

www.qconferences.com





Guagua: An Iterative Computing Framework on Hadoop

Zhang Pengshan(David), PayPal

AGENDA

- Introduction
- Distributed Neural Network Algorithm
- What is Guagua?
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

AGENDA

- Introduction
- Distributed Neural Network Algorithm
- What is Guagua?
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

付款 转账

收款方：*春茂

1.50 元

订单详情

余额宝可支付 0 元 立即转入

账户余额 188.10 元 支付 1.50 元

安全设置检测成功！数字证书正在保护中。

支付宝支付密码： 忘记密码？

确认付款

付款方式： 储蓄卡

快捷支付 72小时100%赔付

付款过程安全流畅

招商银行

选择其他

网上银行：需跳转银行页面，将不享受免费保障服务 [如何购买](#)

中国工商银行

中国建设银行

中国农业银行

中国银行股份有限公司

交通银行

招商银行

中国银行

选择其他

下一步

付款遇到问题：

1. 什么是快捷支付(含卡通)？

答：快捷支付是支付宝联合各大银行推出的全新的支付方式。只要您有银行卡，就可以在支付宝付款。付款时无需登录网上银行，凭支付宝支付密码和手机校验码即可完成付款。（原卡通用户自动升级为快捷支付用户）。

2. 没有网上银行，怎样用储蓄卡付款？

答：您可使用储蓄卡快捷支付轻松完成付款。目前暂时支持工行、农行、交行，其他银行储蓄卡快捷支付即将推出。

3. 我有网上银行，如何在线开通储蓄卡快捷支付（含卡通）？

答：已支持工行、招行、建行、中行、农行、北京银行在线开通。在支付宝填写信息后，根据页面引导完成开通。[马上开通](#)

4. 为什么银行卡已经扣款，但支付宝中没有显示？

答：有可能因为网络原因导致，将在第二个工作日恢复。

[更多帮助](#) [进入我的支付宝](#)

Send Money

Secure Transaction

Review your payment and send

Please check the following information before continuing.

Description	Amount	
Payment To ritalea78@hotmail.com	\$1.00 USD	Change

[Apply coupon/gift certificate](#)
Your total charge: \$1.00 USD

Send Money

Payment Method

☒ Credit Card : MasterCard XXXX XX XXXX 9876 543210
[View card details](#)
This credit card transaction will appear on your bill as "PAYALEA78".
[More funding options](#)

☐ Bank Account : Chase Bank **** *
[View bank details](#)

☐ PayPal : paypal@paypal.com
[View paypal details](#)

Shipping Address

☒ Shipping to address? (choose shipping address if you're paying for services, choose "no shipping required" if you're not)
[Choose Shipping Address](#)

☐ No shipping required

NOTE: We recommend that you provide your shipping address in English if your seller cannot read your local language.

Email to recipient

Subject:

Message (optional):

0 characters entered - 300 maximum

Send Money

Q: Where is risk control in PayPal?

A: Risk control is everywhere in paypal.com.

FRAUD TYPES IN PAYPAL

Fraud Types in PayPal

Account Take Over



Stolen Financials



Credit Cards

INR/SNAD

ebay™
ITEMS NOT RECEIVED (INR).


**Win a FREE Ride in a
Police Car just by
shoplifting from this store.**

Lucky Winners will
also get their name
in the newspaper for
their friends and
family to see!

Won't Mom & Dad be proud?

INR: Item Not Received
SNAD: Significantly Not as Described

RISK CONTROL IN PAYPAL



RISK MODELING IN PAYPAL

MODELING CHALLENGES

Thousands
of Features

Algorithms
(LR, NN, DT)

Big
Training
Data

SLA
(Online)

Simulation

RISK MODELING IN PAYPAL

MODELING CHALLENGES

Thousands
of Features

Algorithms
(LR, NN, DT)

Big
Training
Data

SLA
(Online)

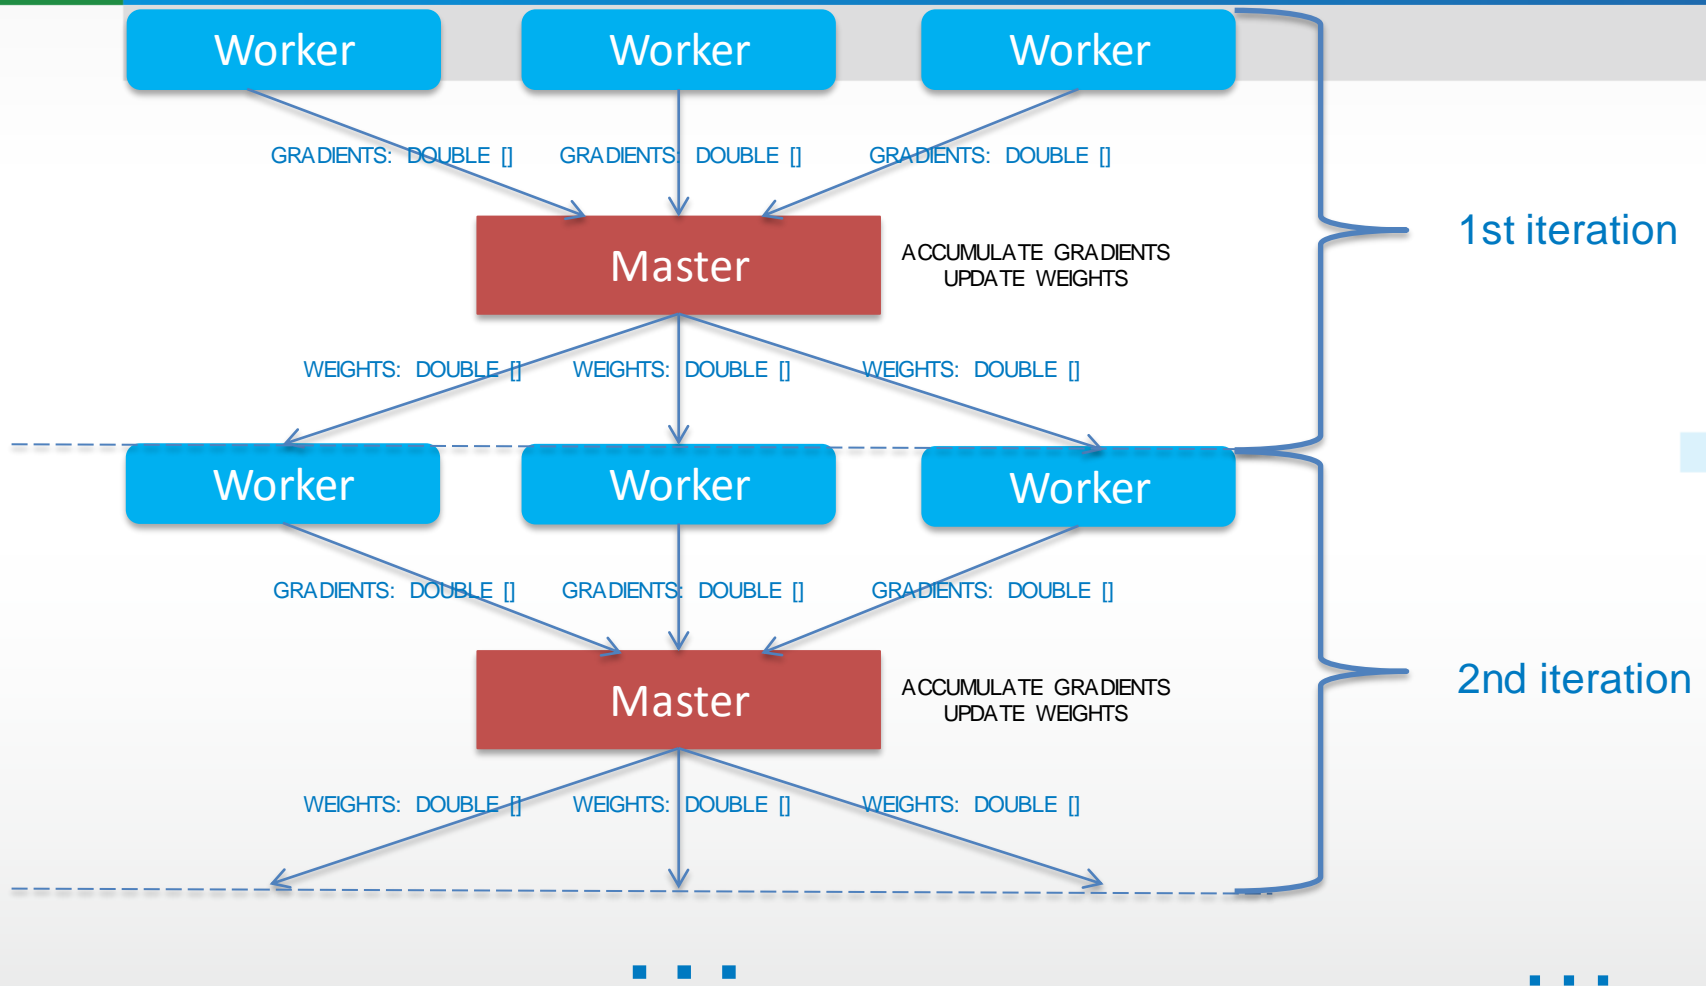
Simulation

Q: How to train models with TB data and thousands of features?

