

*I-DLV-sr: A Stream Reasoning System based on I-DLV**

FRANCESCO CALIMERI, MARCO MANNA, ELENA MASTRIA,
MARIA CONCETTA MORELLI, SIMONA PERRI and JESSICA ZANGARI

*Department of Mathematics and Computer Science,
University of Calabria, Rende, Italy*

(e-mails: francesco.calimeri@unical.it, marco.manna@unical.it, elena.mastria@unical.it,
maria.morelli@unical.it, simona.perri@unical.it, jessica.zangari@unical.it)

submitted 18 May 2021; revised 14 July 2021; accepted 31 July 2021

Abstract

We introduce a novel logic-based system for reasoning over data streams, which relies on a framework enabling a tight, fine-tuned interaction between *Apache Flink* and the \mathcal{I}^2 -DLV system. The architecture allows to take advantage from both the powerful distributed stream processing capabilities of *Flink* and the incremental reasoning capabilities of \mathcal{I}^2 -DLV, based on overgrounding techniques. Besides the system architecture, we illustrate the supported input language and its modeling capabilities, and discuss the results of an experimental activity aimed at assessing the viability of the approach.

KEYWORDS: stream reasoning, stream processing, overgrounding, knowledge representation and reasoning, answer set programming

1 Introduction

Stream Reasoning (SR) (Dell’Aglio *et al.* 2017) consists in the application of inference techniques to data streams. Recently, SR has been studied in several fields, and became more and more relevant in diverse application scenarios, such as IoT, Smart Cities, Emergency Management. Hence, different approaches have been proposed (Barbieri *et al.* 2010; Phuoc *et al.* 2011; Hoeksema and Kotoulas 2011; Pham *et al.* 2019) in contexts such as Complex Event Processing (CEP), Semantic Web and Knowledge Representation and Reasoning (KRR). Among declarative KRR paradigms, Answer Set Programming (ASP) (Brewka *et al.* 2011) is a well-established proposal that gained attention also outside of the academia thanks to the availability of robust and efficient implementations (Gebser *et al.* 2018). ASP is acknowledged as a particularly attractive basis for SR. Indeed, some important steps forward in this direction have been taken; we mention here the Logic for Analytic Reasoning over Streams (LARS) framework (Beck *et al.* 2018)

* This work has been partially supported by the project “MAP4ID - Multipurpose Analytics Platform 4 Industrial Data”, N. F/190138/01-03/X44 and by the Italian MIUR Ministry and the Presidency of the Council of Ministers under the project “Declarative Reasoning over Streams” under the “PRIN” 2017 call (CUP H24I17000080001, project 2017M9C25L.001).

and its implementations (Bazoobandi *et al.* 2017; Beck *et al.* 2018; Eiter *et al.* 2019; Ren *et al.* 2018) StreamRule (Mileo *et al.* 2013), C-ASP (Pham *et al.* 2019) and others (Do *et al.* 2011; Gebser *et al.* 2011). However, ASP-based stream reasoners appear not mature enough with respect to the desirable requirements for SR (Dell’Aglio *et al.* 2017). Hence, there is still room for improvements, especially when dealing with real applications: for instance, some systems are weakly usable in practice (e.g. enforce strict assumptions in input programs) and others suffer from efficiency/scalability issues.

The aim of our work is to obtain a novel, reliable ASP-based stream reasoner; besides efficiently scaling over real-world application domains, it should support a language which inherits the highly declarative nature and ease of use from ASP, while being also easily extendable with new constructs that are relevant for practical SR scenarios. In this paper we present the prototype of *I-DLV-sr*, a system that relies on the proper integration of two well-established solutions in the field of Stream Processing and ASP, respectively. *I-DLV-sr* is based on a continuous cooperation between two components: a custom designed application that leverages on *Apache Flink* (Hueske and Kalavri 2019), a powerful stream processor for efficiently managing data streams, and \mathcal{I}^2 -DLV (Ianni *et al.* 2020), an ASP grounder and a full-fledged deductive database system that enables incremental ASP evaluation via overgrounding techniques (Calimeri *et al.* 2019). Currently, the supported language basically consists in normal stratified ASP enriched with a set of constructs allowing to reason over streams. We tested the system with the aim of assessing its reliability, performance, and scalability, and also of exploring ease of modeling and reasoning capabilities. The results of a number of experiments, conducted on both real-world and synthetic domains, are encouraging, proving the viability of the approach and the robustness of the implementation.

The remainder of the paper is structured as follows. In Section 2 we describe the two main components of *I-DLV-sr*; Section 3 defines syntax and semantics of the currently supported language, and illustrates how it can be used for modeling SR problems via proper examples; in Section 4 the architecture of *I-DLV-sr* is presented and discussed in detail; Section 5 illustrates the experimental settings and the results; related works are discussed in Section 6, and, eventually, conclusions and future works are reported in Section 7.

2 *I-DLV-sr* components

I-DLV-sr mainly consists of two components: a custom *Flink*-based Java application and \mathcal{I}^2 -DLV. *Flink* is a distributed Stream Processing system for both batch and real-time stream data processing with high throughput and low latency. Applications relying on *Flink* are built by designing, using the exposed APIs, ad hoc dataflow graphs that make use of a number of different operators; besides pre-defined ones, custom operators can also be implemented. Each operator transforms one or more input data streams into a new data stream, that, in turn, can be the part of the input to subsequent operators. A dataflow graph is basically a directed acyclic graph, where nodes represent the operators, incoming arcs represent the input data streams, and outgoing arcs represent the intermediate data streams resulting from operator applications. Each dataflow graph must have one or more sources, from which the data streams originate, and at least one sink, that can either emit the final output or persistently store it. Computation is automati-

cally distributed and parallelized on the basis of dependencies among operators (Carbone et al. 2015).

Incremental I-DLV system, namely \mathcal{I}^2 -DLV (Ianni et al. 2020), is the recently presented incremental version of the deductive database system and ASP grounder *I-DLV* (Calimeri et al. 2017). It incorporates overgrounding techniques (Calimeri et al. 2019) to accommodate incremental executions over different inputs. More in detail, \mathcal{I}^2 -DLV works in a server-like mode: given a fixed input program, it remains “listening” for input facts. Every time new such facts arrive, it computes a ground program by properly updating the one resulting from previous “shots”. In other words, it automatically maintains a ground program, semantically equivalent to the input one and monotonically growing over time. It is worth noting that, just like *I-DLV*, in case of normal and stratified w.r.t. negation ASP programs, \mathcal{I}^2 -DLV computes the full semantics (i.e. returns the unique answer set). Typically, an overgrounded program, after a number of shots, converges to a propositional theory general enough to be reused together with possible future inputs, with little or no further update required. This makes overgrounding very attractive in SR contexts, as grounding activities in later iterations tend to be virtually eliminated.

3 *I-DLV-sr* language

In this section, we introduce the syntax and the semantics of *I-DLV-sr* programs and then we show the usage of *I-DLV-sr* language for Knowledge Representation via some example scenarios.

Syntax. We assume to have finite sets V , C and P consisting of *variables*, *constants* and *predicate names*, respectively; we constrain V and C to be disjoint. A *term* is either a variable in V or a constant in C . A *predicate atom* has the form $p(t_1, \dots, t_n)$, where $p \in P$ is a predicate name, t_1, \dots, t_n are terms and $n \geq 0$ is the arity of the predicate atom; a predicate atom $p()$ of arity 0 can be also denoted by p . A predicate atom is *ground* if none of its terms is a variable. We denote as G the set of all ground predicate atoms constructible from predicate names in P and constants in C . Given a *predicate atom* a , a constant $c \in C \cap \mathbb{N}^+$, a term $t \in C \cup V$ (*counting term*), and a non-empty set of numbers $D = \{d_1, \dots, d_m\} \subset \mathbb{N}$, we define three types of *streaming atoms*:

a **at least** c **in** $\{d_1, \dots, d_m\}$ a **always in** $\{d_1, \dots, d_m\}$ a **count** t **in** $\{d_1, \dots, d_m\}$

A streaming atom α (resp., **not** α) is said to be a *positive streaming literal* (resp., *negative streaming literal*), where **not** denotes *negation as failure*. A streaming literal is said to be *ground* if none of its terms is a variable. For a set L of streaming literals, $\text{preds}(L)$ denotes the set of predicates appearing in L . The following shortcuts are admitted:

- a **in** $\{d_1, \dots, d_m\}$ in place of a **at least** 1 **in** $\{d_1, \dots, d_m\}$;
- a in place of a **at least** 1 **in** $\{0\}$ (this is called “degenerate” form of a streaming literal);
- a **at most** c **in** $\{d_1, \dots, d_m\}$ in place of **not** a **at least** c' **in** $\{d_1, \dots, d_m\}$ where $c' = c + 1$.

