# FUNCTIONAL PEARL <br> Building a consensus: A rectangle covering problem 

RICHARD S.BIRD<br>Computing Laboratory, Oxford University<br>Wolfson Building, Parks Road, Oxford, OX1 3QD, UK<br>(e-mail: bird@comlab.ox.ac.uk)

## 1 Introduction

The other day, over a very pleasant lunch in the restaurant of Oxford's recently renovated Ashmolean Museum, ${ }^{1}$ Oege de Moor gave me a problem about rectangles. The problem is explained more fully later, but roughly speaking one is given a finite set of rectangles $R S$ and a rectangle $R$ completely covered by $R S$. The task is to construct a single rectangle covering $R$ among the elements of a larger set of rectangles associated with $R S$, called the saturation of $R S$. The saturation of $R S$ is the closure of $R S$ under so-called consensus operations, a term coined in (Quine, 1959), in which two rectangles are combined in two distinct ways to form new rectangles. The rectangle problem is a simplified version of containment-checking, a crucial component in a type inference algorithm for Datalog programs (Schäfer \& de Moor, 2010). In the Schäfer-de Moor algorithm the problem is generalised to cubes in $n$-space rather than rectangles in two-space, the components of each cube are given by propositional formulae rather than by intervals on the real line, and certain equality and inhabitation constraints are taken into account. Oege felt that the central proof, Lemma 15 in (Schäfer \& de Moor, 2010), deserved to be simplified so he posed the rectangle problem as a special case. This pearl was composed in response to the challenge.

## 2 The problem

Consider Figure 1, which depicts four rectangles $A, B, C$ and $D$, covering another rectangle $X$ framed with dotted lines. It happens in this example that none of the four rectangles can be removed without uncovering some part of $X$, but that is not essential to the problem. Next, consider two binary operations, $\oplus_{0}$ and $\oplus_{1}$, on rectangles. Assuming $A$ and $B$ overlap, the operation $A \oplus_{0} B$ forms a new rectangle by taking the union of the horizontal intervals defining $A$ and $B$ and the intersection of the vertical intervals. The operation $A \oplus_{1} B$ takes the intersection of the horizontals

[^0]

Fig. 1. Four rectangles covering the dotted rectangle.


Fig. 2. The dotted rectangles are $A \oplus_{0} D$ and $A \oplus_{1} D$.
and the union of the verticals. These two operations are the consensus operations on $A$ and $B$. For example, Figure 2 shows the two consensus rectangles for the rectangles $A$ and $D$ of Figure 1. The consensus operations also apply when the two rectangles do not overlap, though the result is either the empty rectangle or two rectangles.

Formally we define a rectangle $A$ to be a pair $\left(X_{A}, Y_{A}\right)$ of intervals on the real line. ${ }^{2}$ The points in $A$ are the elements in the cartesian product $X_{A} \times Y_{A}$. The two consensus operations are defined by

$$
\begin{aligned}
A \oplus_{0} B & =\left(X_{A} \cup X_{B}, Y_{A} \cap Y_{B}\right) \\
A \oplus_{1} B & =\left(X_{A} \cap X_{B}, Y_{A} \cup Y_{B}\right)
\end{aligned}
$$

Both $\oplus_{0}$ and $\oplus_{1}$ are associative, commutative and idempotent; moreover each operation distributes over the other. These properties follow from the same properties of union and intersection. By definition, the saturation of a set of rectangles $R S$ is the least set $X S$ containing $R S$ and containing $A \oplus_{0} B$ and $A \oplus_{1} B$ for all $A$ and $B$ in $X S$. Assuming $R S$ covers $R$, our task is to construct a single rectangle $X$ in the saturation of $R S$ that covers $R$. We do so by building a consensus expression that describes $X$. A consensus expression is a binary tree whose leaves are labelled with suitable names for the rectangles in $R S$ and whose nodes are labelled with one or other of the two consensus operations. Failure to construct a consensus means that the rectangles $R S$ do not cover $R$. To appreciate the problem, the reader is invited to construct a consensus expression over the four rectangles $A, B, C, D$ of

[^1]Figure 1 that covers the fifth rectangle $X$. The answer is given at the beginning of the following section.

