CONTENTS

1

Introduction

On a late evening in March 2009, a group of people stood near the beach on the Atlantic coast of Florida. As they exchanged smiles, they nervously monitored the skies to the south. In the distance, a new space telescope was perched atop a large rocket set to lift off. This long-awaited launch was the culmination of twenty-five years of work from an ever-expanding group of scientists and engineers, some of whom, including me, were standing on that beach looking toward the launchpad. Our mission would help answer some of the most profound questions that humanity has ever posed: Are there other worlds like ours out in the emptiness of space? Are we alone?

This craft was slated to make its observations for three and a half years, though the more optimistic among the group thought it might fly for nearly a decade. Designed to detect the presence of planets orbiting distant Sun-like stars, the hope was to measure the number of Earth-like planets circulating in Earth-like orbits within our galaxy. If the mission functioned roughly as expected, it would fundamentally change how we viewed our home planet, and ourselves.

The distant flash and rising column of light and smoke marked the transition from concept to reality for this mission. Those who were gathered to see its launch knew that these were the frightening few moments where things could go terribly wrong. As the seconds ticked by, the milestones for the launch came and went with the rocket and its payload performing “nominally”—a good sign. Cheers, hugs, and slapped backs grew more frequent as the rocket flew south over the Caribbean and disappeared. Now, there was nothing left for us to do but return to our hotel rooms and wait to hear the fate of the spacecraft. It still had miles to go before it was safely in its orbit and NASA's Kepler mission could start taking data.

Eight years earlier, in 2001, I set foot on campus at the University of Washington in Seattle as a new graduate student in the physics department. I had turned down offers from several famous schools back east (including my dream school) to attend the "U-Dub." After visiting the other potential graduate programs, and after a lot of discussion, my wife and I felt that Seattle was the place for us. The campus and city were familiar sites since, a few summers prior, I interned at the University of Washington in the Institute for Nuclear Theory, where I worked on solar neutrinos. This was one of those big decisions couples often make before their first anniversary.

We spent the months prior to our move working out the details of my graduate education—choosing what classes I would take, where we would live, and with whom I would work on my research. The plan was to study cosmology. It sounded pretty cool, and I was told that there was a world-renowned cosmologist in the department. For the next half decade, I would work with him, doing whatever it was cosmologists did. Unbeknownst to me, this plan was foiled long before I trekked north from Salt Lake City. The world-renowned cosmologist was promoted to the dean of the college and stopped taking new students. With my bags still packed from the move, my imagined advisor and research area were off the table and I needed to find an alternative advisor and an alternative area of study.

During my first year in Seattle, I tagged along with a research group that studied the properties of distant galaxies, looking at the kinds of stars that they held and how those stars moved about. Then, for another year, I worked with a physics research group, building experiments to test alternative theories of gravity using a glass *torsion pendulum*. While these topics were interesting, and the professors were smart and engaging, neither project quite matched the romantic view that I had rattling around my head for what I would study, and where I would make my contribution to the volume of human knowledge.

Eventually, one professor suggested I look into the new guy, Eric Agol, who had just been hired by the Astronomy Department. He was filling the spot vacated by the recently promoted dean. His work sounded sufficiently cosmology-ish. So I reached out, we exchanged a couple of emails, and, sight unseen, we agreed to work together upon his arrival in the fall of 2003.

From our emails, we already knew what my dissertation project would be, and it was awesome. We were going to make pictures of the supermassive black holes that are found deep in the centers of galaxies. At the time, there was talk in bigwig astronomy circles about imaging those black holes using a network of radio telescopes that spanned the globe. With our project, we would be poised to produce the first-ever image of one—an incredible thought to consider as a graduate student.

While I was thrilled with the plan, there was one small issue. I hadn't yet taken general relativity, a class about Einstein's theory of gravity, which would provide essential background information for researching black holes and their environs. So, to pass the time, my advisor outlined a different problem, a "practice problem." It was a scenario he had discussed with a friend, relating to a recent discovery from a different astronomical discipline. Apparently, a couple of years prior (in the year 2000), astronomers had been observing the star HD 209458 and saw the signal of an orbiting planet as it transited, or passed in front of the star. My advisor thought there might be interesting things we could learn about such planetary systems by looking at small changes in the motions of the planets. It was the first time such an event had been seen, and there were lots of unexplored questions that transiting planets might answer.

Indeed, this entire area of astronomy was new to the scene. In late 1995, when *Toy Story* was showing on the big screen and "Gangsta's Paradise" first hit the airwaves, there were only nine known planets—Mercury, Venus, Earth, Mars, Jupiter, Saturn, Uranus, Neptune, and 51 Pegasi b. The last planet in that list, 51 Pegasi b (or "51-Peg"), was the first planet discovered outside our solar system that orbited a star like the Sun. A pair of Swiss astronomers found it circling a star in the constellation of Pegasus some fifty light-years (three hundred trillion miles) away [1]. The planet had a mass roughly half that of Jupiter, or a bit more than Saturn, and the star it orbited was similar to the Sun in size, mass, and temperature.

