

Big Data Analytics

Lucas Rego Drumond

Information Systems and Machine Learning Lab (ISMLL)
Institute of Computer Science
University of Hildesheim, Germany

GraphLab

Outline

1. Introduction
2. GraphLab
 - 2.1 GraphLab Data Model
 - 2.2 Update Function
 - 2.3 Sync Mechanism
 - 2.4 PageRank Example
 - 2.5 Execution Model and Data Consistency

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example

2.5 Execution Model and Data Consistency

Overview

Part III

Machine Learning Algorithms

Part II

Large Scale Computational Models


Part I

Distributed Database


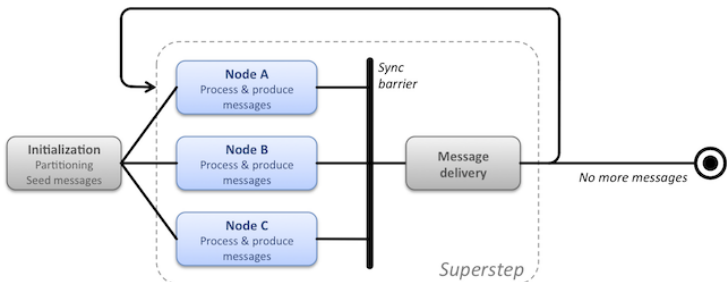
Distributed File System

MapReduce - Review

1. Each mapper transforms a set key-value pairs into a list of output keys and intermediate value pairs
2. all intermediate values are grouped according to their output keys
3. each reducer receives all the intermediate values associated with a given keys
4. each reducer associates one final value to each key



Pregel - Review



Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example

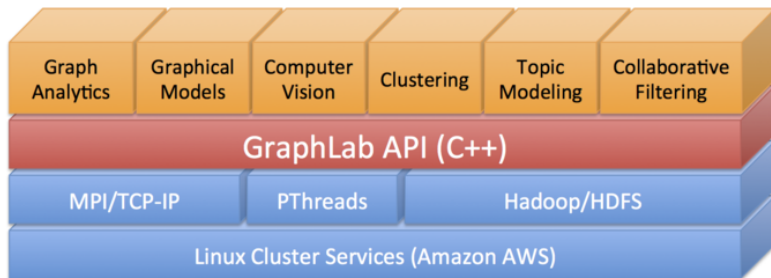
2.5 Execution Model and Data Consistency

GraphLab

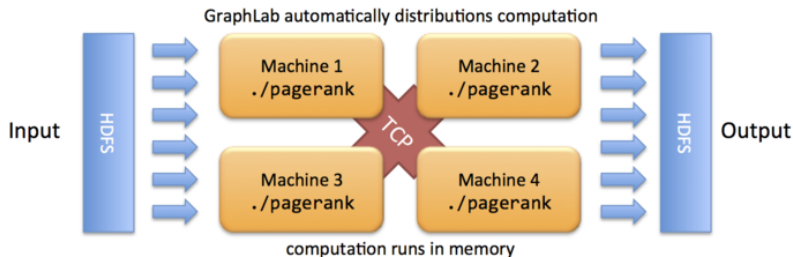
- ▶ A new framework for parallel machine learning
- ▶ Represents data as a directed graph
- ▶ Blocks of data stored in vertices and edges
- ▶ Globally shared data table
- ▶ Shared memory (GraphLab 1.0) distributed (GraphLab 2.1)

Available for download at: <http://graphlab.org>

GraphLab Software Stack



GraphLab System Overview



GraphLab Abstraction

The **GraphLab abstraction** consists of the following main components:

- ▶ A **Data Graph** provides a high level representation for data and complex computational dependencies
- ▶ Configurable **consistency models** for automatically guaranteeing data consistency
- ▶ **Scheduling Primitives** able to express iterative parallel algorithms
- ▶ **User defined computation** consisting of:
 - ▶ **Update Functions**: local computation
 - ▶ **Sync Mechanism**: global aggregation

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example

2.5 Execution Model and Data Consistency

Graphlab Data Model

Consists of two parts:

1. A directed data graph:

$$G := (V, E)$$

where V is the set of vertices and $E \subseteq V \times V$ the directed edges

2. A globally shared data table:

$$\mathbf{T}[\text{Key}] \rightarrow \text{Value}$$

which maps keys to arbitrary blocks of data

Graphlab Data Model

Users can associate with each graph vertice and edge:

- ▶ Arbitrary **data blocks**
- ▶ Program or model **parameters**

Some notation:

Data associated with a vertex v : D_v

Data associated with an edge (u, v) : $D_{u,v}$

Set of all *outbound* edges from v : $(v \rightarrow *)$

Set of all *inbound* edges from v : $(* \rightarrow v)$

Neighboring vertices of v : $\mathcal{N}_v := \{u | (u, v) \in E \vee (v, u) \in E\}$

GraphLab - User defined computation

There are two types of user defined computation on GraphLab:

- ▶ **Update function:**

- ▶ local computation
- ▶ executed on small neighborhoods
- ▶ different functions may access and modify overlapping contexts

- ▶ **Sync Mechanism:**

- ▶ global aggregation
- ▶ runs concurrently with the update functions

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example


2.5 Execution Model and Data Consistency

Update Function - Scope of a Vertex

The **Update Function** is a *stateless* user defined function

It operates on the **scope** \mathcal{S}_v of a vertex v :

$$\mathcal{S}_v := \{v\} \cup (v \rightarrow *) \cup (* \rightarrow v) \cup \mathcal{N}_v$$



Once executed on a vertex v , it may trigger the execution of the update functions on other vertices \mathcal{T}

Update Function

Be \mathcal{S}_v the scope of a vertex v and \mathbf{T} the globally shared table.

The application of an update function f to a vertex v can be defined as

$$(D_{\mathcal{S}_v}, \mathcal{T}) \leftarrow f(D_{\mathcal{S}_v}, \mathbf{T})$$

Where:

- ▶ $D_{\mathcal{S}_v}$ is the data associated with the scope of v
- ▶ f has read-only access to \mathbf{T}
- ▶ \mathcal{T} is a set of vertices

Update Function and Vertex Programs

In the most recent version of the GraphLab abstraction, Update Functions are called **Vertex Programs**

A Vertex Program implements the *Gather-Apply-Scatter (GAS)* model, and has three phases:

- ▶ **Gather:** executed in parallel on each edge of the current vertex (usually to read data)
- ▶ **Apply:** atomic function that modifies the vertex data
- ▶ **Scatter:** Executed in parallel on each edge of the current vertex (usually to update edge data and signaling neighboring nodes)

This decomposition allows us to execute a single vertex program on several machines simultaneously and move computation to the data.

Structure of an Update Function

```
1: procedure UPDATEFUNCTION
   input: vertex  $v$ , scope  $\mathcal{S}_v$ , data  $D_{\mathcal{S}_v}$ , shared table  $\mathbf{T}$ 

2:   GatherType:  $g$                                 ▷ stores the results of the gather phase
3:   for  $(u, v) \in (* \rightarrow v)$  do                      ▷ In Parallel
4:      $g \leftarrow g + \text{gather}((u, v))$ 
5:   end for
6:    $\text{apply}(g, D_v)$ 
7:   for  $(v, u) \in (v \rightarrow *)$  do                      ▷ In Parallel
8:      $\text{scatter}((v, u))$ 
9:   end for
10: end procedure
```

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example

2.5 Execution Model and Data Consistency

Sync Mechanism

Aggregates data accross all vertices in the graph

Associates the result with a particular entry in \mathbf{T}

The user provides:

- ▶ a key k
- ▶ an initial value r_k^0
- ▶ a fold function: $r_k^{i+1} \leftarrow \text{Fold}_k(D_v, r_k^i)$
- ▶ an optional merge function: $r_k^l \leftarrow \text{Merge}_k(r_k^i, r_k^j)$
- ▶ an apply function: $\mathbf{T}[k] \leftarrow \text{Apply}_k(r_k^{|V|})$

Sync Mechanism

1: **procedure** SYNC ALGORITHM

input: key k , vertices V , data $\{D_v\}_{v \in V}$, initial value r_k^0 , shared table \mathbf{T}

