Extending No-Regret Hopping in FMCW Radar Interference Avoidance

Yunian Pan
New York University
370 Jay Street
Brooklyn, NY, USA
yp1170@nyu.edu

Quanyan Zhu
New York University
370 Jay Street
Brooklyn, NY, USA
qz494@nyu.edu

ABSTRACT

Nonlinear frequency hopping combined with no-regret learning has emerged as a promising approach for mitigating interference in automotive FMCW radar systems. However, in dense traffic scenarios, these techniques are challenged by the constraint that each radar has to sample at least one subband per chirp. To enhance range resolution and interference mitigation, we extend No-Regret Hopping by expanding the strategy space along the time dimension, such that every radar can utilize the full available bandwidth—even when the number of radars exceeds the number of partitioned subbands.

1. INTRODUCTION

High-resolution frequency modulated continuous wave (FMCW) automotive radars are crucial for advanced driver assistance system (ADAS) and autonomous driving (AD) because they reliably measure a target's range, velocity, and angle under all weather conditions [4, 6, 12]. However, their widespread deployment has raised concerns over radar-to-radar interference, which degrades target detection performance. To mitigate interference, methods such as proactive Tx-side techniques (e.g., frequency hopping [11], minimum variance beamforming [2]) and network-based medium access control (MAC) protocols [1, 3] have been proposed, though uncoordinated channel parameter assignments limit these approaches' adaptability.

A recent approach models radars as players in a game where their frequency hopping strategies determine their utility [7], and game-theoretic learning is employed to steer these strategies toward a *Coarse Correlated Equilibrium* (CCE). At CCE, the hopping start frequencies are selected so that no radar can improve its expected signal-to-interference-plus-noise ratio (SINR) by deviating. This method, known as *No-regret Hopping*, facilitates distributed frequency scheduling that mitigates interference and enhances range resolution. Its rationale is based on a property of no-regret algorithms [8, 9]: as regret diminishes, the empirical distribution of hopping strategies converge to the CCE.

However, partitioning the frequency bandwidth into a finite number of subbands poses a challenge when many radars share the same spectrum—for instance, an automotive vehicle equipped with six sensors would require at least six subbands for that single vehicle. To address this limitation, we propose augmenting no-regret hopping by expanding the strategy space. In our approach, radars adopt a stochastic round-robin scheme and apply a strategic time shift at the start of each pulse repetition interval (PRI), ensuring full bandwidth utilization even in high-density scenarios.

2. GAME FORMULATION

Consider a scenario in which a set of FMCW automotive radars, denoted by $\mathcal{N} = \{1, \ldots, N\}$, detects surrounding targets in real time. Each radar transmits linear frequency modulated (LFM) pulses over a total bandwidth B with the central frequency f_c.

We target the performance over $\mathcal{T}$ time frames, indexed by $\tau \in \{1, \ldots, \mathcal{T}\}$, each spanning a total duration of T. Within each frame, every radar $i \in \mathcal{N}$ transmits K^i consecutive chirps. The pulse repetition interval (PRI) for radar i is defined as $T^i_{\text{PRI}} = \frac{T}{K^i}$, which includes both the active period T^i_a and the idle period T^i_d, with $T^i_{\text{PRI}} = T^i_a + T^i_d$ where $T^i_d > 0$. Consequently, the chirp sweeping slope for radar i is $\alpha^i = \frac{B_a}{T^i_a}$. For simplicity, we assume that $K^i = K$, $T^i_d = T_d$, and $T^i_a = T_a$ for all $i \in \mathcal{N}$,

We partition:

- the total bandwidth B into A_1 subbands, $\{f_1, \ldots, f_{A_1}\}$.
- the idle duration T^i_d into A_2 time slots $\{t_1, \ldots, t_{A_2}\}$

Thus, each subband is defined by a starting frequency $f_a = f_c + (a-1)B_a$, for $a \in \{1, \ldots, A_1\}$; during each chirp the LFM pulse sweeps only a portion B_a of the total bandwidth B, across time interval $[t, t + T^i_a]$, $t \in \{t_1, \ldots, t_{A_2}\}$. We denote by Cartesian product $\mathcal{A} = \{f_1, \ldots, f_{A_1}\} \times \{t_1, \ldots, t_{A_2}\}$ the action space.

