Lecture 6 Kleene Algebra

Consider a binary relation on an n element set represented by an $n \times n$ Boolean matrix E. Recall from the last lecture that we can compute the reflexive transitive closure of E by divide-and-conquer as follows: partition E into four submatrices A, B, C, D of size roughly $\frac{n}{2} \times \frac{n}{2}$ such that A and D are square:

$$E = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right].$$

By induction, construct the matrices D^*, $F = A + BD^*C$, and F^*, then take

$$E^* = \left[\begin{array}{c|c} F^* & F^*BD^* \\ \hline D^*CF^* & D^* + D^*CF^*BD^* \end{array}\right]. \tag{1}$$

We will prove that the matrix E^* as defined in (1) is indeed the reflexive transitive closure of E, but the proof will be carried out in a more abstract setting which will allow us to use the same construction in other applications. For example, we will be able to compute the lengths of the shortest paths between all pairs of points in a weighted directed graph using the same general algorithm, but with a different interpretation of the basic operations.

How did we come up with the expressions in (1)? This is best motivated by considering a simple finite-state automaton over the alphabet $\Sigma = \{a, b, c, d\}$ with states s, t and transitions $s \xrightarrow{a} s$, $s \xrightarrow{b} t$, $t \xrightarrow{c} s$, $t \xrightarrow{d} t$:

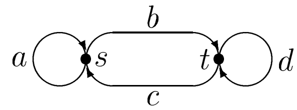

For each pair of states u, v, consider the set of input strings in Σ^* taking state u to state v in this automaton. Each such set is a regular subset of Σ^* and is represented by a regular expression corresponding to the expressions appearing in (1):

$$\begin{aligned} s \to s &: f^* \\ s \to t &: f^*bd^* \\ t \to s &: d^*cf^* \\ t \to t &: d^* + d^*cf^*bd^* \ , \end{aligned}$$

where $f = a + bd^*c$. (See [3, §9.1, pp. 318-319] for more information on finite automata and regular expressions.)

6.1 Definition of Kleene Algebras

The appropriate level of abstraction we are seeking is *Kleene algebra*. This concept goes back to Kleene [61], but received significant impetus from the work of Conway [21]. The definition here is from [63].

Definition 6.1 A *(*-continuous) Kleene algebra* is any structure of the form

$$\mathcal{K} = (S, +, \cdot, {}^*, 0, 1)$$

where S is a set of elements, $+$ and $\cdot$ are binary operations $S \times S \to S$, * is a unary operation $S \to S$, and 0 and 1 are distinguished elements of S, satisfying the axioms

$$\begin{aligned} a + (b + c) &= (a + b) + c && \text{(+ is associative)} && (2) \\ a + b &= b + a && \text{(+ is commutative)} && (3) \\ a + a &= a && \text{(+ is idempotent)} && (4) \\ a + 0 &= 0 + a = a && \text{(0 is an a identity for +)} && (5) \\ a \cdot (b \cdot c) &= (a \cdot b) \cdot c && \text{($\cdot$ is associative)} && (6) \\ a \cdot 1 &= 1 \cdot a = a && \text{(1 is an a identity for $\cdot$)} && (7) \\ 0 \cdot a &= a \cdot 0 = 0 && \text{(0 is an annihilator for $\cdot$)} && (8) \\ a \cdot (b + c) &= a \cdot b + a \cdot c && \text{($\cdot$ distributes over +)} && (9) \\ (b + c) \cdot a &= b \cdot a + c \cdot a && && (10) \end{aligned}$$

plus the following axiom to deal with the * operator, which will require further explanation:

$$ab^*c = \sup_{n \geq 0} ab^nc \qquad (11)$$

where

$$\begin{aligned} b^0 &= 1 \\ b^{n+1} &= b \cdot b^n . \end{aligned}$$

□

Axioms (2–5) say that the structure $(S,\ +,\ 0)$ is an *idempotent commutative monoid*. Axioms (6–7) say that $(S,\ \cdot,\ 1)$ is a *monoid*. Axioms (8–10) describe how these two monoid structures interact. Altogether, Axioms (2–10) say that $\mathcal{K}$ is an *idempotent semiring*.