AGENDA

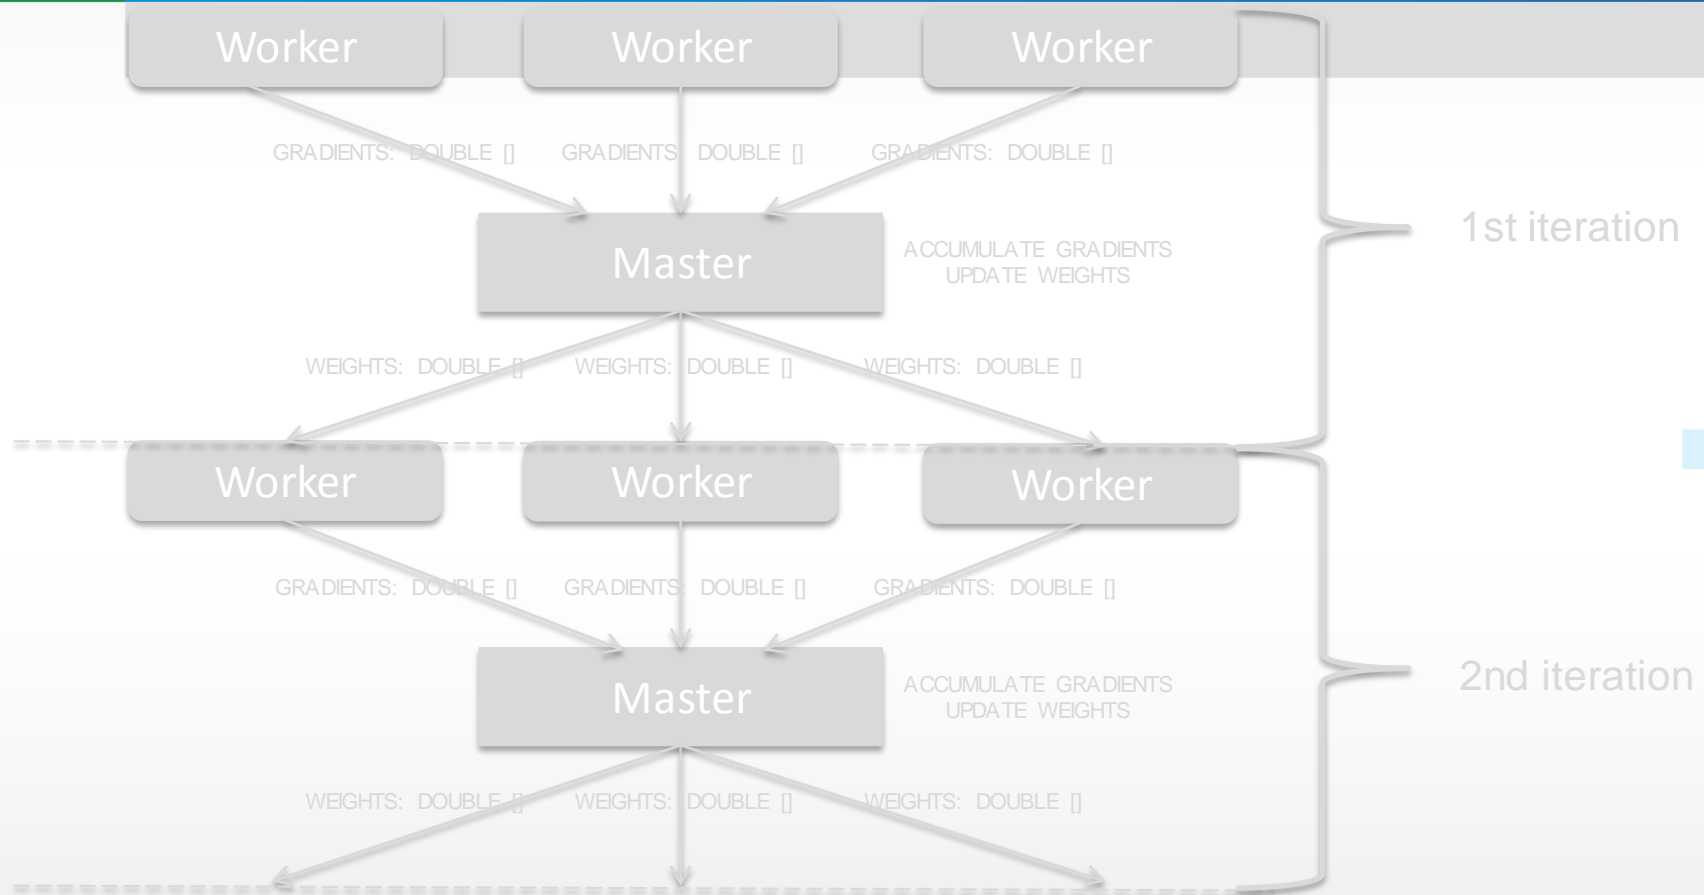
- Introduction
- **Distributed Neural Network Algorithm**
- What is Guagua?
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

DISTRIBUTED NEURAL NETWORK ALGORITHM*



* Distributed batch gradient descent algorithm.

DISTRIBUTED NEURAL NETWORK ALGORITHM



Q: How to implement it?

WHY NOT MAHOUT OR SPARK?

Mahout

- No distributed logistic regression & neural network.
- Iterative through Hadoop jobs, bad performance.

Spark

- No independent Spark cluster.
- Hadoop cluster is still 1.0 based, not YARN.

Q: How to implement it in Hadoop?

