DENSITY, FORCING, AND THE COVERING PROBLEM

ADAM R. DAY AND JOSEPH S. MILLER

Abstract. We present a notion of forcing that can be used, in conjunction with other results, to show that there is a Martin-Löf random set X such that $X \not\geq_T \emptyset'$ and X computes every K-trivial set.

1. Introduction

In 1975, Solovay [10] constructed a noncomputable set whose initial segments have minimal prefix-free (Kolmogorov) complexity, in other words, whose initial segments are no more complex than those of the empty set. We call such sets *K-trivial* because K is usually used to denote prefix-free complexity. The class of K-trivial sets remained an obscure curiosity until the early 2000's, when a series of seminal papers established its importance in the study of Martin-Löf randomness, in part by giving several nontrivial characterizations. For example, $A \in 2^\omega$ is K-trivial if and only if

- It is *low for randomness*, i.e., every Martin-Löf random set is Martin-Löf random relative to A (Nies [8]).
- It is a *base for randomness*, i.e., there is an $X \geq_T A$ that is Martin-Löf random relative to A (Hirschfeldt, Nies and Stephan [7]).

Both of these properties were introduced and studied years before they were proved to be equivalent to K-triviality.

The characterizations above isolate the K-trivial sets as being weak when used to relativize Martin-Löf randomness. Another line of inquiry asks how the K-trivial sets interact with the (incomplete) Martin-Löf random sets in the Turing degrees. For example, in [5] the authors of the present paper proved that a set A is not K-trivial if and only if there is an incomplete Martin-Löf random set X such that $A \oplus X \geq_T \emptyset'$. This answered a question of Kučera. The proof came out of the study of Lebesgue density for Π^0_1 classes, which was started by Bienvenu, Hölzl, Miller and Nies [3, 4]. The present paper is a continuation of the effective study of Lebesgue density, with the purpose of answering another question about how the K-trivial sets interact with Martin-Löf random sets.

Hirschfeldt, Nies and Stephan [7] proved that if $A \in 2^\omega$ is c.e. and there is a Martin-Löf random $X \geq_T A$ that does not computes $\emptyset'$, then A is K-trivial. Stephan asked if this gives a characterization of the c.e. K-trivial sets. Each K-trivial is computable from a c.e. K-trivial, so this amounts to asking:

Date: December 11, 2014.

2010 *Mathematics Subject Classification.* Primary 03D32; Secondary 68Q30, 03D30.

The first author was supported by a Miller Research Fellowship in the Department of Mathematics at the University of California, Berkeley. The second author was supported by the National Science Foundation under grant DMS-1001847.

If A is K-trivial, is there a Martin-Löf random $X \not\geq_T \emptyset'$ that computes A?

More of the history of this question, known as the *covering problem*, is presented in a summary paper by the authors of this paper and Bienvenu, Greenberg, Kučera, Nies and Turetsky [1]. The present paper, combined with theorems of Bienvenu, Greenberg, Kučera, Nies and Turetsky [2], and Bienvenu, Hölzl, Miller and Nies [3, 4], gives a strong affirmative answer to the covering problem:

(a) There is a Martin-Löf random $X \not\geq_T \emptyset'$ that computes every K-trivial.

Furthermore, we get two interesting refinements:

(b) There is a Martin-Löf random $X <_T \emptyset'$ that computes every K-trivial.
(c) If $\langle A_n \colon n \in \omega\rangle$ is a countable sequence of non-K-trivial sets, then there is a Martin-Löf random $X \not\geq_T \emptyset'$ that computes every K-trivial but no A_n.

By (c), for example, there is an incomplete Martin-Löf random set X such that the Δ^0_2 sets computed by X are precisely the K-trivial sets. This X and Chaitin's Ω are Martin-Löf random sets that form an exact pair for the ideal of K-trivial sets (i.e., $A \leq_T X, \Omega$ if and only if A is a K-trivial set).