2.1 Hopping Strategies

To mitigate interference, the radars strategically select channel and time shift sequences $(f^i_1, t^i_1, \ldots, f^i_K, t^i_K)$ across the chirps. We denote the channel selection strategy of radar i by $\sigma^i(\cdot) : \{1, \ldots, K\} \to \mathcal{A}$; for example, $\sigma^i(k) = (f, t)$ indicates that radar i selects subband $[f, f + B_a]$ at chirp k, sending one LFM pulse between $[t, t + T_a]$. Let Σ denote the common strategy set shared by all radars. Note that Σ need not include every possible mapping; for instance, it may be defined as a set of periodic mappings with an offset n, that is, $\Sigma := \{\sigma \mid \sigma_1(k) = k + n \bmod K,\ n \in \mathbb{N}\}$, in which case the frequency hopping is periodically linear with respect to time. In general, the design of the strategy set is not restricted to such configurations.

Then, the radar-to-radar interference scenario can be modeled as a $\mathcal{T}$-stage repeated game $\mathcal{G} = \{\mathcal{N}, \Sigma, \{U_i\}_{i \in \mathcal{N}}\}$, where the radars, acting as players in $\mathcal{N}$, share a common strategy set Σ. The (stage-independent) utility functions $U_i : \Sigma^N \to \mathbb{R}$ endows this game with the anti-coordination property, where the sub-band sequence strategies $\sigma^i, \sigma^{-i} \in \Sigma$ employed by players are from a common strategy set, but whenever two or more strategies agree, say, $\sigma^i(k) = \sigma^j(k)$ at chirp k for some radars i and j, the utilities of radar i and j decrease. In this paper, we give the following definition for players' utility functions:

$$U_i(\sigma^i, \sigma^{-i}) = 10 \log_{10}(\text{SINR}^i_\tau), \quad \tau = 1, \ldots, \mathcal{T} \tag{1}$$

where SINR_{τ}^{i} represents the average SINR of time frame $\tau \in \{1, \ldots, \mathcal{T}\}$, at radar i's receiver, determined by the joint strategies. In the absence of interference, the SINR reduces to the signal-to-noise ratio (SNR). In the sequel, we formally define $U_{i,\tau}$

2.2 Signal Utility Model

Let radar i's strategy σ^i encapsulate a sequence of starting frequencies $(f_1^i, \ldots, f_K^i)$ and a sequence of starting time shifts $(t_1^i, \ldots, t_K^i)$. Consider a set of targets, indexed by $m = 1, \ldots, M$, within its field of view, with range vector $(r_1^i, \ldots, r_M^i)$ and velocity vector $(\dot{r}_1^i, \ldots, \dot{r}_M^i)$; note that the velocities are negative for approaching targets.

The transmitted signal at the k-th chirp is given by

$$s^i[t,k] = e^{j2\pi\left(f_k^i(t-t_k^i)+\frac{1}{2}\alpha^i(t-t_k^i)^2\right)}, \quad t_k^i \le t < t_k^i + T_a.$$

After mixing, the received signal reflected from target m can be written as

$$\begin{aligned} y_m^i[t,k] &= a_m^i\, s^i\left[t - \Delta_{m,k}^i,\, k\right] s^{i*}[t,k] \\ &\approx a_m^i\, e^{-j2\pi\left(f_k^i\,\Delta_{m,k}^i + \alpha^i\,\Delta_{m,k}^i\,(t-t_k^i)\right)}, \end{aligned}$$

where a_m^i is the complex target coefficient. The round-trip delay $\Delta_{m,k}^i$ of the k-th chirp reflected from target m is approximated by $\Delta_{m,k}^i = \frac{2}{c}\left(r_m^i + k\,\dot{r}_m^i\,T_{\mathrm{PRI}}^i\right)$, with c denoting the speed of light.

