

# Asynchronous Peer-to-peer Data Mining with Stochastic Gradient Descent

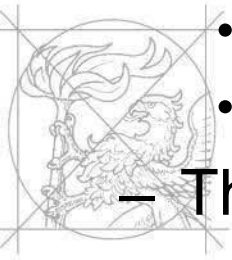
Róbert Ormándi, István Hegedűs  
and Márk Jelasity

EuroPar-2011



# Motivation

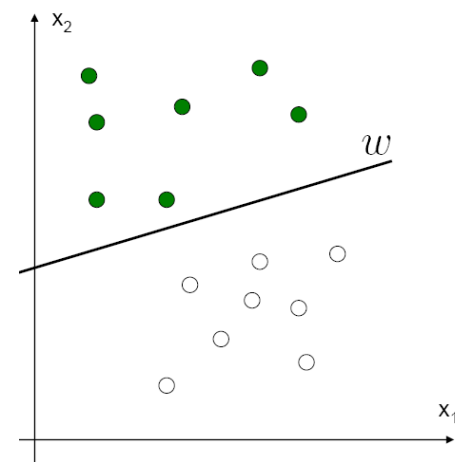
- There are P2P algorithms that compute statistics on a distributed dataset
  - Min, max, avg ...
- Our goal: using **sophisticated machine learning (ML)** algorithms in a P2P way
  - For supporting
    - spam filtering
    - opinion mining
    - personalized rank, search and recommendation
  - That is simple, cheap and robust
    - We use **gossip-based** techniques



# Our System and Data Model

- Given a network of computers (nodes)
- The database is distributed in the network
  - Every node has **exactly one** training sample → training set size = network size
- Every node can get the address of a randomly selected node from the network
  - using the NewsCast peer sampling service
- Every node can send messages to another node if its address is available
- Every node should **use locally** the computed model

# Classification



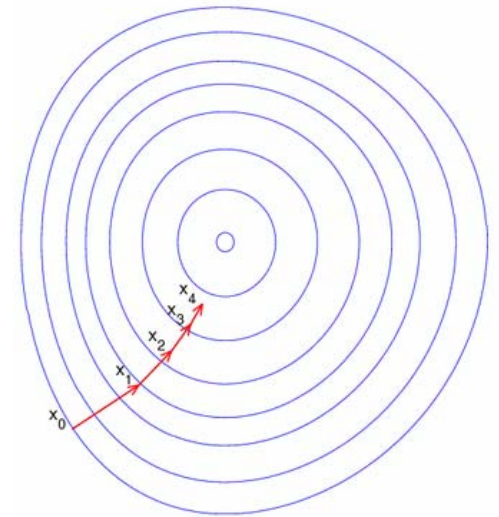
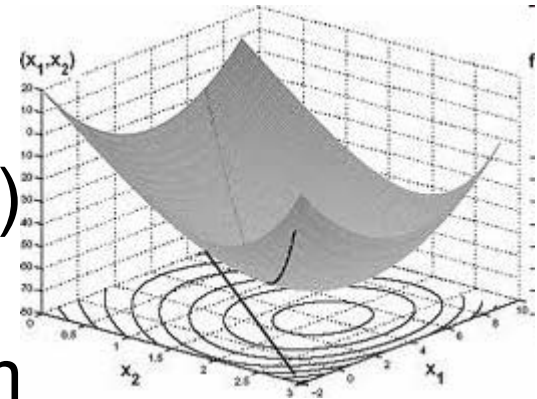
- Binary classification
  - Given  $(x_1, y_1), \dots, (x_n, y_n)$  training samples, where  $x_i \in \mathbb{R}^d$  and  $y_i \in \{-1, 1\}$
  - Task: looking for a model  $f : \mathbb{R}^d \rightarrow \{-1, 1\}$  that correctly separates the samples from different classes  
→ **minimization problem** of the formula
$$\min_f \sum_i (f(x_i) - y_i)^2 \quad i = 1, \dots, n$$
  - In the linear case the model is a  $d$  dimensional hyperplane ( $w$ )
- **Stochastic Gradient Descent** is such a kind of method, that can find this model



# Stochastic Gradient Descent (SGD)

## centralized case

- The SGD is an optimization algorithm for finding the minimum of an objective function  $((f(x_i) - y_i)^2)$  using gradient steps
- Iteratively chooses a **single** uniform random sample from the training set for updating the model based on the gradient step
- Why SGD: Uses **only one** training sample at a time instead of the whole training set



# Stochastic Gradient Descent (SGD)

## centralized case

- Assume the error is defined as
- Then the gradient is
- So the gradient update is
- But SGD makes update based on only one sample

$$Err(w) = \sum_{i=1}^n Err(w, x_i)$$

$$\frac{\partial Err(w)}{\partial w} = \sum_{i=1}^n \frac{\partial Err(w, x_i)}{\partial w}$$

$$w(t+1) = w(t) - \lambda(t) \sum_{i=1}^n \frac{\partial Err(w, x_i)}{\partial w}$$

$$w(t+1) = w(t) - \lambda(t) \frac{\partial Err(w, x_i)}{\partial w}$$

# P2P SGD

- Basic idea: models **jump from node to node** and they update themselves based on the local training sample
- We have the same SGD algorithm in a turned up manner
  - if we can guarantee the uniformity of peer sampling then our SGD models will converge to the optimum
- Privacy aspect: the data never leaves the node, just the model

