

Stochastic Methods for Minimal Surfaces

Robert W. Neel

ABSTRACT. We survey recent applications of stochastic analysis to the theory of minimal surfaces in \mathbb{R}^3 , with an eye toward intuitive explanations of the stochastic techniques involved. We begin by giving several points of view on Brownian motion on a Riemannian manifold and explaining how minimality manifests itself in terms of this Brownian motion. We then explain how the parabolicity and area growth of minimal ends have been studied using universal superharmonic functions and outline an alternative approach, yielding stronger results, based on Brownian motion (following [13]). In a different direction, we recall the idea of coupling two Brownian motions and its use in studying harmonic functions and spectral gaps on manifolds. Next, we describe a way of coupling Brownian motion on two minimal surfaces (possibly the same surface) and describe how this coupling can be used to study properties related to the intersection of two minimal surfaces and to study the non-existence of bounded, non-constant harmonic functions on properly embedded minimal surfaces of bounded curvature (following [14]). We also discuss an application to proving that non-planar minimal graphs are parabolic (following [12]).

1. Introduction

These notes are based on a seventy-five minute talk given by the author on July 1, 2010 at the Santaló Summer School on Geometric Analysis, held at the University of Granada, Spain. The aim of the talk was to briefly indicate how minimal submanifolds, and specifically minimal surfaces in \mathbb{R}^3 , have their minimality reflected in the behavior of Brownian motion on the submanifold, and then to illustrate how these properties of Brownian motion can be used to prove geometric theorems for minimal surfaces. That is, the theorems themselves generally have no random aspect (or can be easily translated into non-random language); they are purely theorems in geometric analysis. Probability is introduced solely as a tool. Many of the theorems proven in this way do not have known non-stochastic proofs. Indeed, the main aim of this program of research is contribute new results, as well as new intuition, to minimal surface theory, not merely to re-prove existing results.

The examples illustrating this use of stochastic methods fall into two categories, those looking at a single Brownian motion in order to prove parabolicity or an extrinsic form of quadratic area growth, and those using a pair of coupled Brownian motions to address intersection theorems, Liouville theorems, and parabolicity. In

2010 *Mathematics Subject Classification*. Primary 53A10; Secondary 58J65 60H30.

©0000 (copyright holder)

all cases, the idea was to describe the results and techniques and provide intuition for why stochastic methods are a natural approach for certain questions and for how they work, as opposed to providing rigorous proofs, which can be found elsewhere. This approach was motivated by the fact that the occasion was a single talk at a summer school and by the fact that stochastic analysis falls well outside the standard background for people working in geometric analysis, making discussion of probabilistic technicalities unhelpful.

It seems appropriate to follow that philosophy in these notes as well. Namely, they provide an opportunity to elaborate somewhat beyond what can be said in seventy-five minutes, and we shall choose to elaborate by way of providing additional background and examples, referring the reader to the literature for detailed proofs. Thus, these notes (like the talk they are based upon) are essentially a survey of recent applications of stochastic analysis to the geometry of minimal surfaces.

I would like to thank the organizers, especially Joaquín Pérez, for the invitation to speak at the School Lluís Santaló. I would also like to thank the Spanish Royal Mathematical Society and the International University Menéndez-Pelayo for organizing and funding this annual event. In addition, I am grateful to Sarah Gíoe for creating the figures for both the talk and these notes. Mathematically, the previous work of the author referenced in these notes owes a debt to many people, who are mentioned in the original papers. However, special mention should be made of Dan Stroock, Bill Meeks, and Ioannis Karatzas.

2. Basic notions

2.1. Brownian motion on manifolds. It has long been known that there is a relationship between stochastic processes and second-order operators, specifically, between the processes and the elliptic and parabolic PDEs associated to the operators. The fundamental example of this is the relationship between Brownian motion and the Laplacian, specifically between Brownian motion and harmonic functions and the heat equation. The same relationship holds for operators and processes on manifolds. For a thorough introduction to stochastic techniques on manifolds, primarily those involving Brownian motion, the reader is referred to two books, by Stroock [16] and Hsu [5]. These books form the background for the following discussion of Brownian motion on a Riemannian manifold.

Let M be a (smooth) complete Riemannian manifold of dimension n . There are several ways to introduce Brownian motion on M , and for our purposes, we discuss three of them. First, it's the continuous version of an isotropic random walk on a (Riemannian) manifold. This is analogous to the well-known fact that Brownian motion on \mathbb{R} is the limit of rescaled simple symmetric random walks. To imitate this construction on a manifold, first take a starting point x_0 . Next, choose a direction from the unit tangent bundle at x_0 according to the uniform probability measure (given by the natural identification of the unit tangent bundle with the sphere of the appropriate dimension) and follow the corresponding geodesic distance 1 in time 1 (hence with speed 1). This determines a function (or a path) $x_t : [0, 1] \rightarrow M$. Now choose a direction from the unit tangent bundle at x_1 according to the uniform probability measure (and independent of the earlier choice of direction at x_0), and again follow the corresponding geodesic distance 1 in time 1. This extends the path x_t to the interval $t \in [0, 2]$. Iterating this process gives a continuous path $x_t : [0, \infty) \rightarrow M$, which is piecewise geodesic. More to the point,

since the path is random, this gives a probability measure on the space of continuous paths on M (equipped with the topology of uniform convergence on compacts and the corresponding Borel σ -algebra).

This measure is a coarse approximation to (the distribution of) Brownian motion on M . To get a finer approximation, we perform a similar procedure, except that we follow each geodesic distance $1/2$ in time $1/4$ (hence with speed 2). For the n th order approximation, we follow each geodesic distance $1/n$ in time $1/n^2$ (hence with speed n). In the limit, the size of the steps goes to zero and the process converges (weakly, in the space of measures on continuous paths on M) to Brownian motion on M . Actually, this statement requires a caveat. As n increases, there might be paths which travel farther and farther from x_0 in some fixed amount of time, so that in the limit they explode. To say that a path explodes means that it exits every compact subset of M in finite time, and thus is only defined up to some finite explosion time. Depending on M , it might happen that Brownian motion explodes in finite time with positive probability. We can allow for this by letting ζ be the explosion time; note that ζ depends on the path and may be infinite, if the path doesn't explode. Then Brownian motion on M is a random variable taking values in the space of continuous, possibly exploding paths $x_t : [0, \zeta) \rightarrow M$, having the right distribution on the space of such paths (we call this distribution characterizing Brownian motion Weiner measure).

If $\zeta = \infty$ almost surely, then we say that M is stochastically complete. Otherwise, we say that M is stochastically incomplete. This mirrors the more usual notion of geodesic completeness, geodesic completeness meaning that all geodesics can be extended for all time and stochastic completeness meaning that (almost) all Brownian paths can be extended for all time. Neither notion implies the other. By putting a metric on \mathbb{R}^2 such that the curvature goes to $-\infty$ sufficiently quickly, one can produce a (geodesically) complete but stochastically incomplete surface; see Section 8.4.2 of [16] for the details. The complimentary result is easier; just consider $\mathbb{R}^2 \setminus (0, 0)$. This is obviously (geodesically) incomplete. However, planar Brownian motion never strikes the origin (almost surely), and thus $\mathbb{R}^2 \setminus (0, 0)$ is stochastically complete. We will have more to say about stochastic completeness later, but for now we continue our introduction to Brownian motion.

The proceeding point of view on Brownian motion is not especially well-suited to computations or to proving theorems, but it is probably the most intuitively appealing. We demonstrate this point with an example which is not specific to minimal surfaces but which illustrates the local interaction between Brownian motion and curvature. Consider a point $x \in M$, and a small ball of radius r around x . If we want to understand how quickly Brownian motion started at x leaves this small ball, we can develop an intuition as follows. We consider a “two step” random walk approximation. If we first take a small step in a random direction from x and then take a second step in a random direction (as described above), how far are we from x ? The component of the second step which is parallel to the first contributes in an obvious and uninteresting way, so for simplicity we assume that the second step is perpendicular to the first. If the steps are small, then they approximately span a (tangent) plane, and the distance from x to the final point depends on the sectional curvature of this plane. In particular, zero sectional curvature means that the distance should be (in the small step approximation) given by the usual Pythagorean theorem, while if the sectional curvature is positive the

distance will be smaller than this, and larger if the curvature is negative (this is just the law of cosines for a Riemannian manifold). Averaging over the possible second steps shows that the average distance from x if the first step lies in a given direction should be given by the Ricci curvature in this direction (here we think of the Ricci curvature as a quadratic form on tangent vectors). Continuing, averaging over the possible first steps then indicates that the distance from x achieved by this two step approximation should be given by the scalar curvature (the trace of the Ricci tensor) at x . Expressing this intuition in terms of Brownian motion on M , we expect that how quickly Brownian motion from x leaves a small ball around x is given by the scalar curvature at x .

One way of seeing that this intuition is accurate is as follows. Let $B_x(r)$ be the ball of radius r around x , and let $f(r)$ be the expected first exit time of Brownian motion (started at x) from $B_x(r)$. Also, let S_x be the scalar curvature at x . Then

$$f(r) = \frac{1}{n}r^2 + \frac{S_x}{6n^2(n+2)}r^4 + O(r^5);$$

see Section 9.4.1 of [16] for a derivation. Obviously, our intuitive argument is not refined enough to guess the constants involved in this expansion, etc., but at some level it explains why there should be a relationship between exit times and scalar curvature. If one wishes to push this further, it seems that when the Brownian motion leaves, it should be more likely to have done so via a first step which is in a direction of relatively large Ricci curvature. This turns out to be the case. That is, the distribution of where the Brownian motion exits $B_x(r)$ encodes the traceless Ricci tensor, as is also described in Section 9.4.1 of [16].

