3.1. Definition. An *energy game* G consists of a game graph $\langle Q, E\rangle$ (with $E \subseteq Q \times Q$) and a *weight function* $w : E \to \mathbb{Z}$. The *energy level* of a play $\rho = q_0\, q_1\, q_2 \dots$ at position $k \geq 0$ is $\mathsf{EL}(\rho, k) = \Sigma_{i=0}^{k-1} w(q_i, q_{i+1})$.

A strategy σ_1 for player 1 is winning from state q with initial credit c_0 for the *energy* objective if for all strategies σ_2 the play ρ produced from q by σ_1, σ_2 is such that $c_0 + \mathsf{EL}(\rho, k) \geq 0$ for all $k \geq 0$.

The *fixed initial credit problem* asks to decide, given an energy game, an initial state q and initial credit c_0, whether there exists a winning strategy for player 1 for the energy objective.

The *unknown initial credit problem* asks to decide, given an energy game and an initial state q whether there exists an initial credit and a winning strategy for player 1 for the energy objective.

3.2. Example. Player 1 (round states) has a winning strategy in the following energy game, with initial credit 7 in q_1, 2 in q_2, 13 in q_3, 6 in q_4, and 0 in q_5.

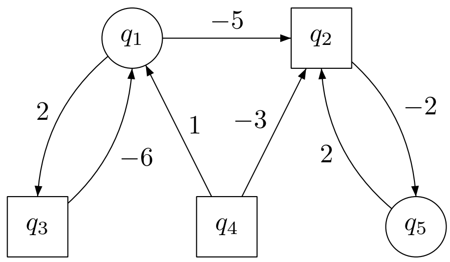

3.3. Memoryless strategies.

We will solve the fixed initial credit problem and the unknown initial credit problem by computing a function $c_0 : Q \to \mathbb{N} \cup \{\infty\}$ that gives the *minimum initial credit* necessary to win from each state of the game. The special value ∞ is assigned to losing states.

The function c_0 satisfies the following conditions:

in all player-1 states $q \in Q_1$, we have $c_0(q) + w(q, q') \geq c_0(q')$ for some q' such that $(q, q') \in E$;
in all player-2 states $q \in Q_2$, we have $c_0(q) + w(q, q') \geq c_0(q')$ for all q' such that $(q, q') \in E$.

A memoryless winning strategy σ_1 for player 1 is defined as follows. For all $q \in Q_1$ with $c_0(q) \neq \infty$, let $\sigma_1(q) = q'$ where $(q, q') \in E$ and $c_0(q) + w(q, q') \geq c_0(q')$. The strategy σ_1 is winning for player 1 from all states that have a finite minimum initial credit (why?). Note that for all strategies σ_2 for player 2, the play $\rho = q_0\, q_1\, q_2 \dots$ produced by σ_1, σ_2 is such that $c_0(q_0) + \mathsf{EL}(\rho, k) \geq c_0(q_k)$ for all $k \geq 0$. It follows that once the strategy σ_1 is fixed in G, all cycles in the resulting graph have nonnegative weight.

We show that memoryless winning strategy exists for player 2 from all states with infinite minimum initial credit. The proof is by induction on the number m of player-2 states with more than one outgoing edge. The case $m = 0$ is trivial. For $m > 0$, let $q \in Q_2$ be a state with infinite minimum initial credit ($c_0(q) = \infty$), and with outgoing edges[1] $(q, q_l) \in E$ and $(q, q_r) \in E$. Let G_l and G_r be the energy games obtained from G by deleting the edge (q, q_r) and (q, q_l) respectively, and let c_l and c_r be the corresponding minimum initial credit functions. If $c_l(q) = \infty$ or $c_r(q) = \infty$, then the result follows from the induction hypothesis. Otherwise, $c_l(q) \neq \infty$ and $c_r(q) \neq \infty$. Then, from winning strategies σ_l and σ_r for player 1 in G_l

[1] W.l.o.g. we assume that every state has at most two outgoing edges.

and G_r, we construct a winning strategy σ for player 1 in G with initial credit $\max(c_l(q), c_r(q))$, showing that this case is impossible since $c_0(q) = \infty$. Assume that $\max(c_l(q), c_r(q)) = c_l(q)$ (the other case is symmetric). The strategy σ plays according to σ_l as long as the play does not visit q. On each visit to q, if player 2 uses the edge (q, q_r), then player 1 suspends σ_l and switches to the strategy σ_r. When player 2 uses the edge (q, q_l) again, player 1 resumes playing σ_l.

In other words, if q is visited in the current play prefix ρ, and player 2 has used the edge (q, q_r) after the last visit to q, then player 1 plays like the strategy σ_r applied to the play prefix $\hat{\rho}$ obtained from ρ by taking out the cycles over q that use the edge (q, q_l). The definition is analogous if player 2 has used the edge (q, q_l) after the last visit to q.

Since σ_l and σ_r are winning strategies (in G_l and G_r respectively), on all visits to q, the energy level is at least $c_l(q)$: this holds on the first visit to q, and on the subsequent visits the energy level gets higher since all cycles produced by winning strategies are nonnegative.

Theorem 3A. The finite initial credit problem and the unknown initial credit problem are in NP $\cap$ coNP.

To see this, it suffices to guess a memoryless strategy for player 1 (resp., player 2), and to check in polynomial time that all cycles it produces are nonnegative (resp., negative).

3.4. Algorithm. The minimum initial credit function can be computed by iterating the operator $f' = \delta(f)$ on functions $f : Q \to \mathbb{N}$, starting with $f_0(q) = 0$ for all $q \in Q$, until a fixpoint is reached:

$$\delta(f)(q) = \begin{cases} \max(f(q), \min\{f(q') - w(q, q') \mid (q, q') \in E\}) & \text{if } q \in Q_1 \\ \max(f(q), \max\{f(q') - w(q, q') \mid (q, q') \in E\}) & \text{if } q \in Q_2 \end{cases}$$

The iterations for Example 3.2 are given in the following table.

