

Tempe: Live Scripting for Live Data

Robert DeLine, Danyel Fisher, Badrish Chandramouli, Jonathan Goldstein,
Michael Barnett, James Terwilliger, and John Wernsing
Microsoft Research
Redmond, WA, USA

Abstract—Data scientists are increasingly working with live streaming data, for example, business telemetry and signals from wearable devices and the Internet of Things. Unfortunately, current tools for exploratory data analysis provide poor support for streaming data. This paper presents Tempe, a data science environment for temporal and streaming data. Tempe’s extensible scripting environment allows for live programming, displays interactive, continually updating visualizations, and provides a uniform query language for both stored and live data. We discuss the streaming features of Tempe and evaluate our design choices with a deployment study at Microsoft with a product team who used Tempe continuously for six months.

Keywords—Programming environments; live programming; data analysis; data mining; data science; web applications; interactive visualization; streaming data.

I. INTRODUCTION

Data scientists are increasingly interested in analyzing streaming data in near real time, using platforms like Storm [1] and Spark Streaming [2]. Some streaming data analysis is behind the scenes, for example, software companies troubleshooting server performance by tracking telemetry. Other streaming data analyses form the functionality of products. Some examples include fitness applications that process real-time sensor signals from wearable devices and smart-home applications that combine real-time signals from multiple devices on the Internet of Things. All these examples require the ability to analyze both historical data (for example, to train a classifier) and real-time streaming data (for example, to apply the classifier to a live signal) [3].

Creating such applications requires exploratory data analysis over streaming data. For instance, a data scientist will not know a priori which classifier works best for a given set of signals, nor which features to use for training the classifier. Making these choices requires an exploration of the data. Unfortunately, existing data science tools do not provide much support for exploratory data analysis over streaming data, much less a combination of historical and live data.

This paper presents the design and evaluation of Tempe, an integrated environment for analyzing streaming temporal data. Tempe provides direct support for data exploration across live data streams. Tempe uses the Trill streaming engine [4] to provide a unified query language for both stored data (of arbitrary size) and live data. For both types of data, Tempe displays interactive visualizations to summarize the

results of the user’s Trill queries. Tempe also provides live programming to keep the script contents and results mutually consistent as the user edits [5]. To reflect changes to live data, Tempe’s visualizations update progressively once per second to show the latest query results. Tempe also provides a button to turn a script’s visualizations into a live dashboard.

This paper makes two contributions. First, we describe several design decisions we made within Tempe to support data science on streaming data. Second, we evaluate Tempe’s major design choices based on a deployment with Microsoft’s 343 Industries, as they prepared to release a new game in the *Halo* franchise. This deployment was a multi-month dialog with the Tempe team to incorporate their feedback and to add features they needed. 343 has used Tempe continuously for over six months to study and monitor server performance and customer behavior. Through interviews, 343 provided feedback on the value and the downsides in Tempe’s design.

II. BACKGROUND

Today’s data science workflows are clerical and awkward. Because of data volume and velocity, a typical data science workflow starts with collecting a snapshot of historical temporal data in a map/reduce system [6] like [Hadoop](#), using a relational query language like Sawzall [7] or Pig Latin [8]. Next, a data scientist often transfers files to a scripting environment like R or Python, for exploration and modeling work. There may be other file transfers to data exploration software, like Tableau, or presentation software.

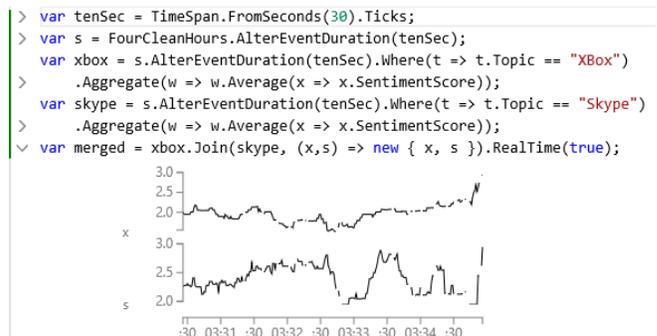


Figure 1. A Tempe script computing a live stream of two different Twitter feeds. The visualization updates as new data arrives.