2: $t \leftarrow r_k^0$

3: **for** $v \in V$ **do**

4: $t \leftarrow \text{Fold}_k(D_v, t)$

5: **end for**

6: $\mathbf{T}[k] \leftarrow \text{Apply}_k(t)$

7: **end procedure**

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function


2.3 Sync Mechanism

2.4 PageRank Example

2.5 Execution Model and Data Consistency

Pagerank: Data Model

- ▶ $G := \{w_1, w_2, w_3, w_4\}$
- ▶ $E := \{(w_1, w_2), (w_1, w_3), (w_3, w_1), (w_2, w_3), (w_4, w_3)\}$
- ▶ $D_v := PR(v)$
- ▶ $D_{v,u} := \emptyset$



Pagerank: Update function

1: **procedure** PAGERANKUPDATEFUNCTION

input: vertex v , scope \mathcal{S}_v , ingoing edges $(* \rightarrow v)$

2: $\mathcal{T} \leftarrow \emptyset$

3: $D_v^{old} \leftarrow D_v$

4: $p \leftarrow 0$

5: **for** $u \in \{u \mid (u, v) \in (* \rightarrow v)\}$ **do**

6: $p \leftarrow p + \frac{D_u}{|(u \rightarrow *)|}$

7: **end for**

8: $D_v \leftarrow (1 - \beta) + \beta * p$

9: **if** $|D_v - D_v^{old}| > \epsilon$ **then**

10: $\mathcal{T} \leftarrow \mathcal{N}_v$

11: **end if**

12: **return** $(D_{\mathcal{S}_v}, \mathcal{T})$

13: **end procedure**

PageRank: Vertex Program

```

1: procedure PAGERANKGATHER
  input: vertex  $v$ , scope  $\mathcal{S}_v$ ,
  ingoing edge  $(u \rightarrow v)$ 

2:   return  $\frac{D_u}{|(u \rightarrow *)|}$ 

3: end procedure
  
```

```

1: procedure PAGERANKAPPLY
  input: vertex  $v$ , scope  $\mathcal{S}_v$ ,
  gather result  $p$ ,

2:    $D_v^{old} \leftarrow D_v$ 
3:    $D_v \leftarrow (1 - \beta) + \beta * p$ 
4:   if  $|D_v - D_v^{old}| > \epsilon$  then
5:     Signal to perform scatter
6:   end if

7: end procedure
  
```

Pagerank: Scatter

1: **procedure** PAGERANKSCATTER

input: outgoing edge ($v \rightarrow u$), set of vertices to be rescheduled \mathcal{T}

2: Add u to \mathcal{T}

3: **end procedure**

PageRank Implementation: Data Types

```
struct web_page {  
    std::string pagename;  
    double pagerank;  
  
    explicit web_page(std::string name):  
        pagename(name), pagerank(0.0) { }  
  
    void save(graphlab::oarchive& oarc) const {  
        oarc << pagename << pagerank;  
    }  
    void load(graphlab::iarchive& iarc) {  
        iarc >> pagename >> pagerank;  
    }  
};  
  
typedef graphlab::distributed_graph  
    <web_page, graphlab::empty> graph_type;
```

PageRank Implementation: Reading the Data

```
bool line_parser(graph_type& graph,
                 const std::string& filename, const std::string&
                 std::stringstream strm(textline);
                 graphlab::vertex_id_type vid;
                 std::string pagename;

                 strm >> vid >> pagename;

                 graph.add_vertex(vid, web_page(pagename));

                 while(1){
                     graphlab::vertex_id_type other_vid;
                     strm >> other_vid;
                     if (strm.fail()) return true;
                     graph.add_edge(vid, other_vid);
                 }
                 return true;
            }
```

PageRank Implementation: Vertex Program

```
class pagerank_program :  
    public graphlab::ivertex_program<graph_type, double>,  
    public graphlab::IS_POD_TYPE {  
  
    public:  
        bool perform_scatter;  
  
        edge_dir_type gather_edges(icontext_type& context,  
                                    const vertex_type& vertex);  
        edge_dir_type scatter_edges(icontext_type& context,  
                                     const vertex_type& vertex);  
  
        double gather(icontext_type& context,  
                      const vertex_type& vertex, edge_type& edge);  
        void apply(icontext_type& context, vertex_type& vertex,  
                  const gather_type& total);  
  
        void scatter(icontext_type& context,  
                    const vertex_type& vertex, edge_type& edge);  
};
```

