Minimal-connectivity circuit for analogue sorting

S.Rovetta and R.Zunino

Abstract: A CMOS circuit for sorting analogue current-mode quantities is presented. The highly modular architecture integrates several elementary cells operating at the local level. The VLSI-oriented approach minimises wiring and silicon area, because very few devices are involved. The sorting process is completed in $O(n)$ time. Simulations at the VLSI layout level prove the effectiveness of the approach in neural-network training applications.

1 Architecture

The sorting of analogue quantities is a crucial task in several signal-processing applications, such as median filtering, neural-network training, etc. The literature offers a variety of HW sorting circuits, whose complexity lies mainly in the need for circulating information among a system's components. Neural-network approaches to sorting use Hopfield models [1]; a dynamic structure has recently been presented in [2]. VLSI sorting circuits often limit connections with hierarchical pairwise comparisons [3]. Local comparisons make for better matching and enhance robustness by higher accuracy [4]; on the other hand, a system's area complexity tends to increase. A circuit approach providing a reduction in corner and offset errors is described in [5]. The literature ultimately indicates that sorting robustness somehow relates to interconnection complexity. In fact, the simplest architecture for HW rank extraction requires one wire connecting elementary cells; such a minimal-connectivity solution is most favourable for VLSI systems, thanks to its easier implementation. Therefore a single-wire connection is adopted in the present article.

The paper presents a novel sorting circuitry specifically aimed at CMOS implementation. The circuit supports partial sorting and operates iteratively: the length of the sorted list K can be digitally programmed, and sorting is completed in $O(K)$ time. The basic schema exploits Lazzaro's winner-takes-all (WTA) structure [6]; the modular approach, however, makes it possible to include alternative schemata improving accuracy or speed [7]. Sorted quantities are represented in current mode and individually processed by simple cells. Design considerations indicate that the overall circuit performance is notably independent of process fluctuations.

The architecture has been developed for integration in a neural hardware training system, for which basic modules have already been realised [8]. The sorting circuit can handle up to 64 operands and its performance has been thoroughly assessed by simulation of the final VLSI layout.

IEE Proceedings online no. 19990323

DOI: 10.1049/ip-cds:19990323

Paper received 15th October 1998

The authors are with the University of Genoa, Department of Biophysical and Electronic Engineering, Via all'Opera Pia 11a, 16145 Genova, Italy

2 Sorting operation

Each input current is handled by a single cell; the single cell-interconnection wire provides an analogue, voltage-mode representation of output values. The circuit operates iteratively and follows a straightforward principle: at each iteration the largest value on the list is identified, represented at the analogue output, and removed from the input set. A clock signal times the process iterations. Each iteration includes the following steps:

(i) (competition) the current largest input value is detected by a WTA schema; such a value is linearly represented at the output voltage V_{out}

(ii) (winner selection) each cell locally detects the output of the competition

(iii) (winner removal) the winning cell disables its WTA subcircuit and thus removes itself from the list of competitors.

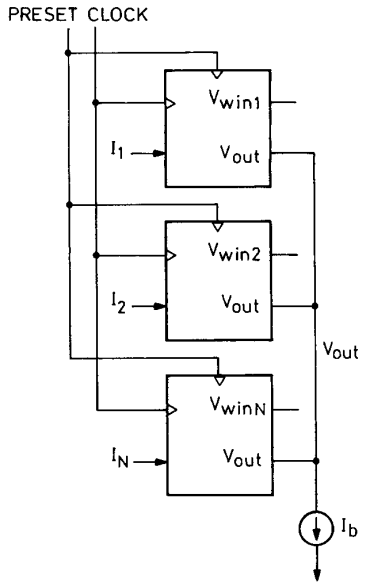

Fig. 1 *Circuit cell interconnections*

The cell-interconnection schema is sketched in Fig. 1; Fig. 2 presents the circuit of an elementary cell that integrates analogue and digital circuitries. Each cell has an internal digital status variable, stored by the charge in

capacitor C_2. A null voltage on C_2 indicates that the cell has not yet won a competition, hence it represents an active competitor; otherwise, the cell is disabled. Likewise, a high-level potential on capacitor C_1 indicates that the cell is the winner of the current iteration. The output wire V_{out} connects the gates of transistors M_1 of all the cells and is biased by a constant current I_b. At startup, a PRESET pulse forces the STATUS voltage to be high in all the cells; all units are enabled and switch M_{sw0} is closed everywhere.

