

Bazaar: Strengthening user reputations in online marketplaces

Ansley Post^{†*}

Vijit Shah[‡]

Alan Mislove[‡]

[‡]*Northeastern University*

[†]*MPI-SWS/Rice University*

^{*}*Now at Google*

March 31, 2011, NSDI'11

Online marketplaces



Online marketplaces:

Sites allowing users to buy/sell goods

Among **most successful Web sites**

E.g., eBay, Overstock, Amazon Marketplace
eBay alone: \$60B in 2009



Allows buyers and sellers to connect

Regardless of location

Enable esoteric products to find a market

Democratized commerce



But, **known to suffer from fraud**

Identities and reputations

Sites support reputations for identities

Feedback from others interacted with

Buyers use reputations

Reputable sellers get better prices



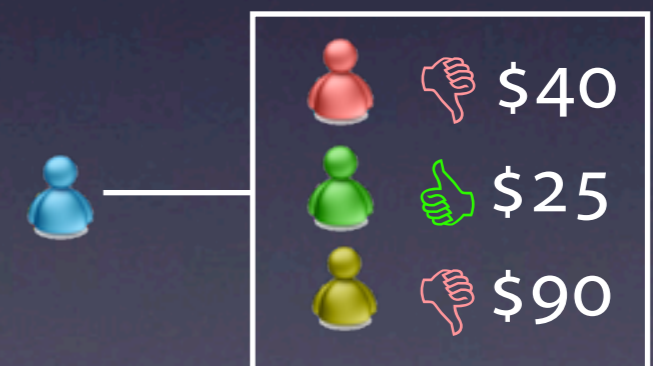
Feedback profile

Complicating detail:

Accounts often "free" to create

Requires only solving CAPTCHA

Can be used to defraud...



Feedback profile

Manipulating reputations for fraud

Can create identities to

Whitewash (erase bad behavior)

Collude (with other attackers)

Sybil attacks (create multiple accounts)

Can observe fraud taking place

Search for “positive feedback guaranteed”

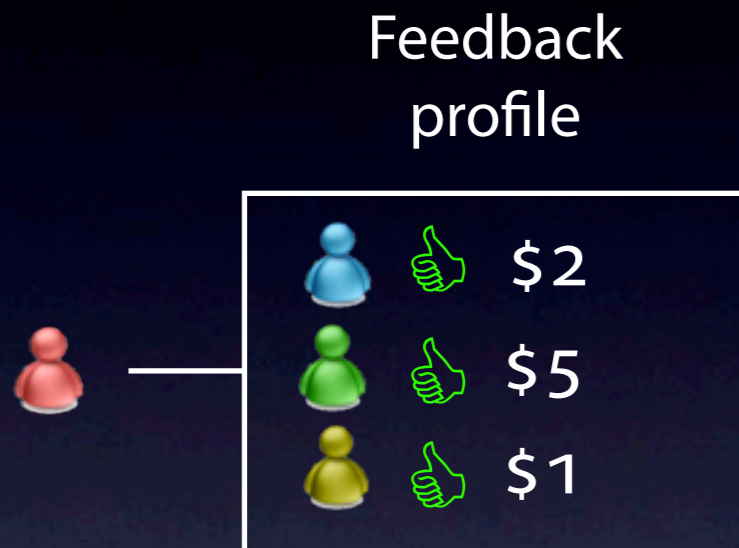
Undermines usefulness of marketplace

Significant **monetary losses**

Recent arrest of malicious user

Stole \$717k from 5,000 users

Used >250 accounts



Manipulating reputations for fraud

Can create identities to

Whitewash (erase bad behavior)

Collude (with other attackers)

Sybil attacks (create multiple accounts)

Can observe fraud taking place

Search for “positive feedback guaranteed”


Undermines usefulness of marketplace

Significant **monetary losses**

Recent arrest of malicious user

Stole \$717k from 5,000 users

Used >250 accounts



Manipulating reputations for fraud

Can create identities to

Whitewash (erase bad behavior)

Collude (with other attackers)

Sybil attacks (create multiple accounts)

Can observe fraud taking place

Search for “positive feedback guaranteed”

Undermines usefulness of marketplace

Significant **monetary losses**

Recent arrest of malicious user

Stole \$717k from 5,000 users

Used >250 accounts

Feedback
profile



		\$2
		\$5
		\$1
		\$300
		\$90
		\$50

Alternate approaches

amazon.com
Marketplace

uBid

Make joining difficult
Limits applicability, usefulness

ebay

Using brokers, escrow
Only feasible for expensive items

craigslist

Requiring in-person transaction
Restricts buyer/seller population

ebay

Providing insurance
Spreads cost of fraud to all users

Others in paper...

