Proof Pearl: The Marriage Theorem

Dongchen Jiang[1,2] and Tobias Nipkow[2]

[1] State Key Laboratory of Software Development Environment, Beihang University
[2] Institut für Informatik, Technische Universität München

Abstract. We describe two formal proofs of the finite version of Hall's Marriage Theorem performed with the proof assistant Isabelle/HOL, one by Halmos and Vaughan and one by Rado. The distinctive feature of our formalisation is that instead of sequences (often found in statements of this theorem) we employ indexed families, thus avoiding tedious reindexing of sequences.

1 Introduction

This paper describes two machine-checked proofs [6] of the marriage theorem, also known as Hall's theorem. The theorem was first proved by Hall in 1935 [4]. It provides a necessary and sufficient condition for the ability of selecting distinct elements from a collection of sets.

The standard statement of the theorem is phrased in terms of a finite sequence of sets $A_1, \ldots, A_n$. A sequence of elements $x_1, \ldots, x_n$ is called a *system of distinct representatives* (or *SDR* for short) for $A_1, \ldots, A_n$ iff

1. $x_i \in A_i$ for all $1 \leq i \leq n$, and
2. $x_i \neq x_j$ for all $1 \leq i, j \leq n$ such that $i \neq j$.

Now we can formulate the *Marriage Theorem*:

> A sequence of finite sets $A_1, \ldots, A_n$ (which need not be distinct) has an SDR iff the union of any $m \leq n$ of the A_i contains at least m elements.

The condition for the existence of an SDR is called the *marriage condition*. Note that we restrict ourselves to the finite version. Hall proved it for arbitrary sets A_i. Later work also relaxed the finiteness of the number of sets A_i — see Rado [8] for details.

We started our formalisation with "Proofs from the Book" [1], surely the ultimate reference for beautiful proofs, which also treats the finite version only. This lead to a lengthy proof that required additional concepts and lemmas about sequences. Then it dawned on us that sequences are a complication: the order of the sets is irrelevant, it only matters that there are only finitely many. Hence we replaced sequences by functions with a finite domain. This reduced the length of the proof by a factor of more than 2 to 140 lines. When we went back to the literature we discovered that, as far as the representation is concerned, we had ended up with the indexed families model by Everett and Whaples [3] (although they

call them sequences). Sequences are easy to understand, even for laymen, which may be why Aigner and Ziegler chose them. Hall himself had used sequences, too. However, sequences are inconvenient as a formal model, as we discovered in our first formalisation and as we shall detail later.

Alerted by a referee, we also formalised the proof by Rado [8]. This proof is already phrased in terms of a family of sets. Its formalisation was shorter again and required only 80 lines. Neither of our two proofs require additional definitions or lemmas beyond what is in the library.

After a review of related work in the next section, Section 3 gives a brief introduction to our logical language, explains the model and states the theorem formally. Section 4 presents our formalisation of the proof by Halmos and Vaughan and we compare our model with the one based on sequences. Section 5 presents our formalisation of Rado's proof, which is very close to Rado's text, and we compare it with the Mizar formalisation by Romanowicz and Grabowski.

2 Related Work

A number of different proofs of Hall's theorem have appeared in the literature [4, 3, 8], usually by some form of induction (for the finite case) that corresponds to an algorithm constructing an SDR. Aigner and Ziegler [1] follow the proof by Halmos and Vaughan [5], which is beautifully written, avoids almost all technical terminology (not just sequences), and takes a mere 13 lines.[3]

Romanowicz and Grabowski [9] formalised Rado's proof [8] (which is of the order of 15 lines) in the Mizar theorem prover. Their formalisation requires more than 1000 lines; for more details about their proof see Section 5.1. Initially we had ignored Rado's proof because of the length of its Mizar formalisation. But a referee pointed out that this was not Rado's fault: the referee had formalised it in his favourite theorem prover in 40 lines. This was the motivation for our own formalisation of this proof.

The proof on Wikipedia [10] also follows Halmos and Vaughan, but is phrased in terms of "collection of sets" (rather than sequences), and employs set-theoretic notation. If collections are interpreted as sets, the proof does not quite work and the statement is weaker than Hall's. If collections are interpreted as multisets, it works, but drags in multisets gratuitously.