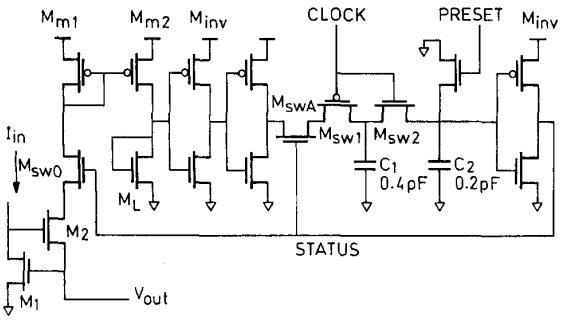

Fig. 2 *Circuit of basic sorting cell*

2.1 WTA competiton

The largest input value at each iteration is detected by an augmented, above-threshold version of Lazzaro's WTA subcircuit [6]. The sets of transistor pairs M_1, M_2 of all the cells constitute the basic WTA subcircuit, which is sketched in Fig. 3. Thanks to the Early effect on MOS devices, positive feedback propagates through line V_{out}, and the entire bias current I_b flows in transistor M_2 of the cell with the largest input current. If the input currents in transistors M_1 are suitably designed [8], V_{out} also yields an analogue, linear, representation of the winning value.

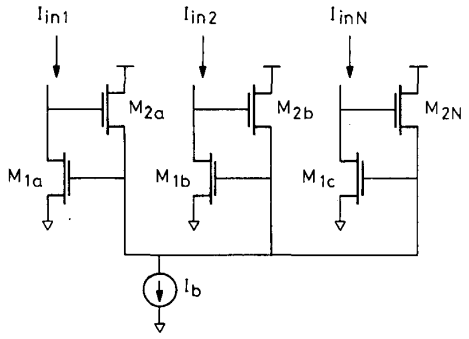

Fig. 3 *Lazzaro's WTA subcircuit*

2.2 Winner detection

The winner-detection subcircuit M_{m1}, M_{m2}, ML exploits the nonuniform distribution of the bias current I_b. As long as switch M_{swA} is kept closed by a high STATUS voltage (default condition), the pair M_{m1}, M_{m2} and the load transistor M_L drive voltage V_x. If the cell is not the winner, no current is mirrored into M_L, and V_x is low. Therefore V_x is high only in the cell that wins the current competition. The double inverter stage M_{inv} converts V_x to a digital level; the output voltage V_{win} is the actual digital indicator of the currently winning cell.

2.3 Synchronous winner removal

CMOS logic exploits this information to disable the winner. As long as CLOCK is low, M_{sw1} is closed and V_{win} is stored in capacitor C_1. The rising edge of CLOCK opens M_{sw1} and transfers this information through M_{sw2} to the cell internal status in capacitor C_2. At each iteration, only the winning cell switches its status (voltage on C_2) from low to high; the consequent low STATUS level turns M_{sw0} off, thus the winner cell will remain unaffected at the following iterations until the next PRESET pulse. At the same time, the low STATUS level opens the switch M_{swA}, thus forcing the drain of M_2 to a high-impedance state and removing the cell from future competitions. The falling edge to the next CLOCK cycle freezes this situation and triggers a new competiton among the remaining cells.

To sum up, the highest input value is mapped at V_{out} at the first iteration; as soon as the winning cell exits the competition, the second highest value is enabled to show up, and so forth. The eventual sorted list is represented by the sequence of potentials V_{out} at the consecutive clock cycles. The process can be stopped in different ways, depending on the mode of operation. Programmable partial sorting (extracting the largest K elements from the input set) is accomplished by presetting a counter to issue a PRESET pulse after K clock cycles ($K < N$). Otherwise, an automated stop is easily implemented by inserting a dummy cell with the smallest possible input current, and by using the associate winner signal to issue a PRESET pulse. This allows the designer to bypass the need for predetermining the number of input lines and for providing additional circuitry to count clock cycles.

Table 1: Aspect ratios and sizes

Device	*W/L* ratio
M_1, M_2, $M_{Drv}(p)$	40/2
$M_{Drv}(n)$	10/2
M_{mx}, M_{sw1}, $M_{inv}(p)$	5/2
ML	1/10
M_{sw0}, M_{sw2}, M_{swA}, M_{PRESET}, $M_{inv}(n)$	2/2
C_1; C_2	25 × 25; 16 × 16