Given a streaming atom of any type, if $D = \{n \in \mathbb{N} \mid 0 \leq n \leq w \wedge w > 0\}$ we indicate it simply as $[w]$; for example, we write a **always in** $[3]$ instead of a **always in** $\{0, 1, 2, 3\}$.

Table 1: Entailment of ground streaming literals.

α	$\Sigma \models \alpha$	$\Sigma \models \text{not } \alpha$
a at least c in $\{d_1, \dots, d_m\}$	$ \{A \in O(\Sigma, D) : a \in A\} \geq c$	$ \{A \in O(\Sigma, D) : a \in A\} < c$
a always in $\{d_1, \dots, d_m\}$	$\forall A \in O(\Sigma, D), a \in A$	$\exists A \in O(\Sigma, D) : a \notin A$
a count c in $\{d_1, \dots, d_m\}$	$ \{A \in O(\Sigma, D) : a \in A\} = c$	$ \{A \in O(\Sigma, D) : a \in A\} \neq c$

A rule can be of one out of the two forms: (1) $a :- l_1, \dots, l_b$. or (2) **#temp** $a :- l_1, \dots, l_b$, where a is a predicate atom, $b \geq 0$ and l_1, \dots, l_b represent a conjunction of streaming literals. For a rule r , we say that the *head* of r is the set $H(r) = \{a\}$, whereas the set $B(r) = \{l_1, \dots, l_b\}$ is referred to as the *body* of r . A program P is a finite set of rules; P is *flat* if all rules contain only streaming literals in the degenerate form; P is *restricted* if only rules of form (1) occur in it. We say that a rule r is *safe* if all variables appearing in $H(r)$ or in a negative streaming literal of $B(r)$ also appear in a positive streaming literal of $B(r)$. A program is safe if all its rules are safe. We require programs to be safe.

A streaming literal is said to be *harmless* if it has form a **at least** c **in** $\{d_1, \dots, d_m\}$ or a **always** **in** $\{d_1, \dots, d_m\}$; otherwise, it is said to be *non-harmless*. A program P is *stratified* if there is a partition of disjoint sets of rules $P = \Pi_1 \cup \dots \cup \Pi_k$ (called strata) such that for $i \in \{1, \dots, k\}$ both these conditions hold: (i) for each harmless literal in the body of a rule in Π_i with predicate p , $\{r \in P | H(r) = \{p(t_1, \dots, t_n)\}\} \subseteq \bigcup_{j=1}^i \Pi_j$; (ii) for each non-harmless literal in the body of a rule in Π_i with predicate p , $\{r \in P | H(r) = \{p(t_1, \dots, t_n)\}\} \subseteq \bigcup_{j=1}^{i-1} \Pi_j$. We call Π_1, \dots, Π_k a *stratification* for P and P is stratified by Π_1, \dots, Π_k . An *I-DLV-sr* program is always stratified.

Semantics. We provide next an operational semantics of *I-DLV-sr* programs. We start by introducing the notion of a stream. A *stream* Σ is a sequence of sets of ground predicate atoms $\langle S_0, \dots, S_n \rangle$ such that for $0 \leq i \leq n$, $S_i \subseteq G$. Each natural number i is called *time point*. A ground predicate atom $a \in S_i$ is true at the i -th time point. Given two streams $\Sigma = \langle S_0, \dots, S_n \rangle$ and $\Sigma' = \langle S'_0, \dots, S'_n \rangle$, $\Sigma = \Sigma'$ iff $S_i = S'_i$ for each $i \in \{0, \dots, n\}$. For a stream $\Sigma = \langle S_0, \dots, S_n \rangle$, a *backward observation* identifies ground predicate atoms that are true at some time points preceding the n -th time point. More formally, given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and a set of numbers $D \subset \mathbb{N}$, we define the *backward observation* of Σ w.r.t. D as the set $\{S_i \mid i = n - d \text{ with } d \in D \wedge i \geq 0\}$, and we denote it as $O(\Sigma, D)$. Given $w \in \mathbb{N}$, a backward observation of Σ w.r.t. $[w]$ is called *window*.

A backward observation allows to define the truth of a ground streaming literal at a given time point. Given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$, $D = \{d_1, \dots, d_m\} \subset \mathbb{N}$, $c \in C \setminus \{0\}$ and the backward observation $O(\Sigma, D)$, Table 1 reports when Σ *entails* a ground streaming atom α (denoted $\Sigma \models \alpha$) or its negation ($\Sigma \models \text{not } \alpha$). If $\Sigma \models \alpha$ ($\Sigma \models \text{not } \alpha$) we say that α is true (false) at time point n .

Example 1

In the stream $\Sigma = \langle \{a(2), b(5)\}, \{a(3), c(7)\}, \{b(5)\}, \{a(3)\} \rangle$, atom $b(5)$ is true at time points 0 and 2. For $D = \{0, 1, 3\}$, the backward observation of Σ w.r.t. D is $O(\Sigma, D) = \{\{a(3)\}, \{b(5)\}, \{a(2), b(5)\}\}$. We have that $\Sigma \models b(5)$ **at least 2 in** $\{0, 1, 3\}$. Indeed, the cardinality of the set $\{A \in O(\Sigma, D) : b(5) \in A\}$ is 2.

We define now the notions of substitution and applicability of a rule to a stream. A *substitution* σ is a mapping from the set V of variables to the set of constants C . Given a predicate atom a and a *substitution* σ , $\sigma(a)$ is the ground predicate atom obtained by replacing each occurrence of a variable v in a by $\sigma(v)$. Given a streaming literal l , $\sigma(l)$ is the ground streaming literal obtained by applying σ to the predicate atom appearing in l and to the counting term t possibly appearing in l if $t \in V$ and $\sigma(t) \neq 0$. Given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$, a rule r is *applicable* on Σ if there exists a substitution σ such that $\Sigma \models \sigma(b)$ for all $b \in B(r)$. In such a case, r is *applicable* on Σ via σ . Roughly, applicability of rules identifies new ground predicate atoms that are true at time point n ; that is, a rule r fires implying the truth of the ground predicate atom $\sigma(a)$, with $a \in H(r)$. Given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and an *I-DLV-sr* program P , a *trigger for P on Σ* is a pair $\langle r, \sigma \rangle$ where σ is a substitution and $r \in P$ is applicable on Σ via σ . An *application* of $\langle r, \sigma \rangle$ to Σ returns the stream $\Sigma' = \langle S_0, \dots, S_{n-1}, S_n \cup \sigma(a) \rangle$, with $a \in H(r)$. A trigger application is denoted as $\Sigma \langle r, \sigma \rangle \Sigma'$.

Example 2

Consider again Σ as in Example 1 and let P_1 be as follows:

```
r1: c(X) :- b(X) at least 2 in { 0, 1, 3 }.
```

Let us consider a substitution σ such that $X \mapsto 5$. The rule r_1 is applicable on Σ via σ and the application of $\langle r_1, \sigma \rangle$ to Σ returns $\Sigma' = \langle \{a(2), b(5)\}, \{a(3), c(7)\}, \{b(5)\}, \{a(3), c(5)\} \rangle$.

We therefore establish the order of application of the rules of an *I-DLV-sr* program P according to a stratification for P , and finally introduce the concept of streaming model of an *I-DLV-sr* program on a stream.

