most likely” corresponds to the input feature vector sequence. The decoder 204 computes the likelihood of a text sentence given the input feature vectors using various speech models shown in FIG 2. Intuitively, the acoustic HMM provides the likelihood of a feature vector being generated as a (portion of a) phoneme (HMM state). These likelihoods are combined in the decoding process, to provide the likelihood of the entire sentence hypothesis.

[0027] For this purpose, the decoder uses the C,L,G finite state transducers (FST’s) as described in the paper M. Mohri, F. Periera, and M. Riley, “Weighted finite-state transducers in speech recognition,” *Computer, Speech and Language*, pp. 16(1): 69-88, 2002, incorporated herein by reference. Informally, an FST is a directed graph consisting of “nodes” (or states) connected by “arcs”. There are designated “initial” and “final” states. Allowed paths or arc sequences must begin from the initial state and terminate in a final state. A generic arc may have an “input” and “output” symbol. In the FST operation, a sequence of symbols is matched against the input symbols of the transducer arcs thus producing a certain arc sequence or path. The arc sequence induces a sequence of output symbols, thus a “transduction” of input symbols into output symbols. Arcs may have probability scores (often “costs” or negative-log-probabilities), that allow for the computation of the likelihood of sequences of input/output symbols. In the

application of FST's to ASR, the L (lexicon or dictionary) FST is designed so that the sequences of arcs (with phonemes as input symbols) correspond to the pronunciation of words (arc output symbols). Therefore, L describes how words consist of phonemes, and it allows to transform phoneme sequences into words. Fig. 4 shows a L transducer for the words "dew" and "data". The path from the initial "0" state, to the final state "6", represents the word "dew" as the sequence of phonemes "d" and "uw". The different paths from state "0" to final state "4" represent different pronunciations of the word "data", like for example the phoneme sequence "d", "ae", "dx", "ax". The probabilities of the L transducer arcs, allows the decoder to compute the probability of different pronunciations of "data". Note the states 402, 404, 408, 410, 412, 414, and final states 414 (if the word is "dew") and 416 (if the word is "data"). The G (grammar or language model) FST uses words as arc input symbols, with paths that correspond to the word sequences allowed by the particular ASR application. For example, the G transducer of Fig. 5 shows the sentence "using intuition is better" as grammatical, because there is a path from the initial state "0" 502 to the final state "5" 512 with arc input labels that match the sentence words. The arc costs of the G transducer allow for the computation of the probability of the sentence, e.g. "using intuition is better" in Fig 5. Note states 504, 506, 508 and 510 in Fig. 5. Language model transducers for general English are complex structures that may contain millions of arcs. An important phenomenon in spoken language is that the pronunciation of a phoneme is affected by the adjacent phonemes (co-articulation). In the ASR field co-articulation is modeled by the C FST that transforms sequences of input context-dependent phonetic symbols into the linguistic phonemes.

[0028] In summary, the C FST transforms a sequence of context dependent phonetic symbols in phonemes, the L FST transforms the phonetic sequence into words, and the G FST allows for the word sequences that are grammatical in the ASR application. In

the ASR field, the FST “composition” operator 206 in Fig. 2 is used to combine the C, L, G FST’s into a single CLG FST that transforms directly sequences of context dependent phonetic symbols into grammatical sentences. Therefore, the CLG transducer can be used by the decoder 204 to derive the probability of sentence hypotheses from the probabilities of the input feature vectors, as provided by the acoustic hidden Markov model (HMM) for the different phonemes. The decoder 204 then outputs the sentence hypothesis with the highest probability (or lowest cost). As can be appreciated, the basic principles here can apply outside of speech processing contexts. The formal mathematical equations known to the ASR practitioners are summarized next.

[0029] Using an LVCSR map decoder, given an observation sequence $\mathbf{O} = (o_1, \dots, o_T)$, the maximum a-posteriori decoded word sequence $\Omega = (\omega_1, \dots, \omega_M)$ is chosen to maximize $p(\Omega | \mathbf{O})$. Those of skill in the art will recognize the well known steps to perform this process:

$$\hat{\Omega} = \arg \max_{\Omega} \mathcal{F}(\Omega), \text{ with decoding function } \mathcal{F} :$$

$$\mathcal{F} = \ln(\mathcal{L}) + \alpha \ln(\mathcal{A}) + \beta \ln(\mathcal{D}) \quad , \quad (1)$$

$\mathcal{L}$: likelihood of language model,
 $\mathcal{A}$: likelihood of the HMM states (acoustic model),
 $\mathcal{D}$: likelihood of the HMM state durations,
 α, β : state and state-duration model multipliers.

