Trust, Reputation and eCommerce

Neel Sundaresan

Sr. Director & Head

eBay Research Labs

Research Labs - Areas of Work

- Search and Ranking
- Machine Learning, DM, NLP, Text Mining, Recommender Systemms
- Social Network Anal. & Apps
- Semantic Sciences
- Large Scale Complex Event Processing and Stream Processing
- Trust, Reputation, Fraud, Info Security
- MicroEconomics
- Platform, Services and Cloud Computing
- Large Scale Scientific Computing
- Systems Modeling, Management and Future Data Centers
- Information Visualization and Analytics
- User Experience and Alternative Interfaces

Topics

- Information Assymetry
- Trust and Reputation
- Feedback as a Trust System
- Trust and Reputation Models
- Trust Propagation
- Trust Marketplace
- Importance of Negative Trust
- Identity
- Tagging
- Summary

Information Assymetry

- Sellers have better knowledge of the goods for sale; buyers don't
- Seller is incentivised to pass of poor quality goods through sale as buyer has no way of verifying
- Creates market inefficiencies, guarantees are indefinites, and such markets disappear

  George Akerloff, "The Market for Lemons: Quality Uncertainty and Market Mechanism" (1970)
- Bad drives out the Good

  Gresham's Law: "Bad Money (Counterfeits) Drives Good Money out of Circulation"

Criteria for a Lemon Market

- Asymmetry of information
- Buyers have no way to assess value before sale
- Sellers have a way to share value before sale
- Seller has incentive to pass off low quality items as high quality ones (continuum of seller quality)
- Sellers with quality items have no way of revealing that information
- No reputation mechanism or regulation to ensure quality
- No *effective* guarantees / warranties

Examples

- Used car market
- Used computer market
- Milk in India in 70s
- Credit in Bangladesh
- Maghribi traders in the Mediterrannian in 11th century
- Rubber Market in South East Asia
- Online: eBay, Craigslist, Y!, Amazon…

- Counter Example: Rice Market

  No Information Assymetry : Hence traded in Open Market!

Solutions to the Lemon Market

- Co-operatives – Milk Market
- Coalition – Maghribi traders, Bangladesh Credit
- Long term contracts – Rubber
- Used Car Market – Branding and Manufacturer certification

- All have a definition of “Reputation”

Reputation

- Sharing of reputation lowers the ability of dishonest agent to profit in the future
  - Dishonest agents will have to seek new partners – who will pay only discounted Lemon price : (Trust Discount)
  - Dishonest agents can still trade outside the coalition boundaries
- Private Reputation vs Public (Shared) Reputation

Solutions to Online "Lemon" Markets

- Reputation Systems

  Improving the Lemons Market with a Reputation System: An Experimental Study of Internet Auctioning" – Toshio Yamagishi

- Vast quantity of cheaply available reputation information in online trades offsets the lack of quality and reliability of reputation

  - \- Resnick&Zackhauser 2001

Online vs Offline Markets

- Offline Markets have closed boundaries whereas Online Markets are open
- Incentive for Shared Reputation not clear in Online market
  - eBay in the early days vs now
- Existence of Negative Reputation and Positive Reputation
  - Positive reputation is more effective in solving the "lemons" problem (Kollock '99)
  - dishonest agents can move to a different market without paying penalty or exit/entrace cost
- Stability of Identity
  - dishonest agents can change identity

Yamagishi Experiments

- Conclusions
  - Information Asymmetry leads to lemon markets
    - Lower quality goods are traded and opportunity for higher quality goods gone
  - Reputation alleviates lemon markets where traders identities are permanent
  - Power of reputation reduced by identity changes and/or cancel reputations
  - Negative reputation vulnerable to identity changes where positive reputation is not vulnerable to it
- Properly designed reputation mechanism should resolve lemons problem

Online Strangers and Trust

- Requirements
  - Buyers/sellers should be able to distinguish between trustworthy and non-trustworthy sellers/buyers
  - Encourage sellers/buyers to be trustworthy
  - Discourage participation from non-trustworthy sellers/buyers
- Note that requirements on buyers is significantly lower than the sellers
  - Sellers hold items till money sent
  - Sellers don't control who they sell to

eBay Feedback

- History
  - Before 1999
    - Anybody could leave feedback for anybody
  - Now
    - Feedbacks are per-transaction between seller and winning bidder
- Accumulative
  - Positive (+1), Negative (-1), Neutral(0)
  - 1 line of qualitative textual feedback
- Feedback Profile is public
  - Any prospective buyer can see all per-transaction feedback with scores and text
- Feedback 2.0
  - Revealed on multiple aspects of the user feedback
- Today
  - Only buyers can leave feedbacks, sellers can only leave positive feedbacks

eBay Feedback

- Most Feedbacks are positive (Pollyanna effect)
  - Negatives are fewer
  - Fear of retaliation
  - Satisfaction of receipt
  - No Feedbacks instead
  - Positive with text specifying negative experiences
  - High courtesy Equilibrium (Resnick/Zeckhauser ’01)
  - Mutually negative feedbacks may represent misplaced blame
  - Who goes first?
    - Since buyer typically pays first, expect seller to go first
    - Buyer goes first twice as often (Resnick/Zeckhauser ’01)

eBay Feedback

- Sound of Feedback Silence (Dellarocas '07)
  - 57% give feedback, 41% are silent
  - Also looked at who goes first, period before feedback
  - Buyer satisfaction(79,29.3,0.7), Seller satisfaction(86,13.5,0.5)
- More Recently
  - Only buyers can leave feedback
  - Sellers can leave only positive feedback
  - What's the impact of an asymmetric reputation system?

eBay Categories ▾ | eBay Motors | eBay Express

Home > Community > Feedback Forum > Feedback Profile

Feedback Profile

chateaurugs (3182 ★) me

Member since May-24-02 in United States

Contact member | View items for sale | View seller's Store | More options

Feedback Score:	**3182**
Positive Feedback:	**99.1%**
Members who left a positive:	3211
Members who left a negative:	30
All positive Feedback:	4047

Find out what these numbers mean

Recent Feedback Ratings (last 12 months)

	1 month	6 months	12 months
Positive	137	1272	2580
Neutral	0	4	23
Negative	2	5	19

Detailed Seller Ratings (since May 2007)

Criteria	Average rating	Number of ratings
Item as described	★★★★½	83
Communication	★★★★½	82
Shipping time	★★★★★	83
Shipping and handling charges	★★★★	82

Feedback as a seller | Feedback as a buyer | **All Feedback** | Feedback left for others

Ratings mutually withdrawn: **0**

4,133 Feedback received — Page 1 of 166

Feedback / Item	From / Price	Date / Time
Beautiful rug. Very professional transaction. OVER 70 YEARS 8'4x4'2 Shiraz Persian Rugs W-6787 (#130113542627)	Buyer: yellowlayla (580 ★) US $51.00	May-30-07 13:06 View Item
Very fast shipping. Great communication. EXTRA FINE RARE 11'5x7'10 Meshkabad Persian Rug Carpet (#130049705651)	Buyer: yellowlayla (580 ★) US $3,515.00	May-30-07 13:02 View Item
Beautiful rug. Very professional transaction. ANTIQUE LAVAR! 13'0x9'8 Kerman Persian Rugs S-11169 (#130113371918)	Buyer: yellowlayla (580 ★) US $3,799.00	May-30-07 12:58 View Item
I was not happy with the rugs coloring. Paid $166 for something I wont use. HANDMADE SQUARE 8'0x8'0 Agra Oriental Rugs MI-838 (#130112707930)	Buyer: mobjackbay777 (5) US $71.00	May-29-07 19:25 View Item

Internet

Analysis of Feedback Text (Sundaresan et al 2007)

eBay Feedback – Rated Aspect Summary (Lu, Sundaresan, Zhang – 2009)

