

# Semantic Email

Luke McDowell, Oren Etzioni, Alon Halevy, and Henry Levy  
Department of Computer Science and Engineering  
University of Washington  
Seattle, WA 98195 U.S.A  
{lucasm,etzioni,alon,levy}@cs.washington.edu

## ABSTRACT

This paper investigates how the vision of the Semantic Web can be carried over to the realm of email. We introduce a general notion of semantic email, in which an email message consists of an RDF query or update coupled with corresponding explanatory text. Semantic email opens the door to a wide range of automated, email-mediated applications with formally guaranteed properties. In particular, this paper introduces a broad class of *semantic email processes*. For example, consider the process of sending an email to a program committee, asking who will attend the PC dinner, automatically collecting the responses, and tallying them up. We define both logical and decision-theoretic models where an email process is modeled as a set of updates to a data set on which we specify goals via certain constraints or utilities. We then describe a set of inference problems that arise while trying to satisfy these goals and analyze their computational tractability. In particular, we show that for the logical model it is possible to automatically infer which email responses are acceptable w.r.t. a set of constraints in polynomial time, and for the decision-theoretic model it is possible to compute the optimal message-handling policy in polynomial time. Finally, we discuss our publicly available implementation of semantic email and outline research challenges in this realm.<sup>1</sup>

## Categories and Subject Descriptors

H.4.3 [Communications Applications]: Electronic mail;  
F.2 [Analysis of Algorithms and Problem Complexity]: Miscellaneous;  
H.2.8 [Database Applications]: Miscellaneous;  
G.3 [Probability and Statistics]: Probabilistic algorithms

## General Terms

Algorithms, Design, Theory

## Keywords

Semantic web, formal model, satisfiability, decision-theoretic

## 1. INTRODUCTION

While the WWW is a rich information space in which we spend significant amounts of time, many of us spend even more time on email. With the exception of the generic header fields associated with each message, email messages typically do not have semantic features. While the majority of email will remain this way, this

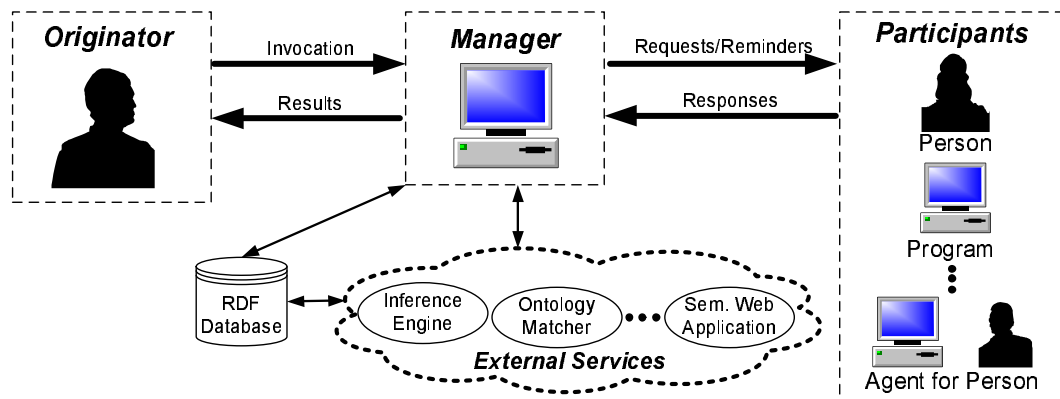
<sup>1</sup>See <http://www.cs.washington.edu/research/semweb/email> for a publicly accessible server (no installation required); source code is also available from the authors.

paper argues that adding semantic features to email offers opportunities for improved productivity while performing some very common tasks. To illustrate the promise, consider several examples:

- **Information Dissemination:** In the simplest case, suppose you send a talk announcement via email. With suitable semantics attached to the email, sending the announcement can also result in automatically (1) posting the announcement to a web calendar, and (2) sending reminders a day before the talk.
- **Event Planning:** Imagine you are organizing a program committee meeting and you want to know which PC members will stay for dinner after the meeting. Currently, you must send out the question and compile the replies *manually*, leafing through emails one by one. With semantic email, the PC members can provide the reply in a way that can be automatically interpreted and compiled, enabling such planning to scale to much larger numbers of people. In addition, after a few days, unresponsive PC members can be automatically reminded to respond, and those who have said they're not coming to the PC meeting need not be bothered with this query at all. Alternatively, suppose that you are organizing a *balanced potluck*, where people should bring either an appetizer, entree, or dessert, and you want to ensure that the meal is balanced. Here semantic email can help ensure that the potluck is indeed balanced by examining the replies and requesting changes where necessary.
- **Report Generation:** Suppose you need to collect projected budget numbers from a large set of managers. With semantic email, you could send a single email request and have the system automatically tabulate the responses, possibly requiring the values to satisfy certain individual or aggregate constraints. The system could then easily generate a spreadsheet report or integrate this data with other sources (e.g., prior budgets).
- **Auction/Giveaway:** Imagine you want to give away concert tickets that you cannot use. You would like to send out an announcement and have the semantic email system give out (or auction) the tickets to the first respondents. When the tickets are gone, the system should respond politely to later requests.

These examples are of course illustrative rather than exhaustive. In general, there are at least three ways in which semantics can be used to streamline aspects of our email habitat:

1. **Update:** We can use an email message to add data to some source (e.g., a web page, as in our first example).
2. **Query:** Email messages can be used to *query* other users for information. Semantics associated with such queries can then be used to automatically answer common questions (e.g., seeking my phone number or directions to my office).
3. **Process:** We can use semantic email to manage simple but tedious processes that we currently handle manually.



**Figure 1:** The invocation and execution of a SEP. The originator is typically a person, but also could be an automated program. The originator invokes a SEP via a simple web interface, and thus need not be trained in the details of SEPs or even understand RDF.

Because email is not set up to handle these tasks effectively, accomplishing them by hand can be tedious, time-consuming, and error-prone. The techniques needed to support the first two uses of semantic email depend on whether the message is written in text by the user or formally generated by a program on the sender's end. In the user-generated case, we would need sophisticated methods for extracting the precise update or query from the text (e.g., [12, 20]). In both cases, we require some methods to ensure that the sender and receiver share terminologies in a consistent fashion.

This paper focuses on the third use of semantic email to streamline processes, as we believe it has the greatest promise for increasing productivity and is where users currently feel the most pain. These processes support the common case of *asking* people a set of questions, *collecting* their responses, and *ensuring* that the results satisfy some set of goals. Some hardcoded email processes, such as the meeting request feature in Outlook, invitation management via *Evite*, and contact management via *GoodContacts*, have made it into popular use. Each of these commercial applications is limited in its scope, but validates our claim about user pain. Our goal in this paper is to sketch a *general* infrastructure for semantic email processes, and to analyze the inference problems it needs to solve to manage processes effectively and guarantee their outcome.

Workflow and collaboration systems such as Lotus Notes/Domino and Zaplets offer scripting capabilities and some graphical tools that could be used to implement sophisticated email processes. However, these systems lack support for reasoning about data collected from a number of participants (e.g., as required to balance a potluck or ensure that a collected budget satisfies aggregate constraints). In addition, such processes are constructed from arbitrary pieces of code, and thus lack the formal properties that our declarative model provides. We describe these properties and the limitations of existing workflow systems in more detail in Sections 3 and 4. Finally, messages in such systems lack the RDF content of semantic email, precluding automated processing by the recipient (e.g., to decline invitations for unavailable times).

Our work is the first to articulate and implement a general model of *semantic email processes* (SEPs). Our technical contributions are the following. Section 2 introduces a formalization for semantic email processes. The formalization specifies the meaning of semantic email processes and exposes several fundamental reasoning problems that can be used by the semantic email *manager* to facilitate SEP creation and execution. In particular, a key challenge is to decide when and how the manager should direct the process toward an outcome that meets the originator's goals. Section 3 addresses this challenge by describing a model of *logical* SEPs (L-SEPs) and demonstrating that it is possible to automatically infer which

email responses are acceptable with respect to a set of ultimately desired constraints in polynomial time. Section 4 describes a model of *decision-theoretic* SEPs (D-SEPs) that alleviates several shortcomings of the logical model, and presents results for the complexity of computing optimal policies for D-SEPs. Finally, Section 5 discusses implementation issues that arise for semantic email and how we have addressed these in our system, and Section 6 contrasts our approach with related work.

