# Hiding Absence of Solution for a Distributed Constraint Satisfaction Problem 

Marius-Călin Silaghi<br>Florida Institute of Technology


#### Abstract

A distributed constraint satisfaction problem (DisCSP) is defined by a set of agents, trying to find assignments for a set of variables with known domains and subject to secret constraints. They can model applications like auctions, distributed team-making, scheduling, and configuration. Secure multiparty computations (MPC) can pick and reveal randomly one of the solutions of any distributed constraint satisfaction and optimization problem. Our previous algorithms are based on an arithmetic circuit for selecting a random element out of the elements with a given value in a secret array (Silaghi 2004a). Here we show an improvement based on a very simple and elegant (optimized) version of the involved functions and on the usage of CSP solvers to exploit public constraints. The technique can hide the absence of solutions.


The algorithm proposed here is called MPC-DisCSP4. Compared to previous techniques, MPC-DisCSP4 is faster and can hide the absence of solutions.

Definition $1 A$ Distributed CSP is defined by five sets $(A, X, D, C, O) . \quad A=\left\{A_{1}, \ldots, A_{n}\right\}$ is a set of agents. $X=$ $\left\{x_{1}, \ldots, x_{m}\right\}$ is a set of variables and $D=\left\{D_{1}, \ldots, D_{m}\right\}$ is a set of domains such that $x_{i}$ can take values only from $D_{i}=\left\{v_{1}^{i}, \ldots, v_{d_{i}}^{i}\right\} . C=\phi_{0} \cup\left\{\phi_{1}, \ldots, \phi_{c}\right\}$ is a set of constraints. $\phi_{i}$ involves an ordered subset $X_{i}=\left\{x_{i_{1}}, \ldots, x_{i_{k_{i}}}\right\}$ of the variables in $X, X_{i} \subseteq X$, and constrains the legality of each combination of assignments to the variables in $X_{i}$. An assignment is a pair $\left\langle x_{i}, v_{k}^{i}\right\rangle$ meaning that variable $x_{i}$ is assigned the value $v_{k}^{i}$. Each constraint $\phi_{i}, i>0$, is known only by one agent, being the secret of that agent. There may exist $c_{0}$ public constraints in $C, \phi_{0}=\left\{\phi_{0}^{1}, \ldots, \phi_{0}^{c_{0}}\right\}$.

A tuple is an ordered set. The projection of a tuple $\epsilon$ of assignments over a tuple of variables $X_{i}$ is denoted $\epsilon_{\left.\right|_{X_{i}}} . A$ solution of a DisCSP $(X, D, C)$ is a tuple of assignments $\epsilon$ with one assignment for each variable in $X$ such that each $\phi_{i} \in C$ is satisfied by $\epsilon_{\left.\right|_{X_{i}}}$ (written $\phi_{i}\left(\epsilon_{\left.\right|_{X_{i}}}\right)$ ).

MPC-DisCSP4. Figure 1 gives an overview of MPCDisCSP4, also shown in Algorithm 1. MPC-DisCSP4 is composed of a CSP solver, a mix-net, and a set of functions

[^0]

Figure 1: Pictorial view for MPC-DisCSP4. The reconstructions of secrets at the end is optional, as it will not be performed when the solver is integrated in other techniques (optimization, auctions). Dark gray show shared secrets.
(arithmetic circuits) that are much more elegant and simple than those used in previous versions.

MPC-DisCSP4 starts by sharing the encoded constraints with the Shamir secret sharing scheme. A complete CSP solver (e.g., backtracking) is used to generate the vector: $S^{\prime \prime}=\left\{\epsilon \mid \phi_{0}(\epsilon)\right\}$. I.e. $S^{\prime \prime}$ contains all tuples $\epsilon$ that satisfy $\phi_{0}$. Each tuple of assignments candidating as solution to the input problem is refered by an index (e.g. see the first array in Figure 1, with elements from 1 to 10). The agents share the secret values of their constraints. Let $\epsilon_{k}$ denote the $k^{t h}$ tuple in the lexicographic order. They compute for each such tuple $\epsilon_{k}$, a shared secret whose unknown value is 1 when the tuple satisfies all agents and 0 otherwise. This is done by using (Ben-Or, Goldwasser, \& Widgerson 1988) to securely simulate the evaluation of the arithmetic circuit: $p\left(\epsilon_{k}\right)=\prod_{i=1}^{c} \phi_{i}\left(\epsilon_{k}\right)$. The results form the second array of the figure, $S^{\prime}, S^{\prime}[i]=p\left(S^{\prime \prime}[i]\right)$ (each participant obtains a distinct vector with a share for each element of $S^{\prime}$ ). $S^{\prime}$ is now shuffled and the shares are randomized with a mix-net whose details are given later. We define:

$$
\begin{aligned}
h_{1}(P) & =1 \\
h_{i}(P) & =h_{i-1}(P) *\left(1-S^{\prime}[i-1]\right)
\end{aligned}
$$

The lexicographically first solution can be isolated:

$$
\begin{equation*}
S[i]=S^{\prime}[i] * h_{i}(P) \tag{1}
\end{equation*}
$$

A vector, $S$, is computed with Equation 1. It is then decoded by traversing the mix-net in the inverse direction and

## procedure MPC-DisCSP4 do

1. Shares the $\{0,1\}$ encoded secret constraints.
2. Compute all tuples $\epsilon$ returned by the CSP solver for $\phi_{0}$. Place the results lexicographically in vector $S^{\prime \prime}$.
3. Compute a vector $S^{\prime}, S^{\prime}[k]=p\left(S^{\prime \prime}[k]\right)$.
4. The mix-net shuffles the vectors of shares, $S^{\prime}$, randomizing the shares at each permutation (by adding shares of 0) exploiting homomorphic encryption.
5. Compute vector $S$ with Equation (1).
6. The mix-net decodes $S$ randomizing the shares and inversing permutations of Step 5.
7. Compute Equations 2.
8. Reveal assignments to the interested agents.

## Algorithm 1: MPC-DisCSP4

with the inverse permutations, randomizing the shares as at shuffling. The solution is the tuple of $S^{\prime \prime}$ situated at the same index as the only element of $S$ that is different from 0 .

The solution can also be transformed into a set of indexes of values of the variables by the following method. Assume the value of the $u^{t h}$ variable in the $t^{t h}$ tuple of the search space is denoted $\eta_{u}(t)$. The values in the solution are computed with the arithmetic circuits in Equation (2).

$$
\begin{equation*}
f_{i}(P)=\sum_{t=1}^{\left|S^{\prime \prime}\right|}\left(\eta_{i}(t)+1\right) * S^{\prime}[t-1] \tag{2}
\end{equation*}
$$

Each variable $x_{i}$ is assigned in the solution to the value in $D_{i}$ at index given by the functions $f_{i}$, and can be revealed.

MPC-DisCSP4's mix-net for reordering vectors of shared secrets. Each agent $A_{i}$ chooses a random secret permutation $\pi_{i}$, picked with a uniform distribution over the set of possible permutations: $\pi_{i}:\left[1 . .\left|S^{\prime \prime}\right|\right] \rightarrow\left[1 . .\left|S^{\prime \prime}\right|\right]$. Each agent chooses a pair of keys for a $(+, \times)$-homomorphic public encryption scheme and publishes the public key. The secret shares, of the values computed in the vector $S^{\prime}$, are encrypted by each $A_{i}$ with her public key and are serialized. The serialized encrypted vectors are sent to $A_{1}$. $A_{1}$ shuffles the serialized vectors according to her permutation $\pi_{1}$, then passes them to $A_{2}$ which applies $\pi_{2}$, etc., until the agent $A_{n}$ which applies $\pi_{n}$. $A_{n}$ sends each vector to its owner.

