
UNAWARE ADVERSARIES: A FRAMEWORK FOR CHARACTERIZING EMERGENT CONFLICT BETWEEN NON-COORDINATING AGENTS

Scott VanRavenswaay
scottvr@gmail.com

ABSTRACT

In many complex systems, constituent agents engage in adversarial behavior without explicit awareness or intent. These “unaware adversaries” operate based on localized objectives and feedback, yet their interactions within a shared environment result in mutual disruption, resource competition, or functional conflict. This paper formally defines this phenomenon and situates it within existing literature on systems theory. We present and analyze several case studies—including closed-loop, open-loop, and designed conflicts—to build a taxonomy of adversarial typologies. By demonstrating how this framework can be used as a generative tool for analysis and design, we derive novel mitigation strategies for problems in fields as diverse as generative modeling and internet routing. The result is a framework intended to help identify, classify, and resolve a pervasive yet under-examined class of system failures, and to propose new avenues for research.

Keywords unaware adversaries · emergent conflict · system dynamics · feedback control loops · multi-agent systems · generative adversarial networks · open-loop systems · BGP · network stability · protocol design.

1 Introduction

Consider a common real-world scenario: an office room equipped with two independent climate control systems. A radiator, governed by a building-wide thermostat, provides heat, while a window-mounted air conditioning unit, with its own separate controls, provides cooling. Each system operates according to its own local feedback loop. If an occupant turns on the A/C to cool a stuffy room while the building’s heating system is simultaneously trying to maintain a minimum winter temperature, the two agents enter a state of persistent, mutually negating work—a thermodynamic conflict that neither is designed to recognize. This scenario serves as an intuitive archetype for a class of interactions we term “unaware adversaries.”

In this paper, we formalize the study of systems wherein agents function as antagonists without semantic awareness of their opposition. While the term “adversary” traditionally implies intentionality, our focus is on interactions where conflict is an emergent property of environmental coupling and feedback dynamics. By analyzing these phenomena, we can better understand and mitigate unintended inefficiencies and instabilities in complex, decentralized systems.

2 Background and Related Work

The concept of unintended systemic conflict is not new, and the framework of unaware adversaries builds upon several established ideas from economics, systems theory, and public policy. In economics, the notion of a “negative externality” describes a cost imposed on a third party by an economic activity. While related, externalities are often unidirectional and do not necessarily involve the reciprocal feedback loop central to our definition of unaware adversaries. In systems theory, Peter Senge’s work on system archetypes is highly relevant [1]. The “Fixes that Fail” archetype and “Shifting the Burden” both describe the outcomes of such conflicts. Similarly, the study of “policy resistance” in public administration explores why policies often fail or backfire due to the complex reactions of the system they are intended to manage. Our framework contributes to this body of work by focusing specifically on the

agent-level *mechanism* of the conflict: two or more non-coordinating agents, each pursuing a local objective, coupled through a shared environment. We propose that by classifying the nature of the agents, their coupling, and the resulting impact, we can create a more targeted diagnostic and prescriptive tool than a general archetype provides.

3 A Framework for Unaware Adversaries

We formally define unaware adversaries as a set of two or more agents or systems that satisfy the following criteria:

Independent Objectives Each agent pursues a distinct objective function based on its own state and local environmental variables.

Shared Environment The agents co-exist and operate within a common environment, defined by a set of shared physical, informational, or resource-based state variables.

Indirect Coupling The actions of one agent affect the state of the shared environment, thereby influencing the sensory inputs and subsequent actions of the other agent(s). There is no explicit communication protocol for deconfliction.

Antagonistic Effect The collective behavior of the agents results in a net negative impact on the fulfillment of their respective objectives or the overall system efficiency.

