WikiRep: Digital Reputations in Virtual Communities

Mikalai Sabel Anurag Garg Roberto Battiti
[msabel,garo,battiti]@dit.unitn.it

Dipartimento di Informatica e Telecomunicazioni
Università degli Studi di Trento
This work is funded by Provincia Autonoma di Trento through the WILMA Project

Congresso Annuale AICA 2005

Outline

1. Motivation, context, and assumptions
   - Virtual communities and digital reputations
   - Virtual environment case study: Wiki
2. WikiRep — key ideas
   - Structure of a wiki-page: tree model
   - Comparing page versions: adoption coefficients
   - Rating inheritance
   - Feedback allocation
3. WikiRep — implementation details
   - Calculating adoption coefficients
   - Representing digital reputations

Virtual Communities and Environments

Definition from Wikipedia

A virtual community is a group whose participants are engaged in a dialog by means of information technologies, typically the Internet, to share information and values.
(also *online community* and *mediated community*)

Non-Collaborative

- forum/message-board
- electronic mailing list
- peer review system
- chat
- online games

Collaborative

- wiki
- CVS
- co-authoring
- online games

Digital reputations

Definition from Wikipedia

Reputation is the general opinion (social evaluation) of the public toward a person, a group of people, or an organization.
(we also use interchangeably with rating or *opinion toward an object*)

Digital reputation is a reputation handled by virtual environment, measured and expressed as number(s)

- reputation is measured by analyzing feedbacks

Advantages of using digital reputations:

- automate maintenance, require less human effort
- better scalability
- affordable for a wider range of situations
- provide motivation for active participation

(`slashdot.org`, `epinions.com`, `ebay.com`, `amazon.com`)

Digital reputations: what is missing

Available systems assume:

- each object has a single responsible 'owner'
- objects are independent from each other
- objects are static, complete

This does not hold for a collaborative environment!

- each object contains contributions of many users
- versions are closely related objects
- dynamic nature of content:
  often modifications $\longrightarrow$ evaluations become irrelevant

We suggest:

- *allocate feedback* among (co-)authors (*author allocation factors*)
- keep track of evaluations for individual versions:
  - allocate feedback among versions (*content allocation factors*)
  - reuse previous evaluations (*rating inheritance*)

Reputations: how to digitize

- Single scalar value is not enough!
  - not only ‘good’ vs. ‘bad’,
  - but also ‘unknown’ vs. ‘time-proved’
- Pair of numbers may be ok
  - we use pair (*value*; *quality*)
  - *value* = the evaluation of object’s merit (‘good-bad’ scale)
  - *quality* = significance of the evaluation

quality

- weight, normalized to the $[0, 1]$ interval
- 0 means a completely unreliable evaluation (never considered)
- 1 means a fully reliable evaluation (e.g. the source is trusted)

- Other approaches exist with 2 or more numbers

Case study virtual environment: Wiki

Wiki

web-site with pages created and maintained by visitors

- Collaborative editing tool
  - each *page* has many *version*
  - each *version* has an *author*
  - the last *version* is shown *by default*
  - but any *version* is accessible
- Concept and first implementation by Ward Cuningham in 1995
- Many implementations and installations exist
  - most known and successful: Wikipedia
- None employs digital reputations for usability and content quality
  - *why?*

