Thatched Partial Orders

Bill Mitchell
w.mitchell@surrey.ac.uk
Department of Computing
University of Surrey
Guildford, Surrey GU2 7XH, UK

May 28, 2009

1 Geometric Thatching

A thatch is a type of partial ordering which is most naturally described as a geometric construction. A thatching consists of a set of vertical and horizontal lines where the vertical lines can have any finite number of nodes attached to them. A thatch is constructed inductively as follows.

The initial thatch is any number of parallel vertical lines. Any open interval in a vertical line can be replaced by two parallel horizontal lines, placed so that the upper line occurs at the beginning of the interval, and the lower line occurs at the end of the interval. The new lines must not intersect any other line already in the thatch. Next these horizontal lines must be connected by a finite non-zero number of new parallel vertical lines which start at the upper horizontal line and terminate at the lower horizontal line. They must not intersect any other lines in the thatch. The two horizontal lines are defined to be dual to each other.

This process may iterate any finite number of times on any of the vertical lines in the current thatch. Finally attach finitely many nodes to distinct points on the vertical lines.

For example if we start with a single vertical line we may choose to insert four new vertical lines as shown here.

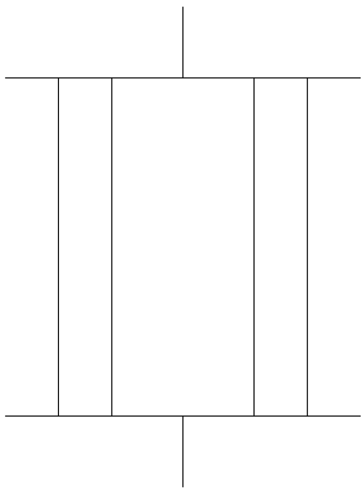

We may then choose to add a further two vertical lines as shown here

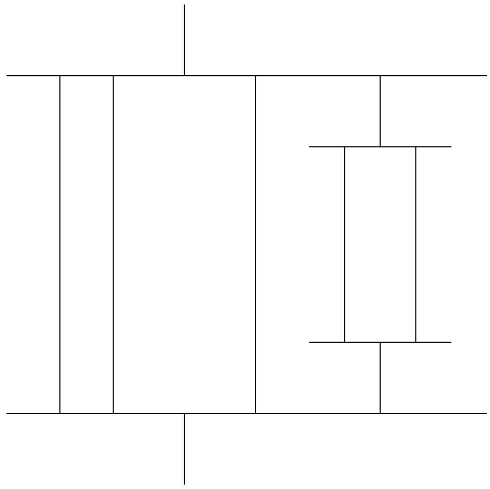

Then we may choose to attach nodes as described here:

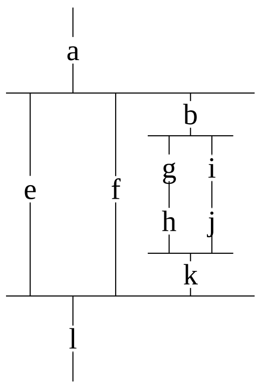

Thatch 1

A thatch describes a partial order of the nodes attached to the vertical lines. Define $x < y$ for nodes x and y if there is a continuous path from x to y along the vertical and horizontal lines contained in the thatch which never travels upwards when traversed from x to y. That is the direction of travel must always be downwards or sideways. Such a path is called a non-upwards path. Any partial order that can be defined by a thatch is called a thatch ordering.

The region defined by a dual pair of horizontal lines is the set of nodes x, where x is passed through by some non-upwards path from the top horizontal line to the bottom horizontal line.

In the above example the partial order given by the thatch is described in the more usual way by the following directed acyclic graph. The ordering is depicted with the order increasing down the page.

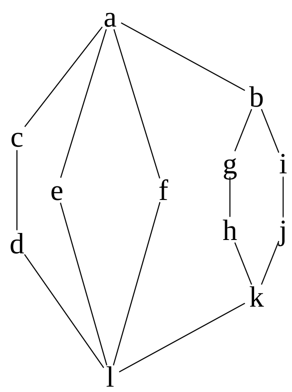

For example, $a < c < d < l$, and $a < b < g < h < k < l$.

