



Automating Incremental and **Asynchronous** Evaluation for Recursive Aggregate Data Processing

Qiange Wang, Yanfeng Zhang, Hao Wang, Liang Geng,
Rubao Lee, Xiaodong Zhang, Ge Yu

Northeastern University, China
The Ohio State University, USA



What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

Datalog perfectly expresses RAPs:

Single Source Shortest Path(SSSP)

$sssp(X, dx) \quad :- \quad X='a', dx=0.$

$sssp(Y, \min[dy]) \quad :- \quad sssp(X,dx), \text{edge}(X,Y,dxy), dy=dx+dxy.$

What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

Datalog perfectly expresses RAPs:

Single Source Shortest Path(SSSP)

$\text{sssp}(X, dx) \quad :- \quad X='a', dx=0.$ \rightarrow Initialization Rule

$\text{sssp}(Y, \min[dy]) \quad :- \quad \text{sssp}(X,dx), \text{edge}(X,Y,dxy), dy=dx+dxy.$

What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

Datalog perfectly expresses RAPs:

Single Source Shortest Path(SSSP)

$sssp(X, dx) \quad :- \quad X='a', dx=0.$

→ Initialization Rule

$sssp(Y, \min[dy]) \quad :- \quad sssp(X, dx), \text{edge}(X, Y, dxy), dy=dx+dxy.$

→ Recursive Rule

What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

Datalog perfectly expresses RAPs:

Single Source Shortest Path(SSSP)

$sssp(X, dx) \quad :- \quad X='a', dx=0.$

→ Initialization Rule

$sssp(Y, \mathbf{min[dy]}) \quad :- \quad sssp(X, dx), \text{edge}(X, Y, dxy), dy=dx+dxy.$

→ Recursive Rule

↓
Aggregate Function

What is a Recursive Aggregate Program(RAP)?

For a given algorithm defined by a function, it is a RAP if it satisfies:

- (1) The function consists of a basic rule for an **initiation** and a small set of other rules for **iteratively self-calling** (recursion).
- (2) The recursive part of the function includes **aggregate** operators, such as MIN, MAX, AVG, etc.

Datalog perfectly expresses RAPs:

Single Source Shortest Path(SSSP)

`sssp(X, dx) :- X='a', dx=0.`

→ Initialization Rule

`sssp(Y, min[dy]) :- sssp(X,dx), edge(X,Y,dxy), dy=dx+dxy.`

→ Recursive Rule

Aggregate Function

Non-aggregate Functions

Recursive Aggregate Programs are Everywhere

□ Connected Components

```
cc(X,X)      :- edge (X,_).
```

```
cc(Y,min [v]) :- cc(X,v), edge (X,Y).
```


Recursive Aggregate Programs are Everywhere

□ Connected Components

```
cc(X,X)      :- edge (X,_).  
cc(Y,min [v]) :- cc(X,v), edge (X,Y).
```

□ PageRank

```
rank(i+1, Y, sum[ry]) :- node(Y), ry=0.2;  
                        :- rank(i, X, rx), edge(X,Y),  
                        :- degree(X, d), ry=0.8 rx/d.
```

Recursive Aggregate Programs are Everywhere

□ Connected Components

```
cc(X,X)      :- edge (X,_).  
cc(Y,min [v]) :- cc(X,v), edge (X,Y).
```

□ PageRank

```
rank(i+1, Y, sum[ry]) :- node(Y), ry=0.2;  
                        :- rank(i, X, rx), edge(X,Y),  
                        :- degree(X, d), ry=0.8 rx/d.
```

□ Graph Convolution Neural Network

```
GCN (j+1,Y, sum [g1]) :- GCN (j,X,g), A(X,Y,w), Para (p),  
                        p1 = relu(g*p) *w
```



□ And many others

Evaluating a Recursive Aggregate Program

Single Source Shortest Path (SSSP)

$sssp(X, dx) \quad :- \quad X='a', dx=0.$

$sssp(Y, \min[dy]) \quad :- \quad sssp(X,dx), \text{edge}(X,Y,dxy), dy=dx+dxy.$

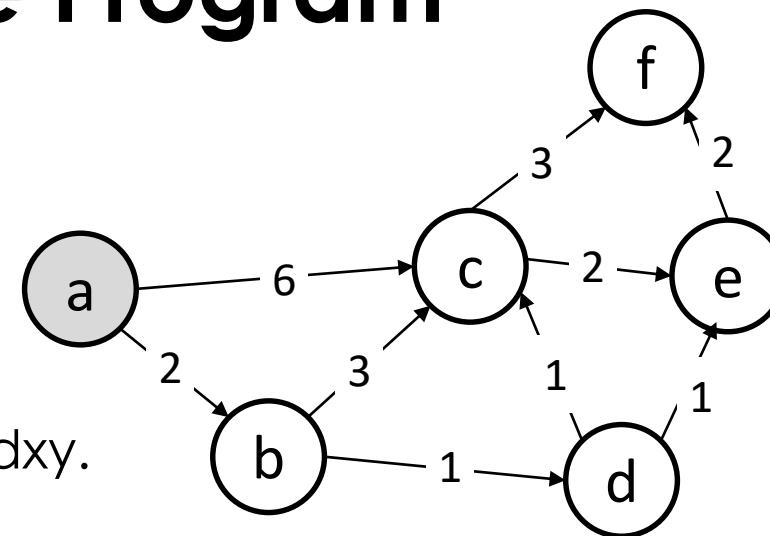


Evaluating a Recursive Aggregate Program

Single Source Shortest Path (SSSP)

$sssp(X, dx) \text{ :- } X='a', dx=0.$

$sssp(Y, \min[dy]) \text{ :- } sssp(X, dx), \text{edge}(X, Y, dxy), dy=dx+dxy.$



a->a	0	0	0	0	0
a->b	∞	2	2	2	2
a->c	∞	6	5	4	4
a->d	∞	∞	3	3	3
a->e	∞	∞	8	4	4
a->f	∞	∞	9	8	6

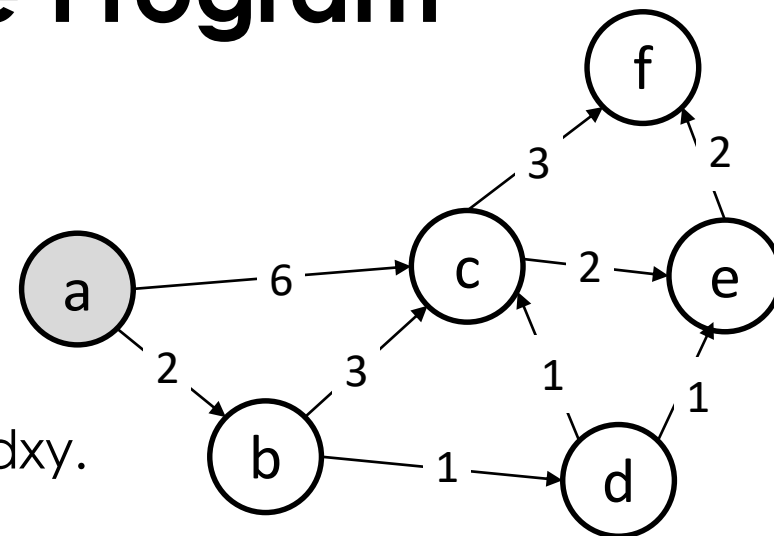
Iteratively performing **aggregate** and **non aggregate** function to derive result until the program reach the convergence.

Evaluating a Recursive Aggregate Program

Single Source Shortest Path (SSSP)

$sssp(X, dx) \text{ :- } X='a', dx=0.$

$sssp(Y, \min[dy]) \text{ :- } sssp(X, dx), \text{edge}(X, Y, dxy), dy=dx+dxy.$



a->a	0	0	0	0	0
a->b	∞	2	2	2	2
a->c	∞	6	5	4	4
a->d	∞	∞	3	3	3
a->e	∞	∞	8	4	4
a->f	∞	∞	9	8	6


Iteratively performing **aggregate** and **non aggregate** function to derive result until the program reach the convergence.

Evaluating a Recursive Aggregate Program

Single Source Shortest Path (SSSP)

$sssp(X, dx) \text{ :- } X='a', dx=0.$

$sssp(Y, \min[dy]) \text{ :- } sssp(X, dx), \text{edge}(X, Y, dxy), dy=dx+dxy.$



a->a	0	0	0	0	0
a->b	∞	2	2	2	2
a->c	∞	6	5	4	4
a->d	∞	∞	3	3	3
a->e	∞	∞	8	4	4
a->f	∞	∞	9	8	6

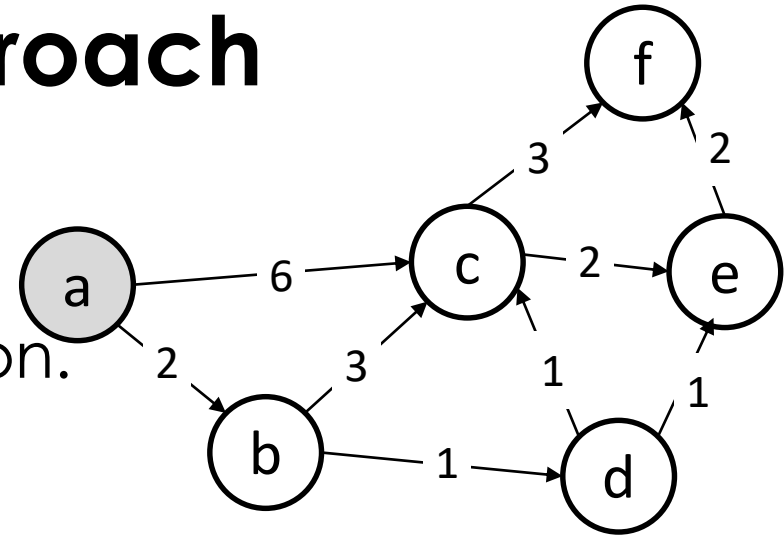
Iteratively performing **aggregate** and **non aggregate** function to derive result until the program reach the convergence.

Naïve Evaluation

Semi-naïve Evaluation

Naïve Evaluation: A Direct Approach


Naïve Evaluation is a direct and intuitive solution.



	X^0
a->a	0
a->b	∞
a->c	∞
a->d	∞
a->e	∞
a->f	∞

Naïve Evaluation on SSSP Algorithm

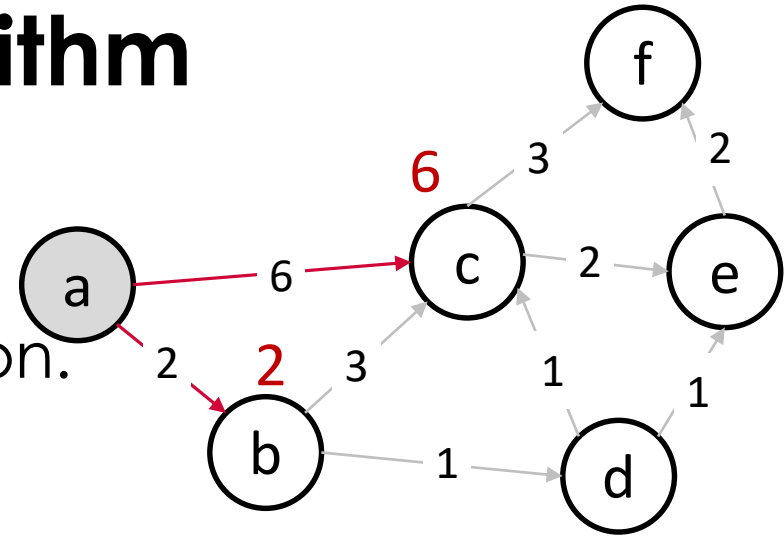
Naïve Evaluation is a direct and intuitive solution.



	X^0	$F(X^0)$	
a->a	0	0	a->a
a->b	∞	2	a->b
a->c	∞	6	a->c
a->d	∞		
a->e	∞		
a->f	∞		

Naïve Evaluation on SSSP Algorithm

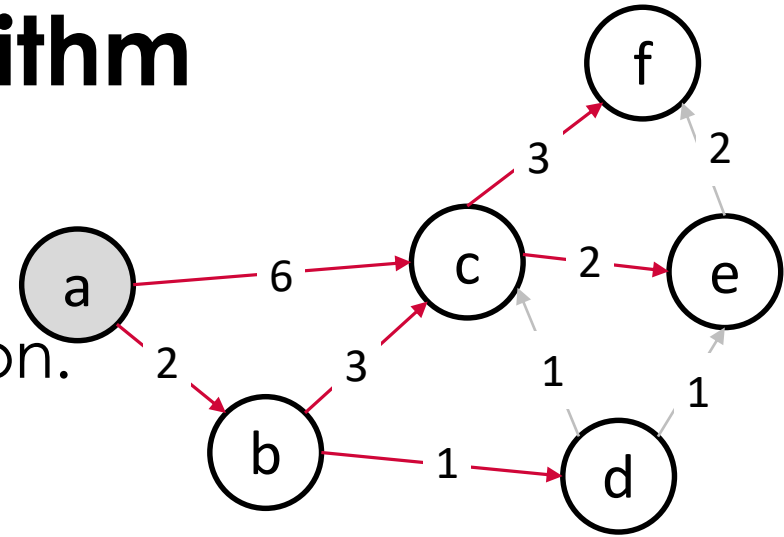
Naïve Evaluation is a direct and intuitive solution.



	X^0	$F(X^0)$	$X^1 = G(F(X^0))$
a->a	0	0	0
a->b	∞	2	2
a->c	∞	6	6
a->d	∞		∞
a->e	∞		∞
a->f	∞		∞

Naïve Evaluation on SSSP Algorithm



Naïve Evaluation is a direct and intuitive solution.



