

Speculative Constraint Processing with Iterative Revision for Disjunctive Answers

Martine Ceberio¹, Hiroshi Hosobe², and Ken Satoh²

¹ University of Texas at El Paso,
500 West University Avenue, El Paso, Texas 79968-0518, USA
mceberio@cs.utep.edu

² National Institute of Informatics,
2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan
{ksatoh, hosobe}@nii.ac.jp

Abstract. In multi-agents systems, incompleteness, due to either communication failure or response delay, is a major problem to handle. To face incompleteness, frameworks for speculative computation were proposed (see references [5, 6, 7]). The idea developed in such frameworks is to allow the asking agent, while waiting for the slave agents to reply, to reason using default beliefs until replies are sent.

In particular, K. Satoh and K. Yamamoto [7] proposed a framework that allows an agent not only to perform speculative computation, but also to accept iterative answer revision for yes/no questions. In this paper, we present an extension of the framework for more general types of questions using constraint logic programming (CLP).

1 Introduction

Multi-agent systems are very fashionable and convenient, because they make it possible, for instance, to take advantage of multi-processor machines, and also make it possible to design human-like efficient organizations of agents. The main limitation to such an approach is that, as also arises in human organizations, communication may be an issue: delayed or broken, it leads to incompleteness of the information in the reasoning structure.

This is a concrete concern when we consider distributed systems such as the Internet, in which communication is indeed not guaranteed, and even if we could guarantee it, communication may either take time, or agents themselves may delay their sending information.

For such non-ideal, but as we believe, practical situations, when problem-solving is at stake, frameworks for speculative computations were proposed: first for yes/no questions only [6], and then for general questions [5] using constraints.

K. Satoh et al. [6] and K. Satoh, P. Codognet, and H. Hosobe [5] only provided the possibility for the master agent to perform speculations and a returned answer from the slave agent is final and with no possibility of a change in answers. However, if we let every agent perform speculative computation, the asked agent may revise his answer since the previous answer sometimes depends on the asked

agent's belief, which might turn out to be false. Therefore, a chain reaction of belief revision among agents might occur, which was firstly observed by K. Satoh and K. Yamamoto [7], and they provide a revisable speculative computation method for yes/no questions. The essential part of their work is a dynamic iterative belief revision mechanism that can handle a revision of an answer for query even during the execution.

Belief revision is indeed very important for both the sake of flexibility (information is processed before it is complete), and speed of computation (time is saved in case prior information is later entailed).

We combine the methods proposed by K. Satoh et al. [5, 7], and extend them, so that we can handle iterative answer revision for a query with constraints. We also complete these methods with the ability to incorporate disjunctive answers. So, the main contribution from this paper is the definition of the framework that enables to perform speculative computations on constraints, while handling belief revision, and that handles disjunctive answers as well. In particular, the main challenges with this framework are the following:

- First, processing speculative constraints, as shown by K. Satoh et al. [5], is manageable when belief revision is not considered. In our research, belief revision is made possible because it enables more speculative computation in multi-agent systems. This makes the problem much harder: the process management needs to be modified to enable changes in the computation at any time, while maintaining a reasonable balance between not being too space-consuming, and not loosing too much time (*i.e.*, we don't want to start from scratch all the time). The process management is presented in detail in this paper, as well as the results on the space complexity of our operational model.
- The second challenging point described in this paper is the way disjunction is now handled in the framework we propose. Indeed, considering the situation where each agent's behavior is specified as a CLP program, we need to handle alternative answers, since these answers may come from different derivations in CLP. By manipulating such alternative answers, we face another complication, in that we need to distinguish a revised answer of a previous answer, from an answer derived from an alternative derivation path. To solve this problem, we devise an answer entry that keeps track of the usage status of the answer in processes. This new feature impacts the way processes are managed, as described in Section 3, and therefore makes the problem more complicated.

For an iterative belief revision, many proposals have been described. As far as we know, existing frameworks separate reasoning and belief revision, except those by K. Satoh et al. [5, 6, 7] and that by F. Sadri and F. Toni [4]. Our framework in this paper is along the line of the works of K. Satoh et al. in a more general setting. F. Sadri and F. Toni proposed an abductive logic programming proof procedure, called LIFF, that enables the interleaving of belief revision and reasoning. The advantage of LIFF is that it allows the addition and deletion of rules, while our framework processes only the addition and deletion of constraints. However, our framework allows predicate cases, while LIFF handles

only propositional cases. In addition, our framework does not require recomputation for constraint narrowing, whereas LIFF needs to recompute goals related to updated rules. Also, our framework performs computation along plausible paths by using default rules, while LIFF does not adopt such explicit control.

There are works on the formalization of an agent in terms of logic programming, such as that conducted by R. A. Kowalski and F. Sadri [3]. Although these works are important in their own right, our paper pursues another branch of investigation in the context of speculative computation.

The most closely related research would be constraint programming languages, such as Andorra Kernel Language (AKL) [2] and Oz [8], which perform a kind of speculative computation. AKL allows local speculative variable bindings in a guard of each clause until one of the guards succeeds, and Oz can control multiple computation spaces, each of which represents an alternative path of constraint processing. As far as we understand, however, speculative computation used in these languages is mainly meant for or-parallel computing, where all alternative paths of computation are executed in parallel, until one of the paths eventually succeeds. On the other hand, we regard a speculative computation as a default computation where the most plausible paths of computation are executed. Moreover, they do not consider any revision of the answers.

The structure of the paper is as follows. We firstly define the framework for speculative constraint processing and semantics of the framework. Then, we describe an operational model, show an example of an execution, and state correctness of our model. Finally, we discuss space complexity issues, before concluding.

2 Speculative Constraint Processing

In this section, we provide a framework of speculative constraint computation based on the CLP framework [1]. This framework is designed so that an agent not only performs speculative constraint processing but also accepts revised and alternative answers. We then define the semantics of this framework, in Sub-section 2.2.

2.1 Framework Definition

Definition 1. Let Σ be a finite set of constants. We call an element in Σ a slave agent identifier. An atom is of the form either $p(t_1, \dots, t_n)$ or $p(t_1, \dots, t_n)@S$, where p is a predicate, $t_i (1 \leq i \leq n)$ is a term, and S is in Σ .

We call an atom with an agent identifier an “askable atom”, and an atom without an identifier a “non-askable atom”.

