Consistency of natural relations on sets

Hirotaka KOIZUMI[†], Akira MARUOKA[††] and Mike PATERSON[†††]

† 2nd Development Dept. 1st Basic Software Development Division, NEC Corporation
4-14-22, Shibaura, Minato-ku, Tokyo 108, Japan

†† Faculty of Engineering, Tohoku University, Sendai 980, Japan

††† Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK

Abstract

Five natural relations for sets, such as inclusion, disjointness, intersection, etc., are introduced in terms of the emptiness of the subsets defined by Boolean combinations of the sets. Let N denote $\{1, 2, \ldots, n\}$ and $\binom{N}{2}$ denote $\{(i, j) \mid i, j \in N \text{ and } i < j\}$. A function μ on $\binom{N}{2}$ specifies one of these relations for each pair of indices. Then μ is said to be *consistent on* $M \subseteq N$ if and only if there exists a collection of sets corresponding to indices in M such that the relations specified by μ hold between each associated pair of the sets. In this paper it is proved that if μ is consistent on all subsets of N of size three then μ is consistent on N. Furthermore, conditions that make μ consistent on a subset of size three are given explicitly.

Key words. inclusion and exclusion relations, subsets, consistency, locally computable

AMS(MOS) subject classification. 04A20, 05A05, 06A07, 68R05

1 Introduction

Let n be a natural number and N denote the set $\{1, \ldots, n\}$. Suppose that we are given some combinatorial object μ defined on N and that for any subset M of N, the object, denoted by μ_M, is obtained by restricting μ to the subset M. Let P denote some predicate defined on all of the objects μ_M, where M is a subset of N. We

A part of the work was done while the second author was visiting at Brown University. The third author gratefully acknowledges support from the British Council while visiting Tohoku University.

consider predicates P that are inheritable, in the sense that if P holds on μ_M then P holds on $\mu_{M'}$ for any $M' \subset M$. Such a predicate P often turns out to be "locally computable", i.e., if P holds on all objects μ_M with M satisfying some conditions then P holds on the whole object μ_N. One typical example of such a predicate is given by Helly's Theorem [H]. For example, in two dimensions, this states that a family of compact convex planar sets has a nonempty intersection if and only if every triple of the sets has a nonempty intersection.

In this paper we give another locally computable predicate for which we are only required to check the predicate on all objects μ_M for "small" M. First we describe five *natural relations* between sets, denoted by $\subset$, $\supset$, $\|$, $\perp$, $\bowtie$, representing strict inclusion in each direction, disjointness, covering the universe, and the general case, respectively, Each can be defined in terms of the emptiness or otherwise of Boolean combinations of the sets, and the set of these five relations is denoted by $\mathcal{R}$. Let $\binom{N}{2}$ denote $\{(i,j) \in N^2 \mid i < j\}$. Each object μ we deal with is an assignment of a natural relation to each pair (i,j) in $\binom{N}{2}$, i.e., a function from $\binom{N}{2}$ to $\mathcal{R}$.

A collection of sets $S_1, \ldots, S_n$ is *compatible with* μ if the relation $\mu(i,j)$ holds between S_i and S_j for all (i,j) in $\binom{N}{2}$. If μ is compatible with some such collection then μ is said to be *consistent*. For any subset $M \subseteq N$, the object μ_M is simply the restriction of μ to $\binom{M}{2}$. If μ_M is consistent then μ is said to be *consistent on* M. Our main result is that if μ is consistent on every subset of N of size three, then μ is consistent. Conditions that make μ consistent are given explicitly in terms of the natural relations that may hold for any three subsets.

The problem of characterizing a predicate on graphs, or equivalently a family of graphs satisfying the predicate, is also discussed in [FL] in a different context. And it is pointed out that to find a finite number of combinatorial structures to characterize a family is crucial. Some combinatorial aspects of inclusion and exclusion and their relation to Boolean complexity are also discussed in [LN].