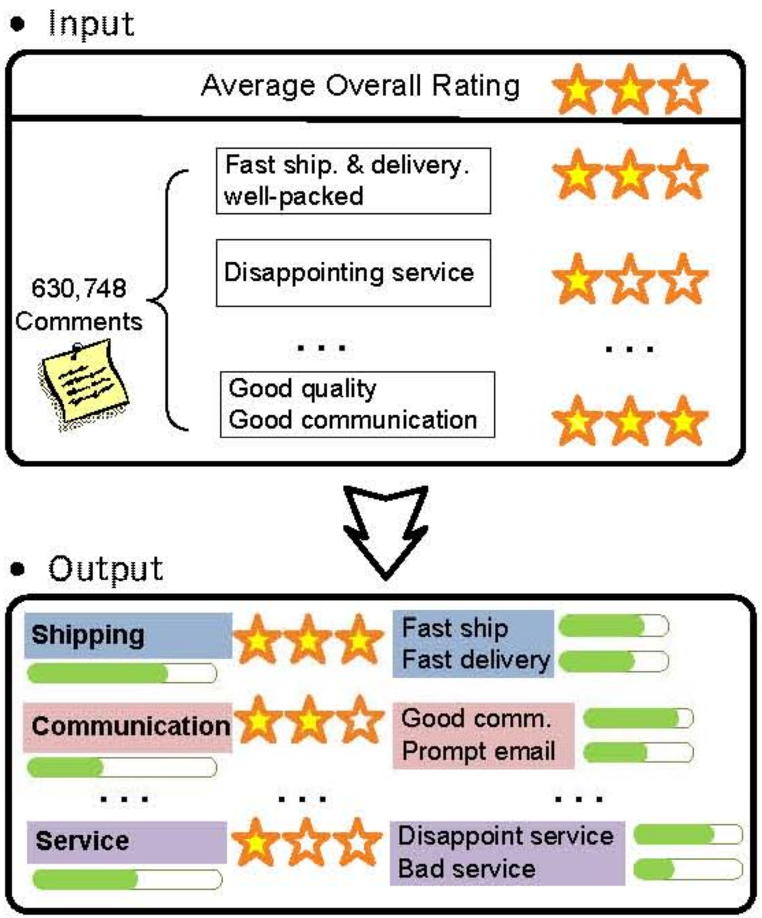

• Break down comments into head terms (aspects) and qualifiers (opinions)

•Phase 1: Identify k interesting aspects and cluster data into these – k-means / PLSA / Structured PLSA

•Use priors (Dirichlet) act as training to bias clustering results

•Then use MAP (Max. A Posteriori) to estimate all the parameters

• Phase 2: Identify rating functions for the k aspect clusters using local (per-user) or global rating information

•WWW09 – Lu, Sundaresan, Zhang, “Rated Aspect Summarization” /

Expression/Sentiment

- Expression is a metaphor for Trust
- When Trust is expressed through feedback or textual communication it influences mutual Trust and future Trust
- Factors to take into account
  - What a user usually says
  - What is the change in what the user usually says
  - How does what someone says affects what the next user says
  - Only those who have significantly significant things to say do say anything at all

Is Reputation Rewarded?

- How reputation impacts buying/selling decisions?
  - Do buyers pay a higher prices for items from higher reputation sellers?
    - Is reputation an indicator of future performance?
  - Do sellers list items at a higher (reserve) price based on their reputation?

- Tricky to study
  - Correlation between reputation and quality of items or listings
  - eBay's "One of a kind" nature of items (harder to standardize on quality)
    - Good Will Hunting (Dellarocas '00) approach to feedback quality of products revealed from sellers to buyers resulting in better behavior
      - reveal (new, NIB, NWoB, NWoT, used, refurb…)

Is Reputation Rewarded? (contd)

- Regression analysis used to study the impact of reputation on price and probability of sale (Resnick/Zeckhauser ’01)

    No significant impact on price

    Significant impact on probability of sale (almost doubles from low feedback to v high feedback)

Do Sellers care about their feedback?

- Sellers can respond to a negative feedback
  - The text of the response is displayed below the feedback text
  - More than a third of the sellers respond to negative feedback
  - Sound of Silence relates somewhat to fear of retaliatory feedback
- Other studies
  - Behavioral changes after a negative feedback
    - Improved behavior vs non-participation
    - Retaliation

So far…

- Trust and Reputation have been loosely used to imply “goodness measure” that sustain quality transactions in marketplaces
- Feedback is an expression of Trust
- Trust and Reputation are sometimes interchangeably used, sometimes confused, or differently defined
- We need these measures as user takes risks based on prior performance when there is no way to “test before buy”

Trust

- Trust (Josang et al 2007)
  - Reliability Trust: (Gambetta 1988)
    - Subjective Probability by which an actor A expects that another actor B performs an action on which its welfare depends
    - There is a dependence/reliance on the trusted party by the trusting party
  - Decision Trust (Broader defiintion: McKnight & Chervany 1996)
    - Extent to which one actor is willing in a given situation with relative security
      - Negative consequences are possible
      - Utility attached -- positive utility resulting from positive outcome and negative utility resulting from negative outcome
      - Risk emerges from Decision Trust when the value of the transaction is high and the probability of failure is non-negligible

Reputation

- Reputation is what is generally said or believed about an actor or item's character or standing
- It's a "global" measure

Trust and Reputation

- Trust is subjective, Reputation is objective
- Trust is relative, Reputation is global
- Trust is personal, Reputation is collective

A trusts B because B has a good reputation

A trusts B in spite of not knowing B's reputation

A trusts B in spite of B's bad reputation

- Reputation may change as Trust between agents change
  - though Reputation measures cannot be oversensitive to trust changes

Mathematical Equivalence Properties of Trust

- Reflexivity

  a T a

- Symmetry

  a T b ⇔ b T a

- Transitivity

  a T b and b T c => a T c

Transitivity is called derived trust

Derived Trust is also important when certifiers or market makers are involved.

Trust Transitivity and Recommendation

- Sometimes transitivity is strengthened by recommendation

  a T b and b T c and b R a => a T c

Trust and Security/Safety

- Purpose of Security is to provide protection against malicious actors
  - Trust and Reputation can be used as soft security mechanisms
  - System specified security rules/flags override user-subjective trust
  - A Trust provider can provide a secure communication path between trusted parties.
    - Notion of privacy and encryption come into place
    - Identity Trust (e.g., PGP)

Recommender Systems

- Collaborative Filtering

  2 Actors may share taste and may rate items similarly. They are neighbors in the recommendation space.

  This information can be used to recommend items that one actor likes to that actor's neighbors.

  Items may be replaced by actors

Recommender vs Reputation Systems

- Reputation systems provide collaborative sanctioning (Montashemi '01) to provide a common judging mechanism for actors
- Recommender (CF) systems use taste as input for rating, whereas reputation system is insensitive to taste.
- CF systems take an optimistic view (all participants trustworthy but different tastes) whereas reputation systems are objective

Combining Recommender and Reputation Systems

- Combining recommender with reputation systems
  - Damiani '02 (P2P systems)
    - E.g. Amazon rating system
  - Collaborative behaviors can be used to weight trust measures which in turn used for reputation
  - Recommender systems first identifies neighborhoods of actors and makes recommendation to an actor in a neighborhood based upon liking for items by others in the neighborhoods and the actor in question
  - Trust models (si trusts sj) can be used to seed recommendations to new entrants in the system

Reputation System Implementation

- Centralized system

  Central authority uses a centralized reputation computation engine

  E.g. eBay, Amazon, Slashdot,…

- Distributed system

  P2P system. The purpose of reputation system is

  Phase 1 (Search phase): to identify which servents (server-clients) are most reliable at offering the best quality resources. This may be centralized (Napster)

  Phase 2 (Download phase): to identify which servent provides the most reliable info

  E.g. KaZaa(Skype), Napster, Gnutella, Freenet,…

Reputation Computation Engines

- Accumulative

  eBay’s feedback system

    Total Positives – Negatives = Feedback score

    Total Postives/Total = Feedback percentage

  Simple and transparent but gameable

  Enhanced: weighted schemes based on rater trustworthiness/reputation, rating age, distance between rating and current score etc.

Rating Computation Engines (contd.)

- Bayesian Systems

  Take binary ratings as input (+ve, -ve)

  Scores computed by updating beta PDF (probability density functions)

  A posteriori (updated) reputation computed by combining a priori (previous) reputation score with the new rating

  Let (α,β) representing +ve and –ve scores.

  The beta-family of distributions is a continuous family of functions indexed by parameters α and β.

Bayesian Systems - Contd.

Beta-PDF beta(p| α,β) can be expressed using a Γ function as:

beta(p| α,β) =($\Gamma(\alpha+\beta)/(\Gamma(\alpha)\Gamma(\beta))) p^{\alpha-1}(1-p)^{\beta-1}$

With the restriction that p != 0 if $\alpha < 1$ and p ! = 1 if $\beta < 1$

Expectation value of beta distribution is given by

$E(p) = \alpha/(\alpha + \beta)$

Reputation can be defined as a function of E(p)

The PDF expresses uncertain probability that future interactions will be +ve.

Example: Assume a priori distribution of $\alpha = 1$, $\beta = 1$.