[0030] The inventors encode the language model probabilities by the arc costs of a C o L o G transducer. A finite-state transducer is a finite automaton whose state transitions are labeled with both input and output symbols. A path through the transducer encodes a mapping from an input symbol sequence to an output symbol sequences. These transitions can also have weights associated with them in addition to the input and output symbols. These weights may encode probabilities, durations, penalties or any other quantity that can accumulate along a path to compute an overall weight associated

with a mapping of an input sequence to an output sequence. For more background information, see the Mohri and Periera paper incorporated above and M. Mohri, “Weighted Finite-State Transducer Algorithms An Overview”, In Carlos Martín-Vide, Victor Mitrana, and Gheorghe Paun, editors, *Formal Languages and Applications*, volume 148, VIII, 620 p. Springer, Berlin, May, 2004, both papers incorporated herein by reference. The C o L o G transducer comprises “C” represents a context-dependency transducer from context-dependent phones to context-independent phones, and “G” is a grammar and “L is a pronunciation lexicon. The C o L o G transducer represents the language model, either a word N-gram or a deterministic grammar, the lexicon, with pronunciation probabilities, and the phonetic context dependencies. For generality, by $\ln(L)$, the inventors denote the negative total cost of the transducer path, induced by Ω :

$$\ln(\mathcal{L}) = - \sum_{arc \in path(\Omega)} cost_{arc} \quad (2)$$

[0031] Next is discussed the HMM state likelihoods. The generic HMM state “s” is a weighted mixture of N_s Gaussians with diagonal covariances (σ denotes the vector of stdv’s):

$$P(o|s) = \sum_{i=1}^{N_s} w_{s,i} \mathcal{N}(o; \mu_{s,i}, \sigma_{s,i}) \approx \max_{i=1, N_s} w_{s,i} \mathcal{N}(o; \mu_{s,i}, \sigma_{s,i})$$

[0032] Given the sequence $S = (s_1, \dots, s_T)$ of states aligned to O , the total state likelihood contribution to (1) is $\propto \ln(A)$:

$$\sum_{t=1}^T \frac{\alpha}{2} \max_{i=1, N_{st}} \left(2c_{s,t,i} + \sum_{j=1}^d \left(\left(o_t^j - \mu_{s,t,i}^j \right) \sigma_{s,t,i}^{j-1} \right)^2 \right) \quad (3)$$

$$\begin{aligned} d &: \text{feature vector dimension} \\ o_t^j, \mu_{s,i}^j, \sigma_{s,i}^j &: j^{th} \text{ component of } o_t, \mu_{s,i}, \sigma_{s,i} \\ c_{s,i} &: \ln(w_{s,i}) - \sum_{j=1}^d \ln(\sigma_{s,i}^j \sqrt{2\pi}) \end{aligned}$$

[0033] The state duration model is discussed next. The inventors denote with $p(\delta|\psi)$ the probabilities of the duration of δ -frames for state ψ . They are estimated (assuming state independence) as gamma p.d.f's, and stored in look-up tables. Given the states $\Psi=(\psi_1, \dots, \psi_\Theta)$, with durations $\Delta=(\delta_1, \dots, \delta_\Theta)$, the contribution of the duration model to (1), is:

$$\beta \ln(\mathcal{D}) = \beta \ln(P(\Delta|\Psi)) = \sum_{\theta=1}^{\Theta} \beta \ln(p(\delta_\theta|\psi_\theta)) \quad (4)$$

[0034] Using fixed-point arithmetic, a decimal number x is stored in a computer word as an integer, where the p_x least-significant bits contain the fractional part of x , and the most-significant bits the integer part. This representation is said to have $Q-p_x$ format. This is a particular format used for manipulating binary data using fixed-point arithmetic. In general, the choice of p_x is a trade-off between truncation errors and overflow problems. Fixed-point arithmetic is based on integer operations. In $z = x \pm y$, the relation $p_x = p_y = p_z$ applies. The product $z = x \times y$ is in format $Q - (p_x + p_y)$. Arithmetic shifts may be used to change the Q fixed-point format to reduce truncation errors and to avoid overflow. While using the Q fixed-point format is preferable, other fixed-point number representations may be used in connection with the invention.

