

CASCADE : CASCADED SCOPED COMMUNICATION FOR MULTI-AGENT RE-PLANNING IN DISRUPTED INDUSTRIAL ENVIRONMENTS

Mingjie Bi

State Key Laboratory of General Artificial Intelligence,
Beijing Institute for General Artificial Intelligence (BIGAI)
Beijing, China
bimingjie@bigai.ai

ABSTRACT

Industrial disruption replanning demands multi-agent coordination under strict latency and communication budgets, where disruptions propagate through tightly coupled physical dependencies and rapidly invalidate baseline schedules and commitments. Existing coordination schemes often treat communication as either effectively free (broadcast-style escalation) or fixed in advance (hand-tuned neighborhoods), both of which are brittle once the disruption footprint extends beyond a local region. We present CASCADE, a budgeted replanning mechanism that makes communication scope explicit and auditable rather than fixed or implicit. Each agent maintains an explicit knowledge base, solves role-conditioned local decision problems to revise commitments, and coordinates through lightweight contract primitives whose footprint expands only when local validation indicates that the current scope is insufficient. This design separates a unified agent substrate (Knowledge Base / Decision Manager / Communication Manager) from a scoped interaction layer that controls who is contacted, how far coordination propagates, and when escalation is triggered under explicit budgets. We evaluate CASCADE on disrupted manufacturing and supply-chain settings using unified diagnostics intended to test a mechanism-design claim—whether explicit scope control yields useful quality–latency–communication trade-offs and improved robustness under uncertainty—rather than to provide a complete algorithmic ranking.

1 INTRODUCTION

Disruptions are routine in industrial systems (Fu et al., 2025). Machine breakdowns, capacity losses, lead-time inflation, and demand shocks can invalidate pre-planned schedules, and material flows within minutes (Tang, 2006). Because these effects propagate through tightly coupled temporal, capacity, and flow dependencies, effective response requires *distributed* replanning across heterogeneous entities—resources, process stages, suppliers, carriers, and distributors. Even localized shocks can induce system-wide losses through networked dependencies, motivating resilience-oriented decision making under time pressure (Snyder & Shen, 2016; Ivanov, 2019). A central challenge is coordination under limited communication. Broadcast-style negotiation may recover feasibility, but it floods the network and consumes scarce response time. Strictly local coordination is cheap, but it can fail silently when disruption effects cross neighborhood boundaries. What is missing is a replanning mechanism that explicitly controls the communication footprint: *who* is involved, *how far* coordination propagates, and *when* escalation is warranted.

Industrial control context and positioning. Distributed scheduling and manufacturing control must balance solution quality against negotiation and information-exchange overhead (Fu et al., 2021). More broadly, recent work in cyber-physical manufacturing and industrial multi-agent control has moved toward modular learning-and-control architectures, dynamic resource reallocation, and resilient distributed decision-making under disruptions (Kovalenko et al., 2022; Bi et al., 2021). Related supply-chain studies have likewise developed model-based and network-aware multi-agent approaches for agile disruption response (Bi et al., 2022). These works establish the need for

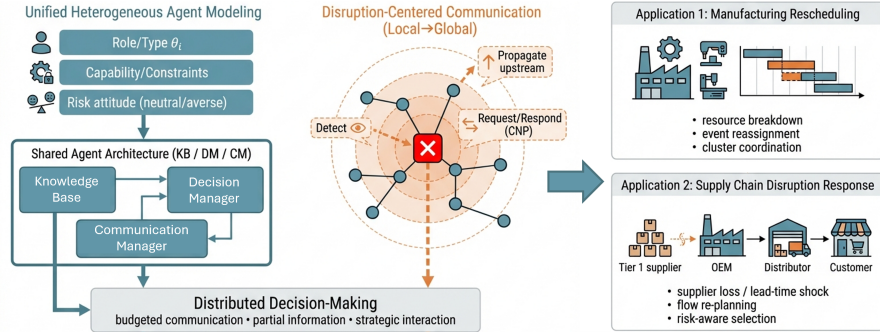


Figure 1: A unified agent substrate (KB/DM/CM) supports heterogeneous local modeling, while gate-triggered scoped communication escalates from local to broader coordination only when needed.

decentralized commitment revision in industrial environments, but they still leave the communication scope largely as an implementation choice rather than an explicit mechanism variable.

Communication-constrained coordination. Related communities increasingly treat communication as a limited decision resource rather than a free side channel. In communication-constrained MARL, recent methods learn when and how to exchange compressed or context-dependent messages under bandwidth limits, as illustrated by context-aware protocols such as CACOM (Li & Zhang, 2024). In parallel, task-oriented communication studies optimize messages for downstream control utility under explicit bit budgets, rather than for faithful reconstruction, including task-oriented data compression for multi-agent coordination (Mostaani et al., 2024). These lines strongly motivate explicit, utility-aware communication, but they are rarely studied in disruption-centered industrial replanning, where commitments, feasibility, and escalation budgets are tightly coupled.

Taken together, the literature points to a missing mechanism layer for industrial disruption response: communication should be *explicitly budgeted and interpretable*, its *scope should adapt* to the evolving disruption footprint, and evaluation should expose systematic *quality–latency–communication* trade-offs under tight operational budgets. However, industrial disruption replanning differs from standard scheduling or generic distributed optimization in two under-addressed ways: (i) feasibility depends on role-dependent commitments under coupled temporal, capacity, and flow constraints, often under uncertainty; and (ii) communication is itself operationally budgeted, so naive escalation can dominate response time. This leaves a core question: *how should disruption replanning regulate communication scope as an explicit control variable while still making progress when local repair fails?*

Contributions. CASCADE answers this question by treating communication scope as a decision variable. Agents attempt local resolution first and expand coordination only when validation gates indicate that the current scope is insufficient. The contributions are: (i) a disruption-replanning formulation over coupled physical and communication graphs under explicit communication and time budgets; (ii) CASCADE itself—a unified agent substrate (Knowledge Base / Decision Manager / Communication Manager) together with contract-style negotiation and gate-triggered scoped propagation; and (iii) a unified diagnostic view that exposes quality–latency–communication trade-offs, propagation footprint, and robustness across disrupted manufacturing and supply-chain environments.

2 PROBLEM SETTING

A disrupted industrial environment is represented by two coupled graphs. The *physical dependency graph* $G_p = (V, E_p)$ captures material, operation, and flow dependencies. Each node $i \in V$ is an entity such as a resource, process stage, supplier, transporter, or distributor, and each typed edge $(i \rightarrow j, k) \in E_p$ denotes a commitment to provide item/type k with quantity and timing constraints. In parallel, agents coordinate over an *agent communication graph* $G_a = (A, L)$, where A indexes decision-making entities and $(a \leftrightarrow b) \in L$ indicates that a and b can exchange messages directly.

