Extended Abstracts of the 1994 International Conference on Solid State Devices and Materials, Yokohama, 1994, pp. 391-393

The PHRASES Neuro-Coprocessor Cost Effective Handwriting Recognition for Personal Digital Assistant

Marc HERVIEU and Jean-Yves BRUNEL

Laboratoires d'Electronique Philips,
22, avenue Descartes
94453 Limeil-Brévannes Cedex (France)

In the context of dynamic handwriting recognition, i.e. handwriting acquired from pen motion, we present the PHRASES system. It generates a cost effective handwriting recognizers chip using an Artificial Neural Network symbol classifier that is backed up by a dictionary search at word level. It uses the chip-synthesis ALMA methodology to produce embeddable implementations which are tuned to application specific requirements such as the resolution of the input tablet or the dictionary size.

1. INTRODUCTION

The race for portable intelligent devices aiming at the mass consumer market has started. The customers of these **Personal Digital Assistants** (PDA) and **Smart Cards**, expect low price, small size and long battery life. They also expect intuitive use, stressing the need for handwriting recognition in particular. However, products offered today on the market yield poor recognition score with a speed below the acceptable level[1].

Much progress have to be made. The strategy is straightforward : it is based upon the use of precise vocabularies, indeed even of simple grammars in restricted contexts, to remove the ambiguities that remain at stroke level. These improvements require more real-time computing power that can be offered by general purpose 30 MIPS controllers.

This has motivated the design of several dedicated circuits that can be used for speeding up symbol recognition or dictionary search [2]. However, these off-the-shelf circuits implement at most one out of the tasks that build up a complete recognition. They are actually designed process several thousands characters per second. The extra hardware cost that must be paid for such a speed is acceptable for institutional markets, must be drastically reduced to address consumer electronics where low price, hence chips count, is at a premium.

The solution relies in providing scalable recognizers that may fit on chip with the core processor of a **D**eeply **E**mbedded system. This integration achieves a price effective one-chip solution, with excellent MIPS-per-watt ratings.

Within this framework, this paper describes the PHRASES project. It relies on the ALMA[3] synthesis-based methodology to generate components for handwriting recognition in a vast range of consumer products. However, the core market of PHRASES is focused on the Personal Digital Assistants, to which it offers a much better recognition score and speed at the expense of a few mm^2 of silicon. As an example, the base version reaches a 60 acceleration factor over an PC486 DX33 implementation. This is obtained for less than 25 mm^2 in 0.8 μm technology.

2. ARCHITECTURE

In contrast to of-the-shelf circuits to date, PHRASES chip includes all or part of the time-consuming recognition tasks. Each component has been identified from a quantitative analysis of the handwriting recognition software that has been developed at LEP[4]. The generic architecture is shown in figure 1. It consists of the composition of synchronous processing elements that cooperate using an

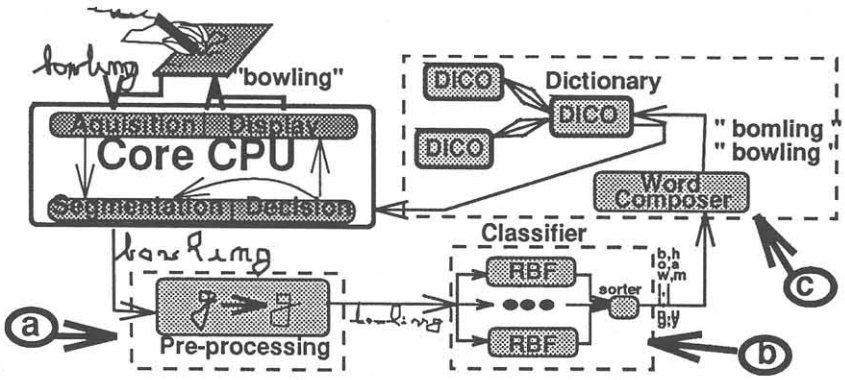

Fig 1: Generic Architecture

asynchronous interface protocol. The strengths of this micro-pipelined architecture are particularly important for generating an embeddable processor implementation which is tuned to the requirements of a specific application.

The complete subsystem receives segments of the "electronic ink" that is issued by the handwriting acquisition module. It returns the list of words of a vocabulary that most closely match the submitted input.

2.1. Ink Preprocessing

The electronic ink consists of a flow of variable length vectors of (x,y) coordinates representing the trajectory of the pen. As shown in figure 1(a), a preliminary processing-element resamples and rescales the input symbol before submitting it to the classification stage. It allows the user to write as big and as fast as he wants.

2.2. Symbol classification

The outputs of the preprocessing module are candidate symbols that are submitted to the classifier. This component is based upon a "Radial Basis Functions" algorithm for learning and recognition vectors of coordinates. Learned symbols, called prototypes, are stored in memory as vectors of (x,y) coordinates. The recognition mode is illustrated in figure 1(b). It returns the N classes of symbols that best match the incoming candidate.

