TU Delft

Delft University of Technology

Document Version
Final published version

Licence
Dutch Copyright Act (Article 25fa)

Citation (APA)
Mălan, A., Guzella, T., Decouchant, J., & Chen, L. (2026). CCBNet: Confidential Collaborative Bayesian Networks Inference. In C. Garman, & P. Moreno-Sanchez (Eds.), *Financial Cryptography and Data Security - 29th International Conference, FC 2025, Revised Selected Papers* (pp. 383-400). (Lecture Notes in Computer Science; Vol. 15752 LNCS). Springer. https://doi.org/10.1007/978-3-032-07035-7_23

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

CCBNet: Confidential Collaborative Bayesian Networks Inference

Abele Mălan[1(✉)], Thiago Guzella[2], Jérémie Decouchant[3], and Lydia Chen[1,3]

[1] University of Neuchâtel, Neuchâtel, Switzerland
{abele.malan,yiyu.chen}@unine.ch
[2] ASML, Eindhoven, Netherlands
thiago.guzella@asml.com
[3] Delft University of Technology, Delft, Netherlands
j.decouchant@tudelft.nl

Abstract. Effective large-scale process optimization in manufacturing industries requires close cooperation between different human expert parties who encode their knowledge of related domains as Bayesian network models. For instance, Bayesian networks for domains such as lithography equipment, processes, and auxiliary tools must be conjointly used to effectively identify process optimizations in the semiconductor industry. However, business confidentiality across domains hinders such collaboration, and encourages alternatives to centralized inference. We propose CCBNet, the first **C**onfidentiality-preserving **C**ollaborative **B**ayesian **Net**works inference framework. CCBNet leverages secret sharing to securely perform analysis on the combined knowledge of party models by joining two novel subprotocols: (i) CABN, which augments probability distributions for variables across parties by modeling them into secret shares of their normalized combination; and (ii) SAVE, which aggregates party inference result shares through distributed variable elimination. We extensively evaluate CCBNet via 9 public Bayesian networks. Our results show that CCBNet achieves predictive quality that is similar to the ones of centralized methods while preserving model confidentiality. We further demonstrate that CCBNet scales to challenging manufacturing use cases that involve 16–128 parties in large networks of 223–1003 variables, and decreases, on average, computational overhead by 23%, while communicating 71k values per request. Finally, we showcase possible attacks and mitigations for partially reconstructing party networks in the protocol.

1 Introduction

Improving productivity and quality standards in manufacturing demands effectively expressing complex interactions between domain items. Bayesian networks

Work partly done while graduating from TU Delft and interning at ASML.

© International Financial Cryptography Association 2026
C. Garman and P. Moreno-Sanchez (Eds.): FC 2025, LNCS 15752, pp. 383–400, 2026.
https://doi.org/10.1007/978-3-032-07035-7_23

(BNs) are commonly adopted to graphically model causality in manufacturing [20], with nodes representing variables and directed edges showing dependencies. An essential trait of these models is their ability to specify arbitrary input and output variables for each query instead of having them fixed.

In the semiconductor industry, the pursuit of smaller chips at a high yield [28] entails cooperation among specialized parties protecting their trade secrets. More specifically, fab operators encode their insight into settings dictating the production process in a BN. Similarly, vendors of equipment like scanners, construct BNs that describe the inner workings of their machines. Pooling together parties' knowledge allows to better optimize production environments, leading to new business opportunities. Simultaneously, the need to protect intellectual property expressed within party models calls for confidential collaboration.

Existing studies on collaborative inference for BNs **ignore model confidentiality** or limit the ability to merge party information. We consider confidentiality preserved if parties cannot infer from others the identity and structural connections of new variables, or parameters of other probability tables, even for variables they also define themselves. Most previous works detail a centralized combination of local BNs into a larger one [10,11] without protecting confidential knowledge within the input networks and global output. Models get stitched together based on common nodes, and remain in the final representation as largely unaltered submodels whose encoded knowledge is easily inspectable. Pavlin et al. [22] propose a distributed combination variant that partially preserves confidentiality by maintaining the locality of combined party models. However, it leaks information when combining them and does not allow merging inner graph nodes having both parents and children. Tedesco et al. [27] maintain the confidentiality of how nodes are linked within parties but only allow propagating information between them in a fixed sequential order. The approach of Kim and Ghahramani [17] has similar confidentiality properties but even greater compatibility restrictions by requiring models to have identical inputs and outputs.

In this paper, we propose `CCBNet`, the first confidential, collaborative BNs inference framework that combines knowledge of multiple parties involved in inference queries through a novel secret sharing scheme. `CCBNet` does not require a trusted third party, and protects confidentiality at both the levels of party models and data instances. The two key components of `CCBNet` are: (i) confidential sharing of aggregated variable probability distributions across all overlapping parties; and (ii) distributed inference based on variable elimination for aggregating party results. The novelty of the augmentation procedure lies in constructing discrete conditional probability distributions for all variables present in more than one party. These variables represent secret shares of a combined and normalized distribution from a centralized scenario without exposing any party's initial probability function. To evaluate `CCBNet`, we simulate knowledge compartmentalization over different public BNs. Moreover, we demonstrate possible attacks against the framework, their limitations, and possible mitigations.

In summary, we make the following contributions:

- We devise a confidential, collaborative framework for BNs, `CCBNet`, to satisfy the needs of industry use cases like process optimization in manufacturing.
- We design a novel secret sharing-based protocol, `CABN`, to confidentially augment the conditional probability function of overlapping variables in parties.
- We define `SAVE`, a method backed by variable elimination for performing distributed secret sharing inference on augmented party models.
- We evaluate `CCBNet` over various scenarios based on nine public BNs. Our results indicate that `CCBNet`'s predictive accuracy is similar to those of non-confidential centralized alternatives and, for many collaborators in large networks, that `CCBNet` decreases the computational overhead by 23% on average when 71k values are communicated per request.
- We discuss possible attacks and mitigations on the two `CCBNet` subprotocols.