# Our Data Mining Framework

---

**Algorithm 1** P2P Stochastic Gradient Descent Algorithm

---

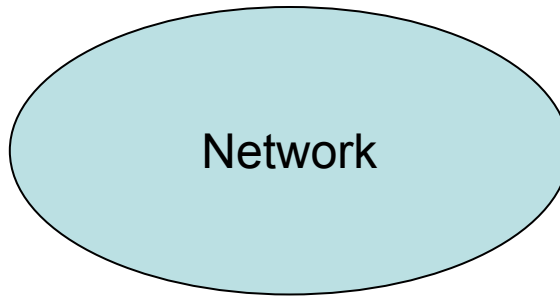
```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

Node



Network



# Our Data Mining Framework

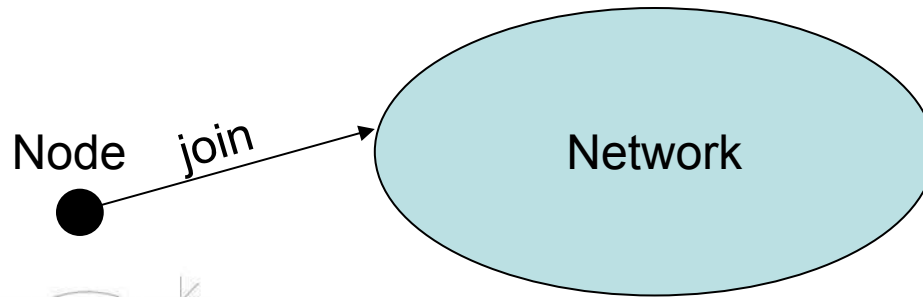
---

**Algorithm 1** P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---



# Our Data Mining Framework

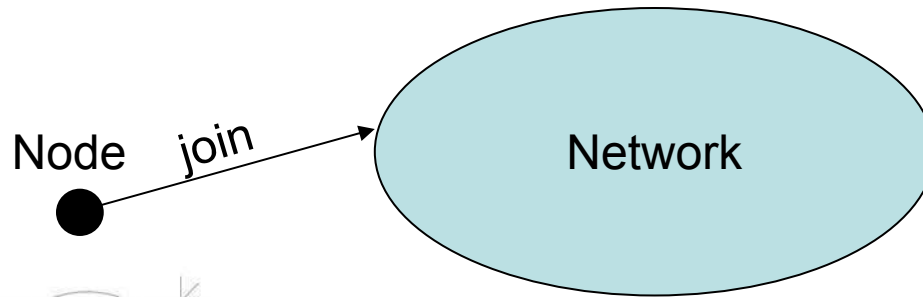
---

**Algorithm 1** P2P Stochastic Gradient Descent Algorithm

---

1: <code>initModel()</code>	6: <b>procedure</b> <code>ONRECEIVEMODEL(<math>m</math>)</code>
2: <b>loop</b>	7: $m \leftarrow \text{updateModel}(m)$
3: <code>wait(<math>\Delta</math>)</code>	8: <code>currentModel</code> $\leftarrow m$
4: $p \leftarrow \text{selectPeer}()$	9: <code>modelQueue.add(<math>m</math>)</code>
5:   send <code>currentModel</code> to $p$	

---



Model Initialization

# Our Data Mining Framework

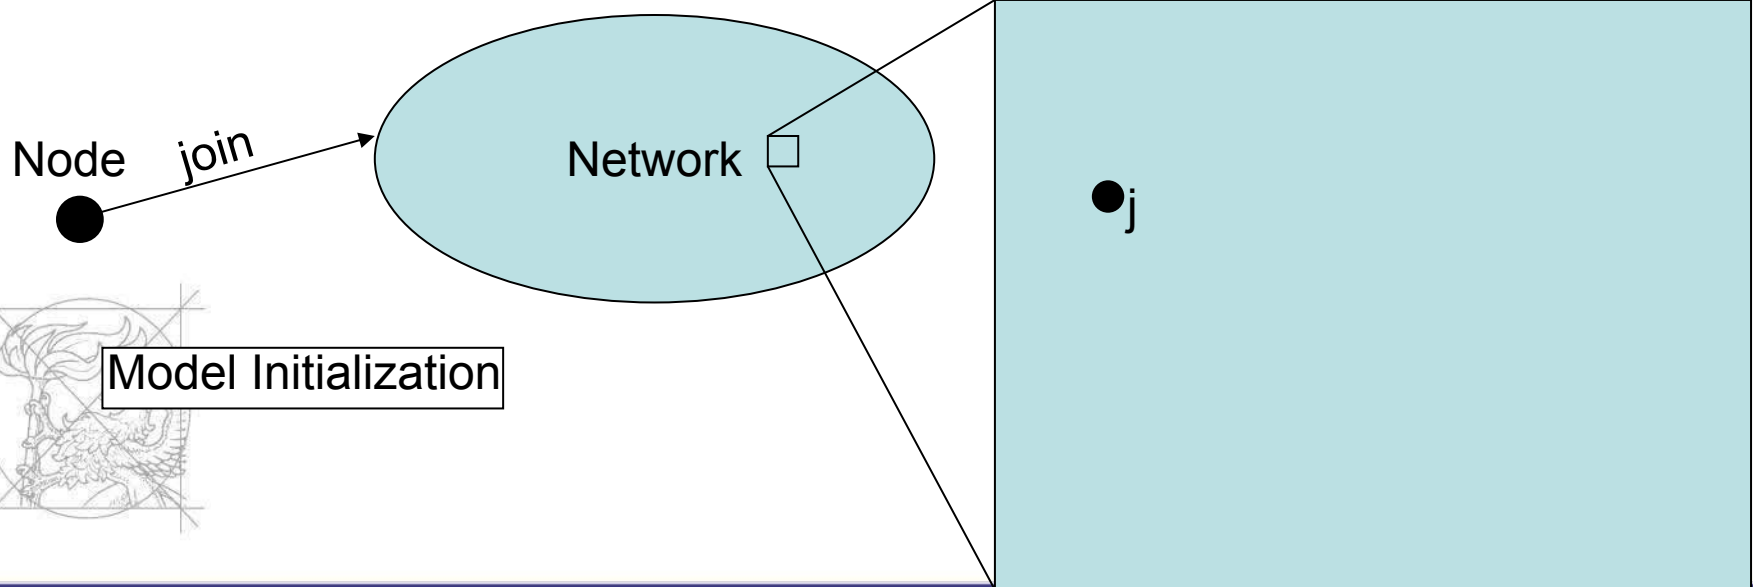
---

**Algorithm 1** P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()  
2: loop  
3:   wait( $\Delta$ )  
4:    $p \leftarrow \text{selectPeer}()$   
5:   send currentModel to  $p$   
6: procedure ONRECEIVEMODEL( $m$ )  
7:    $m \leftarrow \text{updateModel}(m)$   
8:   currentModel  $\leftarrow m$   
9:   modelQueue.add( $m$ )
```

