Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Joint Optimization of Sequential Elements and Metal Layer Usage for Advanced Semiconductor Physical Design Using Artificial Intelligence and Machine Learning

Vidhu Shekhar Bajpai
Advanced Micro Devices, USA
Corresponding Author: Vidhu Shekhar Bajpai, **E-mail**: vidhusbajpai@gmail.com

| ABSTRACT

Technology nodes below 2nm pose a whole new set of challenges, primarily concerning power, performance, area (PPA) optimization, and routing complexity. Sequential elements such as flip-flops and latches are the leading contributors to dynamic power consumption and are among the major timing bottlenecks, while dense metal stacks cause resistance and congestion issues that limit manufacturability. This work presents a unified machine learning framework that enables optimum sequential elements use and metal layer assignment with a design automation process. The framework uses graph neural networks (GNN) to learn the topological patterns of netlist structures for property extraction and applies reinforcement learning agents to make instantaneous decisions for real-time metal allocation over voltage/clocks, routing congestion maps, IR drop limitations, and timing criticality weights. For high-dimensional embeddings of raw feature data, simulation traces, toggle patterns ,and physical layout characteristics are converted to a graph node structure through feature engineering techniques and utilized for semantic clustering. The joint cost function mitigates timing degradation, IR drop penalties, register count, and routing congestion. Adaptive weights allow for distinct tuning for joins and improve usability. Showing implementations in TSMC N2 technology means that there were significant improvements in the following: flip-flop reductions, setup violations severity, improved routing congestion, overall design robustness, and complete management of metal layer assignments. The framework works with existing EDA tool chains through further TCL scripting or Python APIs without altering functional correctness or other manufacturing alteration elements.

| KEYWORDS

Sequential element optimization, metal layer allocation, graph neural networks, reinforcement learning, semiconductor physical design.

| ARTICLE INFORMATION

ACCEPTED: 01 August 2025 **PUBLISHED:** 29 August 2025 **DOI:** 10.32996/jcsts.2025.7.9.18

1. Introduction and Problem Statement

1.1 Technology Scaling Challenges at 2nm Node and Below

The semiconductor industry's shift to sub-2nm technology nodes is a significant transition that adds an additional high level of complexity in the design optimization process. Sophisticated process technologies, including gate-all-around transistor architectures, increase the potential for variability and parasitic coupling effects while increasing the accuracy needed in design closure processes [1]. These scaling difficulties, along with traditional design optimization processes, are conducive to a more complicated optimization platform that aisneeded to address the more complicated interdependencies between power, performance, area, and reliability constraints.

1.2 Sequential Element Redundancy and Power Overhead in Advanced Designs

Sequential logic elements proliferate in modern digital designs due to pipelining requirements, control logic implementations, and interface protocol specifications, often resulting in unnecessary redundancy. This proliferation creates excessive clock tree loading, elevated dynamic power consumption, and increased timing complexity that becomes particularly problematic in

advanced nodes where power density constraints are stringent and timing margins are reduced. Traditional design flows lack the intelligence to identify and eliminate non-essential sequential elements while maintaining functional correctness and timing performance.

1.3 Metal Layer Congestion and Electromigration Issues

Metal layer allocation in advanced technology nodes presents multifaceted challenges encompassing routability, thermal management, and electromigration reliability. Dense local routing on lower metal layers exacerbates congestion bottlenecks, while improper power network strapping across metal layers increases IR drop risks and thermal hotspots. The selection of optimal metal layers requires comprehensive awareness of net criticality, layout density patterns, and electrical constraints, with electromigration concerns becoming particularly acute in narrow metal lines carrying high current densities.

1.4 Limitations of Traditional Heuristic-Based Design Flows

Conventional electronic design automation tools rely primarily on rule-based heuristics and predetermined optimization sequences that struggle to adapt to dynamic interactions present in advanced technology nodes. These traditional methodologies operate with fixed priority schemes and limited visibility into global design impacts, resulting in suboptimal solutions when faced with complex tradeoff scenarios. The heuristic-based approaches lack the capability to learn from design patterns and typically address sequential optimization and metal layer planning as independent problems, missing opportunities for synergistic improvements.

