Turning Personal Calendars into Scheduling Assistants

Preliminary report

Jacob Bank
Stanford University
Stanford, CA 94305
jbank@stanford.edu

Zachary Cain
Stanford University
Stanford, CA 94305
zcain@stanford.edu

Yoav Shoham
Stanford University
Stanford, CA 94305
shoham@stanford.edu

Caroline Suen
Stanford University
Stanford, CA 94305
cysuen@stanford.edu

Dan Ariely
Duke University
Durham, NC 27708
dandan@duke.edu

CHI'12, May 5–10, 2012, Austin, TX, USA.
ACM 978-1-4503-1016-1/12/05.

Abstract
Personal calendars have long played a major role in time management, but they have evolved little over the years, and their contribution to productivity has stagnated. Inspired by logical theories of intention as well as experimental results on human productivity, and leveraging the power of optimization algorithms, we seek to reinvent the digital calendar. First, we increase the expressive power of calendar systems by deriving new entity types that go beyond simple events to better represent human intentions, plans, and goals. Next, we build on social psychological research to characterize the properties of a schedule best engineered for human productivity. Finally, we develop an optimization framework and algorithm to generate these schedules from a set of entities. With these tools combined, we transform the digital calendar from a passive repository into an active scheduling assistant.

Author Keywords
Agents and Intelligent Systems; Office and Workplace

ACM Classification Keywords
F.4.1 [Mathematical Logic And Formal Languages]: Mathematical Logic; I.2.1 [Artificial Intelligence]: Applications and Expert Systems

Introduction

Despite their widespread importance in time management, personal calendars have evolved surprisingly little over the years. In the pre-electronic paper calendars, a user had to manually specify every event with an exact time. Digital calendars have since provided modest improvements by allowing recurring events, integrated invitations, and event reminders, but these systems remain essentially static lists of manually specified events. We seek to improve the usability and value of digital calendars along three dimensions: (a) increasing their expressive power, (b) automating some of the scheduling using optimization algorithms, and (c) incorporating lessons from psychology in the optimization. When coupled with sensible UI design, the overall effect is to transform the passive calendar into an active, intelligent scheduling assistant.

We start by drawing on logical theories of intention to improve the expressive power of digital calendar systems. Most current systems clumsily pair a calendar of concrete events with a simple TODO list, with the goal of capturing both events and tasks. But this simplistic approach still maps quite poorly to how we manage our intentions, plans, and goals. Some human intentions are firmly anchored in time, while others are completely open-ended, and others sit somewhere in between. By creating new types of calendar entities that succinctly capture the major different types of human intention, we allow a natural and easy expression of everything a user needs to accomplish. With these entities, the calendar's representation of intentions can correspond accurately to the user's internal mental representation.

This increase in expressiveness comes at a potential cost; it is harder cognitively to manually manage the more complex repertoire of calendar entries. But this potential shortcoming becomes an additional benefit when we enlist the power of optimization algorithms to perform the scheduling.

In performing this optimization we draw on research in behavioral psychology, which demonstrates clearly what we know from everyday experience: people make systematic mistakes in time management. We procrastinate; we delay important tasks; we waste time travelling needlessly, and much more. By allowing the digital calendar to create the schedule, we can not only reduce work for the user, we can also produce better schedules. Besides the general background work in logic, psychology and optimization, there has also been prior work specifically geared towards evolving the calendar in similar directions [11]. We discuss prior work in a separate section.

The structure of the paper is as follows. First, we describe the types of entities our calendar system supports for representing intentions. Next, we discuss in more depth the psychological considerations of human performance that guide the creation of productive schedules. Then, we formulate the optimization problem and describe the algorithm used to solve it. We next survey the relevant prior work in psychology, intention, and optimization that we build upon to create this system. Finally, we conclude by summarizing the contributions of the work, and describing some of the improvements and extensions we plan to make in the future.

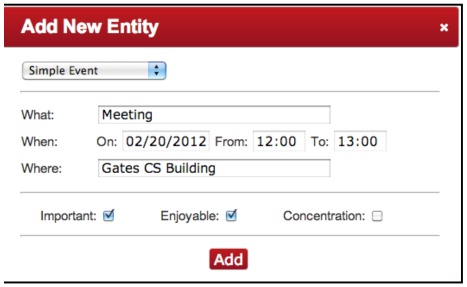

Figure 1: Simple Event dialog.