---



# Our Data Mining Framework

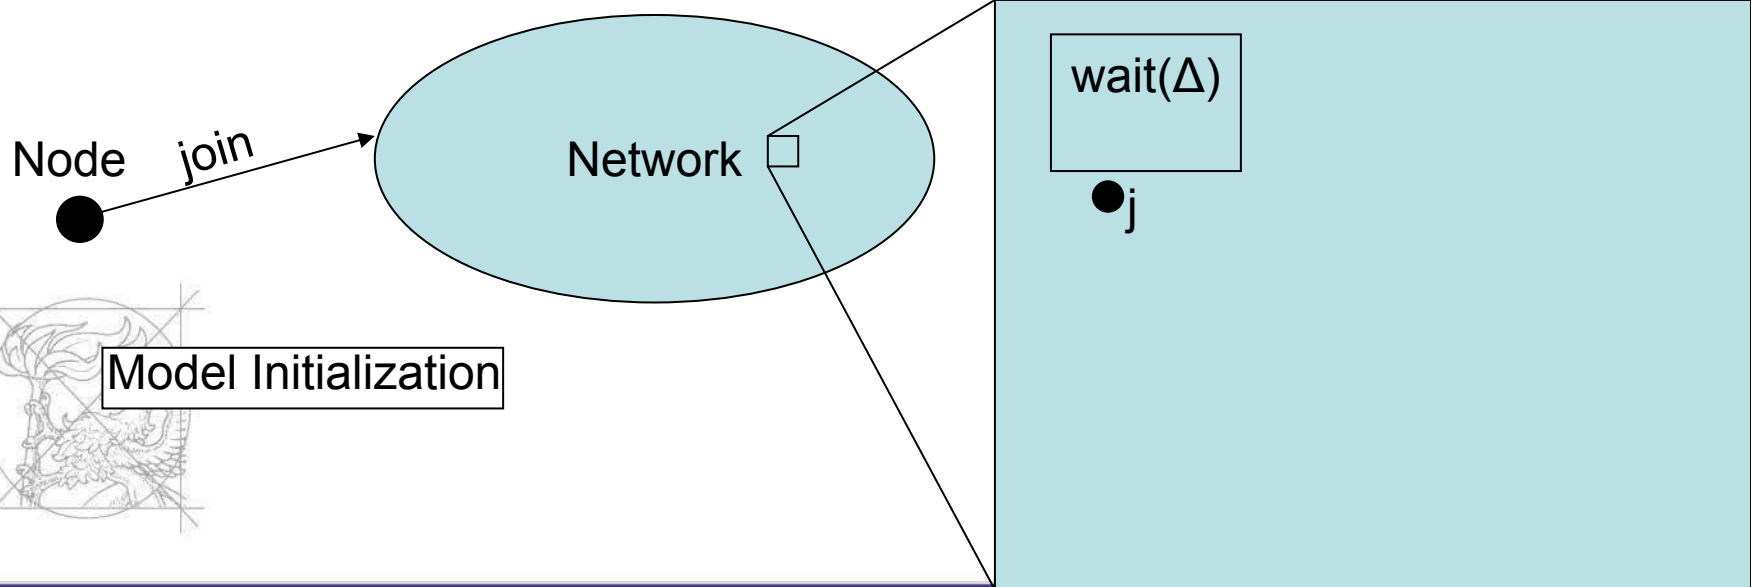
---

**Algorithm 1** P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()  
2: loop  
3:   wait( $\Delta$ )  
4:    $p \leftarrow \text{selectPeer}()$   
5:   send currentModel to  $p$   
6: procedure ONRECEIVEMODEL( $m$ )  
7:    $m \leftarrow \text{updateModel}(m)$   
8:   currentModel  $\leftarrow m$   
9:   modelQueue.add( $m$ )
```

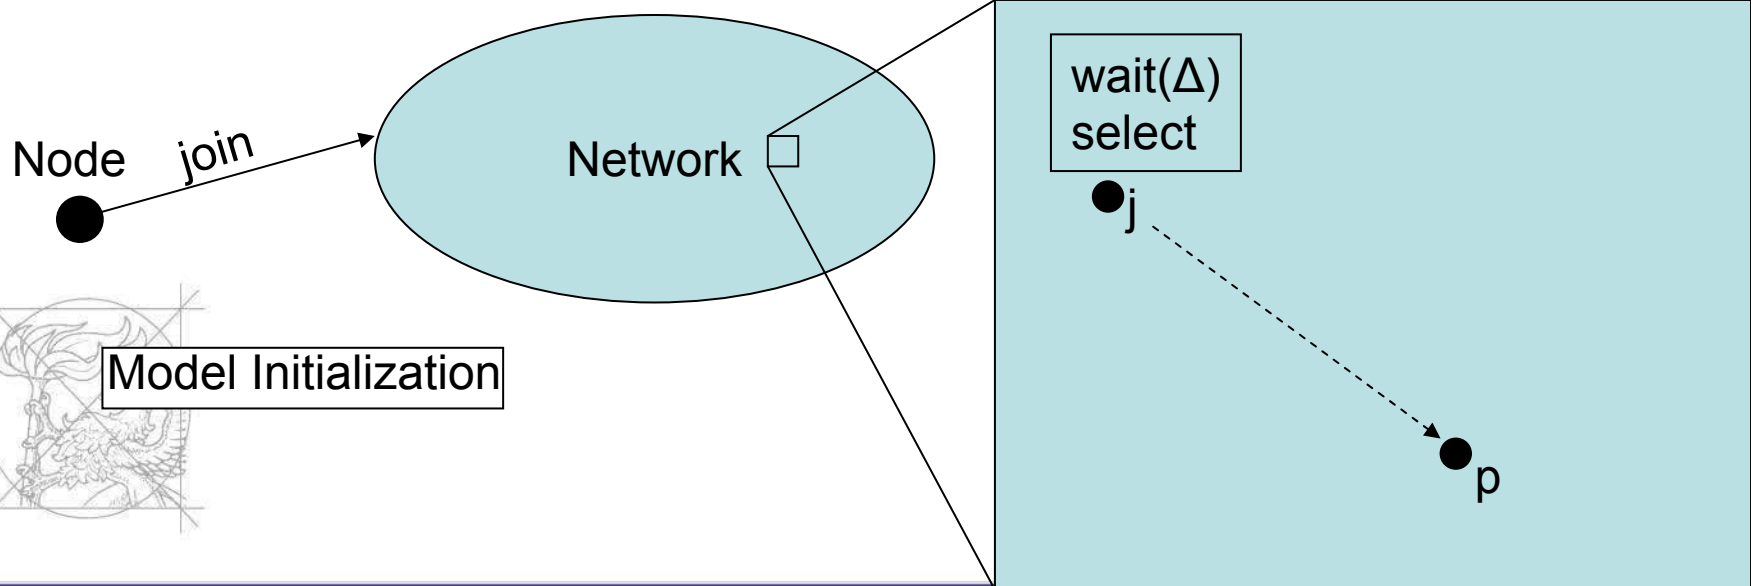
---



# Our Data Mining Framework

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

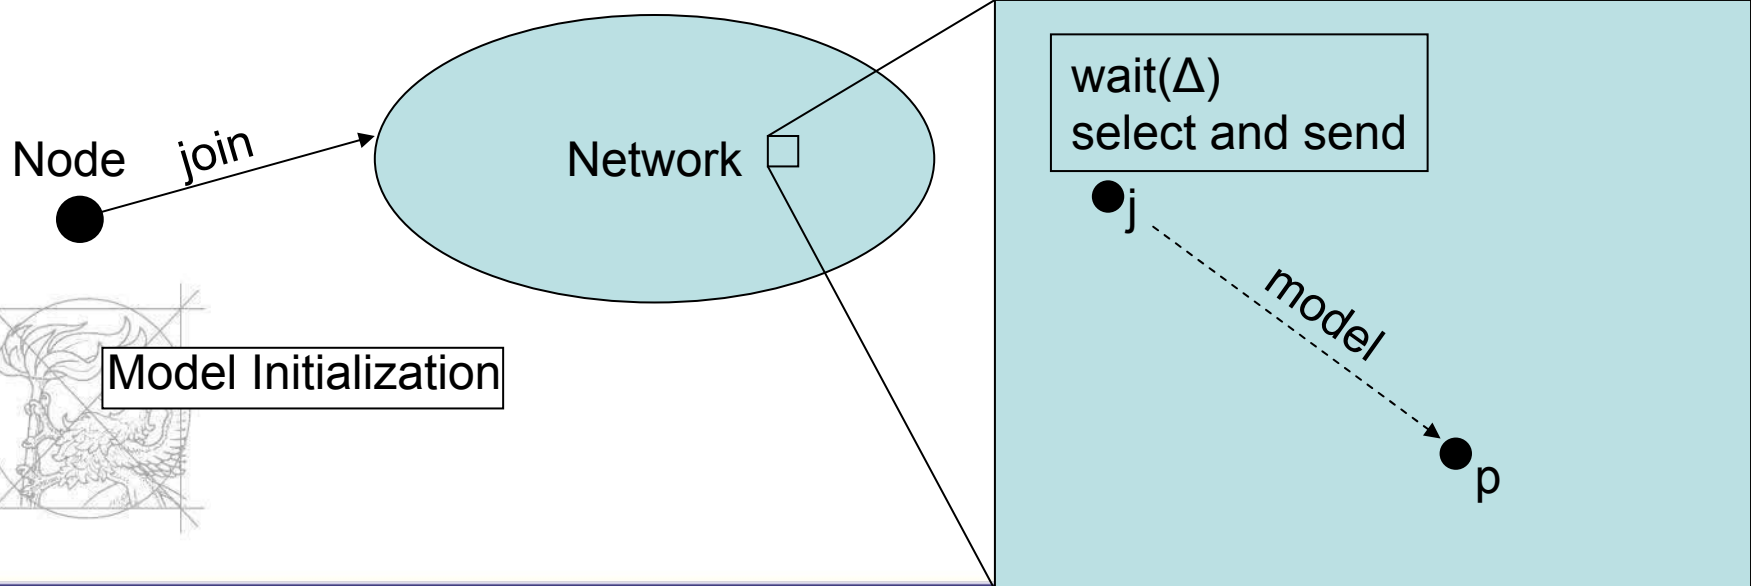
```
1: initModel()  
2: loop  
3:   wait( $\Delta$ )  
4:    $p \leftarrow \text{selectPeer}()$   
