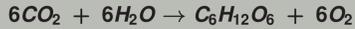


## 1. Challenges of Parallel and Distributed Programming

- ▶ A notoriously laborious and difficult endeavor
  - ▶ Wide range of technical difficulties (e.g. deadlock, atomicity, fault-tolerance).
  - ▶ Traditional computational problems (e.g. correctness, completeness, termination).
  - ▶ While ensuring scalability and performance effectiveness.
- ▶ Open research problem:
  - ▶ Distributed programming frameworks (e.g. Map reduce [DG08], Graph Lab [LGK<sup>+</sup>10], Pregel [MAB<sup>+</sup>10], Mizan [KKAJ10])
  - ▶ Distributed programming languages (e.g. Erlang [AV90], X10 [SSvP07], NetLog [GW10], Meld [CARG<sup>+</sup>12])
  - ▶ High-level programming abstractions (e.g. Join Patterns [TR11], Parallel CHR [LS11])
- ▶ We seek an approach that is *declarative*, based on *logical foundations*, *expressive and concise*.
- ▶ Motivated by chemical reaction equations:



## 2. Introducing Rule-Based Multiset Rewriting

- ▶ Constraint Handling Rules (CHR) [Frü98]
  - ▶ Rule-based constraint logic programming language.
  - ▶ Based on multiset rewriting over first order predicate terms.
  - ▶ Concurrent, committed choice and declarative.
- ▶ CHR programs consist of a set of CHR rules of the following form:
 
$$r : P \setminus S \iff G \mid B$$
  - ▶ Informally means: If we have  $P$  and  $S$  such that  $G$  is satisfiable, replace  $S$  with  $B$ .
- ▶ Example: Greatest common divisor (GCD)

```

base : gcd(0) ⇔ true
reduce : gcd(N) \ gcd(M) ⇔ 0 < N ∧ N ≤ M | gcd(M-N)

{gcd(9), gcd(6), gcd(3)}   reduce : gcd(6) \ gcd(9) ⇔ 0 < 6 ∧ 6 ≤ 9 | gcd(3)
→ {gcd(3), gcd(6), gcd(3)} reduce : gcd(3) \ gcd(6) ⇔ 0 < 3 ∧ 6 ≤ 9 | gcd(3)
→ {gcd(3), gcd(3), gcd(3)} reduce : gcd(3) \ gcd(3) ⇔ 0 < 3 ∧ 6 ≤ 9 | gcd(0)
→ {gcd(0), gcd(3), gcd(3)} base : gcd(0) ⇔ true
→ {gcd(3), gcd(3)}         reduce : gcd(3) \ gcd(3) ⇔ 0 < 3 ∧ 6 ≤ 9 | gcd(0)
→ {gcd(0), gcd(3)}         base : gcd(0) ⇔ true
→ {gcd(3)}
    
```

## 3. MSRE, Distributed Multiset Rewriting for Ensembles

- ▶ Elements are *distributed* across distinct locations ( $k_1, k_2$ , etc..), each possessing its own multiset of elements.
 
$$\{edge(k_2, 1), \dots\}@k_1 \iff \{edge(k_1, 2), edge(k_3, 8), \dots\}@k_2$$

$$\downarrow$$

$$\{edge(k_1, 10)\}@k_3$$
- ▶ Rewrite rules explicitly reference the *relative location* of constraints:
  - base rule* :  $[X] edge(Y, D) \setminus \cdot \iff [X] path(Y, D)$ .
  - elim rule* :  $[X] path(Y, D1) \setminus [X] path(Y, D2) \iff D1 < D2 \mid true$ .
  - trans rule* :  $[X] edge(Y, D), [Y] path(Z, D') \iff X \neq Z \mid [X] path(Z, D + D')$ .
- ▶ Rewrite rules can specify “local” rewriting:
 
$$\{edge(k_2, 1), path(k_2, 1), path(k_2, 10)\}@k_1 \dots$$

$$\rightarrow \{edge(k_2, 1), path(k_2, 1)\}@k_1 \dots \quad [k_1] path(k_2, 1) \setminus [k_1] path(k_2, 10) \iff 1 < 10 \mid true$$
- ▶ Rewrite rules can specify link-restricted rewriting:
 
$$\{edge(k_2, 1), \dots\}@k_1 \iff \{path(k_3, 8), edge(k_1, 2), edge(k_3, 8), \dots\}@k_2$$

$$\downarrow$$

$$\{edge(k_1, 10)\}@k_3$$

$$\rightarrow$$

$$\{edge(k_2, 1), path(k_3, 9), \dots\}@k_1 \iff \{path(k_3, 8), edge(k_1, 2), edge(k_3, 8), \dots\}@k_2$$

$$\downarrow$$

$$\{edge(k_1, 10)\}@k_3$$

$$[k_1] edge(k_2, 1), [k_2] path(k_3, 8) \iff k_1 \neq k_3 \mid [k_1] path(k_3, 9)$$

## 4. Example: Parallel Mergesort

Parallel mergesort: Assumes tightly coupled ensembles (multicore, shared memory, etc..)

```

[X] unsorted([I]) ⇔ [X] sorted([I]).
[X] unsorted(Xs) ⇔ len(Xs) > 2 | exists Y. exists Z. let (Ys, Zs) = split(Xs).
    [Y] parent(X), [Y] unsorted(Ys), [Z] parent(X), [Z] unsorted(Zs).
[X] sorted(Xs), [X] parent(Y) ⇔ [Y] unmerged(Xs).
[X] unmerged(Xs1), [X] unmerged(Xs2) ⇔ [X] sorted(merge(Xs1, Xs2))
    
```

- ▶ New locations “dynamically” created to solve sub-problems.
- ▶ completed sub-problems are transmitted to the “parent” location.

