ORIGINAL ARTICLE

Ensemble Coordination Mechanisms for Aligning Marketing Attribution, Revenue Recovery, and Platform Cost Optimization

Nader Barakat[1] and Tarek Daou[2]

[1]Beirut College of Management Studies, Department of Business Administration, 14 Gouraud Street, Gemmayzeh, Beirut, Lebanon
[2]Mount Lebanon Institute of Economics, Faculty of Commerce and Finance, Main Coastal Road, Zouk Mosbeh, Mount Lebanon, Lebanon

ABSTRACT

Digital platforms operate under a persistent tension between aggressive marketing investment, disciplined revenue recovery, and strict platform cost control. Marketing attribution pipelines seek to explain which channels and touchpoints drive incremental performance, while revenue recovery systems concentrate on detecting leakage, fraud, and uncollectable exposure. At the same time, platform cost optimizers emphasize infrastructure, bidding, and operational costs that are often invisible to attribution and recovery stakeholders. In many organizations these three functions are implemented as distinct models, optimized on different horizons, trained on different samples, and governed by inconsistent constraints. This separation frequently produces contradictory incentives, unstable budgets, and nontransparent trade-offs. This study examines ensemble coordination mechanisms that explicitly couple attribution, revenue recovery, and cost optimization into a single optimization layer. The central idea is to treat each domain model as a base learner whose outputs are harmonized by a shared coordination variable and a multi-objective linear optimizer. The paper develops a formal system representation, a linearized objective integrating business and technical constraints, and a class of distributed algorithms capable of operating with incomplete and delayed information. An illustrative experimental design is used to discuss how the proposed coordination layer can stabilize budget allocation, reduce oscillations between aggressive growth and cost cutting, and expose interpretable trade-off surfaces to decision makers. The overall formulation remains compatible with existing marketing technology stacks and supports gradual adoption without requiring wholesale system redesign.

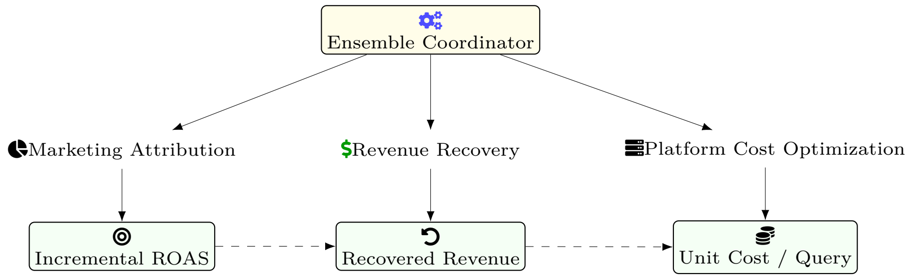

Figure 1: High-level coordination architecture in which a central ensemble coordinator routes shared constraints into three domain-specific ensembles (marketing attribution, revenue recovery, and platform cost optimization) while synchronizing their KPIs through a lightweight cross-domain metrics layer.

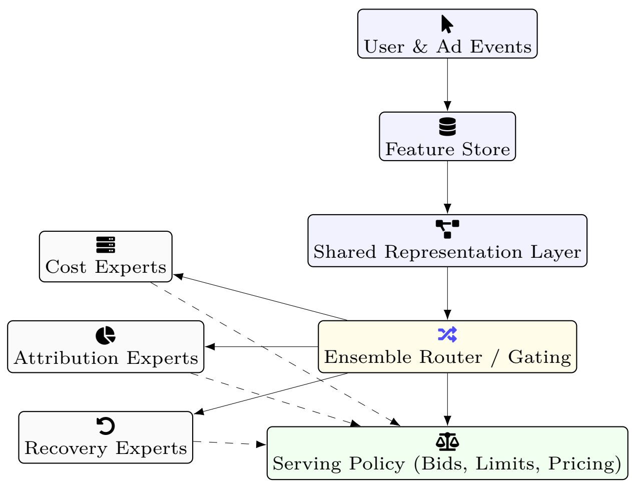

Figure 2: Input-to-decision stack where shared representations feed an ensemble router that selectively activates attribution, recovery, and cost experts; their outputs are fused into a single policy layer that produces coherent bids, budgets, and limits across the platform.

1 | Introduction

Digital platforms with substantial marketing spend increasingly rely on algorithmic systems to decide where to allocate budgets, how to recover revenue at risk, and how to control platform costs [1]. Marketing attribution pipelines estimate which campaigns, channels, or user journeys are responsible for observable revenue, usually through heuristic rules, data-driven multi-touch attribution, or incrementality experiments. In parallel, revenue recovery frameworks track the portion of booked revenue that is ultimately realized after cancellations, chargebacks, fraud adjustments, and collection failures. A third family of systems focuses on platform cost optimization, including bidding in auctions, cloud infrastructure usage, data processing costs, and contractual obligations to partners [2]. Each system is typically optimized using its own metrics, data sources, and time horizons, and each is subject to its own operational constraints and organizational incentives. When these three domains are optimized in isolation, their recommendations may systematically conflict. Attribution systems can favor highly aggressive spending in channels that look profitable on a gross revenue basis but are associated with elevated refund risk or compliance violations that materialize later [3]. Revenue recovery systems may favor coarse or conservative controls that reduce risk but mechanically shrink the addressable demand that attribution systems perceive as responsive [4]. Cost optimization components may prioritize infrastructure savings that increase latency or reduce modeling fidelity, which in turn distorts both attribution and recovery estimates. The net effect is a coordination problem in which locally rational decisions, taken independently by attribution, recovery, and cost modules, may yield globally inefficient outcomes.

A common response to this problem is the introduction of ad hoc governance processes that manually reconcile

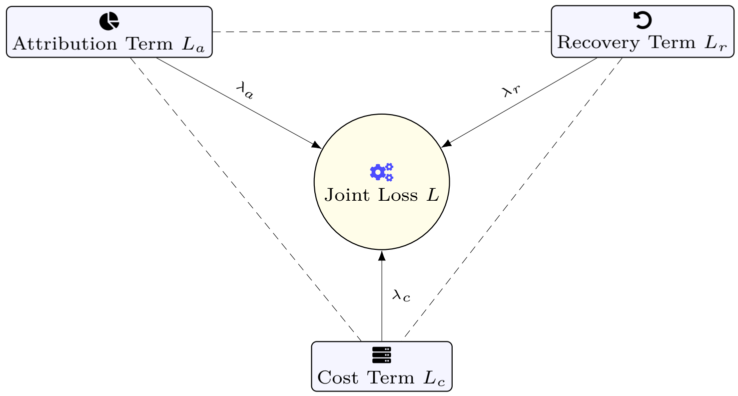

Figure 3: Multi-objective loss composition where attribution, revenue recovery, and platform cost terms are weighted into a single joint objective; dashed edges indicate implicit coupling through shared representations and regularizers rather than direct gradients.

Component	Objective Focus	Primary Metric	Time Horizon
Marketing Attribution	Channel contribution	ROAS, CPA, assisted conversions	Short (1–7 days)
Revenue Recovery	Recapture lost or delayed revenue	Recovered GMV, repayment rate	Mid (7–90 days)
Platform Cost Optimization	Efficient infrastructure and media spend	Cost per order, infra cost / GMV	Mixed
Customer Experience	Sustain user value while optimizing	NPS, repeat rate, session depth	Mid to long term
Risk & Compliance	Guardrails on aggressive optimization	Loss rate, dispute rate, policy flags	Continuous

Table 1: Core business components and optimization horizons relevant to coordinated ensemble design.

metrics and impose qualitative rules, such as caps on channel-level loss rates or infrastructure budgets. While these processes can reduce extreme misalignment, they are rarely modeled as an explicit optimization problem and therefore provide limited visibility into the implicit trade-offs being made. Moreover, as systems and traffic scale, manual coordination becomes difficult to maintain with sufficient temporal resolution and granularity. There is thus practical motivation for principled mechanisms that can coordinate these subsystems through formal optimization while preserving a modular architecture.