In Section 2, after natural relations on sets have been introduced, some constraints on natural relations of μ are given which guarantee that μ is consistent. In order to prove the statement, a set of vectors is used as a model for μ. It is also shown that there exists a feasible algorithm which, given a partial function μ from $\binom{N}{2}$ to $\mathcal{R}$, decides whether or not μ can be extended to obtain a consistent total function. In Section 3, a graph based on the inclusion relations of μ is introduced as another intuitive model for μ, and the same result as in Section 2 is verified using the model. In Section 4, some open problems are presented.

2 Consistency conditions for natural relations

Let the universe U be nonempty. A *natural relation* for any set or sets is one that is defined in terms of the emptiness or otherwise of the subsets defined by Boolean combinations of the sets. For one set, there are four cases, depending on the emptiness of the set and its complement. If both are empty then $S = \phi$, a case we have excluded. The remaining three cases correspond to the set being empty, equal to the universe and proper, respectively. A subset S of U is called *proper* if neither $S = \phi$ nor $S = U$. We will allow only the proper subsets of U.

For two sets A and B, there are formally 16 possible relations. Under our assumptions that the universe is nonempty and both sets are proper, there remain just seven cases. One is $A = B$, another is $A = \overline{B} = U \neg B$. Both of these cases are special in that if they hold then one of the sets can be eliminated by substitution from the remaining relations. The remaining five natural relations constitute $\mathcal{R}$. Table 1 defines these five relations in terms of the emptiness, denoted by 0, or nonemptiness, denoted by 1, of four subsets. In the Table, (a, b) indicates the subset $A^a \cap B^b$, where a and b are in $\{0, 1\}$, $S^1 = S$ and $S^0 = \overline{S}$.

	relations				
subsets	$\subset$	$\supset$	$\parallel$	$\perp$	$\bowtie$
(1,1)	1	1	0	1	1
(1,0)	0	1	1	1	1
(0,1)	1	0	1	1	1
(0,0)	1	1	1	0	1

Table 1: Five natural relations

Let Σ denote $\{0, 1\}$. For v in Σ^n, let $v^{(i)}$ denote the i^{th} component of v. Given n subsets $S_1, \ldots, S_n$ of U, we can determine whether each subset of the form

$$S_1^{v^{(1)}} \cap S_2^{v^{(2)}} \cdots \cap S_n^{v^{(n)}}$$

is empty or not, where $v = (v^{(1)}, \ldots, v^{(n)})$ is in Σ^n. Let V_S denote the set of vectors v in Σ^n such that $S_1^{v^{(1)}} \cap \cdots \cap S_n^{v^{(n)}}$ is nonempty. Furthermore let T_i for $1 \leq i \leq n$ denote the set of vectors v in V_S such that $v^{(i)} = 1$. Then it is easy to see that consistency of $S_1, \ldots, S_n$ for $\mathcal{R}$ is the same as that of $T_1, \ldots, T_n$. In other words, for any natural relation α, $S_i \alpha S_j$ holds if and only if $T_i \alpha T_j$ holds. So, without loss of generality, we can consider a set of vectors V rather than a collection of subsets as far as the consistency problem is concerned.

Relation α in Table 1 is considered to be a function from Σ^2 to Σ in the obvious way: for (a, b) in Σ^2, $\alpha(a, b) = 1$ if $A^a \cap B^b$ is nonempty in the relation, and $\alpha(a, b) = 0$ otherwise. Relation α can also be considered to be the set of vectors in Σ^2 for which the function takes value 1, e.g., $\subset = \{(1, 1), (0, 1), (0, 0)\}$. Let $v^{(i,j)}$ denote $(v^{(i)}, v^{(j)})$, and $V^{(i,j)}$ denote $\{v^{(i,j)} \mid v \in V\}$. This notation can be generalized in an obvious way to the case of more indices. Let μ be a function from $\binom{N}{2}$ to $\mathcal{R}$, and let $M \subseteq N$. We say that μ is *compatible with* V on M if and only if $V^{(i,j)} = \mu(i, j)$ for all (i, j) in $\binom{M}{2}$, and that μ is *consistent on* M if and only if there exists a subset V of Σ^n that is compatible with μ on M. In particular, when $M = N$, the phrase "on M" in the definition may be dropped.

