## A Hyperparameters and Infrastructure

When comparing the performance of NL-CG against DCG [4] and QMIX [22] on the MACO benchmark, we adopted the following hyper-parameter settings. For all tasks, we used a discount factor $\gamma=0.99$ and $\epsilon$-greedy exploration, where $\epsilon$ was linearly decayed from 1 to 0.05 within the first 50,000 training time steps. The replay buffer stored the last 5000 episodes, from which we uniformly sampled batches of size 32 for training. The target network was updated every 200 episodes. The learning rate of RMSprop was set to $5 \times 10^{-4}$, except for the Pursuit experiment, which used $1 \times 10^{-2}$ for faster convergence. Every 10000 time steps we paused training and evaluated the model with 300 greedy test trajectories sampled with $\epsilon=0$.
Each agent processed its local action-observation history at each time step using a linear layer of 64 neurons, followed by a ReLU activation, a GRU of the same dimensionality, and finally a linear layer of $|A|$ neurons. The output served as an individual feature vector and was fed into the utility and payoff function. The parameters of the utility and payoff function were shared amongst agents, who were identified by a one-hot encoded ID in the input. All message passing procedures in NLCG iterated for 4 rounds. The weights and bias of the mixing network in NL-CG were generated by a 2-layer hyper-network with a hidden layer of 64 ReLU neurons, except for the Aloha experiment, which uses one linear layer.
Our implementation of NL-CG was based on the PyMARL [23] framework. We used NVIDIA GeForce RTX 3090 GPUs for training and evaluation.

## B Weighted Max-Sum

Algorithm 1 and Algorithm 2 find global and approximate greedy actions for a non-linear coordination graph, respectively. Both of these two algorithms rely on the weighted Max-Sum algorithm to find local optimal action on a linear piece. We present the weighted Max-Sum algorithm in Algorithm 3

## C Representational Capability

In Sec. 4. we use a matrix game to show the representational capability of non-linear coordination graphs. For conventional coordination graphs on State $2 B$ of this game, since the value function should be permutation invariant, there are five unknowns $q_{1}=q_{i}\left(s_{2 B}, A\right), q_{2}=q_{i}\left(s_{2 B}, B\right), q_{3}=$ $q_{i j}\left(s_{2 B}, A A\right), q_{4}=q_{i j}\left(s_{2 B}, A B\right)$, and $q_{5}=q_{i j}\left(s_{2 B}, B B\right)$ and five equations:

$$
\left\{\begin{align*}
4 q_{1}+6 q_{3} & =0  \tag{10}\\
3 q_{1}+q_{2}+3 q_{3}+3 q_{4} & =-0.1 \\
2 q_{1}+2 q_{2}+q_{3}+4 q_{4}+q_{5} & =0.1 \\
q_{1}+3 q_{2}+3 q_{4}+3 q_{5} & =0.3 \\
4 q_{2}+6 q_{5} & =8
\end{align*}\right.
$$