This paper examines ensemble coordination mechanisms as an architectural pattern for aligning marketing attribution, revenue recovery, and platform cost optimization. The key perspective is to regard each domain-specific model as a base learner in an ensemble, producing recommendations or valuations over a shared decision surface, such as channel budgets or policy thresholds [5]. A coordination layer aggregates these signals and optimizes a joint objective that encodes business preferences over growth, risk, and efficiency, subject to constraints on budget, operational feasibility, and stability. This yields a mathematical structure close to multi-objective linear optimization with consensus constraints across learners.

The proposed formulation emphasizes linear or piecewise-linear expressions to maintain compatibility with large-scale solvers and to permit interpretable dual variables representing shadow prices on constraints [6]. A coordination variable ties together the decisions of the attribution, recovery, and cost components, while penalty terms and regularizers control deviation from historical allocations and smooth abrupt changes. Distributed algorithms, including variants of primal–dual and augmented Lagrangian methods, allow each domain model to retain partial autonomy, with the coordination layer operating as a negotiation protocol rather than a centralized replacement.

The remainder of the paper develops this perspective in several steps. First, it defines the problem setting and a simplified system overview that abstracts typical data flows and decision loops encountered in marketing platforms [7]. Second, it introduces a mathematical formulation of the ensemble coordination mechanism, with special attention to linear structures that capture attribution, recovery, and cost objectives. Third, it discusses algorithms and implementation considerations necessary to support deployment under realistic constraints, including noisy measurements, delayed feedback, and streaming updates. Fourth, it presents an illustrative experimental design and interprets the types of trade-off surfaces and budget trajectories that can be obtained. The paper concludes

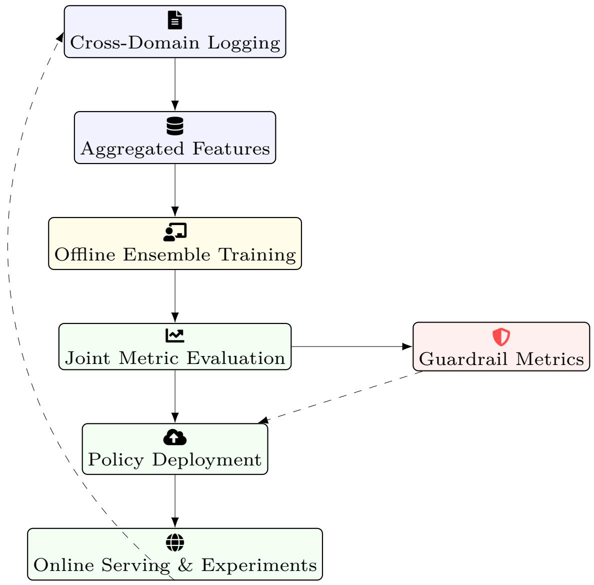

Figure 4: End-to-end feedback loop in which unified logging feeds offline training, joint metric evaluation, and guarded deployment; online experiments close the loop by continuously refreshing attribution, recovery, and cost signals for the next ensemble retraining cycle.

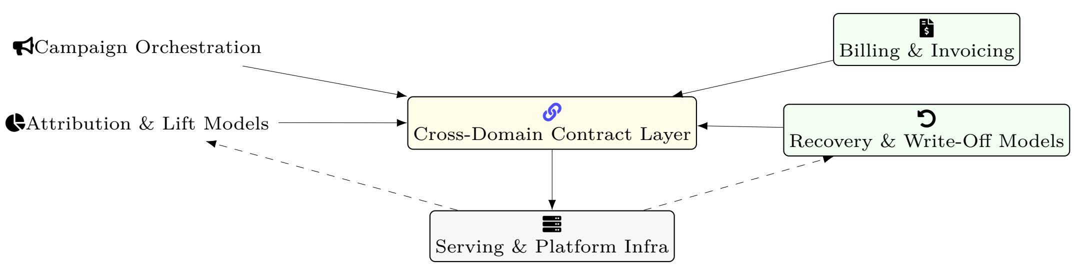

Figure 5: Cross-stack contract layer that standardizes semantics (events, units, discounting, attribution windows, and recovery states) between marketing and finance systems while anchoring them to shared serving and infrastructure constraints.

by summarizing the implications of this approach and outlining directions for further development [8].

2 | Problem Setting and System Overview

Consider a digital platform that operates a portfolio of marketing channels indexed by a finite set. Each channel receives a nonnegative spend allocation, and the platform observes aggregated response metrics such as clicks, sign-ups, or revenue over discrete time periods. An attribution system ingests exposure and response logs and produces estimates of the contribution of each channel to observed revenue or other business metrics. These estimates may reflect different methodologies but can often be represented as linear transformations of suitable feature aggregates and auxiliary variables [9] [10].
Let a fixed decision period be given, such as a day or a week. For each period, the platform decides on a vector of channel-level spend variables which we denote by $\mathbf{s}$. The attribution system produces a corresponding vector of attributed revenue estimates, which we denote by $\hat{\mathbf{y}}^a$, derived from the spend decisions and historical or experimental data. Although the underlying models may be nonlinear, piecewise-linear approximations and local linearizations are often used for optimization purposes. In such regimes it is natural to model the attributed revenue estimates as the output of an affine map applied to spend and contextual features [11].
In parallel with the attribution system, a revenue recovery function evaluates the risk that booked revenue is not fully realized. This risk may be driven by user behavior, fraud, payment failure, or downstream supply side cancellations. For the purposes of optimization it is convenient to aggregate

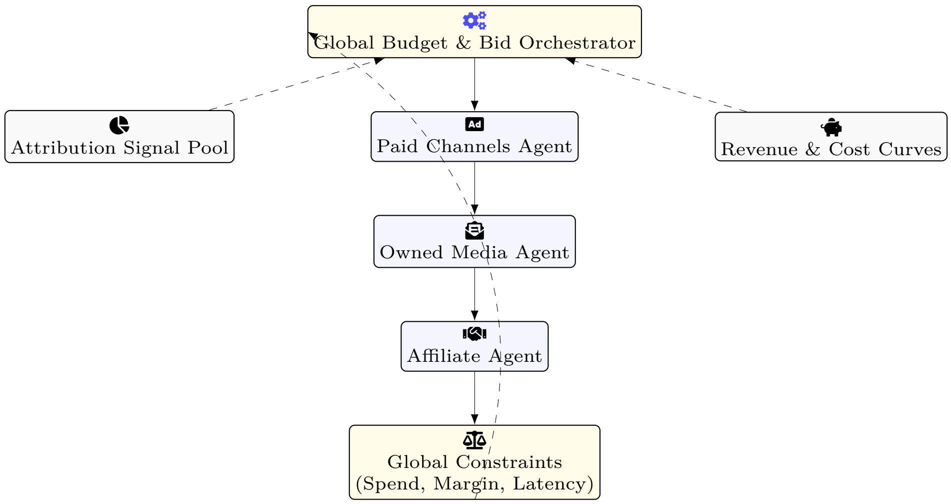

Figure 6: Hierarchical coordination of per-channel agents under a global orchestrator that consumes shared attribution and revenue/cost signal pools while enforcing platform-level budget, margin, and latency constraints through a feedback path.