	X^0	$F(X^0)$	X^1	$F(X^1)$	
a->a	0	0	0	0	a->a
a->b	∞	2	2	2	a->b
a->c	∞	6	6	6	a->{c, b->c}
a->d	∞		∞	3	a->b->d
a->e	∞		∞	8	a->c->e
a->f	∞		∞	9	a->c->f

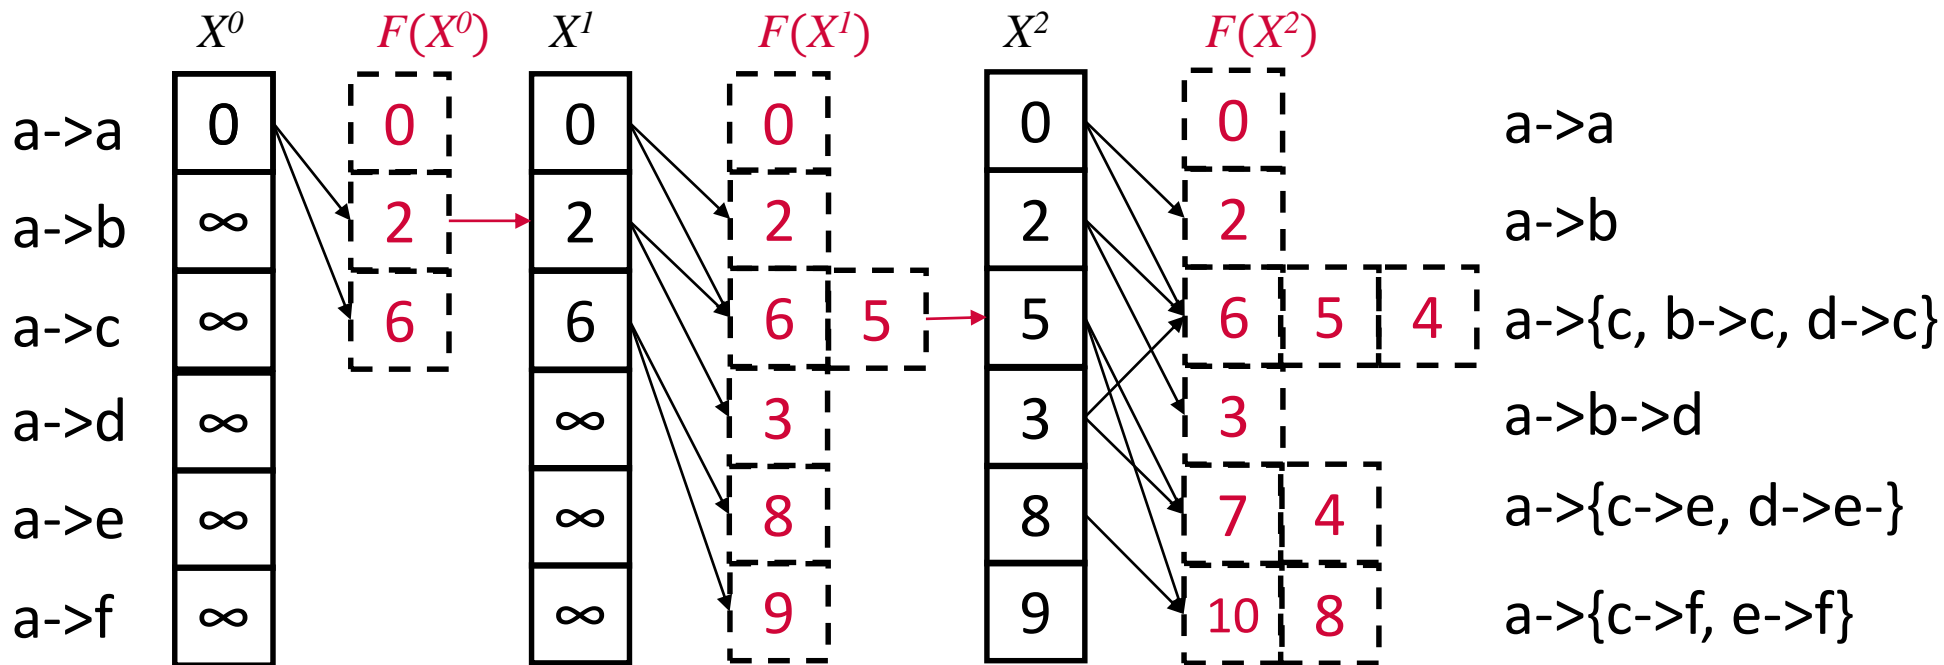

Naïve Evaluation on SSSP Algorithm

Naïve Evaluation is a direct and intuitive solution.



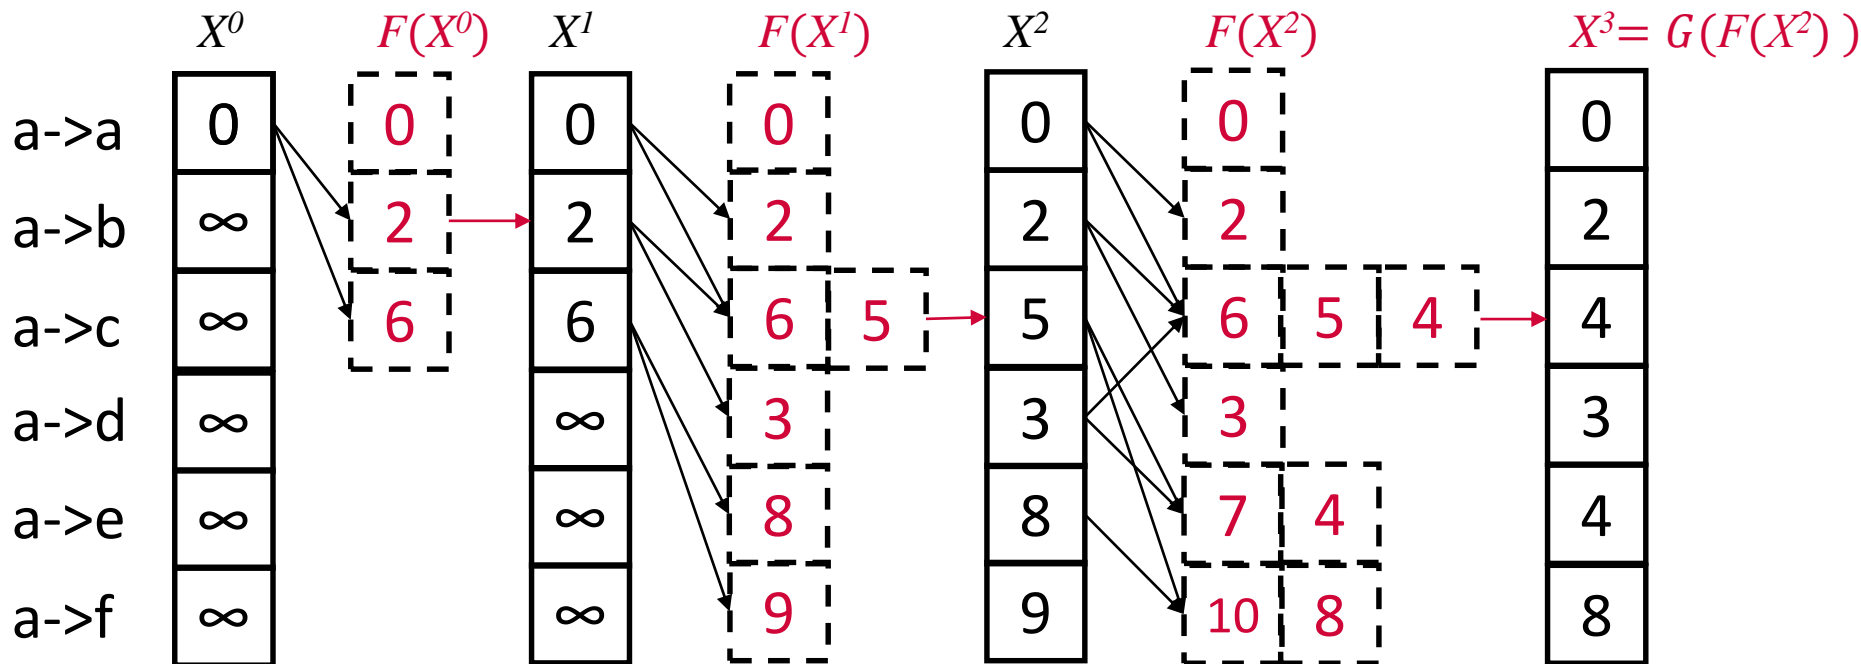

Naïve Evaluation on SSSP Algorithm

Naïve Evaluation is a direct and intuitive solution.





Naïve Evaluation on SSSP Algorithm

Naïve Evaluation is a direct and intuitive solution.





Naïve Evaluation on SSSP Algorithm

Naïve Evaluation is a direct and intuitive solution.




Naïve Evaluation on SSSP Algorithm

Naïve Evaluation is a direct and intuitive solution.




Problems on Naïve Evaluation


Naïve Evaluation is a direct and intuitive solution.



The **converged** nodes still need to involve.




Problems on Naïve Evaluation



Naïve Evaluation is a direct and intuitive solution.

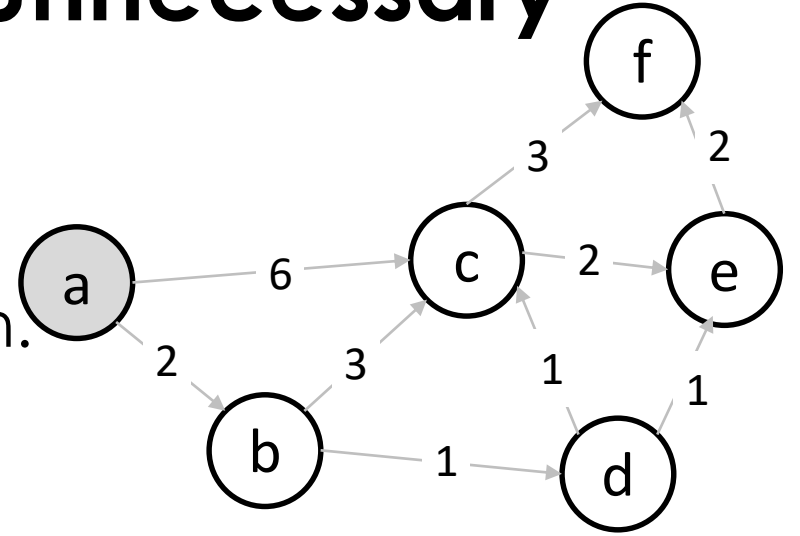
The **converged** nodes still need to involve.



Naïve evaluation brings **unnecessary** computation.

Semi-naïve Evaluation: Avoiding Unnecessary computation



Semi-naïve Evaluation is an incremental approach.



	X^0
a->a	0
a->b	∞
a->c	∞
a->d	∞
a->e	∞
a->f	∞



Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.



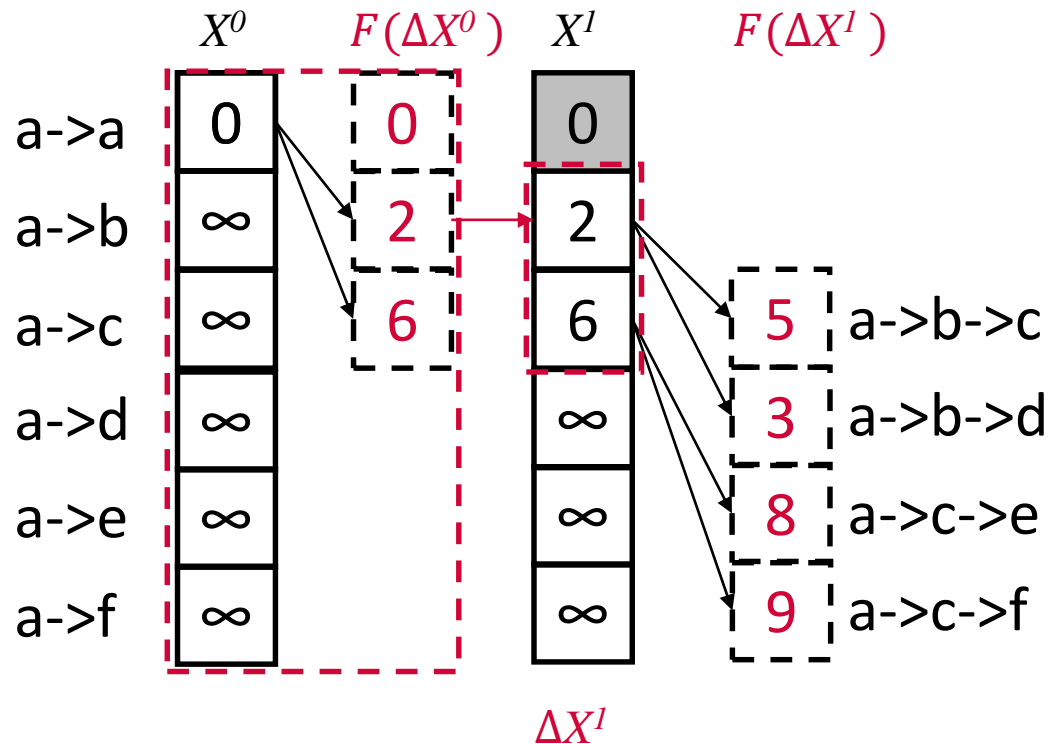

Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.




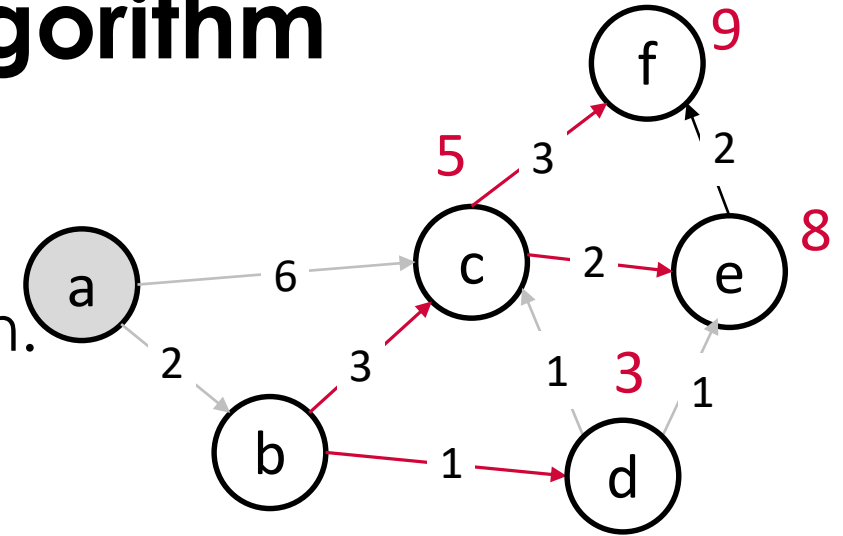
Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.





Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.





Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.





Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.





Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.



