An encoder for vector quantization neural networks

Fabio Ancona, Stefano Rovetta, and Rodolfo Zunino
Department of Biophysics and Electronic Engineering
University of Genova
Via all'Opera Pia 11a - 16145 Genova (Italy)
Phone +39 10 353 2268 - Fax +39 10 353 2175 - E-mail zunino@dibe.unige.it

Abstract— **Large-scale parallelism and analog computation are exploited to obtain a neural module, suitable for both functioning and training, since appropriate signal lines are provided. The VQ encoder is self-contained and therefore can be embedded into any system, either analog or digital. It implements efficiently the vector matching operations, therefore it can be exploited in systems based on any vector quantization algorithm, with good throughput.**

I. Introduction

Vector quantization circuits are at the heart of many neural systems (those based on "competitive" algorithms, including Kohonen's SOMs [1], the Neural Gas model by Martinetz *et al.* [2], and many more). Digital architectures are usually adopted when implementing vector quantization. However, some operations are best implemented in a highly parallel analog system, since they require a high number of non-interacting, simple processing steps. A typical instance of this class of operations is the distance computation and matching in nearest-neighbor-based algorithms, such as vector quantization.

This paper presents an alternate approach exploiting analog VLSI. The relatively small number of components needed to implement the distance and matching functions make it possible to obtain a good level of parallelism. Modular design allows the resulting circuit to be exploited in a highly scaleable chipset, such that the number of neurons can be incremented by simple addition of new chips, while the throughput keeps constant and independent of the number of neurons.

II. Vector-quantization neural networks

Neural models based on vector quantization are usually represented by a collection of reference vectors (or prototypes), such that each neuron stores one prototype. The neurons enter a competition to select which one is the nearest to the input vector; the "winner" remains on, while the other neurons are shut off. The output is encoded by the position of the winner. Accordingly, the vector quantizer replicates several instances of the following sub-blocks: computation of distance between two vectors; selection of the nearest neighbor; memory for prototype vector.

The first two blocks are realized in analog VLSI, paying special attention to two parameters: constant operation with respect to circuit parameters and environmental conditions, and simple configuration. The third block is based on a digital RAM, permanently storing the neuron components. The digital values refresh analog local representations of the neuron components, used to actually perform the distance computations. The paper describes in more detail the circuital realization of these modules.

The overall scheme involves an array of neurons, each implemented by a set of values (the vector components) and a distance computation block. All neurons compute their distance from the input vector in parallel. Then, the distances enter the best-match block (or winner-take-all) which yields the circuit output.

At the time of writing, the chip is being produced in a configuration composed of 40 neurons with 16 inputs each. The size of the chip is about 2mm by 2mm.

III. The vector-distance subsystem

The distance computation implements the following operation:

$$d = \sum_i (x_i - w_i^{(k)})^2$$

where x_i is the i-th component of the input vector and w_i the corresponding component of the k-th reference vector, corresponding to the k-th neuron. The circuit reads x_i as a voltage levels. This ensures easy interfacing with other circuits. However, its output is in the current form, so that the operation of summation is straightforward and is obtained by simple wiring. Each component is processed by a circuit that computes a squared difference, shown in Figure 1. The standard Gilbert configuration has been discarded, since its wide-range version implies a higher number of transistors.

This scheme yields the desired output plus an additive term which is compensated for by the subcircuit M7...M9. However, we choose not to eliminate completely the additive constant, because it is useful for biasing the subsequent block.

0-7803-3694-1/97/$10.00 © 1997 IEEE

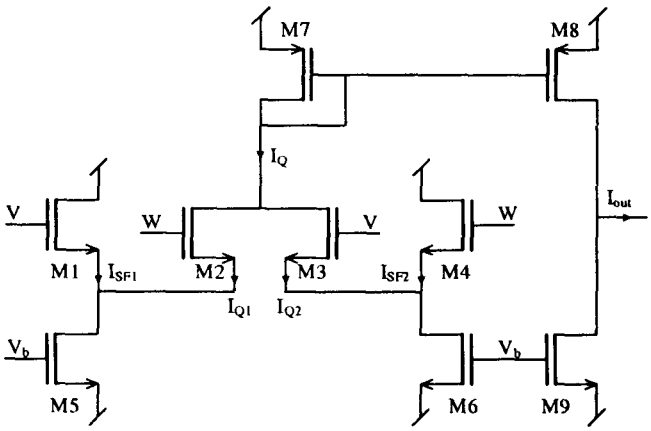

Fig. 1. Diagram of the square–of–difference circuit.