Bazaar: A new approach

New approach to strengthening user reputations

Provides strong bounds on fraud

Works in conjunction with existing marketplace

Assumes **same feedback system as today**

No additional monetary cost

No strong identities

Insight: Successful transactions represent shared risk

Buyer and seller more likely to enter into future transactions

Outline

~~1. Motivation~~

2. Bazaar design

3. Challenges faced

4. Evaluation

Risk network

Reputations calculated using **risk network**

Buyer satisfied → two identities linked

Weighted by amount of transaction

Multiple transactions additive

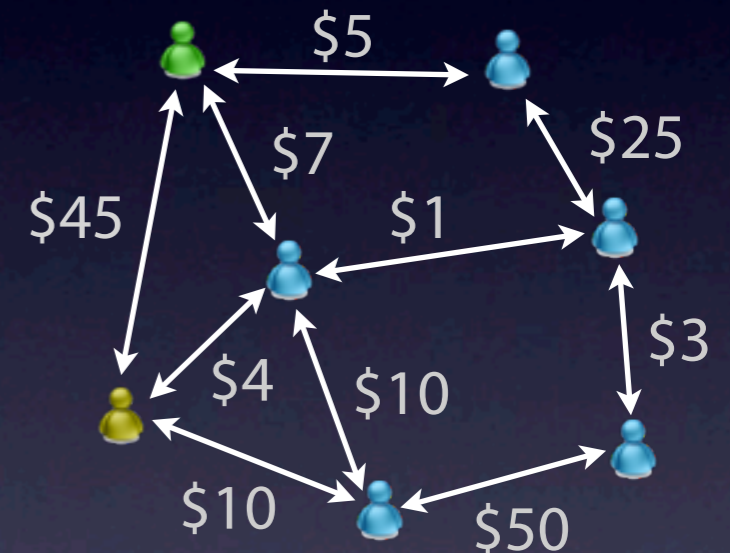
Risk network automatically generated

Users need not even know about it

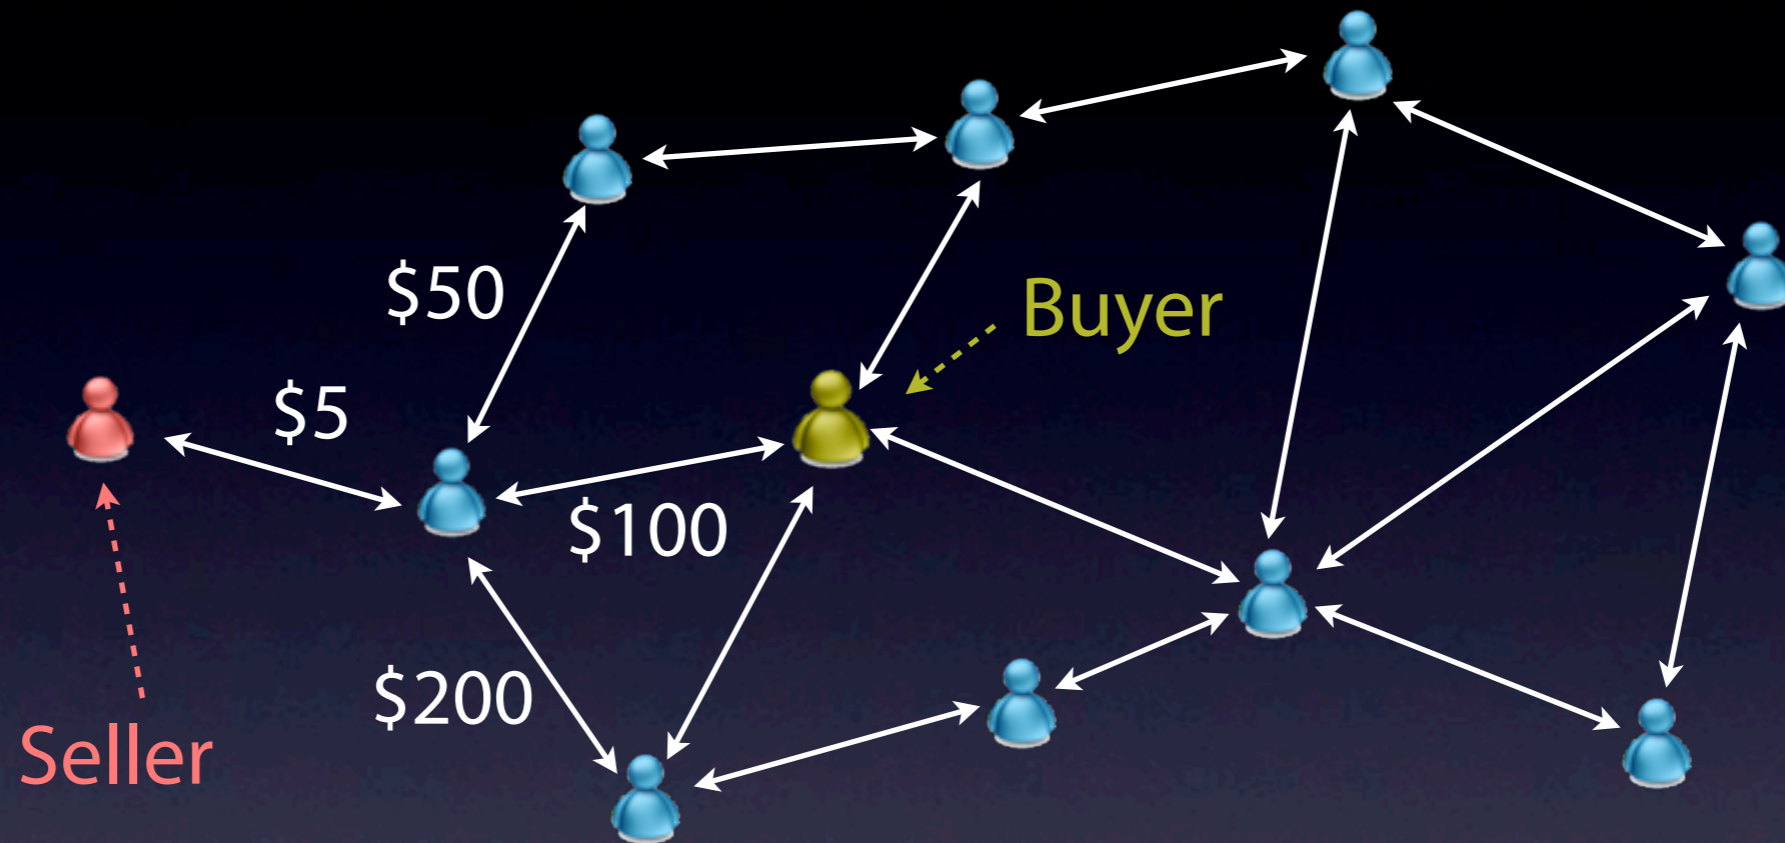
Site operator maintains risk network

Can be used to gauge risk between identities

Model: **Query Bazaar before purchase**



Fraud detection with max-flow



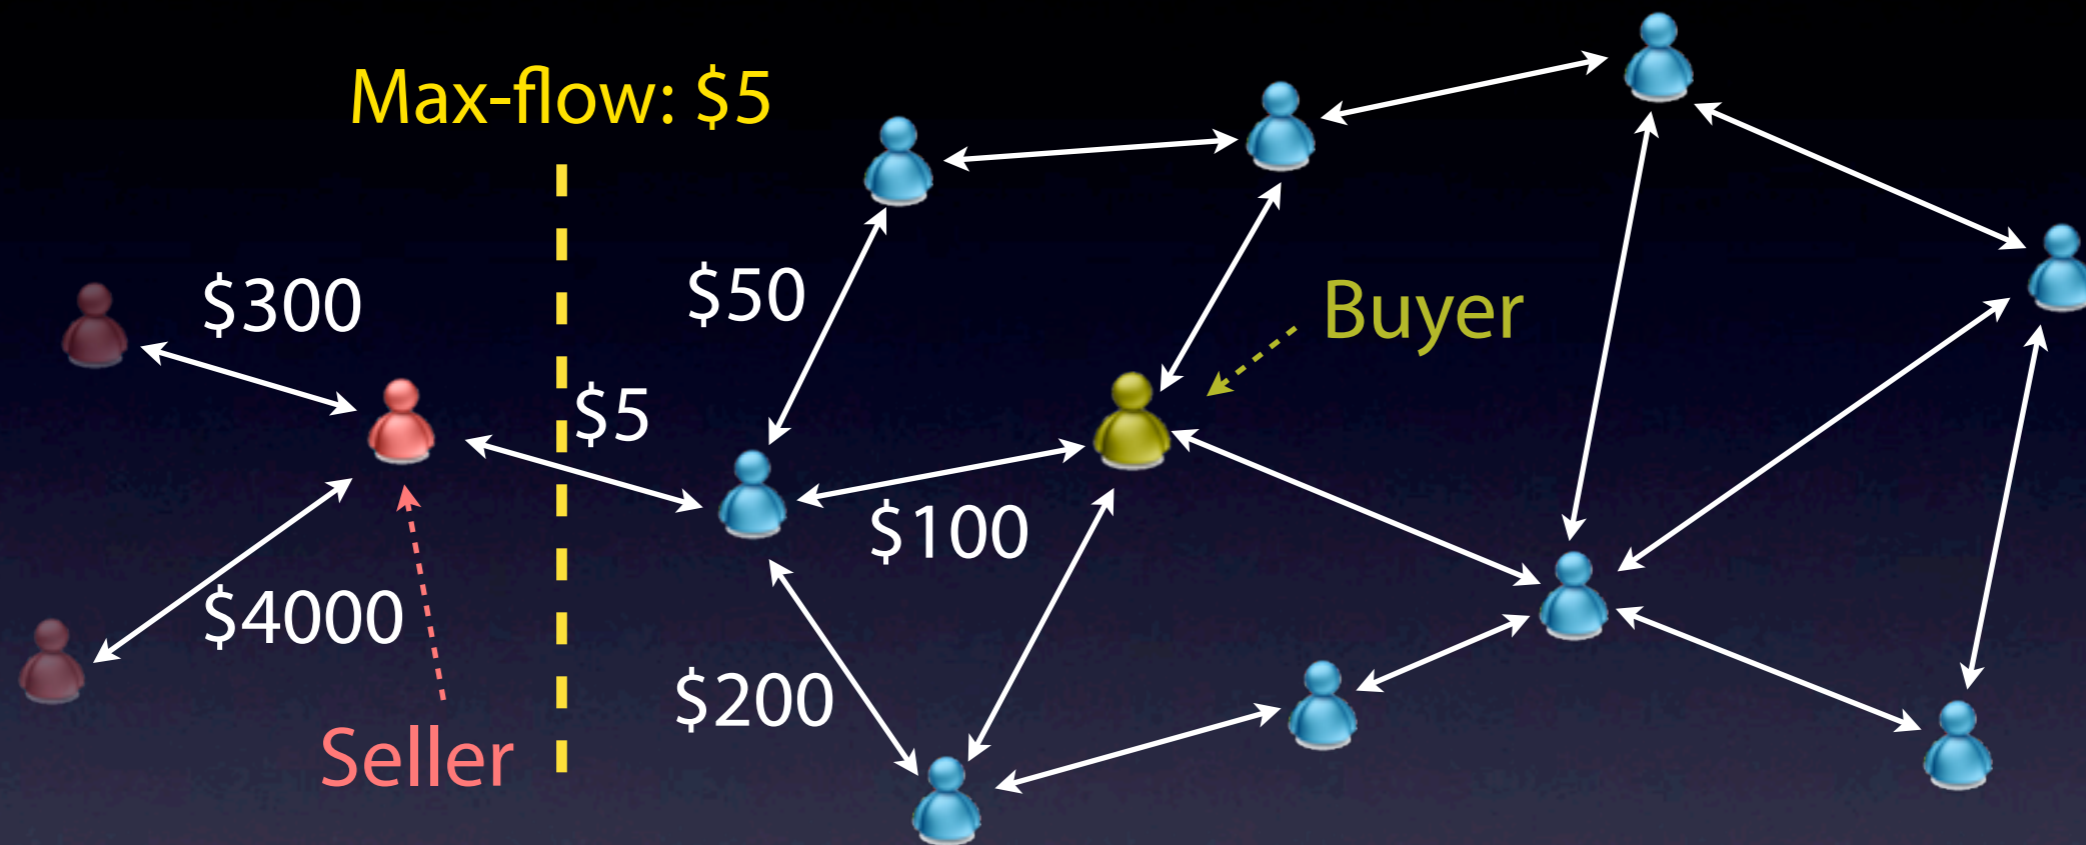
Site operator queries Bazaar before purchase