The rectangle problem admits an immediate generalisation to $n$-cubes, which are cubes in $n$-dimensional space. An $n$-cube $A$ is defined by an $n$-tuple of sets $\left[X_{A}, Y_{A}, \ldots, Z_{A}\right]$. The points in $A$ are the elements of $X_{A} \times Y_{A} \times \cdots \times Z_{A}$, and there are $n$ consensus operations $\oplus_{j}$ for $0 \leqslant j<n$. The value of $A \oplus_{j} B$ is the cube in which the $k$ th component for $k \neq j$ is the intersection of the $k$ th components of $A$ and $B$ for $0 \leqslant k<n$, and the $j$ th component is the union of the $j$ th components of $A$ and $B$.

Finally, let us briefly relate the above consensus operations to those of same name in the Quine-McCluskey method for finding prime implicants. Consider a boolean formula in conjunctive normal form over $n$ propositional variables. For example, taking $n=3$ we might have the expression

$$
a^{\prime} b^{\prime} c+a b c^{\prime}+a^{\prime} b c+a b c
$$

in which the propositional variables are $a, b$ and $c$, and dashed letters denote negation. By definition, the consensus of $a^{\prime} b^{\prime} c$ and $a^{\prime} b c$ is the formula $a^{\prime} c$, equivalently $a^{\prime}\left(b+b^{\prime}\right) c$. If we interpret a term $a^{\prime} b^{\prime} c$ as a cube $\left[a^{\prime}, b^{\prime}, c\right]$ in which $a^{\prime}$ and $b^{\prime}$ are the complements of $a$ and $b$ with respect to some fixed universe, then the consensus of $a^{\prime} b^{\prime} c$ and $a^{\prime} b c$ is just $\left[a^{\prime}, b^{\prime}, c\right] \oplus_{1}\left[a^{\prime}, b, c\right]$.

## 3 Proof and first construction

One way of covering $X$ in Figure 1 is with the rectangle

$$
\left(B \oplus_{0} C\right) \oplus_{1}\left(A \oplus_{0} C\right) \oplus_{1}\left(A \oplus_{0} D\right)
$$

but there are many other expressions that also do the job.
Now, generalising from rectangles to $n$-cubes, how shall we prove that there is a consensus expression over a set of $n$-cubes $X S$ that covers a given cube $X$ if and only if $X S$ does? To answer this question let us first make a simplifying assumption, namely that each component of each cube is a finite set of integers. We will see in Section 4 that this assumption entails no loss of generality. Given that each component has finite size, we can prove the theorem by induction on the size of $X$, defined to be the product of the sizes of the components of $X$. A cube $X$ of size 1 has singleton components and so contains just a single point. A single point can clearly be covered by a single cube in $X S$ if and only if $X S$ covers $X$. This establishes the base case. Otherwise, by way of induction, assume that all cubes of size less than $s$ can be covered by a single cube in the saturation of $X S$. Let $X$ be a cube of size $s>1$, so at least one component of $X$ is a non-singleton set. Say the $k$ th component is not a singleton. Split $X$ into two cubes $Y$ and $Z$ by splitting the $k$ th component into two smaller sets, leaving the other components unchanged. Each of $Y$ and $Z$ has size smaller than $s$, so by induction there exist cubes $A$ and $B$ in the saturation of $X S$ that cover $Y$ and $Z$, respectively. Now $X$ is covered by $A \oplus_{k} B$.

