A Proof of the Locality Lemma

Codex based on Yanjing Wang's Slides

This note gives a fully rigorous proof of the locality lemma needed for Rosen's theorem, while keeping the same game strategy and the same bound as in Otto's sketch and in the slides:

$$\ell = 2^q - 1 \qquad \text{where } q = \mathrm{qr}(\alpha).$$

Setting and notation. We work with finite pointed Kripke models over a finite relational vocabulary consisting of one binary accessibility relation and finitely many unary predicates.

Fix a first-order formula $\alpha(x)$ and let

$$q := \mathrm{qr}(\alpha), \qquad \ell := 2^q - 1.$$

For a finite pointed model $(\mathcal{M}, w)$ let

$$\mathcal{T} := \mathrm{Unr}(\mathcal{M}, w)\restriction_\ell$$

be the depth-ℓ truncation of the full unravelling, rooted at $\langle w \rangle$. Let

$$\mathrm{PUnr}_\ell(\mathcal{M}, w)$$

be the finite partial unravelling from the slides. Thus $\mathrm{PUnr}_\ell(\mathcal{M}, w)$ is finite, bisimilar to $(\mathcal{M}, w)$, and the radius-ℓ neighbourhood of its root is exactly $\mathcal{T}$.

Distances are measured in the underlying undirected graph. If two points lie in different connected components, their distance is taken to be ∞. For a point u and a set X of nodes, let

$$\mathrm{dist}(u, X) := \min\{\mathrm{dist}(u, x) \mid x \in X\}.$$

For a set X of nodes and an integer $r \geq 0$, let

$$\mathrm{Ball}_r(X) := \{u \mid \mathrm{dist}(u, X) \leq r\}.$$

Lemma 1 (Locality lemma for Rosen's theorem)**.** *Let $\alpha(x)$ be first-order and invariant under bisimulation on finite models. Then for*

$$\ell = 2^{\mathrm{qr}(\alpha)} - 1$$

and every finite pointed model $(\mathcal{M}, w)$,

$$\mathcal{M}, w \models \alpha(x) \quad \iff \quad \mathrm{Unr}(\mathcal{M}, w)\restriction_\ell, \langle w \rangle \models \alpha(x).$$

Proof. Fix a finite pointed model $(\mathcal{M}, w)$ and abbreviate

$$\mathcal{T} := \mathrm{Unr}(\mathcal{M}, w)\restriction_\ell .$$

Step 1: the two finite comparison models. Let $\mathcal{M}^*$ be the disjoint union of

- q copies of $\mathcal{T}$, and

- $q+1$ copies of $\mathrm{PUnr}_\ell(\mathcal{M}, w)$,

and let w^* be the root of one of the $\mathrm{PUnr}_\ell(\mathcal{M}, w)$-copies.

Let $\mathcal{N}^*$ be the disjoint union of

- q copies of $\mathrm{PUnr}_\ell(\mathcal{M}, w)$, and
- $q+1$ copies of $\mathcal{T}$,

and let v^* be the root of one of the $\mathcal{T}$-copies.

Then $\mathcal{M}^*$ and $\mathcal{N}^*$ are finite, and

$$(\mathcal{M}^*, w^*) \leftrightarrow (\mathcal{M}, w), \qquad (\mathcal{N}^*, v^*) \leftrightarrow (\mathcal{T}, \langle w \rangle).$$

Indeed, $(\mathcal{M}^*, w^*)$ is just the distinguished root of a copy of $\mathrm{PUnr}_\ell(\mathcal{M}, w)$ inside a disjoint union, while $(\mathcal{N}^*, v^*)$ is literally the distinguished root of a copy of $\mathcal{T}$.

Since $\alpha(x)$ is bisimulation-invariant on finite models, it is enough to prove

$$(\mathcal{M}^*, w^*) \equiv_q^{\mathrm{FO}} (\mathcal{N}^*, v^*).$$

For then

$$\mathcal{M}, w \models \alpha \iff \mathcal{M}^*, w^* \models \alpha \iff \mathcal{N}^*, v^* \models \alpha \iff \mathcal{T}, \langle w \rangle \models \alpha.$$

Step 2: Duplicator's strategy in the q-round EF game. We describe a winning strategy for Duplicator in

$$G_q(\mathcal{M}^*, w^*; \mathcal{N}^*, v^*).$$

For $m = 0, 1, \ldots, q$ put

$$s_m := 2^{q-m}, \qquad r_m := s_m - 1 = 2^{q-m} - 1.$$

So $r_0 = \ell$ and $r_q = 0$.

