Distributed generative CSP framework for multi-site product configuration

Alexander Felfernig[1], **Gerhard Friedrich**[1], **Dietmar Jannach**[1], **and Markus Zanker**[1]

Abstract. Today's configurators are centralized systems and do not allow manufacturers to cooperate on-line for offer-generation or sales-configuration. However, supply chain integration of configurable products requires the cooperation of the configuration systems from the different manufacturers that jointly offer solutions to customers. As a consequence, there is a high potential for methods that enable the computation of such configurations by independent specialized agents. Several approaches to *centralized* configuration tasks are based on constraint satisfaction problem (CSP) solving. Most of them extend the traditional CSP approach in order to comply to the specific expressivity and dynamism requirements for configuration and similar synthesis tasks.
The distributed generative CSP (DisGCSP) framework proposed here builds on a CSP formalism that encompasses the *generative* aspect of variable creation and extensible domains of problem variables. It also builds on the distributed CSP (DisCSP) framework, allowing for approaches to configuration tasks where the knowledge is distributed over a set of agents. Notably, the notions of constraint and nogood are generalized to an additional level of abstraction, extending inferences to types of variables. The usage of the new framework is exemplified by describing modifications to some complete algorithms for DisCSP when targeting DisGCSPs.

1 Introduction/Background

The paradigm of mass-customization allows customers to tailor (configure) a product or service according to their specific needs, i.e. the customer can select between several features and options that should be included in the configured product and can determine the physical component structure of the personalized product variant. Typically, there are several technical and marketing restrictions on the legal parameter constellations and the physical layout. This led manufacturers to develop support for checking the feasibility of user requirements and for computing a consistent solution. This functionality is provided by product configuration systems (configurators), whereby they have shown to be a successful application area for different AI techniques [15] such as description logics [8], or rule-based [1] and constraint-based solving algorithms. [4] describes the industrial use of constraint techniques for the configuration of large and complex systems such as telecommunication switches and [7] is an example of a powerful tool based on Constraint Satisfaction available on the market.
However, companies find themselves in dynamically determined coalitions with other highly specialized solution providers that jointly offer customized solutions. This high integration aspect of todays digital markets implies that software systems supporting the selling and configuration task must no longer be conceived as standalone systems. A product configurator can be therefore seen as an agent with private knowledge that acts on behalf of its company and cooperates with other agents to solve a configuration task. This paper abstracts the *centralized* definition of a configuration task in [16] to a more general definition of a *generative* CSP that is also applicable to the wider range of synthesis problems. Furthermore, we propose a framework that allows to address distributed configuration tasks by extending DisCSPs with the innovative aspects of local generative CSPs:

1. The constraints (and nogoods) are generalized to a form where they can depend on types rather than on identities of variables. This also enables an elegant treatment of the next aspects.
2. The number of variables of certain types that are active in the local CSP of an agent, may vary depending on the state of the search process. In the DisCSP framework, the external variables existing in the system are predetermined, but here the set of variables defining the problem is determined dynamically.
3. The domain of the variables may vary dynamically. Some variables model possible connections and they depend on the existence of components that could become connected.

We also describe the interesting impact of the previously mentioned changes on asynchronous algorithms. In the following we motivate our approach with an example, Section 3 defines a generative CSP and in Section 4 distributed generative CSP is formalized and extensions to current DisCSP frameworks are presented.

2 Motivating example

For the purpose of illustration of our approach we chose as example domain the well known N-queens problem. The characteristics of a distributed configuration problem or similar distributed synthesis tasks are integrated into our N-queens scenario: (a) parts of the problem (i.e., variables) are shared among agents and (b) the problem is dynamically extended (i.e., N is increased), if no solution can be found. Adding additional problem variables leads to domain extensions and thus to a larger search- and solution space. The goal is to place N queens on distinct squares in an $N \times N$ chess board, where no two queens threaten each other [17]. We formalize the problem by making each row of the board a problem variable x_i, where the subscript i ensures unique variable names. In a distributed setting we employ three agents, each owning a fraction of the constraints necessary to solve the N-queens problem. Furthermore, we want to show the *generative* aspect of problem solving in the example, where agents start with a representation of a 0-queens problem and specific requirements on the final solution coming from outside.