The axiom (11) asserts the existence of the *supremum* or *least upper bound* of a certain set with respect to a certain partial order. In any idempotent semiring, there is a natural partial order defined by

$$a \leq b \quad \leftrightarrow \quad a + b = b . \tag{12}$$

It follows easily from the axioms (2–5) that $\leq$ is indeed a partial order; *i.e.*, it is

- reflexive: $a \leq a$
- antisymmetric: if $a \leq b$ and $b \leq a$ then $a = b$; and
- transitive: if $a \leq b$ and $b \leq c$ then $a \leq c$.

If A is a set of elements of a partially ordered set, the element y is said to be the *supremum* or *least upper bound* of the set A (notation: $y = \sup A$) if

- y is an upper bound for A; *i.e.*, $x \leq y$ for all $x \in A$;
- y is the least such upper bound; *i.e.*, for any other upper bound z for A, $y \leq z$.

The supremum of any pair of elements x, y exists and is equal to $x + y$. It follows that the supremum of any finite set $\{a_1, \ldots, a_n\}$ exists and is equal to $a_1 + \cdots + a_n$ (parentheses are not necessary because $+$ is associative). In general, the supremum of an infinite set need not exist, but if it exists then it is unique. The axiom (11) asserts that the supremum of the set $\{ab^nc \mid n \geq 0\}$ exists and is equal to ab^*c.

The postulate (11) captures axiomatically the behavior of reflexive transitive closure of a binary relation. It also captures the behavior of the Kleene * operator of formal language theory. In addition, there are many nonstandard examples of Kleene algebras that are useful in various contexts. We will give several examples below.

Instead of Kleene algebras, many authors (such as [3, 78]) use so-called *closed semirings*. These structures are strongly related to Kleene algebras, but are defined in terms of a countable summation operator $\sum$ instead of a supremum. In closed semirings, the * operator is not a primitive operator but is defined in terms of $\sum$ by

$$b^* \;=\; \sum_{n\geq 0} b^n \;.$$

The countable summation operator $\sum$, which sums a countably infinite sequence of elements, is postulated not to depend on the order of the elements in the sequence or their multiplicity, and thus is essentially a supremum. The operator $\sum$ is also postulated to satisfy an infinite distributivity property that we get for free for all suprema of interest by stating the axiom as we did in (11).

The main drawback with closed semirings is that the suprema of all countable sets are required to exist, which is too many. Although every closed semiring is a Kleene algebra, there are definitely Kleene algebras that are not closed semirings. The most important example of such a Kleene algebra is the family $\mathbf{Reg}_{\boldsymbol{\Sigma}}$ of regular subsets of Σ^*, where Σ is a finite alphabet (Example 6.2 below). This example is important because it is the free Kleene algebra freely generated by Σ, which essentially says that an equation between regular expressions over Σ holds in all Kleene algebras if and only if it holds in $\mathbf{Reg}_{\boldsymbol{\Sigma}}$. We will find this fact very useful in reducing arguments about Kleene algebras in general to arguments about regular subsets of Σ^*.

Kleene algebras were studied extensively in the monograph of Conway [21]. It is possible to axiomatize the equational theory of Kleene algebras in a purely finitary way [65]. The precise relationship between Kleene algebras and closed semirings is drawn in [64].

6.2 Examples of Kleene Algebras

Kleene algebras abound in computer science. Here are some examples.

Example 6.2 Let Σ be a finite alphabet and let $\mathbf{Reg}_{\boldsymbol{\Sigma}}$ denote the family of regular sets over Σ with the following operations:

$$\begin{aligned}
A + B &= A \cup B \\
A \cdot B &= \{xy \mid x \in A,\ y \in B\} \\
A^* &= \{x_1 x_2 \cdots x_n \mid n \geq 0 \text{ and } x_i \in A,\ 1 \leq i \leq n\} \\
&= \bigcup_{n\geq 0} A^n \\
0 &= \emptyset \\
1 &= \{\epsilon\}
\end{aligned}$$

where A^n is defined inductively by

$$\begin{aligned}
A^0 &= \{\epsilon\} \\
A^{n+1} &= A \cdot A^n
\end{aligned}$$

and ϵ is the empty string. Under these operations, $\mathbf{Reg}_\Sigma$ is a Kleene algebra, and a very important one indeed: it is the *free Kleene algebra on free generators* Σ, which essentially means that any equation $\alpha = \beta$ between regular expressions holds in all Kleene algebras iff it holds in $\mathbf{Reg}_\Sigma$. □