Our code is available at: https://github.com/AbeleMM/ccbnet.

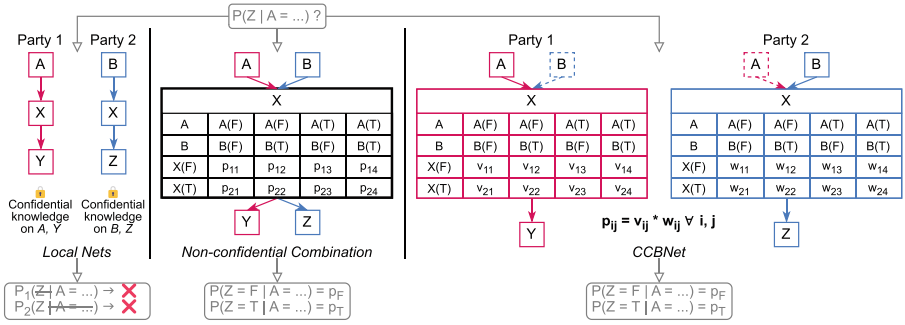


Fig. 1. Bayesian networks collaboration comparison: Variables are X (shared), A, B, Y, and Z. They all have two states, **F**alse and **T**rue. Entries p_{ij} are conditional probability values (e.g., $p_{11} = P(X = F \mid A = F,\ B = F)$). v_{ij} and w_{ij} are secret shares which multiplied reconstruct p_{ij}. After observing A, inferring Z's updated posterior $P(Z \mid A = ...)$ requires propagating information between A, B, and X, which is impossible by averaging model outputs. Meanwhile, the non-confidential approach reveals the combined graph and probability tables to parties. `CCBNet` exchanges minimal structural information among parties and runs distributed inference with secretly shared confidential values.

2 Background and Related Studies

2.1 Background

Bayesian Networks are probabilistic graphical models that maintain explicit conditional probability distributions (CPDs) for variables which form a directed acyclic graph [26]. Shown in Fig. 1, variables and the influences between them give the graph nodes and edges, respectively. Principally for performance and

interpretability, variables in practical applications are generally discrete, with CPDs specifically embodying conditional probability tables. Learning may be driven by data, human experts, or both [8,18], as with other human-readable models like decision trees. Automated learning discovers the graph structure and then populates CPD parameters from training data. Manual learning is desirable when incorporating concepts with known governing rules that need no approximation from observations.

Computation in discrete BNs relies on a few base operations for propagating information: normalization, reduction, marginalization, and products [18]. We outline them with help from the non-confidential combination in Fig. 1. Normalizing a CPD divides its entries by their column sum. Thus column summations, like $p_{11} + p_{21}$, would become 1. Reduction and marginalization remove variables from a CPD by fixing their states or, respectively, summing them out. Reducing or marginalizing A from X leaves B as its sole parent. Previous operations apply to both representations. Products operate on CPDs of the same variable or factors and create a new CPD/factor over the input variables' union, where each entry is the multiplication of the corresponding ones in the original representations. The product of X and Z's factors, thus, also contains A and B.

In the considered discrete scenario, a CPD with target variable X, and parents $A, B, ...$ is denoted as an $(S_X \times S_A \times S_B \times ...)$ array, accompanied by a mapping between dimensions to variables and indices to states. S_V is the number of states for variable V. Each entry specifies a probability value for the state combination it represents. The corresponding factor is identical but does not distinguish between a target and parents. Any continuous variables would need to have a globally agreed-upon discretization applied to them before usage.

Inference in BNs finds updated posterior probabilities for the states of target variables, given the observed states of any evidence variables [23,24]. As general inference is NP-Hard, approximate algorithms help decrease computation costs compared to exact ones while sacrificing some precision in the result. The main exact inference techniques are variable elimination (VE) and junction tree belief propagation, which decomposes the network into a tree of variable clusters, runs VE within them and then disseminates updates between neighbors by message-passing [18]. In Fig. 1's query, Z is the target, given some observed state of A. The VE algorithm performs three main steps over the factors corresponding to the model's CPDs. It first reduces the evidence from input factors. Reduction merely discards entries in the factor where evidence variables have a different state than specified. Then, it iteratively performs products between factors and marginalizes them for each non-query variable. Lastly, it takes the product between the leftover factors to get the final result factor.

Markov random fields (MRFs) are a sibbling model of BNs, backed by undirected graphs, into which every BN is easily transformable via moralization [19,25]. Apart from lacking acyclicity constraints, MRFs directly define parameters as factors and can deal with scenarios where edge directionality is unspecified, but BNs are more compact and efficient for generative use. For

inference, BN properties and algorithms remain applicable. Thus, when aciclicity constraints are unsatisfiable, MRFs can still perform inference like in BNs.

2.2 Prior Art

We identify two high-level categories of collaborative analysis for BNs: single- and multi-model. The first creates one global model, while the other keeps local models separate, merging their results.

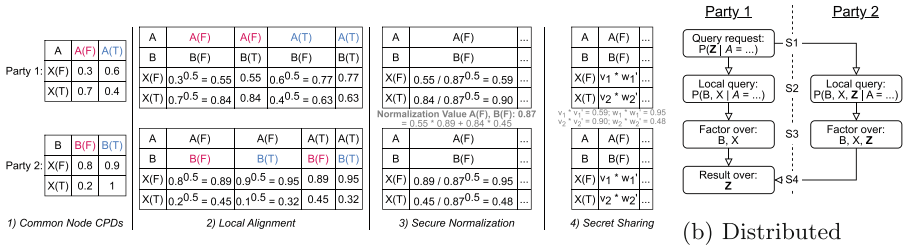
Single-model approaches harness party data instances or models but neglect model confidentiality. Federated learning discovers models [1,7,12,15,21] from private party instances with a coordinator. Fully decentralized formulations also exist [4,13]. Direct network combination fuses structure [2,10] and parameters [11] from party models.

