Proactive Drive Failure Prediction for Large Scale Storage Systems

Bingpeng Zhu, Gang Wang, Xiaoguang Liu, Dianming Hu, Sheng Lin, Jingwei Ma

Nankai-Baidu Joint Laboratory

Parallel and Distributed Software Technology Lab

Outline

- What is the problem?
- How to solve the problem?
  - BP neural network
  - Support Vector Machine (SVM)
- Experimental Results

The problem

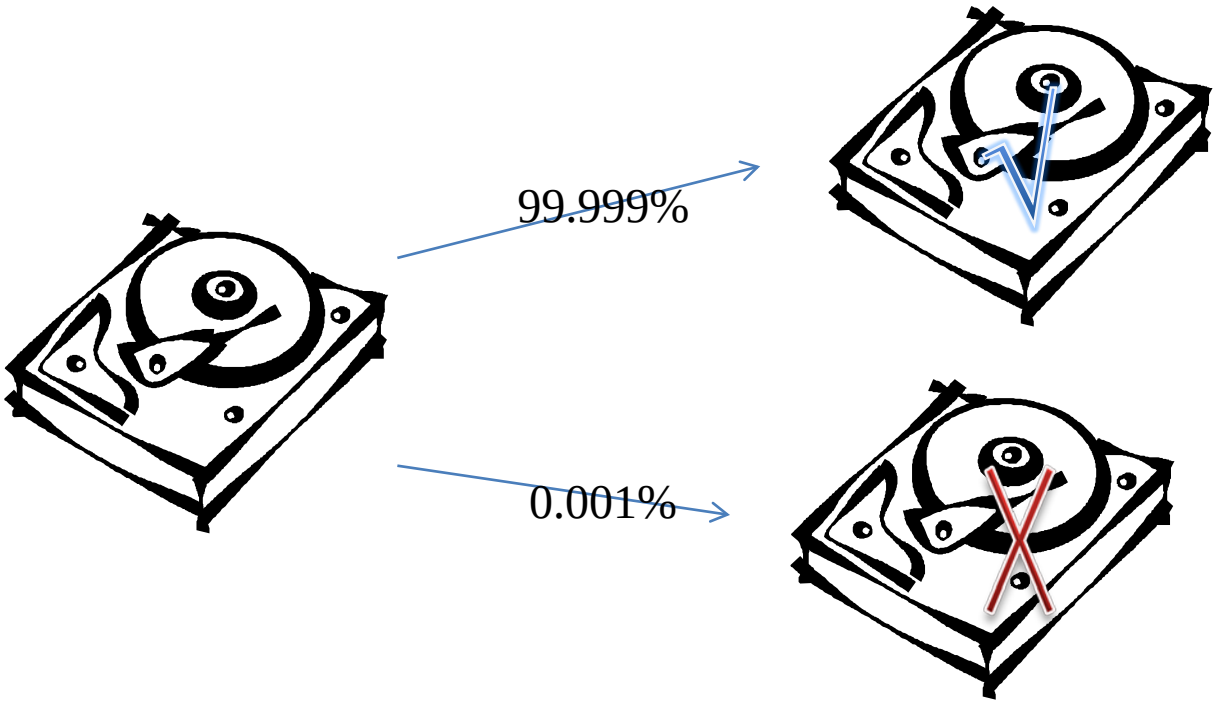

The problem