Note that not any arrangement of vertical and horizontal lines forms a thatch. In particular note that the horizontal lines have to pair up in the correct way. Here is an example which is *not* a thatch.

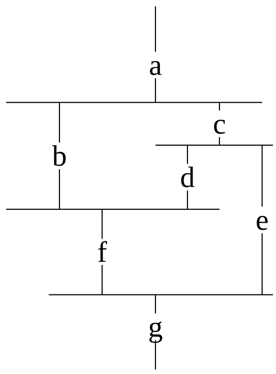

Define horizontal line l_1 to be contained between horizontal lines l_2 and l_3 if there is a non-upwards path which leads from some point on l_2 to some point on l_3 which passes through some point on l_1.

By construction horizontal lines are introduced in pairs. Moreover this is done in such a way that each pair contains an even number of other horizontal lines between them. Also if a pair contains no horizontal lines (which is still even of course) then every vertical line descending from the top horizontal line must terminate at the lower horizontal line. By inspection of the above diagram we can see that it is not possible to pair up the horizontal lines in this way, and so it does not constitute a thatch.

2 Set Theoretic Thatching

We can characterise a thatch partial order in a set theoretic manner. That is we can give a purely set theoretic characterisation of those partial orders which can be constructed as a thatch. Note that there may not be a unique thatch that corresponds to a particular partial order. To simplify matters we assume that every partial order has a unique top and bottom element. Given a partial order which does not, we can simply add two new elements to make this true. Without this restriction it becomes much more complex to succinctly characterise a thatch in a set theoretic manner.

In order to give the set theoretic characterisation we must first introduce some notation. Let $<$ be a partial order on a set S.

Definition 2.1

- $X^* = \{y \mid \forall x \in X,\, y > x\}$
- $X_* = \{y \mid \forall x \in X,\, y < x\}$
- $X^c = \{y \in X^* \mid \nexists z \in X^* : y > z\}$, this set is the ceiling of X
- $X_f = \{y \in X_* \mid \nexists z \in X_* : y < z\}$, this set is the floor of X
- For $x \in S$, let $\mathsf{nxt}(x) = \{x\}^c$, these are the next elements after x.
- For $x \in S$, let $\mathsf{lst}(x) = \{x\}_f$ these are the last elements before x.
- For a set S, define X to be a region of S if there are Y, $Z \subseteq S$ such that $X = Y^* \cap Z_*$
- $X^{\mathsf{ub}} = \{y \notin X \mid \exists x \in X : y \in \mathsf{nxt(x)}\}$, this is the upper boundary of X.
- $X_{\mathsf{lb}} = \{y \notin X \mid \exists x \in X : y \in \mathsf{lst(x)}\}$, this is the lower boundary of X.

Theorem 2.2
A partial order $<$ over a set S is a thatch ordering if and only if for every region $X \subseteq S$,

$$\begin{array}{rcl} X^c & = & X^{\mathsf{ub}} \\ X_f & = & X_{\mathsf{lb}} \end{array}$$

Given a thatch partial order, each region corresponds to a geometric region of the thatch. The above characterisation therefore gives a mechanism for constructing a particular thatch which describes the partial order. The ceilings and floors together with the regions describe how to collect together the nodes to form a thatch. If we are given a partial order which is not a thatch, it will be difficult to prove that this is the case using the above characterisation. However there is another characterisation, which we will discuss later, which is appropriate when disproving that a partial order is a thatch.

Here is a partial order which is a thatch. Note it has a unique top and bottom element.