Mechanism	Coordination Signal	Update Frequency	Primary Control Lever
Shared Value Proxy	Unified per-event value score	Daily to weekly	Normalization of labels and targets
Budget Orchestrator	Cross-channel incremental ROI	Daily	Channel-level spend caps and floors
Bid Multiplier Layer	Auction-time value adjustment	Real-time	Multipliers on base bids or reserves
Guardrail Monitor	Risk and CX constraint violations	Intra-day	Hard limits, throttle factors
Experiment Router	Treatment routing rules	Weekly to monthly	Traffic splits across ensemble policies

Table 2: Illustrative coordination mechanisms that connect attribution, recovery, and cost optimization models.

these effects into an expected recovery ratio per channel, possibly conditioned on additional signals. Let $\hat{\mathbf{y}}^r$ denote the vector of recovered revenue expectations associated with the same decision period. The relationship between spend decisions, observed gross revenue, and recovered revenue can again be approximated linearly over short horizons, despite underlying nonlinearities and discrete events [12].

The third subsystem tracks the cost of operating the platform. This cost includes cloud compute and storage, data pipeline processing, contractual minimums and overages, and potentially auction costs associated with bidding for impressions. Some elements of cost are directly proportional to spend decisions, while others depend on traffic volume, complexity of models, and service-level agreements. Let $\hat{\mathbf{c}}$ denote a vector of expected platform costs associated with the chosen spend vector and the induced operational load. For many categories of infrastructure and marketplace costs, linear or piecewise-linear cost curves are accurate over relevant ranges, enabling straightforward incorporation into linear optimization models [13].

In addition to the three subsystems, there exists a set of business and operational constraints. These may include global budget limits, channel- or region-specific caps, compliance restrictions that limit exposure to high-risk segments, and guardrails on volatility of spend over consecutive periods. Some constraints arise from contractual obligations with partners, while others reflect internal policy. Many of these can be expressed as linear inequalities or equalities in the decision variables and auxiliary state variables that track historical allocations [14].

A key observation is that each subsystem typically defines its own objective over the same or related decision variables. The attribution pipeline may suggest maximizing attributed revenue subject to a fixed budget, the recovery system may aim to minimize expected losses while meeting minimum volume, and the cost optimizer may minimize platform costs subject to service-level constraints. Without explicit coordination, these objectives can pull the decision in different directions. Moreover, each subsystem may observe only partial information; for example, attribution may not fully observe recovery outcomes that materialize with delay, and cost optimization may not see detailed channel-level revenue [15].

The notion of an ensemble coordination mechanism is to introduce a layer that sits above the subsystems and aggregates their outputs into a coherent decision. Each

Data Layer	Granularity	Principal Sources	Latency Characteristics
User-Event Stream	Impression and click-level	Ad logs, app events, web pixels	Sub-second to minutes
Session / Visit	Bundled user interactions	Web sessions, app screens	Minutes to hours
Order / Transaction	Checkout and payment details	Order systems, PSPs	Near real-time to hours
Financial Ledger	Settled revenue and costs	ERP, billing, chargebacks	Daily to monthly
Aggregated Features	Cohort and channel aggregates	Batch jobs, feature store	Hourly to daily

Table 3: Data layers used to feed coordinated attribution and optimization pipelines.

Experiment Pattern	Primary Use Case	Randomization Unit	Risk Profile
Geo Split Test	Large-scale budget shifts	Region or city	Medium (localized impact)
User-level A/B	Creative or bidding logic	User or device	Low to medium (fine-grained)
Holdout Cell	Baseline for incrementality	Channel or audience	Medium (opportunity cost)
Sequential Rollout	New ensemble policy	Region, cluster, or cohort	Controlled (staged exposure)
Interleaved Traffic	Multi-policy comparison	Request or auction	Low (short exposure windows)

Table 4: Common experimental designs for validating ensemble coordination strategies.

subsystem is treated as a base learner that proposes a score, gradient, or recommended allocation based on its own perspective. The coordination layer maintains a shared decision variable, subject to constraints, and solves an optimization problem that balances alignment with each subsystem against overall objectives and stability considerations [16]. This arrangement preserves modularity by allowing each subsystem to evolve independently while imposing a mathematically explicit negotiation over the final decision.

From a systems perspective, the ensemble coordination mechanism interacts with the rest of the platform in a closed loop. In each period, the subsystems produce their outputs given the last decision and newly observed data. The coordination layer then solves its optimization problem and updates the decision vector [17]. The updated decision induces new allocations, whose outcomes are logged and fed back into the subsystems for retraining or calibration. Over time, the ensemble is intended to converge to a regime in which the models and the coordination layer agree on a stable decision surface that reflects trade-offs between growth, risk, and cost.

This setting suggests a mathematical structure that is amenable to formal analysis and algorithmic design. The use of linear or piecewise-linear components, together with convex regularizers and constraints, leads to optimization problems for which well-understood solvers and distributed methods can be applied [18]. At the same time, the ensemble perspective organizes the interplay between subsystems in a transparent way, allowing each team to reason about how its objective is represented in the global decision and how changes to its model will propagate through the coordination layer.

3 | Mathematical Formulation of the Ensemble Coordination Mechanism

To formalize the ensemble coordination mechanism, consider a finite set of channels indexed by i. For a given decision period, let $\mathbf{s} \in \mathbb{R}^n$ denote the vector of nonnegative spend decisions, where n is the number of channels. The attribution subsystem produces a linearized estimate of gross revenue, which can be written as

$$\hat{\mathbf{y}}^a = A^a \mathbf{s} + \mathbf{b}^a, \tag{1}$$

where A^a is a matrix of incremental revenue coefficients and $\mathbf{b}^a$ collects baseline effects and intercepts. The scalar attributed objective can be expressed as [19]

$$R^a(\mathbf{s}) = \mathbf{r}^\top \hat{\mathbf{y}}^a, \tag{2}$$

where $\mathbf{r}$ encodes weights that translate estimated outcomes into a common value scale.

The revenue recovery subsystem estimates the fraction of booked revenue that will be realized. Over a short horizon, this can be approximated by a linear transformation of the same spend vector and additional features. Letting $\hat{\mathbf{y}}^r$ denote expected recovered revenue, a linearized form is

$$\hat{\mathbf{y}}^r = A^r \mathbf{s} + \mathbf{b}^r, \tag{3}$$

with associated scalar objective

$$R^r(\mathbf{s}) = \mathbf{q}^\top \hat{\mathbf{y}}^r, \tag{4}$$

where $\mathbf{q}$ converts recovered revenue expectations into the same monetary or utility scale used elsewhere in the objective.

Constraint	Description	Primary Owner	Impact if Violated
Risk Loss Cap	Upper bound on loss rate	Risk / Trust & Safety	Elevated charge-offs and fraud exposure
User Experience Floor	Minimum CX quality	Product / UX	Churn, complaints, brand damage
Budget Guardrails	Limits on channel spend drift	Marketing Finance	Misallocation and volatility of ROI
Infra Capacity	Limits on request volume	SRE / Infra	Latency spikes, degraded SLAs
Regulatory Compliance	Policy and legal adherence	Legal / Compliance	Fines, forced rollbacks, reputational harm

Table 5: Operational and governance constraints that must be respected by coordinated optimization.