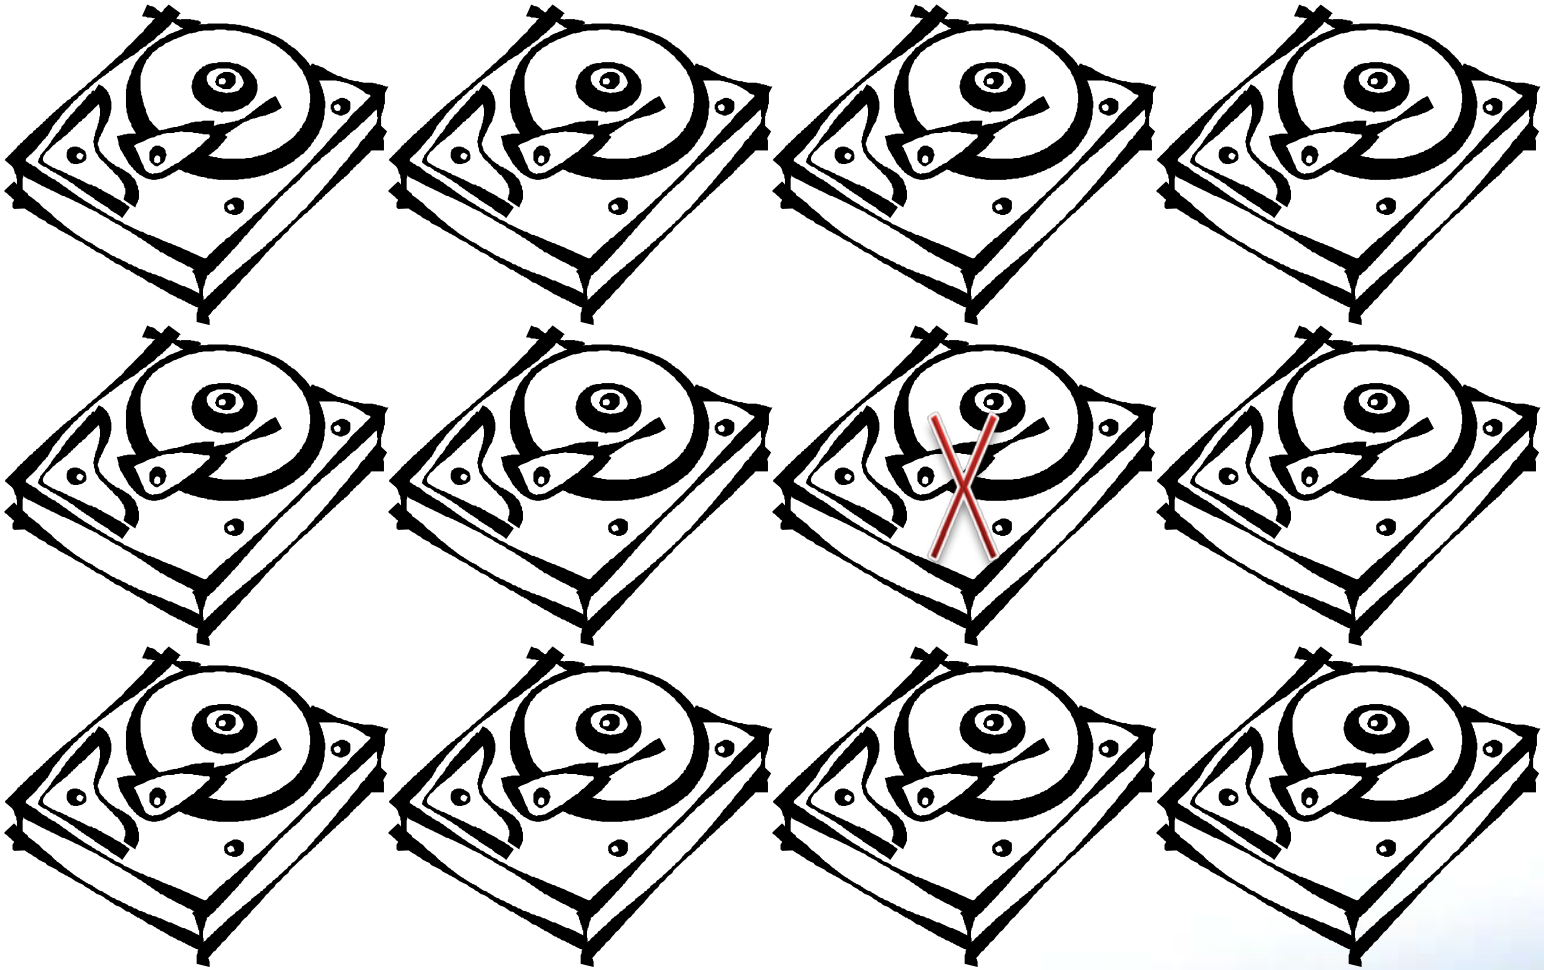

The problem

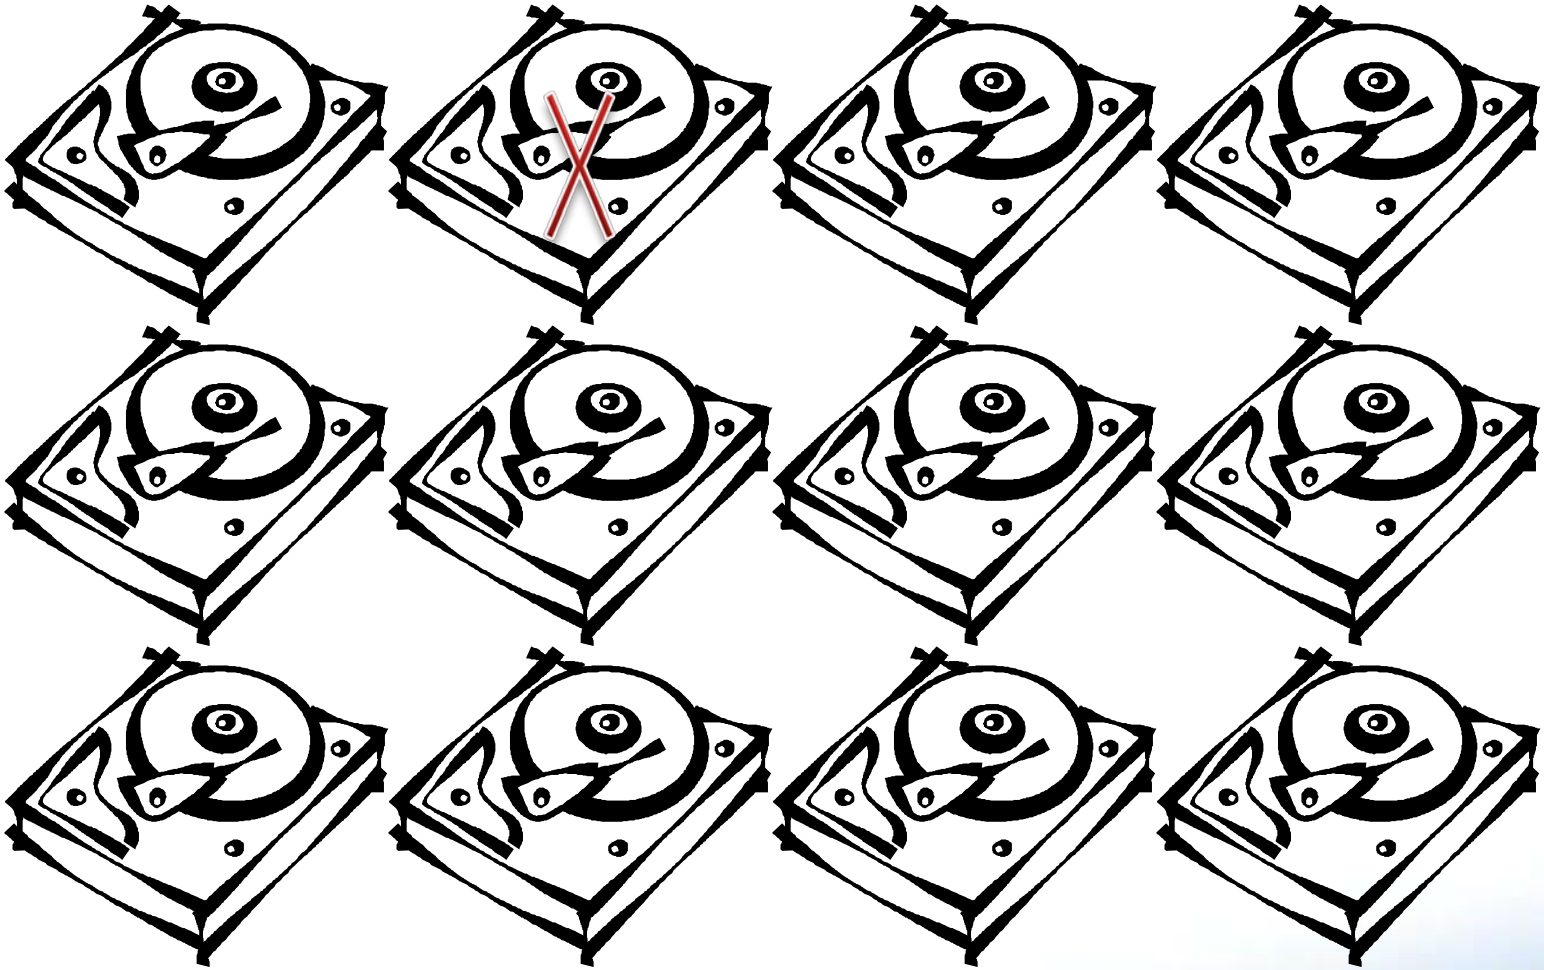

Code and Recovery

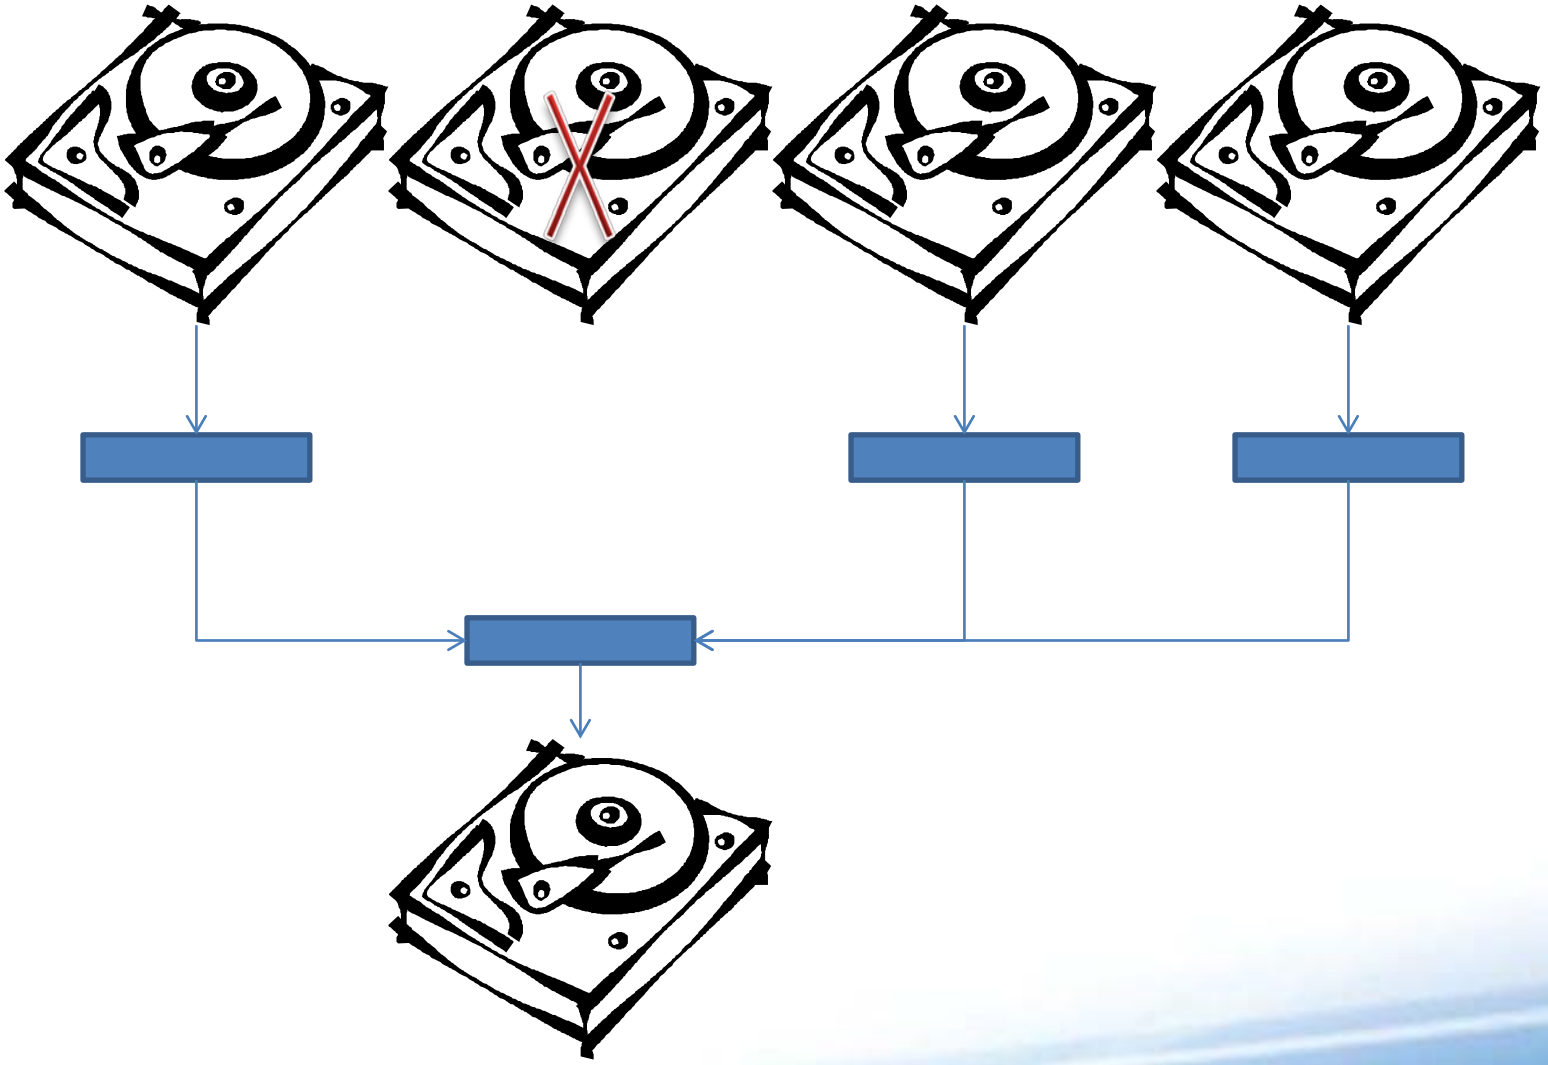

Failure prediction

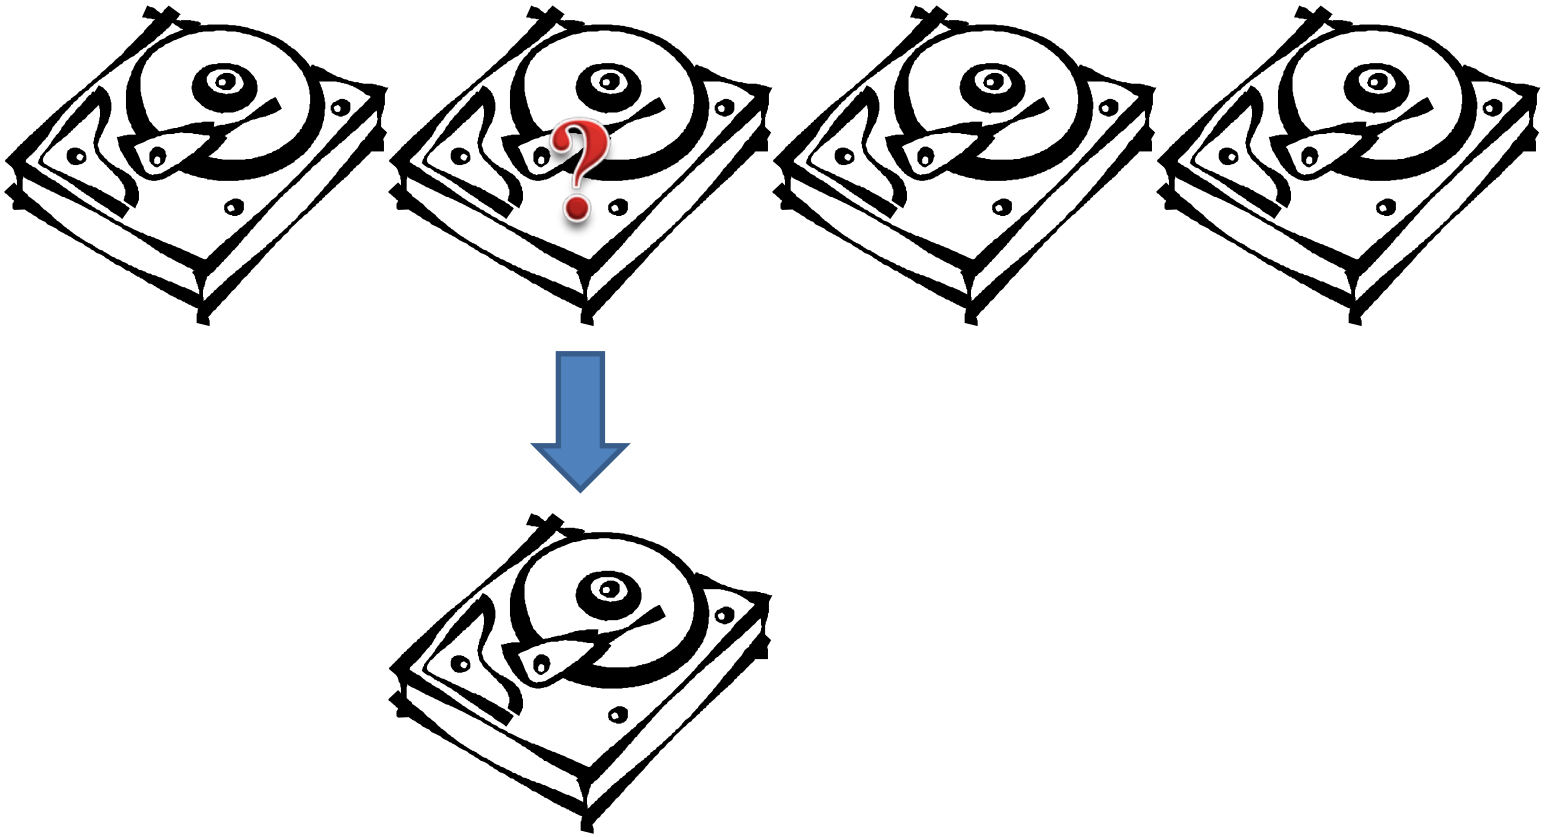

How to solve the problem

- SMART alone
  - Failure detection rate of 3-10% with 0.1% false alarm rate (FAR)
- Our methods
  - Back propagation (BP) neural network
  - Improved Support Vector Machine (SVM)

