

Towards a silicon primordial soup: A fast approach to hardware evolution with a VLSI transistor array

Jörg Langeheine, Simon Fölling, Karlheinz Meier, Johannes Schemmel

Adress of principle author: Heidelberg University, Kirchhoff-Institut für Physik,
Schröderstr. 90, D-69120 Heidelberg, Germany, ph.: ++49 6221 54 4359
`langehei@kip.uni-heidelberg.de`

Abstract. A new system for research on hardware evolution of analog VLSI circuits is proposed. The heart of the system is a CMOS chip providing an array of 16×16 transistors programmable in their channel dimensions as well as in their connectivity. A genetic algorithm is executed on a PC connected to one or more programmable transistor arrays (PTA). Individuals are represented by a given configuration of the PTA. The fitness of each individual is determined by measuring the output of the PTA chip, yielding a high test rate per individual. The feasibility of the chosen approach is discussed as well as some of the advantages and limitations inherent to the system by means of simulation results.

1 Introduction

Analog circuits can be much more effective in terms of used silicon area and power consumption than their digital counterparts, but suffer from device imperfections during the fabrication process limiting their precision ([1]). Furthermore any progress in design automation seems to be much harder to achieve for analog circuits than for digital ones. Recently the technique of genetic algorithms (GA) has been applied to the problem of analog design with some promising results. The capability of artificial evolution to exploit the device physics available has been demonstrated for example in [2].

The various attempts to the evolution of analog circuits range from purely extrinsic evolution that simulates the circuits composed by the genetic algorithm (GA) (see e.g. [4]) over breeding of analog circuitry on FPGA's designed for digital applications ([2], [3]) and the use of external transistors (e.g. [5]) as well as the optimisation of parameters of an otherwise human designed circuit ([6]) to the use of chips designed with certain design principles in mind ([7]).

However, one of the most elementary devices for analog design is given by the transistor, which can also be used to form resistors and capacitors. Therefore the development of a hardware evolution system employing a programmable transistor array (PTA) to carry out the fitness evaluation intrinsically in silicon is proposed.

The evolutionary system aimed at will hopefully yield the following advantages compared to the approaches described above: Except for dc-analyses the evaluation of the fitness of the according individual can be carried out much faster than in simulation, hopefully yielding a cycle time of less than 10 ms. In contrast to simulations the GA has to deal with all the imperfections of the actual dice produced and thus can be used to produce circuits robust against these imperfections (a discussion on the evolution of robust electronics can be found in [3]). Furthermore, circuits evolved on a custom made ASIC can be analyzed more easily than on a commercial FPGA, since in general the encoding of the loaded bitstrings is not documented for commercial FPGAs. On the contrary for custom made ASICs fairly well suited models can be used to simulate the circuit favoured by the GA. Analysis of the attained circuits will be further enhanced by the integration of additional circuitry allowing to monitor voltages and currents in the transistor array itself as well as the die temperature. Finally, a lot of the structures designed for the PTA could be reused to set up a whole family of chips that merely differ in the elementary devices used for evolution. The programmable transistors could be replaced, for example, by programmable transconductance amplifiers, silicon neurons, or operational amplifiers combined with a switched capacitor network. Thereby different signal encodings (voltages, currents, charge, pulse frequencies and so on) could be investigated and compared with regard to their 'evolvability'.

The rest of the paper is organized in a bottom up fashion: Section 2 describes the structure of the programmable transistor cell. In section 3 the PTA chip as a whole is discussed, while section 4 presents the embedding of the chip in a system performing the artificial evolution. Finally in section 5 some simulation results will be given demonstrating the feasibility as well as the limits of the proposed approach, before the paper closes with a summary.