A baseline plan π_0 encodes local decisions and cross-agent commitments. In unified form, a plan π can be written as

$$\pi = \{(c_{ij}^k, \tau_{ij}^k)\}_{(i \rightarrow j, k) \in E_p} \cup \{u_a\}_{a \in A}, \quad (1)$$

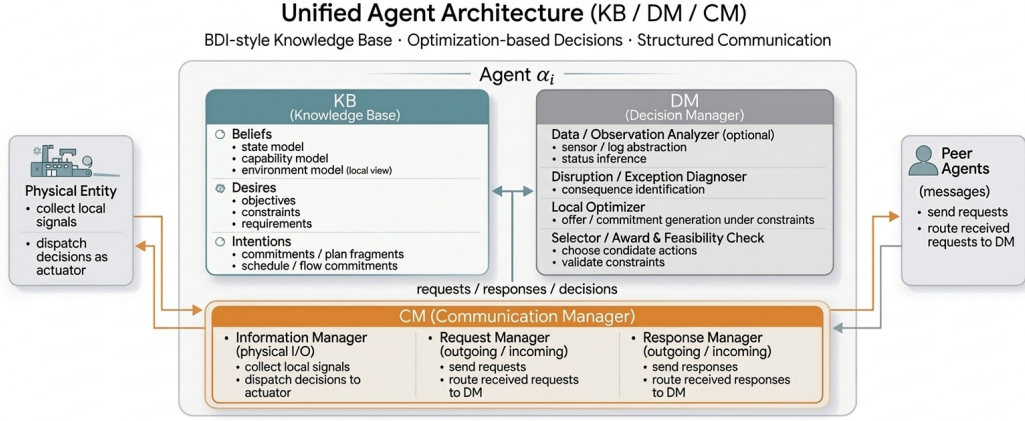


Figure 2: **Unified agent architecture (KB/DM/CM)**. KB maintains local beliefs, priorities, intentions, and uncertainty; DM diagnoses disruptions and computes local commitments; CM routes physical I/O and peer messages.

where (e_{ij}^k, τ_{ij}^k) are committed transfers and u_a denotes agent-local allocations. A disruption δ modifies local feasibility by changing capacities, lead times, costs, or demands, i.e., $\mathcal{C}_a \leftarrow \mathcal{C}_a(\delta)$ for affected agents, and may induce unmet needs that propagate through G_p .

The goal is to compute a revised plan π that restores feasibility while balancing service loss, operational cost, plan deviation, and communication:

$$J(\pi) = J_{\text{service}}(\pi) + \lambda_c J_{\text{cost}}(\pi) + \lambda_\Delta d(\pi, \pi_0) + \lambda_{\text{msg}} M(\pi), \quad (2)$$

where $d(\pi, \pi_0)$ measures deviation from the baseline plan and $M(\pi)$ measures communication on G_a . Replanning is subject to hard budgets: message B_{msg} or round limits B_{round} and a wall-clock deadline B_{time} . The resulting mechanism problem is therefore not only to restore feasibility, but to do so while explicitly regulating communication footprint under operational constraints.

3 CASCADE : ARCHITECTURE AND MECHANISM

This section separates *what lives inside an agent* from *how agents interact*. Figure 2 specifies the unified KB/DM/CM substrate inside an individual agent. Figure 3 specifies the CASCADE mechanism layered on top of that substrate: scoped negotiation, validation, escalation, and propagation on the coupled graphs. KB/DM/CM is the internal representation-and-decision blueprint; CASCADE is the interaction mechanism that controls communication footprint under explicit budgets. The unified substrate improves modularity and interoperability, but the paper’s main mechanism contribution is the scoped interaction policy rather than the substrate alone.

3.1 AGENT ARCHITECTURE: KB/DM/CM

Each agent implements three internal modules: (i) a **Knowledge Base (KB)** for structured local beliefs, priorities, intentions, and uncertainty; (ii) a **Decision Manager (DM)** for role-conditioned local commitment reasoning; and (iii) a **Communication Manager (CM)** as an interface layer for physical I/O and peer messaging. This factorization supports heterogeneous local models while preserving a common interaction interface.

Knowledge Base. Each agent $a \in A$ maintains

$$\text{KB}_a = \{\mathcal{B}_a, \mathcal{D}_a, \mathcal{I}_a\}, \quad \mathcal{B}_a = \{\mathcal{X}_a, \mathcal{C}_a, \mathcal{E}_a, \mathcal{R}_a\}. \quad (3)$$

\mathcal{X}_a stores observable local state; $\mathcal{C}_a(\delta)$ stores local feasibility constraints and capability envelopes under disruption δ ; \mathcal{E}_a stores the local view of G_p and G_a ; and \mathcal{R}_a stores uncertainty and risk attitude. \mathcal{D}_a summarizes local priorities such as service, cost, and robustness preference. \mathcal{I}_a encodes the current local plan fragment induced by π and its relation to π_0 .

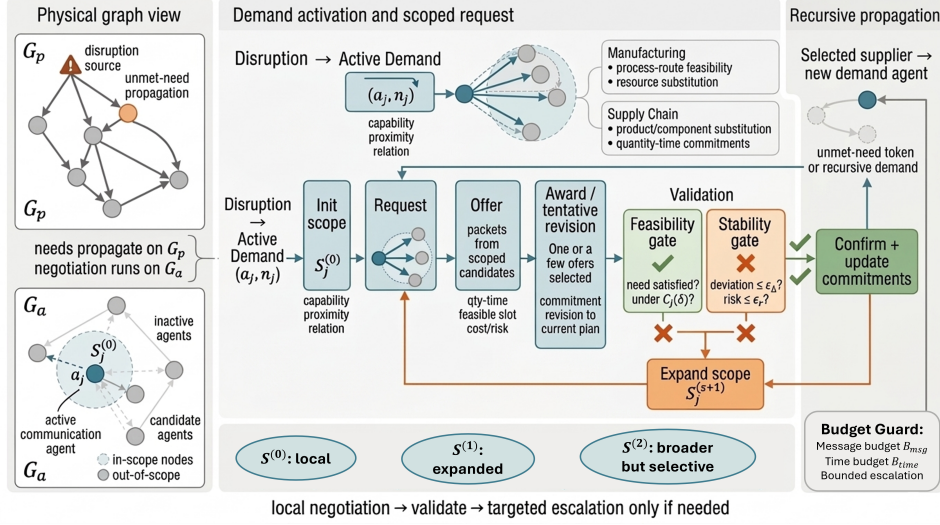


Figure 3: **Gate-triggered scoped communication propagation.** Unmet needs propagate along G_p while coordination is executed on G_a through contract primitives; feasibility and stability gates trigger targeted scope expansion and role transitions without global flooding.

Module Examples: 1) each agent maintains a local capability set $\mathcal{M}_a = \{m_{a,1}, \dots, m_{a,q}\}$, where each capability represents an executable local transformation, transport, or supply commitment. In supply-chain settings, a capability can be instantiated as (o, k) , i.e., the ability to perform operation o for product k . In manufacturing settings, capabilities could be process feasibility, resource class, and local reachability, e.g., supported process types for machines and reachable cells/buffers for mobile robots. 2) The uncertainty and risk component can be represented explicitly as $\mathcal{R}_a = \{\mathcal{U}_a, \rho_a\}$, where \mathcal{U}_a denotes the local uncertainty model and ρ_a denotes the local risk functional or risk attitude. Depending on the setting, \mathcal{U}_a may be represented by stochastic operation times, disruption-dependent response uncertainty, or scenario samples. The risk functional ρ_a can be instantiated as expectation, conservative worst-case evaluation, or other tail-sensitive functionals such as CVaR.

Decision Manager. The DM manages various role-conditioned local decision problems, such as request evaluation, offer generation, award selection, or local repair. In a supplier/resource role, the DM produces feasible commitments under local constraints and uncertainty; in a demand role, it selects or mixes received offers to satisfy unmet needs while controlling local deviation and risk.

A local decision-making module may be instantiated as an optimization model, a trained policy network, or a language-guided decision module, provided that it returns commitments in the same typed format. In the current study, the instantiated DMs are optimization-centered, with uncertainty- and risk-aware local decision models in both manufacturing and supply-chain settings. LLM-based agent architectures provide another compatible realization of such modular decision layers.

Communication Manager. The CM routes incoming and outgoing peer messages and mediates between the DM and the physical interface. It collects local signals, dispatches DM decisions to actuators, and standardizes peer interaction. The question of *when* and *how far* communication should expand is handled by the CASCADE mechanism below.