A more analytic point of view on Brownian motion on M is provided by the observation that it is the (unique) continuous Markov process with transition density given by the heat kernel on M . This makes the connection with PDEs, most obviously the heat equation on M , clearer. (We note that, in keeping with the conventions of probability, we take the heat equation to be $\partial_t u_t(x) = (\Delta/2)u_t(x)$.) It also allows us to give a first demonstration of how Brownian motion is related to global properties of M . Suppose M is compact. Then Δ has a discrete spectrum which we write as $-\lambda_m$ with $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$. Estimating λ_1 (often referred to as the spectral gap) in terms of the geometry of M is a standard problem. Expanding the heat kernel in terms of the eigenfunctions of Δ shows that, for large time, it converges to a constant exponentially quickly and that the rate of convergence is given, up to a constant, by λ_1 . From a stochastic point of view, this means that, in the large time limit, Brownian motion distributes itself evenly over M , or equivalently, that Brownian motion “forgets” its starting point and converges to an equilibrium measure. The rate at which Brownian motion (or, more generally, a Markov process) converges to equilibrium can be studied via a technique called coupling, and the upshot of the proceeding is that the spectral gap of M can thus be studied by coupling Brownian motions on M . We will have more to say about this below, but for the moment we simply remark that this gives, as claimed, a relationship between Brownian motion and the global geometry of M .

Of course, stochastic completeness also arises in this context. We see that M is stochastically complete if and only if the heat kernel is conservative, namely, it integrates to 1 at any time. Naturally, if the heat kernel fails to integrate to 1, the “lost mass” corresponds to Brownian paths that have already exploded.

Our third and final description of Brownian motion on a Riemannian manifold is the one most closely aligned with the techniques of stochastic analysis, and not coincidentally, the one which is least intuitive from the point of view of standard geometric analysis. Brownian motion on M is the solution to a martingale problem, in particular, the martingale problem associated to half the Laplacian. This means that, if B_t is Brownian motion on M , then, for any smooth, compactly supported function f , we have that $f(B_t) - f(x_0) - \int_0^t \frac{1}{2} \Delta f(B_s) ds$ is a martingale. This requires explaining what a martingale is. In general, it means a process which is “conditionally constant,” in the sense that the expected value of the process, conditioned on its history up to time t , is just its value at time t . In the present context, because both f and B_t are continuous, we also know that $m_t = f(B_t) - f(x_0) - \int_0^t \frac{1}{2} \Delta f(B_s) ds$ is continuous. In this case, we have a theorem, originally due to Lévy, stating that a continuous, real-valued martingale is a time-changed (real-valued) Brownian motion. A time-changed process is one in which the time parameter is transformed by a continuous, non-decreasing map which depends on the path. To be more concrete, in the case of the martingale $m_t = f(B_t) - f(x_0) - \int_0^t \frac{1}{2} \Delta f(B_s) ds$, the time change is given by integrating $|\nabla f|^2$ along the Brownian paths. This means that m_t evolves on \mathbb{R} like a Brownian motion, except that time is infinitesimally dilated by a factor of $|\nabla f|^2$, so that m_t evolves faster or slower than a Brownian motion, according to whether $|\nabla f|^2$ is larger or smaller than 1. That the size of the gradient of f should govern how quickly m_t changes is unsurprising. Further, it is important to note that, if the integral of $|\nabla f|^2$ along a path stays bounded, then m_t will follow the path of a Brownian motion run until some finite time. In particular, along such paths m_t will converge as $t \rightarrow \infty$ almost surely.

From this description, some standard analytic facts can be recovered and expanded on. For example, if h is a harmonic function on M then $h(x_0)$ is given by the expectation of $h(B_\sigma)$ where B_t is Brownian motion started at x_0 and σ is a (bounded) stopping time. (Intuitively, a stopping time is the time of the occurrence of an event with the property that whether it has occurred by time t only requires knowledge of the history up to time t . For example, the first hitting time of a closed set is a stopping time, but not the last exit time of a set.) Two familiar examples of this sort of representation of a harmonic function are integrating h against the heat kernel (in which case $\sigma = t$ for some t) and integrating h against harmonic measure (in which case σ is the first hitting time of the boundary of the domain in question and we get around the requirement that σ be bounded by using the dominated convergence theorem). As we will see, one of the advantages of stochastic methods is that they allow other choices of stopping time beyond these two.

2.2. Minimal immersions. We start in a general context. Suppose $i : M \rightarrow N$ is an isometric immersion. We identify M with its image, self-intersections notwithstanding. We now give three closely related ways to define what it means for M to be minimal (in N). Geometrically, M is minimal if its mean curvature vector vanishes identically. In the case when $N = \mathbb{R}^m$, we have the analytic characterization that M is minimal if the restrictions of the coordinate functions to M are harmonic. From a stochastic point of view, M is minimal if Brownian motion on M (viewed as a process on N under the immersion) is an N -martingale. In the case when $N = \mathbb{R}^m$, this means that each coordinate process, that is, $x_{i,t} = x_i(B_t)$, is a martingale. In general, a process is an N -martingale if it satisfies an infinitesimal version of this condition. Since we will deal exclusively with the case when

$N = \mathbb{R}^m$ (and almost exclusively with $m = 3$, for that matter), we won't discuss N -martingales here. Instead, we refer the interested reader to Chapter 2 of [5].

One can see some elementary things fairly easily even at this level of generality. For example, both Stroock and Hsu, in their books just mentioned, observe that Brownian motion on a proper minimal immersion into \mathbb{R}^m never explodes, or equivalently, such manifolds are stochastically complete. In addition, Stroock (see Theorem 5.23 of [16]) takes a moment to prove that, if M is a minimal immersion into Euclidean space with $\dim(M) \geq 3$, then M is transient.

Nonetheless, it appears (to the author, at least) that relatively little has been done to exploit stochastic analysis as a tool in minimal submanifold geometry. In what follows we give two types of examples of stochastic methods in classical minimal surface theory:

- We establish parabolicity and (extrinsic) quadratic area growth of minimal surfaces-with-boundary constrained to lie in certain subsets of \mathbb{R}^3 . This is possibly the less interesting of the two, but it is also the less technical and thus makes a better introductory example of stochastic methods in minimal surface theory.
- We describe a coupling of Brownian motions on two minimal surfaces and explain how it can be used to study halfspace theorems, a conjecture of Sullivan (about a Liouville theorem for properly embedded minimal surfaces), and parabolicity of minimal graphs. Again, we gloss over the technicalities here and concentrate on the driving ideas.

3. Parabolicity and area growth

3.1. Underlying ideas. We begin by defining the properties of interest.

DEFINITION 3.1. A surface M with non-empty boundary ∂M is *parabolic* if any of the following equivalent conditions hold:

- Any bounded harmonic function on M is determined by its boundary values on ∂M .
- There exists a point x in the interior of M such that Brownian motion started from x hits ∂M almost surely.
- Brownian motion started from any interior point hits ∂M almost surely.

Note that we are reserving the term parabolic for surfaces with non-empty boundary. In particular, if we assert that a surface is parabolic, this includes the assertion that it has non-empty boundary. This is at odds with the usual use of the term in the theory of Riemann surfaces, but it is convenient for us here. The first item in this definition makes the geometric interest in parabolicity clear, since it affects the potential theory of the surface and thus its conformal structure. The other two items give a natural stochastic formulation of parabolicity and help motivate the use of stochastic techniques.

The other property of interest is quadratic area growth with respect to the extrinsic distance (that is, the \mathbb{R}^3 -distance). We will refer to this simply as quadratic area growth, even though this phrase generally refers to area growth with respect to intrinsic distance. Again, this usage is more convenient for us, as well as fairly common within the minimal surface literature. We now give a precise definition.

DEFINITION 3.2. We say a surface $M \subset \mathbb{R}^3$ has *quadratic area growth* if, for some $C > 0$ and $A > 0$, we have

$$\text{Area} \left(M \cap \left\{ \sqrt{x_1^2 + x_2^2 + x_3^2} < R \right\} \right) \leq CR^2$$

for all $R > A$.

Unlike in the case of parabolicity, there is no immediately obvious reformulation of this geometric notion as a stochastic notion. Thus it is a welcome surprise that we are able to use stochastic methods to establish quadratic area growth in the result that follows.

The goal is to prove that certain subsets of \mathbb{R}^3 have the property that any minimal surface-with-boundary (with some other mild condition, such as being stochastically complete or properly immersed) which is contained in the subset is parabolic and/or has quadratic area growth. The larger the subsets one can establish such a property for, the better.