Our contribution to the solution of the covering problem comes out of a careful analysis of Lebesgue density for Π^0_1 classes. Let μ be the uniform measure on Cantor space. If $\tau \in 2^{<\omega}$ and P is a measurable set in Cantor space, then we define

$$\mu_\tau(P) = \frac{\mu(P \cap [\tau])}{\mu([\tau])}.$$

Given any measurable set P and $X \in 2^\omega$, we define $\rho(P \mid X) = \liminf_i \mu_{X\restriction i}(P)$. We call $X \in 2^\omega$ a *density-one point* if for every Π^0_1 class P it is the case that

$$X \in P \implies \rho(P \mid X) = 1.$$

If for every Π^0_1 class P we have $X \in P \implies \rho(P \mid X) > 0$, then X is called a *positive density point*. In Section 2, we present a notion of forcing that separates density-one from positive density on the Martin-Löf random sets. In other words, if X is a sufficiently generic set for this notion of forcing then:

(1) X is Martin-Löf random,
(2) X is not a density-one point,
(3) X is a positive density point.

Properties (1), (2) and (3) of generic sets will be established by Claims 2.1, 2.2 and 2.3, respectively. This forcing notion, in conjunction with the following two theorems, provides a solution to the covering problem.

Theorem 1.1 (Bienvenu, Hölzl, Miller and Nies [3, 4])**.** *If $X \geq_T \emptyset'$ and Martin-Löf random, then there exists a Π^0_1 class P such that $X \in P$ and $\rho(P \mid X) = 0$.*

We should note that Bienvenu, et al. prove Theorem 1.1 for density on the unit interval. However, the Cantor space version follows immediately from the proof given in [3, Theorem 20].

Theorem 1.2 (Bienvenu, Greenberg, Kučera, Nies and Turetsky [2])**.** *If $X \in 2^\omega$ is Martin-Löf random and not a density-one point, then X computes every K-trivial set.*

The original proof of Theorem 1.2, given in [2], involves several steps. A direct proof, though one relying on more of the theory of K-triviality, is given by Bienvenu, Hölzl, Miller and Nies [4].

By Theorem 1.1, properties (1) and (3) imply that X does not compute $\emptyset'$. By Theorem 1.2, properties (1) and (2) imply that X computes all K-trivial sets. This shows (a). In Claim 2.4, we show that if A is not K-trivial and X is sufficiently generic for our notion of forcing, then $X \not\geq_T A$. This gives us (c); in a sense, our forcing notion is perfectly tuned to constructing incomplete Martin-Löf random sets that compute all K-trivial sets. To show (b), we effectivize the forcing notion in Section 3 to show that there is a Δ^0_2 set X with properties (1), (2) and (3).

2. The forcing notion

Fix a nonempty Π^0_1 class $P \subseteq 2^\omega$ that contains only Martin-Löf random sets. Our forcing partial order $\mathbb{P}$ consists of conditions of the form $\langle \sigma, Q\rangle$, where

- $\sigma \in 2^{<\omega}$,
- $Q \subseteq P$ is a Π^0_1 class,
- $[\sigma] \cap Q \neq \emptyset$,
- There is a $\delta < 1/2$ such that $(\forall \rho \succcurlyeq \sigma)\ [\rho] \cap Q \neq \emptyset \implies \mu_\rho(Q) + \delta \geq \mu_\rho(P)$.

We say that $\langle \tau, R\rangle$ extends $\langle \sigma, Q\rangle$ if $\tau \succcurlyeq \sigma$ and $R \subseteq Q$. Let λ be the empty string. Note that $\langle \lambda, P\rangle \in \mathbb{P}$, with $\delta = 0$, so $\mathbb{P}$ is nonempty.

If $G \subseteq \mathbb{P}$ is a filter, let $X_G = \bigcup_{\langle\sigma,Q\rangle\in G} \sigma$. In general, $X_G \in 2^{\leq\omega}$. The following claim is trivial to verify and it establishes that if G is sufficiently generic, then X_G is infinite and, in fact, a Martin-Löf random set.

Claim 2.1.

1. If $\langle \sigma, Q\rangle \in \mathbb{P}$ and $\tau \succcurlyeq \sigma$ is such that $[\tau] \cap Q \neq \emptyset$, then $\langle \tau, Q\rangle \in \mathbb{P}$.
2. If $G \subseteq \mathbb{P}$ is sufficiently generic, then $X_G \in P$ (hence it is a Martin-Löf random set).

Claim 2.2. If $G \subseteq \mathbb{P}$ is sufficiently generic, then $\rho(P \mid X_G) \leq 1/2$, so X_G is not a density-one point.