Given a stream Σ and an *I-DLV-sr* program P stratified by $\Pi_1 \dots \Pi_k$, a *stratum application* of Π_s for $s \in \{1, \dots, k\}$ on Σ is a finite sequence of streams $\Sigma_0, \dots, \Sigma_h$ with $\Sigma_0 = \Sigma$ and $h \geq 0$ such that:

- for each $0 \leq i < h$, there is a trigger $\langle r_i, \sigma_i \rangle$ for Π_s on Σ_i such that $\Sigma_i \langle r_i, \sigma_i \rangle \Sigma_{i+1}$;
- for each $0 \leq i < j < h$, if $\Sigma_i \langle r_i, \sigma_i \rangle \Sigma_{i+1}$, $\Sigma_j \langle r_j, \sigma_j \rangle \Sigma_{j+1}$ and $r_i = r_j$, then $\sigma_i \neq \sigma_j$;
- there is no trigger $\langle r, \sigma \rangle$ for Π_s on Σ_h such that $\langle r, \sigma \rangle \notin \{ \langle r_i, \sigma_i \rangle \}_{0 \leq i \leq h}$.

Intuitively, starting from a stream Σ , we apply distinct triggers, considering only different substitutions for the same rule, as long as there is a new trigger. In general, for each stratum Π_s more than one stratum application exists. However, for any two stratum applications of Π_s on Σ , their last streams coincide. This intuition is proved by the following proposition whose proof is given in the supplementary material corresponding to this paper at the TPLP archives.

Proposition 1

Given a stream Σ , an *I-DLV-sr* program P stratified by $\Pi_1 \dots \Pi_k$, a stratum Π_s with $s \in \{1, \dots, k\}$ and two stratum applications of Π_s on Σ , $\Sigma_0, \dots, \Sigma_h$ and $\Sigma'_0, \dots, \Sigma'_t$, we have that $\Sigma_h = \Sigma'_t$.

We call the last stream in a stratum application *outcome of Π_s on Σ* , denoted as *outcome*(Π_s, Σ). We now define the outcome over strata of P on Σ , obtained, starting from the outcome of the first stratum Π_1 of P and considering, one after the other, subsequent strata according to the stratification.

Given an *I-DLV-sr* program P stratified by $\Pi_1 \dots \Pi_k$ and a stream $\Sigma = \langle S_0, \dots, S_n \rangle$, we compute the sequence of streams defined as follows: $\Sigma_{\Pi_1} = \text{outcome}(\Pi_1, \Sigma)$ and $\Sigma_{\Pi_i} = \text{outcome}(\Pi_i, \Sigma_{\Pi_{i-1}})$ for $1 < i \leq k$. We define *outcome over strata* of an *I-DLV-sr* program P on the stream Σ , denoted $\mathcal{R}(P, \Sigma)$, the n -th element in the stream Σ_{Π_k} . Note that $\mathcal{R}(P, \Sigma)$ is independent from the chosen stratification.

Streaming model of I-DLV-sr restricted programs. For simplicity, we first define the notion of streaming model for programs where only rules of the form (1) appear.

Let $\Sigma = \langle S_0, \dots, S_n \rangle$ be a stream, P an *I-DLV-sr* restricted program stratified by $\Pi_1 \dots \Pi_k$. Let $\Sigma' = \langle S'_0, \dots, S'_{n-1}, S_n \rangle$ be such that $S'_0 = \mathcal{R}(P, \langle S_0 \rangle)$ and each S'_i for $i \in \{1, \dots, n-1\}$ is defined as $\mathcal{R}(P, \langle S'_0, \dots, S'_{i-1}, S_i \rangle)$. We define $\mathcal{R}(P, \Sigma')$ as *streaming model of P on Σ* .

Intuitively, the streaming model of P for a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ is the set of the ground predicate atoms derived as true at the n -th time point. These latter come from the evaluation of P over a stream that iteratively accumulates the evaluations at previous time points.

Example 3

Let $\Sigma = \langle \{b(5)\}, \{c(7)\} \rangle$ and P_2 be the following program:

```

r1: c(X) :- b(X).      r2: d(X) :- c(X) in [1].
    
```

Note that P_2 is stratified by the single stratum $\{r_1, r_2\}$; thus, the outcome over strata of P_2 on a stream Σ , $\mathcal{R}(P_2, \Sigma)$, is the n -th element in the stream $\text{outcome}(P_2, \Sigma)$. We have $S'_0 = \mathcal{R}(P_2, \langle S_0 \rangle) = \mathcal{R}(P_2, \langle \{b(5)\} \rangle) = \{b(5), c(5), d(5)\}$ and $\Sigma' = \langle S'_0, S_1 \rangle = \langle \{b(5), c(5), d(5)\}, \{c(7)\} \rangle$. Since $\mathcal{R}(P_2, \Sigma') = \mathcal{R}(P_2, \langle \{b(5), c(5), d(5)\}, \{c(7)\} \rangle) = \{c(7), d(7), d(5)\}$, the streaming model of P_2 for Σ is the set $\{c(7), d(7), d(5)\}$.

Streaming model of I-DLV-sr programs. We now introduce the streaming model also for *I-DLV-sr* non-restricted programs. To this aim, we need the notion of persistent outcome and a slightly different definition of streaming model that takes into account the eventual presence of **#temp** in rule heads. Such marking stands for temporary: indeed the truth of heads of rules in the form (2) is in a sense, limited to the current time point. We are given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and an *I-DLV-sr* program P stratified by $\Pi_1 \dots \Pi_k$. Let $P_{(1)} = \{r \in P \mid r \text{ is of the form (1)}\}$, we define *persistent outcome over strata* of P on Σ the set of ground predicate atoms $\{a \in \mathcal{R}(P, \Sigma) \mid a \in S_n \vee (\exists r \in P_{(1)} \text{ s.t. } r \text{ is applicable to } \Sigma_{\Pi_k} \text{ via } \sigma \wedge \sigma(h) = a \wedge h \in H(r))\}$ and we denote it with $\mathcal{P}(P, \Sigma)$. Basically, for an input stream $\Sigma = \langle S_0, \dots, S_n \rangle$, $\mathcal{P}(P, \Sigma)$ identifies the set of ground predicate atoms in $\mathcal{R}(P, \Sigma)$ deriving from heads of rules of the form (1) or in S_n .

Given a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and an *I-DLV-sr* program P stratified by $\Pi_1 \dots \Pi_k$, let Σ' be the stream $\Sigma' = \langle S'_0, \dots, S'_{n-1}, S_n \rangle$ where $S'_0 = \mathcal{P}(P, \langle S_0 \rangle)$ and each S'_i for $i \in \{1, \dots, n-1\}$ is defined as $\mathcal{P}(P, \langle S'_0, \dots, S'_{i-1}, S_i \rangle)$. We define $\mathcal{R}(P, \Sigma')$ as *streaming model of P on Σ* .

The definition of the streaming model of an *I-DLV-sr* non-restricted program differs from the one given for restricted programs as it considers the persistent outcome over strata for time points up to $n-1$ and the outcome over strata for n including heads of rules in the form (2).

Example 4

Let Σ be the stream of the Example 3 and P_3 be the following program:

```
r1: #temp c(X) :- b(X).          r2: d(X) :- c(X) in [1].
```

Since the first rule of P_3 is of the form (2), we have that $S'_0 = \mathcal{P}(P_3, \langle S_0 \rangle) = \mathcal{P}(P_3, \langle \{b(5)\} \rangle) = \{b(5), d(5)\}$ and $\Sigma' = \langle S'_0, S_1 \rangle = \langle \{b(5), d(5)\}, \{c(7)\} \rangle$. The streaming model of P_3 on Σ is the set $\mathcal{R}(P_3, \Sigma') = \mathcal{R}(P_3, \langle \{b(5), d(5)\}, \{c(7)\} \rangle) = \{c(7), d(7)\}$.

