

Reward-based Blockchain Infrastructure for 3D IC Supply Chain Provenance

Sulyab Thottungal Valapu[†], Aritri Saha*, Bhaskar Krishnamachari[‡], Vivek Menon**, and Ujjwal Guin*

[†]Dept. of Computer Science, University of Southern California

*Dept. of Electrical and Computer Engineering, Auburn University

[‡]Dept. of Electrical and Computer Engineering, University of Southern California

**National Reconnaissance Office

Emails: {thottung,bkrishna}@usc.edu, {aritri.saha,ujjwal.guin}@auburn.edu, and menonviv@nro.mil

Abstract—In response to the growing demand for enhanced performance and power efficiency, the semiconductor industry has witnessed a paradigm shift toward heterogeneous integration, giving rise to 2.5D/3D chips. These chips incorporate diverse chiplets, manufactured globally and integrated into a single chip. Securing these complex 2.5D/3D integrated circuits (ICs) presents a formidable challenge due to inherent trust issues within the semiconductor supply chain. Chiplets produced in untrusted locations may be susceptible to tampering, introducing malicious circuits that could compromise sensitive information. This paper introduces an innovative approach that leverages blockchain technology to establish traceability for ICs and chiplets throughout the supply chain. Given that chiplet manufacturers are dispersed globally and may operate within different blockchain consortiums, ensuring the integrity of data within each blockchain ledger becomes imperative. To address this, we propose a novel dual-layer approach for establishing distributed trust across diverse blockchain ledgers. The lower layer comprises of a blockchain-based framework for IC supply chain provenance that enables transactions between blockchain instances run by different consortiums, making it possible to trace the complete provenance DAG of each IC. The upper layer implements a multi-chain reputation scheme that assigns reputation scores to entities while specifically accounting for high-risk transactions that cross-blockchain trust zones. This approach enhances the credibility of the blockchain data, mitigating potential risks associated with the use of multiple consortiums and ensuring a robust foundation for securing 2.5D/3D ICs in the evolving landscape of heterogeneous integration.

Index Terms—3D IC, heterogeneous integration, blockchain, Byzantine Fault Tolerance.

I. INTRODUCTION

The increasing demands for computational power in high-performance computing (HPC), data centers, cloud computing, and machine learning have surpassed the capabilities of the current System-on-Chip (SoC) paradigm. Meeting the challenges posed by issues such as latency, power consumption, throughput, and fabrication yield requires innovative solutions. In response to these demands, heterogeneous integration (HI) with 2.5D/3D packaging has emerged as a transformative strategy. These groundbreaking approaches involve both horizontal and vertical stacking of multiple dies within a single package or chip. The departure from the monolithic IC architecture seen in traditional SoC designs represents a paradigm shift in the semiconductor industry, with HI and 2.5D/3D packaging

offering innovative solutions to the challenges of latency, power consumption, throughput, and IC fabrication yield [1], [2]. The active exploration of these technologies is evident in the ongoing research and development efforts of various industry players. Notable initiatives include TSMC's 3DFabric™, which focuses on 3D silicon stacking and advanced packaging technologies [3], and Samsung's 3D-TSV (12 layers) DRAM Chip [4]. The establishment of ubiquitous interconnect standards, e.g., Universal Chiplet Interconnect Express (UCIe) [5], at the package level has facilitated die-to-die communication in these advanced packaging approaches.

Unfortunately, the globalization of the semiconductor supply chain exposes critical vulnerabilities, including untrusted electronic products, counterfeit ICs, and devices compromised with hardware Trojans. These threats stem from malicious third-party IP vendors, untrusted manufacturing facilities, rogue distributors, and various untrusted entities in the supply chain. Bloomberg notably reported hardware hacks in 2018 and 2021, involving covertly placed extra tiny chips on boards capable of compromising sensitive data from US companies [6], [7]. While the reported hacks targeted pre-heterogeneous integration (HI) hardware, the potential for adversaries to execute similar attacks by incorporating malicious die(s) or chiplets inside 2.5D/3D packages persist. Compromised hardware, whether unsecured or unreliable, creates opportunities for mounting software attacks. This includes exploiting firmware and software vulnerabilities to gain unauthorized access to the device or system by bypassing existing security measures implemented at the software level.

The security community continually devises diverse solutions to combat these intricate hardware security threats [8], [9]. Unfortunately, the universal adoption of these solutions remains elusive. Firstly, establishing integrity is paramount to ensure that all participating entities can rely on the supply chain. Secondly, providing incentives becomes imperative to encourage entities in the supply chain to implement robust security measures. Lastly, the assessment of trust and security measures must be based on observable and quantifiable metrics. For the industry to effectively address these challenges, there is a pressing need for the adoption of observability and traceability in the supply chain, aligning these measures with business interests.

The widespread interest in incorporating blockchain to ensure supply chain provenance is evident across academia, industry, and government. This attention is driven by the inherent properties and features of blockchain, which have the potential to improve the traceability, transparency, and reliability of the supply chain [10]–[16]. However, while numerous challenges within the realm of supply chain have been addressed, the adoption of these solutions has yet to be accepted for electronic parts. A crucial question persists regarding the motivation of an entity in the electronics supply chain to participate in these blockchain initiatives. The underlying incentive or reward-based mechanism for motivating participation and ensuring trustworthiness in the blockchain framework still remains widely unexplored. The key challenge lies in establishing trust between blockchains run by separate, autonomous *consortiums*. Areas of concern include data quality (transactions may be falsified or omitted), lack of transparency into the KYC process (sybil identities may exist), and trustworthiness of miners. Addressing these concerns is essential for relevant entities to perceive the added overhead of blockchain-based supply chain provenance as beneficial and worthwhile.

