Tidepool

A disruptive, open approach to diabetes care.

A Motivation to Disrupt

Steve McCanne and Howard Look are “D-Dads with Attitudes.” Both have 13 year old daughters with Type 1 Diabetes. Both are fed up with the sorry state of diabetes care technology and have a vision to disrupt the industry as we know it. Both are successful technology leaders and entrepreneurs: Steve was CTO at Inktomi and founder of Riverbed Technologies. Howard was VP of Software and User Experience at TiVo, Pixar and Amazon. Steve invented core technologies that power today’s Internet and provide solutions to optimize network performance. Howard invented disruptive experiences that forever changed how people watch TV.

In founding Tidepool, Steve and Howard are joined by Dr. Saleh Adi, Director of Pediatric Diabetes at UCSF. Dr. Adi and his team believe that the use of technology is essential for managing diabetes. Their patients are encouraged to use insulin pumps and continuous glucose monitors (CGMs) and to share their data with their clinic providers frequently. They share the vision for a better way to manage diabetes, and bring the clinical expertise to ensure that Tidepool’s offerings are effective and safe.

Also joining Tidepool are other world-class team members with a motivation to disrupt and make a difference to people with type 1 diabetes. Please read on...

The Elevator Pitch

We are building the “Mozilla of Diabetes Platforms.” We will disrupt the current model of diabetes technology and care by delivering a full, end-to-end software stack using an open development model. Our first application, “blip,” enables more accessible and intuitive visualizations and more fluid patient/doctor communication and insights. Other applications include remote monitoring, smart meal bolusing and “Endo in the Sky” recommendations. Our platform will also enable improvements in Artificial Pancreas development, helping researchers to focus on algorithms, trials and effectiveness; we'll handle the middleware, cloud monitoring and visualization software.

The project will be developed in the open on GitHub. The platform and code will be available for free to non-profits and non-commercial researchers. Within three years we will generate self-sustaining revenue through licensing fees to commercial devices makers; from the other side, patients and clinics will demand that their devices work with our platform. We may also generate revenue from payer subsidies or clinic subscriptions. We will be agile and pivot as necessary to figure this out.

Howard Look, howard@tidepool.org

The Problem

The problem to be solved is clear to everyone who has diabetes or who has a child with diabetes: managing diabetes sucks. It is a 24 by 7 grind of finger pricks, insulin doses, carb counts, dosing for meals, basal rates, correction factors, pump settings, infusion set changes, adjusting for exercise, travel, illness, and much more.

A quick note for those not aware: Type 1 Diabetes is an auto-immune disease caused when the immune system attacks insulin-producing cells in the pancreas. "T1Ds" cannot make enough of their own insulin and are dependent on external insulin to live. Type 2 Diabetes is brought on by resistance to insulin, associated with many factors including poor diet, obesity, lack of physical activity, and genetics. About 30% of T2Ds are also treated with insulin. All of Tidepool's solutions will apply to both groups, but we are initially motivated by the more acute and complex needs of T1Ds. With little or no pancreatic function, treatment for T1D is to "act like a pancreas," relying on as much data and feedback as possible to make treatment adjustments on a regular basis. T2D typically does not have such acute pressure.

Managing diabetes is daunting, time-consuming, and cumbersome. Parents of kids with T1D (like Steve and Howard) live in a constant state of helping their kids to stay "in range," keeping their blood glucose ("BG") between about 80 and 150 (normal for non-diabetics is 70-110, and up to 140 after meals). It is said that "going high kills you slowly." Persistent hyperglycemia leads to nerve damage, blindness and kidney failure. But "going low kills you quickly" – hypoglycemia can quickly cause seizure, coma and death. Many parents of T1D kids get up every night during the middle of the night to check BG levels to ensure their kids are in a safe range.