Multi-model methods have parties work together during inference to produce a complete analysis result. Tedesco et al. [27] chain model without exchanging their contents but only allows using them one at a time in a predetermined order. Less confidential but more flexible, the work of Pavlin et al. [22] fuses party networks based on common nodes but still requires them to be roots or leaves in the party's directed acyclic graph. Ypma et al. [29] patent a collaborative solution for industrial processes, but only mention data confidentiality preservation by anonymization. Kim and Ghahramani [17] run models autonomously and only average their final outputs, maintaining confidentiality but expecting models to share inputs and outputs.

Single-model solutions break confidentiality by centralizing knowledge. Multi-model ones trade modeling power for confidentiality. `CCBNet` addresses both.

3 `CCBNet`

We propose, `CCBNet`, a framework for secure distributed analysis over a related set of confidential and discrete BNs. `CCBNet` is composed of two key steps, (i) `CABN` augments the BNs of parties through overlapping variables, and (ii) `SAVE` performs joint inference on them.



(a) Key Steps of computing augmented CPDs

(b) Distributed inference flow

Fig. 2. `CABN` & `SAVE` steps for Fig. 1 scenario.

The assumptions we make are that variables from different parties have the same name only if they represent the same concept, and the independence, across parties, of distinct parents for the same node reasonably approximates the ground truth. These are shared by previous BN combination works. Thus, names identify the overlapping (common) nodes between models, giving the contact points for graph fusion. Since variables modeled by parties may be any subset of those from the full domain, modeling direct interactions between their non-overlapping variables requires great amounts of often unavailable information.

Our adversarial model includes semi-honest parties that follow the protocol while trying to abuse gained information [14] but do not collude. No trusted third party exists. The goal is to protect all network parameters and only disclose structure/state-name information among parties modeling the same variables.

Algorithm 1. `CABN`

Input: $P = \{p_1,\ p_2,\ \ldots,\ p_n\}$ – parties,
$CPD_p[T \mid A,\ B,\ \ldots]$ – CPD in party p for target var T with parent vars $\{A,\ B,\ \ldots\}$,
$weight_p$ – weight assigned by p to its BN (defaults to 1)
Output: updated $CPD_p[O \mid \ldots]$, $\forall$ overlap var O modeled by party p

for each var V of parties from P **do**
  $parties_V \leftarrow \varnothing$ {list of parties containing V}
  $states_V \leftarrow \varnothing$ {list of overlap states for V}
for each pair of parties $(p,\ p')$ from P **do**
  for each var $O \in PSI$(vars of p, vars of p') **do**
    $parties_O \leftarrow parties_O \cup \{p,\ p'\}$
for each var O s.t. $parties_O \neq \varnothing$ **do**
  for each var set $\mathcal{V} = \{X, Y, \ldots\}$ s.t. $(p \in parties_O \wedge CPD_p[O \mid \mathcal{V}]$ exists) **do**
    $states_V \leftarrow states_V\ \cup\ \{v_0,\ v_1,\ \ldots\}$ {where v_i is a state of var V}
  $idCPD \leftarrow$ CPD for target var O and $states_V$ with all entries 1
  wSum $\leftarrow \sum_{p \in parties_O} weight_p$
  for each party $p \in parties_O$ **do**
    $CPD_p[O \mid \ldots] \leftarrow CPD_p[O \mid \ldots] \cdot idCPD$ {factor prod}
    $CPD_p[O \mid \ldots] \leftarrow CPD_p[O \mid \ldots]^{weight_p / wSum}$ {element-wise exp}
  for each parent config $\mathcal{S} = (X = x_i,\ Y = y_j,\ \ldots)$ s.t. $idCPD[O \mid \mathcal{S}]$ exists **do**
    $prodHE \leftarrow \odot_{p \in parties_O} HE_{enc}(CPD_p[O \mid \mathcal{S}])$ {enc element-wise prod}
    $norm \leftarrow HE_{dec}(||prodHE||_1)^{1/|parties_O|}$ {dec normalization scalar}
    for each party $p \in parties_O$ **do**
      $CPD_p[O \mid \ldots] \leftarrow CPD_p[O \mid \ldots]\ /\ norm$
  for each party $p \in partiesO$ **do**
    $CPD_p^0[O \mid \ldots],\ CPD_p^1[O \mid \ldots],\ \cdots \leftarrow$ secret share $CPD_p[O \mid \ldots]$
  for each party $p_i \in partiesO$ **do**
    $CPD_{p_i}[O \mid \ldots] \leftarrow \prod_{p \in parties_O} CPD_p^i[O \mid \ldots]$
return $CPD_p[O \mid \ldots]$ $, \forall\ parties_O \neq \varnothing,\ p \in parties_O$

3.1 Confidentially Augmented Bayesian Networks

We now present the CABN[1] protocol, which updates local CPDs for overlap variables to hold secret shares of their combination, protecting the initial probabilities. Algorithm 1 details the four steps of the protocol, illustrated in Fig. 2 a: (i) private common node identification; (ii) local alignment; (iii) secure normalization; and (iv) secret sharing. The protocol updates parties when the number of changes in local networks passes a set threshold.

Overview. Structurally, CABN imitates a union of the involved networks of Del Sagrado and Moral [10]. Parameter-wise, it follows Feng et al. [11] but replaces the superposition operator with the geometric mean. We use the union instead of the more complex ruleset of Feng et al. to decide which overlapping node parents to retain. The more straightforward allows for combining more than two BNs at a time, lowering the number of communication rounds. It also reduces the attack surface area due to fewer checks on party data. For fusing probabilities, the geometric mean enables a multiplication-based secret sharing scheme in CABN. Reconstruction happens automatically during distributed inference when merging local results from parties. The mean also outperforms the superposition in our centralized tests.