2 Design of the basic transistor cell

2.1 Choosing the transistor dimensions and types

The core of the proposed chip consists of an array of 16×16 programmable transistor cells, each acting as one transistor in the programmed circuit. Each cell contains an array of 20 transistors providing 5 different lengths and 4 different widths. The 5 different lengths are decoded by three bits and vary logarithmically from $0.6 \mu\text{m}$ to $8 \mu\text{m}$. Since a transistor of a given width W can be approximated fairly well by using two transistors of the same length and width $W_1 + W_2 = W$ in parallel, 15 different widths can be chosen ranging from 1 to $15 \mu\text{m}$. Relatively small steps between adjacent lengths and widths are chosen to attain a fitness landscape smooth with regard to variations of the channel dimensions of the used transistors. A rather large number of different lengths is chosen, because the characteristics of a MOS transistor do not only depend on the aspect ratio (W/L), but also on its actual length (cf. [8]).

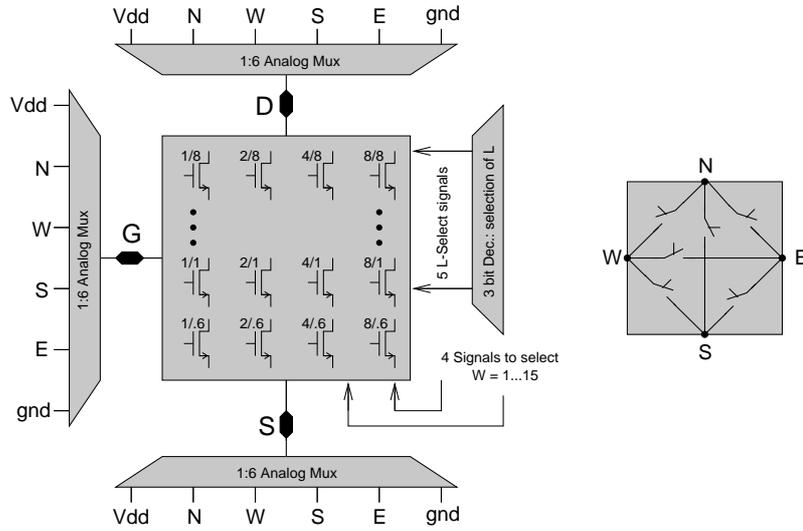


Fig. 1. Basic NMOS cell: **Left:** The cell contains 4×5 different transistors. **Right:** Routing capability of one cell.

2.2 Connecting the chosen transistor

The chosen transistors are connected to the three terminals D (drain), G (gate) and S (source) through transmission gates (Fig. 1 shows an NMOS cell. Since the process used is based on p-doped silicon, there is an additional terminal B for the bulk connection of the PMOS cell). Each of these 'global' transistor terminals can be connected to either of the 6 terminals listed at the output of the 1:6 analog multiplexer. The terminals represent connections to the four adjacent transistor cells, power and ground. In this architecture the signal has to pass two transmission gates from the actually chosen transistor to the outside of the cell, resulting in twice the on-resistance of one transmission gate. On the other hand the two-fold switching saves a lot of transmission gates ($20 + 6$ instead of 20×6 for each 'global' terminal) and thereby a lot of parasitic capacitance. On the right hand side of Fig. 1 the routing capability of the cell is shown. Signals can be connected from any of the four edges of the cell to any of the three remaining edges via one transmission gate. In total

$$4 (W) + 3 (L) + 3 * 3 (\text{Multiplexing of D,G,S}) + 6 (\text{Routing}) = 22 \quad (1)$$

bits are needed for the configuration of one NMOS cell. In case of the PMOS cell another three bits must be added to multiplex the bulk terminal.