Managing insulin doses is hard work and is both a proactive and a reactive affair. Every meal or snack requires a calculation and insulin dose. Time of day affects insulin needs – day time, night time, late evening, midnight, after midnight, early morning – all require different dosing schema. Exercise, illness, fever, medication, stress, immunizations, and menstrual cycles all effect insulin doses. And then, what works today must be tweaked again in a few days, weeks or months as both children and adults grow and change. Patients and parents must collect as much data as possible about everything they do in their daily life – food, exercise, sleep patterns, work/school hours – and send the data to their doctors to analyze, interpret, and recommend changes. Most of this is done manually and intermittently. People with diabetes (PWDs) already do so much to manage the disease on a daily basis. Asking them to do more is not practical.

Doctors, patients and parents do their best with the technology that exists today, but the experience makes everyone cringe. Clinics see their PWDs for 15-30 minutes every three months, and rely on messy charts from disparate data sources like insulin pumps, blood glucose meters (BGMs) and continuous glucose monitors (CGMs). While the data from devices can be uploaded, they are presented in widely varied formats that are often not intuitive. There is no standard for data integration from multiple devices.

Current diabetes technology is developed by largely stodgy, old-school companies that are hampered by their existing business models and products lined up for long FDA cycles. Today's products are closed, proprietary systems that are often the antithesis of "user friendly" and "patient centered." Most companies hoard the data generated by their devices to protect their ecosystem.

To add insult to injury, only about 20% of T1Ds use insulin pumps[1], and fewer than 10% use continuous glucose monitors, despite clear evidence that these devices lead to fewer complications and healthier outcomes.[2] [3] Of those prescribed pumps, about half use default or sub-optimal pump settings.[4]

Tidepool is turning this existing model and state of affairs on its head.

The Solution

Tidepool will develop and deliver an **open, cloud-based platform** that provides:

- Secure, validated patient data using well-known cryptographic techniques.
- APIs for social network integration, so PWDs, parents and caregivers can interact in a fun, productive and ongoing way.
- An open architecture that supports extensions for data visualization and machine learning.
- RESTful APIs for integrating data from existing devices and data repositories.
- RESTful APIs for transferring data to and from existing Electronic Health Record (EHR) systems.

We will also establish an **open source development organizing body** that leverages a vibrant, motivated development community, including:

- People with Diabetes (PWDs) and their loved ones.
- Members of the "Quantified Self" movement.
- Medical informaticists.
- Data visualization gurus.
- Mobile health enthusiasts.
- Proponents of "open health data."
- People who are demanding better diabetes care tools.

We will also create a new license for our source code is "corporate-friendly" like MIT and Apache, but requires that a patient's device data be kept secure, accessible and interoperable.

Finally, we will deliver **sample and reference applications** that are built on top of the Tidepool platform. These applications include functionality like:

- Intuitive, interactive visualizations of data from multiple sources, including insulin pumps, CGMs and activity monitors.

[1] http://chir.asu.edu/sites/default/files/PredictorsInsulin.pdf
[2] http://link.springer.com/article/10.1007%2Fs00125-012-2708-9
[3] http://www.ncbi.nlm.nih.gov/pubmed/22777524
[4] http://www.diabetesnet.com/pdfs/DiabTech2007Poster.pdf

- Lightweight, ongoing conversations between health care providers and their diabetes patients using Facebook and other messaging systems.

- A "smart meal" bolus app that remembers how the PWD reacted to that same meal last time and recommends trying a better insulin dose this time.

- Remote monitoring for parents of their T1D kids from cell phones and other devices.

- In-the-cloud machine learning and pattern recognition techniques from aggregate diabetes data, creating the "Endo in the Sky" that allows caregivers to make better and more frequent recommendations on insulin pump settings.

- A reference device implementation for the Artificial Pancreas remote monitoring system, helping to accelerate trials of the AP.

The first two items are encompassed in the first prototype Tidepool application called "blip." blip is currently being used in a small-scale test by some of Dr. Adi's patients, with a larger-scale trial under development. The value proposition for blip is simple: It helps patients improve blood sugar control while simplifying and improving their lifestyle. It helps clinics and hospitals deliver better care, more efficiently, with fewer errors at lower cost.

