24th European Conference on Artificial Intelligence - ECAI 2020
Santiago de Compostela, Spain

# From Eigentrust to a Trust-measuring Algorithm in the Max-Plus Algebra

**Juan Afanador**[1] and **Nir Oren**[2] and **Murilo Baptista**[3] and **Maria Araujo**[4]

**Abstract.** Eigentrust is a simple and popular method for trust computation, which uses both direct and indirect information about individual performance to provide a global trust rating. This final trust value is based on eigenvectors computed through the Power Method. However, under certain network topologies, the Power Method cannot be used to identify appropriate eigenvectors. After characterising these cases, we overcome Eigentrust's limitations by extending the algorithm's core ideas into the Max-Plus Algebra. An empirical evaluation of our new approach demonstrates its superiority to Eigentrust.

## 1 Introduction

Eigentrust is a popular approach to the calculation of trust with local information [10]. The functioning of the algorithm presupposes the free flow of interaction experiences across a network of agents (i.e., the transfer of reputation information across the system), thus improving or augmenting the information set (based on direct experience) of the querying agent. Eigentrust views trust as a normalised score of reliability computed through the querying agent's own (direct) experience, and the indirect experience information gathered from other agents in the system, i.e., reputation.

The query may be replicated by any other agent within the initial vicinity, and their neighbours' neighbours. This information is aggregated into a single value where the credibility parameters specified by the agent posing the query serve as weights. The resulting measure, computed using the Eigentrust algorithm, globally converges to the level of trust placed in the agent about which the query was made, and the updated trust values of all other agents in the system.

At the technical level, whenever a query is issued, Eigentrust stores local trust scores in a matrix which is then multiplied by an initial vector of trust values. The resultant vector is again multiplied by the trust matrix, and the operation repeated until the vector converges to a fixed-point. This final vector is expected to coincide with the dominant eigenvector of the original trust matrix, and the final vector is treated as an accurate global trust value for the agents in the system.

Due to its simplicity, theoretical foundations, and resulting empirical behaviour, Eigentrust is widely used in domains such as P2P systems [13], Internet-of-Things architectures [2], and Ad-hoc Sensor Networks [18]. While Eigentrust performs particularly well in networks of homogeneous agents such as P2P systems, it does not work as well in other environments.

[1] University of Edinburgh, Scotland, email: jafanado@ed.ac.uk
[2] University of Aberdeen, Scotland, email: n.oren@abdn.ac.uk
[3] University of Aberdeen, Scotland, email: m.baptista@abdn.ac.uk
[4] University of Aberdeen, Scotland, email: r02maa16@abdn.ac.uk

In this paper we study the performance of Eigentrust in networks with various degrees of connectivity, and describe the cases where the algorithm accurately predicts global trust scores, others where it is somehow inaccurate but useful (e.g., where it can spot malicious behaviour), and those where it may be misleading. We provide a theoretical characterisation of all these cases to build on Eigentrust's core ideas toward a more generally applicable procedure.

Our goal is to formulate an algorithm that operates across more diverse environments than Eigentrust does. We argue that Eigentrust performs poorly in precisely those cases where convergence to the dominant eigenvector does not occur. By framing the trust-measuring problem within a different algebraic structure — the Max-Plus Algebra [7] — we are able to obtain informative trust ratings in those situations where Eigentrust fails. We argue that this occurs —- in part — due to Eigentrust's conflation of the case where an agent has no trust in another, and the case where an agent is unable to interact with another, and the fundamental idea of our approach is to differentiate between these two situations.

Our core contributions are therefore as follows: (1) A characterisation of where, and why, Eigentrust performs poorly; (2) A new Eigentrust-like algorithm based on the Max-Plus algebra which overcomes Eigentrust's weaknesses; (3) An empirical evaluation of the new algorithm's performance compared to Eigentrust.

The remainder of this paper is structured as follows. The next section describes the Eigentrust algorithm and the algebraic assumptions it is founded upon. Section 3 describes the Max-Plus Algebra, and characterises the corresponding trust measuring problem. In Section 4 we present the results of our approach. The last two sections discuss our findings and provide suggestions for future work.

## 2 Eigentrust

### 2.1 The Eigentrust Algorithm

Eigentrust considers interactions between pairs of agents and observes the corresponding outcome [10]. If we let $s_{ij}$ denote the difference between the number of successful and unsuccessful interactions between agents $i$ and $j$, then $c_{ij} = max(s_{ij}, 0)/\sum_k max(s_{ik}, 0)$ can be viewed as a normalised measure of trust between $i$ and $j$. Eigentrust assumes that trust is transitive, i.e., $i$'s trust in $k$ can be computed from its level of trust in $j$, and $j$'s trust in $k$:

$$t_{ik} = \sum_j c_{ij} c_{jk} \tag{1}$$

this property means that trust in every agent, within a connected component, can be computed in a similar manner.

The set of $c_{ij}$ values can be represented as a *trust matrix* $\mathbf{C}$ of local trust scores. To capture the transitive nature of trust, $\mathbf{C}$ must be applied to a vector of initial trust values $\mathbf{r}$, depending on the relative position of the agent about whom the query was made. Repeating this multiplication incorporates direct and indirect trust information, about all relevant neighbours, into the production of a vector of stable trust values for every agent in the system.

Repeated multiplication of $\mathbf{C}$ reduces the importance of the precise contents of $\mathbf{r}$, which can in fact (effectively) be random. The repeated multiplication of $\mathbf{C}$ then guarantees that the final trust ranking converges to the "true" distribution of trust values as given by the matrix's dominant eigenvector — provided it is unique. The procedure through which the dominant eigenvector is calculated is called the *Power Method* [5]. Algorithm 1 summarises the Eigentrust approach to quantifying trust. We note that the final trust value associated with agent $i$ occurs at index $i$ within the eigenvector.

**Algorithm 1** Eigentrust

```
1: t^(0) ← r
2: repeat
3:     t^(k+1) ← C^T t^(k)
4:     δ ← |t^(k+1) − t^(k)|
5: until δ < ε
```

## 2.2 The Algebraic Conditions Behind Eigentrust

Since Eigentrust will only operate correctly when the Power Method converges, we now consider one case where Eigentrust performs poorly due to the nonexistence of the dominant eigenvector, and the outright unattainability of convergence. We then characterise one situation where the Power Method is applicable but uninformative, and another where it is fully applicable.