Definition 2. A framework for speculative constraint computation, in a master-slave system, is a triple $\langle \Sigma, \Delta, \mathcal{P} \rangle$, where:

- Σ is a finite set of constants;
- Δ is a set of rules of the following form, called default rule w.r.t. $Q@S$:

$$Q@S \leftarrow C||,$$

where $Q@S$ is an askable atom, each of whose arguments is a variable, and C is a set of constraints, called default constraint for $Q@S$;

- \mathcal{P} is a constraint logic program, that is, a set of rules R of the form:

$$H \leftarrow C \parallel B_1, B_2, \dots, B_n,$$

where:

- H is a non-askable atom; we refer to H as the head of R , denoted as $\text{head}(R)$;
- C is a set of constraints, called the constraint of R , and denoted as $\text{const}(R)$;
- each B_i of B_1, \dots, B_n is either an askable atom or a non-askable atom, and we refer to B_1, \dots, B_n as the body of R denoted as $\text{body}(R)$.

Note that a default is not necessarily specified for every askable atom. Moreover, we allow multiple defaults for the same askable atom.

Example 1. We consider the following example of a hotel room reservation. There is a master agent m : m asks travelers a and b . If both travel, m reserves a twin room. If only one of them travels, m reserves a single room. Agent m has default information about the status of a and b for days 1, 2 and 3, but the real status will be obtained directly from a and b , and the status is therefore likely to be changed.

This example can be represented as the following multi-agent system $\langle \Sigma, \Delta, \mathcal{P} \rangle^1$:

- Σ is the set of slave agents. Here, there is one master agent, m , and two slave agents, a and b . Therefore, $\Sigma = \{a, b\}$.
- Δ is the set of default information (default rules), assumed by the master agent. In particular, let us suppose that m assumes that a is free on days 1 and 2, but busy on day 3, and that b is free on day 2, and busy on day 1. Then the corresponding set Δ is as follows:

$$\Delta = \{ d_1 : fr(D)@a \leftarrow D=1 \parallel, \\ d_2 : fr(D)@a \leftarrow D=2 \parallel, d_3 : bs(D)@a \leftarrow D=3 \parallel, \\ d_4 : fr(D)@b \leftarrow D=2 \parallel, d_5 : bs(D)@b \leftarrow D=1 \parallel. \}$$

Let us remark that it is not necessary for default information to exist for all cases. In particular, m has no default information concerning the status of b on day 3.

- \mathcal{P} is a constraint logic program, to be solved by agent m . In our case of the hotel room reservation with the two travelers, it is made of the following set of rules:

¹ A string beginning with an upper-case letter represents a variable and a string beginning with a lower-case letter represents a constant. We abbreviate “free” as fr , “busy” as bs , “travel” as $trvl$, “reserve” as rsv , “twin room” as tr , and “single room” as sr .

$$\begin{aligned}
rsv(R, L, D) &\leftarrow R = tr, L = [a, b] \parallel fr(D)@a, fr(D)@b. \\
rsv(R, L, D) &\leftarrow R = sr, L = [a] \parallel fr(D)@a, bs(D)@b. \\
rsv(R, L, D) &\leftarrow R = sr, L = [b] \parallel bs(D)@a, fr(D)@b.
\end{aligned}$$

In order to solve this constraint satisfaction problem, agent m will have to ask agents a and b about $fr(D)@a$, $bs(D)@a$, $fr(D)@b$, $bs(D)@b$.

2.2 Semantics of Speculative Constraint Processing

For the semantics of the above framework, we index the semantics of a constraint logic program by a *reply set*, which specifies a reply for an askable atom.

Definition 3. A reply set is a set of rules in the form:

$$Q@S \leftarrow C \parallel,$$

where $Q@S$ is an askable atom, each of whose arguments is a variable, and C is a constraint over these variables.

Let $\langle \Sigma, \Delta, \mathcal{P} \rangle$ be a framework for speculative constraint computation, and \mathcal{R} be a reply set. A belief state w.r.t. \mathcal{R} and Δ is a reply set defined as:

$$\mathcal{R} \cup \{ "Q@S \leftarrow C \parallel" \in \Delta \mid \neg \exists C' \text{ s.t. } "Q@S \leftarrow C' \parallel" \in \mathcal{R} \}$$

and denoted as $BEL(\mathcal{R}, \Delta)$.

We introduce the above belief state since, if the answer is not returned, we use a default rule for an unreplied askable atom.

Definition 4. A goal is of the form $\leftarrow C \parallel B_1, \dots, B_n$, where C is a set of constraints and the B_i 's are atoms. We call C the constraint of the goal and B_1, \dots, B_n the body of the goal.

Definition 5. A reduction of a goal $\leftarrow C \parallel B_1, \dots, B_n$ w.r.t. a constraint logic program \mathcal{P} , a reply set \mathcal{R} , and an atom B_i , is a goal $\leftarrow C' \parallel B'$ such that:

- there is a rule R in $\mathcal{P} \cup \mathcal{R}$ s.t. $C \wedge (B_i = \text{head}(R)) \wedge \text{const}(R)$ is consistent²;
- $C' = C \wedge (B_i = \text{head}(R)) \wedge \text{const}(R)$;
- $B' = \{B_1, \dots, B_{i-1}, B_{i+1}, \dots, B_n\} \cup \text{body}(R)$.

Definition 6. A derivation of a goal $G = \leftarrow C \parallel Bs$ w.r.t. a framework for speculative constraint computation $\mathcal{F} = \langle \Sigma, \Delta, \mathcal{P} \rangle$ and a reply set \mathcal{R} is a sequence of reductions " $\leftarrow C \parallel Bs$ ", ..., " $\leftarrow C' \parallel \emptyset$ "³ w.r.t. \mathcal{P} and $BEL(\mathcal{R}, \Delta)$, where in each reduction step, an atom in the body of the goal in each step is selected. C' is called an answer constraint w.r.t. G , \mathcal{F} , and \mathcal{R} . We call a set of all answer constraints w.r.t. G , \mathcal{F} , and \mathcal{R} the semantics of G w.r.t. \mathcal{F} and \mathcal{R} .

² A notation $B_i = \text{head}(R)$ represents a conjunction of constraints equating the arguments of atoms B_i and $\text{head}(R)$.

³ \emptyset denotes an empty goal.