Before going on, we take a moment to discuss a well-known example that illustrates the necessity of assuming stochastic completeness. In [11], Nadirashvili constructed a minimal conformal immersion of the (open) unit disk into a ball in \mathbb{R}^3 such that the induced metric on the surface is (geodesically) complete. With this in mind, we let M be a minimal surface (without boundary) contained in a ball. We don't assume that M is complete, although that might be the most interesting case. Suppose without loss of generality that M is contained in the ball around the origin of radius $R > 0$. We let ρ be the distance from the origin (in \mathbb{R}^3). Further, let B_t be Brownian motion on M (defined up to a possible explosion time ζ) started at a point on M distance ρ_0 from the origin. Then if we write, as usual, $\rho_t = \rho(B_t)$, an easy application of Ito's formula (see Section 5.2.2 of [16]) shows that

$$(3.1) \quad R^2 \geq \mathbb{E} [\rho_{t \wedge \zeta}^2] = \rho_0^2 + 2\mathbb{E} [t \wedge \zeta].$$

Here the expectation is with respect to B_t and the first inequality comes from the fact that ρ is bounded by R on M , by assumption. Letting $t \rightarrow \infty$, we see that

$$\mathbb{E} [\zeta] \leq \frac{R^2 - \rho_0^2}{2}.$$

This more than shows that ζ is almost surely finite and thus that M is stochastically incomplete; in fact we have an explicit estimate on the expectation of ζ . In particular, the example of Nadirashvili is stochastically incomplete (despite being complete). Stated differently, we see that any stochastically complete minimal surface cannot be contained in any ball. Since any compact Riemannian manifold is stochastically complete, and any compact surface in \mathbb{R}^3 is contained in some ball, this gives the familiar fact that there are no compact minimal surfaces.

The example of Nadirashvili is obviously not parabolic (it has no boundary), and moreover, it has the conformal structure of the disk. This shows that, unless we also assume stochastic completeness (or something even stronger like being properly immersed), there is no hope of finding a subset of \mathbb{R}^3 with the property that any complete minimal surface-with-boundary in this subset is parabolic. We also note, perhaps unnecessarily, that the example of Nadirashvili does not have quadratic area growth.

To continue with this line of thought, we now consider a minimal surface M which is stochastically complete, but which is also allowed to possibly have boundary (so M is actually a minimal surface-with-boundary). In this case, $\zeta \equiv \infty$, but we also need to stop our Brownian motion at the boundary. We let ξ be the hitting time of the boundary, so that Brownian motion on M is stopped at ξ . If we assume that M is contained in the ball of radius R around the origin, then the computations above still hold, with ζ replaced throughout by ξ . In particular,

$$\mathbb{E}[\xi] \leq \frac{R^2 - \rho_0^2}{2}.$$

Thus M is parabolic, and we've shown that any ball has the property that a stochastically complete minimal surface-with-boundary contained in the ball must be parabolic. On the other hand, a ball is, essentially, the smallest (and therefore least interesting) possible subset of \mathbb{R}^3 for which one could try to establish this property. Thus our goal is to do better and find larger subsets of \mathbb{R}^3 with this property.

Before moving on, we mention the complimentary project of finding the smallest region(s) where it is possible to properly immerse a minimal surface with arbitrary conformal structure. For recent results in this direction, we refer the interested reader to work of Alarcón and López [1].

3.2. Universal superharmonic functions. One way to address parabolicity and quadratic area growth is through the use of universal superharmonic functions. Indeed, this is the technique of [3] and [8]. In this sub-section we explain this approach and the results which have previously been obtained by it.

DEFINITION 3.3. Let U be a non-empty, open subset of \mathbb{R}^3 . A function $f : U \rightarrow \mathbb{R}$ is a *universal superharmonic function on U* if the restriction of f to any minimal surface (possibly with boundary) in U is superharmonic (that is, its Laplacian is everywhere non-positive).

Simply having a universal superharmonic function on some region isn't enough. For example, any of the coordinate functions is a universal superharmonic function on all of \mathbb{R}^3 (and even a "universal harmonic function"), but this certainly doesn't imply parabolicity or quadratic area growth for minimal surfaces-with-boundary in \mathbb{R}^3 . The following elementary lemma indicates what additional properties one looks for.

LEMMA 3.4. *If there is a universal superharmonic function on U which is positive and proper (relative to U), then any stochastically complete minimal surface-with-boundary contained in U is parabolic (and in particular, has non-empty boundary).*

We briefly sketch a stochastic proof of this lemma. Hopefully this will clarify the interaction between this analytic approach and direct consideration of Brownian motion on the surface. Let M be a minimal surface and f a universal superharmonic function as described in the lemma. Let B_t be Brownian motion on M , started from any point. Because M is stochastically complete, almost every Brownian path either hits the boundary in finite time or continues forever. Our goal is to show that, almost surely, B_t hits the boundary. The fundamental observation is that Brownian motion composed with a superharmonic function gives a supermartingale. Further, because f is positive, we see that $f(B_t)$ is a supermartingale which is bounded from below. A supermartingale bounded from below almost surely converges (this is a

standard result in stochastic analysis), and thus $f(B_t)$ converges. Since f is proper relative to U , this is only possible if B_t eventually enters a bounded subset of \mathbb{R}^3 (which depends on the particular path) and never leaves, almost surely. Since we previously saw that if B_t runs forever it leaves every bounded subset of \mathbb{R}^3 , this is only possible if B_t almost surely hits the boundary. This establishes the lemma.

As a result of this lemma, one is lead to look for positive, proper universal superharmonic functions on various subsets of \mathbb{R}^3 (which should be as large as possible). There are two such examples which play a role in the literature. Let $r = \sqrt{x_1^2 + x_2^2}$. Both examples start from the basic relationship, valid for any minimal surface M ,

$$(3.2) \quad |\Delta_M \log r| \leq \frac{|\nabla_M x_3|^2}{r^2} \quad \text{on } M \setminus \{r = 0\}$$

For convenience, we provide a quick derivation of this inequality. The chain rule shows that, for any surface, not just minimal ones,

$$\Delta_M \log r = \frac{1}{r^2} \left(-2|\nabla_M r|^2 + |\nabla_M x_1|^2 + |\nabla_M x_2|^2 + x_1 \Delta_M x_1 + x_2 \Delta_M x_2 \right).$$

Because M is minimal, $\Delta_M x_1 = \Delta_M x_2 = 0$, and so we see that $\Delta_M \log r$ depends only on first-order quantities at a point. Next, note that $\{\nabla_{\mathbb{R}^3} x_1, \nabla_{\mathbb{R}^3} x_2, \nabla_{\mathbb{R}^3} x_3\}$ and $\{\nabla_{\mathbb{R}^3} r, r \nabla_{\mathbb{R}^3} \theta, \nabla_{\mathbb{R}^3} x_3\}$ are both orthonormal frames for \mathbb{R}^3 (at least when $r \neq 0$ for the second). It follows that

$$\begin{aligned} |\nabla_M x_1|^2 + |\nabla_M x_2|^2 + |\nabla_M x_3|^2 &= 2 \\ \text{and } |\nabla_M r|^2 + |r \nabla_M \theta|^2 + |\nabla_M x_3|^2 &= 2. \end{aligned}$$

This gives

$$\Delta_M \log r = \frac{1}{r^2} \left(|\nabla_M x_3|^2 + 2|r \nabla_M \theta|^2 - 2 \right).$$

From here, the upper bound $\Delta_M \log r \leq |\nabla_M x_3|^2 / r^2$ follows from $|r \nabla_M \theta|^2 \leq 1$. To get the lower bound, that is, $\Delta_M \log r \geq -|\nabla_M x_3|^2 / r^2$, note that it is enough to show that $2|\nabla_M x_3|^2 + 2|r \nabla_M \theta|^2 - 2 \geq 0$. Using the above formula for the sum of the squared norms of an orthonormal frame, this in turn is equivalent to $2 - 2|\nabla_M r|^2 \geq 0$, which follows from $|\nabla_M r|^2 \leq 1$. This establishes Equation (3.2).

Using this basic estimate, one can compute that:

- for any $c > 0$, we can find $\alpha > 0$ so that $\log r - x_3^2 + \alpha$ is a positive, proper universal superharmonic function on $\{r > 1/\sqrt{2} \text{ and } |x_3| < c\}$; and
- for any $c > 0$, we can find $\alpha > 0$ so that $\log r - x_3 \arctan x_3 + \frac{1}{2} \log(x_3^2 + 1)$ is a positive, proper universal superharmonic function on the region $\{|x_3| < c \log r \text{ and } r > \alpha\}$.

In light of Lemma 3.4, this shows that stochastically complete minimal surfaces-with-boundary in either region are parabolic.

At this point, we should explain something of the context of these results. The most interesting source of minimal surfaces-with-boundary contained in these rotationally symmetric regions is ends of complete minimal surfaces (often properly immersed). Ends (or more accurately, representatives of ends, but we will generally abuse terminology by conflating ends with their representatives) contained in regions of the first sort in the above list are called planar ends, while ends contained in regions of the second sort are called catenoidal ends, both for obvious reasons. Thus we see that planar and catenoidal ends are parabolic. All planar ends are

catenoidal ends (perhaps after adjusting the representative of the end), and thus the second universal superharmonic function above supersedes the first, at least in the context of ends. On the other hand, all of the computations are easier in the first case, and thus we prefer to discuss them both.

To continue, it turns out that the gradient and Laplacian of each of these universal superharmonic functions are geometrically significant. In particular, an argument which rests on computing these quantities and using integration by parts allows one to show that both planar and catenoidal ends have quadratic area growth (we refer the interested reader to [3] or Section 6 of [8] for details).

3.3. The question. Meeks has asked whether there exists a positive, proper universal superharmonic function on the exterior of a sufficiently large cone, that is, the region $\{|x_3| < cr\}$ for small enough c , minus a ball around the origin. The reason that c needs to be small is that otherwise the region will contain a catenoid, which has no boundary (although it is recurrent and has quadratic area growth). The idea is to prove parabolicity and quadratic area growth for minimal surfaces-with-boundary contained in such a region. Here we discuss an alternative, probabilistic approach that bypasses universal superharmonic functions and delivers these results for regions within a sub-logarithmic factor of the exterior of a cone. The reference for this material, and the remainder of this section, is [13].

