

Decentralized Execution of Constraint Handling Rules for Ensembles*

Edmund S. L. Lam and Iliano Cervesato

Apr 2013

CMU-CS-13-106
CMU-CS-QTR-118

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Carnegie Mellon University, Qatar campus.

The author can be reached at sllam@qatar.cmu.edu or iliano@cmu.edu.

Abstract

CHR is a declarative, concurrent and committed choice rule-based constraint programming language. In this paper, we adapt CHR to provide a decentralized execution model for parallel and distributed programs. Specifically, we consider an execution model consisting of an ensemble of computing entities, each with its own constraint store and each capable of communicating with its neighbors. We extend CHR into CHR^e , in which rewrite rules are executed at one location and are allowed to access the constraint store of its immediate neighbors. We give an operational semantics for CHR^e , denoted ω_0^e , that defines incremental and asynchronous decentralized rewriting for the class of CHR^e rules characterized by purely local matching (0-neighbor restricted rules). We show the soundness of the ω_0^e semantics with respect to the abstract CHR semantics. We then give a safe encoding of the more general 1-neighbor restricted rules as 0-neighbor restricted rules, and discuss how this encoding can be generalized to all CHR^e programs.

* Funded by the Qatar National Research Fund as project NPRP 09-667-1-100 (Effective Programming for Large Distributed Ensembles)

Keywords: Distributed Programming, Constraint Logic Programming, Multiset Rewriting

Contents

1	Introduction	1
2	Preliminaries	2
2.1	Notations	2
2.2	CHR Language and Semantics	3
3	The CHR^e Language	4
4	Semantics of CHR^e	7
4.1	ω_α^e Abstract Semantics	7
4.2	ω_0^e Operational Semantics	8
5	Encoding 1-Neighbor Restricted Programs for ω_0^e	14
5.1	Basic Encoding Scheme	14
5.2	Optimizations	20
6	Generalized Encoding for n-Neighbor Restricted Rules	22
7	Related Works	27
8	Conclusion and Future Works	27
A	Proofs	28
A.1	Proofs for the ω_α^e Semantics	28
A.2	Proofs for the ω_0^e Semantics	30
A.3	Proofs for 1-Neighbor Restricted Rule Basic Encoding Scheme	36
A.4	Proofs for 1-Neighbor Restricted Rule Optimized Encoding Scheme	41
A.5	Proofs for n -Neighbor Restricted Rule Encoding Scheme	43

List of Figures

1	Distributed All Shortest Path	1
2	Constraint Handling Rules, Language and Semantics	3
3	Abstract syntax of CHR^e	5
4	ω_α^e Abstract Semantics of CHR^e	7
5	CHR Interpretation of CHR^e	8
6	ω_0^e Ensemble States	9
7	ω_0^e Operational Semantics for CHR^e	10
8	Abstract Ensemble State Interpretation of Operational Ensemble States	12
9	Concurrent ω_0^e Derivation Steps	13
10	Color Swapping Example: 1-Neighbor Restricted Rule	15
11	Color Propagation Example: Primary Propagated 1-Neighbor Restricted Rule	16
12	Example of Flawed Encoding of Rule with Non-persistent Propagated Head	17
13	Basic Encoding of 1-Neighbor Restricted Rules	18
14	Optimized Encoding for Neighbor Persistent Rule	20
15	Optimized Encodings for Primary Persistent Rules	21
16	Graph Sum Example: 2-Neighbor Restricted Rule	23
17	0-Neighbor Restricted Encoding for n -Neighbor Restricted Rules	24

1 Introduction

In recent years, we have seen many advances in distributed systems, multicore architectures and cloud computing, drawing more research interest into better ways to harness and coordinate the combined power of distributed computation. While this has made distributed computing resources more readily accessible to main-stream audiences, the fact remains that implementing distributed applications that can exploit such resources via traditional distributed programming methodologies is an extremely difficult task. As such, finding effective means of programming distributed systems is more than ever an active and fruitful research and development endeavor.

In this paper, we propose an extension of the constraint programming language CHR [Frü94] (Constraint Handling Rules). This language, which we call CHR^e is designed specifically as a high-level distributed and parallel programming language for developing applications that operate in a decentralized manner over an *ensemble* of distributed computing entities, referred to as locations. Each rewrite rule executes at a location, enabling this location to read and write data held by its immediate neighbors. Specifically, we are interested in rules that can read data from up to n of their immediate neighbors for various values of n , but writes to any number of neighbors. We call them *n-neighbor restricted rules*. Such n -neighbor restricted rules are a generalization of link restriction [LCG⁺06, ARLG⁺09] adapted to the context of multiset rewriting. This gives us a highly expressive model for programming complex behaviors involving distributed ensembles in a declarative and concurrent manner. Figure 1 shows an example CHR^e program that contains three rewrite

$$\begin{aligned} \text{base} &: [X] \text{edge}(Y, D) \Longrightarrow [X] \text{path}(Y, D) \\ \text{elim} &: [X] \text{path}(Y, D) \setminus [X] \text{path}(Y, D') \iff D < D' \mid \text{true} \\ \text{trans} &: [X] \text{edge}(Y, D), [Y] \text{path}(Z, D') \Longrightarrow X \neq Z \mid [X] \text{path}(Z, D + D') \end{aligned}$$

Figure 1: Distributed All Shortest Path

rules. This program computes all shortest paths of a graph in a distributed manner. An edge of length D from a location X to Y is represented by a constraint $\text{edge}(Y, D)$ found in X 's constraint store, written $[X] \text{edge}(Y, D)$. Similarly $[X] \text{path}(Y, D)$ expresses a path of length D from X to Y . The location of a constraint is modeled by the $[l]$ operator that prefixes all constraints in the rewrite rules. The rules *base* and *elim* are *0-neighbor restricted* rules because their left-hand sides involve constraints from exactly one location, X . Rule *trans* is a *1-neighbor restricted* rule since its left-hand side involves X and a neighbor Y . We designate X as the *primary* location of this rewrite rule because it references Y in an argument (here, $[X] \text{edge}(Y, D)$). This means that Y is a topological neighbor of X ¹. Neighbor restriction brings aspects of the topology of the ensemble (X has an edge to Y) in the rewrite rules. Rule *base* adds $\text{path}(Y, D)$ at location X for every instance of constraint $\text{edge}(Y, D)$ at the same location. Rule *elim* looks for any instances of a pair of constraints, $\text{path}(Y, D)$ and $\text{path}(Y, D')$ involving the same location Y at location X and removes the longer of the two path ($\text{path}(Y, D')$ for $D < D'$). Rule *trans* adds $\text{path}(Z, D + D')$ at location X whenever there is $\text{edge}(Y, D)$ at location X and a matching $\text{path}(Z, D')$ at location Y for $X \neq Z$. This program is *declarative* because the programmer focuses on *which* distributed computations to synchronize (e.g. $[X] \text{edge}(Y, D), [Y] \text{path}(Z, D')$) to produce *what* results ($[X] \text{path}(Z, D + D')$), rather than *how* synchronization is achieved. It is *concurrent* because while a rewrite rule applies to a fragment of the ensemble, many other rewritings can occur asynchronously in the rest of the ensemble.

In this technical report, we present CHR^e , a distributed programming language that extends CHR with *located constraints* and *n-neighbor restricted rewrite rules*, a generalized notion of link-restriction [LCG⁺06]. We define the ω_α^e abstract semantics of CHR^e and prove its soundness with respect to the standard semantics of CHR. Following this, we extend the original CHR refined operational semantics [DSdIBH04] to

¹Any location reference Y that appears in an *edge* predicate in this manner is considered a neighbor to X . This is how neighbor restriction enforces topological information to be embedded in predicates.

support decentralized incremental multiset matching for 0-neighbor restricted rules obtaining the ω_0^e operational semantics, and prove its soundness and the exhaustiveness of rule application, with respect to the ω_α^e semantics. We then give an optimized encoding of 1-neighbor restricted rules into 0-neighbor restricted rules of ω_0^e , prove the soundness of this encoding. Following this, we generalize this encoding work for n -neighbor restricted rules and prove its soundness.

The main contributions in this paper are as follows:

- We present CHR^e , a distributed programming language that extends CHR with *located constraints* and *n -neighbor restricted rewrite rules*, a generalized notion of link-restriction [LCG⁺06] in the context of multiset rewriting.
- We define the ω_α^e abstract semantics of CHR^e and prove its soundness with respect to the CHR abstract semantics.
- We present the ω_0^e operational semantics (based on the refined CHR operational semantics [DSdlBH04]) that supports decentralized incremental multiset matching for 0-neighbor restricted rewrite rules. We prove its soundness with respect to the ω_α^e abstract semantics.
- We give an optimized encoding of 1-neighbor restricted rules in ω_0^e and prove the soundness of this encoding.
- We generalize this encoding for n -neighbor restricted rules and prove the soundness of this encoding.

The rest of this report is organized as follows: Section 2 recalls notions used in this paper and the traditional CHR language. Section 3 defines the CHR^e language while Section 4 gives the ω_α^e abstract semantics, the ω_0^e operational semantics and relates them. Section 5 encodes 1-neighbor restricted rules into ω_0^e . In Section 6, we show the encoding of the full CHR^e language (n -neighbor restricted rules) into ω_0^e . We review related work in Section 7 and conclude in Section 8. Proof of all lemmas and theorems in this report are consolidated in Appendix A.

2 Preliminaries

In this section, we present the notations that will be used throughout this paper. We also give a brief introduction to the abstract CHR language and its semantics.

2.1 Notations

Let o be a generic syntactic construct or runtime object of our language. We write \bar{o} for a multiset of objects o and \vec{o} for a sequence of objects o . Relying on these notational accents for disambiguation, we write ‘,’ to represent both multiset union and sequence concatenation, with \emptyset as the identity element. For instance, ‘ o, \bar{o} ’ is a multiset while ‘ o, \vec{o} ’ is a sequence (the former is commutative while the latter is not). A set is treated as a multiset with single occurrences of each element. We use standard notations for sets: $\bar{o}_1 \cup \bar{o}_2$ for union, $o \in \bar{o}$ for membership, and $\bar{o}_1 \subseteq \bar{o}_2$ for subset. Given a set \mathcal{I} of labels, we write $\biguplus_{i \in \mathcal{I}} o_i$ to denote the multiset of objects o_i , for $i \in \mathcal{I}$. It is inductively defined as follows:

$$\left(\biguplus_{i \in \{j\} \cup \mathcal{I}} o_i\right) = o_j, \left(\biguplus_{i \in \mathcal{I}} o_i\right) \quad \left(\biguplus_{i \in \emptyset} o_i\right) = \emptyset$$

The substitution of all occurrences of variable x in o with the term expression t is denoted as $[t/x]o$. It is inductively defined on all syntactic constructs and runtime objects o in the usual manner. We assume that substitution is capture-avoiding and rely on implicit α -renaming. Given a sequence of terms \vec{t} and a sequence of variables \vec{x} , we write $[\vec{t}/\vec{x}]o$ as the simultaneous substitution of each variable in \vec{x} with the corresponding term in \vec{t} . We will use this notion with sets as well (i.e. \bar{t} and \vec{x}). We write $\text{FV}(o)$ for the set of free-variables

CHR Abstract Syntax

	Variables x	Values v	Predicate names p	Rule names r	Rule guard G
Term		$t ::= x \mid v \mid \dots$			
Constraint		$c ::= p(\vec{t})$		Rewrite Rule $R ::= r : H \setminus H \iff G \mid B$	
Rule Heads	$H ::= \cdot \mid c, H$			Program $\mathcal{P} ::= R \mid R \mathcal{P}$	
Rule Body	$B ::= \exists \bar{x}. D$			Store $\bar{S} ::= \emptyset \mid \bar{S}, c$	
Rule Body Elements	$D ::= true \mid c, D$				

Simplification and Propagation Rule Short-hands

Simpagation Rule	$\left\{ \begin{array}{l} r : P \setminus S \iff G \mid B \\ r : P \setminus S \iff B \end{array} \right. = r : P \setminus S \iff true \mid B$
Propagation Rule	$\left\{ \begin{array}{l} r : P \implies G \mid B \\ r : P \implies B \end{array} \right. = r : P \setminus \cdot \iff G \mid B$ $= r : P \setminus \cdot \iff true \mid B$
Simplification Rule	$\left\{ \begin{array}{l} r : S \iff G \mid B \\ r : S \iff B \end{array} \right. = r : \cdot \setminus S \iff G \mid B$ $= r : \cdot \setminus S \iff true \mid B$

 ω_α Abstract Semantics of CHR

$$\frac{r : P' \setminus S' \iff G \mid B \in \mathcal{P} \quad \models \theta G \quad P = \theta P' \quad S = \theta S'}{\mathcal{P} \triangleright (\bar{S}, P, S) \mapsto_{\omega_\alpha} (\bar{S}, P, \text{NF}(\text{Inst}(\theta B)))}$$

Figure 2: Constraint Handling Rules, Language and Semantics

in o and say that o is *ground* if $\text{FV}(o) = \emptyset$. We write meta level operators in upright font (e.g. $\text{FV}(-)$) and CHR constraint predicates in italics (e.g. *edge*, *path*).

2.2 CHR Language and Semantics

Figure 2 illustrates the abstract syntax of CHR. CHR is a high-level multiset rewriting language built on top of a *term language*, which we will keep mostly abstract. We shall only assume that it has variables and that each term t has a normal form, denoted $\text{NF}(t)$. A rule guard G is a set of relations among term expressions, called *built-in constraints*. We assume that built-in constraints contain equality. The judgment $\models G$ decides the validity of ground built-in constraint G .

A CHR constraint $p(\vec{t})$ is a first-order predicate symbol p applied to a sequence of terms \vec{t} . A CHR rule $r : P \setminus S \iff G \mid B$ is a rewrite rule with a unique name r , such that P , S and B are multisets of CHR constraints and G is a set of built-in constraints. Each CHR rule specifies a rewriting that can occur over fragments of a multiset of constraints, known as the *CHR store*. Specifically, a rule like this states that we can replace any matching instance of ‘ P, S ’ in the CHR store with the corresponding instance of B , if G is valid when applied to the matching substitution ($\models \theta G$). We denote *true* as the identity built-in constraint such that $\models true$ is universally true. We will refer to P and S as the *propagate* and *simplify* rule heads, G as the rule *guard* and B as the *rule body*. This general form of CHR rule is known as a *simpagation rule*.

The short-hands $r : P \implies G \mid B$ and $r : S \iff G \mid B$, known as *propagation* and *simplification* rules, have empty simplified and propagated rule heads respectively. We will omit the rule guard component if it is empty as well. We take a few benign deviations from traditional treatments of CHR (e.g. [Frü94]): we require that all constraints in a CHR store be ground and that built-in constraints only appear in the guards G . This allows us to focus on the distributed multiset rewriting problem in this paper and set aside features like explicit built-in constraints as orthogonal extensions. Also, a rule body B can be prefixed by zero or more existential variable declarations ($\exists \bar{x}$). Such variables, which do not appear in the scope of the rule heads, are instantiated to new constants (that do not appear in the store) during rule application, allowing the creation of new *destinations* in destination passing-style programming [Pfe04]. For a rule body $\exists \bar{x}. D$ such that \bar{x} is an empty set, we simply write it as D . A CHR rule $r : P \setminus S \iff G \mid B$ is *well-formed* if rule guards and rule body are grounded by the rule heads ($\text{FV}(G) \cup \text{FV}(B) \subseteq \text{FV}(P) \cup \text{FV}(S)$ ²) and all term expressions that appear in the rule are well-formed. A CHR program \mathcal{P} is *well-formed* if all CHR rules in \mathcal{P} are well-formed and have a unique rule name. A CHR store \bar{S} is a multiset of constraints. It is *well-formed* if $\text{FV}(\bar{S}) = \emptyset$ and all terms that appear in constraints are well-formed.

We inductively extend the normalization function to rule bodies. Given a rule body that contains no existentials D , $\text{NF}(D)$ denotes the normalized rule body with all term expressions in D in normal form. Since this operation only applies to rule bodies that contain no existentials, we define a complimentary meta operator that instantiates existential variables to fresh constants: Given a rule body $\exists \bar{x}. D$, $\text{Inst}(\exists \bar{x}. D)$ denotes an instance of D such that each existential variable \bar{x} that appear in D is replaced by a fresh constant a . These meta operation is inductively defined as follows:

$$\begin{array}{l} \text{Normalize} \quad \left\{ \begin{array}{l} \text{NF}(p(\vec{t}), D) = p(\text{NF}(\vec{t}), \text{NF}(D)) \\ \text{NF}(\text{true}) = \text{NF}(\text{true}) \end{array} \right. \\ \\ \text{Instantiate} \quad \text{Inst}(\exists \bar{x}. D) = [\bar{a}/\bar{x}]D \quad \text{for each } a \in \bar{a} \text{ is a fresh constant} \end{array}$$

Figure 2 also shows the abstract semantics ω_α of CHR. It defines the transitions of well-formed CHR stores via the derivation step $\mathcal{P} \triangleright \bar{S} \mapsto_{\omega_\alpha} \bar{S}'$ for a given well-formed CHR program \mathcal{P} . A derivation step models the application of a CHR rule: it checks that rule heads P' and S' match fragments P and S of the store under a matching substitution θ and that the corresponding instance of guard G is valid ($\models \theta G$), resulting to the rewrite of S with θB with existential variables instantiated and then all term expressions evaluated to normal form (i.e. $\text{NF}(\text{Inst}(\theta B))$). Derivation steps preserves the well-formedness of constraint stores. We denote the reflexive and transitive application of derivation steps as $\mathcal{P} \triangleright \bar{S} \mapsto_{\omega_\alpha}^* \bar{S}'$. A CHR store \bar{S} is *terminal* for \mathcal{P} if no derivation steps of ω_α applies in \bar{S} .

While the semantics in Figure 2 models CHR rewritings abstractly, actual CHR implementations implement some variant of the refined operational semantics [DSdlBH04]. This semantics computes new matches incrementally from newly added constraints. Additionally, rule heads are implicitly ordered by an *occurrence index* (typically in textual order of appearance) and matches are attempted in that order, thus providing programmatic idioms that assume a textual ordering of rule application. More details of these will be provided in Section 4.2.

3 The CHR^e Language

In this section, we introduce the CHR^e language. CHR^e extends CHR in several ways and here we focus on the syntactic ramifications of these extensions. *Locations* l in CHR^e rewriting rules are name annotations to a CHR constraint c , written as $[l]c$ and read as ‘ c is located at l ’. We call $[l]$ the *localization operator*. Locations can be variables and in this case are subjected to substitution and other free variable meta operations as if they were an additional argument in constraints. Operationally, $[l]c$ indicates that constraint c is held at location l (details in Section 4.2). Figure 3 shows the abstract syntax of CHR^e rules and programs.

²Recall that the existential variables are not free.

Location names k			
Location	$l ::= x \mid k$		
Rule Heads	$H ::= \cdot \mid [l]c, H$	Rewrite Rule	$R ::= r : H \setminus H \iff G \mid B$
Rule Body	$B ::= \exists \bar{x}. D$	Program	$\mathcal{P} ::= R \mid R \mathcal{P}$
Rule Body Elements	$D ::= true \mid [l]c, D$		

Figure 3: Abstract syntax of CHR^e

All constraints in a rule are now explicitly localized by the operator $[l]$. A localization operator in a rule head indicates the location where the constraint is to be matched, while a localization operator in a rule body indicates the location where that constraint is to be delivered. We call the former locations *matching locations* of the rule while the latter are *forwarding locations*.

CHR^e rules and programs are *well-formed* similarly to CHR rules and programs. Like all variables, location variables of rule bodies must appear as term arguments or localization operators of rule heads, or otherwise be existentially quantified. These “existential locations” represent new locations to be created during rule application and allows the specification of dynamically growing ensembles. We will only consider well-formed CHR^e rules from now on. Given a multiset of located constraints H , $\text{Locs}(H)$ denotes the set of locations that appear in the localization operators of H (e.g. $\text{Locs}([X]a(Y, Z), [Y]b(Z, 2)) = \{X, Y\}$) and $\text{Args}(H)$ denotes the set of all term arguments of constraints in H (e.g. $\text{Args}([X]a(Y, Z), [Y]b(Z, 2)) = \{Y, Z, 2\}$). We write $H|_l$ for the *location restriction* on H , which denotes the multiset of all constraints in H that is located at l (e.g. $([X]a(Y, Z), [Y]b(Z, 2))|_X = a(Y, Z)$). We write $[l]H$ for the multiset of all constraints in H each prefixed with $[l]$. These operators are inductively defined as follows:

$$\begin{array}{l}
 \text{Retrieve Locations} \\
 \text{Retrieve Arguments} \\
 \text{Location Restrict (Rule Heads)}
 \end{array}
 \left\{ \begin{array}{l}
 \text{Locs}([l]c, H) = l \cup \text{Locs}(H) \\
 \text{Locs}(\cdot) = \emptyset \\
 \\
 \text{Args}([l]p(\vec{t}), H) = \vec{t} \cup \text{Args}(H) \\
 \text{Args}(\cdot) = \emptyset \\
 \\
 ([l]c, H)|_l = c, H|_l \\
 ([l']c, H)|_l = H|_l \quad \text{if } l \neq l' \\
 \cdot|_l = \emptyset
 \end{array} \right.$$

Link restriction in distributed rule based languages [CARG⁺12, LCG⁺06] constrains how locations are used. We generalize it to the notion of n -neighbor restriction. A CHR^e rule $r : P \setminus S \iff G \mid B$ is *n -neighbor restricted* (where $n = |\text{Locs}(P, S)| - 1$) if we can select $l \in \text{Locs}(P, S)$ such that it is *directly connected* to each other n locations that appear in the rule heads. Furthermore, rule heads of locations other than l are *isolated* in that they do not contain common variables that do not appear at l . Rule guards also need to be *isolated* in a manner such that each atomic rule guard is grounded by the rule heads of location l and at most one other location of the rule head. Such a matching location l is called the *primary location* of the n -neighbor restricted rule, while all other matching locations of the rule are called *neighbor locations*. We assume that rule guards G have no side-effects and hence are pure boolean assertions. Formally, a CHR^e rule $r : P \setminus S \iff G \mid B$ is *n -neighbor restricted* (for $n = |\text{Locs}(P, S)| - 1$) if there exists $l \in \text{Locs}(P, S)$ such that the following three conditions are satisfied:

- *Directly connected*: $\text{Locs}(P, S) \subseteq \{l\} \cup \text{Args}((P, S)|_l)$

- *Neighbor isolated rule heads*: For all other distinct $l', l'' \in \text{Locs}(P, S)$, for each $x \in \text{FV}((P, S)_{|l'}, (P, S)_{|l''})$, we have $x \in \text{FV}((P, S)_{|l})$.
- *Neighbor isolated rule guards*: For each guard condition $g \in G$, either we have $\text{FV}(g) \subseteq \text{FV}((P, S)_{|l})$ or for some other $l' \in \text{Locs}(P, S)$, we have $\text{FV}(g) \in (\text{FV}((P, S)_{|l}) \cup \text{FV}((P, S)_{|l'}))$.

n -neighbor restricted rules characterizes multiset rewritings across constraint stores of $n + 1$ connected locations in a ‘star’ topology with the *primary* location in the center, directly connected to each n neighbors. The neighbor isolation conditions (for rule heads and guards) are defined so as to make the matching problem decomposable into partial match problems between the primary location l and each of its n -neighbors separately (details in Section 6). In general, an n -neighbor restricted rule is of the following form:

$$r : (\biguplus_{i \in \mathcal{I}_n} [k_i] P_i) \setminus (\biguplus_{i \in \mathcal{I}_n} [k_i] S_i) \iff G \mid \exists \bar{x}. (\biguplus_{i \in \mathcal{I}_n} [k_i] D_i), (\biguplus_{j \in \mathcal{I}_m} [k_j] D_j), (\biguplus_{l \in \mathcal{I}_e} [k_l] D_l)$$

where k_i for $i \in \mathcal{I}_n$ are *matching locations*.

k_j for $j \in \mathcal{I}_m$, $k_j \in (\text{FV}(\biguplus_{i \in \mathcal{I}_n} P_i) \cup \text{FV}(\biguplus_{i \in \mathcal{I}_n} S_i))$ are *non-matched forwarding locations*.

k_l for $l \in \mathcal{I}_e$, $k_l \in \bar{x}$ are *existential forwarding locations*.