4 Case Studies

4.1 Thermostatic Conflict: The Heater-A/C Dyad

Our canonical example involves a room containing two physically distinct and uncoordinated devices: a radiator governed by a building-level thermostat and a window air-conditioning unit with its own integrated controls. This situation is common in older buildings or offices where supplemental climate control is added. Conflict arises when the setpoints are in opposition; for example, the building heat is set to maintain 20°C while the window A/C is set by an occupant to cool the room to 22°C. If the room temperature drifts to 21°C, both systems will activate, leading to a direct conflict. We can model the temperature dynamics and energy cost as follows:

$$\frac{dT}{dt} = -k(T - T_{\text{ext}}) + P_H \cdot u_H(t) - P_C \cdot u_C(t)$$

$$\text{Cost}(J) = \int_0^{T_{\text{final}}} (C_H \cdot u_H(t) + C_C \cdot u_C(t)) dt$$

Here, $u(t)$ is the control signal (1 if on, 0 if off) for the heater (H) and cooler (C), P_H and P_C are their respective thermal powers, and C_H and C_C are their electrical power consumption rates. The total energy cost, J , in the adversarial configuration is orders of magnitude higher than in a stable configuration, for no improvement in temperature stability.

4.2 Email Deliverability Conflict (Open-Loop)

A compelling open-loop example occurs in corporate IT systems. Consider two agents sending email from the same domain ('@company.com'):

- **Agent A (Marketing Platform):** Its objective is to maximize email sends for a campaign, blasting 500,000 emails in an hour.
- **Agent B (Transactional System):** Its objective is to ensure reliable delivery of critical password reset or receipt emails.

The shared environment is the company's domain reputation with email providers (e.g., Google, Microsoft). Agent A's action (the email blast) degrades this reputation, causing Agent B's critical emails to be marked as spam or blocked. The feedback loop is open because the failure of the password reset emails provides no direct signal back to the marketing platform. Agent A's performance metrics are unaffected by the harm caused to Agent B, so it continues its behavior without correction.

4.3 From Designed Conflict to a Novel Research Hypothesis: The Pedagogical GAN

The standard Generative Adversarial Network (GAN) [2] provides a powerful case study for our framework. It is a system of two agents, a Generator (G) and a Discriminator (D), locked in a designed, zero-sum game. This adversarial

dynamic, however, is notoriously unstable and suffers from practical issues like vanishing gradients, where D becomes too proficient, leaving G with no learning signal. The original authors’ first solution was the heuristic "non-saturating" loss, an immediate modification that sought a stronger, more reliable gradient for G . This established the central challenge in the field: managing the adversarial dynamic for stable and efficient training.

In the years since, the dominant paradigm for GAN stabilization has become one of **gradient control**. Landmark models like Wasserstein GAN (WGAN) [3] and its successor WGAN-GP [4] diagnosed the problem as being rooted in the geometry of the loss landscape. Their solution, which now represents the state-of-the-art, is to **tame and constrain** the discriminator’s function (e.g., by enforcing a Lipschitz condition) to guarantee that it always provides a smooth and informative gradient to the generator. This philosophy is about preventing conflict from becoming destructive by carefully limiting the power of the adversary.

Our framework of unaware adversaries prompts a different line of inquiry. Instead of asking, "How do we control the conflict?", we ask, "Can we redesign the agents’ objectives to make the conflict more productive?" This leads us to propose a novel approach that stands in philosophical opposition to gradient control. We term this the **Pedagogical GAN**.

The core idea of the Pedagogical GAN is to change the generator’s objective from simply fooling the discriminator to actively teaching it as efficiently as possible. We formalize this by proposing that the generator should seek to **maximize** the discriminator’s learning signal. The generator’s objective function becomes:

$$\max_G \|\nabla_D \mathcal{L}(D, G)\|_2$$

Here, $\mathcal{L}(D, G)$ is the standard discriminator loss. The generator is now explicitly incentivized to find samples that lie on the steepest parts of the discriminator’s loss landscape. It becomes a “Socratic tutor” that seeks to **weaponize** the gradient for accelerated learning, not suppress it.