Over time, we also intend to deliver an **open device stack** that includes:

- Data importers for all popular blood glucose meters (BGMs), continuous glucose monitors (CGMs), and insulin pumps.

- Reference code software and reference hardware implementations for Insulin Pumps, Continuous Glucose Monitors, and an Artificial Pancreas. These reference implementations will enable existing device makers to easily integrate Tidepool functionality into their products. Code will include working solutions atop commonly available hardware (e.g. Arduino, BeagleBoard and RaspberryPi), but will be architected with a hardware-agnostic abstraction layer.

A few analogies that may resonate:

- The Tidepool cloud platform enables functionality similar to Mint.com. It aggregates data from disparate sources and provides feedback and recommendations based on analysis and visualization of the data. The platform provides RESTful APIs that make it easy to integrate with other systems and create extensions, such as Facebook integration and "Endo in the Sky" recommendations from doctors.

- The Tidepool device stack is similar to the Android Open Source project. In the same way that Amazon, HTC and Samsung build products on top of AOSP, medical device makers may choose to build on top of some or all of the Tidepool device stack.

- The Tidepool development model is similar to that of the Mozilla Foundation. We will enlist the open development community to help us develop and extend the platform. "Area owners" will lead development area specialties, such as data visualization or machine learning. We will also have a core of full-time developers and UI designers.

Tidepool will deliver open source code and open reference designs, but ***we do not currently intend to perform human trials***. Trials will be conducted by adoptees of the Tidepool platform, such as UCSF, and may use Tidepool reference implementations or develop their own applications. Other non-profit and for-profit entities will build products on top of the Tidepool platform which they will submit for FDA and other regulatory approval. The Tidepool Platform will include extensive documentation and automated test harnesses which will help reduce the submittal burden. We also intend to establish a Master File with the FDA that sponsors can refer to with their submissions. We will also collaborate closely with sponsors to support their submittals based on our platform.

The Non-Profit Model and Opportunity

Tidepool will license its technology for free to non-profit, academic and research institutions. Revenue will be generated from these sources:

- Device makers will license the Tidepool platform because it makes their device more attractive. Patients and caregivers will demand it, driven by the social network effects and success of applications like blip and improved outcomes seen in trials.
- Payers will reimburse entities that use Tidepool-based devices and applications because they lower the ongoing cost of diabetes treatment and improve health outcomes over time.
- Going forward, Tidepool may also choose to deliver productized versions of some of its applications like blip.

Any revenue generated by Tidepool will help offset it's operating costs and reduce the need to do philanthropic fundraising. If revenues exceed costs, the difference will be used to further the mission.

Although we expect Tidepool to generate revenue and be self-sustaining, we choose to be a non-profit entity for several reasons:

- Tidepool is working on a mission towards a specific goal – improving the lives of our loved ones and other patients – rather than working towards profit or being beholden to ROI or liquidity expectations.
- The open development community is not motivated to assist for-profit entities, and often views their motives with skepticism. We want to leverage the talents of the broader community.
- The market is relatively small. We are initially focused on the unique and more acute needs of people with Type 1 Diabetes, which is a far smaller market than Type 2 Diabetes. While most everything we will do will also apply to the Type 2 market, and the broader medical care market as a whole, we will address the needs of Type 1 Diabetes first.

More detail on our non-profit mission can be found in our IRS Form 1023 for 501c3 exemption. (Submission is pending legal review, as of May 15, 2013.)

The Tidepool Advantage

We believe Tidepool will succeed for these reasons:

- Our founders and team are not in this for the money. We are in this because of our personal connection to type 1 diabetes and desire to make life better for our family members and other patients.

- Despite being a non-profit, we will operate as an Agile/Lean startup. We will pivot as necessary as we learn from patients, caregivers and partners. Since we have no product line or FDA approval pipeline to protect, we can iterate often and move quickly.