Example 6.3 Let X be a set and let $\mathcal{R}$ be any family of binary relations on X closed under the following operations:

$$\begin{aligned}
R + S &= R \cup S \\
R \cdot S &= \{(x,z) \mid \exists y \in X\ (x,y) \in R \text{ and } (y,z) \in S\} \\
R^* &= \bigcup_{n \geq 0} R^n \\
&= \text{the reflexive transitive closure of } R \\
0 &= \emptyset \\
1 &= \{(x,x) \mid x \in X\}\ .
\end{aligned}$$

where R^n is defined inductively by

$$\begin{aligned}
R^0 &= \{(x,x) \mid x \in X\} \\
R^{n+1} &= R \cdot R^n\ .
\end{aligned}$$

Under these operations, $\mathcal{R}$ forms a Kleene algebra. Kleene algebras of binary relations are used to model programs in Dynamic Logic and other logics of programs. □

Example 6.4 The set $\{0,1\}$ of Boolean truth values forms a Kleene algebra under the operations

$$\begin{aligned}
a + b &= a \vee b \\
a \cdot b &= a \wedge b \\
a^* &= 1
\end{aligned}$$

and 0 and 1 as named. This is the smallest nontrivial Kleene algebra. □

Example 6.5 The family of $n \times n$ Boolean matrices forms a Kleene algebra under the operations

$$\begin{aligned}
A + B &= A \vee B \\
A \cdot B &= \text{Boolean matrix product} \\
A^* &= \text{reflexive transitive closure} \\
0 &= \text{the zero matrix} \\
1 &= \text{the identity matrix.}
\end{aligned}$$

This is essentially the same as Example 6.3 above for an n-element set X. □

Example 6.6 The following rather bizarre example will be useful in computing all-pairs shortest paths in a weighted graph. We will have to be a little more explicit with notation than usual to avoid confusion.

Let $\mathcal{R}_+$ denote the family of nonnegative real numbers, and let ∞ be a new element. Let $+_\mathcal{R}$ denote ordinary addition in $\mathcal{R}_+ \cup \{\infty\}$, where we define

$$a +_\mathcal{R} \infty \quad = \quad \infty +_\mathcal{R} a \quad = \quad \infty$$

for all $a \in \mathcal{R}_+ \cup \{\infty\}$. Let $\leq_\mathcal{R}$ denote the natural order in $\mathcal{R}_+ \cup \{\infty\}$, with $a \leq_\mathcal{R} \infty$ for all $a \in \mathcal{R}_+$. Let $\min_\mathcal{R}\{a, b\}$ denote the minimum of a and b with respect to this order. Let $0_\mathcal{R}$ denote the real number 0.

Define the Kleene algebra operations $+_\mathcal{K}$, $\cdot_\mathcal{K}$, ${}^{*_\mathcal{K}}$, $0_\mathcal{K}$, and $1_\mathcal{K}$ on $\mathcal{R}_+ \cup \{\infty\}$ as follows:

$$\begin{aligned}
a +_\mathcal{K} b &= \min_\mathcal{R}\{a, b\} \\
&= \begin{cases} a\ , & \text{if } a \leq_\mathcal{R} b \\ b\ , & \text{otherwise} \end{cases} \\
a \cdot_\mathcal{K} b &= a +_\mathcal{R} b \\
a^{*_\mathcal{K}} &= 0_\mathcal{R} \\
0_\mathcal{K} &= \infty \\
1_\mathcal{K} &= 0_\mathcal{R}\ .
\end{aligned}$$

If this appears confusing, don't worry, it really is. To make sense of it, just keep in mind that the symbols on the left hand side of these equations refer to the Kleene algebra operations being defined, whereas those on the right hand side refer to the natural operations of $\mathcal{R}_+ \cup \{\infty\}$. Note that the zero element of the Kleene algebra is ∞, the identity for $\min_\mathcal{R}$, and the multiplicative identity 1 of the Kleene algebra is the real number 0, the identity for addition in $\mathcal{R}_+ \cup \{\infty\}$ (which is multiplication in the Kleene algebra). The worst part is that the natural partial order $\leq_\mathcal{K}$ in the Kleene algebra as defined by (12) is the reverse of $\leq_\mathcal{R}$; that is, $a \leq_\mathcal{K} b$ iff $b \leq_\mathcal{R} a$.