After observing some r positive and s negative outcomes, the posteriori distribution is $\alpha = r+1$, $\beta = s+1$

given r=7, s=1, E(p)=8/10=0.8 meaning that relative frequency of positive outcome in the future is most likely to be 0.8

Discrete Trust Model

- Actor's trustworthiness is measured as fixed enumerated values (Very Trustworthy, Trustworthy, UnTrustworthy, Very UntrustWorthy). (Abdul-Rahman et al 2000)

- Referrals are weighted based upon the referring actor's trustworthiness (referring actor's rating of actor x can be compared with the relying actor's own rating of x. Based upon this the referrals from referring party may be downgraded!

Belief Systems

- Based on Belief theory where the sum of the probabilities of possible outcomes is not necessary 1, the residue is identified as uncertainty. (Josang 1999)

Belief/trust metric called Opinion is denoted by

w(x, A) = <b,d,u,a>

where b, d, u represent belief,disbelief,uncertainty,

a represents base rate probability in the absence of evidence and a is used for computing an opinion's probability expectation value

E(w(x,A)) = b+au

Example

- A trusts B and asks B for a recommendation who recommends C
- A trusts D and asks D for a recommendation who recommends C
- Derived trust from A => C is built via B and C by combining the trust paths A->B->C and A->D->C using a consensus operator (say, using Dempster's rule)
- The consensus operator is equivalent to the Bayesian updating as opinions can be uniquely mapped to Beta PDFs

Fuzzy Models

- Use fuzzy inferences to handle uncertainties, fuzziness, and incompleteness.
- Based on the idea that in a P2P transaction system evaluation and dissemination of trust can't be effectively done and actors rely on collection of other's opinions. Global reputation computation is time consuming
- 2 Major inference steps
  - Local Trust Inference
  - Global Reputation Computation

Trust and Reputation Inference

- Buyer's local trust score
  = f(payment method, payment time)
- Seller's local trust score
  = g(shipping time, goods quality)
- Global Reputation weight
  = h(peer's trust score, transaction a/m, transaction date)
  Where f, g, h are fuzzy inference functions

Reputation Weights

- If transaction is new, and amount is high then weight is high
- If transaction is old, amount is low then weight is low
- If peer's reputation good, transaction amount is high then weight is high
- If peer's reputation good, transaction amount is low then weight is medium
- If peer's reputation bad, weight is low

Reputation Calculation

- $R_i = \sum_{j \varepsilon S}(w_j / \sum_{j \varepsilon S}(w_j))t_{ji}$
  $= \sum_{j \varepsilon S}(w_j t_{ji}) / \sum_{j \varepsilon S}(w_j)$

Where R_i is the reputation score for the Peer i, t_{ji} is the trust score of peer i by peer j and w_j is the aggregation weight of t_{ji}

The global reputation computation is an iterative process and converges over multiple iterations as a stable reputation score for peer i

Overlay Computation

- DHT (Distributed Hash table) algorithm (Yideu Mei et al 2008)
  - Each peer maintains 2 tables: a transaction record table and the peers' trust scores.
    - The transaction record information is used for computing weights
  - To make the algorithm scalable an aggregation threshold is maintained and peers whose weight contributions are below this threshold are not queried for trust scores.

PowerTrust (Zhu, Hwang 2006)

- Uses the same architecture as FuzzyTrust

  discovers and uses Power Law matters in the trust system. Uses power trust scores to aggregate efficiently.

  Uses lookahead random walk and locality preserving hash in DHT to perform Reputation Aggregation

PeerTrust (Liong, Xiu 2004)

- Trust score of a peer is computed as the average of the scores weighted by the feedback of the peers
- Scores based on 5 factors – peer record, credibility, transaction context, community context, and scope

SmallTrust (Sakurai Lab, Kyushu univ)

- Based on Small World phenomena
  - 2 actors in the network are connected by a short path of acquaintance actors

Flow Models

- Compute trust and reputation scores through loops and chains called flow models
- E.g. PageRank, Advogato, EigenTrust
  - Models like PageRank assume that the trust/reputation weight for the entire system is a constant and members of the community can increase their trust/reputation at the cost of others.
  - In PageRank increased in-links (incoming flow) to a page increase its ranks and increased outlinks (outgoing flow) from a page decreases it
  - EigenTrust doesn't require all sums of scores to be a constant. It computes the agent trust scores through repeated iterative multiplication aggregation of trust scores along transitive chains till convergence

Static Web (PageRank)

- Let P be a set of hyperlinked web pages and let u and v denote web pages in P. Let $N^-(u)$ denote the set of web pages pointing to u and $N^+(v)$ set of web pages that v points to. Let be some vector over P that gives an initial rank.
- Then the pageRank of a page u is given by:

$$R(u) = c\, E(u) + c \sum_{v \varepsilon N-(u)} (R(v)/|\, N^+(v))|$$

Where c is chosen such that $\sum_{u \varepsilon P} R(u) = 1$

  - PageRank applies transitivity of trust to the extreme as trust scores flow through long chains of links.

- Personalized PageRank: Vote pages based upon queries: Assigning initial votes based upon the topic of the query (Haveliwala, 2002)

Static Web (HITS)

- WebHITS/Clever (Kleinberg '97)
  - Starting with a query a web subgraph is identified to define Hub and Authority pages

    Hub: Pages that link to authoritative pages

    Authority: Pages linked to by hub pages
  - Mutually recursive definition results in solving a simultaneous matrix equation to compute the 2 vectors by computing a principal eigen vector.
  - Higher order eigen vectors reveal dense micro communities related to the query

TrustRank(Gyongji '04)

- Enhances PageRank to separate good pages from spam pages on the web
  - Start with a seed set of pages which are marked "good" or "bad" by experts
  - As you propagate starting from the good pages reduce the trust level by applying a damping factor
  - For multiple incoming links the trust can be the average of incoming trusts
  - For outlinks the trust can be propagated by dampening based on the number of outlinks

EigenTrust System (Kamvar et al)

- Global reputation for each actor is given by the local trust values aassigned to the peer by other peers.
- Normalized local trust values

  To avoid collusion/malbehavior

  $c_{ij} = \max(s_{ij},0)/\sum_j(\max(s_{ij},0)$ where s_{ij} represents actor i's subjective trust on j

  Note that this is equal normalization does not take into account the trust values of the peers themselves to weight

EigenTrust

- Local Trust Value Transitivity

  $t_{ik} = \sum_j (c_{ij} c_{jk})$

  If $C = [c_{ij}]$, $t_i^{->}$ is the vector of t_{ik}'s then

  $t_i^{->} = C^T c_i^{->}$

  This is trust transitivity by actor i asking only his peers.

  To expand to friends' friends $t = (C^T)^2 c_i^{->}$

  . And so on…

  $t = (C^T)^n c_i^{->}$ for large n

  For large n trust vector $t_i^{->}$ will converge to the same vector for every peer i. Namely it will converge to the left principal eigenvector of C. In other words $t^{->}$ is a global trust vector in this model. Its elements tj quantify how much trust the system as a whole places on peer j.

  At the most basic level one could iterate $t_i^{->(x)} = C^T t_i^{->(x-1)}$

  Where x = 0, 1,… k times till the distance between $t^{(k)}$ and $t^{(k-1)}$ is less than some predecided ε

EigenTrust

- In a practical scenario one has to take into account
  - Idle actors
  - Pre-trusted peers
  - Malicious collectives
    - this is accounted for by requiring each peer place some trust in someone outside the collective
    - $t_i^{\to(x)} = (1-a)C^T t_i^{\to(x-1)} + ap^{\to}$ where p is the distribution of pre-trusted peers.