### 2.2.1 Diagonalisable Square Matrices

Suppose the matrix $\mathbf{C}$ satisfies the following two conditions:

1. $\mathbf{C}$ is diagonalisable, i.e., there exists an invertible matrix $\mathbf{P}$ such that $\mathbf{D} = \mathbf{P}^{-1}\mathbf{C}\mathbf{P}$. Here, $\mathbf{D}$ is a diagonal matrix; and
2. $\mathbf{C}$ has a dominant eigenvalue $\lambda_0$. That is, if $\lambda_0, \ldots, \lambda_{n-1}$ are the eigenvalues of $\mathbf{C}$, then it is the case that $|\lambda_0| > |\lambda_i|$ for $i = 1, \ldots, n-1$. The eigenvector associated with the dominant eigenvalue is termed the *dominant eigenvector*.

In such a situation, the Power Method converges to the dominant eigenvector of $\mathbf{C}$. Note that if the matrix $\mathbf{C}$ satisfies the two conditions above, then these results also apply to its transpose.

**Example 1** *The following matrix depicts a situation where agents $a_0$, $a_1$ and $a_2$ have limited information about one another.*

$$A = \begin{array}{c} \begin{array}{ccc} a_0 & a_1 & a_2 \end{array} \\ \begin{pmatrix} 0.75 & 0 & 0.25 \\ 0 & 1 & 0 \\ 0.25 & 0 & 0.75 \end{pmatrix} \end{array}$$

*Despite being diagonalisable, the three eigenvalues of $\mathbf{A}$ are $\lambda_0 = \lambda_1 = 1$ and $\lambda_2 = 0.5$. As $\mathbf{A}$ does not have a dominant eigenvalue, the convergence of the Power Method cannot be guaranteed. The lack of connectivity between agents induces an unstable outcome.*

*While Eigentrust could — potentially — be applied to each connected component in such a graph, this would require additional knowledge of the network structure. Furthermore, in dynamic situations questions arise as to how trust across connected components should be merged when the topology of the network changes.*

**Example 2** *Consider the following trust matrix.*

$$\mathbf{B} = \begin{array}{c} \begin{array}{ccc} a_0 & a_1 & a_2 \end{array} \\ \begin{pmatrix} 0.1 & 0.55 & 0.35 \\ 0 & 0.8 & 0.2 \\ 0 & 0 & 1 \end{pmatrix} \end{array}$$

*This matrix is upper triangular with distinct diagonal entries, and has eigenvalues $\lambda_0 = 1$, $\lambda_1 = 0.8$ and $\lambda_2 = 0.1$. Therefore, for a random vector $\mathbf{v} \in \mathbb{R}^3$ it is the case that $\lim_{k \longrightarrow \infty} \left(\mathbf{B}^T\right)^k \cdot \mathbf{v} = \vec{\pi}$, where $\vec{\pi} = (0, 0, 1)^T$ is the eigenvector associated with the dominant eigenvalue $\lambda_0 = 1$. According to this matrix, and by extension to Eigentrust, only agent $a_2$ can be trusted. A decision-maker views all other options as equally irrelevant, which may not be informative enough in some applications (e.g., when $a_2$ is unable to provide a service).*

### 2.2.2 Positive Stochastic Matrices

Suppose $\mathbf{C}$ is a square, positive, and stochastic matrix. Using the Perron-Frobenius Theorem [15] and the Jordan decomposition of $\mathbf{C}$, it is possible to show that there exists a unique dominant eigenvector and that the limit $\lim_{k \longrightarrow \infty} \left(\mathbf{C}^T\right)^k \cdot \mathbf{v}$ exists and converges to the same value for any initial random vector $\mathbf{v}$. Note that the existence of a dominant eigenvector is a consequence of the theorem's conditions. Lemma 1 and Proposition 1 below, which consider a more general case, are built upon these observations.

**Example 3** *Consider the following matrix.*

$$\mathbf{C} = \begin{array}{c} \begin{array}{ccc} a_0 & a_1 & a_2 \end{array} \\ \begin{pmatrix} 0.15 & 0.55 & 0.3 \\ 0.41 & 0.53 & 0.06 \\ 0.18 & 0.62 & 0.2 \end{pmatrix} \end{array}$$

*This matrix's dominant eigenvalue is $\lambda_0 = 1$. Thus, for any random ranking $\mathbf{v}$ we have $\left(\mathbf{C}^T\right)^k \cdot \mathbf{v} \longrightarrow \vec{\pi}$, as $k \longrightarrow \infty$, where $\vec{\pi} = (0.3, 0.6, 0.1)^T$, and the most trusted agent within the system is $a_1$, followed by $a_0$ and then $a_2$. Here, the positive, square and stochastic nature of such a matrix ensures that Eigentrust works as expected.*

Upper triangularity with distinct diagonal entries in Example 2 guarantees diagonalisability, and hence convergence of the Power Method. In Example 1, however, convergence is not attained despite diagonalisability via symmetry. Convergence in Example 3 is guaranteed via the Perron-Frobenius Theorem for positive (stochastic) matrices. We argue that Eigentrust's performance depends on whether a dominant unique eigenvector does or does not exist. In many useful cases, we argue that a dominant unique eigenvector does not exist.

The Perron-Frobenius Theorem has been generalised to cater for non-negative and irreducible matrices [17]. Our aim is to build on these results in the context of the Max-Plus Algebra, providing a new trust-measuring procedure also applicable to reducible matrices. In the following, we will introduce a version of the Perron-Frobenius Theorem for irreducible matrices, and consider its implications to the Eigentrust algorithm.