3.2 CASCADE MECHANISM: GATE-TRIGGERED SCOPED PROPAGATION

CASCADE is a closed-loop replanning mechanism that treats communication scope as an explicit control variable. Mechanism execution couples the physical dependency graph G_p (where unmet needs propagate) and the communication graph G_a (where negotiation occurs). A replanning episode proceeds as a sequence of scoped negotiations and commitment revisions; scope expands *only* when validation gates fail. Algorithm 1 summarizes CASCADE as a budgeted loop over (i) scoped negotiation on G_a , (ii) tentative local commitment revision, (iii) gate-based validation and escalation, and (iv) role-transition-driven propagation on G_p after confirmation.

Algorithm 1 CASCADE : gate-triggered scoped communication for disruption replanning

```

1: Input: baseline plan  $\pi_0$ , disruption  $\delta$ , graphs  $G_p, G_a$ , budgets  $B_{\text{msg}}, B_{\text{time}}$ 
2: Output: revised plan  $\pi$ 
3:  $\pi \leftarrow \pi_0$ ;  $\mathcal{D} \leftarrow \text{INFERDEMANDS}(\delta, \pi_0, G_p)$ 
4: while  $\mathcal{D} \neq \emptyset$  and budgets not exhausted do
5:   Select an active demand  $(a_j, n_j) \in \mathcal{D}$ 
6:    $s \leftarrow 0$ ; initialize scope  $\mathcal{S}_j^{(0)}$  via Eq. (5)
7:   repeat
8:     (Request) Send REQUEST( $n_j$ ) to candidates in  $\mathcal{S}_j^{(s)}$ ; collect offers
9:     (Award) Construct a tentative revised plan  $\pi^{(s)}$  from the received offers
10:    Evaluate  $g_{\text{feas}}$  and  $g_{\text{stab}}$  via Eqs. (7)–(8)
11:    if both gates pass then
12:      (Confirm) Send award/confirmation to selected suppliers; commit and set  $\pi \leftarrow \pi^{(s)}$ 
13:      (Role transition) If commitments induce upstream needs, add new demands to  $\mathcal{D}$ 
14:    else
15:      Expand scope via Eq. (6);  $s \leftarrow s + 1$ 
16:    end if
17:  until both gates pass OR expansion budget exhausted
18:  Remove  $(a_j, n_j)$  from  $\mathcal{D}$  if satisfied; otherwise terminate if budgets are exhausted
19: end while
20: return  $\pi$ 

```

Contract primitives. CASCADE uses compact, auditable primitives: REQUEST, OFFER, AWARD, and CONFIRM. These primitives make message usage and negotiation rounds explicit, budgeted resources rather than hidden side effects of local search.

Scope policy. Given a demand (a_j, n_j) , the mechanism initializes a bounded scope $\mathcal{S}_j^{(0)} \subseteq A$ on G_a using capability match, structural proximity, and relation similarity:

$$w_{j \rightarrow b} = \alpha \text{sim}_{\text{cap}}(n_j, \text{KB}_b) + \beta \text{prox}_{G_a}(a_j, b) + \gamma \text{sim}_{\text{rel}}(a_j, b), \quad (4)$$

$$\mathcal{S}_j^{(0)} = \{b : \text{dist}_{G_a}(a_j, b) \leq h_0\} \cap \text{TopK}(w_{j \rightarrow \cdot}, K_0). \quad (5)$$

If escalation is needed, scope expands according to a schedule $(h_s, K_s)_{s \geq 0}$:

$$\mathcal{S}_j^{(s+1)} = \mathcal{S}_j^{(s)} \cup \left(\{b : \text{dist}_{G_a}(a_j, b) \leq h_{s+1}\} \cap \text{TopK}(w_{j \rightarrow \cdot}, K_{s+1}) \right). \quad (6)$$

The main text keeps this formulation domain-agnostic; concrete environment-specific instantiations are deferred to the appendix. In manufacturing, capability match corresponds to process-route feasibility and local resource substitutability, while proximity reflects resource adjacency and transport reachability in the production system. In the supply chain, capability match corresponds to feasible product/component substitution and commitment compatibility, while proximity reflects dependency closeness on the production–logistics network. Relation similarity in both settings captures whether two agents occupy structurally or historically compatible coordination roles, allowing the same score to prioritize plausible local substitutes without collapsing immediately to broad broadcast.

Validation gates. Let $\pi^{(s)}$ denote the tentative plan after negotiation within scope $\mathcal{S}_j^{(s)}$. Two gates regulate acceptance and escalation:

$$g_{\text{feas}}(a_j, n_j, \pi^{(s)}, \delta) = \mathbf{1} \left[\exists u_j \text{ s.t. } (u_j, \pi^{(s)}) \in \mathcal{C}_j(\delta) \wedge n_j \text{ is satisfied} \right], \quad (7)$$

$$g_{\text{stab}}(a_j, \pi^{(s)}, \pi_0) = \mathbf{1} \left[d_j(\pi^{(s)}, \pi_0) \leq \varepsilon_\Delta \wedge \rho_j(\pi^{(s)}) \leq \varepsilon_r \right]. \quad (8)$$

The feasibility gate checks whether the current scope can actually satisfy the induced need under disruption-aware constraints. The stability gate checks whether the induced revision remains within acceptable deviation and risk. Confirmation is final after both gates pass; before that, offers are treated as tentative and may be discarded when scope expands via Eq. (6).

Role transition and propagation. If an awarded offer induces upstream needs (e.g., additional materials, transport, or substitute capacity), the selected supplier becomes a new demand agent. This

Table 1: Compact summary of the two experimental settings used in the main text. Detailed environment definitions, disruption-generation rules, and metric formulas are provided in the appendix.

Setting	System and disruption regime	Coordination	Main reported signals / protocol
Manufacturing	Modified Intel Mini-Fab; two product types with fixed process routes; disruptions are machine breakdowns with operation-time uncertainty.	Centralized vs. CASCADE	Cycle time, communication M , runtime T ; incident and risk statistics in the disruption-aware case; 5 trials.
Supply chain	Automotive cockpit supply-chain network. Two disruption regimes: (i) single supplier-loss on the nominal flow plan; (ii) lead-time disruption with different added time attributes.	Centralized vs. CASCADE ; neutral vs. averse	Overage cost O , communication M , and footprint proxies (N_c, N_a) ; delayed-product distribution and cost/lateness (C_{dm}, L_{dm}, J_{dm}) .

role transition drives propagation on G_p , while communication on G_a remains explicitly scoped. The mechanism, therefore, implements a bounded loop: local negotiation \rightarrow validation \rightarrow targeted escalation only if needed.

4 EXPERIMENTS

This section evaluates CASCADE in two industrial domains: (i) manufacturing disruption replanning (rescheduling and resource reallocation), and (ii) supply-chain disruption response (flow and commitment revision). Across both domains, agents operate on a physical dependency graph G_p where unmet needs propagate and exchange messages on a communication graph G_a where negotiation is executed. A disruption δ changes local capabilities $C_a(\delta)$ and may induce unmet needs that propagate through dependencies.

4.1 ENVIRONMENTS AND DISRUPTION REGIMES

Manufacturing. We study a tightly coupled production environment with 20 machine agents and 6 mobile-robot agents (Bi et al., 2024b). Disruptions such as machine breakdowns and availability changes invalidate a baseline production schedule, forcing coordinated operation reassignment under feasibility constraints. The manufacturing setting additionally includes uncertainty and a risk-aware variant used for robustness analysis.

Supply chain. We study an interdependent automotive-cockpit production–logistics network (Bi et al., 2023; 2024a). Two disruption regimes are reported on this shared network. The first is a network-structured supplier-loss setting, used to analyze how structural attributes affect propagation footprint and communication. The second is a stochastic lead-time disruption setting with heterogeneous risk attitudes, used to analyze robustness under uncertainty. Exact system specifications, disruption protocols, and graph definitions are deferred to the appendix.

4.2 EXPERIMENTAL PROTOCOL AND METRICS

We compare **Centralized** replanning (global coordination) against **Distributed** replanning implemented by CASCADE . Table 1 summarizes the evaluation regimes and the scale at which each diagnostic is constructed.