In the above definition, we only consider the most recent reply set, whereas a reply set might be varied during execution according to the slave agent's answer revision. We use the most recent reply set because it reflects the current situation of the slave agents. Let us remark that the order of reply messages is assumed to be preserved; that is, reply messages are always received by the master agent in the order that they are sent by the slave agent.

3 Operational Model for Speculative Computation with Iterative Answer Revision

3.1 Overview of the Operational Model

The execution of the speculative framework is based on two phases: a *process reduction phase* and a *fact arrival phase*. The process reduction phase is a normal execution of a program in a master agent, and the fact arrival phase is an interruption phase when an answer arrives from a slave agent.

For the operational model, we use the following two kinds of objects: a *process* and an *answer entry*.

Each *process* represents an alternative way of computation. Processes are created when a choice point of computation is encountered, such as case splitting, default handling, and answer arrival. A process becomes a finished process when the body of the associated goal with the process becomes empty. A process fails when some used default constraints are found to contradict the newly returned answer.

An *answer entry* is used to distinguish alternative answers and to detect which old answer corresponds to the newly revised answer. This detection is done by attaching an ID to each answer. If a new answer with an ID different from any existing answer comes, it is an alternative answer. Otherwise, the new answer is considered as a revised answer to the old answer with the same ID.

Figures 1–4 intuitively explain how processes are updated according to askable atoms. In the tree, each node represents a process, but we only show constraints associated with the process. The top node represents a constraint for the original process, and the other nodes represent added constraints for the reduced processes. Let us note that we specify *true* for non-top nodes without added constraints, since the addition of the *true* constraint does not influence the solutions of existing constraints. The leaves of the process tree represent the current processes. Therefore, the processes that are not in the leaves are deleted processes.

Figure 1 shows a situation of the processes represented as a tree when an askable atom, whose reply has not yet arrived, is executed in the process reduction phase. In this case, the current process, represented by the processed constraints C , is split into two different kinds of processes: the first one is a process using default information, C_d , and is called *default process*⁴; and the other one is the current process C itself, called *original process*, suspended at this point.

⁴ In this figure, we assume that there is only one default for brevity.




Fig. 1. When $Q@S$ is processed during the process reduction phase




Fig. 2. When first answer C_f for $Q@S$ arrives

Note that, if there are multiple definitions of defaults, we will have more than one default process, but still only one suspended process. In addition, let us note that the reason for suspending the processes (which is, keeping them in memory), is that in case of a contradictory revision of the default, or the arrival of later alternative answers, it is essential to remember the original processes to be able to restore them.

When, after some reduction of the default processes (represented in Fig. 2 by dashed lines), the first answer comes from a slave agent, expressing constraint C_f for this askable literal, we update the default processes as well as the original suspended process as follows:

- Default processes are reduced to two different kinds of processes: the first kind is a process adding C_f to the problem to solve, and the other is the current process itself which is suspended at this point⁵.
- The original process is reduced to two different kinds of processes as well: the first kind is a process adding $\neg C_d \wedge C_f$, and the other is the original process, suspended at this point.

Let us remark that although the tree of processes grows, only the leaves are kept in memory.

To intuitively explain the correctness of the above process update, we define the *frontier*, which represents the computation status of all alternative derivations. A *frontier* w.r.t. a goal $\leftarrow C \parallel Bs$, a framework for speculative constraint computation $\langle \Sigma, \Delta, \mathcal{P} \rangle$, and a reply set \mathcal{R} , is a set of goals defined as follows:

⁵ Let us remark that this splitting process is similar to the splitting process above-described for the case of a first default used.

1. The set consisting of the initial goal, $\{\leftarrow C \parallel Bs\}$ is a frontier.
2. Let F be a frontier w.r.t. the above initial goal, the framework, and the reply set. If a goal G is in F , B is an atom in G , and $RGs = \{G' \mid G' \text{ is a reduction of } G \text{ w.r.t. } \mathcal{P}, BEL(\mathcal{R}, \Delta), \text{ and } B\}$, then $F \setminus \{G\} \cup RGs$ is a frontier.

Then we have the following properties.

Lemma 1. *Let $\leftarrow C \parallel Bs$ be a goal, F be a frontier of this goal, and C' be a constraint. If we add C' to the constraints of every goal in F , then the disjunctions of all answer constraints of these modified goals is logically equivalent to the disjunction of all answer constraints of the goal $\leftarrow C \wedge C' \parallel Bs$.*

Lemma 2. *Let $\leftarrow C \parallel Bs$ be a goal, \mathcal{R} be a reply set, and C' be a constraint. Then, the disjunction of answer constraints of $\leftarrow C \wedge C' \parallel Bs$ and $\leftarrow C \wedge \neg C' \parallel Bs$ is logically equivalent to the disjunction of all answer constraints of $\leftarrow C \parallel Bs$.*

Let $\leftarrow C \parallel Bs$ be a goal containing $Q@S$. Suppose that it is reduced into $\leftarrow C \wedge C_d \parallel Bs \setminus \{Q@S\}$ by a default rule “ $Q@S \leftarrow C_d$ ”. Let F be a frontier of $\leftarrow C \wedge C_d \parallel Bs \setminus \{Q@S\}$ when the first reply “ $Q@S \leftarrow C_f$ ” is returned. Since our semantics considers the most recent replies, at this point, we should consider:

$$\leftarrow C \wedge C_f \parallel Bs \setminus \{Q@S\},$$

instead of:

$$\leftarrow C \wedge C_d \parallel Bs \setminus \{Q@S\}.$$

One possibility to implement this change is that we just discard F and invoke a new goal $\leftarrow C \wedge C_f \parallel Bs \setminus \{Q@S\}$. However, in this case, we throw every computation away before F is obtained. To retain the previous computation as much as possible, we propose the following execution.

1. We add C_f to the constraint of every goal in F . Let us remark that the disjunction of all answer constraints from this new frontier is logically equivalent to the disjunction of all answer constraints of $\leftarrow C \wedge C_d \wedge C_f \parallel Bs \setminus \{Q@S\}$, as Lemma 1 states. This computation keeps the previous computation, which is consistent with the new reply (C_f).
2. In addition to the above computation, we also start computing a new goal:

$$\leftarrow C \wedge \neg C_d \wedge C_f \parallel Bs \setminus \{Q@S\}$$

to guarantee completeness. This is because the disjunction of all answer constraints derived from $\leftarrow C \wedge C_d \wedge C_f \parallel Bs \setminus \{Q@S\}$ and $\leftarrow C \wedge \neg C_d \wedge C_f \parallel Bs \setminus \{Q@S\}$ is logically equivalent to the disjunction of all answer constraints derived from $\leftarrow C \wedge C_f \parallel Bs \setminus \{Q@S\}$, as Lemma 2 states.