While the discovery of 51-Peg was clearly groundbreaking, most astronomers believed that with the right instrument and the right set of observations, the discovery of a planet orbiting a distant star, or an *exoplanet,* was just a matter of time. However, the astronomy community was still taken aback by the 51-Peg discovery, not because they didn't expect there to be other planets, and not because its discovery wasn't a major breakthrough, but because they didn't expect the planet to look like what was found. It was like some inexplicable creature had emerged from the swamp and knocked on our door.

The presumption was that a planet with a mass in the range of Jupiter and Saturn would be similar in structure and composition to these gigantic bodies—mostly gaseous material with a rocky or metallic core. Unlike Jupiter and Saturn, which respectively take twelve and twenty-nine years, to orbit the Sun, 51-Peg orbits its host star every four days. In the solar system, no planet orbits that close to the Sun, and those that come nearest are dense, rocky planets with thin atmospheres. This planet circles its star at one-tenth the distance of Mercury's orbit—along with its huge atmosphere of volatile material that shouldn't have been able to condense under the intense heat.

This discovery upended our astronomical logic. The territory for such giant planets should be the outer parts of a planetary system, where the temperatures are cold. For instance, Jupiter, the closest to the Sun of the solar-system giant planets, is five times farther from the Sun than the Earth. It receives only four percent of the light that we receive. This relatively cold environment allowed the forming Jupiter's gravity to trap lighter elements in its atmosphere. But 51-Peg orbits in less than a week. The blazing temperatures that are found that close to the star should have prevented it from forming. A *hot Jupiter* like 51-Peg ought not to exist, yet there it was—contradicting centuries of theoretical and observational work on how planets form, where they form, and how they evolve with time.

To understand just how strange these planets are, we need a little more background. Since the 1600s, astronomers studied the planetary bodies in our solar system, trying to understand their origins. Improvements in our understanding of physics, and developments in our ability to observe the planets, led to a growing expectation that the solar system shouldn't be unique, that the same rules that governed our beginnings were universal—being in force around the billions of stars across the galaxy.

While lacking the technology to observe a distant planetary system, people have long speculated on their possible existence. In 1584, the Dominican Friar Giordano Bruno published his *On the Infinite Universe and Worlds,* wherein he suggested that the stars in the sky were Suns of their own, and could harbor their own planets. This belief was one of many that brought him before the Roman Inquisition and contributed to his being burned at the stake. While Bruno died, his idea persisted—though it would be nearly four hundred years before his hypothesis of distant worlds would be vindicated.

In the 1700s, Emanuel Swedenborg, Immanuel Kant, and Pierre-Simon Laplace theorized that a star, like the Sun, could form from the collapse of an enormous cloud of gas, or *nebula*. If that initial cloud started with a small amount of rotation, some of the gas would flatten into a disk of circulating material surrounding the star, much like the dress of a spinning dancer flattens into a disk. Dense material would settle to the midplane of the disk, and would eventually coagulate into a set of planets.

Over the subsequent two centuries, this *nebular theory* of planet formation was tested, revised, and tested again using observations of the planets in the solar system. This simple theoretical premise provides a consistent explanation for a lot of the properties that we observe. For example, it explains why all of the planets in the solar system orbit in roughly the same plane—a fact known to humans since prehistoric times. The theory predicted that all of the planets would orbit the Sun in the same direction, as was also known. The nebular theory explains these trends as a consequence of the flat disk from which the planets formed and its rotation around the Sun. Subsequent discoveries in the solar system, like the 1781 observation of Uranus and the later discovery of Neptune in 1846, held true to these predictions.

The nebular theory also explains why Jupiter emerged where it did and how it got so big. High temperatures in the inner solar system prevented volatile compounds that are rich in hydrogen from condensing to make planets. Those same materials could condense out where the giant planets are located, giving more food for the growing Jupiter to feed upon. This also explains why Earth, Venus, Mercury, and Mars formed when and where they did, and why they are made from heavier stuff. The nebular theory had an excellent track record when it came to explaining what we saw in the solar system.

Until recently, virtually all of our knowledge about planet formation was based on observations of the solar system, since that is all that astronomers were able to see. Nevertheless, nothing in the story of our origins was deemed to be particularly unique to the conditions surrounding our Sun. We expected the essential elements of our theories to apply almost everywhere, and there was every reason to expect planetary systems much like ours to form around distant stars, since they would be subject to the same physical principles. The discovery of 51-Peg showed that perhaps these carefully crafted rules may not apply to other stars after all—that either their formation, or subsequent dynamical histories, diverged from the prevailing paradigm.

Many of the early discoveries of exoplanets ran counter to key predictions of the nebular theory. Planets with the mass of Jupiter, and presumably made of the same gaseous material as Jupiter, were supposed to form out in the frigid hinterland of the system, not right next to the star, where it would be subjected to intense radiation. Yet as new exoplanet discoveries accumulated, they continued to defy explanation. Even today, thirty years after they were first seen, we are still trying to understand the origins of hot Jupiters.