Proposition 1. Let μ be consistent. Let V' denote the set $\{v \in \Sigma^n \mid \exists\, (i, j) \in \binom{N}{2}, \mu(i, j)(v^{(i)}, v^{(j)}) = 0\}$. Then μ is compatible with $\Sigma^n \neg V'$.

The proof is immediate from the definitions of consistency and compatibility. Note that, with μ and V' as in Proposition 1, $\Sigma^n \neg V'$ is the largest set of vectors that is compatible with μ in the sense that if μ is compatible with V, then we have $V \subseteq \Sigma^n \neg V'$.

For u in Σ^n and $A \subseteq \Sigma^n$, let $u \oplus A = \{u \oplus v \mid v \in A\}$, where $u \oplus v$ denotes the vector obtained by taking the bit-wise "exclusive or" of u and v. For u in Σ^n, the transformation φ_u on the set of functions from $\binom{N}{2}$ to $\mathcal{R}$ is defined as $\varphi_u(\mu)(i, j) = u^{(i,j)} \oplus \mu(i, j)$. Note that, in the definition, $\mu(i, j)$ is thought of as a subset of Σ^2.

Clearly we have the next proposition.

Proposition 2. Let u be in Σ^n and $V \subseteq \Sigma^n$. Then μ is compatible with V if and only if $\varphi_u(\mu)$ is compatible with $u \oplus V$.

In view of Proposition 2, we have the next proposition which says that the transformation φ_u preserves the consistency of μ for any u in Σ^n.

Proposition 3. Let u be in Σ^n. Then μ is consistent if and only if $\varphi_u(\mu)$ is consistent.

Before proceeding to the main theorem, we show in Figure 1 how the five natural relations are transformed by φ_u for $u = (1, 0)$ and $(0, 1)$.

It is convenient to extend the definition of any object μ to $\{(j, i) \mid i < j\}$ in the obvious way, so that $\mu(i, j) = \subset$ if and only if $\mu(j, i) = \supset$, and $\mu(i, j) = \mu(j, i)$ if $\mu(i, j) \in \{\|, \perp, \bowtie\}$.

If μ is consistent then the transitivity of inclusion implies that the following constraint holds for any distinct indices i, j and k.

$(*)$ if $\mu(i, j) = \subset$ and $\mu(j, k) = \subset$ then $\mu(i, k) = \subset$.

By applying Proposition 3 for various choices of u, we can transform the constraint

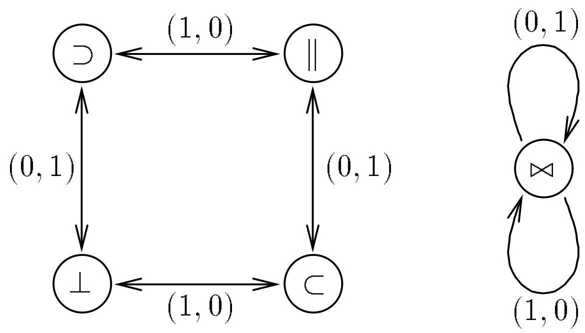

Figure 1: Transformations on natural relations

($*$) in various ways. For example, let $u^{(i,j,k)} = (1,0,1)$ and $\mu' = \phi_u(\mu)$. Now, if $\mu(i,j) = \perp$ and $\mu(j,k) = \parallel$ then $\mu'(i,j) = \subset$ and $\mu'(j,k) = \subset$, and so $\mu'(i,k) = \subset$, which implies that $\mu(i,k) = \supset$. In Table 2 we show the eight *transitivity constraints* which are obtained. If μ satisfies these eight constraints it is said to be *transitive.*