[1] Computer Science and Manufacturing, Universität Klagenfurt, Universitätsstrasse 65-67, 9020 Klagenfurt, Austria. e-mail: {felfernig, friedrich, jannach, zanker}@ifit.uni-klu.ac.at

Once the agents determine that a solution cannot be found, they extend the problem space by adding an additional row which in consequence enlarges the domain of row variables by one. Since the exact number of problem variables is not known from the beginning, constraints cannot be directly formulated on concrete variables. Instead, comparable to programming languages, variable types exist that allow to associate a newly created variable with a domain and we can specify relationships in terms of *generic constraints*. [16] define a generic constraint γ as a constraint schema, where meta-variables V^t act as placeholders for concrete variables of a specific type t[2]. In our example three types of problem variables exist, representing the even (t_e) and the uneven rows (t_u) as well as a type (t_c) of counter variables (x_{type}) for the number of instantiations of each type, which allows us to distribute the N-queens constraints among the agents. Therefore, each agent a_i posesses a set of private constraints Γ^{a_i}, i.e., $\Gamma^{a_1} = \{\gamma_1, \gamma_2, \gamma_3, \gamma_8, \gamma_9\}$, $\Gamma^{a_2} = \{\gamma_4, \gamma_5, \gamma_8, \gamma_9\}$ and $\Gamma^{a_3} = \{\gamma_6, \gamma_7, \gamma_8, \gamma_9\}$, that are defined as follows:
$\gamma_1 : val(x_{t_u}) = val(x_{t_e}) \vee val(x_{t_u}) = val(x_{t_e}) + 1$, where $val(x)$ is a predicate that gives the assigned value of variable x.
Informally, the number of uneven rows may exceed the number of even rows by one.
$\gamma_2 : val(V^{t_u}) \neq val(V^{t_e})$
No two queens on an even and an uneven row are allowed to take the same column value.
$\gamma_3 : abs(2 \times (index(V^{t_u}) - index(V^{t_e})) - 1) \neq abs(val(V^{t_u}) - val(V^{t_e}))$, where $index(x)$ returns a number i indicating that x is the i^{th} variable of its type and $abs(n)$ is a predicate that returns the absolute value of n.
No two queens on an even and an uneven row are allowed to be on the same diagonal.
$\gamma_4 : V_1^{t_u} \neq V_2^{t_u} \rightarrow val(V_1^{t_u}) \neq val(V_2^{t_u})$.
No two queens on uneven rows are allowed to take the same column value.
$\gamma_5 : V_1^{t_u} \neq V_2^{t_u} \rightarrow abs(2 \times (index(V_1^{t_u}) - index(V_2^{t_u}))) \neq abs(val(V_1^{t_u}) - val(V_2^{t_u}))$.
No two queens on uneven rows are allowed to be on the same diagonal.
$\gamma_6 : V_1^{t_e} \neq V_2^{t_e} \rightarrow val(V_1^{t_e}) \neq val(V_2^{t_e})$.
No two queens on even rows are allowed to take the same column value.
$\gamma_7 : V_1^{t_e} \neq V_2^{t_e} \rightarrow abs(2 \times (index(V_1^{t_e}) - index(V_2^{t_e})) \neq abs(val(V_1^{t_e}) - val(V_2^{t_e}))$.
No two queens on even rows are allowed to be on the same diagonal.
$\gamma_8 : val(V^{t_u}) \leq x_{t_e} + x_{t_u}$. $\gamma_9 : val(V^{t_e}) \leq x_{t_e} + x_{t_u}$.
The latter two constraints delimit the domain of row variables to the total number of rows.
Figure 1 depicts the initial situation, with a 0-queens problem. The customer requests agent a_1 to satisfy the requirement of finding a solution containing at least two uneven rows:
$\gamma_{cust} : x_{t_u} \geq 2$.
Having added γ_{cust} to the set of private constraints of agent a_1, the search process starts and the solution space is continuously extended by the instantiation of additional problem variables, until a solution is found for a 4-queens problem that satisfies all local constraints of the agents. The links between two agents indicate that they share variables, which is described in more detail later on. Thus, a solution to a generative constraint satisfaction problem requires not only finding valid assignments to variables, but also determining the exact size of the problem itself. In the sequel of the paper we define a model for the local configurators and we detail extensions to DisCSP algorithms.