2.3. Dictionary Search

The next component of PHRASES localizes the words of a vocabulary that most closely match the sequence of symbols that is delivered by the recognizer. A preliminary word composer uses the activities of the recognized symbols to build up the list of the most probable candidate words. The mono-processing version can localize up to 40 words per second within a list of 40 000 references.

This component is fully cascadable[5] and can have several children as illustrated in figure 1(c).

3. CHIP SYNTHESIS METHODOLOGY

Each processing element is described using ALMA that supports the synthesis of control dominated applications, and that is a front-end of CAD packages, such as COMPASS design automation that was invoked to generate the examples of this paper. The input of the description of a processing element made of a Finite State Machine and a Data-path[3]. The output is a layout. The structure of the data-path is explicitly described by the designer, given a library of generic resources that range from simple logic gates to complex functional units such as a DRAM controller. The behavior is described by a sequential program, without any explicit timing reference. User can specify loops, conditional branches and subprograms.

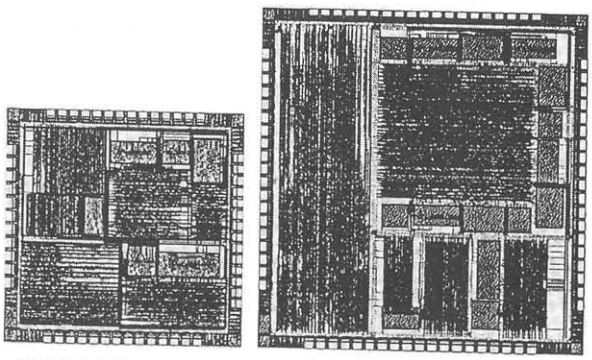

Fig. 2 : Layout of PH1D1 and PH4D3 (COS 0.8 μm)

		Software		PHRASES		NESTOR Ni1000	IBM ZISC	INPG Galatea	IRISA API69
		SPARC 10	PC486D X33	PH1D1	PH4D3				
char	K transistors/chip	-	-	44	114	3,700	400	330	-
	calc. units/chip	-	-	1	4	512	36	4	-
	symbol/sec	80	5	220	880	10,000	1,300	1,000	-
char	K transitors/chip	-	-	72	144				150
	calc. units/chip	-	-	1	3				34
	word/sec	0.5[b] ..22[a]	...2[a]	110[a]	330[a]				50[b]

Table 1 : PHRASES results for the PH1D1 and PH4D3 chips
(a) substitution edit-distance only, (b) substitution, insertion edit-distances

4. RESULTS

We have synthesized the layouts of two chips. They implement the pre-processing, symbol classification and dictionary stages, from the following specifications :

* 4096 x 4096 points input tablet ;
* upper and lower characters (success rate > 93 %);
* 40 000 words dictionary.

In the figure 2, PH1D1 exemplifies a low-cost circuit with one symbol classifier and one dictionary unit. PH4D3 uses four classifiers and three dictionary units that can work in parallel.

Table 1 displays the position of PH1D1 and PH4D3 with respect to existing chips. ZISC, Ni1000 and Galatea are symbol classifiers AP169[2] is a high performance chip for spelling correction.

Performance figures of two software implementations, that originate from our divisions, are given as a reference. They were measured on the same data set that was used for the PHRASES'chips. Due to the lack of a widely accepted benchmark, the performances of the other circuits are only estimated.

The results of PH1D1 and PH3D4 show that the hardware functions are small enough for being integrated on one chip with a core processor. For example, the size of PH1D1 is less than 25 mm^2 in 0.8 μm technology. Per contra, the throughput has been reduced. However, the performances correspond to what is expected for on-line handwriting recognition. Moreover, the results show that the structural parameters may be used to adapt the recognition to the application. This is a key issue for designing a deeply embedded system.

5. CONCLUSION

This paper presented integratable functions for handwriting recognition in the context of embedded systems. Results and comparisons with of-the-shelf circuits as well as with high-end software implementations are encouraging. They demonstrate the strengths of the synthesis method of ALMA that guarantees actual silicon implementation of parameterized architectures.

Eventually we shall use the acceleration that is obtained at symbol and word levels to tackle the syntactic analysis of short sentences in restricted specialisation domains. A first target application is identified for hospital admittance forms.

6. REFERENCES

1) T.R. Halfhill, Byte magazine, 66-86,October 1993.
2) D. Lavenier, Euro ASIC 92, June 1992.
3) J.Y. Brunel, I. Augé et al, 5th Annual European Computer Conference, 596-598, May 1991.
4) P. Gentric, 6th Int. Conf. on Handwriting and Drawing, July 13-15, 1993.
5) J.Y. Brunel, I. Augé et al, IEEE Int. Work. on Rapid System Prototyping, July 1994.