5:   send currentModel to  $p$   
6: procedure ONRECEIVEMODEL( $m$ )  
7:    $m \leftarrow \text{updateModel}(m)$   
8:   currentModel  $\leftarrow m$   
9:   modelQueue.add( $m$ )
```



# Our Data Mining Framework

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

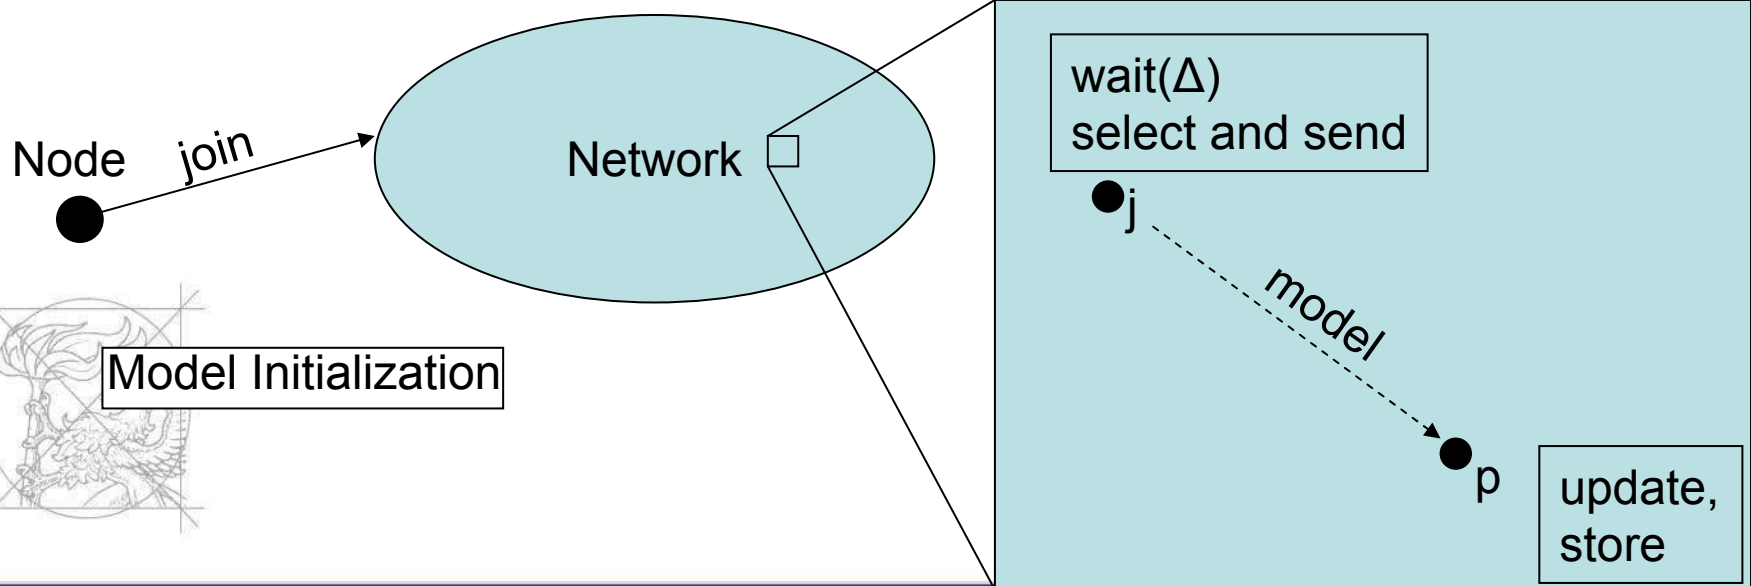
```
1: initModel()  
2: loop  
3:   wait( $\Delta$ )  
4:    $p \leftarrow \text{selectPeer}()$   
5:   send currentModel to  $p$   
6: procedure ONRECEIVEMODEL( $m$ )  
7:    $m \leftarrow \text{updateModel}(m)$   
8:   currentModel  $\leftarrow m$   
9:   modelQueue.add( $m$ )
```



# Our Data Mining Framework

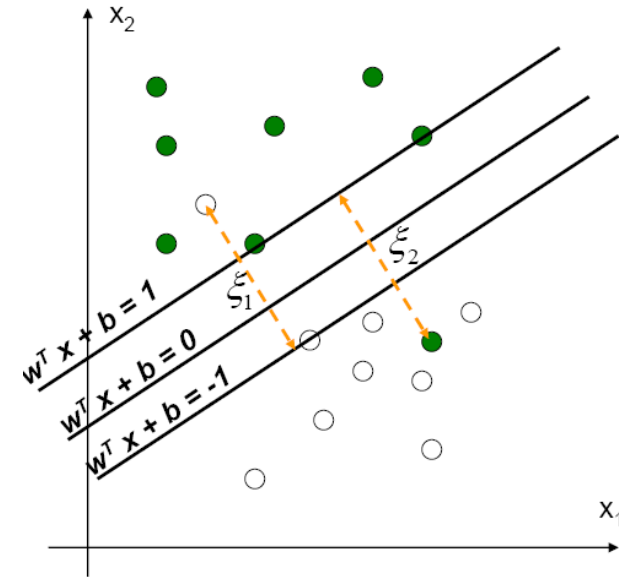
## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