3 The PTA chip

For the chip an array is formed out of 16×16 programmable transistor cells as indicated in Fig. 2. PMOS and NMOS transistors are placed alternately

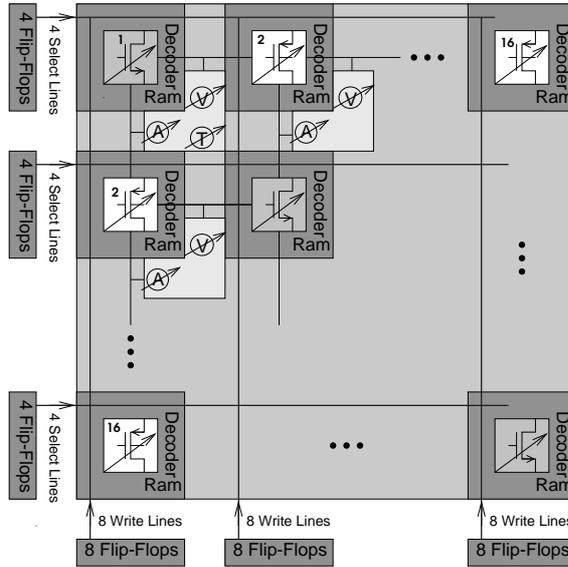


Fig. 2. Array of programmable transistor cells including S-RAM cells for configuration storage and readout electronics allowing to monitor voltages and currents inside the network as well as the die temperature. The digital parts of the array are shaded in darker gray than the analog ones. PMOS cells are distinguished from the NMOS cells by their white shading.

resulting in a checkerboard pattern. With 256 transistors in total, circuits of fairly high complexity should be evolvable.

3.1 Configuration of the programmable transistor cells

The configuration bits for determining the W and L values and the routing of each cell are stored in 32 static RAM bits¹ per cell, yielding $256 * 32 = 8192$ bits for the configuration of the whole chip. In order to test individuals with a frequency of 100 Hz, the configuration time should be small compared to a cycle time of 10 ms, e.g. $100 \mu\text{s}$. Writing 8 bits at a time as shown in Fig.2 thus results in a writing frequency of about 10 MHz, which is easily achieved with the $0.6 \mu\text{m}$ process used. As shown in Fig. 3 the actual configuration can be read back, which is useful to control the writing of the configuration data as well as their stability in the possibly noisy environment. The digital parts of each cell are laid out around the analog parts to allow fast replacement of the core part of the cell enabling the testing of different 'elementary' devices as already discussed in section 1.

¹ The choice of 32 instead of the needed 25 is taken with regard to future implementations of different corecells into the same surrounding infrastructure.

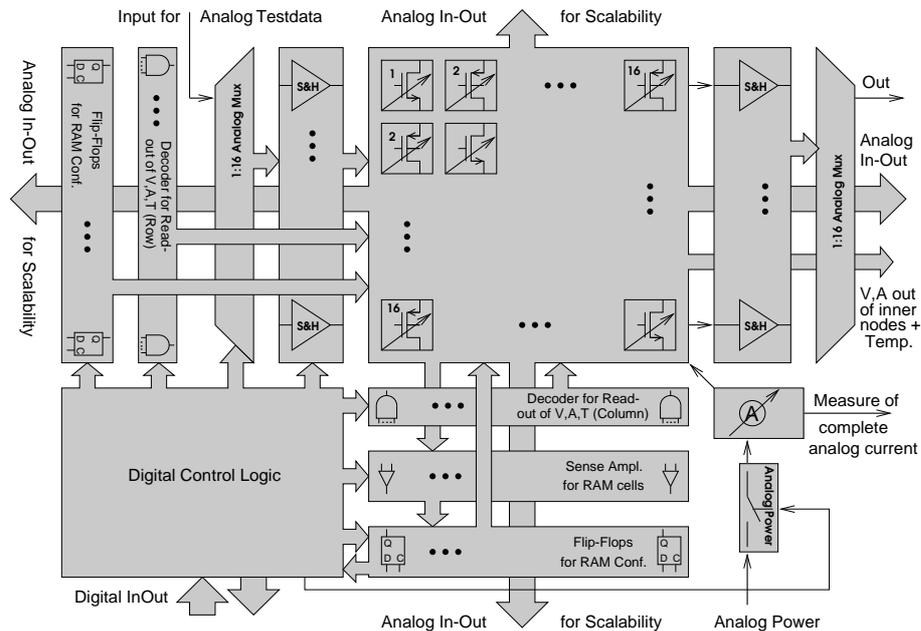


Fig. 3. Block diagram of the PTA chip. The analog multiplexer and the sample and hold stages for in- and output of the border cells are shown on one side of the transistor array only for simplicity. Actually all 64 border cells can be read out and written to.