Budgets and accounting. Communication is measured by the number of point-to-point message exchanges M . Each REQUEST, OFFER, AWARD, and CONFIRM sent to one recipient counts as one exchange; multicast messages are counted per recipient. Centralized coordination is accounted for separately using a domain-specific aggregation of information-collection and update-notification steps, whereas distributed coordination counts only the explicitly executed negotiation messages needed to produce the revised plan. Latency T is measured as replanning runtime under the same implementation setting. All communication results should therefore be read as message-exchange accounting rather than bit-level or semantic-content accounting. The detailed domain-specific protocol, scenario counts, and accounting conventions are provided in the appendix.

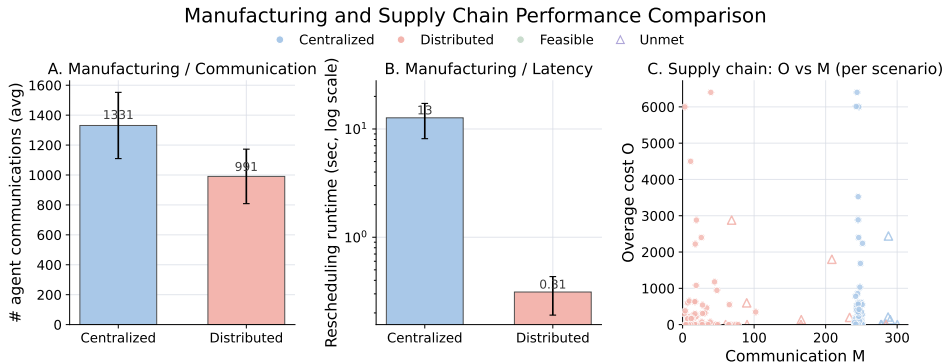


Figure 4: R1: Quality–latency–communication frontier. The axes are domain-specific (cycle-time-based recovery proxy in manufacturing; overage cost with feasibility labels in supply chain), but the shared question is how much recovery quality is achieved per unit communication and latency.

Implementation contingency. In the manufacturing implementation, when the local candidate set is empty, the environment permits escalation to a central controller or human manager as an external fallback. We report this behavior explicitly in the appendix; it is an implementation contingency rather than the intended scoped interaction path and is not treated as a mechanism contribution.

Quality and feasibility. In manufacturing, the primary quality proxy is disruption-recovery quality measured through cycle-time behavior together with incident and risk statistics. In supply chain, quality is measured by overage cost O and feasibility status, with additional lateness and cost–delay metrics in the stochastic setting. Formal definitions are collected in the appendix.

4.3 UNIFIED DIAGNOSTICS

We report three complementary diagnostics that expose mechanism-level trade-offs rather than only task-level averages. The three views are deliberately chosen to answer different questions raised by budgeted industrial coordination. R1 asks whether better recovery quality can be obtained without paying the full latency and communication cost of broad coordination. R2 asks whether the induced communication footprint follows interpretable structural properties of the disrupted network rather than expanding arbitrarily. R3 asks whether risk-aware acceptance and stabilization policies improve tail behavior under stochastic disruption, rather than merely shifting average nominal cost.

R1: Quality–latency–communication frontier. Figure 4 summarizes the relationship between task quality, communication M , and latency T . In manufacturing, the quality axis is a cycle-time-based recovery proxy; in supply chain, it is scenario-level overage cost O together with feasibility labels. The purpose of R1 is not to force identical units across domains, but to ask the same mechanism question in both settings: how much recovery quality is achieved per unit communication and latency. In manufacturing, CASCADE reduces replanning latency substantially while requiring fewer communications than centralized coordination. In the supply chain, the O – M scatter exposes scenario-level trade-offs and feasible versus unmet cases.

R2: Propagation footprint and network structure (supply chain). Figure 5 links propagation footprint to structural attributes in the environment model. We instantiate the footprint using proxy counts such as commitment changes N_c , added commitments N_a , communication M , and quality impact O . These proxies are appropriate because they capture how far a local disruption response must spread in order to restore feasibility: more changed commitments, more activated entities, and more messages all indicate deeper cascade depth. The four groups (LL/LH/HL/HH) correspond to a low/high stratification of connectivity and redundancy. This diagnostic explains *when* deeper cascades become necessary: tightly coupled or low-redundancy regions induce larger footprint and communication, while redundant neighborhoods admit earlier termination.

R3: Robustness under uncertainty via risk-aware stabilization. Figure 6 summarizes robustness under stochastic disruptions. In manufacturing, robustness is reflected by incident counts and schedule-risk summaries with and without risk assessment. In the supply chain, robustness is reflected by




Figure 5: R2: Propagation footprint vs. network structure. Footprint proxies (N_c , N_a), communication M , and quality impact O are stratified by network-attribute groups (LL/LH/HL/HH). The grouped analysis is restricted to the 73 supplier-loss scenarios in which both centralized and distributed recovery are feasible; group sizes are LL= 53, LH= 16, HL= 2, and HH= 13.

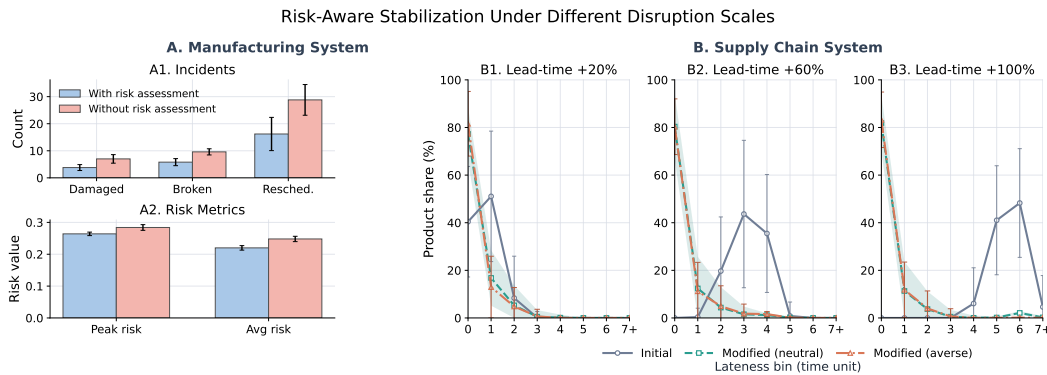


Figure 6: R3: Robustness under uncertainty. The manufacturing panel summarizes five trials, while the supply-chain panel reports the fixed disrupted-subnetwork study. Across both settings, risk-aware stabilization improves tail behavior and reduces brittle recovery under stochastic disruptions.

delayed-product distributions and cost–lateness trade-offs under increasing disruption intensity. The common mechanism interpretation is that risk-aware stabilization changes the acceptance policy for local commitments, improving tail behavior even when the nominal cost may increase slightly.

5 DISCUSSION

Why scoped cascades fit industrial replanning. Industrial replanning often crosses organizational boundaries and must obey strict latency budgets. In such settings, communication is not a free software detail; it is an operational resource. By exposing scope schedules and gate thresholds as explicit knobs, CASCADE supports auditable escalation under practical constraints such as privacy, supplier contracts, and bounded negotiation windows.

Where the mechanism matters most. The diagnostics suggest that scoped cascades are most valuable in the intermediate regime where disruptions can propagate beyond immediate neighbors but do not require full-network mobilization in every case. In tightly coupled or low-redundancy regions, local fixes often fail, and deeper propagation becomes necessary. In more redundant neighborhoods, the same mechanism can terminate early and avoid unnecessary communication. Conversely, when the feasible substitute set is already contained in a very small neighborhood, fixed local coordination may already suffice; when almost the entire network must be mobilized, the mechanism approaches broad escalation. This boundary is important: CASCADE matters most when *adaptive* scope control is the deciding factor.