```
1: initModel()  
2: loop  
3:   wait( $\Delta$ )  
4:    $p \leftarrow \text{selectPeer}()$   
5:   send currentModel to  $p$   
6: procedure ONRECEIVEMODEL( $m$ )  
7:    $m \leftarrow \text{updateModel}(m)$   
8:   currentModel  $\leftarrow m$   
9:   modelQueue.add( $m$ )
```



# Support Vector Machines (SVM) through SGD

- Pegasos SVM
  - SGD based ML algorithm
  - It is looking for a separating hyperplane ( $w$ ) that separates examples of the two classes and maximizes the margin



$$\min_{w, b, \xi_i} \quad \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i$$

$$\text{s.t.} \quad y_i(w^T x_i + b) \geq 1 - \xi_i \quad \text{and} \quad \xi_i \geq 0 \quad (\forall i : 1 \leq i \leq n)$$

# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define updateModel and initModel methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATEMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define **updateModel** and initModel methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATERMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define **updateModel** and initModel methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATERMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define updateModel and **initModel** methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATEMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define updateModel and **initModel** methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATEMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# P2Pegasos SVM Algorithm

---

## Algorithm 1 P2P Stochastic Gradient Descent Algorithm

---

```
1: initModel()
2: loop
3:   wait( $\Delta$ )
4:    $p \leftarrow \text{selectPeer}()$ 
5:   send currentModel to  $p$ 
6: procedure ONRECEIVEMODEL( $m$ )
7:    $m \leftarrow \text{updateModel}(m)$ 
8:   currentModel  $\leftarrow m$ 
9:   modelQueue.add( $m$ )
```

---

- Define **updateModel** and **initModel** methods for different machine learning algorithms
- For Pegasos:

---

## Algorithm 2 P2Pegasos

---

```
1: procedure UPDATEMODEL( $m$ )
2:    $\eta \leftarrow 1/(\lambda \cdot m.t)$ 
3:   if  $y \langle m.w, x \rangle < 1$  then
4:      $m.w \leftarrow (1 - \eta\lambda)m.w + \eta yx$ 
5:   else
6:      $m.w \leftarrow (1 - \eta\lambda)m.w$ 
7:    $m.t \leftarrow m.t + 1$ 
8:   return  $m$ 
9: procedure INITMODEL
10:   $m.t \leftarrow 0$ 
11:   $m.w \leftarrow (0, \dots, 0)^T$ 
12:  send model( $m$ ) to self
```

---



# Local Prediction

- Every node has a bounded queue of the last few received models for making prediction
  - Using a single model

```
1: procedure PREDICT( $x$ )  
2:    $w \leftarrow \text{currentModel}$   
3:   return sign( $\langle w, x \rangle$ )
```

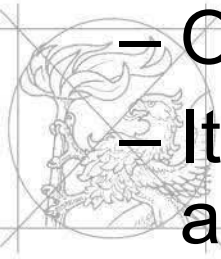
- Using multiple models (majority voting)