The proof easily converts into a constructive procedure for computing the required expression. Turning to Haskell, we first declare

$$
\begin{array}{ll}
\text { type Cube } & =[\text { Component }] \\
\text { type Component } & =[\text { Int }]
\end{array}
$$

A cube is a list of $n$ components and each component is a finite, possibly empty list of integers in strictly increasing order. We allow empty components because the result of a consensus operation might be the empty cube. Each cube represents a set of points in $n$-space:

| type Point | $=[$ Int $]$ |
| :--- | :--- |
| points | $::$ Cube $\rightarrow[$ Point $]$ |
| points $[x s]$ | $=[[x] \mid x \leftarrow x s]$ |
| points $(x s: x s s)$ | $=[x: y s \mid x \leftarrow x s, y s \leftarrow$ points $x s s]$ |

A set $x c s$ of cubes covers a cube $x c$ if for each point of $x c$ there is some member of xcs that contains it:

$$
\begin{array}{ll}
\text { covers } & :: \quad[\text { Cube }] \rightarrow \text { Cube } \rightarrow \text { Bool } \\
\text { covers } x c s ~ x c ~ & =\text { and }[\text { contains } x c s ~ \\
\text { coints } x c] \\
\text { contains } x c s p & =o r[p \in \text { points } x c \mid x c \leftarrow x c s]
\end{array}
$$

Any set of cubes trivially covers the empty cube because the empty cube contains no points.

The operation $\oplus_{k}$ is implemented as a function consensus $k$ :

$$
\begin{aligned}
& \text { consensus } \quad:: \quad \text { Int } \rightarrow \text { Cube } \rightarrow \text { Cube } \rightarrow \text { Cube } \\
& \text { consensus } k x c y c=\text { zipWith cap xc } 1 \text { yc } 1+[\text { cup xs ys }]+\text { zipWith cap } x c 2 y c 2 \\
& \text { where }(x c 1, x s: x c 2)=\text { splitAt } k x c \\
& (y c 1, y s: y c 2)=\text { splitAt } k y c
\end{aligned}
$$

Definitions of cup and cap, the union and intersection functions on ordered lists, are omitted.

Next, we declare a suitable type for consensus expressions:

$$
\begin{aligned}
& \text { data Exp }=\text { Val Label |Op Int Exp Exp } \\
& \text { type Label }=\text { Int }
\end{aligned}
$$

The evaluation function eval is defined by

$$
\begin{array}{ll}
\text { eval } & :: \quad[\text { Cube }] \rightarrow \text { Exp } \rightarrow \text { Cube } \\
\text { eval xcs }(\text { Val } j) & = \\
\text { eval xcs }(O p k u v) & =\text { consensus } k(\text { eval } x c s u)(\text { eval xcs } v)
\end{array}
$$

Cubes are labelled by their positions in the given list of cubes.

The function cover 0 is a straightforward implementation of the induction proof (later on, we consider other versions, cover 1 , cover 2 , and so on):

```
cover \(0 \quad:: \quad[\) Cube \(] \rightarrow\) Cube \(\rightarrow\) Exp
cover 0 xcs xc
        \(=\) if null yss
        then \(\operatorname{Val}(\) head \([j \mid(j, y c) \leftarrow z i p[0\).. \(] x c s\), covers \([y c] x c])\)
        else \(O p\) (length xss) (cover \(0 x c s y c)(\operatorname{cover} 0 x c s z c)\)
        where
        \((x s s, y s s)=\) span single \(x c\)
        \(y c \quad=\quad x s s+[[\) head (head yss) \(]]+\) tail yss
        \(z c \quad=\quad\) xss \(+[\) tail (head yss) \(]+\) tail yss
```

The test single returns True on a singleton list and False otherwise. The expression span single $x c$ splits the components of $x c$ into two lists, xss and yss, where xss consists of singleton sets only and yss is either empty or begins with a non-singleton set. If $y s s$ is empty, then $x c$ consists of a single point. By construction,

$$
\text { covers } x c s x c \quad \Rightarrow \text { covers }[\text { eval xcs }(\text { cover } 0 x c s x c)] x c
$$