[0035] In the floating-point recognizer (Fig. 2) equations (1-4) are computed with floating-point numbers. The process of decoding speech with fixed-point numbers can be exemplified by FIG 3, where “quantizers” transform the HMM parameters and FSM costs into a fixed-point representation, and where the decoding equations (1-4) are computed in fixed-point.

[0036] Next is discussed further fixed-point arithmetic and linear quantization. The function *nearest_integer* ($2^p x$) gives the Q-p fixed-point format of decimal x. Suppose that one wants to quantize the range of decimal values $[a, b]$, using m bits, e.g. to the range of integers $[-2^{m-1}, 2^{m-1}]$. The following is an example procedure that could be followed:

i. Optional. *Demean* decimal values, by subtracting $\frac{a+b}{2}$:

$$[a, b] \Rightarrow \left[-\frac{a+b}{2}, \frac{a+b}{2} \right]$$

ii. Find the largest integer p, such that

$$-2^{m-1} \leq 2^p x < 2^{m-1}, x \in [a, b]$$

iii. Quantize $x \in [a, b]$ by $y = \text{nearest\_integer}(2^p x)$.

m	: bits for the d quantizers of means $\mu_{s,i}^j$, $j = 1, d$.
v	: bits for the d quantizers of $\sigma_{s,i}^{j-1}$, $j = 1, d$.
e	: fixed-point $Q-e$ format for
	– normalized error: $(o_t^j - \mu_{s_t,i}^j) \sigma_{s_t,i}^{j-1}$
	and $Q-2e$ format for:
	– HMM state log-likelihoods,
	– Acoustic duration model log-likelihoods,
	– $C \circ L \circ G$ fsm costs,
	– Cumulative log-probabilities during decoding,
	and related parameters, such as beam threshold

Table 1. Fixed-point parameters m , v and e .

[0037] Step iii in the procedure above yields a fixed-point format of x with scale--

invariant average truncation error, because of the choice of p in ii.

[0038] The core concepts of the invention are discussed next for a fixed-point

implementation of a decoding function. These ideas relate to the fixed-point

implementation of the decoding function for applications such as speech processing.

This may apply to any of the speech processing modules discussed above or to any other

basic transducers or modules. By way of example, the implementation will be discussed

with reference to speech recognition. The inventors' fixed-point design of (1) is

parametrized by e , m , and v , as summarized in Table 1 above. A role in the state

likelihoods (3) is played by, for example, the Mahalanobis distance,

$$\sum_{j=1}^d \left(\left(o_t^j - \mu_{s_t,i}^j \right) \sigma_{s_t,i}^{j-1} \right)^2 \quad (5)$$

and, its fixed-point format is specified preferably as $Q-2e$. Since (3) accumulates (5) into

the state log-likelihoods, an example representation of all the decoder log-likelihoods is

preferably in the $Q-2e$ format. Other formats may be utilized as well. Therefore, the

HMM log-terms $2c_{s,i}$ in (3), are also Q-2e, and, the product by $\alpha/2$ maintains the Q-2e format through an appropriate arithmetic shift. Similarly, the duration log-probabilities of (4) are Q-2e fixed-point, and when multiplying by β , the Q-2e format is maintained by arithmetic shift.

[0039] Also, the cost_{arc} of (2) are converted to the Q-2e format. Note that, in the fixed-point composition $C \circ L \circ G$, a simple integer addition implements the $\otimes$ operator in the tropical semiring, under the condition that the costs of the C, L and G are all Q-2e. Thus, instead of quantizing the costs of the CLG FST (as in Fig. 3), one can quantize the costs of the individual C, L, G FST's, and use the delayed composition feature referenced in the Mohri and Pereira article above as implemented in the floating-point decoder. This feature minimizes run-time memory and is useful in applications such as various embedded speech recognition platforms.