Proof. Fix n. We will show that the conditions forcing

$$(\exists l \geq n)\ \mu_{X_G \upharpoonright l}(P) < 1/2 \tag{2.1}$$

are dense in $\mathbb{P}$. Let $\langle \sigma, Q\rangle$ be any condition and let δ witness that $\langle \sigma, Q\rangle \in \mathbb{P}$. Take m such that $2^{-m} < 1/2 - \delta$. Let Z be the left-most path of $[\sigma] \cap Q$. The set Z is Martin-Löf random and consequently contains arbitrarily long intervals of 1's. Take $\tau \succcurlyeq \sigma$ such that $\tau 1^m \prec Z$ and $|\tau| \geq n$. Because Z is the left-most path in Q it follows that $\mu_\tau(Q) \leq 2^{-m}$ and so

$$\mu_\tau(P) \leq \mu_\tau(Q) + \delta < 2^{-m} + \delta < 1/2.$$

Hence the condition $\langle \tau, Q\rangle$ extends $\langle \sigma, Q\rangle$ and forces (2.1). □

Claim 2.3. Let $S \subseteq 2^\omega$ be a Π^0_1 class and let $\langle \sigma, Q\rangle \in \mathbb{P}$. There is an $\varepsilon > 0$ and a condition $\langle \tau, R\rangle$ extending $\langle \sigma, Q\rangle$ such that either

- $[\tau] \cap S = \emptyset$, or
- If $X \in R$, then $\rho(S \mid X) \geq \varepsilon$.

Therefore, if $G \subseteq \mathbb{P}$ is sufficiently generic, then X_G is a positive density point.

Proof. If there is a $\tau \succcurlyeq \sigma$ such that $[\tau] \cap S = \emptyset$ and $[\tau] \cap Q \neq \emptyset$, then let $\langle \tau, Q\rangle$ be our condition.

Otherwise, it follows that $S \cap [\sigma] \supseteq Q \cap [\sigma]$. In this case let δ witness that $\langle \sigma, Q\rangle \in \mathbb{P}$. Take ε to be a rational greater than 0 and less than $\min\{1/2-\delta, \mu_\sigma(Q)\}$. (Note that $\mu_\sigma(Q) > 0$ because $[\sigma] \cap Q$ is a non-empty Π^0_1 class containing only Martin-Löf random sets.) Consider the Π^0_1 class

$$Q^\varepsilon_\sigma = \{X \in Q \cap [\sigma] \colon (\forall n \geq |\sigma|)\ \mu_{X\restriction n}(Q) \geq \varepsilon\}.$$

We will show that $\langle \sigma, Q^\varepsilon_\sigma\rangle$ is the required condition.

Let M be the set of minimal strings in $\{\rho \succcurlyeq \sigma \colon \mu_\rho(Q) < \varepsilon\}$. Then M is prefix-free and $Q^\varepsilon_\sigma = Q \cap [\sigma] \smallsetminus Q \cap [M]$. Summing over M gives us $\mu_\sigma(Q \cap [M]) < \varepsilon$. Hence $\mu_\sigma(Q^\varepsilon_\sigma) > \mu_\sigma(Q) - \varepsilon > 0$. This proves that $[\sigma] \cap Q^\varepsilon_\sigma \neq \emptyset$.

If $\tau \succcurlyeq \sigma$ and $[\tau] \cap Q^\varepsilon_\sigma \neq \emptyset$, we can use the same argument to show that $\mu_\tau(Q^\varepsilon_\sigma) > \mu_\tau(Q) - \varepsilon$. Because $[\tau] \cap Q \neq \emptyset$,

$$\mu_\tau(P) \leq \mu_\tau(Q) + \delta < \mu_\tau(Q^\varepsilon_\sigma) + \varepsilon + \delta.$$

Hence $\varepsilon + \delta < 1/2$ witnesses that $\langle \sigma, Q^\varepsilon_\sigma\rangle$ is a condition.