Bazaar **calculates max-flow between buyer and seller**

If max-flow lower than potential transaction, **flag as fraudulent**

Otherwise, wait for feedback from buyer

Fraud detection with max-flow



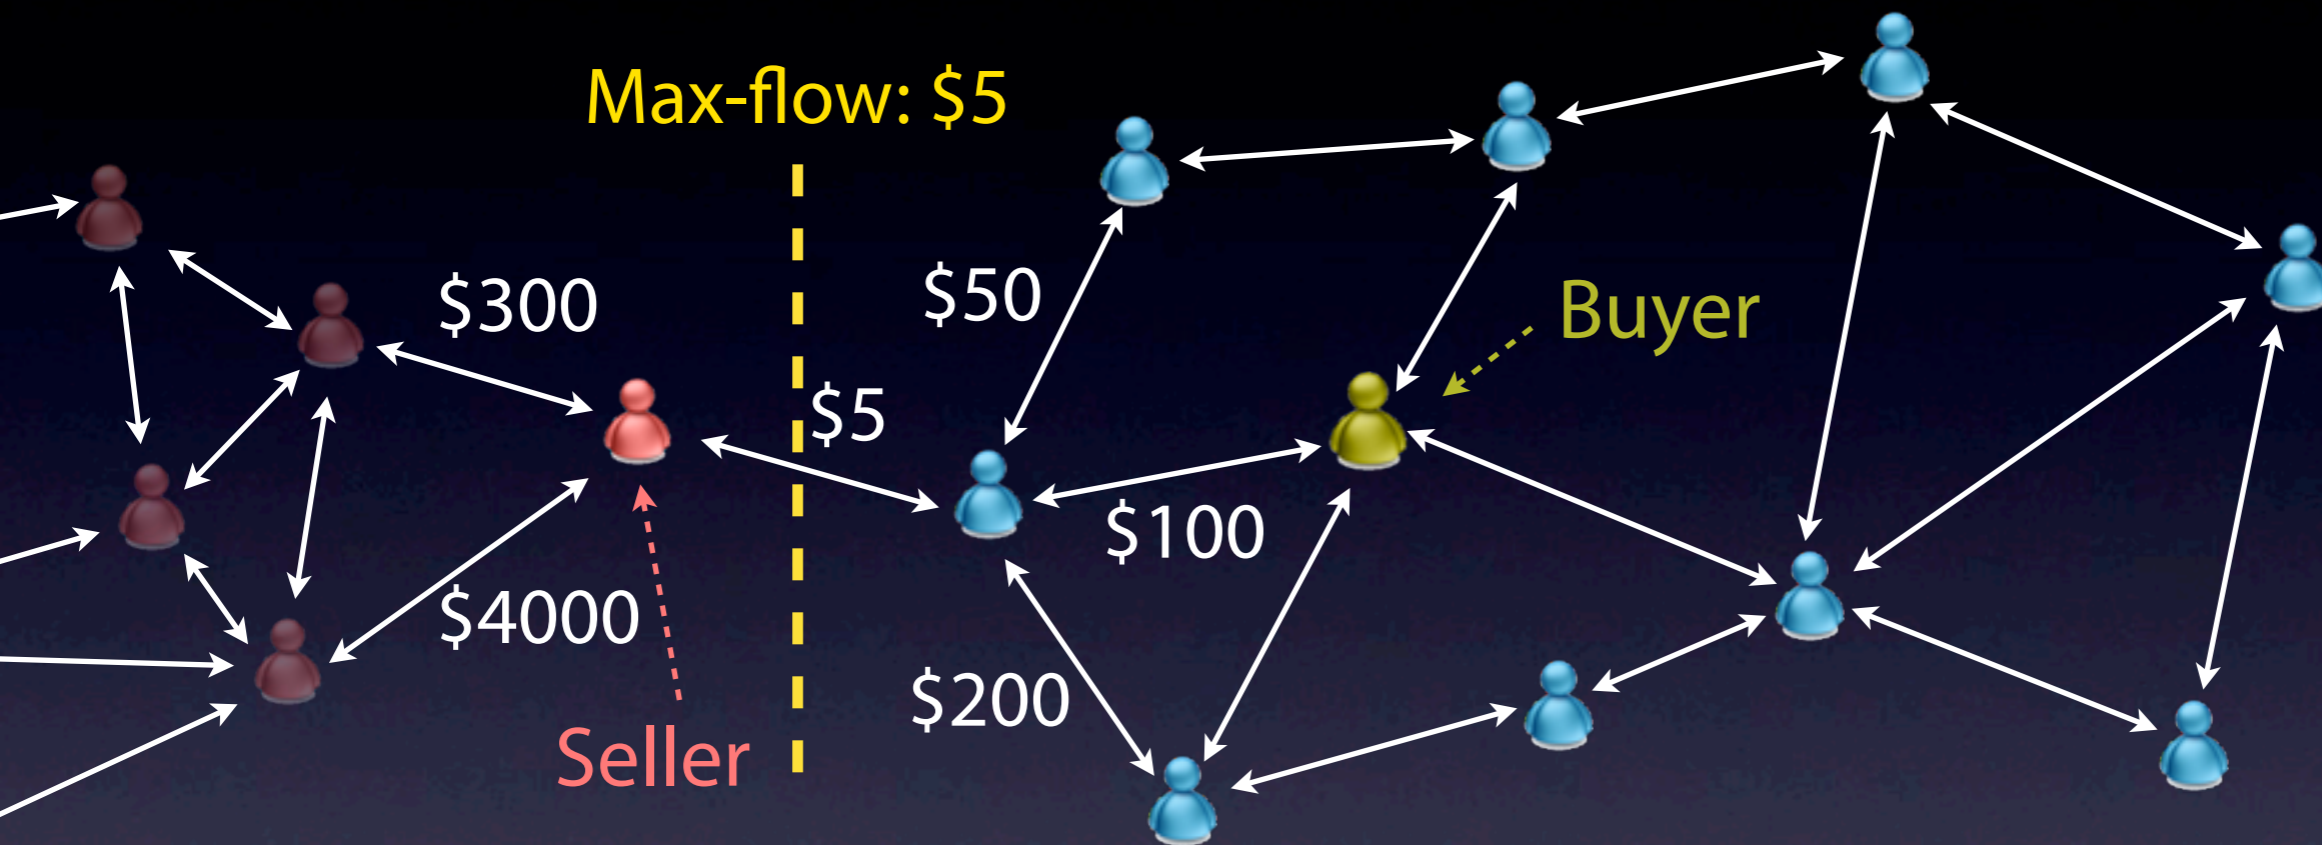
Site operator queries Bazaar before purchase