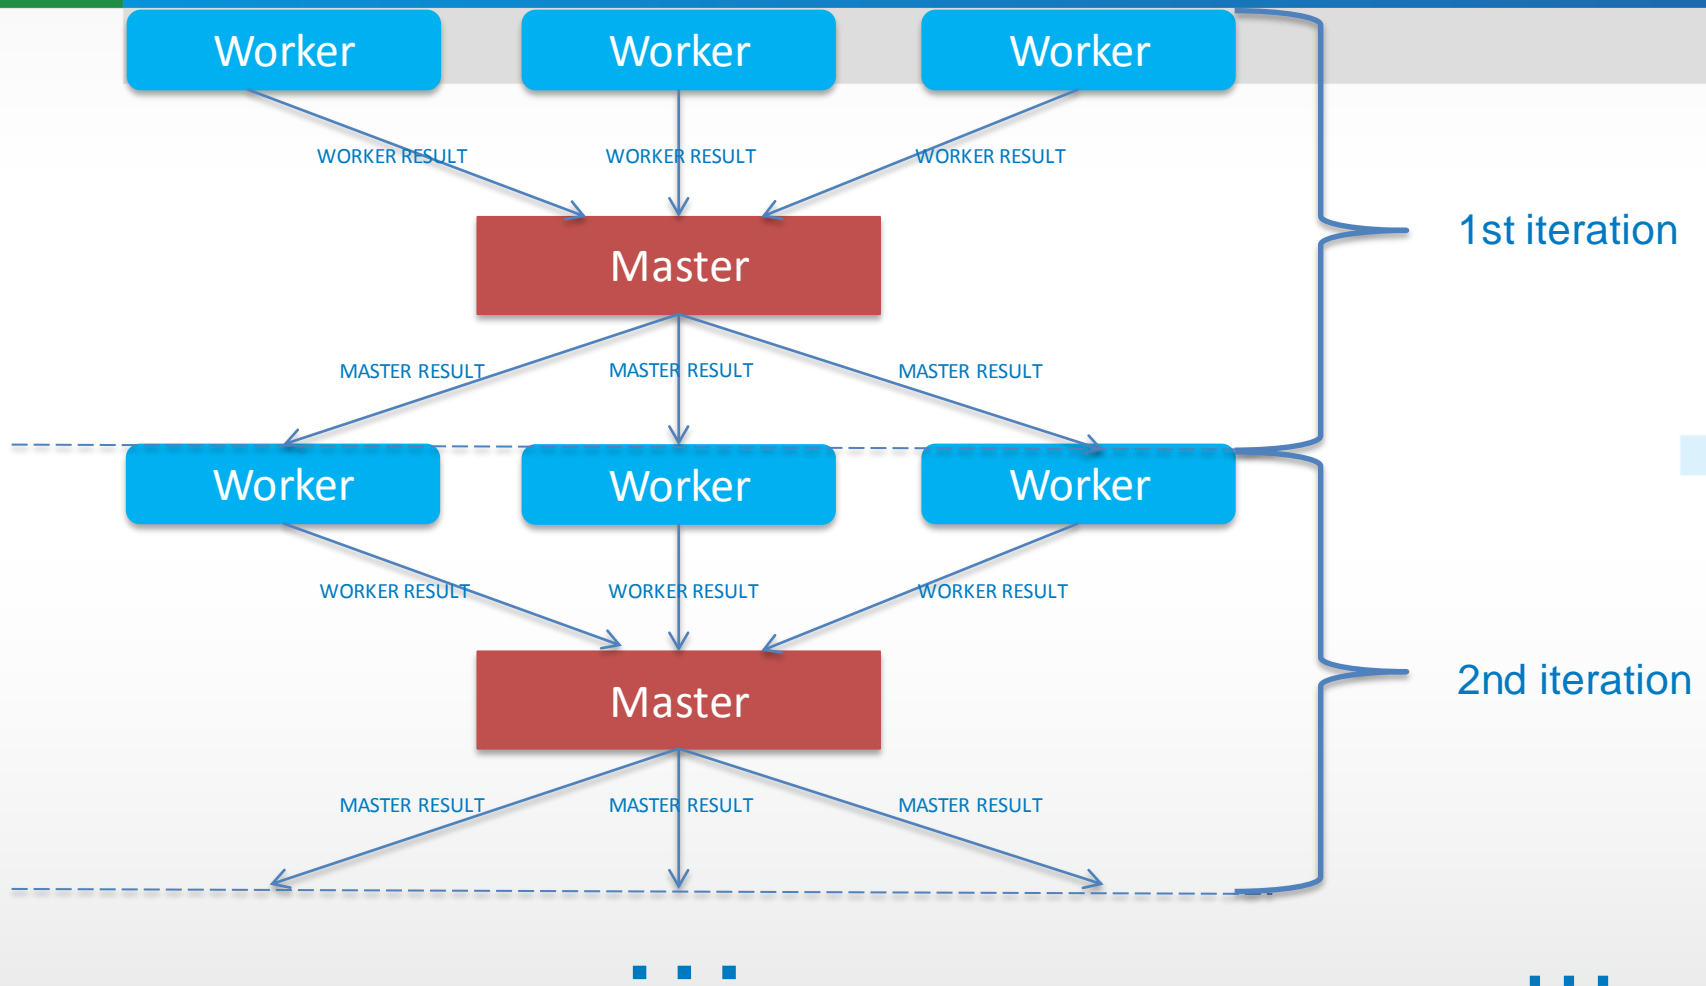
POSSIBLE SOLUTIONS

	Hadoop YARN	Hadoop MapReduce
Pros	Flexible framework for framework	Works well on all Hadoop versions
	Self resource management	Mature computing model
		Internal fault tolerance, splits, UI ...
Cons	2.0.3-Alpha	Different computing model
	PayPal Clusters: Hadoop 0.20.2	How to do iterative coordination?
	Extra fault tolerance, splits, UI ...	

AGENDA

- Introduction
- Distributed Neural Network Algorithm
- **What is Guagua?**
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

ITERATIVE COMPUTING MODEL IN GUAGUA



GUAGUA API

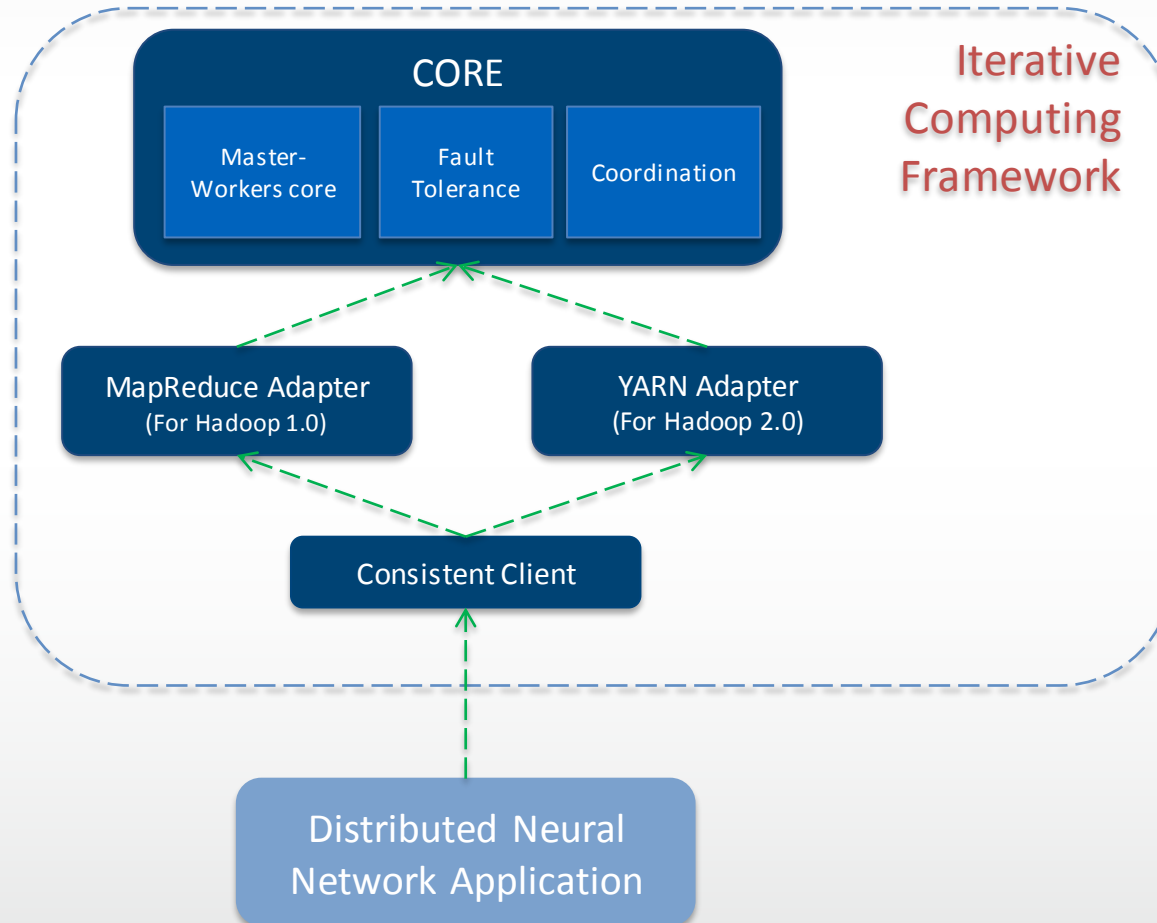
MasterComputable

```
public interface MasterComputable<MASTER_RESULT extends Byteable, WORKER_RESULT extends Byteable> {  
    /**  
     * Master computation for each iteration.  
     *  
     * @param context  
     *         the master context instance which includes worker results and other useful parameters.  
     * @return  
     *         the master result of each iteration.  
     */  
    MASTER_RESULT compute(MasterContext<MASTER_RESULT, WORKER_RESULT> context);  
}
```