3.4. Parabolicity. We discuss parabolicity first, and the basic idea behind our approach is as follows. Consider some (non-empty) open subset $A \subset \mathbb{R}^3$. Suppose we show that any Brownian motion on a minimal surface that runs for all time almost surely leaves A , or equivalently, that any Brownian motion on a minimal surface that stays in A is stopped in finite time, almost surely. Then it follows that any stochastically complete minimal surface-with-boundary M contained in A is parabolic. This is because Brownian motion on M must be stopped in finite time, almost surely, and since M is stochastically complete, Brownian motion on M is stopped only when it hits the boundary. Recall that this is a more general formulation of the approach we used to show that any stochastically complete minimal surface-with-boundary contained in a ball is parabolic.

As already indicated, we are interested in rotationally symmetric domains of the form

$$A = \{|x_3| < f(r) \text{ and } r > e^L\}$$

for some positive, increasing $f(r)$ and positive integer L . This certainly includes the planar, catenoidal, and (exterior) conical domains mentioned above. Some positive lower bound on r is necessary, as we can see from the fact that otherwise A could contain the x_1x_2 -plane, which is not parabolic (in our sense, because it has no boundary). However, we will generally choose L for convenience, rather than trying to make it as small as possible. This is already seen in our restriction of L to be integer-valued, which is purely to simplify our later computations. This attitude is largely justified by the fact that we're primarily interested in minimal ends, and it's ultimately the "large r " asymptotic behavior of the ends that matters. Further, in the case when our minimal surface-with-boundary is properly immersed, changing L (for a fixed function $f(r)$) only changes the surface by a compact set, and neither parabolicity nor quadratic area growth is affected by addition or removal of a compact set from the surface (assuming the boundary remains non-empty).

We now explain the basic technique that we'll use to establish parabolicity, beginning with a non-technical explanation of semi-martingales and how to estimate certain hitting probabilities associated with them (for the details, any standard textbook on stochastic calculus can be consulted, such as [6]). A continuous real-valued martingale has the property that, if it starts at some $x_0 \in \mathbb{R}$, then the probability that it hits $x_0 - a$ before $x_0 + a$ for some positive a is exactly $1/2$, and vice versa, assuming that it almost surely hits one or the other in finite time. In other words, the martingale leaves the interval $(x_0 - a, x_0 + a)$ symmetrically, assuming that it starts in the middle and almost always leaves. This symmetry is a manifestation of the "conditionally constant" nature of martingales mentioned earlier. A semi-martingale is a process consisting of a martingale plus a process of bounded variation, frequently referred to as the "drift." Thus, a semi-martingale need not leave an interval symmetrically, and the amount of asymmetry can be controlled by controlling the drift (or more precisely, by controlling the expectation of the drift, since the drift itself is a process).

We now consider what this means in the case of Brownian motion. Specifically, consider Brownian motion on a Riemannian manifold, such as a minimal surface, which we denote B_t . Then if g is smooth function on the manifold, $g_t = g(B_t)$ is a (continuous) real-valued semi-martingale. Moreover, the drift of g_t is given, for each path, as the integral of half the Laplacian of g along the path. That is, the drift at time t is given by

$$\int_0^t \frac{1}{2} (\Delta g)(B_\tau) d\tau.$$

Now consider a minimal surface-with-boundary M contained in a region A of the above form for some $f(r)$ and L . We suppose that we start Brownian motion on M at a point with $\log r = n$ for some $n > L$, and we're interested in understanding which endpoint of $(n - 1, n + 1)$ the process $\log r_t$ exists via. We first note that a variant of the argument we used to show that Brownian motion on a minimal surface always exists a ball in finite time shows that Brownian motion on a minimal surface always exits the annular region $\{n - 1 < \log r < n + 1\}$ in finite time, unless it is stopped first (and, for future use, we note that the average time until it exits or is stopped can also be estimated from above). Thus, if $\log r_t$ were a martingale, the probability of exiting via either endpoint would be no greater than $1/2$. However, $\log r_t$ is not, in general, a martingale (it is for the x_1x_2 -plane, but not otherwise), and thus we need to estimate the expected integral of $(\Delta \log r)/2$ (with respect to time) along paths of the Brownian motion. The key to doing this is again the basic relationship

$$|\Delta \log r| \leq \frac{|\nabla x_3|^2}{r^2} \quad \text{on } M \setminus \{x_3\text{-axis.}\}$$

Because $\Delta x_3 = 0$, we can compute that $\Delta(x_3^2) = 2|\nabla x_3|^2$. Thus, the expected drift of $\log r_t$ is controlled by the expected drift of $x_{3,t}^2$. Finally, for $r_t \in (e^{n-1}, e^{n+1})$, we know that x_3^2 is bounded from above by $f(e^{n+1})^2$, which implies a bound for the expected drift of $x_{3,t}^2$. Putting this all together, we can estimate the probability that $\log r_t$ exits $\{n - 1 < \log r < n + 1\}$ via $n + 1$ in terms of $f(r)$. We let p_n be the upper bound on this quantity; it depends only on f (and n , of course). In particular, the smaller the growth rate of $f(r)$, the faster p_n decays to $1/2$.

So now we assume that we have a Brownian motion B_t on a (stochastically complete) minimal surface-with-boundary contained in A , and that $r_t = r(B_t)$.




FIGURE 1. A sample trajectory of r_t corresponding to one step of the random walk.

For convenience, we choose L so that p_n is always strictly between 0 and 1 (which of course depends on f). If we iterate the above procedure, that is, supposing that $\log r_t$ starts at n , we wait to see whether it proceeds to $n - 1$ or $n + 1$ and then we ask the analogous question starting from $n \pm 1$, we get a random walk on the integers $\{L + 1, L + 2, \dots\}$, which possibly stops at some time. One (possible) step of this walk is illustrated in Figure 1.

Now consider the Markov chain Y_m on the integers $\{L, L + 1, \dots\}$ with transition probabilities

$$\mathbb{P}[Y_{m+1} = a | Y_m = n] = \begin{cases} p_n, & \text{if } a = n + 1 \\ 1 - p_n, & \text{if } a = n - 1 \\ 0, & \text{otherwise,} \end{cases}$$

where $n = \{L + 1, L + 2, \dots\}$ and the random walk is stopped when it hits L . Then the paths of Markov chain are almost surely “above” the paths of the random walk we’ve associated to $\log r_t$ (in the sense that if the walk associated to $\log r_t$ is at site n after the step m , then $Y_m \geq n$, once we compare the sample paths in the right way), at least until $\log r_t$ is stopped. Now, $\log r_t$ cannot hit L . Therefore, if Y_m hits L almost surely we know that $\log r_t$ must be stopped almost surely, which means that B_t has finite lifetime. Random walks like Y_m are well-studied, and determining a condition on p_n (and thus on $f(r)$) such that Y_m almost surely hits L allows us to prove the following.

THEOREM 3.5. *Suppose that M is a stochastically complete minimal surface-with-boundary contained in*

$$A_1 = \left\{ r > e^L \text{ and } |x_3| < \frac{cr}{\sqrt{\log r \log(\log r)}} \right\}$$

for some $c > 0$ and sufficiently large L (with respect to c). Then M is parabolic.

3.5. Area growth. It is more surprising that these techniques can be used to study area growth. The key is this:

LEMMA 3.6. *If M is a properly immersed, parabolic minimal surface with compact boundary, then there is a compact curve and a probability measure on it such that Brownian motion started from this measure has occupation density equal to a positive constant, outside of a compact set.*

This means that quadratic area growth is equivalent to the occupation time (of sets of the form $\{r \leq a\}$) growing quadratically (in a). To understand the idea behind this lemma, we discuss the simplest case. Let M be a plane with an open disk of radius one removed. For concreteness, we may as well assume that M is the set $\{x_3 = 0 \text{ and } r \geq 1\}$. Then M is certainly a parabolic minimal surface-with-boundary, with the boundary being the unit circle $\{x_3 = 0 \text{ and } r = 1\}$. Consider the following function on M ,

$$h(x_1, x_2) = h(r) = \begin{cases} (1/\pi) \log r & \text{if } 1 \leq r < e \\ 1/\pi & \text{if } r \geq e. \end{cases}$$

Then h is continuous (including at the boundary) and zero at the boundary. Further, h has measure-valued Laplacian supported on the circle $\{r = e\}$, which is given by $1/\pi$ times the length measure on this circle, as can be verified by an elementary computation. Said differently, one-half the Laplacian of h is the uniform probability measure on the circle $\{r = e\}$ (where uniformity is with respect to the standard length measure). Thus, h is the Green's function on M for this measure (and where our Green's function is relative to one-half the Laplacian, which is the convention most appropriate for probability). Moreover, this means that h is also the occupation density (with respect to Lebesgue measure) for Brownian motion on M started from the uniform probability measure on $\{r = e\}$ and stopped at the boundary. That the occupation density is constant on $\{r \geq e\}$ might seem surprising at first. The intuitive explanation is that, while it is unlikely for the Brownian motion to make it out to where r is large, if it does make it out there it tends to wander around for a long time before hitting the boundary. As it turns out, these two effects balance one another exactly, leading to a constant occupation density. Finally, the idea behind the proof of the lemma is that this same approach works more generally, although in that case one doesn't expect an explicit formula for either the probability measure on the curve or the resulting Green's function.