Bazaar **calculates max-flow between buyer and seller**

If max-flow lower than potential transaction, **flag as fraudulent**

Otherwise, wait for feedback from buyer

Fraud detection with max-flow



Site operator queries Bazaar before purchase

Bazaar **calculates max-flow between buyer and seller**

If max-flow lower than potential transaction, **flag as fraudulent**

Otherwise, wait for feedback from buyer

Handling feedback



Modify risk network when buyer provides feedback


Positive: Create new link

Neutral: Make no changes

Negative: **Remove flow** from network

Malicious sellers **punished if they defraud**

Handling feedback



Positive feedback



Original state

Transaction amount: \$4

Modify risk network when buyer provides feedback


Positive: Create new link

Neutral: Make no changes

Negative: **Remove flow** from network

Malicious sellers **punished if they defraud**

Handling feedback



Positive feedback



Original state

Transaction amount: \$4



Neutral feedback

Modify risk network when buyer provides feedback


Positive: Create new link

Neutral: Make no changes


Negative: **Remove flow** from network

Malicious sellers **punished if they defraud**

Handling feedback



Positive feedback



Original state

Transaction amount: \$4



Neutral feedback



Negative feedback

Modify risk network when buyer provides feedback

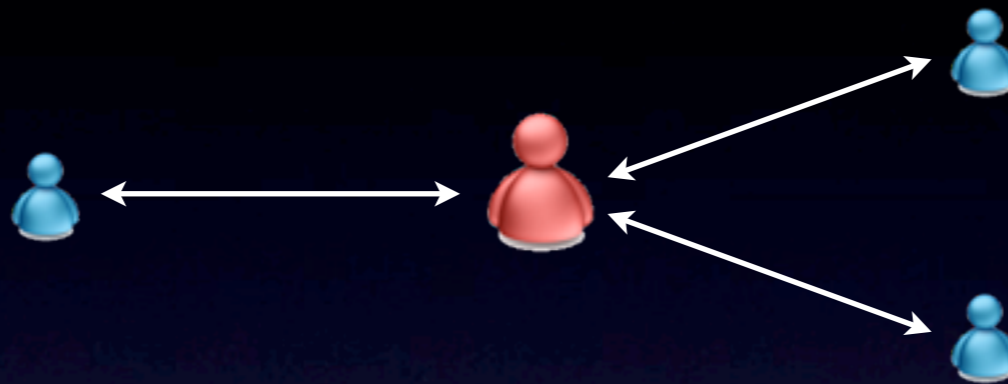
Positive: Create new link

Neutral: Make no changes

Negative: **Remove flow** from network

Malicious sellers **punished if they defraud**

Guarantees




What is the **per-user bound** on defrauding?