## 2. SEMANTIC EMAIL PROCESSES

Our formalization of SEPs serves several goals. First, the formalization captures the exact meaning of semantic email and the processes that it defines. Second, it clarifies the limitations of SEPs, thereby providing the basis for the study of variations with different expressive powers. Finally, given the formalization, we can pose several reasoning problems that can help guide the creation of semantic email processes as well as manage their life cycle. We emphasize that the *users* of SEPs are not expected to understand a formalization or write specifications using it. Generic SEPs are written by trained *authors* (who create simple constraints or utility functions to represent the goal of a process) and *invoked* by untrained users. The semantic email system then coordinates the process to provide the formal guarantees we describe later.

Figure 1 illustrates the three primary components of a SEP:

- **Originator:** A SEP is initiated by the *originator*, who is typically a person, but could be an automated program or agent.
- **Manager:** The originator invokes a new SEP by sending a message to the semantic email *manager*. The manager sends email messages to the *participants*, handles responses, and requests changes as necessary to meet the originator's goals. The manager stores all data related to the process in an RDF supporting data set, which may be configured to allow queries by external services (or other managers). To accomplish its tasks, the manager may also utilize external services such as inference engines, ontology matchers, and other Semantic Web applications, as described further below. The manager may be a shared server or a program run directly by the originator.
- **Participants:** The participants respond to messages received about the process. A participant may be a person, a standalone program (e.g., to represent a resource such as a room), or a software agent that acts on behalf of a person (e.g., to automatically respond to requests where possible, deferring others to the person). We assume that email addresses uniquely determine individuals or sets of potential participants in the process.

Informally speaking, the execution of a process is achieved by the supporting data set and the set of data updates that email recip-

ients make as they respond. In the model we describe, data is represented as a set of relations (i.e., the relational database model). However, as the application domains get more complex, we expect to use a richer representation language. To enable these future extensions as well as interactions with other Semantic Web applications, our system implements this representation using the Jena RDF storage system [25].

We illustrate our formalization with the running example of a “balanced potluck.” The originator invokes a process to announce the potluck and ask everyone whether they are bringing an appetizer, entree, or dessert. The originator also expresses a set of goals for the potluck. For example, he may specify that the number of appetizers, entrees, or desserts should differ by at most two. Note that analogous problems occur in many other contexts (e.g., distributing  $N$  persons among  $K$  committees or time slots).

The manager seeks to expedite the execution of this process and to achieve the originator’s goals. There are a number of ways in which reasoning can enhance the manager’s operation:

- **Predicting responses:** The manager may be able to infer the likely response of some participants even before sending any requests. For instance, the manager could employ another semantic web application or data source to detect that a suggested meeting time is unacceptable for a certain participant, based on information from calendars, course schedules, or other processes. The manager could use this information either to warn the originator as the process is being created, or to serve as a surrogate response until a definitive answer is received. Also, the manager could add a helpful annotation to the request sent to the participant, indicating what time is likely to be a conflict and why. As suggested above, this same reasoning could also be profitably employed on the participant’s end, where an agent may have additional information about the participant’s schedule.
- **Interpreting responses:** Typically, the originator will provide the participants with a finite set of choices (e.g., `Appetizer`, `Entree`, `Dessert`). However, suitable reasoning could enable substantially more flexibility. For instance, we could allow a potluck participant to respond with any value (either in plain text or in some formal language). Then, the manager could use a combination of information extraction or wrapper techniques (e.g., [12, 20]) and/or ontology matching algorithms (e.g., [11, 10]) to map the participant’s response into the potluck’s ontology. There are several interesting outcomes to this mapping. First, the response may directly map to one of the original potluck choices (e.g., “Cake” is an instance of `Dessert`). Second, the response may map to multiple choices in our ontology (e.g., “Jello salad” may be both an `Appetizer` and a `Dessert`). In this case, the manager might consider the response to be half of an appetizer and a dessert, or postpone the decision to a later time and classify it as is most convenient.<sup>2</sup> Third, the response may not map to any given choice, but may still be a subclass of `Food` (e.g., a “Sorbet” is a `Palette Cleanser`); here the manager might accept the response but exclude it from the goal calculations. Fourth, the response may map to a known ontology element that is not `Food` (e.g., “A hat”). Finally, the response may not map to any known element. In these latter two cases, the manager may either reject the response or notify the originator.
- **Recommending interventions:** Reasoning can also assist the manager with directing the process towards outcomes consistent with the originator’s goals. For instance, if the manager

<sup>2</sup>This is a very simple form of semantic negotiation; more complex techniques could also be useful [36].

detects that a potluck process is becoming unbalanced, it could refuse to accept certain responses, request changes from some participants, or warn the originator that further action is needed. In this case reasoning is needed to deduce the likely outcome of a process from the current state, and the likely effects of possible interventions.

In this work we focus on using reasoning for recommending interventions, leaving the other two items for future work. Specifically, we provide two different approaches for modeling the originator’s goals and when to intervene. In the logical model (Section 3), the originator specifies a set of *constraints* over the data set that should be satisfied by any process outcome, while in the decision-theoretic model (Section 4) the originator provides a function representing the *utility* of possible process outcomes. Below we consider each in turn, discuss possible variants, and present results for fundamental reasoning tasks that can determine how and when the manager should intervene.

### 3. LOGICAL MODEL OF SEPS

We now introduce our model of a logical semantic email process (L-SEP) and analyze important inference problems for this model.

#### 3.1 Definition of L-SEPs

A L-SEP is a 5-tuple  $\Lambda(P, D, R, M, C_D)$  with parts as follows:

**Participants  $P$ :** the set of participants in the process. Note that  $P$  may include the originator.

**Supporting data set  $D$ :** the set of relations that holds all data related to the process. The initial contents of  $D$  are specified by the originator (usually to be a set of default values for the columns). With each relation in  $D$  we associate a schema that includes:

- a relation name and names, data types, and range constraints for the attributes. A special data type is `emailAddress`, whose values are the set  $P$ . Attributes may have default values.
- possibly a distinguished from attribute, of type `emailAddress`, which means that rows in the relation whose `from` value is  $p$  can only result from messages from the participant  $p$ . The `from` attribute may be declared unique, in which case every participant can only affect a single row in the table.

**Responses  $R$ :** the set of possible responses to the originator’s email.  $R$  is specified as follows:

- **Attributes:** the set of attributes in  $D$  that are affected by responses from participants. This set of attributes *cannot* include any from attributes.
- **Insert or Update:** a parameter specifying whether participants can only add tuples, only modify tuples, or both. Recall that if there is a `from` field then all changes from  $p$  pertain only to a particular set of tuples.
- **Single or Many:** a parameter specifying whether participants can send a single response or more than one. As we explain in the next section, some responses may be *rejected* by the system. By single, we mean one non-rejected message.

**Messages  $M$ :** the set of messages that the manager may use to direct the process, e.g., to remind the participants to respond or to *reject* a participant’s response.

**Constraints  $C_D$ :** the set of constraints for every relation in  $D$ . These constraints  $C_D$  are specified in a language that includes conjunction and disjunction of atomic predicates. Atomic predicates compare two terms, or a term with a set. We allow comparison predicates ( $=$ ,  $\neq$ ,  $<$ ,  $\leq$ ), `LIKE`, and  $\in$ ,  $\notin$  between a constant and an enumerated finite set. A term may be a constant, an attribute

variable (the value of a specific attribute in a row), an expression combining two terms with any arithmetic operator, or an aggregate applied to a column of a relation (or to a subset of the rows that satisfy an (in)equality predicate).