Semi-naïve Evaluation on SSSP Algorithm



Semi-naïve Evaluation is an incremental approach.




Semi-naïve Evaluation on SSSP Algorithm

Semi-naïve Evaluation is an incremental approach.

The converged nodes do not involved in the computation.




Naïve Evaluation vs. Semi-naïve Evaluation



A **fully computation** is performed.

The **converged** nodes still need to involve.





Only the necessary **computation** is performed.

The **converged** nodes do not need to involve.



Monotonicity Requirement for Semi-naïve Evaluation

The monotonicity property




SSSP: A Monotonic Example

The monotonicity property





Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$




SSSP: A Monotonic Example

The monotonicity property





Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$




SSSP: A Monotonic Example

The monotonicity property





Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$




SSSP: A Monotonic Example

The monotonicity property





Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$




SSSP: A Monotonic Example

The monotonicity property




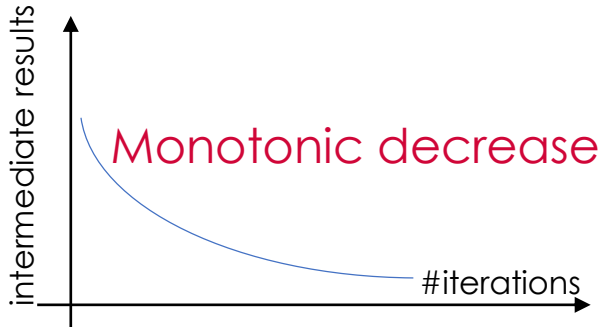
Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$



SSSP: A Monotonic Example


The monotonicity property



Single Source Shortest Path



$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$

Distance a→c **monotonically decreases** iteratively.



SSSP: A Monotonic Example

The monotonicity property




Single Source Shortest Path

$$\text{sssp}(Y, \min[\text{dy}]) \text{ :- } \text{sssp}(X, \text{dx}), \text{edge}(X, Y, \text{dxy}), \text{dy} = \text{dx} + \text{dxy}.$$




Distance a→c **monotonically decreases** iteratively.

The SSSP algorithm can be **semi-naïve** evaluated.




PageRank: A Non-monotonic Example

The monotonicity property



PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- $\text{node}(Y), ry=0.2;$
:- $\text{rank}(i, X, rx), \text{edge}(X,Y),$
:- $\text{degree}(X, d), ry=0.8 rx/d.$

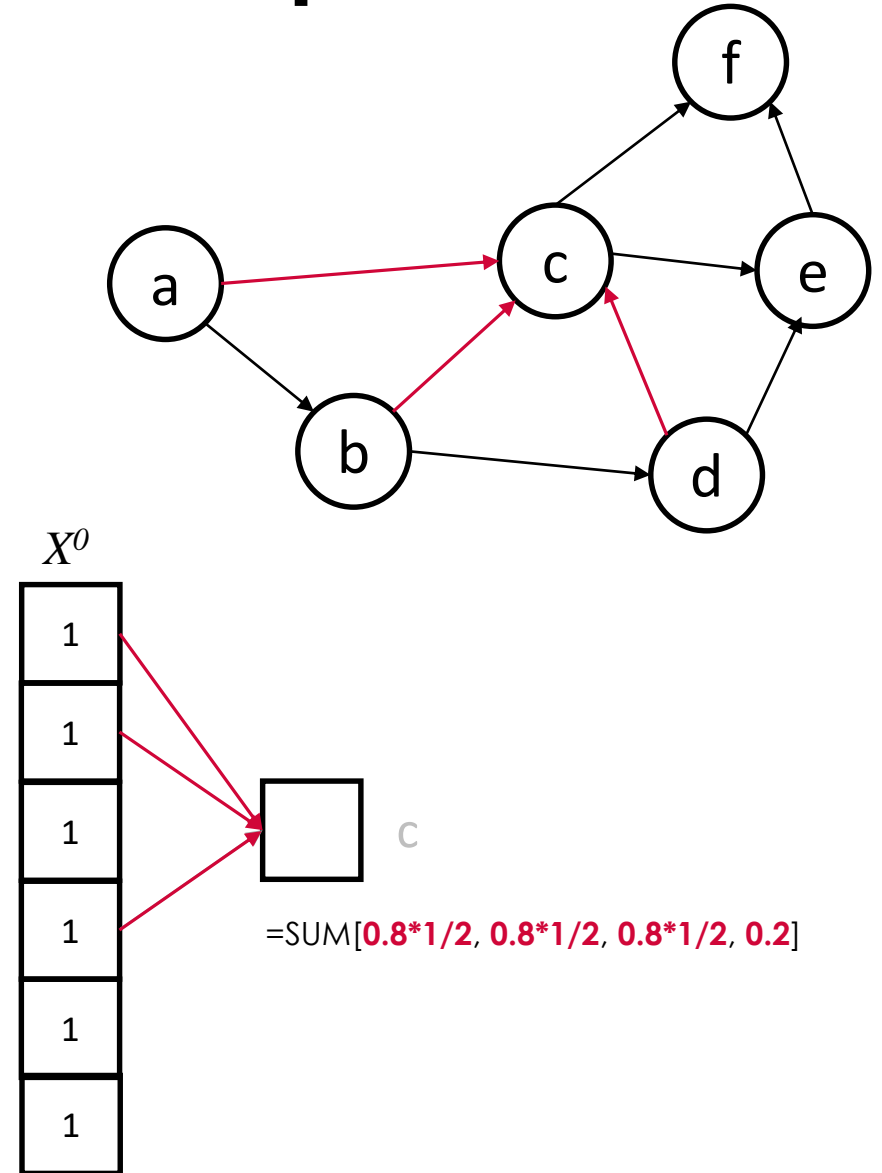


PageRank: A Non-monotonic Example

The monotonicity property

PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- $\text{node}(Y), ry=0.2$;
:- $\text{rank}(i, X, rx), \text{edge}(X,Y)$,
:- $\text{degree}(X, d), ry=0.8 rx/d$.

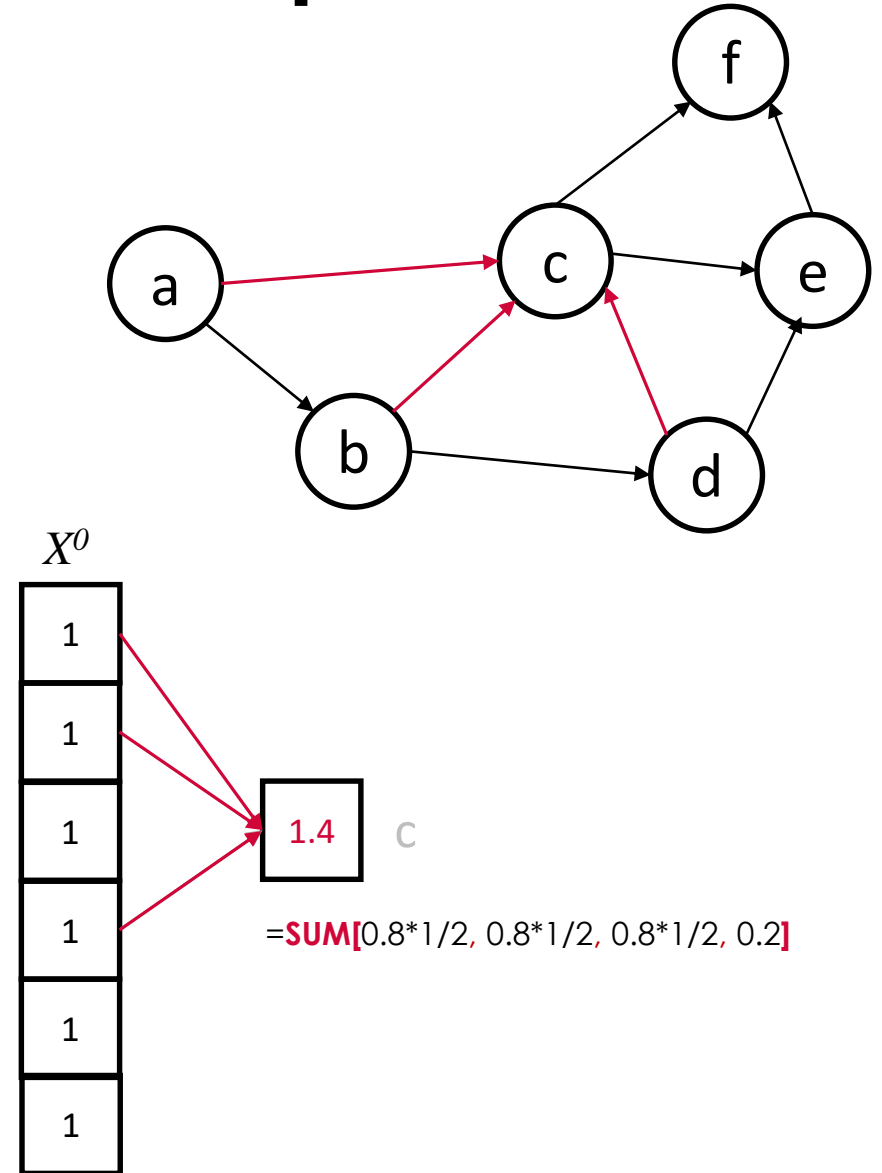


PageRank: A Non-monotonic Example

The monotonicity property


PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- node(Y), $ry=0.2$;
:- rank(i, X, rx), edge(X,Y),
:- degree(X, d), $ry=0.8 \text{ rx}/d$.




PageRank: A Non-monotonic Example

The monotonicity property



PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- $\text{node}(Y), ry=0.2;$
:- $\text{rank}(i, X, rx), \text{edge}(X,Y),$
:- $\text{degree}(X, d), ry=0.8 rx/d.$

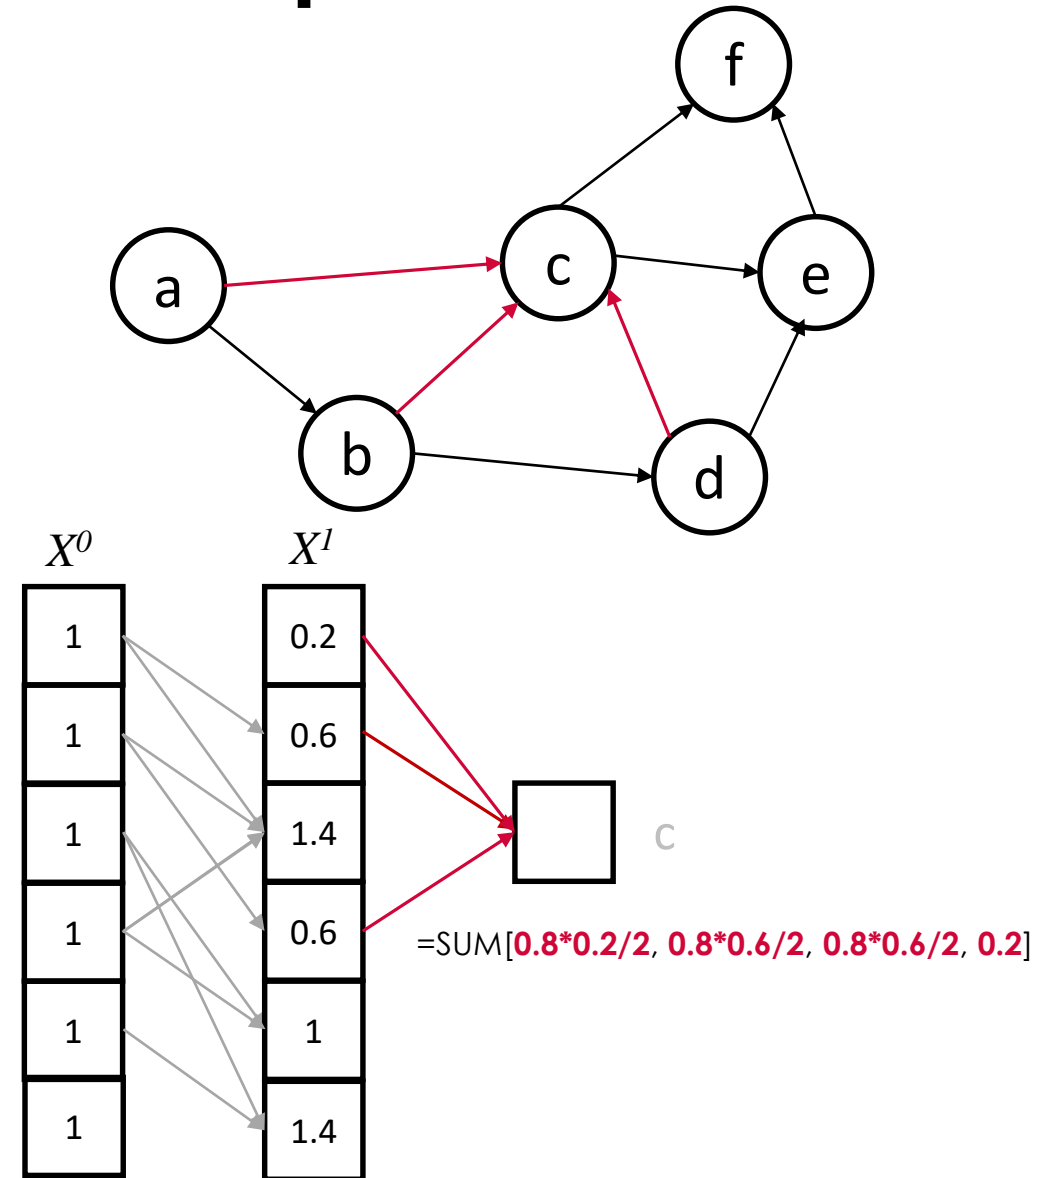


PageRank: A Non-monotonic Example

The monotonicity property

PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- $\text{node}(Y), ry=0.2;$
 :- $\text{rank}(i, X, rx), \text{edge}(X,Y),$
 :- $\text{degree}(X, d), ry=0.8 \text{ rx}/d.$