The matching locations of the rule r are the set of locations k_i labeled by $i \in \mathcal{I}_n$. These are the locations that will be involved in rule matching and correspond to the set of locations $\text{Locs}((\biguplus_{i \in \mathcal{I}_n} [k_i] P_i, [k_i] S_i))$. We call each “ P_i, S_i ” belonging to a location k_i the *matching obligations* of k_i and assume that each location k_i has a non-empty matching obligation (i.e, either $P_i \neq \emptyset$ or $S_i \neq \emptyset$). All locations that appear in localization operators of the right-hand side of rule r are called forwarding locations. We classify them into three types: first we have *matched forwarding locations* in the rule body $(\biguplus_{i \in \mathcal{I}_n} [k_i] D_i)$ are such that each k_i is a matching location as well. Next, *non-matched forwarding locations* within $(\biguplus_{j \in \mathcal{I}_m} [k_j] D_j)$ are such that k_j is not a matching location, but appear as a term argument of some constraint in the rule head³. Finally, *existential forwarding locations* within $(\biguplus_{l \in \mathcal{I}_e} [k_l] D_l)$ are such that $k_l \in \bar{x}$, the set of existential variables. Note that this implicitly means that k_l neither is a matching location nor appear as a term argument of a rule head, hence it is a reference to a *new* location. We assume that the body of matched forwarding locations (i.e, D_i) may be empty but that of non-matched and existential forwarding locations must be non-empty (i.e, D_j and D_l). This assumption provide a more concise notation for n -neighbor restricted rules without superfluous references to locations that are not involved in the rule application⁴. We will refer to a n -neighbor restricted program as a general CHR^e program, or simply a CHR^e program. A well-formed CHR^e rule where heads are located at exactly one location are, by definition, 0-neighbor restricted. We call these *local rules*. A CHR^e program \mathcal{P} is *n-neighbor restricted* if each rule in \mathcal{P} are m -neighbor restricted for $m \leq n$.

As an example, consider the example program in Figure 1. Rules *base* and *elim* are examples of 0-neighbor restricted rules or local rules. This is because the each have rule heads that specify constraints from a single location. Rule *trans* however, has rule heads from two distinct locations X and Y . Specifically, rule heads $P = [X] \text{edge}(Y, D), [Y] \text{path}(Z, D')$ and $S = \emptyset$, satisfying the directly connected condition because we have $\text{edge}(Y, D)$ located at X , while satisfying the neighbor isolation condition by default because the rule only has one neighboring matching location Y . Hence it is a 1-neighbor restricted rule. An interested observation is that the 1-neighbor restriction corresponds directly to the link restriction of [CARG⁺12, LCG⁺06].

³The *directly connected* condition of neighbor restriction dictates that the primary matching location would possess one such constraint.

⁴Note that a matched forwarding location k_i is allowed to have an empty rule body D_i because its presence is justified by its appearance as a matching location.

$$\begin{array}{l}
\text{Constraint Store } \bar{S} ::= \emptyset \mid \bar{S}, c \\
\text{Abstract Ensemble } \mathcal{A} ::= \emptyset \mid \langle \bar{S} \rangle_k, \mathcal{A} \\
\\
\begin{array}{l}
r : (\bigsqcup_{i \in \mathcal{I}_n} [k'_i] P'_i) \setminus (\bigsqcup_{i \in \mathcal{I}_n} [k'_i] S'_i) \iff G \mid \exists \bar{x}. \bar{D} \in \mathcal{P} \\
\vdash \theta G \quad (\bigsqcup_{i \in \mathcal{I}_n} P_i) = \theta (\bigsqcup_{i \in \mathcal{I}_n} P'_i) \quad (\bigsqcup_{i \in \mathcal{I}_n} S_i) = \theta (\bigsqcup_{i \in \mathcal{I}_n} S'_i) \quad (\bigsqcup_{i \in \mathcal{I}_n \cup \mathcal{I}_m} k_i) = \theta (\bigsqcup_{i \in \mathcal{I}_n \cup \mathcal{I}_m} k'_i) \\
\left(\bigsqcup_{i \in \mathcal{I}_n} [k_i] D_i, \bigsqcup_{j \in \mathcal{I}_m} [k_j] D_j, \bigsqcup_{l \in \mathcal{I}_e} [k_l] D_l \right) = \text{NF}(\text{Inst}(\theta(\exists \bar{x}. \bar{D})))
\end{array} \\
\hline
\mathcal{P} \triangleright \mathcal{A}, \left(\begin{array}{l} \bigsqcup_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, S_i \rangle_{k_i}, \\ \bigsqcup_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j} \end{array} \right) \mapsto_{\omega_\alpha^e} \mathcal{A}, \left(\begin{array}{l} \bigsqcup_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, D_i \rangle_{k_i}, \\ \bigsqcup_{j \in \mathcal{I}_m} \langle \bar{S}_j, D_j \rangle_{k_j}, \\ \bigsqcup_{l \in \mathcal{I}_e} \langle D_l \rangle_{k_l} \end{array} \right) \\
\\
\text{where } \begin{array}{l}
k'_j \text{ for } j \in \mathcal{I}_m \quad \text{such that } k'_j \in (\text{FV}(\bigsqcup_{i \in \mathcal{I}_n} P_i, \bigsqcup_{i \in \mathcal{I}_n} S_i)) \\
k'_l \text{ for } l \in \mathcal{I}_e \quad \text{such that } k'_l \in \bar{x} \\
\bar{D} = \left(\bigsqcup_{i \in \mathcal{I}_n} [k'_i] D'_i \right), \left(\bigsqcup_{j \in \mathcal{I}_m} [k'_j] D'_j \right), \left(\bigsqcup_{l \in \mathcal{I}_e} [k'_l] D'_l \right)
\end{array}
\end{array}$$

Figure 4: ω_α^e Abstract Semantics of CHR^e

4 Semantics of CHR^e

We define the ω_α^e abstract semantics of CHR^e in Section 4.1, which introduces a decentralized execution of CHR^e programs, and prove its soundness with respect to the ω_α semantics. Using ω_α^e as a stepping stone, Section 4.2 defines and proves the soundness of the ω_0^e operational semantics that provides a more operational view of decentralized execution of CHR^e programs. Although this operational semantics only supports 0-neighbor restricted rules, Section 5 will show how we encode the arbitral CHR^e programs into 0-neighbor restricted programs. Note that in this work, we assume a lossless network, meaning that communication between locations are always eventually delivered and never lost.

4.1 ω_α^e Abstract Semantics

In this section, we introduce ω_α^e , an abstract decentralized semantics for CHR^e. This semantics accounts for the distributed nature of CHR^e, where each location has its own constraint store. This abstract semantics models a state transition system between abstract ensemble states, which are multisets of local stores $\langle \bar{S} \rangle_k$ where \bar{S} is a constraint store and k a location name. An abstract ensemble states \mathcal{A} is well-formed if all constraint stores \bar{S} that appear in it are well-formed and location names k are unique.

Figure 4 shows the ω_α^e semantics. Given a CHR^e program \mathcal{P} , we write a derivation step of ω_α^e as $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e} \mathcal{A}'$ for abstract states $\mathcal{A}, \mathcal{A}'$. A derivation step defines the application of an n -neighbor restricted rules: Each of the $n + 1$ locations k_i for $i \in \mathcal{I}_n$ provides a partial match P_i and S_i in their respective stores to their respective matching obligations (i.e, P'_i and S'_i), resulting in the combined substitution θ . If guard θG is valid, we apply the rule instance by removing S_i from the respective store of matching location k_i and replace them with the respective normalized rule body fragment (D_i of matching location k_i). For non-matching forwarding locations k_j for $j \in \mathcal{I}_m$, we simply add rule body fragment D_j to their stores. Finally, for existential forwarding locations, we create new location names k_l that contains just the rule body fragments D_l ⁵. Since we assume that rule guards G have no side effects, $\vdash \theta G$ is a global assertion. Note that 0-neighbor restriction is the special case where matching is localized (\mathcal{I}_n is a singleton set). We denote the reflexive and transitive application of derivation steps by $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$. ω_α^e derivation steps

⁵New location names are implicitly created by the instantiation of existential variables (i.e, $\text{Inst}(-)$) applied to the rule body after substitution θ is applied.

Abstract Ensembles	$\llbracket \mathcal{A}, \langle \bar{\mathcal{S}} \rangle_k \rrbracket = \llbracket \mathcal{A} \rrbracket, \llbracket \bar{\mathcal{S}} \rrbracket^k$	$\llbracket \emptyset \rrbracket = \emptyset$
CHR Program	$\llbracket R \mathcal{P} \rrbracket = \llbracket R \rrbracket \llbracket \mathcal{P} \rrbracket$	$\llbracket \cdot \rrbracket = \cdot$
CHR Rule	$\llbracket r : P \setminus S \iff G \mid B \rrbracket = r : \llbracket P \rrbracket \setminus \llbracket S \rrbracket \iff G \mid \llbracket B \rrbracket$	
Constraints	$\llbracket p(\vec{t}) \rrbracket^l = p(l, \vec{t})$	Body $\left\{ \begin{array}{l} \llbracket \exists \bar{x}. D \rrbracket = \exists \bar{x}. \llbracket D \rrbracket \\ \llbracket [l]c, D \rrbracket = \llbracket c \rrbracket^l, \llbracket D \rrbracket \\ \llbracket true \rrbracket = true \end{array} \right.$
Stores	$\left\{ \begin{array}{l} \llbracket c, \bar{\mathcal{S}} \rrbracket^l = \llbracket c \rrbracket^l, \llbracket \bar{\mathcal{S}} \rrbracket^l \\ \llbracket \emptyset \rrbracket^l = \emptyset \end{array} \right.$	Head $\left\{ \begin{array}{l} \llbracket [l]c, H \rrbracket = \llbracket c \rrbracket^l, \llbracket H \rrbracket \\ \llbracket \cdot \rrbracket = \cdot \end{array} \right.$

Figure 5: CHR Interpretation of CHR^e

preserves the well-formedness of states \mathcal{A} , provided that CHR^e programs are well-formed. Lemma 1 states this property of ω_α^e derivations.

Lemma 1 (Well-Formedness Preservation of ω_α^e Derivations) *Given a well-formed CHR^e program \mathcal{P} , a well-formed state \mathcal{A} and state \mathcal{A}' , if $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ then \mathcal{A}' is well-formed.*

From here on, we will implicitly assume the well-formedness of \mathcal{P} , \mathcal{A} and \mathcal{A}' , when writing $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$.

We now relate the CHR^e abstract semantics ω_α^e to the CHR abstract semantics ω_α . Figure 5 defines a function $\llbracket - \rrbracket$ that inductively traverses the structure of a CHR^e syntactic construct and translates it into a CHR construct that represents its CHR interpretation. For a CHR^e rule, this function translates located constraints $[l]p(\vec{t})$ to standard CHR constraints by inserting location l as the first (leftmost) term argument of the predicate (i.e., $p(l, \vec{t})$). We call these *location interpreted CHR constraints*. An abstract ensemble state \mathcal{A} is interpreted in CHR simply by collapsing all constraint stores in \mathcal{A} into a single global constraint store, containing location interpreted CHR constraints. The translation function is defined so that given a well-formed CHR^e syntactic object o , $\llbracket o \rrbracket$ is a well-formed syntactic object of CHR. Lemma 2 states this property of the translation function $\llbracket - \rrbracket$.

Lemma 2 (Well-Formedness Preservation of $\llbracket - \rrbracket$ Translation) *Given a well-formed CHR^e object o , $\llbracket o \rrbracket$ is a well-formed CHR object.*

Theorem 3 states the soundness of ω_α^e , namely the translation function $\llbracket - \rrbracket$ preserves derivability.

Theorem 3 (Soundness of ω_α^e) *Given a CHR^e program \mathcal{P} and abstract states \mathcal{A} and \mathcal{A}' , if $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$, then $\llbracket \mathcal{P} \rrbracket \triangleright \llbracket \mathcal{A} \rrbracket \mapsto_{\omega_\alpha}^* \llbracket \mathcal{A}' \rrbracket$.*

4.2 ω_0^e Operational Semantics

This section introduces the operational semantics ω_0^e . Similarly to ω_α^e , this semantics specifies a distributed execution for CHR^e programs. Unlike ω_α^e , it is *operational* in that it describes the execution of CHR^e programs in a procedural manner with a clear and concise execution strategy for each location of the ensemble. As such, it comprises of more derivation rules, each of which is dedicated to a specific sub-task of decentralized multiset rewriting. The ω_0^e semantics is an extension of the refined CHR operational semantics [DSdlBH04], adapted to describe an execution model for decentralized and incremental multiset matching⁶.

⁶It is incremental in that multiset matches are processed incrementally from new information (constraints). This is a property that is crucial for any effective execution model for distributed rule-based systems.

	Stored constraint id d		Rule Occurrence Index i
Goals	$g ::= b \mid c \mid c\#d : i$	Goals	$\vec{G} ::= \emptyset \mid g, \vec{G}$
Ids	$\vec{D} ::= \emptyset \mid d, \vec{D}$	Numbered Store	$\vec{S} ::= \emptyset \mid \vec{S}, c\#d$
Buffers	$\vec{U} ::= \emptyset \mid c, \vec{U}$	History	$\vec{H} ::= \emptyset \mid \vec{H}, (\vec{D})$
Operational Ensembles $\Omega ::= \emptyset \mid \Omega, \langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k$			
Retrieve Locations	$\begin{cases} \text{Locs}(\Omega, \langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k) & = & \text{Locs}(\Omega) \cup \{k\} \\ \text{Locs}(\emptyset) & = & \emptyset \end{cases}$		

Figure 6: ω_0^e Ensemble States

As such, it shares many meta constructs with the refined CHR operational semantics (e.g. rule head occurrence index, goals, numbered constraints, history). We refer the interested reader to [DSdlBH04] for a more detailed treatment of the refined operational semantics of CHR. The ω_0^e semantics applies only to 0-neighbor restricted rules. While we could easily have included n -neighbor restricted rule execution as a derivation step of ω_0^e , doing so would not capture the operational challenges of synchronizing multiple locations during rule application. Instead we compile n -neighbor restricted rules for $n > 1$ into 0-neighbor restricted rules in Section 5.

Figure 6 defines the states of the ω_0^e semantics. An ω_0^e ensemble Ω is a set of tuples of the form $\langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k$ which represent the state of the computing entity at location k . The *buffer* \vec{U} is a sequence of the constraints that have been sent to location k . The *goals* \vec{G} is a sequence of the constraints c or active constraints $c\#d : i$. The *numbered store* \vec{S} is a multiset of numbered constraints $c\#d$. The index d serves as a reference link between $c\#d$ in the store and an active constraint $c\#d : i$ in the goals. The *history* \vec{H} is a set of indices where each element is a unique set of constraint ids (\vec{D}) . Note the use of accents to explicitly indicate the nature of each collection of objects. Buffers is a novel extension of this semantics while goals, numbered store and history are artifacts of the refined operational semantics of CHR. Also, like the refined operational semantics of CHR, we assume that each rule head in a CHR program have a unique rule occurrence index i , representing the sequence (typically textual order of appearance) in which the rule head is matched to an active constraint. We write occurrence index i as a subscript of the rule head (i.e., $[i]c_i$). A state Ω is well-formed if each $\langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k \in \Omega$ has a unique location name k , all objects in \vec{U} , \vec{G} and \vec{S} are ground and all term expressions that appear in them are well-formed. Furthermore we require that each goals \vec{G} has at most one active constraint ($c\#d : i$) found at the head of \vec{G} with a corresponding $c\#d$ found in the store \vec{S} . Ω is *initial* if all $\vec{U} = \emptyset$, $\vec{S} = \emptyset$ and $\vec{H} = \emptyset$ and *terminal* if all $\vec{U} = \emptyset$ and $\vec{G} = \emptyset$. We extend the meta operation $\text{Locs}(-)$ to the domain of operational states Ω , namely that $\text{Locs}(\Omega)$ returns the set of all locations k that appear in Ω .

We define three meta operations that will be used in the ω_0^e semantics: Given a CHR numbered store \vec{S} , $\text{DropIds}(\vec{S})$ returns the multiset of all constraints in \vec{S} without their numbered ids; $\text{Ids}(\vec{S})$ returns the set of all constraint ids that appear in \vec{S} ; given a CHR program \mathcal{P} , $\text{OccIds}(\mathcal{P})$ returns the set of all rule head occurrence indices that appear in each rule $R \in \mathcal{P}$. These meta operations are inductively defined as follows:

$$\begin{array}{l}
\text{(Flush)} \quad \frac{\vec{U} \neq \emptyset}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; \emptyset; \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \emptyset; \vec{U}; \bar{S}; \bar{H} \rangle_k} \\
\text{(Loc 1)} \quad \frac{}{\mathcal{P} \triangleright \Omega, \left(\begin{array}{l} \langle \vec{U}; ([k']c, \vec{G}); \bar{S}; \bar{H} \rangle_k, \\ \langle \vec{U}'; \vec{G}'; \bar{S}'; \bar{H}' \rangle_{k'} \end{array} \right) \mapsto_{\omega_0^e} \Omega, \left(\begin{array}{l} \langle \vec{U}; \vec{G}; \bar{S}; \bar{H} \rangle_k, \\ \langle (\vec{U}', [c]); \vec{G}'; \bar{S}'; \bar{H}' \rangle_{k'} \end{array} \right)} \\
\text{(Loc 2)} \quad \frac{}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; ([k]c, \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (c, \vec{G}); \bar{S}; \bar{H} \rangle_k} \\
\text{(Loc 3)} \quad \frac{k \neq k' \quad k' \notin \text{Locs}(\Omega)}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; ([k']c, \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; \vec{G}; \bar{S}; \bar{H} \rangle_k, \langle c; \emptyset; \emptyset; \emptyset \rangle_{k'}} \\
\text{(Act)} \quad \frac{d \text{ is a fresh id}}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; (p(\vec{t}), \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (p(\vec{t})\#d : 1, \vec{G}); (\bar{S}, p(\vec{t})\#d); \bar{H} \rangle_k} \\
\text{(Simp)} \quad \frac{r : [l]P' \setminus [l](S', c'_i, S'') \iff G \mid B \in \mathcal{P} \quad \models \theta \wedge G \quad k = \theta l \quad \text{DropIds}(P) = \theta P' \quad \text{DropIds}(S) = \theta(S', S'') \quad c = \theta c'}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); (\bar{S}, P, S, c\#d); \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (\text{NF}(\text{Inst}(\theta B)), \vec{G}); (\bar{S}, P); \bar{H} \rangle_k} \\
\text{(Prop)} \quad \frac{r : [l](P', c'_i, P'') \setminus [l]S' \iff G \mid B \in \mathcal{P} \quad \models \theta \wedge G \quad k = \theta l \quad \text{DropIds}(P) = \theta(P', P'') \quad \text{DropIds}(S) = \theta S' \quad c = \theta c' \quad (d, \text{Ids}(P, S)) \notin \bar{H}}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); (\bar{S}, P, S, c\#d); \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (\text{NF}(\text{Inst}(\theta B)), c\#d : i, \vec{G}); (\bar{S}, P, c\#d); (\bar{H}, (d, \text{Ids}(P, S))) \rangle_k} \\
\text{(Next)} \quad \frac{\text{(Simp) and (Prop) do not apply for } c\#d : i}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (c\#d : (i+1), \vec{G}); \bar{S}; \bar{H} \rangle_k} \\
\text{(Drop)} \quad \frac{i \notin \text{OccIds}(\mathcal{P})}{\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; \vec{G}; \bar{S}; \bar{H} \rangle_k}
\end{array}$$

Figure 7: ω_0^e Operational Semantics for CHR^e

$$\begin{array}{l}
\text{Drop Ids} \quad \begin{cases} \text{DropIds}(\bar{S}, c\#d) & = \text{DropIds}(\bar{S}), c \\ \text{DropIds}(\emptyset) & = \emptyset \end{cases} \\
\text{Retrieve Ids} \quad \begin{cases} \text{Ids}(\bar{S}, c\#d) & = \text{Ids}(\bar{S}), d \\ \text{Ids}(\emptyset) & = \emptyset \end{cases} \\
\text{Retrieve Occ Indices} \quad \begin{cases} \text{OccIds}(R \mathcal{P}) & = \text{OccIds}(R), \text{OccIds}(\mathcal{P}) \\ \text{OccIds}(r : P \setminus S \iff G \mid B) & = \text{OccIds}(P), \text{OccIds}(S) \\ \text{OccIds}([l]c_i, H) & = i, \text{OccIds}(H) \\ \text{OccIds}(\cdot) & = \emptyset \end{cases}
\end{array}$$

Figure 7 defines the ω_0^e operational semantics. Given a 0-neighbor restricted CHR^e program \mathcal{P} , a ω_0^e derivation step expresses a transition between ensemble states, written $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$. Execution in a

location $\langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k \in \Omega$ is mainly driven by the goals \vec{G} which function as a stack of procedures waiting to be executed. By contrast \vec{U} buffers the constraints sent to the location. The (Flush) step states that constraints in a non-empty buffer \vec{U} are to be moved into the goals if the current goal is empty. (Loc 1), (Loc 2) and (Loc 3) model the delivery of body constraint $[k']c$ to forwarding location k' : (Loc 1) applies if the leading goal is $[k']c$ and k' is distinct from the origin location k , and sends c to the buffer \vec{U}' of k' . For (Loc 2) the forwarding location is the origin location, hence no actual transmission occur and the localization operator $[k]$ is simply stripped away. Finally for (Loc 3) with localization operator $[k']$, k' does not appear anywhere in the ensemble, hence we create a new location k' with just c in the buffer and all other collections empty⁷. (Act) applies to a leading goal of the form $p(\vec{t})$. It introduces $p(\vec{t}) : d$ into the store, where d is a fresh constraint identifier, and puts the active constraint $p(\vec{t})\#d : 1$ as the new leading goal. This represents the initialization of rule matching rooted at $p(\vec{t})\#d$ starting from the rule head in the program \mathcal{P} that corresponds to the first occurrence index. The last four derivation steps apply to leading goals of the form $c\#d : i$. (Prop) and (Simp) model the application of a 0-neighbor restricted CHR rule instance $R \in \mathcal{P}$ such that the i^{th} rule head occurrence matches c and respective partner constraints are found in the store. (Prop) additionally enforces a history check similar to the traditional CHR semantics [DSdlBH04, Sch05]. The purpose of this is to disallow multiple applications of rule instances that originate from the same multiset of constraints in the store⁸. The derivation (Next) increments the occurrence index of the active constraint, while (Drop) removes the active constraint from the goal once it has been tried on all rule head occurrences. We define the transitive and reflexive application of ω_0^e derivation steps as $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$. A state Ω' is *reachable* by a program \mathcal{P} under the ω_0^e semantics if there exists some initial state Ω such that $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.

We are interested in a class of CHR^e programs known as *locally quiescent* CHR^e programs. We say that a CHR^e program \mathcal{P} is locally quiescent if given any well-formed reachable state Ω , we cannot have any infinite derivation sequences that *does not include* the (Flush) derivation step. This specifically means that each location k in a Ω must always (eventually and asynchronously) execute to a state where its goals are empty, during which the (Flush) step is applicable and the constraints in the buffer will be pushed into the goals. Locally quiescent programs have the property that constraints sent across locations are not left to “starve” in a location’s buffer, because local execution of a location k is guaranteed to be terminating. Hence it is a form of progress guarantee that each location will eventually process (activate, store and match) each constraint delivered to it. While we expect that it is reasonable that distributed applications can possibly behave in a non *globally quiescent* manner⁹, we expect that each location executes in a locally quiescent manner as described here. We will explicitly consider locally quiescent programs in Section 5.

Lemma 4 states the property that ω_0^e derivations preserves the well-formedness of the operational states Ω .

Lemma 4 (Well-Formedness Preservation of ω_0^e Derivations) *Given a well-formed 0-neighbor restricted CHR^e program \mathcal{P} , a well-formed state Ω and state Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ then Ω' is well-formed.*

Given a state Ω , we define a meta operation $\text{Goals}(\Omega)$ that denotes the consolidated sequence of all goals in Ω . We extend the notion of location retrieval and location restriction on goals: $\text{Locs}(\vec{G})$ denotes the set of distinct locations k such that $[k]c \in \vec{G}$ for some c . Given location k , $\vec{G}|_k$ denotes the sequence containing all constraints c where $[k]c \in \vec{G}$. These are inductively defined by the following:

⁷New locations are introduced by existential forwarding locations in rule bodies.