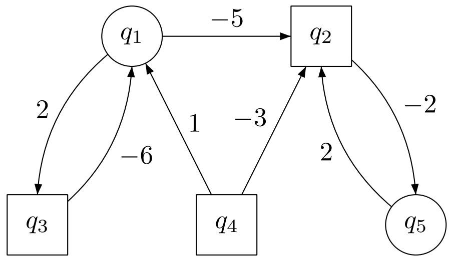

q_1	q_2	q_3	q_4	q_5	
0	0	0	0	0	
0	**2**	**6**	**3**	0	
4	2	6	**5**	0	
4	2	**10**	5	0	
7	2	6	5	0	
7	2	**13**	**6**	0	
7	2	13	6	0	$\leftarrow$ fixpoint

To ensure termination of this algorithm, we use the remark that winning strategies induce only nonnegative cycles, and thus it follows that the minimum initial credit in a (winning) state is at most $|Q| \cdot W$ where W is the largest weight (in absolute value) of the energy game. In the algorithm, values greater than $|Q| \cdot W$ are replaced by ω (such that $\omega + n = \omega$ and $\omega > n$ for all $n \in \mathbb{N}$). They correspond to losing states.

3.5. Correctness. The correctness of Algorithm 3.4 is established as follows:

1. If player 1 has a strategy to maintain the energy level nonnegative for k steps when the initial credit is given by the function f, then player 1 has a strategy to maintain the energy level nonnegative for $k+1$ steps when the initial credit is given by the function $\delta(f)$.

2. It follows that the function $f_k = \delta^k(f)$ obtained after k iterations gives the minimum initial credit necessary to maintain the energy level nonnegative during k steps.

3. Therefore, the (least) fixpoint function f such that $f = \delta(f)$ is the minimum initial credit to win the energy game.

3.6. Complexity. Each iteration of Algorithm 3.4 can be computed in time $O(|E|)$. The value of the credit function increases by 1 in at least one state (otherwise the fixpoint is reached) and it is bounded by $|Q| \cdot W$. Therefore, a naive implementation of Algorithm 3.4 has time complexity $O(|E| \cdot |Q|^2 \cdot W)$.

Consider the improved algorithm in pseudo-code given below. For $q \in Q$, let $\mathsf{pre}(q) = \{q' \mid (q', q) \in E\}$ and $\mathsf{post}(q) = \{q' \mid (q, q') \in E\}$. Efficient implementation of Algorithm 3.4 uses a list of states for which the credit value needs to be updated. When the value of a state q is updated, only the predecessors $q' \in \mathsf{pre}(q)$ such that $f(q') < f(q) - w(q', q)$ may need to be updated and included in the list. In fact all such player-2 states q' need to be updated; but if q' is a player-1 state, then q' may still not need to be updated, e.g., if some other successor r of q' satisfies $f(q') + w(q', r) = f(r)$. We use a function $\mathsf{count} : Q \to \mathbb{N}$ to remember the number of such successors, and we include q' in the list if the counter value $\mathsf{count}(q')$ drops to 0.

Since each state q is updated at most $|Q| \cdot W$ times, and each update of a state requires to scan its successors and predecessors, we get the following time complexity:

$$O\Big(\sum_{q \in Q} (|\mathsf{post}(q)| + |\mathsf{pre}(v)|) \cdot |Q| \cdot W\Big) = O(|E| \cdot |Q| \cdot W)$$

Algorithm 1: Value-iteration algorithm for energy games.

Input : A game graph $\langle Q, E \rangle$ with partition $Q_1 \cup Q_2 = Q$ and a weight function $w : E \to \mathbb{Z}$.

Output: The minimum initial credit function $f : Q \to \mathbb{N} \cup \{\omega\}$.

begin
- $L \leftarrow \{q \in Q_1 \mid \forall (q, q') \in E : w(q, q') < 0\}$;
- $L \leftarrow L \cup \{q \in Q_2 \mid \exists (q, q') \in E : w(q, q') < 0\}$;
- **foreach** $q \in Q$ **do**
  - $f(q) \leftarrow 0$;
  - **if** $q \in Q_1 \cap L$ **then** $\mathsf{count}(q) \leftarrow 0$;
  - **if** $q \in Q_1 \setminus L$ **then** $\mathsf{count}(q) \leftarrow |\{q' \in \mathsf{post}(q) \mid f(q) = f(q') - w(q, q')\}|$;
- **while** $L \neq \emptyset$ **do**
  - Pick $q \in L$;
  - $L \leftarrow L \setminus \{q\}$; $oldq \leftarrow f(q)$;
  - $f(q) \leftarrow \delta(f)(q)$;
  - **if** $f(q) \geq |Q| \cdot W$ **then** $f(q) \leftarrow \omega$;
  - **if** $q \in Q_1$ **then** $\mathsf{count}(q) \leftarrow |\{q' \in \mathsf{post}(q) \mid f(q) = f(q') - w(q, q')\}|$;
  - **foreach** $q' \in \mathsf{pre}(q)$ **such that** $f(q') < f(q) - w(q', q)$ **do**
    - **if** $q' \in Q_1$ **then**
      - **if** $f(q') \geq oldq - w(q', q)$ **then** $\mathsf{count}(q') \leftarrow \mathsf{count}(q') - 1$;
      - **if** $\mathsf{count}(q') \leq 0$ **then** $L \leftarrow L \cup \{q'\}$;
    - **if** $q' \in Q_2$ **then** $L \leftarrow L \cup \{q'\}$;
- **return** f ;

end