Modeling SR Problems. We next show how practical problems that require reasoning over streams can be modeled via the *I-DLV-sr* language. Note that, besides streaming literals, *I-DLV-sr* also supports built-in atoms and aggregate literals as defined in the ASP-Core-2 standard (Calimeri et al. 2020); currently, the only restriction is that aggregate elements cannot feature (non-degenerate) streaming literals. For the sake of readability, we omitted their description in the language syntax; the following simple program briefly shows their usage: in a scenario where the total number of cars passed now or in the previous 20 time points must be computed, rule r_1 counts the number of passes for each car, while r_2 uses an ASP aggregate to compute the total.

```
r1: carPassing(C, N) :- car(C) count N in [20].
r2: tot(T) :- #sum{N, C: carPassing(C, N)}=T.
```

Photo-voltaic system. Suppose that we need to build an Intelligent Monitoring System (IMS) for a photo-voltaic system (PVS) to promptly detect malfunctions. Without going into technical details, for the sake of simplicity, let us suppose that the PVS is composed by a grid of interconnected panels via solar cables and each panel is provided with a sensor that measures the amount of energy produced and continuously sends data to the IMS. Each panel continuously produces energy to be transferred to a Central Energy Accumulator (CEA), directly or via a path between neighbor panels across the grid. Let us assume that a time point corresponds to a second. A panel is working if it is known to have produced an amount of energy greater than a given threshold within the last 4 seconds, and, in addition, if it is reachable by the CEA (i.e. there exists a path of working panels linking it to the CEA). If some unreachable working panels have been detected more than 2 times in the last 3 seconds, an alert must be raised for an identified malfunction. Furthermore, the IMS must request a maintenance intervention if the failure is continuously observed for 5 seconds. This scenario can be modeled via an *I-DLV-sr* program as reported next.

```
r1: workingPanel(P) :- energyDelivered(P, W) at least 1 in [4],
                        energyThreshold(Et), W>=Et.
r2: reachable(cea, P2) :- link(cea, P2), workingPanel(P2).
r3: reachable(P1, P3) :- reachable(P1, P2), link(P2, P3), workingPanel(P3).
r4: unlinked :- workingPanel(P), not reachable(cea, P).
r5: regularFunctioning :- unlinked at most 2 in [3].
r6: alert :- not regularFunctioning.
r7: callMaintenance :- alert always in [5].
```

The predicates `link` and `energyThreshold` represent the PVS configuration and the threshold defining a working panel for the given reasoning interval; these data do not change during such interval. The predicate `energyDelivered` represents the amount of power produced by each panel; at each time point within the reasoning interval, the current values are sent to the IMS, thus producing a stream. Rule r_1 defines a panel

as working if it transmitted an amount of energy greater than the threshold at least once in the interval from the current time point to the previous 4 consecutive ones, that is, the last 4 seconds. Rules r_2 and r_3 recursively define the set of reachable working panels starting from the CEA. Rule r_4 detects if there are unlinked working panels and r_5 defines proper functioning by checking that the atom `unlinked` appeared no more than two times in the last 3 seconds. Eventually, r_6 raises up an alert if there is not a regular functioning and r_7 asks to call the maintenance if an alert has been raised in all the last 5 seconds.

Underground Traffic Monitoring. Let us imagine we want to build a monitoring system for the underground trains in the city of Milan. In this example, we suppose that a time point corresponds to a minute; given a station, passengers expect to see a train stopping every 3–6 minutes, during the rush hours. The following *I-DLV-sr* program models a simple control system that warns passengers when this regularity is broken to several extents (i.e. mild/grave irregularity). This would allow, for example, to properly mark each station on a map of a mobile/web app.

```

r1: irregular    :- train_pass, train_pass at least 1 in {1, 2}.
r2: irregular    :- not train_pass in [6].
r3: #temp num_anomalies(X) :- irregular count X in [30].
r4: mild_alert   :- num_anomalies(X), X>2, X<=5.
r5: severe_alert :- num_anomalies(X), X>5.

```

Rules r_1 and r_2 are used to detect irregularities: trains arriving too early or too late. Rule r_3 counts the number of irregular situations in the last half an hour, producing an instance of the `num_anomalies`; r_4 and r_5 raise the proper warning. Note that `num_anomalies` is intended to be an auxiliary predicate whose instance is used for determining irregularities only watching at the current time point; thus, it is marked as **#temp**, so that its instance contributes to the current streaming model but it is no longer considered in next time points.

4 I-DLV-sr architecture

The system takes as input an *I-DLV-sr* program P and a stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and iteratively builds a stream $O = \langle O_0, \dots, O_n \rangle$ such that each O_t contains the result of the evaluation at the time point t . Figure 1 depicts the high-level system architecture that consists of three main modules: *Execution Manager*, *Stream Manager* and *Subprogram Manager*, all making an ad hoc use of *Flink* APIs. Details about each module are given below.

4.1 Execution Manager

The *Execution Manager* is in charge of setting up the evaluation of an *I-DLV-sr* program P . First, it computes from P a flat program P' and determines a mapping τ from streaming atoms to predicate atoms. Moreover, it divides P into subprograms taking into account dependencies among all rules in P caused by streaming atoms so that each subprogram can be separately processed, limiting the interplay between *Flink* and \mathcal{I}^2 -DLV. On the basis of such program splitting, it constructs the *Flink* dataflow graph

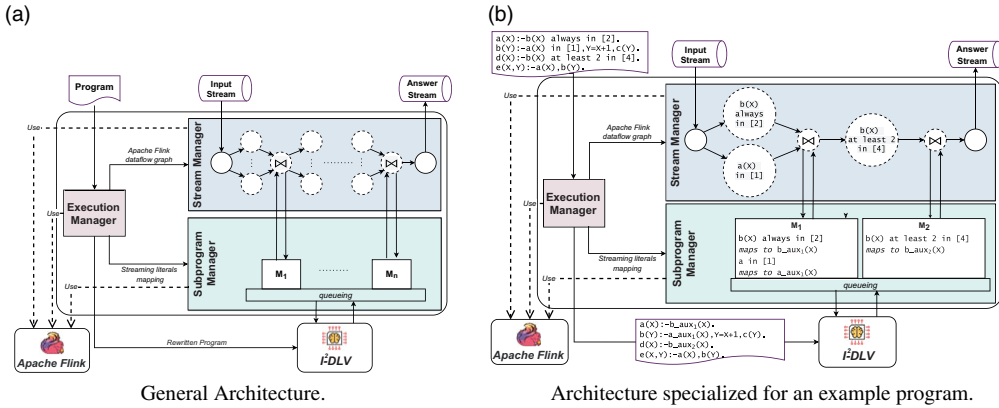


Figure 1: System architecture.

(cf. Section 2). Finally, it provides: the *Stream Manager* with the dataflow graph, the *Subprogram Manager* with τ , and I^2 -DLV with P' . The *Execution Manager* tasks are detailed below.

Program Rewriting. The *Execution Manager* produces a flat program P' from P . Each streaming atom $p(t_1, \dots, t_n) \diamond \text{in } \{d_1, \dots, d_m\}$, with $\diamond \in \{\text{at least } c, \text{ at most } c, \text{ always}, \text{count } c\}$, which is not in the degenerate form and so that c is a constant, is replaced by $p'(t_1, \dots, t_n)$; each streaming atom $p(t_1, \dots, t_n) \text{count } X \text{ in } \{d_1, \dots, d_m\}$ where X is a variable, is replaced by $p'(t_1, \dots, t_n, X)$; in both cases, p' is a fresh predicate name. Such replacements are stored in a mapping τ . Ground instances of fresh predicates will be generated by the *Subprogram Manager* on the basis of the evaluation of corresponding streaming atoms performed by the *Stream Manager*.

Example 5

Below are reported a program P_4 (left) and the flat program P'_4 obtained by rewriting it (right).

$r_1 : a(X) :- b(X) \text{ always in } [2].$ $r_2 : b(Y) :- a(X) \text{ in } [1], Y=X+1, c(Y).$ $r_3 : d(X) :- b(X) \text{ at least } 2 \text{ in } [4].$ $r_4 : e(X, Y) :- a(X), b(Y).$	$r'_1 : a(X) :- b_aux1(X).$ $r'_2 : b(Y) :- a_aux1(X), Y=X+1, c(Y).$ $r'_3 : d(X) :- b_aux2(X).$ $r'_4 : e(X, Y) :- a(X), b(Y).$
---	--

The mapping of the replacements is $\tau = \{b(X) \text{ always in } [2] \mapsto b_aux1(X), a(X) \text{ in } [1] \mapsto a_aux1(X), b(X) \text{ at least } 2 \text{ in } [4] \mapsto b_aux2(X)\}$.