Step 1: Private Common Node Identification. CABN starts with party pairs identifying their common nodes like in the central case, albeit privately. We use a private set intersection (PSI) protocol [9] to attain confidentiality (ll. 1-6). Only parties that have updates need to recalculate their intersection with the others. Outside the private intersection context, node and state names are communicated obfuscated to prevent information leakage about which parties model which nodes. Parties can choose any unique representation for non-overlapping nodes, but involved parties agree on an obfuscated representation for overlapping ones. Our implementation uses a commutative elliptic curve ElGamal cypher.

Step 2: Local Alignment. After parties know which local nodes overlap with which peers, they start solving overlaps by exchanging structure and weight information about their local CPDs and independently updating local representations accordingly (ll. 7-14). First, the obfuscated union of CPD nodes and states for overlapping CPDs is determined (ll. 8-9). From it, an identity CPD containing all the parents across parties gets created for the union (l. 10). An identity CPD (or factor) has all entries equal to 1, so its product with another replicates the later's columns over their joint state space. The initial CPD gets replaced by the product with the identity in each party, giving all overlap CPDs the same shape (l. 13), as seen in Fig. 2 a. Parties have a public weight representing confidence in their BN, which in the default unweighted case is 1. They compute the sum of their weights (l. 11) and individually raise the entries of their CPD to the ratio between their weight and the sum, computing the partial geometric mean (l. 14). By the exponent product rule $(XY)^k = X^k Y^k$, the CPDs' product already yields the unnormalized CPD of the central combination.

Model Weighting. As previously mentioned, CABN allows weighting CPDs through the geometric mean, unlike previous BN works that cover parameter

[1] **C**onfidentially **A**ugmented **B**ayesian **N**etworks.

fusion. A natural integration of unequal weighting of inputs is another advantage of using a geometric mean instead of Feng et al.'s superposition. We implement weights at the model level as 0–1 values, encoding the human expert's confidence in the network or data availability for algorithmic learning. Nevertheless, weighting can be applied at the CPD level.

Step 3: Secure Normalization. Homomorphic encryption (HE) [5] allows to privately compute column normalization values (ll. 15-19). One party is elected to generate the encryption key pair and another to perform the encrypted computation. All parties receive the public key and send their encrypted columns to the party that calculates normalization value ciphers (l. 16). The private key party decrypts the values and shares them with the rest (l. 17). A column normalization value is the sum of entries obtained by multiplying matching party columns element-wise. Letting P_{ij} denote party i CPD column j, we thus find $|| \odot_i P_{ij}||_1$. Then, the K overlapping parties individually divide each column by the K-th root of the appropriate normalization value (l. 19), so their factor product is the normalized geometric mean. Figure 2 a shows an example. Our implementation relies on the CKKS [5] HE scheme for floating-point addition and multiplication.

Because local columns no longer sum to 1 after exponentiation, even in a two-party overlap where the variable has only two possible states, a party cannot reconstruct the other's entries by only knowing the normalization values and its own entries. Furthermore, vital for HE schemes in practice, we know that the number of consecutive multiplications needed for each column is equal to the party count, which allows configuring the scheme accordingly. Functional encryption, in which completing the desired computation also decrypts the output [3], can be a more viable choice, but existing implementations have overly stringent limits on the number of inputs and complexity of the applied functions. Using a secret sharing scheme (SSS) [6] instead of HE is also possible, but we favor decreasing the communication count over computing overhead for this step. An SSS has the advantage of requiring fewer computational resources and being more robust against collusion. However, it requires communication for each multiplication operation and, depending on the scheme, the presence of a third party. Despite the expectation that `CABN` needs to run more rarely than inference, we still favor optimizing for message count. High communication latency is likelier to be a bottleneck than processing for envisioned deployments.

The only extra role of the two elected parties, K and C, is to manage the private/public key pair and compute with the encrypted data, respectively. We elect both parties from those in the overlap, so their CPDs are also part of the calculation. They can be picked based on any criteria, like the order of joining the network. Party K forwards the public key to all others in the overlap, C excluded. Parties, K included, send their encrypted data to C, which calculates the normalization value. The result gets forwarded to K, which decrypts and shares it with all overlap parties. Apart from the normalization value, party K only ever has the chance to decrypt data it owns. The elected parties may also be outside the overlap as long as they are different entities.

Algorithm 2. SAVE

Input: $P = \{p_1,\ p_2,\ \dots\}$ – parties,
$\mathcal{Q} = \{X,\ Y,\ \dots\}$ – query vars, $\mathcal{E} = \{a_i,\ b_j,\ \dots\}$ – evidence var states
Output: factor over $\mathcal{Q}$

1: **for each** party $p \in P$ **do**
2: {receive $\mathcal{Q}$ & $\mathcal{E}$ from querying party}
3: $inFacts_p \leftarrow$ set of CPDs in p as factors
4: $\mathcal{Q}_p \leftarrow$ all overlap vars of p and their parents
5: $outFact_p \leftarrow VE(\mathcal{Q} \cup \mathcal{Q}_p,\ \mathcal{E},\ inFacts_p)$ {exec by p}
6: {send $outFact_p$ to querying party}
7: **return** $VE(\mathcal{Q},\ \{\},\ \bigcup_{p \in P} outFact_p)$ {exec by querying party}

Step 4: Secret Sharing. Finally, to combat party parameter leaks at inference, we secret share [16] the CPD entries of parties in each overlap through a multiplication-based scheme (ll. 20-23). The scheme allows using the product operation needed for inference while exchanging a similar number of messages to HE and achieving much better computation scaling. The common shape of updated local CPDs facilitates the procedure. In the classic additive secret sharing scheme, a secret value is split into shares distributed among parties whose sum is the secret. It allows efficient and secure computation of expressions summing multiple secret values and applying other operations involving non-secret values. Parties perform the computation with their local share of each secret and all aggregate their results to reconstruct the answer. The utilized scheme functions similarly but uses multiplication as the base operation instead. Reconstruction happens implicitly during inference, with no extra overhead since other factors containing partial results from parties get incorporated into the final result via the same product operations. Figure 2 a exemplifies share splitting in CPDs.