Exoplanet discoveries, while concentrated in the last few decades, were enabled by several technological advances that occurred over the last few centuries. One instrument, in particular, that emerged as a capable workhorse was the spectrograph. First employed in the late 1800s, spectrographs take light from a distant source, and break it into its spectrum of constituent colors, like a prism spreading the light from the Sun into a rainbow. When mounted on a telescope and pointed at a star, the spectrum of light from that star similarly spreads into its array of colors.

We can learn quite a lot about stars with a spectrograph because, when light from either the Sun or a distant star passes through the upper layers of the star's atmosphere, it shows a pattern of dark spaces or gaps called *spectral lines*, where specific colors are missing. These spectral lines arise because of the different chemical elements in the stellar atmosphere. The structure of the atoms of an element, or the structure of the molecules of a compound, has a unique set of energy levels for the orbiting electrons. (Imagine each element or compound having a ladder of states, where the positions of the rungs of the ladder are different for each substance—like a fingerprint for that material.) The elements in the stellar atmosphere absorb the specific wavelengths of light that correspond to those energy levels.

This absorption causes the dark spectral lines, and allows scientists to measure the star's chemical composition. For example, in 1868, Jules Janssen and Joseph Lockyer independently used a spectrograph to discover a new element in the atmosphere of the Sun. After shining the Sun's light through their instrument, Lockyer identified the fingerprints of several known elements, along with one set that hadn't been seen before. He named the unknown substance responsible for this new fingerprint "helium" after the Greek word for the Sun—"Helios."

As spectrographs improved, astronomers used them to study the properties of distant stars—learning about their composition and how they compared to the Sun. Spectrographs became so precise that new, more subtle signals could be gleaned from the stellar spectra. If the star was moving relative to the Earth, then the wavelengths of light would be stretched or compressed by that motion. This is exactly the same phenomenon, the Doppler effect, that causes the sound from a passing siren to produce a high pitch with shorter sound waves when it approaches, and a lower pitch with longer sound waves when it recedes. Here, using a spectrograph to measure the stellar spectral lines, and comparing those lines to what we see in a laboratory setting, we can measure how they are stretched or compressed by the Doppler shift. This shows how the stars move along our line of sight, or their *radial velocity*.

Initially, we saw only the average, large-scale motion of the stars in the galaxy, but as these measurements became more and more precise, astronomers started to see slight variations in the speed with which stars moved—periodic shifts superposed on their otherwise constant motion. The cause of these variations was the presence of companion stars that orbited the primary star. As the two stars moved about each other, the Doppler effect would cause the spectral lines from those stars to periodically shift from longer to shorter wavelengths and back.

Astronomers had long known that many stars in the sky were actually multiple stars that orbited each other—they could trace the orbits of the stars in the sky. Some stellar pairs take centuries to circle each other, and the wide separation of these stellar binaries could readily be seen using the telescopes of the late 1700s and early 1800s. Now, with the spectrograph, astronomers could study the orbits with much greater precision. Stellar pairs were soon found whose orbits were only a few days—orbits too close together to see with a telescope alone. This type of system was unknown until the spectrograph uncovered its existence.

As the precision of our spectrographs improved, our ability to measure the Doppler effect caused by the orbits of smaller and smaller objects also improved. It reached a point in the 1950s where the Russian–American astronomer Otto Struve proposed using a spectrograph to detect planets orbiting distant stars. Planets are thousands of times less massive than stars, so the Doppler signal from one would be a thousand times smaller than what was seen for binary star systems. In a 1952 article in *The Observatory* entitled “Proposal for a Project of High-Precision Stellar Radial Velocity Work,” Struve noted that the cutting-edge spectrographs would be capable of detecting Jupiter-mass planets if those planets happened to orbit close to their host stars, with orbits of roughly one day [2]. Planets that massive and orbiting that close could cause a Doppler signal large enough to be seen as periodic shifts in the wavelengths of the dark lines of the stellar spectrum.

His claim was largely true on the technology side, but the premise of the proposed observations was shaky on theoretical grounds. Everyone "knew" that Jupiter-like planets wouldn't exist that close to the host star. The nebular theory predicted that they would form at large distances, where it was cool enough for hydrogen-rich compounds like water, methane, and ammonia to condense. Our theories gave no reason to expect a Jupiter-mass planet to orbit its host star in one day. Even if we were able to see its tiny Doppler signal, it shouldn't be there in the first place. Struve didn't propose a mechanism to form such massive planets on these short orbits. He simply argued that since some binary stars were seen to have orbits that small, planets might have them too. He turned out to be right—such planets indeed exist, notwithstanding the predictions of planet-formation theory.