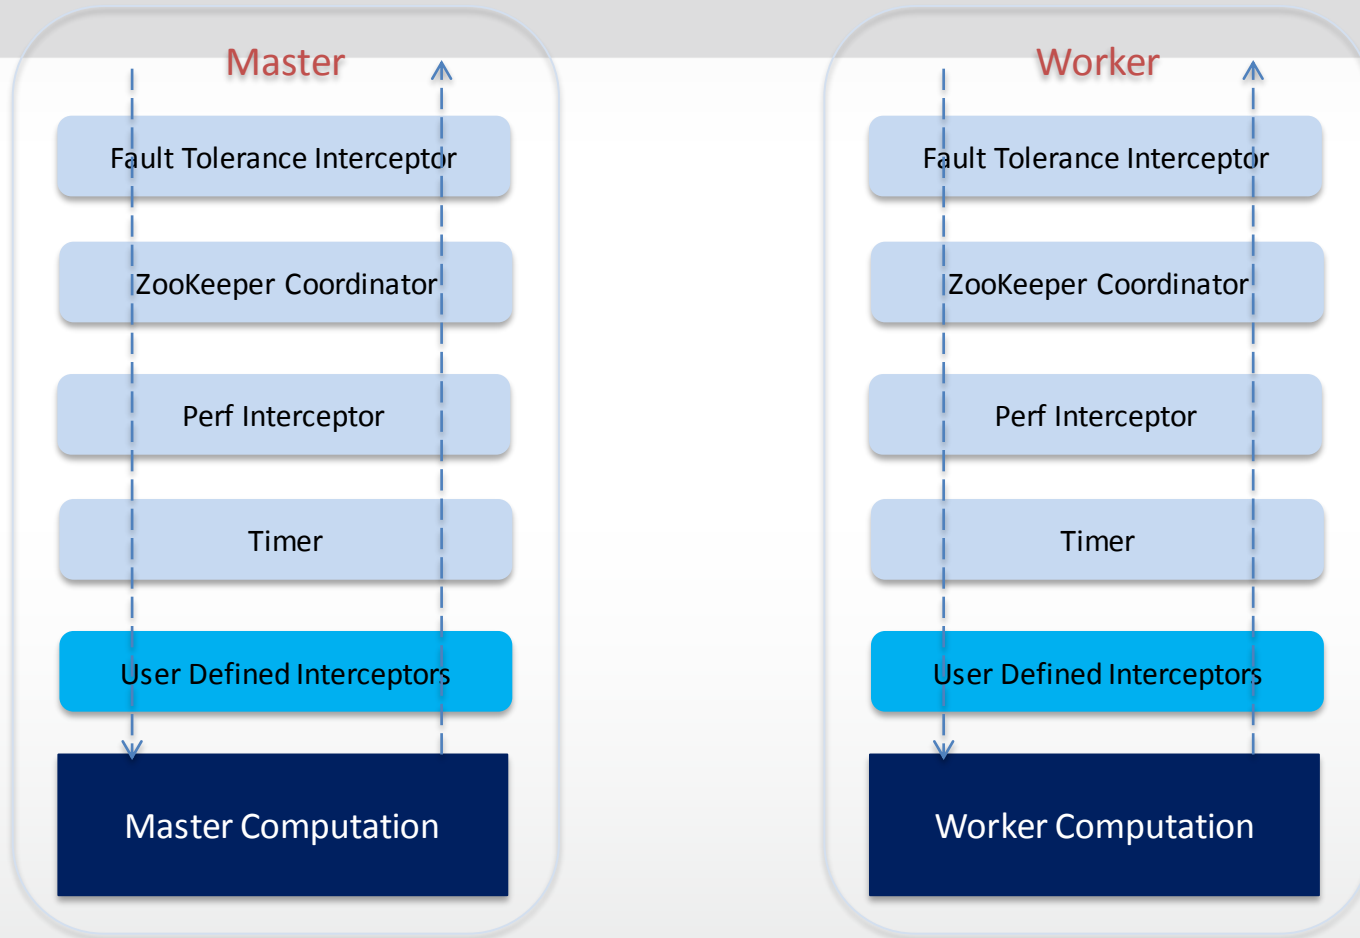
WorkerComputable

```
public interface WorkerComputable<MASTER_RESULT extends Byteable, WORKER_RESULT extends Byteable> {  
    /**  
     * Worker computation for each iteration.  
     *  
     * @param context  
     *         the worker context instance which includes worker info, master result of last iteration or other  
     *         useful info for each iteration.  
     * @return  
     *         the worker result of each iteration.  
     * @throws IOException  
     *         if any io exception in computation, for example, IOException in reading data.  
     */  
    WORKER_RESULT compute(WorkerContext<MASTER_RESULT, WORKER_RESULT> context) throws IOException;  
}
```

GUAGUA OVERVIEW



PLUGGABLE, SCALABLE INTERCEPTORS



* These two graphs are aspects for each iteration.

GUAGUA RUNTIME

Master: Mapper (Container)

Worker: Mapper (Container)

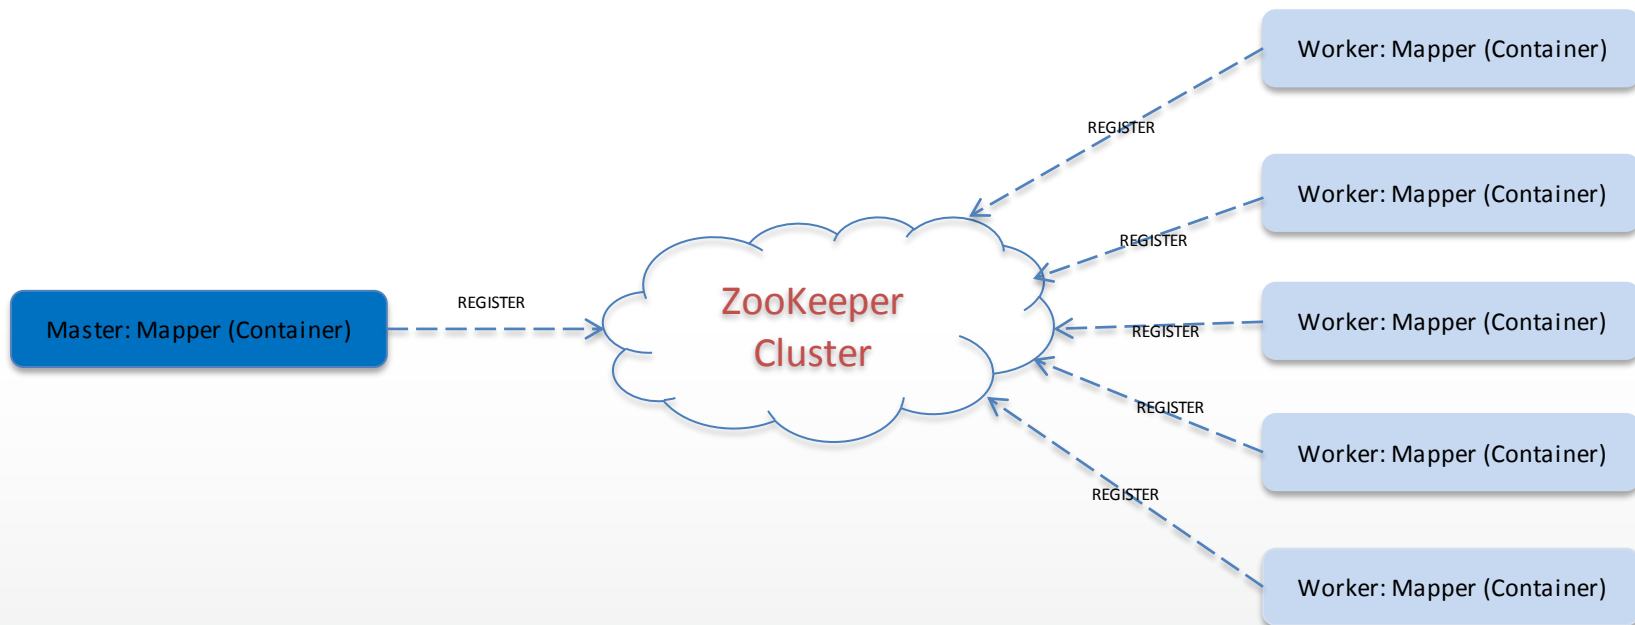
Worker: Mapper (Container)

Worker: Mapper (Container)

Worker: Mapper (Container)

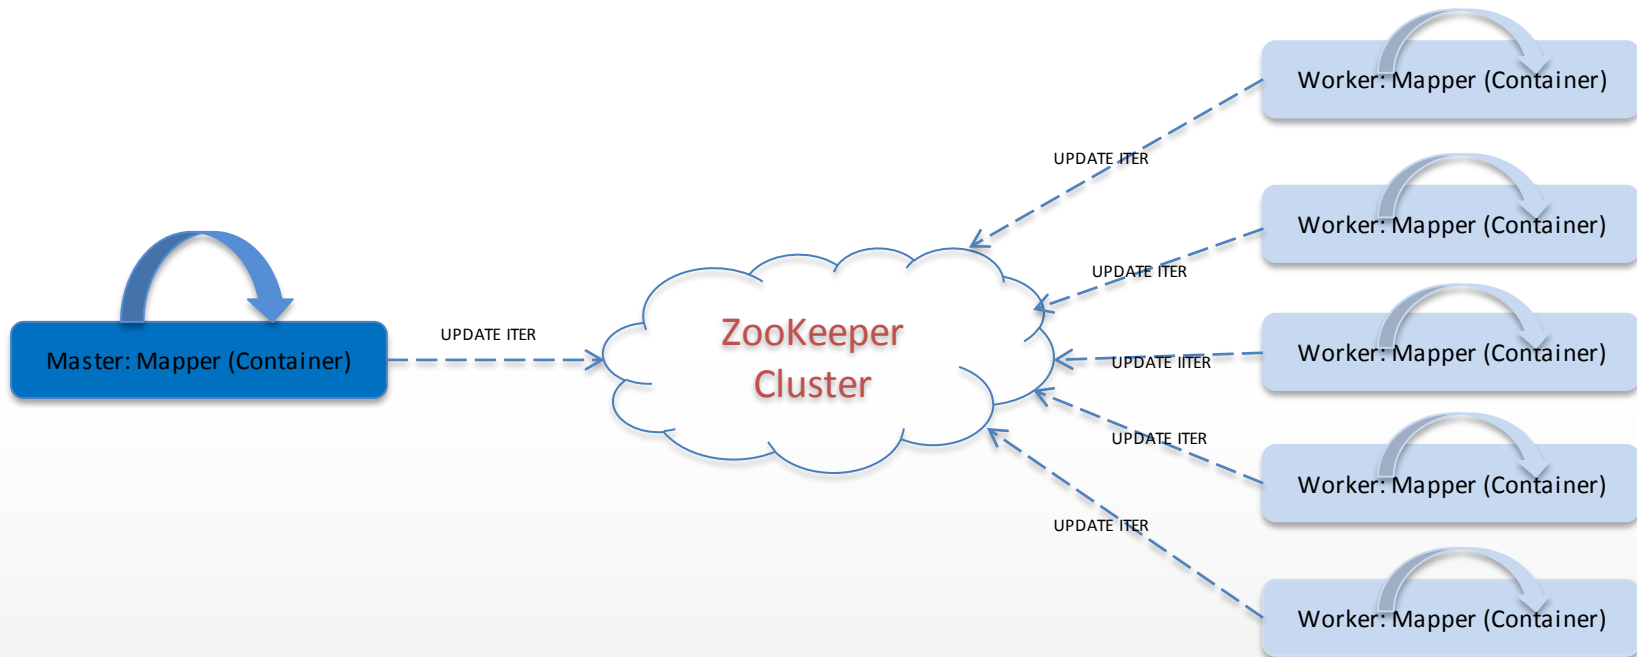
Worker: Mapper (Container)

