Wire Topology Optimization for Low Power (TopCool and TopCoolViewer)*

Paul Zuber, Robert Hartl, Thomas Ilnseher and Walter Stechele

TU München, Institute for Integrated Systems, Arcisstr. 21, 80290 München, Germany
{Paul.Zuber, Walter.Stechele}@tum.de

Abstract

TopCool*, activity-driven local wire spacing for up to 9% less chip-wide wire power is demonstrated. Routed layouts are im-/exported from/to EDA tools. Area and timing are not adversely affected.*

1. Introduction

Despite the increasing importance of capacitive power consumption, today's design tools do not offer switching activity driven layout synthesis for low power digital CMOS. *TopCool* was developed to reduce the wire power consumption of detail-routed circuits by locally modifying the wire topology. Globally relaxing the routing pitch [1] or specifying the temporal routing sequence of the nets [3] shows limited success on many circuits due to routability problems. There exists related work focused to optimizing the wire power consumption of buses [2]. Its principle: individual, non-uniform wire spacing driven by activities, is leveraged and applied to entire circuits.

2. TopCool and TopCoolViewer

TopCool searches a given layout for all instances of groups of parallel wire segments that can be shifted laterally. The wires are then spaced through integer-convex programming. The result depends on activity and availability of surrounding whitespace. Unlike prior spacing techniques, the presented approach routes connections between the old and new wire endings. Finally, the best subset of all groups is selected. A detailed description can be found in [4].

Switching activities and layout data are accessed through appropriate interfaces to commercial physical design tools. Figure 1 illustrates the flow of the proposed method.

The OpenGL-based viewer *TopCoolViewer* was developed to investigate on-chip interconnects. It supports wire coloring according to activity, 3-D operations, starting the optimization and visualizing its results. The new model can be overlaid on the reference for comparison. Figure 2 was created with this tool.

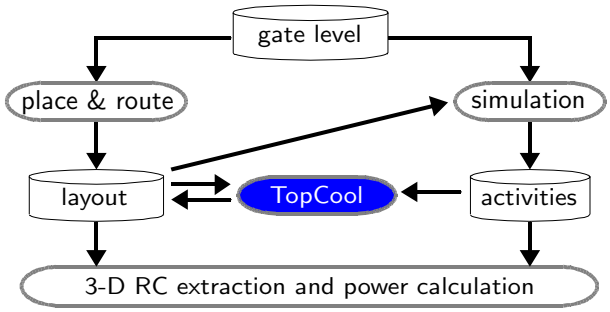

Figure 1. Illustration of how the commercial design flow is tapped.

3. Results

Benchmark circuits were synthesized from RT-level, placed and routed with a commercial tool, BlastChip by Magma*. Two types of timing and area constraints were used, moderate, and hard, cf. the table. To obtain switching activities, high-effort propagation including gate- and wire delays was used. Power was evaluated by BlastChip after full-chip 3-D capacitance extraction with QuickCap.

The table compares the interconnect switching power of the routing results of the commercial layout synthesis tool to the six techniques *spreading*, *Sanyo [1]*, *Intel [3]*, *TopCool*, and combinations of *TopCool* with [1] and [3].

Both wire spreading and *TopCool* are post routing operations. They can be applied to any design, no matter how congested the routing area is. This is not true for [3] and [1] which modify the routing parameters of an existing router. This can save remarkable amounts of power but fails on larger or more tightly constrained designs due to congestion problems, limiting their universal applicability.

*The authors would like to thank Magma Eindhoven for their support.

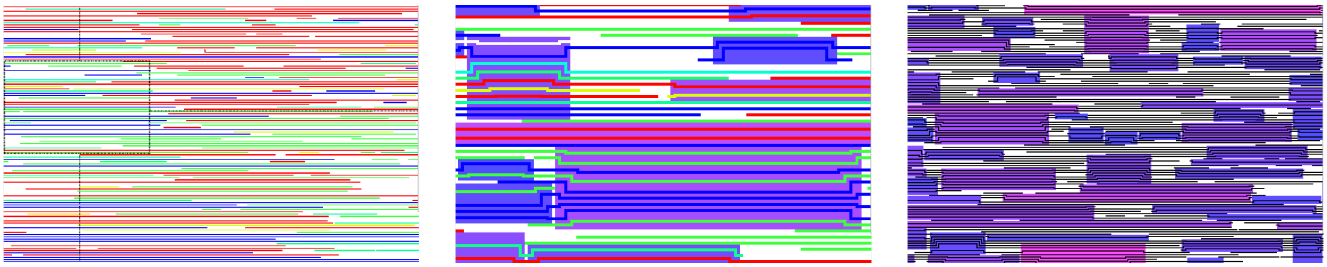

Figure 2. Left: Routing layer with wires of different activity. Middle: Activity-driven in-place wire spacing in groups (zoomed). Right: Committed optimization areas in entire layer.