This approach represents a significant conceptual departure. It is distinct from other cooperative frameworks like Unrolled GANs [5], which use strategic foresight, or other non-antagonistic models that alter loss functions to escape the zero-sum game [6]. Instead, it can be viewed as the principled and extreme conclusion of the line of thinking that began with the very first non-saturating GAN loss. Our literature review suggests that while the raw intuition for cooperative training has been informally discussed, this specific mechanism of maximizing the discriminator’s gradient norm appears to be a formally unexplored, high-risk, high-reward avenue for GAN research.

4.4 Systemic Opposition in Urban Planning

Perhaps the most societally significant examples of unaware adversaries occur in socio-technical systems like urban planning. Consider two municipal departments with laudable but conflicting mandates:

- **The Transportation Department:** Tasked with mitigating traffic congestion, its primary Key Performance Indicator (KPI) is maintaining a high Level of Service (LOS) on major arterial roads, a standard defined in the *Highway Capacity Manual* [8]. A high LOS corresponds to low traffic density and high travel speeds. The department’s primary tool is increasing road capacity (K) by adding lanes.
- **The Urban Planning Department:** Tasked with addressing housing shortages and promoting sustainability, its goal is to encourage dense, walkable, Transit-Oriented Development (TOD).

These two locally-optimized strategies become potent unaware adversaries. The core of the conflict can be formally captured using standard models from transportation engineering. A widely used relationship to model the effect of congestion on travel time is the Bureau of Public Roads (BPR) formula [9]:

$$t = t_0 \left[1 + \alpha \left(\frac{V}{K} \right)^\beta \right]$$

Here, t is the actual travel time on a road segment, t_0 is the free-flow travel time (with no traffic), V is the traffic volume (vehicles per hour), and K is the practical capacity of the road. The parameters α (typically ≈ 0.15) and β (typically ≈ 4) determine the severity of congestion effects. The Transportation Department’s objective is to minimize t by increasing K .

Concurrently, the Planning Department approves dense residential projects. The immediate effect is an increase in the base number of trips originating in the area, contributing to a higher traffic volume V . However, the more insidious coupling occurs through the phenomenon of **induced demand**. The Transportation Department’s expansion of a highway temporarily reduces travel time t , making driving faster and more attractive. This, in turn, fuels suburban growth and causes more people to drive, increasing the total traffic volume V .

This effect has been empirically verified and is known as the "Fundamental Law of Road Congestion," which states that the total vehicle-kilometers traveled (VKT) in a metropolitan area increases in direct proportion to the available lane-kilometers of roadways [10]. We can model this relationship simply as:

$$V = V_{\text{base}} + \gamma K$$

where V_{base} represents the volume from existing development and γK represents the new volume induced by the capacity expansion itself.

The adversarial loop is now clear. Substituting the induced demand model into the BPR formula reveals the systemic trap:

$$t = t_0 \left[1 + \alpha \left(\frac{V_{\text{base}} + \gamma K}{K} \right)^\beta \right] = t_0 \left[1 + \alpha \left(\frac{V_{\text{base}}}{K} + \gamma \right)^\beta \right]$$

When the Transportation Department increases capacity (K), the term $\frac{V_{\text{base}}}{K}$ decreases, which is the intended effect of alleviating existing congestion. However, the presence of the induced demand constant γ means that as K becomes very large, the ratio inside the parentheses does not approach zero, but rather approaches γ . The travel time t does not return to the free-flow time t_0 , but instead approaches a new, congested equilibrium of $t_0(1 + \alpha\gamma^\beta)$.

Each department successfully meets its own local KPIs—lanes are built, and housing is approved—yet the system as a whole suffers from persistent congestion despite massive capital expenditure. The agents are locked in a state of policy-induced homeostasis, where each one's logical, well-intentioned actions systemically negate the efforts of the other.

