Optimizing ECO Efficiency and Precision: an alternative approach for Implementing Digital-Intensive RTL Module Restructuring Through Metal Changes in IC Design

Carmen Galotta, Dario D'Afiero, Marco Leo, Nicola Imperato, Gabriele Faraone, Eugenio Serianni, Giovanna Antonella Franchino
STMicroelectronics (Ltd)
Analog, Power & Discrete, MEMS and Sensor Group R&D, via C. Olivetti 2, 20864 Agrate Brianza, Italy.
Phone: +39 366 57 615 10, Email: carmen.galotta@st.com

***Abstract*— Engineering Change Orders (ECOs) enable late-stage chip modifications without a full redesign, saving time and cost. Metal ECOs are preferred over full mask ECO for their flexibility but are limited by the size of changes and availability of spare components. Small updates are straightforward, while larger ones risk timing issues and design rule violations, making metal ECO not feasible. Current tools often lack awareness of the chip's clock system and struggle to reuse existing logic efficiently. This study presents a new method applied to a real chip redesign, successfully achieved using a metal ECO to handle complex design changes while improving timing and resource use.**

Keywords—ECO, metal mask revision, logic reuse

I. INTRODUCTION

In integrated circuit (IC) design, an Engineering Change Order (ECO) is an important process that allows designers to implement device changes at a later stage of the development cycle without requiring a complete redesign of the chip. This method is crucial within the IC design flow as it provides a cost-effective and time-efficient solution to address design errors, performance issues, or specification changes that arise late in the design cycle. There are two types of ECOs: the first one is a full mask revision, which requires changes to the entire masks set used in fabrication process and it's necessary when there is a considerable change in the design specification or a partial redesign of a major block of the circuit that requires new logic gates. The second type of ECO is a metal mask only, suitable for smaller fixes, in which only the metal layer masks of the chip fabrication process are modified, therefore allowing reuse of most design components and gates already present in the original design.

Metal only revision takes significantly less time and incur substantially lower cost when compared to full mask revision; while full mask changes can cost up to millions of dollars, metal-only updates typically amount to only tens of thousands [1]. This greater flexibility and reduced impact on both the design and manufacturing schedule, make metal-only ECOs the preferred solution, whenever feasible

In this context, the number of changes is crucial in determining the feasibility of the metal fix. For minor changes, ECO fixes are typically introduced manually on a post-layout netlist, and the number of logic gates modified from the original design is low. This implies minimal adjustments on the post-route database to ensure the new functionality does not disrupt performance [2], and that Static Timing Analysis (STA) and Design Rule Checking (DRC) remain clean. If, however, the design changes are more extensive and involve a significant update to the RTL code, the scope of modifications to the original logic gates increases, potentially complicating the metal fix implementation. In such cases, a key consideration is whether the available logic gates and spare cells are sufficient to accommodate the new functionality, ensuring that the design can be effectively modified without necessitating additional resources or significant alterations to the existing layout.

This paper focuses on the methodology employed to implement a metal ECO on a digital module within a top-level circuit. The process involved a significant redesign aimed at optimizing the module's architecture to enhance performance, reduce complexity, and improve overall efficiency, ultimately resulting in a complete rewrite of the module's RTL code. Initially, the standard ECO flow was applied and thoroughly analyzed, revealing significant challenges and limitations, that brought to the conclusion that a metal-ECO was not feasible. In response, we employed an alternative approach designed to better support the ECO tool's capability, enabling successful implementation of the metal ECO and achieving the desired design objectives.

II. CONTEXT

The original architecture achieves an effective trade-off between area efficiency and computational performance by employing a single counter composed of enlarged half-adder cells. Control and operation sequencing are managed by a Finite State Machine (FSM), which orchestrates the computational process as shown in Figure 1. This approach enables streamlined hardware utilization while maintaining the desired functional throughput, making it suitable for resource-constrained digital systems.

