

Qohere – Constraint satisfaction engine

Qohere private limited, Bangalore

Authors: Madhava Bekkem, Vaibhav Patil

Qohere Constraint satisfaction engine reduces SSP Constraint Management Costs by 40%

The Hidden Cost of Constraint Violations in Programmatic Advertising

Executive Summary

Supply-Side Platforms (SSPs) face ongoing challenges with constraint violations that cost mid-to-large platforms \$500K-\$2M annually in penalties, make-goods, and operational overhead. Current machine learning and rule-based systems achieve 97-98% constraint satisfaction, leaving room for meaningful improvement.

The **Qohere Constraint Satisfaction Engine** introduces an advanced optimization approach that achieves **99.7%** constraint satisfaction while maintaining **sub- 10 ms latency** at scale. Through sophisticated algorithms and intelligent caching, Qohere reduces constraint violations by **85%** compared to standard systems.

This white paper examines the real cost of constraint violations, the technical challenges involved, and how advanced optimization provides measurable improvements.

Table of Contents

1. The Constraint Violation Challenge
2. Understanding SSP Constraints
3. The Mathematical Challenge
4. Why Current Solutions Have Limitations
5. The Advanced Optimization Approach
6. The Qohere Solution
7. Performance and Results
8. Business Impact
9. Implementation
10. Conclusion

The Constraint Violation Challenge

The Real Cost Structure

For a typical mid- size SSP processing 50- 100 billion bid requests monthly, constraint violations create measurable costs:

- **Direct Penalties:** \$100- 300 K annually in contractual adjustments
- **Make- Goods:** \$200- 400 K in compensatory impressions
- **Operational Costs:** \$150- 250 K managing violation investigations and remediation
- **Opportunity Cost:** \$100- 200 K from reduced advertiser confidence

The Compliance Reality

Constraint violations create genuine business challenges:

- **Contractual Risk:** SLA violations and penalty clauses
 - **Relationship Management:** Advertiser satisfaction issues
 - **Operational Burden:** Time spent on violation investigation
-

Understanding SSP Constraints

Types of Constraints

SSPs must simultaneously satisfy multiple constraint types:

1. *Private Marketplace (PMP) Guarantees*

- Contractual obligations to deliver specific impression volumes
- Typical violation rate: **2- 3%**
- Primary cause: Inventory availability fluctuations

2. *Frequency Caps*

- User- level impression limits (hourly, daily, weekly)
- Typical violation rate: **1- 2%**
- Primary cause: Cross- device tracking gaps

3. *Competitive Separation*

- Minimum time gaps between competing advertisers
- Typical violation rate: **0.5- 1%**
- Primary cause: Real- time processing delays

4. *Brand Safety*

- Content- category restrictions
- Typical violation rate: **0.3- 0.5%**

- Primary cause: Classification errors

5. *Budget Caps*

- Daily spend limits and pacing
- Typical violation rate: **1- 1.5%**
- Primary cause: Distributed- system synchronization

The Scale Challenge

For a mid- size SSP:

- **500** active campaigns
 - **20- 30** constraints per campaign
 - **5 000** bid requests per second
 - **150 million** constraint checks *per second*
-

The Mathematical Challenge

The Constraint Satisfaction Problem (CSP)

The SSP constraint problem remains computationally complex, though manageable with modern approaches.

Practical Complexity

While theoretically NP- hard, real- world SSP constraints exhibit structure that enables optimization:

Constraint Clustering

- **70%** of constraints are independent
- **25%** have limited interaction
- **5%** require complex resolution

Time Windows

- Most constraints operate on predictable cycles
- Historical patterns provide optimization opportunities
- Caching strategies can reduce computation by **60%**

Mathematical Formulation

Mathematical Formulation

The SSP constraint management problem is formulated as an integer linear program (ILP) to maximize revenue while minimizing penalties. While NP-hard in general, real-world structure (e.g., 70% independent constraints) enables efficient solving via advanced methods like FALQON for feedback-based optimization.

Objective Function

maximize:

$$\sum_{i=1}^N (\text{CPM}_i \cdot x_i) - \lambda \cdot \sum_{j=1}^M (\text{Penalty}_j \cdot v_j)$$

Where:

- $x_i \in \{0,1\}$: Serve impression i (1) or not (0)
- $v_j \in \{0,1\}$: Violation of constraint j (1) or not (0)
- CPM_i : Revenue per mille for impression i
- Penalty_j : Cost of violating constraint j
- $\lambda = 10$: Weight factor prioritizing constraint satisfaction
- N : Bid requests (e.g., 5,000/second)
- M : Constraints (e.g., 20-30/campaign \times 500 campaigns)

Subject to Constraints

1.PMP Delivery Guarantees (2-3% typical violation rate)

$$\sum_{i \in \text{PMP}_k} x_i \geq \text{Guaranteed}_k \quad \forall k \in \text{PMPs}$$

2. Frequency Caps (1-2% violation rate)

$$\sum_{i \in I(u,a,t)} x_i \leq \text{Cap}_{a,t} \quad \forall u \in \text{Users}, \forall a \in \text{Advertisers}, \forall t \in \text{TimeWindows}$$

3. Competitive Separation (0.5-1% violation rate)

$$x_i + x_j \leq 1 \quad \forall (i,j) \text{ where:}$$

- $\text{advertiser}(i) \in \text{Competitors}(\text{advertiser}(j))$
- $\text{slot}(i) = \text{slot}(j)$
- $|\text{time}(i) - \text{time}(j)| < \text{MinSeparation}$

4. Budget Caps (1-1.5% violation rate)

$$\sum_{i \in A(a)} \text{CPM}_i \cdot x_i \leq \text{Budget}_a(t) \quad \forall a \in \text{Advertisers}$$

Why Current Solutions Have Limitations

Machine Learning Approaches

Performance: 97- 98% constraint satisfaction—good for standard cases; struggles with edge cases.