4.5 Interdomain Routing Instability and the YouTube/Pakistan BGP Hijack

The dynamic of independent agents pursuing local optimizations at the expense of global system health extends from urban planning to the core architecture of the internet itself. The Border Gateway Protocol (BGP) [13] governs interdomain routing across the Internet, enabling thousands of autonomous systems (ASes) to exchange reachability information. While BGP is highly flexible, its decentralized and policy-driven design creates opportunities for emergent instability. This subsection presents a general case of unaware adversarial interaction in interdomain traffic engineering, followed by a concrete, high-impact example: the 2008 YouTube/Pakistan BGP hijack incident.

General Case: Unaware Traffic Engineering Between ASes

Consider two neighboring ASes:

- **AS A (Transit Provider):** Uses internal link metrics and MEDs [14] to load-balance traffic and minimize internal congestion.
- **AS B (Content Provider):** Optimizes for inbound performance by modifying BGP community tags and selectively announcing prefixes.

Both systems reactively adjust their routing announcements based on local utility functions—CPU load, latency, throughput—but do so without coordination. Their interaction takes place through BGP advertisements and path selection behavior, producing a shared environment susceptible to instability.

Each AS modulates its route announcements as a function of observed performance:

$$R_A(t + \Delta t) = f_A(R_B(t), L_A(t)) \quad \text{and} \quad R_B(t + \Delta t) = f_B(R_A(t), L_B(t))$$

This coupling forms a delayed feedback loop. The system may enter oscillatory regimes or route flap storms, degrading routing stability and increasing global convergence latency [11].

Proposed Mitigation: Gradient-Paced BGP Advertiser (GPBA) Inspired by the Pedagogical GAN's gradient-based objective modification, we propose a novel mitigation strategy for interdomain routing: the **Gradient-Paced BGP Advertiser (GPBA)**. Rather than impose fixed rate limits, GPBA adapts the pace of route advertisement based on the environmental sensitivity observed at the peer:

$$\Delta t_{\text{propagation}} \propto \left\| \frac{dR_{\text{remote}}}{dR_{\text{local}}} \right\|$$

This principle discourages volatile updates by throttling agents whose announcements produce outsized ripples in routing state. Unlike traditional filters, GPBA uses a reflexive meta-policy: it applies only when sensitivity is high and relaxes when peers demonstrate robustness or stability. This maintains responsiveness while discouraging feedback-amplifying behavior.

Case Study: The YouTube/Pakistan Hijack (2008)

In February 2008, Pakistan Telecom (AS17557) attempted to implement a government-ordered block of YouTube by announcing a more-specific prefix (208.65.153.0/24) intended for internal blackholing. This announcement was mistakenly propagated to PCCW (AS3491), a major transit provider, which further distributed the route across the global Internet.

Due to BGP’s path selection rules, the hijacked prefix—which was more specific than YouTube’s legitimate announcement—was globally preferred, redirecting worldwide traffic to a null route in Pakistan. YouTube was rendered inaccessible for several hours [12].

This incident demonstrates classic unaware adversarial dynamics:

- **Independent Objectives:** Pakistan Telecom sought national censorship; PCCW aimed to maintain routing continuity and performance.
- **Shared Environment:** The global BGP routing table.
- **Indirect Coupling:** PCCW’s propagation choices were influenced by AS17557’s unexpected prefix origin.
- **Antagonistic Effect:** Global denial of service for YouTube, triggered by local policy and lack of propagation checks.

Could GPBA Have Mitigated This? A GPBA-aware edge at PCCW could have detected that:

1. The origin of 208.65.153.0/24 (AS17557) was highly anomalous.
2. The route produced a high volatility score in terms of downstream reachability change.

Rather than immediately propagate the route, the GPBA heuristic would have increased the propagation delay or required additional validation. This passive reflexivity—without requiring global RPKI or semantic understanding of the policy intent—could have bought time for human intervention or for automated anomaly detectors to act.