Handling potential cycles. To avoid compatibility restrictions between combinable BNs, if solving overlaps creates a cycle, the distributed global network gets treated as an MRF, with no changes to the inference, which operates on factors regardless. Edges that form cycles in the BN are effectively incorporated into the moralized MRF and treated as undirected. Since the main target is to not share the complete combined network, the readability advantages of BNs are not a concern in the joint global model. Deciding which edge to remove from a cycle often requires unavailable information and threatens confidentiality. Alternatives like treating all nodes within a cycle as a single node [29] are coarse-grained and threaten confidentiality.

3.2 Share Aggregation Variable Elimination

SAVE[2] is the inference protocol wherein all parties run VE locally before aggregating their outputs into the final factor. Algorithm 2 describes its steps, and

[2] **S**hare **A**ggregation **V**ariable **E**limination.

Fig. 2 b visualizes an example query. Parties execute most of the process in parallel (algo. l. 1). They first receive the obfuscated target variables $\mathcal{Q}$ and evidence $\mathcal{E}$ from the querying party (algo. l.2, S1 in fig.). Then, they prepare all their CPDs as factors (algo l. 3). They extend the queried variables set with local overlap variables and their parents (algo. l. 4). Extending the set avoids illegal marginalization operations within our SSS. Although the individual operations have no overhead via our scheme, the changed order of operations is less efficient, increasing the cost of VE. Parties each run VE over their factors and extended query set (algo. l. 5, S2 in fig.). To reduce overall communication strain, our implementation defaults to stopping VE early and returning a list of disjoint smaller factors instead of a single large one. Consequently, each party sends its result to the querying party (algo. l. 6, S3 in fig.). Finally, the querying party runs VE with the original set of query variables and received factors to get the query result factor (algo. l. 7, S4 in fig.).

Applying a *log* transform to values and using additive secret sharing instead of the multiplication-based variant would also work for products but cause even more marginalization issues. Allowing multiplications between shares and plaintext would require taking the *log* of all values, not just shares. Consequently, plaintext marginalization becomes impossible since addition has no correspondence under *log*. Fixed-point arithmetic keeps the scheme's theoretical guarantees by mapping fractional CPD values to integers. Within it, rounding avoids overflows from multiplication chains.

Queryable Nodes. To maintain confidentiality, parties can only specify modeled variables in inference by default, even if the result still reflects the effect of prior knowledge about others, so we propose mechanisms for expanding the set of possible queries. The first involves all parties that own a node agreeing to expose its unobfuscated name and states with select others to use as a target or evidence. Doing so only requires revealing a node's existence, not its place in the network(s). The other mechanism implicitly enhances evidence with the help of some key shared between parties (e.g., timestamp, product batch identifier). If a query request also includes a value for the shared key, parties incorporate any observations for the key's value as evidence during their local inference step. Parties do not have to disclose the value of the observed data, but the query output is the same as if it had been part of the initial evidence.

Communication Properties. The number of messages exchanged within an inference request is of magnitude $O(N)$ (where N is the number of parties), but the size of the messages varies. Regarding count, the requester sends out $N-1$ messages and receives the same amount of replies adding up to $2N-2$. The size of the messages, particularly replies, varies greatly depending on the number and complexity of the responder's overlaps and the query itself.

4 Performance Evaluation

Our main results show `CCBNet` achieving the predictive ability of centralized solutions in the confidential, collaborative setting and examine its ability to maintain reasonable computation and communication costs.

4.1 Setup

Table 1. Evaluation datasets with node/edge/parameter counts

Class	Name	#Nodes	#Edges	#Params
Small (< 20 Nodes)	ASIA	8	8	18
Medium (20–49 Nodes)	CHILD	20	25	230
	ALARM	37	46	509
	INSURANCE	27	52	1008
Large (50–99 Nodes)	WIN95PTS	76	112	574
Very Large (101–999 Nodes)	ANDES	223	338	1157
	PIGS	441	592	5618
	LINK	724	1125	14211
Massive (≥ 1000 Nodes)	MUNIN2	1003	1244	69431

Our experiments evaluate average predictive performance, computation overhead, and communication cost of inference in single-machine simulations. We measure prediction quality based on the Brier Score ($= \frac{1}{N}\sum_{t=1}^{N}\sum_{i=1}^{R}(f_{ti}-o_{ti})^2$, where N is the number of queries, R is the number of target variables state combinations, while f and o are predicted and reference probabilities). Note that, different from the original multi-category definition, we have probability values for all reference entries o instead of a binary value differentiating an observed (1) state from all others (0). We report the total processing time ratio between examined methods and the ground truth network for computation overhead. For communication we consider the count of factor values exchanged per query. We do not evaluate `CABN` time or communication as we expect it to be amortizable.

We test on public networks[3] (Table 1), and consider two variable splitting methods alongside multiple overlap variable ratios. **Related splits** assign to parties variables connected in the ground truth network. **Random splits** ensure parties have equal variable counts and share the same overlaps. Test sets contain 2000 queries, each with one random overlap variable as the target and 60% of the others fixed as evidence. Related splits attempt to simulate a realistic deployment in which the experts within clients have an incomplete but close-to-the-ground truth view of the interactions between their modeled variables (which also differ in number). Random splits aim for a worst-case scenario where the variables within clients (of roughly equal number) are not subgraphs of the complete network, generally giving more densely connected local networks and overlaps.

[3] https://www.bnlearn.com/bnrepository/.

4.2 Baselines

As the closest prior work is strictly centralized, our baselines are:

Centralized Combination (CC) iteratively combines parties by the method of [11] and treats the network as an MRF if cycles form.

Centralized Union (CU), the approach confidentially mimicked by `CCBNet`, is structurally based on the union of [10], combines parameters via geometric mean, and also applies the MRF principle.

Decentralized Output Mean (DOM) is a naïve approach that takes the geometric mean for each target variable's state probabilities over independently operating contained parties, trading long-range effect modeling for greater safety.