When an alternative answer, with the constraint C_a , comes from a slave agent (Fig. 3), we need to follow the same procedure as when the first answer comes (Fig. 2), except that now the processes handling only default information are suspended. So, this is done by splitting the suspended default process(es), in order




Fig. 3. When alternative answer C_a for $Q@S$ arrives

to obtain the answer constraints that are logically equivalent to the answer constraints of:

$$\leftarrow C \wedge C_d \wedge C_a \parallel Bs \setminus \{Q@S\},$$

as well as by splitting the suspended original process, in order to obtain the answer constraints that are logically equivalent to the answer constraints of $\leftarrow C \wedge \neg C_d \wedge C_a \parallel Bs \setminus \{Q@S\}$ (Fig. 3). By gathering these answer constraints, we can compute all answer constraints for the alternative reply.

On the other hand, when a revised answer with the constraint C_r arrives, all processes using the first (or current) answer are split, in order to obtain the answer constraints that are logically equivalent to the answer constraints of:

$$\leftarrow C \wedge C_f \wedge C_r \parallel Bs \setminus \{Q@S\},$$

and the suspended original process is split as well, in order to obtain the answer constraints that are logically equivalent to the answer constraints of $\leftarrow C \wedge \neg C_f \wedge C_r \parallel Bs \setminus \{Q@S\}$ (Fig. 4). By gathering these answer constraints, we can override the previous reply by the revised reply.

3.2 Preliminary Definitions

A process is either an *ordinary process* or a *finished process*. An *ordinary process* P is an expression of the form $\langle PID, C, GS, WA, AA \rangle$, where:

- PID : the ID for a process denoted as $pid(P)$;
- C : the current constraint in the goal denoted as $pconst(P)$;
- GS : the body in the goal denoted as $gs(P)$;
- WA : a set of pairs $\langle Q@S, WAID \rangle$, where $Q@S$ is an askable atom and $WAID$ is the ID of an answer entry whose answer is awaited by the process. We denote WA as $wa(P)$;
- AA : a set of pairs $\langle Q@S, AAID \rangle$, where $Q@S$ is an askable atom and $AAID$ is the ID of an answer entry whose answer is used in the process. We denote AA as $aa(P)$.




Fig. 4. When revised answer C_r for $Q@S$ arrives

A *finished process* FP is an expression of the form $\langle Query, FPID, C \rangle$, where:

- *Query*: an initial query for this process. It is used to send an answer to the asking agent;
- *FPID*: the ID for a process. This is also used when this answer is returned to the asking agent;
- *C*: the current constraint in the process.

For simplicity, an ordinary process is sometimes just called a process.

An *answer entry* A is an expression of the form $\langle Q@S, AID, C, UPIDs \rangle$, where:

- $Q@S$: the query given to the other agent denoted as $aq(A)$;
- AID : the ID for an answer entry denoted as $aid(A)$. We have the special IDs, “ o ” for the answer entry created when this query is firstly asked, and “ d_1 ”, ... for default answers. We call an answer entry with the ID “ o ” an *original answer entry* for $Q@S$, an answer entry with an ID of “ d_1 ”, ... a *default answer entry*, and other answer entries *ordinary answer entries*;
- C : the most recent answer constraint for $Q@S$ for answer entry A denoted as $aconst(A)$. The constraint of the original answer entry is defined as *true*;
- $UPIDs$: the set of IDs of processes using an answer in A denoted as $ups(A)$.

3.3 Process Reduction Phase

In the process reduction phase, we process the constraints in a regular CLP way. The only difference is that we may have to consider default information, or answers. In this subsection, we describe how we manage processes, following the above-given definitions.

We do the following until no more processes can be processed.

- When a query $Q_{init}@S_{self}$ is asked from another agent S' , where S_{self} is the ID for this agent, we record Q_{init} as the initial query and S' as the asking agent. We then create a new process $\langle PID, true, Q_{init}, \emptyset, \emptyset \rangle$, where PID is a new process ID.

- If there is an ordinary process P such that $gs(P) = wa(P) = \emptyset$,
 1. We send an answer to the asking agent S' that is of the form: $\langle Q_{init}@S_{self}, pid(P), pconst(P) \rangle$.
 2. We change this process into a finished process of the form: $\langle Q_{init}@S_{self}, pid(P), pconst(P) \rangle$.
- Else if there is a process P such that $gs(P) \neq \emptyset$ and $wa(P) = \emptyset$, then we select an atom L in $gs(P)$ and reduce L as follows:
 - If L is a non-askable atom,
 1. For every rule R such that $pconst(P) \wedge (L = head(R)) \wedge const(R)$ is consistent, we do the following:
 - (a) We create a new process $\langle newPID, newC, GS, \emptyset, AA \rangle$, where
 - * $newPID$ is a new process ID;
 - * $newC := pconst(P) \wedge (L = head(R)) \wedge const(R)$;
 - * $GS := body(R) \cup gs(P) \setminus \{L\}$;
 - * $AA := aa(P)$.
 - (b) For every answer entry A s.t. $\langle aq(A), aid(A) \rangle$ in $aa(P)$, $ups(A) := ups(A) \cup \{newPID\}$.
 2. For every answer entry A s.t. $\langle aq(A), aid(A) \rangle$ in $aa(P)$, $ups(A) := ups(A) \setminus \{pid(P)\}$.
 3. We delete P .
 - If L is an askable atom $Q@S$,
 1. We do either of the following according to non-arrival/arrival of the answer.
 - * If there is no ordinary answer entry of the form $\langle Q@S, AID, C, UPIDs \rangle$, then for each default “ $Q@S \leftarrow C_d$ ” such that $pconst(P) \wedge C_d$ is consistent, we do the following:
 - (a) We create a new process $\langle newPID, newC, GS, \emptyset, AA \rangle$, where
 - $newPID$ is a new process ID;
 - $newC := pconst(P) \wedge C_d$;
 - $GS := gs(P) \setminus \{Q@S\}$;
 - $AA := aa(P) \cup \{ \langle Q@S, d \rangle \}$, where d is an ID for this default.
 - (b) We associate the newly created process with a default d of $Q@S$ as follows:
 - If there is a default answer entry $A_d = \langle Q@S, d, C_d, UPIDs_d \rangle$, then $ups(A_d) := UPIDs_d \cup \{newPID\}$.
 - Else if there is no default answer of the form $\langle Q@S, d, C_d, UPIDs_d \rangle$, we create an answer entry $\langle Q@S, d, C_d, \{newPID\} \rangle$.
 - (c) For every answer entry A s.t. $\langle aq(A), aid(A) \rangle$ in $aa(P)$, $ups(A) := ups(A) \cup \{newPID\}$.
 - * Else if there is an ordinary answer entry of the form $\langle Q@S, AID, C, UPIDs \rangle$, then for each ordinary answer entry $\langle Q@S, AID, C_a, UPIDs \rangle$ s.t. $pconst(P) \wedge C_a$ is consistent, we do the following:

- (a) We create a new process $\langle newPID, newC, GS, \emptyset, AA \rangle$, where
 - $newPID$ is a new process ID;
 - $newC := pconst(P) \wedge C_a$;
 - $GS := GS \setminus \{Q@S\}$;
 - $AA := aa(P) \cup \{\langle Q@S, AID \rangle\}$.
- (b) For every answer entry A s.t. $\langle aq(A), aid(A) \rangle$ in $aa(P)$,
 $ups(A) := ups(A) \cup \{pid(P)\}$.
2. We associate P with $Q@S$ as follows:
 - * If there is an original answer entry
 $A_o = \langle Q@S, o, true, UPIDs_o \rangle$, then
 $ups(A_o) := UPIDs_o \cup \{pid(P)\}$.
 - * Else if there is no original answer entry of the form
 $\langle Q@S, o, true, UPIDs \rangle$, we create an answer entry
 $\langle Q@S, o, true, \{pid(P)\} \rangle$, and send a question Q to S .
3. $wa(P) := \{\langle Q@S, o \rangle\}$.

3.4 Fact Arrival Phase

Suppose that an answer is returned from an agent S for a question $Q@S$ of the form $\langle Q@S, AID, C \rangle$. Then, we do the following after one step of process reduction is finished.

- If there is no answer entry of the form $\langle Q@S, AID, C_f, UPIDs' \rangle^6$,
 1. We create an answer entry $\langle Q@S, AID, C, UPIDs \rangle$, where $UPIDs$ is initially set to \emptyset , but will be incremented as shown below.
 2. For every default answer entry for a default d of the form $\langle Q@S, d, C_d, UPIDs_d \rangle$ and for every process P_d such that $pid(P_d) \in UPIDs_d$, we do the following:
 - If P_d is a finished process of the form $\langle Q_{init}@S_{self}, PID, C_{Final} \rangle$ s.t. $C \wedge C_{Final} \neq C_{Final}$, we send an answer of the form $\langle Q_{init}@S_{self}, PID, C \wedge C_{Final} \rangle$ to the asking agent S' .
 - If P_d is an ordinary process, we do the following:
 - (a) $wa(P_d) := wa(P_d) \cup \{\langle Q@S, d \rangle\}$.
 - (b) $aa(P_d) := aa(P_d) \setminus \{\langle Q@S, d \rangle\}$.
 - (c) If $C \wedge pconst(P_d)$ is consistent, we do the following:
 - i. We create a new process $\langle newPID, newC, GS, WA, AA \rangle$, where
 - * $newPID$ is a new process ID;
 - * $newC := C \wedge pconst(P_d)$;
 - * $GS := gs(P_d)$;
 - * $WA := wa(P_d) \setminus \{\langle Q@S, d \rangle\}$;
 - * $AA := aa(P_d) \cup \{\langle Q@S, AID \rangle\} \setminus \{\langle Q@S, d \rangle\}$.
 - ii. $UPIDs := UPIDs \cup \{newPID\}$.
 3. Pick up the original answer entry of the form $\langle Q@S, o, true, UPIDs_o \rangle$.
 4. For every process P_o such that $pid(P_o) \in UPIDs_o$ and $C \wedge pconst(P_o) \wedge \bigwedge_{(Q@S \leftarrow C_d) \in \Delta} \neg C_d$ is consistent, do the following:

⁶ This means that the arriving answer is a first or alternative answer to the query $Q@S$.

- (a) We create a new process $\langle newPID, newC, GS, WA, AA \rangle$, where
- $newPID$ is a new process ID;
 - $newC := C \wedge pconst(P_o) \wedge \bigwedge_{(Q@S \leftarrow C_d) \in \Delta} \neg C_d$;
 - $GS := gs(P_o)$;
 - $WA := wa(P_o) \setminus \{ \langle Q@S, o \rangle \}$;
 - $AA := aa(P_o) \cup \{ \langle Q@S, AID \rangle \}$.
- (b) $UPIDs := UPIDs \cup \{ newPID \}$.
- Else if there is an answer entry of the form $\langle Q@S, AID, C_f, UPIDs' \rangle^7$,
1. We change $\langle Q@S, AID, C_f, UPIDs' \rangle$ into $\langle Q@S, AID, C, UPIDs \rangle$, where $UPIDs := UPIDs'$ initially but will be incremented/decremented as shown below.
 2. For every process P such that $pid(P) \in UPIDs'$ do the following:
 - If P is a finished process of the form $\langle Q_{init}@S_{self}, PID, C_{Final} \rangle$ s.t. $C \wedge C_{Final} \neq C_{Final}$, we send an answer of the form $\langle Q_{init}@S_{self}, PID, C \wedge C_{Final} \rangle$ to the asking agent S' .
 - If P is an ordinary process, we do the following:
 - * If $C \wedge pconst(P)$ is consistent, $pconst(P) := C \wedge pconst(P)$.
 - * Otherwise, delete P and $UPIDs := UPIDs \setminus \{ pid(P) \}$.
 3. Pick up the original answer entry of the form $\langle Q@S, o, true, UPIDs_o \rangle$.
 4. For every process P_o such that $pid(P_o) \in UPIDs_o$ and $C \wedge pconst(P_o) \wedge \neg C_f$ is consistent, we do the following:

(a) We create a new process $\langle newPID, newC, GS, WA, AA \rangle$, where

 - $newPID$ is a new process ID;
 - $newC := C \wedge pconst(P_o) \wedge \neg C_f$;
 - $GS := gs(P_o)$;
 - $WA := wa(P_o) \setminus \{ \langle Q@S, o \rangle \}$;
 - $AA := aa(P_o) \cup \{ \langle Q@S, AID \rangle \}$.