Support Vector Machine

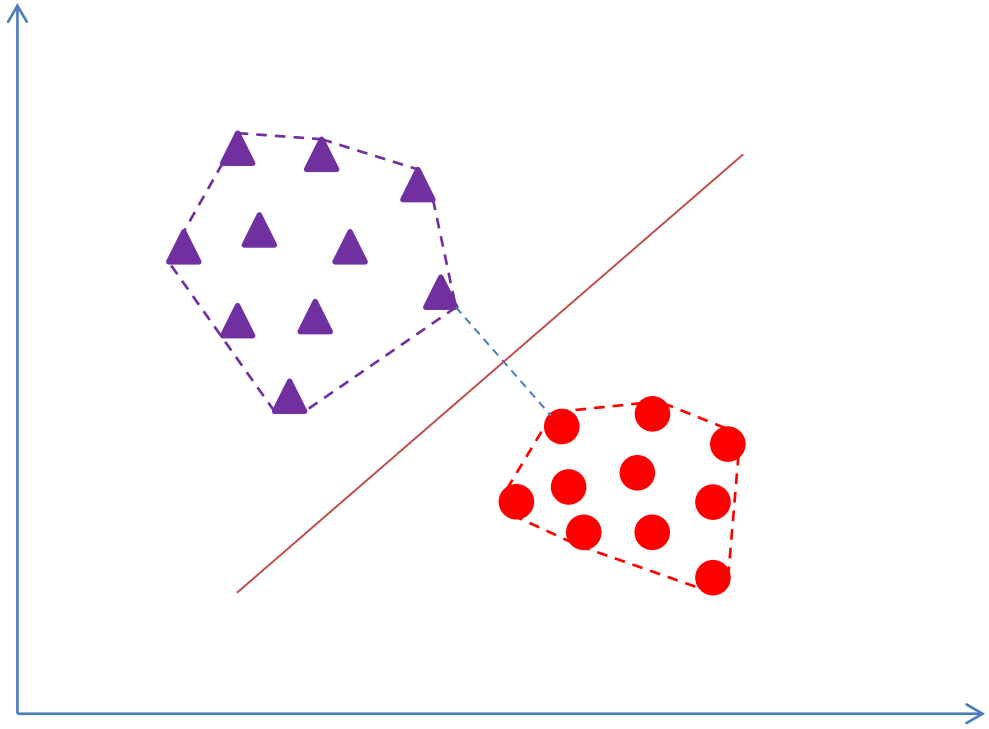

BP neural network

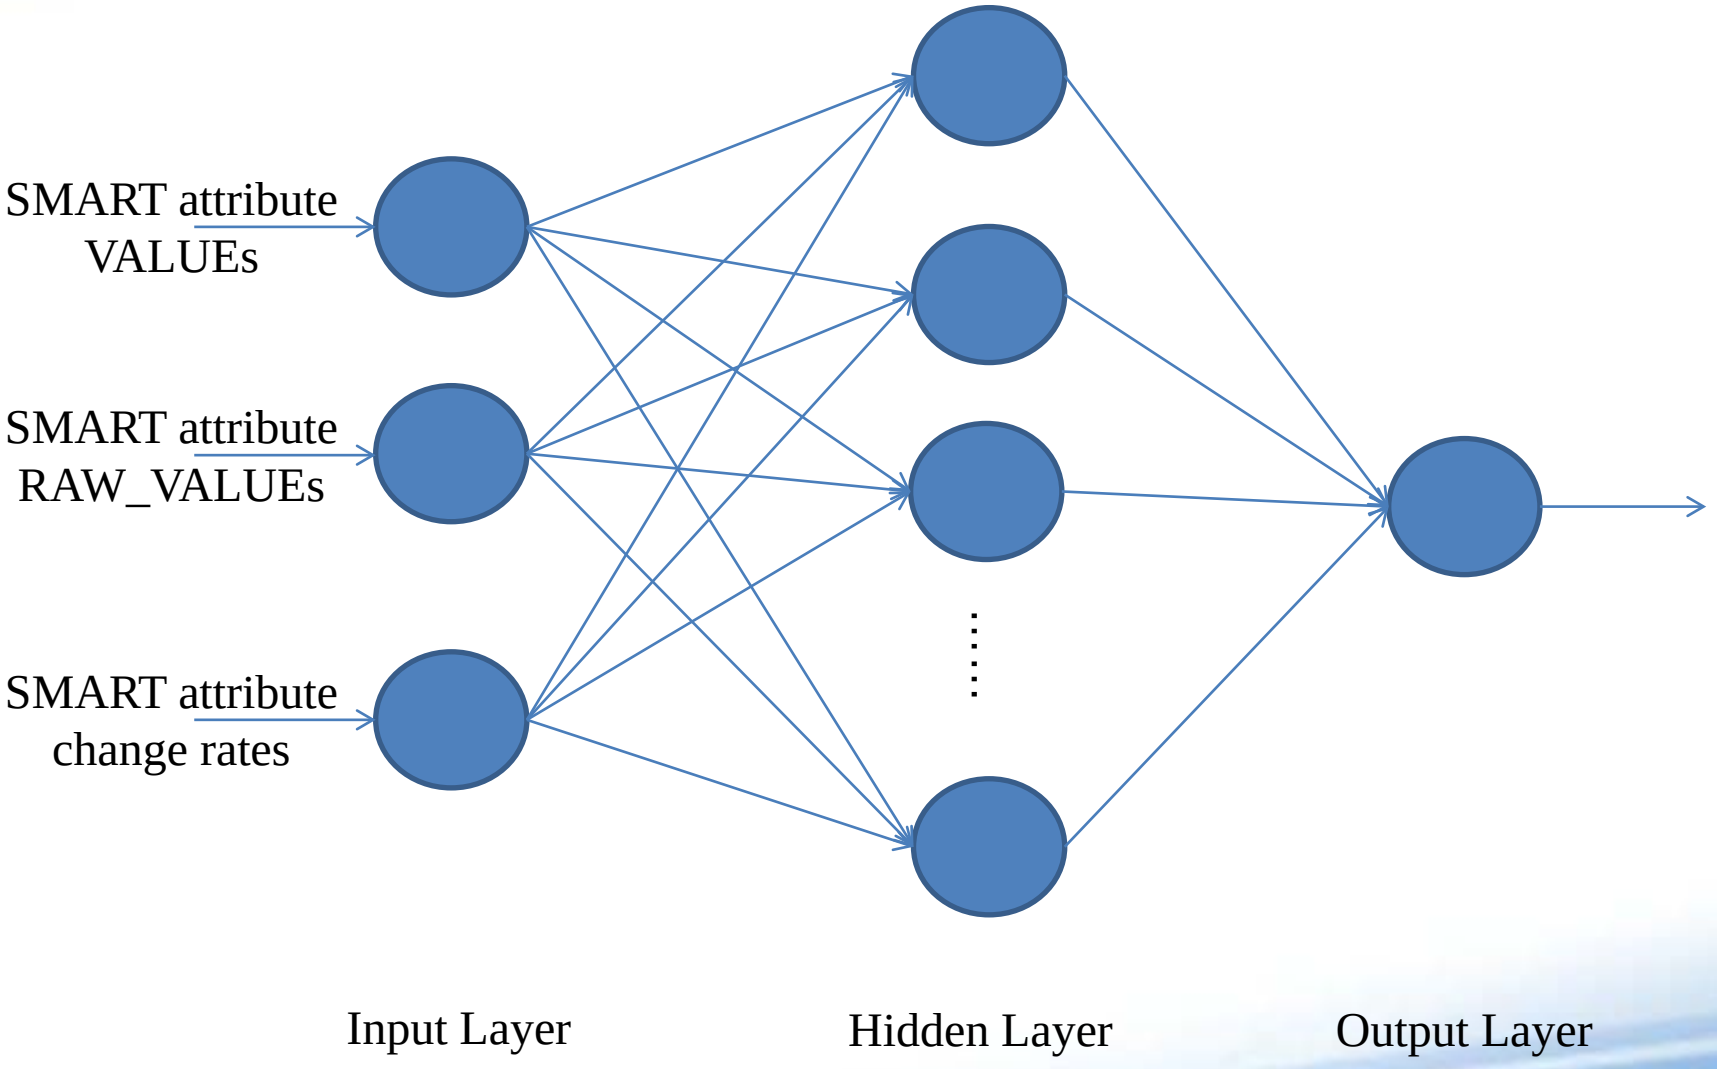

Features

Attributes	
ID #	**Attribute Name**
1	Raw Read Error Rate
3	Spin Up Time
5	Reallocated Sectors Count
7	Seek Error Rate
9	Power On Hours
187	Reported Uncorrectable Errors
189	High Fly Writes
194	Temperature Celsius
195	Hardware ECC Recovered
197	Current Pending Sector Count