1.5 Opportunity for AI/ML-Driven Optimization Approaches

Recent advancements in artificial intelligence and machine learning present unique opportunities to disrupt semiconductor design optimization through intelligent automation [2]. Graph neural networks have the ability to learn complex topological and other structural patterns in netlist representations, and reinforcement learning can be utilized for adaptive decision-making based on evolving design criteria. Furthermore, machine learning models can process the vast amounts of design data to spot underlying patterns and dependencies, which traditional heuristical approaches often fail to discover, allowing optimization algorithms to learn continuously, and without limits, consistently adapting to the complexities of designs and technologies.

2. Prior Work and Technical Foundation

2.1 Sequential Circuit Optimization: Formal Verification, Retiming, and Clock Control

Sequential circuit optimization relies on three primary techniques that have shaped modern digital design practices. Formal verification methods employ satisfiability checking algorithms to guarantee functional equivalence during circuit modifications, ensuring design integrity throughout optimization processes. Retiming algorithms systematically move storage elements across logic boundaries to minimize critical path delays and reduce overall register requirements [3]. Clock control mechanisms insert conditional gating logic to prevent unnecessary switching in idle circuit portions, directly reducing dynamic power consumption. These foundational approaches encounter scalability limitations when applied to contemporary designs containing extensive sequential logic networks.

Technique	Primary Objective	Computational Complexity	Design Stage	Functional Verification Required
Formal Methods	Equivalence Verification	Exponential	Post-Synthesis	Yes
Retiming	Critical Path Optimization	Polynomial	Synthesis/Physical	Yes
Clock Gating	Dynamic Power Reduction	Linear	RTL/Synthesis	Partial
Sequential Pruning	Area/Power Optimization	NP-Complete	Physical Design	Yes

Table 1: Sequential Logic Optimization Techniques Comparison [3]

2.2 Physical Routing Strategies in Commercial Design Software

Current industry-standard design automation platforms implement metal layer assignment through algorithmic rules that prioritize completion rates and manufacturing compliance. Power grid networks receive allocation to thick upper metal layers, while signal routing is distributed across the remaining layers based on local congestion measurements. Vertical connection minimization guides inter-layer routing decisions to control parasitic effects and fabrication complexity. Design rule constraints specify geometric limitations, including minimum dimensions and spacing requirements that restrict routing flexibility. These deterministic approaches cannot accommodate dynamic optimization requirements or balance competing design objectives effectively.

2.3 Computational Intelligence in Circuit Design Automation

Artificial intelligence integration within circuit design tools addresses complexity challenges that exceed traditional optimization capabilities [4]. Supervised learning models predict circuit delays, estimate power consumption, and forecast routing congestion with enhanced accuracy compared to analytical methods. Classification algorithms support design exploration by identifying optimal parameter combinations based on performance criteria. Deep neural networks model complexnonlinearr relationships between design variables and circuit characteristics. Pattern recognition techniques group similar design configurations to inform optimization strategies. These computational methods demonstrate significant improvements but generally target isolated design problems rather than comprehensive system optimization.

Application Domain	ML Technique	Input Features	Output Prediction	Accuracy Level
Congestion Prediction	Decision Trees	Routing Density, Net Distribution	Congestion Hotspots	High
Power Estimation	Neural Networks	Activity Factors, Gate Counts	Dynamic/Static Power	Very High
Delay Modeling	Regression Models	Process Parameters, Layout Geometry	Propagation Delays	High
Design Space Exploration	Classification	PPA Metrics, Design Constraints	Feasible Solutions	Medium
Placement Optimization	Reinforcement Learning	Netlist Topology, Timing Constraints	Component Positions	High

Table 2: Machine Learning Applications in EDA Domain Classification [4]

2.4 Network-Based Learning and Decision Algorithms for Circuit Design

Graph-structured neural networks naturally represent circuit topologies by encoding components as vertices and connections as directed edges. Neighborhood aggregation operations enable information propagation across circuit structures while preserving topological relationships. Feature learning mechanisms extract relevant patterns from local circuit configurations to guide global optimization decisions. Decision-making algorithms model circuit optimization as sequential choice problems where automated agents develop optimal strategies through environment interaction. Multi-objective reward structures encode competing design goals, including performance, area efficiency, power consumption, and manufacturing feasibility. These advanced computational frameworks outperform traditional heuristic methods in handling complex combinatorial optimization challenges.