This paper provides a comprehensive solution to these challenges by presenting a reward-based blockchain system designed to facilitate the seamless movement of chiplets and ICs across interconnected consortium blockchains. The proposed infrastructure introduces a novel approach, leveraging a reputation-based system to instill trustworthiness in the chain and serve as a motivation for organizations to take part in our framework. Notably, when the nodes consistently exhibit trustworthy behavior, they are rewarded which reinforces our previous point of incentive-based blockchain architecture. Furthermore, our architecture extends beyond the individual consortium chains, establishing a solid foundation of trust that spans across interconnected networks, whether the chiplet fabrication is located on-shore(trusted) or off-shore(untrusted).

Contributions.

- *Reputation Scheme:* We propose a novel multi-chain reputation scheme specifically tailored for real-world IC supply chains. Our scheme is designed to operate with independent blockchain consortiums, reflecting real-world business and administrative trust domains where a unified trust metric may not be acceptable due to a lack of mutual trust. In our dual-layer approach, the lower layer enables IC supply chain provenance by facilitating transactions between consortium-run blockchains, while the upper layer dynamically assigns reputation scores to entities using an additive increase and multiplicative decrease approach. ***To the best of our knowledge, we are the first to introduce a comprehensive method for assigning reputation scores to blockchain members by rigorously evaluating their performance as they engage in transactions across various stages of the semiconductor supply chain.*** This approach not only enhances the accountability of individual participants but also strengthens the overall resilience and credibility of

blockchain-based systems in the semiconductor supply chain by promoting ethical behavior and reducing the risk of malicious activities. Our scheme, therefore, represents a significant advancement in managing trust and reputation within decentralized environments.

- *Broader Blockchain Adoption:* One of the significant challenges hindering the adoption of blockchain in the semiconductor supply chain is the lack of sufficient incentives for participating entities. Our proposed blockchain infrastructure is specifically designed to address this issue by offering tangible rewards to members who maintain high reputation scores. Entities with a strong reputation will benefit from easier and more efficient sales processes because their trustworthiness is recognized across the network, and thus reducing transaction friction. The system is not only scalable, as it allows seamless growth as more members join, but it is also structured to be user-friendly, making integration into the consortium straightforward for new participants. This combination of incentives, ease of entry, and scalability ensures that the blockchain network remains robust, efficient, and attractive to a broad spectrum of stakeholders.
- *Simulation Platform for Semiconductor Supply Chain:* We have developed a comprehensive simulation platform specifically designed for the complex semiconductor supply chain. ***To the best of our knowledge, this is the first instance where a directed acyclic graph (DAG) has been proposed as a model for representing the intricate dependencies and processes within the supply chain.*** By introducing this novel approach, we aim to provide the research community with a powerful tool to facilitate a deeper understanding of semiconductor supply chain dynamics. We believe this platform will enable more accurate modeling, offer opportunities for optimizing various supply chain-related applications, and allow researchers to test ideas without needing to implement them in real-world supply chains, which is often impractical in academic settings.

The rest of the paper is organized as follows. Section II covers prior works on using blockchain for supply chain provenance. In Section III, we present our proposed blockchain architecture for supply chain provenance, including the chiplet ecosystem. We present our proposed reputation scheme in Section IV. We analyze the effectiveness of our proposed framework in Section V. In Section VI, we discuss the strengths and limitations of our proposed scheme and highlight potential areas for future research. Finally, we conclude the paper in Section VII.

II. RELATED WORK

Blockchain-based reputation schemes have been widely explored in research for diverse applications ranging from e-commerce [17] and crowdsourced ratings/reviews [18] to quality control for IoT data [19] and reputation-aware routing for satellite constellations [20]. Reputation calculation methods are typically tailored to suit the properties and require-

ments of specific applications, such as scale of deployment, computational capacity of nodes, and the availability and trustworthiness of reputation “signals”. Bellini et al. [21] provides an excellent survey on the subject.

Blockchain research has been actively carried out across different supply chains, such as pharmaceutical industries [22], [23], agriculture [24] and the clothing industry [25], [26], addressing both security and management challenges to streamline the process of the movement of assets across the supply chain. Few research on blockchains has stepped foot onto the realm of vendor-managed inventory [27], [28], where the vendor observes the end users’ supply and sales and makes decisions regarding stocking of supplies accordingly.

Over the years, much work has been extended to implement a blockchain framework to enhance security in IoT devices and the semiconductor supply chain. Guin et al. [15] proposed a blockchain-enabled framework that assured the authenticity of devices using an unclonable identifier (ID) generated from an SRAM PUF. Cui et al. proposed a confirmation-based ownership transfer of devices across the chain using Hyperledger fabric [11], while Zhong et al. proposed a modular framework, utilizing many of the functions proposed in [11], for protection of trade secrets across different supply chain lifecycles. Additionally, to keep up with Moore’s Law, the introduction of 3D ICs has exposed the supply chain to several attacks at different stages of die manufacturing and assembly due to the nature of decentralized production. Calzada et al. have proposed a blockchain framework to establish authentication of a SiP throughout its lifecycle using a Chiplet hardware Security Module [29].

While our paper draws inspiration from the three research areas summarized above, to the best of our knowledge, we are the first to introduce a blockchain-based reputation framework for semiconductor supply chains to support inter-operation between blockchains controlled by independent consortiums.

III. BLOCKCHAIN FOR SUPPLY CHAIN PROVENANCE

The widespread interest in incorporating blockchain to ensure supply chain provenance is evident across academia, industry, and government. This attention is driven by the inherent properties and features of blockchain, which have the potential to greatly improve the traceability, transparency, and reliability of the supply chain [10]–[16]. Guin et al. presented a traceability framework that utilizes a layered, scalable, and permissioned blockchain [10], [11]. Tailored for versatility across diverse industries, this framework automates the tracking and traceability of electronic parts and systems from the design phase to end-of-life. The proposed infrastructure guarantees not only traceability but also safeguards privacy, protects trade secrets, and ensures the integrity of the electronic parts. In this section, we delve into a comprehensive analysis outlining the steps and considerations involved in constructing a provenance framework for chiplets manufactured on a global scale.