[0040] Different approaches are suitable for the implementation of (5) in Q-2e fixed-point. For example, one way is to square and accumulate terms:

$$\left(o_t^j - \mu_{s_t,i}^j \right) \sigma_{s_t,i}^{j-1}, \text{ in Q-e format,} \quad (6)$$

that are computed using suitable integer representations of the HMM parameters. To account for the different dynamic ranges of the Gaussian mean components, one way to perform this is to build a quantizer for every j^{th} ($j = 1, d$) component, as in i, ii and iii referenced above for the decimal range:

$$\left[\min_{\text{Gaussian } i, \text{ state } s} \mu_{s,i}^j, \quad \max_{\text{Gaussian } i, \text{ state } s} \mu_{s,i}^j \right]$$

```

sum = 0 , j = 1
while(j ≤ d) {
    tmp = (otj - μst,ij) σst,ij-1 // Q-(pj + rj)
    tmp = tmp >> shiftj // change to Q-e
    sum = sum + tmp * tmp // sum is Q-2e
    j = j + 1
}

```

Table 2. Pseudo-code for fixed-point implementation of (5)

[0041] Parameter m specifies the number of bits of the quantizers. The inventors denote by $Q\text{-}p^j$ the fixed-point format of $\mu_{s,i}^j$, induced by the j^{th} quantizer, which also defines the fixed-point format of o_t^j , in (6). Similarly, one way to build another set of d quantizers, one for every $\sigma_{s,i}^{j-1}$ (the j^{th} inverse stdv component) is using steps ii and iii referenced above. Parameter v specifies the number of bits, output range $[0, 2^v)$, of these quantizers. The inventors denote by $Q\text{-}r^j$ the fixed-point format of $\sigma_{s,i}^{j-1}$, induced by its quantizer. The fixed-point format of the integer product (6) is therefore $Q\text{-}(p^j + r^j)$, that is changed to $Q\text{-}e$ with a right arithmetic shift of $\text{shift}_j = (p^j + r^j - e)$ bits, as shown in Table 2 (negative shift_j implies a left shift). In practice, e , m and v can be chosen so that $\text{shift}_j \geq 0$, $j = 1, d$. It might also be worthwhile for higher computational speed, to implement (5) through multiply-add operations. Other ways of implementing the approaches above than those presented.

[0042] FIG. 2 illustrates a method embodiment of the invention. An exemplary method relates to processing data associated with speech processing. The method uses a speech recognition decoder that operates or uses fixed-point arithmetic. The exemplary method comprises representing arc costs associated with at least one finite-state

transducer (FST) in fixed-point (602), representing parameters associated with a hidden Markov model (HMM) in fixed-point (604) and processing speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters (606). The method may also include computing at the decoder sentence hypothesis probabilities with fixed-point arithmetic as type Q-2e numbers. The term “representing” is a broad term which encompasses many aspects of how to represent parameters, such as processing data so that the parameters may be represented in a certain format, performing computations on data such that the data may be represented in a certain format, applying a look-up table, or any other conceptual way of arriving at a result of having parameters that are in the particular format.

[0043] Other aspects of this method include features such as the method being applied for large vocabulary automatic speech recognition, although other contexts as well may be applied. The HMM is preferably a phonetic context-dependent acoustic HMM and the at least one FST may be either one of a single CLG FST or distinct C, L and G FSTs.

[0044] Representing arc costs that are associated with at least one finite-state transducer (FST) in fixed-point may further comprise converting floating-point costs of the at least one FST to fixed-point by quantization. In another aspect of the invention, representing parameters associated with a hidden Markov model (HMM) in fixed-point may involve converting HMM state observation probability density functions to fixed-point format by quantization and computing HMM state likelihoods as Q-2e fixed-point numbers.

Representing parameters associated with a hidden Markov model (HMM) in fixed-point may also comprise converting HMM state duration probabilities to fixed-point by quantization using a fixed-point format of type Q-2e.

[0045] Another aspect of the invention involves a method comprising generating a fixed-point speech recognition transducer by converting a floating point speech recognition transducer via quantization into a fixed-point format, receiving input speech to the fixed-

point speech recognition transducer and generating a best hypothesis of the received input speech using the fixed-point speech recognition transducer.

[0046] The invention utilizes existing typical algorithms associated with speech processing such as an automatic speech recognition module, and applies the concepts herein to the main components such as the HMM and language models to quantize costs and other parameters within these components to obtain a suitable fixed-point representation of the speech processing components. Speech processing can then be applied to the new fixed-point speech processing components and tested to insure that their accuracy is acceptable relative to the floating point versions. The best sentence hypothesis is returned. In the case of speech recognition, after the conversion of the modules occurs, the system will use HMM and language models applying a fixed-point computation.