#### 2.2.3 Non-negative Stochastic Matrices

**Definition 1 (Irreducible matrix [17])** *An $n \times n$ matrix $\mathbf{A}$ is said to be* irreducible *if there exists no permutation of coordinates such that:*

$$\mathbf{PAP}^T = \begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{0} & \mathbf{A}_{22} \end{pmatrix} \tag{2}$$

*where $\mathbf{P}$ is an $n \times n$ permutation matrix with each row and column having a single entry equal to one and the rest full of zeros; while $\mathbf{A}_{11}$ and $\mathbf{A}_{22}$ are non-trivial (i.e., their size is greater than 0) square matrices. In other words, an irreducible matrix cannot be converted into a block upper-triangular matrix via simultaneous row/column permutations. A matrix is* reducible *if it is not irreducible.*

**Theorem 1 (Perron - Frobenius Theorem [15])** *If $\mathbf{C} = (c_{ij})$ is an $n \times n$ irreducible non-negative matrix with spectral radius[5] $\rho(\mathbf{C}) = \lambda_0$ , then:*

1. *$\lambda_0 \in \mathbb{R}^+$ is a simple eigenvalue of $\mathbf{C}$, called the Perron-Frobenius eigenvalue.*
2. *$\lambda_0$ can be associated with unique (up to a constant) and strictly positive left and right eigenvectors.*

**Lemma 1 ([3])** *Let $\mathbf{C}$ denote an irreducible non-negative stochastic matrix. The set of eigenvalues of $\mathbf{C}$ has a maximal eigenvalue equal to 1, and an associated left eigenvector describing a probability distribution $\vec{\pi}$ over the set of interacting nodes.*

Given the properties of the trust matrix described before and Lemma 1, it trivially follows that:

**Proposition 1** *If $\mathbf{C}$ is an irreducible non-negative trust matrix, then the Eigentrust algorithm yields the right eigenvector associated to the Perron-Frobenius eigenvalue of $\mathbf{C}$.*

## 3 Trust Measuring in the Max-Plus Algebra

Eigentrust does not distinguish the lack of interactions between two nodes from the impossibility of such interactions due to the network's topology. We redefine $\mathbf{C} = (C_{ij})$ over the corresponding set of edges ($E$) to differentiate between these two cases as follows:

$$C_{ij} = \begin{cases} c_{ij} & \text{if } (i,j) \in E \\ -\infty & \text{otherwise} \end{cases} \tag{3}$$

This differentiation gives rise to a distinct algebraic structure that facilitates the measuring of trust. The algebraic structure is an idempotent, commutative semiring (dioid) known as the *Max-Plus Algebra* [7], or *max-plus* for short. Within max-plus, we extend the Eigentrust algorithm to cater for multi-agent systems with reducible trust matrices. To this end, we must consider how basic operations can be performed within max-plus, before examining trust computation within the algebra.

[5] Given a matrix $\mathbf{M}$ with eigenvalues $\lambda_0, \dots \lambda_n$, the spectral radius is defined as $\rho(\mathbf{M}) = max\{|\lambda_0|, \dots, |\lambda_n|\}$.

### 3.1 Summing and Multiplying in Max-Plus

**Definition 2 (Max-Plus [7])** *Let $\mathbb{R}_{max} = \mathbb{R} \cup \{\varepsilon\}$ be the union of the set of real numbers $\mathbb{R}$ and $\varepsilon = -\infty$. Given $x, y \in \mathbb{R}_{max}$, we define the following two operations.*

$$\begin{aligned} x \oplus y &= \max(x,y) \\ x \otimes y &= x + y \end{aligned} \tag{4}$$

*The set $(\mathbb{R}_{max}, \oplus, \otimes)$ constitutes a semiring commonly known as the* Max-Plus Algebra.

Since $x \oplus \varepsilon = x$ and $x \otimes 0 = x$ for every $x \in \mathbb{R}_{max}$, $\varepsilon$ and 0 are the neutral elements of the $\oplus$ and $\otimes$ operations, respectively. The term $e$ is preferred for referring to the latter, as to avoid confusion with $0 \in \mathbb{R}$. Also note that the Max-Plus Algebra is an idempotent semiring in relation to $\oplus$, as $x \oplus x = x$ for any $x \in \mathbb{R}_{max}$.

Addition and multiplication in max-plus can be naturally extended to matrices by replacing the usual "$+$" and "$\cdot$" operators with $\oplus$ and $\otimes$. The $m \times n$ *zero matrix* is denoted by $\mathcal{E}$, such that $\mathcal{E}_{ij} = \varepsilon$ for all $i, j$. The $n \times n$ *identity matrix*, $\mathbf{E}_n$, takes the form:

$$[E_n]_{ij} = \begin{cases} e & \text{if} \quad i = j \\ \varepsilon & \text{if} \quad i \neq j \end{cases}$$

The power of a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}_{max}$ is inductively defined as $\mathbf{A}^{\otimes^0} \equiv \mathbf{E}_n$, and $\mathbf{A}^{\otimes^k} = \mathbf{A} \otimes \mathbf{A}^{\otimes^{k-1}}$ for $k > 0$. Eigenvalues and eigenvectors can be described within the Max-Plus Algebra as follows.

**Definition 3 (Eigenvalues and eigenvectors)** *Let $\mathbf{A} \in \mathbb{R}^{n \times n}_{max}$, and consider the scalars $\lambda \in \mathbb{R}_{max}$, and vectors $\mathbf{v} \neq (\varepsilon, \varepsilon, \dots, \varepsilon) \in \mathbb{R}^n_{max}$ satisfying:*

$$\mathbf{A} \otimes \mathbf{v} = \lambda \otimes \mathbf{v} \tag{5}$$

*$\lambda$ and $\mathbf{v}$ are referred to as the eigenvalues and eigenvectors of $\mathbf{A}$, respectively.*

Eigenvalues may be equal to $\varepsilon$. Both eigenvalues and eigenvectors may not be unique. The derivation of eigenvectors, i.e., the solutions to Equation (5), can be expressed more readily through a linear optimization problem: $\max_j(a_{ij} + v_j) = \lambda + v_i$, where $\mathbf{A} = (a_{ij})$ and $\mathbf{v} = (v_1, v_2, \dots, v_n)$.