In light of this lemma, we need to find conditions that give quadratic growth for the occupation time of Brownian motion on M in sets of the form $\{r \leq a\}$. Basic stochastic analysis gives that the occupation time can be estimated from $1 \leq \frac{1}{2} \Delta r^2 \leq 2$ and knowledge of the p_n . Translating the resulting condition on the p_n into a condition on $f(r)$ gives the following.

THEOREM 3.7. *Suppose that M is a properly immersed minimal surface with compact, non-empty boundary, contained in the region*

$$A_2 = \left\{ r > e^L \text{ and } |x_3| < \frac{cr}{\sqrt{\log r \log(\log r)}} \right\}$$

for some $c > 0$ and sufficiently large L (with respect to c). Then M has quadratic area growth.

Note that $A_2 \cap \{r > a\}$ is contained in $A_1 \cap \{r > a\}$ for sufficiently large a (and the same value of c). Thus, any properly immersed minimal surface with compact, non-empty boundary contained in A_2 is necessarily parabolic. This explains why we don't need to assume parabolicity in this theorem, despite it being used in the above lemma. It also shows that we require stronger conditions to prove quadratic area growth than we do to prove parabolicity, not only in that we assume the minimal surface is properly immersed with compact boundary, but also in that the region in which the surface needs to be contained is asymptotically smaller.

We close this section by placing these results in context and suggesting further questions. Of the various cases just discussed, the most important is that of catenoidal ends, which was already established in [3]. This is what is needed for the study of properly embedded minimal surfaces with two limit ends. The improvement from catenoidal ends to ends contained in A_1 or A_2 as above doesn't immediately have an application to larger issues in minimal surface theory, at least as far as we know. Nonetheless, we would like to note several positive features of this development. First, even for catenoidal ends, this approach *might* seem better motivated than cooking up a universal superharmonic function. Further, the results represents progress toward, and evidence for, Meeks' conjecture for cones. Finally, this approach introduces a new method, which demonstrates the applicability of stochastic methods to minimal surface theory (especially since the results available via this stochastic approach currently surpass those otherwise available) and might suggest methods for related questions.

There are several natural questions raised by the above. In the spirit of comparing this stochastic approach to the earlier approach using universal superharmonic functions, it would be helpful to know if there exists a positive, proper universal superharmonic function on A_1 or A_2 . Obviously, one wants to know if these ideas could be refined in order to give parabolicity and/or quadratic area growth for minimal surfaces-with-boundary in the exteriors of certain cones and thus establish Meeks' conjecture. Moreover, one might try to modify this approach in order to establish these properties for (minimal surfaces-with-boundary in) the region above the bottom half of a catenoid, which would establish another conjecture of Meeks and come closer to the complement of the sorts of regions studied by Alarcón and López, as mentioned above.

4. Coupling

4.1. Background. Coupling two stochastic processes is a common technique in probability. Given two processes, say x_t and y_t , with given distributions, a coupling is a joint distribution (x_t, y_t) such that the two marginals each have the desired distribution. Obviously, one possibility would be to take x_t and y_t to be independent. However, the goal of a coupling is generally to get x_t and y_t to meet, that is, to get $x_t = y_t$ at some time t (where we're assuming that x_t and

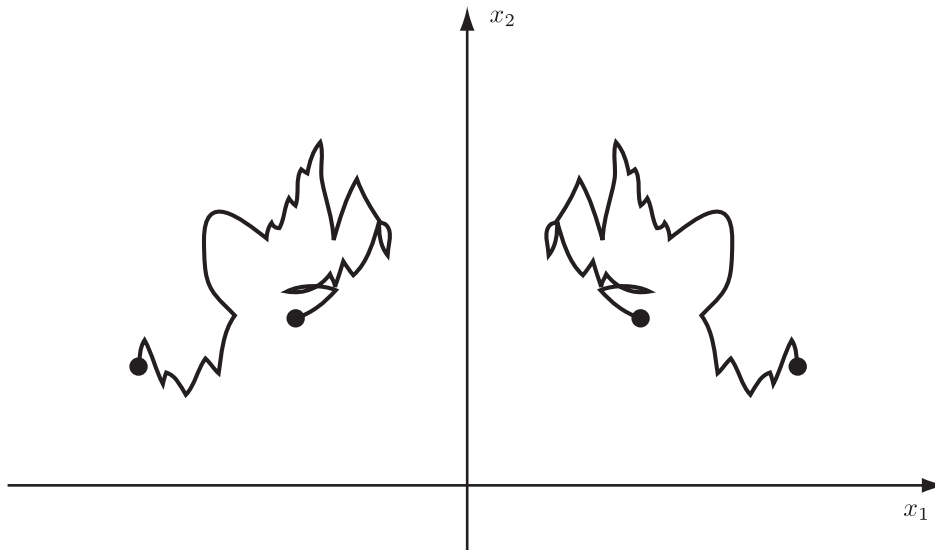


FIGURE 2. A planar Brownian motion and its mirror image over a finite interval of time.

y_t are defined on the same space), with as high a probability and/or as quickly as possible, and this generally requires the processes to be far from independent. Once the processes meet, we will stop them, so that we're only worried about defining the coupling up to their first meeting time. In a slightly awkward bit of terminology, when the processes meet they are often said to have coupled, and the time at which this happens is referred to as the coupling time; for clarity, we will stick with the term "meet."

We now give perhaps the simplest non-trivial example of the above ideas, which should serve to clarify matters. Suppose that we consider two Brownian motions in the plane (with coordinates (x_1, x_2)), y_t and z_t , which we write in coordinates as $y_t = (y_{1,t}, y_{2,t})$ and $z_t = (z_{1,t}, z_{2,t})$. Further, suppose for convenience that we start the processes at $y_0 = (1, 1)$ and $z_0 = (-1, 1)$ (where we've started them off of the x_1 -axis only in order to make the figure a bit cleaner). If we let y_t and z_t be independent, then $y_t - z_t$ is an \mathbb{R}^2 -valued Brownian motion sped up by a factor of $\sqrt{2}$ and started at $(2, 0)$. Because planar Brownian motion never hits the origin, almost surely, it follows that these independent Brownian motions will never meet. In order to get these two Brownian motions to meet, we instead wish for them to be very dependent. In particular, let z_t be completely determined from y_t by $(z_{1,t}, z_{2,t}) = (-y_{1,t}, y_{2,t})$. This is called the mirror coupling, for the obvious reason, namely that z_t is obtained by reflecting y_t across the x_2 -axis. A sample path of this coupled process (over a finite interval of time) is illustrated in Figure 2. In this case, $y_t - z_t$ can be written as $(2W_t, 0)$ where W_t is a one-dimensional Brownian motion started at 1. Since a one-dimensional Brownian motion almost surely hits the origin in finite time, we see that y_t and z_t , coupled in this way, almost surely meet in finite time.

Coupling two planar Brownian motions in order to make them meet is fairly straight-forward, since the mirror coupling admits an easy global description. Next,

consider the case of two Brownian motions on a compact Riemannian manifold M . Again, we wish to make them meet as quickly as possible. The idea is that, infinitesimally, we know what we should do. Suppose that, at some time t , z_t is not in the cut locus of y_t (and recall that “being in the cut locus” is a symmetric relationship). Then (since we’re interested in the case when the processes haven’t met yet) there is a unique minimal geodesic from y_t to z_t . Infinitesimally, we want z_t to evolve as the reflection of y_t across the plane (in $T_{z_t}M$) perpendicular to this minimal geodesic. Thus, we’re mimicking the mirror coupling, as well as possible, on an infinitesimal scale. This determines a stochastic differential equation (SDE) which we would like y_t and z_t to (jointly) satisfy. The cut locus creates an obvious problem for this scheme. However, it turns out that the amount of time the marginal processes spend in each other’s cut loci has measure zero, so one can get around that potential worry. In particular, this SDE can be solved, and the resulting joint process (y_t, z_t) , for whatever starting points we’re considering, is called the Kendall-Cranston mirror coupling (see Sections 6.5-6.7 of [5] for the both the construction of this coupling and details of the applications we give below).

Under the Kendall-Cranston mirror coupling, two Brownian motions on M will meet almost surely. Further, if the Ricci curvature is non-negative they do so quickly, in the sense that the average distance between the particles decreases exponentially at a rate that can be estimated. As an aside, this is of interest because this rate gives a lower bound on the spectral gap of M (that is, the first non-zero eigenvalue of minus the Laplacian on M , where this choice of sign gives a positive operator on $L^2(M)$). In particular, the Kendall-Cranston coupling gives the sharp estimate in two cases of interest, which we now state. For a compact Riemannian manifold M , let $\lambda_1(M)$ be this first non-zero eigenvalue, let $\text{Ric}_M(x)$ be the Ricci curvature of M , and let $\text{diam}(M)$ be the diameter of M . The following combines Theorems 6.7.3 and 6.7.4 of [5].

THEOREM 4.1. *Let M be a smooth, compact Riemannian manifold of dimension n .*

- i) If $\text{Ric}_M(x) \geq (n - 1)K > 0$ for a (positive) constant K and all $x \in M$, then $\lambda_1(M) \geq nK$.*
- ii) If M has non-negative Ricci curvature, then $\lambda_1(M) \geq \pi^2 / \text{diam}(M)^2$.*

4.2. A coupling for minimal surfaces. Suppose that M and N are stochastically complete minimal surfaces. Recall that a coupled Brownian motion is a process (x_t, y_t) on $M \times N$ such that x_t and y_t are Brownian motions on M and N , respectively. Note that x_t and y_t are also martingales in \mathbb{R}^3 . Our goal is to couple x_t and y_t in such a way as to make them meet in \mathbb{R}^3 (in finite time). Intuitively, we want to do this by developing an extrinsic analogue of the Kendall-Cranston mirror coupling. That we are concerned with the extrinsic (\mathbb{R}^3) distance is not too surprising in light of the fact that minimal surfaces are defined in terms of an extrinsic condition. The reference for the remainder of this section is [14], except for Theorem 4.8, for which [12] should be consulted.