General Insight The YouTube/Pakistan incident underscores the vulnerability of loosely coordinated protocols like BGP to unaware adversarial dynamics. GPBA-style mechanisms offer a path forward by embedding environmental reflexivity into protocol behavior. Rather than controlling agent behavior through centralized rules, GPBA designs the feedback surface itself to discourage destabilizing actions. This reframing parallels the shift in GAN stabilization from adversarial suppression to cooperative learning seen in the Pedagogical GAN. Rather than controlling agent behavior through centralized rules, GPBA designs the feedback surface itself to discourage destabilizing actions. In the language of our design principles, it represents a sophisticated form of the Observability Principle, where the system observes its own sensitivity to change and creates feedback to regulate it.

5 A Taxonomy of Unaware Adversarial Systems

To demonstrate the analytical power of this framework, we now apply the taxonomy to our diverse case studies. As shown in Table 1, this classification not only organizes the examples but also reveals critical insights into the underlying dynamics of each type of conflict.

Table 1: Taxonomy Analysis of Case Studies

Case Study	Domain	Coupling	Feedback	Symmetry
Heater vs. A/C	Mechanical	Physical	Closed-Loop	Symmetric
Email Deliverability	Digital	Informational	Open-Loop	Asymmetric
Pedagogical GAN	Digital	Informational	Coordinated	Asymmetric
Urban Planning	Social	Environmental	Closed-Loop	Asymmetric
BGP Hijack / GPBA	Network Infrastructure	Informational	Semi-Open Loop	Asymmetric

The taxonomy reveals a critical distinction: in the open-loop conflict, the harm is asymmetric and the adversarial agent receives no corrective feedback, suggesting such problems can persist undetected for long periods. This contrasts with the noisy, symmetric inefficiency of the closed-loop thermostat conflict.

The BGP case introduces a "semi-open loop," where an agent receives operational feedback that its actions have propagated (the route was accepted) but remains blind to the wider, negative systemic consequences of that action (the denial of service).

6 Implications for System Design: From Mitigation to Architecture

An understanding of unaware adversaries does more than explain a class of failures; it provides a powerful set of principles for designing robust, efficient, and intelligent systems. By consciously analyzing the potential for emergent conflict, designers can move from a reactive posture of debugging and mitigation to a proactive one of thoughtful system architecture. The framework allows us to ask critical questions at the design stage: Who are the agents in this system? What environment do they share? How are they coupled? And most importantly, how could their local objectives lead to global conflict? The answers lead to a set of clear design principles.

6.1 The Observability Principle: Make System-Wide Effects Locally Apparent

This principle directly targets the most insidious conflicts: the open-loop, asymmetric cases where one agent operates completely unaware of the harm it causes. The solution is to create a feedback channel where one does not naturally exist.

The Problem: As seen in the Email Deliverability case, the marketing platform's performance metrics (e.g., emails sent) are decoupled from the negative externality it creates (degraded domain reputation). The agent is flying blind to the true system-wide cost of its actions.

The Design Solution: Instrument the system to make the shared state visible and, crucially, to integrate that state into each agent's local decision-making context. Don't just provide the marketing team with a dashboard showing their send volume; create a shared, top-level "Domain Health" or "Deliverability Score" that is a KPI for *both* the marketing team and the IT department. When the marketing platform initiates a blast, it should see a real-time dip in a metric it is also responsible for. This makes the invisible visible and forces the unaware adversary to reckon with the consequences of its behavior.

6.2 The Incentive Principle: Redefine "Winning" to Induce Cooperation

For conflicts where agents have fundamentally misaligned goals, imposing external controls can be brittle and inefficient. A more elegant and robust solution is to redesign the agents' objective functions so that cooperation becomes an emergent property of them pursuing their own, revised goals.

The Problem: In the Urban Planning and standard GAN examples, each agent "winning" by its local definition leads to collective failure (gridlock) or instability (vanishing gradients). Their objectives are locked in zero-sum or mutually destructive patterns.