Within the Max-Plus Algebra, the Perron-Frobenius Theorem takes on a more succinct form:

**Theorem 2 (Perron-Frobenius Theorem in Max-Plus [1])** *An* irreducible *matrix $\mathbf{A} \in \mathbb{R}^{n \times n}_{max}$ has a unique dominant eigenvalue such that:*

$$\lambda_0 = \bigoplus_{i=1}^{n} tr(A^i)^{1/i}$$

Within our adapted version of Eigentrust, we will seek to find eigenvalues for reducible matrices. Such matrices can be rewritten in max-plus in *normal form*.

**Definition 4 ( [17])** *Let $\mathbf{A} \in \mathbb{R}^{n \times n}_{max}$ be a reducible matrix, then its* normal form *is the upper triangular matrix:*

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} & \dots & \dots & \mathbf{A}_{1n} \\ \mathcal{E} & \mathbf{A}_{22} & \dots & \dots & \mathbf{A}_{2n} \\ \mathcal{E} & \mathcal{E} & \mathbf{A}_{33} & \dots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathcal{E} & \mathcal{E} & \mathcal{E} & \dots & \mathbf{A}_{nn} \end{pmatrix} \tag{6}$$

*where $\mathbf{A}_{nn}$ is irreducible and the matrices $\mathbf{A}_{ii}$ are either irreducible or equal to $\varepsilon$, for all $1 \leq i \leq n$. The remaining block matrices in Equation* (6) *are all different from $\mathcal{E}$.*

## 3.2 Measuring Trust in Max-Plus

We focus now on line 3 of the Eigentrust algorithm (Algorithm 1), situating it within max-plus. This operation updates the vector of trust values $\mathbf{t}$, effectively describing the evolution of a discrete system throughout $k$ iterations:

$$\mathbf{t}(k+1) = \mathbf{C}^T \mathbf{t}(k) \tag{7}$$

If $\mathbf{C}$ is irreducible, the above equation cannot be further expanded, and Eigentrust would yield a satisfactory result on account of the Perron-Frobenius Theorem. The reducible case, on the other hand, leads to a more elaborate recurrence relation hindering Eigentrust's performance. Given $\mathbf{D} \equiv \mathbf{C}^T$ and considering its normal form we can rewrite Equation (7) as [11]:

$$\mathbf{t}(k+1) = \mathbf{D}_{ii} \otimes \mathbf{t}_i(k) \oplus \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \otimes \mathbf{t}_j(k), \forall k \leq 0 \tag{8}$$

where $\mathbf{D}_{ii}$ are irreducible or equal to $\epsilon$, for $i \leq n$; and $\mathbf{D}_{ij} \neq \mathcal{E}$, for $j = i+1, i \in \{0, 1, \ldots, n-1\}$.

Provided $\mathbf{D}$ is reducible, there exist finite vectors $v_1, v_2, \ldots, v_n \in \mathbb{R}_{max}^{n \times 1}$ and scalars $\xi_1, \xi_2, \ldots, \xi_n \in \mathbb{R}$ producing a solution to Equation (8). More specifically, we have the following result [11].

**Theorem 3** *The solution to the discrete dynamic system in equation* (8) *is given by:*

$$t_i(k) = v_i \otimes \xi_i^{\otimes k} \tag{9}$$

*for all $k \geq 0$ and $i \in \{1, 2, \ldots, n\}$. The vectors $v_1, v_2, \ldots, v_n \in \mathbb{R}_{max}^{n \times 1}$ are finite, and the scalars $\xi_1, \xi_2, \ldots, \xi_n \in \mathbb{R}$ can be derived from the eigenvalues $\lambda_i$ of the irreducible block matrices $\mathbf{D}_{ij}$:*

$$\xi_i = \bigoplus_{j \in \mathcal{H}} \xi_j \oplus \lambda_j \tag{10}$$

*where $\mathcal{H} = \{j \in \{1, 2, \ldots, n\} : j > i, \mathbf{D}_{ij} \neq \mathcal{E}\}$.*

**Proof:** Set $\mathbf{r} = \mathbf{w} \otimes \xi$ as our initial vector of trust values for $\xi \in \mathbb{R}^n$ and a random vector $\mathbf{w} \in \mathbb{R}_{max}^n$, and also let $\mathbf{D}_{ii}$ be a matrix. $\mathbf{D}_{ii}$ is irreducible, hence Theorem 2 guarantees the existence of an eigenvalue $\lambda \in \mathbb{R}$ with eigenvector $\mathbf{v} \in \mathbb{R}^n$, chosen in accordance with the initial vector of trust scores: $\mathbf{v} \otimes \lambda \geq \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{r}_j$. Whenever $\lambda > \xi_j$ for $j \in \{i+1, \ldots, q\}$, $\mathbf{v}$ satisfies both of the following relations.

$$\mathbf{v} \otimes \lambda^{\otimes k+1} = \mathbf{D}_{ii} \otimes \mathbf{v} \otimes \lambda^{\otimes k}, \mathbf{v} \otimes \lambda^{\otimes k} \geq \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{w}_j \otimes \xi_j^{\otimes k}$$

This, in turn, implies the following equation.

$$\mathbf{v} \otimes \lambda^{\otimes k+1} = max\{\mathbf{D}_{ii} \otimes \mathbf{v} \otimes \lambda^{\otimes k}, \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{w}_j \otimes \xi_j^{\otimes k}\}$$

or, equivalently, if we set $\mathbf{t}(k) \equiv \mathbf{v} \otimes \lambda^{\otimes k}$ equation (8) is obtained. Note that this procedure is also applicable if the diagonal blocks are scalars, by making $\lambda = \varepsilon$ and $\mathbf{v} = \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{w}_j$.