Table I summarizes the various entities and their roles in the electronics supply chain. The above-mentioned entities

TABLE I: PARTICIPATING ENTITIES IN THE BLOCKCHAIN-BASED TRACEABILITY FRAMEWORK.

Entity	Description
Chiplet Manufacturer	Fabricates chiplets. Chiplet manufacturers can directly send chiplets to the IC manufacturers. However, they can send chiplets to their authorized distributors as well.
Chiplet Distributor	Distributes chiplets and can be untrusted.
IC Manufacturer/ Foundry/Fab	Owns a foundry and fabricates ICs. Most IC design houses typically do not possess their own foundry but instead contract the fabrication of ICs to a fabrication facility (fab). Note that design houses can be excluded as they are not directly involved in handling chiplets.
IC Distributor	Distributes ICs, and can be untrusted.
System Integrator (SI)	Designs electronic systems. SIs use ICs to build complex systems.
End User (EU)	Plays a crucial role in utilizing ICs for the construction and maintenance of complex systems. SIs can function as a category of EUs. Note that EUs may not actively participate in the blockchain framework and can serve as an off-chain entity.

represent the design, manufacturing, integration, assembly, and distribution phases in the supply chain. We typically consider end users and system integrators as trusted entities, while distributors may fall into both trusted and untrusted categories. Adversaries have the potential to create cloned devices, integrate recycled or used devices, or introduce tampered devices with hardware Trojans or malware into ICs.

Figure 1 shows the overall provenance of chiplet and IC ecosystem. Chiplet and IC manufacturers have the exclusive authority to register device types (chiplets or ICs) on the blockchain. Each device type must undergo separate registration within the blockchain, a process facilitated by the use of smart contracts. IC manufacturers utilize the `ICCreateReg()` function to formally register new IC types. For instance, TSMC may create an entry for the Qualcomm Snapdragon 800 Processor type to track all the processors manufactured. On the other hand, chiplet manufacturers are restricted to registering only their chiplet types using `ChipletCreateReg()` function. The creation and registration of these device types are permanently recorded on the blockchain and visible to all downstream chain members and major identified upstream participants who require notifications. It is important to note that only chiplet and IC manufacturers possess the capability to register device types within the blockchain system. Other participants, such as distributors, system integrators, or end users, are not permitted to register device types due to the specified blockchain policy.

After registering device types in the blockchain, chiplet or IC manufacturers proceed to register individual devices for traceability. To ensure traceability, a unique device ID is essential, achievable through the integration of an ECID [30], PUF [31]–[34], or another unique identification method. Rather than directly placing the ID in the blockchain, we advocate storing the hash of the ID for enhanced security. This approach prevents the determination of the original ID unless one possesses the actual devices. While all members