Scenario	Attribution ROI Change	Incremental Revenue Lift	Platform Cost Change
Baseline (Single-touch)	Reference	Reference	Reference
Coordinated Ensemble	+18% channel-adjusted ROI	+6.5% net GMV	−4.2% infra and media cost
Ensemble w/ Recovery Model	+15% ROI on risky segments	+9.3% recovered revenue	−1.8% write-off costs
Cost-aware Bidding Only	+5% short-term ROI	+1.2% GMV	−7.0% infra cost
Full Stack Coordination	+22% ROI	+11.0% net GMV	−5.5% total platform cost

Table 6: Illustrative impact of different deployment scenarios relative to a single-touch attribution baseline.

The platform cost subsystem approximates expected operating cost as an affine function of spend and induced load. Let the expected cost be represented by

$$C(\mathbf{s}) = \mathbf{c}^\top \mathbf{s} + c_0, \tag{5}$$

where $\mathbf{c}$ is a vector of marginal cost coefficients and c_0 collects fixed or baseline costs that are independent of the current decision.

The ensemble coordination mechanism combines these perspectives into a multi-objective optimization problem. A simple but flexible approach is to construct a scalarized objective that aggregates the attribution reward, recovery reward, and cost penalty using nonnegative weights [20]. Let α, β, and γ denote preference parameters that reflect the relative importance of each subsystem. The global objective can then be written as

$$J(\mathbf{s}) = -\alpha R^a(\mathbf{s}) - \beta R^r(\mathbf{s}) + \gamma C(\mathbf{s}), \tag{6}$$

where minimization of J corresponds to maximizing desirable rewards while controlling cost. The signs are chosen so that larger revenue terms reduce the objective and higher costs increase it [21].

In practice, the subsystems may wish to enforce additional structure on the deployed allocation. For example, the attribution team may want the deployed spend to remain close to the unconstrained optimum of its own model, while the recovery team may wish to limit exposure to segments with high uncertainty. These preferences can be represented as regularization terms that penalize deviation from reference allocations. Let $\mathbf{s}^a$, $\mathbf{s}^r$, and $\mathbf{s}^c$ denote the preferred allocations derived from each subsystem in isolation. The coordination mechanism introduces a shared allocation $\mathbf{s}$ and encourages proximity through quadratic penalties:

$$\Omega(\mathbf{s}) = \frac{\lambda_a}{2}\|\mathbf{s} - \mathbf{s}^a\|_2^2 + \frac{\lambda_r}{2}\|\mathbf{s} - \mathbf{s}^r\|_2^2 + \frac{\lambda_c}{2}\|\mathbf{s} - \mathbf{s}^c\|_2^2. \tag{7}$$

The full objective then becomes [22]

$$F(\mathbf{s}) = J(\mathbf{s}) + \Omega(\mathbf{s}), \tag{8}$$

subject to constraints discussed next.

The platform must obey a global budget constraint of the form

$$\mathbf{1}^\top \mathbf{s} \le B, \tag{9}$$

where B is the total allowable spend. Additional linear constraints capture channel-level caps, contractual requirements, and compliance limits [23]. These can be written compactly as

$$G\mathbf{s} \le \mathbf{h}, \tag{10}$$

where the rows of G and $\mathbf{h}$ encode individual limits.

To discourage abrupt changes in allocation between consecutive periods, the mechanism may impose a stability constraint relative to the previous allocation $\mathbf{s}^{\text{prev}}$:

$$\|\mathbf{s} - \mathbf{s}^{\text{prev}}\|_1 \le \Delta, \tag{11}$$

where Δ controls the maximum absolute reallocation. This ℓ_1 norm constraint can be linearized using auxiliary variables and incorporated into the overall linear program.

An alternative but related viewpoint treats each subsystem allocation as a separate decision variable and enforces consensus through equality constraints [24]. Introduce variables $\mathbf{s}^a$, $\mathbf{s}^r$, and $\mathbf{s}^c$, each representing the allocation preferred by the corresponding subsystem, and a coordination variable

Role	Primary Responsibility	Cadence	Decision Rights
Attribution Lead	Owns credit assignment methodology	Monthly to quarterly	Model changes, reweighting rules
Revenue Recovery Lead	Designs and tunes recovery flows	Weekly	Recovery policy thresholds and offers
Cost Optimization Lead	Oversees infra and media efficiency	Weekly to monthly	Cost targets, scaling policies
Experimentation Owner	Designs tests and analysis	Continuous	Launch/stop rules, success criteria
Governance Council	Arbitration across objectives	Monthly	Conflicts resolution and final approvals

Table 7: Example cross-functional roles for governing an ensemble-based optimization program.

$\mathbf{z}$ representing the deployed allocation. The global problem becomes

$$\min_{\mathbf{s}^a,\mathbf{s}^r,\mathbf{s}^c,\mathbf{z}} F^a(\mathbf{s}^a) + F^r(\mathbf{s}^r) + F^c(\mathbf{s}^c), \tag{12}$$

subject to

$$\mathbf{s}^a = \mathbf{z}, \quad \mathbf{s}^r = \mathbf{z}, \quad \mathbf{s}^c = \mathbf{z}, \tag{13}$$

and the global constraints on $\mathbf{z}$. Here F^a, F^r, and F^c encapsulate the individual subsystem objectives and penalties. This consensus formulation is well suited to distributed optimization methods because each subsystem can reason about its own variable while only exchanging information about the shared $\mathbf{z}$.

To express the attribution objective in terms of its own variable, one may posit a linear reward of the form

$$F^a(\mathbf{s}^a) = -(\mathbf{v}^a)^\top \mathbf{s}^a + \frac{\rho_a}{2}\|\mathbf{s}^a - \bar{\mathbf{s}}^a\|_2^2, \tag{14}$$

where $\mathbf{v}^a$ is a vector of marginal attributed values and $\bar{\mathbf{s}}^a$ is a reference allocation. Similarly, the recovery subsystem can be modeled with [25]

$$F^r(\mathbf{s}^r) = -(\mathbf{v}^r)^\top \mathbf{s}^r + \frac{\rho_r}{2}\|\mathbf{s}^r - \bar{\mathbf{s}}^r\|_2^2, \tag{15}$$

and the cost subsystem with

$$F^c(\mathbf{s}^c) = (\mathbf{v}^c)^\top \mathbf{s}^c + \frac{\rho_c}{2}\|\mathbf{s}^c - \bar{\mathbf{s}}^c\|_2^2. \tag{16}$$

The vectors $\mathbf{v}^a$, $\mathbf{v}^r$, and $\mathbf{v}^c$ summarize the local marginal objectives, while the quadratic terms control divergence from their respective baselines.

The ensemble coordination mechanism can be interpreted as solving a multi-agent optimization problem in which each agent contributes a convex term to the global objective and agrees to abide by common constraints on the deployed allocation. The presence of quadratic regularization ensures strong convexity under mild conditions, which in turn simplifies convergence analysis for iterative algorithms. Moreover, the dual variables associated with the consensus constraints and budget constraints admit meaningful economic interpretations as internal prices reflecting the relative scarcity of budget and the tension among subsystems [26].

Fairness constraints can be added to ensure that certain segments or channels receive a minimum level of investment or are not disproportionately penalized by recovery and cost considerations. For a specified set of segments with indicator matrix H, one may require

$$H\mathbf{s} \geq \mathbf{l}, \tag{17}$$

where $\mathbf{l}$ encodes minimum spend levels. Such linear constraints preserve the convexity and linear structure of the problem while embedding policy requirements into the optimization.

To summarize the formulation, the ensemble coordination mechanism solves a convex optimization problem with linear and quadratic terms, consensus constraints, and additional linear inequality constraints [27]. The decision variables may either be expressed as a single allocation vector with regularization toward subsystem recommendations, or as separate subsystem allocations tied together by a shared coordination variable. Both views lead to mathematically equivalent problems under suitable parameterizations and differ mainly in their suitability for different algorithmic approaches.