Program Splitting and Processing Order. Given an I -DLV- sr program P , we build a directed labeled graph called *Stream Dependency Graph* and denoted G_P^{SD} , whose nodes are the predicates in rule heads of P and for each pair of nodes p and q , there is an arc (p, q) if there exists a rule $r \in P$ such that $preds(H(r)) = \{q\}$ and $p \in preds(B(r))$. In case p occurs in a streaming literal that is not in the degenerate form, then the arc is labeled with “ $<$ ”; no label is added otherwise. On the basis of G_P^{SD} , we define an additional directed labeled graph, called *Stream Component Graph* and denoted G_P^{SC} : its nodes are the strongly connected components of G_P^{SD} (i.e. sets of predicates), and there is an arc from component A to component B if there exists an arc (p, q) in G_P^{SD} such that $p \in A$ and $q \in B$. Each arc (A, B) in G_P^{SC} is labeled with “ $<$ ” if there exists at least one arc (p, q) labeled with “ $<$ ” in G_P^{SD} such that $p \in A$ and $q \in B$; no label is added otherwise.




Figure 2: Stream dependency (left) and component graphs (right) of Example 5.

Relying on G_P^{SC} , rules of P that can be processed together are grouped, and a processing ordering among groups of rules is established. For any pair of nodes A and B of G_P^{SC} , we say that A precedes B (denoted $A \prec B$) if there exists a path in G_P^{SC} from A to B containing at least one arc labeled with “ $<$ ”; we say A is alongside B (denoted $A \approx B$) otherwise. We identify an ordering C_1, \dots, C_n of the nodes of G_P^{SC} such that, for each $i < j$, it does not hold that $C_j \prec C_i$. According to such an ordering C_1, \dots, C_n , we collect the predicates of some consecutive nodes in order to form a macro-node as follows: for each pair of nodes C_i and C_k , with $i \leq k \leq n$ we construct the macro-node $M = \bigcup_j C_j$ with $i \leq j \leq k$ such that either $i = k$ (i.e. the macro-node actually consists of a single node), or for each $j \neq k$, $C_j \approx C_{j+1}$ (i.e. the macro-node consists of nodes that are all alongside each other). We can therefore define an ordered sequence of all the maximal macro-nodes such that for any pair of macro-nodes $M_1 = \bigcup_{j_1} C_{j_1}$ with $i_1 \leq j_1 \leq k_1 \leq n$ and $M_2 = \bigcup_{j_2} C_{j_2}$ with $i_2 \leq j_2 \leq k_2$, if M_1 precedes M_2 in the sequence, then $k_1 < i_2$. This ordered sequence induces both a splitting of P into subprograms and a processing order for them: for each macro-node M , the subprogram P_M is the set of all rules r of P such that $preds(H(r)) \subseteq M$, that is, the predicate occurring in $H(r)$ belongs to M ; the processing order for the subprograms coincides with the ordering in the sequence of the corresponding macro-nodes. Note that for an *I-DLV-sr* program several orderings of the nodes of G_P^{SC} might exist, in general and therefore, different processing orders might be obtained. In the following, we will refer to one of these processing orders as Ord_P .

Eventually, a rule r within a subprogram P_M is *streaming-recursive* if there is a cycle in the Stream Dependency Graph G_P^{SD} among two nodes p and q where $p \in preds(B(r))$ and $q \in preds(H(r))$ such that there is at least an arc labeled with “ $<$ ”.

Example 6

Let us consider the program P_4 of Example 5. Figure 2 shows the Stream Dependency and Component Graphs of P_4 . We can observe that $G_{P_4}^{SD}$ has three strongly connected components: $\{a, b\}$, $\{d\}$ and $\{e\}$. According to $G_{P_4}^{SC}$, we have: $\{a, b\} \prec \{d\}$, $\{a, b\} \approx \{e\}$, $\{e\} \approx \{a, b\}$, $\{d\} \approx \{e\}$ and $\{e\} \approx \{d\}$. As a consequence, there are four orderings: $o_1 = \{\{e\}, \{a, b\}, \{d\}\}$; $o_2 = \{\{a, b\}, \{e\}, \{d\}\}$; $o_3 = \{\{a, b\}, \{d\}, \{e\}\}$; $o_4 = \{\{a, b\}, \{e\}, \{d\}\}$. Let us focus on o_1 . The maximal macro-nodes that we can build are $\{a, b, e\}$ and $\{d\}$ and the resulting program splitting is: $P_{4\{a,b,e\}} = \{r_1, r_2, r_4\}$, $P_{4\{d\}} = \{r_3\}$. Moreover, r_1 and r_2 in $P_{4\{a,b,e\}}$ are streaming-recursive. Note that, o_2 would induce the same maximal macro-nodes and in turn, the same splitting of o_1 and in both cases, $\{r_1, r_2, r_4\}$ are processed before $\{r_3\}$.

Flink dataflow graph creation. As described above, given an *I-DLV-sr* program, the *Execution Manager* splits it into subprograms and computes a processing order Ord_P . Then, according to Ord_P , it builds an *Flink* dataflow graph by adding one or more nodes for each subprogram and for each of its streaming atoms. As shown in Figure 1, there are two kinds of nodes represented as empty or \bowtie -filled dashed circles. Empty circles

correspond to operators needed for evaluating a streaming atom; by means of *Flink* APIs, we defined some operators in order to implement the semantics of *in*, *always*, *count*, *at least*, and *at most*. \bowtie -filled circles correspond to custom operators intended to join results of the evaluation of all streaming atoms in a subprogram and to set up the evaluation of such subprogram. We implemented within such operators, the low-level transformations needed for filtering and joining outcomes of linked circles. Empty circles are connected with directed arcs to the \bowtie -filled one associated to the subprogram in which the corresponding streaming atoms occur. \bowtie -filled circles are linked to the empty ones associated to the streaming atoms appearing in the subsequent subprogram in Ord_P .

4.2 Stream Manager and Subprogram Manager

As depicted in Figure 1, the *Stream Manager* executes the scheduled operators in the *Flink* dataflow graph and interacts with the *Subprogram Manager* that, in turn, mediates communications with \mathcal{I}^2 -DLV responsible for the evaluation of the flat program P' , which in fact, is always a normal and stratified w.r.t. negation ASP program. Let us assume to be at the time point n . The *Stream Manager* handles the input stream $\Sigma = \langle S_0, \dots, S_n \rangle$ and the stream $O = \langle O_0, \dots, O_n \rangle$ where O_0, \dots, O_{n-1} collect the results obtained from the evaluations in previous time points and O_n , initially empty, is filled in while subprograms are evaluated.

More in detail, let P_M be the first subprogram in Ord_P . The *Stream Manager* executes the operators in the dataflow graph corresponding to the streaming atoms in P_M over Σ and O . Then, it passes to the *Subprogram Manager* S_n , O_n and the set of holding streaming atoms H , that is, those evaluated as true, at the time point n . The *Subprogram Manager*: (i) according to the mapping τ , generates from H the set H' of holding ground instances of the fresh predicates introduced by the *Execution Manager*; (ii) provides $S_n \cup O_n \cup H'$ as input to \mathcal{I}^2 -DLV that incrementally evaluates P' over such input and computes the corresponding unique answer set A ; (iii) receives A (i.e. a set of ground (predicate) atoms) and adds A to O_n . The *Stream Manager* takes control back, evaluates the next nodes in the dataflow graph and interacting with the *Subprogram Manager*, performs the evaluation of the next subprogram in Ord_P . The process continues until all subprograms are evaluated according to Ord_P ; the output at n is $S_n \cup O_n$. Note that the smaller it is the number of subprograms, the smaller it is the number of iterations between \mathcal{I}^2 -DLV and *Flink*: indeed, maximal macro-nodes are limiting such interplay.

In case a subprogram contains some streaming-recursive rules R , the *Stream Manager* has to repeatedly evaluate the streaming atoms in R , that is, every time O_n is enriched by the *Subprogram Manager*. When no more new ground (predicate) atoms can be derived for the predicates in the heads of R , the above described process goes on with the next subprogram in Ord_P .

Note that, the set $S_n \cup O_n$ actually coincides with the streaming model of P on Σ . Indeed, the streaming model is independent from the stratification of choice (see Section 3); thus, if we consider the stratification where each stratum is the smallest possible, then each subprogram consists of rules that might belong to one or more consecutive strata; when evaluating a subprogram P_M , \mathcal{I}^2 -DLV by design takes into account dependencies among the strata and processes rules in P_M accordingly. Thus, the output computed for P_M coincides with the output that we would obtain when evaluating, one after the other, the strata contained in P_M as well as in all subprograms preceding P_M , according to the

definition of streaming model. In other words, the output computed for P_M coincides with the streaming model of the *I-DLV-sr* program constituted by all subprograms up to P_M on Σ . It is worth noting that the *Stream Manager* implements a queuing mechanism for the subtasks that have to be performed for evaluating P on Σ exploiting *Flink* APIs. Thanks to this, *I-DLV-sr* is able to manage backpressure (Karimov *et al.* 2018) guaranteeing no loss of data, and thus the correctness of the output.