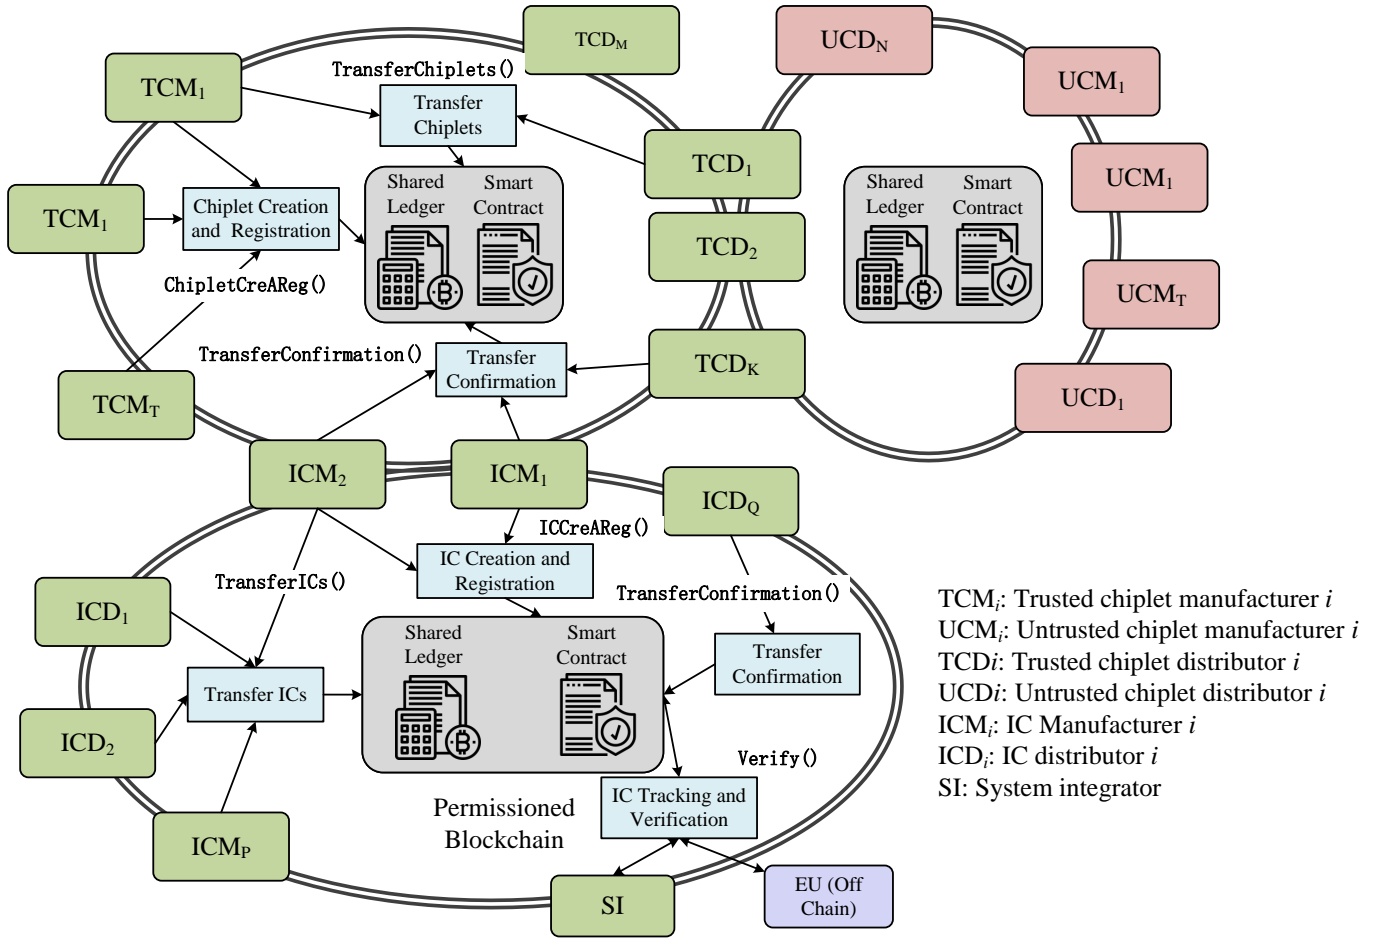


Figure 1: Proposed blockchain-based traceability framework for chiplet and IC supply chain provenance.

are aware of the uploaded hashes and the number of entries, the use of hashing ensures that actual IDs are not disclosed.

It is imperative to record the transfer of devices (e.g., chiplets and ICs) among different entities in the supply chain to ensure traceability. This can be achieved through the utilization of **TransferICs()** and **TransferChiplets()**. Consider the following example: Entity X intends to transfer a quantity N_1 of chiplets to Entity Y from its possession of N chiplets. Entity X initiates a device transfer transaction function with **TransferChiplets(chiplet_name, N_1 , {ID_{N1}}, {SP_{N1}}Y)**. It is important to note that, depending on the implementation, transferring N_1 entries in the blockchain may involve N_1 transactions, each denoting the transfer of one chiplet or IC, containing one hashed ID, or alternatively, a single transaction (or several) could include all the hashed IDs. It is permissible for manufacturers and distributors to initiate transactions for device transfers, provided they possess a specific quantity of devices. However, it is crucial to understand that the actual ownership of the device declared to be transferred in the device transfer transaction does not transition to the new owner until a confirmation transaction is received to address in-transit thefts.

When an entity (such as a manufacturer or distributor) delivers a specified quantity of electronic devices to a new

owner, the recipient must initiate a confirmation transaction. The device transfer process remains incomplete and unverified until the confirmation of the transfer is officially acknowledged. The trace and ownership of the device are transferred within the smart contract only after the confirmation has been successfully processed. The confirmation transaction function, denoted as **TransferConfirmation()**, receives the following inputs: (chiplet_name/IC_name, N_1 , {ID_{N1}}). The invoking of the confirmation process, indicates a successful transaction.

Upon physically receiving a device (IC/chiplet), participants are obligated to verify its identity (ID) using **Verify()**, which is present (hashed) in the blockchain. The verification process necessitates retrieving the unique device ID, accessible through the JTAG interface [35] or similar methods. If the ID is not found in the system, a flag will be raised, marking the device as suspicious. It is important to note that the verification and tracking procedure does not modify the data stored in the blockchain. Consequently, no actual transaction occurs, rendering the entire process highly efficient.

IV. MULTI-CHAIN REPUTATION SCHEME

This section presents a comprehensive design of a multi-chain reputation scheme, which builds upon the blockchain traceability framework introduced earlier in Section III. To

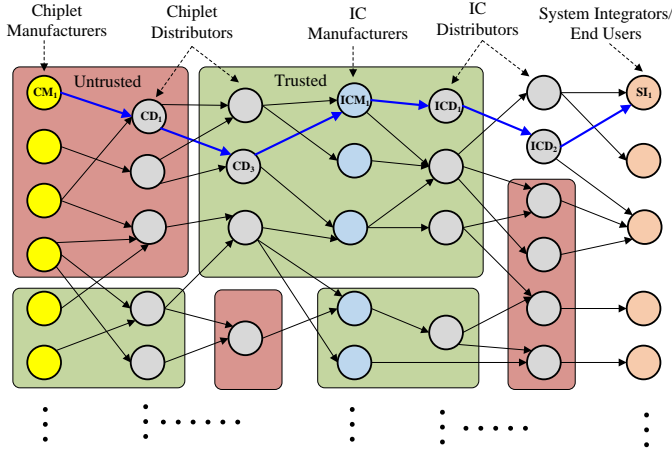


Figure 2: Semiconductor supply chain represented as a DAG.

begin with, it is crucial to recognize that trust and reputation are inherently subjective concepts, which differ significantly from the objective nature of provenance. For instance, one blockchain (B_X), might perceive another blockchain (B_Y) as untrustworthy based on its criteria and experiences, while blockchain B_Y could hold a similarly skeptical view of blockchain B_X . This subjectivity underscores the complexity of establishing a universally accepted reputation system across multiple blockchains. Given this, our subsequent discussion on reputation is explicitly framed from the perspective of the blockchain of which the System Integrator (SI) or End User (EU) is a part. This approach allows us to develop a reputation scheme that is adaptable to the diverse and decentralized nature of multi-chain environments, ensuring that trust can be managed in a way that reflects the unique characteristics and requirements of each blockchain network involved.

A. Directed Acyclic Graph (DAG) of Supply Chain

We have proposed modeling the semiconductor supply chain with a directed acyclic graph (DAG) to represent the intricate dependencies within the supply chain for computing the trust score, as shown in Figure 2. The arrows represent the connection between the supply chain entities as parts travel from seller to buyer. For example, a chiplet travels from its manufacturer (CM_1) through many distributors (CD_1 , and CD_3), before finally reaching an IC manufacturer (ICM_1), where it is consumed into an IC. The chips then travel across different distributors (ICD_1 and ICD_2) and finally are integrated into the systems by SI_1 . Note that the blue arrows trace the provenance path of a single chiplet produced by CM_1 , integrated into an IC by ICM_1 , and finally used by SI_1 . In addition, the background colors indicate whether the corresponding blockchain is trusted (green) or untrusted (red). A key goal of our scheme is to establish a trust measure applicable to ICs whose provenance graph crosses trust boundaries similar to the one shown.