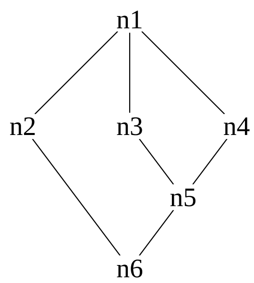

The regions for this ordering are

$$\begin{array}{lcl} \{n2, n3, n4, n5\} & = & \{n1\}^* \cap \{n6\}_* \\ \{n3, n4\} & = & \{n1\}^* \cap \{n5\}_* \\ \{n5\} & = & \{n3, n4\}^* \cap \{n6\}_* \end{array}$$

By direct calculation we can see that

$$\begin{array}{lclcl} \{n2, n3, n4, n5\}^c & = & \{n6\} & = & \{n2, n3, n4, n5\}^{\mathsf{ub}} \\ \{n2, n3, n4, n5\}_f & = & \{n1\} & = & \{n2, n3, n4, n5\}_{\mathsf{lb}} \\ \{n3, n4\}^c & = & \{n5\} & = & \{n3, n4\}^{\mathsf{ub}} \\ \{n3, n4\}_f & = & \{n1\} & = & \{n3, n4\}_{\mathsf{lb}} \\ \{n5\}^c & = & \{n6\} & = & \{n5\}^{\mathsf{ub}} \\ \{n5\}_f & = & \{n3, n4\} & = & \{n5\}_{\mathsf{lb}} \end{array}$$

Hence this is a thatch. Using the regions together with their floor and ceilings to reconstruct the thatch we can describe the partial order as follows.

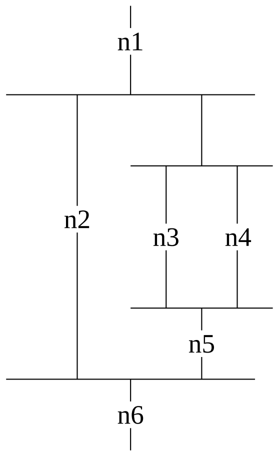

As another example consider Thatch 1. Here the regions are

$$\begin{array}{lcl} \{b,c,d,e,f,g,h,i,j,k\} & = & \{a\}^* \cap \{l\}_* \\ \{g,h,i,j\} & = & \{b\}^* \cap \{k\}_* \end{array}$$

As a final example illustrating theorem 2.2 consider this thatch.

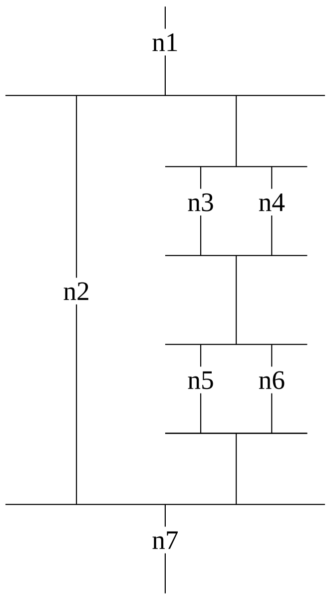

Here the regions are

$$\begin{array}{lcl} \{n2,n3,n4,n5,n6\} & = & \{n1\}^* \cap \{n7\}_* \\ \{n3,n4\} & = & \{n1\}^* \cap \{n5,n6\}_* \\ \{n5,n6\} & = & \{n3,n4\}^* \cap \{n7\}_* \end{array}$$

The floors and ceilings for these regions are:

$$\begin{array}{lcl}
\{n2,n3,n4,n5,n6\}^c & = \{n7\} & = \{n2,n3,n4,n5,n6\}^{\mathsf{ub}} \\
\{n2,n3,n4,n5,n6\}_f & = \{n1\} & = \{n2,n3,n4,n5,n6\}_{\mathsf{lb}} \\
\{n3,n4\}^c & = \{n5,n6\} & = \{n3,n4\}^{\mathsf{ub}} \\
\{n3,n4\}_f & = \{n1\} & = \{n3,n4\}_{\mathsf{lb}} \\
\{n5,n6\}^c & = \{n7\} & = \{n5,n6\}^{\mathsf{ub}} \\
\{n5,n6\}_f & = \{n3,n4\} & = \{n5,n6\}_{\mathsf{lb}}
\end{array}$$

The interesting point in this example is that $\{n3, n4\}$ is the floor of $\{n5, n6\}$, and $\{n5, n6\}$ is the ceiling of $\{n3, n4\}$.

There is an alternative characterisation of a thatch ordering which is easier to apply when proving that a given ordering is not a thatch ordering. This characterisation is more difficult to use when we wish to prove that a given partial order is a thatch ordering. Again we introduce some notation before giving the characterisation.