Then, the true target range for radar i at chirp k is given by $r_m^i + k\,\dot{r}_m^i\,T_{\mathrm{PRI}}^i = \bar{r}_m^i + \left(\epsilon_0 + k\,\dot{r}_m^i\,T_{\mathrm{PRI}}^i\right)$, where $\bar{r}_m^i$ denotes the coarse range center of the target with a bin width of $\frac{c}{2B_a}$, and $\epsilon_0 \in \left[-\frac{c}{4B_a}, \frac{c}{4B_a}\right]$ represents the fine quantization, having a bin width of $\frac{c}{2B_a}$. Assuming that the coarse range undergoes negligible migration [11], the received echo from target m, originally given in (2.2), can be reformulated as

$$y_m^i[t,k] = \tilde{a}_m^i\, e^{-j2\pi f_r t'}\, e^{j2\pi f_d k}\, e^{-j2\pi\left[\frac{2\bar{r}_m^i}{c} + \frac{2\left(\epsilon_0 + k\,\dot{r}_m^i\,T_{\mathrm{PRI}}^i\right)}{c}\right]\Delta f_k^i}, \tag{2}$$

where $t' = t - t_k^i$, and all constant factors have been absorbed into the complex target coefficient $\tilde{a}_m^i$. Here, the coarse range frequency is defined as $f_r = \frac{2r_m^i\,\alpha^i}{c}$, and the Doppler frequency as $f_d = -\frac{2\dot{r}_m^i\,T_{\mathrm{PRI}}^i\,f_c}{c}$. The final exponential term implements phase compensation across slow time due to frequency hopping; in this term, $\Delta f_k^i = f_k^i - f_c$ represents the known frequency shift at both the transmitter and receiver. Without frequency hopping ($\Delta f_k^i = 0$), the received signal would exhibit the conventional form; with frequency hopping, however, the additional phase term in (2) facilitates a higher range resolution.

Once the received signal is processed through a low-pass filter (LPF) and sampled by an analog-to-digital converter (ADC), a range FFT along the time index t extracts the coarse range information. Subsequently, a modified Doppler FFT applied over the index k in conjunction with the nonlinear frequency hopping sequence Δb_k^i enables the retrieval of both velocity and fine range details [11].

We assume that at radar i's receiver the target echo is mixed with signals transmitted by other radars $o \in \mathcal{N} \setminus \{i\}$, each transmitting at its own carrier frequency f_k^o. Interference occurs when an interfering radar transmits at the same frequency as radar i, i.e., when $f_k^o = f_k^i$. After demodulation and dechirping, the interference signal from radar o can be approximated by, e.g., when $t_k^o = t_k^i = 0$,

$$y_m^o[t,k] \approx \tilde{a}_m^o\, e^{j\pi(\alpha^i - \alpha^o)t^2},$$

where a_m^o is the complex interference coefficient (normalized as $\tilde{a}_m^o$) and α^o is the chirp slope of radar o. Thus, the interference appears as a new chirp with a slope given by $\alpha^i - \alpha^o$. In practice, the exact start time and frequency of the interference depend on the chirp parameters of both the victim and the interfering radars, the location of the interferer, and other conditions [5].

The overall received signal at radar i is the sum of the target signal, the interference signal, and additive white Gaussian noise:

$$\hat{\mathbf{y}}_k^i = \mathbf{y}_k^i + \mathbf{y}_k^o + \mathbf{e}_k^i.$$

Here, $\mathbf{y}_k^i$ and $\mathbf{y}_k^o$ are the low-pass filtered ADC samples corresponding to $\sum_m y_m^i[t,k]$ and $\sum_{o\in\mathcal{N}\setminus\{i\}} \sum_m y_m^o[t,k]$, respectively, while $\mathbf{e}_k^i$ denotes the noise component.