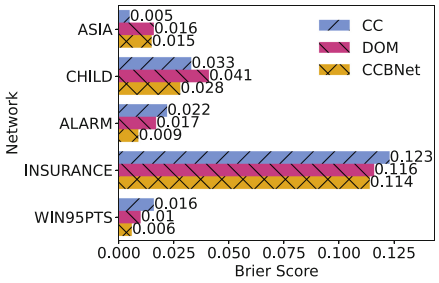
`CCBNetJ` is a degenerate `CCBNet` variant that stores fully combined central CPDs for overlaps in singular parties, trading some safety for faster inference.

4.3 `CCBNet` Performance Overview

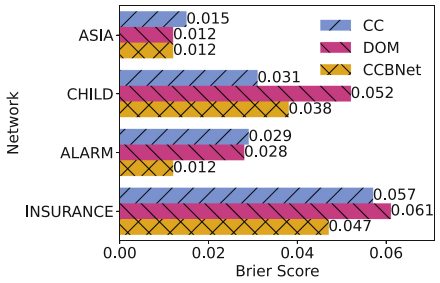
In Fig. 3, we summarize Brier score, computation time, and communication cost results under two splitting methods, four parties, and a 30% overlap ratio.

Predictive Performance. Regardless of split type, `CCBNet` predictions often outperform or match the classic CC and always beat the naïve DOM. Figure 3 a gives results for related splits, `CCBNet` only scores worse than CC on the smallest network (0.015 versus 0.005) and gains a significant advantage over the larger ones (0.022 versus 0.009, 0.123 versus 0.114, and 0.006 versus 0.016). `CCBNet` has an advantage over DOM in all scenarios. Examining the random splits in Fig. 3 b, `CCBNet` still outscores CC in all but one of the smaller networks (0.031 vs 0.038) and maintains its lead over DOM in all but the smallest network, where the two tie. Since `CCBNet` yields the same predictive ability as its centralized counterpart CU, the differentiating points with CC are the structure combination policy and parameter combination operator. Over both splits, the pairing of the union selection and geometric mean aggregation in `CCBNet` fares better in most cases than the nondeterministic selection and superposition aggregation of CC, apart from some instances within small networks.

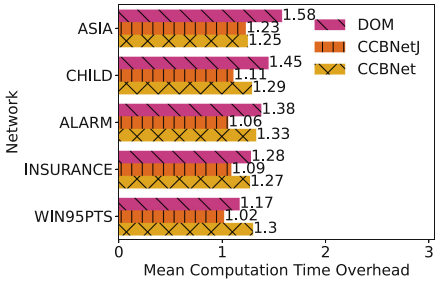
Computation Overhead. Regarding computation cost relative to centralized inference on the original network, average slowdowns are 1.63x for DOM, 1.15x for `CCBNetJ`, and 1.6x for `CCBNet`. Communication latency is unaccounted for as it can vary greatly based on the deployment. Still, we overestimate wall time by summing computation time across parties, as much processing would happen concurrently in reality. In the related splits from Fig. 3 c, across the board, `CCBNetJ` is the fastest, but others follow closely. `CCBNet` is the second best in all except the largest networks, where it fares worse than DOM (1.17x vs 1.3x slower). The networks also examined for random splits in Fig. 3 d show a similar trend, with somewhat higher general overhead, especially for `CCBNet`. Thus, as expected, `CCBNetJ` is faster than `CCBNet`, but both perform reasonably in most scenarios, even if the latter suffers more as overlaps increase in complexity. The DOM implementation does has a higher base overhead, but the gap to the other



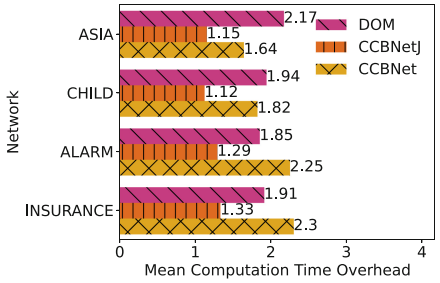
(a) Predictive performance, related splits



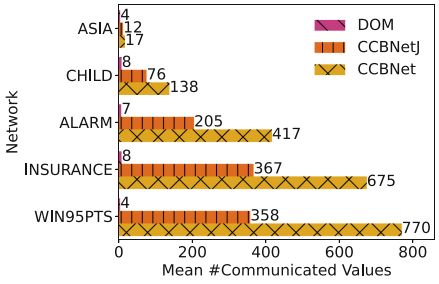
(b) Predictive performance, random splits



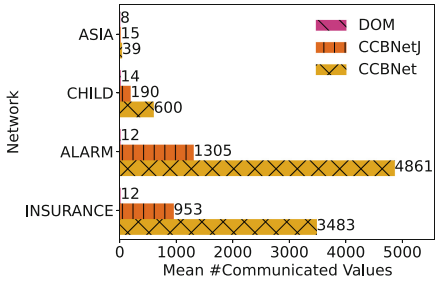
(c) Computation overhead, related splits



(d) Computation overhead, random splits



(e) Communication cost, related splits



(f) Communication cost, random splits

Fig. 3. Performance of decentralized methods relative to ground truth (lower is better); 4 parties, 30% of vars in > 1 party.

algorithms is often relatively contained. As for absolute total computation time, queries are slowest on the largest related splits dataset, as DOM averages 2.9 ms/query, `CCBNetJ` 2.5 ms, and `CCBNet` 3.2 ms.

Communication Cost. DOM averages merely 9 CPD values communicated per query, while `CCBNetJ` and `CCBNet` need orders of magnitude more at 387 and 1222, respectively. Since the number of communicated values depends on which party initiates a query, the reported figures include communication within the querying party to eliminate variability but overestimate reality. The number of messages to complete a query is the same for all methods. Furthermore, the

raw data transmitted in bytes remain in the low megabyte range for hundreds of thousands of values before compression. Figure 3 e shows the mentioned discrepancy over related splits for all but the smallest network, in which the three methods are comparable. DOM merges complete party outputs and cannot propagate evidence between parties. Thus, it does not increase communication with the number of overlaps, and parties that do not contain any target variables send empty replies. The situation for random splits, illustrated in Fig. 3 f, is very similar, although the disadvantage of `CCBNet` over `CCBNetJ` widens considerably.