Online Implementations: eCommerce

- eBay
  - Feedback (+ve, -ve, neutral)
  - Most are positive
    - Reciprocation of +ve and retaliation of –ves
  - Research has shown correlation between feedback scores and sell-throughs
  - (refer to original slides early on)

Product Reviews

- Epinions
  - Members can provide reviews on goods, products and services
    - Textual PLUS ratings of 1-5 stars on various aspects
  - Other members rate reviewers as Very Helpful, …, Not Helpful
  - Accumulated ratings of a member over a period make that reviewer an Advisor, Top Reviewer or a Category Lead
    - Top reviewers are automatically chosen and advisors are similarly chosen at lower thresholds
    - Category leads are chosen by the company based on member nominations

Epinons Web of Trust

- Members can decide to 'trust' or 'block' other members
  - A members trusted circle of members is its personal Web of Trust
  - Trust and Block have +ve, and –ve impact on a member's qualification as a Top Reviewer

Epinions Incentive System

- The company makes money from businesses based upon click-throughs and lead generation
- Through their Income Share Program members can earn money
  - Based upon usefulness of reviews (both positive and negative)
- Other early dot.com incentives like cash for member signups

Bizrate

Consumer driven merchant rating service

- Merchants are Bizrate certified if enough members rate Bizrate listed merchants on various dimensions.
- Incentives to members is discount at the stores
  - Positive bias since frustrated customers never finish
- Also a Product rating service as Epinions

Amazon

- Of items, of reviewers, of members, of businesses
  - Items rated, final 'item rating' aggregate average of all ratings
  - Reviews include text and ratings
    - Reviews can also be rated and graduates people to "Top 1000" reviewer etc.
  - Favorite People. Influence ranking of reviews in favorites list.
- Incentives
  - None from Amazon
  - Publishers could incent reviewers
- Negatives
  - Ballot stuffing, badmouthing by top reviewers
  - Top reviewer may not be an individual (has to have read more books than everyone else)
  - Entering the elite circle triggers negative feedback
  - Ratings are cookie-based so can game the system by working around that

Online Implementations: Discussion Space

- Slashdot.org
  - Automatic moderator selection
  - 2 layered moderation scheme: M1 for moderating articles, M2 for moderating moderators
  - The system regularly picks moderators,gives them points to moderate comments. Positive/negative moderations to comments influence the comments and the author positively/negatively.
  - Users have Karma attached to them, karma increases as users' comments are positively moderated, decreases as they are negatively moderated.
  - Comments by users with high karma start at a score of 2, low Karma starts at 0 or -1.
  - Points given to moderators when they are selected is high or low depending on their karma levels.
  - To address unfair moderations, Slashdot has layer 2 moderators or M2.
  - Any user can metamoderate several time per day. They will be asked to metamoderate on randomly selected postings. This moderation affects the Karma of M1 moderators (which in turn impacts their future ability to be moderators)

Digg

- Community submits stories. Once a story gets enough diggs, it is relevant enough to show up on the top page.
- Stories with fewer diggs or that are marked as spam are kept in the "digg all" area to be eventually removed.
- Negatives
  - Top 100 diggers control 56% content
  - Just 20 users have submitted top 25% content
- System changed due to negative experiences with the current algorithm

Advogato

- A community of open-source programmers
- Uses a trust scheme to manage peer review process based on PageRank style algorithm (based on a Flow model)
  - Models a flow network (members as nodes and referrals as edges).
  - Members refer each other as Apprentice, Journeyer, Master.
  - A separate flow graph is generate for each level
  - A member reachable by the highest level flow graph has that rating

The Reputation Market

- “thelandseller” case study (Brown, Morgan 2006)

  “Riddle for a penny! No shipping – Positive Feedback” for a penny

  - ok: selling a joke

  - suspicious: title spam “feedback”

  - suspicious: total price < cost of listing

  212 jokes sold (to 172 buyers) at a loss of $87.42

  At feedback 598 (100%) the seller actually selling land in Texas

Reputation Market (contd)

- New entrants need to start somewhere and might be participant to such offers
  (see later)
- Preparing for a larger blow (big sale, or fraud) by padding reputation
- Take Volume based Reputation
  - Sale of a Car different from cookie recipe
- Reputation score gets less transparent as factors added in
  - Can be opaque to catch violators

Need for Negative Reputation and Complaints

- Lack of complaints make reputation implementations weaker (Resnick 2002)
- Lack of penalizing or reducing reputation mechanisms helps create market for trading recommendations.(Clausen 2004)
  - SearchKing is a matchmaker of PageRanks (those who have it with those who want it)

Multiple Identities: Sybil Attack on Reputation

- Sybil Attack: Single person voting many times (Douceur 2002) with multiple identities
- So, what's the cost of an attack on a reputation system?

Cost of Attack on Reputation System

- PageRank Attack (Clausen 2004)
- Assume that the web graph into 2 parts – the good part and the one controlled by the attacker
- The cost can be computed based upon the cost to register a domain name (traditionally root web pages are assigned initial page rank votes, anyway)
- Cost is computed at a particular page rank g and is given by $z = g.\sum_{v \varepsilon V} c(v) / \sum_{p \varepsilon P} c(p)$

  where V is the part of the web controlled by the attacker and P is the web graph

Costs and Payoffs

- For lower pageranks the estimate is tens of dollars and for high over 100K
  - This compares to what SearchKing charges for PageRank
  - Attacker could buy unmaintained/stale sites for cheap
  - Other strategies could be to take over high pagerank sites

- High cost of acquiring sites to rip people off may not make sense. However, once acquired site could scam people with the lack of mechanism for complaint

Trust and Distrust Propagation (Guha et al 2004)

- We can store trust and distrust in 2 different matrices T = [tij], D = [dij].
- B is the belief matrix B = T – D in simple cases
- Propagation – let M be the operator, t be the trust operator

  Atomic (1-step transitivity: i t j, j t k => i t k) so $B.M = B^2$

  Co-citation - i_1 t j_1 and j_2, and i_2 t j_2 then i_1 t j_2. This operator is B^TB, so $B.M = BB^TB$

  Transpose i trusts j => j trusts i. Here the operator is B^T

  Coupling i trusts j => i trusts k because j and k trust actors in common. Operator is BB^T

  Let $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ be a weight vector combining these 4 propagation schemes. Then we can capture all propagations into a single combined matrix $C_{B,\alpha} = \alpha_1 B + \alpha_2 B^T B + \alpha_3 B^T + \alpha_4 BB^T$

Propagating both Trust and Distrust

- Let $C_{B,\alpha}$ show beliefs should flow from i to j via an atomic propagation step. (if the entry is 0 then nothing can be concluded in an atomic step).
- Let k be a +ve integer and $P^{(k)}$ a matrix whose i,j-th entry indicates the k propagation operations.
- Three models that to define B (the belief matrix)

  Trust only: $B = T$, and $P^{(k)} = C_{B,\alpha}^{(k)}$

  One-step Distrust: distrust propagates one step only $B = T$, and $P^{(k)} = C_{B,\alpha}^{(k)}.(T-D)$

  Propagated Distrust. In this case, $B = T – D$. and $P^{(k)} = C_{B,\alpha}^{(k)}$

Reaching the Final Value

- 2 approaches

  Eigenvalue propagation

    Let K be a chosen integer. The final matrix F is given by $P^{(K)}$

  Weighted Linear Combinations. To penalize longer chains over shorter chains choose γ (smaller than the largest eigen value of $C_{B,\alpha}$ and let K be a chosen integer.)

    Then $F = \sum_{k=1,K} \gamma^k . P^{(K)}$

Interpreting F

- To interpret F as trust or distrust

  Various threshold at local, global, or at majority level can be used to partition trust and distrust.

Is Distrust Transitive?

- A distrusts B, B distrusts C, then we can think of 2 models
  - Additive: A > B, B > C, A >> C
  - Multiplicative: A distrusts B, B distrusts C, A trusts C. This might have the negative implication of A distrusting A.
  - Distrust is not a negating function. For instance, if A distrusts B, A should distrust B's actions that include distrusting C.

Qualifying Reputation Score

- In a Marketplace like eBay a seller to successfully sell or a buyer to win an auction has to be of certain capability
  - There might be a translation from this to the reputation
  - The fact that there is a market for reputation implies this as well

- In eBay different categories are different when it comes to motifs of transaction

- We can look at Feedback as an approximation for reputation and compute the qualifying feedback score