Let $\mathrm{P}(\cdot)$ be the average power function. Divide the hopping process into episodes, with the τ^{th} episode defined as $\boldsymbol{\tau} := \{k_{\tau-1}+1, \ldots, k_\tau\}, \quad \tau = 1, \ldots, \mathcal{T} \quad (k_0 = 0)$. For each episode τ, Then the theoretical average signal-to-interference-plus-noise ratio (SINR) at radar i is given by

$$\mathrm{SINR}_{\tau}^{i} = \frac{1}{|\boldsymbol{\tau}|} \sum_{k\in\boldsymbol{\tau}} \frac{\mathrm{P}(\mathbf{y}_k^i)}{\mathrm{P}(\mathbf{y}_k^o) + \mathrm{P}(\mathbf{e}_k^i)}.$$

In practice, interference detection methods—such as thresholding [5]—are used to separate the received signal $\hat{\mathbf{y}}_k^i$ into an interference component $\tilde{\mathbf{y}}_k^o$ and a clean signal $\tilde{\mathbf{y}}_k^i$. However, due to the randomness of the signals and the unknown subband choices of other radars, we estimate the SNR and SINR for each subband periodically using a sliding window. The SINR for subband f and time shift t is estimated by $\overline{\mathrm{SINR}}_{\tau}^{i}(f,t) = \frac{1}{\sum_{k\in\boldsymbol{\tau}} \mathbb{1}\{f_k^i=f, t_k^i=t\}} \sum_{k\in\boldsymbol{\tau}} \frac{\mathrm{P}(\tilde{\mathbf{y}}_k^i)\,\mathbb{1}\{f_k^i=f, t_k^i=t\}}{\mathrm{P}(\tilde{\mathbf{y}}_k^o)+\mathrm{P}(\mathbf{e}_k^i)}$, when interference is present. If interference-free chirps occur, the SNR is computed as $\overline{\mathrm{SNR}}_{\tau}^{i}(f,t) = \frac{1}{\sum_{k\in\boldsymbol{\tau}} \mathbb{1}\{f_k^i=f, t_k^i=t, \tilde{\mathbf{y}}_k^i=\mathbf{y}_k^i\}} \sum_{k\in\boldsymbol{\tau}} \frac{\mathrm{P}(\mathbf{y}_k^i)\,\mathbb{1}\{f_k^i=f, t_k^i=t\ \tilde{\mathbf{y}}_k^i=\mathbf{y}_k^i\}}{\mathrm{P}(\mathbf{e}_k^i)}$.

3. NO-REGRET HOPPING

3.1 CCE and Regret Minimization

DEFINITION 1. *For the K-stage repeated game $\mathcal{G}$, a joint probability distribution $\pi \in \mathcal{P}(\Sigma^N)$ is said to be a Coarse Correlated Equilibrium (CCE) if the following inequalities are satisfied:*

$$\mathbb{E}_{\sigma\sim\pi}[U_i(\sigma^i, \sigma^{-i})] \ge \mathbb{E}_{\sigma\sim\pi}[U_i(\sigma^{i,\prime}, \sigma^{-i})], \tag{3}$$

for all $\sigma^{i,\prime} \in \Sigma$. The interpretation is that whenever radar i deviates from a joint strategy σ sampled from a CCE π by selecting another sequence of sub-bands $\sigma^{i,\prime}$, its expected utility cannot be improved.

Computing a CCE is feasible through some learning dynamics that minimize individual regrets, a metric defined as follows.

DEFINITION 2. *Suppose the strategy sequence is $\{\sigma_\tau^i, \sigma_\tau^{-i}\}_{\tau=1}^{\mathcal{T}}$. Then, the (external) regret of hopping process for radar i is, as defined in* (4),

$$\mathcal{R}_i(\mathcal{T}) = \max_{\sigma^i\in\Sigma} \sum_{\tau=1}^{\mathcal{T}} [U_i(\sigma^i, \sigma_\tau^{-i}) - U_i(\sigma_\tau^i, \sigma_\tau^{-i})]. \tag{4}$$

The well-established result (Lemma 1) from algorithmic game theory [10] bridges the regret analysis and the CCE. Let the empirical joint distribution of strategies be

$$\overline{\pi}(\sigma^i, \sigma^{-i}) := \frac{\sum_{\tau=1}^{\mathcal{T}} \mathbb{1}\{\sigma_\tau^i = \sigma^i, \sigma_\tau^{-i} = \sigma^{-i}\}}{\mathcal{T}} \tag{5}$$

for all $(\sigma^i, \sigma^{-i}) \in \Sigma^N$.