WikiRep := Wiki + digital reputations

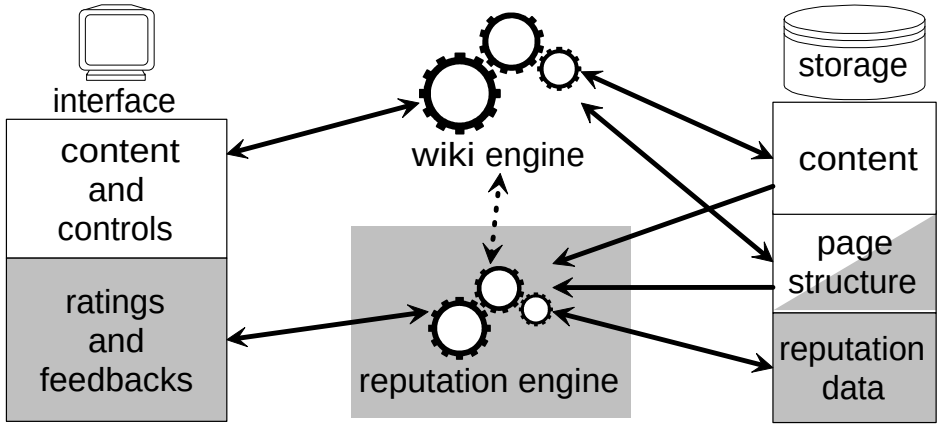

- Reputation mechanism is an add-on to the base Wiki system:
  - reputation data and additional information about pages structure
  - interface elements to visualize reputations and collect feedbacks
  - engine to manage reputations and interact with the base system
- Stored reputation data:
  - 'local' member-about-member opinions
  - 'global' object reputations (page ratings)
- Not stored:
  - global user reputations are derived from local opinions
  - each user can evaluate each page version only once

