

Incorporating Blockchain into RDF Store at The Lightweight Edge Devices

Anh Le-Tuan, Darshan Hingu, Manfred Hauswirth and Danh Le Phuoc | Technische Universität Berlin

RDF4LED - An RDF Engine for Edge Devices

Incorporating Blockchain into RDF4Led

Why Edge Devices ???

Edge Computing - Internet of Things

- reduce network overhead.
- reduce latency for real-time applications.
- improve scalability.
- better privacy control.

Linked Data

- enable data integration of heterogeneous sources.
- enable data federation over edge nodes.

Missions


Moving semantic data processing task away from centralised cloud for the IoT

- How much semantic data on small devices ?
- How to scale data federation over small devices on edge systems ?
- How to encourage people sharing the data from their edge devices ?

RDF4LED - An RDF Engine for Edge Devices

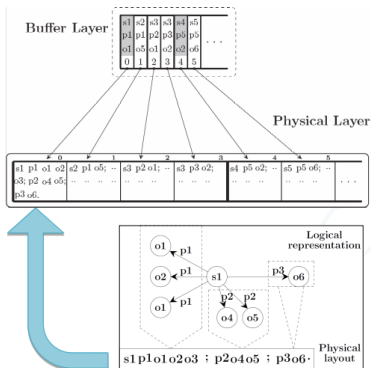
System Design

- **Physical storage:** Introduce a *two-layers index schema* for triples to adapt the *flash I/O behaviors*.
- **Buffer Manager:** Introduce a *buffer replacement policy* to reduce *the number of overwritten*.
- **Query Executor:** Introduce an *adaptive strategy* for *iterative join execution* to avoid caching *intermediated results* that *minimizing memory usage*.



RDF4LED (cont.)

Physical RDF Organisation: Two-Layers Index



Buffer Layer:

- managing triples **in memory**.
- acting as **an index** for the Physical Layer by keeping **the first triple** of a page and mapping it to **the physical address** of the page.
- organising the **atomic triples** into pages.
- grouping the **dirty pages** into writing blocks.

Physical Layer:

- Managing triples **on flash**.
- Using **molecule-based data structure** to avoid storing **repetitive values**.
- Keeping the molecules continuously sorted in pages, and keeping the pages sorted in blocks.

RDF4LED (cont.)

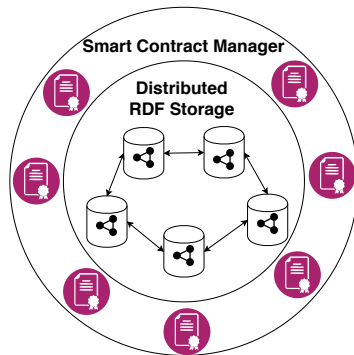
Targeted Small Devices



	Devices		
	Pi0	BBB	GII
CPU	ARM 11, 1.0 GHz, 1-core	ARM A8, 1.0 GHz, 1-core	x86 Quark, 0.4 GHz, 1-core
RAM	512 MB	512 MB	256 MB
Storage	Transcend MicroSD 16GB class 10 (40MB/s)		
OS	Raspbian	Debian 7.0	Yocto