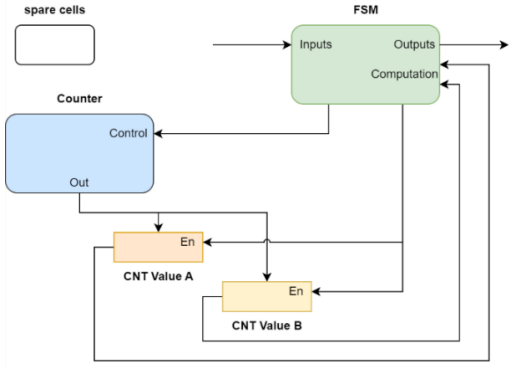

Fig. 1. Original RTL module architecture

During on-silicon testing, it was observed that the input signal dynamics exceeded initial expectations, resulting in computational latency introduced by the Finite State Machine (FSM) that proved unacceptable in certain operational scenarios. This latency constraint highlighted a critical limitation in the original architecture, where the FSM-based control introduced delays incompatible with faster input transitions. To address this issue, a redesign was required involving the elimination of the FSM and the introduction of a second counter, as shown in Figure 2. This modification aimed to overcome the hardware bottlenecks and meet the stringent timing requirements imposed by the accelerated signal input behavior.

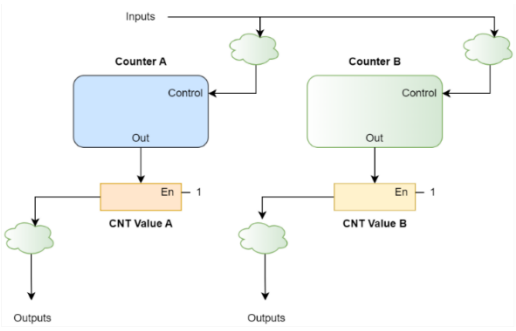

Fig. 2. New RTL module architecture with required modifications

The deployment of two parallel timer instances effectively enhanced the circuit's performance, successfully overcoming previous hardware limitations and meeting the required timing constraints.

III. STANDARD ECO FLOW

To facilitate these changes, during the first implementation of the device, additional unused spare cells have been added to the layout to accommodate any subsequent design changes that may occur after the tape-out process. For some technologies, it is possible to use mask-programmable cells [3], or gate array filler cells, which are pre-designed standard cells with the capability to alter their functionality to any basic function (e.g., AND, OR) by changing only metal layers, instead of base layers. However, for other technologies where these types of cells are unavailable, normal logic gates are used as spare cells and various techniques are used to select them [4].

In this device, the number and type of spare cells were selected based on a statistical analysis aimed to identify the most frequently used cells in the design, prioritizing those most effective in generating new logic functions [4]. These spare cells were inserted early in the Place and Route flow (Fig 3.a) and distributed evenly across the floorplan [5], targeting approximately 3-5% of the total design area. Their inputs are tied to a constant logic level, either zero or one, and their outputs are left floating.

For all subsequent changes, that are found after the tape-out phase, either during pre-silicon or post-silicon validation, a specialized EDA tool called Conformal ECO from Cadence is employed. This tool helps to facilitate the metal-only ECO by automating much of the process, significantly reducing the time and manual effort required.

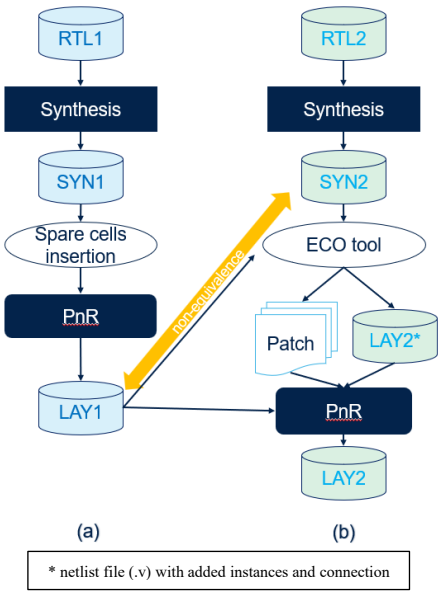

Fig. 3. Pre and post-layout ECO flow for metal fixes.

This tool, as illustrated in Figure 3.b, takes as inputs:

- the post-layout netlist (LAY1), representing the current state of the design, and
- the new RTL code (RTL2), containing the desired modifications and/or enhancements.

compares the original and revised designs, analyzes all non-equivalences and differences, and generates the necessary gate-level updates to the post-layout database. By automating the synthesis and the identification of the changes, the tool produces a set of commands, often in the form of TCL scripts, which guide the Place and Route tools to implement the modifications directly on the post-route database.