WikiRep Look-and-Feel

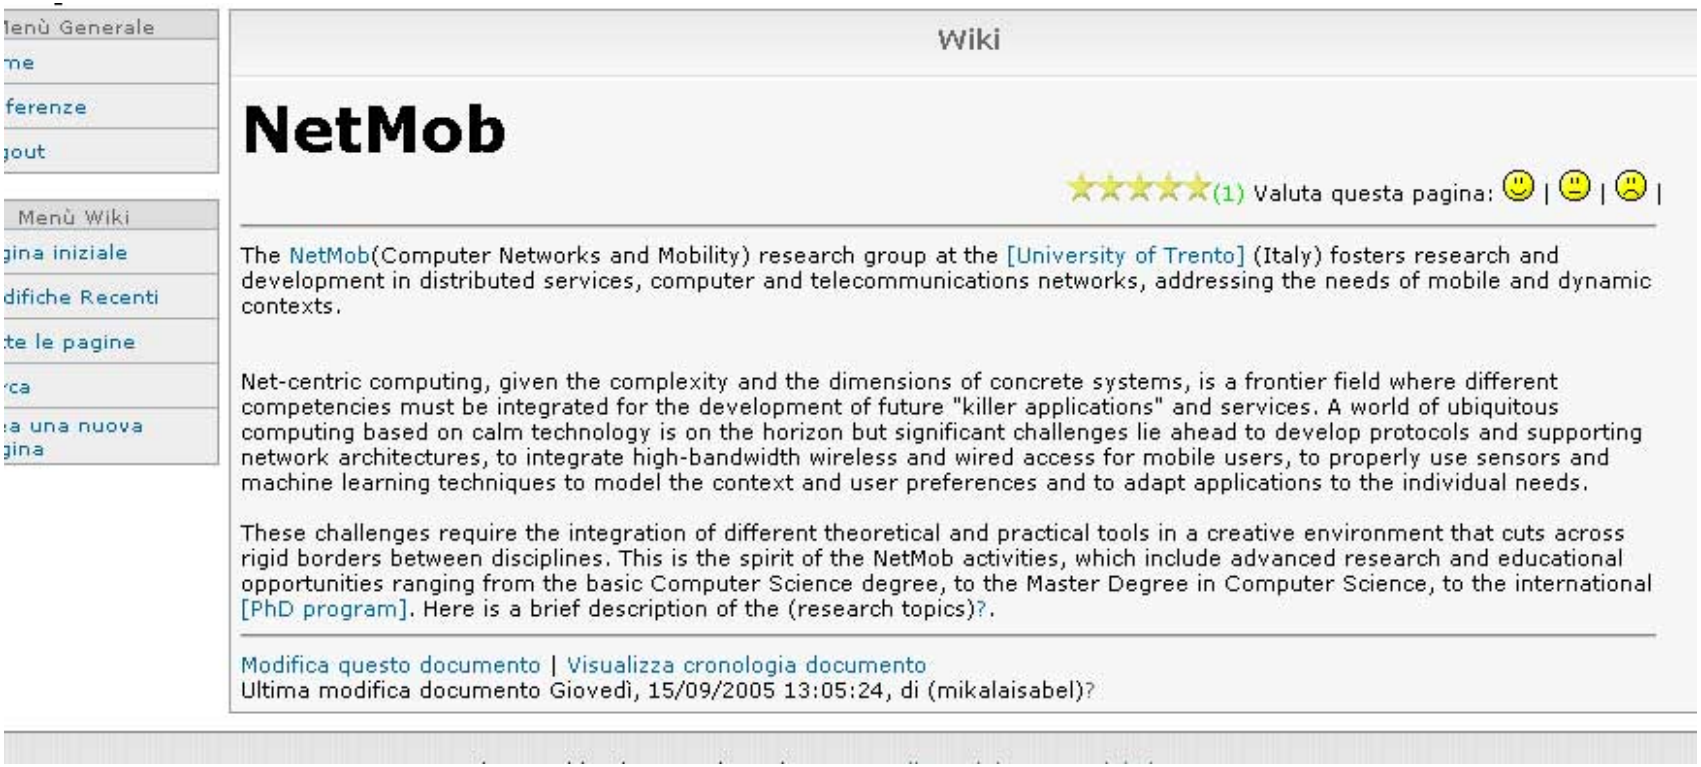

- based on eGroupWare
- modifications made do not lower the ease of use of Wiki
  - we aim at simplicity and high usability
- only two elements are added to the page visualization interface
  - a symbolic representation of page rating
  - buttons to leave feedback

Wiki page structure: tree model

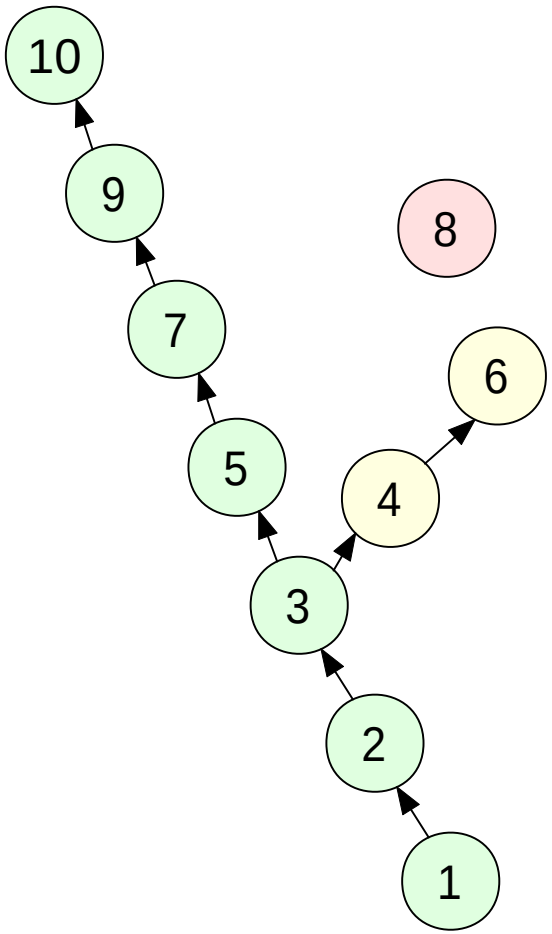

- Versions of the same document have similar content
  - we need a model for the structure
- Tree model reasonably simple and captures a typical Wiki well
  - *linear model* is too simple
  - *grid model* is too complex
- Each version j has at most one $parent(j)$
- A completely original version has no parent