4.4 Party Weighting

Table 2 shows the weighted version of the proposed method having better predictive performance than the unweighted one in almost all scenarios with random splits. In weighting tests, we reduce the overall amount of data used for learning local BNs to ensure more variance and alternatingly assign parties a smaller or larger fraction of training data. Over the one scenario where the unweighted version performs better (0.079 versus 0.104), parties with lower data get overly punished for their perceived imprecision. Since each CPD has a single weight, all parents of a node within the party are still treated uniformly according to that value, even if there is a mismatch between the weight and the actual quality of the estimates. Similarly, if a party with lower overall confidence is the only one to model a highly influential parent, its importance is underrepresented in the final result. Nevertheless, despite these phenomena, which can adversely affect performance, weighting has a positive overall impact in tested scenarios.

Table 2. Unweighted **(UW)** & Weighted **(W)** Brier scores for `CCBNet` relative to original network – random splits, imbalanced learning data (lower is better)

#Parties		2				4			
Vars in >1 Party		10%		30%		10%		30%	
Method		UW	W	UW	W	UW	W	UW	W
Network	ASIA	0.028	**0.018**	0.029	**0.009**	0.036	**0.020**	0.031	**0.011**
	CHILD	**0.067**	**0.067**	0.088	**0.047**	0.084	**0.067**	0.097	**0.048**
	ALARM	0.063	**0.058**	0.061	**0.050**	0.103	**0.092**	0.061	**0.057**
	INSURANCE	**0.079**	0.104	**0.051**	**0.051**	0.112	**0.088**	0.135	**0.098**

4.5 Large Networks and Many Parties

Table 3. Metrics for large networks & many parties – related splits, 10% of vars in >1 party (lower is better)

Networks/	Brier Score			Avg Comp Time Overhead			Avg #Comm Values		
Parties	CC	DOM	`CCBNet`	DOM	`CCBNetJ`	`CCBNet`	DOM	`CCBNetJ`	`CCBNet`
ANDES/16	0.040	0.039	**0.035**	0.58	0.62	0.68	4	317	445
PIGS/32	0.103	0.092	**0.076**	0.39	0.51	1.34	7	4202	37838
LINK/64	0.144	**0.125**	0.126	0.25	0.35	0.39	6	2402	4549
MUNIN2/128	0.016	**0.015**	**0.015**	0.20	0.41	0.67	11	88581	242788

Lastly, in Table 3, our tests for challenging use cases with large networks (223–1003 variables), many parties (16–28), and related variable splits, but lower overlaps reconfirm prediction/communication trends, yet computation improves over original networks. In larger networks, `CCBNet`'s predictions outperform CC, and communication size increases with parties and network size, averaging 71k values per request. Computation overhead is always <1, for a mean speedup of 21%. The hardness of inference makes approximating a big network by splitting it into chunks faster, even before considering parallel party solving.

5 Attacks on `CCBNet`

Carelessly combining related BNs into a single global model has a high risk of leaking confidential information to all parties that can access it. As illustrated in the middle of Fig. 1: (i) at a purely structural level the centralized combination can contain a large amount of the local party information; and (ii) at a parameter level, probability functions for any non-overlap nodes remain unmodified. Since BNs are human-readable, inspection can compromise sensitive information even before any inference. We review two attacks that respectively reconstruct CPDs during `CABN` and `SAVE`, along with their implications in `CCBNet` and `CCBNetJ`. The attacks do not bypass the obfuscation of unowned variable names and states but still expose potentially sensitive information via the recovered probability values. We execute these attacks as visualized in Fig. 4 on a corner case of **two-party ASIA and CHILD network instances**.

The `CABN` **attack** allows a party to reconstruct a peer's CPD for an overlap variable without inference, assuming no other parties are involved in the overlap, the attacker holds the combined CPD, and the protocol is `CCBNetJ`. The attacker first runs `CABN` to compute the combined CPD. Then, it removes its contribution from the geometric product that yielded the CPD, marginalizes any parents that should not be present in the victim's version, and normalizes to get the final result. With any secure computation scheme, if one input parties knows

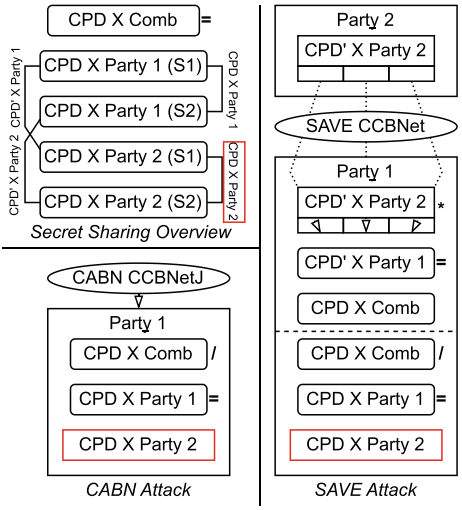


Fig. 4. `CABN` & `SAVE` attacks visualization.

the output and the applied operations are reversible, it can find the other input. With three or more parties, the attacker can only reconstruct an aggregation of the other involved CPDs. Even when the attack is possible, name obfuscation still hides the real-world meaning of the variables. Should any party outside the overlap exist, in the assumed no collusion setting, having it compute and store the combined CPD instead gives a minimum-change fix. `CCBNet` is not vulnerable, as it does not join shares before applying inference operations.