To avoid that agents get a chance to learn the final permutation by matching final shares with the ones that they encrypted, a randomization step is also applied at each shuffling. Each agent applies a randomization step on the set of shares for each element of $S^{\prime}$, by adding corresponding shares of zero. Since operands are encrypted, to perform this summation we exploit the $(+, \times)$-homomorphic property. For each secret in $S^{\prime}$, a new set of 0 's Shamir shares are generated, and $\forall i, i \leq n$, the $0^{\prime} s i^{t h}$ share is encrypted with the public key of $A_{i}$, then it is multiplied to the corresponding $A_{i}$ 's encrypted share of the secret (resulting in resharing the secret). This assumes $\mu>\nu(n+1)$, for the decryption to be correct in $\mathbb{Z}_{\nu}$. $\left(\mathbb{Z}_{\mu},+\right)$ is the group with possible plaintexts for the homomorphic encryption. $\left(\mathbb{Z}_{\nu},+, \times\right)$ is the field in which that arithmetic circuits are evaluated.
Example 1 Let us see an example of how MPC-DisCSP4 is applied to the Example 2. $p(P, W)$ is not computed $\left(\phi_{0}\right)$.
$p(P, T)=1, p(Q, T)=0, p(Q, W)=1$.
$S^{\prime \prime}=((P, T),(Q, T),(Q, W)) S^{\prime}=(1,0,1)$
Shuffle $(1,0,1)$, (assume it remains unchanged)
$h_{1}(P)=1, h_{2}(P)=0, h_{3}(P)=0$.
With Equation 1 we get vector $S=\{1,0,0\}$.
Unshuffle $S=(1,0,0): S=(1,0,0)$
Using Equation 2: $\eta_{1}(1)=0, \eta_{1}(2)=1, \eta_{1}(3)=1 . \quad \eta_{2}(1)=0$, $\eta_{2}(2)=0, \eta_{2}(3)=1 . f_{1}(P)=1, f_{2}(P)=1$.
Therefore the chosen solution is $x_{1}=$ Paris, $x_{2}=$ Tuesday.
Besides the mixnet, the most expensive operation - requiring exchanges of messages - is the muliplication of two shared secrets. MPC-DisCSP4 has $2\left(\left|S^{\prime \prime}\right|-1\right)$ such multiplications. $\left|S^{\prime \prime}\right|-1$ multiplications to compute the $h$ series, and $\left|S^{\prime \prime}\right|-1$ multiplications to obtain $S$. The number of such multiplications in MPC-DisCSP3 is $5\left(\left|S^{\prime \prime}\right|-1\right)$. This corresponds to a 2.5 times improvement. Preparing the vectors $S^{\prime}$ has a cost proportional to $c\left|S^{\prime \prime}\right|$.

In contrast to previous techniques, the most expensive computation in MPC-DisCSP4 is due to the mix-nets. Their efficiency was improved by the use of the CSP solver to prune tuples that are inconsistent with public constraints.

Hiding the (in)existance of a solution. To hide the fact that a solution does not exist for a problem, the participants may prefer to miss an existing solution with some probability $p$. Then, the fact that no solution is found does not prove that no solution exist, and certain leaks of secrets induced by the lack of a solution are avoided. This can be achieved by generating unknown shared number(s) $K$ that with probability $p$ equal 0 , and otherwise equal 1 (Silaghi 2004b).

Now, each of the secret elements of the vector $S$ (or all the secret values returned by $f_{i}$ ) are multiplied with such a number $K$, losing the solution with probability $p$.

Conclusions. MPC-DisCSP4 can offer a privacy that is stronger than what is offered by the previous methods for DisCSPs. The solution is picked randomly over the set of existing solutions, hiding statistical information about whether other tuples are solutions or not. If needed, one may choose to hide the information about whether a solution exists at all, by probabilistically discarding the solution.

The computation cost on shuffled problems is 2.5 times faster with the new technique than with MPC-DisCSP3, which has the closest performance in privacy. As with MPCDisCSP3, by taking advantage of public constraints, the efficiency of the most expensive operations (public key encryptions), is reduced proportionally with their tightness.

## References

Ben-Or, M.; Goldwasser, S.; and Widgerson, A. 1988. Completeness theorems for non-cryptographic faulttolerant distributed computating. In STOC, 1-10.
Silaghi, M.-C. 2004a. Meeting scheduling system guaranteeing $\mathrm{n} / 2$-privacy and resistant to statistical analysis (applicable to any DisCSP). In IC on Web Intel., 711-715.
Silaghi, M. C. 2004b. Secure multi-party computation for selecting a solution according to a uniform distribution over all solutions of a general combinatorial problem. Cryptology e-print Archive 2004/333.


[^0]:    Copyright © 2005, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

