3 WINQUIST, F.,   SPETZ, A.,   ARMGARTH, M.,   NYLANDER, C.,   and
   LUNDSTRÖM, I.: 'Modified palladium metal-oxide-semiconductor
   structures with increased ammonia gas sensitivity', *Appl. Phys.
   Lett.*, 1983, **43**, pp. 839–841

4 CRISTOFIDES, C., and MANDELIS, A.: 'Solid-state sensors for trace
   hydrogen gas detection', *J. Appl. Phys.*, 1990, **68**, pp. R1–R30

5 GAGGOITTI, G.,   GALDIKAS, A.,   KAĆIULIS, S.,   MATTOGNO, G.,   and
   ŠETKUS, A.: 'Surface chemistry of tin oxide based gas sensors', *J.
   Appl. Phys.*, 1994, **76**, pp. 4467–4471

6 FUKUDA, H.,   SEO, H.,   KASAMA, K.,   ENDOH, T.,   and   NOMURA, S.:
   'Highly sensitive MOSFET gas sensors with porous platinum gate
   electrode', *Jpn. J. Appl. Phys.*, 1998, **37**, pp. 1100–1102

CMOS/BiCMOS mixed design using tabu search

S.M. Sait and H. Youssef

The problem of optimising mixed CMOS/BiCMOS circuits is
solved using a tabu search. Only gates on critical sensitisable
paths are considered. This strategy leads to a circuit speed
improvement of > 20% with a < 3% increase in the overall circuit
capacitance. Comparison and experimental results are presented.

Introduction: BiCMOS technology combines the high-speed and
driving capabilities of Bipolar with the low area and low power
consumption advantages of CMOS. For a mixed CMOS/BiCMOS
design, there exist 2^m solutions for a circuit with m gates. Hence,
full/brute force design space exploration is infeasible [1]. In this
Letter, we discuss the problem of optimising mixed CMOS/BiC-
MOS circuits in terms of delay and power.

Problem definition: Given a circuit consisting of only CMOS cells,
some of those cells are replaced by their equivalent in BiCMOS,
such that the entire delay of the circuit is decreased while not
exceeding a power threshold constraint. Only gates belonging to
the longest sensitisable paths are considered.

Given a circuit with m nodes and K sensitisable paths, we first
extract all nodes i that are included in the sensitisable paths and
satisfy the inequality $F_i > C_i$, since a BiCMOS gate has a smaller
delay than the corresponding CMOS gate only if its fanout load-
ing capacitance F_i is greater than a certain threshold C_i [2]. Let A
$= \{1, 2, ..., n\}$ be the set of such nodes, with n its size, and let d_i^c
and d_i^b be the delay of gate i implemented in CMOS and BiC-
MOS, respectively. The circuit delay gain incurred by changing
gate i from CMOS to BiCMOS is $g_i = d_i^c - d_i^b$. Obviously, for all i
$\in A$, $g_i \geq 0$. Let C_i^c and C_i^b be the total capacitance of the circuit
when cell i is CMOS and BiCMOS, respectively, with all other
gates implemented in CMOS. Since the power of CMOS and BiC-
MOS gates is proportional to their capacitive load, we express the
changes in power in terms of changes in capacitance. In reality,
the power dissipation of a gate is a function of both its output
capacitance and the average number of transitions. Since we are
interested in relative rather than absolute accuracy, for the pur-
pose of quick estimation, as reported in [3], it is sufficient to esti-
mate the effect on the overall loading capacitance of the circuit.

The problem is to find a subset S of A which, when the nodes
of S are implemented in BiCMOS, will lead to a maximum reduc-
tion in T_{max}, while satisfying a threshold constraint on capacitance,
i.e. $\Sigma_{i \in S} \Delta C_i \leq C_T$; C_T is a given threshold which represents the max-
imum allowable total capacitance increase ($\Delta C_i = C_i^b - C_i^c$). Let
T_{max} and C_o be the initial operational clock period and the total
capacitance of the circuit when all gates are in CMOS. Assume
that the optimisation algorithm reduces the overall delay by ΔD
and increases the overall capacitance by ΔC. Then, the final oper-
ational clock and capacitance of the circuit are $T_{clock} = T_{max} - \Delta D$
and $C = C_o + \Delta C$. Furthermore $C \leq C_T$. The circuit optimisation
problem as defined above is NP-hard.

Proposed solution: To identify the gates of set S, we proceed by: (i)
generation of critical paths of the input circuit; (ii) elimination of
false paths; and (iii) application of the TS algorithm to select the
gates (subset S) among those covered by the sensitisable critical
paths and to satisfy the load constraint ($F_i > C_i$).

Critical path enumeration: Path enumeration is achieved via a
PERT-like trace of the circuit graph [4]. A path π is classified as
critical if its total delay T_π is very close to its latest required arrival
time ($LRAT_\pi$). If $T_\pi > LRAT_\pi$, path π is a long path. The path
delay consists of two components: the logic delay and the inter-
connect delay. The interconnect capacitance is a key element in
the total interconnect delay.