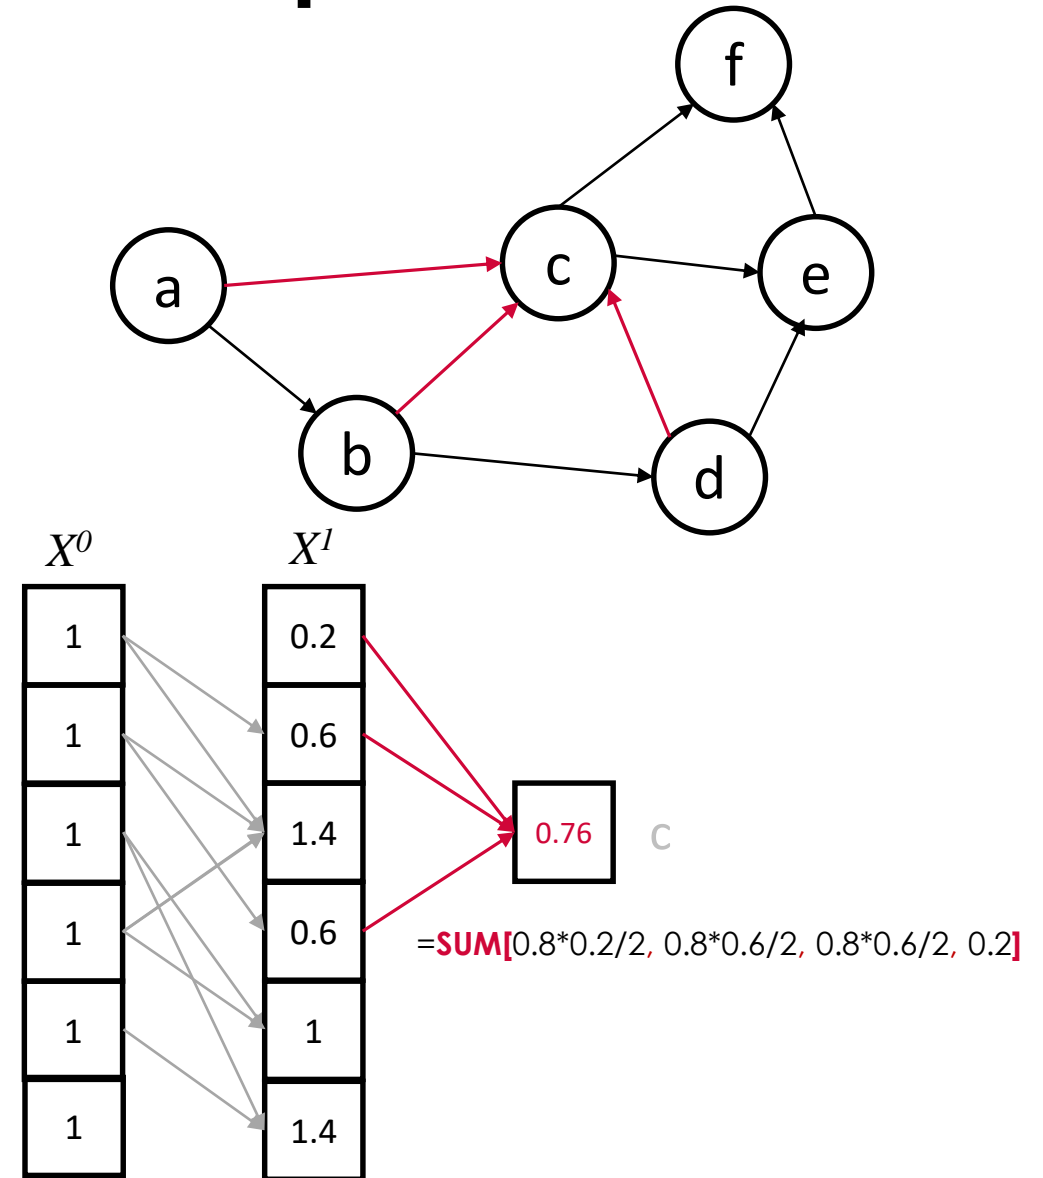


PageRank: A Non-monotonic Example

The monotonicity property


PageRank

$\text{rank}(i+1, Y, \text{sum}[ry])$:- node(Y), $ry=0.2$;
 :- rank(i, X, rx), edge(X,Y),
 :- degree(X, d), $ry=0.8 \text{ rx}/d$.




PageRank: A Non-monotonic Example

The monotonicity property




PageRank

$\text{rank}(i+1, Y, \text{sum}[ry]) \text{ :- node}(Y), ry=0.2;$
 $\text{:- rank}(i, X, rx), \text{edge}(X,Y),$
 $\text{:- degree}(X, d), ry=0.8 rx/d.$



PageRank: A Non-monotonic Example

The monotonicity property



PageRank


$\text{rank}(i+1, Y, \text{sum}[ry])$:- $\text{node}(Y), ry=0.2;$
 :- $\text{rank}(i, X, rx), \text{edge}(X,Y),$
 :- $\text{degree}(X, d), ry=0.8 rx/d.$



X^0	X^1	X^2	X^3	
1	0.2	0.2	0.2	a
1	0.6	0.28	0.28	b
1	1.4	0.76	0.57	c
1	0.6	0.44	0.31	d
1	1	1	0.68	e
1	1.4	1.56	1.30	f

PageRank: A Non-monotonic Example

The monotonicity property




PageRank

$\text{rank}(i+1, Y, \text{sum}[ry]) \text{ :- node}(Y), ry=0.2;$
 $\text{:- rank}(i, X, rx), \text{edge}(X,Y),$
 $\text{:- degree}(X, d), ry=0.8 rx/d.$




Rank values are **not monotonically** increase/decrease.



PageRank: A Non-monotonic Example

The monotonicity property



PageRank

$\text{rank}(i+1, Y, \text{sum}[ry]) \text{ :- node}(Y), ry=0.2;$
 $\text{:- rank}(i, X, rx), \text{edge}(X,Y),$
 $\text{:- degree}(X, d), ry=0.8 rx/d.$



Rank values are **not monotonically** increase/decrease.

PageRank is a **non-monotonic** algorithm.

X^0	X^1	X^2	X^3	
1	0.2	0.2	0.2	a
1	0.6	0.28	0.28	b
1	1.4	0.76	0.57	c
1	0.6	0.44	0.31	d
1	1	1	0.68	e
1	1.4	1.56	1.30	f



Existing Works on Monotonic Conditions

- **Ross and Sagiv** [PODS'92], **Socialite** [ICDE'13] and **Myria** [VLDB'15] formalize the conditions for using **monotonic aggregate** in recursive query.

Existing Works on Monotonic Conditions

- **Ross and Sagiv** [PODS'92], **Socialite** [ICDE'13] and **Myria** [VLDB'15] formalize the conditions for using **monotonic aggregate** in recursive query.

As presented in **Myria** [VLDB'15], a recursive aggregate program can be semi-naïve evaluated if

1. It require the aggregate function to be bag-monotonic.
2. The non-aggregate function is monotonic w.r.t. aggregate function.

- Checking on the monotonicity for arbitrary **recursive aggregate programs** is still sophisticated.

Existing Works on Monotonic Conditions

- **Ross and Sagiv** [PODS'92], **Socialite** [ICDE'13] and **Myria** [VLDB'15] formalize the conditions for using **monotonic aggregate** in recursive query.

Existing works **Maiter** [TPDS'14], **GRAPE** [SIGMOD'17] demonstrate that some **non-monotonic** algorithm can also be incrementally executed by using a similar approach to semi-naïve evaluation.

1. PageRank
2. SimRank
3. Belief Propagation

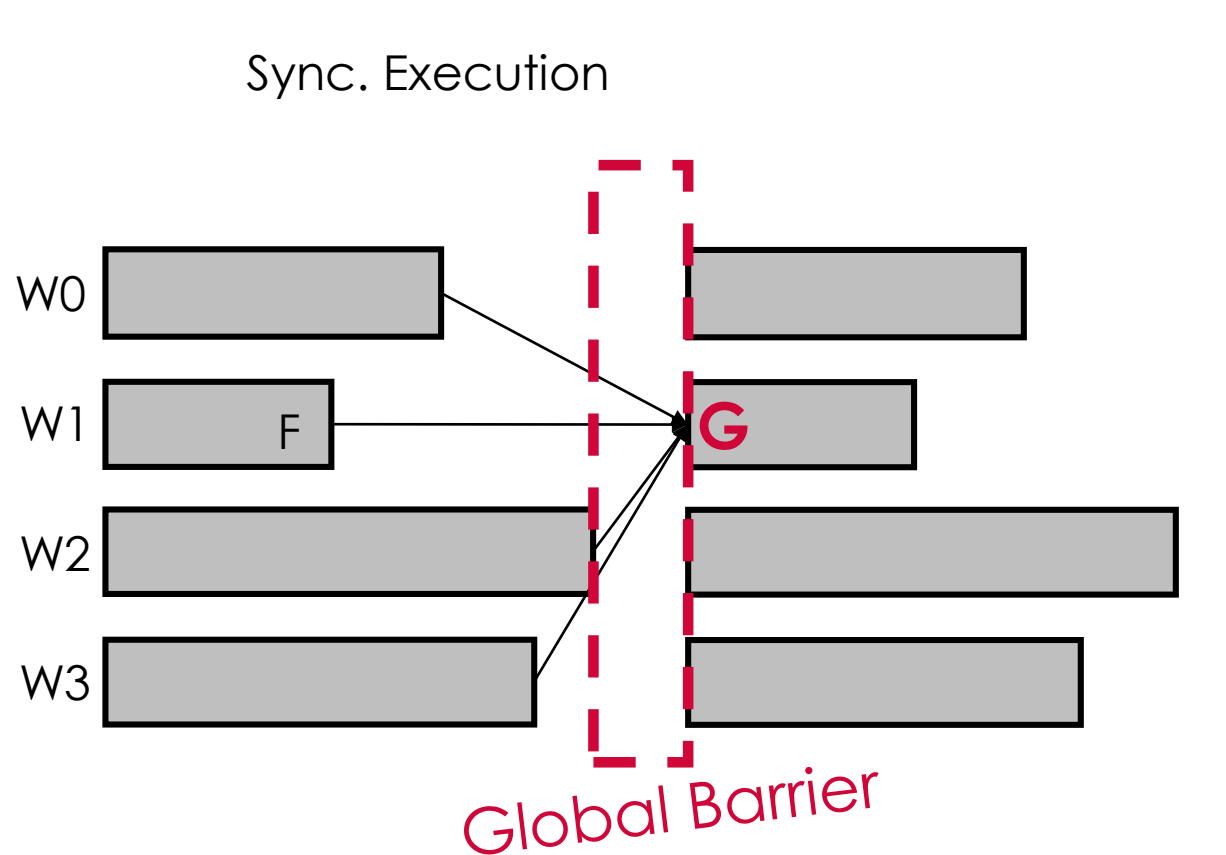
- The application scope of existing conditions is narrow.

Two Research Questions:

- Under what **conditions** can a Recursive Aggregate Program be evaluated with **semi-naïve** evaluation?
- How to **automatically** verify the conditions for evaluating a recursive aggregate program with semi-naïve evaluation?

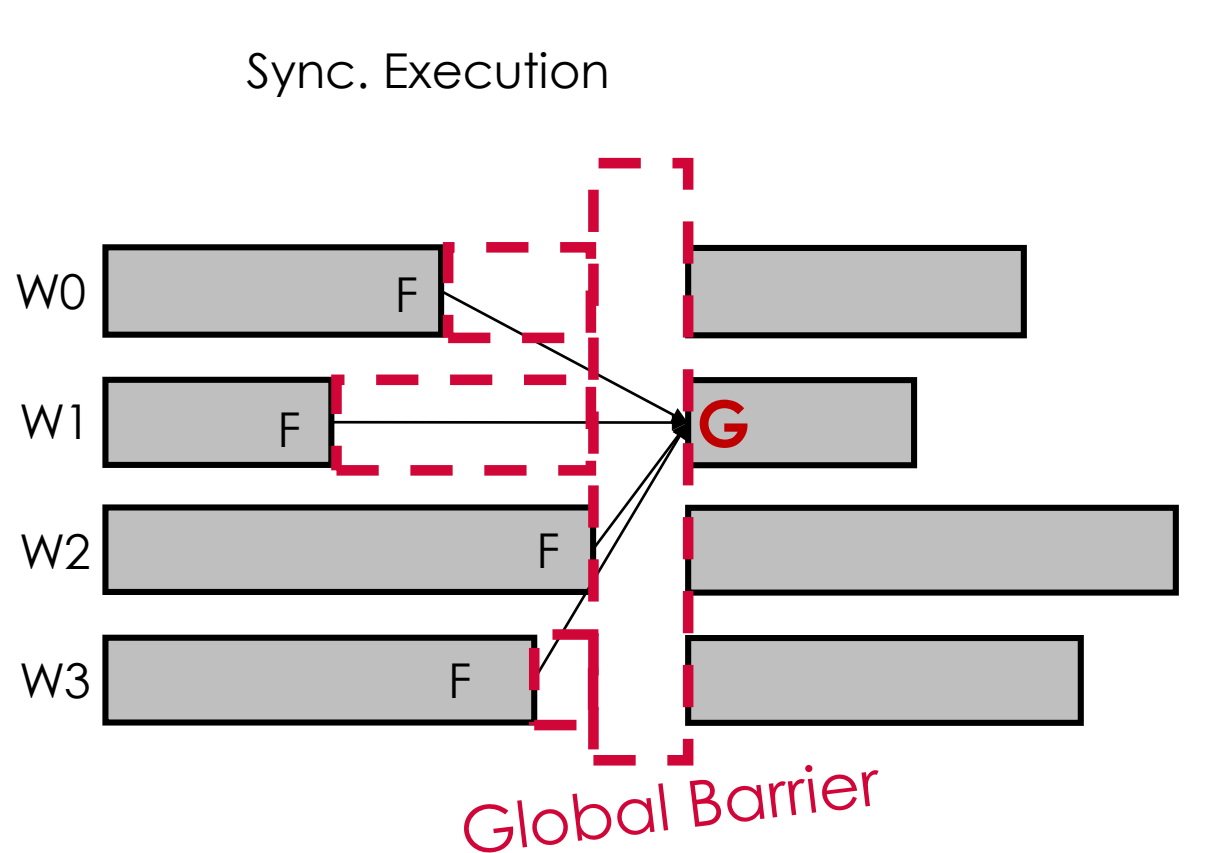
Async. Execution

The **aggregate** operations require communications among workers.
The global barrier of sync execution can guarantee the **correctness**.



Async. Execution


The **aggregate** operations require communications among workers.
The global barrier of sync execution can guarantee the **correctness**.