Finally, deploying models or logic to a working service often means rewriting them from a scripting notation like R to the system’s implementation language, with potentially different runtime environments and data organization.

Recent interviews with data scientists document the pain points of these clerical work practices [9] [10]. Given the lengthy waits of batch systems and the clerical work involved, researchers have been working on improving its various stages. For instance, some research has focused on avoiding the long waits of batch processing, in favor of more interactive query processing [11] [12] [13]; others have automated the tedium of individual steps like data cleaning [14] [15] or capturing the workflow artifacts [16].

To support exploratory work, Tempe provides a live programming experience. Live programming was first created for visual programming languages, focused on specific domains like physics simulation [17] and image processing [18]. This domain focus continues today, e.g. *live coding*, in which a programmer-musician continuously updates a MIDI program to create a live music performance [19] [20]. Other live programming systems, such as Alvis Live [21], were designed to help students learn programming and proposed multiple variations of the live model. Flogo II, focused on end-user robot programming, provides *live text*, in which the current execution behavior is reflected in the program text, like graying out untaken branches [22]. AgentSheets provides a live spreadsheet representation of rewrite rules that encode the logic of a simulation [23].

III. DESIGN AND RATIONALE

The Tempe prototype is the result of a two-year development effort from a multidisciplinary team spanning HCI, information visualization, software engineering, programming languages, machine learning and databases. Although we reused major components where possible, we also designed the technology stack “from disk to pixels,” allowing us to coordinate implementation choices. In particular, the members of the Trill team, who created the underlying streaming engine, were also members of the Tempe team.

A. Scripting in a Notebook Model

Our goal was to design an environment for open-ended data cleaning, shaping, and exploration, for which scripting is a good fit. Data exploration tools based on direct manipulation, like Tableau, provide a large, but nonetheless fixed, space of operations. Given the popularity of scripting languages for data science—R, MATLAB, Python, and most recently Scala—C# would seem like an odd choice. However, a key goal is to allow a user to explore data in the context of a deployed system, a workflow called monitor-manage-mine (M^3) [3]. That is, the user could brainstorm new logic in Tempe, then directly copy/paste the resulting code into the deployed system; or, dually, a Tempe script could reuse logic from the deployed system, either by copy/paste or by linking. This goal is easiest to achieve when the scripting language is

the same as the programming language of the deployed system. Further, C# offers several other advantages: (1) familiarity to our user community; (2) C#’s LINQ feature [24] which embeds SQL-like queries in C# code; and (3) language safety, for running scripts on a server.

We implemented Tempe as a web-based notebook, where a central service executes the scripts. This allows clients to close local windows while keeping a dashboard script running and to share in-progress results with each other. Tempe uses a distinct URL for each script; thus, sharing a script can be as simple as sending a URL to a collaborator.

One aspect of combining a scripting experience with a notebook model is the need to show the current state of a script’s computation. A notebook page always shows its most recent results, even if they were produced in some long-ago run. This allows notebook pages to be shared as reports, even after the script finishes. The user may click the *Restart* button to re-run a script, which is always possible because a script’s complete context (necessary data and libraries) are stored as part of the notebook page, along with the script’s content. A stripe next to the script shows the current state of the computation: green, for successfully finished; orange, for finished with errors; striped gray, for in progress; and blank for not running.

B. Visualizations and Dashboards

Tempe creates inline visualizations as the user writes code. The choice of visualization for a value depends on its runtime type, for example, timelines for Trill streams, tables for generic collections, and print strings for scalars. These interactive visualizations allow the user to explore their query results. Figure 1 shows a script whose last result is visualized as a timeline. The user collapsed the visualizations for the previous four results.

When a script analyzes live data, its visualizations (typically, timelines) continually update as the queries produce new results. This means that a script can be used to monitor data in real-time. To better support this task, we provide a *Dashboard* button at the top of a notebook page. Clicking the button produces a dashboard version of the notebook page, with its own URL for sharing. The dashboard page shows every non-collapsed visualization from the corresponding notebook page, tiled to fill the screen.