⁸History checking is only required on propagation rules ($r : P \setminus S \iff G \mid B$ where $S = \emptyset$). For brevity, we conservatively apply history checking to all states that applies to the (Prop) derivation step.

⁹Globally quiescent: collective execution of the ensemble terminates. In other words, there exists a reachable state where all locations of the ensemble reaches quiescence.

$$\begin{array}{l}
\text{Ensembles} \quad \left\{ \begin{array}{l} [\Omega] = [\Omega]^{\vec{\mathcal{G}}'}, \uplus_{k \in (\text{Locs}(\vec{\mathcal{G}}') - \text{Locs}(\Omega))} \langle \vec{\mathcal{G}}'_{|k} \rangle_k \quad \text{where } \vec{\mathcal{G}}' = \text{Goals}(\Omega) \\ [\Omega, \langle \vec{\mathcal{U}} ; \vec{\mathcal{G}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k]^{\vec{\mathcal{G}}'} = [\Omega]^{\vec{\mathcal{G}}'}, \langle [\vec{\mathcal{U}}], [\vec{\mathcal{G}}], [\vec{\mathcal{S}}], \vec{\mathcal{G}}'_{|k} \rangle_k \\ [\emptyset]^{\vec{\mathcal{G}}'} = \emptyset \end{array} \right. \\
\\
\begin{array}{ll} \text{Identity} & [\emptyset] = \emptyset \\ \text{Buffers} & [c, \vec{\mathcal{U}}] = c, [\vec{\mathcal{U}}] \\ \text{Stores} & [c\#d, \vec{\mathcal{S}}] = c, [\vec{\mathcal{S}}] \end{array} & \text{Goals} \quad \left\{ \begin{array}{ll} [c, \vec{\mathcal{G}}] & = c, [\vec{\mathcal{G}}] \\ [c\#d : i, \vec{\mathcal{G}}] & = [\vec{\mathcal{G}}] \\ [[k]c, \vec{\mathcal{G}}] & = [\vec{\mathcal{G}}] \end{array} \right.
\end{array}$$

Figure 8: Abstract Ensemble State Interpretation of Operational Ensemble States

$$\begin{array}{ll}
\text{Consolidate Goals} & \left\{ \begin{array}{ll} \text{Goals}(\Omega, \langle \vec{\mathcal{U}} ; \vec{\mathcal{G}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k) & = \text{Goals}(\Omega), \vec{\mathcal{G}} \\ \text{Goals}(\emptyset) & = \emptyset \end{array} \right. \\
\\
\text{Retrieve Locations (Goals)} & \left\{ \begin{array}{ll} \text{Locs}(c, \vec{\mathcal{G}}) & = \text{Locs}(\vec{\mathcal{G}}) \\ \text{Locs}([k]c, \vec{\mathcal{G}}) & = \{k\} \cup \text{Locs}(\vec{\mathcal{G}}) \\ \text{Locs}(\emptyset) & = \emptyset \end{array} \right. \\
\\
\text{Location Restrict (Goals)} & \left\{ \begin{array}{ll} c, \vec{\mathcal{G}}_{|k} & = \vec{\mathcal{G}}_{|k} \\ [k']c, \vec{\mathcal{G}}_{|k} & = \vec{\mathcal{G}}_{|k} \quad \text{if } k \neq k' \\ [k]c, \vec{\mathcal{G}}_{|k} & = c, \vec{\mathcal{G}}_{|k} \\ \emptyset_{|k} & = \emptyset \end{array} \right.
\end{array}$$

Figure 8 gives the translation function $[-]$ that inductively traverses the structure of an operational ensemble state Ω and translates it to a corresponding fragment of an abstract state \mathcal{A} . At the top-most level, $[\Omega]$ is equivalent to $[\Omega]^{\vec{\mathcal{G}}'}, \uplus_{k \in (\text{Locs}(\vec{\mathcal{G}}') - \text{Locs}(\Omega))} \langle \vec{\mathcal{G}}'_{|k} \rangle_k$ such that $\vec{\mathcal{G}}' = \text{Goals}(\Omega)$ is the consolidated sequence of all goals in the ensemble. The first component $([\Omega]^{\vec{\mathcal{G}}'})$ takes each location in Ω and collapses $\vec{\mathcal{U}}$, $\vec{\mathcal{G}}$ and $\vec{\mathcal{S}}$ together while dropping the history $\vec{\mathcal{H}}$ entirely. Kept as a superscript is the sequence of all goals $\vec{\mathcal{G}}'$. Each location k also extracts from $\vec{\mathcal{G}}'$ all located constraints that belongs to k (i.e. $\vec{\mathcal{G}}'_{|k}$). Buffer $\vec{\mathcal{U}}$ is interpreted as a multiset of constraints. For $\vec{\mathcal{G}}$, we discard active constraints $c\#d : i$ because well-formed states have a corresponding $c\#d$ in $\vec{\mathcal{S}}$. We also discard located constraints $[k]c$. For $\vec{\mathcal{S}}$, we strip away constraint ids $\#d$. We implicitly convert sequences ($\vec{\mathcal{U}}$ and $\vec{\mathcal{G}}$) into multisets. The second component $\uplus_{k \in (\text{Locs}(\vec{\mathcal{G}}') - \text{Locs}(\Omega))} \langle \vec{\mathcal{G}}'_{|k} \rangle_k$ retrieves the set of all locations k that are referenced in the goals but are not known locations in Ω , and “creates” these new locations each containing their respective fragment of the goals (i.e. $\vec{\mathcal{G}}'_{|k}$). Lemma 5 states the property that this translation maps a well-formed operational state into a well-formed abstract state.

Lemma 5 (Well-Formedness Preservation of $[-]$ Translation) *Given a well-formed ω_0^e operational state Ω , $[\Omega]$ is a well-formed ω_α^e abstract state.*

Theorem 6 states that ω_0^e derivations can be mapped to corresponding ω_α^e derivations via the $[-]$ interpretation of operational states.

Theorem 6 (Soundness of ω_0^e) *Given 0-neighbor restricted CHR^e program \mathcal{P} and states Ω and Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$.*

Note that a terminal state Ω (where all $\vec{\mathcal{U}}_i = \emptyset$ and $\vec{\mathcal{G}}_i = \emptyset$) is in a state of *quiescence* where no ω_0^e

$$\begin{array}{c}
\text{(Single)} \quad \frac{\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'}{\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'} \\
\\
\text{(Concurrent)} \quad \frac{\mathcal{P} \triangleright (\Omega_1, \Omega_2) \mapsto_{\omega_0^e} (\Omega'_1, \Omega_2) \quad \mathcal{P} \triangleright (\Omega_1, \Omega_2) \mapsto_{\omega_0^e} (\Omega_1, \Omega'_2)}{\mathcal{P} \triangleright (\Omega_1, \Omega_2) \mapsto_{\omega_0^e} (\Omega'_1, \Omega'_2)}
\end{array}$$

Figure 9: Concurrent ω_0^e Derivation Steps

derivation steps can apply. The ω_0^e semantics guarantees that when we reach quiescence from a well-formed initial state, all rules in the program have been *exhaustively applied*.

To prove this guarantee of exhaustive rule application in quiescence of ω_0^e derivations, we first define the notion of rule applicability: Given a 0-neighbor restricted rule $R = r : P' \setminus S' \iff G \mid B$, we say that R is *applicable* in a state Ω , if there exists some $\langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k \in \Omega$ and some substitution θ where there exists $P, S \in \vec{S}$ such that $\text{DropIds}(P) = \theta P'$ and $\text{DropIds}(S) = \theta S'$, $\models \theta G$ and $\text{Ids}(P, S) \notin \vec{H}$. The rule instance of R in Ω is said to be *active* if for its rule head instance, P and S , there exists some ‘ $c\#d \in P, S$ ’ that matches the j^{th} occurrence index of R , such that $c\#d : i \in \vec{G}$ and $i \leq j$. Lemma 7 states the property that all rule instances in a reachable state are active.

Lemma 7 (Always Active Rule Head Instances) *Given a 0-neighbor restricted CHR^e program \mathcal{P} and an initial state Ω , for any reachable state Ω' such that $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, all rule instances in Ω' must be active.*

Theorem 8 states the exhaustiveness of rule application in the ω_0^e semantics. Specifically, it states that starting from an initial state Ω , if we derive a terminal state from Ω , this terminal state has no rule instances. This is an important property because it means that exhaustive execution of ω_0^e derivations of a CHR^e program \mathcal{P} to a state of quiescence results to the exhaustive application of all rule instances of \mathcal{P} .

Theorem 8 (Exhaustiveness of Rule Application in ω_0^e) *Given a 0-neighbor restricted CHR^e program \mathcal{P} and reachable states Ω and Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ and Ω' is terminal, then there exists no rule instance $R \in \mathcal{P}$ such that R is applicable.*

Figure 9 define the concurrent ω_0^e derivation step, denoted $\mapsto_{\omega_0^e}^{\parallel}$. It explicitly defines concurrent execution of ω_0^e semantics derivation steps that modifies non-overlapping portions of an operational state. The (Single) rule states that a standard ω_0^e derivation (i.e., $\mapsto_{\omega_0^e}$) maps naively to a concurrent derivation step, $\mapsto_{\omega_0^e}^{\parallel}$. The rule (Concurrent) states that we can compose concurrent derivation steps together, as long as they operate on non-overlapping portions of an operational state. We write $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$ for the reflexive and transitive application of $\mapsto_{\omega_0^e}^{\parallel}$ derivation steps. The composibility of ω_0^e concurrent derivation steps stems from a more fundamental property of the abstract CHR semantics known as *monotonicity* [Frü94]. Specifically, it states that CHR derivations are valid even in a larger context of constraints. The ω_0^e semantics has this property as well, but in a slightly different context: ω_0^e derivations are *monotonic* with-respect-to locations, namely in a larger context of locations¹⁰. Lemma 9 states this property in the ω_0^e semantics in its various forms. We assume implicit α -renaming of locations when composing ensemble states together.

Lemma 9 (Monotonicity of ω_0^e semantics) *Given a 0-neighbor restricted program \mathcal{P} , and reachable states $\Omega, \Omega', \Omega''$,*

1. *if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e} \Omega', \Omega''$*

¹⁰Note that the refined operational semantics [DSdlBH04] does not have this property with-respect-to constraint stores, but it is monotonic with-respect-to its goals.

2. if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^* \Omega', \Omega''$
3. if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^{\parallel} \Omega', \Omega''$
4. if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^{\parallel*} \Omega', \Omega''$

With the monotonicity property stated in Lemma 9, we can prove Theorem 10, that states that every concurrent ω_0^e derivation can be simulated using sequential ω_0^e derivation steps.

Theorem 10 (Serializability of Concurrent ω_0^e Derivations) *Given a 0-neighbor restricted program \mathcal{P} , and reachable states Ω, Ω' , if we have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$ then we must have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$*

Theorem 10 has far reaching implications on the concurrent behavior of ω_0^e : Specifically that 0-neighbor restricted CHR^e rule application *need not be* executed entirely *atomically*, since the sequence of ω_0^e derivation steps that models a rule application can be interleaved between multiple ω_0^e derivation steps. Rather, *rule head observation and deletion* (observing presence of all rule head constraints in the store and deleting the matching simplified rule heads) is the critical component of rule application that needs to be atomic¹¹. More concretely, it is an application of either the (Simp) or (Prop) ω_0^e derivation step followed by a sequence of (Act), (Next) and (Drop) derivation steps that incrementally “moves” the body B of the rule instance into the constraint store. Between this sequence of derivation steps, we can possibly execute other concurrent derivation steps that involve rewritings on non-overlapping portions of the local constraint store or ensemble. Yet these concurrent derivations of non-atomic rule applications can be serialized to interleaving sequential ω_0^e derivation steps. We will exploit this observation in our encoding of 1-neighbor restricted rules into ω_0^e (Section 5).

5 Encoding 1-Neighbor Restricted Programs for ω_0^e

In this section, we define a translation that transforms a 1-neighbor restricted program into a 0-neighbor restricted program. Given a 1-neighbor restricted rule $r : [X]P_x, [Y]P_y \setminus [X]S_x, [Y]S_y \iff G \mid B$, we designate X as the primary location and Y as the neighbor location. We define two properties of a 1-neighbor restricted rule r , namely that r is *primary propagated* if $S_x = \emptyset$ and that r is *neighbor propagated* if $S_y = \emptyset$. We call P_x and S_x the *primary matching obligations*, while P_y and S_y are the *neighbor matching obligations*. We only consider locally quiescent CHR^e programs from this section.

5.1 Basic Encoding Scheme

Figure 10 illustrates our basic encoding scheme, $\rightsquigarrow_{\text{INb}}^{\text{basic}}$ on an example. It applies to a 1-neighbor restricted rule that is neither primary nor neighbor propagated. We assume that constraints of the predicate *neighbor* are never deleted, i.e, no other rules in a program with the *swap* rule have a *neighbor* constraint as a simplified head. We call such a constraint a *persistent* constraint. A constraint c is *persistent* if there is no substitution θ such that $\theta c = \theta c'$ for some rule $r : P \setminus S \iff G \mid B \in \mathcal{P}$ and constraint $c' \in S$. Note that the rule heads demand that we must *atomically observe* that in some location X we have *neighbor*(Y) and *color*(C), while in Y we have *color*(C'). This observation must be *atomic* in the sense that the observation of X 's and Y 's matching obligations should not be interrupted by an interleaving concurrent derivation.

The 0-neighbor restricted encoding of the *swap* rule (Figure 10) recovers this atomicity: In *swap_1* X sends a swap request *swap_req*(X, Y, C) to Y if it possesses the primary matching obligation of the

¹¹This property is first observed in [SL08, LS11] and we call this atomic fragment of rule application *atomic rule head verification*.

$$\begin{array}{c}
\text{swap} : \left(\begin{array}{c} [X] \text{neighbor}(Y), \\ \cdot \end{array} \right) \setminus \left(\begin{array}{c} [X] \text{color}(C), \\ [Y] \text{color}(C') \end{array} \right) \iff \left(\begin{array}{c} [X] \text{color}(C'), \\ [Y] \text{color}(C) \end{array} \right) \\
\downarrow \text{basic} \\
\text{LNB} \\
\left(\begin{array}{l}
\text{swap_1} : [X] \text{neighbor}(Y), [X] \text{color}(C) \implies [Y] \text{swap_req}(X, Y, C) \\
\text{swap_2} : [Y] \text{color}(C') \setminus [Y] \text{swap_req}(X, Y, C) \iff [X] \text{swap_match}(X, Y, C, C') \\
\text{swap_3} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(C), [X] \text{swap_match}(X, Y, C, C') \\
\quad \iff [Y] \text{swap_commit}(X, Y, C, C') \\
\text{swap_4a} : [Y] \text{swap_commit}(X, Y, C, C'), [Y] \text{color}(C') \iff [X] \text{color}(C'), [Y] \text{color}(C) \\
\text{swap_4b} : [Y] \text{swap_commit}(X, Y, C, C') \iff [X] \text{color}(C)
\end{array} \right)
\end{array}$$

Figure 10: Color Swapping Example: 1-Neighbor Restricted Rule

swap rule. In *swap_2* if Y observes this request together with a $\text{color}(C')$, it responds to X by sending $\text{swap_match}(X, Y, C, C')$. Note that in this example, $\text{swap_req}(X, Y, C)$ is simplified since X 's matching obligation is not primary propagated¹². In *swap_3*, X must observe that it has a response from Y ($\text{swap_match}(X, Y, C, C')$) and that its matching obligations are still valid¹³, then it sends a commit request to Y $\text{swap_commit}(X, Y, C, C')$. This is the point where X actually *commits* to the match by consuming $\text{color}(C)$. From here there are two possibilities: *swap_4a* considers that, if Y still possesses the matching instance of $\text{color}(C')$, we complete the execution of *swap* by delivering its rule body. *swap_4b* considers the alternative case where Y no longer has $\text{color}(C')$ hence we cannot commit to the rule instance. Hence we *roll back* X 's commitment by returning $\text{color}(C)$ to X . Note that ω_0^e 's sequencing of rule occurrences is vital to ensure that given a *swap_commit* instance, rule matching for *swap_4a* is always attempted before *swap_4b*. We call the predicates introduced in the encoding (like *swap_req*, *swap_match*) *synchronizing predicates* and constraints they form *synchronizing constraints*.

The five 0-neighbor restricted rules in Figure 10 implement an *asynchronous* and *optimistic* synchronization protocol between two locations of the ensemble. It is asynchronous because neither primary X nor neighbor Y ever “blocks” or busy-waits for responses. Rather they communicate asynchronously via the synchronizing constraints, while potentially interleaving with other derivation steps. It is optimistic because non-synchronizing constraints are only ever consumed after both X and Y have independently observed their respective fragment of the rule head instance.

Figure 11 defines another example of a 1-neighbor restricted rule that differs from the *swap* example in two ways: It is *primary propagated* and it contain a primary propagated head (namely $\text{color}(C)$) that is not *persistent*. We highlight in boxes the fragments which differ by the fact that the rule is *primary propagated*, and with an underline the fragments which differ by the fact that propagated head $\text{color}(C)$ is not *persistent*.

Since *prop* is primary propagated, it is possible that a single instance of this rule head fragment be applied to multiple instances of $[Y] \text{color}(C')$ for a particular location Y . If we follow a similar encoding to the *swap_2* rule for the previous example (Figure 10) we cannot guarantee exhaustiveness of 1-neighbor restricted rule application. Therefore, as highlighted in a box in figure 11, rule *prop_2* is defined such that the synchronizing constraint $\text{prop_req}(X, Y, C)$ is propagated as opposed to being simplified (highlighted in a box).

¹²If it was primary propagated (by making $[X] \text{color}(C)$ of the *swap* rule propagated instead), we would have to propagate $\text{swap_req}(X, Y, C)$ instead, since X 's matching obligation can match and apply to multiple instances of *swap* (see Figure 11 for such an example).

¹³This revalidation ensures that Y 's observation of its matching obligation has not been invalidated by an interleaving rule application that consumed any part of X 's obligations.

$$\begin{array}{c}
\text{prop} : \left(\begin{array}{c} [X] \text{neighbor}(Y), \underline{[X] \text{color}(C)} \\ \cdot \end{array} \right) \setminus \left(\begin{array}{c} \boxed{\cdot} \\ [Y] \text{color}(C') \end{array} \right) \iff [Y] \text{color}(C) \\
\sim_{\text{basic}}^{\text{1Nb}} \\
\left(\begin{array}{l}
\text{prop1} : [X] \text{neighbor}(Y), [X] \text{color}(C) \implies [Y] \text{prop_req}(X, Y, C) \\
\text{prop2} : \boxed{[Y] \text{color}(C'), [Y] \text{prop_req}(X, Y, C)} \implies [X] \text{prop_match}(X, Y, C, C') \\
\text{prop3} : [X] \text{neighbor}(Y) \setminus \underline{[X] \text{color}(C)}, [X] \text{prop_match}(X, Y, C, C') \\
\iff [Y] \text{prop_commit}(X, Y, C, C') \\
\text{prop4a} : [Y] \text{prop_commit}(X, Y, C, C'), [Y] \text{color}(C') \iff \underline{[X] \text{color}(C)}, [Y] \text{color}(C) \\
\text{prop4b} : [Y] \text{prop_commit}(X, Y, C, C') \iff \underline{[X] \text{color}(C)}
\end{array} \right)
\end{array}$$

Figure 11: Color Propagation Example: Primary Propagated 1-Neighbor Restricted Rule

Since *prop* has a non persistent propagated head, in order to safely guarantee that the observation of X and Y matching obligations are done independently, *prop3* commits its obligation by deleting its non persistent propagated constraint(s), in this case $\text{color}(C)$, while in *prop4a* and *prop4b*, this constraint is returned to X . While this possibly introduces additional overhead, it is crucial to ensuring the safety of this rule application. Figure 12 that shows an example of the same rule *prop* of Figure 11 but with an encoding that treats non-persistent propagated head $\text{color}(C)$ as persistent. This encoding is flawed because we can derive *prop* rule applications from the program encoding \mathcal{P}_0 , while \mathcal{P}_1 has no valid derivation. Specifically, Figure 12 highlights the steps that leads to a wrong derivation. We consider a contrived program with two additional rules *cheat1* and *cheat2* that shuffles a constraint color between two neighboring locations. The initial state starts with only location k_1 having a $\text{color}(\text{blue})$ constraint, but with proper timing of the *cheat* rules, we can shuffle the single color constraint and derive the application of rule *prop*.