- We believe that great software and experiences come about by maximizing iterations in response to feedback. We will use simulators, broad automated test harness, and continuous integration and deployment practices to achieve this.

- The open community has already started down this road, and is highly motivated. There are many loosely coupled initiatives that can benefit from a more concerted, organized approach.[5] [6] [7]

- We believe that Design Matters. We will invest in user-centered UI design practices.

- A strong, collaborative relationship with UCSF (and more clinics in the future) will ensure timely validation through clinical trials and continuous feedback from clinical end users.

The Founding Team

Howard Look was on the founder's team at TiVo where as VP of Software and User Experience he led the efforts that made TiVo as easy to use as it was disruptive. He was also VP of Software at Pixar, where he led the team developing Pixar's proprietary film-making system, and at Amazon where he ran a secret project developing devices that leverage cloud services. At Linden Lab, he led the team that delivered the open-sourced Second Life Viewer 2.0 project. His 13 year old daughter has T1D. Howard will serve as President and CEO of the Tidepool.

Steve McCanne was CTO at Inktomi and founder and CTO at Riverbed Technologies. His 13 year old daughter has T1D. Steve has personally architected and coded large portions of the Tidepool platform and will serve part-time as CTO of the Tidepool. Steve and his wife Tami have a passionate commitment to helping the diabetes community, most notably through their anchor donation to the Madison Clinic for Pediatric Diabetes at UCSF and now through Steve's involvement with Tidepool.

Dr. Saleh Adi is a pediatric endocrinologist, diabetes specialist and Director of the Madison Clinic for Pediatric Diabetes at UCSF Children's Hospital and the UCSF Diabetes Center. Dr. Adi will serve as lead medical advisor for Tidepool.

Dr. Aaron Neinstein is an adult Endocrinologist at UCSF. He is a noted expert in the intersection between technological innovations and system improvement in healthcare. Dr. Neinstein will oversee the integration efforts with UCSF's existing EHR system.

---

[5] https://github.com/medevice-users/diabetes/
[6] http://www.opensourcediabetes.org/
[7] http://www.kickstarter.com/projects/nialg/the-diabetic-journal

Dr. Jenise Wong is an Assistant Professor of Pediatric Endocrinology at UCSF, a clinician at the Madison Clinic for Pediatric Diabetes, and a scholar specializing in clinical research in childhood Type 1 Diabetes. She will oversee design and implementation of clinical trials at UCSF.

Dr. Yao Sun is an Associate Professor of Clinical Pediatrics at UCSF and Director of Neonatal Clinical Programs. Dr. Sun also holds a PhD in Computer Science from MIT, and is a noted expert in the field of Medical Informatics. Dr. Sun is designing the Tidepool data analysis, machine learning and pattern recognition system.

Bryan Roberts, PhD, is a partner at Venrock and focuses on a broad range of healthcare investments. Bryan is currently Chairman of the Board of Directors of Achaogen, Castlight Health and Ironwood Pharmaceuticals. Bryan is serving on the Tidepool board of directors, along with Howard Look, Steve McCanne and Dr. Adi. Bryan also has a son with T1D.

Drs. Adi, Neinstein, Wong and Sun comprise the Tidepool Medical Advisory Board.

Funding Model

As we stated previously, we believe that it is important for Tidepool to operate as a non-profit entity. Our promise to funders is not measured in return on investment. The promise is that this team will deliver a disruptive technology platform and application experience, and that our efforts will lead to a healthier life less burdened by diabetes management for our loved ones and others.

That said, Tidepool will generate revenue. Any income beyond operating costs will be funneled back into Tidepool. There will be no stock or equity issued in Tidepool. The goal of generating revenue is to cover operating expenses and fund future initiatives.