C. Streaming query language

In order to express these scripts, the user needs a language that makes writing streaming code straightforward. We designed Tempe to work with the Trill streaming engine [4]. Trill’s query language is designed for temporal data, but its notion of time is abstract. This allows us to use the same query language for stored data by choosing a suitable notion of time, e.g. row number. This also allows the user to write queries that combine (join) live and stored data. For example, a live Twitter stream can have tweets with GPS coordinates; a stored table can have mappings from GPS coordinates to

country names. Joining these data streams allows the user to group tweets by country.

Trill’s query language is based on LINQ [24], but supplements it with temporal constructs like windowed aggregates. Trill queries use a publish/subscribe pattern and push new results to query subscribers. Because Trill is much faster than previous streaming systems [4], these pushes typically occur at a much faster frequency than the one-second interval at which Tempe updates its visualizations.

D. Live Programming

According to Tanimoto’s liveness taxonomy [18], a programming environment can offer four levels of liveness: (level 1) a user’s edits have no effect on the computation (e.g. a typical text editor); (level 2) a user explicitly submits edits to cause updates to the computation (e.g. a command loop); (level 3) a user’s edits automatically trigger any necessary re-computation (e.g. a spreadsheet); and (level 4) a user’s edits trigger updates to ongoing computations. Tempe implements level-3 liveness for stored data and level-4 liveness for live. That is, a user’s edits cause the scripting service to compute deltas and re-run any effected queries data [5]. When running a query on stored data, we choose to run the query again from “the beginning”, that is, from the first row. With live data, there is no “beginning”: running a query on live data runs the query on the next row, whenever it arrives.

With a live programming experience, script execution is asynchronous with respect to scripting editing. When the user’s editing forms a correct and complete query, Tempe’s scripting service starts executing it. If the query runs on live data, then the query runs indefinitely, until the user stops the script or overwrites the query’s content.

Despite the similar names, live programming and live data interact subtly to create the potential for confusion. Consider the following script pseudocode:

```
var x = Data.Count();  
var y = Data.Where(e => e.Action=="Open").Count();
```

Logically, one would expect the progressive value of x always to exceed the value of y, since y counts a subset of events that x counts (those whose Action field is “Open”). This expectation holds if both queries start at the same time on the same live data stream. However, if the user edits the first line after the script is already running, our live programming algorithm restarts that query (to keep its result up-to-date). This means that x’s query starts later in the live stream than y’s query. Hence, x could have a smaller value than y.

To support live scripting on live data, there is a tradeoff between maintaining consistency and retaining valuable live results. Because Tempe builds a dependency graph for each script, it could restart a whole dependency graph whenever any statement in the graph changes. This would keep query results mutually consistent, but would throw away results from ongoing queries. Unfortunately, previous query results over live streams cannot be recomputed—the past has

passed. We currently handle this tradeoff by allowing inconsistency, but providing a *Restart* button that forces all queries on a page to re-execute at the same time.

IV. DEPLOYMENT CASE STUDY

To understand the extent to which our design choices help software teams accomplish their data science goals, we wanted to engage with teams as they first adopted Tempe. To recruit teams, we released Tempe across Microsoft in March 2014. Since then, we have engaged with several teams, one of which agreed to be a case study for this paper. The team continues to use Tempe and provides ongoing feedback to improve its design.

343 Industries is a Microsoft game studio that produces the computer game series *Halo*. Their team has hundreds of engineers producing games played by tens of millions of players worldwide. *Halo* games run on consumer game consoles, but communicate with a backend service for achievements, game updates, and network play. Members of 343 have used Tempe for over six months to understand both the performance of their backend service as well as customer behavior. “Mark” and “Brad” (pseudonyms) are both software developers on the team responsible for the backend service. Mark’s primary responsibility is gameplay data; Brad focuses on performance data. Because the backend data reveal insights about customer usage, Mark often shares information with both the product’s Business Intelligence team and several of its project managers. As Mark put it, “*We’re this bridge between being a source of truth for the data that’s been returned to us from the [product] and also being a source of analysis for what’s come through the [product], and we kind of split that task with our BI team. We do it more in real-time, and they do it more in the background and over larger portions of data.*” Mark and Brad are good examples of our intended Tempe user: they are trained as developers, but because they work on a data platform, they often fill the role of data scientist.