Lemma 1. *The empirical joint distribution of frequency hopping strategies* $\bar{\pi}$ *within* $\mathcal{T}$ *time frames is a* ε*-CCE, i.e.,*

$$\mathbb{E}_{\sigma\sim\bar{\pi}}[U_i(\sigma^i,\sigma^{-i})] \geq \mathbb{E}_{\sigma\sim\bar{\pi}}[U_i(\sigma',\sigma^{-i})] - \varepsilon, \tag{6}$$

for all $\sigma' \in \mathcal{A}$*, if all radars follow the* ε*-no-regret learning dynamics, i.e.,* $\frac{1}{\mathcal{T}}\mathcal{R}_i \leq \varepsilon$ *for all* $i \in \mathcal{N}$*.*

3.2 The Algorithm

Oftentimes the overall strategy space Σ^N is prohibitively large, so instead of defining mixed strategies over Σ, we introduce two independent, time-varying mixed strategy vectors: $p^i_{f,\tau}$ for frequency selection and $p^i_{t,\tau}$ for time shifts. Under this stochastic Round-Robin policy, at the first chirp of each episode (i.e., at $k_{\tau-1}+1$), radar i samples its operating subband as $f^i_{k_{\tau-1}+1} \sim p^i_{f,\tau}(\cdot)$, so that $f^i_{k_{\tau-1}+1} = f_a$ for some subband a. For the subsequent chirps within the episode, indexed by $k_\tau + u + 1 \in \boldsymbol{\tau}$, the frequency is determined cyclically as $f^i_{k_\tau+u+1} = f_{u \bmod A_1}$. Meanwhile, at every chirp $k \in \boldsymbol{\tau}$ the time shift is independently sampled as $t^i_k \sim p^i_{t,\tau}(\cdot)$. We illustrate the hopping process in Figure 1.

Algorithm 1 `No-Regret Hopping`

Input: initialize $p^i_{f,1} = \text{unif}(\{f_1,\ldots,f_{A_1}\})$, parameters γ_τ, η_τ, loss vectors $\hat{L}^i_{0,f}(f') = 0$ for all $f' \in \{f_1,\ldots,f_{A_1}\}$, $\hat{L}^i_{0,t}(t') = 0$ for all $t' \in \{t_1,\ldots,t_{A_2}\}$. We illustrate the hopping process in 1.
for $\tau = 1 : \mathcal{T}$ **do**
    for all $i \in \mathcal{N}$ in parallel **do**
        Sample a strategy sequence with stochastic Round-Robin scheme, $(f^i_k, t^i_k)_{k\in\boldsymbol{\tau}}$;
        Sweep $[f^i_k, f^i_k + B_a]$ at $[t^i_k, t^i_k + T_a]$ for each chirp k;
        Estimate average $\overline{\text{SINR}}^i_\tau$ for all f, t;
        Update $\hat{L}^i_{\tau,f}$ and $\hat{L}^i_{\tau,t}$:

$$\hat{L}^i_{\tau,f}(f^i_k) = \hat{L}^i_{\tau-1,f}(f^i_k) - \frac{10\log_{10}(\overline{\text{SINR}}^i_\tau(f^i_k,\cdot))}{p^i_{\tau,f}(f^i_k)}$$

$$\hat{L}^i_{\tau,t}(t^i_k) = \hat{L}^i_{\tau-1,t}(t^i_k) - \frac{10\log_{10}(\overline{\text{SINR}}^i_\tau(\cdot,t^i_k))}{p^i_{\tau,t}(t^i_k)}$$