[0047] The quantization process involves applying algorithms to a language model and an HMM to represent in fixed-point the costs associated with different hypothesis with different word sequences. The acoustic model with its mean and variance parameters and Gaussian weighting components are also represented in fixed-point. During recognition, the decoder instructs the HMM to provide likelihood values for the input speech and these values are computed in fixed-point using fixed-point arithmetic. Similarly during recognition, the decoder instructs the language model to return the likelihood of word sequences, which response is provided as fixed-point numbers. The decoder accumulates these numbers to test different hypothesis of sentences and come up with the best fit.

[0048] The speech decoder algorithm is preferably a Viterbi algorithm (which may or may not utilize floating point numbers) but the computation of word hypothesis or sentence hypothesis scores is performed in fixed-point.

[0049] An HMM comprises a series of states and there are typically three states for each phoneme. If one considers the states as a subunit of a phoneme, a state model represents the probability of the duration of the states or the probability of a portion of a phoneme to have a particular duration. It is these states or the probabilities of the duration of these states that is quantized in the process of conversion to fixed-point. For example, in an HMM, the context dependent phonemes may each be represented by three states. For example, different states may represent an initial portion of a phoneme, a middle and a last portion of phonemes. The transition between states corresponds to the duration of a phoneme which also relates to how long the speech is aligned with a particular state. So the duration with a speech aligned or assigned to a particular state is associated with a probability.

[0050] The HMM is represented by Gaussian mixture models shown in equation (3) above. The probabilities of a certain instant of speech belonging to a certain state are given by equation (3) as well. These states are connected in various ways and speech assigned to a certain duration delta which relates to the probability of that delta depends on the state and the expression of that is provided in equation (4).

[0051] The inventors also experimented to prove the concepts above. Important to the design is the fixed-point representation of (6), whose statistics are largely independent of the ASR task: after HMM maximum likelihood estimation (6) are Gaussians with zero mean and unit variance. The concept is that the decoder fixed-point parameterization does not need task-specific calibrations, which can be verified in recognition tests on:

[0052] **RM:** Darpa Naval Resource Management, with word-pair grammar, speaker-independent task,

[0053] **SBCC:** fluent telephone speech from a customer-care application, with word tri-gram language model (perplexity of 60), vocabulary of 7,000 words, and up to 50 words/sentence,

[0054] **SWBD**: Darpa Switchboard task, tested on the 2003 real-time test set

(recognition from first-pass only), with feature types:

[0055] **MFCC**: mel-frequency cepstrum coefficients,

[0056] **PLP**: perceptual linear prediction cepstra, and with feature transformations:

[0057] **DD**: cepstra with 1st and 2nd differentials, 39 components,

[0058] **HDA**: discriminative linear transformation, 60 components,

[0059] **VTNL**: vocal tract length normalization.

ASR System	Word Accuracy (%)	
	Floating	Fixed
RM_MFCC-HDA_ML	96.4	96.4
RM_MFCC-DD_ML	95.7	95.6
RM_PLP-DD_ML	95.6	95.5
SBCC_MFCC-HDA-VTLN_MMI	80.5	80.6
SBCC_MFCC-HDA_MMI	78.4	78.4
SWBD_MFCC-HDA_MMI	59.2	59.1
SWBD_MFCC-HDA_ML	56.7	56.5
SWBD_PLP-HDA_ML	55.7	55.6

Table 3. Floating and fixed-point decoder accuracy. Fixed-point parameters: $e = 5$, $m = v = 8$

[0060] The HMMs are, for example, context-dependent triphonic models, estimated either by maximum likelihood (ML) or maximum mutual information (MMI). Training of the SBCC and SWBD HMM's use 170 and 300 hours of audio, respectively. For example, the 1st pass of the AT&T RT-03 Switchboard system, with MFCC features and discriminative transformation, and MMI-trained HMM, is denoted by SWBD_MFCC--HDA_MMI. Table 3 shows that the accuracies of the fixed-point and the floating-point recognizers (equal beam width), are the same, within 0.1%, for all tasks. Means and variances can be linearly quantized to 5 bits, without significant loss of accuracy (Table 4). Additional compression may be obtained by nonlinear quantization, at the cost of