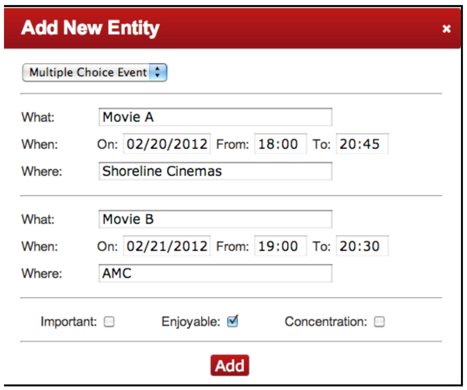

Figure 2: Multiple Choice dialog.

Calendar Entities

To improve the expressive power of digital calendars, we derive four new types of calendar entities from logical theories of intention [4,5,12]. The logical theories distinguish (among other things) between simple and complex intentions, and the temporal dimension of intension. Inspired by that literature, we introduce the following types of calendar entry:

1. *Simple Events* are standard calendar entries with a manually specified time and duration, and can optionally be recurring. This entity type captures events like: "I have a meeting from 12 to 1 on Monday".

2. Multiple Choice Events are sets of simple events of which only one needs to be scheduled, and the user is indifferent between them. For example: "I want to see Movie A on Monday at 7 or Movie B on Tuesday at 8".

3. *Floating Events* are uninterruptible events of a known duration that can occur at any time in a specified window or set of windows. This construct captures intentions like: "I need to do laundry sometime on Monday and it takes 2 hours."

4. *Tasks* are interruptible intentions with a known or unknown duration that may or may not have a deadline. These capture a wide variety of intentions, ranging from "I have a paper due on January 9th that will take 50 hours" to "I want to learn French."

The three properties at the bottom of each entity type- important, enjoyable, and concentration- will be used by our scheduling algorithm. In the future, we plan to allow ranked preferences of multiple choice event options and floating event windows.

Human Productivity Considerations

One could imagine that once a user has specified her intentions using this repertoire, it would be hard to actually commit to a particular set of concrete activities, given all the combinations possible. Indeed, this is hard for the human, but easy for algorithms. The question is which algorithm to apply. To decide that, we tap into insights from experimental psychology to develop the properties of a productive schedule:

1. Important tasks prioritized – Procrastination disproportionally impacts important tasks, so a schedule should place important tasks at better times.

2. No wasted travel time – A productive schedule should combine trips to minimize total travel time.

3. Difficulty of work follows circadian rhythms – Our cognitive ability varies throughout the day, so a schedule should align task difficulty with capability.

4. Task breakup – Splitting tasks into segments can improve performance, and tasks with different properties have different optimal splitting schemes.

5. Splitting for creative tasks – Creative tasks benefit from "incubation time", gaps between segments that allow ideas to grow. Also, since creative tasks require more time to fully engage, segments should be longer.

6. Ordering of tasks for best state of mind – A positive state of mind leads to better results on creative tasks, thus a schedule should order tasks based on task properties to leave the user with the correct affect to best complete the next task.

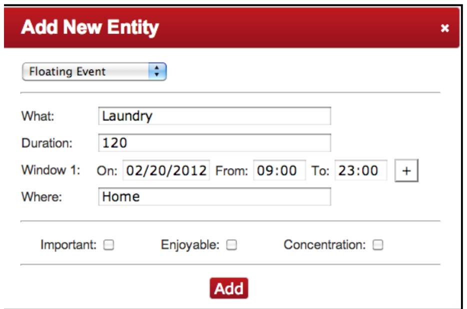

Figure 3: Floating Event dialog.

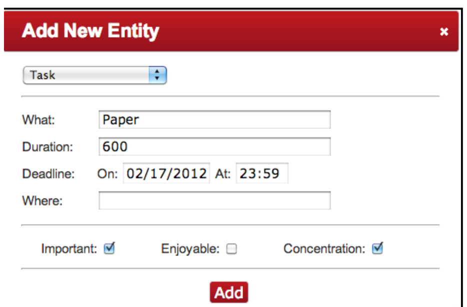

Figure 4: Task dialog.