In order to allow all possible configurations to be loaded into the PTA without causing the chip to destroy itself, two precautions are taken: Firstly all metal lines are made wide enough to withstand the highest possible currents expected for this particular connection. Secondly the analog power for the transistor array can be switched off, if the die temperature exceeds a certain limit.

3.2 Analog in- and output signals

The input for the analog test signals (meant to be voltages) can be multiplexed to the inputs of any of the 16 cells of any of the four edges of the array (For simplicity this is shown for one edge only in Fig. 3). The signal for each input is therefore maintained by means of a sample and hold unit. Similarly all of the 64 outputs at the border of the array are buffered by sample and hold circuits that can be multiplexed to the analog output (Again only shown for one edge in Fig. 3). That way the in- and output(s) of the circuit being evolved can be chosen freely, such that different areas of the chip can be used for the same experiment without loss of symmetry for the signal paths.

The output amplifiers are designed to have a bandwidth higher than 10 MHz, which should be sufficient regarding the expected bandwidth of the transistor array (cf. section 5).

All 64 outputs can be accessed directly via bond pads (referred to as 'Analog In-Out for Scalability' in Fig. 3). That way the outputs can be connected to external loads, and experiments using more than 256 programmable transistor cells can be set up by directly connecting an array of dice via bond wires.

3.3 Monitoring of node voltages, intercellular currents and temperature

As already mentioned in section 1 the PTA chip offers the possibility to read out the voltages of all intercellular nodes as well as the currents flowing through the interconnection of two adjacent cells. While the former one can be read out with several MHz, current measurements are limited to bandwidths much smaller than the bandwidth of the transistor array in order to limit the area occupied by measurement circuitry to a reasonable percentage of the chip area. Since one cannot (and maybe does not even want to) prevent the GA from evolving circuits wasting power or oscillating, the environment in the PTA may be very noisy. Accordingly the node voltages have to be buffered as closely to their origin as possible. For the same reason the current and temperature signals are locally amplified and transformed into differential currents that are multiplexed to the edges of the transistor array, where they are transformed into voltages mapped to the according pads (Fig. 3 indicates the multiplexing).

The power net for the transistor array is separated from the analog power of the rest of the chip to be able to measure the overall current consumption as well as to provide different supply voltages allowing for example the evolution of low voltage electronics.

4 The system around the chip

The architecture of the evolution system is shown in Fig. 4: The GA software

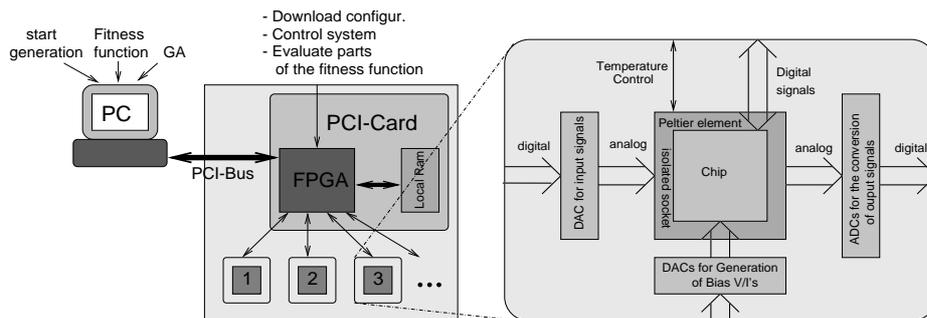


Fig. 4. Architecture of the system: **Left:** Motherboard housing several daughterboards connected to the PC via a PCI card. **Right:** One of the daughterboards presented in closer detail.