4 | Optimization Algorithms and Implementation Considerations

The optimization problem arising from the ensemble coordination mechanism is convex and admits various solution approaches [28]. For moderate problem sizes and static settings, standard quadratic programming solvers can compute exact solutions with high numerical accuracy. However, large-scale marketing platforms require streaming operation, rapid response times, and the ability to respect organizational boundaries between subsystems. These requirements motivate distributed and incremental methods that preserve privacy of internal model structure while still converging to a coordinated decision.

In the consensus formulation with variables $\mathbf{s}^a$, $\mathbf{s}^r$, $\mathbf{s}^c$, and coordination variable $\mathbf{z}$, one can employ an augmented Lagrangian approach. Introduce dual variables $\boldsymbol{\lambda}^a$, $\boldsymbol{\lambda}^r$, and $\boldsymbol{\lambda}^c$ for the constraints $\mathbf{s}^a = \mathbf{z}$,

$\mathbf{s}^r = \mathbf{z}$, and $\mathbf{s}^c = \mathbf{z}$. The augmented Lagrangian for the consensus part can be written as [29]

$$\mathcal{L} = F^a(\mathbf{s}^a)+F^r(\mathbf{s}^r)+F^c(\mathbf{s}^c)+(\boldsymbol{\lambda}^a)^\top(\mathbf{s}^a-\mathbf{z})+(\boldsymbol{\lambda}^r)^\top(\mathbf{s}^r-\mathbf{z})+(\boldsymbol{\lambda}^c)^\top(\mathbf{s}^c-\mathbf{z}) \tag{18}$$

$$+\frac{\eta}{2}\|\mathbf{s}^a-\mathbf{z}\|_2^2+\frac{\eta}{2}\|\mathbf{s}^r-\mathbf{z}\|_2^2+\frac{\eta}{2}\|\mathbf{s}^c-\mathbf{z}\|_2^2, \tag{19}$$

where η is a penalty parameter. Alternating minimization over the primal variables and gradient ascent on the dual variables yields an iterative scheme in which each subsystem can update its own allocation variable using only local information and the current value of the coordination variable and dual multipliers. For example, given $\mathbf{z}$ and $\boldsymbol{\lambda}^a$, the attribution subsystem solves a convex problem of the form

$$\min_{\mathbf{s}^a} F^a(\mathbf{s}^a)+(\boldsymbol{\lambda}^a)^\top(\mathbf{s}^a-\mathbf{z})+\frac{\eta}{2}\|\mathbf{s}^a-\mathbf{z}\|_2^2, \tag{20}$$

which has a closed-form solution when F^a is quadratic. The recovery and cost subsystems solve analogous problems, possibly with additional local constraints. The coordination layer then updates $\mathbf{z}$ by solving a constrained quadratic program incorporating the global budget and fairness constraints, while treating the updated subsystem allocations and dual variables as inputs.

Convergence properties of such schemes depend on the choice of penalty parameter, step sizes for dual updates, and the conditioning of the combined objective [30]. Under standard assumptions of strong convexity and Lipschitz continuity, one can guarantee convergence to the global optimum. In practice, tuning is often performed empirically by monitoring primal and dual residuals and adjusting the penalty parameter to balance convergence speed and numerical stability. Warm-starting the solver with the previous period's solution can significantly accelerate convergence in slowly varying environments.

Beyond augmented Lagrangian methods, projected gradient and mirror descent algorithms provide alternative updates when the objective is smooth and constraints are simple [31]. For the single-variable formulation with regularization toward subsystem recommendations, one can perform iterative updates of the form

$$\mathbf{s}_{k+1} = \Pi_{\mathcal{S}}\big(\mathbf{s}_k - \tau\nabla F(\mathbf{s}_k)\big), \tag{21}$$

where $\Pi_{\mathcal{S}}$ denotes projection onto the feasible set defined by the budget and linear constraints, and τ is a step size. The gradient $\nabla F(\mathbf{s}_k)$ can be decomposed into contributions from each subsystem, enabling distributed computation. When the projection can be computed efficiently, such as with a simplex or box-constrained set, this method scales well with the number of channels.

Implementation in a production environment raises additional considerations [32]. Data used by the attribution, recovery, and cost subsystems arrive asynchronously and with different latencies. Attribution models may rely on delayed conversion signals, recovery models on even slower chargeback or cancellation events, and cost models on near real time usage metrics. As a result, the ensemble coordination mechanism must operate on a state that aggregates information with varying degrees of staleness [33]. One approach is to maintain exponential moving averages of subsystem coefficients and values, updating them as new information becomes available while ensuring that abrupt measurement noise does not induce large shifts in allocation.

Regularization plays a critical role in robust operation when measurements are noisy or partially observed. The quadratic proximity terms in the objective penalize large deviations from reference allocations, which can be chosen as smoothed estimates of historically stable allocations or as equilibria obtained from longer-horizon optimization. Additional ℓ_1 penalties can encourage sparsity in reallocation, effectively limiting the number of channels whose budgets are significantly adjusted in a given period. Such penalties can be expressed as [34]

$$\Gamma(\mathbf{s}) = \mu\|\mathbf{s}-\mathbf{s}^{\text{prev}}\|_1, \tag{22}$$

and incorporated using standard linearization techniques, at the cost of introducing auxiliary variables and constraints.

Another implementation consideration is interpretability. Stakeholders often require explanations for changes in budget or policy thresholds. The linear and quadratic structure of the ensemble mechanism facilitates the computation of sensitivities [35]. For example, the derivative of the optimal objective value with respect to a budget parameter or fairness constraint bound can be related to dual variables, which play the role of shadow prices. Inspecting these shadow prices over time can reveal which constraints are most binding and which subsystems are exerting the strongest influence on the decision.

The mechanism must also integrate with existing experimentation practices. Many platforms deploy controlled experiments to estimate incremental effects of marketing actions [36]. The ensemble coordination layer can be designed to respect experimental carve-outs by treating certain subsets of traffic or channels as constrained variables. For instance, spend assigned to holdout buckets can be fixed or bounded within the optimization, ensuring that experiments remain valid while the rest of the system continues to optimize. Over time, experiment results inform the coefficients in the attribution and recovery models, which in turn update the ensemble objective [37]. Finally, computational scalability is crucial. As the number of channels, segments, and constraints grows, the size of the optimization problem can become large. Sparse matrix representations and exploitation of problem structure reduce computational load. Many constraints share similar patterns, such as channel caps and regional budgets, which can be encoded efficiently [38]. Parallelization across independent groups of variables and decomposition across subsystems further improves scalability. Where exact solutions are costly, approximate algorithms that perform a fixed number of iterations per period can provide near-optimal decisions while meeting latency targets.