B. Trusted Authority (TA)

We assume the presence of one or more Trusted Authorities (TA) within each blockchain network. A TA can be a govern-

mental agency or a group consisting of one or more members of the blockchain consortium, designated to make binding and final decisions on reports. We do not define a specific method for selecting TAs; instead, this decision is left to the discretion of the blockchain consortium members. TAs are responsible for evaluating reports submitted by part (i.e., chiplets or ICs) manufacturers, and SIs or EUs. Upon receiving a report, one of the TAs validates the authenticity of the report that contains information of chiplets or ICs that are faulty, counterfeit, or compromised. We can collectively refer to it as “defective”.

C. Chiplet and IC Life Cycles

Our reputation scheme differentiates between the life cycles of chiplets and ICs. The life cycle of a chiplet begins with its manufacturer and ends with an IC manufacturer, whereas an IC’s life cycle starts at its manufacturer and ends with a SI or EU. There are two primary reasons for this separation. Firstly, we conservatively wait until a verification process is completed before we reward or penalize the reputation of the entities involved. These verification processes occur at IC manufacturers for chiplets and at SIs/EUs for ICs. IC manufacturers are expected to check the quality of chiplets before incorporating them into IC fabrication. Similarly, SIs/EUs verify ICs through intensive testing or manual use. Based on these findings, we can reward or penalize the chiplet or IC manufacturers and their distributors. The second reason is that the time interval between a chiplet’s fabrication and its incorporation into an IC, which then reaches the SI/EU, can span months or even years. Dividing this interval into two stages helps shorten the feedback delay between transactions and reputation updates.

D. Cross-blockchain transactions

Transactions that cross trust boundaries carry high risk. Our reputation scheme assigns this risk to the buyer who is responsible for a product to cross from an untrusted blockchain (where the seller belongs) to the trusted blockchain. We achieve this by considering the buyer as a member of the untrusted blockchain for the purpose of calculating reputation penalties associated with such transactions. The exact mechanism of how this is done is described in Section IV-E.

Furthermore, to identify chains that frequently sell defective products, it is useful to assign *chain reputation* for each blockchain that a given chain transacts with. We achieve this by introducing *meta-entities* that signify cross-blockchain trust. Let UB and TB be an untrusted and trusted blockchain, respectively. Then, the meta-entity X_{TB}^{UB} acts as a logical intermediary for transactions between a seller (S) in UB and a buyer (B) in TB. Thus, a transaction (S, B, amt) is split into (S, X_{TB}^{UB}, amt) and (X_{TB}^{UB}, B, amt). The reputation calculation functions consider meta-entities like any other entity, and the reputation of X_{TB}^{UB} can be used for purposes such as making administrative decisions.

E. Reputation Scheme

Our proposed reputation scheme follows two basic principles:

- 1) *Purchase equals endorsement*: We consider a sale by entity S to entity B for an amount of d dollars as an endorsement of S by B for a magnitude of d units.
- 2) *Reward conservatively, penalize liberally*: We base the reputation gain/loss calculations based on the additive increase, multiplicative decrease (AIMD) approach.

The reputation update procedure begins upon the receipt of a chiplet or IC by an IC manufacturer or SI/EU (hereafter collectively termed as *verifier*) respectively. After inspecting and running checks on the product, the verifier themselves, or through a third party, invokes the `Report()` smart contract call. `Report()` takes a binary result flag as input which denotes the outcome of the verification process, with 0=pass and 1=fail. Naturally, the former results in reputation gain, whereas the latter potentially causes reputation loss. We now separately describe each process in detail.

1) *Reputation Gain*: If `result = 0`, the `Report()` call internally invokes `RewardReputation()` which is responsible for distributing the reputation rewards. To do this, `RewardReputation()` first walks through the DAG (see Figure 2) and retrieves the trace associated with the product. Each edge is represented as a 3-tuple of the form (buyer, seller, amount). Then, for each edge, the seller is rewarded with a reputation gain which is equal to the value of amount and shown as follows:

$$\text{seller.reputation} += \text{cur_convert}(\text{amount}) \quad (1)$$

where, `cur_convert` is a utility function that converts the sale amount to a standard currency for uniformity.

We choose cost as the positive reward metric because it is both simple and objective, while effectively reflecting real-world business and economic dynamics. Higher IC complexity or demand typically leads to increased per-chiplet/IC costs, while deeper supply chains drive costs up due to cumulative markups. Although alternative metrics may be useful in specific contexts, we consider sale amount to be the most broadly applicable and easily quantifiable metric for our framework.

2) *Reputation Loss*: If `result = 1`, the `Report()` call flags the transaction for review by the TA. Note that the procedure by which the TA validates or rejects a report is beyond the scope of this work. Instead, we consider only the outcome of this process (which is recorded as a transaction on the blockchain), namely, zero or more chiplets/ICs judged to be defective. Although zero chiplets/ICs judged as defective implies that the report was rejected by the TA, it is nevertheless recorded in the blockchain so that IC manufacturers and SI/EUs that frequently raise false alarms can be identified.