The game starts from the fixed initial pebble pair (w^*, v^*). We count the q EF-rounds after this initial pair.

Clusters. After m rounds, the pebbled points on each side are partitioned into *clusters*. Initially there is just one cluster, consisting of the distinguished root. When a new pebble is played in round $m+1$:

- if it is within distance s_{m+1} of some already pebbled point in a cluster C, it joins C;
- otherwise it starts a new singleton cluster.

Because the critical distances shrink by a factor of 2, a point cannot be within distance s_{m+1} of two distinct old clusters: if it were within distance s_{m+1} of pebbles c_1, c_2 in different clusters, then

$$\mathrm{dist}(c_1, c_2) \le 2s_{m+1} = s_m,$$

contradicting the induction invariant below.

Invariant after round m. Duplicator maintains the following three conditions.

(i) The current pebble matching is a partial isomorphism on the pebbled points.

(ii) Any two pebbled points belonging to different clusters on the same side are at distance strictly greater than s_m.

(iii) For each corresponding cluster pair (C, D) there is an isomorphism

$$h_{C,D} \colon \mathcal{M}^* \restriction \mathrm{Ball}_{r_m}(C) \to \mathcal{N}^* \restriction \mathrm{Ball}_{r_m}(D)$$

which sends every pebbled point of C to its mate in D.

Base case $(m = 0)$**.** Before Spoiler makes any EF move, there is only the distinguished pair (w^*, v^*), hence only one cluster on each side. Since the radius-ℓ neighbourhood of the root in $\mathrm{PUnr}_\ell(\mathcal{M}, w)$ is exactly $\mathcal{T}$, we have

$$\mathcal{M}^* \restriction \mathrm{Ball}_{r_0}(w^*) \cong \mathcal{T} \cong \mathcal{N}^* \restriction \mathrm{Ball}_{r_0}(v^*),$$

because $r_0 = \ell$. So the invariant holds at $m = 0$.

Inductive step. Assume the invariant holds after round $m < q$. We show that Duplicator can answer so that it also holds after round $m + 1$.

By symmetry, suppose Spoiler chooses a point a in $\mathcal{M}^*$.

Case 1: the move is local. Assume that

$$\mathrm{dist}(a, C) \le s_{m+1}$$

for some old cluster C on the left. By the uniqueness remark above, this cluster is unique. Let D be the corresponding cluster on the right, and let

$$h := h_{C,D} \colon \mathcal{M}^* \restriction \mathrm{Ball}_{r_m}(C) \to \mathcal{N}^* \restriction \mathrm{Ball}_{r_m}(D)$$

be the isomorphism given by the induction hypothesis.

Since

$$r_m = s_{m+1} + r_{m+1},$$

every point within distance r_{m+1} of a lies within distance r_m of C. Hence

$$\mathrm{Ball}_{r_{m+1}}(C \cup \{a\}) \subseteq \mathrm{Ball}_{r_m}(C),$$

and in particular $a \in \mathrm{Ball}_{r_m}(C)$. Duplicator therefore responds with

$$b := h(a).$$

The new pebble pair is added to the corresponding clusters C and D.

Condition (iii) for the enlarged cluster pair is witnessed by the restriction of h to the smaller neighbourhoods

$$\mathrm{Ball}_{r_{m+1}}(C \cup \{a\}) \quad \text{and} \quad \mathrm{Ball}_{r_{m+1}}(D \cup \{b\}).$$

For all other cluster pairs, just restrict the old isomorphisms from radius r_m to radius r_{m+1}.