The Design Solution: This is the most transformative principle. Instead of just adding a coordinating layer, change what it means for an agent to succeed. Our proposal for the **Pedagogical GAN** is the canonical example: we shift the Generator's objective from "fool the Discriminator" to "maximize the Discriminator's learning rate." This reframes the adversary not as an opponent to be defeated, but as a student to be taught efficiently. In the Urban Planning context, it would mean replacing department-level KPIs like "Level of Service" and "Housing Units Approved" with a single, shared, system-level objective, such as "Minimize average citizen commute time" or "Maximize accessibility score." This forces the departments to find a cooperative, globally optimal solution.

6.3 The Hierarchy Principle: Appoint a Single Arbiter for Contested States

When agent objectives are simple, directly oppositional, and tightly coupled through a physical state, the most effective solution is often to remove autonomy and establish a clear control hierarchy.

The Problem: The Heater-A/C dyad is a perfect example. Two simple agents are fighting a direct, symmetric battle over a single, continuous variable: room temperature. Their control logic is too simple to negotiate a truce.

The Design Solution: Institute a higher-level arbiter or a "master controller" that has authority over both agents and manages the contested resource. The solution is a single, smart thermostat that controls both the heat and the A/C. It becomes the sole source of truth for the room's temperature and can make intelligent decisions (e.g., defining a deadband between heating and cooling setpoints) that are impossible for the independent agents. This principle applies to any system fighting over a fungible resource, from CPU cycles in a computer to water rights in a river basin.

By applying these principles, the framework of unaware adversaries becomes a generative tool. It equips a designer to anticipate failure modes and architect systems that are not only resilient to conflict but are designed from the ground up for emergent cooperation.

7 Discussion: The Architecture of Interaction

The primary contribution of this paper is a framework that serves not only as a descriptive and diagnostic tool but as a generative engine for system design. We demonstrated this generative capacity twice over, deriving novel solutions in two disparate domains: the **Pedagogical GAN** for machine learning and the **Gradient-Paced BGP Advertiser (GPBA)** for network protocol design. Both concepts emerged directly from applying our framework's lens to a well-known problem. In each case, instead of accepting the standard adversarial dynamic as given, the framework prompted a fundamental question: "How can we redesign the agent's interaction rules to make the conflict less destructive or even productive?" The proposed solutions—one that reframes training as teaching, the other that paces communication based on systemic sensitivity—are deliberate architectural choices about the nature of the agents' interaction. They represent a significant philosophical departure from conventional approaches that focus on adversarial suppression or fixed rate-limiting.

While the Pedagogical GAN is a high-risk, high-reward proposal that requires empirical validation, its conceptual origin is proof of the framework's value. It is the direct result of applying the **Incentive Principle**, one of the three core design strategies that emerge from our analysis. The other principles are equally direct consequences of the framework. The **Observability Principle** is the clear solution to the open-loop, asymmetric harm seen in the Email Deliverability case, while the **Hierarchy Principle** is the obvious remedy for the simple, symmetric conflict of the Heater-A/C dyad. The framework, therefore, does not just identify problems; it naturally leads to a specific class of solutions tailored to the type of conflict.

The concept of unaware adversaries is especially salient today. As we continue to build ever more complex, decentralized systems—from microservice architectures and the Internet of Things to multi-agent AI ecosystems and smart cities—the potential for emergent, unintended conflict grows exponentially. Our traditional engineering and debugging practices, which focus on verifying the correctness of individual components, are insufficient for addressing failures that arise from the negative space *between* those components. The problem is not that the transportation department's models are wrong or that the marketing platform's code has a bug; the problem is that the system's architecture creates incentives for globally destructive, locally optimal behavior.