Unit channel length = 1μm

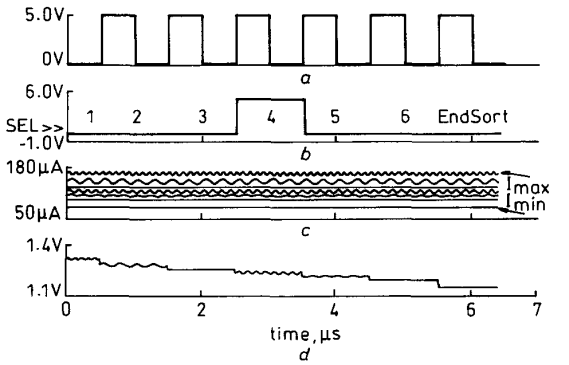

Fig. 4 *Sample simulation results*
a Clock signal
b Digital output V_{win} of cell 4; graph marks activation interval of each cell
c Input currents
d Analogue output voltage

3 Experimental results

3.1 Implementation details

The MOST aspect ratios for the cell subcircuit are given in Table 1. The power supply is 5V and the input currents may be in the range [50, 200] μA to ensure compatibility with the external neural circuitry [8]. The bias current I_b = 30μA; the standard clock frequency is 2MHz, but clock rates higher than 10MHz are admissible. The circuit was

simulated at the layout level by using HSPICE level 13 with 1μm ES2 technology and operated correctly under stressing conditions (64 inputs). Fig. 4 presents an example of the circuit operation; for simplicity and without loss of generality, a demonstration involving six cells is shown. The sequence of analogue outputs at V_{out} witnesses the correct functioning of the circuit and the linear mapping of input/output values.

3.2 Robustness and power consumption issues

An appealing feature of the circuit is its insensitivity to design parameters. Digital logic devices are very small, and the range of frequencies considered does not impose any strict constraints on their precision; the winner-detection circuitry M_{m1}, M_{m2}, M_L processes digital currents (the drain current in M_{m1} is either I_b or null), hence exact mirroring is not required. The only devices subject to matching are the WTA pairs M_1, M_2. In fact, this requirement depends on the specific WTA schema, which has been chosen for its simplicity, even though it may be replaced with many alternative ones [8], without affecting the general validity of the approach.

CMOS structures notably limit power dissipation in the digital subparts; however, the schema proves effective in its analogue components as well: during each iteration, only one branch (M_2 – M_{m1}) throughout the entire circuit pulls the bias current I_b; in addition, the mirroring pair Mm1, M_{m2} might be suitably designed to attenuate the digital current in M_L. Thus, disregarding the dissipation due to digital-logic switching, the power consumption is due to two terms

$$P_{TOT} = V_{cc}(I_b + I_L) + V_{cc}\sum_{n=1}^{N} I_n \qquad (1)$$

which shows that the contribution of the sorting circuitry to the total power dissipation is constant and independent of the number of inputs. In the present implementation a total power dissipation of 30mW was measured. A final remark concerns the total occupied area, which is kept to a minimum thanks to both the relatively small capacitances and the digital circuitry involved.

4 References

1 WANG, J.: 'Analysis and design of an analog sorting network', *IEEE Trans. Neural Netw.*, 1995, **6**, (4), pp. 962–971

2 DRANGER, T.S., and PRIEMER, R.: 'Collective process circuit that sorts', *IEE Proc. Circuits Devices Syst.*, 1997, **144**, (3), pp. 145–148

3 URAHAMA, K., and NAGAO, T.: 'Direct analog rank filtering', *IEEE Trans. Circuits and Syst. I. Fundam. Theory and Appl.*, 1995, **42**, (7), pp. 385–388

4 DEMOSTHENOUS, A., SMEDLEY, S., and TAYLOR, J.: 'A CMOS analog winner-take-all network for large-scale applications', *IEEE Trans. Circuits Syst. I, Fundam. Theory and Appl.*, 1998, **45**, (3), pp. 300–304

5 OPRIS, I.E.: 'Analog rank extractors', *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, 1997, **CAS–44**, (12), pp. 1114–1121

6 LAZZARO, J., RYCKEBUSH, R., MAHOWALD, M.A., and MEAD, C.: 'Winner take all networks of O(n) complexity' *in* 'NIPS 2' (Morgan Kaufmann, 1989), pp.703–711

7 SANCHEZ-SINENCIO, E., and GUNAY, Z.S.: 'CMOS winner-take-all circuits: a detail comparison'. Proceedings of ISCAS'97, Hong Kong, 1997, Vol. 1, pp. 41–44

8 ANCONA, F., ODDONE, G., ROVETTA, S., UNEDDU, G., and ZUNINO, R.: 'A vector quantization circuit for trainable neural networks'. Proceedings of 3rd international conference on *Electr. circuits and systems*, ICECS'96, Rodos, 1996, Vol. 2, pp. 1131–1134