is executed on a commercial PC and communicates with the PTA chip using an FPGA on a PCI card. The FPGA gathers PTA output data and does basic calculations, the software then uses the preprocessed data to calculate the circuit's fitness. The PTA board and the software operate asynchronously to minimize cycle time. For further increase of the evaluation frequency the software is designed for maximum scalability in terms of additional processors, boards or additional computer/board systems. The FPGA distributes the digital signals for the DACs and the ADCs to the daughterboards shown on the right side of Fig. 4. A daughterboard contains a DAC for the evaluation input signal generation, DAC's providing some bias voltages and currents and ADCs for the conversion of the analog output signals of the chip (i.e. the output of the transistor array, the measured total current consumption, the signals representing die temperature and intercellular node voltages and currents.). Furthermore, a temperature control via a Peltier element will be implemented to control the temperature of the chip.

As was already pointed out by [3], the usefulness of evolved circuits strongly depends on their robustness against variations of the environment experienced by the chip. In order to provide a possibility for the evolution of robust electronics the system will be capable of testing the same individuals in parallel under different conditions, e.g. different temperatures or different dice (maybe even wafers).

5 Simulation results

As already mentioned in section 2 parasitic resistance as well as parasitic capacitance is introduced with every switch used for choosing a special transistor

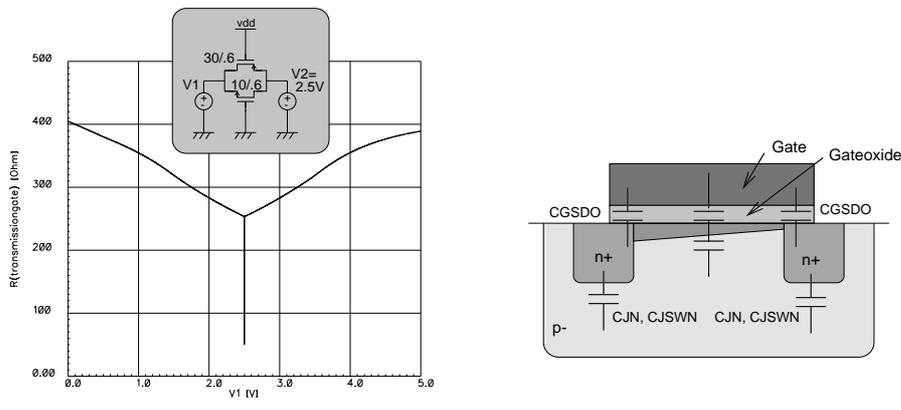


Fig. 5. Parasitic properties of the transmission gates **Left:** Simulated resistance (setup shown in the inset) **Right:** Two dimensional cut through a MOS transistor indicating its different capacitances.

geometry or routing possibility. Since increasing the switch size to obtain lower resistance values will increase its capacitance, a tradeoff has to be found. The on-resistance of a switch realized as a transmission gate with one node connected to 2.5 V is shown on the left side of Fig. 5, together with the setup for the according simulation (The glitch at 2.5 V is due to the limited computational accuracy.). For the chosen transistor geometries the on resistance is of the order of 300 to 400 Ω .

5.1 Resistance and capacitance of the switches used

The right part of Fig. 5 shows the cross section of an NMOS transistor and its capacitances. For an open switch capacitances include the gate source/drain overlap as well as the $n^+ p^-$ capacitance per area (CJN) and per width of the transistor (CJSWN). For the transmission gate simulated in the left part of Fig. 5 these add up to 54 fF. For the multiplexing of one 'global' D,G,S or B terminal (cf. Fig. 1) the 5 open switches result in a node capacitance of 0.27 pF, while for the 19 open switches for multiplexing the D,G,S and B terminals of the actually chosen transistor inside the cell to the global terminals a node capacitance of 1.03 pF is obtained.