Table 2: R3 robustness summary across the two industrial domains. Panel A reports the effect of risk assessment in manufacturing recovery (five trials). Panel B reports the effect of risk attitude in the supply-chain lead-time disruption setting. Negative Δ indicates improvement for loss/risk metrics; negative ΔJ indicates that the Averse policy improves the objective relative to Neutral.

Panel A: Manufacturing robustness				
Metric	With risk	Without	Δ	$\Delta\%$
# damaged products	3.80±1.10	7.00±2.12	-3.20	-45.7%
# broken machines	4.52±1.23	7.84±2.49	-3.32	-42.3%
# rescheduled processes	4.05±1.19	7.20±2.47	-3.15	-43.8%
Peak risk value	0.02±0.04	0.03±0.04	-0.01	-33.3%
Average risk value	0.03±0.02	0.06±0.04	-0.03	-50.0%
Panel B: Supply-chain robustness				
Rate	Risk	C_{dm}	L_{dm}	$J_{dm} / \Delta J$
Lead time +0%	Neutral	31,120	0	31,120 / -
Lead time +20%	Neutral	28,386	1	128,386 / -
	Averse	28,416	1	128,416 / +30 (+0.02%)
Lead time +60%	Neutral	28,386	6	628,386 / -
	Averse	29,342	5	529,342 / -99,044 (-15.76%)
Lead time +100%	Neutral	28,792	8	828,792 / -
	Averse	29,208	3	329,208 / -499,584 (-60.28%)

How the current evidence should be read. The present results support a mechanism-design claim rather than a complete algorithmic ranking. The diagnostics are intended to isolate the value of explicit scope control and gate-triggered escalation across two industrial domains, not yet to substitute for a full benchmark study against all distributed coordination baselines. Fixed-scope distributed negotiation and ungated schedule-based expansion are the most important missing comparators for testing whether adaptive gated scope is genuinely necessary, while stronger DCOP-style baselines would help separate scoped control from decentralized replanning more broadly. Likewise, the current study exposes scope schedules and gate thresholds as auditable knobs, but does not yet provide a systematic sensitivity sweep across (h_s, K_s) and $(\varepsilon_\Delta, \varepsilon_r)$ to test deployment stability.

Remaining limitations. First, the evaluation spans two representative industrial domains, but it is not yet a fully standardized public benchmark suite. Second, communication is accounted for at the level of message exchanges rather than bit-level or content-level cost; integrating task-level compression and richer network constraints remains an important extension. Third, the stochastic supply-chain analysis focuses on a single disrupted supplier and a limited heterogeneous-risk setting; broader topological regimes and richer combinations of conflicting agent risk attitudes remain open.

6 CONCLUSION

CASCADE is a budgeted industrial replanning mechanism that makes the communication footprint explicit and controllable. By combining a unified KB/DM/CM agent substrate with contract-style negotiation and gate-triggered scope expansion, it enables stepwise, auditable coordination without immediate global flooding. Across manufacturing rescheduling and supply-chain disruption response, the unified diagnostics provide mechanism-level evidence for a common quality-latency-communication frontier, relate propagation footprint to network structure, and show that risk-aware stabilization can improve robustness under stochastic disruptions.

The present results should be read as support for a design principle rather than a complete benchmark verdict. The main lesson is that communication scope should not be treated as fixed or free in industrial disruption response. It should be made explicit, validated before confirmation, and escalated only when the current coordination neighborhood is insufficient.

REFERENCES

- Mingjie Bi, Ilya Kovalenko, Dawn M Tilbury, and Kira Barton. Dynamic resource allocation using multi-agent control for manufacturing systems. *IFAC-PapersOnLine*, 54(20):488–494, 2021.
- Mingjie Bi, Gongyu Chen, Dawn M Tilbury, Siqian Shen, and Kira Barton. A model-based multi-agent framework to enable an agile response to supply chain disruptions. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pp. 235–241. IEEE, 2022.
- Mingjie Bi, Dawn M Tilbury, Siqian Shen, and Kira Barton. A distributed approach for agile supply chain decision-making based on network attributes. *IEEE Transactions on Automation Science and Engineering*, 21(3):2223–2236, 2023.
- Mingjie Bi, Juan-Alberto Estrada-Garcia, Dawn M Tilbury, Siqian Shen, and Kira Barton. Heterogeneous risk management using a multi-agent framework for supply chain disruption response. *IEEE Robotics and Automation Letters*, 9(6):5126–5133, 2024a.
- Mingjie Bi, Ilya Kovalenko, Dawn M Tilbury, and Kira Barton. Dynamic distributed decision-making for resilient resource reallocation in disrupted manufacturing systems. *International Journal of Production Research*, 62(5):1737–1757, 2024b.
- Bo Fu, Mingjie Bi, Shota Umeda, Takahiro Nakano, Youichi Nonaka, Quan Zhou, Takaharu Matsui, Dawn M Tilbury, and Kira Barton. Digital twin-based smart manufacturing: Dynamic line reconfiguration for disturbance handling. *IEEE Transactions on Automation Science and Engineering*, 2025.
- Yaping Fu, Yushuang Hou, Zifan Wang, Xinwei Wu, Kaizhou Gao, and Ling Wang. Distributed scheduling problems in intelligent manufacturing systems. *Tsinghua Science and Technology*, 2021. doi: 10.26599/TST.2021.9010009.
- Dmitry Ivanov. Disruption tails and revival policies: A simulation analysis of supply chain design and production-ordering systems in the recovery and post-disruption periods. *Computers & Industrial Engineering*, 127:558–570, 2019.
- Ilya Kovalenko, James Moyne, Mingjie Bi, Efe C Balta, Wenyuan Ma, Yassine Qamsane, Xiao Zhu, Z Morley Mao, Dawn M Tilbury, and Kira Barton. Toward an automated learning control architecture for cyber-physical manufacturing systems. *IEEE Access*, 10:38755–38773, 2022.
- Xinran Li and Jun Zhang. Context-aware communication for multi-agent reinforcement learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, pp. 1156–1164, 2024. doi: 10.5555/3635637.3662972. URL <https://www.ifaamas.org/Proceedings/aamas2024/pdfs/p1156.pdf>. IFAAMAS, Auckland, New Zealand.
- Arsham Mostaani, Thang X. Vu, Hamed Habibi, Symeon Chatzinotas, and Björn Ottersten. Task-oriented communication design at scale. *IEEE Transactions on Communications*, 73(1):378–393, 2024. doi: 10.1109/TCOMM.2024.3416898. URL https://orbilu.uni.lu/bitstream/10993/62307/1/Task-Oriented_Communication_Design_at_Scale.pdf.
- Lawrence V. Snyder and Zuo-Jun Max Shen. Or/ms models for supply chain disruptions: A review. *IIE Transactions*, 48(2):89–109, 2016. doi: 10.1080/0740817X.2015.1067735.
- Christopher S. Tang. Perspectives in supply chain risk management. *International Journal of Production Economics*, 103(2):451–488, 2006. doi: 10.1016/S0925-5273(06)00040-5.

A FULL MECHANISM SPECIFICATION

A.1 UNIFIED AGENT BLUEPRINT

Each agent implements three internal modules: a *Knowledge Base* (KB), a *Decision Manager* (DM), and a *Communication Manager* (CM). The KB stores structured local beliefs, desires, and intentions,

$$\text{KB}_a = \{B_a, D_a, I_a\}, \quad B_a = \{X_a, C_a, E_a, R_a\}, \quad (9)$$

where X_a denotes observable local state, $C_a(\delta)$ denotes local feasibility constraints and capability envelopes under disruption δ , E_a denotes the local view of the physical dependency graph and the communication graph, and R_a stores uncertainty and risk attitude. Desires D_a summarize local priorities such as service level, cost sensitivity, and robustness preference, while intentions I_a encode the current local plan fragment induced by the baseline plan π_0 and the revised plan π .