$u^{(i,j,k)}$	$\mu(i,j)$	$\mu(j,k)$	$\mu(i,k)$
(0,0,0)	$\subset$	$\subset$	$\subset$
(0,0,1)	$\subset$	$\parallel$	$\parallel$
(0,1,0)	$\parallel$	$\perp$	$\subset$
(0,1,1)	$\parallel$	$\supset$	$\parallel$
(1,0,0)	$\perp$	$\subset$	$\perp$
(1,0,1)	$\perp$	$\parallel$	$\supset$
(1,1,0)	$\supset$	$\perp$	$\perp$
(1,1,1)	$\supset$	$\supset$	$\supset$

Table 2: The eight transitivity constraints

The next theorem says that these conditions that are necessary to make μ consistent on any set of three indices turn out to be sufficient conditions to make μ consistent on the set of all indices.

Theorem 4. If μ is transitive then μ is consistent.

Proof. We shall prove the statement of the theorem by induction on n. The statement holds trivially when $n = 2$. Assume that the statement holds for $n - 1$, where $n \geq 3$.

Let $N' = \{1, 2, \ldots, n-1\}$. In view of Figure 1, it is easy to see that there exists some u in Σ^n such that $\varphi_u(\mu)(i,n) \in \{\subset, \supset, \bowtie\}$ for all i in N'. By Proposition 3 it suffices to show that $\varphi_u(\mu)$ is consistent. Furthermore, it is easy to see that if μ satisfies the conditions of the theorem then $\varphi_u(\mu)$ also satisfies the conditions. So, denoting $\varphi_u(\mu)$ again by μ, we may assume that μ satisfies the conditions of the theorem and that $\mu(i,n) \in \{\supset, \subset, \bowtie\}$ for all i in N'. We will show that μ is consistent.

Let

$$\begin{aligned}
V_1 &= \{v \in \Sigma^n \mid \exists\, (i,j) \in \tbinom{N'}{2}, \mu(i,j)(v^{(i,j)}) = 0\}, \\
V_2 &= \{v \in \Sigma^n \mid \exists\, i \in N',\ \mu(i,n)(v^{(i,n)}) = 0\}, \\
V' &= \Sigma^n \neg V_1, \\
V &= \Sigma^n \neg V_1 \neg V_2.
\end{aligned}$$

Let μ' be the function obtained by restricting μ to $\binom{N'}{2}$. Then by the induction hypothesis and Proposition 1, μ' is compatible with V'.

Fact 1. For any v in V' at least one of $(v^{(1)}, \ldots, v^{(n-1)}, 0)$ and $(v^{(1)}, \ldots, v^{(n-1)}, 1)$ is in V.

Proof. Assume to the contrary that there exists v in V' such that both of $(v^{(1)}, \ldots, v^{(n-1)}, 0)$ and $(v^{(1)}, \ldots, v^{(n-1)}, 1)$ belong to V_2. Then, since $\mu(k,n) \in \{\supset, \subset, \bowtie\}$ for any k in N', there exist i and j in N' such that

$$\begin{aligned}
\mu(i,n) &= \subset, \\
\mu(j,n) &= \supset, \\
v^{(i)} &= 1, \\
v^{(j)} &= 0.
\end{aligned}$$

Hence by transitivity we have $\mu(i,j) = \subset$, which, together with $v^{(i)} = 1$ and $v^{(j)} = 0$, implies that v belongs to V_1, contradicting the assumption. $\square$

By Fact 1 we have

$$V^{(i,j)} = V'^{(i,j)}$$

for any (i,j) in $\binom{N'}{2}$. On the other hand, by the induction hypothesis and Proposition 1 we have $V'^{(i,j)} = \mu'(i,j) = \mu(i,j)$ for any (i,j) in $\binom{N'}{2}$. Hence we have $V^{(i,j)} = \mu(i,j)$ for any (i,j) in $\binom{N'}{2}$.