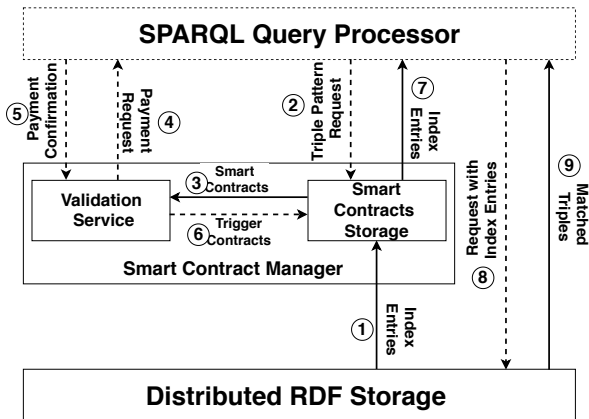
Incorporating Blockchain into RDF4Led

System Overview



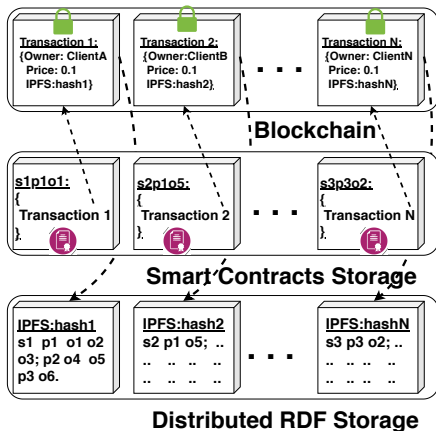
Incorporating Blockchain into RDF4Led

System Workflow



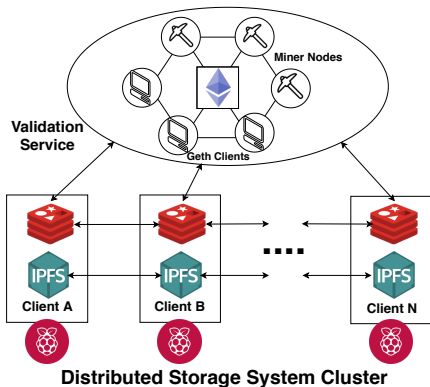

Incorporating Blockchain into RDF4Led

Physical Organisation



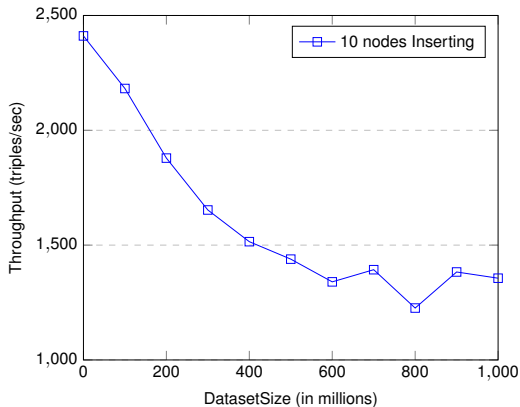
Incorporating Blockchain into RDF4Led

System Deployment



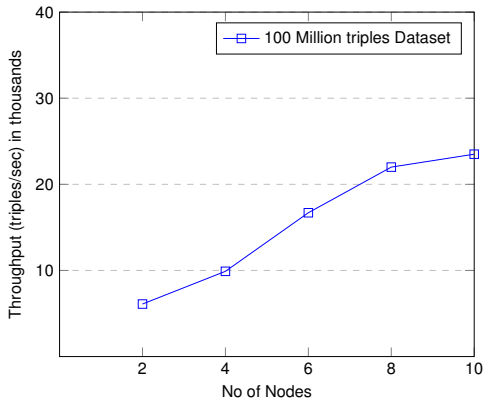
Incorporating Blockchain into RDF4Led

Evaluation - Input (1) Acc. Throughput on Static Cluster Sizes



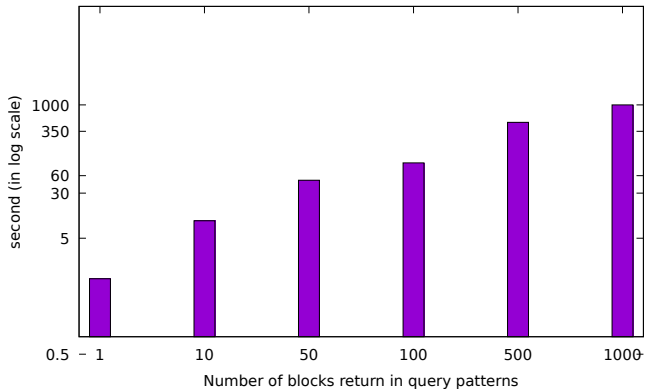
Incorporating Blockchain into RDF4Led

Evaluation - Input (2) Acc. Throughput on Varying Cluster Sizes



Incorporating Blockchain into RDF4Led

Evaluation - Query




Thank You !!!


RDF4LED (cont.)

Write Management (1): Clustering example

Buffer Layer



Persistent Layer



RDF4LED (cont.)

Write Management (2): Cache Management

Replacement policy:

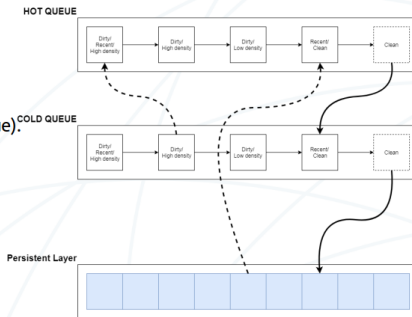
- data access frequency.
- probability of data being changed in future.

Cache organisation:

- using 2 buffer queues (hot queue/cold queue).
- the recent active blocks are kept in the hot queue.
- the inactive blocks are kept in the cold part.

Blocks are sorted in queue with following priority:

- (i) the clean/unmodified blocks;
- (ii) the number of atomic triples in a block
- (iii) the higher density blocks :
- (iv) the last recently access block.



Replacement Schema

RDF4LED (cont.)