2 RAW Attributes
5
197

Change Rates	
1	CR of Attribute
5	CR of Attribute
187	CR of Attribute
195	CR of Attribute
197	CR of Attribute
5	CR of RAW Attribute
197	CR of RAW Attribute

Experimental Results

- Dataset
  - Seagate ST31000524NS
  - 433 failed drives
  - 22962 good drives
- Experimental setup
  - 70% in training set
  - 30% in test set

SVM result

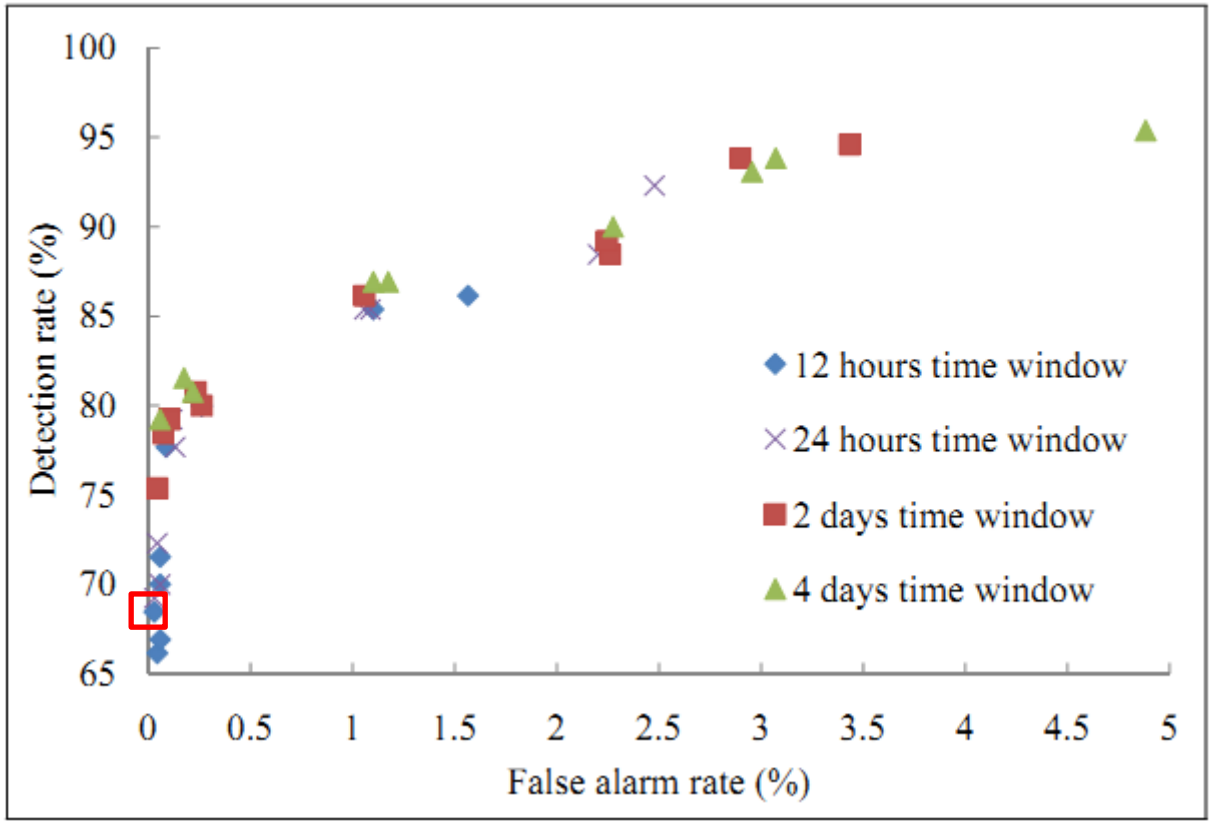

Failure predication performance of SVM models

The lead time

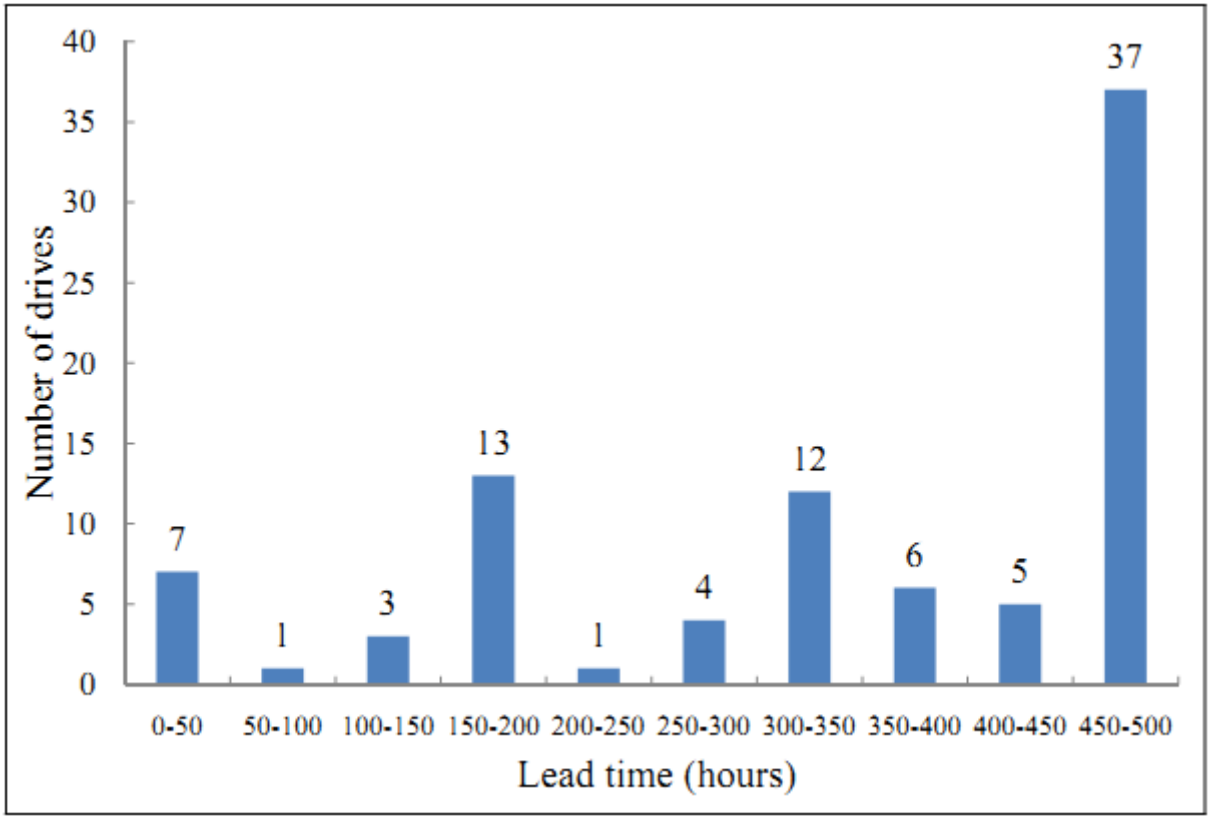

Distribution of lead time of SVM model

Effectiveness of change rate features

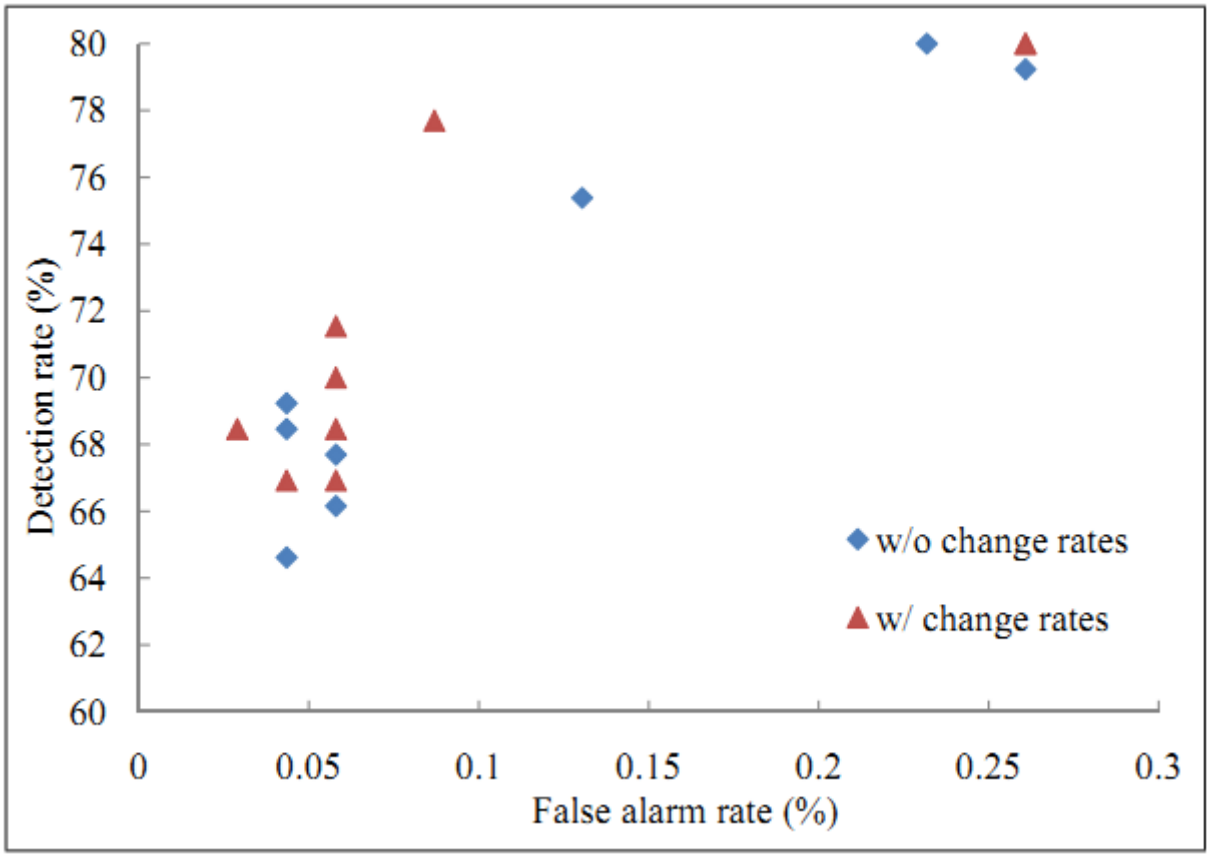

The effectiveness of change rate features for SVM

Results of BP neural network

Prediction results of BP neural network models

Time window	FAR (%)	Detection rate (%)	Lead time (hours)
12 hours	0.48	94.62	360.4
24 hours	1.14	97.69	355.9
2 days	1.39	99.23	357.0
4 days	2.26	100.0	356.8

Voting based failure detection

Algorithm 1 Voting-based failure detection algorithm

Input: The sample set $S[1..t]$ of the drive, the BP prediction model $BP()$ which returns 0 if the input sample is classified as good and 1 otherwise, and the voter turnout N

Output: good or failed

1: **Begin**
2: $C[1..N] = 0$
3: **for** $i = 1$ to t **do**
4: $\quad C[((i-1) \mod N) + 1] = BP(S[i])$
5: $\quad$ **if** $\sum_{j=1}^{N} C[j] > N/2$ **then**
6: $\quad\quad$ **return** failed
7: $\quad$ **end if**
8: **end for**
9: **return** good
10: **End**