Despite these capabilities, the process remains complex and often requires multiple iterations to produce a clean database that satisfies both performance targets [2] and Design Rule Checking (DRC) constraints [6]. Specifically, the challenges encountered during this process are described below.

A. Unused cone logic

A key problem arises from the tool's inability to recognize all unused logic cones that become available after the changes are made. This results in underutilization of existing cells within the design, forcing the tool to rely predominantly on spare cells. Since spare cells may be limited or poorly located, this inefficiency complicates the implementation and can increase design complexity. Furthermore, most of the unused logic remains connected and it can still consume power, thus becoming completely unproductive and contributing to unnecessary leakage and dynamic power dissipation within the design.

B. Spare cells availability

Linked to the previous point, is the availability and distribution of spare cells within the design, especially around the restructured module. It has been found that the number of

spare cells was insufficient, or placed too far apart from the restructured module, which restricted the tool's ability to implement changes effectively. This limitation forced the ECO process to rely on suboptimal fixes, both due to the tool's lack of dynamic spare-cell aware technology remapping [7] [8], but also due to the limited availability of complex standard cells needed to implement advanced digital logic functions. Consequently, a large number of basic cells were required, leading to increased cell count, routing congestion, and a detrimental effect on overall design performance.

C. *Clock tree structure*

The clock tree structure was another critical area of concern, as the changes made to the clock by the ECO tool, did not accurately reflect the real clock tree topology. For example, new flip-flops were added to trunk nets instead of leaf nets, disrupting the intended clock distribution hierarchy. This misalignment resulted in unbalanced clock paths that cause timing failures such as setup and hold violations, which then require manual adjustments to the clock tree structure to restore proper synchronization and performance. These manual corrections are error-prone and often require multiple iterations to identify the correct clock connections, a task typically managed by complex algorithms during the place-and-route stage.

Additionally, the newly introduced flip-flops lacked clock gating logic on their clock inputs, resulting in continuous clock activity. This absence of clock gating leads to increased dynamic power consumption, as the flip-flops are toggled unnecessarily even when their operation is not required.

D. *DFT connection*

Design-for-Test (DFT) connectivity also poses additional challenges. Newly inserted flip-flops are not automatically connected to the existing scan chain, necessitating manual intervention. This manual connection is time-consuming and prone to errors, and failure to properly integrate new flip-flops can reduce overall DFT coverage, thereby affecting the testability and reliability of the final product. Furthermore, based on the location of the spare flops and their surrounding scan chains, the integration of the additional flops can extend certain chains disproportionately and can lead to imbalanced scan chain length. This imbalance can adversely affect test efficiency by increasing overall scan shift time and potentially delaying time to market.

E. *DRV*

Design Rule Violations (DRV) were a consequence of badly optimized metal-only ECOs. Changes and suboptimal choice of spare cells introduced a huge number of DRV violations such as exceeding the net maximum capacitance or maximum transition limits, which degrade signal integrity and timing performance. Correcting these violations is a manual process, time consuming and often requires the insertion of numerous buffers [9], which can further complicate timing closure and increase power consumption. Furthermore, the number of buffers is limited for each spare module; therefore, once these are exhausted, more distant buffer cells must be used, which often fail to resolve the issue effectively.

F. *DRC*

Lastly, a substantial number of Design Rule Check (DRC) violations were present, including metal shorts, spacing violations, and other layout infractions. Such violations critically impact manufacturability and yield if not rigorously identified and mitigated throughout the ECO iteration process. Despite executing additional routing optimization and congestion alleviation steps, these measures proved insufficient to fully eliminate the violations, indicating the necessity to require more comprehensive intervention.

Collectively, these limitations (insufficient spare cells, the presence of metal shorts and timing violations, and the extensive RTL-level modifications) led to the conclusion that a metal-only ECO was not feasible with the current implementation, ultimately requiring consideration of a full mask set ECO to accommodate the scope of the changes.

IV. PROPOSED ECO FLOW

To overcome these challenges, an alternative methodology has been adopted: instead of using the original post-layout netlist from the first tape-out of the project, a revised post-layout netlist is supplied to the ECO tool, as illustrated in Figure 4.