GUAGUA RUNTIME



1. Master is listening znodes of workers.
2. Workers are listening znode of master.

GUAGUA RUNTIME

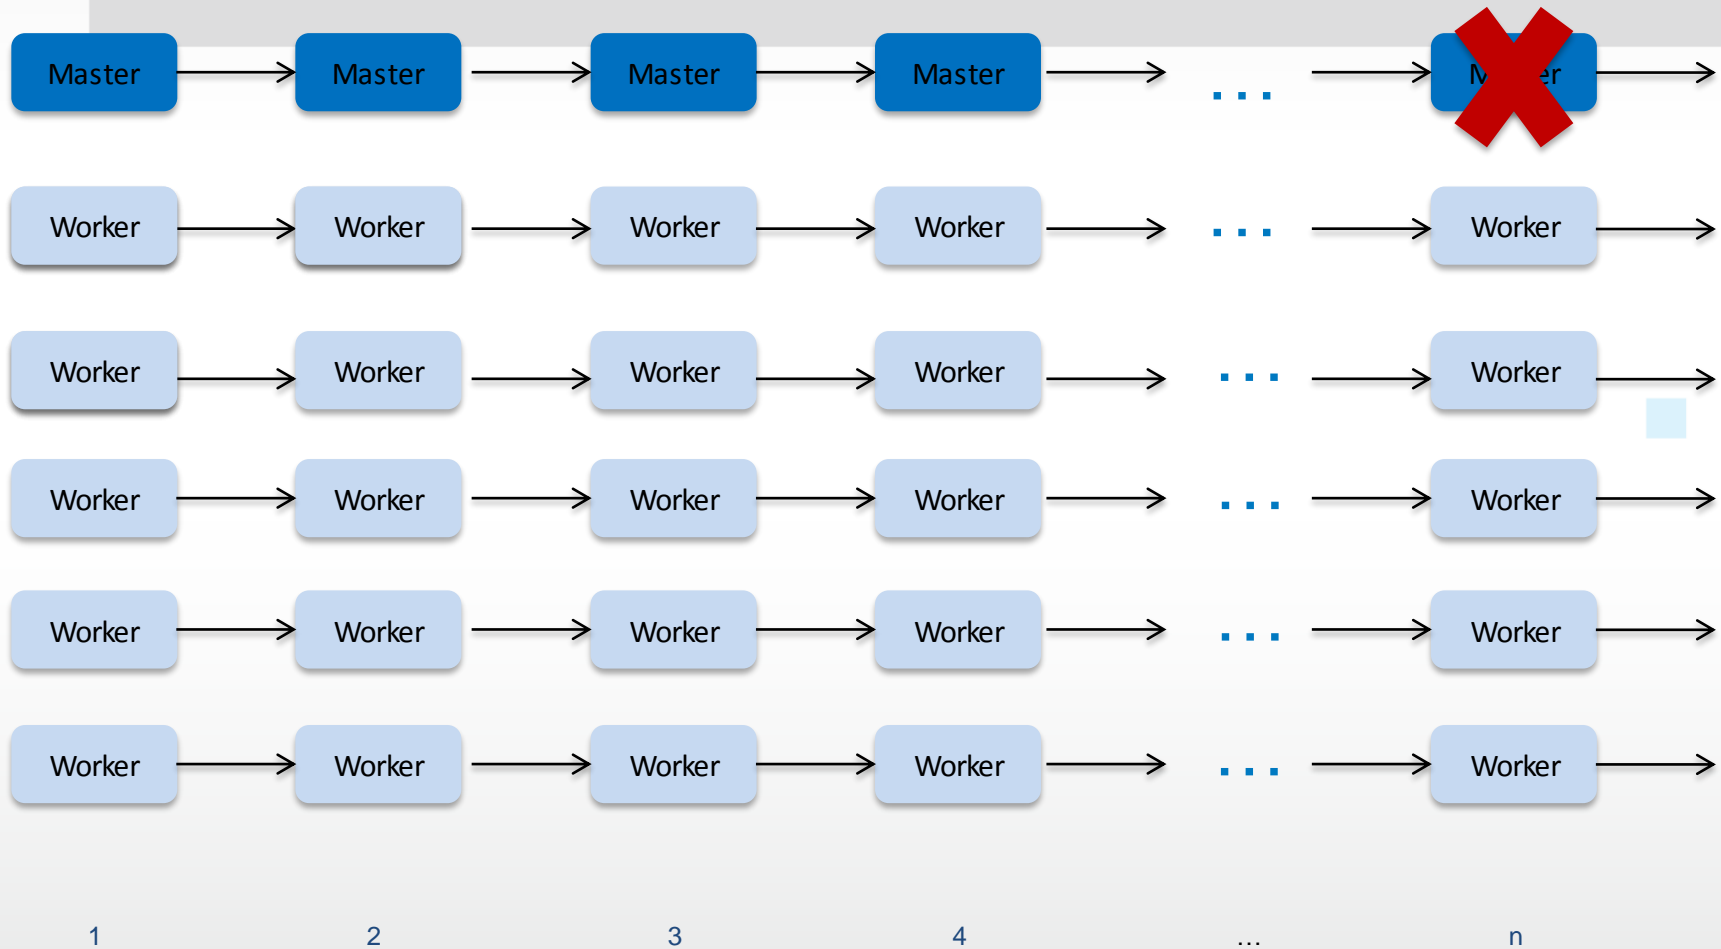


1. Data is loaded in worker memory in the first iteration.
2. Whole process is done when reaches maximal iteration or halt condition is triggered.