[2] The exact semantics of generic constraints is given in Definition 2 in Section 3.

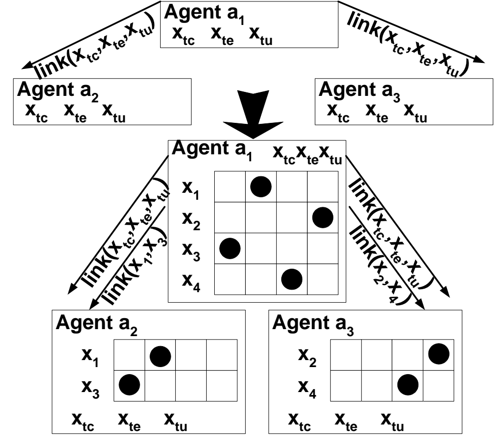

Figure 1. Motivating example

3 Generative Constraint Satisfaction

In many applications, solving is a *generative* process, where the number of involved components (i.e., variables) is not known from the beginning. To represent these problems we employ an extended formalism that complies to the specifics of configuration and other synthesis tasks where problem variables representing components of the final system are *generated* dynamically as part of the solution process because their total number cannot be determined beforehand. The framework is called *generative* CSP (GCSP) [5, 16]. This kind of dynamicity extends the approach of dynamic CSP (DCSP) formalized by Mittal and Falkenhainer [9], where all possibly involved variables are known from the beginning. This is needed because the activation constraints reason on the variable's activity state. [10] propose a conditional CSP to model a configuration task, where structural dependencies in the configuration model are exploited to trigger the activation of subproblems. Another class of DCSP was first introduced by [3] where constraints can be added or removed independently of the initial problem statement. The dynamicity occuring in a GCSP differentiates from the one described in [3] in the sense that a GCSP is extended in order to find a consistent solution and the latter has already a solution and is extended due to influence from the outside world (e.g., additional constraints) that necessitates finding a new solution. Here we give a definition of a GCSP that abstracts from the configuration task specific formulation in [16] and applies to the wider range of synthesis problems.

Definition 1 (Generative constraint satisfaction problem (GCSP)) *A generative constraint satisfaction problem is a tuple GCSP(X, Γ, T, Δ), where:*

- X *is the set of problem variables of the GCSP and* $X_0 \subseteq X$ *is the set of initially given variables.*
- Γ *is the set of generic constraints.*
- $T = \{t_1, \ldots, t_n\}$ *is the set of variable types* t_i*, where* $dom(t_i)$ *associates the same domain to each variable of type* t_i*, where the domain is a set of atomic values.*

- *For every type $t_i \in T$ exists a counter variable $x_{t_i} \in X_0$ that holds the number of variable instantiations for type t_i. Thus, explicit constraints involving the total number of variables of specific types and reasoning on the size of the CSP becomes possible.*
- *Δ is a total relation on $X \times (T, N)$, where N is the set of positive integer numbers. Each tuple $(x, (t, i))$ associates a variable $x \in X$ with a unique type $t \in T$ and an index i, that indicates x is the i^{th} variable of type t. The function $type(x)$ accesses Δ and returns the type $t \in T$ for x and the function $index(x)$ returns the index of x.*

By generating additional variables, a previously unsolvable CSP can become solvable, which is explained by the existence of variables that hold the number of variables.
When modeling a configuration problem, variables representing named connection points between components, i.e., *ports*, will have references to other ports as their domain. Consequently, we need variables whose domain varies depending on the size of a set of specific variables [16].