Limitations:

1. Training- data requirements: 3- 6 months of history
2. Black- box decisions: difficult to debug violations
3. Adaptation lag: 24- 48 h to adjust to new patterns

Rule- Based Systems

Performance: 95- 97% constraint satisfaction—predictable behavior; higher latency.

Limitations:

1. Maintenance overhead: ~200 h/yr of rule updates
2. Performance impact: +15- 20 ms latency
3. Conflict resolution: manual intervention required

Hybrid Approaches

Current best practice combines both:

- 98- 98.5% satisfaction rate
- 10 ms average latency
- Still leaves **\$500K- 1M** in annual costs

The Advanced Optimization Approach

Multi- Stage Optimization

The Qohere approach uses sophisticated but practical techniques:

Stage 1: Preprocessing

- Constraint clustering & categorization
- Historical- pattern analysis
- Priority- queue construction

Stage 2: Real- Time Resolution

- Parallel constraint checking
- Intelligent caching (**80%** hit rate)
- Fast- path for common cases

Stage 3: Conflict Resolution

- Graph- based constraint analysis
- Penalty- minimization algorithms
- Guaranteed resolution in **< 10 ms**

Key Algorithms

Constraint Graph Optimization

- Represents constraints as directed graphs
- Identifies independent subgraphs for parallel processing
- Reduces effective complexity by **70%**

Adaptive Caching

- LRU cache for recent decisions
- Pattern- based prefetching
- **80%** cache hit rate in production

Smart Conflict Resolution

- Priority- based resolution
 - Minimal- penalty path finding
 - Backtracking with memoization
-

The Qohere Solution

Architecture Overview

The Qohere Constraint Satisfaction Engine provides:

1. **Advanced Optimization Core**
 - Multi- stage constraint processing
 - Intelligent caching layer
 - Parallel execution engine
2. **Real- Time Processing**
 - **5 ms** average latency
 - **10 ms** P99 latency
 - **50 000 QPS** per instance
3. **Integration Layer**
 - REST API
 - Prebid.js adapter
 - Real- time monitoring

Key Features

High Accuracy

- **99.7%** constraint satisfaction
- **85%** reduction in violations
- Configurable strictness levels

Performance

- Horizontal scaling
- Sub- 10 ms latency
- Minimal infrastructure requirements

Observability

- Real- time violation tracking
- Detailed constraint analytics
- Performance dashboards

Performance and Results

Benchmark Results

Metric	Qohere	ML- Based	Rule- Based
Satisfaction Rate	99.7%	98%	97%
Latency (P99)	10 ms	15 ms	20 ms
Annual Penalties (mid- size SSP)	\$50- 100 K	\$200- 400 K	\$300- 500 K
ROI	250%	120%	80%

Production Results

Large E- commerce Event (Black Friday)

- 2x normal traffic
- 800+ active constraints
- **0.3%** violation rate (vs 2% baseline)
- **\$180 K** in penalties avoided

Premium Publisher Network

- 50 PMP campaigns
 - Complex competitive separation
 - **99.5%** delivery rate achieved
 - **\$75 K** reduction in make- goods
-

Business Impact

Financial Analysis for Mid- Size SSP

Annual Cost Savings:

- Reduced Penalties: **\$150- 250 K**
- Fewer Make- Goods: **\$100- 200 K**
- Operational Efficiency: **\$50- 100 K**
- **Total Savings: \$300- 550 K**

Investment:

- Qohere License: **\$25 -100 K/yr**
- **Payback: 3-4 months**

Operational Benefits

1. **Reduced Violations:** 85% fewer incidents
2. **Advertiser Satisfaction:** Improved NPS scores
3. **Team Efficiency:** 50% less time on violations
4. **Competitive Positioning:** Premium advertiser attraction

Case Study: Regional SSP

Metric	Before Qohere	After Qohere (6 mo)
Violation Rate	2%	0.3%
Annual Costs	\$800 K	\$120 K
FTE on Violations	1.0	0.2
Advertiser Feedback	Quarterly complaints	Positive

Results: \$680 K savings · 85% violation reduction · Team redeployed to revenue- gen · 2 new premium advertisers onboarded

Implementation

Integration Process

Phase	Duration	Key Activities
1. Assessment	Week 1	Violation analysis, constraint audit, ROI calc
2. Integration	Week 2	API hookup, test env, config
3. Pilot	Week 3	5% traffic, monitoring, tuning
4. Rollout	Week 4	Gradual ramp- up, success tracking

Technical Requirements

- REST API (2- 3 days)
- Compatible with major SSPs
- No infra changes needed
- Optional Prebid.js adapter

Support & Training

- Integration assistance
- Ops training (4 h)
- Monthly optimization reviews
- 24/7 monitoring dashboard

Conclusion

Constraint violations remain a costly operational challenge for SSPs, even with modern ML and rule systems. The Qohere Constraint Satisfaction Engine delivers **99.7%** accuracy with **sub- 10 ms** latency, reducing violations by **85%** and yielding a typical **3- year ROI of 250%**. By transforming constraint management from a cost center to a strategic advantage, Qohere helps SSPs cut costs, satisfy advertisers, and unlock new revenue.

About Qohere

Qohere specializes in advanced optimization solutions for programmatic ad tech , logistics and supply chain management, Blockchain, drug discovery, finance, material sciences. Built by experts in distributed systems and ad- tech optimization, Qohere delivers practical solutions to complex operational challenges.

Contact

- **Email:** info@qohere.in
- **Web:** <https://www.qohere.in>