        Calculating mixed strategy:

$$p^i_{f,\tau+1}(\cdot) = \left(\frac{\exp(-\eta_\tau \hat{L}_{\tau,f})}{\sum_{f'}\exp(-\eta_\tau \hat{L}_{\tau,f}(f'))}\right)_{f\in\{f_1,\ldots,f_{A_1}\}}$$
$$p^i_{t,\tau+1}(\cdot) = \left(\frac{\exp(-\eta_\tau \hat{L}_{\tau,t})}{\sum_{t'}\exp(-\eta_\tau \hat{L}_{\tau,t}(t'))}\right)_{t\in\{t_1,\ldots,t_{A_2}\}} \tag{7}$$

    end for
end for

4. CONCLUSION

In this paper, we have extended no-regret hopping methods by expanding the strategy space along the time dimension. Our proposed approach leverages a stochastic Round-Robin policy for frequency selection, augmented with strategic time shifts such that each radar can fully exploit the available bandwidth. Future work will focus on refining the learning dynamics, conducting comprehensive assessments in realistic environments, and exploring further extensions of the strategy space.

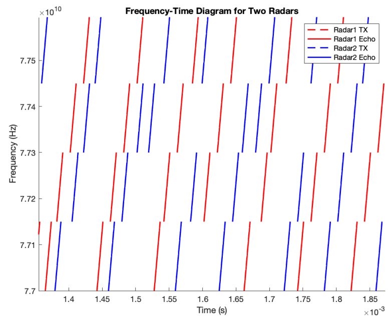

Figure 1: A scenario with two radars, four subbands, and five time slots where both radars employ a stochastic Round-Robin scheme to cyclically utilize all four subbands without interfering with each other.

5. REFERENCES

[1] S. Jin and S. Roy. FMCW radar network: Multiple access and interference mitigation. *IEEE Journal of Selected Topics in Signal Processing*, 15(4):968–979, 2021.

[2] S. Jin, P. Wang, P. Boufounos, P. V. Orlik, R. Takahashi, and S. Roy. Spatial-domain mutual interference mitigation for MIMO-FMCW automotive radar. *IEEE Transactions on Vehicular Technology*, pages 1–16, 2024.

[3] J. Khoury, R. Ramanathan, D. McCloskey, R. Smith, and T. Campbell. RadarMAC: Mitigating radar interference in self-driving cars. In *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 1–9, 2016.

[4] J. Li, R. Wu, I.-T. Lu, and D. Ren. Bayesian linear regression with cauchy prior and its application in sparse mimo radar. *IEEE Transactions on Aerospace and Electronic Systems*, 59(6):9576–9597, 2023.

[5] J. Li, J. Youn, R. Wu, J. Overdevest, and S. Sun. Performance evaluation and analysis of thresholding-based interference mitigation for automotive radar systems. In *2024 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, pages 204–208, 2024.

[6] M. Markel. *Radar for Fully Autonomous Driving*. Boston, MA: Artech House, 2022.

[7] Y. Pan, J. Li, L. Xu, S. Sun, and Q. Zhu. A game-theoretic approach for high-resolution automotive fmcw radar interference avoidance, 2025.

[8] Y. Pan, T. Li, and Q. Zhu. On the resilience of traffic networks under non-equilibrium learning. In *2023 American Control Conference (ACC)*, pages 3484–3489. IEEE, 2023.

[9] Y. Pan, T. Li, and Q. Zhu. On the variational interpretation of mirror play in monotone games. *arXiv preprint arXiv:2403.15636*, 2024.

[10] T. Roughgarden. *Algorithmic Game Theory*, volume 53. ACM New York, NY, USA, 2010.

[11] Y. Stettiner and N. Arkind. FMCW automotive radar incorporating nonlinear frequency hopping sequence of fractional bandwidth multiband chirps, Mar. 21 2023. US Patent 11,609,303.

[12] L. Xu, S. Sun, K. V. Mishra, and Y. D. Zhang. Automotive FMCW radar with difference co-chirps. *IEEE Transactions on Aerospace and Electronic Systems*, 59(6):8145–8165, 2023.