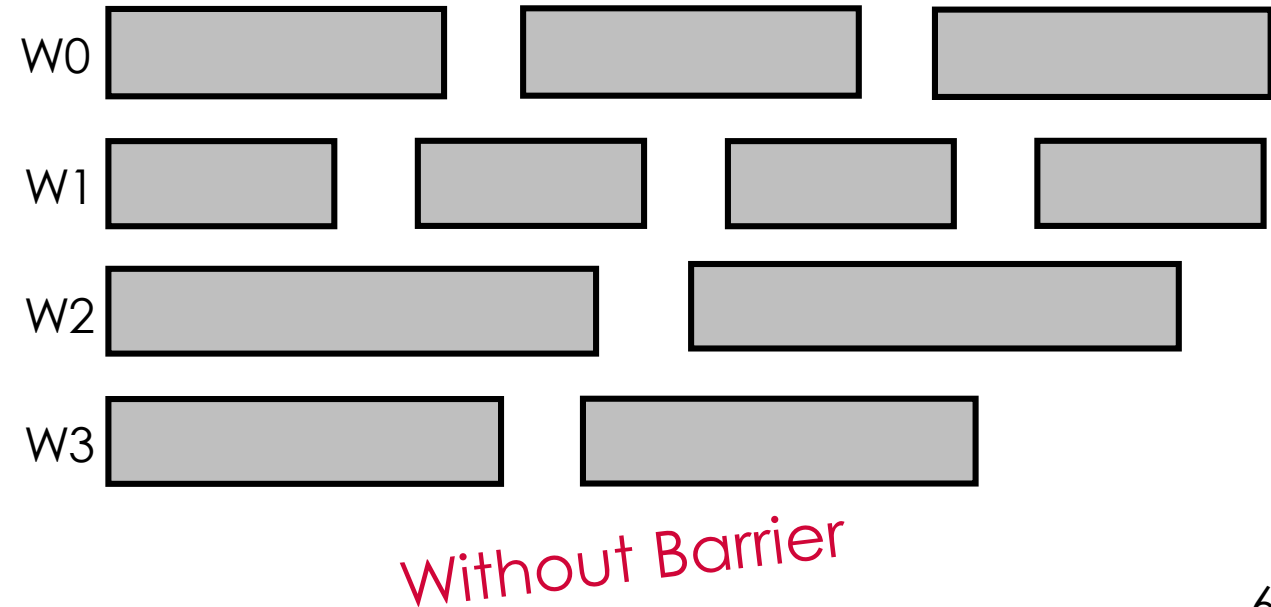
Async. Execution

Recently, The asynchronous parallel processing in distributed environment have emerged in the past few years, e.g., **Maiter**[TPDS'14], **Giraph++**[VLDB'15] and **Myria** [VLDB'15].

Sync. Execution



Async. Execution



Conditions for Async. Execution

- **Maiter** [TPDS'14] and **Grape+** [SIGMOD'18] propose the correctness conditions for graph algorithms with asynchronous execution.
- The conditions are not general enough for recursive aggregate programs.


Two more Research Questions:

- Under what **conditions** can a recursive aggregate programs be **asynchronously** executed ?
- How to **automatically** verify the conditions for executing a recursive aggregate program with async model?


Neither Sync nor Async Execution is Perfect

Neither Sync nor async execution can outperform each other because

- (1) sync processing may be **over-controlled** (too much idle time)
- (2) async processing may be **under-controlled** (stale computation)



The Same algorithm on different Datasets




The Same dataset on different Algorithms


Neither Sync nor Async Execution is Perfect

Neither Sync nor async execution can outperform each other because

- (1) sync processing may be **over-controlled** (too much idle time)
- (2) async processing may be **under-controlled** (stale computation)



The Same algorithm on different Datasets



The Same dataset on different Algorithms

We develop a **unified sync-async** engine to realize **properly controlled** processing.

Our Work

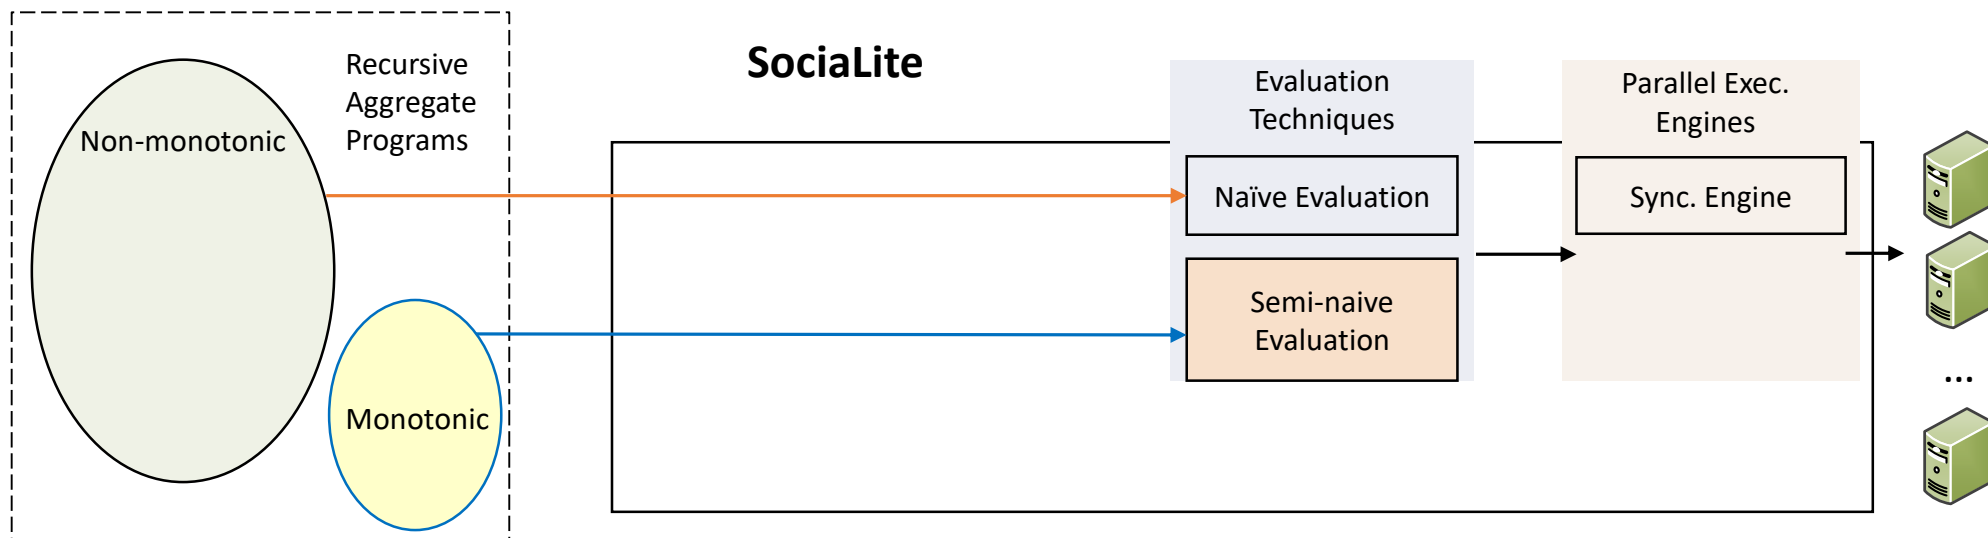
We develop and implement **PowerLog**, a high-performance distributed Datalog system based on **Socialite** [ICDE'13].

Our Work

We develop and implement **PowerLog**, a high-performance distributed Datalog system based on **Socialite** [ICDE'13].

Socialite

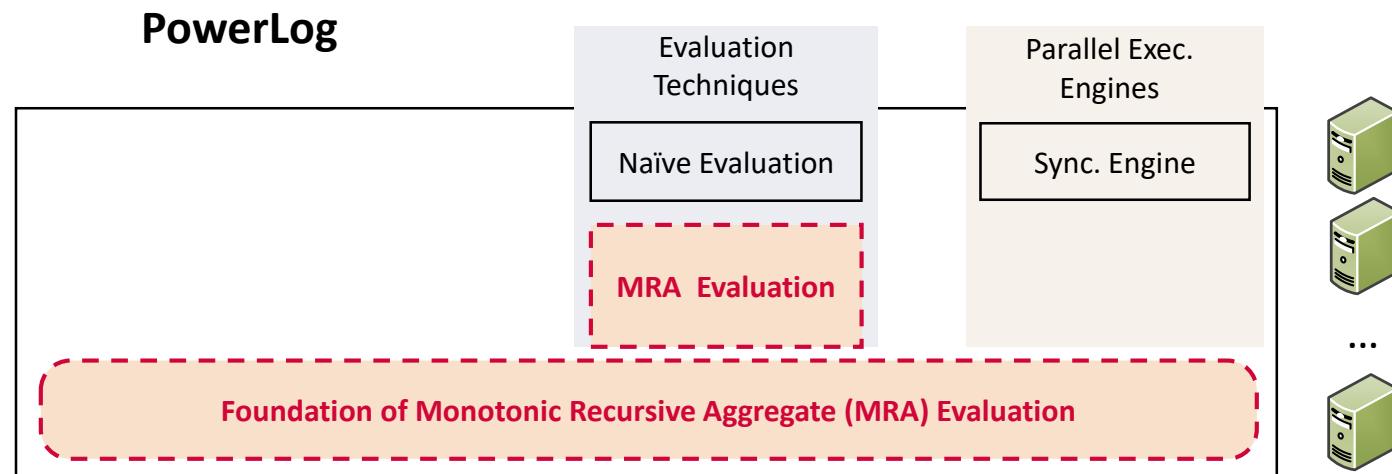

- (1) processes non-monotonic recursive aggregate programs with Naïve Evaluation
- (2) processes monotonic recursive aggregate programs with Semi-naïve Evaluation
- (3) executes both evaluation methods with a synchronous engine



Contributions of PowerLog

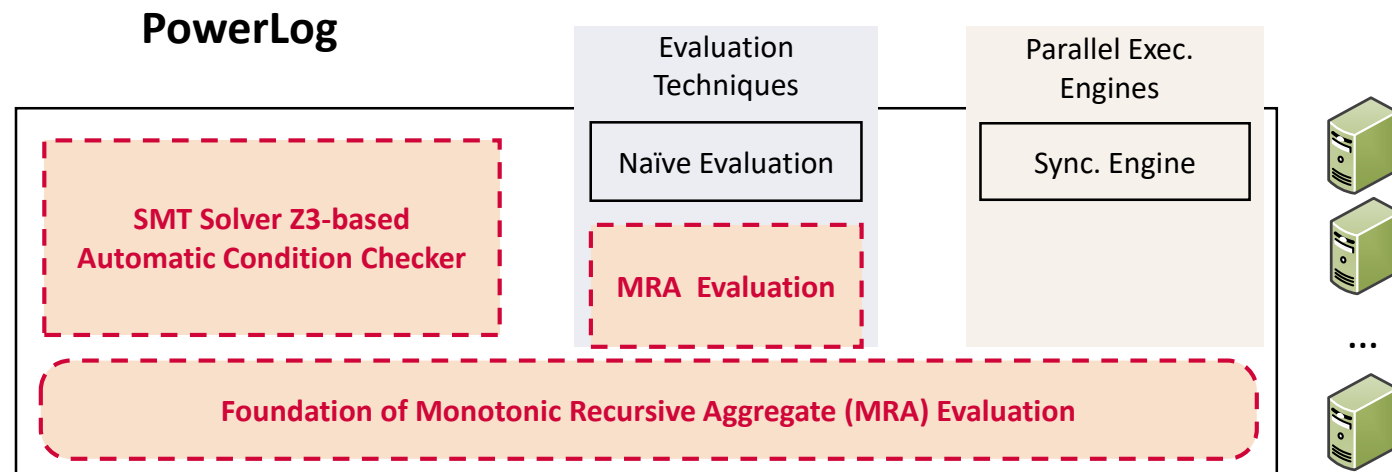
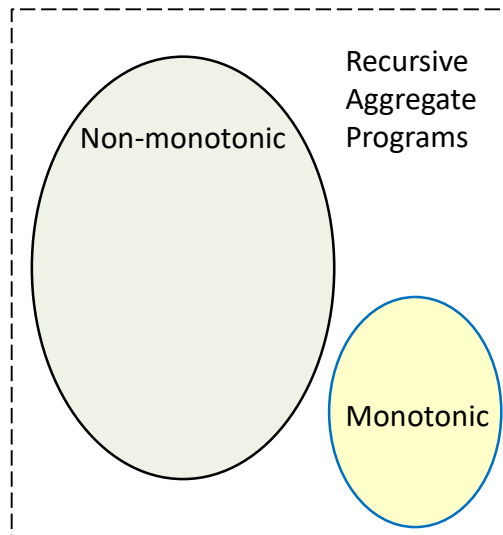
We propose **M**onotonic **R**ecursive **A**ggregate evaluation (**MRA Evaluation**), a variant of Semi-naïve Evaluation to answer.