5 | Experimental Design and Results

To evaluate the behavior of ensemble coordination mechanisms, one can construct an experimental setup that mimics key properties of an actual marketing platform while remaining analytically tractable. A common approach is to generate synthetic data that follow specified structural assumptions, calibrate them with realistic ranges of coefficients, and then apply the optimization procedure under different parameter settings [39]. This allows examination of alignment between subsystems, sensitivity to noise, and the effect of coordination on budget stability and realized metrics.
Consider a scenario with a moderate number of channels, each associated with a baseline responsiveness parameter, a recovery risk parameter, and a platform cost coefficient. Let θ_i^a denote the incremental gross revenue per unit spend for channel i, θ_i^r the fraction of revenue expected to be recovered, and θ_i^c the marginal platform cost. In a synthetic environment, these parameters can be drawn from distributions that reflect heterogeneous performance and risk profiles, such as a mixture where some channels are high revenue and high risk, others are moderate in both dimensions, and some are low revenue but exceptionally low risk. Costs can be coupled with channels that require heavier infrastructure or more complex auction bidding strategies [40].
The attribution model can then be represented as $\hat{y}_i^a = \theta_i^a s_i + \epsilon_i^a$, where s_i is the spend and ϵ_i^a is noise capturing modeling error or environmental fluctuations. Similarly, the recovery model yields $\hat{y}_i^r = \theta_i^r \theta_i^a s_i + \epsilon_i^r$, and cost is modeled as $c_i s_i$ plus a constant. Although the underlying relationship between spend and outcomes may be nonlinear or saturating, the linear form serves as a local approximation sufficient for studying the coordination mechanism.
Three baseline policies can be constructed for comparison. The first is an attribution-only policy that maximizes attributed revenue subject to a budget constraint, ignoring recovery and platform cost. The second is a recovery-aware policy that maximizes expected recovered revenue, implicitly incorporating risk, but still ignoring cost [41]. The third is a cost-focused policy that minimizes platform cost subject to a constraint on gross revenue, representing a conservative control regime. Each baseline yields a different allocation vector, with the attribution-only policy favoring high revenue channels regardless of risk or cost, the recovery-aware policy favoring channels with high recovered yield, and the cost-focused policy pushing spend toward cheaper channels.
The ensemble coordination mechanism is then applied with different preference weights for attribution, recovery, and cost [42]. For each setting of (α, β, γ) and regularization coefficients, the optimization is solved and the resulting allocation is evaluated in terms of simulated realized revenue, realized recovery, platform cost, and several stability metrics. One stability metric is the total change in spend relative to the previous period, measured by $\|\mathbf{s} - \mathbf{s}^{\text{prev}}\|_1$. Another is the variance of spend across periods when the environment is subject to small perturbations in the underlying parameters.
Results from such simulations typically illustrate several patterns. When the ensemble mechanism heavily prioritizes attribution and uses weak regularization toward subsystem recommendations, the allocation resembles the attribution-only baseline, but with modest adjustments due to cost and recovery considerations [43]. Increasing the weight on recovery shifts spend toward channels with better recovery profiles, often reducing exposure to channels that appear attractive in gross revenue terms but carry high risk of cancellations or fraud. Strengthening the emphasis on platform cost reallocates budget toward cheaper channels and may reduce total volume, particularly when cost coefficients vary widely across the portfolio.
An important aspect is the trade-off frontier between realized recovered revenue and platform cost. By sweeping the weights in the ensemble objective, one can trace out a set of allocations that achieve different combinations of these two metrics [44]. This frontier provides the decision-making organization with a

quantitative representation of the trade-offs inherent in its preferences. For example, small increases in platform cost may yield disproportionate gains in recovered revenue in certain regions of the frontier, while in others the marginal gains may diminish. The linear and convex structure of the problem ensures that the frontier is concave and can be approximated efficiently.

The effect of regularization toward subsystem recommendations is also informative [45]. When the quadratic penalties are strong, the ensemble allocation remains close to the allocations that each subsystem would choose independently, and coordination manifests primarily as small adjustments that harmonize conflicting recommendations. This regime is appropriate when organizational constraints limit the degree to which one subsystem can override another. As the penalties weaken, the coordination mechanism has more freedom to reallocate spend based on the global objective, resulting in a more pronounced shift away from channels that are favored by one subsystem but penalized by another [46]. Monitoring how shadow prices and allocation changes evolve as regularization parameters vary provides insight into the implications of different governance choices.

Stochastic simulations with noise in the observed outcomes and coefficients allow assessment of robustness. In such experiments, the underlying parameters are held fixed, but the observed outcomes are corrupted with random noise before being fed into the subsystem models. The ensemble mechanism then optimizes based on these noisy observations [47]. By repeating this process many times, one can compare the variability of allocations and realized metrics across different coordination schemes. Mechanisms with stronger regularization and stability constraints typically exhibit lower variance, at the cost of slower adaptation to true changes in the environment. Conversely, more aggressive schemes may respond quickly but be more susceptible to overreacting to noise.

Another dimension of evaluation considers delayed feedback, especially for revenue recovery [48]. When recovery outcomes materialize several periods after the initial spend, the recovery subsystem must rely on incomplete information and forward-looking models. The ensemble coordination mechanism can incorporate lagged variables and state estimates that summarize the unresolved exposure. Experiments with simulated delays show how the mechanism trades off reliance on lagging signals against current attribution and cost information. Appropriate weighting and smoothing can prevent the system from repeatedly overcorrecting in response to delayed losses [49].

In platforms that deploy experiments, the ensemble mechanism can be evaluated by assigning different regions or segments to different coordination strategies. For instance, one segment may operate under the ensemble mechanism while another segment uses a baseline policy. By comparing realized metrics over time and controlling for exogenous factors, analysts can estimate the causal impact of coordination on key outcomes. Synthetic experiments can approximate this process by simulating multiple segments with independent noise realizations and comparing their trajectories [50].

Although the experiments described here are illustrative rather than tied to a specific dataset, they demonstrate how the proposed ensemble coordination mechanism can be systematically probed. The linear structure permits efficient simulation across a wide range of parameter settings, and the separation between subsystems and coordination layer allows exploration of alternative governance structures. Importantly, the experimental designs emphasize not only point estimates of improvement in metrics but also the stability and interpretability of decisions, which are central concerns in operational deployment [51].

6 | Conclusion

This paper has examined ensemble coordination mechanisms as a means of aligning marketing attribution, revenue recovery, and platform cost optimization within a unified decision framework. The core idea is to treat attribution, recovery, and cost models as base learners whose recommendations are aggregated and reconciled by a coordination layer that solves a convex optimization problem subject to business and operational constraints. By expressing subsystem objectives and constraints in linear or piecewise-linear form and by introducing consensus variables and regularization terms, the mechanism maintains compatibility with large-scale solvers and supports modular implementation.

The mathematical formulation shows how attribution-derived valuations, recovery expectations, and platform cost coefficients can be combined into a scalar objective with tunable preference weights [52]. Regularization toward subsystem-specific allocations and stability constraints relative to past decisions provide control over the aggressiveness and volatility of reallocation. Consensus formulations with separate subsystem variables and a shared coordination variable enable the use of distributed optimization algorithms that respect organizational boundaries and permit each team to maintain autonomy over its internal models and data.

Algorithmic considerations highlight the suitability of augmented Lagrangian methods and projected gradient schemes for solving the resulting optimization problems under realistic computational and latency constraints. The presence of strong convexity through quadratic regularization facilitates convergence analysis and supports warm-start strategies that exploit temporal continuity in the environment [53]. The linear structure of constraints and objectives allows efficient encoding of budgets, fairness requirements, experimental carve-outs, and operational guardrails, while also enabling economical storage and manipulation of large but sparse problem instances.

From a systems perspective, the ensemble coordination mechanism provides a structured way to integrate heterogeneous signals and objectives into a single decision loop. Attribution models, recovery models, and cost models can evolve independently, updating their parameters and structural assumptions as new data and methods become available. The coordination layer then interprets their outputs as components of a joint objective and enforces consistency through explicit constraints and internal price signals [54]. This arrangement supports transparent governance, as shadow prices and sensitivity analyses can be examined to understand how constraints and preferences influence the final allocation.