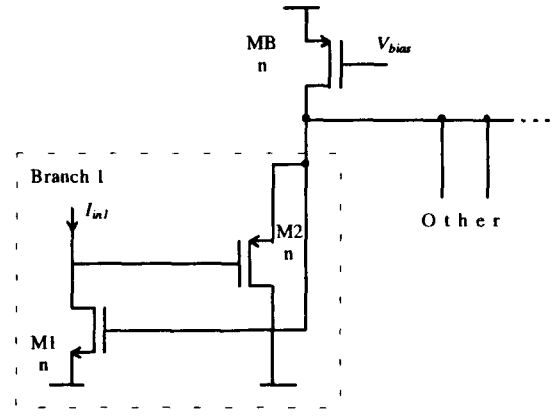

Fig. 2. Basic circuit for the minimum-selection function.

IV. The best–match subsystem

There are many approaches to the competition function, or WTA, which is here used to select the best–matching codevector. The scheme adopted in the present work [3] is based on the model by Lazzaro *et al.* [4], which is the most well-known although many other exist. The scheme by Lazzaro *et al.* is simple and smart, requiring only two transistors per input branch. Many drawbacks of the original circuit have been here compensated for, resulting in a circuit that is a bit more complex but behaves as required.

This configuration may suffer from a reduced resolution (sensitivity to differences in input currents). Moreover, a large number of input branches causes interaction, and the circuit fails to work properly. The input impedance is quite large. If the input current is generated by a current mirror, this generates excessive voltage variations, with loss of accuracy of the mirror. Other shortcomings of this scheme are not relevant for the present application, so they should not be taken into account. An example is the constant current consumption, which is negligible when compared to the whole chip since it grows as the number of neurons, not as the number of weights.

The schematic of the circuit adopted is shown in Figure 3. To achieve a better input impedance it is sufficient to modify the operating point of M1 [4][5], which amounts to adding a bias current to the input. This is obtained by a proper sizing of the output stage of each square-of-difference block, such that, with zero input, their sum yields the required bias current (around 90μA). To improve the performance of the current-mode input, we adopt a cascoded current-mirror configuration (Figure 3: M5–M7).

The new operating point causes the transistors to operate in a linear fashion, without changing operating zone. As a result, the common node presents voltage variations proportional to the winning current variations: $V_{out} = R\, I_{winner}$. Therefore, these variations can simply be read at the output (through subcircuit M18–M19). The proportionality coefficient is quite constant and reasonably insensitive to parameter variations [3]. The linear output allows the use of the vector quantizer in on–line adaptive circuits, since the neural updating rules always involve this quantity. Moreover, it allows comparison of the output of several similar circuits, which allows parallel operation by means of a single, additional WTA chip to process all outputs.

The circuit configuration shown in Figure 2 illustrates the principle of minimum selection. Because of the p-channel MOS M2, this configuration enhances the current on the branch with minimum input: the minimum input current causes the minimum drain-source voltage of M1 (Early effect), which in turn causes an increase of the drain-source voltage of M2. From now on, the behavior is perfectly dual to that of the standard configuration. This circuit section can be identified with M5–M7 in Figure 3.

Since this configuration is not guaranteed to possess a feasible set of operating conditions (compatible with the desired operation), it is necessary to relax some constraints by allowing the common node voltage and the output of each branch to be different. This is done through a level-shifter. The voltage difference (1.5V) is easily achieved. However, a simple source-follower is not sufficient for reasons of speed, as a frequency analysis points out. Consequently, the level shifter is somewhat more complex (Figure 3: M8–M11). This helps overcoming the effects of the long capacitive structure implementing the common node in the scheme, without need for a buffer.

Finally, M12–M17 implement a current comparator to convert the 1-out-of-n current output into a digital voltage