wire power (mW) ([1] μW) **bold: best in row** ′ best in column * failed (over-congested)							
moderate: $t_{\mathrm{clk}} = 2t_{\mathrm{min}}$, rowutil. $= 70\%$							
model	ref	spread	[1]	[3]	*TopCool*	[1]+*TC*	[3]+*TC*
mlite	2.366	2.343	1.940	2.054	2.233	**1.836′**	2.032
eth	1.420	1.391	1.422	1.316	1.325	1.354	**1.283**
dct	12.82	12.74	12.43	11.91	12.16	12.03	**11.68**
3des	20.61	20.40	19.77	18.96	19.72	19.25	**18.59**
iuL2	12.90	12.43	*	*	**12.07**	*	*
1cL3	11.22	10.80	*	10.63	10.38	*	**10.23**
2cL3	20.59	19.71	*	19.00	18.96	*	**18.22**
4cL3	28.57	27.28	*	29.04	**26.33**	*	27.96
8cL3	55.49	52.78′	*	*	**50.91**	*	*
b14[1]	92.46	88.83	83.90	77.69′	83.76′	79.89	**74.84′**
b15	.2297	.2244	.2121	.2002	.2120	.2031	**.1937**
b17	.4856	.4752	.4399	.4370	.4501	**.4227**	.4275
b18	1.842	1.791	1.667′	1.581	1.698	1.590	**1.542**
b19	5.152	5.010	4.973	4.436	4.724	4.762	**4.382**
max (-%)		*4.88*	*18.0*	*16.0*	*9.41*	**22.4**	*19.1*
avg (-%)		2.83	4.59	8.13	7.16	7.08	**10.2**
hard: $t_{\mathrm{clk}} = 1.25t_{\mathrm{min}}$, rowutil. $\geq 80\%$							
mlite	4.396	4.396	*	*	**4.065**	*	*
eth	1.388	1.361	1.403	1.263	1.296	1.342	**1.232**
dct	39.20	38.67	37.13	36.07	37.03	35.87	**35.23**
3des	35.71	35.28	33.86	32.42	34.16	33.04	**31.80**
iuL2	16.80	16.24	*	*	**15.88**	*	*
1cL3	12.71	12.23	*	*	**11.82**	*	*
2cL3	22.06	21.15	*	*	**20.48**	*	*
4cL3	38.04	36.25′	*	*	**35.20**	*	*
8cL3	69.81	66.61	*	*	**64.68**	*	*
b14	.2394	.2312	.1991′	.2041	.2184′	**.1882′**	.1966
b15	.4365	.4248	.3953	.3477′	.4017	.3761	**.3366′**
b17	.7442	.7258	.7350	*	**.6918**	.7092	*
b18	3.939	3.826	3.657	3.430	3.656	3.513	**3.360**
b19	10.91	10.62	10.45	*	10.16	**10.10**	*
max (-%)		*4.71*	*16.8*	*20.3*	*8.77*	*21.4*	**22.9**
avg (-%)		2.80	3.45	5.30	**6.88**	5.53	6.27

TopCool saves about twice the amount of power as spreading. Additionally, *TopCool* achieves the highest average power savings of all methods with hard constraints. The quality of the optimization is largely independent of the circuit size and the constraints set. With moderate constraints, the largest circuit considered (eight Leon 3 cores with 1.2 million wires) shows a power reduction of 8.25%. *TopCool* cooperates well with [3] and [1]. All of the highest power savings are achieved with either *TopCool* or one of the combinations. Timing is not adversely affected.

Our method operates about five times faster than commercial spreading on average, and up to two orders of magnitude faster than [3] and [1]. It adheres to quasilinear runtime complexity in the number of objects.

4. Limitations and Improvements

TopCool improves the yield limit of the metal layers but does not achieve the same values as spreading. A hybrid optimization for yield and low power is feasible by not rejecting wire groups with low power reduction as currently done. Improvements in the power saving results are possible by relocating vias or ordering wires prior to spacing. More research is required, in this regard.

5. Conclusions

Activity-driven wire spacing applied in-place in already detail-routed layouts has been demonstrated. *TopCool* reads layouts from existing physical design tools and switching activities from simulation. Wires are then relocated depending on the activities and the amount of surrounding whitespace available. Optimization is fast as multiple wires are treated simultaneously with a tailored integer-convex program. After optimization, a new layout is exported and can be read back by EDA tools. Reductions of up to 9.41% of a modern benchmark design were achieved after 3-D RC-extraction without affecting area or timing. Commercial spreading which does not consider activities achieves only half of the savings and is four to eight times slower. *TopCoolViewer* allows the user to investigate the resulting interconnect power and the optimization steps.

References

[1] A. Sakai et al. Reduction of coupling effects by optimizing the 3-d configuration of the routing grid. In *IEEE Trans. VLSI Systems*, Oct. 2003.

[2] L. Macchiarulo et al. Wire placement for crosstalk energy minimization in address buses. In *DATE*, 2002.

[3] N. Magen et al. Interconnect-power dissipation in a microprocessor. In *SLIP*, 2004.

[4] P. Zuber. *Wire topology optimisations for low power CMOS*. PhD thesis, Technische Universität München, to appear in 2007.