When $\lambda \leq \xi_j$ for $j \in \{i+1, \ldots, q\}$ we could still obtain $\mathbf{v}$ as before given that $\bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{w}_j$ has at least one finite element. This choice, however, would also involve the following inequality.

$$\mathbf{v} \otimes \bigoplus_{j=i+1}^{q} \xi_j \geq \mathbf{D}_{ii} \otimes \mathbf{v} \otimes \lambda \oplus \bigoplus_{j=i+1}^{q} \mathbf{D}_{ij} \mathbf{w}_j \xi_j$$

Again, this leads to equation (8), if we let $\mathbf{t}(k) \equiv \mathbf{v} \otimes \mu^{\otimes k}$ and $\mu_i = \bigoplus_{j \in \mathcal{H}} \xi_j \oplus \lambda_j$ for all $k \geq 0, i \in \{1, 2, \ldots, n\}$, and $\mathcal{H}$ defined as in the statement of the theorem. □

**Algorithm 2** Power Method for regular irreducible matrices in Max-Plus

**Input:** $\mathbf{r}$: Arbitrary vector of trust values, $\mathbf{C}$: Trust Matrix.
**Output:** $\lambda$: Dominant eigenvalue of $\mathbf{C}$, $\mathbf{v}$ dominant eigenvector of $\mathbf{C}$.

**procedure** MAX_POWER
  $p \leftarrow 0$
  $\mathbf{v}_p \leftarrow \mathbf{r}$
  **repeat**
    $\mathbf{v}_{p+1} \leftarrow \mathbf{C}^T \mathbf{v}_p$
    $p \leftarrow p + 1$
  **until** There is some $q \in [0, \ldots, p]$ such that $\mathbf{v}_q = c \otimes \mathbf{v}_p$ for some $c \geq 0$.
  $\lambda \leftarrow \frac{c}{p-q}$
  $\mathbf{v} \leftarrow \bigoplus_{i=1}^{p-q} \left( \lambda^{\otimes(p-q-i)} \otimes \mathbf{v}_{q+i-1} \right)$
**return** $\lambda, \mathbf{v}$

Theorem 3 indicates that trust can be measured over reducible matrices as prescribed by the Eigentrust algorithm, invoking the spectral properties of the irreducible blocks of its normal form. Thus, we can recover the main graph-theoretic asymptotic traits of the system by looking into the connected components of the underlying network. Based on this, we introduce the MaxTrust Algorithm (Algorithm 3).

**Algorithm 3** Trust-Measuring Algorithm in Max-Plus

**Input:** $\mathbf{C}$: Regular Trust Matrix, $\mathbf{w}$: Vector of initial trust values, $T$: Terminal time.
**Output:** $\mathbf{t}$: trust ranking of agents at terminal time.

**procedure** $MaxTrust$
  $\mathbf{D} \leftarrow$ GET_NORMAL_FORM($\mathbf{C}$)
  $\lambda_n, \mathbf{v}_n \leftarrow$ MAX_POWER($\mathbf{D}_{nn}$)
  $\xi_n \leftarrow \lambda_n$
  $j \leftarrow n - 1$;
  **while** $j > 1$ **do**
    $\lambda_j \leftarrow$ MAX_POWER($\mathbf{D}_{jj}$).
    **if** $\lambda_j > \xi_{j+1}$ **then**
      $\xi_j \leftarrow \lambda_j$
      $\mathbf{v}_j \leftarrow \bigoplus_{k=1}^{n} \mathbf{D}_{jk} \otimes \mathbf{w}_k \otimes \lambda_j^{\otimes j-1}$
    **else**
      $\xi_j \leftarrow \lambda_{j+1}$
      $\mathbf{v}_j \leftarrow (\xi_j)^{-1} \otimes \bigoplus_{k=1}^{n} \mathbf{D}_{jk} \otimes \mathbf{w}_k \otimes \lambda_j^{\otimes j-1}$
    $j \leftarrow j - 1$
  **return** $\mathbf{t} \leftarrow \mathbf{v} \otimes \xi^{\otimes T}$

As with Eigentrust, the MaxTrust Algorithm uses a vector of initial trust values $\mathbf{w}$ which can be selected randomly. After converting the trust matrix to normal form, a straightforward max-plus adaptation of the Power Method — shown in Algorithm 2 — is applied to the last irreducible block of $\mathbf{D}$, obtaining its corresponding eigenvalues and eigenvectors. A similar operation is carried out for the rest of the diagonal blocks in $\mathbf{D}$, serving as the basis for the eigenvectors of the supradiagonal blocks (lines 10 and 13). Lines 8 and 15 mirror Equation (10) in the conventional algebra, and Equation (9) in max-plus, respectively.

## 4 Evaluation

We performed an empirical evaluation of MaxTrust in a simple simulated peer-to-peer network scenario. We begin by describing our experimental setup before detailing our results.

It is important to note that we did not use existing trust corpora such as epinions or the Ciao dataset. We justify this decision by noting that such datasets capture trust relationships between entities, but do not differentiate between lack of trust and lack of connectivity. Augmenting the datasets to capture such features would require making (unjustifiable) decisions resulting in an arbitrary dataset. Instead, we believe that the experimental setup described below is (somewhat) realistic, as it is based on ideas from peer-to-peer networking, and better illustrates the advantages and disadvantages of our approach and of Eigentrust.

### 4.1 Experimental Setup

Our experiments evaluated how trust propagates across a peer-to-peer network of routers, whose goal is to make routing decisions for data by deciding which of their neighbours such data should be transmitted to. Routers in the network interact with each other by exchanging connectivity information, consisting of trust measures regarding the network. This trust measure mirrors $t_{ij}$ — if router $i$ is broadcasting such a trust measure to its neighbours, index $j$ of the vector will contain either the level of trust $i$ ascribes to $j$; or $\varepsilon$ if $i$ has no knowledge of $j$ through direct or indirect experiences (in the case of MaxTrust), or 0 (in the case of Eigentrust).

Routers could either be trustworthy or malicious. The former broadcast trust measures correctly, while the latter transmit either a 0 or a value which begins at 0.5 and decays towards 0 as the router repeatedly interacts with others (to simulate the malicious router trying to undermine the network more actively). Initially, each trustworthy router began by imputing a random level of trust to all of its neighbours (with $\varepsilon$ in the remaining indices of its trust vector).

Experiments were run over 100 time steps. In each time step, all routers were given 10 opportunities to exchange information with their most trusted neighbour. After each such interaction, the trust they ascribed to their neighbour was either increased (if the neighbour was a trustworthy router) or decreased (if the neighbour was a malicious router) by 0.0001. There was also a 0.0025 chance of trust decreasing (effectively due to a mis-catagorisation of the neighbour).