Example 7

Consider again P_4 of Example 5. Figure 1 (b) illustrates the evaluation process of P_4 by referring to the general architecture. We assume to be at the n -th time point and that $\Sigma = \langle S_0, \dots, S_n \rangle$ is the input stream, while $O = \langle O_0, \dots, O_n \rangle$ with $O_n = \emptyset$ is the stream iteratively built so far.

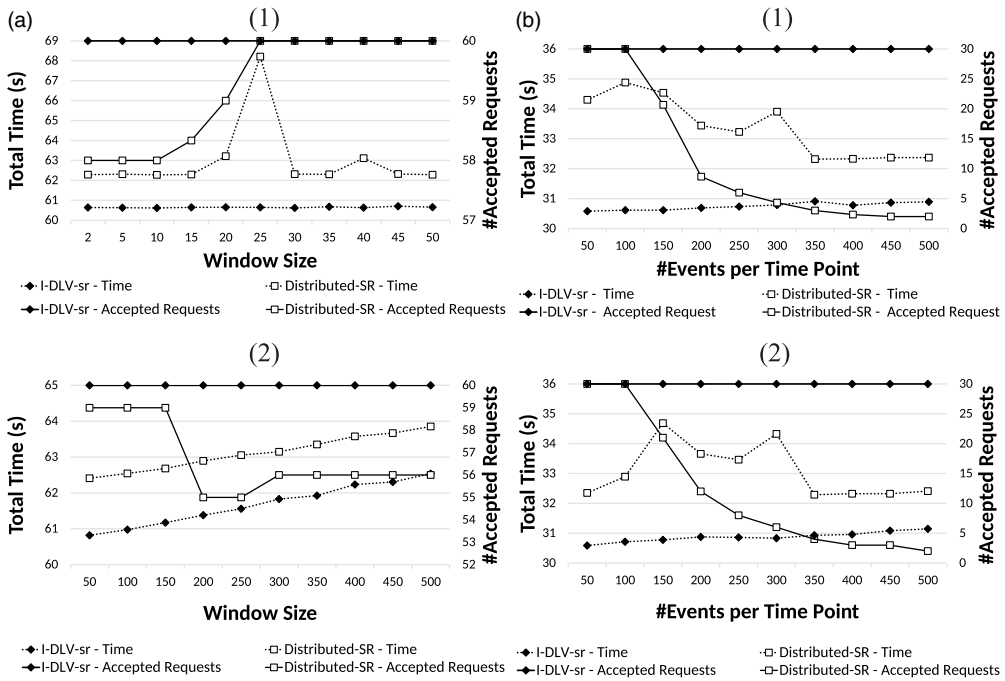
The *Execution Manager* provides \mathcal{I}^2 -DLV with the flat program P'_4 obtained by rewriting P_4 , splits P_4 and identifies the processing order for its subprograms. Suppose as shown in Example 6 that the splitting is $P_{4\{a,b,e\}} = \{r_1, r_2, r_4\}$ and $P_{4\{d\}} = \{r_3\}$ and that $P_{4\{a,b,e\}}$ precedes $P_{4\{d\}}$. The *Execution Manager* accordingly creates the dataflow graph and passes it to the *Stream Manager*.

The dataflow graph contains three empty dashed circles for the operators needed for evaluating streaming atoms in $P_{4\{a,b,e\}}$ and $P_{4\{d\}}$, and two \bowtie -filled ones that receive results from the linked empty circles and forward them to *Subprogram Manager*. More in detail, the *Stream Manager* first evaluates the empty dashed circles relative to $P_{4\{a,b,e\}}$ and then, the *Subprogram Manager* is required to take into account $P_{4\{a,b,e\}}$ hence, producing the ground instances for `a_aux1` and `b_aux1` and properly invoking \mathcal{I}^2 -DLV. The result received by the *Subprogram Manager* is forwarded back to the *Stream Manager* that updates O_n . Such loop between *Stream Manager* and *Subprogram Manager* for processing of $P_{4\{a,b,e\}}$ continues until nothing new can be inferred for the predicates $\{a, b, e\}$ and only after $P_{4\{d\}}$ is processed.

5 Experimental evaluation

In order to assess reliability and performance of *I-DLV-sr*, we carried out an experimental activity over different SR problems and settings. All experiments have been performed on a NUMA machine equipped with two 2.8GHz AMD Opteron 6320 CPUs, with 16 cores and 128GB of RAM. For the sake of reproducibility, both the system and the whole set of experiments are available at <https://demacs-unical.github.io/I-DLV-sr>.

We conducted two kind of analysis. First, we wanted to compare the performance of *I-DLV-sr* with some other available logic-based stream reasoner. A number of implementations are available (see Section 6); however, they considerably differ in syntax and/or semantics, and also from an architectural/implementative point of view; thus, a fair comparison is rather difficult. Given the distributed nature of *I-DLV-sr*, we compared it to Distributed-SR (Eiter *et al.* 2019), which is the most recent LARS-based implementation, supports a large set of features and relies on a distributed architecture. The latest available version has been executed. For the former analysis, we considered two benchmarks: *Content Caching* and *Heavy Join*. *Content Caching* (Beck *et al.* 2017; Eiter *et al.* 2019) is a real-world benchmark that requires to manage the caching policy of a video content over an incoming stream that describes the evolving popularity level



Content Caching

Heavy Join

Figure 3: Results on *Content Caching* and *Heavy Join*.

of the content. Besides the original problem (Beck et al. 2017), we considered a slightly different version that deals with more than one event per time point (we refer to a true atom at a time point in the stream as an “event”). Therefore, the encoding is adapted to handle more than one video content, and the incoming stream contains a number of instances, representing the popularity levels for each content, ranging from 50 to 500. *Heavy Join* is an artificial problem conceived in order to test scalability. It consists of the single rule: $a(X, Y) :- b(X, Z) \text{ in } [w], c(Z, Y) \text{ in } [w]$. where w is 2 or 20, depending on the experiment. The input streams feature an equal number of instances of predicates b and c whose total ranges from 50 to 500. In addition, in order to test *I-DLV-sr* specific features, we conducted a performance analysis on the problem introduced in Section 3, namely *Photo-voltaic System*, over grids of increasing size ranging from 20×20 panels with 11,970 links up to 30×30 panels with 60,682 links. In this data-intensive domain, we performed a further analysis to check the advantages of incrementality on *I-DLV-sr* performance.

Results are reported in Figures 3 and 4; in each plot, a marker in a line corresponds to the average of the results of three executions; in each execution, the tested system received inputs for 60 time points for *Content Caching* and *Photo-voltaic System* and 30 for *Heavy Join*. In the comparison with Distributed-SR, input arrives at a frequency fixed to 1 time point per second, while in the *Photo-voltaic System* experiments, frequency varies. Our analysis focuses on three measures: **total time**, **#accepted requests** and **latency**: **total time** represents the total elapsed time for each execution, excepting the initial time spent for the set up; **#accepted requests** is the number of accepted incoming requests computed by checking the system logs, and counting the number of time points