Example Given t_{mod} as the type of variables representing ports of *modules* and t_{port} as the type of *port* variables that are allowed to connect to *modules*, then the domain of the *port* variables $dom(t_{port})$ must contain references to *modules*. This is specified by defining $dom(t_{port}) = \{1, \ldots, ub\}$, where ub is an upperbound on the number of variables of type t_{mod}, and formulating an additional generic constraint that restricts all variables of type t_{port} using the counter variable for the total number of variables having type t_{mod}, i.e., $val(V^{t_{port}}) \leq x_{t_{mod}}$. With the help of the $index()$ function concrete variables can then be referenced.
Referring to our introductory example we can formalize the local GCSP of agent a_1 (initially consisting only of counter variables x_{t_i}, their type t_c, and the types of row variables) as $X^{a_1} = \{x_{t_c}, x_{t_e}, x_{t_u}\}$, $\Gamma^{a_1} = \{\gamma_1, \gamma_2, \gamma_3, \gamma_8, \gamma_9\}$, $T^{a_1} = \{t_c, t_e, t_u\}$ and $\Delta^{a_1} = \{(x_{t_c},(t_c, 1)),(x_{t_e},(t_c, 2)), (x_{t_u},(t_c, 3))\}$. The domain for even and uneven row variables is consequently defined as $dom(t_e) = dom(t_u) = dom(t_c) = \{1, \ldots, ub\}$, where the domains for the row variables are limited by the domain constraints (i.e., γ_8, γ_9).

Definition 2 (Generic constraint) *A generic constraint $\gamma \in \Gamma$ formulates a restriction on the meta-variables $M_a, \ldots, M_k$. A meta-variable M_i is associated a variable type $type(M_i) \in T$ and must be interpreted as a placeholder for all concrete variables x_j, where $type(x_j) = type(M_i)$.*

Note, that generic constraints can also formulate restrictions on specific initial variables from X_0 by employing the $index()$ function.
Consider the GCSP(X, Γ, T, Δ) and let $\gamma \in \Gamma$ restrict the meta-variables $M_a, \ldots, M_k$, where $type(M_i) \in T$ is the defined variable type of the meta variable M_i.

Definition 3 (Consistency of generic constraints) *Given an assignment tuple θ for the variables X, then γ is said to be satisfied under θ, iff*
$\forall x_a, \ldots, x_k \in X : type(x_a) = type(M_a) \wedge \ldots \wedge type(x_k) = type(M_k) \rightarrow \gamma[M_a|_{x_a}, \ldots, M_k|_{x_k}]$ is satisfied unter θ, where $M_i|_{x_i}$ indicates that the meta-variable M_i is substituted by the concrete variable x_i.

Thus a *generic* constraint must be seen as a constraint scheme that is expanded into a set of constraints after a preprocessing step, where meta-variables are replaced by all possible combinations of concrete variables having the same type, e.g., given a GCSP of agent a_1 (excluding counter variables) with $X^{a_1} = \{x_1, x_2, x_3\}$, $T^{a_1} = \{t_u, t_e\}$ and $\Delta^{a_1} = \{(x_1, (t_u, 1)), (x_2, (t_e, 1)), (x_3, (t_u, 2))\}$, the satisfiability of the generic constraint γ_2 is checked by testing the following conditions: $val(x_1) \neq val(x_2)$. $val(x_3) \neq val(x_2)$.

Definition 4 (Solution for a generative CSP) *Given a generative constraint satisfaction problem GCSP(X_0, Γ, T, Δ_0), then its solution encompasses the finding of a set of variables X, type and index assignments Δ and an assignment tuple θ for the variables in X, s.t.*

1. *for every variable $x \in X$ an assignment $x = v$ is contained in θ, s.t. $v \in dom(type(x))$ and*
2. *every constraint $\gamma \in \Gamma$ is satisfied under θ and*
3. *$X_0 \subseteq X \wedge \Delta_0 \subseteq \Delta$.*