For example, cover 0 xcs [[1, 2, 3], [1, 2], [1, 2]] produces the expression

$$
\begin{aligned}
& \left(\left(111 \oplus_{2} 112\right) \oplus_{1}\left(121 \oplus_{2} 122\right)\right) \oplus_{0} \\
& \left(\left(211 \oplus_{2} 212\right) \oplus_{1}\left(221 \oplus_{2} 222\right)\right) \oplus_{0} \\
& \left(\left(321 \oplus_{2} 312\right) \oplus_{1}\left(311 \oplus_{2} 322\right)\right)
\end{aligned}
$$

in which the term $a b c$ represents the label of the first cube in xcs that covers the point $[a, b, c]$.

## 4 Digitization

We now show why we can assume that the components of cubes are finite sets of integers. An interval on the real line is defined by a pair $(a, b)$ of numbers representing the set of real numbers $x$ in the range $a \leqslant x<b$. The length of $(a, b)$ is $b-a$, so a pair $(a, a)$ denotes an empty interval. There is no loss in generality in requiring $a$ and $b$ to be integers, but, even so, there is an infinite number of points in each nonempty interval.

The solution is to digitize real intervals as integer intervals. Consider for example the three rectangles of Figure 3. Suppose their component intervals are as follows:

$$
A=[(0,50),(40,80)], B=[(30,70),(10,60)], C=[(40,80),(0,50)]
$$

The dotted lines shown in the figure extend each horizontal and vertical line to the $x$ and $y$ axes and partition the axes into five horizontal segments

$$
H_{0}=(0,30), H_{1}=(30,40), H_{2}=(40,50), H_{3}=(50,70), H_{4}=(60,80)
$$

and five vertical segments

$$
V_{0}=(0,10), V_{1}=(10,40), V_{2}=(40,50), V_{3}=(50,60), V_{4}=(60,80)
$$



Fig. 3. Dividing the axes into segments.
By construction, the components of $A, B$ and $C$ can each be expressed as the union of segments. For example, $X_{A}=H_{0} \cup H_{1} \cup H_{2}$ and $Y_{A}=V_{2} \cup V_{3} \cup V_{4}$. Going one step further, we can now represent each rectangle by the lists of integer labels of the horizontal and vertical segments that make up its coordinates. That gives new rectangles

$$
A=[[0 . .2],[2 \text {.. } 4]], B=[[1 \text {.. } 3],[1 \text {.. 3]], } C=[[2 \text {.. 4], }[0 \text {.. 2]] }
$$

whose components are intervals of integers. Moreover, containment-checking on the original problem returns the same result as on the digitized version. However, digitization is possible only because each cube is assumed to have components that are intervals.

To implement digitization we first declare the types

$$
\begin{aligned}
& \text { type RealCube }=[\text { Interval }] \\
& \text { type Interval }=(\text { Int }, \text { Int })
\end{aligned}
$$

We digitize each list of corresponding components separately:

$$
\begin{aligned}
& \text { digitize }:: \quad[\text { RealCube }] \rightarrow[\text { Cube }] \\
& \text { digitize }
\end{aligned}=\text { transpose } \cdot \text { map digitizeRow } \cdot \text { transpose }
$$

The list of cubes is transposed to bring each corresponding component into the same row. Each row is then digitized and the result is transposed back to give a new list of cubes. To implement digitizeRow, we first sort the boundaries of each interval into strictly increasing order, discarding duplicates:

$$
\begin{array}{lrl}
\text { boundaries } & :: & {[\text { Interval }] \rightarrow[\text { Int }]} \\
\text { boundaries } & = & \text { sort } \cdot \operatorname{concatMap}(\lambda(a, b) \rightarrow[a, b])
\end{array}
$$