additional indirections in the computation. The inventors wanted to quantize HMM means and variances to no more than 8 bits, to reduce the HMM memory storage to a relatively small fraction of total run-time memory use. The accuracy is unchanged (within 0.1%) for $1 < e < 7$ (Table 4). Accuracy suffers from truncation errors for $e \leq 1$, and from overflow problems for $e \geq 7$. Larger e 's require normalization of the cumulative log-likelihoods in the Viterbi search. In any case, the decoder operates correctly over a wide range of e . The fixed-point experiments on the Switchboard and SBCC tasks were run on a Pentium processor (32-bit integers); but the good performance for e as small as 2, suggests that the implementation is suitable for 16-bit CPU's with 32-bit accumulators.

[0061] The inventors benchmarked the RM task on a StrongARM-1100, running at 206 MHz, embedded in a desk-top telephone set prototype, with 30 Mbytes of RAM, and Linux 2.6.6.

[0062] StrongARM executables are cross-compiled on the PC, with the GNU-toolchain and gcc3.4.2. For testing ASR on the device, access to executables and to fixed-point speech feature files is through ethernet-NFS. Any computation on the StrongARM produces the exact same results as on a Pentium, but, in a longer time. Real-time recognition of the RM_MFCC-DD_ML fixed-point system on the StrongARM-1100 is shown in Figure 3, that plots 700 the accuracy as a function of time (normalized by duration of input speech), for different beam-widths (runtime memory use of 7.5 MBytes). Modern embedded systems, with clock rates of 650 MHz, and more than 128 MBytes of RAM, could run in real-time more complex tasks, such as SBCC.

[0063] The fixed-point implementation is necessary for embedded ASR: the inventors tested the floating-point decoder, cross-compiled with floating-point software emulation, and it was ≈ 50 times slower.

[0064] The presented fixed-point implementation of the LVCSR decoder, including language models and search, is as accurate as our best floating-point recognizer, in medium and large vocabulary continuous speech recognition tasks, as tested on the Pentium and StrongARM-1100 CPU's.

[0065] In addition to the more important motivation discussed in the Introduction, the algorithms are useful to prototype ASR applications in embedded systems. In fact, the decoder fixed-point parameters do not need critical task-dependent calibrations, and the LVCSR language and acoustic models, trained with the standard floating-point algorithms, can be automatically ported to the required fixed-point representation.

	$v=8$	$v=7$	$v=6$	$v=5$	$v=4$	$v=3$
$m=8$	95.6	95.7	95.6	95.5	95.5	93.2
$m=7$	95.7	95.7	95.7	95.6	95.4	93.2
$m=6$	95.5	95.6	95.6	95.6	95.4	93.1
$m=5$	95.4	95.4	95.5	95.4	95.3	92.9
$m=4$	94.7	94.6	94.6	94.7	94.7	91.9
$m=3$	92.4	92.4	92.4	92.6	92.1	87.7

(a)

	$v=8$	$v=7$	$v=6$	$v=5$	$v=4$	$v=3$
$m=8$	80.6	80.5	80.5	80.2	76.6	45.6
$m=7$	80.5	80.6	80.4	80.2	76.5	45.7
$m=6$	80.4	80.5	80.2	80.1	76.4	45.3
$m=5$	80.3	80.2	80.1	80.0	76.3	45.0
$m=4$	77.8	78.0	77.6	77.8	74.4	44.1
$m=3$	56.4	56.6	57.0	59.9	53.3	43.5

(b)

Table 4. Word accuracy (%) as function of m and v ($e=5$) of fixed-point systems RM_MFCC-DD_ML (a) and SBCC_MFCC-HDA-VTLN_MMI (b).

	$e=8$	$e=7$	$e=6$	$e=5$	$e=4$	$e=3$	$e=2$	$e=1$	$e=0$
(a)	94.8	94.9	95.6	95.6	95.7	95.7	95.6	95.3	92.1
(b)	69.2	79.9	80.5	80.6	80.5	80.5	80.5	80.2	75.6

Table 5. Word accuracy (%) as function of e ($m = v = 8$) of fixed-point systems RM_MFCC-DD_ML (a) and SBCC_MFCC-HDA-VTLN_MMI (b).

[0066] Embodiments within the scope of the present invention may also include computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code means in the form of computer-executable instructions or data structures. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or combination thereof) to a computer, the computer properly views the connection as a computer-readable medium. Thus, any such connection is properly termed a computer-readable medium. Combinations of the above should also be included within the scope of the computer-readable media.