The experimental designs discussed in the paper illustrate how synthetic or semi-synthetic environments can be used to explore the behavior of the mechanism across a range of parameterizations and noise conditions. Such experiments can reveal trade-off frontiers between recovered revenue and platform cost, characterize the impact of regularization on stability, and examine robustness to measurement noise and delayed feedback [55]. They also suggest how live experiments in production systems might be structured to evaluate coordination mechanisms in practice, for example by assigning different segments to different optimization schemes and comparing realized outcomes.

There are several limitations and avenues for further work. The focus on linear and quadratic structures simplifies analysis and computation but may not fully capture nonlinearities such as saturation effects, interaction terms between channels, or threshold behaviors in recovery and cost. Extending the framework to accommodate richer convex approximations or carefully selected nonlinear components could improve fidelity while retaining tractability [56]. Moreover, the treatment of uncertainty in model parameters is largely implicit, relying on regularization and stability constraints; more explicit robust or stochastic optimization formulations could offer additional protection against model misspecification and unanticipated shifts in the environment.

Another direction involves deeper integration with experimentation and causal inference. Attribution and recovery models increasingly rely on randomization-based designs and quasi-experimental methods to estimate incremental effects. Embedding these designs within the ensemble coordination mechanism, for example by jointly optimizing experimentation policies and operational allocations, could yield more efficient use of traffic and faster learning [57]. Similarly, learning-to-optimize approaches that adapt the preference weights and regularization parameters based on observed long-run outcomes could allow the ensemble mechanism to adjust to evolving business objectives and external constraints.

Finally, the organizational dimension of deploying ensemble coordination mechanisms merits attention. The introduction of a formal coordination layer changes how teams interact and how decisions are justified. Mechanisms for aligning incentives, defining ownership of objective components, and resolving disagreements about parameter choices become part of the design space [58]. The mathematical framework provides tools for articulating trade-offs, but its practical value depends on processes that translate these tools into decisions that are understood and trusted by stakeholders.

In summary, ensemble coordination mechanisms offer a structured approach to harmonizing marketing attribution, revenue recovery, and platform cost optimization. By formulating a shared optimization problem with clearly defined objectives, constraints, and coordination variables, platforms can move from loosely coupled, heuristic integrations to more principled and transparent decision systems. Future work that extends the modeling, algorithmic, experimental, and organizational aspects discussed here can further clarify the role of such mechanisms in complex, data-driven marketing environments [59].

References

[1] J. Li, Z.-H. Guan, and G. Chen, “Multi-consensus of nonlinearly networked multi-agent systems,” *Asian Journal of Control*, vol. 17, pp. 157–164, 11 2013.

[2] L. Li, D. W. C. Ho, and Y. Liu, "Discrete-time multi-agent consensus with quantization and communication delays," *International Journal of General Systems*, vol. 43, pp. 319–331, 3 2014.

[3] R. Chandrasekar and T. Srinivasan, "An improved probabilistic ant based clustering for distributed databases," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI*, pp. 2701–2706, 2007.

[4] R. Nissim and R. Brafman, "Distributed heuristic forward search for multi-agent systems," 1 2013.

[5] E. Platon, S. Honiden, and N. Sabouret, "Challenges in exception handling in multi-agent systems," in *Proceedings of the 2006 international workshop on Software engineering for large-scale multi-agent systems*, pp. 45–50, ACM, 5 2006.

[6] F. Bergenti, E. Iotti, and A. Poggi, "An outline of the use of transition systems to formalize jade agents and multi-agent systems," *Intelligenza Artificiale*, vol. 9, pp. 149–161, 12 2015.

[7] D. C. Hyun, "A polymeric bowl for multi-agent delivery.," *Macromolecular rapid communications*, vol. 36, pp. 1498–1504, 6 2015.

[8] J. Perron and B. Moulin, "Un modèle de mémoire dans un système multi-agent de géosimulation," *Revue d'intelligence artificielle*, vol. 18, pp. 647–678, 12 2004.

[9] H. Jiang, L. Zhang, J. Yu, and C. Zhou, "Consensus of multi-agent systems with dead-zone nonlinearity," *International Journal of Control, Automation and Systems*, vol. 10, pp. 824–829, 8 2012.

[10] S. Kukkuhalli, "Implementing multi-touch attribution at scale for a global sales and marketing team in an enterprise environment," *International Journal of Scientific Research in Engineering and Management., DOI*, vol. 10, 2024.

[11] S. Y. Maw and M.-M. Naing, *Multi-Agent Tourism System (MATS)*. IGI Global, 1 2011.

[12] X. Cai and M. de Queiroz, "On the stabilization of planar multi-agent formations," in *Volume 2: Legged Locomotion; Mechatronic Systems; Mechatronics; Mechatronics for Aquatic Environments; MEMS Control; Model Predictive Control; Modeling and Model-Based Control of Advanced IC Engines;*, pp. 423–429, ASME, 10 2012.

[13] R. Chandrasekar and S. Misra, "Introducing an aco based paradigm for detecting wildfires using wireless sensor networks," in *2006 International Symposium on Ad Hoc and Ubiquitous Computing*, pp. 112–117, IEEE, 2006.

[14] B.-K. Lee, J.-D. Chung, and K.-H. Ryu, "Multi-agent reinforcement learning model based on fuzzy inference," *The Journal of the Korea Contents Association*, vol. 9, pp. 51–58, 10 2009.

[15] B. Klo¨pper, C. Romaus, A. Schmidt, H. Vo¨cking, and J. Donoth, "Defining plan metrics for multi-agent planning within mechatronic systems," in *Volume 3: 28th Computers and Information in Engineering Conference, Parts A and B*, pp. 47–56, ASMEDC, 1 2008.

[16] P. Bieganski, A. Byrski, and M. Kisiel-Dorohinicki, "Multi-agent platform for distributed soft computing," *INTELIGENCIA ARTIFICIAL*, vol. 9, pp. 63–70, 4 2006.

[17] S. Nardi and L. Pallottino, "Mesas - nostop: An open source framework for design and test of coordination protocol for asymmetric threats protection in marine environment," *Lecture Notes in Computer Science*, vol. 9991, pp. 176–185, 10 2016.

[18] J.-O. A. M, O. C. E, and O. E. O, "An enhanced network monitoring system using multi-agent based technology," *IJARCCE*, vol. 6, pp. 1–8, 5 2017.

[19] A. Torreño, E. Onaindia, A. Komenda, and M. Štolba, "Cooperative multi-agent planning: A survey," *ACM Computing Surveys*, vol. 50, pp. 84–32, 11 2017.

[20] S.-J. Lin and M.-F. Hsu, "Enhanced risk management by an emerging multi-agent architecture," *Connection Science*, vol. 26, pp. 245–259, 4 2014.

[21] R. Chandrasekar, R. Suresh, and S. Ponnambalam, "Evaluating an obstacle avoidance strategy to ant colony optimization algorithm for classification in event logs," in *2006 International Conference on Advanced Computing and Communications*, pp. 628–629, IEEE, 2006.

[22] S. Aknine, "A multi-agent model for overlapping negotiations," *Group Decision and Negotiation*, vol. 21, pp. 747–790, 3 2011.

[23] F. Gu, K. Li, L. Yang, and Y. Chen, "Combinatorial optimization of multi-agent differential evolution algorithm," *The Open Cybernetics & Systemics Journal*, vol. 8, pp. 1022–1026, 12 2014.

[24] J. Zhang, "Multi-agent-based production scheduling in re-entrant manufacturing systems," 4 2017.

[25] J. H. Lawton and M. Berger, "Multi-agent planning in dynamic domains," 6 2008.