(b) $UPIDs := UPIDs \cup \{ newPID \}$.

3.5 Execution Trace Example

We show a part of an execution trace for a question $rsv(R, L, D)$ in Example 1. In this trace, we consider a scenario that highlights process updates upon arrivals of an alternative answer and a revised answer. We firstly give the initial process $\langle p_0, true, \{ rsv(R, L, D) \}, \emptyset, \emptyset \rangle$.

1. Select process p_0 and reduce it to p_1, p_2, p_3 .

Processes:

$$\begin{aligned} &\langle p_1, \{ R = tr, L = [a, b] \}, \{ fr(D)@a, fr(D)@b \}, \emptyset, \emptyset \rangle, \\ &\langle p_2, \{ R = sr, L = [a] \}, \{ fr(D)@a, bs(D)@b \}, \emptyset, \emptyset \rangle, \\ &\langle p_3, \{ R = sr, L = [b] \}, \{ bs(D)@a, fr(D)@b \}, \emptyset, \emptyset \rangle. \end{aligned}$$

⁷ This means that the arriving answer is a revised answer of one of the previous answers to the query $Q@S$.

2. Select p_1 , and ask a question $fr(D)@a$, and create answer entries for $fr(D)@a$ and new processes p_4, p_5 for default answers.

Answer entries:

$$\begin{aligned} &\langle fr(D)@a, o, true, \{p_1\} \rangle, \\ &\langle fr(D)@a, d_1, \{D = 1\}, \{p_4\} \rangle, \\ &\langle fr(D)@a, d_2, \{D = 2\}, \{p_5\} \rangle. \end{aligned}$$

Processes: p_2, p_3 ,

$$\begin{aligned} &\langle p_4, \theta_{tr} \cup \{D = 1\}, \{fr(D)@b\}, \emptyset, \{\langle fr(D)@a, d_1 \rangle\} \rangle^8, \\ &\langle p_5, \theta_{tr} \cup \{D = 2\}, \{fr(D)@b\}, \emptyset, \{\langle fr(D)@a, d_2 \rangle\} \rangle, \\ &\langle p_1, \theta_{tr}, \{fr(D)@b\}, \{\langle fr(D)@a, o \rangle\}, \emptyset \rangle. \end{aligned}$$

3. Suppose that $\langle fr(d)@a, a_1, \{D = 2\} \rangle$ is returned from agent a . We suspend p_4 and p_5 since they use a default answer and then create new processes p_6 from p_5 since the default answer used in p_5 is consistent with the returned answer. Note that we create no new process from p_1 since the returned answer contradicts one of the negations of default answers.

Answer entries: $fra_o, fra_{d_1}, fra_{d_2}$ ⁹,

$$\langle fr(D)@a, a_1, \{D = 2\}, \{p_6\} \rangle.$$

Processes: p_1, p_2, p_3 ,

$$\begin{aligned} &\langle p_6, \theta_{tr2}, \{fr(D)@b\}, \emptyset, \{\langle fr(D)@a, a_1 \rangle\} \rangle, \\ &\langle p_4, \theta_{tr1}, \{fr(D)@b\}, \{\langle fr(D)@a, d_1 \rangle\}, \emptyset \rangle, \\ &\langle p_5, \theta_{tr2}, \{fr(D)@b\}, \{\langle fr(D)@a, d_2 \rangle\}, \emptyset \rangle^{10}. \end{aligned}$$

4. Suppose that $\langle fr(D)@a, a_2, \{D = 3\} \rangle$ is returned from the agent a . Since this has a different answer ID from the previous answer in the last step, this answer is an alternative answer. Then, we create a new process from p_1 that is the original process for query $fr(D)@a$. Note that we create no new process from the processes created by default answers for $fr(D)@a$ since this answer contradicts the defaults.

Answer entries: $fra_o, fra_{d_1}, fra_{d_2}, fra_{a_1}$ ¹¹,

$$\langle fr(D)@a, a_2, \{D = 3\}, \{p_7\} \rangle.$$

Processes: $p_1, p_2, p_3, p_4, p_5, p_6$,

$$\langle p_7, \theta_{tr} \cup \{D = 3, D \neq 1, D \neq 2\}, \{fr(D)@b\}, \emptyset, \{\langle fr(D)@a, a_2 \rangle\} \rangle.$$

5. Suppose that $\langle fr(D)@a, a_1, \{D = 1\} \rangle$ is returned from the agent a . The ID a_1 for the returned answer indicates that this answer is a revised answer for “ $D = 2$ ”. Therefore, we revise every process using a_1 , which is recorded in the answer entry fra_{a_1} . This is p_6 , but its associated constraint is contradictory to the returned answer, and therefore we kill this process. Then, we create a new process p_8 from p_1 .

⁸ $\theta_{tr} = \{R = tr, L = [a, b]\}$.

⁹ $fra_o = \langle fr(D)@a, o, true, \{p_1\} \rangle,$
 $fra_{d_1} = \langle fr(D)@a, d_1, \{D = 1\}, \{p_4\} \rangle,$
 $fra_{d_2} = \langle fr(D)@a, d_2, \{D = 2\}, \{p_5\} \rangle.$

¹⁰ $\theta_{tr2} = \theta_{tr} \cup \{D = 2\}$ and $\theta_{tr1} = \theta_{tr} \cup \{D = 1\}$.

¹¹ $fra_{a_1} = \langle fr(D)@a, a_1, \{D = 2\}, \{p_6\} \rangle.$

Answer entries: $fra_o, fra_{d_1}, fra_{d_2}, fra_{a_2}$ ¹²,
 $\langle fr(D)@a, a_1, \{D = 1\}, \{p_8\} \rangle$.

Processes: $p_1, p_2, p_3, p_4, p_5, p_7,$

$\langle p_8, \theta_{tr} \cup \{D = 1, D \neq 2\}, \{fr(D)@b\}, \emptyset, \{\langle fr(D)@a, a_1 \rangle\} \rangle$.