Condition (ii) is also preserved. Let Z be any other cluster and $z \in Z$. Pick $c \in C$ with $\mathrm{dist}(a, c) \le s_{m+1}$. By the induction hypothesis,

$$\mathrm{dist}(c, z) > s_m = 2s_{m+1}.$$

Hence

$$\mathrm{dist}(a, z) \ge \mathrm{dist}(c, z) - \mathrm{dist}(a, c) > 2s_{m+1} - s_{m+1} = s_{m+1}.$$

So the enlarged cluster remains more than s_{m+1} away from every other cluster.

Finally, condition (i) remains true. Inside the enlarged cluster this follows from the local isomorphism h; between the new pebble and pebbles in other clusters there is no edge, because those pebbles are at distance $> s_{m+1} \ge 1$ from a.

Case 2: the move is far. Assume that

$$\mathrm{dist}(a, S_m) > s_{m+1},$$

where S_m is the set of all previously pebbled points on the left. Then a is farther than s_{m+1} from every old cluster, so it starts a new singleton cluster.

Let K be the connected component of $\mathcal{M}^*$ containing a. Then K is either a copy of $\mathcal{T}$ or a copy of $\mathrm{PUnr}_\ell(\mathcal{M}, w)$. Duplicator chooses a *fresh*, previously untouched connected component K' on the right of the same isomorphism type, together with an isomorphism

$$g \colon K \to K'.$$

She responds with

$$b := g(a).$$

This is always possible: besides the distinguished components, each side has exactly q spare copies of $\mathcal{T}$ and q spare copies of $\mathrm{PUnr}_\ell(\mathcal{M}, w)$, and at most one fresh component is needed in each of the q rounds.

Condition (iii) for the new singleton cluster is now witnessed by

$$g \restriction \mathrm{Ball}_{r_{m+1}}(a) \colon \mathcal{M}^* \restriction \mathrm{Ball}_{r_{m+1}}(a) \to \mathcal{N}^* \restriction \mathrm{Ball}_{r_{m+1}}(b).$$

For all old cluster pairs, again just restrict the old isomorphisms from radius r_m to radius r_{m+1}.

Condition (ii) is immediate for the new cluster, since a is already more than s_{m+1} away from all previously pebbled points.

Condition (i) is also immediate: the new pebble a has no edge to any old pebble, because it is at distance $> s_{m+1} \geq 1$ from each of them, and the chosen point b lies in an isomorphic fresh component.

This completes the induction. So the invariant holds after every round $m \leq q$.

Step 3: Duplicator wins the q-round EF game. After the last round,

$$r_q = 0 \qquad \text{and} \qquad s_q = 1.$$

By condition (iii), for each corresponding cluster pair the matching between the pebbled points of that cluster extends to an isomorphism on those pebbled points themselves. By condition (ii), pebbled points from different clusters are at distance > 1, so there is no accessibility edge between different clusters.

Therefore the union of the clusterwise bijections is a partial isomorphism on *all* pebbled points. So Duplicator wins the q-round EF game, and hence

$$(\mathcal{M}^*, w^*) \equiv_q^{\mathrm{FO}} (\mathcal{N}^*, v^*).$$

By Step 1 and bisimulation invariance of $\alpha(x)$, we conclude

$$\mathcal{M}, w \models \alpha(x) \iff \mathrm{Unr}(\mathcal{M}, w) \restriction_\ell, \langle w \rangle \models \alpha(x).$$

This proves the locality lemma. □

Core ideas. The proof keeps the same strategy and the same bound as the slide sketch, but it uses the right invariant.

1. Use
$$\ell = 2^q - 1 = 2^{q-1} + 2^{q-2} + \cdots + 2 + 1.$$
This is exactly the total local "budget" available over the remaining rounds.

2. In round $m+1$, a Spoiler move is *local* if it lands within distance $2^{q-(m+1)}$ of an old cluster; otherwise it is *far* and starts a new cluster.

3. The correct induction does *not* require one global map on the union of all protected neighbourhoods. It only requires one map for each cluster pair, on the radius-r_m neighbourhood of that cluster.

4. Distinct clusters stay more than 2^{q-m} apart, so later rounds can never merge two old clusters. This is where the shrinking critical distances do the real work.

5. Only in the last round do we take the union of the clusterwise maps. At that stage $r_q = 0$, so each map only concerns the pebbled points themselves, and different clusters are already more than 1 apart.