Let $\pi = \{v_1, v_2, ..., v_p\}$ be a path in the circuit graph where v_1
and v_p are the source and sink cells. The total delay on π is

$$T_\pi = \sum_{i=1}^{p-1}(CD_{v_i} + ID_{v_i}) \tag{1}$$

where CD_{v_i} is the switching delay of cell v_i and ID_{v_i} is the intercon-
nect delay of the net driven by cell v_i which is estimated as follows:

$$ID_{v_i} = LF_{v_i} \times C_{v_i} \tag{2}$$

where LF_{v_i} is the load factor of the output pin of the driving cell v_i,
expressed in time per unit of capacitance, and C_{v_i} is the total inter-
connect capacitance (area + fringe) of the net driven by cell v_i.
These capacitances are determined after placement assuming all
gates are CMOS. Let T_{max} be the expected delay of the longest
path in the circuit:

$$T_{max} = \max_{\pi \in \Pi}(T_\pi) \tag{3}$$

where Π is the set of all paths in the circuit graph G.

Definition: A path π is α-critical iff: $T_\pi \geq \alpha T_{max}$. The parameter
α, a positive number < 1, is supplied by the user and is interpreted
as a confidence level. Then, all paths that satisfy the above condi-
tion are reported as the α-critical paths.

Elimination of false paths: Following the generation of critical
paths, false paths are extracted. This is necessary to focus the opti-
misation of only sensitisable paths. We use the algorithm of [5]
with a more accurate delay model [4] and modifications related to
the handling of new events.

Circuit optimisation: We replace some of the CMOS gates on
these paths by BiCMOS gates to optimise the circuit for delay
without exceeding a capacitance threshold constraint. The output
is a mixed CMOS/BiCMOS circuit with optimum cost. Optimisa-
tion is achieved using a tabu search heuristic [6] as explained
below.

Tabu search (TS) implementation: Tabu search is an optimisation
heuristic which operates by incorporating flexible memory func-
tions to forbid transitions (moves) between solutions that reinstate
certain attributes of past solutions. Attributes that are not permit-
ted to be reinstated are called tabu, and are maintained in a tabu
list.

Initial solution: The initial solution is a set A of nodes of type
CMOS only, which are covered by the sensitisable α-critical paths.
As the search proceeds, neighbour solutions are generated by
changing a CMOS gate to a BiCMOS gate or vice versa. The gates
are selected randomly for change.

Aspiration criterion: The notion of aspiration criterion is used to
override the tabu status of moves whenever appropriate. The best
solution aspiration criterion (AS_1) is used to override the tabu
restriction if the move produces a new best solution [6].

Evaluation function: The objective is to maximise $\Sigma_{i \in S} g_i$, which
minimises T_{max}, subject to capacitance constraints. The evaluation
function $E(s) = \Sigma_{i \in S} g_i - X$, where X (a small positive number) is a
penalty factor that is included if capacitance constraints are vio-
lated [1].

Long-term memory/diversification: In many combinatorial opti-
misation problems, application of short-term memory alone may
not produce a good solution. Long-term memory functions can be
very important for obtaining improved results [6]. Here, we apply
frequency-based diversification. In this strategy, we keep track of
the number of times a move has been made. At the point when
diversification is to be applied, we penalise those moves that have
been most frequent, thereby taking the search to unvisited areas
[6]. Therefore, when the short-term TS algorithm hits a local
optima, the following actions are taken:
(i) All BiCMOS nodes, denoted by NUM_OF_BiCMOS, are
changed to CMOS.
(ii) Nodes with the lowest frequency of changes are searched and
replaced by BiCMOS. The search replacement process continues

until the objective load threshold is reached or until the number of replaced nodes equals NUM_of_BiCMOS. Some of the BiCMOS nodes that were changed to CMOS may be changed back again to BiCMOS. Hence, the search process is transferred to another region from where it might lead to a better solution.

Then, the short-term TS is restarted, and search continues until another local optimum is reached.

Results and discussion: For the nine ISCAS-85 circuits used, the percentage of false paths reported from among the critical paths ranged from 0 to 24%. In all but one case the maximum delay of the circuit did not change due to removal of false paths, but the number of gates on the sensitisable paths was reduced. For all experiments, we used a capacitance threshold that was 10% of the total capacitance, the required reduction in the delay being 30%.

Table 1: Results of TS with long-term memory and aspiration criterion AS_1

Circuit name	Maximum delay	Delay reduction	Percentage delay reduction	Total capacitance	Capacitance increase	Percentage capacitance increase	Number of BiCMOS gates
		ns	%	pF	pF	%	
c432	171.911	43.342	25.20	109.260	1.478	1.00	19
c499	65.344	13.357	20.00	101.717	1.384	1.00	21
c880	125.506	25.089	20.00	215.111	2.148	1.00	23
c1355	109.860	26.615	24.20	399.564	7.242	2.00	63
c3540	185.201	48.645	26.30	683.401	1.796	0.26	18
struct	121.894	26.130	21.40	750.081	2.534	0.34	26
highway	32.438	8.815	27.20	15.15	0.762	5.00	16
fract	76.575	22.497	29.38	43.405	2.691	6.20	23