- (1) under what conditions a recursive aggregate program can be executed with Semi-naïve Evaluation (**incrementally**)
- (2) under what conditions a recursive aggregate program can be executed **asynchronously**



Contributions of PowerLog

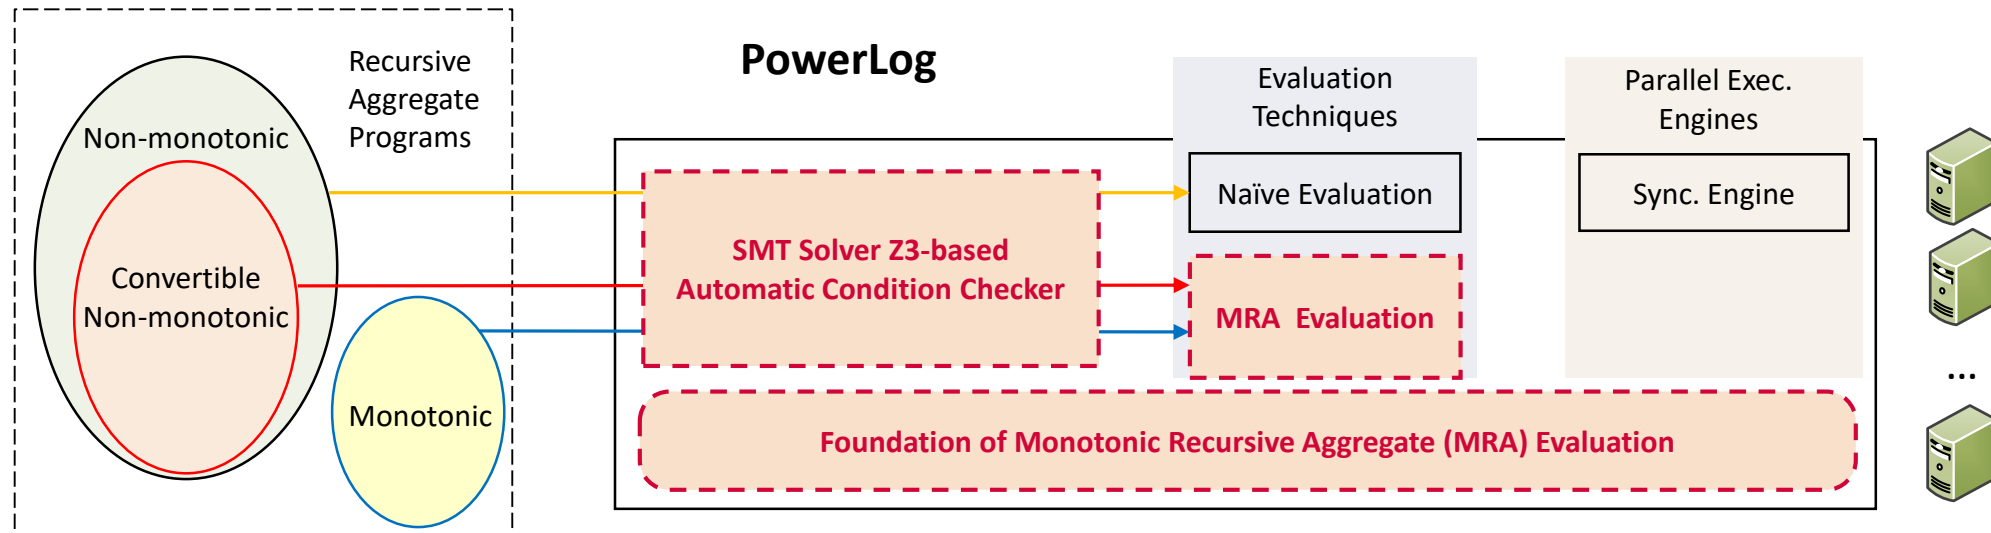
We design and implement **a Z3-based condition checker** to **automatically** check if a recursive aggregate program can satisfy the MRA conditions.



Contributions of PowerLog

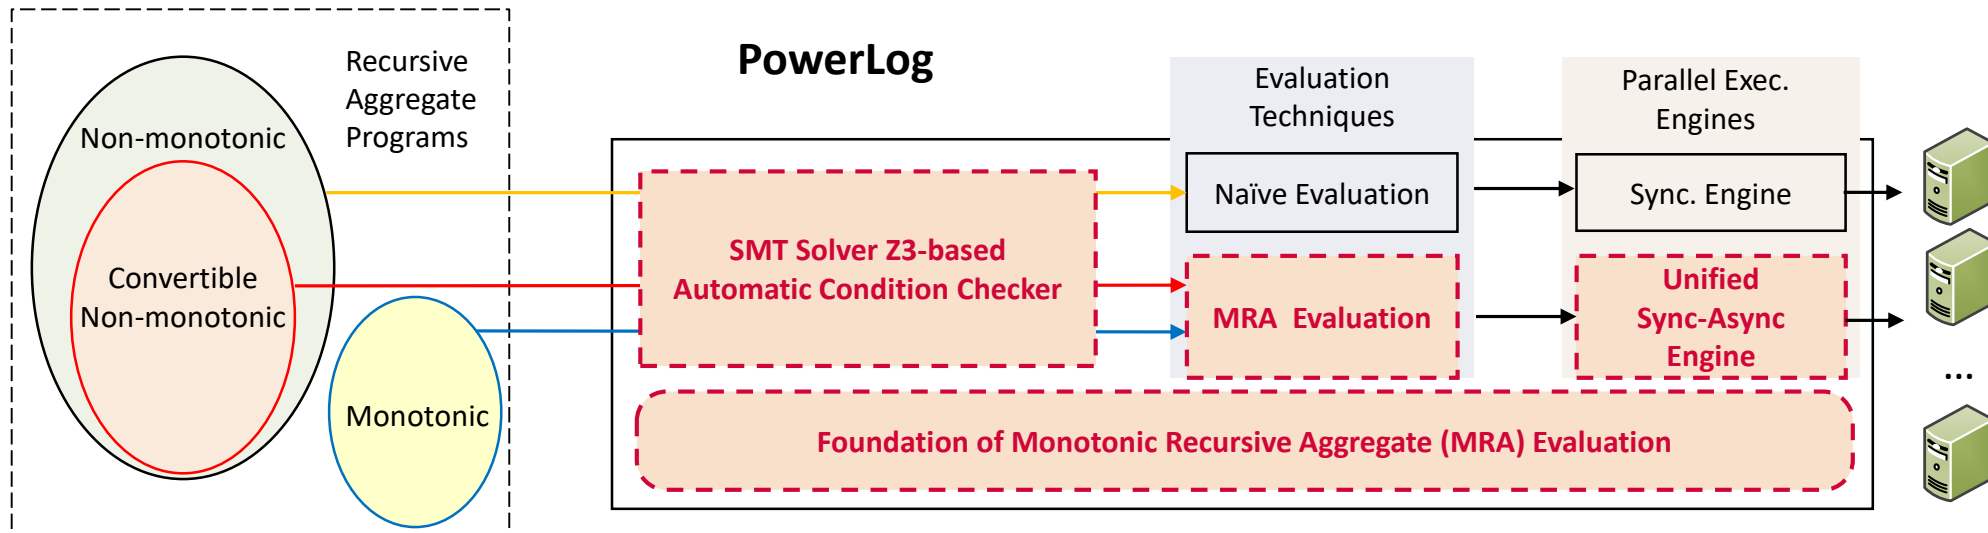
We design and implement **a Z3-based condition checker** to **automatically** check if a recursive aggregate program can satisfy the MRA conditions.

- (1) monotonic programs will be executed with MRA Evaluation
- (2) convertible non-monotonic programs that pass the condition check, e.g., PageRank, will be executed with MRA Evaluation
- (3) others that cannot pass the condition check will be executed with Naïve Evaluation



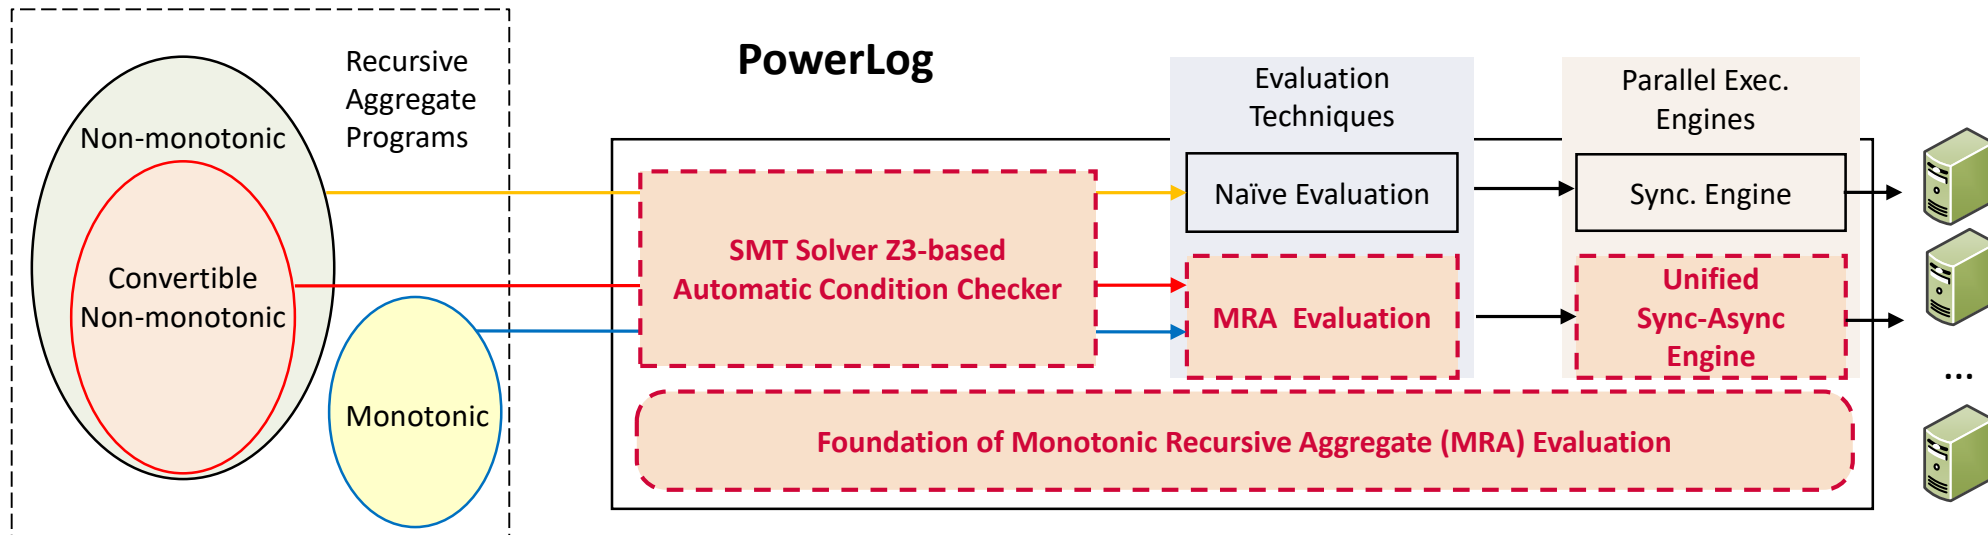
Contributions of PowerLog

We design and implement **a unified sync-async engine** to execute the programs with MRA Evaluation for high performance.



Contributions of PowerLog

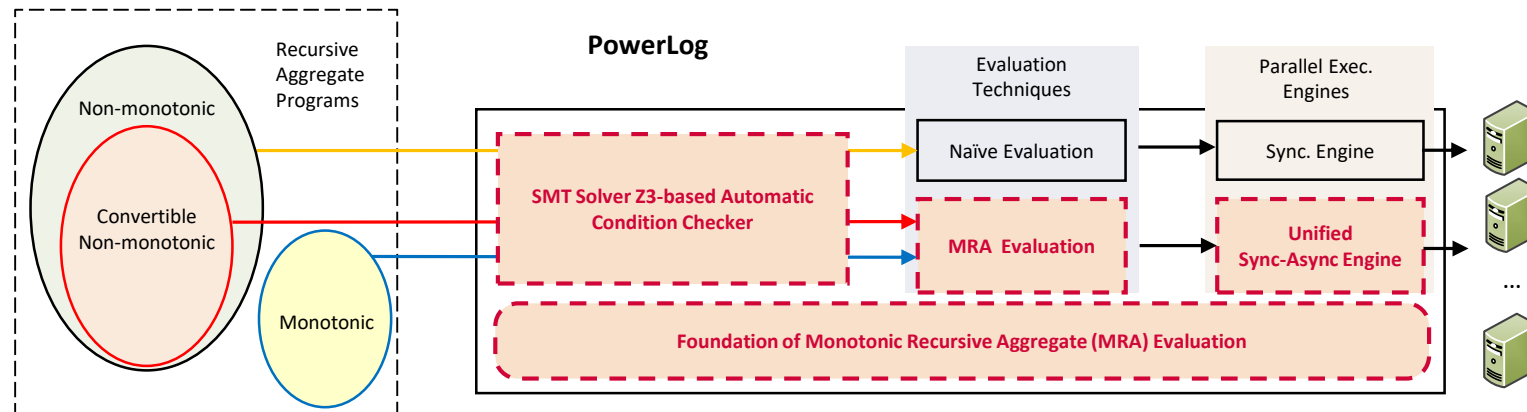
Putting all these together, we develop **PowerLog** (MRA Evaluation + a Z3-based condition checker + a unified sync-async engine).