Note, that we do not impose a minimality criterium on the number of variables in our solution, because in practical applications different optimization criteria exist, such as total cost or flexibility of the solution, thus non-minimal solutions can be preferred over minimal ones.
The calculated solution (excluding counter variables) for the local GCSP of agent a_1 consists of $X^{a_1} = \{x_1, x_2, x_3, x_4\}$, $\Delta^{a_1} = \{(x_1, (t_u, 1)),(x_2, (t_e, 1)),(x_3, (t_u, 2)),(x_4, (t_e, 2))\}$ and the assignment tuple $x_1 = 2$, $x_2 = 4$, $x_3 = 1$ and $x_4 = 3$. Thus, $x_1, \ldots, x_4$ are the names of *generated* variables.
Note, that names for generated variables are unique and can be randomly chosen by the GCSP solver implementation and therefore constraints must not formulate restrictions on the variable names of generated variables. Consequently, substitution of any generated variable (i.e., $x \in X \setminus X_0$) by a newly generated variable with equal type, index and value assignment has no effect on the consistency of generic constraints. Our GCSP definition extends the definition from [16] in the sense that a finite set of variable types T is given and during problem solving variables having any of these types can be generated, whereas in [16] only variables of a single type, i.e., component variables, can be created. Current CSP implementations of configuration systems (e.g., [7] [4]) use a type system for problem variables, where new variable instances, having one of the predefined types, are dynamically created. This is only indirectly reflected in the definition of [16] by the domain definition of component variables, which we explicity represent as a set of types. Furthermore, the definition of *generic constraints* does not enforce the use of a specific constraint language for the formulation of restrictions. Examples are the LCON language used in the COCOS project [16], or the configuration language of the ILOG Configurator [7].
Note, that the set of variables X can be theoretically infinite, leading to an infinite solution space. For practical reasons, solver implementations for a GCSP put a limit on the total number of problem variables to ensure decidability and finiteness of the search space. This way a GCSP is reduced to a dynamic CSP and in further consequence to a CSP. A DCSP models each search state as a static CSP, where complex activation constraints are required to ensure the alternate activation of variables depending on the search state. These constraints need to be formulated for every possible state of the GCSP, which leads to combinatorial explosion of concrete constraints. Furthermore, the formulation of large configuration problems as a DCSP is merely impractical from the perspective of knowledge representation, which is crucial for knowledge-based applications such as configuration systems.

4 DisCSP Framework

In our framework, we are interested only in algorithms that guarantee a good/optimal solution. The first asynchronous complete search algorithm is Asynchronous Backtracking (ABT) [18]. [2] shows how

ABT can be adapted to networks where not all agents can directly communicate to one another. [6] makes the observation that versions of ABT with polynomial space complexity can be designed. The extension of ABT with asynchronous maintenance of consistencies, and asynchronous dynamic reordering is described in [12, 14]. [11] achieves an increased level of abstraction in DisCSPs by letting nogoods (i.e. certain constraints) consist of aggregates (i.e. sets of variable assignments), instead of simple assignments.

We show how the basic DisCSP framework for ABT with extensions for private constraints [11] can be applied to a scenario of distributed product configuration. Therefore, improving the performance of ABT with extensions as referenced above is straightforward. We summarize in the following the properties of the ABT algorithm that guarantee its correctness and completeness [18]. Then we apply this DisCSP framework to a scenario where each agent locally solves a generative constraint satisfaction task. Each time an agent extends the solution space of his local GCSP by creating an additional variable, the DisCSP setting is transformed into a new DisCSP setting, which again has all properties required by asynchronous search to correctly function.