Auroral Diagrams (Shen, Sundaresan 07) - Across All Categories

Buyers

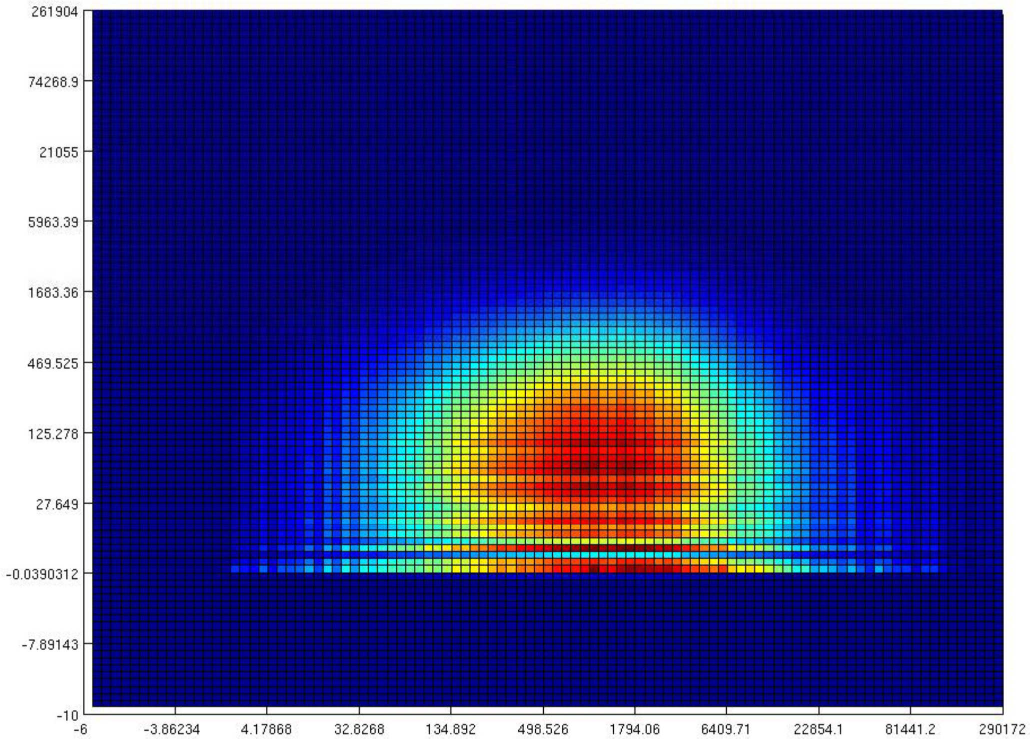

Sellers

Auroral Diagram: Arts and Craft

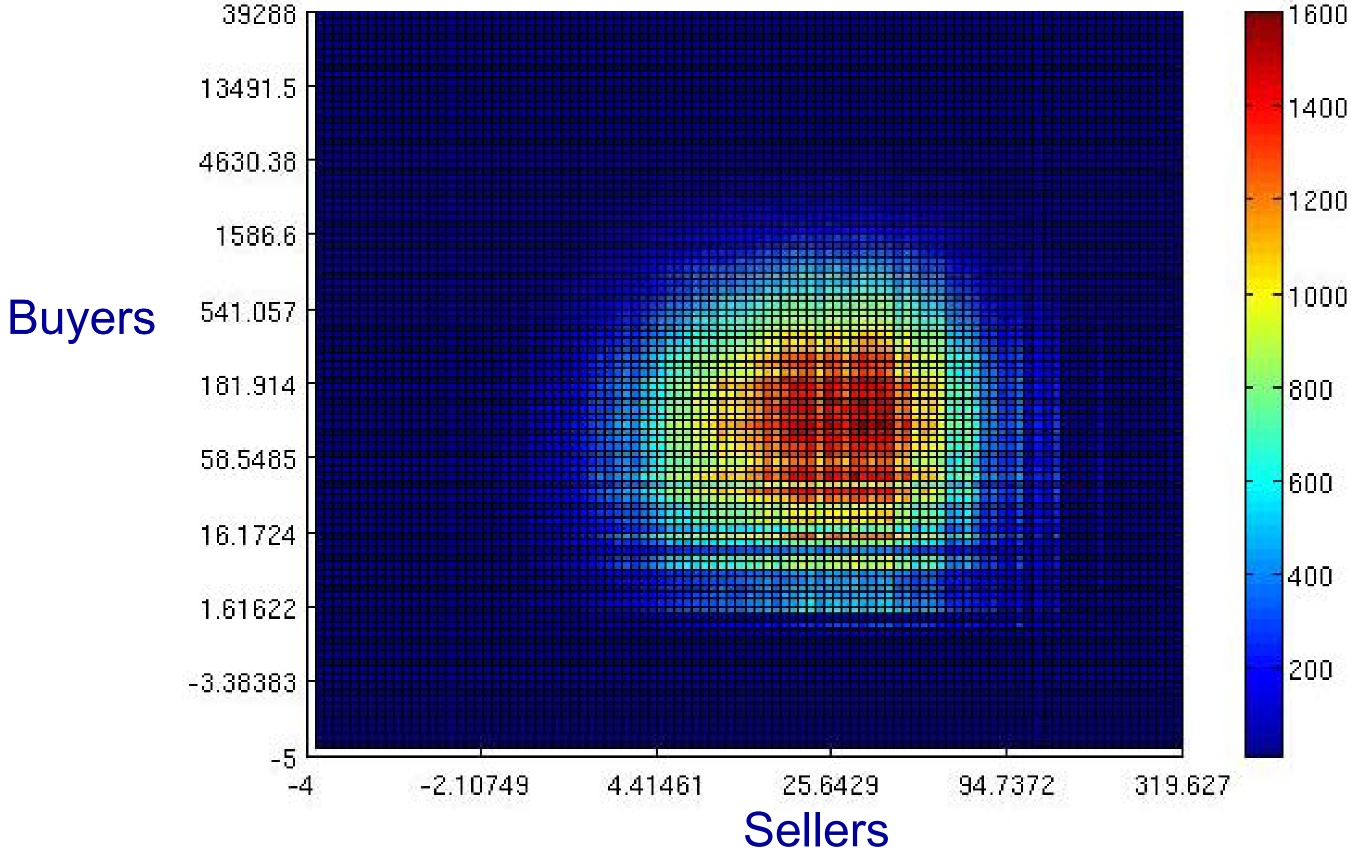

Auroral Diagram: Collectibles

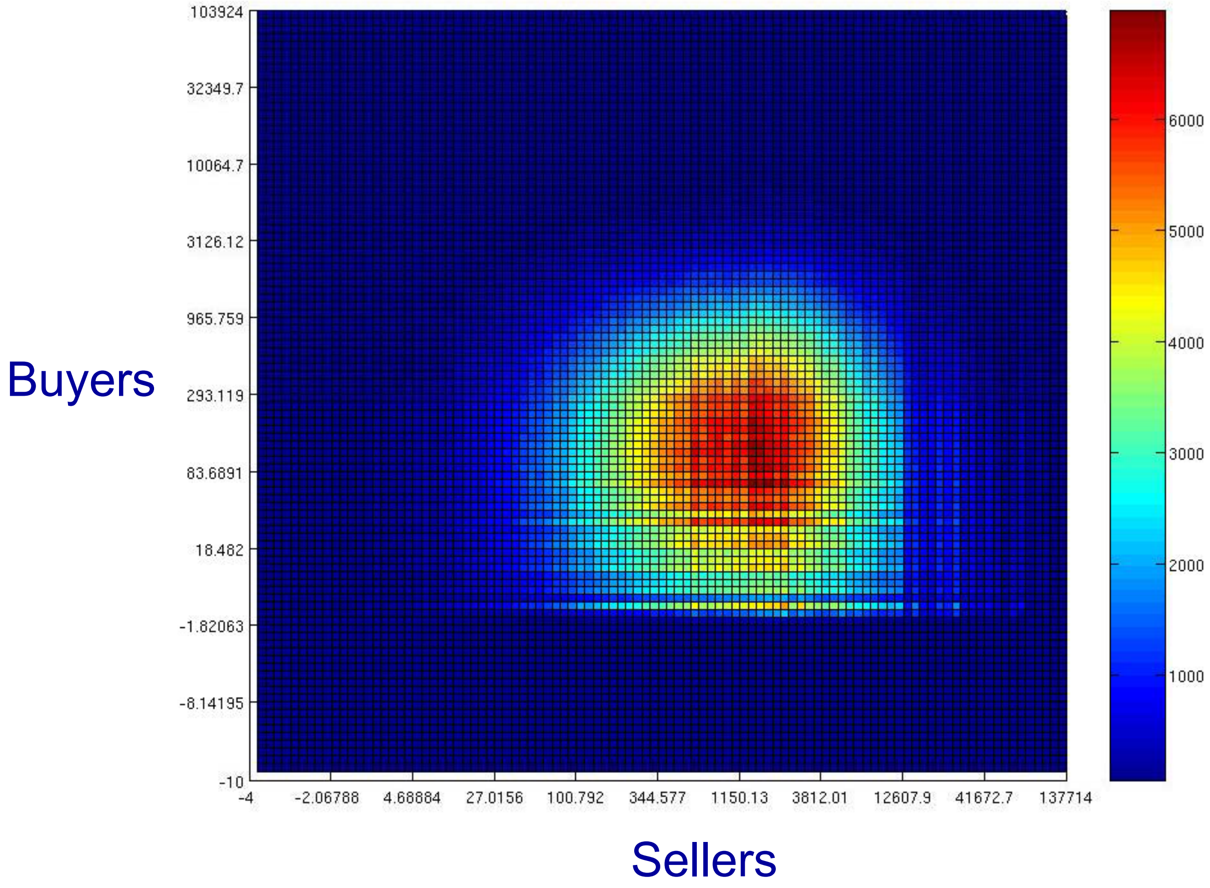

Motivation for Dynamic Reputation (Shen, Sundaresan 07)