The DM solves role-conditioned local optimization problems. When an agent acts as a supplier/resource, it generates offers by solving a constrained local problem over feasible commitments; when it acts as a demand agent, it selects offers to satisfy unmet needs while controlling deviation and risk. The CM provides the interface to the physical layer and to peer agents. It collects local signals, dispatches DM decisions to actuators, and routes incoming and outgoing peer messages.

A.2 OFFER GENERATION AND AWARD SELECTION

Given an incoming request, a supplier-side agent solves a local risk-aware commitment problem of the form

$$\min_{z \in Z_i} \rho_{\omega \sim R_i} [J_i(z; \omega)] \quad \text{s.t.} \quad z \in C_i(\delta), \quad (10)$$

where z parameterizes the proposed commitment (for example, quantity–time commitments and auxiliary local actions), and ρ is a risk functional such as expectation or CVaR. The resulting OFFER includes feasibility and robustness summaries.

After collecting offers $O_j^{(s)}$ from the current scope $S_j^{(s)}$, the demand-side agent solves an award problem

$$\min_y \sum_{b \in S_j^{(s)}} y_{jb} J_{jb}(o_{jb}) + \lambda_\Delta d_j(\pi^{(s)}(y), \pi_0) + \lambda_r \rho_j(\pi^{(s)}(y)) \quad (11)$$

subject to

$$\text{NEEDSAT}(n_j, y, O_j^{(s)}) = 1, \quad \text{LOCALFEAS}(y; C_j(\delta)) = 1, \quad (12)$$

with $y_{jb} \in \{0, 1\}$ for single-award selection or $y_{jb} \in [0, 1]$ for splittable commitments. Here $J_{jb}(o_{jb})$ is the local evaluated utility/cost of selecting offer o_{jb} , $d_j(\cdot, \cdot)$ is local plan deviation relative to π_0 , and $\rho_j(\cdot)$ is a robustness/risk score.

A.3 SCOPED COMMUNICATION POLICY

CASCADE treats communication scope as an explicit control variable on the communication graph $G_a = (A, L)$. For a demand (a_j, n_j) , the initial scope is selected using a weighted score

$$w_{j \rightarrow b} = \alpha \text{sim}_{\text{cap}}(n_j, \text{KB}_b) + \beta \text{prox}_{G_a}(a_j, b) + \gamma \text{sim}_{\text{rel}}(a_j, b), \quad (13)$$

where the three terms respectively capture capability match, structural proximity, and relationship similarity. The initial scope is

$$S_j^{(0)} = \{b : \text{dist}_{G_a}(a_j, b) \leq h_0\} \cap \text{TopK}(w_{j \rightarrow \cdot}, K_0). \quad (14)$$

If escalation is needed, the scope expands according to a schedule $(h_s, K_s)_{s \geq 0}$:

$$S_j^{(s+1)} = S_j^{(s)} \cup \left(\{b : \text{dist}_{G_a}(a_j, b) \leq h_{s+1}\} \cap \text{TopK}(w_{j \rightarrow \cdot}, K_{s+1}) \right). \quad (15)$$

Operational constraints such as organizational boundaries, privacy restrictions, or supplier-policy masks can be enforced by removing candidate agents or communication edges from G_a before ranking.

A.4 VALIDATION GATES, PROPAGATION, AND STOPPING

For the tentative plan $\pi^{(s)}$ generated from the current scope $S_j^{(s)}$, CASCADE evaluates two gates:

$$g_{\text{feas}}(a_j, n_j, \pi^{(s)}, \delta) = 1 \iff \exists u_j \text{ s.t. } (u_j, \pi^{(s)}) \in C_j(\delta) \wedge n_j \text{ is satisfied}, \quad (16)$$

$$g_{\text{stab}}(a_j, \pi^{(s)}, \pi_0) = 1 \iff d_j(\pi^{(s)}, \pi_0) \leq \varepsilon_\Delta \wedge \rho_j(\pi^{(s)}) \leq \varepsilon_r. \quad (17)$$

If either gate fails, the mechanism expands scope via Eq. (15). If an awarded offer induces upstream needs (e.g., additional materials, transport, or substitute capacity), the selected supplier becomes a new demand agent. In this way, unmet needs propagate on the physical dependency graph G_p , while negotiation remains explicitly scoped on G_a .

A replanning episode terminates when all active demands are satisfied, when the expansion budget is exhausted, or when hard message/time budgets are exhausted. In the current paper, the operational budgets are the message budget B_{msg} and wall-clock deadline B_{time} .

B MANUFACTURING ENVIRONMENT AND DISRUPTION MODEL

B.1 FACILITY LAYOUT AND AGENTIZATION

The manufacturing environment is a modified Intel Mini-Fab instance implemented in Repast Symphony. The simulated facility contains two infinite-capacity buffers (`ENTRY` and `EXIT`), 20 machine agents, and 6 mobile-robot agents. Machines support six process types, P1–P6. Operation costs are measured in RepastS ticks and range from 110 to 200 depending on the process and resource.

Each machine agent is associated with a subset of process capabilities and a workspace class (small or large). Each robot agent is associated with a reachability set over cells and buffers. In the unified AIMS view, this environment is treated as a tightly coupled production graph with 26 resource agents.

B.2 PRODUCTS AND NOMINAL SCHEDULE

Two product types are used:

- **S-product:** P1 → P2 → P3 → P6.
- **L-product:** P1 → P3 → P4 → P5.

Machines labeled `L` can process both L-products and S-products, while machines labeled `S` process only S-products. Products enter from `ENTRY` and leave through `EXIT` after completing their required routes.

The nominal workload contains 50 L-products and 50 S-products, released alternately every 30 ticks starting at tick 10. The initial production schedule is generated at approximately 50% resource utilization so that local rescheduling has available slack under disruptions.

B.3 DISRUPTION REGIME

The manufacturing disruption regime models machine breakdowns together with uncertainty in operation time. Machine breakdown probabilities at the start of simulation range from 3.3% to 10%. When a machine breaks down, a rescheduling process is triggered. The mean time to repair ranges from 1000 to 1500 ticks. If a breakdown occurs while a machine is processing a product, that product is damaged and cannot be recovered.

B.4 RESCHEDULING SEMANTICS

The local manufacturing replanning problem replaces an affected event sequence s_d by a new event sequence s_{new} such that the relevant product transition is preserved:

$$\text{Tr}(x_{\text{prior}}, s_{\text{new}}) = x_{\text{post}}. \quad (18)$$

Only the shortest affected sequence is replaced, which keeps replanning local and preserves as much of the nominal schedule as possible. Transportation events may need to be inserted when a transformation operation is moved to a different machine.

Given a candidate resource schedule, the earliest feasible start time heuristic $H(I, t, \delta, e)$ inserts a new event e into a resource idle-interval set I while allowing at most one local shift of an already scheduled operation. This heuristic is used to keep rescheduling local and to reduce induced deviation from the original schedule.

B.5 MANUFACTURING RISK MODEL

The risk-aware manufacturing setting evaluates two risks for a candidate replacement sequence:

- **Operational-delay risk** R_1 : the likelihood that inserting a new event causes downstream delays for other products already assigned to the selected resource.
- **Breakdown risk** R_2 : the likelihood that assigning the new event increases the probability of a future resource breakdown.

The total risk is represented by a weighted sum,

$$R(s_{\text{new}}) = w_1 R_1(s_{\text{new}}) + w_2 R_2(s_{\text{new}}), \quad (19)$$

with the reported case using $w_1 = 0.2$ and $w_2 = 0.8$. This weighting emphasizes avoidance of brittle recovery sequences that push machines close to failure.