It would take a few decades, but eventually our observations caught up to Struve's imagination, and his approach produced most of the first exoplanet discoveries. His speculative paper has been cited over a hundred times in the scientific literature—over half of them in just the last ten years, and all but six of them since the search for exoplanets began picking up steam in the 1980s. By the late 1980s and early 1990s, there were a few telescopes around the world with spectrographs powerful enough to measure a star's motion to within a few meters per second.

This is an incredible level of precision. Imagine being able to observe a professor pacing back and forth in front of a class ten miles away, and measuring the speed of the motion by looking at changes to the color of the laser pointer in their hand. Or imagine looking at a star that is a hundred times the size of the Earth, located a trillion miles away, and seeing it move at the speed of a person leisurely strolling down the sidewalk. These instruments are pretty sensitive, and coupled with a few years of observations were good enough to start harvesting the low-hanging fruit—systems that cause the most significant Doppler effects on the host star. For exoplanets, that fruit is giant planets on short orbits.

Throughout the 1990s, following closely on the heels of the discovery of 51-Peg, the Swiss were joined by a collection of American astronomers from California, Texas, and Massachusetts (all working independently). Planet discoveries started gracing the pages of scientific journals and newspaper headlines. Each new planet was more weird than the last. Some were several times more massive than Jupiter, and barely classified as planets. Others had highly elongated, *eccentric* orbits that plunge toward their host star, passing within distances only a few times larger than the star itself, before being flung back out to the hinterland. Had these planets been in our solar system, they would have crossed the orbits of all the inner planets—eventually smashing into them or ejecting them from the system altogether. On top of this madness, all these newly discovered planets were orbiting too close to their stars.

The observations were like a parade of counterexamples to the predictions from planet-formation theory (namely that small planets should be near the host star, with large ones more distant, and all on circular orbits). Clearly, the universe of possibilities was larger than anticipated. Despite these discoveries, astronomers were loath to throw away the nebular theory, not because of some nostalgic affection for it, but because it did such a good job matching the solar system, and because it only relied on a few, unremarkable assumptions. The theory's assumptions were so generic, and predictions so straightforward, that they should be seen virtually everywhere—except, apparently, everywhere we looked.

As the number of exoplanets mounted, each piece of information we could extract was certain to provide new insights into their formation and subsequent history, and how their past was different from the solar system. Some of the most valuable information about planets is the sizes and shapes of their orbits, and the sizes and masses of the planets themselves. The Doppler measurements of these systems can determine the orbital size and shape, and also can give a rough estimate for the planet masses. But, in order to really compare exoplanets with the planets in the solar system, we needed to measure their sizes—specifically, their radii—and spectrographs simply can't provide that information. We need a different kind of measurement to do that.

A planet's radius, if it can be found, gives you a lot of information about the planet. You can find the planet's volume, and with the volume you can determine its density. The density measurement gives insights into the types of materials that compose the planet. For example, a large volume for a planet with a given mass implies that it is made of lighter material. On the other hand, a small volume for the same mass requires more dense material. Low-density planets, like Jupiter and Saturn, would be made primarily from gases or light elements. High-density planets, like the Earth or Mercury, would be made from rocks and metals.

Despite a variety of potential detection methods, distant planets are just too small and are hard to see outright. This limits what we can learn about them since, in order to know what planets are made of, we still need a way to determine their sizes. In the year 2000, we got what we needed. Two teams of astronomers observed a planet that happened to pass in front of its host star. That is, the planet *transited* the star. As it passed, the less-than-memorably-named HD 209458b, blocked a portion of the starlight. This planetary transit was a small signal—the star changed its brightness by only one percent as the orbiting Jupiter-sized planet swept across the stellar disk—but it was an unmistakable blip that stood out from the ever-present noise that appears in any observational data.

Planetary transits, while rare, happen regularly in the solar system. In the early 1600s, Johannes Kepler published his work on the orbits of the solar-system planets. He predicted in 1608 that both Venus and Mercury would transit the Sun, and would be visible from the Earth. It was around the time of his prediction that the telescope was invented—precisely the type of instrument that would enable this kind of observation.

Within a few years, Mercury was seen to transit the Sun for the first time, followed shortly thereafter by a transit of Venus. Seeing Venus transit was fortunate timing, since it would not transit again for another hundred years. Mercurial transits happen about once per decade. Venusian transits happen in pairs, about once per century, with the most recent pair being in 2004 and 2012. (The next time will be in 2117—so if you want to see it when it happens next, be sure to eat well and exercise.) Observing these transits of solar-system bodies based on Kepler's predictions was a remarkable feat. While humans had known of both Mercury and Venus for eons, tens of thousands of years passed before anyone saw them silhouetted against the Sun.

Despite having observed planetary transits in the solar system for nearly four centuries, the transit of HD 209458b in the year 2000 marked the first time we had seen this effect outside the solar system. These transits allowed us to measure the radius of that planet because, during the transit, the amount of light it blocked was equal to the relative sizes of the star and the planet. Given the size of the star, which we can estimate from computer models, we can determine the size of the planet that is transiting that star—the crucial piece of information missing from Doppler and other measurements.