$$\sum_{l \in L} w_l$$

Link weight

Set of risk network links

Guarantees for groups



$$\sum_{l \in N} w_l$$

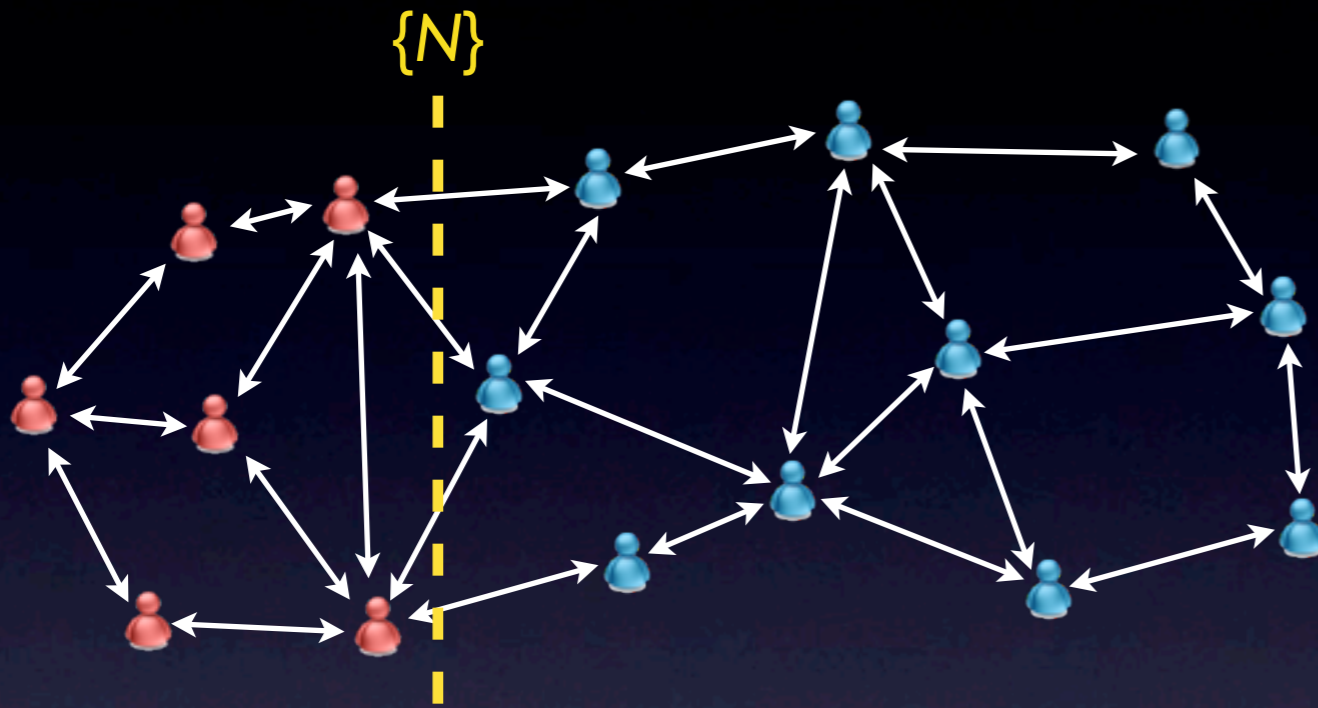
Analysis is same **for any subgraph**

Only way to defraud more: Participate in real transactions

Provides bound on fraud

Result: Collusion, Sybil attacks, white-washing doesn't help

Guarantees for groups



$$\sum_{l \in N} w_l$$

Analysis is same **for any subgraph**

Only way to defraud more: Participate in real transactions
Provides bound on fraud

Result: Collusion, Sybil attacks, white-washing doesn't help

Outline

~~1. Motivation~~


~~2. Bazaar design~~

3. Challenges faced

4. Evaluation

Challenge 1: Feedback delay

Transaction amount: \$4



Buyer cannot immediately determine if fraudulent

Could be used as “window of vulnerability”

Malicious seller could defraud many users quickly

Address by putting credit “on hold”


Set of paths with **flow equal to transaction amount**

Cannot be used by any other transactions

Restore if positive/neutral feedback received

Challenge 1: Feedback delay

Transaction amount: \$4



Buyer cannot immediately determine if fraudulent

Could be used as “window of vulnerability”

Malicious seller could defraud many users quickly

Address by putting credit “on hold”

Set of paths with **flow equal to transaction amount**

Cannot be used by any other transactions

Restore if positive/neutral feedback received


Challenge 2: Bootstrapping



New users have zero max-flow

How to securely bootstrap new users?

New user 



Option 1: **Use social network**

Users can “vouch” for friends, create links

Put their own links at risk

Option 2: **Provide link escrow service**

New user “escrows” for links

Can later ask for escrow back


Links removed; no money returned if lost

Challenge 2: Bootstrapping



New users have zero max-flow

How to securely bootstrap new users?



Option 1: **Use social network**

Users can “vouch” for friends, create links

Put their own links at risk

Option 2: **Provide link escrow service**

New user “escrows” for links

Can later ask for escrow back

Links removed; no money returned if lost


Challenge 2: Bootstrapping



New users have zero max-flow

How to securely bootstrap new users?

New user  \$15



Option 1: **Use social network**

Users can “vouch” for friends, create links

Put their own links at risk

Option 2: **Provide link escrow service**

New user “escrows” for links

Can later ask for escrow back

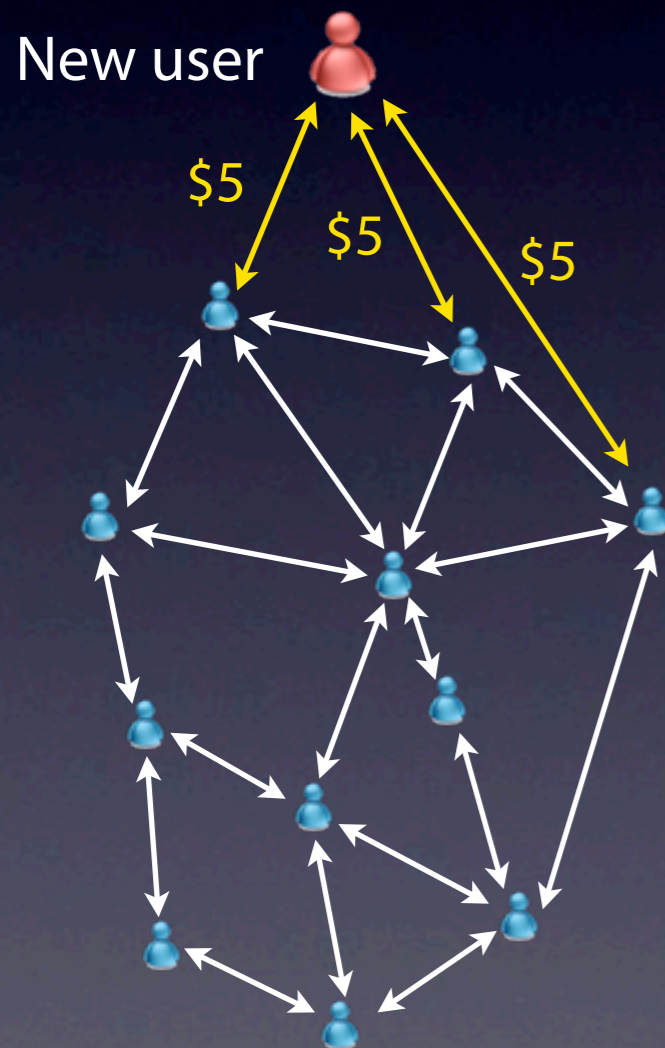
Links removed; no money returned if lost

Challenge 2: Bootstrapping



New users have zero max-flow

How to securely bootstrap new users?