3 Language and Formalisation

Our work was performed with the help of the theorem prover Isabelle/HOL [7], whose set theoretic language is close to that of standard mathematics, with a

[3] Aigner and Ziegler state that Halmos and Vaughan merely rediscovered the proof by Easterfield [2]. This is technically correct, but because Easterfield did not even realise himself that he had proved Hall's theorem, and did not phrase the lemma as abstractly as Hall did, we refer to Halmos and Vaughan for the proof.

few minor exceptions. Set difference is written as $X - Y$, and the image of a function f over a set X, i.e. $\{f\ x \mid x \in X\}$, is written as $f\ `\ X$.

HOL is a typed logic with type variables (α, β, etc), function types ($\alpha \Rightarrow \beta$) and set types ($\alpha\ set$). To express that x is of some type τ we write $x :: \tau$. Predicate *finite* expresses that a given set is finite. The fact that some function f is injective on some set X is written as $f\ inj\text{-}on\ X$. Updating a function f at argument x with new result value y is written $f(x := y)$.

Earlier on we stated that we would model the collection of sets A_i as a function with a finite domain, thus avoiding sequences with their irrelevant order. In HOL, we express this as a function $A :: \alpha \Rightarrow \beta\ set$ together with a finite set $I :: \alpha\ set$. We call I the *index set*. This model subsumes sequences (let I by the set $\{1, \ldots, n\}$) but is more flexible: we can remove arbitrary subsets from I without the need to renumber the result. Of course, mathematically speaking, renumbering is trivial, but in formal proofs it requires additional machinery and proof steps (see Section 4.1). The marriage condition (for A and I) can now be expressed as follows:

$$\forall J \subseteq I.\ |J| \leq |\bigcup_{i \in J} A\ i|$$

To avoid unnecessary index variables we will write $\bigcup_J A$ instead of $\bigcup_{i \in J} A\ i$.

An SDR (for A and I) is formalised as a function $R :: \alpha \Rightarrow \beta$ that returns the representative for each index and satisfies the following conditions:

1. $\forall i \in I.\ R\ i \in A\ i$, and
2. $R\ inj\text{-}on\ I$.

Thus the marriage theorem can be stated as follows in Isabelle:

assumes *finite* I **and** $\forall i \in I.\ finite(A\ i)$
shows $(\exists R.\ \forall i \in I.\ R\ i \in A\ i \wedge R\ inj\text{-}on\ I)$
$\longleftrightarrow\ (\forall J \subseteq I.\ |J| \leq |\bigcup_J A|)$

Necessity of the marriage condition is easy (and takes us 13 lines to formalise). Let R be an SDR for A and I, and let $J \subseteq I$. Hence $R\ `\ J \subseteq \bigcup_J A$ because $\forall i \in I.\ R\ i \in A\ i$. Thus $|R\ `\ J| \leq |\bigcup_J A|$. Because R is injective on $I \supseteq J$, we also have $|J| = |R\ `\ J|$. Combining the two cardinality facts yields the desired $|J| \leq |\bigcup_J A|$.

We will now present the formalisation of two proofs of sufficiency of the marriage condition: we assume the marriage condition and construct an SDR for A and I.

4 The Proof by Halmos and Vaughan

The proof is by induction on the finiteness of I: we may assume that the proposition holds for all proper subsets of I and we have to show it for I. This is slightly more convenient than a proof by induction on the cardinality of I. The

proposition to be proved must now include all assumptions of the theorem about I, including the marriage condition, and becomes

$$\begin{array}{c} \mathit{finite}\ I \longrightarrow (\forall i \in I.\ \mathit{finite}(A\ i)) \longrightarrow (\forall J \subseteq I.\ |J| \leq |\bigcup_J A|) \\ \longrightarrow (\exists R.\ \forall i \in I.\ R\ i \in A\ i \wedge R\ \mathit{inj\text{-}on}\ I) \end{array} \tag{1}$$

This is what we prove by induction on *finite* I. The case $I = \emptyset$ is trivial. Otherwise assume $I \neq \emptyset$ and make a case distinction on whether there is a *critical family* (as Aigner and Ziegler call it), i.e. a nonempty $K \subset I$ such that $|K| = |\bigcup_K A|$.