Figure 13 defines the basic encoding scheme for 1-neighbor restricted rule. It is denoted by $R_1 \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}_0$, where R_1 is a well-formed 1-neighbor restricted rule while \mathcal{P}_0 a well-formed 0-neighbor restricted program. The propagated rule heads of the primary location X are split into two, namely P_x containing all rule heads which are persistent in \mathcal{P}_1 and P'_x containing all those which are non-persistent. P'_x and S_x will be consumed in r_3 when the primary location commits, this effectively “locks” the primary matching obligation. The neighbor location Y applies r_4a to complete the rule application r , specifically adding the rule body D and “unlocking” P'_x , the primary propagated rule heads which are not persistent. If it is unable to complete this rule application, r_4b is applied to roll-back location X ’s commit attempt by returning P'_x and S_x to X . $\text{MatchRule}(\cdot)$ characterizes the fragment of the encoding unique to primary propagated 1-neighbor restricted rules, while $\text{MatchRule}(S_x)$ for a non-empty S_x represents the corresponding fragment of non-primary propagated rules. Xs and Ys are the set of variables of the primary and neighbor location rule heads while Rs is the union of the two. Rule guards are divided into two parts, namely G_x the primary rule guards and G_y the neighbor rule guards. Primary rule guards G_x are all the guard conditions that are grounded by Xs , while G_y are the rest of the guards. We refer to these rules (r_i) as *encoding rules*. We define the meta operator $\text{DropSynchs}(\bar{S}; \mathcal{P})$ that returns the multiset of all constraints that appear in \bar{S} that are not synchronizing constraints of \mathcal{P} :

$$\begin{aligned}
\mathcal{P}_1 &= \begin{cases} \text{prop} : \left(\begin{array}{c} [X] \text{neighbor}(Y), [X] \text{color}(C) \\ \cdot \end{array} \right) \setminus \left(\begin{array}{c} \boxed{} \\ [Y] \text{color}(C') \end{array} \right) \iff [Y] \text{color}(C) \\ \text{cheat1} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(\text{blue}) \iff [Y] \text{color}(\text{red}) \\ \text{cheat2} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(\text{red}) \iff [Y] \text{color}(\text{blue}) \end{cases} \\
\mathcal{P}_0 &= \begin{cases} \text{prop1} : [X] \text{neighbor}(Y), [X] \text{color}(C) \implies [Y] \text{prop_req}(X, Y, C) \\ \text{prop2} : [Y] \text{color}(C'), [Y] \text{prop_req}(X, Y, C) \implies [X] \text{prop_match}(X, Y, C, C') \\ \text{prop3} : [X] \text{neighbor}(Y), [X] \text{color}(C) \setminus [X] \text{prop_match}(X, Y, C, C') \\ \quad \iff [Y] \text{prop_commit}(X, Y, C, C') \\ \text{prop4a} : [Y] \text{prop_commit}(X, Y, C, C'), [Y] \text{color}(C') \iff [Y] \text{color}(C) \\ \text{prop4b} : [Y] \text{prop_commit}(X, Y, C, C') \iff \text{true} \\ \text{cheat1} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(\text{blue}) \iff [Y] \text{color}(\text{red}) \\ \text{cheat2} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(\text{red}) \iff [Y] \text{color}(\text{blue}) \end{cases} \\
\mathcal{P}_0 \triangleright & \langle \boxed{\text{neighbor}(k_2), \text{color}(\text{blue})} \rangle_{k_1}, \langle \text{neighbor}(k_1) \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{prop1}} & \langle \boxed{\text{neighbor}(k_2), \text{color}(\text{blue})} \rangle_{k_1}, \langle \text{neighbor}(k_1), \text{prop_req}(k_1, k_2, \text{blue}) \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{cheat1}} & \langle \text{neighbor}(k_2) \rangle_{k_1}, \langle \text{neighbor}(k_1), \boxed{\text{prop_req}(k_1, k_2, \text{blue}), \text{color}(\text{red})} \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{prop2}} & \langle \text{neighbor}(k_2), \text{prop_match}(k_1, k_2, \text{blue}, \text{red}) \rangle_{k_1}, \langle \boxed{\text{neighbor}(k_1)}, \text{prop_req}(k_1, k_2, \text{blue}), \boxed{\text{color}(\text{red})} \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{cheat2}} & \langle \boxed{\text{neighbor}(k_2), \text{prop_match}(k_1, k_2, \text{blue}, \text{red}), \text{color}(\text{blue})} \rangle_{k_1}, \langle \text{neighbor}(k_1), \text{prop_req}(k_1, k_2, \text{blue}) \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{prop3}} & \langle \boxed{\text{neighbor}(k_2), \text{color}(\text{blue})} \rangle_{k_1}, \langle \text{neighbor}(k_1), \text{prop_req}(k_1, k_2, \text{blue}), \text{prop_commit}(k_1, k_2, \text{blue}, \text{red}) \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{cheat1}} & \langle \text{neighbor}(k_2) \rangle_{k_1}, \langle \text{neighbor}(k_1), \text{prop_req}(k_1, k_2, \text{blue}), \boxed{\text{prop_commit}(k_1, k_2, \text{blue}, \text{red}), \text{color}(\text{red})} \rangle_{k_2} \\
\mapsto_{\omega_\alpha^e}^{\text{prop4a}} & \langle \text{neighbor}(k_2) \rangle_{k_1}, \langle \text{neighbor}(k_1), \text{prop_req}(k_1, k_2, \text{blue}), \text{color}(\text{blue}) \rangle_{k_2}
\end{aligned}$$

Figure 12: Example of Flawed Encoding of Rule with Non-persistent Propagated Head

$$\begin{aligned}
\text{SyncPred}((r : P \setminus S \iff G \mid B) \mathcal{P}) &= r_req, r_match, r_commit, r_abort, \text{SyncPred}(\mathcal{P}) \\
\text{SyncPred}(\cdot) &= \emptyset
\end{aligned}$$

$$\text{DropSyncs}(p(\bar{t}), \bar{\mathcal{S}}; \mathcal{P}) = \begin{cases} p(\bar{t}), \text{DropSyncs}(\bar{\mathcal{S}}; \mathcal{P}) & \text{if } p \notin \text{SyncPred}(\mathcal{P}) \\ \text{DropSyncs}(\bar{\mathcal{S}}; \mathcal{P}) & \text{otherwise} \end{cases}$$

$$\text{DropSyncs}(\emptyset; \mathcal{P}) = \emptyset$$

As illustrated in Figure 13, we generalize the application of this translation to 1-neighbor restricted programs, thus given a 1-neighbor restricted program \mathcal{P}_1 , we have its encoding via $\mathcal{P}_1 \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}_0$, such that the encoding operation is applied to each 1-neighbor restricted rule in \mathcal{P}_1 while 0-neighbor restricted rules are simply left unmodified. All rule encodings (each a 0-neighbor program) are then concatenated into a

$$\begin{array}{c}
 (r : [X]P_x, [X]P'_x, [Y]P_y \setminus [X]S_x, [Y]S_y \iff G_x, G_y \mid \exists \bar{z}. D) \\
 \Downarrow_{\text{basic 1Nb}} \\
 \left(\begin{array}{l}
 r_1 : [X]P_x, [X]S_x \implies G_x \mid [Y]r_req(Xs) \\
 \text{MatchRule}(S_x) \\
 r_3 : [X]P_x \setminus [X]P'_x, [X]S_x, [X]r_match(Rs) \iff [Y]r_commit(Rs) \\
 r_4a : [Y]P_y \setminus [Y]S_y, [Y]r_commit(Rs) \iff \exists \bar{z}. [X]P'_x, D \\
 r_4b : [Y]r_commit(Rs) \iff [X]P'_x, [X]S_x
 \end{array} \right)
 \end{array}$$

where

All $c \in P_x$ are persistent constraints and all $c' \in P'_x$ are non-persistent constraints
 $\text{MatchRule}(\cdot) = r_2 : [Y]P_y, [Y]S_y, [Y]r_req(Xs) \implies G_y \mid [X]r_match(Rs)$
 $\text{MatchRule}(S_x) = r_2 : [Y]P_y, [Y]S_y \setminus [Y]r_req(Xs) \iff G_y \mid [X]r_match(Rs)$ if $S_x \neq \cdot$.
 $Xs = \text{FV}(P_x, P'_x, S_x)$ $Ys = \text{FV}(P_y, S_y)$ $Rs = \text{FV}(P_x, P'_x, S_x, P_y, S_y)$
 $\text{FV}(G_x) \subseteq \text{FV}(P_x, P'_x, S_x)$ $\text{FV}(G_y) \subseteq \text{FV}(P_x, P'_x, S_x, P_y, S_y)$

$$\begin{array}{c}
 \text{(1-neighbor)} \quad \frac{\mathcal{P} \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}' \quad R_1 \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}_0 \quad R_1 \text{ is 1-neighbor restricted}}{(\mathcal{P} R_1) \rightsquigarrow_{\text{1Nb}}^{\text{basic}} (\mathcal{P}' \mathcal{P}_0)} \\
 \text{(0-neighbor)} \quad \frac{\mathcal{P} \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}' \quad R_0 \text{ is 0-neighbor restricted}}{(\mathcal{P} R_0) \rightsquigarrow_{\text{1Nb}}^{\text{basic}} (\mathcal{P}' R_0)} \quad \text{(identity)} \quad \frac{}{\cdot \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \cdot}
 \end{array}$$

Figure 13: Basic Encoding of 1-Neighbor Restricted Rules

single 0-neighbor restricted program \mathcal{P}_0 . We assume that unique rule head occurrence indices are issued in order of rule head appearance in \mathcal{P}_0 . When required for specific discussions, we will denote a (Simp) or (Prop) derivation step that involves the application of a r_i encoding rule instance as $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^E}^{r_i} \Omega'$, where Ω and Ω' are the states before and after the application of r_i .

It is possible that a partial sequences of encoding rules (r_i) is executed in a ω_0^E derivation. For instance, primary location X can apply r_1 but never receives a reply from neighbor location Y with r_2 because Y does not possess the matching obligations required to complete the rule instance. Or similarly, Y can apply r_2 in response to X 's instance of r_1 , but never receives a reply from X with r_3 because X no longer possess matching obligations. Such partial sequences of execution are the side-effect of asynchrony in this synchronization protocol, and are *benign* in that they do not rewrite (delete or insert) non-synchronizing constraints. Furthermore, their only observable effects are the introduction of synchronizing constraints r_req and r_match whose only purpose and effect is the sequencing and staging of the flow of consensus building between locations X and Y . Lemma 11 states this property that derivation steps of each r_i encoding rule must have been preceded by derivation steps of matching instances of encoding rules that come before i .

Lemma 11 (Prefix Executions of Encoding) *Given a locally quiescent 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{1Nb}}^{\text{basic}} \mathcal{P}_0$, if we have reachable states Ω and Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^E} \Omega'$, then we have the following:*

1. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^E}^{r_2} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^E}^{r_1} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^E}^* \Omega$
2. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^E}^{r_3} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^E}^{r_2} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^E}^* \Omega$

3. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r-4a} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r-3} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
4. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r-4b} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r-3} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$

We now consider the soundness of this encoding. Specifically, the soundness condition that we need is that the ω_0^e derivations of the 0-neighbor restricted encodings of a 1-neighbor restricted program \mathcal{P} derive valid states computable by \mathcal{P} in the ω_α^e semantics. However, not all states derived by our encodings are such valid states: It is possible that the ω_0^e derivations of the 0-neighbor restricted encodings derive intermediate states in which a 1-neighbor restricted rule instance is partially applied. Specifically, after application of $r-3$ only the X matching obligation is consumed, hence the rule instance is only partially applied. For this reason, the encodings are defined such that these intermediate states always contain the r_commit synchronizing constraint. We call such states *non-commit free* states and *commit free* states are all other states that do not contain r_commit . A state Ω is *commit free* if and only if for all $\langle \bar{S} \rangle_k \in [\Omega]$, all $p(\vec{t}) \in \bar{S}$ is such that $p \neq r_commit$. Note we use $[\Omega]$ for the convenience of collapsing the contents of buffers, goals and stores of the ω_0^e semantics into one abstract store of the ω_α^e semantics. By definition of the $[-]$ translation, this effectively means that all buffers, goals and store in the operational state Ω must not contain any commit synchronizing constraints in order for Ω to qualify as a commit-free state. An important property of this encoding is that non-commit free states can always be eventually returned to a commit free state by applying either $r-4a$ or $r-4b$, resulting to the complete execution of r or a roll-back to the state before its execution attempt, respectively. Note that this is only true for CHR^e programs which are locally quiescent: The reason is because if a location's execution is not locally quiescent, it might never push buffered constraints into the goals (via the (Flush) derivation step of ω_0^e semantics) and hence might remain non-commit free indefinitely.

Lemma 12 states that the encoding operation $\rightsquigarrow_{\text{INb}}^{\text{basic}}$ preserves local quiescence of CHR^e programs.

Lemma 12 ($\rightsquigarrow_{\text{INb}}^{\text{basic}}$ Preserves Local Quiescence) *Given a locally quiescent 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}}^{\text{basic}} \mathcal{P}_0$, then \mathcal{P}_0 is also locally quiescent.*

Lemma 13 states the property that non-commit free states can always eventually derive commit free state.

Lemma 13 (1-Neighbor Commit-Free Reachability) *Given a locally quiescent 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}}^{\text{basic}} \mathcal{P}_0$ and states Ω reachable by \mathcal{P}_0 , if Ω is not commit-free, then there exists some commit-free state Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.*

Lemma 14 states that given any ω_α^e derivation between two commit free state $\mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$, we can safely permute it such that all applications of $r-3$ encoding rules are immediately followed by either $r-4a$ or $r-4b$. This lemma is proven by using the monotonicity property of the ω_α^e semantics (Lemma 9).

Lemma 14 (Basic Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that:*

1. For some basic encoding rule instances $r-3$ and $r-4a$, we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r-3} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r-4a} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 and \mathcal{A}'_3 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r-3} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^{r-4a} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_4$
2. For some basic encoding rule instances $r-3$ and $r-4b$, we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r-3} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r-4b} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 and \mathcal{A}'_3 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r-3} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^{r-4b} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_4$

Theorem 15 asserts the soundness of the basic encoding: It states that ω_0^e derivations between commit free states of 0-neighbor restricted encodings have a mapping to ω_α^e derivations of its original 1-neighbor restricted program.

$$\begin{array}{c}
blend : \left(\begin{array}{l} [X] \text{neighbor}(Y), \\ [Y] \text{pallette}(C') \end{array} \right) \setminus \left(\begin{array}{l} [X] \text{color}(C) \\ \boxed{} \end{array} \right) \iff [Y] \text{color}(C + C') \\
\downarrow \sim_{\text{1NB}}^{\text{n-persist}} \\
\left(\begin{array}{l} \text{blend1} : [X] \text{neighbor}(Y), [X] \text{color}(C) \implies [Y] \text{blend_req}(X, Y, C) \\ \text{blend2} : [Y] \text{pallette}(C') \setminus [Y] \text{blend_req}(X, Y, C) \iff [X] \text{blend_match}(X, Y, C, C') \\ \text{blend3} : [X] \text{neighbor}(Y) \setminus [X] \text{color}(C), [X] \text{blend_match}(X, Y, C, C') \iff [X] \text{color}(C + C') \end{array} \right) \\
(r : [X] P_x, [Y] P_y \setminus [X] S_x \iff G_x, G_y \mid B) \\
\downarrow \sim_{\text{1NB}}^{\text{n-persist}} \\
\left(\begin{array}{l} r_1 : [X] P_x, [X] S_x \implies G_x \mid [Y] r_req(Xs) \\ \text{MatchRule}(S_x) \\ r_3 : [X] P_x \setminus [X] S_x, [X] r_match(Rs) \iff B \end{array} \right)
\end{array}$$

where

All $c \in P_y$ are persistent constraints

$\text{MatchRule}(\cdot) = r_2 : [Y] P_y, [Y] r_req(Xs) \implies G_y \mid [X] r_match(Rs)$

$\text{MatchRule}(S_x) = r_2 : [Y] P_y \setminus [Y] r_req(Xs) \iff G_y \mid [X] r_match(Rs)$ if $S_x \neq \cdot$.

$Xs = \text{FV}(P_x, S_x)$ and $\text{FV}(G_x) \subseteq \text{FV}(P_x, S_x)$ and $\text{FV}(G_y) \subseteq \text{FV}(P_x, P_y, S_y)$

Figure 14: Optimized Encoding for Neighbor Persistent Rule

Theorem 15 (Soundness of Basic Encoding) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{1NB}}^{\text{basic}} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega]; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega']; \mathcal{P}_1)$.*

5.2 Optimizations

We define optimized encoding schemes for two special cases. Consider the basic encoding scheme in Figure 13. Suppose that we apply to it a 1-neighbor restricted rule that is neighbor propagated (i.e, $S_y = \emptyset$) and furthermore P_y is persistent. We call such rules *neighbor persistent*. This encoding executes an unnecessary indirection of sending r_commit to neighbor Y in r_3 and completing the rule instance r with either r_4a or r_4b . However if $S_y = \emptyset$, primary location X can immediately complete the rule instance at r_3 without further communications with Y . This specialized encoding scheme is denoted by $\rightsquigarrow_{\text{1NB}}^{\text{n-persist}}$ is shown in Figure 15. We also show an example rule *blend* which has this property ($S_y = \emptyset$ and we assume constraints *pallette*(C) are persistent).

The next optimized encoding scheme applies specifically to 1-neighbor restricted rules which are not only primary propagated, but furthermore all propagated matching obligations are persistent¹⁴. We call such rules *primary persistent* rules. An example of this is the *trans* rule of the program of figure 1: Its primary matching obligation consists of only one propagated constraint *edge*(Y, D) which is never deleted by any rule of the program. Figure 15 shows this specialized encoding as the function $\rightsquigarrow_{\text{1NB}}^{\text{P-persist}}$. This optimized

¹⁴This corresponds to a rule in Figure 13 such that $P_x = \emptyset$ and $S_x = \emptyset$, hence we only have P'_x the persistent propagated rule heads.

$$\begin{array}{c}
 \text{trans} : \boxed{[X] \text{edge}(Y, D)}, [Y] \text{path}(Z, D') \Longrightarrow X \neq Z \mid [X] \text{path}(Z, D + D') \\
 \uparrow \sim_{\text{1Nb}}^{\text{p-persist}} \\
 \left(\begin{array}{l}
 \text{trans_1} : [X] \text{edge}(Y, D) \Longrightarrow [Y] \text{trans_req}(X, D) \\
 \text{trans_2} : [Y] \text{trans_req}(X, D), [Y] \text{path}(Z, D') \Longrightarrow X \neq Z \mid [X] \text{path}(Z, D + D')
 \end{array} \right) \\
 (r : [X] P_x, [Y] P_y \setminus [Y] S_y \iff G_x, G_y \mid B) \\
 \uparrow \sim_{\text{1Nb}}^{\text{p-persist}} \\
 \left(\begin{array}{l}
 r_1 : [X] P_x \Longrightarrow G_x \mid [Y] r_req(Xs) \\
 r_2 : [Y] r_req(Xs), [Y] P_y \setminus [Y] S_y \iff G_y \mid B
 \end{array} \right) \\
 \text{where } \begin{array}{l}
 \text{All } c \in P_x \text{ are persistent constraints} \\
 Xs = \text{FV}(P_x) \text{ and } \text{FV}(G_x) \subseteq \text{FV}(P_x) \text{ and } \text{FV}(G_y) \subseteq \text{FV}(P_x, P_y, S_y)
 \end{array}
 \end{array}$$

Figure 15: Optimized Encodings for Primary Persistent Rules

translation scheme is very similar to *rule localization* of link-restricted rule in distributed Datalog [LCG⁺06] and indeed it is a special case of rewriting where left-hand sides are not removed as a result of rule application.

Given a 1-neighbor restricted program \mathcal{P}_1 , we define the translation function $\mathcal{P}_1 \rightsquigarrow_{\text{1Nb}} \mathcal{P}_0$, where \mathcal{P}_0 is the 0-neighbor restricted program encoding in which we apply the optimized encodings where possible: We apply $\sim_{\text{1Nb}}^{\text{n-persist}}$ for 1-neighbor restricted rules in \mathcal{P}_1 that are neighbor persistent, while we apply $\sim_{\text{1Nb}}^{\text{p-persist}}$ for those that are primary persistent, and $\sim_{\text{1Nb}}^{\text{basic}}$ for all other 1-neighbor restricted rule. This is inductively defined by the following rules:

$$\begin{array}{c}
 \text{(Program)} \quad \frac{R \rightsquigarrow_{\text{1Nb}} \mathcal{P}_0 \quad \mathcal{P}_1 \rightsquigarrow_{\text{1Nb}} \mathcal{P}'_0}{(R \mathcal{P}_1) \rightsquigarrow_{\text{1Nb}} (\mathcal{P}_0 \mathcal{P}'_0)} \quad (0 - \text{N}) \quad \frac{R \text{ is 0-neighbor restricted}}{R \rightsquigarrow_{\text{1Nb}} R} \\
 \\
 \text{(N-Persist)} \quad \frac{\begin{array}{l} R \text{ is 1-neighbor restricted} \\ R \text{ is neighbor persistent} \end{array} \quad R \rightsquigarrow_{\text{1Nb}}^{\text{n-persist}} R'}{R \rightsquigarrow_{\text{1Nb}} R'} \quad (\text{P-Persist}) \quad \frac{\begin{array}{l} R \text{ is 1-neighbor restricted} \\ R \text{ is primary persistent} \end{array} \quad R \rightsquigarrow_{\text{1Nb}}^{\text{p-persist}} R'}{R \rightsquigarrow_{\text{1Nb}} R'} \\
 \\
 \text{(Basic)} \quad \frac{\begin{array}{l} R \text{ is 1-neighbor restricted} \\ (\text{N-Persist}) \text{ and } (\text{P-Persist}) \text{ does not apply} \end{array} \quad R \rightsquigarrow_{\text{1Nb}}^{\text{basic}} R'}{R \rightsquigarrow_{\text{1Nb}} R'}
 \end{array}$$

We now consider the soundness of the optimized encoding (the soundness of the basic encoding has been shown in the previous section). In general, both the neighbor and primary persistent encoding schemes exploit a similar property: one location's matching obligation M of the 1-neighbor restricted rules is completely persistent and we do not need to lock and reserve any constraint in M . This simplifies the application of such 1-neighbor restricted rules and as long as we observe that persistent matching obligations M are fulfilled, the rule application can be completed entirely at the other location's discretion. Lemma 16 and 17

states this property for the neighbor persistent and primary persistent encoding schemes respectively.

Lemma 16 (Neighbor Persistent Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1Nb} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that for some matching neighbor persistent encoding rule instances r_2 and r_3 , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_4$*

Lemma 17 (Primary Persistent Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1Nb} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that for some matching primary persistent encoding rule instances r_1 and r_2 , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_1} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_1} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_4$*

Theorem 18 states the soundness of the optimized encoding.

Theorem 18 (Soundness of Optimized Encoding) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1Nb} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_1 \triangleright \text{DropSyncs}(\lceil \Omega \rceil; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}(\lceil \Omega' \rceil; \mathcal{P}_1)$.*

6 Generalized Encoding for n -Neighbor Restricted Rules

In this section, we show a generalized encoding scheme for n -neighbor restricted rules. The basic 1-neighbor restricted encoding of Figure 13 implements a consensus protocol between two nodes. Specifically, this encoding implements a *two-phase commit protocol* [SS83] lead by an initial round of matching. Rules r_1 and r_2 represent the *matching phase*, while r_3 the *voting phase*, and r_4a and r_4b the *commit phase*. The encoding of n -neighbor restricted rules is then an implementation of a general consensus protocol that establishes consensus of a rule application among the primary location X (acting as the *coordinator* of the consensus) and n other directly connected and isolated neighbor locations Y_i for $i \in [1, n]$ (acting as the *cohorts* of the consensus).

We first show an example of a 2-neighbor restricted rule encoding into 0-neighbor restricted rules. Figure 16 shows an example of a 2-neighbor restricted rule that sums up all values in an acyclic graph and distributes the value to all nodes of the graph. The primary location X must synchronize with neighbor locations Y and Z to attempt to “atomically” observe each location’s matching obligations and unanimously decide to commit or abort the rule instance. The translation function \rightsquigarrow_{nNb} decomposes the rule into a sequence of 0-neighbor restricted rules that implement the consensus protocol that coordinates the application of rule s in the following manner: Rule s_1 models the broadcast of the requests (to each matching neighbor Y and Z) to match on rule s should the primary location X possess its rule head obligations. We assume that constraints *child* are persistent and hence never deleted. Note that in the body, we introduce an existential variable E and each request is parameterized by it. E is known as a *destination* [Pfe04], used to group together all synchronizing constraints that originate from a particular request (r_req) from primary location X to synchronize a rule application of s . This provides a form of logical isolation from synchronizing constraints that belongs to other instances of this consensus building¹⁵. Specifically in this context, we will refer to destinations like E as *consensus destinations*. Rules s_2Y and s_2Z implements neighbors Y ’s and Z ’s respective match response to X (s_match_Y and s_match_Z) only if they possess their rule head obligations. Destination E is passed on for reasons discussed earlier. Rule s_3 implements primary location X attempt to *commit* if it has confirmation from each neighbor that each possesses matching instances of their rule

¹⁵For the 1-neighbor restricted encoding detailed in Section 5 we can omit destinations because isolation is achieved simply by passing all free-variables between primary location X and its only neighbor location Y .

$$\begin{array}{c}
s : [X] \text{child}(Y), [X] \text{child}(Z) \setminus [X] \text{val}(A), [Y] \text{val}(B), [Z] \text{val}(C) \\
\iff [X] \text{val}(A+B+C), [Y] \text{val}(A+B+C), [Z] \text{val}(A+B+C) \\
\Downarrow_{\text{qNB}} \\
\left(\begin{array}{l}
s_1 : [X] \text{child}(Y), [X] \text{child}(Z), [X] \text{val}(A) \\
\implies \exists E. [Y] s_req_Y(X, Y, Z, A, E), [Z] s_req_Z(X, Y, Z, A, E) \\
s_2\bar{Y} : [\bar{Y}] \text{val}(\bar{B}) \setminus [\bar{Y}] s_req_Y(\bar{X}, \bar{Y}, \bar{Z}, \bar{A}, \bar{E}) \iff [\bar{X}] s_match_Y(\bar{B}, \bar{E}) \\
s_2Z : [Z] \text{val}(C) \setminus [Z] s_req_Z(X, Y, Z, A, E) \iff [X] s_match_Z(C, E) \\
s_3 : [\bar{X}] \text{child}(\bar{Y}), [\bar{X}] \text{child}(\bar{Z}) \setminus [\bar{X}] \text{val}(\bar{A}), [\bar{X}] s_match_Y(\bar{B}, \bar{E}), [\bar{X}] s_match_Z(\bar{C}, \bar{E}) \\
\iff [X] s_commit_X(Y, Z, A, B, C, E), [Y] s_vote_Y(X, B, E), [Z] s_vote_Z(X, C, E) \\
s_4a\bar{Y} : [\bar{Y}] \text{val}(\bar{B}), [\bar{Y}] s_vote_Y(\bar{X}, \bar{B}, \bar{E}) \iff [\bar{X}] s_commit_Y(\bar{Y}, \bar{B}, \bar{E}) \\
s_4bY : [Y] s_vote_Y(X, B, E) \iff [X] s_abort(E) \\
s_4aZ : [Z] \text{val}(C), [Z] s_vote_Z(X, C, E) \iff [X] s_commit_Z(Z, C, E) \\
s_4bZ : [Z] s_vote_Z(X, C, E) \iff [X] s_abort(E) \\
s_5a : [\bar{X}] s_commit_X(\bar{Y}, \bar{Z}, \bar{A}, \bar{B}, \bar{C}, \bar{E}), [\bar{X}] s_commit_Y(\bar{Y}, \bar{B}, \bar{E}), [\bar{X}] s_commit_Z(\bar{Z}, \bar{C}, \bar{E}) \\
\iff [X] \text{val}(A+B+C), [Y] \text{val}(A+B+C), [Z] \text{val}(A+B+C) \\
s_5b\bar{X} : [\bar{X}] s_abort(\bar{E}) \setminus [\bar{X}] s_commit_X(\bar{Y}, \bar{Z}, \bar{A}, \bar{B}, \bar{C}, \bar{E}) \iff [\bar{X}] \text{val}(\bar{A}) \\
s_5bY : [X] s_abort(E) \setminus [X] s_commit_Y(Y, B, E) \iff [Y] \text{val}(B) \\
s_5bZ : [X] s_abort(E) \setminus [X] s_commit_Z(Z, C, E) \iff [Z] \text{val}(C)
\end{array} \right)
\end{array}$$

Figure 16: Graph Sum Example: 2-Neighbor Restricted Rule

head obligations, and its own obligations are still valid. Upon applying this rule instance, X consumes the simplified components of its rule head obligations (i.e, $[X] \text{val}(A)$), produces the synchronizing constraint s_commit_X as a witness to this partial rule application of s and broadcasts to each neighbor Y and Z a call to vote for commitment to this rule instance. Rules s_4aY and s_4aZ implement successful commits in Y and Z only if each still possesses their respective rule head obligations, producing r_commit_Y or r_commit_Z synchronizing constraints as witnesses to their respective partial rule applications. Rules s_4bY and s_4bZ implements the alternative scenario where either neighbor has to abort the rule instance application. Rule s_5a implements the successful commitment for the rule instance of s . Specifically, if primary location X receives matching commit synchronizing constraints from its neighbors Y and Z , we rewrite with the body of s and thus completing the rule application of s . Finally, rules r_5bX , r_5bY and r_5bZ implements the roll back procedure for the cases that any neighbor Y or Z vote to abort the rule instance because it can no longer fulfill its matching obligation. Specifically, each rule replaces a *commit* synchronizing constraint of a given location (that witnesses the partial rule application of that location) with the simplified rule head obligation of that particular location (e.g. $[X] \text{val}(A)$), thereby rolling back a failed non commit free state into a commit free consistent state.