So it remains to show that

$$V^{(i,n)} = \mu(i,n)$$

holds for any i in N'. For v in Σ^n and relation α, let

$$\begin{aligned}
A_\alpha(v) &= \{i \in N' \mid v^{(i)} = 1, \mu(i,n) = \alpha\}, \\
B_\alpha(v) &= \{i \in N' \mid v^{(i)} = 0, \mu(i,n) = \alpha\}.
\end{aligned}$$

Note that from the assumption we only have to consider the case where α belongs to $\{\subset, \supset, \bowtie\}$.

Fact 2. Let v be in V'. Let $\tau_0(v)$ and $\tau_1(v)$ be vectors in Σ^n such that

$$\tau_0(v)^{(j)} = \begin{cases} 0 & \text{if } j \in A_\subset(v) \text{ or } j = n, \\ v^{(j)} & \text{otherwise,} \end{cases}$$

$$\tau_1(v)^{(j)} = \begin{cases} 1 & \text{if } j \in B_\supset(v) \text{ or } j = n, \\ v^{(j)} & \text{otherwise.} \end{cases}$$

Then $\tau_0(v)$ and $\tau_1(v)$ are in V.

Proof. Let $\tau_0(v)$ and $\tau_1(v)$ be denoted by v_0 and v_1, respectively. It is easy to see that $\mu(i,n)(v_b^{(i,n)}) = 1$ for any b in Σ and any i in N', and hence $v_b \notin V_2$ for any b in Σ. So it suffices to show $v_b \notin V_1$ for any b in Σ. To do so, we assume to the contrary that there exist b in Σ and (i,j) in $\binom{N'}{2}$ such that $\mu(i,j)(v_b^{(i,j)}) = 0$. Since $\mu(i,j)(v^{(i,j)}) = 1$ for all (i,j) in $\binom{N'}{2}$, at least one of $v_b^{(i)} \neq v^{(i)}$ and $v_b^{(j)} \neq v^{(j)}$ holds. Without loss of generality, we may assume that $b = 0$ and $v_0^{(i)} \neq v^{(i)}$. So we only need to consider the following cases.

Case 1. $i \in A_\subset(v), v_0^{(j)} = 0$ and $\mu(i,j) = \perp$.

Since $\mu(j,i) = \perp$ and $\mu(i,n) = \subset$, by transitivity (see Table 2) we have $\mu(j,n) = \perp$, which contradicts the assumption that $\mu(k,n) \in \{\subset, \supset, \bowtie\}$ for any k in N'.

Case 2. $i \in A_\subset(v), v_0^{(j)} = 1$ and $\mu(i,j) = \supset$.

Since $\mu(j,i) = \subset$ and $\mu(i,n) = \subset$, by transitivity we have $\mu(j,n) = \subset$, which contradicts $v_0^{(j)} = 1$. □

Fact 3. Let $\mu(i,n)(v^{(i,n)}) = 0$ for some $i \in N'$ and $v \in V'$. Then for any $(a,b) \in \Sigma^2$ other than $v^{(i,n)}$, there exists $v' \in V$ such that $v'^{(i,n)} = (a,b)$.

Proof. Let i and v be as in the hypothesis. Without loss of generality we may assume that $\mu(i,n) = \subset$, and $v^{(i,n)} = (1,0)$ for some $i \in N'$ and $v \in V'$.

Since $v \notin V$, Fact 1 implies that $w = v \oplus (0, \ldots, 0, 1) \in V$, and $w^{(i,n)} = (1,1)$. Since $\tau_0(v) \in V$ and $i \in A_\subset(v)$, we have $\tau_0(v)^{(i,n)} = (0,0)$. Finally, since $\tau_0(v) \in V \subset V'$, we have $\tau_1(\tau_0(v)) \in V$, and $\tau_1(\tau_0(v))^{(i,n)} = (0,1)$. □

By the induction hypothesis we have that $V'^{(i,n)} = \{(0,0),(0,1),(1,0),(1,1)\}$ for any i in N'. Thus, by Fact 3, $V^{(i,n)} = \mu(i,n)$ holds for any i in N', completing the proof of the theorem. □