Push-based Join Algorithm

Algorithm 1: Join propagation

```

1 Function propagate ( $\mu, \mathbb{P}$ )
   input :  $\mu$  : mapping,  $\mathbb{P}$  : set of triple query patterns
   output:  $\mu$  : mapping
2 if isEmpty( $\mathbb{P}$ ) then
3   return  $\mu$ ;
4  $p \leftarrow$  findNextPattern ( $\mu, \mathbb{P}$ );
5  $p_{key} \leftarrow$  createKey( $\mu, p$ );
6  $\mathbb{T} \leftarrow$  indexScan( $p_{key}$ );
7  $\mathbb{P}' \leftarrow \mathbb{P} \setminus \{p\}$ ;
8 for  $t \in \mathbb{T}$  do
9    $\mu \leftarrow$  bindMapping ( $t, p$ );
10  propagate ( $\mu, \mathbb{P}'$ );
11  $\mu \leftarrow$  resetMapping ( $t, p$ );

```

Algorithm 2: Find the next triple pattern


```

1 Function findNextPattern ( $\mu, \mathbb{P}$ )
   input :  $\mu$  : mapping,  $\mathbb{P}$  : set of triple query patterns
   output:  $P$  : triple query pattern
2  $p_{next} \leftarrow$  null;
3  $s_{min} \leftarrow$  IntegerMax;
4 for  $p \in \mathbb{P}$  do
5   if isShared( $\mu, p$ ) then
6      $p_{key} \leftarrow$  createKey( $\mu, p$ );
7      $\mathbb{I} \leftarrow$  indexLookup( $p_{key}$ );
8      $s \leftarrow$  sizeOf( $\mathbb{I}$ );
9     if  $s < s_{min}$  then
10       $s_{min} \leftarrow$   $s$ ;
11       $p_{next} \leftarrow$   $p$ ;
12 return  $p_{next}$ ;

```

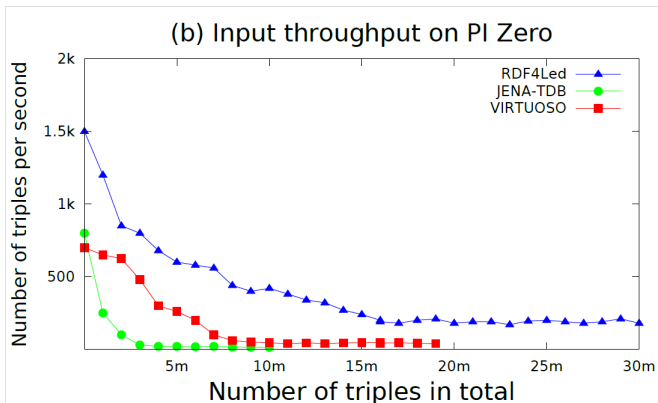
RDF4LED (cont.)

Evaluation Results - Input (1)



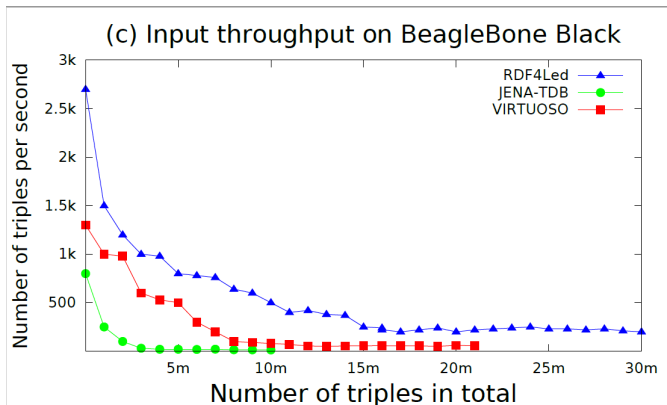
RDF4LED (cont.)

Evaluation Results - Input (2)




RDF4LED (cont.)

Evaluation Results - Input (3)



RDF4LED (cont.)

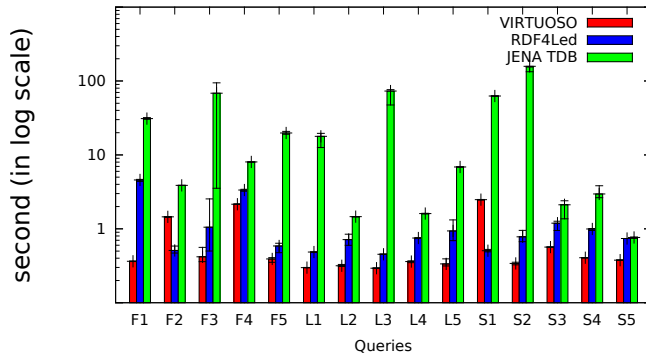
Evaluation Results - Query (1)



RDF4LED (cont.)

Evaluation Results - Query (2)


(b) Query response time on Raspberry Pi Zero



RDF4LED (cont.)


Evaluation Results - Query (3)

(c) Query response time on Beagle Bone Black



RDF4LED (cont.)


Evaluation Results - Query (4)



RDF4LED (cont.)

Evaluation Results - Memory (1)

(a) Memory consumption of inserting of VIRTUOSO, JENA TDB, RDF4Led



RDF4LED (cont.)

Evaluation Results - Memory (2)