C SUPPLY-CHAIN ENVIRONMENT AND DISRUPTION MODELS

C.1 COCKPIT SUPPLY-CHAIN NETWORK

The supply-chain testbed is an automotive cockpit network. Vehicle cockpits are the final products and are assembled from components, parts, and raw materials. The network contains three vehicle models, and each model requires one, two, or three cockpit styles. Each auto assembly plant produces a single vehicle model, while each cockpit assembly plant can produce multiple cockpit styles.

Each cockpit requires ten components. For cockpits of the same vehicle model, the cluster, substrate, glove box, HVAC system, cross-car beam, and steering column are shared. Each cockpit style has a style-specific infotainment system, wiring harness, and bezel combination, while all cockpits share a common airbag type. Upstream component production depends on parts and raw materials, and different components may share part/material suppliers.

The full network contains 117 supplier/customer agents and 413 transportation agents. The supplier/customer layer contains 5 customers, 3 cockpit assembly plants, 31 component suppliers, 62 part suppliers, and 16 raw-material suppliers. Each supplier or transportation agent has its own cost and capacity attributes.

C.2 GRAPH MODEL AND AGENT ROLES

The supply chain is modeled as a physical flow graph $G = (V, E)$ whose vertices are enterprise entities and whose edges are transportation/flow channels. In the agentized view, both vertices and edges can be represented as agents in the communication network. Roles include customer, distributor, OEM, tier supplier, and transporter. In the current AIMS exposition, the mechanism is instantiated over the active disruption-response subnetwork induced by the scenario under study, while keeping the same cockpit-network semantics.

C.3 NETWORK-STRUCTURED DISRUPTION SETTING

The network-structured setting studies a single supplier-loss disruption. The initial flow plan is computed by a centralized optimization model and then perturbed by removing one supplier agent that is active in the nominal plan. The scenario-generation rules are:

- the lost agent must have production tasks in the initial plan;

- exactly one supplier agent is disrupted in each scenario;
- production and transportation can exceed nominal capacity by up to 30% at an additional 50% unit cost.

In the nominal plan, 84 agents carry production tasks, yielding 84 disruption scenarios. Of these, 73 scenarios admit feasible recovery plans for both centralized and distributed coordination and are retained for the attribute-stratified analysis.

C.4 LEAD-TIME DISRUPTION AND HETEROGENEOUS-RISK SETTING

The stochastic supply-chain setting uses the same cockpit network with additional temporal attributes. Lead times are attached to suppliers and cockpit assemblers, and delivery deadlines are attached to customer demands. The network is initialized with an optimal baseline flow plan.

The disrupted supplier is `cluster_sup_3` (denoted S_3), which supplies three cluster products,

$$K_d = \{\text{cluster}_1, \text{cluster}_2, \text{cluster}_3\}. \quad (20)$$

Its downstream demand agents are the three cockpit suppliers

$$A_{\text{dm}} = \{A_1, A_2, A_3\}, \quad (21)$$

and alternative cluster suppliers are denoted by S_1 , S_2 , and S_4 . The disruption increases the lead time of S_3 by 20%, 60%, or 100%.

C.5 UNCERTAINTY MODEL AND DISTRIBUTED DECISION MAKING

The stochastic lead-time case uses a discrete-event simulation. Simulation starts from upstream agents with outflows at time 0. For each outflow, lead time is sampled from a known normal distribution. A downstream agent begins production only after all required components have arrived, so its completion time is determined by the latest required arrival among its incoming components. The process is iterated from upstream to downstream until final cockpits are delivered to all customers. Out-of-sample evaluation uses 300 independent replications.

Risk heterogeneity enters at both offer generation and supplier selection. A supplier agent solves a local stochastic response problem under uncertainty in production capacity and lead time. A demand agent solves a supplier-selection problem using received responses, while modeling trust in a supplier response as uncertainty around the offered quantity and arrival time. In the case study, risk-neutral decision making optimizes the expected objective via sample-average approximation, whereas risk-averse decision making optimizes a conservative worst-case variant.

D METRIC DEFINITIONS

D.1 PLAN DEVIATION AND COMMUNICATION

At the mechanism level, the revised plan is scored by a generic objective

$$J(\pi) = J_{\text{service}}(\pi) + \lambda_c J_{\text{cost}}(\pi) + \lambda_\Delta d(\pi, \pi_0) + \lambda_{\text{msg}} M(\pi), \quad (22)$$

where $d(\pi, \pi_0)$ is plan deviation and $M(\pi)$ is communication on G_a . A common instantiation of plan deviation is the sum of changed commitments plus local edit distances:

$$d(\pi, \pi_0) = \sum_{(i \rightarrow j, k) \in E_p} \mathbf{1}[(c_{ij}^k, \tau_{ij}^k) \neq (c_{ij}^{k,0}, \tau_{ij}^{k,0})] + \sum_{a \in A} d_a(u_a, u_a^0). \quad (23)$$

Communication is the total number of exchanged messages during replanning,

$$M(\pi) = \sum_{m \in \mathcal{M}(\pi)} 1. \quad (24)$$

D.2 MANUFACTURING METRICS

The manufacturing studies use the following metrics:

- **Cycle time:** completion time per product; used as the main recovery-quality proxy in the manufacturing frontier.
- **Communication M :** total request/response/award/confirm style exchanges during the rescheduling process.
- **Latency T :** total wall-clock runtime of the replanning implementation in seconds.
- **Damaged products:** number of products damaged due to breakdown during processing.
- **Broken machines:** number of machine breakdown events.
- **Rescheduled processes:** number of triggered rescheduling episodes.
- **Peak and average risk values:** summary statistics of the selected replacement schedules.

In the unified diagnostics, cycle time is used as the main quality proxy for R1, while damaged products, broken machines, rescheduled processes, and risk values support R3.

D.3 SUPPLY-CHAIN STRUCTURAL ATTRIBUTES

For the supply-chain environment, four structural attributes characterize the role of a disrupted agent:

Connectivity.

$$C_i = \sum_{a_j \in V} b_{ij} + b_{ji}, \quad (25)$$

where $b_{ij} = 1$ if edge (i, j) carries material/product flow and 0 otherwise.

Depth.

$$D_i = \max_{a_j \in \text{Customer}} d(a_i, a_j), \quad (26)$$

where $d(\cdot, \cdot)$ is geodesic distance.

Capability redundancy. For capability m ,

$$R_i(m) = |\{a_j \mid a_j \text{ has capability } m, a_j \in A \setminus \{a_i\}\}|. \quad (27)$$

Production complexity.

$$P_i = \left| \bigcup_{k \in K_i} \{k_f \mid k \in Z(k_f), \forall k_f \in K_f\} \right| + \left| \bigcup_{k \in K_i} Z(k) \right|, \quad (28)$$

where K_i is the product set of agent a_i , K_f is the final product set, and $Z(k)$ maps a product to the required components/materials.

D.4 SUPPLY-CHAIN REPLANNING METRICS

For the supplier-loss setting, the following metrics quantify footprint and cost:

Overage cost.

$$O = \sum_{(i,j) \in E, k \in K} \alpha_{ij} c_{ijk} y_{ijk}^o + \sum_{i \in V, k \in K} \beta_i e_{ik} p_{ik}^o, \quad (29)$$

where y_{ijk}^o and p_{ik}^o denote over-capacity transportation and production, respectively.

Network changes.

$$N_c = |\{a_i \mid p_i \neq p'_i\}| + |\{(i, j) \mid y_{ij} \neq y'_{ij}\}|. \quad (30)$$

Network additions.

$$N_a = \sum_{i \in V} \max\{0, \xi'_i - \xi_i\} + \sum_{(i,j) \in E} \max\{0, \zeta'_{ij} - \zeta_{ij}\}, \quad (31)$$

where $\xi_i, \zeta_{ij} \in \{0, 1\}$ indicate whether an agent or edge is used in the nominal plan and ξ'_i, ζ'_{ij} are the corresponding indicators after replanning.