This transit measurement was a big deal. It opened a completely new arena for investigation. By combining both the size measurements from the transits, and the mass measurements from the Doppler shift, we can get an idea of what materials compose the planets—whether the planets are rocky or gaseous, or somewhere in between. This new transit information allows direct comparisons between the sizes and masses of solar-system planets and the newly discovered extrasolar-system planets that are orbiting distant stars.

Astronomers had considered transits as a possible exoplanet signature for decades—indeed, Struve's 1952 paper outlining the use of spectrographs to find planets also mentioned planetary transits. However, as with any detection method, there are technical challenges one must overcome. For one thing, as we saw with HD 209458, the signal is quite small. A planet like Jupiter, which is one-tenth the size of the Sun, only blocks about one percent of a star's light (a planet with one-tenth the radius of the star would have an area that is a hundred times smaller). This relationship makes the transit signal much smaller for a smaller planet. The Earth, for example, is ten times smaller than Jupiter and would block only one-hundredth of one percent of the light from the Sun. This fact makes finding Earth-sized planets particularly hard because there are other astrophysical effects that muddle the search. There are spots on the Sun's surface that have a larger effect on the Sun's brightness than a transit of the Earth would have. Nevertheless, despite how small this change in brightness would be, from the 1960s through the 1980s transits were deemed a viable method to find planets given the technological capabilities of the time—except for another small problem with this method: we would need to be exceptionally lucky.

Detecting a transit requires that we somehow catch the planet in the act. The chances of having the orbit of a distant exoplanet coincide with our line of sight is already small. A random view of the solar system would line up with the Earth's orbit only about one percent of the time. For Jupiter, whose orbit is five times larger than the Earth's, a chance alignment would only occur one-fifth as often. The unlikely geometrical alignment is only the first strike against this method. We must also consider that Jupiter-like planets (the planets easiest to see) in Jupiter-like orbits (the orbits where our theories say we should find them) only transit once per decade, and for only a few hours.

A transit of a Jupiter-like exoplanet is an incredibly rare event, and designing a campaign to find even a single example led to some prohibitively low probabilities. To realistically capture a planet like Jupiter in the act of transiting its host star, you need to choose a target star. You hope that it has a planet. You hope that the planet is in one of the 0.2 percent of orbits that will geometrically align with the Earth and therefore will be seen to transit. Then you stare at it, without blinking, for ten years, looking for a one-percent change in the star's brightness that lasts a few hours. Sound fun?

Fast forward to the late 1990s, where the discovery of hot Jupiters significantly changed the detection probabilities. Hot Jupiters are much more likely to transit (about ten percent of the time instead of two-tenths of one percent). They also transit much more frequently, roughly once per week instead of once per decade. Had their existence been known twenty years earlier, astronomy in the 1980s and 1990s may have looked much different. But these kinds of planets weren't known until we unexpectedly stumbled across them, so the solar system was our only example to work from, and the properties of the solar system drove both the expectations and the design of campaigns to find exoplanets.

In the 1970s, the NASA Ames Research Center in California's Bay Area, held a series of seminars to discuss the possibility of finding distant planets. Here, the idea of a large brightness, or *photometric*, survey for transiting exoplanets drove a NASA engineer, William Borucki, to investigate the possibility by examining both the pitfalls and the potential of the approach. Having previously worked on the Apollo missions, he felt this was a good challenge to accept in their aftermath. In 1984, he published his first paper on the subject, and began a series of workshops to identify the best technology to use as the primary detector [3].

There are several devices that can be used to measure the brightness of stars, but these *photometers* tend to be bulky, they often need a fiber optic cable to connect them to the telescope, and they usually have to be cooled to reduce the noise in the data they produce. Observing a lot of stars would result in a Medusa-looking system that would be nearly impossible to make work. However, there was at least one new technology developed throughout the 1970s that appeared promising—the digital camera.

Digital cameras use charged-coupled devices, or CCD chips, which are arrays of tiny semiconductor pixels. Each pixel can store electrons that are dislodged from the surrounding material when photons of light strike the surface. Those electrons are then "read out" and counted in order to determine how much light arrived at each pixel's location during the observation. With large CCDs, you can watch lots of stars continuously, and do so for a really long time. CCDs, if used to search for exoplanet transits, would allow you to make up for the rarity of a transit by providing a way to study more stars at once.

Borucki suggested that a digital camera, attached to a large telescope and lofted into space, could continuously observe tens of thousands of stars, recording their brightness every few minutes. Being in space meant that the confounding effects of the atmosphere would be eliminated, and he could detect the transits of smaller planets. Smaller planets were expected to be on smaller orbits, so he wouldn't need to rely on chance transits of the slow, cumbersome orbits of larger gas giants—he could look for Earth-sized planets that had one-year orbits. Instead of needing to observe for several decades, he would only need to look for a few years. If he observed enough stars at once, the sheer number of targets would overcome the small probability of a chance alignment, and the continuous monitoring—which is more easily done from space—would ensure that none of the transits slipped through the cracks.