4 Correctness of the Operational Model

We guarantee that the above operational model gives a correct answer w.r.t. the most recent replies. Let us note that the order of reply messages is assumed to be preserved.

Theorem 1. *Let $\langle \Sigma, \Delta, \mathcal{P} \rangle$ be a framework for speculative constraint computation. Suppose that there is an ordinary process P such that $gs(P) = wa(P) = \emptyset$ for the initial query Q_{init} . Let*

$$\mathcal{R} = \{ \text{“}Q@S \leftarrow C\text{”} \mid \text{there exists an answer entry } \langle Q@S, AID, C, UPIDs \rangle \\ \text{s.t. } \langle Q@S, AID \rangle \in aa(P) \}.$$

Then, there exists an answer constraint C' w.r.t. Q_{init} , the framework, and \mathcal{R} s.t. $\pi_V(pconst(P))$ entails $\pi_V(C')$, where V is the set of the variables that occur in Q_{init} , and π_V is the projection of constraints onto V .

Proof Sketch. See Appendix. □

5 Space Complexity of Our Approach

Our approach, compared to traditional approaches (no belief revision), generates an additional cost in terms of space. In this section, we briefly show that the additional cost in space is linear. This cost is observed based on the size of the set PS of processes related to the revised or alternative answer to handle.

When a revised answer comes, say C_r , as shown in Fig. 4:

- If C_r entails the previous answer, say C_f , PS either remains the same size, or reduces (because some processes in PS may now have inconsistent constraints and therefore be killed);
- If C_r is inconsistent with C_f , then all the processes using C_f in PS are killed, the original suspended processes are duplicated and resumed with C_r , and therefore PS grows by at most, the number of original suspended processes;
- If C_r is consistent with C_f but does not entail it, PS grows by at most, the number of original suspended processes.

These three cases exhibit only linear (or less) behavior.

When an alternative answer comes, say C_a , as shown in Fig. 3, all the processes suspended by the first answer, as well as the original suspended processes, are duplicated and resumed with C_a . Therefore, PS grows by at most, the number of these suspended processes.

¹² $fra_{a_2} = \langle fr(D)@a, a_2, \{D = 3\}, \{p_7\} \rangle$.

As briefly covered here, the growth of the set of processes on the arrival of revised and alternative answers follows a linear behavior.

6 Conclusion

In this paper, we presented an operational model for speculative constraint processing with iterative revision for alternative answers. This paper is a generalization of two previous works; the work of revisable speculative computation for yes/no questions [7] and the work of non-revisable speculative computation for queries with constraints [5].

As for future work, we will prove the correctness and completeness for more general forms of multi-agent systems, where every agent can perform speculative computation. Our current framework is focused on master-slave multi-agent systems, and defines the operational model of the master agents. To handle a more general multi-agent system, we need to guarantee the appropriate computation of the overall system by additionally considering communication paths among agents. For another direction, we will also consider applications for this framework.

References

1. J. Jaffar, M. J. Maher, K. Marriott, and P. J. Stuckey. The semantics of constraint logic programs. *Journal of Logic Programming*, 37(1-3):1–46, 1998.
2. S. Janson and S. Haridi. Programming paradigms of the andorra kernel language. In *Proc. of ISLP'91*, pages 167–186, 1991.
3. R. A. Kowalski and F. Sadri. From logic programming towards multi-agent systems. *Annals of Mathematics and Artificial Intelligence*, 25:391–419, 1999.
4. F. Sadri and F. Toni. Interleaving belief revision and reasoning: Preliminary report. In *Proc. of Convegno Italiano di Logica Computazionale (CILC)*, 2005.
5. K. Satoh, P. Codognot, and H. Hosobe. Speculative constraint processing in multi-agent systems. In *Proc. of PRIMA2003*, volume 2891 of *LNCS*, pages 133–144, 2003.
6. K. Satoh, K. Inoue, K. Iwanuma, and C. Sakama. Speculative computation by abduction under incomplete communication environments. In *Proc. of ICMAS2000*, pages 263–270, 2000.
7. K. Satoh and K. Yamamoto. Speculative computation with multi-agent belief revision. In *Proc. of AAMAS2002*, pages 897–904, 2002.
8. C. Schulte. *Programming Constraint Services: High-Level Programming of Standard and New Constraint Services*, volume 2302 of *LNCS*. 2002.

Appendix

Proof Sketch of Theorem 1. To prove the property described in Theorem 1, we show that a more general property holds for any existing ordinary process at any “step” in the process reduction or fact arrival phase. By a “step”, we mean the execution of operations in the process reduction or fact arrival phase from its beginning to its end, without returning to the beginning, and without transferring to

the other phase. Then the property that we show is the following: at any n -th step, for any ordinary process P' , there exists a sequence of reductions “ $\leftarrow \parallel Q_{init}$ ”, \dots ,

$$“\leftarrow C'' \parallel \{Q@S \mid \langle Q@S, o \rangle \in wa(P')\} \cup gs(P')”$$

w.r.t. \mathcal{P} and $BEL(\mathcal{R}_{P'}^{(n)}, \Delta)$, such that

$$\pi_V(pconst(P')) \text{ entails } \pi_V(C''),$$

where $\mathcal{R}_{P'}^{(n)}$ is the most recent reply set for P' at the n -th step, which is defined in the same way as \mathcal{R} in Theorem 1.

Below we prove this property by induction on the progress of process reduction and fact arrival steps.

Induction base. When a query $Q_{init}@S_{self}$ is asked in the initial step, a process $P' = \langle PID, true, Q_{init}, \emptyset, \emptyset \rangle$ is created. This process corresponds to the initial goal “ $\leftarrow \parallel Q_{init}$ ”. The above property holds since $pconst(P') = true$ and $C'' = true$.

Induction step. Assume that, at the n -th step, the property holds.

Now consider the $(n+1)$ -th step. It is straightforward to show that the property holds for the process reduction phase.

Here we consider the processing of a first or alternative answer in the fact arrival phase. Let the returned answer be $\langle Q@S, AID, C \rangle$. In this case, there is no answer entry in the form $\langle Q@S, AID, C_f, UPID_s' \rangle$.