Definition 2.3

- For a pair a, b write $a|b$ when $a \not< b$ and $b \not< a$. We say that a and b are disjoint.
- A pair a, b are capped above if for every $c \in \{a, b\}^*$, for each d such that $c \not> d$, and $d > a$, then $d > b$.
- A pair a, b are capped below if for every $c \in \{a, b\}_*$, for each d such that $c \not< d$, and $d < a$, then $d < b$.

Theorem 2.4
A partial order $<$ on a set S is a thatch ordering if and only if every disjoint pair is capped below and above.

Here is a partial order, which is also a lattice. This lattice is not a thatch.

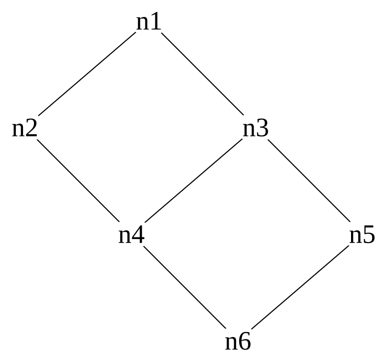

We can prove this using theorem 2.4. Consider n4 and n5, they are disjoint. Note n3 $\in$ { n4, n5 }$_*$, also n3 $\not<$ n2, but n2 $<$ n4. If this were a thatch then theorem 2.4 states that n2 $<$ n5, however in the lattice n2 $\not<$ n5 which is a contradiction. Hence the lattice can not be a thatch ordering.

In the next example we consider another lattice which is not a thatch. With this we illustrate that theorem 2.4 is simpler to use than theorem 2.2 when disproving that a partial order is a thatch.

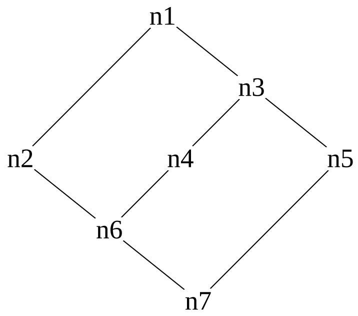

Notice that n6 $\in$ { n2, n3 }*, that n2|n3, n6 $\not>$ n5, and n5 $>$ n3. If this were a thatch then by theorem 2.4 n5 $>$ n2. This is not the case so this lattice can not be a thatch.

We can also use theorem 2.2 to disprove that this is not a thatch. Consider

$$\{n4, n5, n6\} = \{n3\}^* \cap \{n7\}_*$$

Therefore $\{n4, n5, n6\}$ is a region. However,

$$\begin{array}{rcl} \{n4, n5, n6\}_f & = & \{n3\} \\ \{n4, n5, n6\}_{lb} & = & \{n2, n3\} \end{array}$$

Therefore by theorem 2.2 the lattice is not a thatch. This calculation has required us to construct a region, its floor and its lower boundary. This is much more work than the first calculation.

3 Wait-and-see concurrent threads

A wait-and-see concurrent process may include any number of sub-processes, each regarded as a separate thread of execution. However if at any point the process launches a collection of concurrent threads it must wait and see how they terminate before proceeding. Each sub-process may also launch further sub-processes, which are also subject to the wait and see restriction.

A wait-and-see process can be regarded as a thatch. Each atomic line of code can be thought of as a node. Each sub-process corresponds to a vertical line in the thatch. When a process launches some sub-processes this is represented by a horizontal line, where each of the subprocesses corresponds to one of the vertical lines hanging down from the horizontal line. The termination of the subprocesses is represented by their respective thatches being attached at their ends to a common horizontal line.

For example Thatch 1 regarded as a wait-and-see process behaves as follows. After a is executed the process launches four sub-processes. The forth of these after execution of b launches a further two sub-processes. The first of these executes g then h then terminates. The second executes i then j then terminates. After both terminate sub-process four executes k before terminating. After all four sub-processes have terminated the main process executes l and then terminates.

Given a set of executable expressions and a specification of their inter-dependency in the form of a partial order on them, when is it possible to implement this specification in the form of a set of wait-and-see processes? Theorems 2.2 and 2.4 describe exactly when this can occur, and when it can not.