Communication. For supply-chain centralized coordination,

$$M = 1 + 2|V| + N_c + N_a, \quad (32)$$

which accounts for one global replan trigger, information requests and responses with all agents, and update notifications to changed agents and edges. For distributed coordination, M counts all local request, response, award/selection, and confirmation exchanges used to construct the new plan.

D.5 SUPPLY-CHAIN STOCHASTIC ROBUSTNESS METRICS

For the lead-time disruption setting, let \hat{y}_{zjk} be the replanned flow from supplier a_z to demand agent a_j for product k , and let $v_{zjk,i}$ be its realized arrival time in replication i . Flow lateness is

$$\Delta_{zjk,i} = \max\{v_{zjk,i} - t_{jk}, 0\}. \quad (33)$$

The delayed-product percentage at a delay level δ_i is the fraction of total disrupted-demand flow whose realized lateness equals δ_i :

$$\frac{\sum_{a_j \in A_{\text{dm}}, a_z \in Z_j(k), k \in K_d} \hat{y}_{zjk} \mathbf{1}[\Delta_{zjk,i} = \delta_i]}{\sum_{a_j \in A_{\text{dm}}, a_z \in Z_j(k), k \in K_d} \hat{y}_{zjk}}. \quad (34)$$

At the aggregate demand-agent level,

$$C_{\text{dm}} = \sum_{a_j \in A_{\text{dm}}} C_{d,j}, \quad L_{\text{dm}} = \sum_{a_j \in A_{\text{dm}}} \sum_{k \in K_d} \Delta t_{jk}, \quad (35)$$

and the deterministic plan objective is

$$J_{\text{dm}} = C_{\text{dm}} + w_t L_{\text{dm}}, \quad (36)$$

with lateness penalty weight $w_t = 10^5$ in the reported case study.

E UNIFIED DIAGNOSTICS CONSTRUCTION

E.1 R1: QUALITY-LATENCY-COMMUNICATION FRONTIER

R1 combines communication M , latency T , and a domain-specific quality signal into a single cross-domain diagnostic. In the manufacturing domain, the plotted frontier uses cycle time as the primary throughput-quality proxy together with average communication and replanning runtime. In the supply-chain domain, the plotted frontier uses scenario-level overage cost O against communication M , with feasibility additionally marked as *feasible* or *unmet*.

E.2 R2: PROPAGATION FOOTPRINT AND NETWORK STRUCTURE

R2 is reported on the supply-chain network-structured setting. The propagation footprint is represented by the footprint proxies (N_c, N_a) together with communication M and quality impact O . In the AIMS presentation, scenarios are stratified into four groups based on low/high connectivity and low/high redundancy, denoted LL, LH, HL, and HH. The current grouped analysis contains 53, 16, 2, and 13 scenarios in these four groups, respectively.

E.3 R3: ROBUSTNESS UNDER UNCERTAINTY

R3 combines the two stochastic settings. In manufacturing, robustness is captured by incident counts (damaged products, broken machines, rescheduled processes) and risk summaries (peak and average schedule risk). In supply chain, robustness is captured by the delayed-product distribution across disruption intensities and by deterministic cost-latency trade-offs summarized through $(C_{\text{dm}}, L_{\text{dm}}, J_{\text{dm}})$.

E.4 INCLUSION CRITERIA

For the supplier-loss study, not all disruptions admit feasible recovery under local or centralized coordination. Cases with zero redundancy or insufficient residual capacity can become unmet-demand scenarios. The attribute-stratified propagation analysis is therefore restricted to the 73 scenarios in which both centralized and distributed recovery produce feasible plans. The stochastic lead-time analysis is performed on the fixed disrupted-subnetwork centered on the cluster supplier $S3$ and the three affected cockpit assemblers.

F EXPERIMENTAL PROTOCOL, BUDGETS, AND ACCOUNTING

F.1 TRIALS, SCENARIOS, AND REPLICATIONS

The manufacturing study reports averages over five simulation trials. The supplier-loss supply-chain study evaluates 84 single-supplier-loss scenarios, of which 73 are used for feasible-plan comparison and attribute-stratified diagnostics. The stochastic lead-time study evaluates three disruption intensities (20%, 60%, 100% lead-time increase) and uses 300 out-of-sample simulation replications for each setting.

F.2 PLATFORM AND SOLVER NOTES

The manufacturing environment is implemented in Repast Symphony. The stochastic supply-chain simulation is implemented using Python scripts that schedule agent decisions and simulate the resulting discrete-event dynamics. Supplier-response and supplier-selection problems are solved as tractable mixed-integer models under sample-average approximation in the stochastic setting.

F.3 BUDGET DEFINITIONS

The mechanism is evaluated under explicit communication and time budgets. Communication is bounded by a message budget B_{msg} (and optionally a round budget), while latency is constrained by a wall-clock deadline B_{time} . In the presented diagnostics, latency is reported as measured runtime and communication is reported as message exchanges under a fixed accounting rule.

F.4 MESSAGE ACCOUNTING RULE

For the appendix calculations and the unified diagnostics, each point-to-point REQUEST, OFFER, AWARD, and CONFIRM transmission is counted as one message exchange. If a message is sent to multiple recipients, the accounting is per recipient. Centralized coordination is accounted for separately using a domain-specific aggregation of information-collection and update-notification steps; distributed coordination counts only the explicitly executed negotiation messages needed to produce the revised plan.

G ADDITIONAL ALGORITHMIC AND QUALITATIVE DETAILS

G.1 MANUFACTURING FALLBACK BEHAVIOR

If no feasible replacement sequence is found within the local coordination space—for example because no redundant resource is available or hard requirements cannot be met—the distributed manufacturing controller escalates to a central controller or human manager for further action or for relaxation of constraints. This fallback is not the default path; it is only used when the local candidate set is empty.

G.2 SUPPLY-CHAIN PROPAGATION SEMANTICS

Supply-chain communication proceeds in contract-style rounds. A disrupted agent informs directly affected downstream agents of the lost or delayed flow. These agents become demand agents and request support from candidate upstream suppliers. Once a demand agent selects replacement flows,

the selected suppliers update their own local balance; if those commitments create new upstream needs, those suppliers become new demand agents. Propagation stops when all induced needs are balanced or when no additional feasible support exists.

G.3 QUALITATIVE EFFECT OF HETEROGENEOUS RISK ATTITUDES

In the lead-time disruption case, heterogeneous risk attitudes change which backup suppliers are chosen even when all affected agents face the same disrupted supplier. In the reported examples, risk-averse decisions place less reliance on the disrupted supplier $S3$ and shift more quantity toward backup suppliers with better tail behavior, even when that choice incurs a modest increase in nominal cost. This effect is strongest at higher disruption intensity, where conservative commitments reduce severe delay outcomes.

H SCOPE AND REMAINING LIMITATIONS

The current appendix closes several implementation gaps, but it does not claim a fully standardized industrial benchmark suite. The presented evidence covers two representative industrial domains and two supply-chain disruption regimes, but not a full cross-domain benchmark with a single shared simulator and protocol.

The mechanism exposes explicit scope and gate knobs—notably (h_s, K_s) and $(\varepsilon_\Delta, \varepsilon_r)$ —that are important for deployment. The current paper reports their functional role and the resulting trends, but does not provide a systematic sensitivity sweep or automated tuning procedure.

Communication is accounted for at the message-exchange level rather than at the bit or semantic-content level. This is appropriate for the present mechanism-level comparison, but it does not capture compression, variable payload size, or network-layer contention.

Finally, the present stochastic supply-chain analysis focuses on a single disrupted supplier with heterogeneous risk attitudes primarily at the demand side. More diverse agent-risk combinations, simultaneous disruptions, and broader topological regimes remain open extensions.