4.1 Asynchronous Search

We summarize the characteristics of asynchronous search algorithms like ABT [18] and its extensions [11] for private constraints:

1. $A = \{a_1, \ldots a_n\}$ is a set of n totally ordered agents, where a_i has priority over a_j if $i < j$.
2. Each agent a owns a set of local constraints Γ^a and a is *interested in* those variables that are contained in its local constraints, called *local variables*. A *link* exists between two agents if they share a variable, that is directed from the agent with higher priority to the agent with lower priority. A link from agent a_1 to agent a_2 is refered to as an *outgoing link* of a_1 and an *incoming link* of a_2.
3. An aggregate is a triplet (x_j, set_j, h_j), where x_j is a variable, set_j a set of values for x_j and h_j is a *history* of the pair (x_j, set_j), where the history marks the aggregate with the information required for a correct message ordering (a counter in ABT).
4. The *view* of an agent a is a set of the aggregates for those variables agent a is interested in.
5. The agents communicate using the following types of messages, where channels without communication loss are assumed:
   - **ok?** message. Agents with higher priorities communicate via **ok?** messages a proposal for a set of variables to lower priority agents. Each proposal is associated with a history, that allows the recipient to identify the most recent message.
   - **nogood** message. In case an agent cannot find a proposal that does not violate its own constraints and its stored nogoods, it generates an explanation under the form of an explicit nogood $\neg N$. A nogood can be interpreted as a constraint that forbids a combination of value assignments to a set of variables. It is announced via a **nogood** message to the lowest priority agent that has proposed an assignment[3] in N.
   - **addlink** message. It transports a set of variables $vars$, where the receiver agent is informed that the sender is interested in the variables $vars$ and for every variable in $vars$ a *link* is established from the higher priority agent to the agent with lower priority.
6. A *system agent* is a special agent that receives the subscriptions of the agents for the search. Its task is to decide the order of the agents, initialize the links and announce the termination of the search.

[3] aggregate in the Asynchronous Aggregation Search (AAS) algorithm [11]

4.2 Framework for DisGCSP

A distributed configuration problem is a multi-agent scenario, where each agent wants to satisfy a local GCSP and agents keep their constraints private for security and privacy reasons, but share all variables which they are interested in. As constraints employ meta-variables, the *interest* of an agent in variables needs to be redefined:

Definition 5 (Interest in variables) *An agent a_j owning a local $GCSP^{a_j}$(X^{a_j},Γ^{a_j},T^{a_j},Δ^{a_j}) is said to be interested in a variable $x \in X^{a_h}$ of an agent a_h, if there exists a generic constraint $\gamma \in \Gamma^{a_j}$ formulating a restriction on the meta-variables $M_a, \ldots, M_k$, where $type(M_i) \in T^{a_j}$ is the defined variable type of the meta variable M_i, and $\exists M_i \in M_a, \ldots, M_k : type(x) = type(M_i)$.*

Definition 6 (Distributed generative CSP) *A distributed generative constraint satisfaction problem has the following characteristics:*

- *$A = \{a_1, \ldots, a_n\}$ is a set of n agents, whereby each agent a_i owns a local $GCSP^{a_i}$(X^{a_i},Γ^{a_i},T^{a_i},Δ^{a_i}).*
- *All variables in $\bigcup_{i=1}^{n} X^{a_i}$ and all type denominators in $\bigcup_{i=1}^{n} T^{a_i}$ share a common namespace, ensuring that a symbol denotes the same variable, resp. the same type, with every agent.*
- *For every pair of agents $a_i, a_j \in A$ and for every variable $x \in X^{a_j}$, where agent a_i is interested in x, must hold $x \in X^{a_i}$.*
- *For every pair of agents $a_i, a_j \in A$ and for every shared variable $x \in X^{a_i} \cap X^{a_j}$ the same type and index must be associated to x in the local GCSPs of the agents, i.e., $type^{a_i}(x) = type^{a_j}(x) \wedge index^{a_i}(x) = index^{a_j}(x)$.*

For every pair of agents $a_i, a_j \in A$ and for every shared variable $x \in X^{a_i} \cap X^{a_j}$ a *link* must exist that indicates that they share variable x. The *link* must be directed from the agent with higher priority to the agent with lower priority.