Option 1: **Use social network**

Users can "vouch" for friends, create links

Put their own links at risk

Option 2: **Provide link escrow service**

New user "escrows" for links

Can later ask for escrow back

Links removed; no money returned if lost

Challenge 3: Scaling max-flow

Computing max-flow is expensive

Especially on large, dense graphs

Standard approaches (Gomery-Hu, Goldman-Rao) are poor fit

But, can leverage two observations:

Challenge 3: Scaling max-flow

Computing max-flow is expensive

Especially on large, dense graphs

Standard approaches (Gomery-Hu, Goldman-Rao) are poor fit

But, can leverage two observations:

1. **Risk networks tend to have a dense core**

High-weight links form mostly-connected subgraph

Challenge 3: Scaling max-flow

Computing max-flow is expensive

Especially on large, dense graphs

Standard approaches (Gomory-Hu, Goldman-Rao) are poor fit

But, can leverage two observations:

1. **Risk networks tend to have a dense core**

High-weight links form mostly-connected subgraph

2. **Don't need actual max-flow value**

Only need to know if higher than potential transaction amount

Challenge 3: Scaling max-flow

Computing max-flow is expensive

Especially on large, dense graphs

Standard approaches (Gomory-Hu, Goldman-Rao) are poor fit

But, can leverage two observations:

1. **Risk networks tend to have a dense core**


High-weight links form mostly-connected subgraph

2. **Don't need actual max-flow value**

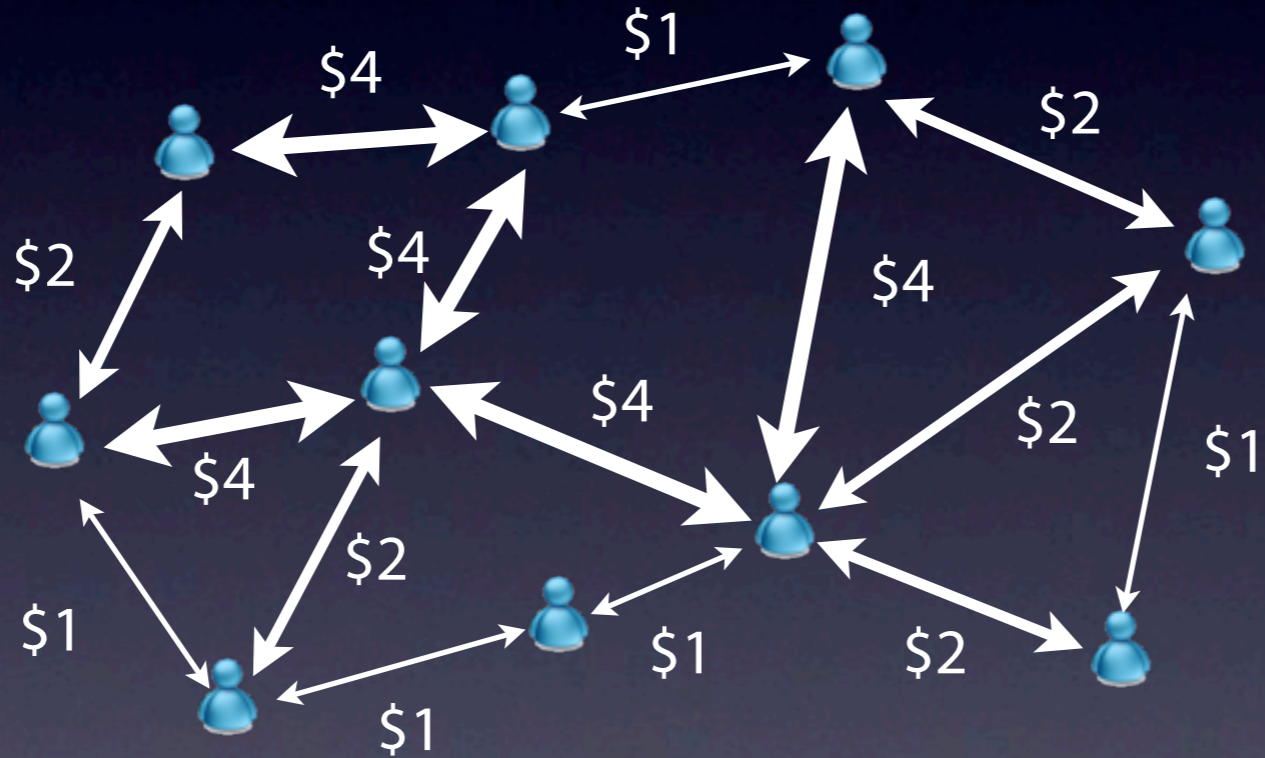
Only need to know if higher than potential transaction amount

Leverage observations with multi-graphs


Multi-graph construction




Multi-graph construction



Multi-graph construction




Multi-graph construction




Normal graph

Multi-graph

Multi-graph construction




Normal graph




Multi-graph

Multi-graph construction




Normal graph




Level 0

Multi-graph

Multi-graph construction




Normal graph




Multi-graph

Multi-graph construction




Normal graph



Multi-graph

Max-flow with multi-graphs



Check for sufficient flow in each level
Starting with the highest

Sufficient flow found → success

Since each level is a subset of the next

Insufficient flow found in all levels → failure

Since Level 0 is entire graph

Possibility of ending quickly

Higher levels have bigger links

Higher levels are smaller networks

Outline

~~1. Motivation~~

~~2. Bazaar design~~

~~3. Challenges faced~~

4. Evaluation

Evaluating Bazaar

Goal: Determine how Bazaar would work in practice

Does it prevent fraud?

How much does it “cost”?

Does it incorrectly flag honest transactions?

Implemented Bazaar in C

Use **multi-graph representation** to store risk network

Run simulations on single processor

How to simulate?