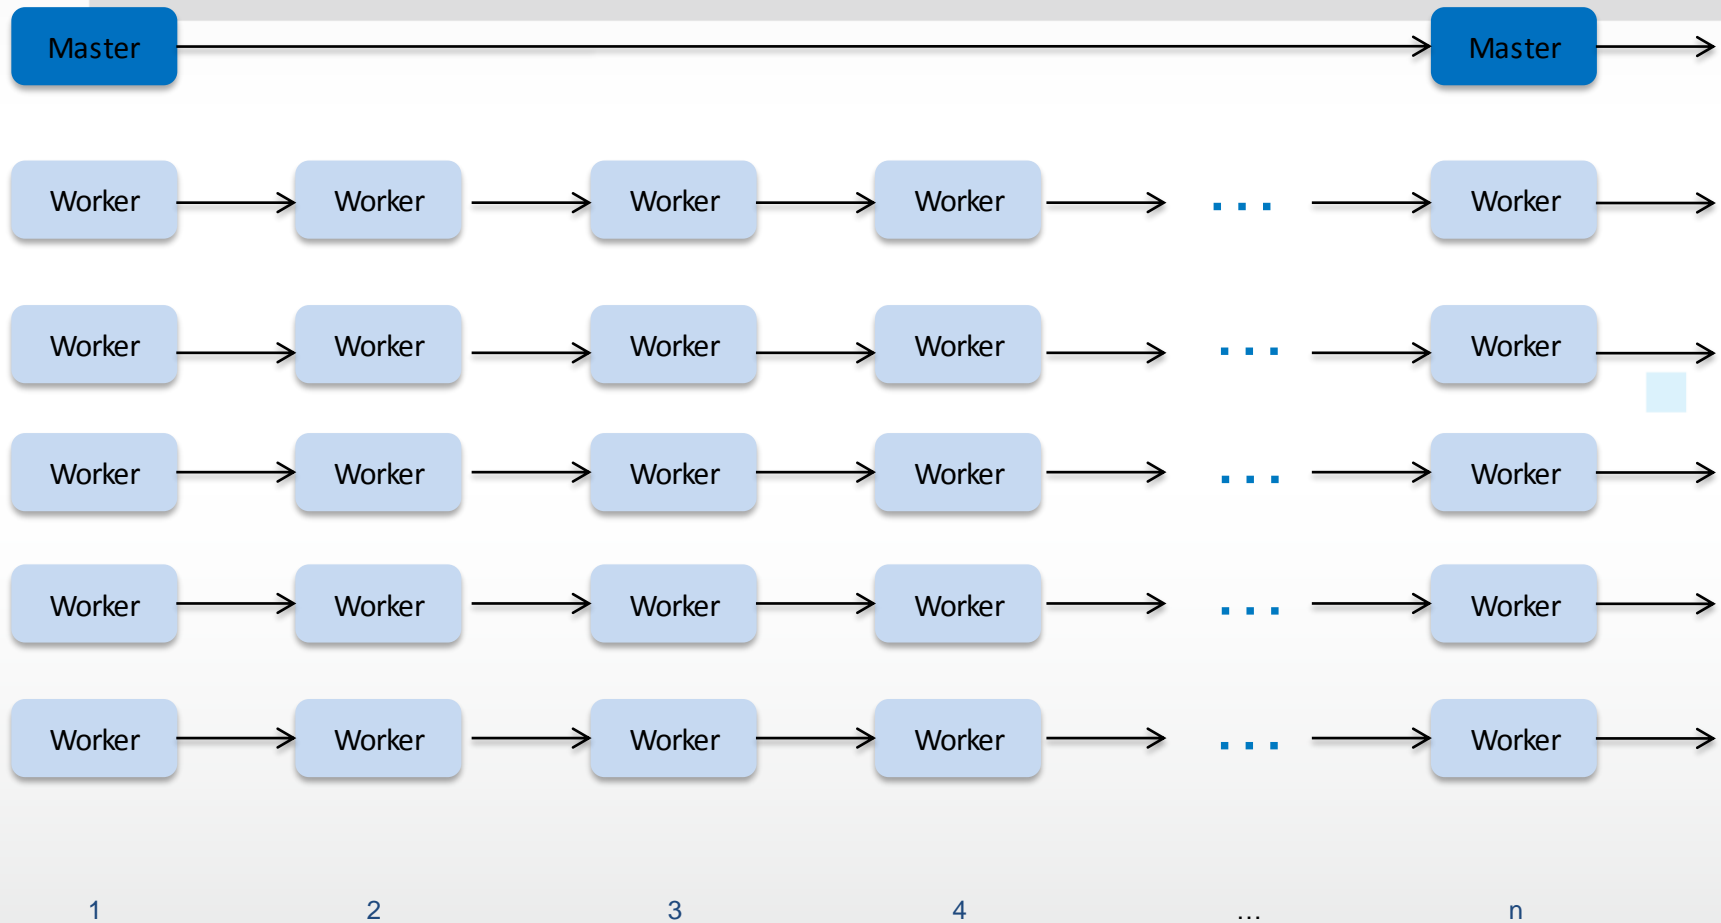
AGENDA

- Introduction
- Distributed Neural Network Algorithm
- What is Guagua?
- **Guagua Advanced Features**
- Shifu on Guagua
- Future Plans