This encoding of rule s is effectively an implementation of a consensus protocol [AD76] in 0-neighbor restricted rules. It establishes a consensus between the primary location X as the *coordinator* of the consensus and neighboring location Y and Z as *cohorts* of the consensus. This encoding implements a *two-phase commit protocol* [SS83] that is lead by an initial round of matching. Specifically, the s_1 rule and s_2x rules collectively represents the *matching phase*, while s_3 rule and s_4xx rules represent the *voting phase*, and s_5a rule and

$$r : [X]P_x, [X]P'_x, (\biguplus_{i \in \mathcal{I}_n} [Y_i]P_i, [Y_i]P'_i) \setminus [X]S_x, (\biguplus_{i \in \mathcal{I}_n} [Y_i]S_i) \iff G_x, (\biguplus_{i \in \mathcal{I}_n} G_i) \mid \exists \bar{Z}. D$$

$$\underbrace{\qquad\qquad\qquad}_{\text{0-neighbor}}$$

$$\left(\begin{array}{l} r_1 : [X]P_x, [X]P'_x, [X]S_x \implies G_x \mid \exists E. (\biguplus_{i \in \mathcal{I}_n} [Y_i]r_req_i(Xs, E)) \\ \hline (\biguplus_{i \in \mathcal{I}_n} \text{MatchRule}(i, S_x)) \\ \hline r_3 : [\bar{X}]P_x \setminus [\bar{X}]P'_x, [\bar{X}]S_x, (\biguplus_{i \in \mathcal{I}_n} [\bar{X}]r_match_i(Ys_i, E)) \\ \iff [X]r_commit_X(Rs, E), (\biguplus_{i \in \mathcal{I}_n} [Y_i]r_vote_i(X, Ys_i, E)) \\ \hline (\biguplus_{i \in \mathcal{I}_n} \text{VoteRules}(i)) \\ \hline r_5a : [X]r_commit_X(Rs, E), (\biguplus_{i \in \mathcal{I}_n} [X]r_commit_i(Y_i, Ys_i, E)) \\ \iff \exists \bar{Z}. [X]P'_x, (\biguplus_{i \in \mathcal{I}_n} [Y_i]P'_i), D \\ \hline r_5bx : [\bar{X}]r_abort(\bar{E}) \setminus [\bar{X}]r_commit_X(\bar{R}s, \bar{E}) \iff [\bar{X}]P'_x, [\bar{X}]S_x \\ \hline (\biguplus_{i \in \mathcal{I}_n} r_5bi : [\bar{X}]r_abort(\bar{E}) \setminus [\bar{X}]r_commit_i(\bar{Y}_i, \bar{Y}s_i, \bar{E}) \iff [\bar{Y}_i]P'_i, [\bar{Y}_i]S_i) \end{array} \right)$$

where

All $c \in P_x$ are persistent constraints and all $c' \in P'_x$ are non-persistent constraints

All $c \in (\biguplus_{i \in \mathcal{I}_n} P_i)$ are persistent constraints and all $c' \in (\biguplus_{i \in \mathcal{I}_n} P'_i)$ are non-persistent constraints

$\text{MatchRule}(i, \cdot) = r_2i : [Y_i]P_i, [Y_i]P'_i, [Y_i]S_i, r_req_i(Xs, E) \implies G_i \mid [X]r_match_i(Ys_i, E)$

$\text{MatchRule}(i, S) = r_2i : [Y_i]P_i, [Y_i]P'_i, [Y_i]S_i \setminus r_req_i(Xs, E) \iff G_i \mid [X]r_match_i(Ys_i, E)$ if $S \neq \cdot$.

$$\text{VoteRules}(i) = \begin{cases} r_4ai : [Y_i]P_i \setminus [Y_i]P'_i, [Y_i]S_i, [Y_i]r_vote_i(X, Ys_i, E) \iff [X]r_commit_i(Y_i, Ys_i, E) \\ r_4bi : [Y_i]r_vote_i(X, Ys_i, E) \iff [X]r_abort(E) \end{cases}$$

$$Xs = \text{FV}(P_x, S_x) \quad Ys_i = \text{FV}(P_i, S_i) \quad Rs = \text{FV}(P_x, S_x, (\biguplus_{i \in \mathcal{I}_n} P_i, S_i)) \quad \text{FV}(G_x) \subseteq Xs \quad \text{FV}(G_i) \subseteq Xs \cup Ys_i$$

Figure 17: 0-Neighbor Restricted Encoding for n -Neighbor Restricted Rules

s_5bx rules the *commit phase*.

The generalized form of this n -neighbor restricted encoding is shown in Figure 17. Similarly to the 1-neighbor restricted encoding shown in Figure 13, this encoding comprises a set of 0-neighbor restricted rules, referred to as *encoding rules*, which collectively implements a n -consensus protocol that allows $n + 1$ locations (one primary and n neighbor locations) to unanimously decide on the application of an instance of a n -neighbor restricted rule. Specifically, each encoding rule implements a sub-routine of this n -consensus protocol in the following manner:

1. Encoding rule r_1 models the initiation of the consensus by the primary location X : If location X fulfills the primary matching obligations (P_x, P'_x, S_x where G_x), it sends a set of requests $r_req_i(Xs, E)$ to n candidate locations which will participate in a round of consensus building, uniquely identified by the existential variable E .
2. Each encoding rule r_2i models the asynchronous response of each neighbor location if its matching obligation is fulfilled (P_i, P'_i, S_i where G_i), replying to primary location X with a $r_match_i(Ys_i, E)$ constraint.
3. Encoding rule r_3 states that if location X 's matching obligation is still available and if X has received

a match response $r_match_i(Y_{s_i}, E)$, from each neighbor involved in consensus E , it commits its non-persistent matching obligations (P'_x and S_x) and calls each neighbor location of consensus E to vote for this rule application ($r_vote_i(X, Y_{s_i}, E)$).

4. Each encoding rule r_4ai states that if the neighbor location still has its matching obligation (P_i, P'_i, S_i where G_i) and receives a call to vote for the rule application of r ($r_vote_i(X, Y_{s_i}, E)$), it commits its non-persistent matching obligations (P'_i and S_i) and sends its decision to commit to location X ($r_commit_i(Y_i, Y_{s_i}, E)$). Otherwise, encoding rule r_4b states that if the neighbor location no longer has its matching obligation, it sends its decision to abort consensus E to location X ($r_abort(E)$).
5. Encoding rule r_5a models the completion of rule application r in the event that all neighbor locations of consensus E have decided to commit to the rule application. Specifically, X receives a full set of commit synchronizing constraints of consensus E , during which it completes the rule application by applying the rule body of the n -neighbor restricted rule r (i.e, $\exists \bar{Z}.D$) and return all non-persistent propagated rule head matching obligations of the rule r (i.e, P'_x and $\biguplus_{i \in \mathcal{I}_n} P'_i$)
6. Encoding rule r_5bx and each encoding rule r_5bi models the abortion of consensus E . Specifically, upon receive $r_abort(E)$ from any neighbor location of consensus E , r_5bx rolls back location X 's commitment to E by return its non-persistent matching obligations (i.e, P'_x and S_x) while each r_5bi does the same of each neighbor location.

Note that each derivation step of an encoding rule that models this consensus protocol is designed so that it can be interleaved with concurrent and non-overlapping derivations without compromising the correctness of the consensus building. Since we assume a lossless network communication between locations, this encoding implements a two-phase commit consensus protocol, which is sufficed to guarantee correctness in a lossless network. Given a n -neighbor restricted program encoding \mathcal{P}_0 , we write $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_j(E)} \Omega'$ as a derivation step that involves the application of an instance of the encoding rule r_j with the consensus destination instantiated to some E . We will refer to such derivation steps as E derivation steps.

We now establish the soundness of the n -neighbor restricted rule encoding. The steps of the proof of the soundness theorem for this encoding are similar to the steps shown in Section 5. Firstly, Lemma 19 states that the encoding operation \rightsquigarrow_{nNb} preserves local quiescence of CHR^e programs.

Lemma 19 (\rightsquigarrow_{nNb} Preserves Local Quiescence) *Given a locally quiescent n -neighbor restricted program \mathcal{P}_n and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$, then \mathcal{P}_0 is also locally quiescent.*

Lemma 20 states the property that all applications of encoding rules are part of a derivation of some prefix sequence of a consensus building routine of a specific consensus destination E .

Lemma 20 (Prefix Executions of n -Neighbor Encoding) *Given a locally quiescent n -neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$, if we have reachable states Ω, Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$, then we have the following:*

1. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_2i(E)} \Omega'$ for r_2i encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r_1(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
2. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_3(E)} \Omega'$ for r_3 encoding rule of some consensus destination E , then there exists some reachable derivation $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^* \Omega'''$ which contains n derivation steps $r_2i(E)$ for each neighbor Y_i of $i \in \mathcal{I}_n$, and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
3. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_4ai(E)} \Omega'$, for r_4ai encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r_3(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$

4. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_4bi(E)} \Omega'$, for r_4bi encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r_3(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
5. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_5a(E)} \Omega'$ for r_5a encoding rule of some consensus destination E , then there exists some reachable derivation $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^* \Omega'''$ which contains n derivation steps $r_2i(E)$ for each neighbor Y_i of $i \in \mathcal{I}_n$, and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
6. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_5bx(E)} \Omega'$ of $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_5bi(E)} \Omega'$ for r_5bx or r_5bi encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r_4bj(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$

Lemma 21 states the property that non-commit free states that contain partial applications of n -neighbor restricted rules, can always eventually derive commit free state that only contain complete applications of n -neighbor restricted rules.

Lemma 21 (n -Neighbor Commit-Free Reachability) *Given a locally quiescent n -neighbor restricted program \mathcal{P}_n and a 0 -neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$ and states Ω reachable by \mathcal{P}_0 , if Ω is not commit-free, then there exists some commit-free state Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.*

Similarly to Lemma 14, Lemma 22 states that given any ω_α^e derivation between two commit free state $\mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ labeled \mathcal{I} , we can safely permute \mathcal{I} so that for each consensus destination E , derivation step of $r_3(E)$ is immediately preceded by all derivation steps of consensus destination E that occur after it in \mathcal{I} . This simply allows us to focus on derivations which are controlled and orderly.

Lemma 22 (n -Neighbor Encoding Rule Serializability) *Given an n -neighbor restricted and locally quiescent CHR^e program \mathcal{P}_n and a 0 -neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}, \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ labeled as derivation \mathcal{I} , if \mathcal{I} contain derivations of a n -neighbor restricted encoding rule instance r of consensus destination E such that:*

- We have a successful application of rule instance r of E , then there exists a valid derivation $\mathcal{I}' = \mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E$ consists the following non-interleaving derivation sequence:
 1. The derivation step $r_3(E)$.
 2. All n derivation steps of $r_4ai(E)$ each of Y_i for $i \in \mathcal{I}_n$.
 3. The derivation step $r_5a(E)$.
- We have an aborted attempt to apply rule instance r of E , then there exists a valid derivation $\mathcal{I}' = \mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E$ consists a non-interleaving derivation sequences of derivation step $r_3(E)$ followed by all n derivation steps of $r_4ai(E)$ or r_4bi each of Y_i for $i \in \mathcal{I}_n$, followed by
 1. The derivation step $r_3(E)$.
 2. All m derivation steps of either $r_4ai(E)$ each of Y_i for $i \in \mathcal{I}_n$ where $m \leq n$.
 3. All $n - m$ derivation steps of either $r_4bi(E)$ each of Y_i for $i \in \mathcal{I}_n$.
 4. The derivation step $r_5b(E)$.
 5. All $n - m$ derivation steps of $r_5bi(E)$ of Y_i for $i \in \mathcal{I}_n$.

Theorem 23 states the soundness of the n -neighbor restricted encoding scheme.

Theorem 23 (Soundness of n -Neighbor Restricted Encoding) *Given an n -neighbor restricted and locally quiescent CHR^e program \mathcal{P}_n and a 0 -neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_n \triangleright \text{DropSyncs}([\Omega]; \mathcal{P}_n) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega']; \mathcal{P}_n)$.*

7 Related Works

An extension of Datalog for implementing network protocols is introduced in [LCG⁺06]. It defines *link restricted Datalog rules* along with the idea of *rule localization* which encodes link restricted rules into local Datalog rules. Our work adapts, expands and generalizes these ideas to the context of distributed multiset rewriting. A distributed and incremental algorithm for computing rule matchings involved in such distributed Datalog systems is presented in [NJLS11] and is analogous to what ω_0^e provides for CHR^e . Our work draws inspiration from the programming language Meld [ARLG⁺09]. Meld introduces a multitude of extensions to Datalog, including linearity, aggregates and comprehensions. Work in [CARG⁺12] adapts Meld for the context of general purpose multicore parallel programming. The ω_0^e semantics extends the refined CHR operational semantics [DSdlBH04]. A parallel execution model of CHR is introduced in [LS11], which assumes a more tightly coupled model of concurrent execution, where multiple execution threads compute rewritings concurrently over a *global constraint store*.

8 Conclusion and Future Works

We introduced CHR^e , an extension of CHR with located constraints and n -neighbor restricted rewrite rules for programming an ensemble of distributed computing entities. We defined the ω_α^e abstract semantics and ω_0^e operational semantics of CHR^e and showed their soundness. We gave an optimized encoding for 1-neighbor restricted rules into 0-neighbor restricted rules. Following this, we generalize this encoding scheme for n -neighbor restricted rules. We have developed a prototype implementation of CHR^e in Python with MPI (Message Passing Interface) as a proof of concept and demonstrated its relative scalability in distributed execution. In the future we intend to develop a more practical and competitive implementation of CHR^e in C or C++, imbued with existing CHR optimization techniques [Sch05]. We also intend to explore an implementation over higher-level distributed graph processing frameworks like Google's Pregel [MAB⁺10]. Additionally, we intend to explore using ω_0^e to serve as an operational semantics that describes the core multiset rewriting fragment of Meld and derive extensions like aggregates and comprehensions as higher-level language encodings into ω_0^e . Our works on encoding n -neighbor restricted rules can be also applied to extend Meld.

References

- [AD76] P. A. Alsberg and J. D. Day. A Principle for Resilient Sharing of Distributed Resources. In *Proc. of the 2nd Int. Conf. on Software Engineering, ICSE '76*, pages 562–570, Los Alamitos, CA, USA, 1976. IEEE Computer Society Press.
- [ARLG⁺09] M. P. Ashley-Rollman, P. Lee, S. C. Goldstein, P. Pillai, and J. D. Campbell. A Language for Large Ensembles of Independently Executing Nodes. In *Proc. of ICLP'09*, volume 5649, pages 265–280. Springer-Verlag, 2009.
- [CARG⁺12] F. Cruz, M. P. Ashley-Rollman, S. C. Goldstein, Ricardo Rocha, and F. Pfenning. Bottom-Up Logic Programming for Multicores. In Vítor Santos Costa, editor, *Proc. of DAMP 2012*. ACM Digital Library, January 2012.
- [DSdlBH04] G. J. Duck, P. J. Stuckey, M. Garcia de la Banda, and C. Holzbaur. The Refined Operational Semantics of Constraint Handling Rules. In *In 20th Int. Conf. on Logic Programming ICLP'04*, pages 90–104. Springer, 2004.
- [Frü94] T. Frühwirth. Constraint handling rules. In *Constraint Programming*, pages 90–107, 1994.

- [LCG⁺06] B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative Networking: Language, Execution and Optimization. In *Proc. of SIGMOD '06*, pages 97–108. ACM, 2006.
- [LS11] E. S. L. Lam and M. Sulzmann. Concurrent Goal-based Execution of Constraint Handling Rules. *TPLP*, 11(6):841–879, 2011.
- [MAB⁺10] G. Malewicz, M. H. Austern, A. J.C Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: A System for Large-Scale Graph Processing. In *Proc. of Int. Conf. on Management of data*, SIGMOD, pages 135–146, USA, 2010. ACM.
- [NJLS11] V. Nigam, L. Jia, B. T. Loo, and A. Scedrov. Maintaining distributed logic programs incrementally. In *Proc. of PPDP'11*, pages 125–136. ACM, 2011.
- [Pfe04] F. Pfenning. Substructural Operational Semantics and Linear Destination-Passing Style (Invited Talk). In *APLAS*, page 196, 2004.
- [Sch05] T. Schrijvers. Analyses, Optimizations and Extensions of Constraint Handling Rules: Ph.D. Summary. In *ICLP*, pages 435–436, 2005.
- [SL08] M. Sulzmann and E. S. L. Lam. Parallel Execution of Multi-set Constraint Rewrite Rules. In *PPDP*, pages 20–31, 2008.
- [SS83] D. Skeen and M. Stonebraker. A Formal Model of Crash Recovery in a Distributed Systems. *IEEE Transactions on Software Engineering*, pages 219–228, 1983.

A Proofs

A.1 Proofs for the ω_α^e Semantics

Lemma 1 (Well-Formedness Preservation of ω_α^e Derivations) *Given a well-formed CHR^e program \mathcal{P} , a well-formed state \mathcal{A} and state \mathcal{A}' , if $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ then \mathcal{A}' is well-formed.*

Proof: We proof by induction on ω_α^e derivation steps. We consider the base case such that $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}$ of zero ω_α^e derivation step, hence trivially we have $\mathcal{A}' = \mathcal{A}$ is well-formed. Now suppose well-formedness preservation holds for derivations $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}''$ of m number of derivation steps, hence state \mathcal{A}'' is well-formed. We consider the successor state \mathcal{A}' such that $\mathcal{P} \triangleright \mathcal{A}'' \mapsto_{\omega_\alpha^e} \mathcal{A}'$. Specifically, we consider this derivation step in its most general form as specified in Figure 4: Since all rules in \mathcal{P} are well-formed, we are guaranteed that all fragments of rule body (i.e, D'_i , D'_j and D'_k) consist of well-formed constraints, hence instantiated rule body fragments D_i , D_j and D_k are well-formed as well. Constraint stores \bar{S}_i and \bar{S}_j extended with well-formed constraints result to well-formed stores as well, hence new components of the state (i.e, $\biguplus_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, D_i \rangle_{k_i}$, $\biguplus_{j \in \mathcal{I}_m} \langle \bar{S}_j, D_j \rangle_{k_j}$, $\biguplus_{l \in \mathcal{I}_m} \langle \bar{S}_l, D_l \rangle_{k_l}$) are all well-formed. We show that new location names k_l (of existential forwarding locations) are unique and do not exist in \mathcal{A} : location variables k'_l are defined such that they do not appear as term arguments of rule heads (i.e, $k'_l \notin (\text{FV}(\biguplus_{i \in \mathcal{I}_n} P_i) \cup \text{FV}(\biguplus_{i \in \mathcal{I}_n} S_i))$) nor are they matching locations (i.e, $\forall i \in \mathcal{I}_n. k'_i \neq k'_i$). Therefore, variables k'_l are treated as existential variables and the meta operation $\text{Inst}(-)$ would simply assign it a fresh constant k_l that has never appeared before. Hence the state \mathcal{A}' is well-formed. \square

Lemma 2 (Well-Formedness Preservation of $\llbracket - \rrbracket$ Translation) *Given a well-formed CHR^e object o , $\llbracket o \rrbracket$ is a well-formed CHR object.*

Proof: We prove by structural induction over well-formed CHR^e objects, where CHR constraints $p(\vec{t})$ are the base elements:

- *Constraints:* $\llbracket p(\vec{t}) \rrbracket^l = p(l, \vec{t})$. $p(\vec{t})$ is well-formed, namely terms \vec{t} are well-formed. l is either a variable or location name (a constant) which is a well-formed term, hence $p(l, \vec{t})$ is well-formed.
- *Constraint Store:* For $\llbracket \emptyset \rrbracket^l = \emptyset$, an empty set which is by default a well-formed store. For $\llbracket c, \bar{S} \rrbracket^l$, suppose $\llbracket c \rrbracket^l$ and $\llbracket \bar{S} \rrbracket^l$ are well-formed constraint stores (via constraint and constraint store cases respectively), then the multiset union $\llbracket c \rrbracket^l, \llbracket \bar{S} \rrbracket^l$ is also a well-formed constraint store.
- *Abstract Ensemble States:* For $\llbracket \emptyset \rrbracket = \emptyset$, an empty set is by default a well-formed constraint store. For $\llbracket \mathcal{A}, \langle \bar{S} \rangle_k \rrbracket = \llbracket \mathcal{A} \rrbracket, \llbracket \bar{S} \rrbracket^k$, suppose that $\llbracket \mathcal{A} \rrbracket$ is a well-formed (this case) and furthermore, it is a constraint store. Also suppose that $\llbracket \bar{S} \rrbracket^k$ is a well-formed constraint store (constraint store case), then the multiset union of the two is also a well-formed constraint store.
- *Rule Body:* At the top-level, $\llbracket \exists \bar{x}. D \rrbracket$ is simply $\exists \bar{x}. \llbracket D \rrbracket$ and is clearly well-formed, if $\llbracket D \rrbracket$ is well-formed. For $\llbracket true \rrbracket = true$, $true$ is by default a well-formed CHR rule body. For $\llbracket [l]c, D \rrbracket = \llbracket c \rrbracket^l, \llbracket D \rrbracket$, suppose that $\llbracket c \rrbracket^l$ is a well-formed CHR constraint (constraint case) and that $\llbracket D \rrbracket$ is a well-formed rule body (this case), hence the composition $\llbracket c \rrbracket^l, \llbracket D \rrbracket$ is indeed a well-formed CHR rule body.
- *Rule Head:* For $\llbracket \cdot \rrbracket = \cdot$, \cdot is by default a well-formed CHR rule head. For $\llbracket [l]c, H \rrbracket = \llbracket c \rrbracket^l, \llbracket H \rrbracket$, suppose that $\llbracket c \rrbracket^l$ is a well-formed CHR constraint (constraint case) and that $\llbracket H \rrbracket$ is a well-formed rule head (this case), hence the composition $\llbracket c \rrbracket^l, \llbracket H \rrbracket$ is indeed a well-formed CHR rule head.
- *Rule:* $\llbracket r : P \setminus S \iff G \mid B \rrbracket = r : \llbracket P \rrbracket \setminus \llbracket S \rrbracket \iff G \mid \llbracket B \rrbracket$, Given multisets of well-formed located constraints P, S and B , we have $\llbracket P \rrbracket, \llbracket S \rrbracket$ and $\llbracket B \rrbracket$ as well-formed CHR constraints. Hence the resultant CHR rule is well-formed.
- *Program:* For $\llbracket \cdot \rrbracket = \cdot$, empty \cdot is by default a well-formed CHR program. For $\llbracket R \mathcal{P} \rrbracket$, suppose $\llbracket R \rrbracket$ is a well-formed CHR rule (rule case) and $\llbracket \mathcal{P} \rrbracket$ is a well-formed CHR program (this case), then the composition of the two is a well-formed CHR program.