The reputation reduction procedure works on a per-defective chiplet/IC basis, i.e., `PenalizeReputation()` is invoked for each chiplet/IC found to be defective. The penalization process for a particular chiplet/IC starts from the manufacturer and traverses the provenance path to the IC manufacturer or SI/EU. For each seller along the path, the fraction of reputation to be reduced is determined based on a *multiplicative decrease parameter* m , as well as a *discounting parameter* d . The

multiplicative decrease parameter $m > 1$ is a global parameter defined at the blockchain level that determines the base reputation penalty rate. The reputation of the chiplet/IC manufacturer is reduced by a factor of m , and described as follows:

$$\text{manufacturer.reputation} /= m \quad (2)$$

The discounting parameter is a per-blockchain parameter used to calculate transitive trust penalties for the remaining sellers along the provenance path. Its value depends on whether the blockchain in which a given transaction took place is trusted (d_t) or untrusted (d_{ut}). d_{ut} is set to 1 so that every member of the sub-path in the untrusted blockchain is penalized equally. This is a necessity since untrusted blockchains may have sybils, i.e., a single entity assuming multiple identities. Furthermore, transactions crossing from an untrusted blockchain to a trusted blockchain also use d_{ut} so that the buyer responsible is penalized at a higher rate to account for the increased risk they introduce by allowing parts to cross blockchain trust boundaries (as discussed in Section IV-D). d_t can be set higher than 1 so that entities appearing later in a provenance path are penalized less. Thus, the reputation of the remaining sellers is reduced according to the following expression:

$$\text{seller.reputation} /= \text{seller.penalty} \quad (3)$$

where,

$$\text{seller.penalty} = \text{parent_txn.seller.penalty}/d \quad (4)$$

To make the above concepts clear, let us consider Fig. 2 where a chiplet produced by CM_1 was judged defective following a report by SI/EU. By inspecting the blue arrows, we can see that CM_1 , CD_1 and CD_3 are penalized equally, whereas ICM_1 , ICD_1 , ICD_2 are penalized with a discounting factor of d_t (denoted as d for brevity). Hence, the reputations of CM_1 , CD_1 and CD_3 are slashed by a factor of m , and the reputations of ICM_1 , ICD_1 and ICD_2 are slashed by m/d and m/d^2 and m/d^3 , respectively.

The above example also illustrates the effect of cross-blockchain transactions: for the purpose of calculating the penalty, CD_3 is essentially considered as a member of the untrusted blockchain (red) even though it actually belongs to the trusted blockchain (green). This is done to emphasize the risk assigned to the entity that is responsible for transacting with the untrusted blockchain. This way, the reputation scheme can capture the high risk of such transactions while keeping the computations simple.

F. Normalized Reputation Score

It is important to note that our reputation scheme assigns absolute reputation scores to entities, reflecting their overall trustworthiness based on their past actions and interactions. However, moving to a normalized reputation score is essential, which could provide additional benefits. For example, a normalized score within the range $[0, 1]$ could offer a more standardized and easily comparable metric across different entities in the supply chain.

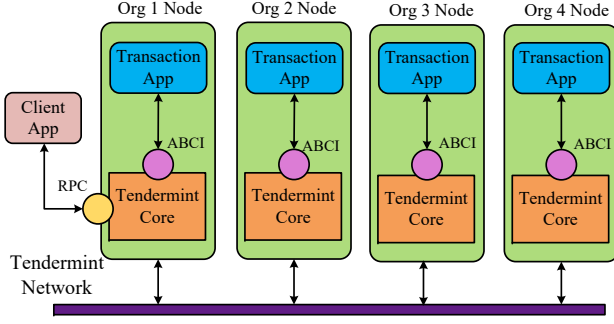


Figure 3: Proposed blockchain-based framework implemented with Tendermint.

Let r be the absolute reputation score of an entity e , and r' be its *ideal* absolute reputation score, i.e., what r would have been if e sold no defective parts. Clearly, $r' \geq r$.

We define the normalized reputation score of e as

$$\begin{cases} 1 & r' = 0 \\ \frac{r}{r'} & r' > 0 \end{cases} \quad (5)$$

Keeping track of the normalized reputation score of an entity is efficient in terms of both computation and memory complexity, since it only requires $O(1)$ additional space and $O(n)$ additional reward computations (where n is the number of transactions).

V. EVALUATION

In this section, we demonstrate our proposed reputation scheme with a reference blockchain implementation. We follow this up with a detailed analysis of the increase or decrease of reputation scores across participating members based on their performance in the chain.

A. Blockchain Implementation

We have implemented our proposed framework using Tendermint [36]. However, it can be extended to other permissioned blockchain frameworks such as Hyperledger ‘fabric’ [37]. Figure 3 demonstrates an overall representation of our Tendermint architecture. The Tendermint Core performs Byzantine fault tolerant (BFT) state machine replication, which is used to create and manage the blockchain network. Tendermint’s ABCI interface accounts for a more modular structure, facilitating the communication between the blockchain and the application’s operational logic. This ultimately ensures that the application can process transactions without interfering with the underlying consensus mechanism.

Our Tendermint application consists of a Trusted IC supply chain and a Python script that acts as the client application for sending and receiving transactions to and from the blockchain server. The client application ensures that the data sent over to the blockchain is formatted correctly and the received data is interpreted according to our requirements. The transactions are submitted to the blockchain network via HTTP requests

to the Tendermint node’s RPC interface [38]. Each transaction is appended with the user’s unique organization ID and their specific role within the supply chain, such as chiplet/IC manufacturer, distributor, or system integrator. As we have demonstrated a proof of concept of our framework, in a real-world implementation, we anticipate that users will be equipped with identity certificates that authenticate their transactions with unique identities and roles within the chain. These are then used to verify access control before invoking certain operations in our operational logic. Once the transaction undergoes the consensus mechanism and is processed according to our Transaction Application logic, as shown in Figure 1, it is stored immutably in our ledger. The implementation of the Report() function further populates our ledger with the reputation scores of individual organizations of the blockchain consortiums that can later be queried by organizations to make an informed decision before transacting with the respective organizations. To demonstrate a multi-computer network set-up, we have deployed our Tendermint nodes in 4 docker containers. These validator nodes are configured to connect to their replicated transaction applications, which also run in their own Docker containers.