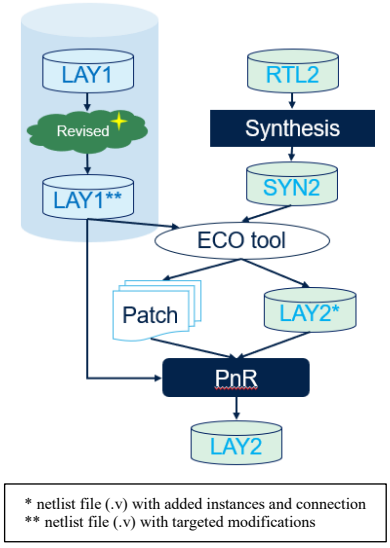

Fig. 4. Revised Flow

This revised version (referred to as LAY1** in Figure 4) incorporates a series of targeted modifications aimed at enhancing the tool's efficiency and effectiveness in resolving all the issues encountered and described in the previous chapter.

Specifically, all flip-flops from the original module are preserved to maintain the integrity of the existing sequential elements. The Design-for-Test (DFT) logic associated to it, is also retained, ensuring that the original scan chain architecture and length remains intact. This is achieved by grounding the functional inputs of the flip-flops while preserving their scan chain inputs (Test-Input and Test-Enable) and outputs (Test-Output), thereby maintaining proper testability without impacting the original coverage. The aim of this modification also has the scope to ensure that all flip flops are re-used by

the tool, eliminating the need for additional flip-flops as spare cells.

To further support this objective, the original clock tree structure is entirely preserved to maintain consistent timing and synchronization among the flip-flops, thereby preventing the ECO tool from modifying the original clock tree. Initial trials using the standard flow demonstrated that allowing the ECO tool to alter the clock tree resulted in suboptimal outcomes, such as newly added flip-flops being connected to clock tree trunks instead of leaves, which led to a significant number of setup and hold violations. By retaining the original clock tree structure, this problem is bypassed and only minimal adjustments are required to resolve any remaining hold violations. The only modification applied to the clock tree cells involves tying the enable signal of the clock gates to a one logic constant value, while keeping the test enable signal connected. This allows the tool to correctly repurpose all the clock gates instances and correctly connect the new logic that drives their enable signals, while preserving their associated DFT logic.

Lastly, all combinational logic between flip-flops withing the original module is left available for reuse and reconfiguration as needed. This allows the tool to repurpose all logic cells from the original module, maximizing cell reuse and avoiding the selection of spare cells that create more basic logic function and are physically distant from the original module in the floorplan. As a result, the original module's cells effectively become "spare" cells, enabling the tool to remap the new logic into these existing resources.

At the RTL level, the redesigned architecture (RTL2 in Fig. 4) was implemented by restructuring the original code to eliminate the FSM and introduce a second counter operating in parallel. This approach leveraged the reuse of the majority of the FSM combinational logic cells to form the new control logic and counters, supplemented by a minimal number of additional spare cells. Careful attention was given to preserving register naming conventions, ensuring that count value registers and control logic registers retained their original identifiers to maintain compatibility with the Logic-Equivalence-Checking (LEC) verification tool.

This alternative methodology has demonstrated significant efficacy in enabling the ECO tool to implement modifications by utilizing the pre-existing design infrastructure, thereby minimizing the use of spare cells. The preservation of the sequential elements and their DFT connectivity has facilitated optimal flip-flop reutilization, eliminating the need to allocate additional flip flops from spare cells resources. Maintaining the complete test infrastructure has ensured that DFT coverage remains unchanged, while the retention of the original clock tree architecture has minimized any adverse impact on timing closure and preserved clock domain integrity.

Furthermore, since all repurposed logic elements originate from the initial module architecture and are all spatially localized in the same area, the resulting netlist from ECO tool exhibits substantially reduced interconnect area and improved routability within the affected region. This significantly limited the number of DRV and DRC violations in the modified design, only requiring minimal adjustment and few eco-route loops to resolve them, ultimately making the metal-ECO a viable route.

Table 1 exhibits a comparative summary of the standard flow and the proposed methodology adopted in this work, demonstrating that the proposed flow has achieved better performances in all aspects.

TABLE I.