[0067] Computer-executable instructions include, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Computer-executable instructions also include program modules that are executed by computers in stand-alone or network environments. Generally, program modules include routines, programs, objects, components, and data structures, etc. that perform particular tasks or implement particular abstract data types. Computer-executable instructions, associated data structures, and program modules represent examples of the program code means

for executing steps of the methods disclosed herein. The particular sequence of such executable instructions or associated data structures represents examples of corresponding acts for implementing the functions described in such steps.

[0068] Those of skill in the art will appreciate that other embodiments of the invention may be practiced in network computing environments with many types of computer system configurations, including personal computers, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. Embodiments may also be practiced in distributed computing environments where tasks are performed by local and remote processing devices that are linked (either by hardwired links, wireless links, or by a combination thereof) through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0069] Although the above description may contain specific details, they should not be construed as limiting the claims in any way. Other configurations of the described embodiments of the invention are part of the scope of this invention. Accordingly, the appended claims and their legal equivalents should only define the invention, rather than any specific examples given.

CLAIMS

We claim:

1. A method of processing data associated with speech processing, the method using a speech recognition decoder using fixed-point arithmetic, the method comprising:
 - representing arc costs associated with at least one finite-state transducer (FST) in fixed-point;
 - representing parameters associated with a hidden Markov model (HMM) in fixed-point; and
 - processing speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters.
2. The method of claim 1, wherein the method is applied for large vocabulary automatic speech recognition.
3. The method of claim 1, wherein the HMM is a phonetic context-dependent acoustic HMM.
4. The method of claim 1, wherein the at least one FST is one of a single CLG FST or distinct C, L and G FSTs.
5. The method of claim 1, wherein representing arc costs associated with at least one finite-state transducer (FST) in fixed-point further comprises:
 - converting floating-point costs of the at least one FST to fixed-point by quantization.

6. The method of claim 5, wherein a format for the fixed-point numbers is type Q-2e.

7. The method of claim 1, wherein representing parameters associated with a hidden Markov model (HMM) in fixed-point further comprises:
converting HMM state observation probability density functions to fixed-point format by quantization; and
computing HMM state likelihoods as Q-2e fixed-point numbers.

8. The method of claim 1, wherein representing parameters associated with a hidden Markov model (HMM) in fixed-point further comprises:
converting HMM state duration probabilities to fixed-point by quantization using a fixed-point format of type Q-2e.

9. The method of claim 1, further comprising:
computing at the decoder sentence hypothesis probabilities with fixed-point arithmetic as type Q-2e numbers.

10. A computing device that processes data associated with speech processing, the device using a speech recognition decoder using fixed-point arithmetic, the computing device comprising:
a module configured to represent arc costs associated with at least one finite-state transducer (FST) in fixed-point;
a module configured to represent parameters associated with a hidden Markov model (HMM) in fixed-point; and

a module configured to process speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters.

11. The computing device of claim 10, wherein the computing device performs large vocabulary automatic speech recognition.

12. The computing device of claim 10, wherein the HMM is a phonetic context-dependent acoustic HMM.

13. The computing device of claim 10, wherein the at least one FST is one of a single CLG FST or distinct C, L and G FSTs.

14. The computing device of claim 10, wherein the module configured to represent arc costs associated with at least one finite-state transducer (FST) in fixed-point further converts floating-point costs of the at least one FST to fixed-point by quantization.

15. The computing device of claim 14, wherein a format for the fixed-point numbers is type Q-2e.

16. The computing device of claim 10, wherein the module configured to represent parameters associated with a hidden Markov model (HMM) in fixed-point further:


converts HMM state observation probability density functions to fixed-point format by quantization; and

computes HMM state likelihoods as Q-2e fixed-point numbers.