**Example:** In our example,  $D$  contains one table named `Potluck` with two columns: `email`, a `from` attribute of type `emailAddress` and declared to be unique, and `bringing`, with the range constraint `Potluck.bringing`  $\in$   $\{\text{not-coming, appetizer, entree, dessert, NULL}\}$ . In addition,  $C_D$  contains a few constraint formulas similar to the abstract one below, specifying that the potluck should be balanced:

$$\begin{aligned} &(\text{COUNT}(\ast) \text{ WHERE bringing} = \text{'dessert'}) \leq \\ &(\text{COUNT}(\ast) \text{ WHERE bringing} = \text{'appetizer'}) + 2 \end{aligned}$$

Finally, the set of messages in our example includes (1) the initial message announcing the potluck and asking what each person is bringing, (2) messages informing each responder whether their response was accepted or not, (3) a reminder to those who have not responded 2 days before the potluck, (4) regular messages to the originator reporting the status of the RSVPs, and (5) a message to the originator in the event that everyone has responded.

### 3.2 Inference for L-SEPs

Given the formal model for an L-SEP we can now pose a wide variety of inference problems, whose results can serve to assist in the manager’s operation. This section describes the first such inference problem, which has different variations.

The core problem we want to address is determining whether an L-SEP will terminate in an *acceptable* state, i.e., a state that satisfies  $C_D$ . The input to the inference problem includes the constraints  $C_D$  and possibly the *current* state of  $D$  along with a response  $r$  from a participant. The output of the inference problem is a condition that we will check on  $D$  and  $r$  to determine whether to *accept*  $r$ . In our discussion, we assume that  $r$  is a legal response, i.e., the values it inserts into  $D$  satisfy the range constraints on the columns of  $D$ ; if not, the manager can respond with error messages until a legal response is received. Our goal is to automatically determine whether to *accept*  $r$  given the current state and  $C_D$ .

The space of possible inference problems is defined by several dimensions:

- **Necessity vs. possibility:** As in modal logics for reasoning about future states of a system [31, 14], one can either look for conditions that guarantee that *any* sequence of responses ends in a desired state (the  $\square$  operator), or that it is possible that *some* sequence ends in a desired state (the  $\diamond$  operator).
- **Assumptions about the participants:** In addition to assuming that all responses are legal, we can consider other assumptions, such as: (1) *all* the participants will respond to the message or (2) the participants are flexible, i.e., if asked to change their response, they will cooperate.
- **The type of output condition:** At one extreme, we may want a constraint  $C_r$  that the manager checks on  $D$  when a response  $r$  arrives, where  $C_r$  is specified in the same language used to specify  $C_D$ . At another extreme, we may produce an arbitrary procedure with inputs  $D$  and  $r$  that determines whether to accept  $r$ . We note that a constraint  $C_r$  will inevitably be weaker than an arbitrary algorithm, because it can only inspect the state of  $D$  in very particular ways. As intermediate points, we may consider constraints  $C_r$  in more expressive constraint languages. Note that in cases where we can successfully derive  $C_r$ , we can use database triggers to implement modifications to  $D$  or to indicate that  $r$  should be rejected.

As a very simple example, consider the case where we want *all* response sequences to end in an acceptable state, we make no assumptions on the participants except that we can elicit a legal response from them, and we are interested in deriving a constraint  $C_r$  that will be checked when a response arrives. If the initial state of  $D$  is an acceptable state, then simply setting  $C_r$  to be  $C_D$  provides a sufficient condition; we only let the data set  $D$  be in states that satisfy  $C_D$ . In the example of the balanced potluck, we would not accept a response with a dessert if that would lead to having 3 more desserts than entrees or appetizers. For an auction process, we would not accept a bid that is less than the current high bid.

In many cases, such a conservative strategy will be overly restrictive. For example, we may want to continue accepting desserts so long as it is still *possible* to achieve a balanced potluck. This leads us to the following inference problem.

### 3.3 Ultimate Satisfiability

We now describe our central result concerning inference for L-SEPs. Our goal is to find necessary and sufficient conditions for accepting a response from a participant. To do that, we cut across the above dimensions as follows. Suppose we are given the data set  $D$  after 0 or more responses have been accepted, and a new response  $r$ . Note that  $D$  does not necessarily satisfy  $C_D$ , either before or after accepting  $r$ . The manager will accept  $r$  if it is *possible* that it will lead to a state satisfying  $C_D$  (i.e., considering the  $\diamond$  temporal operator). We do not require that the acceptance condition be expressed in our constraint language, but we are concerned about whether it can be efficiently verified on  $D$  and  $r$ . Furthermore, we assume that participants can only update their (single) row, and only do so once. Hence, the columns that can be affected by the participants start out with the NULL value, and can get assigned values in  $a_1, \dots, a_n$ .

**DEFINITION 3.1. (ultimate satisfiability)** *Given a data set  $D$ , a set of constraints  $C_D$  on  $D$ , and a response  $r \in R$ , we say that  $D$  is ultimately satisfiable w.r.t.  $r$  if there exists a sequence of responses from the participants, beginning with  $r$ , that will put  $D$  in a state that satisfies  $C_D$ .*  $\square$

In what follows, let  $\mathcal{C}$  be our constraint language where we allow both conjunction and disjunction of atomic predicates. A term in a predicate of  $C_D$  may select a group of rows in an attribute  $A$ , and aggregate the value of the corresponding values in an attribute  $B$ . We say that the predicates in  $C_D$  are *constant-bounded* if the predicates refer to at most  $K$  distinct constants. If, in addition, each predicate only refers to attributes whose domain size is at most some constant  $L$ , then we say that the constraints are *domain-bounded*. We consider the aggregation functions COUNT, MIN, MAX, SUM, and AVERAGE.

**THEOREM 3.1.** *Let  $\Lambda$  be an L-SEP where  $C_D$  is in the language  $\mathcal{C}$ . Then,*

- *If  $C_D$  may contain any predicate allowed by  $\mathcal{C}$ , then ultimate satisfiability is NP-complete in the size of  $D$ .*
- *If the predicates in  $C_D$  are domain-bounded, then ultimate satisfiability is polynomial time in the size of  $D$  and  $C_D$ .*
- *If the predicates in  $C_D$  are just constant-bounded, but the only permitted aggregate is COUNT, then ultimate satisfiability is polynomial time in the size of  $D$  and  $C_D$ .*  $\square$

All of the examples in this paper may be described by constraints that satisfy the constant-bounded, COUNT-only condition above, while the domain-bounded case may be useful for SEPs that require more complex interactions. As an example of applying this theorem to the balanced potluck, suppose a new dessert response arrives. At that point, the inference procedure will (1) determine

the maximal number of people who *may* come to the potluck (i.e., the number of participants minus the number of people who replied *not-coming*), (2) check that even if the dessert response is accepted, then there are still enough people who have not answered such that the ultimate set of dishes could be balanced. Similar reasoning applies to other processes, e.g., to ensure that at least one person will sign up for each spot in a colloquium series.

The theorem is proved by enumerating representative states of the data set, each of which corresponds to a number of potential data sets that are all equivalent with respect to the constraints. The key is to express the constraints in terms of variables representing *aggregates* on the number of participants with each response. Space limitations preclude describing the complete proof here.

In comparison to related work, the challenge here is reasoning about the possible relationships between *aggregate* values (current and future), given a particular state of  $D$ . Reasoning about aggregation has received significant attention in the query optimization literature [33, 21, 9, 16] and some in the description logic literature (e.g., [3]). This body of work considered the problem of optimizing queries with aggregation by moving predicates across query blocks, and reasoning about query containment and satisfiability for queries involving grouping and aggregation. In contrast, our result involves considering the current state of the database to determine whether it can be brought into a state that satisfies a set of constraints. Furthermore, since  $C_D$  may involve several grouping columns and aggregations, they cannot be translated into single-block SQL queries, and hence the containment algorithms will not carry over to our context.