This algebra is often called the min,+ Kleene algebra. □

Lecture 7 More on Kleene Algebra

In this lecture we will see how Kleene algebra can be used in a variety of situations involving *-like operations. The key result that allows these applications is that the $n \times n$ matrices over a Kleene algebra again form a Kleene algebra. Along the way we will establish a central lemma that establishes the importance of the regular sets $\mathbf{Reg}_{\mathbf{\Sigma}}$ over the finite alphabet Σ in reasoning about Kleene algebras in general.

Let

$$\mathcal{K} \;=\; (K,\ +_{\mathcal{K}},\ \cdot_{\mathcal{K}},\ ^{*_{\mathcal{K}}},\ 0_{\mathcal{K}},\ 1_{\mathcal{K}})$$

be a Kleene algebra. Let Σ a set and let $\mathbf{RExp}_{\mathbf{\Sigma}}$ denote the family of regular expressions over Σ (see [3, §9.1, pp. 318-319]). An *interpretation* over $\mathcal{K}$ is a map

$$I : \Sigma \;\rightarrow\; \mathcal{K}$$

assigning an element of $\mathcal{K}$ to each element of Σ. An interpretation can be extended to domain $\mathbf{RExp}_{\mathbf{\Sigma}}$ inductively as follows:

$$\begin{aligned}
I(0) &= 0_{\mathcal{K}} \\
I(1) &= 1_{\mathcal{K}} \\
I(\alpha + \beta) &= I(\alpha) +_{\mathcal{K}} I(\beta) \\
I(\alpha \cdot \beta) &= I(\alpha) \cdot_{\mathcal{K}} I(\beta) \\
I(\alpha^*) &= I(\alpha)^{*_{\mathcal{K}}} \; .
\end{aligned}$$

At the risk of confusing the operator symbols in regular expressions and the corresponding operations in $\mathcal{K}$, we henceforth drop the subscripts $\mathcal{K}$.

For example, the interpretation

$$\begin{aligned} R : \Sigma &\to \mathbf{Reg}_\Sigma \\ a &\mapsto \{a\} \end{aligned}$$

over $\mathbf{Reg}_\Sigma$ extends to the map

$$R : \mathbf{RExp}_\Sigma \to \mathbf{Reg}_\Sigma$$

in which $R(\alpha)$ is the regular set denoted by the regular expression α in the usual sense. The interpretation R is called the *standard interpretation* over $\mathbf{Reg}_\Sigma$.

The following lemma generalizes (11).

Lemma 7.1 *Let $R : \Sigma \to \mathbf{Reg}_\Sigma$ be the standard interpretation over $\mathbf{Reg}_\Sigma$ and let $I : \Sigma \to \mathcal{K}$ be any interpretation over any Kleene algebra $\mathcal{K}$. For any regular expression α over Σ,*

$$I(\alpha) = \sup_{x \in R(\alpha)} I(x) . \tag{13}$$

Note that since $R(\alpha)$ is a regular set of strings over the alphabet Σ, the x in (13) denotes a string. Strings over Σ are themselves regular expressions over Σ, so the expression $I(x)$ makes sense. The equation (13) states that the supremum of the possibly infinite set

$$\{I(x) \mid x \in R(\alpha)\} \subseteq \mathcal{K}$$

exists and is equal to $I(\alpha)$. We leave the proof of Lemma 7.1 as an exercise (Homework 3, Exercise 2).

It follows that for any pair α, β of regular expressions over Σ, the equation $\alpha = \beta$ is a logical consequence of the axioms of Kleene algebra, *i.e.* it holds under all interpretations over all Kleene algebras, if and only if it holds under the standard interpretation R over $\mathbf{Reg}_\Sigma$. A fancy way of saying this is that $\mathbf{Reg}_\Sigma$ is the *free Kleene algebra on free generators* Σ.

Theorem 7.2 *Let α and β be regular expressions over Σ and let R be the standard interpretation over $\mathbf{Reg}_\Sigma$. Then*

$$I(\alpha) = I(\beta)$$

for all interpretations I over Kleene algebras if and only if

$$R(\alpha) = R(\beta) .$$

Proof. ($\rightarrow$) This follows immediately from the fact that $\mathbf{Reg}_\Sigma$ is a Kleene algebra and R is an interpretation over $\mathbf{Reg}_\Sigma$.