First we assume there is no critical family, i.e.

$$\forall K \subset I.\ K \neq \emptyset \longrightarrow |\bigcup_K A| \geq |K| + 1 \tag{2}$$

Because I is nonempty, we obtain an index $n \in I$. We also have $\forall i \in I.\ A\ i \neq \emptyset$ because an empty $A\ i$, $i \in I$, would imply, by the marriage condition, that $1 = |\{i\}| \leq |\bigcup_{\{i\}} A| = 0$, a contradiction. Thus we obtain some $x \in A\ n$, which we take as the representative for $A\ n$. Then we apply the induction hypothesis to the reduced problem A' and I':

$$A' = \lambda i.\ A\ i - \{x\} \qquad I' = I - \{n\}$$

From the assumption that A and I satisfy the marriage condition, it is easy to show that A' and I' still satisfy the marriage condition. Let J be an arbitrary subset of I'. Because we delete the same element x from each $A\ i$, $|\bigcup_J A'|$ can only be 1 smaller than $|\bigcup_J A|$, which, by (2), means that still $|J| \leq |\bigcup_J A'|$. Thus the induction hypothesis actually applies and yields an SDR R' for A' and I'. Because $x \notin A'\ i$ for $i \in I'$, it is easy to prove that the following R is indeed an SDR for A and I:

$$R\ =\ R'(n := x)$$

If there is a critical family, i.e. in the negation of case (2), we obtain a nonempty index set $K \subset I$ such that $|\bigcup_K A| < |K| + 1$. By the marriage condition we have $|K| \leq |\bigcup_K A|$. Together this implies that K is indeed a critical family:

$$|K|\ =\ |\bigcup_K A|$$

Because $K \subset I$, the induction hypothesis applies and we obtain an SDR R_1 for A and K. It remains to find an SDR for $I - K$. We simply remove $\bigcup_K A$ from each $A\ j$:

$$A'\ =\ \lambda j.\ A\ j - \bigcup_K A \qquad I' = I - K$$

As the cardinality of K equals the cardinality of $\bigcup_K A$, the marriage condition still holds for A' and I'. As also $I' \subset I$, the induction hypothesis applies and we obtain an SDR R_2 for A' and I'. Let

$$R\ =\ \lambda i.\ \text{if } i \in K \text{ then } R_1\ i \text{ else } R_2\ i$$

Because we excluded $\bigcup_K A$ from $A' \mathbin{`} I'$, it is clear that $A\ i \neq A'\ j$ for any $i \in K$ and $j \in I'$. Therefore R is an SDR for A and I.

This concludes the inductive proof of (1). The sufficiency of the marriage condition for the existence of an SDR follows trivially.

4.1 Sequences Versus Indexed Families

Both our proof and the one by Aigner and Ziegler are more detailed expositions of the proof by Halmos and Vaughan. The only difference is the underlying model: sequence versus indexed family of sets. Sequences are familiar to everybody and their finiteness is built in. The renumbering necessary in the critical-family case is easy for a human, but tedious for the machine. Not only do we need a function to remove a set of indices from a sequence (to allow us to apply the induction hypothesis to a subsequence) but we also need to compute the function that maps indices of the subsequence back to indices of the original sequence (to allow us to turn the SDR obtained for the subsequence into an SDR for the original sequence). And then we need to prove a number of tedious lemmas about how SDRs stay SDRs when they are lifted from the subsequence to the original sequence. All of this because we have introduced an irrelevant order.

5 Rado's Proof

The proof is by induction on the number of $A\ i$ that contain two or more elements. If $|A\ i| \geq 2$ for some i, then Rado shows that there is an $x \in A\ i$ such that $A(i := A\ i - \{x\})$ still satisfies the marriage condition. If all $A\ i$ are singletons, the marriage condition implies that the $A\ i$ directly yield the desired SDR. We merely present the key step of the induction: if $x_1, x_2 \in A\ i$ and $x_1 \neq x_2$, then $A(i := A\ i - \{x_1\})$ or $A(i := A\ i - \{x_2\})$ must satisfy the marriage condition.