To the best of our knowledge, formalisms for reasoning about workflow [29, 28] or about temporal properties of necessity and possibility have not considered reasoning about aggregation. For instance, Abiteboul et. al [1] define a notion of goal reachability for *relational transducers* that is similar to our ultimate satisfiability (see also [17] for extensions to this model and a survey of other related work). Various restrictions on the model allow decidability of goal reachability in P, NP, or NEXPTIME, but none of these restrictions permit goals involving aggregation. Likewise, workflow formalisms have generally been restricted to reasoning about temporal and causality constraints. Such formalisms could potentially convert aggregation constraints to temporal constraints by enumerating all possible data combinations, but this may result in an exponential number of states. One exception is the recent work of Senkul et al. [34], who extend workflows to include *resource constraints* based on aggregation. Each such constraint, however, is restricted to performing a single aggregation with no grouping (and thus could not express the potluck constraint given in the earlier example). In addition, their solution is based on general constraint solving and thus will take exponential time in the worst case. We have shown, however, that in our domain L-SEPs can easily express more complex aggregation constraints while maintaining polynomial-time inference complexity for bounded constraints.

## 4. DECISION-THEORETIC MODEL

The logical model of SEPs described above supports a number of useful inferences that have both theoretical and practical applications. This model, however, has a number of shortcomings. First, L-SEPs, like logical theories in general, make no distinctions among unsatisfied outcomes. In our example, there is no way for L-SEPs to strive for a “nearly-balanced” potluck, since all unbalanced potlucks are equivalently undesirable. Second, an L-SEP ignores the cost of the actions taken in pursuit of its goals. For instance, a potluck L-SEP will always reject a response that results in unsatisfiable constraints, even if rejecting that response (e.g., from

an important official) may produce far worse effects than a slightly unbalanced potluck. Finally, L-SEPs make a very strong assumption that participants are always willing to change their responses if rejected. For instance, participants in a meeting scheduling process may be somewhat accommodating, but may refuse to modify a rejected response due to other commitments.

To address these limitations, we offer a decision-theoretic approach. We describe the goal of a decision-theoretic SEP (D-SEP) by a utility function over the outcome of the process that takes into consideration the cost of all actions required to achieve that outcome. In addition, instead of *rejecting* responses, the decision-theoretic model sometimes *suggests* that participants modify their choices. For instance, the balanced potluck uses a utility function that measures the extent to which the final meal selection is balanced, minus the costs (social or otherwise) of asking some participants to switch their responses. Below we formalize this model and then examine the tractability of finding optimal policies for it.

### 4.1 Definition of D-SEPs

A decision-theoretic SEP is a 6-tuple,  $\delta(P, S, V, A, U, T)$ . Note that the first five components of this tuple correspond roughly to the five components in our model for L-SEPs.

- **Participants  $P$ :** the set of participants, of size  $N$ .
- **States  $S$ :** the set of possible states of the system. The state describes both the current responses received and the outgoing change requests sent by the system.
- **Values  $V$ :** the set of possible values for participants to choose from (e.g.,  $V = \{\text{appetizer}, \text{entree}, \text{dessert}\}$ ), of size  $K$ .
- **Actions  $A$ :** the set of actions available to the system after sending out the initial message. Actions we consider are *NoOp* (do nothing until the next message arrives),  $SW_v$  (ask a participant to switch their response from  $v$  to something else), or *Halt* (enter a terminal state, typically only permitted when a message has been received from every participant). Other variants of actions are also useful (e.g., ask a participant to switch from  $v$  to a particular value  $w$ ); such additions do not fundamentally change the model or our complexity results.
- **Utilities  $U(s, a)$ :** the utility from executing action  $a$  in state  $s$ . For the potluck example,  $U(s, SW_v)$  is the (negative) utility from making a change suggestion, while  $U(s, Halt)$  is the utility based on the final potluck balance.
- **Transitions  $T(s, a, s')$ :** the probability that the system will transition to state  $s'$  after performing action  $a$  in state  $s$ . However, rather than having to specify a probability for each transition, these are computed from a smaller set of building blocks. For instance,  $\rho_v$  is the probability that a participant will originally respond with the value  $v$ ;  $\rho_{vw}$  is the probability that, when asked to switch from the choice  $v$ , a participant will change their response to  $w$  ( $\rho_{vv}$  is the probability that a participant refuses to switch).

The execution of the process proceeds in discrete steps, where at each step the manager decides upon an action to take (possibly *NoOp*). The outcome of this action, however, is uncertain since the manager is never sure of how participants will respond. The transition function  $T()$  models this uncertainty.

A *policy*  $\pi$  describes what action the manager will take in any state, while  $\pi(s)$  denotes the action that the manager will take in a particular state  $s$ . An *optimal policy*  $\pi^*$  is a policy that maximizes the expected utility  $U(\delta)$  of the process, where

$$U(\delta) = U(s_1, a_1) + U(s_2, a_2) + \dots + U(s_j, a_j)$$

for the sequence of states and actions  $\{(s_1, a_1), \dots, (s_j, a_j)\}$ .

D-SEPs are a special case of *Markov Decision Processes* (MDPs), a well-studied formalism for situations where the outcome of per-

forming an action is governed by a stochastic function and costs are associated with state transitions [32]. Consequently, we could find the optimal policy for a D-SEP by converting it to an MDP and using known MDP policy solvers.<sup>3</sup> However, this would not exploit the special characteristics of D-SEPs that permit more efficient solutions, which we consider below.

## 4.2 Variations of D-SEPs

As with our logical model, the space of possible D-SEPs is defined by several dimensions:

- **Restrictions on making suggestions:** Most generally, the manager may be allowed to suggest changes to the participants at *any time*, and to do so repeatedly. To be more user-friendly, we may allow the manager to make suggestions anytime, but *only once* per participant. Alternatively, if users may be expected to make additional commitments soon after sending their response (e.g., purchasing ingredients for their selected dish), then we may require the manager to respond with any suggestion *immediately* after receiving a message, before any additional messages are processed.
- **Assumptions about the participants:** In addition to the assumed probabilities governing participant behavior, we may also wish to assume that all participants will eventually respond to each message they receive. Furthermore, we might assume that participants will respond *immediately* to any suggestions that they receive (particularly if the manager also responds immediately to their original message), or instead that they can respond to suggestions *anytime*.
- **The type of utility functions:** At one extreme, we might allow complex utility functions based upon the individual responses of the participants (e.g., “+97 if Jay is bringing dessert”). Often, however, such precision is unnecessary. For instance, all potluck outcomes with 8 desserts and 1 entree have the same low utility, regardless of who is bringing what dish.

Below we consider the impact of these variations on the complexity of computing the optimal policy.

## 4.3 Computing the Optimal Policy

In this section we examine the time complexity of computing the optimal policy  $\pi^*$  for a D-SEP. We begin by considering a D-SEP with an arbitrary utility function and then examine how restrictions to the utility function and the permitted quantity and timing of suggestions make computing  $\pi^*$  more tractable. In all cases we assume that the participants will eventually respond to each message and suggestion that they receive. The following theorem is proved by reduction from QBF and the EXPTIME-hard game  $G_4$  [35, 22]:

**THEOREM 4.1.** *Let  $\delta$  be a D-SEP with  $N$  participants where the utility  $U(s, a)$  is any deterministic function over the state  $s$  and the current action  $a$ . If the manager can send only a bounded number of suggestions to each participant, then determining  $\pi^*$  is PSPACE-hard in  $N$ . If the manager can send an unlimited number of suggestions, then this problem is EXPTIME-hard in  $N$ . The corresponding problems of determining if the expected utility of  $\pi^*$  for  $\delta$  exceeds some constant  $\theta$  are PSPACE-complete and EXPTIME-complete, respectively.  $\square$*

<sup>3</sup>Specifically, D-SEPs are “Stochastic Shortest-Path” MDPs where the terminal state is reachable from every state, so an optimal policy is guaranteed to exist [4]. Incorporating additional features from *temporal MDPs* [8] would enable a richer model for D-SEPs (e.g., scheduling a meeting should be completed before the day of the meeting). However, existing solution techniques for TMDPs do not scale to the number of participants required for semantic email.