Note that if $X \in Q^\varepsilon_\sigma$, then $\rho(Q \mid X) \geq \varepsilon$. This implies that $\rho(S \mid X) \geq \varepsilon$ because $S \cap [\sigma] \supseteq Q \cap [\sigma]$, proving the claim. □

The following observation about the proof of Claim 2.3 will be used in the following section. This proof shows that if $\langle \sigma, Q\rangle$ is a condition witnessed by δ, and $\tau \succcurlyeq \sigma$ is such that $[\tau] \cap Q \neq \emptyset$, then for any $\varepsilon > 0$ with $\varepsilon < \min\{1/2 - \delta, \mu_\tau(Q)\}$ we have that $\langle \tau, Q^\varepsilon_\tau\rangle$ is a condition witnessed by $\delta + \varepsilon$.

A difference test is a Π^0_1 class R and a uniform sequence of c.e. open sets $\langle U_n \colon n \in \omega\rangle$ such that for all n, $\mu(U_n \cap R) \leq 2^{-n}$. A set X is captured by such a difference test if $X \in \bigcap_{n\in\omega} U_n \cap R$. We call a set X *difference random* if it is not captured by any difference test. Difference randomness was introduced by Franklin and Ng [6]. They showed that X is difference random if and only if X is Martin-Löf random and $X \not\geq_T \emptyset'$. Hence Claims 2.1 and 2.3 along with Theorem 1.1 establish that if $G \subseteq \mathbb{P}$ is sufficiently generic, then X_G is difference random.

Claim 2.4. Assume that $A \in 2^\omega$ is not K-trivial, $\langle \sigma, Q\rangle \in \mathbb{P}$, and Φ is a Turing functional. There is a $\tau \in 2^{<\omega}$ such that $\langle \tau, Q\rangle$ extends $\langle \sigma, Q\rangle$ and

$$(\forall X \in [\tau] \cap Q)[\, \Phi^X = A \implies X \text{ is not difference random}\,].$$

Therefore, if $G \subseteq \mathbb{P}$ is sufficiently generic relative to A, then X_G does not compute A.

Proof. If there is a $\rho \succcurlyeq \sigma$ and an n such that $\Phi^\rho(n)\downarrow \neq A(n)$ and $[\rho] \cap Q \neq \emptyset$, then take $\tau = \rho$.

Assume that no such ρ and n exist. Define $V_n = \{X \in 2^\omega \colon X \in U_n[\Phi^X]\}$, where $U_n[Z]$ is the nth level of the universal Martin-Löf test relative to Z. If $X \in V_n \cap [\sigma] \cap Q$, then because Φ^X is not incompatible with A, we have $X \in U_n[\Phi^X] \subseteq U_n[A]$. Hence $\mu(V_n \cap [\sigma] \cap Q) \leq \mu U_n(A) \leq 2^{-n}$. In other words, Q and $\langle V_n \cap [\sigma] \colon n \in \omega\rangle$ form a difference test.

Now assume that $X \in [\sigma] \cap Q$ and $\Phi^X = A$. Hirschfeldt, Nies and Stephan [7] showed that because A is not K-trivial, it is not a *base for randomness*. In other words, no set that is Martin-Löf random relative to A can compute A, so X is not random relative to A. Therefore, $X \in U_n[A] = U_n[\Phi^X]$ for all n. This shows that

$X \in \bigcap_{n\in\omega} V_n \cap [\sigma] \cap Q$, so X is not difference random. Hence the claim is satisfied by taking $\tau = \sigma$. □

3. Effectivizing the forcing

In this section we give a construction of a Δ^0_2 set with properties (1), (2) and (3). This construction is an effectivization of the forcing approach. It is conceptually similar to Sacks's construction of a Δ^0_2 minimal degree, which can be seen as an effectivization of Spector's minimal degree construction [9, 11]. The key problem is that $\emptyset'$ cannot determine which of the two outcomes in Claim 2.3 to force, i.e., whether to ensure that X is not a member of some Π^0_1 class S, or to ensure that it is a positive density point of S.