($\leftarrow$) Suppose $R(\alpha) = R(\beta)$. Then

$$\begin{aligned} I(\alpha) &= \sup_{x \in R(\alpha)} I(x) \quad \text{by Lemma 7.1} \\ &= \sup_{x \in R(\beta)} I(x) \quad \text{by the assumption } R(\alpha) = R(\beta) \\ &= I(\beta)\ , \quad \text{again by Lemma 7.1.} \end{aligned}$$

□

7.1 Matrix Kleene Algebras

The collection $M(n, \mathcal{K})$ of $n \times n$ matrices with elements in a Kleene algebra $\mathcal{K}$ again forms a Kleene algebra, provided the Kleene algebra operators on $M(n, \mathcal{K})$ are defined appropriately. We always define $+$ as ordinary matrix addition, $\cdot$ as ordinary matrix multiplication, 0 as the zero matrix, 1 as the identity matrix, and * recursively by equation (1) of the previous lecture. We must show that all the axioms of Kleene algebra are satisfied by $M(n, \mathcal{K})$ under these definitions. For example, in $M(2, \mathcal{K})$ the identity elements for $+$ and $\cdot$ are

$$\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

respectively, and the operations $+$, $\cdot$, and * are given by

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} + \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} a+e & b+f \\ c+g & d+h \end{bmatrix}$$

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \cdot \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} ae+bg & af+bh \\ ce+dg & cf+dh \end{bmatrix}$$

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^* = \begin{bmatrix} f^* & f^*bd^* \\ d^*cf^* & d^* + d^*cf^*bd^* \end{bmatrix}$$

where $f = a + bd^*c$. Note that $A \leq B$ in the natural order on $M(n, \mathcal{K})$ defined by (12) if and only if $A_{ij} \leq B_{ij}$ for all $1 \leq i, j \leq n$.

Most of the Kleene algebra axioms are routine to verify for the structure $M(n, \mathcal{K})$. Let us verify (11) explicitly, assuming all the other axioms have been verified. First we will show that it is true for a particular choice of matrices over a particular Kleene algebra of regular sets, using a combinatorial argument; next we will use Theorem 7.2 to extend the result to all Kleene algebras.

Let $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ be $n \times n$ symbolic matrices with ij^{th} elements $\mathbf{a}_{ij}$, $\mathbf{b}_{ij}$, and $\mathbf{c}_{ij}$, respectively, where the $\mathbf{a}_{ij}$, $\mathbf{b}_{ij}$, $\mathbf{c}_{ij}$ are distinct letters. Let

$$\Sigma \quad = \quad \{\mathbf{a}_{ij}, \mathbf{b}_{ij}, \mathbf{c}_{ij} \mid 1 \leq i, j \leq n\} \ .$$

Build an automaton $M_{\mathbf{B}}$ with n states and transition from state i to state j labeled with the letter $\mathbf{b}_{ij}$. The ij^{th} element of $\mathbf{B}^k$, the symbolic k^{th} power of $\mathbf{B}$, is a regular expression representing the set of strings of length k over Σ taking state i to state j in $M_{\mathbf{B}}$. Moreover, the ij^{th} element of $\mathbf{B}^*$ represents the set of all strings (of any length) taking state i to state j in $M_{\mathbf{B}}$. This follows from a purely combinatorial inductive argument, using the definition of $\mathbf{B}^*$ as given in (1); the partition in (1) corresponds to a partition of the states of $M_{\mathbf{B}}$ into two disjoint sets. We thus have

$$R((\mathbf{B}^*)_{ij}) \quad = \quad \bigcup_{k \geq 0} R((\mathbf{B}^k)_{ij})$$

where R is the standard interpretation.

Let $M_{\mathbf{A}}$ and $M_{\mathbf{C}}$ consist of n states each. Connect state i of $M_{\mathbf{A}}$ with state j of $M_{\mathbf{B}}$ and label the transition $\mathbf{a}_{ij}$. Similarly, connect state i of $M_{\mathbf{B}}$ with state j of $M_{\mathbf{C}}$ and label the transition $\mathbf{c}_{ij}$. Call this new automaton M. Then the regular set over Σ denoted by the ij^{th} element of $\mathbf{A}\mathbf{B}^k\mathbf{C}$ is the set of strings of length $k+2$ taking state i of $M_{\mathbf{A}}$ to state j of $M_{\mathbf{C}}$ in M, and the regular set denoted by the ij^{th} element of $\mathbf{A}\mathbf{B}^*\mathbf{C}$ is the set of all strings (of any length) taking state i of $M_{\mathbf{A}}$ to state j of $M_{\mathbf{C}}$ in M. Therefore

$$R((\mathbf{A}\mathbf{B}^*\mathbf{C})_{ij}) \quad = \quad \bigcup_{k \geq 0} R((\mathbf{A}\mathbf{B}^k\mathbf{C})_{ij}) \ .$$