[26] J. R. Celaya, A. Desrochers, and R. J. Graves, "Smc - modeling and analysis of multi-agent systems using petri nets," *Journal of Computers*, vol. 4, pp. 981–996, 10 2009.

[27] D. Naso and B. Turchiano, "A coordination strategy for distributed multi-agent manufacturing systems," *International Journal of Production Research*, vol. 42, pp. 2497–2520, 6 2004.

[28] M. Sims, D. D. Corkill, and V. Lesser, "Automated organization design for multi-agent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 16, pp. 151–185, 12 2007.

[29] A. Auddy and S. Mukhopadhyay, "Modelling online admission system: A multi- agent based approach," *International Journal of Modern Education and Computer Science*, vol. 6, pp. 26–32, 5 2014.

[30] V. Vijaykumar, R. Chandrasekar, and T. Srinivasan, "An ant odor analysis approach to the ant colony optimization algorithm for data-aggregation in wireless sensor networks," in *2006 International Conference on Wireless Communications, Networking and Mobile Computing*, pp. 1–4, IEEE, 2006.

[31] A. M. Farid, *HoloMAS - Designing Multi-agent Systems for Resilient Engineering Systems*, pp. 3–8. Germany: Springer International Publishing, 8 2015.

[32] X. Qi and J. Chen, "Multi-agent coalition formation tactic based on grid," *Journal of Software*, vol. 5, pp. 36–43, 1 2010.

[33] W. Jiang, Q. Jiang, H. Li, and Y. Su, *Research on a Novel Distributed Multi Agent System Plan Method*, pp. 665–671. Germany: Springer Netherlands, 11 2011.

[34] J. Zhang, "Multi-agent-based material data acquisition," 4 2017.

[35] M. Hannebauer, J. Wendler, and E. Pagello, *Balancing Reactivity and Social Deliberation in Multi-Agent Systems - Balancing Reactivity and Social Deliberation in Multi-Agent Systems - A Short Guide to the Contributions*, pp. 3–8. Germany: Springer Berlin Heidelberg, 7 2001.

[36] T. Han, Z.-H. Guan, M. Chi, B. Hu, T. Li, and X.-H. Zhang, "Multi-formation control of nonlinear leader-following multi-agent systems.," *ISA transactions*, vol. 69, pp. 140–147, 5 2017.

[37] W. Dahhane, J. Berrich, and T. Bouchentouf, *Agent Story: An Agile Requirements Modeling Approach for Multi-agent Paradigm*, pp. 363–371. Germany: Springer International Publishing, 4 2016.

[38] T. Srinivasan, V. Vijaykumar, and R. Chandrasekar, "An auction based task allocation scheme for power-aware intrusion detection in wireless ad-hoc networks," in *2006 IFIP International Conference on Wireless and Optical Communications Networks*, pp. 5–pp, IEEE, 2006.

[39] H. J. LeBlanc and X. Koutsoukos, "Hscc - consensus in networked multi-agent systems with adversaries," in *Proceedings of the 14th international conference on Hybrid systems: computation and control*, pp. 281–290, ACM, 4 2011.

[40] L. He, G. Li, L. Xing, and Y. Chen, "An autonomous multi-sensor satellite system based on multi-agent blackboard model," *Eksploatacja i Niezawodnosc - Maintenance and Reliability*, vol. 19, pp. 447–458, 6 2017.

[41] S. W. Han and J. Kim, "Nossdav - a multi-agent-based toolkit for multimedia-oriented collaboration environments," in *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, pp. 111–112, ACM, 5 2008.

[42] L. Grivault, A. E. Fallah-Seghrouchni, and R. Girard-Claudon, *IDC - Multi-Agent System to Design Next Generation of Airborne Platform*, pp. 103–113. Germany: Springer International Publishing, 10 2017.

[43] R. Dai, X. Xu, C. F. Shi, and Y. Lu, "Research on multi-agent intelligent information retrieval model," *Applied Mechanics and Materials*, vol. 710, pp. 139–143, 1 2015.

[44] M. O. Anderson, C. W. Nielsen, M. D. McKay, D. C. Wadsworth, R. C. Hruska, and J. A. Koudelka, "Using multiple unmanned systems for a site security task," *SPIE Proceedings*, vol. 7350, pp. 735002–, 5 2009.

[45] J. Huang, "Adaptive distributed observer and the cooperative control of multi-agent systems," *Journal of Control and Decision*, vol. 4, pp. 1–11, 11 2016.

[46] W. Guo, "Leader-following consensus of the second-order multi-agent systems under directed topology," *ISA transactions*, vol. 65, pp. 116–124, 8 2016.

[47] R. Chandrasekar, V. Vijaykumar, and T. Srinivasan, "Probabilistic ant based clustering for distributed databases," in *2006 3rd International IEEE Conference Intelligent Systems*, pp. 538–545, IEEE, 2006.

[48] V. S. Melo, A. R. Panisson, and R. H. Bordini, "Meta-information and argumentation in multi-agent systems," *iSys - Brazilian Journal of Information Systems*, vol. 10, pp. 74–97, 9 2017.

[49] L. Wang, S. Sun, and C. Xia, "Finite-time stability of multi-agent system in disturbed environment," *Nonlinear Dynamics*, vol. 67, pp. 2009–2016, 7 2011.

[50] null KleinFlorian and null GieseHolger, "Analysis and design of physical and social contexts in multi-agent systems using uml," *ACM SIGSOFT Software Engineering Notes*, vol. 30, pp. 1–8, 5 2005.

[51] K. V. Hindriks and J. Dix, *Agent-Oriented Software Engineering - GOAL : A multi-agent programming language applied to an exploration game*, vol. 9783642544323. Springer Berlin Heidelberg, 3 2014.

[52] S. Aknine, *AIMDM - Exploiting Social Reasoning of Open Multi-agent Systems to Enhance Cooperation in Hospitals*, pp. 133–137. Germany: Springer Berlin Heidelberg, 6 1999.

[53] G. Miao and T. Li, "Mean square containment control problems of multi-agent systems under markov switching topologies," *Advances in Difference Equations*, vol. 2015, pp. 157–, 5 2015.

[54] E. A. Engel, *Dump Truck Fault's Short-Term Forecasting Based on the Multi-agent Adaptive Fuzzy Neuronet*, pp. 72–78. Germany: Springer International Publishing, 8 2017.

[55] J. Wang and M. Xin, "Multi-agent consensus algorithm with obstacle avoidance via optimal control approach," *International Journal of Control*, vol. 83, pp. 2606–2621, 12 2010.

[56] C. Ramachandran, R. Malik, X. Jin, J. Gao, K. Nahrstedt, and J. Han, "Videomule: a consensus learning approach to multi-label classification from noisy user-generated videos," in *Proceedings of the 17th ACM international conference on Multimedia*, pp. 721–724, 2009.

[57] C. Wang, K. Zhou, L. Li, and S. Yang, "Multi-agent simulation-based residential electricity pricing schemes design and user selection decision-making," *Natural Hazards*, vol. 90, pp. 1309–1327, 11 2017.

[58] P. Skobelev, E. V. Simonova, A. Zhilyaev, and V. Travin, *SOHOMA - Multi-Agent Planning of Spacecraft Group for Earth Remote Sensing.* Germany: Springer International Publishing, 3 2016.

[59] A. Vyavhare, V. Bhosale, M. Sawant, and F. Girkar, "Co-operative wireless intrusion detection system using mibs from snmp," *International Journal of Network Security & Its Applications*, vol. 4, pp. 147–154, 3 2012.