The augmented matrix of this system has a higher rank (4) than its coefficient matrix (3). Therefore, this linear system does not have a solution, which means that a conventional coordination graph cannot represent the accurate value function for this task.
https://github.com/oxwhirl/pymarl

```
Algorithm 3 w-MAX-SUM
\(/ *\) Greedy action selection with \(k\) message passing for one linear piece of the mixing network \(f_{m}\).
This algorithm is called by Algorithm 1 and 2, and the definition of inputs can be found there. \({ }^{* /}\)
    Input: \(\boldsymbol{f}^{\mathbb{V}} \in \mathbb{R}^{|\mathcal{V}| \times A}, \boldsymbol{f}^{\mathrm{E}} \in \mathbb{R}^{|\mathcal{E}| \times A \times A},\left(\boldsymbol{W}_{\mathcal{V}}, \boldsymbol{W}_{\mathcal{E}}\right) \in \mathbb{R}^{|\mathcal{V}|+|\mathcal{E}|}, b \in \mathbb{R}\)
    \(\boldsymbol{f}^{\mathrm{V}}:=\boldsymbol{W}_{\mathcal{V}} \circ \boldsymbol{f}^{\mathrm{V}}\)
    \(\boldsymbol{f}^{\mathrm{E}}:=\boldsymbol{W}_{\mathcal{E}} \circ \boldsymbol{f}^{\mathrm{E}}\)
    \(\boldsymbol{\mu}^{0}, \overline{\boldsymbol{\mu}}^{0}:=\mathbf{0} \in \mathbb{R}^{|\mathcal{E}| \times A}\)
    /*Initialize forward messages ( \(\boldsymbol{\mu}\) ) and backward messages \((\overline{\boldsymbol{\mu}})\).*/
    \(\boldsymbol{q}^{0}:=f^{\mathrm{V}}\)
    /*Initial "Q-value".*/
    \(q_{\text {max }}:=-\infty ; \boldsymbol{a}_{\text {max }}:=\left[\underset{a \in \mathcal{A}^{i}}{\arg \max } q_{a i}^{0} \mid i \in \mathcal{V}\right]\)
    /*Initialize the best found solution.*/
    for \(t \in\{1, \ldots, k\}\) do
    /*k rounds of message passing.*/
    for \(e=(i, j) \in \mathcal{E}\) do
        /*Update forward and backward messages. Subscripts of \(\boldsymbol{f}, \boldsymbol{q}, \boldsymbol{\mu}\) mean indexing.*/
        \(\boldsymbol{\mu}_{e}^{t}:=\max _{a \in \mathcal{A}^{i}}\left\{\left(q_{i a}^{t-1}-\bar{\mu}_{e a}^{t-1}\right)+\boldsymbol{f}_{e a}^{\mathrm{E}}\right\}\)
        /*Forward: maximize sender.*/
        \(\overline{\boldsymbol{\mu}}_{e}^{t}:=\max _{a \in \mathcal{A}^{j}}\left\{\left(q_{j a}^{t-1}-\mu_{e a}^{t-1}\right)+\left(\boldsymbol{f}_{e}^{\mathrm{E}}\right)_{a}^{\top}\right\}\)
        /*Backward: maximize receiver.*/
        \(\boldsymbol{\mu}_{\boldsymbol{e}}^{t} \leftarrow \boldsymbol{\mu}_{\boldsymbol{e}}^{t}-\frac{1}{\left|\mathcal{A}^{j}\right|} \sum_{a \in \mathcal{A}^{j}} \mu_{e a}^{t}\)
        /*Normalize forward messages to ensure converging.*/
        \(\overline{\boldsymbol{\mu}}_{e}^{t} \leftarrow \overline{\boldsymbol{\mu}}_{e}^{t}-\frac{1}{\left|\mathcal{A}^{i}\right|} \sum_{a \in \mathcal{A}^{i}} \bar{\mu}_{e a}^{t}\)
        /*Normalize backward messages to ensure converging.*/
    end for
    for \(i \in \mathcal{V}\) do
        /*Update "Q-value" with messages.*/
        \(\boldsymbol{q}_{i}^{t}:=\boldsymbol{f}_{i}^{\mathrm{V}}+\sum_{e=(\cdot, i) \in \mathcal{E}} \boldsymbol{\mu}_{e}^{t}+\sum_{e=(i, \cdot) \in \mathcal{E}} \overline{\boldsymbol{\mu}}_{e}^{t}\)
        /*Utility plus incoming messages.*/
        \(a_{i}^{t}:=\arg \max \left\{q_{i a}^{t}\right\}\)
            \(a \in \mathcal{A}^{i}\)
        \(/ *\) Select the greedy action of agent \(i . * /\)
    end for
    \(q^{\prime} \leftarrow \sum_{i=1}^{|\mathcal{V}|} f_{a^{i}}^{\mathrm{V}}+\sum_{(i, j) \in \mathcal{E}} \boldsymbol{f}_{a^{i} a^{j}}^{\mathrm{E}}+b\)
    /*Get the Q -value of the greedy action.*/
    if \(q^{\prime}>q_{\text {max }}\) then
        \(\left\{\boldsymbol{a}_{\max } \leftarrow \boldsymbol{a}^{t} ; \quad q_{\max } \leftarrow q^{\prime} \boldsymbol{q} \leftarrow \boldsymbol{q}_{i}^{t}\right\}\)
        \(/ *\) Remember the best action.*/
    end if
end for
return \(q_{\max }, \boldsymbol{a}_{\max } \in \mathcal{A}^{1} \times \ldots \times \mathcal{A}^{|\mathcal{V}|}, \boldsymbol{q}\)
\(/ *\) Return the maximum \(Q\) value, the corresponding action, and utilities/payoffs.*/
```