Now let A, B, C be arbitrary matrices over an arbitrary Kleene algebra $\mathcal{K}$. Let a_{ij}, b_{ij}, c_{ij} denote the ij^{th} elements of A, B, and C, respectively. Let I be the interpretation

$$\begin{aligned} I(\mathbf{a}_{ij}) &= a_{ij} \\ I(\mathbf{b}_{ij}) &= b_{ij} \\ I(\mathbf{c}_{ij}) &= c_{ij} \ . \end{aligned}$$

Then

$$\begin{aligned} (AB^*C)_{ij} &= I((\mathbf{A}\mathbf{B}^*\mathbf{C})_{ij}) \\ &= \sup\{I(x) \mid x \in R((\mathbf{A}\mathbf{B}^*\mathbf{C})_{ij})\} \quad \text{by Lemma 7.1} \\ &= \sup\{I(x) \mid x \in \bigcup_{k \geq 0} R((\mathbf{A}\mathbf{B}^k\mathbf{C})_{ij})\} \\ &= \sup_{k \geq 0} \sup\{I(x) \mid x \in R((\mathbf{A}\mathbf{B}^k\mathbf{C})_{ij})\} \\ &= \sup_{k \geq 0} I((\mathbf{A}\mathbf{B}^k\mathbf{C})_{ij}) \\ &= \sup_{k \geq 0} (AB^kC)_{ij} \ , \end{aligned}$$

therefore

$$AB^*C \quad = \quad \sup_{k \geq 0} AB^kC \ .$$

This establishes (11) for $M(n, \mathcal{K})$.

7.2 Applications

The obvious divide-and-conquer algorithm for computing E^* given by (1) yields the recurrence

$$T(n) \quad = \quad T(\frac{n}{2}) + O(M(n)) \ ,$$

where $M(n)$ is the number of basic operations needed to add or multiply two $n \times n$ matrices over $\mathcal{K}$. Under the quite reasonable assumption that $M(2n) \geq 4M(n)$, this recurrence has solution

$$T(n) \quad = \quad O(M(n)) \ .$$

For most applications, $M(n) = O(n^3)$. Better bounds can be obtained using Strassen's algorithm or other fast matrix multiplication algorithms when $\mathcal{K}$ is a ring.

Reflexive Transitive Closure

Using matrix Kleene algebras, we can prove the correctness of the algorithm for reflexive transitive closure presented in the last lecture. Let $\mathcal{B}$ denote the two-element Kleene algebra described in Example 6.4 above. Let E denote the adjacency matrix of a directed graph G with n vertices. Then $E \in M(n, \mathcal{B})$, and the ij^{th} element of E^k is 1 if and only if there exists a directed path in G from vertex i to vertex j of length exactly k. By the result of the last section, we know that

$$E^* \quad = \quad \sup_{k \geq 0} E^k \ ,$$

so the ij^{th} element of E^* is 1 iff there exists a path of *some* length from i to j. This is the reflexive transitive closure.

All-Pairs Shortest Paths

Here we use the same algorithm, but a different underlying Kleene algebra, namely the min,+ algebra of Example 6.6 above. Supremum in this order is infimum in the usual order on $\mathcal{R}_+ \cup \{\infty\}$. Thus a^* is the real number 0 for all a.

We apply this to the all-pairs shortest path problem. Let E be a matrix over the min,+ algebra containing the edge lengths of a weighted directed graph G. If (i, j) is not an edge in G, set $E_{ij} = \infty$. In E^2, the ij^{th} element will be the minimum over all vertices k of the sum of the lengths of (i, k) and (k, j). That is, it will contain the length of a shortest path of two edges from i to j. It follows by induction that the ij^{th} element of E^k is the length of a shortest path of k edges from i to j. Since

$$E^* \quad = \quad \sup_{k \geq 0} E^k$$

and supremum in the Kleene algebra is infimum in the natural order, E^* gives the length of a shortest path of any number of edges.