After each time step, all routers computed new trust values for the system using Eigentrust or MaxTrust, and the process repeated. We considered three different router topologies: free trees (branching factor of 2); a toroidal network; and a random network of connections. Each network began by containing 4 routers with 8 links between them.

Each topology was evaluated under 3 different scenarios.

- **Scenario 1** Network was unchanged over all 100 trials.
- **Scenario 2** Every 5 interactions, between 2 and 6 new routers were added to the system. Half of the new routers in the system were set to be malicious.
- **Scenario 3** As in Scenario 2, between 2 and 6 new routers were added to the system every 5 interactions. None of the routers in the initial system were malicious, each new router had a 1/3 likelihood of being malicious.

When adding routers, network topology was preserved.

### 4.2 Results

We ran a total of 18 experiments (for each topology, scenario and trust algorithm combination), averaging 100 runs of each experimental condition (over 100 time steps) to obtain the results shown in Figure 1. The vertical axis in each plot compares averaged distance between the dominant eigenvector $v$ obtained by Eigentrust (dashed lines) or MaxTrust (solid lines), with the (actual) dominant eigenvector computed for the trust matrix obtained at the end of each experimental run ($v_{\lambda_0}$). This actual dominant eigenvector was obtained from the corresponding eigenvalue computed with the Newton method [19].

While all methods converge to the final dominant eigenvector as more information is exchanged between routers, it is clear that by differentiating between distrust (i.e., a trust value of 0) and no trust information (i.e., a trust value of $\varepsilon$), MaxTrust significantly outperforms Eigentrust across all topologies and scenarios, converging to the dominant eigenvector more rapidly. Within Scenario 1, trees gave the sparsest connectivity structure, meaning that Eigentrust struggled most in this case. As connectivity increased in Scenario 1, Eigentrust's performance improved, but still converged much more slowly than MaxTrust.

Within Scenario 2, Tori provided more connectivity than trees, leading to improved performance for Eigentrust compared to the latter case, and both outperform random networks. Given the results of the third scenario, we believe that this behaviour is caused by the disruption in the transmission of information due to the introduction of malicious routers. Indeed, when the majority of routers are not malicious, as in Scenario 3, the extra connectivity provided by random networks enhance their performance, while trees and tori induce no considerable changes.

Figure 2 illustrates the effectiveness of the algorithms in identifying potential threats and the dispersion of trust. Generally, Eigentrust underestimates the levels of trust agents accrue, reducing the scores corresponding to good peers computed in the presence of malicious agents interacting within heterogeneous environments. This is least apparent in random topologies, possibly due to the large number of connections in such a network.

The frequency of successful interactions is measured as $max(s_{ij}, 0)/\sum_k max(s_{ij} + f_{ij}, 0)$, where $s_{ij}$ designates the number of correct interactions between $i$ and $j$ — following the notation in Section 2 — while $f_{ij}$ denotes the number of unsuccessful interactions. The vertical axes of Figure 2 represent the differences between the trust values imputed to good peers by the different algorithms and the observed frequency of successful interactions. According to these results, Maxtrust is more stable than Eigentrust, but tends to inflate the level of trust accrued by good peers, thus penalising the rest of the agents.

The degree of reliability associated with the global rankings produced by Eigentrust and Maxtrust agrees with our underlying assumptions. On (potentially) reducible trust matrices (i.e., where the