Hence we have shown that given a well-formed CHR^e object o , $\llbracket o \rrbracket$ is a well-formed CHR object. \square

Theorem 3 (Soundness of ω_α^e) *Given a CHR^e program \mathcal{P} and abstract states \mathcal{A} and \mathcal{A}' , if $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$, then $\llbracket \mathcal{P} \rrbracket \triangleright \llbracket \mathcal{A} \rrbracket \mapsto_{\omega_\alpha^e}^* \llbracket \mathcal{A}' \rrbracket$.*

Proof: We prove by induction on ω_α^e derivation steps: We consider the base case such that $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}$ of zero derivation steps. With Lemma 2, since $\llbracket \mathcal{P} \rrbracket$ and $\llbracket \mathcal{A} \rrbracket$ are well-formed CHR program and store, we trivially have $\llbracket \mathcal{P} \rrbracket \triangleright \llbracket \mathcal{A} \rrbracket \mapsto_{\omega_\alpha^e}^* \llbracket \mathcal{A} \rrbracket$. We now consider the inductive case: Suppose we have $\mathcal{P} \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}''$ such that $\llbracket \mathcal{P} \rrbracket \triangleright \llbracket \mathcal{A} \rrbracket \mapsto_{\omega_\alpha^e}^* \llbracket \mathcal{A}'' \rrbracket$ for some derivation of m steps. We consider the successor derivation step that extends this by one, in its most general form as specified in Figure 4:

$$\mathcal{P} \triangleright \mathcal{A}', \left(\begin{array}{c} \biguplus_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, S_i \rangle_{k_i}, \\ \biguplus_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j} \end{array} \right) \mapsto_{\omega_\alpha^e} \mathcal{A}', \left(\begin{array}{c} \biguplus_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, D_i \rangle_{k_i}, \\ \biguplus_{j \in \mathcal{I}_m} \langle \bar{S}_j, D_j \rangle_{k_j}, \\ \biguplus_{l \in \mathcal{I}_e} \langle D_l \rangle_{k_l} \end{array} \right)$$

$$\text{such that } \mathcal{A}'' = \mathcal{A}', \left(\biguplus_{i \in \mathcal{I}_n} \langle \bar{S}_i, P_i, S_i \rangle_{k_i}, \biguplus_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j} \right)$$

and from this we can infer (for some rule name r and guard G) the rule instance in \mathcal{P} that corresponds to this derivation step is:

$$r : \biguplus_{i \in \mathcal{I}_n} [k_i] P_i \setminus \biguplus_{i \in \mathcal{I}_n} [k_i] S_i \iff G \mid \exists \bar{x}. \biguplus_{i \in \mathcal{I}_n} [k_i] D_i, \biguplus_{j \in \mathcal{I}_m} [k_j] D_j, \biguplus_{l \in \mathcal{I}_e} [k_l] D_l$$

Hence, to complete the induction proof, we need to show that the following holds:

$$\llbracket \mathcal{P} \rrbracket \triangleright \llbracket \mathcal{A}' \rrbracket, \left(\begin{array}{c} \biguplus_{i \in \mathcal{I}_n} \llbracket \bar{S}_i \rrbracket^{k_i}, \llbracket P_i \rrbracket^{k_i}, \llbracket S_i \rrbracket^{k_i}, \\ \biguplus_{j \in \mathcal{I}_m} \llbracket \bar{S}_j \rrbracket^{k_j} \end{array} \right) \mapsto_{\omega_\alpha} \llbracket \mathcal{A}' \rrbracket, \left(\begin{array}{c} \biguplus_{i \in \mathcal{I}_n} \llbracket \bar{S}_i \rrbracket^{k_i}, \llbracket P_i \rrbracket^{k_i}, \llbracket D_i \rrbracket^{k_i}, \\ \biguplus_{j \in \mathcal{I}_m} \llbracket \bar{S}_j \rrbracket^{k_j}, \llbracket D_j \rrbracket^{k_j}, \\ \biguplus_{l \in \mathcal{I}_e} \llbracket D_l \rrbracket^{k_l} \end{array} \right) \quad (A)$$

We apply the translation function on rule instance r , which we know to be an instance of some rule $R \in \mathcal{P}$, hence we have the CHR interpretation of rule instance r :

$$\begin{aligned} & \llbracket r : \biguplus_{i \in \mathcal{I}_n} [k_i] P_i \setminus \biguplus_{i \in \mathcal{I}_n} [k_i] S_i \iff G \mid \exists \bar{x}. \biguplus_{i \in \mathcal{I}_n} [k_i] D_i, \biguplus_{j \in \mathcal{I}_m} [k_j] D_j, \biguplus_{l \in \mathcal{I}_e} [k_l] D_l \rrbracket \\ = & r : \biguplus_{i \in \mathcal{I}_n} \llbracket P_i \rrbracket^{k_i} \setminus \biguplus_{i \in \mathcal{I}_n} \llbracket S_i \rrbracket^{k_i} \iff G \mid \exists \bar{x}. \biguplus_{i \in \mathcal{I}_n} \llbracket D_i \rrbracket^{k_i}, \biguplus_{j \in \mathcal{I}_m} \llbracket D_j \rrbracket^{k_j}, \biguplus_{l \in \mathcal{I}_e} \llbracket D_l \rrbracket^{k_l} \end{aligned}$$

Note that this CHR rule instance corresponds the rule instance applied in CHR derivation (A), thus proving that derivation (A) is valid and thus, ω_α^e is sound and preserves derivability. \square

A.2 Proofs for the ω_0^e Semantics

Lemma 4 (Well-Formedness Preservation of ω_0^e Derivations) *Given a well-formed 0-neighbor restricted CHR^e program \mathcal{P} , a well-formed state Ω and state Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ then Ω' is well-formed.*

Proof: We proof by induction on ω_0^e derivation steps. We consider the base case such that $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega$ of zero ω_0^e derivation step, hence trivially we have $\Omega' = \Omega$ is well-formed. Now suppose well-formedness preservation holds for derivations $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega''$ of m number of derivation steps, hence state Ω'' is well-formed. We consider all possible forms (each rule of Figure 7) of successor derivation step that maps Ω'' to some state Ω' (i.e., $\mathcal{P} \triangleright \mathcal{A}'' \mapsto_{\omega_\alpha^e} \mathcal{A}'$):

- (Flush) Step: This step modifies a location $k \in \text{Locs}(\Omega'')$ by moving its buffer \vec{U} to its empty goals. Since \vec{U} (a sequence of constraints) is by definition a valid sequence of goals, hence successor state Ω' is well-formed.
- (Loc 1) This step modifies two distinct locations $k, k' \in \text{Locs}(\Omega'')$ by moving leading goal in k $[k']c$ into the buffer of k' as c . Hence goals of k and buffer of k' are still well-formed and therefore successor state Ω' is well-formed.
- (Loc 2) Step: This step modifies the goals of $k \in \text{Locs}(\Omega'')$ from $[k]c, \vec{G}$ to c, \vec{G} which is clearly still well-formed. Therefore successor state Ω' is well-formed.
- (Loc 3) Step: This step modifies the goals of $k \in \text{Locs}(\Omega'')$ from $[k]c, \vec{G}$ to c, \vec{G} and extends the successor state Ω' with $\langle c ; \emptyset ; \emptyset ; \emptyset \rangle_{k'}$. c, \vec{G} is clearly still well-formed, and premise of the premise of the rule restricts k' to be unique from any other state that appear in Ω'' . Therefore Ω' is well-formed.
- (Act) Step: This step modifies a location $k \in \text{Locs}(\Omega'')$ by changing its leading goal from $p(\vec{t})$ to $p(\vec{t})\#d : 1$ and adding $p(\vec{t})\#d$ to its store. Since $p(\vec{t})\#d : 1$ is a valid goal and d is a fresh id, successor state Ω' is well-formed.
- (Simp) Step: This step modifies a location $k \in \text{Locs}(\Omega'')$ by removing a fragment ' $S, c\#d$ ' from the store, removing leading goal $c\#d : i$ from the goals and adding rule body instance $\text{NF}(\text{Inst}(\theta B))$ to the goals. Removing elements from store and goal have benign effects on well-formedness and since CHR^e program \mathcal{P} contains only well-formed rules, rule body instance $\text{NF}(\text{Inst}(\theta B))$ is a ground and well-form multiset of constraints. We treat this multiset of constraints as an arbitrarily ordered sequence of constraints when appending to the goals. Therefore successor state Ω' is well-formed.

- (Prop) Step: This step is similar to (Simp) with exception that the active constraint $c\#d : i$ is retained in the goal, $c\#d$ is retained in the store and $(d, \text{Ids}(P, S))$ is added to the history. Since $(d, \text{Ids}(P, S))$ is a valid history element and premise restricts that it must not already appear in the current history (thus enforcing the set semantics of the history), successor state Ω' is well-formed.
- (Next) Step: This step modifies a location $k \in \text{Locs}(\Omega'')$ by changing its leading goal from $c\$d : i$ to $c\#d : i + 1$, which is a valid goal. Hence successor state Ω' is well-formed.
- (Drop) Step: This step modifies a location $k \in \text{Locs}(\Omega'')$ by removing its leading goal. Therefore successor state is well-formed.

Therefore we have shown that all derivation steps of ω_0^e preserves well-formedness. \square

Lemma 5 (Well-Formedness Preservation of $[-]$ Translation) *Given a well-formed ω_0^e operational state Ω , $[\Omega]$ is a well-formed ω_α^e abstract state.*

Proof: We prove by structural induction over well-formed ω_0^e operational state and sub-structural fragments:

- *Ensembles:* $[\Omega]$, it is well-formed assuming that its two components $[\Omega]^{\vec{\mathcal{G}}'}$ and $\biguplus_{k \in (\text{Locs}(\vec{\mathcal{G}}') - \text{Locs}(\Omega))} \langle \vec{\mathcal{G}}' \rangle_k$ are well-formed sets of store and location pairs $(\langle \vec{\mathcal{S}} \rangle_k)$ and all k are unique. Indeed, by definition of both components each are sets of store and location pairs and uniqueness of each k is guaranteed by the assumption that each sub-structure is well-formed and that the second explicitly quantify over locations $k \notin \text{Locs}(\Omega)$ (i.e, $k \in \text{Locs}(\vec{\mathcal{G}}') - \text{Locs}(\Omega)$). For $[\Omega, \langle \vec{\mathcal{U}} \rangle; \vec{\mathcal{G}}; \vec{\mathcal{S}}; \vec{\mathcal{H}}]_k^{\vec{\mathcal{G}}'}$, we show that it is well-formed by assuming that the sub-structure $[\Omega]^{\vec{\mathcal{G}}'}$ is well-formed and that $[\vec{\mathcal{U}}]$, $[\vec{\mathcal{G}}]$, $[\vec{\mathcal{S}}]$ and $\vec{\mathcal{G}}'_k$ are well-formed multisets of constraints. The latter four sub-structures are collapsed into one multiset of constraints, hence $[\Omega, \langle \vec{\mathcal{U}} \rangle; \vec{\mathcal{G}}; \vec{\mathcal{S}}; \vec{\mathcal{H}}]_k^{\vec{\mathcal{G}}'}$ is well-formed. Finally, $[\emptyset]^{\vec{\mathcal{G}}'}$ is the empty set \emptyset which is by default a well-formed abstract state.
- *Buffers:* For $[\emptyset] = \emptyset$, \emptyset is by default a well-formed constraint store. For $[c, \vec{\mathcal{U}}]$, it is well-formed assuming that $[\vec{\mathcal{U}}]$ is a well-formed constraint store. Since c is a valid element of a constraint store, unioning $[\vec{\mathcal{U}}]$ with c forms a well-formed constraint store.
- *Stores:* For $[\emptyset] = \emptyset$, \emptyset is by default a well-formed constraint store. For $[c\#d, \vec{\mathcal{S}}]$, it is well-formed assuming that $[\vec{\mathcal{S}}]$ is a well-formed constraint store. Since we drop d and only retain c , a valid element of a constraint store, unioning $[\vec{\mathcal{S}}]$ with c forms a well-formed constraint store.
- For $[\emptyset] = \emptyset$, \emptyset is by default a well-formed constraint store. For $[c, \vec{\mathcal{G}}]$, it is well-formed assuming that $[\vec{\mathcal{G}}]$ is a well-formed constraint store. Since c is a valid element of a constraint store, unioning $[\vec{\mathcal{G}}]$ with c forms a well-formed constraint store. For $[c\#d : i, \vec{\mathcal{G}}]$ and $[c, \vec{\mathcal{G}}]$, we drop both respective atoms and only rely on the assumption that $[\vec{\mathcal{G}}]$ is well-formed.

Hence we have shown that given a well-formed ω_0^e operational state Ω , $[\Omega]$ is a well-formed ω_α^e abstract state. \square

Theorem 6 (Soundness of ω_0^e) *Given 0-neighbor restricted CHR^e program \mathcal{P} and states Ω and Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$.*

Proof: We proof by induction on ω_0^e derivation steps. We consider the base case such that $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega$ of zero ω_0^e derivation step, hence trivially we have $\Omega' = \Omega$ and that $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega]$ for zero ω_α^e derivation step. Now suppose that for derivations $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega''$ of m number of steps, we have $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$. We consider all possible forms (each rule of Figure 7) of successor derivation step that maps Ω'' to some state Ω' (i.e, $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$), and show that for each instance we have $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$:

- (Flush) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{\mathcal{U}} ; \emptyset ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \emptyset ; \vec{\mathcal{U}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k$$

Since $\lceil \langle \vec{\mathcal{U}} ; \emptyset ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \rceil = \lceil \langle \emptyset ; \vec{\mathcal{U}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \rceil$, hence $\lceil \Omega'' \rceil = \lceil \Omega' \rceil$ and we trivially have $\mathcal{P} \triangleright \lceil \Omega'' \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright \lceil \Omega \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$.

- (Loc 1) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \left(\begin{array}{l} \langle \vec{\mathcal{U}} ; ([k']c, \vec{\mathcal{G}}) ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k, \\ \langle \vec{\mathcal{U}}' ; \vec{\mathcal{G}}' ; \vec{\mathcal{S}}' ; \vec{\mathcal{H}}' \rangle_{k'} \end{array} \right) \mapsto_{\omega_0^e} \Omega, \left(\begin{array}{l} \langle \vec{\mathcal{U}} ; \vec{\mathcal{G}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k, \\ \langle (\vec{\mathcal{U}}', c) ; \vec{\mathcal{G}}' ; \vec{\mathcal{S}}' ; \vec{\mathcal{H}}' \rangle_{k'} \end{array} \right)$$

By definition, $\lceil [k']c, \vec{\mathcal{G}} \rceil = \lceil \vec{\mathcal{G}} \rceil$ and since $k \neq k'$, we have $\text{Goals}(\Omega'')|_k = \text{Goals}(\Omega')|_k$. Thus for location k we have $\langle \lceil \vec{\mathcal{U}} \rceil, \lceil [k']c, \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}} \rceil, \text{Goals}(\Omega'')|_k \rangle_k = \langle \lceil \vec{\mathcal{U}} \rceil, \lceil \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}} \rceil, \text{Goals}(\Omega')|_k \rangle_k$. For location k' , $c, \lceil \vec{\mathcal{U}}' \rceil = \lceil \vec{\mathcal{U}}', c \rceil$ but since $\text{Goals}(\Omega'')|_k = c, \text{Goals}(\Omega')|_k$, we have $\langle \lceil \vec{\mathcal{U}}' \rceil, \lceil \vec{\mathcal{G}}' \rceil, \lceil \vec{\mathcal{S}}' \rceil, \text{Goals}(\Omega'')|_{k'} \rangle_{k'} = \langle \lceil \vec{\mathcal{U}}', c \rceil, \lceil \vec{\mathcal{G}}' \rceil, \lceil \vec{\mathcal{S}}' \rceil, \text{Goals}(\Omega')|_{k'} \rangle_{k'}$. These imply that $\lceil \Omega'' \rceil = \lceil \Omega' \rceil$. We trivially have $\mathcal{P} \triangleright \lceil \Omega'' \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright \lceil \Omega \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$.

- (Loc 2) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{\mathcal{U}} ; ([k]c, \vec{\mathcal{G}}) ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{\mathcal{U}} ; (c, \vec{\mathcal{G}}) ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k$$

By definition, we have $c, \lceil [k]c, \vec{\mathcal{G}} \rceil = \lceil c, \vec{\mathcal{G}} \rceil$ and $\text{Goals}(\Omega'')|_k = c, \text{Goals}(\Omega')|_k$. These implies that we have $\langle \lceil \vec{\mathcal{U}} \rceil, \lceil [k]c, \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}} \rceil, \text{Goals}(\Omega'')|_k \rangle_k = \langle \lceil \vec{\mathcal{U}} \rceil, \lceil c, \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}} \rceil, \text{Goals}(\Omega')|_k \rangle_k$ and hence $\lceil \Omega'' \rceil = \lceil \Omega' \rceil$. We trivially have $\mathcal{P} \triangleright \lceil \Omega'' \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright \lceil \Omega \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$.

- (Loc 3) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{\mathcal{U}} ; ([k']c, \vec{\mathcal{G}}) ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{\mathcal{U}} ; \vec{\mathcal{G}} ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k, \langle c ; \emptyset ; \emptyset ; \emptyset \rangle_{k'}$$

By definition, we have $\lceil [k']c, \vec{\mathcal{G}} \rceil = \lceil \vec{\mathcal{G}} \rceil$. However, since Ω' has the additional location $\langle c ; \emptyset ; \emptyset ; \emptyset \rangle_{k'}$, hence $c, \lceil \Omega'' \rceil^{\text{Goals}(\Omega'')} = \lceil \Omega' \rceil^{\text{Goals}(\Omega')}$. Yet we have the following in the second component of the $\lceil - \rceil$ definition:

$$\biguplus_{k \in (\text{Locs}(\text{Goals}(\Omega'')) - \text{Locs}(\Omega))} \langle k | \text{Goals}(\Omega'') \rangle_k = c, \biguplus_{k \in (\text{Locs}(\text{Goals}(\Omega')) - \text{Locs}(\Omega))} \langle k | \text{Goals}(\Omega') \rangle_k$$

Hence we have $\lceil \Omega'' \rceil = \lceil \Omega' \rceil$ and thus $\mathcal{P} \triangleright \lceil \Omega'' \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright \lceil \Omega \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$.

- (Act) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{\mathcal{U}} ; (p(\vec{t}), \vec{\mathcal{G}}) ; \vec{\mathcal{S}} ; \vec{\mathcal{H}} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{\mathcal{U}} ; (p(\vec{t})\#d : 1, \vec{\mathcal{G}}) ; (\vec{\mathcal{S}}, p(\vec{t})\#d) ; \vec{\mathcal{H}} \rangle_k$$

We have $\lceil p(\vec{t}), \vec{\mathcal{G}} \rceil = p(\vec{t}), \lceil p(\vec{t})\#d : 1, \vec{\mathcal{G}} \rceil$, $p(\vec{t}), \lceil \vec{\mathcal{S}} \rceil = \lceil \vec{\mathcal{S}}, p(\vec{t})\#d \rceil$ and $\text{Goals}(\Omega'')|_k = \text{Goals}(\Omega')|_k$. Hence we have:

$$\langle \lceil \vec{\mathcal{U}} \rceil, \lceil p(\vec{t}), \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}} \rceil, \text{Goals}(\Omega'')|_k \rangle_k = \langle \lceil \vec{\mathcal{U}} \rceil, \lceil p(\vec{t})\#d : 1, \vec{\mathcal{G}} \rceil, \lceil \vec{\mathcal{S}}, p(\vec{t})\#d \rceil, \text{Goals}(\Omega')|_k \rangle_k$$

Hence we have $\lceil \Omega'' \rceil = \lceil \Omega' \rceil$ and thus $\mathcal{P} \triangleright \lceil \Omega'' \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright \lceil \Omega \rceil \mapsto_{\omega_\alpha^e}^* \lceil \Omega' \rceil$.

- (Simp) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); (\bar{S}, P, S, c\#d); \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (\text{NF}(\text{Inst}(\theta B)), \vec{G}); (\bar{S}, P); \bar{H} \rangle_k$$

for some substitution θ . Applying the translation function and simplifying, we need to prove that:

$$\begin{aligned} & \mathcal{P} \triangleright [\Omega]^{\vec{G}''}, \langle [\vec{U}], [\vec{G}], [\bar{S}], [P], [S], c, \vec{G}''_{|k} \rangle_k, \mathcal{A}'' \\ \mapsto_{\omega_\alpha^e} & [\Omega]^{\vec{G}'}, \langle [\vec{U}], \text{NF}(\text{Inst}(\theta B))_{|k}, [\vec{G}], [\bar{S}], [P], \vec{G}'_{|k} \rangle_k, \mathcal{A}' \end{aligned}$$

$$\begin{aligned} \text{where } & \vec{G}'' = \text{Goals}(\Omega'') \text{ and } \vec{G}' = \text{Goals}(\Omega') \\ & \mathcal{A}'' = \bigsqcup_{k \in \text{Locs}(\text{Goals}(\Omega'')) - \text{Locs}(\Omega'')} \langle \text{Goals}(\Omega'')_{|k} \rangle_k \\ & \mathcal{A}' = \bigsqcup_{k \in \text{Locs}(\text{Goals}(\Omega')) - \text{Locs}(\Omega')} \langle \text{Goals}(\Omega')_{|k} \rangle_k \end{aligned}$$

more precisely, we want to show that this ω_α^e derivation step involves the application of the 0-neighbor restricted rule instance:

$$[k]P \setminus [k]S, [k]c \iff G \mid \theta B_k, \bigsqcup_{j \in \mathcal{I}_m} \theta B_j, \bigsqcup_{l \in \mathcal{I}_e} \theta B_l \quad (A)$$

We know that this rule instance (A) with rule heads $[k]P$ and $[k]S$ and rule guard G such that $\models G$ exists in \mathcal{P} because it has been assumed to apply in $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_0^e} \Omega'$. We now need to show that the rule body $\theta B_k, \bigsqcup_{j \in \mathcal{I}_m} \theta B_j, \bigsqcup_{l \in \mathcal{I}_e} \theta B_l$ corresponds to new constraints that appear in $[\Omega']$ but not $[\Omega'']$. Specifically, we need to show the following:

- That $\text{NF}(\text{Inst}(\theta B)) = \theta B_k, \bigsqcup_{j \in \mathcal{I}_m} \theta B_j, \bigsqcup_{l \in \mathcal{I}_e} \theta B_l$ such that $\text{NF}(\text{Inst}(\theta B))_{|k} = \theta B_k$, for each $j \in \mathcal{I}_m$ $\text{NF}(\text{Inst}(\theta B))_{|k_j} = \theta B_j$ and for each $l \in \mathcal{I}_e$ $\text{NF}(\text{Inst}(\theta B))_{|k_l} = \theta B_l$
- For each $j \in \mathcal{I}_m$ such that $\langle \bar{S}_j \rangle_{k_j} \in [\Omega'']^{\vec{G}''}$ for some \bar{S}_j , we have $\langle \bar{S}_j, \theta B_j \rangle_{k_j} \in [\Omega']^{\vec{G}'}$
- For each $l \in \mathcal{I}_e$ such that $\langle \bar{S}_l \rangle_{k_l} \notin \mathcal{A}''$ for any \bar{S}_l , we have $\langle \theta B_l \rangle_{k_l} \in \mathcal{A}'$

By definition of the translation function $[-]$, each translated location component k_j of $[\Omega'']^{\vec{G}''}$, contains $\vec{G}''_{|k_j}$. By comparison of new constraints in Ω' but not in Ω'' , we have for each $j \in \mathcal{I}_m$ $\text{NF}(\text{Inst}(\theta B))_{|k_j} = \theta B_j$. Hence the difference of $[\Omega']^{\vec{G}'}$ and $[\Omega'']^{\vec{G}''}$ is captured by $\bigsqcup_{j \in \mathcal{I}_m} \text{NF}(\text{Inst}(\theta B))_{|j}$. For each $l \in \mathcal{I}_e$ such that $k_l \notin \text{Locs}(\Omega'')$, k_l will appear in $\mathcal{A}' = \bigsqcup_{k \in \text{Locs}(\text{Goals}(\Omega')) - \text{Locs}(\Omega')} \langle \text{Goals}(\Omega')_{|k} \rangle_k$, i.e., $\langle \bar{S}_l \rangle_{k_l}$ for some \bar{S}_l . This is because the only change in goals between Ω'' to Ω' is the introduction of $\text{NF}(\text{Inst}(\theta B))$ and all locations not found in Ω'' will appear in \mathcal{A}' , by its definition. Furthermore, $\bar{S}_l = \text{NF}(\text{Inst}(\theta B))_{|k_l}$. Hence we have shown that $\text{NF}(\text{Inst}(\theta B))$ captures all constraints in $[\Omega']$ that does not appear in $[\Omega'']$, and furthermore with the constraints distributed to the corresponding locations. Since constraints that appear in $[\Omega'']$ but not in $[\Omega']$ is exactly ' S, c ' in location k , these corresponds to the effects of applying rule instance (A) via the ω_α^e derivation step, in other words: $\mathcal{P} \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$.