Phases and Funding Ask

Phase 1: 2013 – Spring 2014 (please note: these are startup phases, not FDA trial phases)
Tidepool is currently raising $2M in seed financing as a non-profit entity. No equity will be issued. As of April 9, 2013, $500k has been raised from Lightspeed Venture Partners (thank you !!!). Additional funding is currently being solicited from other firms and individuals, many of whom have a close connection to Type 1 Diabetes.

During Phase 1, milestones include:

- The open source initiative will be organized and launched, joining forces with other current initiatives. An organized repository will be established along with a new open source license.
- The Tidepool platform will be established, including open APIs and data model.
- Our first sample app, blip, will enter internal clinical trial at UCSF.
- The platform will include minimal device support (Dexcom, Animas, Medtronic), enough to prove out a Minimum Viable Product and conduct the trial.

Funds will be used for operating expenses, including:

- Salary for staff for up to 12 full-time employees for about one year. Since we will not be issuing equity, we expect to pay higher than average startup salary to attract and retain top talent.
- Legal and other operating expenses.

- Communications, branding and web site development.
- Office space, utilities, bandwidth, cloud servers, etc.

We want to spend our time building a great team and code base, not on “tin cupping” as a non-profit. Therefore, we are asking for substantial investments ($500k to $1M) from a small base of donors. Assuming that Tidepool succeeds in delivering value substantially similar to the milestones listed above, we ask that participants in Phase 1 commit to at least doubling their participation in Phase 2.

Phase 2: Spring 2014 – Early 2015
For this phase, Tidepool will raise an additional $6-8M. Milestones during Phase 2 include:

- Launch event and Developer Day, coincident with the 2013 Diabetes Mine Tech Summit and World Diabetes Day.
- v 0.1 of the Tidepool “Big Data in the Sky” Machine Learning platform.
- v 0.1 of the Tidepool Device stack.
- At least one additional trial partnership beyond UCSF, likely Stanford, UCSB or UVA.
- Additional conduits added to the Tidepool platform for other devices, like FitBit and Nike Fuel Band.
- Additional conduits added and to the Tidepool platform for Artificial Pancreas development, allowing researches to monitor and communicate with patients remotely, see Artificial Pancreas telemetry data, and intervene when necessary.
- Launch one additional application based on the Tidepool platform, such as a smart meal bolus system.

Phase 3: Early 2015 through end of 2015

Based on the success in Phases 1 and 2, Phase 3 will raise an additional $3-5M, which will fuel revenue-generating partnerships and licensing agreements. By the end of Phase 3, Tidepool will be a cash flow positive, self-sustaining organization.

Phase 3 may include milestones like:

- blip v2 launched into multi-site trial.
- Remote monitoring and smart bolusing apps enter trial.
- EHR conduits added to the Tidepool platform.
- “Endo in the Sky” recommendations based on big data machine learning.
- NRE and Licensing agreements with partners like Medtronic, Dexcom, Animas and Tandem.

Technology Approach

We are building the Tidepool platform using using modern web and mobile design patterns that abstract away complexity. It's early, and there is still great opportunity for developers joining the effort, either as community developers or employees, to help us shape this architecture.

Design and architecture patterns we are using include:

- Apps may be mobile or web-based. Our current web apps are constructed using the emerging and popular "single-page" client-side JavaScript pattern, interacting with one or more backend services via RESTful APIs.
- Most (if not all) of the application logic is embodied in the client rather than in a complex, multi-tier service architecture. The application state is maintained in the client and any long-term application state is persisted to the backend through simple, RESTful interfaces.
- We leverage client-side JavaScript such as RequireJS, Backbone.js, and Bootstrap.
- Data is represented in JSON, and will be cryptographically signed and validated as a way to ensure data integrity as it passes through the system.
- We are currently using D3.js as a visualization engine, and envision building a "widget" library of diabetes-related visualizations that can be leveraged by other apps built atop the Tidepool platform.
- We currently use a node.js-based RESTful back-end, but are considering using a BaaS.