2.5 Limitations in Unified Design Optimization

Contemporary optimization strategies isolate sequential logic modification from physical implementation planning, creating artificial boundaries between interdependent design decisions. Current methodologies lack comprehensive frameworks that simultaneously consider register placement effects on routing congestion and metal layer utilization impacts on timing performance. Sequential optimization operates without routing awareness, potentially generating configurations that create physical implementation difficulties. Physical planning algorithms ignore sequential element distribution consequences for clock network synthesis and power delivery design. This compartmentalized approach prevents comprehensive optimization that could achieve superior results through coordinated decision-making across design domains.

3. Technical Approach and Implementation Strategy

3.1 Multi-Source Data Collection: Simulation Monitoring, Voltage Analysis, and Route Density Assessment

Comprehensive data gathering operations extract information from diverse design verification sources to construct training repositories for algorithmic development. Dynamic behavior monitoring captures switching patterns across storage elements during representative benchmark execution, generating activity profiles and temporal relationships [5]. Voltage distribution assessment produces spatial maps revealing power delivery stress concentrations and thermal gradient locations throughout chip regions. Route density evaluation generates detailed measurements of interconnect utilization across multiple abstraction levels, pinpointing constraint areas that limit physical realization. Supplementary information streams encompass parasitic modeling outputs, delay verification summaries, and fabrication rule compliance reports that establish complete design characterization datasets.

3.2 Attribute Construction for Storage Components and Interconnect Properties

Data transformation processes convert raw design information into organized formats compatible with learning algorithm requirements through specialized encoding techniques. Storage element characterization includes activity frequency measurements, connectivity distribution patterns, delay sensitivity rankings, and geometric clustering indices calculated from coordinate positions. Interconnect attribute generation produces density distribution charts, layer consumption profiles, connection point statistics, and length measurements spanning routing hierarchies. Integration methods combine storage and interconnect properties through statistical correlation examination, revealing dependencies between register positioning choices and routing resource demands. Sequential encoding captures design modification patterns across optimization cycles to provide learning systems with convergence behavior information.

3.3 Hybrid Learning Architecture: Gradient Boosting with Attention Mechanisms for Storage Optimization

Storage optimization employs combined architectural strategies merging ensemble decision methods with attention-based processing frameworks [6]. Gradient boosting modules process structured attribute tables, including delay margin data, consumption estimates, and connectivity measurements through iterative refinement techniques. Attention processing units handle sequential patterns in activity information and design evolution using self-referencing mechanisms that identify extended relationships between circuit components. Integration layers merge outputs from architectural elements through adaptive weighting systems that modify based on design circumstances and optimization targets. Parallel attention heads enable simultaneous processing of different attribute categories while preserving computational efficiency for extensive circuit representations.

Component	Architecture Type	Input Processing	Feature Learning	Output Generation
XGBoost Module	Gradient Boosting Trees	Structured Tabular Data	Ensemble Feature Selection	Binary Classification
Transformer Module	Attention Mechanism	Sequential Patterns	Self-Attention Encoding	Context Representation
Feature Fusion Layer	Neural Network	Multi-Modal Inputs	Cross-Domain Correlation	Unified Feature Vector
Multi-Head Attention	Parallel Processing	Different Feature Aspects	Simultaneous Learning	Comprehensive Understanding

Table 3: Hybrid Learning Architecture Components [6]

3.4 Policy Learning System for Layer Assignment Decisions

Layer assignment operations utilize strategy-based learning algorithms to execute sequential routing choices according to evolving design conditions and accumulated knowledge. Condition representation encodes current routing status, available layer capacity, density patterns, and electrical violation indicators as organized attribute collections. Decision options include discrete layer selection possibilities for individual connections, constrained by fabrication rules and manufacturing specifications. Objective evaluation balances competing goals, including length reduction, density mitigation, voltage improvement, and connection minimization through proportional combinations. Strategy network development uses actor-critic approaches with experience storage to ensure learning stability and prevent knowledge degradation during optimization sequences.