Tree example: Wikipedia article *Virtual Community*

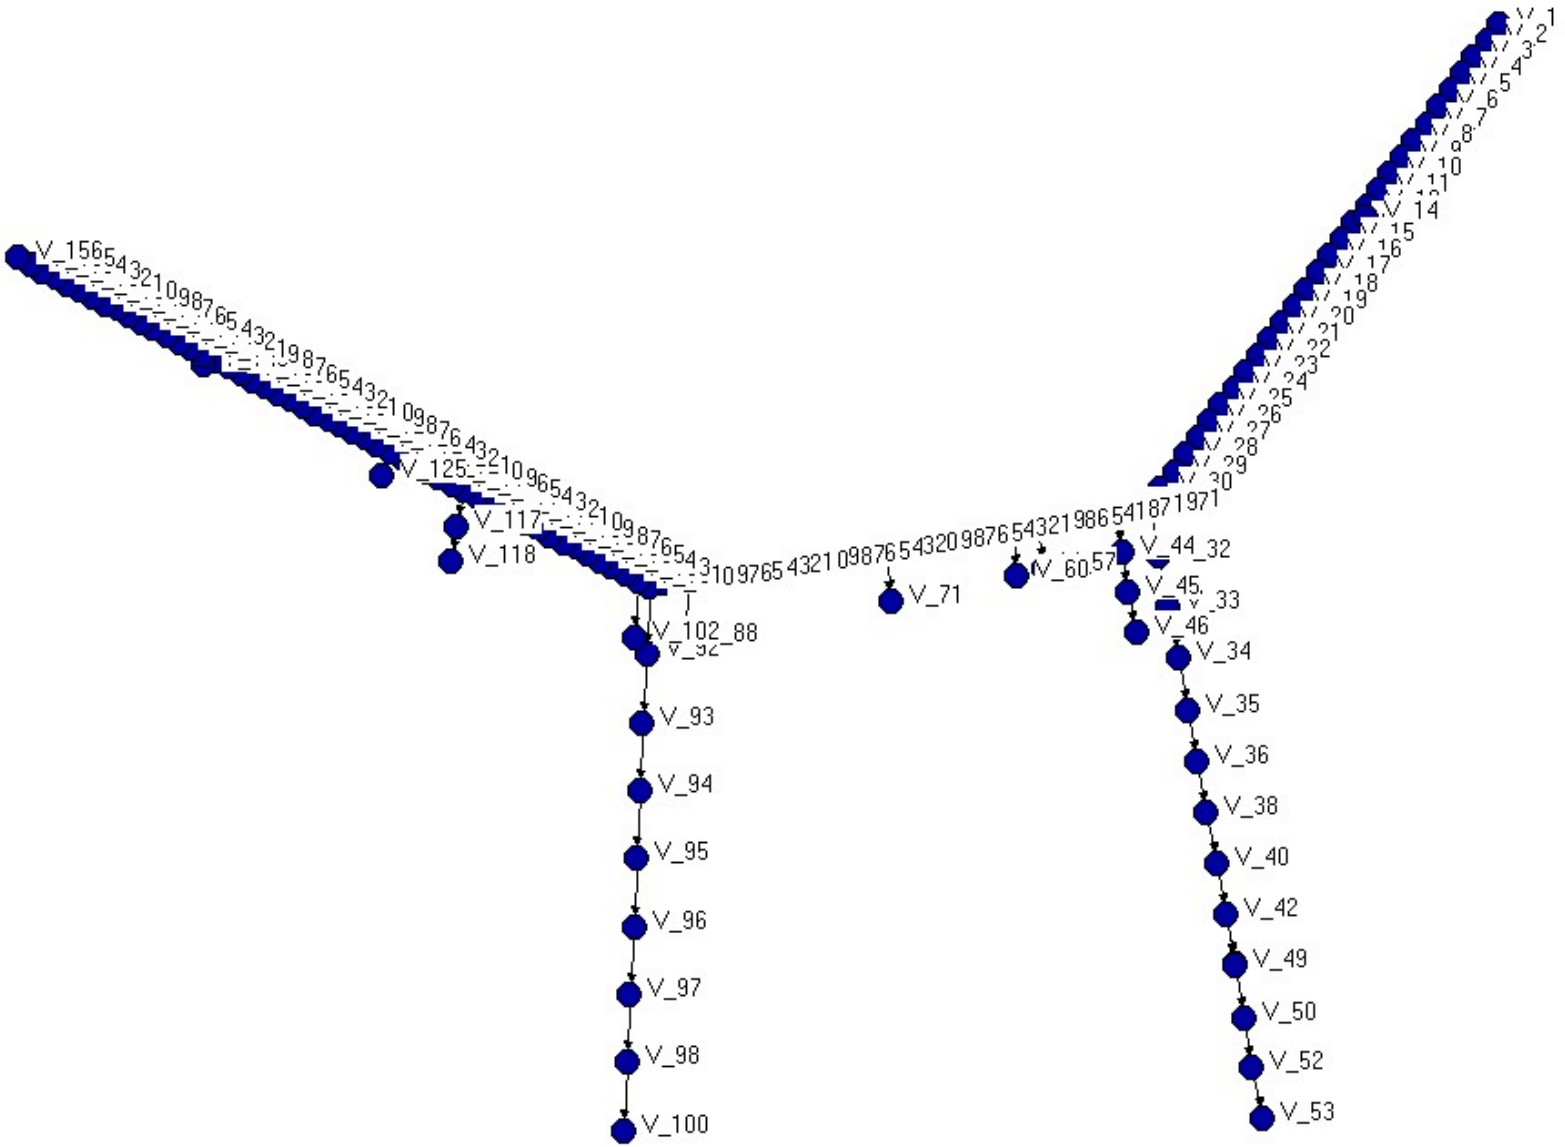

Adoption Coefficients: measure of versions' similarity

Adoption coefficient $a_{i,j}$

characterizes numerically similarity between
version i (older) and version j (newer)