Theorem 3.1. *There is a Δ^0_2 set with properties* (1)*,* (2) *and* (3)*.*

Proof. Let $\langle S_e : e \in \omega\rangle$ enumerate all Π^0_1 classes. Using $\emptyset'$ as an oracle we will define a sequence of conditions $\langle p_i : i \in \omega\rangle$ in the partial order $\mathbb{P}$. If $p_i = \langle \tau, Q\rangle$ and $p_{i+1} = \langle \sigma, R\rangle$ we will ensure that $\sigma \succcurlyeq \tau$. We will construct a Martin-Löf random X that is the limit of the first coordinates of $\langle p_i : i \in \omega\rangle$. The most important aspect in which this construction differs from the previous section is that we will not require that $R \subseteq Q$. Essentially, this means that our oracle construction can make incorrect guesses as to which Π^0_1 classes to use, provided that a correct guess is made eventually. We will define p_s at stage s of the construction. Additionally at stage s we will define a_s to be a finite sequence of triples $\langle Q, \sigma, \varepsilon\rangle$ where Q is a Π^0_1 class, $\sigma \in 2^{<\omega}$, and ε is a rational. The sequence a_s will be used to recover information about previous stages in the construction. We let $l(a_s)$ be the length of the sequence a_s and we define partial functions Q, σ and ε such that if $e < l(a_s)$ then $\langle Q(s, e), \sigma(s, e), \varepsilon(s, e)\rangle$ is the eth element of a_s.

The idea behind the construction is as follows. Fix a Π^0_1 class S_e. First we will try to ensure that X is an element of S_e, and that for some $\sigma \preccurlyeq X$ and ε we have that $\mu_{X\restriction n}(S_e) \geq \varepsilon$ for all $n \geq |\sigma|$ (and hence $\rho(S_e \mid X) \geq \varepsilon$). At stage s, our plan is to achieve this objective by keeping X inside $Q(s, e)$. Our guess for σ will be $\sigma(s, e)$ and our guess for ε will be $\varepsilon(s, e)$. It is possible that our guess is incorrect. However, if this happens, then we will be able to ensure that X is not an element of S_e.

As in the proof of Claim 2.3, given Q a Π^0_1 class, σ a finite string and, ε a positive real number we define

$$Q^\varepsilon_\sigma = \{X \in Q \cap [\sigma] : (\forall n \geq |\sigma|)\ \mu_{X\restriction n}(Q) \geq \varepsilon\}.$$

We shall maintain the following construction invariants for all stages s:

(i) If $i < j < l(a_s)$, then $Q(s, j) \subseteq Q(s, i)^{\varepsilon(s,i)}_{\sigma(s,i)}$ and $\sigma(s, i) \preccurlyeq \sigma(s, j)$.
(ii) If $p_s = \langle \tau, R\rangle$ and $i < l(a_s)$ then $R \subseteq Q(s, i)^{\varepsilon(s,i)}_{\sigma(s,i)}$ and $\sigma(s, i) \preccurlyeq \tau$.

The construction is as follows. At stage 0, let $p_0 = \langle \lambda, P\rangle$ and let a_0 be the empty sequence. Our construction invariants hold trivially.

At stage $s + 1$, given $p_s = \langle \tau, Q\rangle$, we use $\emptyset'$ to find a condition $\langle \sigma, Q\rangle$ such that σ is a strict extension of τ, and $\mu_\sigma(P) < 1/2$. Claim 2.2 established that such a condition exists, and as the value of $\mu_\sigma(P)$ is computable in $\emptyset'$ we can simply search for a suitable σ. At this point we ask the following question. Does there exist $e < l(a_s)$ and ν such that

$$(\tau \preccurlyeq \nu \preccurlyeq \sigma) \wedge ([\tau] \cap S_e \neq \emptyset) \wedge (\mu_\nu(S_e) < \varepsilon(s,e))? \tag{3.1}$$

If not, then we define $p_{s+1} = \langle \sigma, Q_\sigma^{\varepsilon_{s+1}} \rangle$ where ε_{s+1} is chosen to be less than $\min\{2^{-s-3}, \mu_\sigma(Q)\}$. By the observation following the proof of Claim 2.3, this makes p_{s+1} a condition witnessed by $\delta = \sum_{i \leq s} 2^{-s-3} < 1/2$. Define a_{s+1} to be the sequence obtained by appending $\langle Q, \sigma, \varepsilon_{s+1} \rangle$ to the end of a_s. Note that the construction invariants are maintained.