3.5 Unified Objective Formulation and Design Tool Coordination

Integrated optimization coordinates storage and routing choices through composite objective formulations that represent interdependencies across design domains. Objective construction incorporates delay penalty terms, consumption increase factors, overhead measurements, and density violation scores with flexible weighting approaches that adapt according to design stage and priority specifications. Design tool coordination employs scripting protocols and programming interfaces to facilitate information transfer between learning components and commercial design platforms. Response mechanisms capture post-optimization evaluation data to improve model accuracy and modify optimization approaches based on actual implementation results, guaranteeing continuous enhancement of decision effectiveness throughout design processes.

4. Validation Setup and Platform Configuration

4.1 Circuit Repository Assembly: Industrial Designs and Academic Testbenches

Circuit validation employs heterogeneous design collections spanning commercial development projects and research-oriented implementations to establish comprehensive algorithmic assessment frameworks. Industrial circuit samples include communication protocol processors, encryption computation engines, and digital signal transformation blocks originating from production-ready intellectual property libraries with demanding operational specifications [7]. Academic testbench circuits encompass processor core implementations, memory controller architectures, and mathematical computation units derived from educational and research initiatives. Design characteristics vary from control-dominated structures containing extensive state machine logic to arithmetic-intensive implementations featuring complex datapath organizations. Repository diversity guarantees algorithmic robustness evaluation across distinct optimization scenarios, including critical timing constraints, energy efficiency targets, and silicon area limitations.

4.2 Fabrication Process Characterization and Physical Implementation Rules

Validation employs state-of-the-art manufacturing processes incorporating nanosheet device architectures that deliver enhanced gate control capabilities beyond conventional three-dimensional transistor geometries [8]. Process characteristics encompass multiple threshold voltage selections for dynamic power management, advanced metallization schemes for interconnect resistance reduction, and comprehensive layout restriction frameworks governing geometric feature specifications and proximity requirements. Manufacturing limitations include photolithographic resolution boundaries, plasma etching process tolerances, and surface planarization technique constraints that directly influence physical design optimization strategies. Process-dependent considerations encompass mechanical stress phenomena in nanosheet device channels, inter-layer dielectric capacitance fluctuations, and heat dissipation challenges associated with elevated integration densities.

Parameter	TSMC N2 Specification	Design Impact	Constraint Type	Optimization Relevance
Transistor Architecture	Gate-All-Around Nanosheet	Improved Electrostatic Control	Physical	High
Metal Layers	12-Layer Stack (M1-M12)	Routing Resource Availability	Electrical	Very High
Minimum Feature Size	Sub-2nm Critical Dimensions	Layout Density Limitations	Geometric	High
Power Supply Voltage	Multiple VDD Options	Power/Performance Tradeoffs	Electrical	Medium
Design Rule Complexity	Advanced DRC Requirements	Manufacturing Constraints	Process	High

Table 4: Technology Node Specifications and Design Constraints [8]

4.3 Software Tool Environment Assembly: Commercial Implementation and Verification Systems

Implementation platform construction utilizes established industry software solutions configured for advanced technology node design optimization and verification workflows. Logic synthesis and physical realization operations integrate comprehensive rule validation, temporal analysis capabilities, and energy consumption estimation functionalities within cohesive development environments. Verification infrastructure encompasses electromagnetic field computation engines for parasitic parameter extraction, thermal distribution modeling systems for temperature profiling, and durability evaluation platforms for current

density assessment. System configuration incorporates specialized optimization procedures, design limitation specifications, and process library characterization datasets, enabling precise circuit behavior modeling across diverse operational scenarios.

4.4 Algorithm Implementation Environment: Computational Framework Assembly and Training Infrastructure

Learning algorithm deployment utilizes established computational platforms optimized for diverse algorithmic methodologies and hardware acceleration configurations. Deep learning implementations leverage graphics processing unit acceleration for efficient gradient computation and concurrent matrix operations throughout model development cycles. Conventional learning algorithms employ central processing unit-optimized implementations for tabular data manipulation and ensemble technique execution. Platform integration facilitates efficient information transfer between algorithmic components while preserving computational performance during optimization sequences. Training infrastructure encompasses distributed processing capabilities for extensive dataset manipulation and parameter optimization across multiple algorithmic variations.