Let $\langle Q@S, d, C_d, UPID_{s_d} \rangle$ be any default answer entry and P_d be any ordinary process such that $pid(P_d) \in UPID_{s_d}$. By the induction hypothesis, P_d satisfies the above property for some C'' and $\mathcal{R}_{P_d}^{(n)}$; that is, there is a sequence of reductions “ $\leftarrow \parallel Q_{init}$ ”, \dots , “ $\leftarrow C_1 \parallel \{Q@S\} \cup GS$ ”, “ $\leftarrow C_1 \wedge C_d \parallel GS$ ”, \dots , “ $\leftarrow C_1 \wedge C_d \wedge C_2 \parallel \{Q'@S' \mid \langle Q'@S', o \rangle \in wa(P_d)\} \cup gs(P_d)$ ” w.r.t. \mathcal{P} and $BEL(\mathcal{R}_{P_d}^{(n)}, \Delta)$, such that $\pi_V(pconst(P_d))$ entails $\pi_V(C_1 \wedge C_d \wedge C_2)$, where C_1 and C_2 are the constraints obtained before and after processing $Q@S$, respectively.

Assume that $C \wedge pconst(P_d)$ is consistent. Then a process $P' = \langle newPID, C \wedge pconst(P_d), gs(P_d), wa(P_d) \setminus \{\langle Q@S, d \rangle\}, aa(P_d) \cup \{\langle Q@S, AID \rangle\} \setminus \{\langle Q@S, d \rangle\} \rangle$ is created, and we have $\mathcal{R}_{P'}^{(n+1)} = \mathcal{R}_{P_d}^{(n)} \cup \{Q@S \leftarrow C\} \setminus \{Q@S \leftarrow C_d\}$. Then we can consider the sequence of reductions “ $\leftarrow \parallel Q_{init}$ ”, \dots , “ $\leftarrow C_1 \parallel \{Q@S\} \cup GS$ ”, “ $\leftarrow C_1 \wedge C \parallel GS$ ”, \dots , “ $\leftarrow C_1 \wedge C \wedge C_2 \parallel \{Q'@S' \mid \langle Q'@S', o \rangle \in wa(P')\} \cup gs(P')$ ” w.r.t. \mathcal{P} and $BEL(\mathcal{R}_{P'}^{(n+1)}, \Delta)$. Then, $\pi_V(pconst(P'))$ entails $\pi_V(C_1 \wedge C \wedge C_2)$ since $pconst(P') = C \wedge pconst(P_d)$ and $\pi_V(pconst(P_d))$ entails $\pi_V(C_1 \wedge C_d \wedge C_2)$. Thus, the above property holds for P' .

For the processing of a first answer, this step changes P_d by setting $wa(P_d) := wa(P_d) \cup \{\langle Q@S, d \rangle\}$ and $aa(P_d) := aa(P_d) \setminus \{\langle Q@S, d \rangle\}$, and hence we have $\mathcal{R}_{P_d}^{(n+1)} = \mathcal{R}_{P_d}^{(n)} \setminus \{Q@S \leftarrow C_d\}$. In the other case (that is, for processing an alternative answer), P_d is unchanged since $\langle Q@S, d \rangle \in wa(P_d)$ and $\langle Q@S, d \rangle \notin aa(P_d)$ hold for the original P_d , and therefore, we have $\mathcal{R}_{P_d}^{(n+1)} = \mathcal{R}_{P_d}^{(n)}$. In both cases, $BEL(\mathcal{R}_{P_d}^{(n+1)}, \Delta) = BEL(\mathcal{R}_{P_d}^{(n)}, \Delta)$ since “ $Q@S \leftarrow C_d$ ” $\in \Delta$. Therefore, the above property is kept satisfied for P_d .

Next, let $\langle Q@S, o, true, UPIDS_o \rangle$ be the original answer entry and P_o be any ordinary process such that $pid(P_o) \in UPIDS_o$. By the induction hypothesis, P_o satisfies the above property for some C'' and $\mathcal{R}_{P_o}^{(n)}$; that is, there is a sequence of reductions “ $\leftarrow \parallel Q_{init}$ ”, \dots , “ $\leftarrow C'' \parallel \{Q@S\} \cup gs(P_o)$ ” w.r.t. \mathcal{P} and $BEL(\mathcal{R}_{P_o}^{(n)}, \Delta)$, such that $\pi_V(pconst(P_o))$ entails $\pi_V(C'')$. Since this step does not change P_o , the above property is kept satisfied for P_o .

Assume that $C \wedge pconst(P_o) \wedge \bigwedge_{(Q@S \leftarrow C_d) \in \Delta} \neg C_d$ is consistent. Then a process $P' = \langle newPID, C \wedge pconst(P_o) \wedge \bigwedge_{(Q@S \leftarrow C_d) \in \Delta} \neg C_d, gs(P_o), wa(P_o) \setminus \{\langle Q@S, o \rangle\}, aa(P_o) \cup \{\langle Q@S, AID \rangle\} \rangle$ is created, and we have $\mathcal{R}_{P'}^{(n+1)} = \mathcal{R}_{P_o}^{(n)} \cup \{Q@S \leftarrow C\}$. Then we can consider the sequence of reductions “ $\leftarrow \parallel Q_{init}$ ”, \dots , “ $\leftarrow C'' \parallel \{Q@S\} \cup gs(P')$ ”, “ $\leftarrow C'' \wedge C \parallel gs(P')$ ” w.r.t. \mathcal{P} and $BEL(\mathcal{R}_{P'}^{(n+1)}, \Delta)$. Then $\pi_V(pconst(P'))$ entails $\pi_V(C'' \wedge C)$ since $pconst(P') = C \wedge pconst(P_o) \wedge \bigwedge_{(Q@S \leftarrow C_d) \in \Delta} \neg C_d$ and $\pi_V(pconst(P_o))$ entails $\pi_V(C'')$. Therefore, the above property holds for P' .

The above property is kept satisfied for the other processes that are not handled in this case, since those processes and their most recent reply sets are unchanged.

Therefore, the above property holds for any processes after processing a first or alternative answer in the fact arrival phase.

Similarly, we can show that the above property holds for the processing of a revised answer in the fact arrival phase. Thus, the above property holds in all the cases.

Since the property described in Theorem 1 corresponds to the special case of the above property, where $gs(P') = wa(P') = \emptyset$, Theorem 1 holds. \square