## 5. Example: Distributed Hyper-Quicksort

Distributed Hyper-Quicksort: Assumes loosely coupled ensembles (network, message passing interface, etc..)

```

-- “Local” sorting algorithm Parallel merge sort rules
...
-- Distributed Hyper quicksort rules
[X] sorted(Xs), [X] leader() \ [X] leaderLinks(G) ⇔ len(G) > 1 |
    let LG, GG = split(G). [X] leaderLinks(LG),
        [head(GG)] leader(), [head(GG)] leaderLinks(GG),
        {[Y] median(Xs | len(Xs) / 2) | Y in G}
        {[Y] partnerLink(Z) | Y, Z in zip(LG, GG)}
[X] median(M), [X] sorted(Xs) ⇔ let Ls, Gs = partition(Xs, M). [X] leqM(Ls), [X] grM(Gs)
[X] partnerLink(Y), [X] grM(Xs), [Y] leqM(Ys) ⇔ [X] leqM(Ys), [Y] grM(Xs)
[X] leqM(Ls1), [X] leqM(Ls2) ⇔ [X] sorted(merge(Ls1, Ls2))
[X] grM(Gs1), [X] grM(Gs2) ⇔ [X] sorted(merge(Gs1, Gs2))
    
```

- ▶ Data (unsorted numbers) initially distributed across  $2^n$  locations.
- ▶ In termination,  $2^n$  locations are in total order.

## 6. Multiset Comprehensions

- ▶ MSRE language includes *multiset comprehension patterns*
- ▶ Additional form of constraint patterns:  $\{p(\vec{t}) \mid g\}_{\vec{x} \in t}$ 
  - ▶ Collects *all*  $p(\vec{t})$  in the store that satisfy  $g$
  - ▶  $\vec{t}$  is the multiset of all bindings to  $\vec{x}$
- ▶ Example: pivoted swapping via comprehensions

```

swap(X, Y, P)
pivotSwap @ {data(X, D) | D ≥ P}_{D ∈ Xs} ⇔ {data(Y, D)}_{D ∈ Xs}
{data(Y, D) | D < P}_{D ∈ Ys} ⇔ {data(X, D)}_{D ∈ Ys}
    
```

- ▶  $Xs$  and  $Ys$  built from the store — *output*
- ▶  $Xs$  and  $Ys$  used to unfold the comprehensions — *input*

## 7. Preliminary Experiment Results

Program	Input Size	Standard rules only			With comprehensions			Code reduction (lines)		Speedup
		5 preds	7 rules	21 lines	2 preds	1 rule	10 lines	110%		
Swap	(40, 100)	241 vs 290	121 vs 104	vs 104	+OJO	+Mono	+OJO	vs 92	vs 91	33%
	(200, 500)	1813 vs 2451	714 vs 681	vs 670	+Bt		+Uniq	vs 621	vs 597	20%
	(1000, 2500)	8921 vs 10731	3272 vs 2810	vs 2651				vs 2554	vs 2502	31%
	(100, 200)	814 vs 1124	452 vs 461	vs 443				vs 437	vs 432	5%
GHS	(500, 1000)	7725 vs 8122	3188 vs 3391	vs 3061				vs 3109	vs 3005	6%
	(2500, 5000)	54763 vs 71650	15528 vs 16202	vs 15433				vs 16097	vs 15214	2%
HQSort	(8, 50)	1275 vs 1332	1117 vs 1151	vs 1099				vs 1151	vs 1081	10%
	(16, 100)	5783 vs 6211	3054 vs 2980	vs 2877				vs 2916	vs 2702	15%
	(32, 150)	13579 vs 14228	9218 vs 8745	vs 8256				vs 8107	vs 8013	15%

- ▶ Execution times (ms) for various optimizations on programs with increasing input size.
- ▶ Preliminary, but promising: we get code reduction plus some performance improvement.

## 8. Research Outcome, So Far...

- ▶ Prototype Implementation:
  - ▶ Download: <https://github.com/sllam/chrcp>
  - ▶ Online Demo: <http://rise4fun.com/msre>
- ▶ Publications:
  - ▶ “Decentralized Execution of Constraint Handling Rules for Ensembles”, PPDP’13
  - ▶ “Constraint Handling Rules with Multiset Comprehension Patterns”, CHR’14
  - ▶ “Reasoning About Set Comprehensions”, SMT’14
  - ▶ “Optimized Compilation of Multiset Rewriting with Comprehensions”, APLAS’14

## 9. What’s Next?

- ▶ Practical applications for MSRE:
  - ▶ cloud computing applications
  - ▶ P2P apps on mobile devices, Android SDK (Funded by the QNRF, JSREP 4-003-2-001)
- ▶ Logical interpretation of comprehensions:
  - ▶ Logical proof system for MSRE
  - ▶ Linear logic + Fixed Points + Subexponentials

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