4.5 Quality Measurement Standards and Correctness Validation Procedures

Assessment methodology defines comprehensive evaluation criteria capturing optimization effectiveness across multiple design targets and quality characteristics. Quality measurements encompass timing convergence achievement rates, energy consumption decreases, silicon utilization enhancements, and interconnect completion statistics documented across varied circuit implementations. Validation procedures include functional accuracy confirmation through behavioral simulation testing, manufacturing viability assessment using layout rule verification, and performance comparison against established optimization baselines. Statistical reliability analysis ensures dependable performance evaluation while accommodating process fluctuations and measurement variability. Validation approach incorporates cross-validation methodologies, preventing algorithm overfitting and generalization capability assessment across different circuit classifications and process variations.

5. Performance Evaluation and Algorithmic Assessment

5.1 Register Architecture Refinement: Component Reduction and Constraint Resolution

Register architecture optimization yields considerable enhancements in circuit efficiency through systematic component elimination and constraint violation resolution. Component reduction strategies successfully decrease register populations while preserving circuit functionality and maintaining operational specifications [9]. Constraint violations associated with setup timing requirements exhibit substantial resolution through strategic register positioning and selective elimination techniques that identify superfluous storage components without degrading critical path performance. Clock distribution network loading experiences corresponding reductions, contributing to diminished dynamic power requirements and enhanced distribution efficiency. Optimization success rates demonstrate variability across circuit categories, with state machine implementations showing greater component reduction potential than arithmetic processing blocks.

5.2 Metal Stack Resource Management: Utilization, Redistribution, and Density Control

Metal stack resource coordination achieves superior allocation patterns across hierarchical conductor layers, producing enhanced completion rates and mitigated density concentrations. Utilization redistribution transfers routing demands from saturated lower conductors to available upper layers, establishing balanced consumption characteristics throughout the conductor hierarchy [10]. Density control mechanisms eliminate localized bottleneck formations through strategic net assignment procedures that balance electrical characteristics with fabrication requirements. Connection point reduction accompanies enhanced layer allocation, diminishing parasitic impedance and strengthening signal transmission quality. Resource coordination benefits maintain consistency across varied circuit scales and layout arrangements.

5.3 Circuit Electrical Characteristics: Supply Network Enhancement and Fabrication Compliance

Circuit electrical optimization produces meaningful advances in power delivery effectiveness and manufacturing specification adherence across evaluated implementations. Supply network enhancement achieves substantial decreases in maximum voltage drop magnitudes through optimized distribution routing and conductor layer selection strategies. Fabrication compliance improvements emerge from strategic routing selections that reduce spacing infractions, dimension violations, and density specification failures. Current density evaluation reveals considerable reductions in migration risk factors across high-current transmission paths and power distribution structures. Supply network optimization exhibits particular success in performance-demanding implementations with substantial current delivery specifications.

5.4 Circuit Timing Characteristics: Delay Margin Enhancement and Path Optimization

Circuit timing evaluation demonstrates comprehensive advances across multiple temporal metrics, including negative slack reduction and maximum delay minimization. Path optimization accomplishes notable delay decreases through coordinated storage component positioning and conductor layer selection methodologies. Setup timing enhancements emerge from reduced clock distribution variation and optimized signal transmission that minimizes delay fluctuations across timing-sensitive

paths. Hold timing adherence benefits from enhanced clock distribution balance and decreased interconnect impedance in storage component connections. Timing optimization success exhibits strong relationships with circuit complexity levels and initial violation magnitudes.

5.5 Method Reliability Evaluation: Optimization, Balance, and Application Scope

Method reliability assessment reveals critical insights regarding optimization compromises, model applicability, and scaling characteristics across different implementation scenarios. Optimization balance examination shows that storage component reduction may extend combinational processing depth in specific circuit arrangements, necessitating careful equilibrium between component quantity and timing characteristics. Model applicability evaluation demonstrates potential specialization risks when training collections lack adequate implementation variety or operational condition coverage. Scaling assessment indicates method effectiveness increases proportionally with implementation complexity while preserving reasonable computational demands. Cross-verification outcomes demonstrate consistent performance across different fabrication processes and implementation approaches, although effectiveness fluctuates with particular circuit features and constraint hierarchies.