PageRank Implementation: Gather

```
// we are going to gather on all the in-edges
edge_dir_type gather_edges(icontext_type& context,
                           const vertex_type& vertex) const {

    return graphlab::IN_EDGES;

}

// for each in-edge gather the weighted sum of the edge.
double gather(icontext_type& context,
              const vertex_type& vertex,
              edge_type& edge) const {

    return    edge.source().data().pagerank
              / edge.source().num_out_edges();

}
```


PageRank Implementation: Apply

```
// Use the total rank of adjacent pages  
// to update this page  
void apply(icontext_type& context,  
           vertex_type& vertex,  
           const gather_type& total) {  
  
    double newval = total * 0.85 + 0.15;  
    double oldval = vertex.data().pagerank;  
    vertex.data().pagerank = newval;  
    perform_scatter = (std::fabs(prevval - oldval) > 1E-3);  
  
}
```

PageRank Implementation: Scatter

```
edge_dir_type scatter_edges(icontext_type& context ,
                             const vertex_type& vertex) {

    if (perform_scatter) return graphlab::OUT_EDGES;
    else return graphlab::NO_EDGES;

}

void scatter(icontext_type& context ,
             const vertex_type& vertex ,
             edge_type& edge) const {

    context.signal(edge.target());

}
```

Outline

1. Introduction

2. GraphLab

2.1 GraphLab Data Model

2.2 Update Function

2.3 Sync Mechanism

2.4 PageRank Example

2.5 Execution Model and Data Consistency

Execution Model

```
1: procedure GRAPHLABEXEC
   input: Graph  $G := (V, E)$ , data  $D$ , initial vertex set  $\mathcal{T} \subseteq V$ 

2:   while  $|\mathcal{T}| > 0$  do
3:      $v \leftarrow \text{RemoveNext}(\mathcal{T})$ 
4:      $(\mathcal{T}', D_{S_v}) \leftarrow f(D_{S_v}, \mathbf{T})$ 
5:      $\mathcal{T} \leftarrow \mathcal{T} \cup \mathcal{T}'$ 
6:   end while
7: end procedure
```

Serializability

- ▶ The GraphLab abstraction provides a rich *sequential model*
- ▶ Parallel Execution:
 - ▶ Multiple processors to execute the same loop on the same graph
 - ▶ Simultaneous execution of update functions on different vertices
- ▶ We desire a **serializable execution**:

Definition


A GraphLab program is **sequentially consistent** if for every parallel execution, there exists a sequential execution of update functions that produces the same values in the data graph

Data Consistency



To ensure Serializability we need **consistency models**
GraphLab supports three types of consistency models:

- ▶ **Full Consistency:** the scope of two concurrently executing update functions do not overlap
- ▶ **Edge Consistency:** each update function has R/W access to its vertex and adjacent edges but only read access to neighboring vertices
- ▶ **Vertex Consistency:** only ensures that no concurrently update functions are executed on the same vertice

Data Consistency - Exclusion Sets



Data Consistency




Data Consistency


GraphLab *guarantees* sequential consistency under the following three conditions:

- ▶ The **full consistency** model is used
- ▶ The **edge consistency** model is used and update functions do not modify data in adjacent vertices
- ▶ The **vertex consistency** model is used and update functions only access local vertex data

Benchmarks - Horizontal Scaling



Dataset 1



Dataset 2

Dataset	$ V $	$ E $	Size
1	61.2K	50.9 M	655 MB
2	65.6 M	1.8 B	31 GB

Yong Guo, Marcin Biczak, Ana Lucia Varbanescu, Alexandru Iosup, Claudio Martella, Theodore L. Willke. *Towards Benchmarking Graph-Processing Platforms*. Super Computing 2013

Summary

A GraphLab program consists of:

- ▶ A **data graph**
- ▶ An **update function**
 - ▶ Gather Function
 - ▶ Apply Function
 - ▶ Scatter Function
- ▶ A **sync mechanism**
- ▶ A **consistency model**