Ultimately, this framework encourages a crucial shift in perspective for architects and designers. It moves the focus from the design of agents in isolation to the design of the environment they share and the incentive landscape they inhabit. By analyzing the structure of the coupling and the nature of the feedback loops, we can stop being surprised by "policy resistance" or "unintended consequences" and start architecting systems where cooperation is not an afterthought, but an emergent property of the system's fundamental design.

8 Directions for Future Research

This framework opens several avenues for future inquiry, most notably:

- The formal implementation and empirical validation of the proposed **Pedagogical GAN**. This would involve developing the necessary regularization techniques to stabilize the training dynamic and testing its performance against baseline models and gradient control methods like WGAN-GP.
- Analysis of unaware conflict in large-scale distributed AI systems, such as competing recommendation algorithms on a single platform.
- Application of the unaware adversary lens to AI alignment and interpretability, especially in ensembles of independently trained agents.
- The formal specification and empirical validation of the proposed **Gradient-Paced BGP Advertiser (GPBA)**. This would involve simulation in environments like BGPStream to test its efficacy in dampening route flaps.

and mitigating the propagation speed of malicious hijacks compared to existing mechanisms like route flap damping.

9 Conclusion

The concept of unaware adversaries provides a valuable lens for analyzing and resolving a significant class of system failures. By moving beyond debugging individual components to architecting the environment in which they interact, we can mitigate and even prevent emergent, destructive conflict. This paper defined a formal framework for this analysis and demonstrated its utility across mechanical, digital, and socio-technical systems. Crucially, we have shown this framework to be more than just descriptive; it is a generative tool for innovation. The derivation of two distinct and plausible solutions, the Pedagogical GAN and the Gradient-Paced BGP Advertiser, proves that a principled understanding of emergent conflict can inspire novel designs for more robust, stable, and cooperative systems.

Acknowledgments

The author acknowledges the conceptual contributions of the thermostatically-controlled agents in countless offices and laboratories.

References

- [1] Peter M. Senge. *The Fifth Discipline: The Art & Practice of The Learning Organization*. Doubleday, 1990.
- [2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680, 2014.
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- [4] Ishan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved Training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30*, pages 5767–5777, 2017.
- [5] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- [6] Jianwen Xie, Yang Lu, Song-Chun Zhu, and Ying Nian Wu. Cooperative Training of Descriptor and Generator Networks. *arXiv preprint arXiv:1609.09408*, 2016.
- [7] Anthony Downs. *Stuck in Traffic: Coping with Peak-Hour Traffic Congestion*. Brookings Institution Press, 1992.
- [8] Transportation Research Board. *Highway Capacity Manual, 7th Edition: A Guide for Multimodal Mobility Analysis*. National Academies of Sciences, Engineering, and Medicine, 2022.
- [9] Bureau of Public Roads. *Traffic Assignment Manual*. U.S. Department of Commerce, Urban Planning Division, 1964.
- [10] Gilles Duranton and Matthew A. Turner. The Fundamental Law of Road Congestion: Evidence from US Cities. *American Economic Review*, 101(6): 2616–52, 2011.
- [11] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. Delayed Internet Routing Convergence. In *Proceedings of ACM SIGCOMM*, 1997.
- [12] Jim Cowie. *Pakistan Hijacks YouTube*. Renesys Blog, February 24, 2008. Available at: <https://web.archive.org/web/20210506213233/https://www.renesys.com/2008/02/pakistan-hijacks-youtube-1/> (archived from the original).
- [13] Y. Rekhter, T. Li, and S. Hares, editors. *A Border Gateway Protocol 4 (BGP-4)*. Request for Comments 4271, Internet Engineering Task Force, January 2006. <https://datatracker.ietf.org/doc/html/rfc4271>.
- [14] D. McPherson, V Gill, editors. *BGP MULTI_EXIT_DISC (MED) Considerations*. Request for Comments 4451, Internet Engineering Task Force, March 2006. <https://datatracker.ietf.org/doc/html/rfc4451>