B. Basic Simulation

The objectives of our initial evaluation were twofold: (1) to analyze the impact of different system parameters on the proposed reputation scheme, and (2) to assess the suitability of the scheme for real-world scenarios. Among the two system parameters, we focused on the multiplicative decrease parameter m , as the choice of discounting factor d_{ut} is more of an administrative decision. Towards this goal, we simulated how the reputation of a chiplet manufacturer evolves with the number of transactions, considering different values of *defect probability* d (i.e., the likelihood that a chiplet produced by the manufacturer is defective) and the multiplicative decrease factor m . The results of this simulation are shown in Figure 4.

Several key observations can be drawn from the figure. First, entities with very low defect probabilities maintain a normalized reputation close to 1, even after 10^6 transactions. However, this value is influenced by m – a higher m leads to a more rapid decline in normalized reputation for the same defect probability. For instance, for $m = 0.001$, an entity with a defect probability of 0.00001 experiences almost no reputation loss over the course of the simulation. However, when $m = 0.01$, there is a slight but noticeable decline in reputation, similar to what happens with a defect probability of 0.0001 when $m = 0.001$. *This suggests that m should be selected so that an entity with a defect probability no worse than the typical industry fabrication defect rate experiences almost no loss in normalized reputation.*

Secondly, it is evident that the higher the d value of an entity, the more rapidly its normalized reputation declines. For example, when $m = 0.01$, an entity with $d = 0.0001$ experiences a nearly linear decline in normalized reputation to 0.6 over 10^6 transactions. In contrast, an entity with $d = 0.01$ or higher is heavily penalized, with its normalized reputation

dropping sharply to nearly zero within a few thousand transactions. Since this observation aligns with the expected behavior of a reputation scheme, it qualitatively supports the suitability of our proposed scheme for real-world applications. Next, we further investigate this claim through an end-to-end simulation.

C. End-to-end Simulation

To validate the real-world applicability of the proposed reputation scheme, we developed a simulation platform for the entire semiconductor supply chain using Python. The platform consists of two components: a supply chain simulator and a reputation simulator. We now briefly describe each component and explain how they work together to achieve an end-to-end simulation.

1) *Supply Chain Simulator*: The purpose of the supply chain simulator is to generate a sequence of transactions that closely mimic the flow of parts in a real-world semiconductor supply chain. Figure 2 illustrates a directed acyclic graph (DAG) representing an example supply chain involving various chiplet and IC manufacturers, distributors, and system integrators, described in Section IV-A. The flow of parts between entities is represented in the simulation as a sequence of transactions in the following format:

$$T_i = \{PT, S_i, D_i, c, n, \{ID_1, ID_2, \dots, ID_n\}\} \quad (6)$$

where, T_i , PT , S_i , D_i , n , and ID represent i^{th} transaction, part type, source entity, destination entity, cost of part, number of parts and their IDs, respectively.

The supply chain simulator is designed with flexibility in mind, allowing users to configure various parameters such as the number of transactions, the number of entities of each type, the number of trusted and untrusted consortiums, resale markup percentages, and defect probability. While this simulator was developed as part of our end-to-end simulation pipeline, we hope that the research community finds it valuable for other studies involving the semiconductor supply chain.

2) *Reputation Simulator*: The sequence of transactions generated by the supply chain simulator is then fed into the reputation simulator, which tracks the evolution of both absolute and normalized reputations of entities with each transaction. By decoupling the reputation calculation from the supply chain simulation, we can test various reputation system parameters and compare the results for a specific supply chain instance.

We tested the end-to-end simulation pipeline using a sequence of 1 million simulated transactions within a supply chain consisting of 100 chiplet manufacturers, 1000 chiplet distributors, 100 IC manufacturers, and 500 IC distributors. The results are shown in Figure 5. As seen in the figure, the two trusted consortiums consistently outperform the untrusted consortiums in terms of normalized reputation. This suggests that our proposed reputation scheme is effective for real-world semiconductor supply chains. However, this analysis focuses on the overall reputation trends of consortiums, not on the behavior of *individual* entities that choose to engage in malicious activities. We address this question next.

D. Attack Resilience

A major threat to the integrity of reputation systems is “sleepers agents” which act benignly until they build up enough reputation, then launch attacks on the system. In poorly designed reputation schemes, these entities may maintain high reputations even after initiating attacks, by virtue of the reputation accrued beforehand. To evaluate the resilience of our proposed reputation scheme against such attacks, we designed a simulation experiment with three types of behaviors: benign, malicious, and benign-then-malicious. Entities exhibiting benign behavior are assumed to have a defect probability of $d = 0.001$. We consider two levels of malicious behavior, with defect probabilities of $d = 0.0015$ and $d = 0.002$. Benign-then-malicious entities behave benignly for the first 500K transactions, then switch to one of the two malicious behavior levels for the next 500K transactions. The results of this experiment are shown in Figure 6.

In Figure 6, the vertical dashed red line marks the point at which benign-then-malicious entities begin exhibiting malicious behavior. Up until that point, both their absolute and normalized reputations follow a similar trajectory to that of a benign entity. The key observation from the figure is that after a sufficient number of transactions following the onset of malicious behavior, entities exhibiting benign-then-malicious behavior arrive at roughly the same absolute and normalized reputation as if they had been engaging in malicious behavior from the beginning. Thus, we conclude that benign-then-malicious behavior offers no long-term advantage to entities that adopt it. A more detailed analytical investigation of this phenomenon is left for future work.

It is important to note that, while the simulation allows us to know whether a user is malicious (since we predefine them), a real-world deployed system lacks this knowledge. The only actionable metric available to such a system is the defect probability d . Our simulations demonstrate that when the defect probability increases, the reputation system eventually recovers. Moreover, the recovery time can be shortened by increasing the multiplicative decrease factor m , enabling the system to identify malicious entities more quickly. Together, these results highlight the system’s robustness against long-term risks and its ability to mitigate short-term risks through parameter tuning.

We are not aware of any real-world datasets currently available to evaluate the feasibility of our proposed scheme beyond the insights gained from our extensive simulations. Consequently, we leave the analysis using real-world datasets as a direction for future work.