Table 2: Comparison of our work with results of [2]

Circuit name	Results of [2]		Results of proposed solution	
	Percentage speed improvement	Percentage BiCMOS gates	Percentage speed improvement	Percentage BiCMOS gates
	%	%	%	%
c432	40.9	11.1	25.2	4.8
c499	19.1	14.5	20.0	7.4
c880	12.3	21.7	20.1	2.7
c1355	11.2	9.4	24.3	4.3

Table 1 shows the results of the TS with tabu list size of 5, and a candidate list size of 14. Diversification was applied after every 400 iterations. The number of BiCMOS gates required to speed up circuits using this approach is a very small percentage of the total number of gates. For medium to large sized circuits (> 200 gates) this number is < 5%. For most circuits tested, a percentage delay reduction between 20 and 30% has been achieved with a capacitance increase of < 3%. Table 2 shows a comparison of our work with that reported in [2], where all the nodes are considered, and all output nodes are replaced whether they are on critical paths or not. Also, no constraints were placed on power dissipation, and the number of BiCMOS gates was high. In three out of four cases the result is a much better speed performance than in [2], with a smaller number of BiCMOS gates. In all cases the percentage of BiCMOS gates used is smaller. Only for one circuit (c432), owing to the constraint on the total capacitance, the speed improvement was less, but so was the number of BiCMOS gates used.

7 May 1998

References

1 SAIT, S.M., and YOUSSEF, H.: 'Iterative computer algorithms and their applications in engineering'. IEEE Computer Society Press, 1998 (to appear)
2 BABA-ALI, A.R., and BELLAOUAR, A.: 'An optimization tool for mixed CMOS/BiCMOS standard cells circuits', *Arabian J. Sci. Eng.*, 1994, **19**, (4B), pp. 883–888
3 BRAND, D., and VISWESWARIAH, C.: 'Inaccuracies in power estimation during logic synthesis'. IEEE/ACM Int. Conf. Computer Aided Design, November 1996, pp. 388–395
4 YOUSSEF, H., SAIT, S.M., and AL-FARRAH, K.: 'Timing influenced force-directed floorplanning'. European Design Automation Conf. with Euro-VHDL, Euro-DAC'95, Brighton, September 1995, pp. 156–161
5 YEN, S., *et al.*: 'Efficient algorithms for extracting the K-most critical paths in timing analysis'. Proc. 26th Design Automation Conf., 1989, pp. 649–654
6 GLOVER, F., and LAGUNA, M.: 'Tabu search' *in* REEVES, C. (Ed.): 'Modern heuristic techniques for combinatorial problems' (McGraw-Hill Book Co., Europe, 1995)

Integrated optic polymeric components fabricated with microstructured strip-off covers

H. Kragl, R. Hohmann, C. Marheine, W. Pott, G. Pompe, A. Neyer, T. Diepold and E. Obermeier

Fabrication of integrated optic polymeric components by moulding waveguide and fibre alignment grooves to a substrate in a one-step process, and subsequently filling the waveguide grooves by means of a microstructured strip-off cover which is later removed, is presented. The latest measurement results on waveguide attenuation, coupling losses owing to passive fibre-waveguide coupling and excess loss values are given.

Introduction: To fabricate very low cost integrated optics components for singlemode and multimode applications, we use polymer moulding technology [1 – 3]. If the waveguide formation process and the fibre coupling process are carried out simultaneously, technical problems due to mutual perturbation between fibre alignment and waveguide formation occur. The pressure necessary for waveguide formation may disarrange the precisely aligned fibres and the presence of the fibres close to the ends of the waveguide grooves can hinder the waveguide formation process if the cover does not fit perfectly in length to the waveguide grooves area. Owing to these problems, the fabrication steps of waveguide formation and fibre alignment have been separated [4, 5]. Recently, we presented the first results on singlemode integrated optics components fabricated by the 'microstructured strip-off cover' approach, i.e. during the waveguide formation process a removable cover is used which protects fibre alignment grooves and all other structures with the exception of waveguide grooves by filling them with an inverse mechanical structure. Only the waveguide grooves are left open for the insertion of waveguide core material. Fig. 1 shows a schematic drawing of the moulding shim, polymer substrate and corresponding microstructured strip-off cover.

The waveguide formation process was improved and the new components show better excess loss figures of 4.5dB at 1.3µm wavelength for a 3cm long waveguide. This Letter describes the fabrication technique for microstructured strip-off covers and evaluates the fibre-chip coupling loss and waveguide attenuation loss, in order to determine the main obstacle to further improving excess loss figures.

Technical approach: The master structure fabrication was carried out with only one mask in two steps. As an etch mask of both the V-grooves and the waveguides, a silicon dioxide layer was fabricated using reactive ion etching (RIE). In the first step, the V-grooves were etched in aqueous KOH (the waveguides must be covered by a silicon nitride layer during this step). Next the waveguide grooves were etched using RIE (the V-grooves are masked by a silicon dioxide layer). The RIE-etching was optimised