Let A satisfy the marriage condition and let $x_1, x_2 \in A\ i$ be such that $x_1 \neq x_2$. For a contradiction, let $A_k = A(i := A\ i - \{x_k\})$ and assume that neither A_1 nor A_2 satisfy the marriage condition. Hence, for both k there is a $J'_k \subseteq I$ such that $|J'_k| > |\bigcup_{J'_k} A_k|$. Because A satisfies the marriage condition, $i \in J'_k$. Let $J_k = J'_k - \{i\}$. Hence $|J_k| \geq |(\bigcup_{J_k} A) \cup (A\ i - \{x_k\})|$. Let $U_k = \bigcup_{J_k} A$ and $U'_k = U_k \cup (A\ i - \{x_k\})$. This leads to the following contradiction:

$$
\begin{aligned}
|J_1| + |J_2| &\geq |U'_1| + |U'_2| \\
&= |U'_1 \cup U'_2| + |U'_1 \cap U'_2| \\
&= |U_1 \cup U_2 \cup A\ i| + |U'_1 \cap U'_2| \\
&\geq |U_1 \cup U_2 \cup A\ i| + |U_1 \cap U_2| \\
&\geq |\bigcup_{J_1 \cup J_2 \cup \{i\}} A| + |\bigcup_{J_1 \cap J_2} A| \\
&\geq |J_1 \cup J_2 \cup \{i\}| + |J_1 \cap J_2| \qquad (3) \\
&= |J_1 \cup J_2| + 1 + |J_1 \cap J_2| = |J_1| + |J_2| + 1
\end{aligned}
$$

Every step can be justified by set theory and side conditions like $i \notin J_1 \cup J_2$ and $x_1 \neq x_2$. Step (3) holds because A satisfies the marriage condition.

5.1 The Mizar Formalisation

Following Rado, Romanowicz and Grabowski [9] provided the first formal proof of the Marriage Theorem, in the Mizar prover. They used sequences, although Rado used indexed families. However, this should not have a large impact on their proof because there is no need for reindexing, the set I remains fixed throughout the proof. Nevertheless the Mizar proof is much longer than ours. Of course a comparison is difficult because the two provers differ, although at least both computer proofs are declarative and not unreadable proof scripts. The Mizar proof comes to 1600 lines, consisting of 8 definitions and 32 lemmas. Even if we exclude the definitions and lemmas in the Preliminaries and Union of Finite Sequences sections (which can be seen as general background knowledge), there still are more than 1200 lines and 22 lemmas left. But line counts are misleading: after sending those 1200 lines through gzip, 7.3 kB remain, as compared with 1.9 kB for our corresponding proof. This factor of 3.8 probably reflects differences such as automation, the library and sequences vs families.

Acknowledgement We are grateful to the anonymous referee for motivating us to formalise Rado's proof, too.

References

1. Aigner, M., Ziegler, G.M.: Proofs from the Book. Springer-Verlag (2001)
2. Easterfield, T.E.: A combinatorial algorithm. Journal London Mathematical Society 21, 219–226 (1946)
3. Everett, C.J., Whaples, G.: Represetations of sequences of sets. American Journal of Mathematics 71, 287–293 (1949)
4. Hall, P.: On representatives of subsets. Journal London Mathematical Society 10, 26–30 (1935)
5. Halmos, P.R., Vaughan, H.E.: The marriage problem. American Journal of Mathematics 72, 214–215 (1950)
6. Jiang, D., Nipkow, T.: Hall's marriage theorem. In: Klein, G., Nipkow, T., Paulson, L. (eds.) The Archive of Formal Proofs. `afp.sf.net/entries/Marriage.shtml` (Dec 2010), formal proof development.
7. Nipkow, T., Paulson, L., Wenzel, M.: Isabelle/HOL — A Proof Assistant for Higher-Order Logic, Lect. Notes in Comp. Sci., vol. 2283. Springer-Verlag (2002)
8. Rado, R.: Note on the transfinite case of Hall's Theorem on representatives. Journal London Mathematical Society 42, 321–324 (1967)
9. Romanowicz, E., Grabowski, A.: The Hall marriage theorem. Formalized Mathematics 12(3), 315–320 (2004)
10. Wikipedia: Hall's marriage theorem — wikipedia, the free encyclopedia (2011), `en.wikipedia.org/w/index.php?title=Hall%27s_marriage_theorem&oldid=419179777`, [Online; accessed 8-September-2011]