Before discussing the construction of this coupling, we explain the motivation behind looking for such a coupling. If M and N are different and x_t and y_t meet with positive probability, then M and N intersect. This shows that such a coupling gives an approach to halfspace theorems, as various intersection theorems are called in the minimal surface literature. In a different direction, if $M = N$ is embedded and x_t and y_t meet in \mathbb{R}^3 almost surely, then M admits no non-constant bounded

harmonic functions. The point is that if M is embedded, then if x_t and y_t meet in \mathbb{R}^3 they also meet on M (intrinsically). Then it is true in general that if Brownian motions from any two points on a Riemannian manifold can be coupled so as to meet almost surely, the manifold admits no non-constant bounded harmonic functions (this follows directly from the Brownian motion representation of harmonic functions). This shows that such a coupling can be used to study the conformal structure of embedded minimal surfaces and, more specifically, a conjecture of Sullivan that a properly embedded minimal surface admits no non-constant positive harmonic functions. Another general motivation for introducing stochastic methods in minimal surface theory is that one expects that they will give proofs that are relatively easy to adapt to the case when the surfaces are allowed to have boundary.

We begin by indicating what the coupling should look like pointwise. Suppose the particles are at points x_t and y_t , and consider the vector $x_t - y_t \in \mathbb{R}^3$ between them. We are guided by the mirror coupling on \mathbb{R}^3 . In that case, the idea is to have the particles move in opposite directions along the direction of $x_t - y_t$ and move in the same direction in the plane perpendicular to $x_t - y_t$. Here we wish to come as close to that ideal as possible, except that the particles are constrained to move only in the planes $T_{x_t}M$ and $T_{y_t}N$, respectively. This forces us to make a trade-off between the behavior in the $x_t - y_t$ direction and the behavior in the orthogonal complement. Without going into detail, we now describe the result of making this trade-off in the optimal way.

Let $r_t = |x_t - y_t|_{\mathbb{R}^3}$ be the distance between the particles. We wish to make the evolution of r_t as favorable, instantaneously, as possible. In order to understand how well we're doing, infinitesimally, in making the particles meet, that is, in making r_t eventually hit 0, we recall the one-parameter family of Bessel processes. For each $d \in (0, \infty)$, there is a (continuous) sub-martingale on $[0, \infty)$, defined until the first time it hits 0 (at which point, for our purposes, we just stop it), called the d -dimensional Bessel process. When $d \in \{1, 2, 3, \dots\}$, the d -dimensional Bessel process is the distance from the origin of Brownian motion on \mathbb{R}^d . For our purposes, two aspects of the Bessel processes are important. First, they are well-understood with many explicit formulas, making them well-suited for use as comparison processes. Second, and more concrete, the behavior of the d -dimensional Bessel process relative to 0 is as follows. For $d < 2$ the Bessel process hits 0 (in finite time) almost surely, for $d = 2$ the Bessel process almost surely comes arbitrarily close to 0 but never hits it, and for $d > 2$ the Bessel process has positive infimum and diverges to infinity in the large time limit, almost surely.

Infinitesimally, we can make the evolution of r_t look like a time-change of a Bessel process of dimension 2 or less (in the sense that the coefficients of the SDE satisfied by r_t agree with those of such a time-changed Bessel process at time t). Further, both the infinitesimal time-dilation and the dimension of the infinitesimal model depend only on the relative position of three unit vectors, the unit normal vectors to both M and N at the points x_t and y_t and $(x_t - y_t)/r_t$, and each of these matters only up to sign (in particular, the orientability and orientations of M and N don't matter). We continue to discuss the qualitative aspects of the evolution of r_t under this proposed coupling. The time-change is understood in terms of the quadratic variation of r_t . The point is that if r_t accumulates only a finite amount of quadratic variation, it will converge to some finite value, much like a Bessel process run for a finite interval of time. Obviously, if the particles meet we stop the coupled

process and r_t has only accumulated a finite amount of quadratic variation. What we are concerned with is the possibility that r_t converges to some positive value by accumulating only a finite amount of quadratic variation even as $t \rightarrow \infty$. The other concern is that, when r_t is small, it might be “infinitesimally modelled” on Bessel processes of dimension arbitrarily close to 2. In this case, it might come arbitrarily close to 0 without ever hitting it, just like a Bessel process of dimension 2.

Stated positively, these are the only two obstructions to x_t and y_t almost surely meeting (in \mathbb{R}^3 , in finite time). This shows that our infinitesimal control of r_t is borderline. For instance, if we knew that the rate of accumulation of quadratic variation (or equivalently, the instantaneous time dilation) were bounded from below by a positive constant (as opposed to zero) and that the evolution of r_t infinitesimally looked like a time-change of a Bessel process of dimension less than $2 - \epsilon$ for some $\epsilon > 0$, then there would be no question that r_t hit 0. However, there is no way to construct a coupling which satisfies these stronger conditions, even in simple cases, as we will see below. On the other hand, for generic configurations of the unit normal vectors to both M and N at the points x_t and y_t and $(x_t - y_t)/r_t$ the rate of accumulation of quadratic variation is positive and the infinitesimal model has dimension strictly less than 2. Moreover, the exceptional set of configurations where this does not hold can be described explicitly, though we don’t do so here.

This leads to a fairly natural approach to getting x_t and y_t to meet. We impose some global conditions on M and N and try to use these conditions to show that the configuration process doesn’t spend too much time on (or near) the exceptional set just mentioned. If we can do this, then we should be able to rule out both of the obstructions and conclude that x_t and y_t meet in finite time. This situation is somewhat analogous to wanting to show that a function $f : [0, \infty) \rightarrow \mathbb{R}$ with $f(0) > 0$ has a zero. Showing $f' \leq 0$ is the borderline infinitesimal condition, but you also need some global condition so that you don’t get “stuck” before reaching zero.

4.3. Explicit examples. Here we give two cases in which the coupled processes can be given explicitly, which hopefully serve to illustrate the above ideas.

First, consider the case where M and N are parallel (and distinct) planes. We may as well normalize things so that the planes are distance one apart and choose coordinates so that $M = \{x_3 = 0\}$ and $N = \{x_3 = 1\}$. Further, we note that the situation decomposes in a natural way. Namely, choose any two initial points, one on each plane. Then we can rotate and translate our coordinates on \mathbb{R}^3 so that the x_2x_3 -plane contains both initial points and is perpendicular to both M and N . Recall that the x_3 -axis is perpendicular to M and N . Now, to keep the notation more manageable, let a_t be the Brownian motion on M and b_t the Brownian motion on N . Under the coupling described above, $x_1(a_t) = x_1(b_t)$ for all time (until the particles meet and we stop everything, of course). In other words, the particles are coupled simply by tracking one another exactly in the x_1 -direction. Hence, the problem essentially reduces to the lower dimensional case of understanding the evolution of the x_2 coordinate of both processes (note that $x_3(a_t) = 0$ and $x_3(b_t) = 1$ for all t). This is equivalent to implementing the analogous coupling on two parallel lines.

The joint evolution of $x_2(a_t)$ and $x_2(b_t)$ is given by Figure 3. It is important to note that this diagram represents the Cartesian product of the “ x_2 -axis on M ,” which we label y_1 , and the “ x_2 -axis on N ,” which we label y_2 , not the original space

where the planes are embedded. The idea is that the joint process evolves along the dashed lines. More concretely, the dashes should be thought of as vectors, and the joint diffusion is generated by the squares of these vectors (which explains why they have no “heads” or “tails;” after squaring it doesn’t matter which is which). The length of the dashes should be understood to be $\sqrt{2}$ so that the marginal distributions are just one-dimensional Brownian motions (the dashes have slope ± 1 so that the projections onto the two axes each have length 1). It is clear that there are two regions (one of which is disconnected) of this configuration space. On one region (the connected one), $x_2(a_t)$ and $x_2(b_t)$ are coupled so as to move in the same direction, with the result that they differ only by some fixed translation. In particular, the \mathbb{R}^3 -distance between the particles remains fixed if the joint process is in this region. Also note that once the joint process enters this region, it never leaves. On the other region (the disconnected one), $x_2(a_t)$ and $x_2(b_t)$ are coupled so as to move in opposite directions and the distance between the particles clearly evolves (randomly). We also note that if the joint process is in this region, it eventually (almost surely) enters the other region, which occurs when it hits the boundary between these two regions. (It is important to note that this boundary belongs to the first, connected region.) The boundary between the two regions corresponds to those points at which $|x_2(a_t) - x_2(b_t)| = 1$. Said differently, it is those points at which the vector from $x_2(a_t)$ to $x_2(b_t)$ makes a $\pm\pi/4$ angle with the horizontal axis.