The definition of sort is omitted. The function mksegs constructs a new list of intervals out of the boundaries:

$$
\begin{array}{lll}
\text { mksegs } & :: & {[\text { Int }] \rightarrow[\text { Interval }]} \\
\text { mksegs xs } & = & \text { zip xs }(\text { tail xs })
\end{array}
$$

Each new interval is then encoded as a component:

$$
\begin{aligned}
\begin{array}{l}
\text { encode } \\
\text { encode row }(x, y)
\end{array} & ::[\text { Interval }] \rightarrow \text { Interval } \rightarrow \text { Component } \\
\text { where } j & =\text { head }[j \mid(j,(a, b)) \leftarrow j a b s, a=x] \\
k & =\text { head }[j \mid(j,(a, b)) \leftarrow j a b s, b=y] \\
j a b s & =\text { zip }[0 . .] \text { row }
\end{aligned}
$$

That gives

$$
\text { digitizeRow row }=\text { map }(\text { encode }(\text { mksegs }(\text { boundaries row }))) \text { row }
$$

### 4.1 Propositional components

We mentioned briefly in the Introduction that in Datalog containment-checking the components of rectangles were propositional formulae rather than intervals. For example, consider the rectangles

$$
\begin{aligned}
& A=(a+c, a+b c) \\
& B=\left(a, b^{\prime}+c\right) \\
& C=\left(b c, a+b^{\prime} c\right) \\
& D=\left(a+b c, b^{\prime}\right) \\
& X=\left(a+b c, a+b^{\prime}\right)
\end{aligned}
$$

over the propositional variables $a, b$ and $c$. The rectangle $X$ is covered by $A, B, C$ and $D$ in the sense that

$$
\begin{array}{ll}
(a+b c) & \Rightarrow(a+c)+a+b c+(a+b c) \\
a+b^{\prime} & \Rightarrow(a+b c)+\left(b^{\prime}+c\right)+\left(a+b^{\prime} c\right)+b^{\prime}
\end{array}
$$

We can convert this version of the rectangle problem into an integer version by coding formulae using the truth table for $a, b$ and $c$ :

|  | $a b c$ |  | $a b c$ |
| :--- | :--- | :--- | :--- |
| 0 | 000 | 4 | 100 |
| 1 | 001 | 5 | 101 |
| 2 | 010 | 6 | 110 |
| 3 | 011 | 7 | 111 |

Using the labels of the entries in this table, the five rectangles above translate to the integer rectangles

$$
\begin{aligned}
A & =([1,3,4,5,6,7],[3,4,5,6,7]) \\
B & =([4,5,6,7],[0,1,3,4,5]) \\
C & =([3,7],[1,4,5,6,7]) \\
D & =([3,4,5,6,7],[0,1,4,5]) \\
X & =([3,4,5,6,7],[0,1,4,5,6,7])
\end{aligned}
$$

Of course, the components are not intervals, but cover 0 still works because it did not assume interval components.

## 5 Optimisation

Let us now return to cubes whose components are lists of integers. The trouble with cover $0 x c s x c$ is that each point of $x c$ is associated with a cube in $x c s$ and these cubes are then assembled into a potentially very large consensus expression. Because adjacent points of $x c$ may well be associated with the same cube, and each consensus operation is idempotent, there is room for improvement both in speed and expression size.

One way to reduce the size of the final consensus expression is to reduce the size of the cubes using digitization. For example, consider again the rectangles of Figure 3, this time as intervals of integers:

$$
A=[[0 . .49],[40 . .79]], B=[[30 . .69],[10 . .59)], C=[[40 . .79],[0 . .49]]
$$

Each of these rectangles is of size 2000, which is quite large. But digitization converts the problem into one whose rectangles each have size 9 , a substantial reduction. The function cover 1 is defined by

$$
\begin{array}{ll}
\text { cover } 1 & :: \quad[\text { Cube }] \rightarrow \text { Cube } \rightarrow \text { Exp } \\
\text { cover } 1 x c s x c & =\text { cover } 0 y c s y c \\
\text { where } y c: y c s & =\text { digitize }(\text { map }(\text { map toPair })(x c: x c s)) \\
\text { toPair } x s & = \\
(\text { head } x s, \text { last } x s+1)
\end{array}
$$

The function cover 1 , which is a valid optimisation only if the components of the cubes are contiguous intervals of integers, may reduce the size of the consensus expression but not necessarily minimize it.