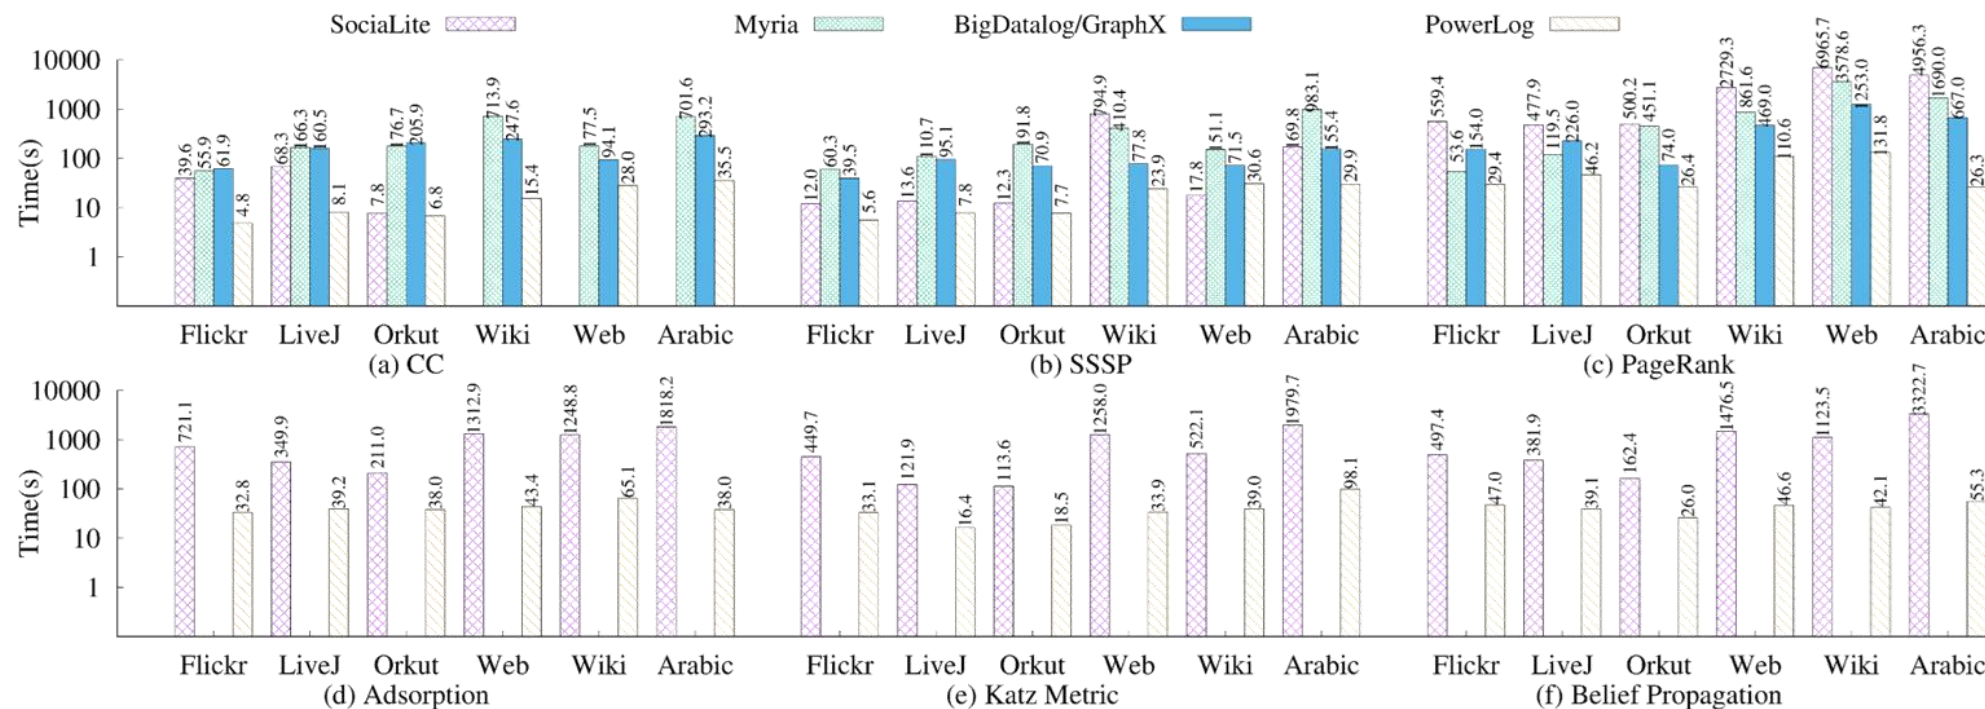
Workload Summary

By using PowerLog, we have checked **14** recursive aggregate programs. **12** of them can be executed incrementally and asynchronously but **2** cannot.

Program	MRA sat.	Aggregator	Program	MRA sat.	Aggregator
SSSP [24]	yes	min	PageRank [39]	yes	sum
CC [24]	yes	min	Adsorption [7]	yes	sum
Katz metric [21]	yes	sum	Belief Propagation [40]	yes	sum
Computing Paths in DAG [50]	yes	count	Cost [50]	yes	sum
Viterbi Algorithm [50]	yes	max	SimRank [20]	yes	sum
Lowest Common Ancestor [44]	yes	min	APSP [50]	yes	min
CommNet [52]	no	sum	GCN-Forward [22]	no	sum



Performance Summary



Compared with three representative Datalog systems (**SocialLite**, **Myria**, **BigDatalog**), PowerLog can achieve **1.1x – 188.3x** speedups on 6 programs and several real datasets.

Summary

The development of PowerLog involves both theory and system development.

- **Scope enhancement**

We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- **Automatic condition verification**

We develop a machine tool to eliminate tedious and error-prone efforts.

- **A fast execution engine**

We implement an highly optimized unified sync-async system.

Summary

The development of PowerLog involves both theory and system development.

- **Scope enhancement**

We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- **Automatic condition verification**

We develop a machine tool to eliminate tedious and error-prone efforts.

- **A fast execution engine**

We implement an highly optimized unified sync-async system.

Questions



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

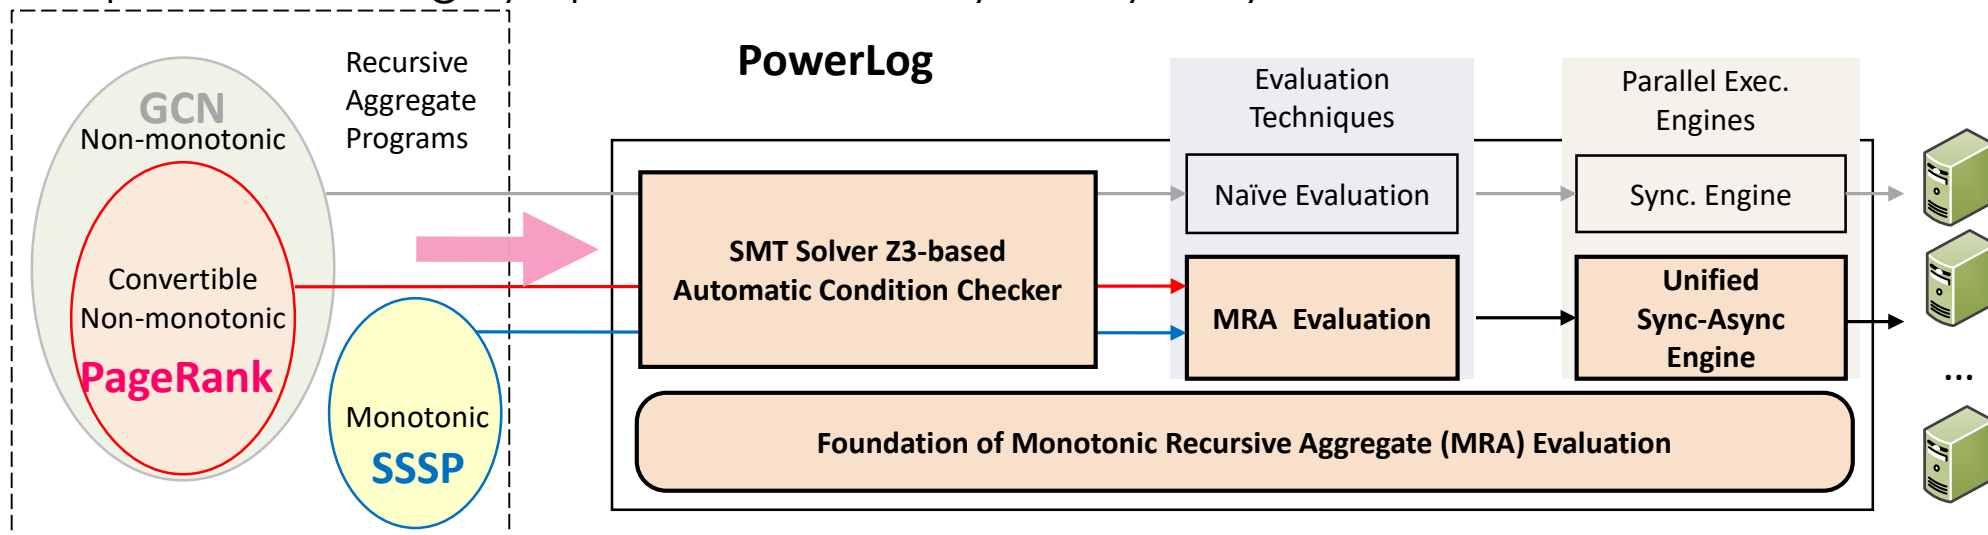
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

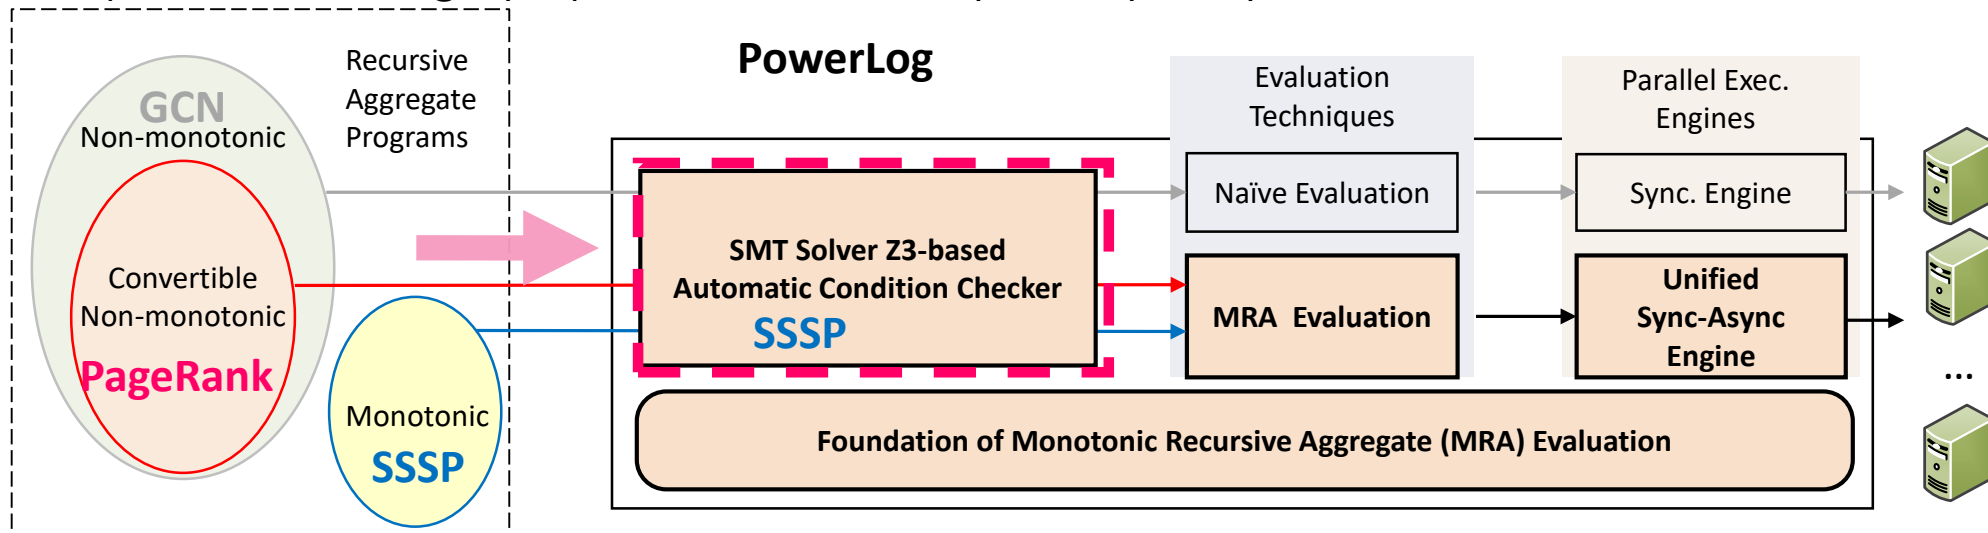
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

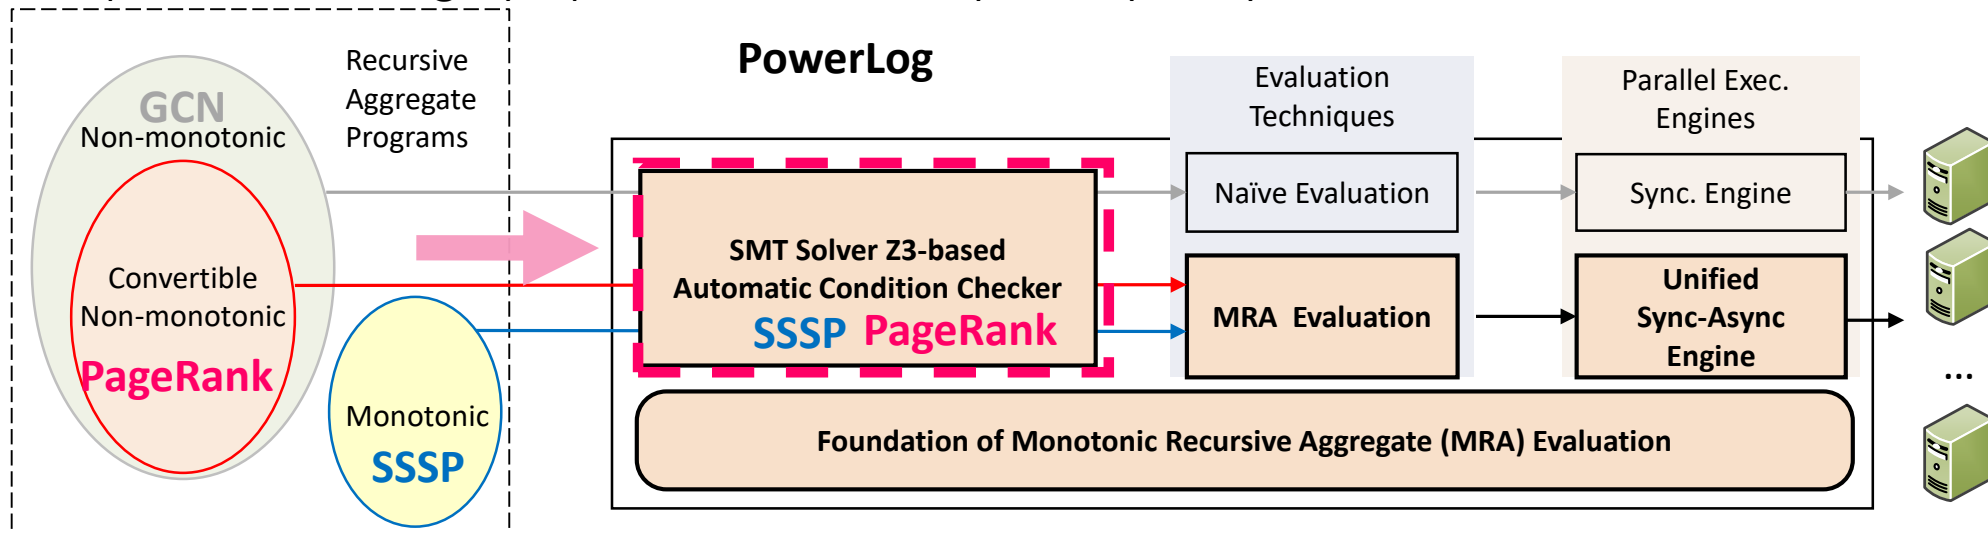
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