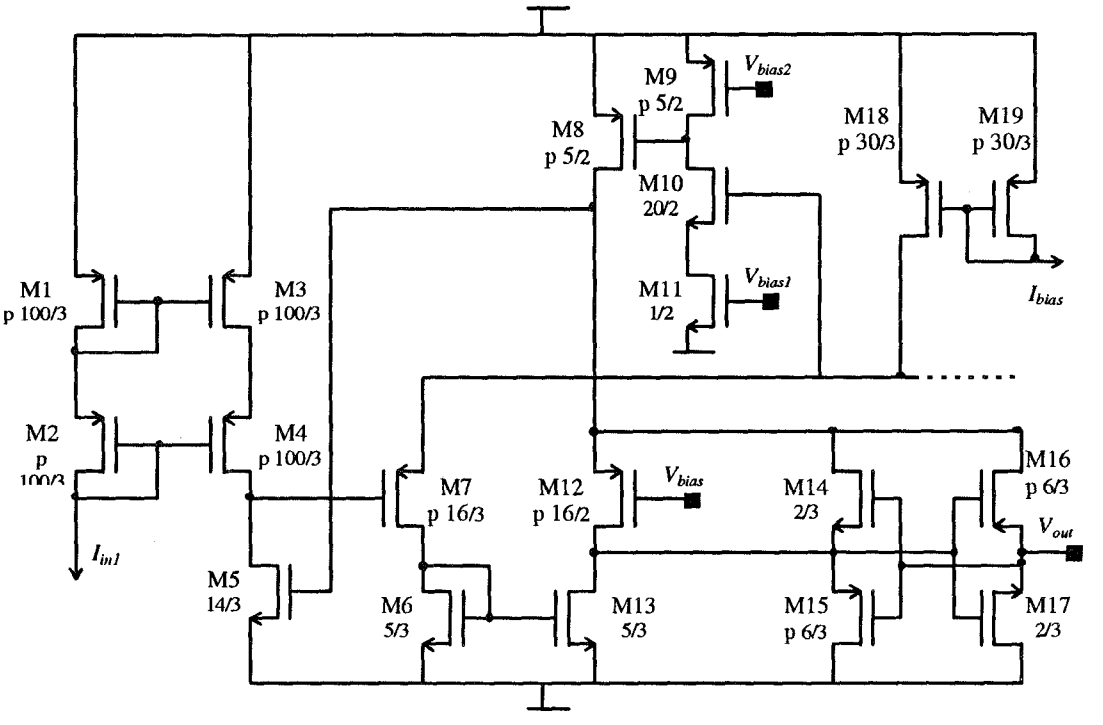

Fig. 3. Complete schematic of the competition and selection circuit.

level.

The digital memory is interfaced with the system by an analog bus driven by an array of digital–to–analog converters. There is a converter for each group of components, and the array is multiplexed over the neurons. The refresh circuitry makes use of dummy transistors [6] to avoid leakage in the commutation phase. This allows a longer refresh periodicity (more than 1ms) and consequently a smaller percentage of idle time for each neuron.

V. The memorization subsystem

The analog memory is realized according to the simple principle of sample-and-hold. Therefore it requires an accurate sizing to minimize the undesired effects of the simple configuration and achieve maximum precision. The principal effects to be taken into account are the static charge loss through switches (leakage currents) and the dynamic charge injection (clock feed-through).

The capacitor is implemented with a simple nMOS with source and drain connected together. The capacity is therefore C_{GS}. This simple approach yields a better linearity when compared with the double-poly solution (not available in the technology adopted). The capacity could be integrated in the distance computation block, but this results in an unbalanced structure, implying asymmetric behaviour.

To estimate the capacitor size required for at least 2-3 pF, a rough computation yields at least $(W/L) = (40\mu/40\mu)$, quite a large area.

To optimize the memory system performance, it is necessary to focus on the realization of the analog switch. The central point is to compensate for clock feed-through [7][8][9]. CFT is caused by presence of charge in the non-ideal switch, the pass transistor, which is a MOS device with its own channel. This charge combines with the stored charge on the capacitor, altering its value. This effect is compensated for by a design including dummy switches [6], shown in Figure 4. A dummy switch, added in parallel to the actual switch and driven by a reversed clock, has no effect on the switch function, but contributes to the charge injection mechanism attenuating the CFT effect.

VI. Auxiliary components

The control system for memory addressing and the output encoding are digital subsystems. They are custom-built using precharged logic technology. This choice helps avoiding an excessive power dissipation and the deep level hierarchy featured by static schemes. The speed performances are not reduced, however, since the speed of the system result from parallel operation, not from high clock frequency; moreover, the various processing phases are

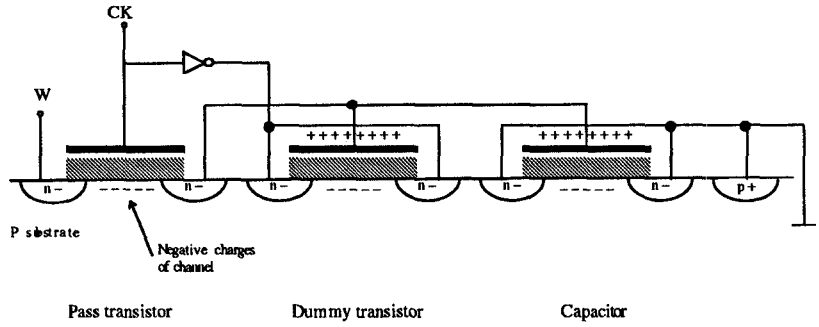

Fig. 4. Layout of devices in the analog memory.

pipelined for better efficiency.