5.2 Simulation of a simple Miller OTA

In order to get an impression of the influence of the parasitic effects described above, a simple CMOS Miller operational amplifier has been implemented using the programmable transistor cells. Fig. 6 shows the implementation in the cell array (referred to as Cell-Op) as well as the equivalent circuit that results from

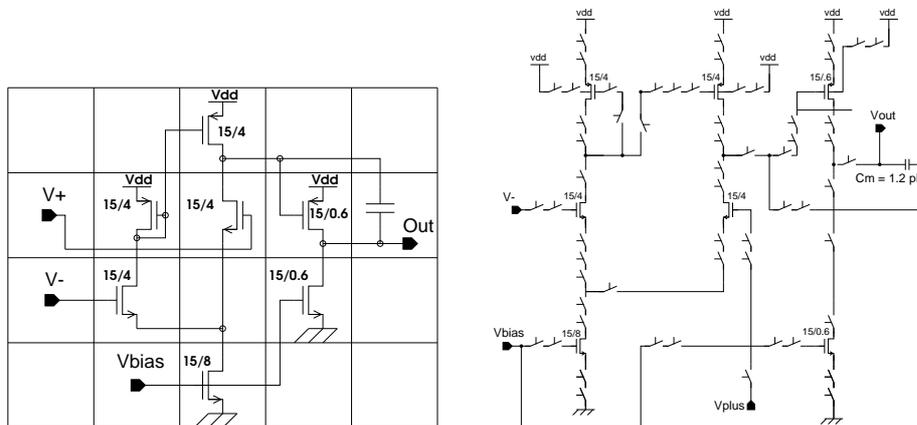


Fig. 6. Simulated operational amplifiers: **Left:** Implementation in the PTA chip. **Right:** Model of the circuit shown on the left including only the closed switches, here drawn as open ones for recognizability.

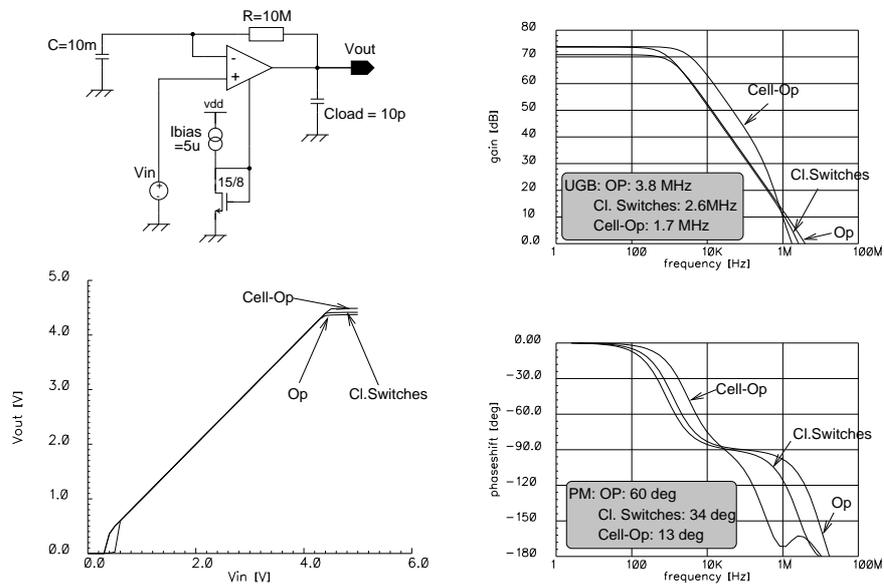


Fig. 7. Simulation results for the circuits shown in Fig. 6 and the original amplifier without switches: **Left: Top:** Test setup for the simulation. **Bottom:** DC-Responses **Right:** Results of the AC analysis: **Top:** Gain versus frequency (Bode plot). **Bottom:** Phase shift versus frequency.

disregarding all opened switches (referred to as Cl.Switches).