Thus, for the case of arbitrary utility functions determining  $\pi^*$  for a D-SEP is impractical for large values of  $N$ . (Conversion to an MDP offers little help, since the MDP would require a number of states exponential in  $N$ .) Note that this is a significant limitation, since for many D-SEPs it is natural to wish to scale to large numbers of participants (e.g., for large meetings or company-wide surveys). Below, we begin to make the calculation of  $\pi^*$  more tractable by restricting the type of utility function:

**DEFINITION 4.1. (K-Partitionable)** *The utility function  $U(s, a)$  of a D-SEP is K-partitionable if it can be expressed solely in terms of the variables  $a, C_1, \dots, C_K$  where  $a$  is the current action chosen by the manager and each  $C_i$  is the number of participants who have responded with value  $V_i$  in state  $s$ .  $\square$*

Intuitively, a utility function is K-partitionable if what matters is the number of participants that belong to each of a fixed number of K groups, rather than the specific participants in each of these groups. For instance, the utility function of our example potluck is *4-partitionable*, because all that matters for evaluating current and future utilities is the current number of participants that have responded `Appetizer`, `Entree`, `Dessert` and `Not-Coming`. In this case a simple utility function might be:

$$\begin{aligned} U(s, Halt) &= -\alpha(|C_A - C_E|^2 + |C_A - C_D|^2 + |C_E - C_D|^2) \\ U(s, SW_v) &= -1 \end{aligned}$$

where  $\alpha$  is a scaling constant and  $C_A, C_E, C_D$  are the numbers of appetizers, entrees, and desserts, respectively. Note that the maximum utility here is zero.

A K-partitionable utility function is analogous to the COUNT-only constraint language of Theorem 3.1. As with Theorem 3.1, we could allow more complex utility functions (e.g., with variables representing the MAX, SUM, etc. of the underlying responses); with suitable restrictions, such functions yield polynomial time results similar to those described below. Note, however, that the simpler K-partitionable definition is still flexible enough to support all of the SEPs discussed in this paper. In particular, a K-partitionable utility function may still distinguish among different types of people by counting responses differently based on some division of the participants. This technique increases the effective value of  $K$ , but only by a constant factor. For instance, the utility function for a meeting scheduling process that desires to have the number of *faculty* members attending ( $C_{yes,F}$ ) be at least three and the number of *students* attending ( $C_{yes,S}$ ) be as close as possible to five, while strongly avoiding asking faculty members to switch, might be:

$$\begin{aligned} U(s, Halt) &= -\alpha \cdot [\max(3 - C_{yes,F}, 0)]^2 - \beta |C_{yes,S} - 5|^2 \\ U(s, SW_{no, faculty}) &= -10 \\ U(s, SW_{no, student}) &= -1 \end{aligned}$$

A D-SEP that may make an unlimited number of suggestions but that has a *K-partitionable* utility function can be represented as an “infinite-horizon” MDP with just  $O(N^{2K})$  reachable states. Consequently, the D-SEP may be solved in polynomial time with the use of linear programming (LP), though alternative methods (e.g., policy iteration, simplex-based LP solvers) that do not guarantee polynomial time may actually be faster in practice due to the large polynomial degree of the former approach [23].

Furthermore, if we also restrict the system to send only one suggestion to any participant (likely a desirable property in any case), then computing the optimal policy becomes even more tractable:

Restrictions	Description of Restrictions	Complexity with arbitrary utility function	Complexity when $K$ -partitionable
AnyUnlimited	Manager may suggest changes at any time, and may send an unlimited number of suggestions to any participant.	EXPTIME-hard	MDP with $O(N^{2K})$ states
AnyOnce	Manager may suggest changes at any time, but only once per participant.	PSPACE-hard	$O(N^{3K})$ time
Immediate	Manager may suggest changes only immediately after receiving a response, once per participant.	PSPACE-hard	$O(N^{2K})$ time
Synchronous	Same as "Immediate", but each participant is assumed to respond to any suggestion before the manager receives any other message.	PSPACE-hard	$O(N^K)$ time

**Table 1: Summary of theoretical results for D-SEPs. The last two columns show the time complexity of finding the optimal policy for a D-SEP with  $N$  participants. In general, this problem is EXPTIME-hard but if the utility function is  $K$ -partitionable then the problem is polynomial time in  $N$ . Adding restrictions on how often the manager may send suggestions makes the problem even more tractable. Note that the size of the optimal policy is finite and must be computed only once, even though the execution of a SEP may be infinite (e.g., with "AnyUnlimited").**

**THEOREM 4.2.** *Let  $\delta$  be a D-SEP with  $N$  participants where  $U(s, a)$  is  $K$ -partitionable for some constant  $K$  and where the system is permitted to send at most one suggestion to any participant. Then  $\pi^*$  for  $\delta$  can be determined in  $O(N^{3K})$  time. (If the system can send at most  $L$  suggestions to any participant, then the total time needed is  $O(N^{(2L+1)K})$ .)  $\square$*

**Proof Sketch:** Although the formalisms are very different, the key observation underlying this proof is similar to that of Theorem 3.1. Here we also create a state space that only models the number of participants in each group, rather than their specific members.

We define a summary state function  $S = \{\bar{C}, \bar{D}, \bar{E}\}$  where

- $\bar{C} = (C_1, \dots, C_K)$  where  $C_i$  is the number of responses  $V_i$  that were received that do *not* have a suggestion pending.
- $\bar{D} = (D_1, \dots, D_K)$  where  $D_i$  is the number of responses  $V_i$  that were received that *do* have a suggestion pending.
- $\bar{E} = (E_1, \dots, E_K)$  where  $E_i$  is the number of responses  $V_i$  that were received as a response to a suggestion.

In what follows, the notation  $\bar{C} - v$  indicates "subtract one from the variable in  $\bar{C}$  specified by value  $v$ ." Given  $S$ , we can define the following transitions (omitting details for states where everyone has already responded):

$$\begin{aligned}
T(\{\bar{C}, \bar{D}, \bar{E}\}, SW_v, \{\bar{C}-v, \bar{D}+v, \bar{E}\}) &= 1 \\
T(\{\bar{C}, \bar{D}, \bar{E}\}, NoOp, \{\bar{C}+v, \bar{D}, \bar{E}\}) &= \rho_o(\bar{C}, \bar{D}, \bar{E}) \cdot \rho_v \\
T(\{\bar{C}, \bar{D}, \bar{E}\}, NoOp, \{\bar{C}, \bar{D}-v, \bar{E}+w\}) &= \rho_{sv}(\bar{C}, \bar{D}, \bar{E}) \cdot \rho_{vw}
\end{aligned}$$

The first equation represents the manager requesting that some respondent switch their response from the value  $v$ ; the state is updated to note that a suggestion has been made (with probability 1). The next two equations handle the uncertainty when the manager decides to wait for the next message to arrive. Specifically, the second equation handles the case when the next message is an original response from a previously unheard from participant (probability  $\rho_o(\bar{C}, \bar{D}, \bar{E})$ ), while the third equation handles the case where the next message is a response to a previously made suggestion to switch from value  $v$  (probability  $\rho_{sv}(\bar{C}, \bar{D}, \bar{E})$ ).

At any time each participant's response is either counted once among the  $K$  variables of each of  $\bar{C}$ ,  $\bar{D}$ , or  $\bar{E}$ , or has not yet been received. The number of possible states is thus the number of ways of dividing  $N$  participants among  $3K + 1$  groups, which is:

$$|S| = \binom{N + 3K}{3K} = O(N^{3K})$$

Because of the restriction to send at most one suggestion to each participant, the graph formed by the transition function over these states is acyclic. Thus, the optimal policy may be computed via a depth-first search over the graph in total time  $O(N^{3K})$ .  $\square$

Table 1 summarizes the results presented above as well as a few other interesting cases ("Immediate" and "Synchronous"). These results rely on two key optimizations. First, we can dramatically reduce the number of distinct states via  $K$ -partitioning<sup>4</sup>; this permits  $\pi^*$  to be found in polynomial time. Second, we can ensure that the state transition graph is acyclic (a useful property for MDPs also noted in other contexts [5]) by bounding the number of suggestions sent to each participant; this enables us to find  $\pi^*$  with simple graph search algorithms instead of with policy iteration or linear programming. Furthermore, this approach enables the use of existing heuristic search algorithms where an exact computation remains infeasible. Consequently, with appropriate restrictions many useful D-SEPs can be efficiently solved in polynomial time.