Before closing the section, we note that using Theorem 4 we can construct a feasible algorithm which, given a partial function μ from $\binom{N}{2}$ to $\mathcal{R}$ decides whether or not μ can be extended to obtain a consistent total function. The algorithm works as follows. Given a partial function μ, check if it satisfies the transitivity constraints. If not, give the answer that μ is not consistent. Otherwise, extend μ

using the transitivity constraints repeatedly until none of these constraints can be applied. In doing this, if there exists a pair to which different relations are assigned then give the answer that μ is not extensible consistently. Otherwise, conclude that μ is extensible consistently. In fact, if we assign $\bowtie$, to any pairs that remain unspecified at the end of the algorithm, we obtain a total function. Clearly the total function obtained in this way is an extension of μ and is consistent in view of Theorem 4.

3 An intuitive model for μ

In this section we shall introduce another, intuitive, model for μ based directly on the natural relations, so that we can give another proof of Theorem 4.

Proposition 5. If μ is transitive then there exists u in Σ^n such that $\varphi_u(\mu)(i,j) \in \mathcal{R}_0 = \{\subset, \supset, \|, \bowtie\}$ holds for any (i,j) in $\binom{N}{2}$.

Proof. Let μ be transitive. Then it is easy to see that $\varphi_u(\mu)$ is also transitive for any u in Σ^n. We shall prove the conclusion of the proposition by induction on n. When $n = 2$, Figure 1 shows the result at once. Let $N' = \{1, 2, \ldots, n-1\}$ and assume as the inductive hypothesis that the result holds for N'. Then there exists u' in Σ^n such that $\varphi_{u'}(\mu)(i,j) \in \mathcal{R}_0$ holds for any (i,j) in $\binom{N'}{2}$. Now at most one of $\supset$ and $\perp$ appears in $\varphi_{u'}(\mu)(1,n)$, $\varphi_{u'}(2,n), \ldots, \varphi_{u'}(n-1,n)$. This is because, if there exist i, j in N' such that $\varphi_{u'}(\mu)(i,n) = \perp$ and $\varphi_{u'}(\mu)(j,n) = \supset$, then by transitivity we have $\varphi_{u'}(i,j) = \perp$, contradicting the assumption. Thus by taking $u = u' \oplus (0, \ldots, 0, 1)$ if there exists i in N' such that $\varphi_{u'}(\mu)(i,n) = \perp$, and taking $u = u'$ otherwise, we see from Figure 1 that $\varphi_u(\mu)(i,j) \in \mathcal{R}_0$ holds for any (i,j) in $\binom{N}{2}$, completing the induction step. $\square$

Let $\varphi_u(\mu)$ be as in Proposition 5. We note that in order to obtain a model for μ it is sufficient by Proposition 2 to obtain a model for $\varphi_u(\mu)$. We denote again $\varphi_u(\mu)$ by μ so that $\mu(i,j) \in \mathcal{R}_0$ for any (i,j) in $\binom{N}{2}$. We shall define a collection of subsets that is compatible with μ. To do this, consider the directed graph G' with vertex set $V' = \{x_1, \ldots, x_n\}$ and edge set $E' = \{(x_i, x_j) \mid \mu(i,j) = \supset\}$. Since μ is transitive, G' is an acyclic graph, and we define S'_i to be the set of x_i and its descendants. It is easy to see that $S'_i \supset S'_j$ if and only if $(x_i, x_j) \in E'$, and if and only if $\mu(i,j) = \supset$, so that G' already gives a model for the set containment relations of μ.