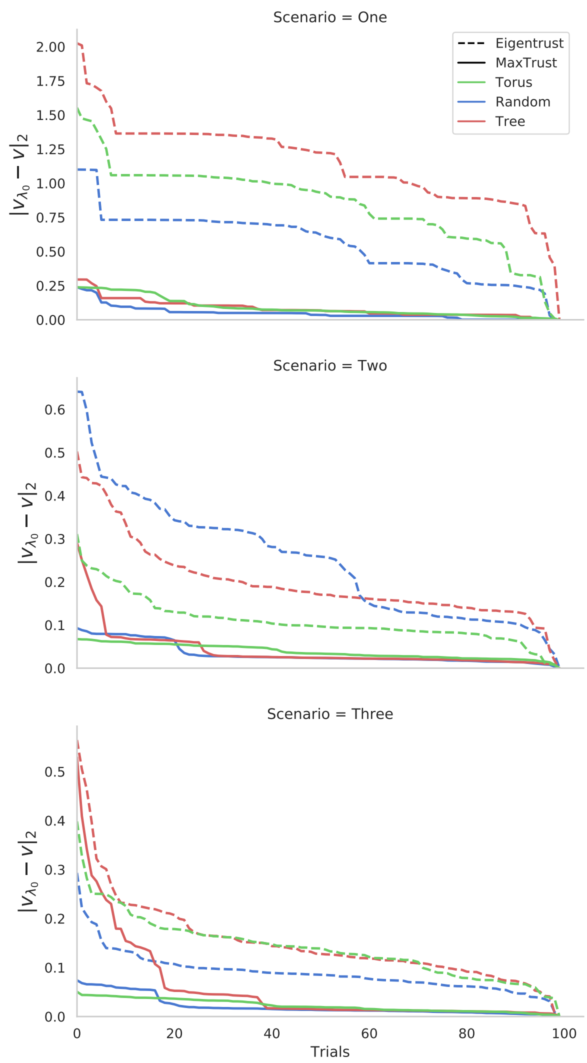


**Figure 1.** Relative Convergence to the Dominant Eigenvalue.

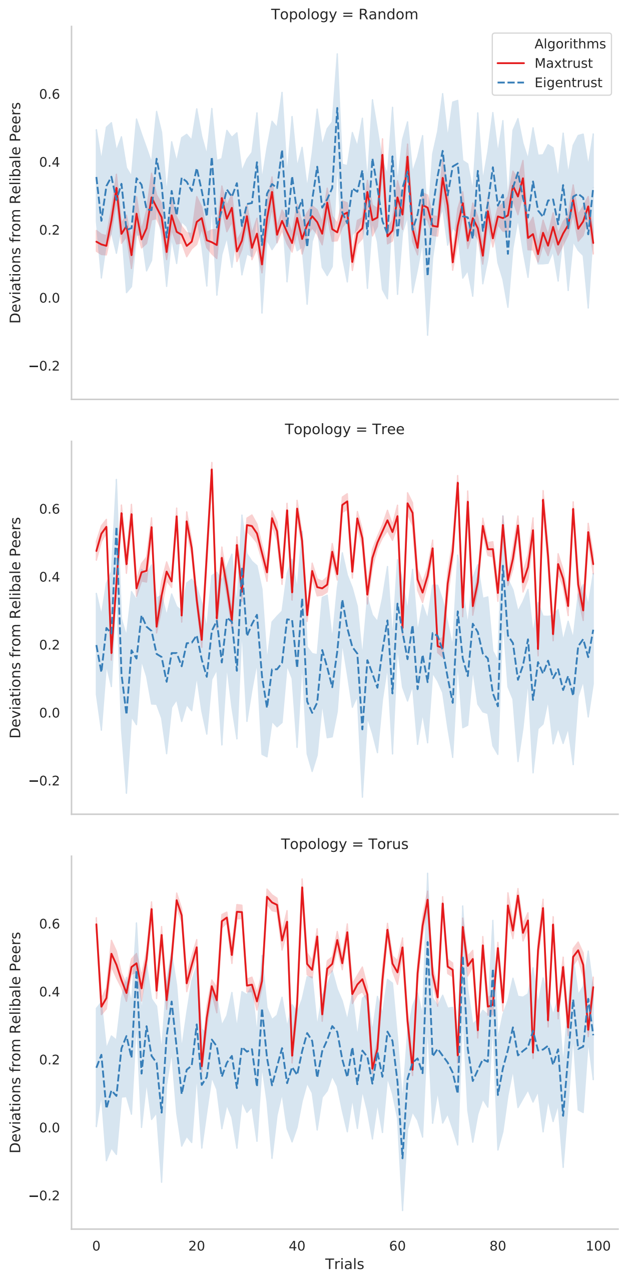


**Figure 2.** Deviations from Trustworthy Peers

network topology prevents the formation of (weakly) connected components), Eigentrust yields less informative rankings. The time series plots with error bands in Figure 2, and the statistical tests in Table 1, provide additional statistical evidence for this result.

Table 1 summarises the essential statistics detailing the posterior distributions of the distance $|v_{\lambda_0} - v|$ for MaxTrust (MT) and Eigentrust (ET). When treated as random variables the distributions of such deviations indicate how different the two results may be. Here, the mean and standard deviation are calculated over all 100 time steps (again over all 100 experiment runs). The low means and standard deviations of MaxTrust across all scenarios and topologies demonstrate its faster rate of convergence when compared to Eigentrust.

Ultimately, our results demonstrate the benefit of a trust and reputation system being able to differentiate between the lack of trust in an agent (i.e., a 0 trust value), and lack of information about trust in an agent (captured via $\varepsilon$). While Eigentrust conflates these two concepts, MaxTrust deals with them separately. This differentiation is particularly important in open dynamic multi-agent systems.

## 5 Discussion

Our experimental results indicate that a trust-measuring procedure framed in the Max-Plus algebra outperforms the standard Eigentrust approach with regards to convergence and stability of trust ratings across different scenarios and topologies. By leveraging max-plus, we are able to differentiate between no trust rating, and the inability of an agent to provide a trust rating, and we are able to extend the Eigentrust approach to the domain of reducible matrices.

Reducible trust matrices can be used to represent domains other than trust, and can capture many situations where the underlying network is not fully-connected.Although Eigentrust is often applied to

| Scenario | Network Structure | Mean | Standard Deviation | 95% HDP |
|---|---|---|---|---|
| One | Random MT | 0.035 | 0.001 | [0.034, 0.037] |
| | Random ET | 0.516 | 0.039 | [0.442, 0.583] |
| | Torus MT | 0.074 | 0.003 | [0.067, 0.079] |
| | Torus ET | 0.802 | 0.066 | [0.662, 0.925] |
| | Tree MT | 0.068 | 0.0015 | [0.065, 0.070 ] |
| | Tree ET | 1.106 | 0.059 | [1.012, 1.225] |
| Two | Random MT | 0.028 | 0.0003 | [0.027, 0.028] |
| | Random ET | 0.221 | 0.012 | [0.203, 0.246] |
| | Torus MT | 0.036 | 0.0001 | [0.035, 0.036] |
| | Torus ET | 0.097 | 0.0012 | [0.095, 0.099] |
| | Tree MT | 0.029 | 0.0002 | [0.029, 0.030] |
| | Tree ET | 0.176 | 0.0029 | [0.171, 0.180] |
| Three | Random MT | 0.016 | 0.0002 | [0.016, 0.017] |
| | Random ET | 0.080 | 0.0006 | [0.078, 0.080] |
| | Torus MT | 0.020 | 0.0001 | [0.020, 0.020] |
| | Torus ET | 0.125 | 0.0019 | [0.121, 0.129] |
| | Tree MT | 0.031 | 0.001 | [0.028, 0.033] |
| | Tree ET | 0.129 | 0.002 | [0.125, 0.135] |

**Table 1.** Summary of results for MaxTrust (MT) and Eigentrust (ET) over all scenarios and topologies.

P2P systems, many software and managerial applications functioning under trust protocols are modelled using non-fully connected graphs. This observation posits the question of whether Eigentrust is suitable for computing trust with reducible trust matrices.

In its original form, the Eigentrust algorithm employs the Power Method to compute the state of the (discrete) multi-agent system induced by a query about some peer's reputation. Max-plus provides an environment amenable to optimisation problems on discrete dynamic systems, enabling trust-measuring while circumventing the stringent assumptions behind the convergence of the Power Method.

Despite the numerous applications of the Max-Plus Algebra [7, 4, 6, 8], our proposal is among the first attempts to bring elements of tropical mathematics into the field of reputation and computational trust. Several works on Petri nets for the verification of cryptographic methods and the operation of autonomous vehicles [14] have already investigated techniques to circumvent the curse of dimensionality. Our work, albeit related, provides an algorithm for the more efficient functioning of trust protocols within an ample range of multi-agent systems.