## 4.4 Discussion

Compared to L-SEPs, the primary advantage of D-SEPs are their ability to balance the utility of the process's goals vs. the cost of additional communication with the participants, and their graceful degradation when goals cannot be completely satisfied. On the other hand, the need to determine suitable utilities and probabilities is an inherent drawback of any decision-theoretic framework. Below we consider techniques to approximate these parameters.


First, the  $\pi^*$  for a D-SEP depends upon the relative value of positive utilities (e.g., having a well-balanced potluck) vs. negative utilities (e.g., the cost of making a suggestion). Our discussion above exhibited a number of simple but reasonable utility functions. In practice, we expect that D-SEPs will provide default utility functions based on their functionality, but would allow users to modify these functions by adjusting parameters or by answering a series of utility elicitation questions [6].

Second, D-SEPs also require probabilistic information about how participants are likely to respond to original requests and suggestions. This information can be determined in a number of ways:

- **User-provided:** The process originator may be able to provide reliable estimates of what responses are likely, based on some outside information or past experience.
- **History-based:** Alternatively, the system itself can estimate probabilities by examining the history of past processes.
- **Dynamically-adjusted:** Instead of or in addition to the above methods, the system could dynamically adjust its probability estimates based on the actual responses received. If the number of participants is large relative to the number of choices, then the system should be able to stabilize its probability estimates well before the majority of responses are received.

Thus, although the need to provide utility and probability estimates is a drawback of D-SEPs compared to L-SEPs, simple tech-

<sup>4</sup>See Boutilier et al. [7] for an alternative (though not guaranteed) use of domain structure to reduce the effective number of states.



**Figure 2:** The creation of a semantic email process. Initially, a SEP template is authored by the “Author”, then is later instantiated by the “Originator”. Typically, a template is authored once and then instantiated many times.

niques can produce reasonable approximations for both. In practice, the choice of whether to use a D-SEP or L-SEP will depend on the target usage and the feasibility of parameter estimation.

## 5. IMPLEMENTATION AND USABILITY

We have implemented a complete semantic email system and deployed it in several applications. In doing so, we faced several challenges. This section describes the desiderata for a usable semantic email system, highlights the challenges to achieving these desiderata, and discusses our particular implementation choices.

### 5.1 Desiderata

To be successful, we argue that any semantic email system (both SEP-based and otherwise) should fulfill the following desiderata:

- **Instant Gratification:** Most importantly, semantic email must provide an immediate, tangible benefit to users. Users must not be expected to annotate outgoing or incoming mail with semantic content for some vague future benefit. Instead, a semantic email system must provide users with existing services that yield immediately obvious benefits. In fact, the notion of instant gratification is key to getting people to invest in adding semantics to their data, and has been the motivation behind our MANGROVE semantic web system [26].
- **Gradual Adoption:** At first, semantic email will be initiated by only a small number of “early adopters.” If semantic email could be profitably exchanged only among these users, it would have very limited applicability. Thus, to succeed, semantic email must be usable even when some or all of the participants have no experience with or software installed for it.
- **Ease of use:** Semantic email must be simple enough for a non-technical person to use. It should not expect such users to understand RDF, disrupt normal email processing, or require email senders or recipients to use a particular email client.

Below we elaborate on the challenges in implementing a system that achieves these goals.

### 5.2 Process Creation and Execution

**Translating SEP theory to real problems:** Applying our SEP theory to real problems requires enabling an originator to easily create an L-SEP or D-SEP model that corresponds to his goals. One option is to build a GUI tool that guides the originator through constructing the appropriate choices, messages, and constraints or utilities for the process. Practically, however, a tool that is general enough to build an arbitrary process is likely to be too complex for untrained users.

Instead, our system is based on the construction of reusable *templates* for specific classes of SEPs. Figure 2 demonstrates this approach. Initially, a trained user *authors* a SEP template using an editor (most likely by modifying an existing template). The template is written in OWL based on a SEP ontology that describes the

possible queries, constraints, and messages for a process. For instance, the “balanced potluck” template defines the general balance constraints for the process, but has variable placeholders for the participants’ addresses, the specific choices to offer, and how much imbalance to permit. To enable an originator to provide these values later, the author also constructs a simple web form that prompts the originator for each variable (see Figure 3).

Our implementation provides a simple tool that can automatically generate such web forms from some additional OWL information describing each required variable. This same tool could also be used to generate a service description for the template, e.g., in WSDL or OWL-S [2]. Then, a program could also serve as an originator by utilizing the service description and template to automatically invoke the process directly.

An untrained originator finds a SEP from a public library of SEP templates and *instantiates* the template by filling out the corresponding web form, yielding a SEP *declaration* (also in OWL). The originator then *invokes* the process by forwarding this declaration to the manager. Given the formal declaration, the manager then executes the process, using appropriate L-SEP and D-SEP algorithms to decide how to direct the process via appropriate message rejections and suggestions.

**Facilitating responses:** Another key challenge is enabling participants to respond to messages in a way that is convenient but that can be automatically interpreted by the manager. A number of different solutions are possible:

- **Client software:** We could provide a custom email client that would present the participant with an interface for constructing legal responses, or automatically respond to messages it knows how to handle (e.g., “Decline all invitations for Friday evenings”). This client-based approach, however, requires all participants in a process to install additional software (conflicting with our gradual adoption goal) and is complicated by the variety of mail clients currently in use.
- **Information extraction:** We could allow participants to respond in a natural language (e.g., “I’ll bring a dessert”). We could then use wrappers or information extraction techniques to attempt to convert this response to one of the offered choices. This approach is promising but risks having the wrapper fail to extract the correct information.
- **Email or web forms:** We could provide participants with a text-encoded form to fill out, or we could send them a link to a suitable web-based form to use for their response. Embedded HTML forms are also attractive, but unfortunately are not handled uniformly by existing email clients.

While web forms have some advantages, we chose to use email text forms instead (see Figure 4) because we feel they fit more naturally with how people typically handle incoming messages. In addition, text forms offer a simple solution that works for any participant. Participants respond by replying to the process message and editing the original form.

**Manager deployment:** Potentially, the manager could be a program run on the originator’s personal computer, perhaps as part of his mail client. This permits an easy transition between authoring traditional mails and invoking SEPs, and can also benefit from direct access to the originator’s personal information (e.g., calendar, contacts). However, as with providing client software for participants, this approach requires software installation and must deal with the wide variety of existing mail clients.

Our implementation instead deploys the manager as a shared server. The server receives invocations from the originator and



**Figure 3:** A web form used to initiate a “balanced collection” process, such as our balanced potluck example. For convenience, clicking submit converts the form to text and sends the result to the server and a copy to the originator. The originator may later initiate a similar process by editing this copy and mailing it directly to the server.

sends out an initial message to the participants. Participants reply via mail directly to the server, rather than to the originator, and the originator receives status and summary messages from the server when appropriate. The originator can query or alter the process via additional messages or a web interface.

**Discussion:** Our server-based approach is easy to implement and meets our gradual adoption and ease-of-use goals since it requires no software installation, works for all email clients, and does not require users (as originators) to read or write RDF/OWL. In addition, this method supports our instant gratification goal by providing untrained users with existing, useful SEPs that can be immediately invoked and yield a tangible output (in the form of messages sent and processed on the users’ behalf). Finally, we believe that divorcing the processing of semantic email (in the server) from the standard email flow (in the client) will facilitate adoption by ameliorating user concerns about privacy<sup>5</sup> and about placing potentially buggy code in their email client.