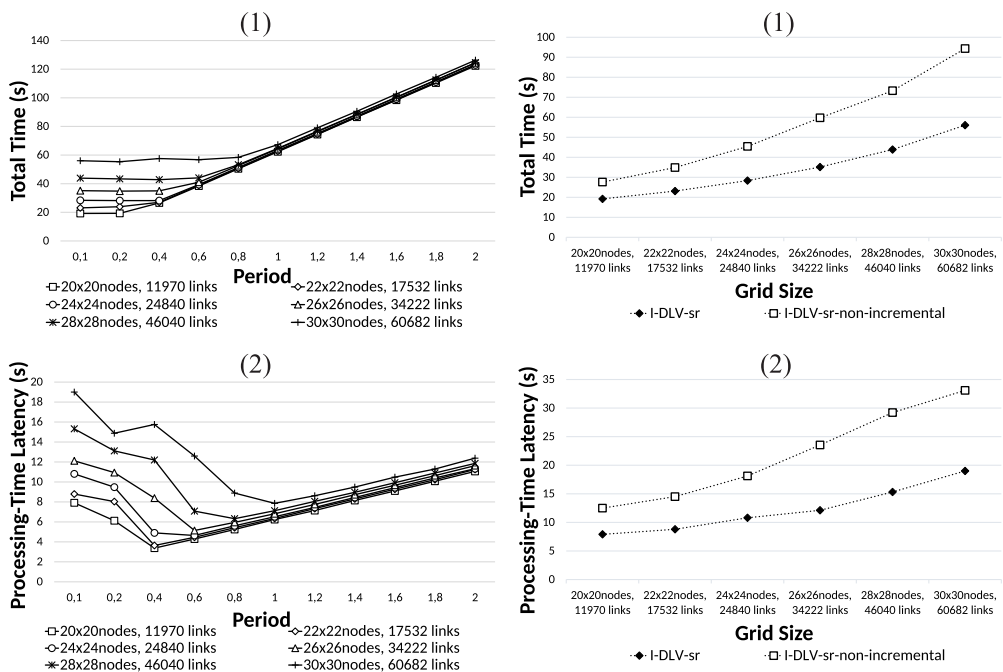


Figure 4: Results on Photo-voltaic System.

whose corresponding input is read by the system; **latency** is the processing-time latency, that is, the interval between the time at which the system receives the input relative to a time point and the time at which the system returns the corresponding output.

The left column of Figure 3 refers to *Content Caching* tests. Plot (a.1) is related to experiments on the original encoding, where each time point features a single incoming event; it reports **total time** (left y-axis) and **#accepted requests** (right y-axis), while window sizes vary from 2s to 50s (x-axis). The two systems appear to have similar behaviors w.r.t. **total time**, slightly above 60s on average, that is the minimum amount of time needed for waiting all the incoming 60 requests (recall that here the frequency is of a time point per second). The picture changes when looking at the right y-axis: while *I-DLV-sr* accepts and correctly returns the expected output for all the 60 requests, Distributed-SR fails in handling them all. Plot (a.2) refers to the modified encoding mentioned above, with window sizes set to 5 seconds. For each execution, the graph plots the two measures considered above; on the x-axis, the number of events per time point ranges from 50 to 500. The two systems show similar trends in total times, that grow along with the number of events for time points. As in the previous case, a loss of incoming requests is reported for Distributed-SR, becoming more evident for greater numbers of events.

Results for *Heavy Join* are displayed in the right column of Figure 3; the analysis is the same of Plot (a.2). We considered window size of 2s in Plots (b.1) and 20s in (b.2), respectively. The two plots depict similar results. In terms of **total time**, it is more evident w.r.t. the *Content Caching* case that *I-DLV-sr* performs better than Distributed-SR, and it maintains the same trend in both cases, while Distributed-SR exhibits a fluctuating behavior. As for **#accepted requests**, Distributed-SR “misses” a relevant

number thereof: the loss is limited when the number of events per time point is small, but rapidly grows with the number of events for time point, so that in the end almost all incoming requests are ignored.

The left column of Figure 4 depicts the results on the *Photo-voltaic System* benchmark where *I-DLV-sr* has been tested over different grids with increasing sizes, one size per line. This is a quite expensive domain as rules at lines 2–3 encode a reachability task among working grid panels that vary over time. Plots (a.1) and (a.2) respectively report **total time** and **latency** when on the x -axis, the period of incoming requests varies ranging from 0.1s to 2s. In this case, **#accepted requests** is not reported as the system correctly processes all the 60 requests. As for **total time**, the six lines in the plot show almost the same trend: constant up to a certain period P , and then linearly growing. Note that, for greater grid sizes, also P is greater: this is expected, and can be explained by observing **latency** in Plot (a.2). Indeed, for each line in the plot, **latency** is smaller when the period is close to its P value beyond which we observe the so-called sustainable throughput (Karimov et al. 2018). This is because for period values smaller than P , *I-DLV-sr* is asked to process requests that come more often, so it starts to enqueue pending requests. When the period is greater than P , instead, *I-DLV-sr* is able to consider a request as soon as it arrives, thus no queuing is needed, and the greater the period is w.r.t. P , the greater it is the idle time between the complete processing of a request and the incoming of the next one. Periods close to P are ideal, as idle time is close to 0 and no queuing occurs. The right column of Figure 4 compares the performance of two different versions of *I-DLV-sr* versions: one relies on the incremental \mathcal{I}^2 -DLV system, while the other relies on the non-incremental \mathcal{I} -DLV engine (v. 1.1.6). Plots (b.1) and (b.2) report **total time** and **latency**, respectively, of the two tested versions of *I-DLV-sr*, when on the x -axis the grid sizes vary; the period of incoming requests is set to 0.1s. We observe that the version relying on the incremental evaluator keeps **total time** and **latency** smaller than the other; the gap between the two lines becomes more and more evident as the grid size increases, suggesting that *I-DLV-sr* scales better when relying on the incremental-based version. Intuitively, the gain is attributable to the incremental computation of the transitive closure of the reachable relation over working panels; indeed, the adopted overgrounding technique maintains and reuses all the previously computed ones, rather than recomputing from scratch all the connections. We conclude by mentioning that the same analysis about incrementality has been also done for *Content Caching* and *Heavy Join*; however, as these domains are not data-intensive, the evaluation time per rule is so small that recomputing everything from scratch each time is not that expensive. Results are reported in the supplementary material corresponding to this paper at the TPLP archives and show that the incremental-based version behaves as the non-incremental one without suffering from potential overheads.

6 Related works

SR has been subject of a number of researches over the latest years. Nevertheless, there are no standardized formalisms nor techniques for SR to date, making the comparison among approaches relying on different semantics and technologies rather difficult. Apart several relevant solutions stemmed in the semantic-web context (Barbieri et al. 2010; Phuoc et al. 2011; Pham et al. 2019; Hoeksema and Kotoulas 2011), the proposals that

relate most with ours are the ones based on ASP. With this respect, one well-established is LARS: Logic for Analytic Reasoning over Streams (Beck *et al.* 2018), a formal framework enriching ASP with temporal modalities and window operators. LARS theoretically consists of a full-fledged non-monotonic formalism for reasoning over streams; indeed, the full language is computationally intractable. Current implementations, such as Laser (Bazoobandi *et al.* 2017), Ticker (Beck *et al.* 2018), and its recent distributed version (Eiter *et al.* 2019), support smaller, yet practically relevant fragments. For performance reasons, Laser handles only negation-stratified and stream-stratified programs, that is, recursion is not supported if it involves negation or windows (Beck *et al.* 2015). Ticker comes with two evaluation modes: one makes use of the state-of-the-art ASP system clingo (Gebser *et al.* 2019) as back-end, and is intended for stratified programs, that is, programs having only a single model; the other mode uses incremental truth maintenance techniques under ASP semantics, and, in case of multiple solutions, computes and maintains one single model randomly chosen. Distributed-SR is an additional version of Ticker, recently released; in order to increase the throughput, it implements an interval-based semantics of LARS that relies on Ticker as internal engine; it supports distributed computation, at the price of disabling the support for recursion through window operators. All Ticker versions require any variable appearing in the scope of a window atom to be “guarded” by some standard atom including it; such variables are grounded upfront in a so-called pre-grounding phase.

I-DLV-sr supports the ASP fragment stratified w.r.t. negation, which is extended with streaming literals over temporal intervals: recursion involving streaming literals is freely allowed as well. The implementation is designed for supporting incremental evaluation of *I-DLV-sr* programs, thanks to the integration with \mathcal{I}^2 -DLV, and parallel/distributed computation, thanks to the integration with *Flink*. From the one hand, incrementality allows to efficiently evaluate logic programs at each time point and to avoid the need for pre-grounding; also, apart from “canonical” safety, no restriction is required for variables in streaming literals. On the other hand, parallelism and distribution pave the way to the efficient evaluation of streaming atoms. During execution, each operator in the *Flink* dataflow can have one or more operator subtasks, that are independent from each other; hence, they can be executed in different threads and possibly on different machines or containers. In case of recursive subprograms, distribution is limited by evaluating the involved streaming atoms via single-threaded, undistributed operators.