```
4: procedure VOTEDPREDICT( $x$ )  
5:   pRatio  $\leftarrow$  0  
6:   for  $m \in \text{modelQueue}$  do  
7:     if sign( $\langle m.w, x \rangle$ )  $\geq$  0 then  
8:       pRatio  $\leftarrow$  pRatio + 1  
9:   return sign(pRatio/modelQueue.size() - 0.5)
```



# Communication Cost

- The cost of communication is exactly the same as the cost of other gossip based algorithms
  - Every node sends a message (model) to a uniform randomly selected node in each  $\Delta$  time period
  - $O(n)$  in the network
  - $O(1)$  for a node
  - It includes the cost of peer sampling service as well



# Algorithm Summary

- Random walk based P2P SGD method
  - Sending model instead of selecting sample
- Nodes collect model(s)
  - Simple and voted prediction
- Local prediction
  - No extra communication cost
- Each node runs the same algorithm
- Asynchronous

# Experimental Setup

- PeerSim as a simulation environment
- NewsCast for peer sampling service
- The 10 latest models for Voting prediction
- Baselines: Pegasos, SVMLight
- Datasets: Iris, Reuters, SpamBase, Malicious
- Scenarios:
  - No failure
  - Drop only: 0.5 probability
  - Delay only: uniform random from  $[\Delta, 10\Delta]$
  - Churn only: from approx. Log-normal dist.
  - All failures

# Database Properties

	Iris1	Iris2	Iris3	Reuters	SpamBase	Malicious10
Training set size	90	90	90	2000	4140	2155622
Test set size	10	10	10	600	461	240508
Number of features	4	4	4	9947	57	10
Classlabel ratio	50/50	50/50	50/50	1300/1300	1813/2788	792145/1603985
Pegasos 20000 iter.	0	0	0	0.025	0.111	0.080 (0.081)
Pegasos 1000 iter.	0	0	0.4	0.057	0.137	0.095 (0.060)
SVMLight	0	0	0.1	0.027	0.074	0.056 (–)

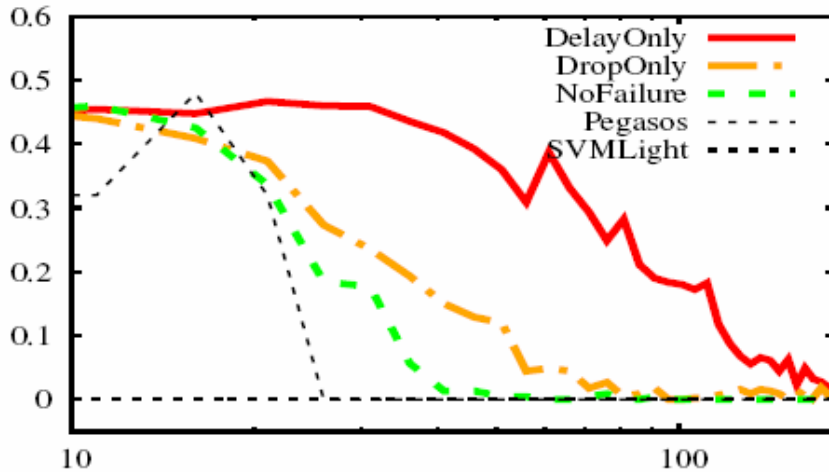
- Different types of datasets:
  - Small – large num. of samples
  - Small – large num. of features
- Split Iris and reduced Malicious features
- Performance of baseline algorithms

# Effects of Different Failures

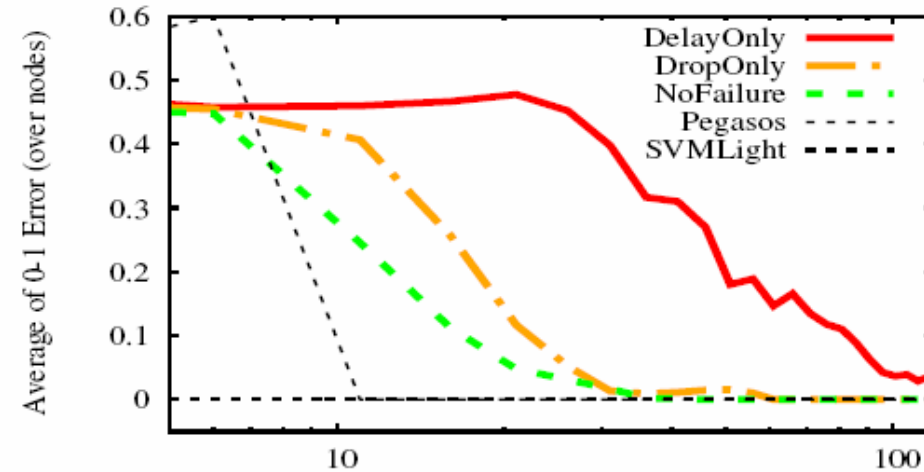
Average of 0-1 Error (over nodes)

al Intelligence

Iris1

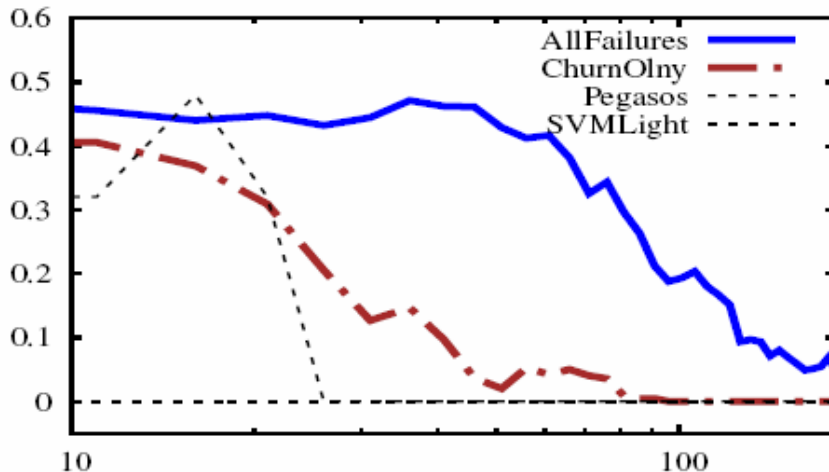


Iris2

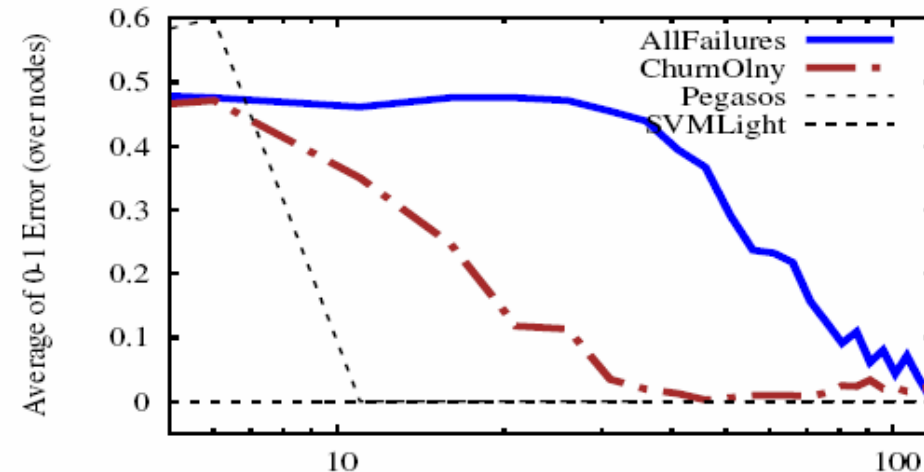


Iris1

Average of 0-1 Error (over nodes)