Like most new ideas, people thought this was crazy. Despite the success of small projects to demonstrate the capabilities of the technology, there were regular calls to put a stop to Borucki's work. After all, this was five years before the launch of the Hubble Space Telescope, ten years before the first exoplanets were discovered, and fifteen years before one was seen to transit. To many, his work was clearly a waste of resources, with one NASA lawyer even questioning its legality. He was given one last opportunity to make his case. NASA assembled a panel of experts to review his idea—chaired by Jill Tarter, who was known for her work on searches for extra-terrestrial life. William was told to either convince them, or to quit. When he finished, a few of the panel members, notably including Jill Tarter and Gibor Basri (a stellar astrophysicist from Berkeley), joined his team.

Despite the green light to continue working on the idea, he faced another challenge from the fact that there was no good way for NASA to fund the whole enterprise. At the time, NASA basically had small "explorer" missions, and large, strategic or "flagship" missions—there was no middle class for ideas of the size and scope that William had in mind. Without the ability to apply for the right-sized pot of money, his idea couldn't get further off the ground than small research projects to test various technologies. However, this state of affairs changed in the 1990s when NASA unveiled its *Discovery Program* for mid-sized missions. This was a program where the available funding matched what would be needed for Borucki's idea. So the group of fellow travelers that he had accumulated over the preceding half decade started their design of the FRESIP mission, an abbreviation for Frequency of Earth-Sized Inner Planets, which would hopefully fly as one of these new discovery-class missions [4, 5].

As with most scientific proposals that appear on NASA's desk (speaking from experience), this one was rejected. It was too risky.

There was hardly any justifying ground for the FRESIP mission to stand on. At the time, people thought that, while they may exist, planets were probably rare and the best ones to find had longer orbits than the duration of the proposed mission. Besides, the technical challenges to detecting an Earth-sized planet orbiting a Sun-like star in a one-year orbit were substantial. How could you know the camera would function well enough to distinguish such a small signal? How could you know that there would be enough planets out there to find the one in a hundred that would pass across the line of sight? How could you be certain you weren't just looking at star spots or other astrophysical events?

The reviewers of the proposal indicated that it would have been the highest-ranked proposal, but they had serious doubts about the untested CCD technology. Before NASA would take the proposal seriously, William had to show that it was possible to do what he proposed. He kept shopping FRESIP around to scientists and engineers both at NASA and at other institutions, adding new people with new expertise to the team, while NASA continued to demur. Over time, a handful of discoveries from different parts of the world would slowly change William's fortunes.

In the early 1990s, Aleksander Wolszczan and David Frail were using the giant Arecibo observatory in Puerto Rico to study pulses of light emanating from a rapidly rotating corpse of a dead star, a so-called millisecond pulsar. Wolszczan found that the timing between pulses coming from the star was changing—spreading apart, then coming back together, then spreading apart and coming back together again. This should not be. Pulsars, especially millisecond pulsars, emit exceptionally stable signals by astronomical standards. In fact, they are regularly used as astronomical clocks. (A set of fifteen pulsars, each with its spin period indicated, was used as a way to pinpoint the location of the Earth on the famous plaques that were attached to the Pioneer and Voyager spacecrafts. Pulsars were also used to discover distortions of spacetime that span huge distances across the galaxy by watching for small changes in their pulsation frequencies.)

While the pulsation periods of pulsars do drift over time, that evolution is gradual and (except under specific conditions) it always slows down. The observed periodic variation was something not seen before, or predicted. Plus, it wasn't just one variation. The time between pulses was fluctuating on two different timescales. Attempts to construct a theory to explain these variations pointed in one direction—planets. Wolszczan had discovered a system of two planets orbiting this pulsar, the first planetary system seen outside the solar system [6].

These pulsar planets were found in 1992. Then came the discovery of 51-Peg in 1995, the first exoplanet orbiting a Sun-like star. This distinction is an important one, as the planets that orbited the pulsar likely formed after the demise of the progenitor star, while the planet orbiting 51-Peg has a formation history that compares more directly to our own. The pulsar planets notwithstanding, 51-Peg is typically hailed the first exoplanet, as seen by the fact that its discovery was awarded the 2019 Nobel Prize in Physics. Nevertheless, as strange as they were, the pulsar-orbiting planets broke a lot of ground for the field of exoplanets.

The pulsar planets also boosted the motivation to look for other planets. William and his ever expanding team (this time including the great Carl Sagan) revised and submitted the FRESIP proposal again in 1994. The rejection this time was because it was deemed too expensive. Then in 1996, rejected, but at least it was "highly meritorious." In 1998, rejected.