Equally relevant to many applications is the robustness of our approach to dealing with malicious agents and the ability to update trust during each round within each trial. These results not only indicate MaxTrust's ability to fend off the effects that deceptively reliable agents have on the overall functioning of the system, but also corroborate the algorithm's accuracy when agents are allowed to revise past trust assignments before a trust-measuring cycle is over. Furthermore, these results point to the possibility of retrieving one of the system's invariant properties through MaxTrust.

## 6 Conclusions

Algorithms for trust measurement and computation are critical for the effective operation of open multi-agent systems. Due to its simplicity and effectiveness, Eigentrust is perhaps the most widely used trust-measuring algorithm. Building on an analysis of the situations where Eigentrust performs poorly (notably in cases where multiple distinct connected components exist), we introduced the MaxTrust algorithm. This algorithm shares the same basic intuitions used to create Eigentrust, but builds on the Max-Plus Algebra, and, in doing so, provides improved convergence to the *ex-post* or actual trust values when compared to Eigentrust.

We note that a vast number of popular trust and reputation systems have been proposed in the literature [9, 16]. Our focus in this work was — due to its popularity and ease of explanation — on the simplest version of (distributed) Eigentrust, and we did not consider other trust and reputation systems. Nor did we consider variants of Eigentrust such as those with pre-trusted peers [10]; or extensions of Eigentrust (e.g., [12]). In future work, we intend to examine whether the gains obtained using the Max-Plus algebra translate to Eigentrust's extensions. We believe, given the similarities between MaxTrust and Eigentrust, that the application of the Max-Plus Algebra to Eigentrust extensions will not be difficult, and given the improved results presented in the current work, we further believe that the application of Max-Plus to these extensions will also yield improved results. A comparison between MaxTrust (and variants) against other trust and reputation systems would therefore be a natural piece of future work.

We also intend to investigate the theoretical properties of MaxTrust in future work. Such properties include identifying guarantees on convergence rates, and the effects of different attacks against the algorithm. Given the shared intuitions between MaxTrust and Eigentrust, we believe that many results will carry through, but stress that a theoretical and empirical evaluation of MaxTrust under different scenarios is critical if the performance improvements it seems to hold, are to be realised in practical applications.

## REFERENCES

[1] Marianne Akian, Ravindra Bapat, and Stéphane Gaubert, 'Asymptotics of the perron eigenvalue and eigenvector using max-algebra', *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, **327**(11), 927–932, (1998).

[2] Muhammad Ajmal Azad, Samiran Bag, and Feng Hao, 'M2m-rep: Reputation of machines in the internet of things', in *Proceedings of the 12th international conference on availability, reliability and security*, p. 28. ACM, (2017).

[3] Abraham Berman and Robert J Plemmons, *Nonnegative matrices in the mathematical sciences*, volume 9, Siam, 1994.

[4] J-L Bouquard, Christophe Lenté, and J-C Billaut, 'Application of an optimization problem in max-plus algebra to scheduling problems', *Discrete Applied Mathematics*, **154**(15), 2064–2079, (2006).

[5] Moody T Chu and J Loren Watterson, 'On a multivariate eigenvalue problem, part i: Algebraic theory and a power method', *SIAM Journal on scientific computing*, **14**(5), 1089–1106, (1993).

[6] Stéphane Gaubert and Ricardo D Katz, 'The minkowski theorem for max-plus convex sets', *Linear Algebra and its Applications*, **421**(2-3), 356–369, (2007).

[7] Stephane Gaubert and Max Plus, 'Methods and applications of (max,+) linear algebra', in *Annual symposium on theoretical aspects of computer science*, pp. 261–282. Springer, (1997).

[8] Bernd Heidergott, Geert Jan Olsder, and Jacob Van Der Woude, *Max Plus at work: modeling and analysis of synchronized systems: a course on Max-Plus algebra and its applications*, volume 48, Princeton University Press, 2014.

[9] Audun Jøsang, Roslan Ismail, and Colin Boyd, 'A survey of trust and reputation systems for online service provision', *Decision support systems*, **43**(2), 618–644, (2007).

[10] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina, 'The eigentrust algorithm for reputation management in P2P networks', in *Proceedings of the 12th international conference on World Wide Web*, pp. 640–651. ACM, (2003).

[11] Zvi Retchkiman Konigsberg, 'A generalized eigenmode algorithm for reducible regular matrices over the max-plus algebra', in *2009 Chinese Control and Decision Conference*, pp. 5598–5603. IEEE, (2009).

[12] Heba A Kurdi, 'Honestpeer: An enhanced eigentrust algorithm for reputation management in p2p systems', *Journal of King Saud University-Computer and Information Sciences*, **27**(3), 315–322, (2015).

[13] Sergio Marti and Hector Garcia-Molina, 'Taxonomy of trust: Categorizing p2p reputation systems', *Computer Networks*, **50**(4), 472–484, (2006).

[14] Ahmed Nait-Sidi-Moh and Maxime Wack, 'Modelling of process of electronic signature with petri nets and (max, plus) algebra', in *Inter-*

*national Conference on Computational Science and Its Applications*, pp. 792–801. Springer, (2005).

[15] S Unnikrishna Pillai, Torsten Suel, and Seunghun Cha, 'The perron-frobenius theorem: some of its applications', *IEEE Signal Processing Magazine*, **22**(2), 62–75, (2005).

[16] Yefeng Ruan and Arjan Durresi, 'A survey of trust management systems for online social communities–trust modeling, trust inference and attacks', *Knowledge-Based Systems*, **106**, 150–163, (2016).

[17] Eugene Seneta, *Non-negative matrices and Markov chains*, Springer Science & Business Media, 2006.

[18] George Theodorakopoulos and John S Baras, 'On trust models and trust evaluation metrics for ad hoc networks', *IEEE Journal on selected areas in Communications*, **24**(2), 318–328, (2006).

[19] James Hardy Wilkinson, *The algebraic eigenvalue problem*, volume 662, Oxford Clarendon, 1965.