17. The computing device of claim 10, wherein the module configured to represent parameters associated with a hidden Markov model (HMM) in fixed-point further:
- converts HMM state duration probabilities to fixed-point by quantization using a fixed-point format of type Q-2e.
18. The computing device of claim 10, further comprising:
- a module configured to compute at the decoder sentence hypothesis probabilities with fixed-point arithmetic as type Q-2e numbers.
19. A computer-readable medium storing instructions for controlling a computing device to process data associated with speech processing, the computing device including a speech recognition decoder using fixed-point arithmetic, the instructions comprising:
- representing arc costs associated with at least one finite-state transducer (FST) in fixed-point;
 - representing parameters associated with a hidden Markov model (HMM) in fixed-point; and
 - processing speech data in the speech recognition decoder using fixed-point arithmetic for the fixed-point FST arc costs and the fixed-point HMM parameters.
20. The computer-readable medium of claim 19, wherein the step of representing arc costs associated with at least one finite-state transducer (FST) in fixed-point further comprises converting floating-point costs of the at least one FST to fixed-point by quantization.

1/4

FIG. 1



2/4

FIG. 2





FIG. 3



3/4

FIG. 4







FIG. 5



4/4

FIG. 6*FIG. 7*

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2007/061911

A. CLASSIFICATION OF SUBJECT MATTER
INV. G10L15/18 G10L15/14

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G10L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	MOHRI M ET AL: "Weighted finite-state transducers in speech recognition" COMPUTER SPEECH AND LANGUAGE, ELSEVIER, LONDON, GB, vol. 16, no. 1, January 2002 (2002-01), pages 69-88, XP004418697 ISSN: 0885-2308 cited in the application abstract page 71, lines 22-39 page 81, lines 27-30 page 83, lines 16-22 page 84; table 1 ----- -/--	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

* & * document member of the same patent family

Date of the actual completion of the international search

21 June 2007

Date of mailing of the international search report

29/06/2007

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Geißler, Christian

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2007/061911

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,X	BOCCHIERI E ET AL: "A Decoder for Lvcscr Based on Fixed-Point Arithmetic" ACOUSTICS, SPEECH AND SIGNAL PROCESSING, 2006. ICASSP 2006 PROCEEDINGS. 2006 IEEE INTERNATIONAL CONFERENCE ON TOULOUSE, FRANCE 14-19 MAY 2006, PISCATAWAY, NJ, USA, IEEE, 14 May 2006 (2006-05-14), pages I-1113, XP010930379 ISBN: 1-4244-0469-X the whole document	1-20
P,X	YUQING GAO ET AL: "IBM Mastor: Multilingual Automatic Speech-To-Speech Translator" ACOUSTICS, SPEECH AND SIGNAL PROCESSING, 2006. ICASSP 2006 PROCEEDINGS. 2006 IEEE INTERNATIONAL CONFERENCE ON TOULOUSE, FRANCE 14-19 MAY 2006, PISCATAWAY, NJ, USA, IEEE, 14 May 2006 (2006-05-14), pages V-1205, XP010931658 ISBN: 1-4244-0469-X abstract pages V-1205, right-hand column, lines 6-12 pages V-1207, left-hand column, lines 8-33 pages V-1207, left-hand column, line 50 - right-hand column, line 1 pages V-1207, right-hand column, lines 8-11	1-20
A	WO 03/036618 A (DSP FACTORY LTD [CA]; SHEIKHZADEH-NADJAR HAMID [CA]; CORNU ETIENNE [CA] 1 May 2003 (2003-05-01) page 11, line 28 - page 12, line 4	1-20
A	US 5 806 034 A (NAYLOR JOE A [US] ET AL) 8 September 1998 (1998-09-08) column 5, line 66 - column 6, line 14 column 6, lines 56-58	1-20
A	US 2002/198706 A1 (KAO YU-HUNG [US] ET AL) 26 December 2002 (2002-12-26) paragraphs [0013], [0014] paragraphs [0022] - [0027] paragraph [0084]	1-20
A	WO 01/20597 A (CONEXANT SYSTEMS INC [US]) 22 March 2001 (2001-03-22) page 5, line 31 - page 6, line 4	1-20
A	US 2004/220804 A1 (ODELL JULIAN J [US]) 4 November 2004 (2004-11-04) paragraphs [0009], [0010] paragraph [0044]	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2007/061911

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 03036618	A	01-05-2003	CA 2359544 A1 EP 1449203 A1 US 2003110033 A1	22-04-2003 25-08-2004 12-06-2003
US 5806034	A	08-09-1998	NONE	
US 2002198706	A1	26-12-2002	NONE	
WO 0120597	A	22-03-2001	TW 521263 B US 2005149337 A1	21-02-2003 07-07-2005
US 2004220804	A1	04-11-2004	NONE	