If (3.1) holds, then we cannot make X extend σ without breaking some attempt to ensure that $\rho(S_e \mid X) \geq \varepsilon(s,e)$. In this case choose e and ν such that e is minimal for which (3.1) holds. Our construction invariants ensure that $Q \subseteq Q(s,e)^{\varepsilon(s,e)}_{\sigma(s,e)}$ and $\nu \succcurlyeq \tau \succcurlyeq \sigma(s,e)$. This implies that $\mu_\nu(Q(s,e)) \geq \varepsilon(s,e)$. Therefore there is some $\xi \succcurlyeq \nu$ such that $[\xi] \cap Q(s,e) \neq \emptyset$ and $[\xi] \cap S_e = \emptyset$. Define $p_{s+1} = \langle \xi, Q(s,e) \rangle$ and define $a_{s+1} = a_s \upharpoonright e$. Observe that construction invariant (i) is maintained because a_{s+1} is a subsequence of a_s, and construction invariant (ii) is maintained because construction invariant (i) held at stage s. This ends the construction.

Let $X = \bigcup \{\tau \colon (\exists s, Q)\ p_s = \langle \tau, Q \rangle\}$. To verify that X has the desired properties, we first show that $\lim_s l(a_s) = \infty$. Assume that for some s_0, for all $s \geq s_0$, $l(a_s) \geq e$. Assume at some stage $s_1 > s_0$, we have that $l(a_s) = e$. This can only occur because (3.1) held for e, and e was the least such value for which it held. Hence if $\langle \tau, Q \rangle = p_{s_1}$ then $[\tau] \cap S_e = \emptyset$. This implies that (3.1) will never again hold for e and hence for all $s > s_1$, $l(a_s) \geq e+1$.

If $l(a_{s+1}) > l(a_s)$, then condition (3.1) does not hold. Hence as $\lim_s l(a_s) = \infty$, for infinitely many stages s, condition (3.1) does not hold. This implies that X has infinite length, hence is a Martin-Löf random set, and $\rho(P \mid X) \leq 1/2$. Now assume that for some e, $X \in S_e$. Let s_0 be a stage such that for all $s \geq s_0$, $l(a_s) > e$. Let $\langle \tau, Q \rangle = p_{s_0}$. It must be that for any finite string ν such that $\tau \preccurlyeq \nu \prec X$, $\mu_\nu(S_e) > \varepsilon(s_0, e)$ because for all $s \geq s_0$ we know that condition (3.1) does not hold for e. Hence $\rho(S_e \mid X) > 0$. □

References

[1] Laurent Bienvenu, Adam R. Day, Noam Greenberg, Antonín Kučera, Joseph S. Miller, André Nies, and Dan Turetsky. Computing K-trivial sets by incomplete random sets. *Bull. Symb. Log.*, 20(1):80–90, 2014.

[2] Laurent Bienvenu, Noam Greenberg, Antonín Kučera, André Nies, and Daniel Turetsky. Coherent randomness tests and computing the K-trivial sets. Submitted.

[3] Laurent Bienvenu, Rupert Hölzl, Joseph S. Miller, and André Nies. The Denjoy alternative for computable functions. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012)*, volume 14 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 543–554, Dagstuhl, Germany, 2012. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[4] Laurent Bienvenu, Rupert Hölzl, Joseph S. Miller, and André Nies. Denjoy, Demuth and density. *J. Math. Log.*, 14(1):1450004, 35, 2014.

[5] Adam R. Day and Joseph S. Miller. Cupping with random sets. *Proc. Amer. Math. Soc.*, 142(8):2871–2879, 2014.

[6] Johanna N. Y. Franklin and Keng Meng Ng. Difference randomness. *Proc. Amer. Math. Soc.*, 139(1):345–360, 2011.

[7] Denis R. Hirschfeldt, André Nies, and Frank Stephan. Using random sets as oracles. *J. Lond. Math. Soc. (2)*, 75(3):610–622, 2007.

[8] André Nies. Lowness properties and randomness. *Adv. Math.*, 197(1):274–305, 2005.

[9] Gerald E. Sacks. *Degrees of Unsolvability.* Princeton University Press, 1963.

[10] Robert M. Solovay. Draft of paper (or series of papers) on Chaitin's work. May 1975. Unpublished notes, 215 pages.
[11] Clifford Spector. On degrees of recursive unsolvability. *Ann. of Math. (2)*, 64:581–592, 1956.

School of Mathematics, Statistics and Operations Research, Victoria University of Wellington, Wellington, New Zealand
E-mail address: `adam.day@vuw.ac.nz`

Department of Mathematics, University of Wisconsin, Madison, WI 53706-1388, USA
E-mail address: `jmiller@math.wisc.edu`