With each rejection came a new list of questions and concerns to be addressed. William and his team needed to show was that it was possible to search for planetary transits by measuring the brightness of tens of thousands of stars simultaneously. With that many targets, spread over a wide field, all sorts of false signals could creep into the data. The mission was supposed to simultaneously look at roughly a hundred thousand stars and find changes in brightness of fewer than a hundred parts per million on each one—which, when they happen, only lasts for a few hours each year. What William was proposing to do would be like looking at Las Vegas from space and detecting a fly buzzing around a streetlight.

To show that this task was possible, he and a group of collaborators made a small, wide-angle telescope at the Lick Observatory in the mountains east of San Jose, California, and pointed it skyward. This *Vulcan* survey didn't find planets, but it did find hundreds of binary stars that eclipsed each other—a signal similar to a planetary transit. Vulcan showed that the team could take and analyze brightness data for the large number of targets expected for the ambitious space telescope. (A copy of the setup was later built in Antarctica, where the long winter night would provide good conditions for continuous observations.)

The last major request from the NASA reviewers was to show that the proposed telescope and instrument design was sensitive enough to see, under realistic observing conditions and with a realistic instrument, the tiny brightness change that would result from the transit of an Earth-sized planet in a year-long orbit around a distant, Sun-like star. NASA needed a successful end-to-end trial that included the drift that would occur while the spacecraft flew, readout noise from the electronics, varying operating temperatures, flight software, and more. In short, William needed to show that every link in the chain, from the photon detection to the analysis software, would work before NASA would be willing to give him the half-billion dollars he was asking for.

A Kepler test facility with all of the trappings was built by Ball Aerospace, the industrial partner who would ultimately build the spacecraft if it were funded.[1] Here, the mission's second-in-command, the Deputy Principal Investigator David Koch (not of wealthy industrialist fame), devised a clever experiment to show that the proposed mission had the necessary sensitivity to accomplish the task at hand. To prove his point, he needed to make an artificial starfield that matched the brightness of some of the dimmest targets that their mission would observe. After that, he needed to somehow change the brightness of these fake stars by only a few parts per million over the course of several hours. Finally, once he had those fake stars and fake planetary transits, he needed to show that the camera could see the tiny effect. (Plates 1 and 2 show William Borucki and David Koch.)

---

[1] As a side note, it was only in 1993 that Ball spun off its century-old canning jar business to focus primarily on aerospace.

They made the starfield by drilling holes in a piece of metal and illuminating it from behind. To mimic the tiny change in brightness from a planetary transit, they stretched a thin wire across the holes. However, it wasn't the presence of the wire itself that made the synthetic planetary signal. Instead, they caused the wire to change size—by a few parts per million—by passing an electric current through it. The current made the wire heat up and expand slightly for each degree it rose in temperature. The expanded wire would block more of the light that would otherwise pass through the holes—producing the small change in brightness that was equivalent to a transiting, Earth-sized planet.

With this setup they ran the test to show whether or not the camera, along with the rest of the hardware and software, was up to the challenge. The camera passed. The collaboration resubmitted their proposal and the mission formerly known as FRESIP, but renamed in 1996 to Kepler, was selected by NASA in December of 2001. It would be the tenth mission for NASA's Discovery Program. Over the next several years, William and his team worked to build the Kepler space telescope and prepare it for launch.

This good news for the Kepler team arrived only a few months after I began my graduate-school career. While they were busy putting their mission together, I was bouncing around between different research groups, and working toward graduation. All the while, planets kept surfacing in the observational data. By the early 2000s, there were enough new results to cause people's heads to turn. More and more astronomers were paying attention to these new celestial objects, including my graduate advisor and me.

With the discoveries of transiting exoplanets, beginning with the detection of HD 209458b, my advisor believed there might be more we could learn about them than just their sizes and orbital periods. Given transit measurements with sufficient precision, he thought we could uncover information about additional planets in the system that other instruments might miss. He went on: If the planet were alone, with no sibling planets, then the planetary transits would occur at regular intervals—once each orbital period. But, if there were a second planet in the system, the two planets would interact with each other gravitationally. These interactions would cause small changes to each planet's orbital period. The changes in the orbital period would show up as deviations in the time interval between successive transits. Sometimes one planet would cause the other to transit earlier than normal, and sometimes it would cause the transit to be late.

He wondered how large these "transit timing variations" might be, and how they might change depending upon the properties and orbits of the other planets in the system. Perhaps we could measure the masses of the planets without the need for expensive, cutting-edge spectrographs (the largest spectrographs can cost millions of dollars to build, and nearly ten thousand dollars per hour to operate, after all). It might be possible for us to detect and characterize unseen planets—ones that didn't transit, or were too small to see or to be deduced with spectrographs. Maybe we could directly measure the size of the star so that we didn't need to rely on computer models when estimating their size and the sizes of the planets that orbit them—as was the case up to that point. Perhaps we could constrain some details about the shape of the planetary orbits. All of this information might come out of an analysis of the variations in the transit timings. My "practice problem," while we waited for me to complete the general relativity course I needed for the black-hole-imaging research, was to explore these possibilities.