Combining these observations, we arrive at the following description of the behavior of the joint process for $x_2(a_t)$ and $x_2(b_t)$. If the particles start at points with $|x_2(a_0) - x_2(b_0)| \leq 1$, then they remain that same fixed distance apart for all time. In fact, a_t and b_t , as processes in \mathbb{R}^3 , are related by translation by the fixed vector $a_0 - b_0$. If the particles start with $|x_2(a_0) - x_2(b_0)| > 1$, then $x_2(a_t)$ and $x_2(b_t)$ evolve as mirror reflections of one another until the distance between them hits 1, which it almost surely does in finite time. Once that happens, the joint process evolves as in the previous case, and in particular, $|x_2(a_t) - x_2(b_t)| = 1$ for all t greater than or equal to the first time when $|x_2(a_t) - x_2(b_t)|$ hits 1. Thus, if a_t and b_t start far apart in \mathbb{R}^3 (more precisely, if the distance between them is greater than $\sqrt{2}$), then the coupling almost surely brings them to distance $\sqrt{2}$ apart in finite time. However, it will not make them any closer. In hindsight, this is perhaps not too surprising. After all, parallel planes do not intersect, and there’s no hope of a_t and b_t meeting. This manifests itself in the fact that, once the particles gets close enough, r_t stops accumulating quadratic variation. However, the coupling performs more or less as well as it can under the circumstances, by at least bringing the particles to distance $\sqrt{2}$.

For our second example, consider the case when M and N are perpendicular planes. Then we choose coordinates for \mathbb{R}^3 so that $M = \{x_3 = 0\}$ and $N = \{x_2 = 0\}$. Again, we let a_t be the Brownian motion on M and b_t the Brownian motion on N , and we also assume that $x_1(a_0) = x_1(b_0) = 0$. In this case, under the coupling described above, we again have that $x_1(a_t) = x_1(b_t)$ for all time (until the particles meet and we stop everything, of course). Hence, the problem essentially reduces to the lower dimensional case of understanding the joint evolution of $x_2(a_t)$ and $x_3(b_t)$ (note that $x_3(a_t) = x_2(b_t) = 0$ for all t). This is equivalent to implementing the analogous coupling on two perpendicular lines.




FIGURE 3. The state space of the coupled Brownian motion on parallel lines. The dashes should be thought of as the vector field generating the diffusion.

The joint evolution of $x_2(a_t)$ and $x_3(b_t)$ is given by Figure 4. As before, this diagram represents the Cartesian product of the “ x_2 -axis on M ,” which we label y_1 , and the “ x_3 -axis on N ,” which we label y_2 , not the original space where the planes are embedded. We interpret the dashes in the same way as before. In the interior of a given quadrant, $x_2(a_t)$ and $x_3(b_t)$ (equivalently, y_1 and y_2) move either in the same direction or the opposite direction. As the joint process crosses an axis, it moves toward the origin according to the local time of the process on the axis. (Recall that both $x_2(a_t)$ and $x_3(b_t)$ are martingales. Equivalently, if we were to recast the martingale problem represented by the diagram as a stochastic differential equation determined by the squares of the dashes, then the SDE would be understood in the Ito sense.) Note that the origin of the diagram corresponds exactly to the line of intersection of M and N , and thus a_t and b_t meet precisely if and when the joint process hits the origin. Since the joint process will continue to cross an axis until the process is stopped, we see that it almost surely hits the origin in finite time, and thus a_t and b_t almost surely meet in finite time. (Diffusions in

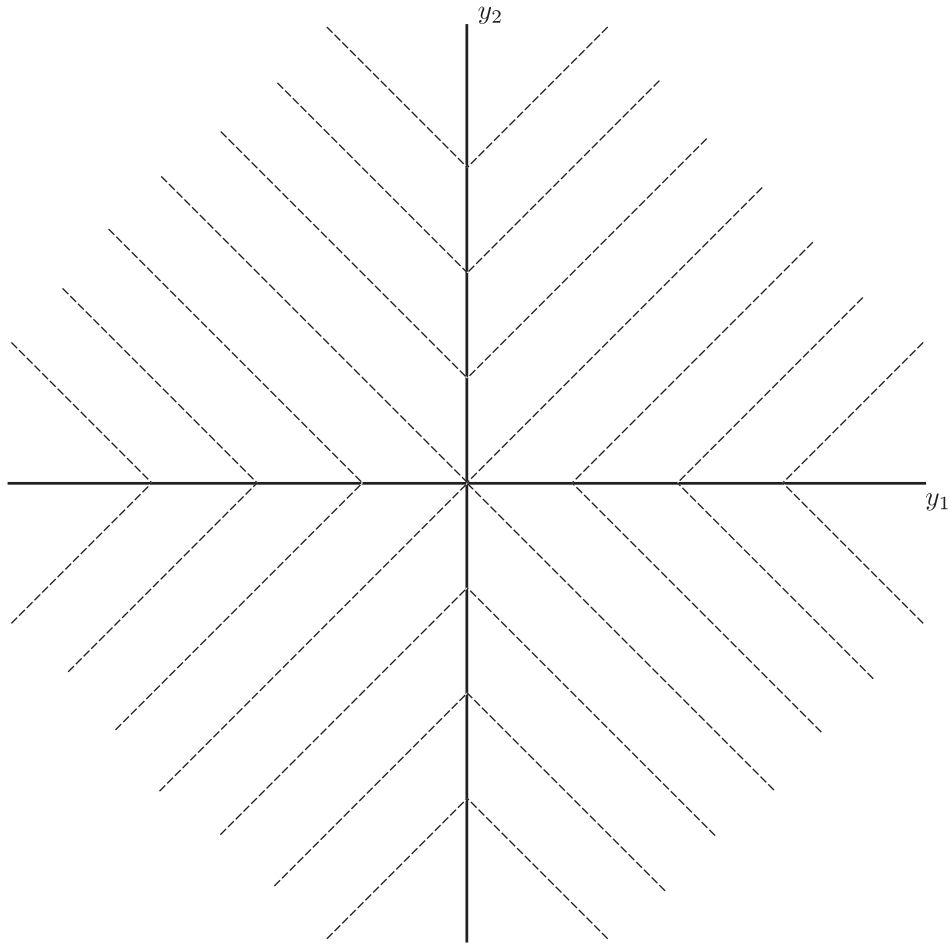


FIGURE 4. The state space of the coupled Brownian motion on perpendicular lines. Again, the dashes should be thought of as the vector field generating the diffusion (in the Ito sense).

the plane with this sort of behavior have appeared before. For example, Sections 5.1 and 5.3 of [2] discuss a process which evolves in the same way, except that it moves outward, away from the origin, according to the local time at the discontinuities.)

Again, we can give a more intuitive description of the joint behavior of $x_2(a_t)$ and $x_3(b_t)$. If $|x_2(a_0)| = |x_3(b_0)|$, then the joint process starts on a diagonal of the diagram. Thinking of the processes as Brownian motions on perpendicular lines, this means they both start at the same distance from the point of intersection. They are coupled so that they both move toward or away from the point of intersection together, and thus they meet at the first time one, and hence both, of them hit the point of intersection. Now suppose that $|x_2(a_0)| \neq |x_3(b_0)|$, and without loss of generality that $|x_2(a_0)| > |x_3(b_0)|$. Then again, the two particles both move toward or away from the point of intersection together. However, in this case $x_3(b_t)$ reaches the point of intersection for the first time before $x_2(a_t)$ does. When it passes through the point of intersection, the directions “toward and away from

the point of intersection” switch, and this is what causes the generator of the joint process to be discontinuous at the axes. At any rate, the two particles continue in this fashion until $x_2(a_t)$ reaches the point of intersection for the first time. The coupling is such that $x_3(b_t)$ is almost surely also at the point of intersection at this time, and thus the particles meet. One thing to take away from this example is that, at least in this simple case, the coupled Brownian motions on M and N succeed in finding points of intersection of the two surfaces almost surely.

Note that these are, in a sense, the two simplest and most explicit examples, and they clearly show that the generator of the coupled process must be discontinuous, as well degenerate.

4.4. The coupling in general. In the simple examples just discussed, a coupling having the desired infinitesimal structure could be constructed by hand. In general, this is more difficult. Because the generator is both degenerate and discontinuous, standard existence theorems do not apply. Nonetheless, it is possible to prove the existence of a coupling satisfying the desired inequalities. Once that is done, the issue becomes gaining enough control of the global behavior of the process in order to show that the particles meet. As the example of parallel planes shows, some kind of global control is clearly necessary in order to avoid the two potential obstacles to coupling mentioned above.

In what follows, we discuss the cases in which we have sufficient global control to draw geometrically interesting conclusions. (The theorems not otherwise attributed are from [14].) Given that the infinitesimal behavior of the coupling is governed by the relative position of the vector connecting the two particles and the unit normal vectors to both surfaces, it is not surprising that control of the evolution of the unit normal vector plays a role. To that end, consider the evolution of the unit normal vector to a minimal surface under Brownian motion, that is, the composition of the Gauss map with Brownian motion on the surface. For convenience, we call the resulting process on \mathbb{S}^2 the Gauss sphere process. It is a basic geometric result that the Gauss map of a minimal surface is conformal, with the conformal dilation given by K , the Gauss curvature of the surface (that K is non-positive means that the Gauss map reverses orientation, or is anti-conformal if we take our angles to be signed). In terms of the Gauss sphere process, this means that it’s time-changed Brownian motion in \mathbb{S}^2 with the time-change given by K . A key result for the Gauss sphere process is the following.

THEOREM 4.2. *Let M be a stochastically complete minimal surface of bounded curvature. Then if M is not flat, the Gauss sphere process almost surely accumulates infinite occupation time in every (non-empty) open subset of \mathbb{S}^2 .*

This result provides sufficient control of the global behavior of the unit normal and thus the coupling to prove the following theorem.