- (Prop) Step: The prove of this step is similar to (Simp), with the exception that active constraint $c\#d$ is kept in the store in successor state $[\Omega']$. This is consistent to the application of a rule instance where c matches with a propagated rule head, hence it does not invalidate the proof. Ω' does still contain the active goal constraint $c\#d : i$, but this is not visible in $[\Omega']$ as defined by the $[-]$ translation of goals. Lastly, history of k in Ω' is extended by $(d, \text{Ids}(P, S))$, but histories are simply discarded in $[\Omega']$. Therefore we still have $\mathcal{P} \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$.
- (Next) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{U}; (c\#d : i, \vec{G}); \bar{S}; \bar{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U}; (c\#d : (i+1), \vec{G}); \bar{S}; \bar{H} \rangle_k$$

Since we have $[c\#d : i, \vec{G}] = [c\#d : (i+1), \vec{G}]$, we have $[\Omega''] = [\Omega']$ and thus $\mathcal{P} \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$.

- (Drop) Step: Successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_\alpha^e} \Omega'$ is of the form:

$$\mathcal{P} \triangleright \Omega, \langle \vec{U} ; (c\#d : i, \vec{G}) ; \vec{S} ; \vec{H} \rangle_k \mapsto_{\omega_0^e} \Omega, \langle \vec{U} ; \vec{G} ; \vec{S} ; \vec{H} \rangle_k$$

Since we have $[c\#d : i, \vec{G}] = [\vec{G}]$, we have $[\Omega''] = [\Omega']$ and thus $\mathcal{P} \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$ of zero ω_α^e derivation step. Therefore $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$.

Therefore, we have shown $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$ for all forms of derivation step $\mathcal{P} \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^* [\Omega']$ and thus completing the inductive proof of $\mathcal{P} \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$. \square

Lemma 7 (Always Active Rule Head Instances) *Given a 0-neighbor restricted CHR^e program \mathcal{P} and an initial state Ω , for any reachable state Ω' such that $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, all rule instances in Ω' must be active.*

Proof: We proof this by induction on derivation steps between initial state Ω and reachable state Ω' . For the base case, we consider $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ of zero derivation step, hence $\Omega' = \Omega$. Since Ω' is an initial state, all stores \vec{S} in Ω' are empty, therefore we have no rule instances in Ω' and the property holds by default. For inductive case, we assume that we have some derivation of m steps to some reachable state Ω'' (i.e, $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega''$) and that all rule instances in Ω'' are active. We will now show that for any successor derivation step $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_0^e}^* \Omega'$, reachable state Ω' is such that all rule instances are active. We show this by considering all possible forms of this successor derivation:

- (Flush) Step: This step modifies location k in Ω'' by moving its buffers into the goals. Constraint store of k remains the same, hence rule instances in Ω' are similar to those in Ω'' . Therefore all rule instances in Ω' are active.
- (Loc 1) Step: This step modifies two locations k and k' in Ω'' by moving leading goal of k , $[k']c$ to goals of k' . Constraint stores of k and k' remain the same, hence rule instances in Ω' are similar to those in Ω'' . Therefore all rule instances in Ω' are active.
- (Loc 2) Step: This step modifies location k in Ω'' by changing leading goal from $[k]c$ to c . Constraint store of k remains the same, hence rule instances in Ω' are similar to those in Ω'' . Therefore all rule instances in Ω' are active.
- (Loc 3) Step: This step modifies location k in Ω'' by removing its leading goal $[k']c$, and creating a new location k' with c in its buffer and everything else empty. Constraint store of k remains the same and constraint store of k' is empty, hence rule instances in Ω' are similar to those in Ω'' . Therefore all rule instances in Ω' are active.
- (Act) Step: This step modifies location k in Ω'' by changing leading goal $p(\vec{t})$ into $p(\vec{t})\#d : i$ and also adding $p(\vec{t})\#d$ into the store \vec{S} . This possibly adds new rule instances R in Ω' which has $p(\vec{t})\#d$ as a rule head. Since $p(\vec{t})\#d : i$ is in the goals, all these newly added rule instances are active. Therefore all rule instances in Ω' are active.
- (Simp) Step: This step modifies location k in Ω'' by applying an active rule instance R which has the leading goal $c\#d : i$ as one of its rule heads. Constraints matching the simplified heads S are removed from the store: the removes rule instances but does not affect activeness of other rule instances that remains. Rule body of R is added to goals, hence does not remove or add new rule instances. Leading goal $c\#d : i$ is removed from the goals, but since constraint $c\#d$ in the constraint store is simplified as well, all active rule instances that depend on $c\#d : i$ to be active will not be found in Ω' . Therefore all rule instances that remain in Ω' are still active.

- (Prop) Step: This step is similar to the (Simp) step. Except that leading goal $c\#d : i$ is matched to a propagated rule head of R . Since constraint $c\#d$ is not removed from the store, hence there may exist other rule instances that contain $c\#d$ and depend on its corresponding goal $c\#d : i$ to be active. But leading goal $c\#d : i$ is retained in Ω' , therefore these rule instances remain active. This step also extends the history with the applied rule instances rule head ids. While this renders R no longer a rule instance in Ω' , no new rule instances are added. Therefore all rule instances that remain in Ω' are still active.
- (Next) Step: This step modifies a location k in Ω by changing leading goal $c\#d : i$ to $c\#d : (i + 1)$. The premise of this rule explicitly asserts that no (Simp) and (Prop) steps must apply for this goal $c\#d : i$. This effectively enforces that this step only have applied because Ω'' contains no rule instances that contain $c\#d : i$ hence $c\#d : i$ can be removed without making any rule instances inactive. Adding $c\#d : (i + 1)$ possibly introduces new rule instances, but since it is part of the goals, all these newly added rule instances are active. Therefore all rule instances that remain in Ω' are still active.
- (Drop) Step: This step modifies a location k in Ω by removing its leading goal $c\#d : i$. The premise of this rule restricts it to only cases where i is not a occurrence index that appear in the program \mathcal{P} , hence there can never be any rule instances that contain $c\#d : i$ and it can be safely removed without rendering any rule instances inactive. Therefore all rule instances that remain in Ω' are still active.

Therefore, we have completed the inductive step proof for $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$. Hence all rule instances in any reachable state Ω' are active. \square

Theorem 8 (Exhaustiveness of Rule Application in ω_0^e) *Given a 0-neighbor restricted CHR^e program \mathcal{P} and reachable states Ω and Ω' , if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ and Ω' is terminal, then there exists no rule instance $R \in \mathcal{P}$ such that R is applicable.*

Proof: We proof this by negation, showing that if we assume otherwise, we contradict Lemma 7. Assume that we have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ and Ω' is terminal, but there exists some rule instance $R \in \mathcal{P}$ such that R is applicable. Suppose this rule instance is applicable in some location $k \in \text{Locs}(\Omega')$. Since Lemma 7 states that any rule instance R must be active, hence there must exists some $c\#d : i$ in the goals of k which matches the i^{th} rule head occurrence of rule instance R . However Ω' is terminal, in other words, its goals and buffer are suppose to be empty. Hence we have a contradiction. Therefore, it must be the case that in Ω' there exists no rule instance $R \in \mathcal{P}$ such that R is applicable. \square

Lemma 9 (Monotonicity of ω_0^e semantics) *Given a 0-neighbor restricted program \mathcal{P} , and reachable states $\Omega, \Omega', \Omega''$,*

1. *if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e} \Omega', \Omega''$*
2. *if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^* \Omega', \Omega''$*
3. *if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^{\parallel} \Omega', \Omega''$*
4. *if $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$, then $\mathcal{P} \triangleright \Omega, \Omega'' \mapsto_{\omega_0^e}^{\parallel*} \Omega', \Omega''$*

Proof:

1. We prove by structural induction over all forms of $\mapsto_{\omega_0^e}$ derivation steps. For the (Loc 3) step, we assume implicit α -renaming of location k' during composition of a larger context Ω'' . All other cases are straight-forward, each simply to show that unmodified fragment of the ensemble state Ω can be extended without affecting the derivation step.

2. We prove by induction on $\mapsto_{\omega_0^e}$ derivation steps and also relying on the proof of (1).
3. We prove by structural induction over the (Single) and (Concurrent) derivation steps of $\mapsto_{\omega_0^e}^{\parallel}$ and also relying on the proof of (1).
4. We prove by induction on $\mapsto_{\omega_0^e}^{\parallel}$ derivation steps and also relying on the proof of (3).

□

Theorem 10 (Serializability of Concurrent ω_0^e Derivations) *Given a 0-neighbor restricted program \mathcal{P} , and reachable states Ω, Ω' , if we have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$ then we must have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$*

Proof: We first proof this for one $\mapsto_{\omega_0^e}^{\parallel}$ derivation step, $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel} \Omega'$: By structural induction on all forms of $\mapsto_{\omega_0^e}^{\parallel}$ derivation step, we consider the base structure (Single). Hence $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel} \Omega'$ consist of a single $\mapsto_{\omega_0^e}$ derivation (i.e, $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$) and we naively have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$. For the (Concurrent) step, we have the derivation step form:

$$\mathcal{P} \triangleright (\Omega_1, \Omega_2) \mapsto_{\omega_0^e}^{\parallel} (\Omega'_1, \Omega'_2) \text{ such that } \Omega = (\Omega_1, \Omega_2) \text{ and } \Omega' = (\Omega'_1, \Omega'_2)$$

Hence, assuming that we have $\mathcal{P} \triangleright \Omega_1 \mapsto_{\omega_0^e}^{\parallel} \Omega'_1$ and $\mathcal{P} \triangleright \Omega_2 \mapsto_{\omega_0^e}^{\parallel} \Omega'_2$ such that $\mathcal{P} \triangleright \Omega_1 \mapsto_{\omega_0^e} \Omega'_1$ and $\mathcal{P} \triangleright \Omega_2 \mapsto_{\omega_0^e} \Omega'_2$ by Lemma 9, we have:

$$\mathcal{P} \triangleright \Omega_1, \Omega_2 \mapsto_{\omega_0^e} \Omega'_1, \Omega_2 \mapsto_{\omega_0^e} \Omega'_1, \Omega'_2$$

Therefore we have $\mathcal{P} \triangleright (\Omega_1, \Omega_2) \mapsto_{\omega_0^e}^* (\Omega'_1, \Omega'_2)$.

We have shown that given $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel} \Omega'$, we have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, now we prove the same for transitive $\mapsto_{\omega_0^e}^{\parallel}$ derivation steps by induction on $\mapsto_{\omega_0^e}^{\parallel*}$ derivation steps. For base case, we have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$ of zero derivation step, hence $\Omega = \Omega'$ and we naively have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$. For the inductive case, we assume that we have the property for concurrent derivations of m steps, i.e, $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega''$ of m steps implies $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega''$. We consider Ω' the successor of Ω'' after one derivation step, $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_0^e}^{\parallel} \Omega'$. Since we show that for a single concurrent derivation step, we have $\mathcal{P} \triangleright \Omega'' \mapsto_{\omega_0^e}^* \Omega'$, hence we also have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.

Therefore, we have shown that given $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^{\parallel*} \Omega'$ then we must have $\mathcal{P} \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$ □

A.3 Proofs for 1-Neighbor Restricted Rule Basic Encoding Scheme

Lemma 11 (Prefix Executions of Encoding) *Given a locally quiescent 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1\text{Nb}}^{\text{basic}} \mathcal{P}_0$, if we have reachable states Ω, Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e} \Omega'$, then we have the following:*

1. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r-2} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r-1} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
2. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r-3} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r-2} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$
3. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r-4a} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r-3} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$

4. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^{r_{-4b}} \Omega'$, then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^e}^{r_{-3}} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^e}^* \Omega$

Proof: These are easily proved by reasoning about the definition of encoding rules of Figure 13:

1. By definition of encoding rule r_{-1} and r_{-2} , the former has $r_{-req}(Xs)$ as a rule body, while the latter has $r_{-req}(Xs)$ as a rule head. Since the synchronizing predicate r_{-req} appears in no where else, Application of r_{-2} encoding rule must always be preceded by application of a matching instance of r_{-1} .
2. By definition of encoding rule r_{-2} and r_{-3} , the former has $r_{-match}(Rs)$ as a rule body, while the latter has $r_{-match}(Rs)$ as a rule head. Since the synchronizing predicate r_{-match} appears in no where else, Application of r_{-3} encoding rule must always be preceded by application of a matching instance of r_{-2} .
3. By definition of encoding rule r_{-3} and r_{-4a} , the former has $r_{-commit}(Rs)$ as a rule body, while the latter has $r_{-commit}(Rs)$ as a rule head. Since the synchronizing predicate $r_{-commit}$ appears in no where else, Application of r_{-4a} encoding rule must always be preceded by application of a matching instance of r_{-3} .
4. By definition of encoding rule r_{-3} and r_{-4b} , the former has $r_{-commit}(Rs)$ as a rule body, while the latter has $r_{-commit}(Rs)$ as a rule head. Since the synchronizing predicate $r_{-commit}$ appears in no where else, Application of r_{-4b} encoding rule must always be preceded by application of a matching instance of r_{-3} .

□

Lemma 12 ($\rightsquigarrow_{\text{INb}}^{\text{basic}}$ Preserves Local Quiescence) *Given a locally quiescent 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}}^{\text{basic}} \mathcal{P}_0$, then \mathcal{P}_0 is also locally quiescent.*

Proof: By definition of the $\rightsquigarrow_{\text{INb}}^{\text{basic}}$ translation operation, encoding rules r_{-1} , r_{-2} , r_{-3} and r_{-4b} are guaranteed to be locally quiescent. This is because their rule bodies only include constraints to be delivered to its partner location, hence the rules are conservatively locally quiescent. Local quiescence of encoding rule r_{-4a} however, depends on the local quiescence of the original rule in \mathcal{P}_1 . Since we assume that \mathcal{P}_1 is locally quiescent, then r_{-4a} must be locally quiescent. □

Lemma 13 (1-Neighbor Commit-Free Reachability) *Given a 1-neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}}^{\text{basic}} \mathcal{P}_0$ and states Ω reachable by \mathcal{P}_0 , if Ω is not commit-free, then there exists some commit-free state Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.*

Proof: We prove this by considering all possible forms of derivations that lead to non-commit free states. We consider all forms of 0-neighbor restricted encodings that contains the r_{-3} , r_{-4a} and r_{-4b} encoding rules that rewrites (add or delete) $r_{-commit}$ synchronizing constraints. By definition of the encodings, the only rule that produces $r_{-commit}$ constraints is r_{-3} , hence derivations that reach a non commit-free state Ω must consist of several applications of the r_{-3} rule. Suppose that one of rule application is contained in the

following derivation:

$$\mathcal{P}_0 \triangleright \Omega, \left(\langle \vec{U}_x ; \vec{G}_x ; \left(\bar{S}_x, \boxed{P_x, S_x, r_match(Rs)\#d_1} \right) ; \bar{H}_x \rangle_x \right) \quad (1)$$

$$\mapsto_{\omega_0^e} \Omega, \left(\langle \vec{U}_x ; \left([y]r_commit(Rs), \vec{G}_x \right) ; (\bar{S}_x, P_x) ; \bar{H}_x \rangle_x \right) \quad (2)$$

$$\mapsto_{\omega_0^e} \Omega, \left(\langle \vec{U}_x ; \vec{G}_x ; (\bar{S}_x, P_x) ; \bar{H}_x \rangle_x \right) \quad (3)$$

$$\mapsto_{\omega_0^e}^* \Omega, \left(\langle \vec{U}'_y ; \left(r_commit(Rs)\#d_2 : j, \vec{G}'_y \right) ; \left(\bar{S}'_y, \boxed{r_commit(Rs)\#d_2} \right) ; \bar{H}'_y \rangle_y \right) \quad (4)$$

Suppose that x and y the are primary and neighbor locations of the r_3 rule instance. By the definition of ω_0^e derivation steps, we are guaranteed that from an initial state such that a rule instance of r_3 is active in location x (highlighted in the derivation step (1) in above derivation) is applied (step (2) via (Simp) or (Prop) ω_0^e derivation step) and delivered to location y (step (3) via (Loc 1) ω_0^e derivation step), there exists some successor state ((4)) such that $r_commit(Rs)\#d_2 : j$ is an active constraint of y . This is guaranteed because we assume that CHR^e programs are locally terminating, hence local execution eventually reach quiescences (goals are emptied) and buffer of y containing $r_commit(Rs)$ is moved into the goals. Eventually, $r_commit(Rs)\#d_2 : j$ will be the active constraint. From (4), there are two possibilities: If matching obligation of y is still valid in (4) (i.e, $P_y, S_y \subseteq \bar{S}'_y$), then we have the following derivation:

$$\mathcal{P}_0 \triangleright \left(\langle \vec{U}'_y ; \left(r_commit(Rs)\#d_2 : j, \vec{G}'_y \right) ; \left(\bar{S}'_y, \boxed{P_y, S_y, r_commit(Rs)\#d_2} \right) ; \bar{H}'_y \rangle_y \right) \quad (4)$$

$$\mapsto_{\omega_0^e} \left(\langle \vec{U}'_y ; \left(B, \vec{G}'_y \right) ; (\bar{S}'_y, P_y) ; \bar{H}'_y \rangle_y \right) \quad (5)$$

This results to a state which is commit free. If matching obligation of y is no longer valid in (4), r_4b applies, consuming the commit synchronizing constraint $r_commit(Rs)$. Since application of r_4b universal in that it always applies if r_4a doesn't, hence any non-commit free state can reach a commit free state by a series of applications of either r_4a or r_4b . \square

Lemma 14 (Basic Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that:*

1. *For some basic encoding rule instances r_3 and r_4a , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_4a} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 and \mathcal{A}'_3 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_4a} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_4$*
2. *For some basic encoding rule instances r_3 and r_4b , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_4b} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 and \mathcal{A}'_3 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_4b} \mathcal{A}'_3 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_4$*

A PROOFS

Proof: Let r be any 1-neighbor restricted rule instance of the general form:

$$r : [X]P_x, [X]P'_x, [y]P_y \setminus [X]S_x, [y]S_y \iff G \mid \exists \bar{z}. [X]D_x, [Y]D_y, \uplus_{j \in \mathcal{I}_m} [k_j]D_j, \uplus_{l \in \mathcal{I}_e} [k_l]D_l \in \mathcal{P}_1$$

where k_j for $j \in \mathcal{I}_m$, $k_j \in \text{FV}(P_x, P'_x, P_y, S_x, S_y)$ are *non-matched forwarding locations*.
 k_l for $l \in \mathcal{I}_e$, $k_l \in \bar{z}$ are *existential forwarding locations*.

for some location X and Y . Hence this instance rule is translated into the following r_3 , r_4a and r_4b encoding rule instances:

$$\begin{aligned} r_3 &: [X]P_x \setminus [X]P'_x, [X]S_x, [X]r_match(Rs) \iff [Y]r_commit(Rs) \\ r_4a &: [Y]P_y \setminus [Y]S_y, [Y]r_commit(Rs) \iff \exists \bar{z} [X]P'_x, [X]D_x, [Y]D_y, \uplus_{j \in \mathcal{I}_m} [k_j]D_j, \uplus_{l \in \mathcal{I}_e} [k_l]D_l \\ r_4b &: [Y]r_commit(Rs) \iff [X]P'_x, [X]S_x \end{aligned}$$

We now will show that derivation steps in the derivation $\mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ can be permuted such that an application of r_3 instance is immediately followed by an application r_4a or r_4b . We first consider any matching pairs of r_3 and r_4a . For this case the derivation $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_4a} \mathcal{A}_4$ is of the form:

$$\begin{aligned} \mathcal{P} \triangleright & \mathcal{A}'', \uplus_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j}, \left(\langle \bar{S}_x, \boxed{P_x, P'_x, S_x, r_match(Rs)} \rangle_X, \langle \bar{S}_y \rangle_Y \right) && \text{Apply } r_3 \text{ instance} \\ \mapsto_{\omega_0^e}^{r_3} & \boxed{\mathcal{A}''}, \uplus_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j}, \left(\langle \bar{S}_x, P_x \rangle_X, \langle \bar{S}_y, r_commit(Rs) \rangle_Y \right) && \text{Unspecified interleaving derivations} \\ \mapsto_{\omega_0^e}^* & \mathcal{A}''', \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j \rangle_{k_j}, \left(\langle \bar{S}'_x, P_x \rangle_X, \langle \bar{S}'_y, \boxed{P_y, S_y, r_commit(Rs)} \rangle_Y \right) && \text{Apply } r_4a \text{ instance} \\ \mapsto_{\omega_0^e}^* & \mathcal{A}''', \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j, D_j \rangle_{k_j}, \uplus_{l \in \mathcal{I}_e} \langle D_l \rangle_{k_l}, \left(\langle \bar{S}'_x, P_x, P'_x, D_x \rangle_X, \langle \bar{S}'_y, P_y, D_y \rangle_Y \right) \end{aligned}$$

We highlight the boxes the portions of the state which is responsible for the subsequent derivation step. By Lemma 9 that states monotonicity of the ω_α^e semantics, we can permute derivation steps that rewrite over non-overlapping portions of the state. Specifically we want to push all derivation steps between r_3 and r_4a either before or after their subsequent derivation steps. Derivation steps that are responsible for introducing Y 's matching obligation for r_4a will be pushed before r_3 while all others will be pushed after. This is illustrated by the following:

$$\begin{aligned} \mathcal{P} \triangleright & \boxed{\mathcal{A}''}, \uplus_{j \in \mathcal{I}_m} \langle \bar{S}_j \rangle_{k_j}, \left(\langle \bar{S}_x, \boxed{P_x}, P'_x, S_x, r_match(Rs) \rangle_X, \langle \bar{S}_y \rangle_Y \right) && \text{Interleaving derivations resulting to } Y\text{'s matching obligation - (I1)} \\ \mapsto_{\omega_\alpha^e}^* & \mathcal{A}''', \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j \rangle_{k_j}, \left(\langle \bar{S}'_x, \boxed{P_x, P'_x, S_x, r_match(Rs)} \rangle_X, \langle \bar{S}'_y, P_y, S_y \rangle_Y \right) && \text{Apply } r_3 \text{ instance} \\ \mapsto_{\omega_0^e}^{r_3} & \boxed{\mathcal{A}'''}, \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j \rangle_{k_j}, \left(\langle \bar{S}'_x, P_x \rangle_X, \langle \bar{S}'_y, \boxed{P_y, S_y, r_commit(Rs)} \rangle_Y \right) && \text{Apply } r_4a \text{ instance} \\ \mapsto_{\omega_0^e}^{r_4a} & \boxed{\mathcal{A}'''}, \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j, D_j \rangle_{k_j}, \uplus_{l \in \mathcal{I}_e} \langle D_l \rangle_{k_l}, \left(\langle \bar{S}'_x, \boxed{P_x, P'_x, D_x} \rangle_X, \langle \bar{S}'_y, P_y, D_y \rangle_Y \right) && \text{Other interleaving Derivations - (I2)} \\ \mapsto_{\omega_0^e}^* & \mathcal{A}''', \uplus_{j \in \mathcal{I}_m} \langle \bar{S}'_j, D_j \rangle_{k_j}, \uplus_{l \in \mathcal{I}_e} \langle D_l \rangle_{k_l}, \left(\langle \bar{S}'_x, P_x, P'_x, D_x \rangle_X, \langle \bar{S}'_y, P_y, D_y \rangle_Y \right) \end{aligned}$$

Consider derivation steps (I1): We highlight *conservatively* the largest portion of the state which possible contributes to the derivation of Y 's matching obligations (i.e, P_y and S_y) that were originally interleaving

between r_{-3} and r_{-4a} . Note that this includes fragments of location X 's store including the P_x matching obligations of encoding rule r_{-3} (underline in the derivation, P_x , P'_x and S_x). Matching obligations P'_x and S_x are excluded because the original derivations are scheduled after r_{-3} where P'_x and S_x do not exist. P_x is included but by definition they are persistent and any derivations that depend on any constraint of P_x are guaranteed not to be simplified. Hence we can safely permute these rules upwards. For derivation steps in $(I2)$, since they involve portions of the state that are not overlapping with r_{-4a} , we can safely push them downwards. Hence, we have proven the lemma for the r_{-4a} case.