- measures how much content of version i is preserved in version j
- on the $[0, 1]$ scale:
  - 0 for independent versions
  - 1 when j is a copy of i
- any automated algorithm to calculate is feasible
  - from naive text comparison
  - up to semantic-aware tools
  - depending on particular environment
- simplified notation for the tree page model: $a_j := a_{parent(j),j}$

AC example: Wikipedia article *Virtual Community*

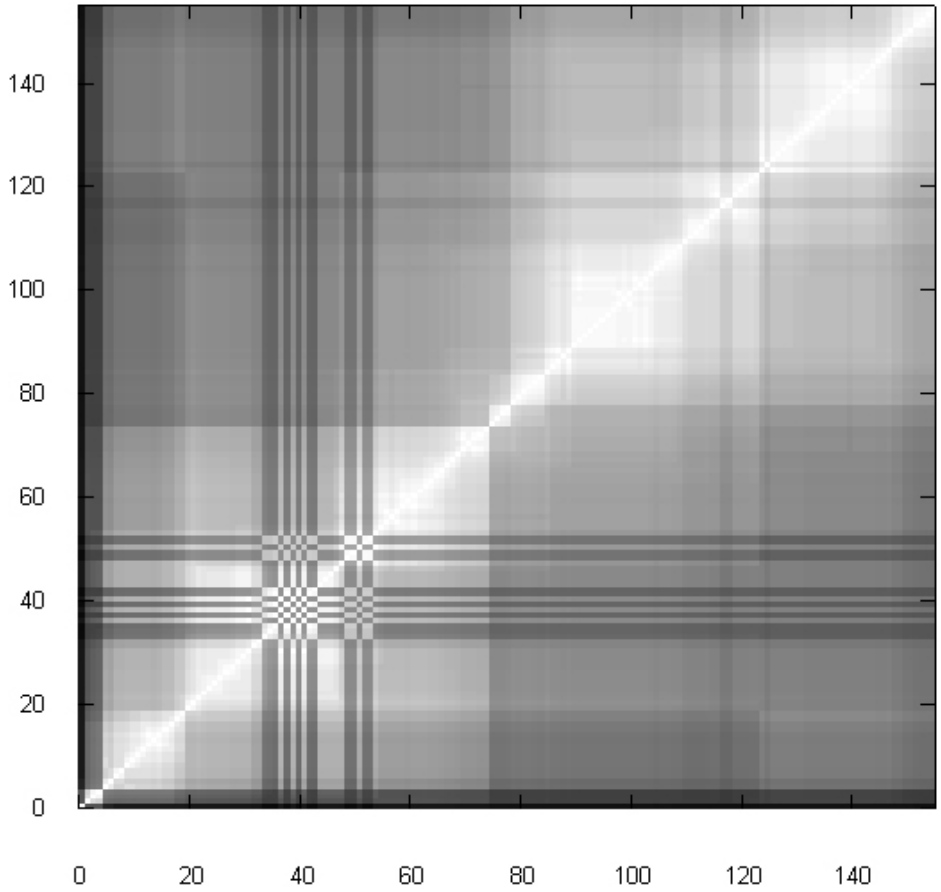

- full AC matrix

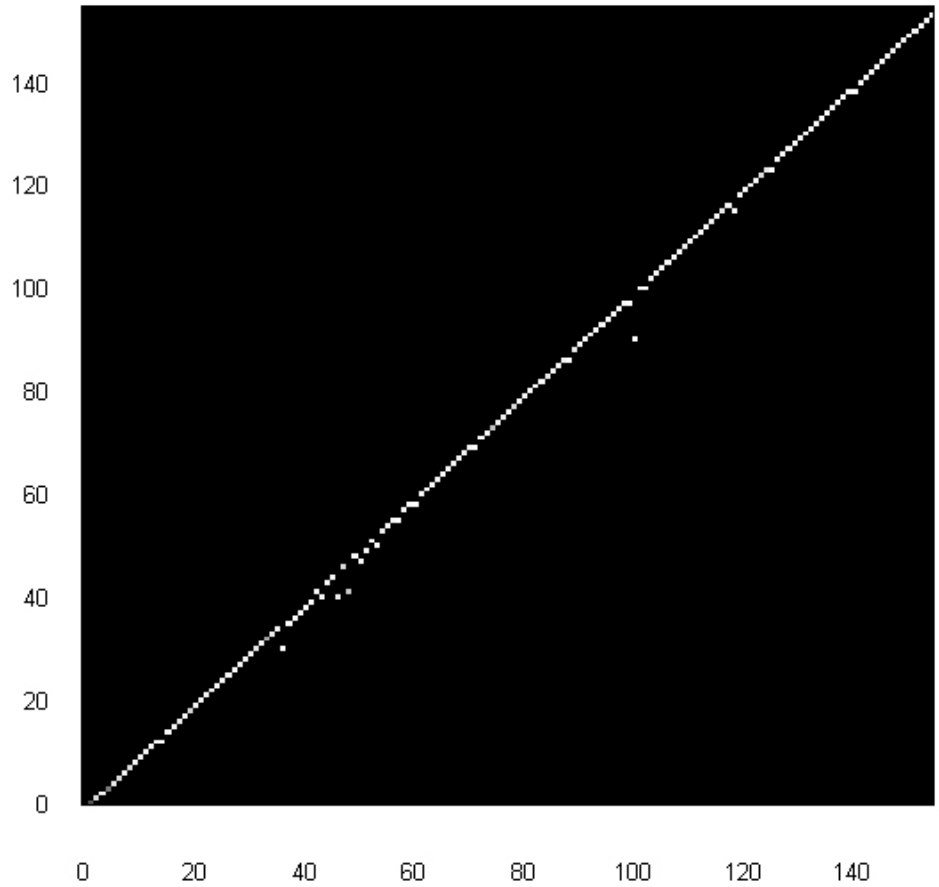

- only parents (stored data)

Rating inheritance

mechanism to reuse existing evaluations for parent version

- New version of a page partially inherits rating of its ancestor
  - proportionally to similarity (adoption coefficient to parent a_j)
  - let parent's rating be $(R_{parent(j)}; quality_{parent(j)})$
  - then new page is initiated with $(R_{parent(j)}; a_j \cdot quality_{parent(j)})$
- The value of the rating stays the same
- But its reliability (*quality*) decreases:
  - only fraction a_j of the content is inherited, and so is rating weight
  - the other part, $(1 - a_j)$, is new, not evaluated, has zero reliability

Feedback allocation

mechanism to distribute credit between versions and authors

- Reader leaves a feedback: mark *M*
  - during reading version *j*
  - clicks one of the buttons
- affects previous versions and authors
  - all involved page versions and authors get 'fractions' of the feedback
  - *contributors*(*j*) = versions (re-)used in *j*
  - the addition is a weighted summation
  - weight $w_{i,j}$ is content allocation factor
  - weight $u_{i,j}$ is author allocation factor
- Each *i* from *contributors*(*j*) gets:
  - $(M; w_{i,j})$ for its content
  - $(M; u_{i,j})$ for its author