Scheduling Problem and Algorithm

Using the properties described above, we can formulate the problem of generating a maximally productive schedule as an optimization problem. Given a set of *N* entities $E=e_1,\dots,e_n$ and a matrix of temporal distances between all pairs of locations, output a schedule *S* of entities, in which each entity e_i is described by p_i, the number of parts, t_{ij}, the start time of each part, and dur_{ij}, the duration of each part to:

$$\text{maximize:} \quad \sum_{e \in E} U(e) - D(e)$$

$$\text{where } U(e) = \begin{cases} \dfrac{R}{travel(e)} & \text{if the entity is flextime or multiple choice} \\[2ex] \dfrac{satisfying(e) * time(e) * priority(e)}{travel(e)} & \text{if the entity is a task} \end{cases}$$

subject to:

$$\text{C1.} \quad \forall T_{ij}, \sum_{j=1}^{p_i} dur_{ij} = duration(e_i)$$

$$\text{C2.} \quad \forall T_{ij}, minpart(e_i) \le dur_{ij} \le maxpart(e_i)$$

$$\text{C3.} \quad \forall T_{ij}, T_{ik}, j < k \Rightarrow t_{ij} + dur_{ij} + incubation(e_i) \le t_{ik}$$

$$\text{C4.} \quad \forall T_{ij}, start(e_i) \le t_{ij} \le deadline(e_i) - dur_{ij}$$

$$\text{C5.} \quad \forall T_{ij}, T_{mn}, T_{ij} \ne T_{mn} \Rightarrow (Dist(l_{ij}, l_{mn}) > 0) \Rightarrow (t_{ij} + dur_{ij} + Dist(l_{ij}, l_{mn}) \le t_{mn}) \vee (t_{mn} + dur_{mn} + Dist(l_{ij}, l_{mn}) \le t_{ij}$$

The optimization problem seeks to maximize the sum of the utilities of the entities. For floating and multiple choice events, the utility is *R* divided by the total travel time, travel(*e*), where *R* is a constant large enough to ensure that floating and multiple choice events will not be displaced by task segments. For tasks, the utility is the product of three functions-satisfying(*e*), time(*e*), and priority(*e*)- divided by travel(*e*). The first function, satisfying(*e*), captures the effect of the emotional state generated by the entity on the productivity of the entities around it, capturing property (6) from above. The second, time(*e*), captures how well the timing of the task fits with circadian rhythms (property (3)), and the final function, priority(*e*), provides a bonus if the task is high priority, leading us to schedule high priority tasks at better times (property (1)). The second term in the summation, *D(e)*, is a penalty term for the amount the entity has moved from the previous schedule. Since the scheduler constantly receives new entities, wild changes would be jarring, so we penalize shifting the timing of entities. The constraints then allow us to capture the remaining properties. C1 states that we must complete every entity we schedule. C2 states that the duration of each part must be between the minimum and maximum part length for the given entity (property (4)). C3 states that the gap between parts of a task must be at least as large as the incubation time required by the entity (property (5)). C4 captures that each entity must be scheduled after its start time and before its deadline. Finally, C5 ensures that the user can travel between all adjacently scheduled entities.

Since the search space for this optimization problem is exponential in size, we use heuristic methods that have been shown to approximately solve optimization problems like this one efficiently. First, since it is unnecessary and expensive to schedule all future entities, the algorithm only considers the entities that need to be fully or partially scheduled immediate two-week period. Next, we follow prior work [10] and use a modified version of the Squeaky Wheel Optimization (SWO) [7] framework to solve the scheduling optimization problem. We discuss this algorithm in detail in the full version of this paper.

Prior Work

As mentioned above, we build on three distinct bodies of prior work: (1) theoretical work on intention, (2) experimental work on human productivity, and (3) algorithmic work on optimization and scheduling. To improve the expressive power of our calendar, we build on the theory of intention, anchored in the work of Cohen and Levesque[4]. Our new entity types are inspired by this seminal work and also two more recent works. In the first [12], Shoham aimed to provide firmer foundations for computational theories of intention. And in the second [5], Icard et al. extended the seminal AGM axiomatic scheme on belief revision [1] to a dynamic theory of belief and intention. These theoretical examinations allowed us to derive the entities that would best represent human intention.

In experimental psychology, human productivity under different conditions has been studied for many years, and we build on a few branches of this work to establish the properties of a productive schedule. As procrastination is one of the most damaging errors in time management, we first reference recent work that has documented the negative effects of procrastination on performance. In one study, O'Donoghue and Rabin [8] showed that people procrastinate more on important tasks, which inspired us to prioritize high priority tasks in scheduling. Motivation for breaking tasks into segments came from work by Ariely and Wertenbroch [2], which demonstrated that a lack of intermediate deadlines decreases performance. In addition to experiments on procrastination, we also drew heavily on work on creative problem solving. In a study on the spacing of work on such tasks, Olton and Johnson [9] showed that incubation time between sessions of work boosts performance. In a study of the effect of emotional state on creative ability, Isen et al [6] showed that a positive state of mind improves results on tasks requiring originality. Finally, we draw on a study from biopsychology in which Carrier and Monk [3] documented how cognitive ability rises and falls throughout the day.