We now consider r_{-3} encoding rule instances that are paired with matching r_{-4b} applications. Specifically, they are of the form:

$$\begin{array}{lcl}
 \mathcal{P} \triangleright & \mathcal{A}'' , \left(\langle \bar{S}_x, \boxed{P_x, P'_x, S_x, r_match(Rs)} \rangle_X, \langle \bar{S}_y \rangle_Y \right) & \text{Apply } r_{-3} \text{ instance} \\
 \mapsto_{\omega_0^e}^{r_{-3}} & \boxed{\mathcal{A}''} , \left(\langle \bar{S}_x, P_x \rangle_X, \langle \bar{S}_y, r_commit(Rs) \rangle_Y \right) & \text{Unspecified interleaving derivations} \\
 \mapsto_{\omega_0^e}^* & \mathcal{A}''' , \left(\langle \bar{S}'_x, P_x \rangle_X, \langle \bar{S}'_y, \boxed{r_commit(Rs)} \rangle_Y \right) & \text{Apply } r_{-4b} \text{ instance} \\
 \mapsto_{\omega_0^e}^* & \mathcal{A}'''' , \left(\langle \bar{S}'_x, P_x, P'_x, S_x \rangle_X, \langle \bar{S}'_y \rangle_Y \right) &
 \end{array}$$

Since matching obligation of r_{-4b} only consist of the constraint $r_commit(Rs)$ we can straightforwardly push all interleaving derivations downwards:

$$\begin{array}{lcl}
 \mathcal{P} \triangleright & \mathcal{A}'' , \left(\langle \bar{S}_x, \boxed{P_x, P'_x, S_x, r_match(Rs)} \rangle_X, \langle \bar{S}_y \rangle_Y \right) & \text{Apply } r_{-3} \text{ instance} \\
 \mapsto_{\omega_0^e}^* & \mathcal{A}'' , \left(\langle \bar{S}_x, P_x \rangle_X, \langle \bar{S}_y, \boxed{r_commit(Rs)} \rangle_Y \right) & \text{Apply } r_{-4b} \text{ instance} \\
 \mapsto_{\omega_0^e}^* & \boxed{\mathcal{A}''} , \left(\langle \bar{S}_x, P_x, P'_x, S_x \rangle_X, \langle \bar{S}_y \rangle_Y \right) & \text{Interleaving Derivations} \\
 \mapsto_{\omega_0^e}^* & \mathcal{A}'''' , \left(\langle \bar{S}'_x, P_x, P'_x, S_x \rangle_X, \langle \bar{S}'_y \rangle_Y \right) &
 \end{array}$$

Hence we have proven this lemma for all cases. \square

Theorem 15 (Soundness of Basic Encoding) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1Nb}^{basic} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_1 \triangleright DropSyncs([\Omega]; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* DropSyncs([\Omega']; \mathcal{P}_1)$.*

Proof: Given that we have $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, by Theorem 6, we have $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_0^e}^* [\Omega']$, labeling this as \mathcal{D} . By Lemma 14, we can safely permute \mathcal{D} such that all applications of r_{-3} are followed immediately by r_{-4a} or r_{-4b} . We shall consider such a permutation of \mathcal{D} . We will now prove this theorem by induction on ω_α^e derivation steps. For base case we consider $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$ of zero derivation step, hence $[\Omega] = [\Omega']$ and we naively have $\mathcal{P}_1 \triangleright DropSyncs([\Omega]; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* DropSyncs([\Omega']; \mathcal{P}_1)$. for $[\Omega']$ being commit free and hence $[\Omega'] = [\Omega'']$. For the inductive case, we assume we have this property for derivations $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega''']$ of m derivation steps. We now consider the successor derivation step $[\Omega''']$ such that $\mathcal{P}_0 \triangleright [\Omega'''] \mapsto_{\omega_\alpha^e}^* [\Omega'']$. We consider all possible forms rules r applied:

- r is a non-encoding rule instance. Therefore we have $r \in \mathcal{P}_1$ as well. Furthermore, since synchronizing

constraints are reserved for encoding rules, and hence r does not have rule heads or rule bodies that are synchronizing constraints. From this, we have $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega''']; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^r \text{DropSyncs}([\Omega'']; \mathcal{P}_1)$

- r is an instance of the encoding rule r_1 . By definition of r_1 encoding rules, this derivation simply adds synchronizing constraint $r_req(Xs)$ to some location x in $[\Omega''']$. Therefore, $\text{DropSyncs}([\Omega''']; \mathcal{P}_1) = \text{DropSyncs}([\Omega'']; \mathcal{P}_1)$ and we naively have $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega''']; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega'']; \mathcal{P}_1)$ of zero derivation step.
- r is an instance of the encoding rule r_2 . By definition of r_1 encoding rules, this derivation simply removes synchronizing constraint $r_req(Xs)$ from some location x add synchronizing constraint $r_match(Rs)$ to some location y in $[\Omega''']$. Therefore, $\text{DropSyncs}([\Omega''']; \mathcal{P}_1) = \text{DropSyncs}([\Omega'']; \mathcal{P}_1)$ and we naively have $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega''']; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega'']; \mathcal{P}_1)$ of zero derivation step.
- r is an instance of the encoding rule r_3 . By definition of r_3 encoding rules, this derivation adds synchronizing constraint $r_commit(Rs)$, hence making Ω'' non commit free. By Lemma 14, This derivation is followed by a derivation $\mathcal{P}_0 \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^{r'} [\Omega''']$ such that r' is either r_4a or r_4b encoding rule instance. Furthermore, these encoding rules are of the same matching instance and of the form:

$$\begin{aligned} r_3 &: [X]P_x \setminus [X]P'_x, [X]S_x, [X]r_match(Rs) \iff [Y]r_commit(Rs) \\ r_4a &: [Y]P_y \setminus [Y]S_y, [Y]r_commit(Rs) \iff \exists \bar{z}. [X]P'_x, D \\ r_4b &: [Y]r_commit(Rs) \iff [X]P'_x, [X]S_x \end{aligned}$$

and these encoding rules corresponding to some 1-neighbor restricted rule instance $r_1 \in \mathcal{P}_1$. Let this rule instance be $r_1 : [X]P_x, [X]P'_x, [Y]P_y \setminus [X]S_x, [Y]S_y \iff G \mid \exists \bar{z}. D$. We now consider the case that $r' = r_4a$: This instance of r_4a removes S_y and $r_commit(Rs)$ from location Y and adds P'_x to location X and D to various locations of the state, while its preceding r_3 derivation step removes S_x and P'_x from location X . Hence we have that Ω'''' is commit free, and the net effect of these two derivation is the application of the r_1 1-neighbor restricted rule instance. Therefore $\mathcal{P}_1 \triangleright \text{DropSyncs}(\Omega'''; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^{r_1} \text{DropSyncs}(\Omega''''; \mathcal{P}_1)$.

We now consider the case that $r' = r_4b$: This instance of r_4b removes $r_commit(Rs)$ from location Y , and adds P'_x and S_x to location X , while its preceding r_3 derivation step removes P'_x and S_x from location X . The net effect of these two derivation is that we return to the state before r_3 was applied. Therefore, $\text{DropSyncs}([\Omega''']; \mathcal{P}_1) = \text{DropSyncs}([\Omega'''']; \mathcal{P}_1)$ and we naively have $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega''']; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega'''']; \mathcal{P}_1)$ of zero derivation step.

Therefore we have proven the soundness of the basic encoding scheme. \square

A.4 Proofs for 1-Neighbor Restricted Rule Optimized Encoding Scheme

Lemma 16 (Neighbor Persistent Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{1NB} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that for some matching neighbor persistent encoding rule instances r_2 and r_3 , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_3} \mathcal{A}_4$*

Proof: By definition, encoding rule r_2 are of the form: $[Y]P_y, [Y]r_req(Xs) \iff G_y \mid [X]r_match(Rs)$ or $[Y]P_y \setminus [Y]r_req(Xs) \iff G_y \mid [X]r_match(Rs)$ Given that P_y are all persistent constraints, and that $r_req(Xs)$ strictly matches to this encoding rule only, by Lemma 9 and since P_y will never be deleted and $r_req(Xs)$ is exclusive for the application of r_2 and furthermore its rule body $r_match(Rs)$ is only ever required for the application of the matching instance of r_3 encoding rule, we can safely push a derivation step of r_2 downwards after any derivation step that comes after it, except for its matching r_3 encoding

rule. \square

Lemma 17 (Primary Persistent Encoding Rule Serializability) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and \mathcal{A}_4 , given that for some matching primary persistent encoding rule instances r_1 and r_2 , we have $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^{r_1} \mathcal{A}_2 \mapsto_{\omega_\alpha^e}^* \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_4$ then there exists some \mathcal{A}'_2 such that $\mathcal{P}_0 \triangleright \mathcal{A}_1 \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_2 \mapsto_{\omega_\alpha^e}^{r_1} \mathcal{A}_3 \mapsto_{\omega_\alpha^e}^{r_2} \mathcal{A}_4$*

Proof: By definition, encoding rule r_1 are of the form: $[X]P_x \Longrightarrow G_x \mid [X]r_req(Xs)$ Given that P_x are all persistent constraints, by Lemma 9 and since P_x will never be deleted and furthermore its rule body $r_req(Xs)$ is only ever required for the application of the matching instance of r_2 encoding rule, we can safely push a derivation step of r_1 downwards after any derivation step that comes after it, except for its matching r_2 encoding rule. \square

Theorem 18 (Soundness of Optimized Encoding) *Given a 1-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_1 and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_1 \rightsquigarrow_{\text{INb}} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_1 \triangleright \text{DropSyncs}([\Omega]; \mathcal{P}_1) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega']; \mathcal{P}_1)$.*

Proof: The proof of this is similar to the proof of Theorem 15. Specifically, we first consider the ω_α^e derivations of \mathcal{P}_0 , i.e, $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$, Next we use Lemma 14, 16 and 17 to consider only $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$ derivations such that

- For basic encoding rules, all applied r_3 rule instances are immediately followed by a matching application of r_4a or r_4b (Lemma 14).
- For neighbor persistent encoding rules, all applied r_2 rule instances are immediately followed by a matching application of r_3 (Lemma 16).
- For primary persistent encoding rules, all applied r_1 rule instances are immediately followed by a matching application of r_2 (Lemma 17).

Given a derivation $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$ with the above property, we prove by induction on the derivation steps: For derivation steps corresponding to the application of basic encoding rules and 0-neighbor restricted rules we defer to Theorem 15 for details, here we consider the inductive step derivation $\mathcal{P}_0 \triangleright [\Omega''] \mapsto_{\omega_\alpha^e}^r [\Omega']$ for neighbor and primary persistent encoding rules:

- r is an instance of the neighbor persistent encoding rule r_2 . By Lemma 16, this derivation is followed by a derivation $\mathcal{P}_0 \triangleright [\Omega''] \mapsto_{\omega_0^e}^{r_3} [\Omega''']$. These encoding rules are of the form:

$$\begin{aligned} r_2 &: [Y]P_y \setminus r_req(Xs) \iff G_y \mid [X]r_match(Rs) \\ r_3 &: [X]P_x \setminus [X]S_x, [X]r_match(Rs) \iff B \end{aligned}$$

and these encoding rules corresponds to some 1-neighbor restricted rule instance $r_1 \in \mathcal{P}_1$. Let this rule instance be $r_1 : [X]P_x, [Y]P_y \setminus [X]S_x \iff G_x, G_y \mid B$ and the net effects of the two derivation step is the application of the r_1 rule instance. Therefore $\mathcal{P}_1 \triangleright \text{DropSyncs}(\Omega''; \mathcal{P}_1) \mapsto_{\omega_0^e}^{r_1} \text{DropSyncs}(\Omega'; \mathcal{P}_1)$.

- r is an instance of the primary persistent encoding rule r_1 . By Lemma 17, this derivation is followed by a derivation $\mathcal{P}_0 \triangleright [\Omega''] \mapsto_{\omega_0^e}^{r_3} [\Omega''']$. These encoding rules are of the form:

$$\begin{aligned} r_1 &: [X]P_x \Longrightarrow G_x \mid [Y]r_req(Xs) \\ r_2 &: [Y]r_req(Xs), [Y]P_y \setminus [Y]S_y \iff G_y \mid B \end{aligned}$$

and these encoding rules corresponds to some 1-neighbor restricted rule instance $r_1 \in \mathcal{P}_1$. Let this rule instance be $r_1 : [X]P_x, [Y]P_y \setminus [Y]S_y \iff G_x, G_y \mid B$ and the net effects of the two derivation step is the application of the r_1 rule instance. Therefore $\mathcal{P}_1 \triangleright \text{DropSyncs}(\Omega''; \mathcal{P}_1) \mapsto_{\omega_0^\epsilon}^{r_1} \text{DropSyncs}(\Omega'; \mathcal{P}_1)$.

Hence we have proven the soundness of the optimized encoding scheme. \square

A.5 Proofs for n -Neighbor Restricted Rule Encoding Scheme

Lemma 19 ($\rightsquigarrow_{n\text{Nb}}$ Preserves Local Quiescence) *Given a locally quiescent n -neighbor restricted program \mathcal{P}_n and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{n\text{Nb}} \mathcal{P}_0$, then \mathcal{P}_0 is also locally quiescent.*

Proof: By the definition form of the encoding rules of $\rightsquigarrow_{n\text{Nb}}$, encoding rules r_{-1} , r_{-2i} , r_{-3} , r_{-4ai} , r_{-4bi} and r_{-5bi} are guaranteed to be locally quiescent. This is because their rule bodies only include constraints to be delivered to its partner location, hence the rules are conservatively locally quiescent. Local quiescence of encoding rule r_{-5a} and r_{-5bx} depends on the local quiescence of the original rule r in \mathcal{P}_1 . Since we assume that \mathcal{P}_1 is locally quiescent, then r_{-5a} and r_{-5bx} must be locally quiescent. \square

Lemma 20 (Prefix Executions of n -Neighbor Encoding) *Given a locally quiescent n -neighbor restricted program \mathcal{P}_1 and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{n\text{Nb}} \mathcal{P}_0$, if we have reachable states Ω, Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon} \Omega'$, then we have the following:*

1. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-2i}(E)} \Omega'$ for r_{-2i} encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^{r_{-1}(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$
2. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-3}(E)} \Omega'$ for r_{-3} encoding rule of some consensus destination E , then there exists some reachable derivation $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^* \Omega'''$ which contains n derivation steps $r_{-2i}(E)$ for each neighbor Y_i of $i \in \mathcal{I}_n$, and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$
3. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-4ai}(E)} \Omega'$, for r_{-4ai} encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^{r_{-3}(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$
4. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-4bi}(E)} \Omega'$, for r_{-4bi} encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^{r_{-3}(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$
5. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-5a}(E)} \Omega'$ for r_{-5a} encoding rule of some consensus destination E , then there exists some reachable derivation $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^* \Omega'''$ which contains n derivation steps $r_{-2i}(E)$ for each neighbor Y_i of $i \in \mathcal{I}_n$, and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$
6. If $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-5bx}(E)} \Omega'$ of $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^\epsilon}^{r_{-5bi}(E)} \Omega'$ for r_{-5bx} or r_{-5bi} encoding rule of some consensus destination E , then there exists some reachable states Ω'' and Ω''' , such that $\mathcal{P}_0 \triangleright \Omega'' \mapsto_{\omega_0^\epsilon}^{r_{-4bj}(E)} \Omega'''$ and $\mathcal{P}_0 \triangleright \Omega''' \mapsto_{\omega_0^\epsilon}^* \Omega$

Proof: The proof for this is similar to Lemma 11's. The main observation is that each item of the lemma essentially states a specific 'input' 'output' dependency between the encoding rules. This dependency is defined by the synchronizing constraints that appear in the encoding rules. For instance, encoding rule r_{-1} adds a unique synchronizing constraint (r_{-req_i}) for each of n neighbors. Each r_{-2i} encoding rule has synchronizing constraint r_{-req_i} in its left-hand side, and given that these synchronizing constraints appear

no where else in the program \mathcal{P}_0 , we can infer (1). For the rest of the lemma, we can repeat this same argument about the definition of the encoding rules. \square

Lemma 21 (*n*-Neighbor Commit-Free Reachability) *Given a locally quiescent *n*-neighbor restricted program \mathcal{P}_n and a 0-neighbor restricted program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{\text{nNb}} \mathcal{P}_0$ and states Ω reachable by \mathcal{P}_0 , if Ω is not commit-free, then there exists some commit-free state Ω' such that $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$.*

Proof: The proof for this is similar to Lemma 13's. For the case of *n*-neighbor restricted rule encoding, the proof of this lemma can be substantiated by the guarantee that the encoding rules have the following form and property: Encoding rule *r_3* and *r_4ai* for each $i \in \mathcal{I}_n$ are the rules ($n + 1$ in total for each *n*-neighbor restricted encoding) in which any form of commit synchronizing constraint will be added to the state. Hence a state will be non-commit free after the application of any of those rules. If each of the *n*-neighbors applies *r_4ai* we have $n + 1$ matching commit synchronizing constraint of the same consensus destination E . With Lemma 19, we have a progress guarantee that each commit synchronizing constraint will eventually be delivered to a location and processed by the location. Hence, these $n + 1$ commit synchronizing constraints can be consumed by a matching application of *5_a*. If not all of the *n*-neighbors applies *r_4ai*, by definition of the *r_4bi* rule, a matching *r_4bi* rule will be applied instead, producing an abort synchronizing constraint. This instead of the abort synchronizing constraint will initiate the application of the *r_5bx* and *r_5bi* encoding rules, that will remove each of the commit synchronizing constraints of consensus destination E . The derivation dependencies of these encoding rules are guarantee by Lemma 20. Hence we have proven this Lemma. \square

Lemma 22 (*n*-Neighbor Encoding Rule Serializability) *Given an *n*-neighbor restricted and locally quiescent CHR^e program \mathcal{P}_n and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{\text{nNb}} \mathcal{P}_0$ and commit free abstract states $\mathcal{A}, \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ labeled as derivation \mathcal{I} , if \mathcal{I} contain derivations of a *n*-neighbor restricted encoding rule instance r of consensus destination E such that:*

- *We have a successful application of rule instance r of E , then there exists a valid derivation $\mathcal{I}' = \mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E$ consists the following non-interleaving derivation sequence:*
 1. *The derivation step $r_3(E)$.*
 2. *All n derivation steps of $r_4ai(E)$ each of Y_i for $i \in \mathcal{I}_n$.*
 3. *The derivation step $r_5a(E)$.*
- *We have an aborted attempt to apply rule instance r of E , then there exists a valid derivation $\mathcal{I}' = \mathcal{P}_0 \triangleright \mathcal{A} \mapsto_{\omega_\alpha^e}^* \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'$ such that $\mathcal{P}_0 \triangleright \mathcal{A}_E \mapsto_{\omega_\alpha^e}^* \mathcal{A}'_E$ consists a non-interleaving derivation sequences of derivation step $r_3(E)$ followed by all n derivation steps of $r_4ai(E)$ or r_4bi each of Y_i for $i \in \mathcal{I}_n$, followed by*
 1. *The derivation step $r_3(E)$.*
 2. *All m derivation steps of either $r_4ai(E)$ each of Y_i for $i \in \mathcal{I}_n$ where $m \leq n$.*
 3. *All $n - m$ derivation steps of either $r_4bi(E)$ each of Y_i for $i \in \mathcal{I}_n$.*
 4. *The derivation step $r_5b(E)$.*
 5. *All $n - m$ derivation steps of $r_5bi(E)$ of Y_i for $i \in \mathcal{I}_n$.*

Proof: The proof of this Lemma is similar to Lemma 14's. Firstly, it relies on Lemma 20 for the guarantee that all the necessary encoding rules exists (For instance, if we have a derivation step of the *r_4ai* encoding rule, then *r_3* must have been applied as some preceding derivation step). The next part of the proof is to show that we can safely push all interleaving derivations between the consensus E derivation

steps away (either before or after all E derivation steps). By Lemma 9, all we need to prove that all such interleaving derivations are non-overlapping with the E derivation steps in question.

For the case that E derivation steps model the successful application of a rule instance r , since the encoding rules r_3 and each r_4ai are defined such that their propagated rule heads are persistent (r_3 propagated heads consist of only P_x while each r_4ai consist of only P_i , all of which are persistent), derivations interleaving between any of these encoding rules must be non-overlapping with all E derivation steps. As such we can push all interleaving derivations before all r_3 and r_4ai E derivation steps. For r_5a , its rule head contains entirely of only commit synchronizing constraints produced by r_3 and r_4ai , hence all other interleaving derivations can be push after it. Hence we have proven this Lemma for the success case.

For the case that E derivation steps model the abortion of an attempt to apply a rule instance of r , encoding rules r_3 and r_4ai are the same as the previous case, while we now also have r_4bi rules. These encoding rules are single headed rules with only one synchronizing constraint r_vote_i as its only rule head. Since this constraint is produced by r_3 , we can push all other non consensus E derivation steps interleaving between them, before the derivation step of r_3 . For encoding rules 5_bx and 5_bi , all their rule heads are synchronizing constraints produced by r_4ai and r_4bi rules. Hence we can push all other interleaving derivations downwards after all E derivation steps. Hence we have proven this Lemma for the abort case. \square

Theorem 23 (Soundness of n -Neighbor Restricted Encoding) *Given an n -neighbor restricted and locally quiescent CHR^e program \mathcal{P}_n and a 0-neighbor restricted CHR^e program \mathcal{P}_0 such that $\mathcal{P}_n \rightsquigarrow_{nNb} \mathcal{P}_0$, for any reachable states Ω and Ω' , if $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, then we have either Ω' is not commit free or $\mathcal{P}_n \triangleright \text{DropSyncs}([\Omega]; \mathcal{P}_n) \mapsto_{\omega_\alpha^e}^* \text{DropSyncs}([\Omega']; \mathcal{P}_n)$.*

Proof: The prove of this theorem is similar to Theorem 15's. Firstly, by Theorem 6, we are guaranteed that from $\mathcal{P}_0 \triangleright \Omega \mapsto_{\omega_0^e}^* \Omega'$, we have $\mathcal{P}_0 \triangleright [\Omega] \mapsto_{\omega_\alpha^e}^* [\Omega']$. By Lemma 22, we can safely permutate any such derivations to an equivalent one, such that all consensus E derivations are compressed together (no concurrent and interleaving non E derivations between each E derivation step). Hence, we can focus our attention only to such orderly derivations. From here, we can prove this Theorem in a manner similar to Theorem 15: by induction on ω_α^e derivation steps to show that each derivation step of \mathcal{P}_0 to a commit free state has a corresponding derivation of \mathcal{P}_1 with all synchronizing constraints removed. \square