The output encoding translates the 1-out-of-64 code of the competition block into a 6-bit digital word. The precharged NOR logic would yield an inverted output. This is avoided by placing inverters at the end of the six lines, which also serve as output buffers (necessary in any case).

The analog output pin (value of the winner) is driven through a buffer amplifier. The amplifier standard cells available in CAD software libraries are not adequate, since they are oversized with respect to the needs of this project. Therefore an original circuit has been designed. The emphasis has been focused on minimizing offset, to ensure a good linearity of the analog signal (required to maintain the necessary precision).

The output buffer amplifier is a wide-range operational transconductance amplifier [10][11], with a cascode-configured output stage. In the present case, the OTA scheme is more appropriate than the simpler Miller configuration, where the load to be driven features large capacitive components. At working frequencies (for instance 10Mhz) the Miller scheme cannot drive capacitive loads larger than about 10-15pF. The wide-range configuration helps in obtaining a linear respones even at the extremes of the output signal range.

The biasing components for the whole circuit are realized using the "resistive interpolation biasing" technique [12], which compensates for variations in parameters due to the extension of long biasing structures across the chip.

Acknowledgments

The authors acknowledge the valuable contribution of G. Oddone, G. Uneddu and A. Novaro to this research, as part of their "Laurea" thesis projects.

Part of this work was funded by a MURST 40% research project. The chip is being produced thanks to a financial contribution of the University of Genova for the research training of graduate and undergraduate students.

References

[1] Teuvo Kohonen, *Self Organization and Associative Memories*, Springer, 3rd edition, 1989.

[2] T.M. Martinetz, S.G. Berkovich, and K.J. Schulten, "'Neural gas' network for vector quantization and its application to time–series prediction", *IEEE Transactions on Neural Networks*, vol. 4, no. 4, pp. 558–569, 1993.

[3] Fabio Ancona, Giorgio Oddone, Stefano Rovetta, Gianni Uneddu, and Rodolfo Zunino, "Enhanced WTA network with linear output and stable transimpedance", *Alta Frequenza*, vol. 8, no. 5, pp. 71–73, September 1996.

[4] J. Lazzaro, R. Ryckebush, M. A. Mahowald, and C. Mead, "Winner-take-all networks of $O(n)$ complexity", in *Advances in Neural Information Processing Systems II*, San Mateo, 1989, pp. 703–711, Morgan Kaufmann.

[5] M.E. Robinson, H. Yoneda, and E. Sánchez-Sinencio, "A modular CMOS design of a Hamming network", *IEEE trans. on neural networks*, vol. 3, no. 3, pp. 444–456, 1992.

[6] Cristoph Eichenberger and Walter Guggenbuhl, "On charge injection in analog CMOS switches and dummy switch compensation techniques", *IEEE Transactions on Circuits and Systems*, vol. 37, no. 2, pp. 256–264, February 1990.

[7] Y. Tsividis and P. Antognetti, *Design of MOS VLSI circuits for telecommunications*, Prentice-Hall, Englewood Cliffs, NJ, 1985, first edition.

[8] W. Wilson, H. Mussoud, R. Swanson, R. George, and R. Fair, "Measurement and modelling of charge feedthrough in n-channel mos analog switches", *IEEE Journal of Solid-State Circuits*, vol. 20, pp. 1206–1213, 1985.

[9] G. Wegmann and E. Vittoz, "Basic principles of accurate dynamic current mirrors", *IEE Proceedings-G*, vol. 22, no. 2, pp. 95–100, 1990.

[10] F. Soin, F. Maloberti, and J. Franca, *Analogue-digital ASICs circuit techniques*, Peter Peregrinus Ltd., London, 1990.

[11] P.R. Gray and R.G Meyer, "MOS operational amplifier design: a tutorial overview", *IEEE Journal of Solid-State Circuits*, vol. 17, pp. 969–982, 1982.

[12] S. Satyanarayana and K. Suyama, "Resistive interpolation biasing: a technique for compensating linear variation in array of MOS current sources", *IEEE Journal of Solid-State Circuits*, vol. 30, pp. 595–598, 1995.