The `SAVE` **attack** reconstructs a peer's CPD for an overlap variable via inference assuming that the overlap contains no other parties but also applies to `CCBNet`. As one of the two parties in the overlap, the attacker begins by querying for the overlap variable as the only target, specifying a state for each of its parents in the evidence. With all parents fixed, no other variables in the victim's network can affect the transmitted result. Thus, the attacker receives one row from the victim's share of the combined target CPD. The attacker repeats the procedure for all other configurations of parent states, building up the victim's complete CPD share. It then obtains the full combined CPD as the product of its share and the one recovered from the other party. Finally, like in the previous attack, it removes the contribution of its whole local CPD from the combination, marginalizes any parents not in the victim, and normalizes to find the result. Also similar to the previous attack, obfuscation limits the damage that can be done, while with the involvement of more than two parties, the attacker is able to compromise their shares, but not the contribution of each, before sharing the secrets. Thus, secret sharing avoids the attack for overlaps with three or more parties. As the attack requires a series of specific queries, redundant in most real applications, a simple defense has parties limit the number of requests serviced that target some node and assign all parents in the evidence.

6 Conclusion

We propose `CCBNet` to address the issue of collaborative analysis for BNs in confidential (manufacturing) settings. The framework allows distributed analysis spanning multiple models without revealing their contents. It has no model compatibility restrictions and allows unequally weighting parties. Results show `CCBNet` outperforming/matching distributed/centralized baselines while maintaining reasonable computation time that decreases to 23% of centralized formulations in large networks with many parties but produces much larger messages the distributed baseline, averaging 71k communicated values per request. Altogether, `CCBNet` offers centralized-like predictive performance in a distributed setting and ensures a base for protecting parties' private information.

Acknowledgments. This research was partly funded by the NWO Perspectief project, DepMAT, and the SNSF project, Priv-GSyn 200021E_229204.

References

1. Abyaneh, A., Scherrer, N., Schwab, P., Bauer, S., Schölkopf, B., Mehrjou, A.: Fed-cd: federated causal discovery from interventional and observational data (2022)
2. Alrajeh, D., Chockler, H., Halpern, J.Y.: Combining experts' causal judgments. Artif. Intell. **288**, 103355 (2020)
3. Boneh, D., Sahai, A., Waters, B.: Functional encryption: definitions and challenges. In: Theory of Cryptography, pp. 253–273 (2011)
4. Campbell, T., How, J.P.: Approximate decentralized bayesian inference (2014)
5. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 409–437. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_15
6. Cramer, R., Damgård, I.B., Nielsen, J.B.: Secure Multiparty Computation and Secret Sharing. Cambridge University Press (2015)
7. van Daalen, F., Ippel, L., Dekker, A., Bermejo, I.: Vertibayes: Learning bayesian network parameters from vertically partitioned data with missing values (2022)
8. Daly, R., Shen, Q., Aitken, S.: Learning bayesian networks: approaches and issues. Knowl. Eng. Rev. **26**(2), 99–157 (2011)
9. De Cristofaro, E., Tsudik, G.: Practical private set intersection protocols with linear complexity. In: Financial Cryptography and Data Security, pp. 143–159 (2010)
10. Del Sagrado, J., Moral, S.: Qualitative combination of bayesian networks. Int. J. Intell. Syst. **18**(2), 237–249 (2003)
11. Feng, G., Zhang, J.D., Shaoyi Liao, S.: A novel method for combining bayesian networks, theoretical analysis, and its applications. Pattern Recogn. **47**(5), 2057–2069 (2014)
12. Gao, E., Chen, J., Shen, L., Liu, T., Gong, M., Bondell, H.: Feddag: federated DAG structure learning (2021)
13. Gholami, B., Yoon, S., Pavlovic, V.: Decentralized approximate bayesian inference for distributed sensor network. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 30, no. 1 (2016)

14. Goldreich, O.: Foundations of cryptography – a primer. Foundations and Trends in Theoretical Computer Science **1**(1), 1–116 (2005)
15. Huang, J., Yu, K., Guo, X., Cao, F., Liang, J.: Towards privacy-aware causal structure learning in federated setting (2022)
16. Kilbertus, N., Gascon, A., Kusner, M., Veale, M., Gummadi, K., Weller, A.: Blind justice: fairness with encrypted sensitive attributes. In: Proceedings of the 35th International Conference on Machine Learning. vol. 80, pp. 2630–2639 (2018)
17. Kim, H.C., Ghahramani, Z.: Bayesian classifier combination. In: Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics. vol. 22, pp. 619–627 (2012)
18. Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. MIT Press (2009)
19. Li, S.Z.: Markov random field modeling in image analysis. Springer Science & Business Media (2009)
20. Nannapaneni, S., Mahadevan, S., Rachuri, S.: Performance evaluation of a manufacturing process under uncertainty using bayesian networks. J. Clean. Prod. **113**, 947–959 (2016)
21. Ng, I., Zhang, K.: Towards federated bayesian network structure learning with continuous optimization (2021)
22. Pavlin, G., de Oude, P., Maris, M., Nunnink, J., Hood, T.: A multi-agent systems approach to distributed bayesian information fusion. Inform. Fusion **11**(3), 267–282 (2010)
23. Pearl, J.: Causality. Cambridge University Press, 2 edn. (2009)
24. Russell, S.J.: Artificial intelligence a modern approach. Pearson Education (2010)
25. Scutari, M., Denis, J.B.: Bayesian networks: with examples in R. CRC Press (2021)
26. Stephenson, T.A.: An introduction to bayesian network theory and usage. Tech. rep, Idiap (2000)
27. Tedesco, R., Dolog, P., Nejdl, W., Allert, H.: Distributed bayesian networks for user modeling. In: Proceedings of E-Learn: World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education 2006, pp. 292–299 (2006)
28. Yang, W.T.: An Integrated Physics-Informed Process Control Framework and Its Applications to Semiconductor Manufacturing. Université de Lyon, Theses (2020)
29. Ypma, A., Koopman, A.C.M., Middlebrooks, S.A.: Methods & apparatus for obtaining diagnostic information, methods & apparatus for controlling an industrial process (2018)