In addition, specifying SEP templates and declarations in OWL has a number of advantages. First, unlike the original version of semantic email [13] (which used process-specific code), a SEP is described entirely by its OWL declaration. This greatly simplifies the deployment of a new SEP, both because no programming is required and because authors need not run their own server (since shared servers can accept and execute OWL declarations from anyone, something they are unlikely to do for arbitrary code). Second, authoring SEPs in OWL enables the use of a variety of automated tools to ensure that a SEP declaration is valid. Finally, OWL specifications could enable future work that automatically composes several SEPs to accomplish more complicated goals.

<sup>5</sup>Only semantic email goes through the server, personal email is untouched. Of course, when the semantic email also contains sensitive information, the security of the server becomes significant.

**Figure 4:** A message sent to participants in a “balanced potluck” process. The bold text in the middle is a form used for human recipients to respond, while the bold text at the bottom is a query that maps their textual response to a formal language (e.g., RDF).

### 5.3 Human/Machine Interoperability

The previous section highlighted how semantic email messages can be handled by either a human or by a program operating on their behalf. Thus, an important requirement is that every message must contain both a human-understandable (e.g., “You’re invited to the potluck on Oct 5...”) and a corresponding machine-understandable portion. For messages sent to a participant, this approach supports gradual adoption by permitting the originator to send the same message to all participants without any knowledge of their capabilities. For responses, a machine-understandable portion enables the manager to evaluate the message against the process constraints/utilities and take further action. The human-readable component provides a simple record of the response if needed for later review.

In our implementation, we meet this interoperability requirement with a combination of techniques. For responses, a human can fill out the included text form (see Figure 4), which is then converted into RDF at the server with a simple mapping from each field to an unbound variable in a RDQL query associated with the message. Alternatively, a machine can respond to the message simply by answering the query in RDF, then applying the inverse mapping in order to correctly fill out the human-readable text form.

For messages to the participants, the challenge is to enable the manager to construct these textual and RDF/RDQL portions directly from the SEP declaration. Here there is a tension between the amount of RDF content that must be provided by the SEP author (in the template) vs. that provided by the SEP originator (when instantiating the template). Very specific SEP templates (e.g., to balance N people among appetizer, entree, and dessert choices) are the easiest to instantiate, because the author can specify the RDF terms needed in advance. General SEP templates (e.g., to balance N people among K arbitrary choices) are much more reusable, but require substantially more work to instantiate (and may require understanding RDF). Alternatively, authors may provide very general templates but make the specification of RDF terms for the choices

optional; this enables easy template reuse but fails to provide semantic content for automated processing by the participants.

In our current system, we offer both highly specialized SEPs (e.g., for meeting scheduling) and more general SEPs (e.g., to give away some type of item). Enabling originators to easily customize general SEPs with precise semantic terms, perhaps from a set of offered ontologies, is an important area of future work.

## 5.4 Integrating with Non-Semantic Messages

Despite the advantages of semantic email, we do not want to create a strict dichotomy in our email habitat. In our potluck example, suppose that one of the participants wants to know whether there is organized transportation to the potluck (and this information affects his decision on what to bring). What should he do? Compose a separate non-semantic email to the originator and respond to the semantic one only later? A better (and easier to use) solution would be to treat both kinds of emails uniformly, and enable the participant to ask the question in replying to the semantic email, ultimately providing the *semantic* response later on in the thread.

Our implementation supports this behavior by supplying an additional *Remarks* field in each response form, where a participant may include a question or comment to be forwarded to the originator. For a question, the originator can reply, enabling the participant to respond to the original semantic question with the included form or pose another question.

## 5.5 Experience

Our semantic email system is deployed and may be freely used by anyone without any software installation; the source code for deploying other instances of the server is also available. So far we have developed simple processes for functions like collecting RSVPs, giving tickets away, scheduling meetings, and balancing a potluck. Our system uses standard ontologies where possible (e.g., RDF Calendar [37]), augmented as needed with a local semantic email schema.

Despite very limited publicity, our semantic email server has seen growing interest over the six months that it has been available. For instance, a DARPA working group has adopted semantic email for all of its meeting scheduling and RSVP needs, students have used semantic email to schedule seminars and Ph. D. exams, and semantic email has been used to organize our annual database group and departmental-wide potlucks. Furthermore, a number of other institutions have expressed interest in deploying copies of semantic email locally at their sites. These are merely anecdotes but lend credence to our claim that semantic email is both useful and practical.

Our prototype is integrated within our larger MANGROVE [26] semantic web system. This provides us with an RDF-based infrastructure for managing email data and integrating with web-based data sources and services. For instance, the MANGROVE web calendar accepts event information via email or from a web page. In addition, MANGROVE provides semantic email with an RDF data source about courses, people, etc. that could be used to support the prediction of likely responses by the manager discussed in Section 2. Likewise, a semantic email client could utilize data from MANGROVE to answer common questions. When previously unknown questions are answered manually by the user, these responses could be stored for future use, thus enabling the automatic acquisition of semantic knowledge over time. Future work will consider additional ways to synergistically leverage data from both the web and email worlds in MANGROVE.

## 6. RELATED WORK

Information Lens [24] used forms to enable a user to generate a single email message with semi-structured content that might assist recipients with filtering and prioritizing that message. Our SEPs generalize this earlier work by enabling users to create an email *process* consisting of a set of interrelated messages, and by extending Information Lens's rule-based message processing to support more complex constraint and utility reasoning based on information from the entire set of messages. Consequently, SEPs support a much broader range of possible applications. More recently, Kalyanpur et al. [18] proposed having users semantically annotate messages to improve mail search, sorting, and filtering. This approach can potentially result in rich semantic content, but requires users to invest significant annotation effort for some *potential* future benefit (e.g., in improved searching for an old email) or primarily for the benefit of the recipient. SEPs instead generate both the semantic content and the text of the email message directly from simple forms, and provide instant gratification by immediately utilizing this content for simple but time-saving email processes.

Our vision for semantic email was initially described in Etzioni et al. [13]. Possible uses of semantic email are similar to those of some existing semantic web systems (e.g., [30, 19, 27], cf., RDF Calendar group discussions [37]). For instance, Rcal [30] uses messages between participants to agree upon meeting times and McIlraith et al. [27] describe an agent that makes travel arrangements by invoking various web services (which could be modeled as participants in a SEP). These systems, however, enable full interaction only between two parties that are both executing domain-specific software. For instance, though Rcal provides a web interface to let anyone schedule an appointment with an installed Rcal user, an Rcal user cannot use the system to request an appointment with a non-“Rcal-enabled” person. Likewise, McIlraith et al.'s agent is designed only to communicate with specific web services, not with humans (such as human travel agents) that could offer the same functionality. Our system instead permits processes to include any user, regardless of their capabilities. An additional, though less critical, distinction is our use of email instead of HTTP or a custom protocol (cf., Everyware [15]). Email provides a convenient transport mechanism because the vast majority of users already have well-known addresses (no additional directories are needed), messages can be sent regardless of whether the recipient has performed any configuration, and existing email clients provide a useful record of messages exchanged. The combination of these factors makes semantic email a lightweight, general approach for automating many tasks that would be impractical with other systems.

## 7. CONCLUSIONS

This paper generalizes the original vision of the semantic web to also encompass email. We have introduced a paradigm for semantic email and described a broad class of semantic email processes. These automated processes offer tangible productivity gains on email-mediated tasks that are currently performed manually in a tedious, time-consuming, and error-prone manner. Moreover, semantic email opens the way to scaling similar tasks to large numbers of people in a manner that is infeasible today. For example, large organizations could carry out surveys, auctions, and complex meeting coordination via semantic email with guarantees on the behavior of these processes.