Iris2

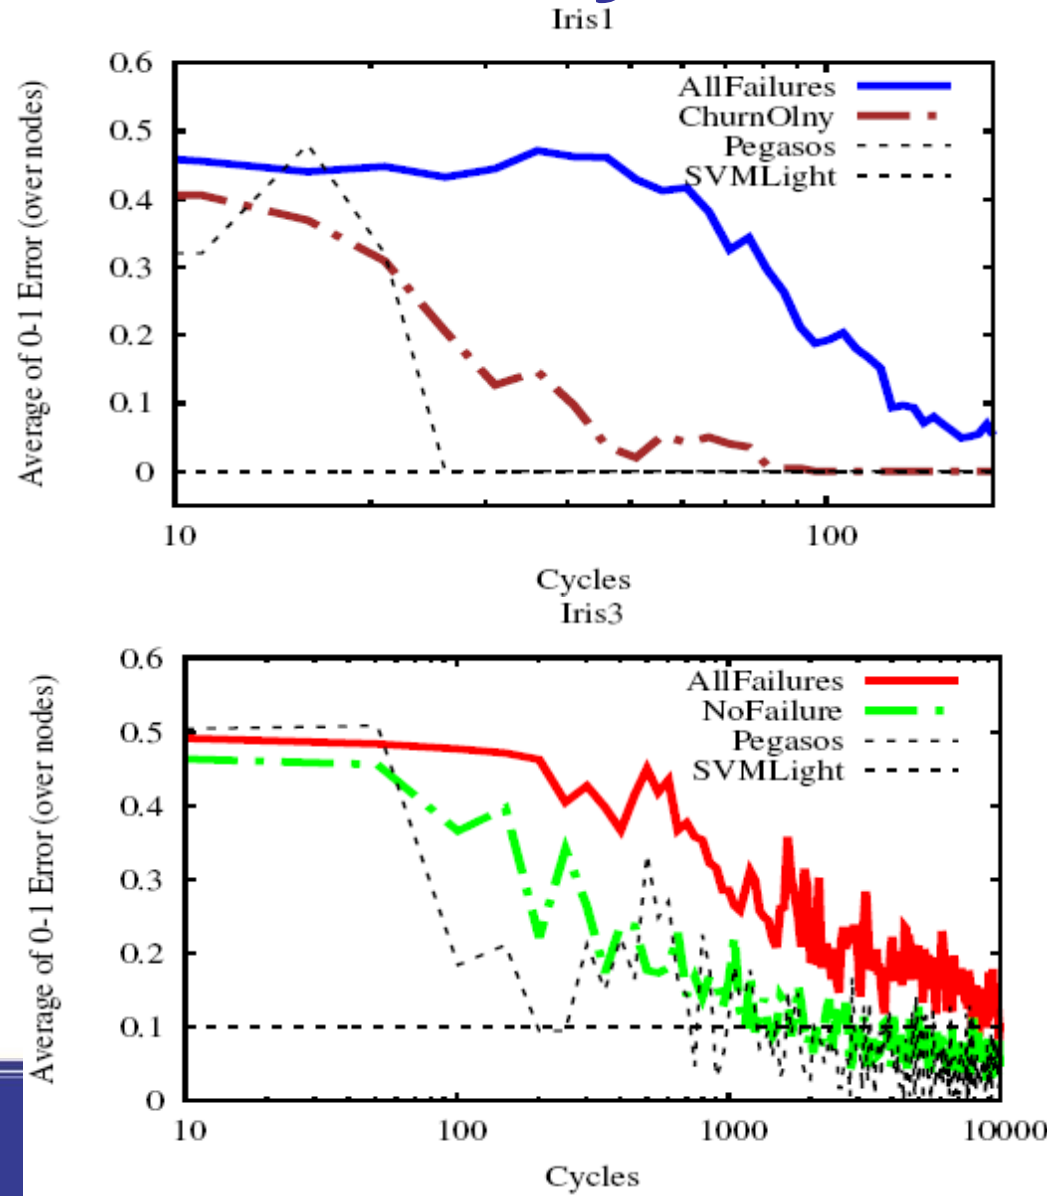


Cycles

Cycles

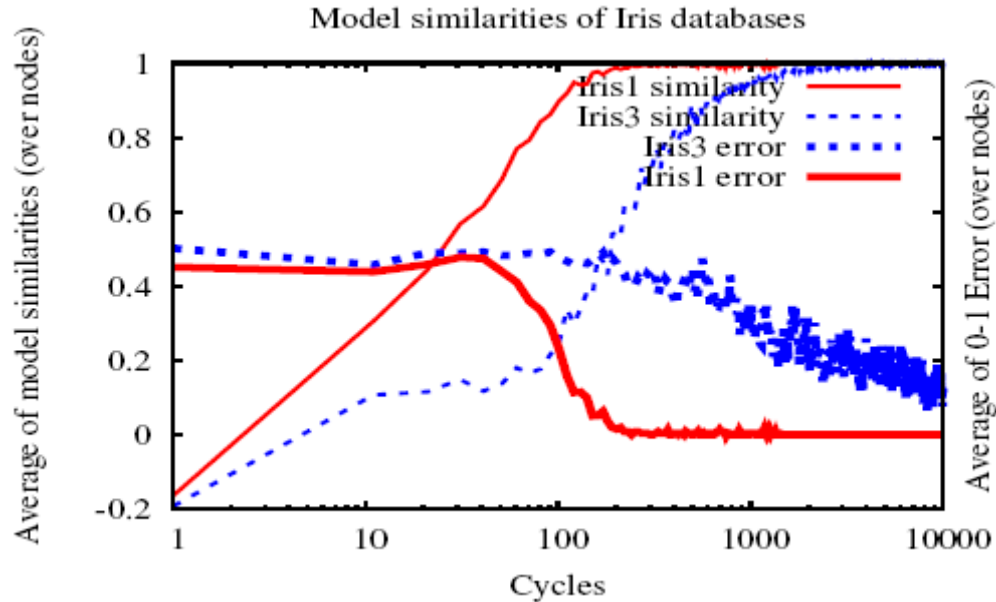
# Size vs. Learnability

- **No relation** between DB size and learnability
- Learnability depends on DB patterns rather than the size



# Convergence

- Clearly shows the correlation between model performance and similarity
- The model similarity grows along with the performance
  - the models converge to the same optimum rather than get only more similar

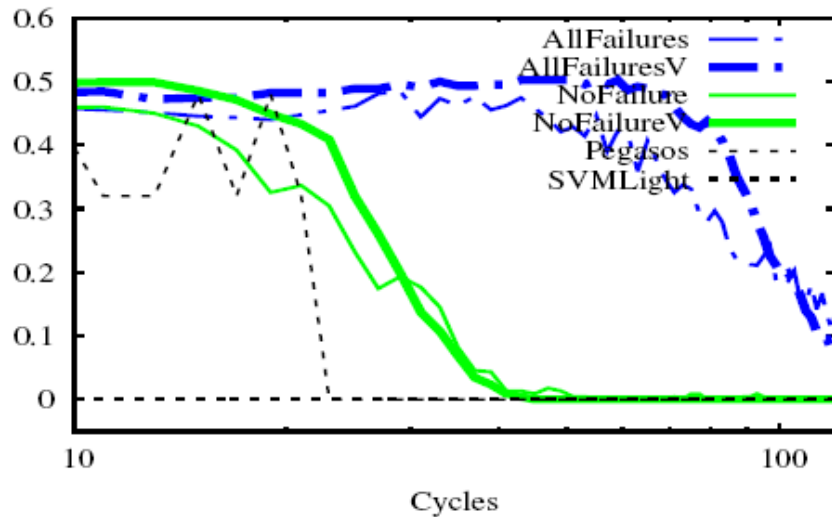


# Results on Large Datasets

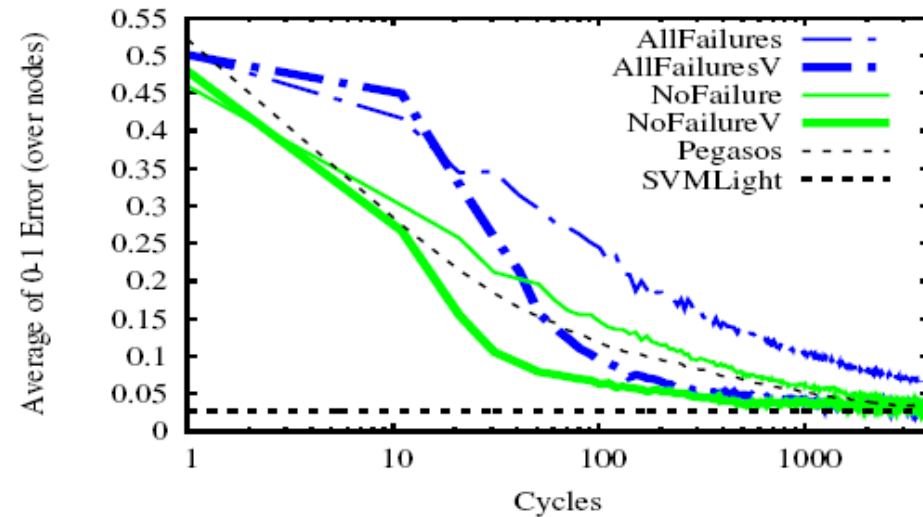
Average of 0-1 Error (over nodes)

ial Intelligence

Iris1

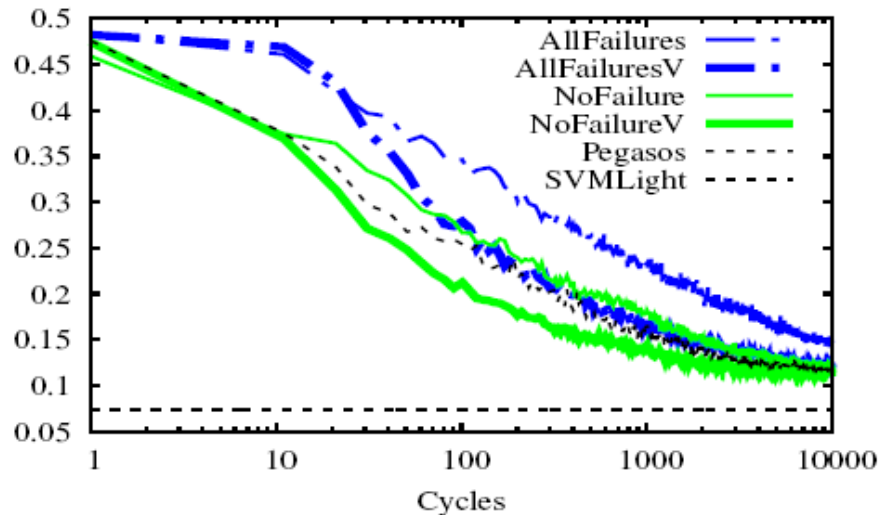


Reuters

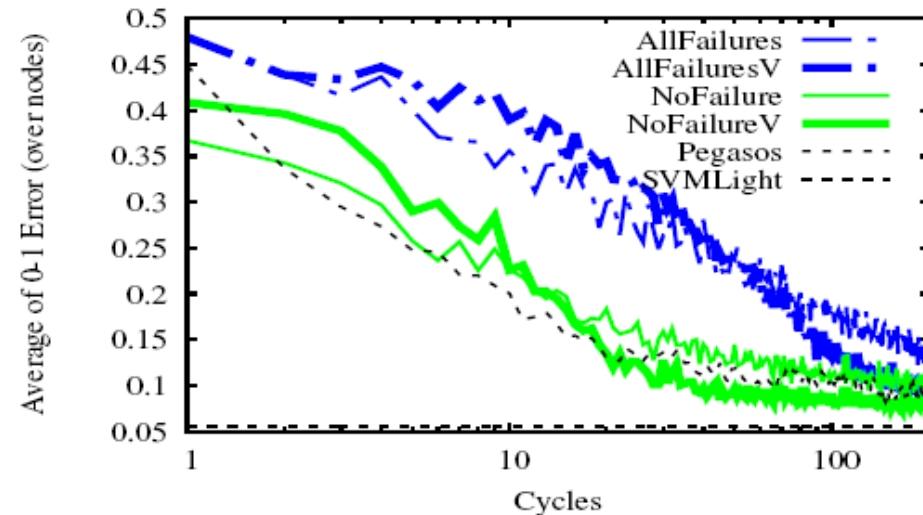


SpamBase

Average of 0-1 Error (over nodes)



Malicious URLs



# Summary

- P2P SGD framework was presented
- P2P Pegasos SVM algorithm was implemented and tested on different datasets
- Almost the same performance in a P2P environment as in a centralized one
- The algorithm works well with extreme communication failures as well