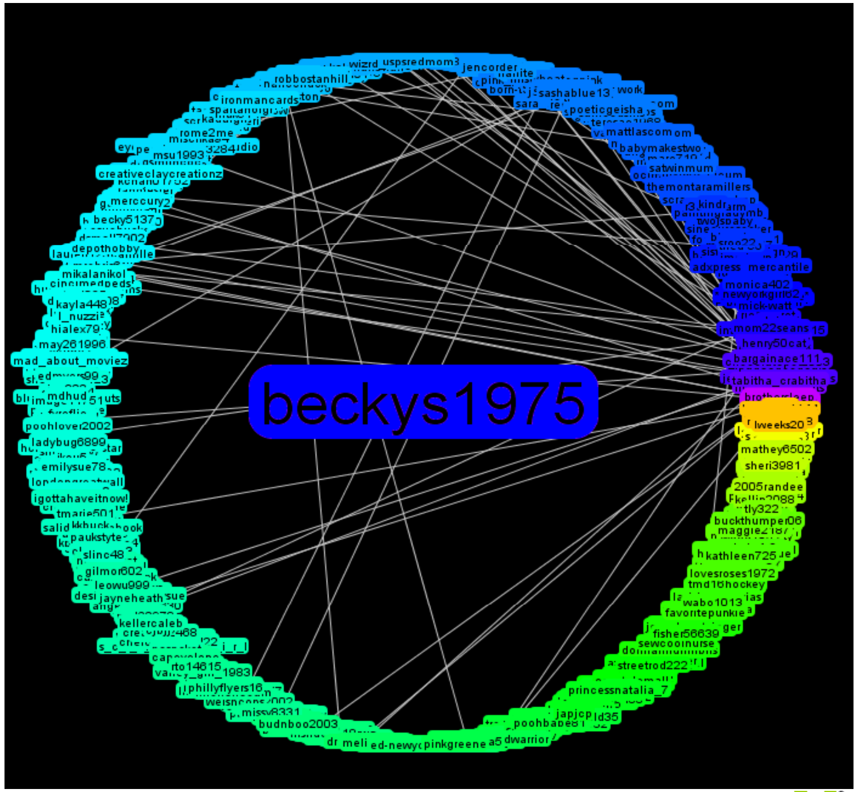

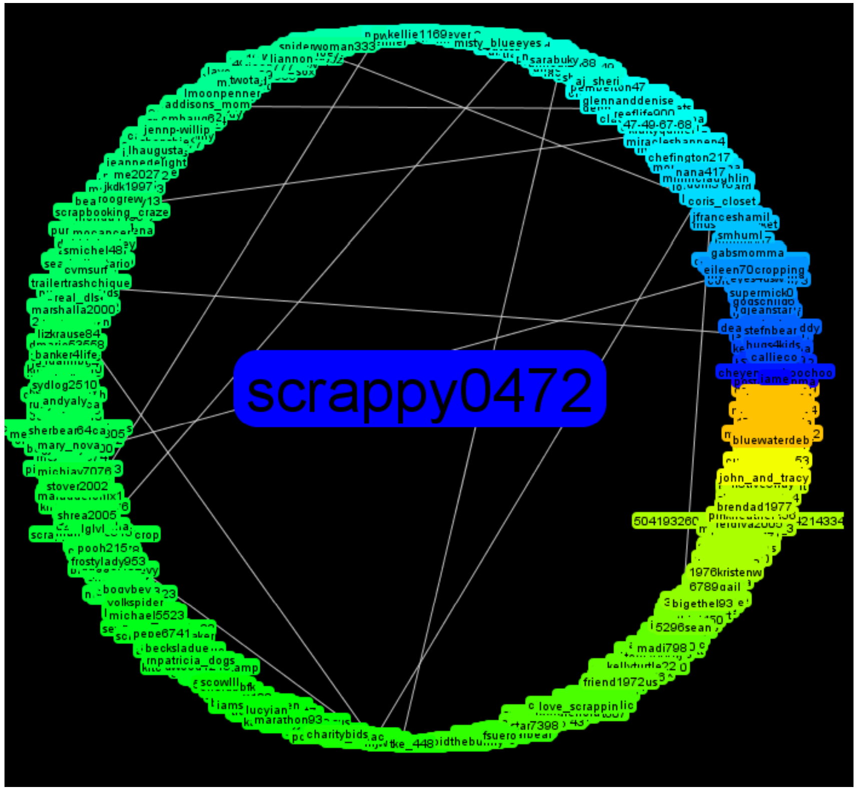

Trust in Different Categories

Stamps vs Antiques

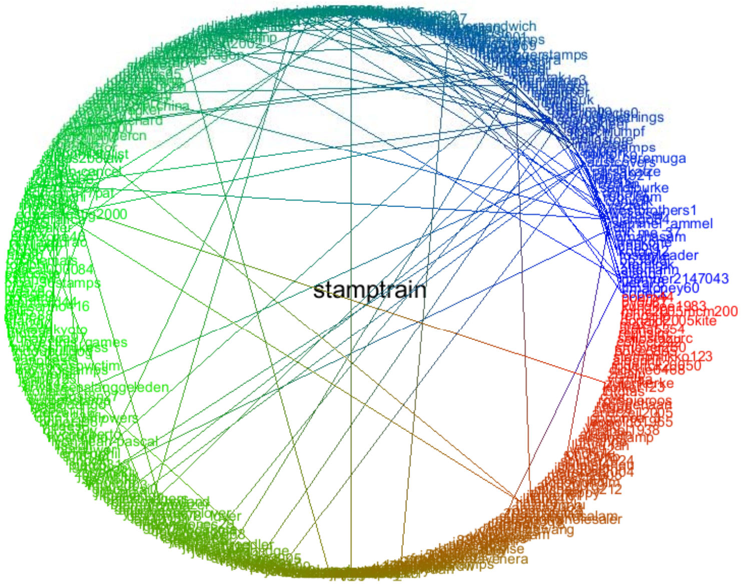

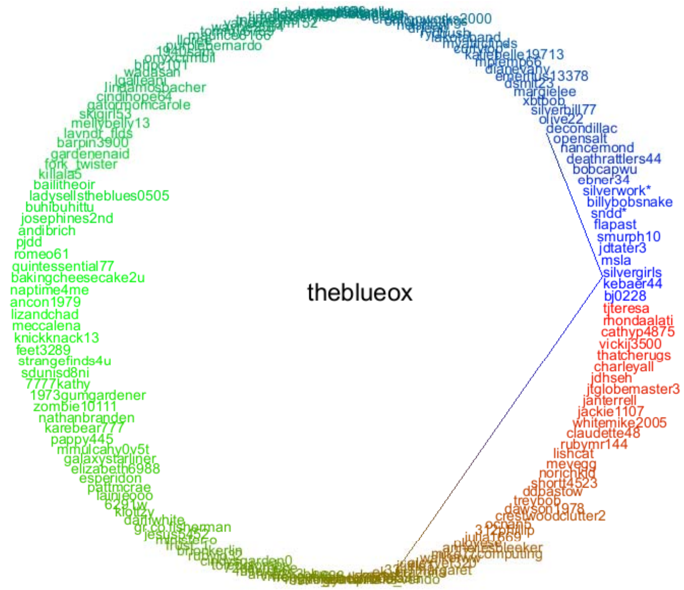

Spread the reputation

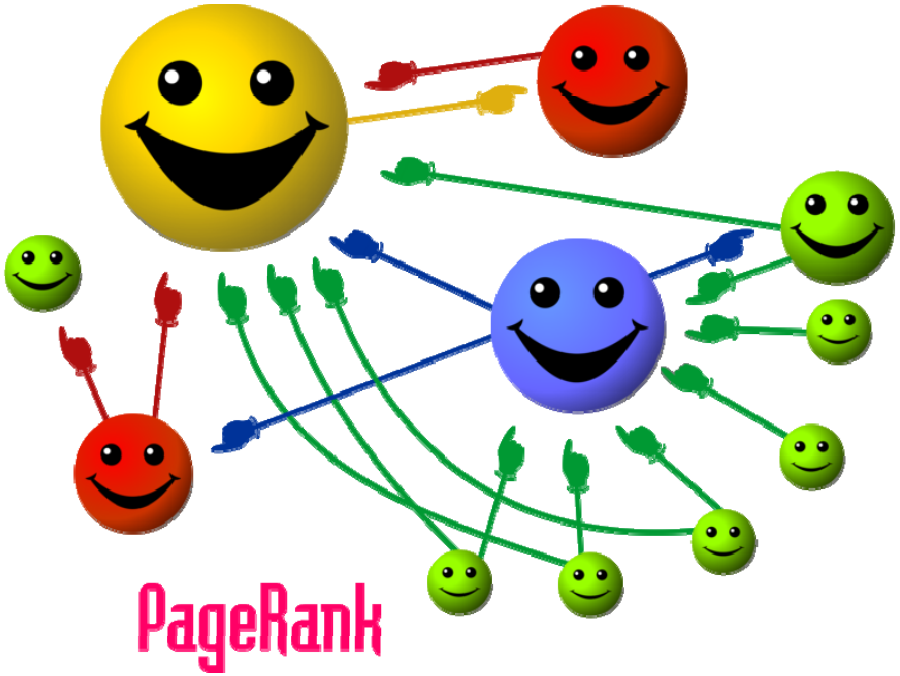

From Wikipedia
http://en.wikipedia.org/wiki/PageRank

The Web

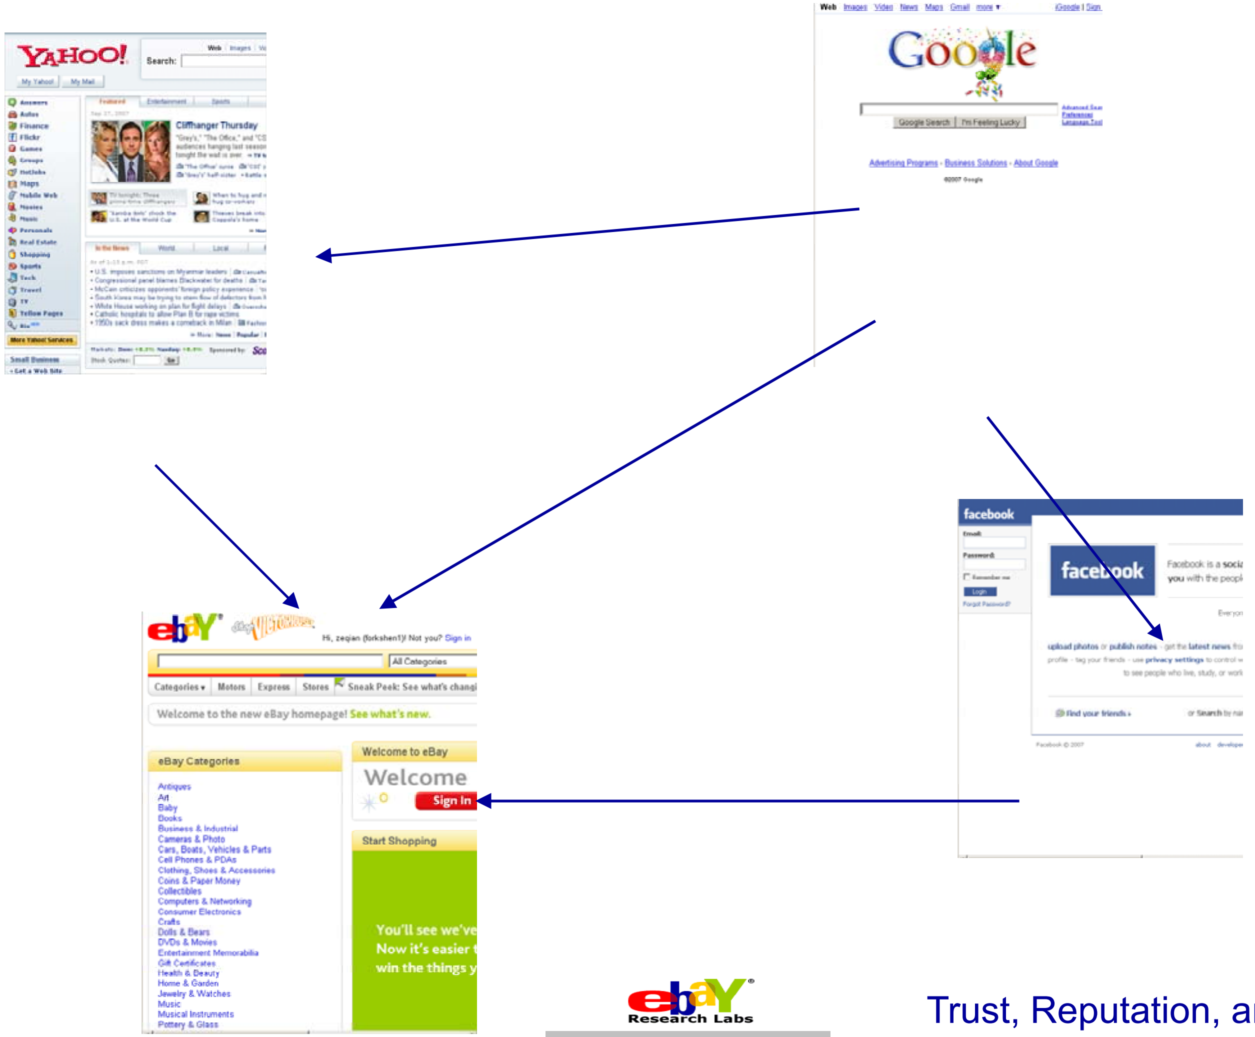