We got started working through the pages and pages (and pages) of algebra. Upon finishing a calculation, we did it again, and again. A constant challenge with this line of work is that a single misplaced minus sign, or one mistake in a summation, and the whole calculation is wrong. After a few months we convinced ourselves that we had done the math correctly. We wrote some computer models that matched our calculations, adding to our confidence. As we worked, something became increasingly apparent. The mutual gravitational interactions of the planets in the system would cause changes to their orbital periods that were surprisingly large. Under the right circumstances, a small planet like the Earth could change the orbital period of a large planet like 51-Peg or some other hot Jupiter (which are over three hundred times heavier than the Earth) by minutes, or even hours. That was a big signal.

With the existing, ground-based technologies, astronomers could measure the moment that a planet transited its host star to within about one minute, so a deviation of an hour would be easy to spot. Through our work, we had indeed uncovered a way to find hidden planets, and our method was exceptionally sensitive. These transit timing variations, soon shortened to the acronym TTVs (the creation of which may be my greatest contribution to exoplanet science), were possibly the best method available to find Earth-mass planets orbiting Sun-like stars. No other options available at the time really compared to TTVs—neither Doppler measurements nor existing transit measurements could find such small planets. We submitted our initial paper on this method for peer review in late 2004, in conjunction with a similar study from Harvard, feeling good about the likely value of our work [7, 8, 9]. To use this technique, however, we would need the right data to analyze.

More than a dozen transiting exoplanets had been seen by this time, but the kind of data that existed in the scientific literature wasn't good enough for us to make use of our methods. Most systems had only one or a few transit measurements, and we needed a series of lots of measurements. Just as we wrapped up our initial work, a group of astronomers published nearly a dozen observations of a transiting planet—a planet called TrES-1b. These were several, good-quality transit times for a single system, spread over a large time interval. They were precisely the data we needed to try our hand at finding unseen planets with TTVs.

With our paper undergoing peer review, and likely to be accepted, we had a decision to make. Do we keep pursuing research involving exoplanets and TTVs? Or do we turn our attention back to imaging supermassive black holes?

When a theorist in physics or astronomy announces a potential signal, such as ours, they often need to wait for years or decades before adequate data are available to search for that signal. The data on TrES-1 appeared at just the moment when we were ready to use them. My advisor said to me, "It isn't very often that the data you need fall in your lap." With these data and our methods, we had a chance to be the first people to detect an Earth-mass planet orbiting a distant star. It was an opportunity not to be missed. With that, we dropped our plans to image black holes, and shifted our efforts entirely to exoplanets. My so-called practice problem became the subject of my PhD dissertation, and over the next two years I would build the methods and software that I needed to analyze transit data. I could then apply them to transiting exoplanet systems, looking for hidden planets in those systems as new observational data started to appear. (For the record, in April 2019 a group of scientists did take that black-hole image using the Earth-sized Event Horizon Telescope, but by then I was far down the exoplanet rabbit hole.)

Now the stage is set. By 2006, planet-formation theory was in shambles as the discoveries of the previous decade kept defying its predictions. In order to understand what was happening, the scientific community needed more data to shed light on the issue, and it would require more than just a few new discoveries here or there—they needed a large, comprehensive survey. After two decades of development, Kepler had been selected for construction and flight. The growing Kepler Science Team now included many of the scientists whose exoplanet discoveries had made the headlines of the previous decade. The engineering team was actively building and calibrating the instrument. The software engineers were coding the analysis tools needed to find planets in the flood of data expected from the spacecraft. And the astronomers were working feverishly through the enormous pile of preparatory work required to make the mission succeed—cataloging, characterizing, and selecting the stars that the telescope would observe.

The mission was also looking for new blood to join the science team. They wanted people whose research would provide novel insights into the discoveries from the Kepler data, or perform some analysis that wasn't envisioned when Kepler was first designed nearly two decades in the past. This request included a specific reference to the "detection of non-transiting planets by timing of the variations of the transit epochs." I had gone to the University of Washington, where my career preconceptions were dead on arrival, but I had just defended my practice-problem-turned-dissertation, “Detecting New Planets in Transiting Systems,” by analyzing the timing of the variations of the transit epochs caused by the mutual gravitational interactions of planets in those systems [10]. My research had been about analyzing precisely the kind of data that Kepler would produce, and at that fleeting moment, I was the only person on Earth who had done so—twice. In short, I won the scientific lottery.

NASA had just asked for someone like me to join their science team. I was a newly minted PhD, working at a national laboratory that let me submit a proposal to the competition. And, when the powers that be at NASA accepted my proposal, I was given the opportunity of a lifetime to work with the best colleagues, in the most exciting scientific field, on the hottest mission in the world. Kepler would revolutionize astronomy, I had a seat on the ride, and was about to have the trip of a lifetime.

INDEX