For a complete model for μ we need to extend G' with extra vertices. Define $G'' = (V'', E'')$, where $V'' = V' \cup \{x_{i,j} \mid \mu(i,j) = \bowtie \text{ and } i < j\}$ and $E'' = E' \cup \{(x_i, x_{i,j}), (x_j, x_{i,j}) \mid x_{i,j} \in V''\}$. As in the case of graph G', let S''_i be the set of x_i and its descendants in the graph G''. If there exist i and j in N such that $S''_i \cup S''_j = V''$ holds, then let $V = V'' \cup \{x_\infty\}$ and $E = E''$. Otherwise, let $V = V''$

and $E = E''$. The final graph G is defined to be (V, E). Now we define S_i to be the set consisting of x_i and its descendants in the graph G. The containment relation on the new sets is the same as in the graph G' and agrees with $\mu^{-1}(\supset)$. Therefore, if $\mu(i, j) = \bowtie$ then neither containment can hold between S_i and S_j but $x_{i,j} \in S_i \cap S_j$. Hence, since $S_i \cup S_j$ cannot be the whole set V, we have $S_i \bowtie S_j$. For the proof of the converse, suppose that $S_i \bowtie S_j$ and so $\mu(i, j) \in \{\|, \bowtie\}$. If there is some x_k in $S_i \cap S_j$ then $k \neq i, j$, so $S_k \subset S_i$ and $S_k \subset S_j$. Hence $\mu(k, i) = \mu(k, j)$ and the transitivity constraints imply that $\mu(i, j) \neq \|$. Otherwise there is some $x_{k,l}$ in $S_i \cap S_j$ where $S_k \subseteq S_i$ and $S_l \subseteq S_j$, and so (I) $k = i$ or $\mu(k, i) = \subset$, and (II) $l = j$ or $\mu(l, j) = \subset$. If $\mu(i, j) = \|$ then the constraints with (I) and (II) yield $\mu(k, l) = \|$. This result is contradicted by the existence of $x_{k,l}$ and this completes the proof that the graph G is a model for μ.

4 Concluding remarks

We investigated the problem of deciding whether or not a function μ that specifies the type of the natural relation for each pair of sets is consistent, and proved that if μ is consistent on any three sets then μ is consistent on the whole collection of sets. So the problem of deciding the consistency of μ for n sets can be computed in time $O(n^3)$. Furthermore, it can be seen [J] that the problem can be solved in time $O(n^{2.37})$. To show this fact, let $M_\subset$ be the matrix whose (i, j) component is 1 if $\mu(i, j) = \subset$, and 0 otherwise. Then the constraint $(*)$ in Section 2 can be written as $M_\subset M_\subset \leq M_\subset$, where the matrix product is done using Boolean sum and product, and "$\leq$ " holds between matrices if and only if "$\leq$ " holds between all the corresponding components in the matrices. Likewise, we can rewrite the remaining transitivity constraints in Table 2 in matrix terms. Since the product of two $n \times n$ matrices can be computed in time $O(n^{2.37})$ [CW], the consistency problem for n sets can be computed in time $O(n^{2.37})$.

Finally we leave the following as an open question. Can the result in the present paper be generalized to "natural relations" with r arguments where $r > 2$? More specifically, can we prove or disprove that if μ is consistent on any subset of size less than or equal to $2r - 1$ of N then μ is consistent on N? A reasonable restriction of natural relations might be to those relations α on r sets in which for every i in $\{1, \ldots, r\}$ there exists b in Σ such that $v^{(i)} = b$ implies $\alpha(v) = 1$. Does this restriction make it easier to resolve the question?

References

[CW] D. Coppersmith and S. Winograd, *Matrix multiplication via arithmetic progressions*, J. Symb. Comput., 9, 3 (1990), pp. 251–280.

[FL] M. Fellows and M. Langston, *Nonconstructive tools for proving polynomial-time decidability*, JACM, 35 (1988), pp. 727–739.

[J] S. Jimbo, Private communication.

[H] E. Helly, *Über Mengen konvexer Körper mit gemeinschaftlichen Punkten*, Jahresber. Deutsch. Math. Verein 32 (1923), pp. 175–176.

[LN] N. Linial and N. Nisan, *Approximate inclusion-exclusion*, Proc. 32nd STOC, (1990), pp. 260–270.