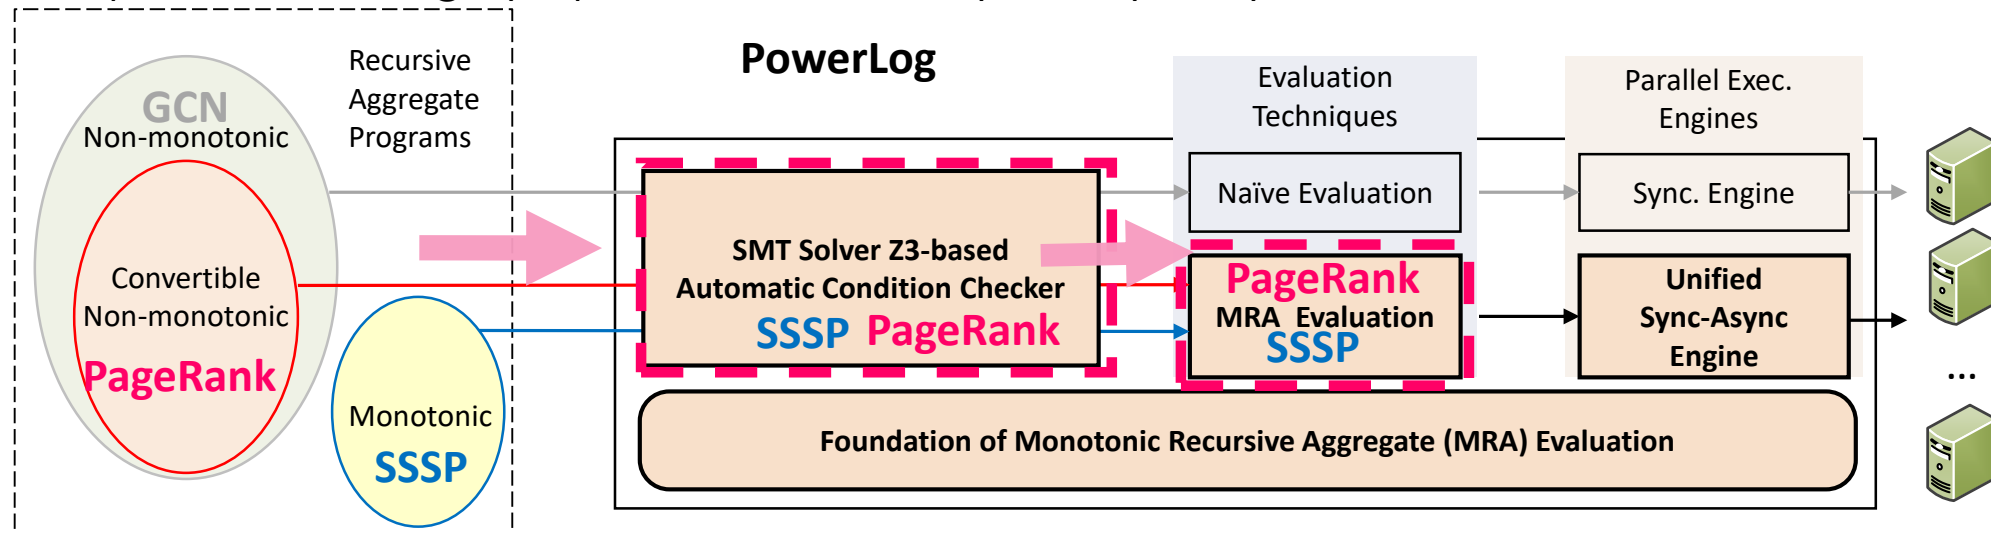
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

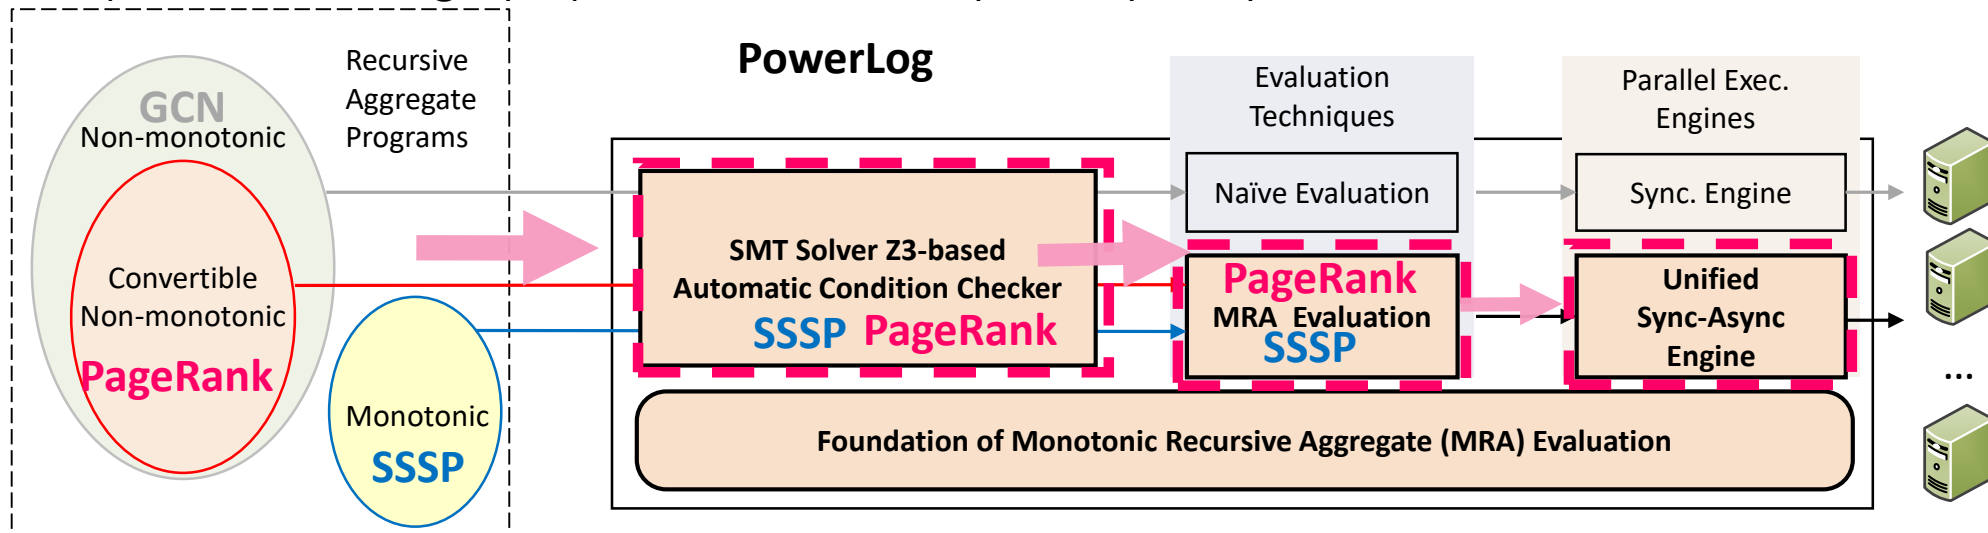
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

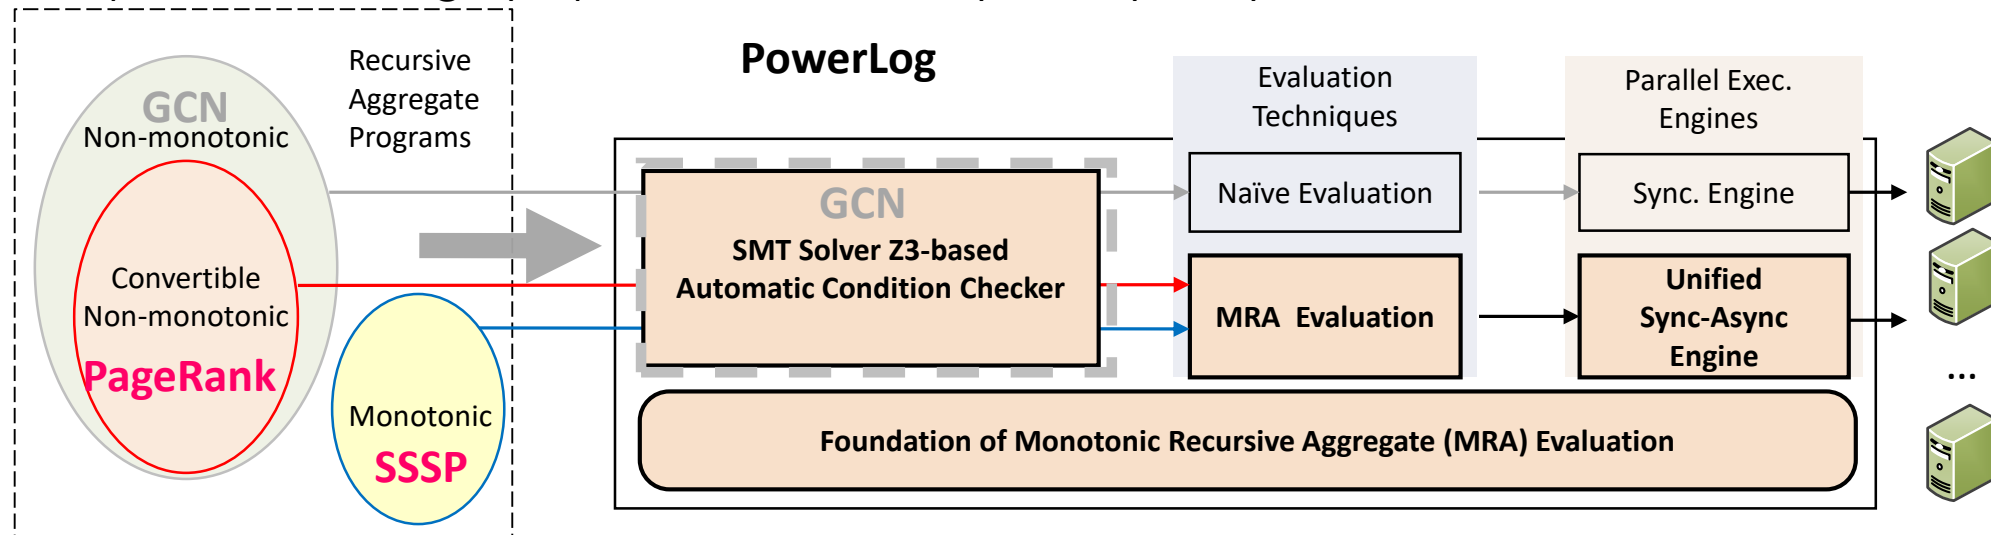
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

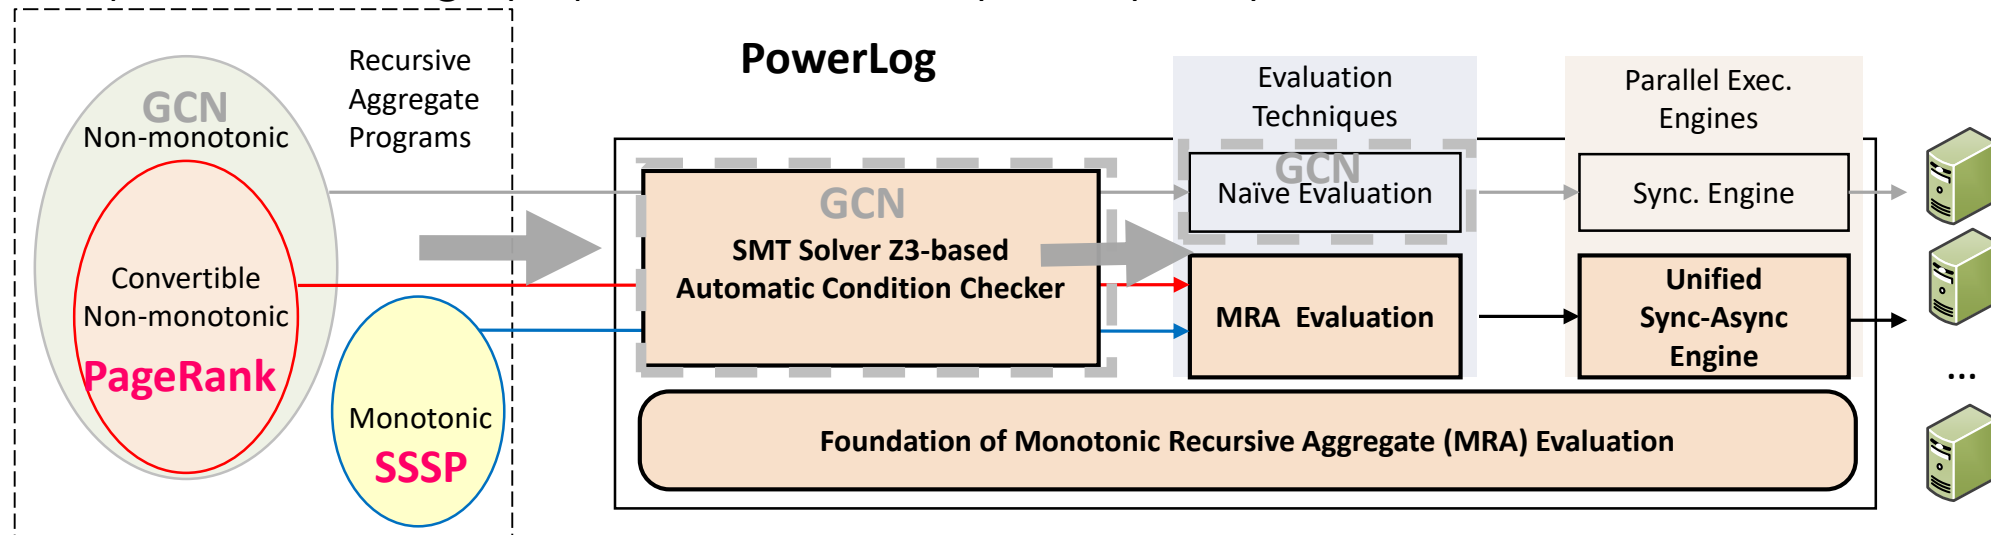
We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.



Contributions:

The development of PowerLog involves both theory and system development.

- Scope enhancement

We lay an analytical foundation to determine the conditions for monotonic and non-monotonic programs for correct execution.

- Automatic condition verification

We develop a machine tool to eliminate tedious and error-prone efforts.

- A fast execution engine

We implement an highly optimized unified sync-async system.