Need **real-world data**

Data from eBay

Category	Purchases	Users	Avg. Price (£)
Clothes	3,311,878	1,436,059	9.74
Collectibles	940,815	454,773	8.90
Computing	964,925	661,285	21.31
Electronics	861,108	652,350	20.67
Home/Garden	2,795,795	1,426,785	16.57

Crawled eBay UK site

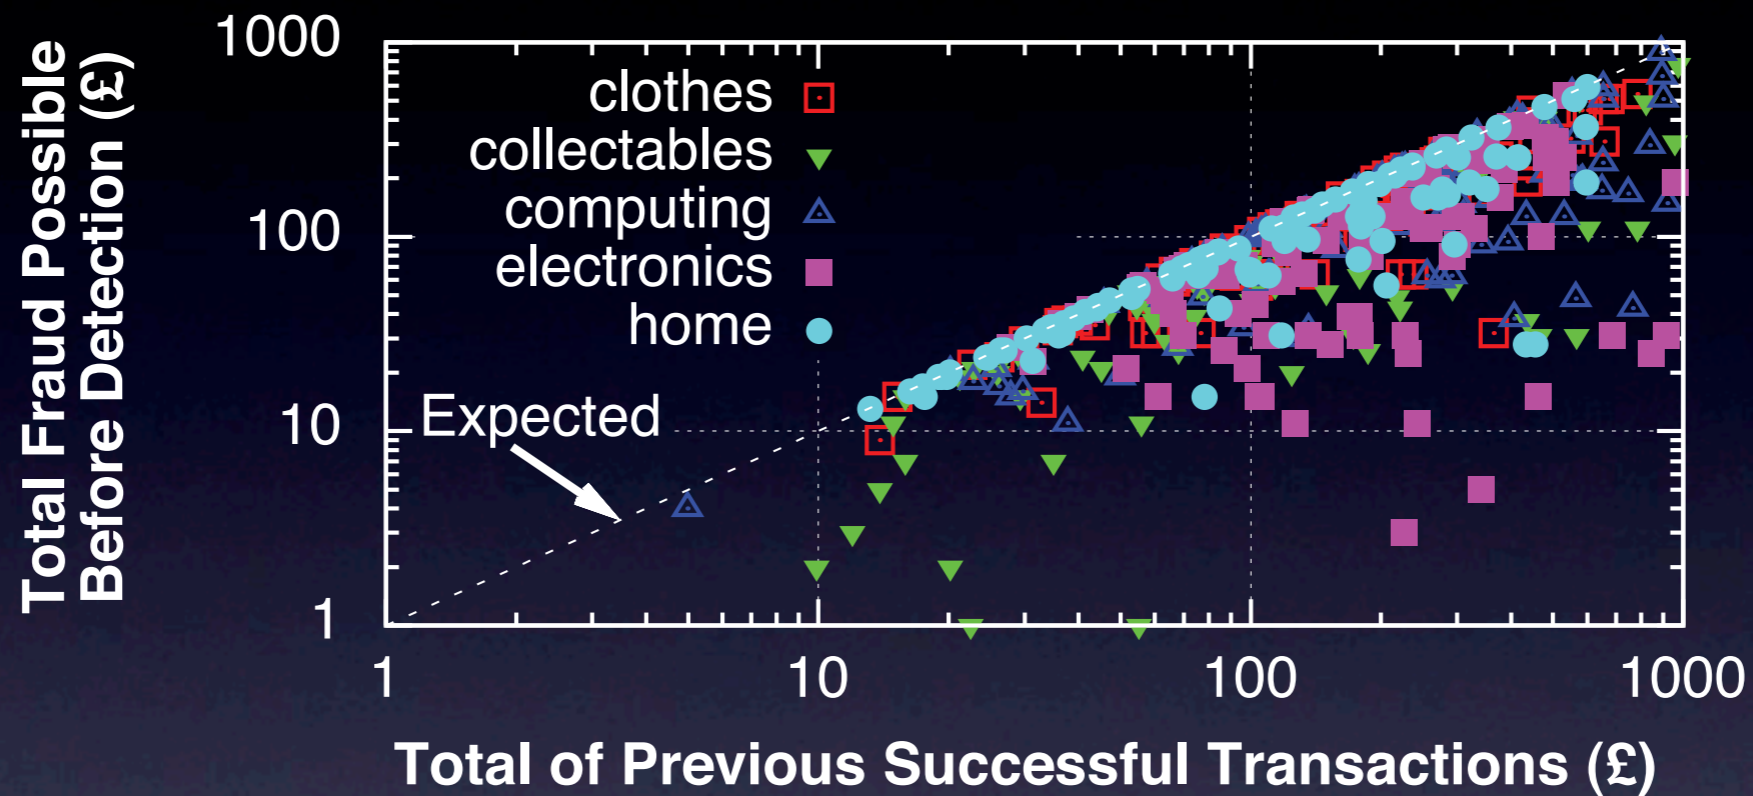
Collected 90-day trace

Focused on five of the most popular categories



Total: Over 8M pieces of feedback

Does Bazaar prevent fraud?



Simulated Bazaar on each eBay category

80% of data creates risk network, remaining is simulated

Random "malicious" users conduct as much fraud as possible

Bazaar **bounds malicious users as expected**

How expensive is Bazaar?

Category	Time (s)		Speedup
	Single	Multi-graph	
Clothes	18.0	6.29	2.86 ×
Collectibles	2.53	1.18	2.14 ×
Computing	3.78	1.66	2.27 ×
Electronics	2.71	1.41	1.92 ×
Home/Garden	11.6	5.34	2.15 ×

What is the **time taken to run max-flow**?

Practical with a few servers provided by site

Can use additional tricks to lower average time

What is the impact on good users?

Category	Fraction of honest transactions incorrectly flagged
Clothes	1.11%
Collectibles	1.12%
Computing	3.23%
Electronics	4.68%
Home/Garden	2.43%

What is Bazaar's **false positive rate**?

Assumes mechanism for "bootstrapping" new users

Less than 5% false positive rate

Summary

Online marketplaces very successful

Democratized commerce, many billions \$ per year

But, known to have significant fraud

Partially due to “free” nature of accounts, reputation manipulation

Bazaar: A **new approach to strengthening reputations**

Leverages risk network between participants

Deployable on sites of today

Were Bazaar deployed during trace

Would have **prevented £164K of negative feedback**

Questions?