THEOREM 4.3. *Let M be a stochastically complete, non-flat minimal surface with bounded curvature, and let N be a stochastically complete minimal surface. Then the distance between M and N is zero.*

The point is that the control of the Gauss sphere process on M that we get from the fact that M has bounded curvature allows us to rule out the possibility that r_t “runs out of quadratic variation” and converges to a positive value. Unfortunately, it does not give sufficient control when r_t is small to show that the particles actually

meet. In other words, we're able to overcome one of the potential difficulties with the coupling, but not the other. This explains why we have an almost-intersection result, instead of the full strong halfspace theorem for minimal surfaces of bounded curvature. In particular, the above result should be compared with the following result of Rosenberg [15].

THEOREM 4.4. *Let M and N be complete minimal surfaces of bounded curvature. Then either M and N intersect, or they are parallel planes.*

In the spirit of extending results to minimal surfaces-with-boundary, we note that the proof of Theorem 4.3 can be easily modified, essentially by stopping the coupled process when either particle hits a boundary, to give

THEOREM 4.5. *Let M and N be stochastically complete minimal surfaces-with-boundary, at least one of which has non-empty boundary, such that $\text{dist}(M, N) > 0$. If M has bounded curvature and is not flat, then*

$$\text{dist}(M, N) = \min\{\text{dist}(M, \partial N), \text{dist}(\partial M, N)\}.$$

Again, this should be compared with the following “full” version of this theorem, due to Meeks and Rosenberg [10] and known as the maximum principle at infinity.

THEOREM 4.6. *Let M and N be disjoint, complete, properly immersed minimal surfaces-with-boundary, at least one of which has non-empty boundary. Then the distance between them satisfies*

$$\text{dist}(M, N) = \min\{\text{dist}(M, \partial N), \text{dist}(\partial M, N)\}.$$

Moving on, recall the conjecture of Sullivan mentioned earlier that a complete, properly embedded minimal surface admits no non-constant, positive harmonic functions. This was previously verified in a couple of cases. Obviously, the conjecture holds if M is recurrent. Also, Meeks, Pérez, and Ros [9] have proved this under additional symmetry assumptions (double or triple periodicity, or various conditions about the quotient by the isometry group having finite topology). Using the coupling, the following partial result can be proved.

THEOREM 4.7. *Let M be a complete, properly embedded minimal surface of bounded curvature. Then M has no non-constant bounded harmonic functions.*

The idea behind the proof is fairly straight-forward. As was the case with Theorem 4.3, the bounded curvature of M means the particles get arbitrarily close, either infinitely often or until they couple. However, unlike in that case, we also have control of the coupling when the particles are close. This comes from Meeks and Rosenberg's tubular neighborhood theorem (also proven in [10] as a consequence of the maximum principle at infinity) which implies that when the particles are close, the situation is uniformly close to the mirror coupling on \mathbb{R}^2 . This allows us to show that the particles eventually meet, almost surely. As discussed earlier, this is enough to conclude that M admits no non-constant bounded harmonic functions.

As a final example, we point out that this coupling can be used to settle a conjecture of Meeks for minimal graphs, as was done in [12] and given in the following theorem. (Weitsman [17] proved this under the additional assumption that the domain is finitely connected. Related results were obtained by López and Pérez [7].)

THEOREM 4.8. *Any minimal graph (that is, a complete minimal surface-with-boundary, the interior of which is a graph over some planar region), other than a plane, is parabolic.*

Here the “global control” used to make the coupling work is that the Gauss sphere process converges on a minimal graph, which gives particularly strong control of both unit normals. Also, we note that any graph is embedded. Using this, it is possible to show that if Brownian motion does not almost surely hit the boundary (which is equivalent to M not being parabolic), then we can choose two starting points such that the coupled processes from those two points meet with high probability. However, we can also choose these two points so that the associated Gauss sphere processes do not intersect with high probability. This contradicts the fact that the particles meet with high probability, since when they meet the corresponding normal unit vectors must be the same. This contradiction establishes the result.

4.5. Further questions. Several of the above results obtained by coupling fall short of what one might ultimately hope for, as already noted. However, rather than simply asking if they can be improved, here we raise some broader questions about the use of these techniques. Regarding the conjecture of Sullivan, it would be nice to extend Theorem 4.7 to exclude non-constant positive harmonic functions. While coupling Brownian motions gives a natural approach to excluding non-constant bounded harmonic functions on any Riemannian manifold, something more is needed for positive harmonic functions. An effective stochastic technique for positive harmonic functions on minimal surfaces would likely be helpful in situations beyond minimal surfaces. (It’s not obvious how to make a natural guess, like h-transforms, work in this situation.)

Many of the above theorems assume bounded Gauss curvature, mainly in order to take advantage of Theorem 4.2 which gives some global control of the coupling. However, it is natural to try to prove theorems for properly immersed (or embedded) minimal surfaces, with no assumption on the Gauss curvature. (Indeed, the strong halfspace theorem was originally proven by Hoffman and Meeks [4] under such hypotheses.) Note that properness of the immersion is a global condition, in contrast to the local condition of bounded curvature (also note that properness implies geodesic completeness). In light of this, it would be desirable to understand how properness manifests itself stochastically. Another important property of minimal surfaces (and of minimal submanifolds more generally) is stability (and related notions like the index of stability). Again, it would be interesting to know how stability interacts with the global behavior of Brownian motion, since this might allow stochastic methods to be developed for studying stability and related concepts.

Finally, as noted at the beginning of Section 2.2, Brownian motion on a minimal immersion in any ambient manifold possesses nice properties, with minimal surfaces in \mathbb{R}^3 being the simplest (interesting) case. Thus a more thorough study of how stochastic methods can be used to understand the geometry of minimal submanifolds seems worthwhile.

References

1. Antonio Alarcón and Francisco J. López, *Minimal surfaces in \mathbb{R}^3 properly projecting into \mathbb{R}^2* , [arXiv:0910.4124v2](https://arxiv.org/abs/0910.4124v2).
2. Václav E. Beneš, Ioannis Karatzas, and Raymond W. Rishel, *The separation principle for a Bayesian adaptive control problem with no strict-sense optimal law*, Applied stochastic

- analysis (London, 1989), Stochastics Monogr., vol. 5, Gordon and Breach, New York, 1991, pp. 121–156. MR 1108420 (92b:93087)
3. Pascal Collin, Robert Kusner, William H. Meeks, III, and Harold Rosenberg, *The topology, geometry and conformal structure of properly embedded minimal surfaces*, J. Differential Geom. **67** (2004), no. 2, 377–393. MR MR2153082 (2006j:53004)
 4. D. Hoffman and W. H. Meeks, III, *The strong halfspace theorem for minimal surfaces*, Invent. Math. **101** (1990), no. 2, 373–377. MR MR1062966 (92e:53010)
 5. Elton P. Hsu, *Stochastic analysis on manifolds*, Graduate Studies in Mathematics, vol. 38, American Mathematical Society, Providence, RI, 2002. MR 2003c:58026
 6. Ioannis Karatzas and Steven E. Shreve, *Brownian motion and stochastic calculus*, second ed., Graduate Texts in Mathematics, vol. 113, Springer-Verlag, New York, 1991. MR MR1121940 (92h:60127)
 7. Francisco J. López and Joaquín Pérez, *Parabolicity and Gauss map of minimal surfaces*, Indiana Univ. Math. J. **52** (2003), no. 4, 1017–1026. MR MR2001943 (2004f:53005)
 8. William H. Meeks, III and Joaquín Pérez, *The classical theory of minimal surfaces*, available at <http://www.ugr.es/~jperez/papers/papers.htm>.
 9. William H. Meeks, III, Joaquín Pérez, and Antonio Ros, *Liouville-type properties for embedded minimal surfaces*, Comm. Anal. Geom. **14** (2006), no. 4, 703–723. MR 2273291 (2007i:53009)
 10. William H. Meeks, III and Harold Rosenberg, *Maximum principles at infinity*, J. Differential Geom. **79** (2008), no. 1, 141–165. MR 2401421 (2009g:53011)
 11. Nikolai Nadirashvili, *Hadamard's and Calabi-Yau's conjectures on negatively curved and minimal surfaces*, Invent. Math. **126** (1996), no. 3, 457–465. MR MR1419004 (98d:53014)
 12. Robert W. Neel, *Brownian motion and the parabolicity of minimal graphs*, arXiv:0810.0669v1.
 13. ———, *On parabolicity and area growth of minimal surfaces*, arXiv:1004.4544v1.
 14. ———, *A martingale approach to minimal surfaces*, J. Funct. Anal. **256** (2009), no. 8, 2440–2472. MR MR2502522
 15. Harold Rosenberg, *Intersection of minimal surfaces of bounded curvature*, Bull. Sci. Math. **125** (2001), no. 2, 161–168. MR MR1812162 (2002b:53007)
 16. Daniel W. Stroock, *An introduction to the analysis of paths on a Riemannian manifold*, Mathematical Surveys and Monographs, vol. 74, American Mathematical Society, Providence, RI, 2000. MR MR1715265 (2001m:60187)
 17. Allen Weitsman, *A note on the parabolicity of minimal graphs*, Complex analysis and dynamical systems III, Contemp. Math., vol. 455, Amer. Math. Soc., Providence, RI, 2008, pp. 411–416. MR 2408185 (2009j:35109)

DEPARTMENT OF MATHEMATICS, LEHIGH UNIVERSITY, BETHLEHEM, PENNSYLVANIA, USA
E-mail address: robert.neel@lehigh.edu