AdaBoost-enhanced BP network with voting-based detection

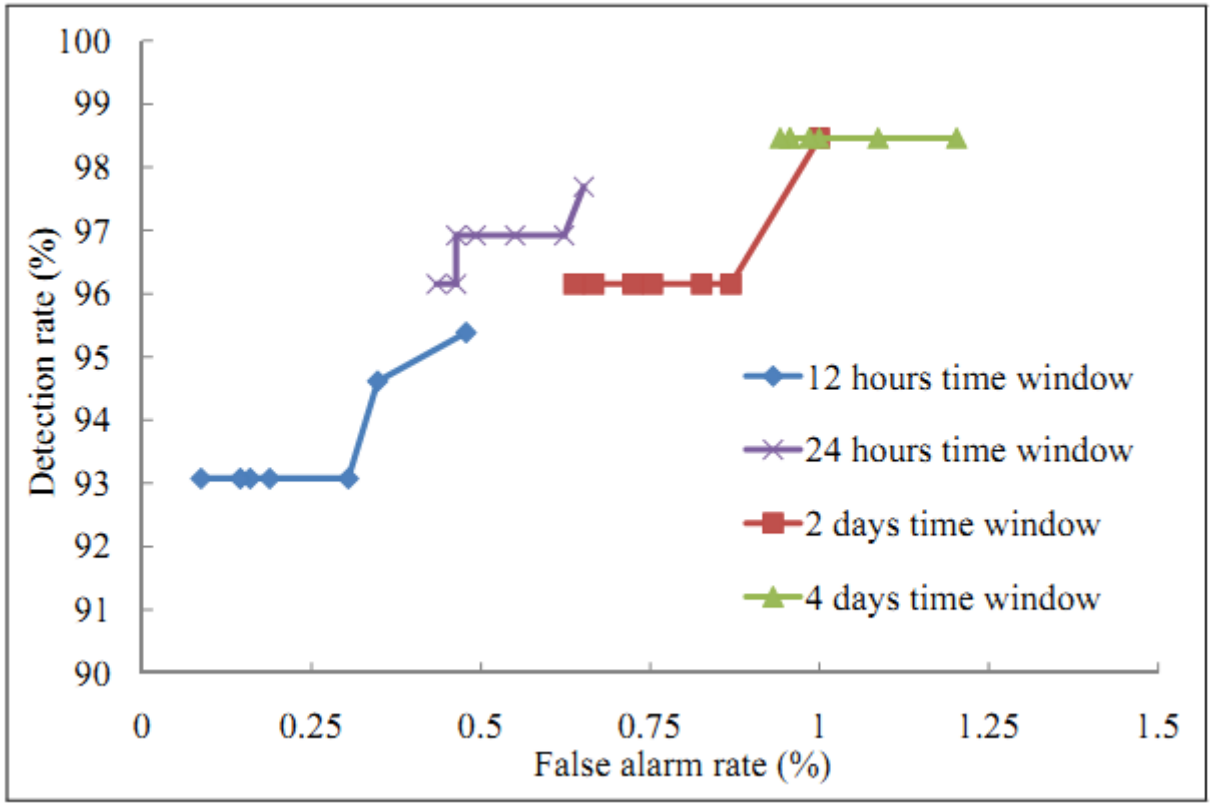

Failure prediction performance of AdaBoost-enhanced BP network models using voting-based detection method

MTTDL

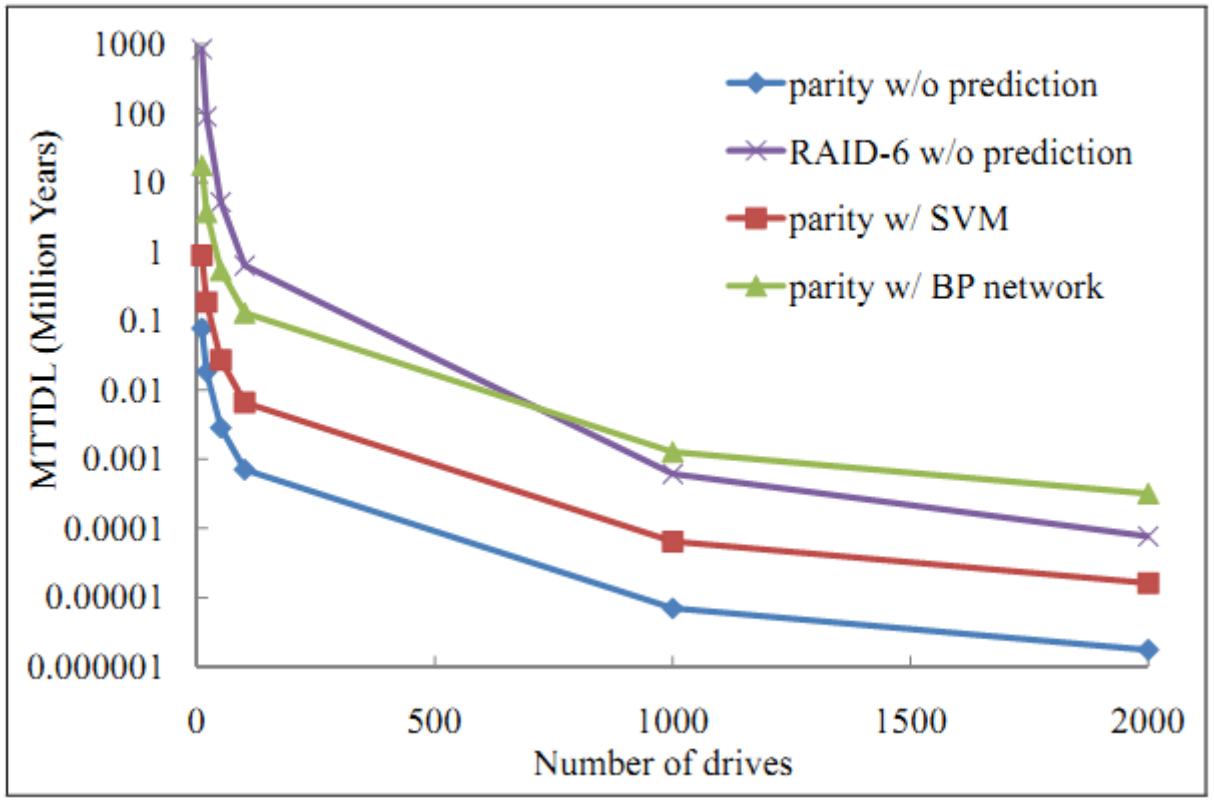

MTTL of RAID with varying sizes

Conclusion

- Two methods to predict disk failures
  - BP neural network
  - Improved SVM
- Our prediction models achieve much higher accuracy
  - FAR (SVM 0.03%), detection rate (BP network: 95%)

Thank You