A distributed stream reasoner that has in common with *I-DLV-sr* the usage of a stream processor is BigSR (Ren *et al.* 2018); it is released in two versions built on top of the state-of-the-art stream processors Spark Streaming (<https://spark.apache.org>) and *Flink*, respectively. BigSR implements internal reasoning algorithms to compute the semantics, and significantly differs from *I-DLV-sr* on the supported language, the input format and the provided features. In particular, it is oriented toward ontology-based reasoning and accepts RDF input streams. Moreover, input programs fall in the positive fragment of plain LARS and can only contain the *in* window operator. In addition, depending on the stream processor, further limitations are required. The version relying on Spark Streaming accepts stratified programs that can be recursive but have only global windows, that is, the rules must share the same window operator. The version built on top of *Flink*, instead, accepts non-recursive stratified programs, but with global windows at rule scope, that is, all literals in a rule must share the same window operator. In this

latter version, such limitations are due to the BigSR usage of the multi-core/distributed nature of *Flink*, that makes the handling of synchronization of clocks, task progress and window trigger mechanisms difficult. Furthermore, *I-DLV-sr* and BigSR differs for the adopted notion of time, as the former relies on *event time* notion while the latter on the *ingestion time* one (Hueske and Kalavri 2019).

When designing *I-DLV-sr*, Spark Streaming and Apache Storm (Toshniwal et al. 2014) were evaluated as alternative stream processors. Spark Streaming was excluded as, when receiving live input streams, it divides the data into batches, which are lately processed by the Spark engine to generate the final stream in batches. On the other hand, *Flink* does not require the a-priori creation of batches and works in real-time, record by record, rather than batch by batch. Apache Storm can handle data processing record by record but, it does not provide the event time processing and does not natively offers the exactly-once semantics as instead *Flink* does. This semantics ensures to *I-DLV-sr* that each incoming event affects the final outcome exactly once and, even in case of a machine or software failure, there is no data duplication nor unprocessed data.

7 Conclusions

We presented *I-DLV-sr*, an ASP-based stream reasoner relying on a tight interaction between \mathcal{T}^2 -DLV and a *Flink* application, that in the experiments showed good performance and scalability. *I-DLV-sr* is easily extendable by design; hence, we plan to add the support to additional language constructs while extending tests over new real-world domains. Furthermore, we plan to study proper means to extend the language for the management of *noise* and *incompleteness*, and further move toward a more complete SR reasoner (Dell’Aglia et al. 2017).

Competing interests declaration: The authors declare none.

Supplementary material

To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/S147106842100034X>.

References

- BARBIERI, D. F., BRAGA, D., CERI, S., VALLE, E. D. AND GROSSNIKLAUS, M. 2010. C-SPARQL: a continuous query language for RDF data streams. *International Journal of Semantic Computing* 4, 1, 3–25.
- BAZOUBANDI, H. R., BECK, H. AND URBANI, J. 2017. Expressive stream reasoning with laser. In *International Semantic Web Conference (1)*. LNCS, vol. 10587. Springer, 87–103.
- BECK, H., BIERBAUMER, B., DAO-TRAN, M., EITER, T., HELLWAGNER, H. AND SCHEKOTIHIN, K. 2017. Stream reasoning-based control of caching strategies in CCN routers. In *IEEE International Conference on Communications, ICC 2017, Paris, France, May 21–25, 2017*. IEEE, 1–6.
- BECK, H., DAO-TRAN, M. AND EITER, T. 2015. Answer update for rule-based stream reasoning. In *IJCAI*. AAAI Press, 2741–2747.

- BECK, H., DAO-TRAN, M. AND EITER, T. 2018. LARS: A logic-based framework for analytic reasoning over streams. *Artificial Intelligence* 261, 16–70.
- BECK, H., DAO-TRAN, M., EITER, T. AND FOLIE, C. 2018. Stream reasoning with LARS. *Künstliche Intell.* 32, 2–3, 193–195.
- BREWKA, G., EITER, T. AND TRUSZCZYNSKI, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54, 12, 92–103.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., MARATEA, M., RICCA, F. AND SCHAUB, T. 2020. Asp-core-2 input language format. *TPLP* 20, 2, 294–309.
- CALIMERI, F., FUSCÀ, D., PERRI, S. AND ZANGARI, J. 2017. I-DLV: the new intelligent grounder of DLV. *Intelligenza Artificiale* 11, 1, 5–20.
- CALIMERI, F., IANNI, G., PACENZA, F., PERRI, S. AND ZANGARI, J. 2019. Incremental answer set programming with overgrounding. *Theory and Practice of Logic Programming* 19, 5–6 (Sep), 957973.
- CARBONE, P., KATSIFODIMOS, A., EWEN, S., MARKL, V., HARIDI, S. AND TZOUMAS, K. 2015. Apache flinkTM: Stream and batch processing in a single engine. *IEEE Data Engineering Bulletin* 38, 4, 28–38.
- DELL’AGLIO, D., VALLE, E. D., VAN HARMELEN, F. AND BERNSTEIN, A. 2017. Stream reasoning: A survey and outlook. *Data Science* 1, 1-2, 59–83.
- DO, T. M., LOKE, S. W. AND LIU, F. 2011. Answer set programming for stream reasoning. In *Canadian Conference on AI*. LNCS, vol. 6657. Springer, 104–109.
- EITER, T., OGRIS, P. AND SCHEKOTIHN, K. 2019. A distributed approach to LARS stream reasoning (system paper). *Theory and Practice of Logic Programming* 19, 5-6, 974–989.
- GEBSER, M., GROTE, T., KAMINSKI, R. AND SCHAUB, T. 2011. Reactive answer set programming. In *LPNMR*. LNCS, vol. 6645. Springer, 54–66.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2019. Multi-shot ASP solving with clingo. *TPLP* 19, 1, 27–82.
- GEBSER, M., LEONE, N., MARATEA, M., PERRI, S., RICCA, F. AND SCHAUB, T. 2018. Evaluation techniques and systems for answer set programming: a survey. In *IJCAI*. ijcai.org, 5450–5456.
- HOEKSEMA, J. AND KOTOULAS, S. 2011. High-performance distributed stream reasoning using s4. In *Ordering Workshop at ISWC*.
- HUESKE, F. AND KALAVRI, V. 2019. *Stream Processing with Apache Flink: Fundamentals, Implementation, and Operation of Streaming Applications*. O’Reilly Media, Incorporated.
- IANNI, G., PACENZA, F. AND ZANGARI, J. 2020. Incremental maintenance of overgrounded logic programs with tailored simplifications. *Theory and Practice of Logic Programming* 20, 5, 719–734.
- KARIMOV, J., RABL, T., KATSIFODIMOS, A., SAMAREV, R., HEISKANEN, H. AND MARKL, V. 2018. Benchmarking distributed stream data processing systems. In *34th IEEE International Conference on Data Engineering, ICDE 2018, Paris, France, April 16-19, 2018*. IEEE Computer Society, 1507–1518.
- MILEO, A., ABDELRAHMAN, A., POLICARPIO, S. AND HAUSWIRTH, M. 2013. Streamrule: A non-monotonic stream reasoning system for the semantic web. In *RR*. LNCS, vol. 7994. Springer, 247–252.
- PHAM, T., ALI, M. I. AND MILEO, A. 2019. C-ASP: continuous asp-based reasoning over RDF streams. In *LPNMR*. LNCS, vol. 11481. Springer, 45–50.
- PHUOC, D. L., DAO-TRAN, M., PARREIRA, J. X. AND HAUSWIRTH, M. 2011. A native and adaptive approach for unified processing of linked streams and linked data. In *International Semantic Web Conference (1)*. LNCS, vol. 7031. Springer, 370–388.

- REN, X., CURÉ, O., NAACKE, H. AND XIAO, G. 2018. Bigsr: real-time expressive RDF stream reasoning on modern big data platforms. In *IEEE BigData*. IEEE, 811–820.
- TOSHNIWAL, A., TANEJA, S., SHUKLA, A., RAMASAMY, K., PATEL, J. M., KULKARNI, S., JACKSON, J., GADE, K., FU, M., DONHAM, J., BHAGAT, N., MITTAL, S. AND RYABOY, D. V. 2014. Storm@twitter. In *SIGMOD Conference*. ACM, 147–156.