The last major pillar of prior research that supports our work is algorithms for scheduling and optimization. The general framework and algorithm we use to solve the scheduling problem is Squeaky Wheel Optimization (SWO) from Joslin and Clements [7]. In using SWO, we follow previous work by Refanidis, Yorke-Smith, and Alexiadis [10,11], which was the first to present scheduling in digital calendars as an optimization problem. In their work, they demonstrated that their slightly modified SWO algorithm solves the scheduling optimization problem effectively and efficiently. Our optimization formulation and SWO solving algorithm are directly inspired by their work [10]. While they deserve credit for their pioneering work, our framework is unique in its particular calendar entity types and the incorporation of psychological considerations.

Conclusion

We believe digital calendars should function less as static, passive repositories of events and more as powerful temporal spreadsheets, or active agents to manage the scheduling of intentions. To achieve this transformation, we first increase the expressive power of the calendar system to capture all of a user's intentions. Then, we combine the psychology of productivity with optimization algorithms to automatically generate schedules designed for human productivity. When combined, these advances lead to a markedly different and more useful model of digital calendar. We have designed and implemented a prototype of this system, which can be found at https://chronos.stanford.edu.

Clearly, much remains to be done. First, user experiments are needed to validate that (a) users will be willing to use these more complex entities and (b) the scheduling algorithm produces better schedules. In addition to testing those two key claims, we envision many extensions to our current system. Since different people perform best under different conditions, we plan to use machine learning to personalize the scheduling algorithm to a given user. We also plan to explore the multi-agent possibilities of such a calendar, including mechanism design for meeting scheduling and task allocation, scheduling of tasks with dependencies on multiple agents, and coordination between agents. Further innovation on digital calendars presents a rich array of possible research in mechanism design, optimization, human-computer interaction, and experimental psychology, and we plan to build on our current system to create a truly useful scheduling assistant.

Acknowledgements
This work was supported by the Department of the Army under Grant No. W911NF-10-1-0250. The first author was supported by an NSF Graduate Fellowship.

References
[1] C. E. Alchourron, P. Garndefors, and D. Makinson, *On the logic of theory change: Partial meet contraction and revision functions*. Journal of Symbolic Logic 50 (2):510-530, 1985.

[2] D. Ariely and K. Wertenbroch, *Procrastination, deadlines, and performance: Self-control by precommitment.* Psychological Science 13 (3):219-244.

[3] J. Carrier and T. Monk. *Circadian rhythms of performance: new trends*. Chronobiology International 17 (6): 719-732, 2000.

[4] P. Cohen and H. Levesque, *Intention is choice with commitment*. Artificial Intelligence 42 (2-3):213-261, March 1990.

[5] T. Icard, E.Pacuit, and Y. Shoham, *Joint Revision of Belief and Intention*. Proc KR, 2010.

[6] A. Isen, K. Daubman, and G. Nowicki, *Positive affect facilitates creative problem solving*. Journal of Personality and Social Psychology 52 (6):1122-1131, 1987.

[7] D. Joslin and D. Clements, *Squeaky Wheel Optimization*. Journal of Artificial Intelligence Research (10):353-373, 1999.

[8] T. O'Donoghue and M. Rabin, *Choice and Procrastination*. The Quarterly Journal of Economics 116 (1):121-160, 2001.

[9] R. Olton and D. Johnson, *Mechanisms of incubation in creative problem solving*. American Journal of Psychology 89 (4):617-630, 1972.

[10] I. Refanidis and N. Yorke-Smith, *A constraint-based approach to scheduling an individual's activities*. ACM TIST 1 (2), 2010.

[11] I. Refanidis and A. Alexiadis, *Deployment and evaluation of Selfplanner, an automated individual task management system*. Computational Intelligence 27 (1):41-59, 2011.

[12] Y. Shoham, *Logical Theories of Intention and the Database Perspective*. Journal of Philosophical Logic 38 (6): 633:648, 2009.