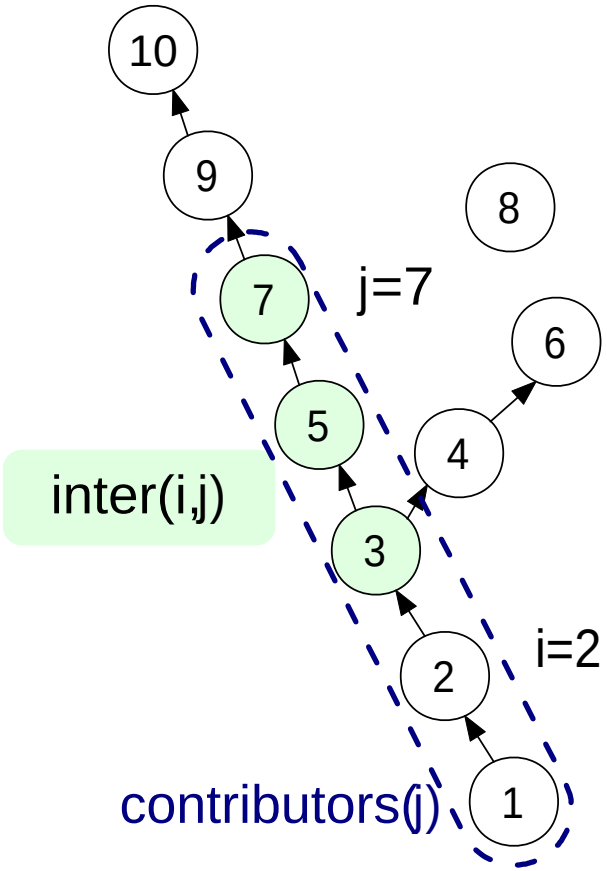

Content allocation factor $w_{i,j}$

$$w_{i,j} = \begin{cases} 0 & \text{if } i \notin \textit{contributors}(j) \\ 1 & \text{if } i = j \\ \prod_{k \in \textit{inter}(i,j)} a_k & \text{if } i \in \textit{contributors}(j) \text{ and } i \neq j \end{cases}$$

- $\textit{inter}(i, j)$ = set of j's ancestors between j and i
  - excluding i and including j
- characterizes merit of the page version as a whole
  - not just the added value beyond its ancestors
- for itself $w_{j,j} = 1$
  - the version that attracted the feedback gets the 'full' weight
  - because it is what the reader expects
- each parent in the chain gets a_i-less evaluation
  - because a_j fraction of the parent's content is reused

Author allocation factor $u_{i,j}$

$$u_{i,j} = \begin{cases} 0 & \text{if } i \notin \textit{contributors}(j) \\ 1 - a_j & \text{if } i = j \\ (1 - a_i) \prod_{k \in \textit{inter}(i,j)} a_k & \text{if } i \in \textit{contributors}(j) \text{ and } i \neq j \end{cases}$$

- Measures the original contribution of version j compared to i

Theorem

Sum of $u_{i,j}$ between a given version and all other versions is 1

$$\sum_i u_{i,j} = 1, \textit{for any } j$$

user can never gain reputation by creating superfluous versions

WikiRep: calculating adoption coefficients

- Whenever a new version of a Wiki page is saved
  - calculate adoption coefficients to all previous versions
  - the largest (latest if a tie) determines the parent
- Each version has:
  - text
  - author
  - modification time
- Adoption coefficients are calculated by comparing texts:
  - the two texts are divided into blocks, separators are `.,:;!?`
  - a form of *edit distance* is found, using the blocks as characters
  - minimal number of insert and delete operations on characters, needed to transform one text into the other: $(N_{inserted} + N_{deleted})$
  - normalize to the [0,1] range
  - subtract from 1 to obtain *similarity*

$$a_{old,new} = 1 - \frac{N_{inserted} + N_{deleted}}{N_{total,new} + N_{deleted}}$$

WikiRep: reputation representation

- based on the ROCQ (Reputations, Opinion, Credibility and Quality) scheme, opinion or reputation is a pair ($value$; $quality$)
- $value$ is a real number in the [0,1] range
- $quality$ characterizes significance of the evaluation
- $quality$ = probability that interval $[value - \Delta_r/2; value + \Delta_r/2]$ holds the actual mean of the underlying mark distribution
  - assume that marks are independent $\rightarrow$ mean has normal distribution (the central limit theorem)
  - Δ_r is a parameter to be chosen

$$quality = 1 - I_{(N-1)/(N-1+t^2)}\left(\frac{N-1}{2}, \frac{1}{2}\right)$$

where $I_x(a, b)$ is incomplete beta function, and

$$t^2 = \left(\frac{\Delta_r}{2}\right)^2 \frac{N^2(N-1)}{N \cdot S_q - (S)^2}$$

S_q is the weighted sum of squares of collected evaluations marks

Conclusions

- Extends reputation algorithms for user-managed virtual communities with multi-authored, multi-versioned objects.
- The algorithm relies on
  - maintaining reputations of individual page versions
  - proper allocating of feedback credit
  - established page reputations are reused
- Implemented using Wiki as the base system
  - the system is currently being deployed