Our technical contributions are as follows. We presented a formalization that teases out the issues involved, and used this formalization to explore several central inference questions. We then defined and explored two useful models for specifying the goals

of a process and formalizing when and how the manager of the process should intervene. For our logical model we showed how the problem of deciding whether a response was ultimately acceptable relative to the constraints could be solved in polynomial time for bounded constraints. With our decision-theoretic model we addressed several shortcomings of the logical model and demonstrated how appropriate restrictions could enable the optimal policy for this model to be computed in polynomial time. In both cases we identified restrictions that greatly improved the tractability of the key reasoning problem while still enabling a large number of useful processes to be represented. Finally, we described our publicly available semantic email system and how it satisfies the implementation desiderata of instant gratification, gradual adoption, and ease of use.

There are a number of interesting directions for future work. First, we want to consider interactions between semantic email and other semantic web applications to support more sophisticated reasoning techniques (e.g., check calendars and other resources to help constrain the number of messages and responses from some participants). We also plan to incorporate our recent work on schema and ontology mapping [11] to support more flexibility in responding to a semantic email message. Finally, we identified specific cases where our reasoning is tractable, but there are many opportunities for studying other cases within the framework we provided for modeling SEPs.

## 8. ACKNOWLEDGEMENTS

This research was supported in part by NSF grant IIS-0312988, DARPA contract NBCHD030010, and ONR grant N00014-02-1-0324 for Oren Etzioni; NSF CAREER grant IIS-9985114 and ITR grant IIS-0205635 for Alon Halevy; and by NSF Graduate Research and Microsoft Endowed Fellowships for Luke McDowell. Jeff Lin assisted with the server implementation. Thanks to Abraham Bernstein, Phil Bernstein, Sandy Liu, Mausam, Matthew Richardson, Stani Vlasheva, Dan Weld, and the anonymous reviewers for their helpful comments on improving this work.

## 9. REFERENCES

- [1] S. Abiteboul, V. Vianu, B. S. Fordham, and Y. Yesha. Relational transducers for electronic commerce. In *Proc. of PODS*, 1998.
- [2] A. Ankolenkar, M. Burstein, J. Hobbs, O. Lassila, D. Martin, S. McIlraith, S. Narayanan, M. Paolucci, T. Payne, K. Sycara, and H. Zeng. DAML-S: A semantic markup language for web services. In *Proc. of the Semantic Web Working Symposium*, pages 411–430, 2001.
- [3] F. Baader and U. Sattler. Description logics with concrete domains and aggregation. In *Proc. of the European Conference on Artificial Intelligence*, pages 336–340, 1998.
- [4] D. Bertsekas. Dynamic programming and optimal control. Athena Scientific, 1995.
- [5] B. Bonet and H. Geffner. Planning with incomplete information as heuristic search in belief space. In *Proc. of Artificial Intelligence Planning Systems*, pages 52–61, 2000.
- [6] C. Boutilier. A POMDP formulation of preference elicitation problems. In *Proc. of AAAI*, pages 239–246, 2002.
- [7] C. Boutilier, R. Dearden, and M. Goldszmidt. Exploiting structure in policy construction. In *Proc. of IJCAI*, 1995.
- [8] J. A. Boyan and M. L. Littman. Exact solutions to time-dependent MDPs. In *Proc. of Advances in Neural Information Processing Systems (NIPS)*, pages 1026–1032, 2000.
- [9] S. Cohen, W. Nutt, and A. Serebrenik. Rewriting aggregate queries using views. In *Proc. of PODS*, pages 155–166, 1999.
- [10] A. Doan, P. Domingos, and A. Halevy. Reconciling schemas of disparate data sources: a machine learning approach. In *Proc. of SIGMOD*, 2001.
- [11] A. Doan, J. Madhavan, P. Domingos, and A. Halevy. Learning to map between ontologies on the semantic web. In *Proc. of WWW*, 2002.
- [12] R. B. Doorenbos, O. Etzioni, and D. S. Weld. A scalable comparison-shopping agent for the world-wide web. In *Proc. of the First International Conference on Autonomous Agents*, 1997.
- [13] O. Etzioni, A. Halevy, H. Levy, and L. McDowell. Semantic email: Adding lightweight data manipulation capabilities to the email habitat. In *Proc. of Sixth Int. Workshop on the Web and Databases*, 2003.
- [14] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. M.I.T Press, 1995.
- [15] F. Ferreira, D. Schwabe, and C. Lucena. Using semantic web services now. In *Proc. of IX Brazilian Symp. on Hypermedia and Multimedia*, 2003.
- [16] S. Grumbach and L. Tininini. On the content of materialized aggregate views. In *Proc. of PODS*, 2000.
- [17] R. Hull, M. Benedikt, V. Christophides, and J. Su. E-Services: A look behind the curtain. In *Proc. of PODS*, 2003.
- [18] A. Kalyanpur, B. Parsia, J. Hendler, and J. Golbeck. SMORE - semantic markup, ontology, and RDF editor. <http://www.mindswap.org/papers/>.
- [19] S. Kumar, A. Kunjithapatham, M. Sheshagiri, T. Finin, A. Joshi, Y. Peng, and R. S. Cost. A personal agent application for the semantic web. In *AAAI Fall Symposium on Personalized Agents*, 2002.
- [20] N. Kushmerick, R. Doorenbos, and D. Weld. Wrapper induction for information extraction. In *Proc. of IJCAI*, 1997.
- [21] A. Y. Levy, I. S. Mumick, and Y. Sagiv. Query optimization by predicate move-around. In *Proc. of VLDB*, pages 96–107, 1994.
- [22] M. L. Littman. Probabilistic propositional planning: Representations and complexity. In *Proc. of AAAI*, 1997.
- [23] M. L. Littman, T. L. Dean, and L. P. Kaelbling. On the complexity of solving Markov decision problems. In *Proc. of UAI*, 1995.
- [24] T. Malone, K. Grant, F. Turbak, S. Brobst, and M. Cohen. Intelligent information-sharing systems. *Comm. of the ACM*, 30(5):390–402, 1987.
- [25] B. McBride. Jena: Implementing the RDF model and syntax specification. In *Proc. of the 2001 Semantic Web Workshop*, 2001.
- [26] L. McDowell, O. Etzioni, S. D. Gribble, A. Halevy, H. Levy, W. Pentney, D. Verma, and S. Vlasheva. Mangrove: Enticing ordinary people onto the semantic web via instant gratification. In *Proc. of Second International Semantic Web Conference*, October 2003.
- [27] S. A. McIlraith, T. C. Son, and H. Zeng. Mobilizing the semantic web with daml-enabled web services. In *Proc. of Semantic Web Workshop*, 2001.
- [28] C. Mohan. Workflow management in the internet age. [www.almaden.ibm.com/u/mohan/workflow.pdf](http://www.almaden.ibm.com/u/mohan/workflow.pdf), 1999.
- [29] S. Mukherjee, H. Davulcu, M. Kifer, P. Senkul, and G. Yang. Logic based approaches to workflow modeling and verification. In *Logics for Emerging Applications of Databases*, 2003.
- [30] T. Payne, R. Singh, and K. Sycara. Calendar agents on the semantic web. *IEEE Intelligent Systems*, 3(11):84–86, 2002.
- [31] A. Pnueli. The temporal logic of programs. In *Proc. of the 18th Annual IEEE Symposium on Foundations of Computer Science*, 1977.
- [32] M. Puterman. Markov decision processes. Wiley Inter-science, 1994.
- [33] K. Ross, D. Srivastava, P. Stuckey, and S. Sudarshan. Foundations of aggregation constraints. In *Principles and Practice of Constraint Programming*. LNCS, 874. Springer Verlag, 1994.
- [34] P. Senkul, M. Kifer, and I. H. Toroslu. A logical framework for scheduling workflows under resource allocation constraints. In *Proc. of VLDB*, 2002.
- [35] L. J. Stockmeyer and A. K. Chandra. Provably difficult combinatorial games. *SIAM Journal on Computing*, 8(2):151–174, 1979.
- [36] V. Tamma, M. Wooldridge, and I. Dickinson. An ontology-based approach to automated negotiation. In *Proc. of AMEC*, 2002.
- [37] W3C. RDF calendar workspace. <http://www.w3.org/2002/12/cal/>.