Search:
Cliffhanger Thursday
Google Search
I'm Feeling Lucky
facebook
Welcome to the new eBay homepage! See what's new.
eBay Categories
Welcome to eBay
Welcome
Sign In
Start Shopping

eBay

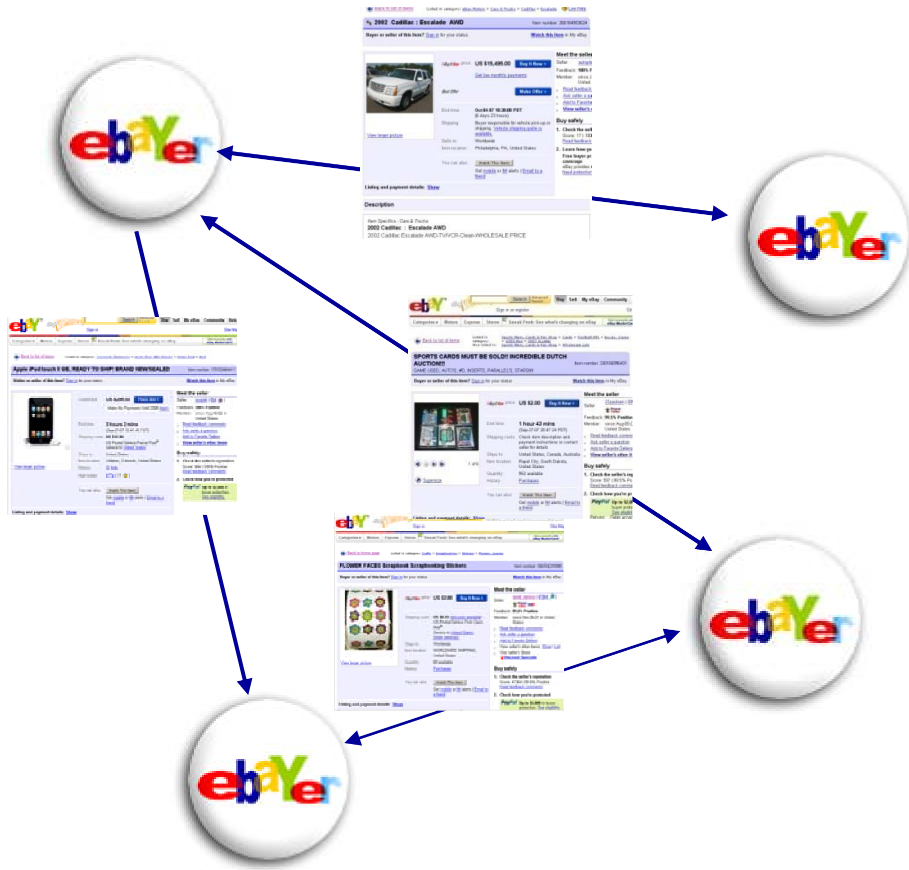

Dynamic Trust and Reputation (Sundaresan 06)

- Trust and (in turn) Reputation are evolving entities and need to be incrementally updated.