Critical aspects	**Flow Comparison**	
	Standard flow	***Proposed flow***
Logic cell reuse	Minimal (~10%)	Maximized (~80%)
Spare cells usage	140 spare cells 21 spare flip flops	20 spare cells 0 spare flip flops
Clock tree	Altered and need manual intervention	Original structure preserved and timing synchronization maintained
DFT (Design-for-Test)	Manual addition of new flip flop to existing scan chain	Original Scan-Chain structure is preserved and DFT coverage maintained
DRV (Design-Rule-Viol)	Long nets → huge number (>100) of Max capacitance and max transitions violations	Shorter nets → low number (~30) of max cap/trans violation that can be fixed manually
DRC (Design Rule Check)	Metal shorts and metal spacing violations ~ 250 violations	No metal shorts
Metal ECO outcome	**Not feasible**	**Feasible**

V. CONCLUSION

In this study, we developed a comprehensive methodology for the full architectural redesign of a digital RTL module embedded within a top-level SoC to resolve hardware bottlenecks identified through on-silicon testing. These bottlenecks critically limited the system's ability to meet stringent timing constraints and target performance metrics. The extensive restructuring of the original RTL architecture raised concerns regarding the feasibility of implementing the required changes through either a metal or full mask ECO at the digital layout level.

Initial attempts employing a conventional post-layout ECO flow revealed multiple critical issues, including timing violations, design rule check (DRC) errors, and design rule violations (DRV), which collectively rendered metal ECO implementation infeasible

In contrast, our proposed methodology maximizes the reuse of existing standard cell instances from the original RTL module within the post-layout netlist. Key elements such as flip-flop instances, the original clock distribution network, and scan chain connectivity were explicitly preserved to maintain critical timing paths and ensure design-for-test (DFT) integrity. This approach also significantly reduced the dependency on additional spare cells and flip-flops, thereby minimizing layout perturbations.

The methodology was validated on a production-grade design, successfully enabling manufacturing tape-out with a substantial reduction in manufacturing costs and without compromising timing closure or test coverage.

REFERENCES

[1] Lee Sungchul , Shin Hyunchul, "Efficient and Low-Cost Metal Revision Techniques for Post Silicon Repair", JSTS:Journal of Semiconductor Technology and Science, Volume 14,Issue 3, Pages 322-330, 2014

[2] S. K. Taumar and S. Agrawal, "An Efficient Metal ECO Methodology for Addressing Timing Violations with 10X Turnaround Time," *2022 3rd International Conference on Smart Electronics and Communication (ICOSEC)*, Trichy, India, 2022, pp. 53-58

[3] Tilak Wadhwa, Gaurav Gupta (NXP Semiconductor, India), "Metal ECO implementation using Mask Programmable cells", *Design Reuse, https://www.design-reuse.com*

[4] Chang, Kh., Markov, I.L., Bertacco, V. (2009). Methodologies for Spare-Cell Insertion. In: Functional Design Errors in Digital Circuits. Lecture Notes in Electrical Engineering, vol 32. Springer, Dordrecht. https://doi.org/10.1007/978-1-4020-9365-4_12

[5] Z.-W. Jiang, M.-K. Hsu, Y.-W. Chang, and K.-Y. Chao. Spare-cellaware multilevel analytical placement. In Proc. DAC, 2009, pp. 430-435

[6] C. -Y. Tan and I. H. -R. Jiang, "Recent research development in metal-only ECO," 2011 IEEE 54th International Midwest Symposium on Circuits and Systems (MWSCAS), Seoul, Korea (South), 2011, pp. 1-4

[7] K. -H. Ho, J. -H. R. Jiang and Y. -W. Chang, "TRECO: Dynamic technology remapping for timing Engineering Change Orders," *2010 15th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Taipei, Taiwan, 2010, pp. 331-336

[8] K. -H. Ho, Y. -P. Chen, J. -W. Fang and Y. -W. Chang, "ECO Timing Optimization Using Spare Cells and Technology Remapping," in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 29, no. 5, pp. 697-710, May 2010

[9] C.-P. Lu, M. C.-T. Chao, C.-H. Lo, and C.-W. Chang. A metal-only-ECO solver for input slew and output loading violations. In Proc. ISPD, pp. 191-198,2009.