In search of an algorithm that does minimize expression size, let us first recast cover 0 in iterative form. To this end, consider the function expand $k$ that splits a cube into $s_{k}$ subcubes, where $s_{k}$ is the size of the $k$ th component:

$$
\begin{aligned}
\text { expand }: & \text { Int } \rightarrow \text { Cube } \rightarrow[\text { Cube }] \\
\text { expand } k x c= & {[x s s+[[y]]+y s s \mid y \leftarrow y s] } \\
& \text { where }(x s s, y s: y s s)=\text { splitAt } k x c
\end{aligned}
$$

Using expand we can define

```
cover \(2 x c s x c=\) pcover \(0 x c\)
where
pcover \(k x c=\) if \(k=n\)
    then \(\operatorname{Val}(\) head \([j \mid(j, y c) \leftarrow z i p[0 .] x c\).\(s , covers [y c] x c])\)
    else foldr \(1(O p k)(\) map \((p c o v e r ~(k+1))(\operatorname{expand} k x c))\)
\(n \quad=\) length \(x c\)
```

We omit the proof that cover $0=$ cover 2 . The function cover 2 is faster than cover 0 , though not dramatically so. However, cover 2 gets us, quite literally, to the nub of the problem. Let $\oplus$ be some associative, commutative and idempotent operation. The critical fact about $\oplus$ is that

$$
\text { foldr } 1(\oplus)=\text { foldr } 1(\oplus) \cdot \text { nub }
$$

where nub is the standard Haskell library function that removes duplicates from a list. The value of foldr $1(\oplus) x s$ depends only on the set of elements in $x s$ and not on the order of the elements or on duplications.

We could minimize the size of the final consensus expression simply by inserting a nub after foldr $1(O p k)$ in the definition of pcover, but nub uses a quadratic number of equality tests and each equality test compares two consensus expressions, so the new version, cover 3 say, may well be slower than cover 2 .

Here is a better way to remove duplicates. The idea is to convert the target cube $x c$ to an $n$-dimensional array of points, then to replace each point by the label of the first cube in xcs which covers it. This array of labels is then pruned, dimension by dimension, by removing all duplicates in each row, then all duplicate rows, all duplicate planes, and so on. To represent an $n$-dimensional array in Haskell we need trees:

$$
\text { data Tree } a=\text { Leaf } a \mid \text { Fork }[\text { Tree } a]
$$

The function convert converts an $n$-cube into a list of trees of points:

```
convert \(\quad:: \quad\) Cube \(\rightarrow\) [Tree Point \(]\)
convert \([x s] \quad=[\) Leaf \([x] \mid x \leftarrow x s]\)
convert (xs:xss) \(=[\) Fork (map (cons \(x\) )ts) \(\mid x \leftarrow x s]\)
        where \(t s=\) convert \(x s s\)
cons \(x\) (Leaf \(x s)=\) Leaf \((x: x s)\)
cons \(x\) (Fork ts) \(=\operatorname{Fork}(\) map (cons \(x) t\) s)
```

The function covering xcs labels a tree of points with the first cube in xcs that covers each point:

```
covering \(::[\) Cube \(] \rightarrow\) Cube \(\rightarrow[\) Tree Label \(]\)
covering xcs \(x c=\) map label (convert \(x c\) )
where
label (Leaf xs) \(=\) Leaf (lookup xs)
label (Fork ts) \(=\) Fork (map label ts)
lookup \(x \quad=\) head \([j \mid(j, x c) \leftarrow z i p[0 .] x c\).\(s , member x x c]\)
```