Both, ac and dc simulations², were carried out for three different implementations of the Miller OTA, namely the two circuits shown in Fig. 6 and the same Opamp without any switches (referred to as Op). The test setup for the simulation is shown in Fig. 7 together with the according dc and ac responses. For the dc analysis, where the Opamps are hooked up in a voltage follower configuration all three show the same desired behaviour for voltages between 0.5 V and 4.5 V (the region of interest).

The results of the ac analysis given in Fig. 7 show that the frequency response is degraded significantly by the introduction of both, the closed and the additional open switches (The according unity gain bandwidths (UGB) and phase margins (PM) are listed in the inset of the frequency plots.). This is due to the parasitic resistors and capacities introduced with the additional switches.

What can be learned from this? First of all, the maximum bandwidth of the PTA chip is of the order of MHz. Secondly, since the GA will have to deal with the parasitic capacitances, the evolved solutions will require a different frequency compensation than their ideal counterparts.

² All simulations were made using the circuit simulator spectreS.

6 Summary

A new research tool for evolution of analog VLSI circuits has been presented. The proposed system featuring a programmable transistor array, which will be designed in a $0.6\ \mu\text{m}$ CMOS process, will be especially suited to host hardware evolution with intrinsic fitness evaluation. Advantages are the fast fitness evaluation and therefore high throughput of individuals as well as the possibility to create a selection pressure towards robustness against variations of the environment and chip imperfections. Furthermore analyzability of the evolved circuits is enhanced by the implementation of local voltage, current and temperature sensors. First simulation results prove the feasibility of the chosen approach. The flexible design allows the test of different signal processing concepts with different chips easily derived from the chip presented here. Submission of the chip is planned for the first quarter of the year 2000.

7 Acknowledgement

This work is supported by the Ministerium für Wissenschaft, Forschung und Kunst, Baden-Württemberg, Stuttgart, Germany.

References

1. M. Loose, K. Meier, J. Schemmel: Self-calibrating logarithmic CMOS image sensor with single chip camera functionality, *IEEE Workshop on CCDs and Advanced Image Sensors, Karuizawa, 1999, R27*
2. Thompson, A.: An evolved circuit, intrinsic in silicon, entwined with physics. In Higuchi, T., & Iwata, M. (Eds.), *Proc. 1st Int. Conf. on Evolvable Systems (ICES'96)*, LNCS 1259, pp. 390-405. Springer-Verlag.
3. Thompson, A.: On the Automatic Design of Robust Electronics Through Artificial Evolution, In: Proc. 2nd Int. Conf. on Evolvable Systems: *From biology to hardware (ICES98)*, M. Sipper et al., Eds. , pp13-24, Springer-Verlag,1998.
4. R. Zebulum, M. Vellasco ,M. Pacheco,: Analog Circuits Evolution in Extrinsic and Intrinsic Modes, In: Proc. 2nd Int. Conf. on Evolvable Systems: *From biology to hardware (ICES98)*, M. Sipper et al., Eds., pp 154-165, Springer-Verlag,1998.
5. Layzell, P.: Reducing Hardware Evolution's Dependency on FPGAs, In *7th Int. Conf. on Microelectronics for Neural, Fuzzy and Bio-inspired Systems (MicroNeuro '99)*, IEEE Computer Society, CA. April 1999.
6. M. Murakawa, S. Yoshizawa, T. Adachi, S. Suzuki, K. Takasuka, M. Iwata, T. Higuchi: Analogue EHW Chip for Intermediate Frequency Filters, In: Proc. 2nd Int. Conf. on Evolvable Systems: *From biology to hardware (ICES98)*, M. Sipper et al., Eds., pp 134-143, Springer-Verlag,1998.
7. Stoica, A.: Toward Evolvable Hardware Chips: Experiments with a Programmable Transistor Array. In *7th Int. Conf. on Microelectronics for Neural, Fuzzy and Bio-inspired Systems (MicroNeuro '99)*, IEEE Computer Society, CA. April 1999.
8. K. R. Laker, W. M. C. Sansen: *Design of analog integrated circuits and systems*, pp 17-23, McGraw-Hill, Inc. 1994