FAULT TOLERANCE

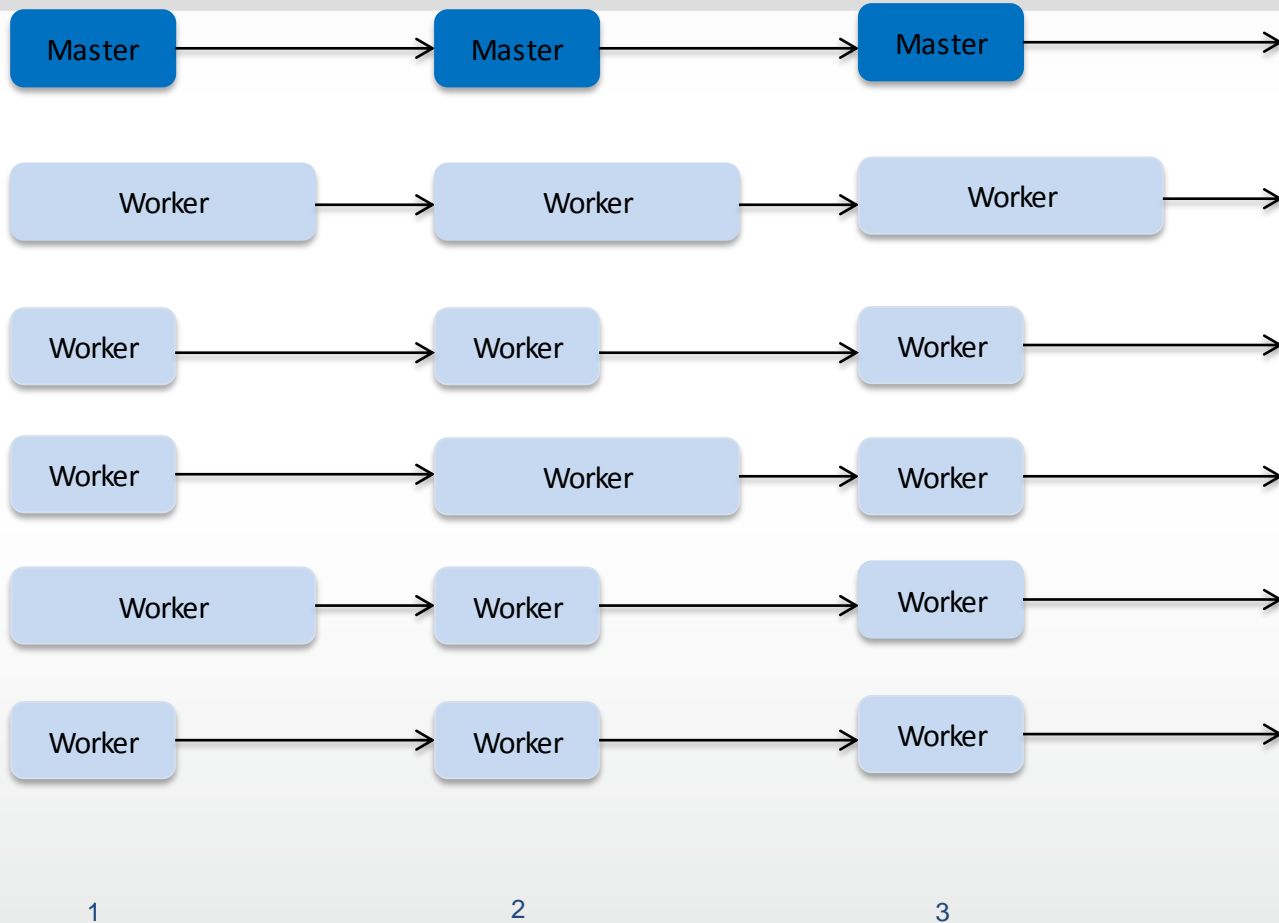


FAULT TOLERANCE

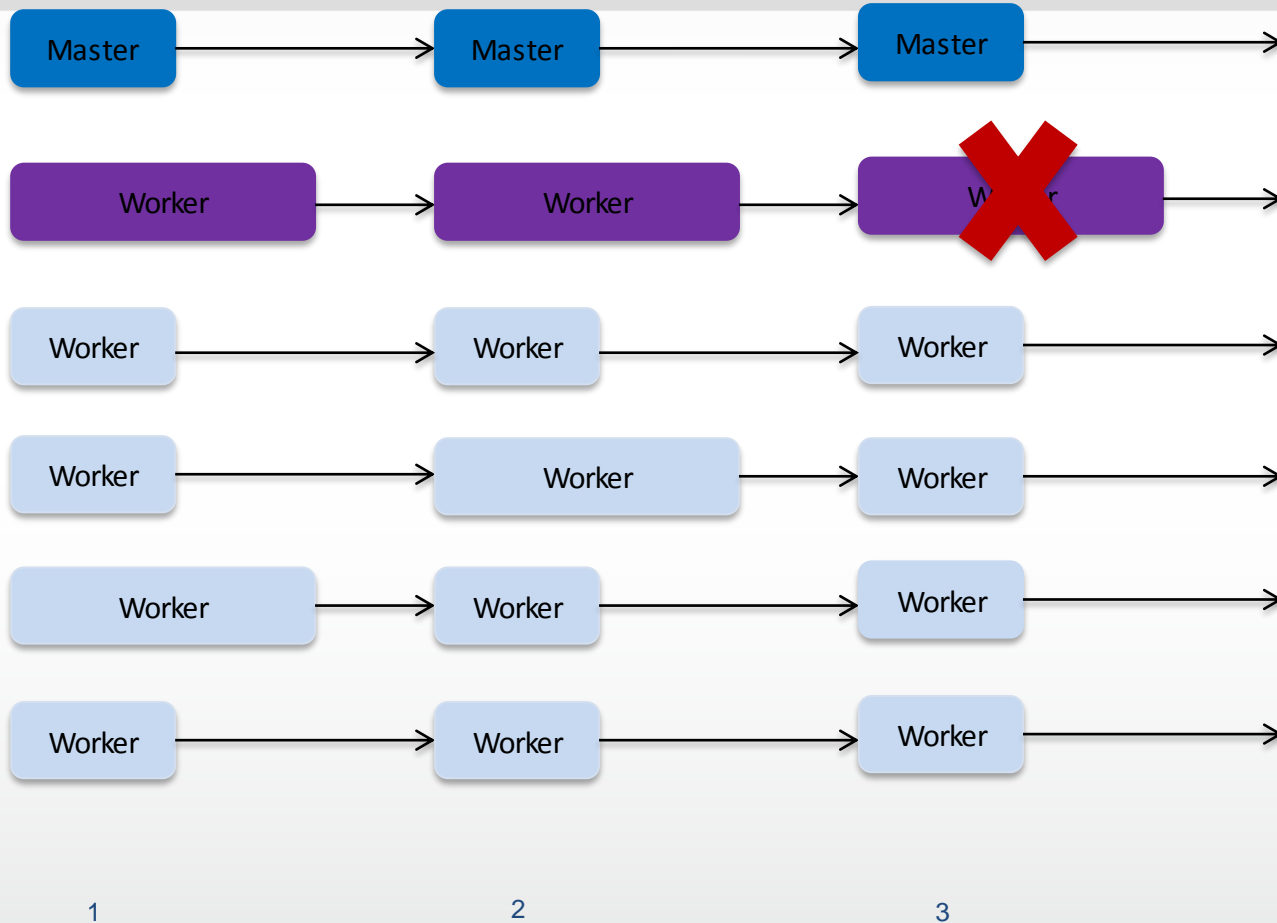


* The same on workers.

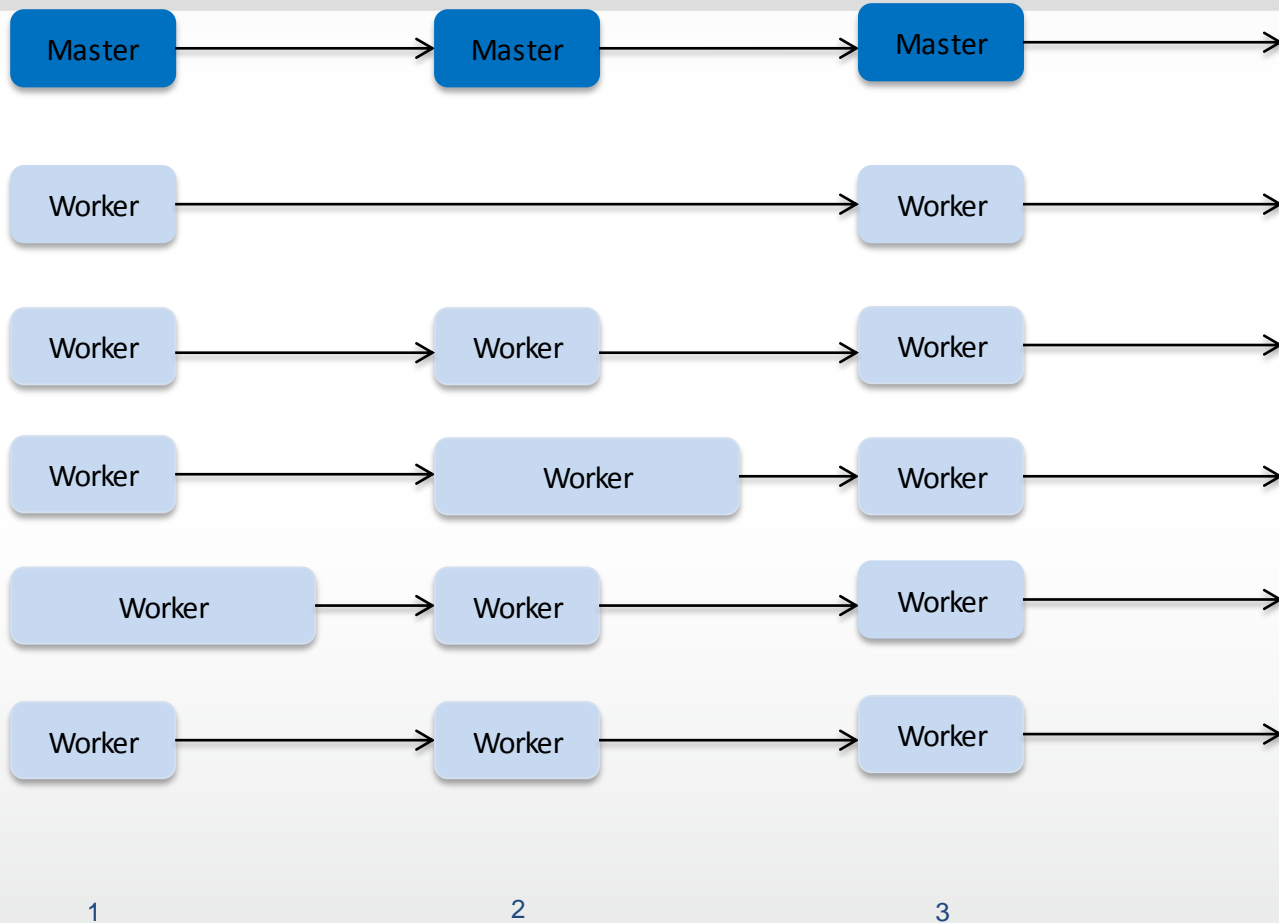
STRAGGLER MITIGATION





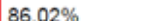


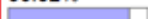
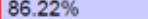
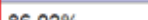


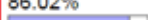




STRAGGLER MITIGATION



STRAGGLER MITIGATION



PROGRESS AND STATE REPORT

Task	Complete	Status	Start Time	Finish Time	Errors	Counters
task_201409031209_114530_m_000000	86.22% 	Start master iteration (433/501), progress 86%	8-Sep-2014 22:56:01			11
task_201409031209_114530_m_000001	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:01			10
task_201409031209_114530_m_000002	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:01			10
task_201409031209_114530_m_000003	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:02			10
task_201409031209_114530_m_000004	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:02			10
task_201409031209_114530_m_000005	86.22% 	Complete worker computing (433/501) progress 86%	8-Sep-2014 22:56:04			10
task_201409031209_114530_m_000006	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:05			10
task_201409031209_114530_m_000007	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:05			10
task_201409031209_114530_m_000008	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000009	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000010	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000011	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000012	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000013	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10
task_201409031209_114530_m_000014	86.02% 	Complete worker computing (432/501) progress 86%	8-Sep-2014 22:56:06			10

0.86 = 432/501

(Current Iteration) / (Total Iteration)