Definition 7 *Given a distributed generative constraint satisfaction problem among a set of n agents then its solution encompasses the finding of a set of variables $X = \bigcup_{i=1}^{n} X^{a_i}$, type and index assignments $\Delta = \bigcup_{i=1}^{n} \Delta^{a_i}$ and an assignment tuple $\theta = \bigcup_{i=1}^{n} \theta^{a_i}$ for every variable in X, s.t. for all agents $a_i : X^{a_i}$, Δ^{a_i} and θ^{a_i} are a solution for the local $GCSP^{a_i}$ of agent a_i.*

Remark A solution to a distributed generative CSP is also a solution to a centralized GCSP($\bigcup_{i=1}^{n} X^{a_i}$, $\bigcup_{i=1}^{n} \Gamma^{a_i}$, $\bigcup_{i=1}^{n} T^{a_i}$, $\bigcup_{i=1}^{n} \Delta^{a_i}$).

Definition 8 (Generic aggregate) *A generic aggregate is a unary generic constraint. It takes the form: $\langle M, i, s, h \rangle$, where M is a meta-variable, i is a set of index values for which the constraint applies, s is a set of values and h is a history of the aggregate.*

Definition 9 (Generic nogood) *A generic nogood takes the form $\neg N$, where N is a set of generic aggregates for distinct meta-variables.*

Given the characteristics of a DisGCSP (see Definition 6) the links can be initialized before the start of the algorithm, due to the common naming space for type denominators and the condition of a unique type and index assignment to variables over all agents.

Value assignments to variables are communicated to agents via **ok?** messages that transport *generic aggregates* in our DisGCSP framework, which represent domain restrictions on variables by unary constraints. Each of these unary constraints in our DisGCSP has attached an unique identifier called constraint reference (cr) [13]. Any inference has to attach the crs associated to arguments into the obtained nogood. We treat the extension of the domains of the variables as a constraint relaxation [13]. For this reason we introduce the next features for algorithm extensions:

- **announce** message broadcasts a tuple (x, t, i), where x is a newly created variable of type t and with index i to all other agents. The receiving agents determine their interest in variable x and react depending on their interest and priority in one of the following ways (a) send an **addlink** message transporting the variable set $\{x\}$ (b) add the sending agent to its outgoing links or (c) discard the message.
- **domain** message broadcasts a set CR of obsolete constraint references. Any receiving agent removes all the nogoods having attached to them a constraint reference $cr \in CR$. The receiver of the message calls then the function *check_agent_view()* detailed in [18], making sure that it has a consistent proposal or that it generates nogoods.
- **nogood** messages transport *generic nogoods* $\neg N$ that contain assignments for meta-variable instances. These messages are multicasted to all agents interested in $\neg N$. An agent A_i is interested in a generic nogood $\neg N$ if it has *interest in* any meta-variable in $\neg N$.
- When an agent needs to revoke the creation of a new variable due to backtracking in his local solving algorithm, he assigns it a specific value from its domain indicating the deactivation of the variable and communicates it via an **ok?** message to all interested agents.

In order to avoid too many messages a broker agent can be introduced that maintains a static list of agents and their interest in variables of specific types comparable to a *yellow pages* service. In this case the agent that created a new variables only needs to request the broker agent for a list of interested agents and does not need to broadcast an **announce** message to all agents.

Theorem 1 *Whenever an existing extension of ABT is extended with the previous messages and is applied to DisGCSPs, the obtained protocols are correct, complete and terminate.*

Proof: Let us consider that we extend a protocol called P.
Completeness: All the generated information results by inference. If failure is inferred (when no new component is available), then indeed no solution exists.
Termination: Without introducing new variables, the algorithm terminates. Since the number of variables that can be generated is finite, termination is ensured.
Correctness: The resulting overall protocol is an instance of P, where the delays of the system agent initializing the search equals the time needed to insert all the variables generated before termination. Therefore the result satisfies all the agents and the solution is correct □

5 Conclusions

Building on the definition of a centralized configuration task from [16], we formally defined a new class of CSP, termed generative CSP (GCSP), that generalizes the approaches of constraint-based configurator applications in use [4, 7]. The innovative aspects include an additional level of abstraction for constraints and nogoods. Constraints and nogoods can refer to types of variables. Furthermore, we extended GCSP to a distributed scenario, where this abstraction adapts well DisCSP frameworks for dynamic configuration problems (but it can be used in static models as well). We have described how this enhancement can be naturally integrated in a large family of existing asynchronous algorithms for DisCSPs.