343’s use of Tempe focused around their two live sources of data. The Tempe team created two new types of data ingress for their performance and gameplay data sources. (These data sources are not specific to their product and are useful to other Tempe users.) Brad wrote temporal queries over the performance data to show timelines of metrics like memory usage, CPU usage and other counters. At 343’s request, we also created the *Dashboard* button, so that these real-time timelines could be monitored without the clutter of the script text. 343 then built a dedicated kiosk to show these dashboards in their team room where the team members could monitor them.

During the same period as the performance monitoring, Mark issued many ad-hoc queries over the real-time customer data, some queries based on his own curiosity, some at the request of the BI team or project managers (PMs). Mark’s queries were typically temporal, like Brad’s, for ex-

ample, windowed averages of gameplay actions. In one instance, they looked for actions that they predicted would be rare—game achievements that were difficult to get—and found that one of the achievements was surprisingly frequent. The root cause was a bug that had previously been undetected.

URL-based sharing of analyses proved to be a useful feature.

I know during testing when we were trying to find out what's our service acting like right now, it was very convenient to be able to just say here's a query I just ran, I'll link it to you, you can take a look at what it's doing right now. ... Usually I'd share with project managers or people that maybe weren't the ones who could actually find this out themselves. I also share with my lead and my lead's lead to give them some insights.

For communicating with stakeholders, both visualizations and real-time update were important.

They were PMs [Program Managers] of PMs and they were concerned more about the overall health of our service and how we were doing especially when we were testing things in a live environment...The questions they were asking were more like how many events are we seeing in our system right now, how many active [customers] right now, and they wanted to see the charts on that and see that it lined up with their estimates...We'd create a visual chart using Tempe—the line graph would be the most common one—and we'd chart over our day what our population looked like...They wanted to see it happening in real time...Anyone could go in and take a look at where we are right now.

Tempe also provided easy transitions between monitoring and ad-hoc data explorations.

Being able to know with some level of confidence that everything is okay is good. And we have multiple ways of monitoring that, but this is ...a more reactive tool in our bag, because we can dig in our code and look for different things when we see an issue.

On the downside, Tempe's focus on scripting and advanced query APIs limited team participation.

We had originally we took the PMs and said here's a quick way to do this. They sort of tried to use it, but they weren't able to, so it fell back to me. It was probably [lack of] familiarity with how to do temporal-type queries...It's not unexpected.

Mark also found the live programming to be distracting and would prefer a mode switch.

Normally I don't care that something immediately starts populating. I much prefer to say when I press a button that says Start it then starts running and in between the states, like when I hit Stop, ... it should wait until I'm done, and then I'll hit start. I don't really care about generating

that in-between data. If I had an idea in mind that I have to create multiple variables to get there because I want to go stepwise into what I'm doing, having it start generating the data and popping open graphs while I'm doing it is a little disruptive. ... It was kind of a hindrance.

He also noted that a mode switch could clear up potential confusion for how live programming interacts with live data.

My mental model is that I expect it to start at wherever point it is in the stream. I don't expect to see any past events. I know it's a live stream. ... If the model is [the script] will immediately evaluate as I hit enter, then I expect that the later queries won't have seen what the previous ones saw. I don't expect it to go back and replay the events that the previous ones saw.... In my ideal world where I hit Start and Stop, then I don't have to worry about that at all. Then I know that when I hit Start the data came in at the top and went all the way through to the bottom.

Finally, Mark validated the utility of moving logic between Tempe and the production system.

It allows us to have a low barrier to testing things before we implement them. Seeing it against live data is also helpful. ... How we can tune [the logic] while it's in production, just to see, maybe if we flip these bits over here, this might be better [logic], without affecting any actual users.

V. CONCLUSIONS

This paper presents the major design choices behind a novel integrated environment for data science. Based on long-term engagement with several teams, some of our design choices clearly fit our users' needs:

- A cloud-hosted environment allows easy sharing via URLs, although read-only sharing would likely suffice.
- Providing visualizations throughout the interface allows users to spot errors, share insights, and monitor behaviors.
- A scripting environment allows developers to fill the role of data scientists, but makes them data gatekeepers by excluding non-programming teammates.
- Users appreciate flexibility of moving logic between historical and live data and between exploration and deployment.
- The value of live programming varies between users and tasks, so the environment should be flexibility about the triggering gesture and update rate.