GUAGUA UNIT

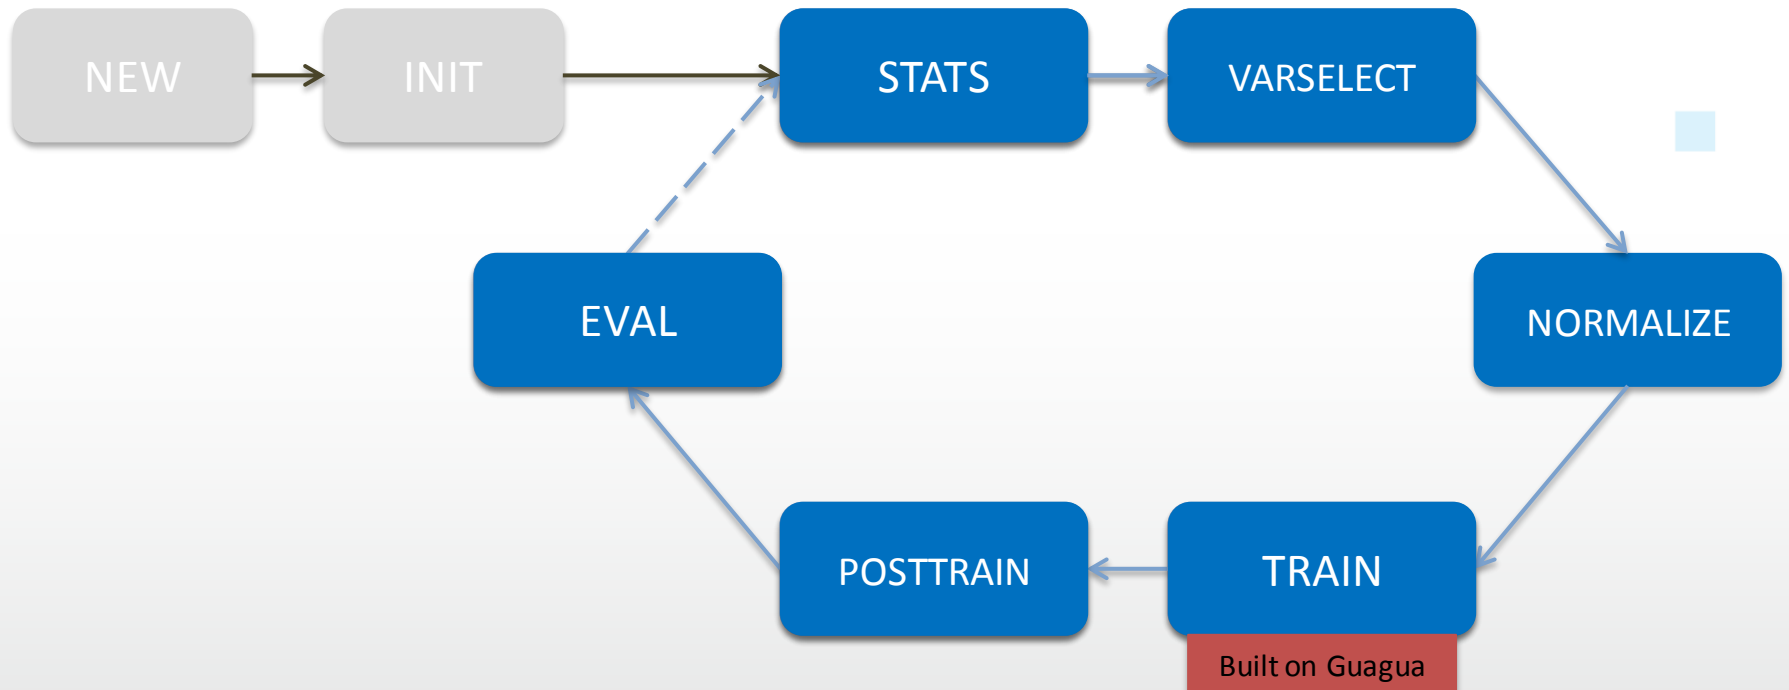
```
public class SumTest {  
    ...  
    private static final String SUM_OUTPUT = "sum-output";  
    ...  
    @Test  
    public void testSumApp() throws IOException {  
        Properties props = new Properties();  
        props.setProperty(GuaguaConstants.MASTER_COMPUTABLE_CLASS, SumMaster.class.getName());  
        props.setProperty(GuaguaConstants.WORKER_COMPUTABLE_CLASS, SumWorker.class.getName());  
        props.setProperty(GuaguaConstants.GUAGUA_ITERATION_COUNT, "10");  
        props.setProperty(GuaguaConstants.GUAGUA_MASTER_RESULT_CLASS, LongWritable.class.getName());  
        props.setProperty(GuaguaConstants.GUAGUA_WORKER_RESULT_CLASS, LongWritable.class.getName());  
        ...  
        props.setProperty(GuaguaConstants.GUAGUA_MASTER_INTERCEPTERS, SumOutput.class.getName());  
        ...  
        props.setProperty(GuaguaConstants.GUAGUA_INPUT_DIR, getClass().getResource("/sum").toString());  
        ...  
        props.setProperty("guagua.sum.output", SUM_OUTPUT);  
        ...  
        GuaguaUnitDriver<GuaguaWritableAdapter<LongWritable>, GuaguaWritableAdapter<LongWritable>> driver =  
            new GuaguaMRUnitDriver<GuaguaWritableAdapter<LongWritable>, GuaguaWritableAdapter<LongWritable>>(props);  
        driver.run();  
        ...  
        Assert.assertEquals(15345 + "", FileUtils.readLines(new File(SUM_OUTPUT)).get(0));  
    }  
    ...  
    @After  
    public void tearDown() {  
        FileUtils.deleteQuietly(new File(System.getProperty("user.dir") + File.separator + SUM_OUTPUT));  
        FileUtils.deleteQuietly(new File(System.getProperty("user.dir") + File.separator + "." + SUM_OUTPUT + ".crc"));  
    }  
}
```

AGENDA

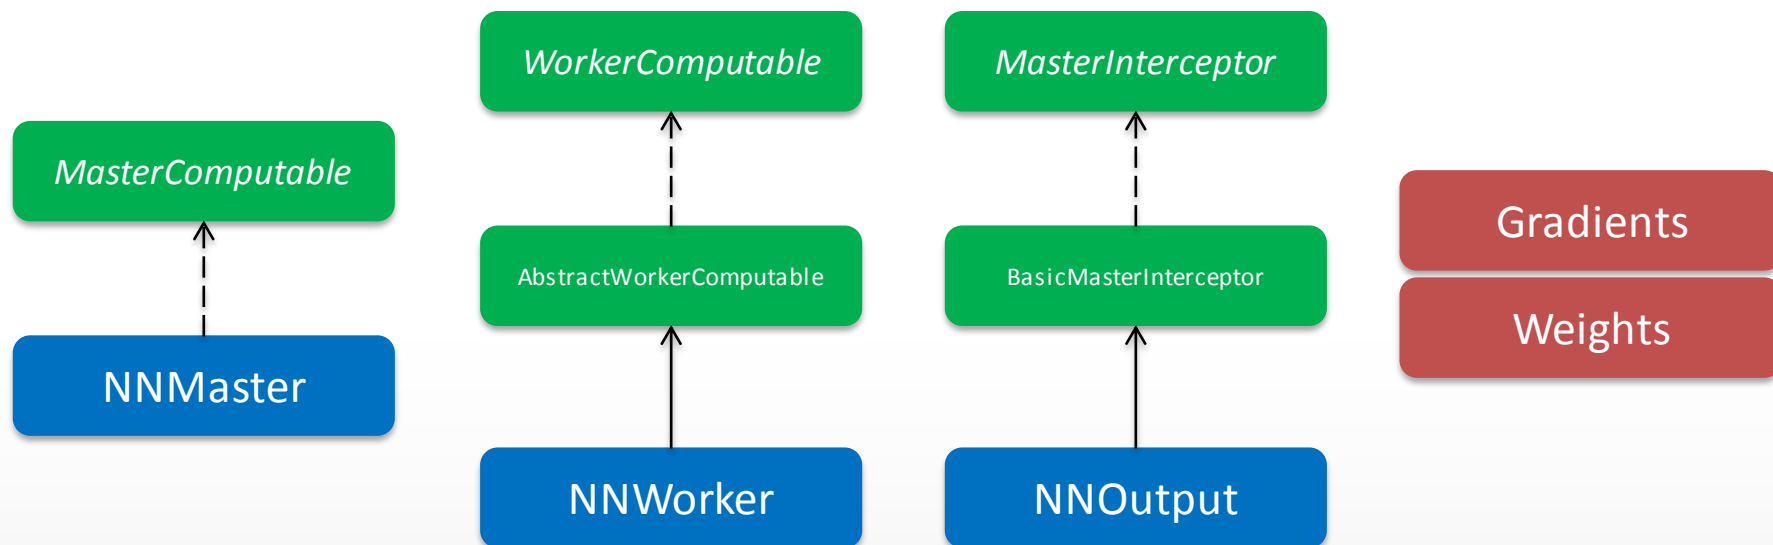
- Introduction
- Distributed Neural Network Algorithm
- What is Guagua?
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

WHAT IS SHIFU?

Shifu* is an open-source, end-to-end machine learning and data mining framework built on top of Hadoop.



SHIFU ON GUAGUA (TRAIN STEP)



GUAGUA API




SHIFU CODE



ENCOG CODE


SHIFU NN vs. SPARK LR

Run Time Comparison



Shifu-NN: 1102*20*1 Network, 319 Mappers * 1G
Spark-LR: 1102 features, 120 executors * 3G

SHIFU NN BENCHMARK RESULTS



All data are located in memory. At most we used 2400 mappers. 20 epochs are used.

AGENDA

- Introduction
- Distributed Neural Network Algorithm
- What is Guagua?
- Guagua Advanced Features
- Shifu on Guagua
- Future Plans

WHAT'S NEXT?

- More open source docs
- Support more (distributed) machine learning algorithms
- Improve YARN (Beta) implementation
- Support more input formats
- Big model support
- Deep learning support

Q&A



APPENDIX


- Website
 - <http://shifu.ml>
 - <http://shifu.ml/docs/guagua/>
- Guagua issue website
 - <https://github.com/shifuml/shifu/issues>
 - <https://github.com/shifuml/guagua/issues>
- Shifu & Guagua source code:
 - <https://github.com/shifuml/shifu/>
 - <https://github.com/shifuml/guagua/>



Brought by **InfoQ**



@InfoQ



infoqchina

软件
正在改变世界!