REFERENCES

[1] V.E. Barker, D.E. O'Connor, J.D. Bachant, and E. Soloway, 'Expert systems for configuration at Digital: XCON and beyond', *Communications of the ACM*, **32(3)**, 298–318, (1989).

[2] C. Bessière, A. Maestre, and P. Meseguer, 'Distributed dynamic backtracking', in *Proc. of 7th Int. Conf. on Principles and Practice of Constraint Programming (CP)*, p. 772, Paphos, Cyprus, (2001).

[3] R. Dechter and A. Dechter, 'Belief Maintenance in Dynamic Constraint Networks', in *Proc. 7th National Conf. on Artificial Intelligence (AAAI)*, pp. 37–42, St. Paul, MN, (1988).

[4] G. Fleischanderl, G. Friedrich, A. Haselböck, H. Schreiner, and M. Stumptner, 'Configuring Large Systems Using Generative Constraint Satisfaction', in *IEEE Intelligent Systems, Special Issue on Configuration*, ed., E. Freuder B. Faltings, volume 13(4), 59–68, (1998).

[5] A. Haselböck, *Knowledge-based configuration and advanced constraint technologies*, Ph.D. dissertation, Technische Universität Wien, 1993.

[6] W. Havens, 'Nogood caching for multiagent backtrack search', in *Proc. of 14th National Conf. on Artificial Intelligence (AAAI), Agents Workshop*, Providence, Rhode Island, (1997).

[7] D. Mailharro, 'A classification and constraint-based framework for configuration', *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, **12(4)**, 383–397, (1998).

[8] D.L. McGuiness and J.R. Wright, 'Conceptual Modeling for Configuration: A Description Logic-based Approach.', *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, **12(4)**, 333–344, (1998).

[9] S. Mittal and B. Falkenhainer, 'Dynamic Constraint Satisfaction Problems', in *Proc. of 8th National Conf. on Artificial Intelligence (AAAI)*, pp. 25–32, Boston, MA, (1990).

[10] D. Sabin and E.C. Freuder, 'Configuration as Composite Constraint Satisfaction', in *Proc. of AAAI Fall Symposium on Configuration*, Cambridge, MA, (1996). AAAI Press.

[11] M.-C. Silaghi, D. Sam-Haroud, and B. Faltings, 'Asynchronous search with aggregations', in *Proc. of 17th National Conf. on Artificial Intelligence (AAAI)*, pp. 917–922, Austin, TX, (2000).

[12] M.-C. Silaghi, D. Sam-Haroud, and B. Faltings, 'ABT with asynchronous reordering', in *Proc. of Intelligent Agent Technology (IAT)*, pp. 54–63, Maebashi, Japan, (October 2001).

[13] M.-C. Silaghi, D. Sam-Haroud, and B.V. Faltings, 'Maintaining hierarchically distributed consistency', in *Proc. of 7th Int. Conf. on Principles and Practice of Constraint Programming (CP), DCS Workshop*, pp. 15–24, Singapore, (2000).

[14] M.-C. Silaghi, D. Sam-Haroud, and B.V. Faltings, 'Consistency maintenance for ABT', in *Proc. of 7th Int. Conf. on Principles and Practice of Constraint Programming (CP)*, pp. 271–285, Paphos, Cyprus, (2001).

[15] M. Stumptner, 'An overview of knowledge-based configuration', *AI Communications*, **10(2)**, (June, 1997).

[16] M. Stumptner, G. Friedrich, and A. Haselböck, 'Generative constraint-based configuration.', *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, **12(4)**, 307–320, (1998).

[17] E. Tsang, *Foundations of Constraint Satisfaction*, Academic Press, 1993.

[18] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara, 'Distributed constraint satisfaction for formalizing distributed problem solving', in *Proc. of 12th Int. Conf. on Distributed Computing Systems (ICDCS)*, pp. 614–621, Yokohama, Japan, (1992).