VI. REFERENCES

- [1] A. Toshiwal et al., "Storm @Twitter," *Proc. SIGMOD*, pp. 147-156, 2014.
- [2] M. Zaharia et al., "Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing," *In Proc. of the 9th USENIX conference on Networked Systems Design and Implementation*.

- [3] B. Chandramouli, M. Ali, J. Goldstein, B. Sezgin, and S.B. Raman, "Data Stream Management Systems for Computational Finance," *IEEE Computer*, no. December, pp. 45--52, 2010.
- [4] B. Chandramouli et al., "Trill: A High-Performance Incremental Query Processor for Diverse Analytics," *VLDB*, 2015.
- [5] R. DeLine and D. Fisher, "Supporting Exploratory Data Analysis with Live Programming," *Proc. of the IEEE Symp. on Visual Languages and Human-Centric Computing*, 2015.
- [6] J. Dean and G. Ghemawat, "MapReduce: simplified data processing on large clusters," *Proc. OSDI*, 2004.
- [7] R. Pike, S. Dorward, R. Griesemer, and S. Quinlan, "Interpreting the data: Parallel analysis with Sawzall," *Sci. Program.*, vol. 13, no. 4, 2005.
- [8] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins, "Pig Latin: a not-so-foreign language for data processing," *Proc. ACM SIGMOD Intl. Conf. on Management of Data*, 2008.
- [9] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer, "Enterprise Data Analysis and Visualization: An Interview Study," in *IEEE Visual Analytics Science & Technology (VAST)*, 2012.
- [10] D. Fisher, R. DeLine, M. Czerwinski, and S. Drucker, "Interactions with big data analytics," *interactions*, vol. 19, no. 3, pp. 50-59, May/June 2012.
- [11] M. Barnett et al., "Stat! - An Interactive Analytics Environment for Big Data," in *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2013.
- [12] T. Condie et al., "MapReduce Online," in *7th USENIX Symposium on Networked Systems Design and Implementation*, 2010.
- [13] C. Jermaine, S. Arumugam, A. Pol, and A. Dobra, "Scalable approximate query processing with the DBO engine," in *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2007, pp. 725-736.
- [14] P.J. Guo, S. Kandel, J. Hellerstein, and J. Heer, "Proactive Wrangling: Mixed-Initiative End-User Programming of Data Transformation Scripts," in *ACM User Interface Software & Technology (UIST)*, 2011.
- [15] S. Gulwani, W. Harris, and R. Singh, "Spreadsheet Data Manipulation using Examples," *Communications of the ACM*, vol. 55, no. 8, pp. 97-105, August 2012.
- [16] J. Guo and M. Seltzer, "BURRITO: Wrapping Your Lab Notebook in Computational Infrastructure," *Proceedings of the 4th USENIX Workshop on the Theory and Practice of Provenance*, 2012.
- [17] R.B. Smith, "The Alternate Reality Kit: An animated environment for creating interactive simulations," in *IEEE Computer Society Workshop on Visual Languages*, 1986.
- [18] S.L. Tanimoto, "VIVA: A visual language for image processing," *Journal of Visual Languages and Computing*, vol. 1, no. 2, pp. 127-139, June 1990.
- [19] A.R. Brown and A. Sorensen, "Interacting with Generative Music through Live Coding," *Contemporary Music Review*, vol. 28, no. 1, pp. 17-29, February 2009.
- [20] S. Aaron, A.F. Blackwell, R. Hoadley, and T. Regan, "A principled approach to developing new languages for live coding," in *Proceedings of New Interfaces for Musical Expression*, 2011, pp. 381-386.
- [21] C.D. Hundhausen and J.L. Brown, "What You See Is What You Code: A "Live" Algorithm Development and Visualization Environment for Novice Learners," *Journal of Visual Languages and Computing*, vol. 18, no. 1, pp. 22-47, 2007.
- [22] C.M. Hancock, *Real-Time Programming and the Big Ideas of Computational Literacy*.: Massachusetts Institute of Technology, 2003.
- [23] A. Repenning, "AgentSheets: An interactive simulation environment with end-user programmable agents," *Interactions*, 2000.
- [24] E. Meijer, B. Beckman, and G. Bierman, "LINQ: reconciling object, relations and XML in the.NET framework," in *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, 2006.