6. Conclusion

This work presents a comprehensive framework that jointly optimizes sequential elements and metal layer usage in advanced semiconductor physical design through artificial intelligence and machine learning techniques. The unified optimization strategy demonstrates substantial improvements in register utilization efficiency, routing congestion mitigation, timing closure enhancement, and electrical performance characteristics across diverse circuit implementations. Graph neural networks effectively capture topological patterns in circuit structures while reinforcement learning agents adapt dynamically to evolving design constraints and optimization objectives. The integration of semantic embeddings for sequential element characterization enables intelligent clustering and elimination of redundant components without compromising functional correctness. Metal layer allocation strategies successfully balance competing objectives, including wirelength minimization, congestion reduction, and IR drop improvement through multi-dimensional reward functions. Implementation results validate the effectiveness of machine learning-driven design automation in managing increasing complexity at advanced technology nodes while preserving manufacturability and reliability requirements. Future extensions include ECO-aware reoptimization capabilities that learn from post-signoff corrections, multi-die integration for chiplet architectures, and real-time embedding within commercial EDA tools for live optimization during design iterations. The demonstrated framework establishes a foundation for intelligent design automation that addresses fundamental challenges in next-generation semiconductor development through coordinated optimization across multiple design domains.

Funding: This research received no external funding

Conflicts of interest: The authors declare no conflict of interest

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Aman A, et al., (n.d) Koios: A Deep Learning Benchmark Suite for FPGA Architecture and CAD Research, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems,. [Online]. Available: https://lca.ece.utexas.edu/pubs/Koios_Benchmarks_IEEE_verified.pdf

[2] Ming Z, et al., (2006) Sequential Element Design With Built-In Soft Error Resilience, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, December 2006. [Online]. Available: https://shanbhag.ece.illinois.edu/publications/ming-tvlsi2006.pdf

[3] Mohammadhossein G et al., (2020) AI-Based Modeling and Data-Driven Evaluation for Smart Manufacturing Processes, IEEE/CAA *Journal of Automatica Sinica,* 2020. [Online]. Available: https://www.ieee-jas.net/en/article/doi/10.1109/JAS.2020.1003114

[4] Peiyu L, et al., (2023) DREAMPlace 4.0: Timing-Driven Placement With Momentum-Based Net Weighting and Lagrangian-Based Refinement, IEEE TCAD, October 2023. [Online]. Available: https://cseweb.ucsd.edu/classes/fa23/cse248-a/papers/placement/dreamplace4.pdf

[5] Peter A. B, Massoud P, (2018) Opportunities for Machine Learning in Electronic Design Automation, 2018 IEEE International Symposium on Circuits and Systems (ISCAS), May 4, 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8351731

[6] Roberto S (2021) Going Beyond Impossible: 2nm Chip Technology, IEEE Future Directions Blog, June 25, 2021. [Online]. Available: https://cmte.ieee.org/futuredirections/2021/06/25/going-beyond-impossible-2nm-chip-technology/

[7] Samuel K. M (2024) TSMC Lifts the Curtain on Nanosheet Transistors, IEEE Spectrum, December 12, 2024. [Online]. Available: https://spectrum.ieee.org/tsmc-n2

[8] Sharad M, et al., (1991) Retiming and Resynthesis: Optimizing Sequential Networks with Combinational Techniques, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, January 1991. [Online]. Available: https://people.eecs.berkeley.edu/~alanmi/courses/2007_290N/papers/rr_malik_tcad91.pdf

[9] Yali Z, et al., (2025) Hybrid LSTM–Transformer Architecture with Multi-Scale Feature Fusion for High-Accuracy Gold Futures Price Forecasting, Mathematics, Volume 13, Issue 10, May 8, 2025. [Online]. Available: https://www.mdpi.com/2227-7390/13/10/1551

[10] Ya-Shian L, et al., (2008) Semiconductor Manufacturing Equipment Data Acquisition Simulation for Timing Performance Analysis, IEEE ISPCS 2008 – International Symposium on Precision Clock Synchronization, September 22–26, 2008. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=33160