- As the actor a_i participates in a transaction c_{ij}^l with another actor a_j with reputation r_j then each entity – the 2 actors and the transaction have attached to them certain reputation.

Let a_i have reputation $r_i^{(l-1)}$ and aj have reputation $r_j^{(l-1)}$ before entering the transaction.

Let t_{ij} be a_i's trust for a_j and t_{ji} be a_j's trust for a_i expressed at this transaction.

The reputation of the transaction itself be r_c^l. Since transactions are all unique we could associate reputation with the aspects of the transactions like price, shipping cost, reputation of the participants, item category, auction format etc. to identify its reputation. This would be the implicit quality of the transaction.

Dynamic Trust and Reputation

- We can compute the new reputation after this transaction for each actor as

  $r_i^k = f(r_i^{k-1}, t_{ji}^{k-1}, \varsigma_{ji}, r_j^{k-1}, r_c^l)$

  $r_j^k = f(r_j^{k-1}, t_{ij}^{k-1}, \varsigma_{ij}, r_i^{k-1}, r_c^l)$

  $r_c^k = g(r_c^{k-1}, r_i^k, r_j^k)$

Where ς is the feedback score that the actors assign each other

Where f and g are bounded functions that appropriately dampen or enhance the reputations based upon the incoming factors.

Benefits of this approach

- Looks at reputation as constant at any observed time but changes as behavior of the actors change
- Can be applied to actors or to any entity within the system as long as it can be characterized based upon the parameters that describe it
- Takes into account up to date reputation measures of participating entities and updates all reputation post-transaction accordingly.

ReputationRank (Shen, Sundaresan 07)

- **Step1**:

  Compute edge weights W_{uv}

  $W_{uv} = F(price, time, \ldots)$

- **Step2**:

  Reputation propagation

  $R(u) = c\, E(u) + c \sum_{v \varepsilon N(u)} W_{vu}\, R(v)$

  Where c is chosen such that $\sum_u R(u) = 1$

  In matrix form

  $R' = c\, E + c\, W\, R$, W is the propagation matrix

The Good, The Bad

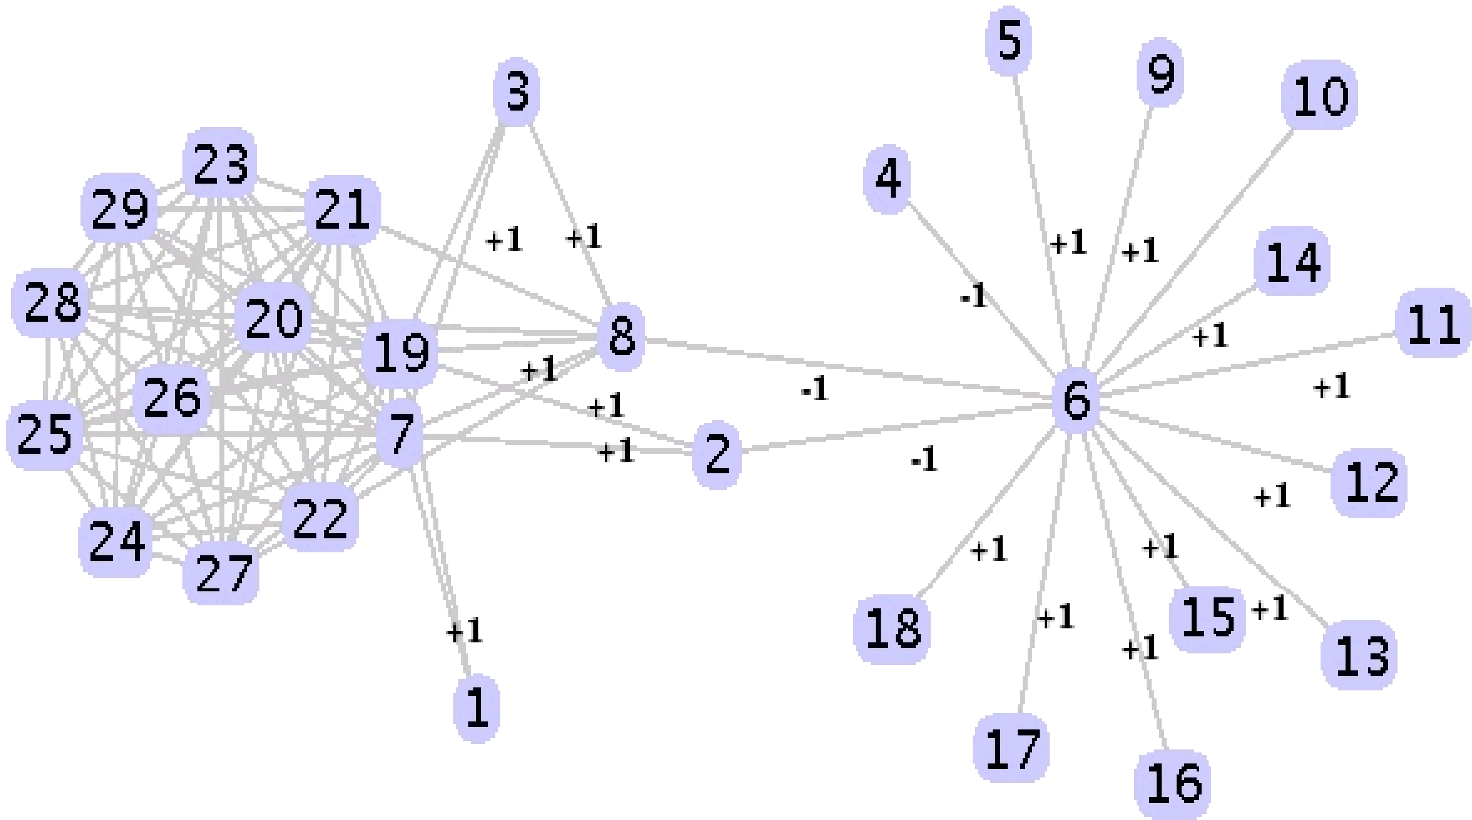

Other Advantages

- The reputation model is opaque and not easy to game with.
- Vector *E* gives us more control of user ranking
  - *Personalized ranking (E can be different for users based on their preferences)*
  - *Commercial interests*

Impact of Reputation

- Customer Support Cost
- User Stickiness

Object-level Trust and Reputation

- Trust and Reputation can be factored into every object that belongs the the environment (actors, transactions, widgets, etc.)
- Trust or relative reputation applies to each one of them
- Reputation is dynamic and is computed based on mutual trust and previous reputation

Transparent vs Opaque Reputation

- Transparency helps understand and improve negative behavior
- Opaque is useful to verify mechanism and also evaluate actors and avoid gaming
- Both are important in a reputation system

Reputation and Relevance Sort

- PageRank makes reputation as integral part of relevance sort
- A Marketplace Search like eBay is complex
  Diverse items, Diverse sellers, Diverse scenarios
- Reputation has to be combined with relevance and other factors like diversity
- Additionally needs to be personalized at some level

Identity

- Dellarocas(2000) showed attacks on reputation systems can be staged
- Resnick(1998) an easily modifiable identity (pseudonym) system creates incentive to misbehave without consequences on reputation

Identity and Reputation Portability

- Can you take your identity and reputation with you?
  - eBay Reputation scores into Amazon
  - Context matters for reputation (great credit score doesn't mean great reviewer!)
  - iKarma
    - Create a profile page, carry around the ikarma seal with you, the reputation is captured, managed, standardized, and used by ikarma
  - Trufina.com, sxip.com
    - Provide managed identity service that can be used anywhere on the net
    - Needs adoption
  - Opinity.com
    - Users can manage reputations
    - Apply reputation profiles for different context

Tagging and Trust

- With the explosion of social network sites, blogs, media content (images, audio, video) tagging is created a huge wave
- As the differential between producers and consumers turns huge the community (consisting of producers, consumers, others) is tapped to bridge the gap using tagging.
- Intention, Incentives and Trust models essential here.

The New Phenomena

- LinkedIn, Facebook, Twitter, …
  - What do connections mean?
  - What does rejection of a connection mean?
  - How do you assess the quality of any network?
  - Beyond glorified address books?

Summary

- Strong Identity and Longevity of actors to build a good trust and reputation system
- Trust is relative, Reputation is Global or Integrated
- Trust can be of different types
- Both Trust and Reputation can be dynamic
- Recommender systems can augment or use Trust systems
- Appropriate Intent and Incentives need to be identified when used to measure trust
- A Reputation system is weak without allowance for "complaints"
- Both Trust and Distrust have to be propagated
- Circles of Trust and Rings of Fraud are complementary