The function member is defined by

$$
\text { member } x x c=\text { and }[a \leqslant y \wedge y<b \mid((a, b), y) \leftarrow \text { zip (map toPair } x c) x]
$$

The function prune removes duplicate labels:

```
prune :: [Tree Label] }->\mathrm{ [Tree Label]
prune ts@(Leaf_: )}== remdups t
prune ts = remdups [Fork (prune vs)| Fork vs \leftarrowts]
remdups }\quad:: Eq a m [a]->[a
remdups = foldr op []
    where opx [] = [x]
        opx (y:ys)= if x=y then y:ys else x:y:ys
```

The function remdups removes adjacent duplicates only, but that is sufficient because in the cubes problem all duplicates will occur together. The function lookup always returns the first cube in xcs that covers the point $x$ and that fact guarantees that duplicate cubes occur together.

Now we can define our final version:

$$
\begin{aligned}
& \text { cover } 4 \quad::[\text { Cube }] \rightarrow \text { Cube } \rightarrow \text { Exp } \\
& \text { cover } 4 x c s x c \quad=\quad \text { mkexp } 0(\text { prune (covering ycs yc)) } \\
& \text { where } y c: y c s=\operatorname{digitize}(\operatorname{map}(\text { map toPair })(x c: x c s)) \\
& \text { mkexpkts@(Leaf_:-) }=\text { foldr } 1(\text { Opk })[\text { Val } x \mid \text { Leaf } x \leftarrow t s] \\
& \text { mkexpkts } \quad=\text { foldr } 1(\text { Opk })[\text { mkexp }(k+1) v s \mid \text { Fork } v s \leftarrow t s]
\end{aligned}
$$

## 6 Experimental results

Using ghci we carried out a brief comparison of the five versions of cover. Recall that cover 0 was the divide and conquer version, cover 1 the version that reduced problem size by digitization, cover 2 was the linearised form of cover 0 , cover 3 used the Haskell nub function, and cover 4 the super-duper version that used both digitization and a tailored implementation for removing duplicates. The results, using three sample sets of cubes, are summarised in the table below. Times are in seconds, and size refers to the size of the consensus expression. The set $x c s 0$ is a set of rectangles of total area 525, the set $x c s 1$ a set of rectangles of total area 161, and $x c s 2$ a set of four-dimensional cubes of area 8097.

|  | $x c s 0$ |  | $x c s 1$ |  | $x c s 2$ |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
|  | time | size | time | size | time | size |
| cover 0 | 0.05 | 241 | 1.64 | 503 | 38.97 | 1295 |
| cover 1 | 0.00 | 49 | 0.02 | 31 | 0.06 | 23 |
| cover 2 | 0.05 | 241 | 1.63 | 503 | 38.89 | 1295 |
| cover 3 | 0.05 | 17 | 1.59 | 3 | 38.75 | 9 |
| cover 4 | 0.02 | 17 | 0.02 | 3 | 0.02 | 9 |

Clearly cover 4 was best by a long way, so it was worth the effort to obtain it.

## Acknowledgments

The author thank Jeremy Gibbons and two anonymous referees for their constructive remarks and suggestions. As a result the pearl has been thoroughly overhauled.

## References

Quine, W. V. (1959) On cores and prime implicants of truth functions. Am. Math. Mon., 66(9), 755-760.
Schäfer, M. \& de Moor, O. (2010) Type inference for datalog with complex type hierarchies. Princ. Program. Lang. 2010. ACM SIGPLAN Not., 45(1), 145-156.


[^0]:    ${ }^{1}$ Well worth a visit if you are ever in Oxford, especially as entry to the museum is free.

[^1]:    ${ }^{2}$ The assumption that $X_{A}$ and $Y_{A}$ are intervals is considered further in Section 4.1.