VI. DISCUSSION

In the previous sections, we introduced and evaluated a novel reputation scheme for entities within the global semiconductor supply chain. Unlike most reputation schemes, our proposed approach allows for a unified reputation metric across multiple blockchain consortiums, while permitting variations in how the reputation scheme is implemented within each consortium. As a result, each consortium will have its own

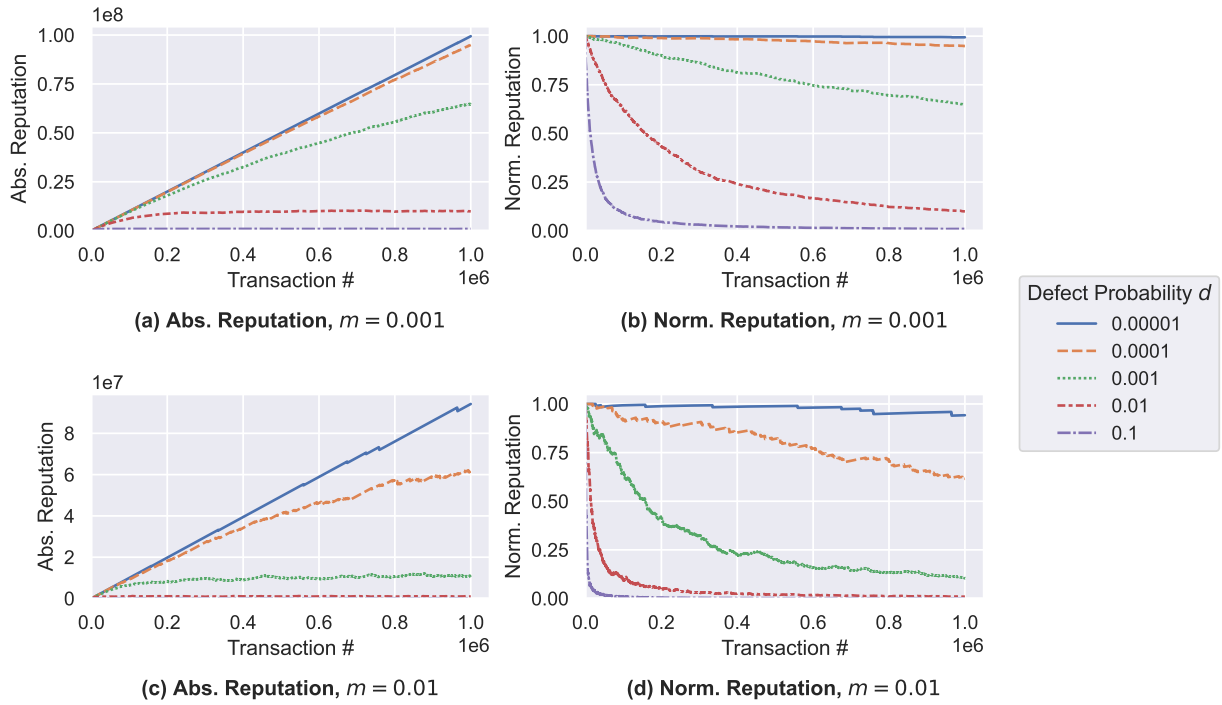


Figure 4: Evolution of absolute and normalized reputation with the number of transactions for different defect probabilities (d) and multiplicative decrease factors (m). Figure (a) & (c) show the absolute reputation that starts at 0 and increases/decreases according to Equations 1-4. Figure (b) & (d) show the normalized reputation that starts at 1.0 and can only decrease, as per Equation 5.

perspective on the reputation of a given entity. It is even possible that a Trusted Authority recognized in one consortium may not be regarded as trusted in another. While this approach might seem counterintuitive, we believe that such flexibility is essential for entities operating in different countries with diverse (and sometimes conflicting) goals to collaborate within a shared blockchain-based provenance system.

We argue that our scheme has a minimal entry barrier in terms of effort and cost, as each consortium is solely responsible for managing transactions involving its own members, and the reputation calculations are computationally lightweight. Moreover, once the infrastructure is established, new entities can integrate seamlessly and benefit from the system without requiring substantial investment.

It is worthwhile to briefly discuss how our proposed scheme handles some common challenges faced by reputation systems. One such challenge is the *cold-start* problem, where an entity that joins a reputation scheme long after its establishment struggles to catch up with entities with longer histories. Our scheme partially mitigates this issue by proposing a normalized reputation. Another challenge is *wash trading*, where a group of entities artificially inflates their sales and reputations by repeatedly buying and selling from each other. Our scheme does not directly address wash trading, and enhancing it to resist such manipulation is left for future work. A third issue is *buyer tampering*, where a buyer intentionally alters a purchased product to damage the reputation of the seller. While solutions like ECID and PUFs may help partially

mitigate this, this issue falls outside the scope of this work. We leave the evaluation of the resilience of our scheme against these and other sophisticated attacks for future research.

VII. CONCLUSION

In today's landscape of horizontally integrated semiconductor supply chains and the advent of 2.5D/3D ICs, it has become increasingly difficult for organizations to accurately assess the behavior and reliability of the organizations with which they are conducting their transactions. To mitigate this issue, we have first detailed a blockchain framework for the supply chain provenance of a 2.5/3D IC. Our paper then introduces a novel evaluation parameter that assigns reputation scores to entities based on their performance within the supply chain. This is accomplished through an additive increase and multiplicative decrease model, which dynamically adjusts reputation scores over time; good actors are rewarded with progressively increasing incentives, while bad actors face proportionate penalties, thereby maintaining a balanced and fair system of accountability. Our paper concludes with a threefold comprehensive analysis of the efficacy of our proposed reward scheme. First, we assess the performance of manufacturers under varying percentages of defective chips; second, we evaluate the performance of different participating entities within both trusted and untrusted consortiums; third, we compare the reputation scheme of entities as they behave maliciously at different points in the supply chain. Our proposed model provides a detailed assessment of the reliability

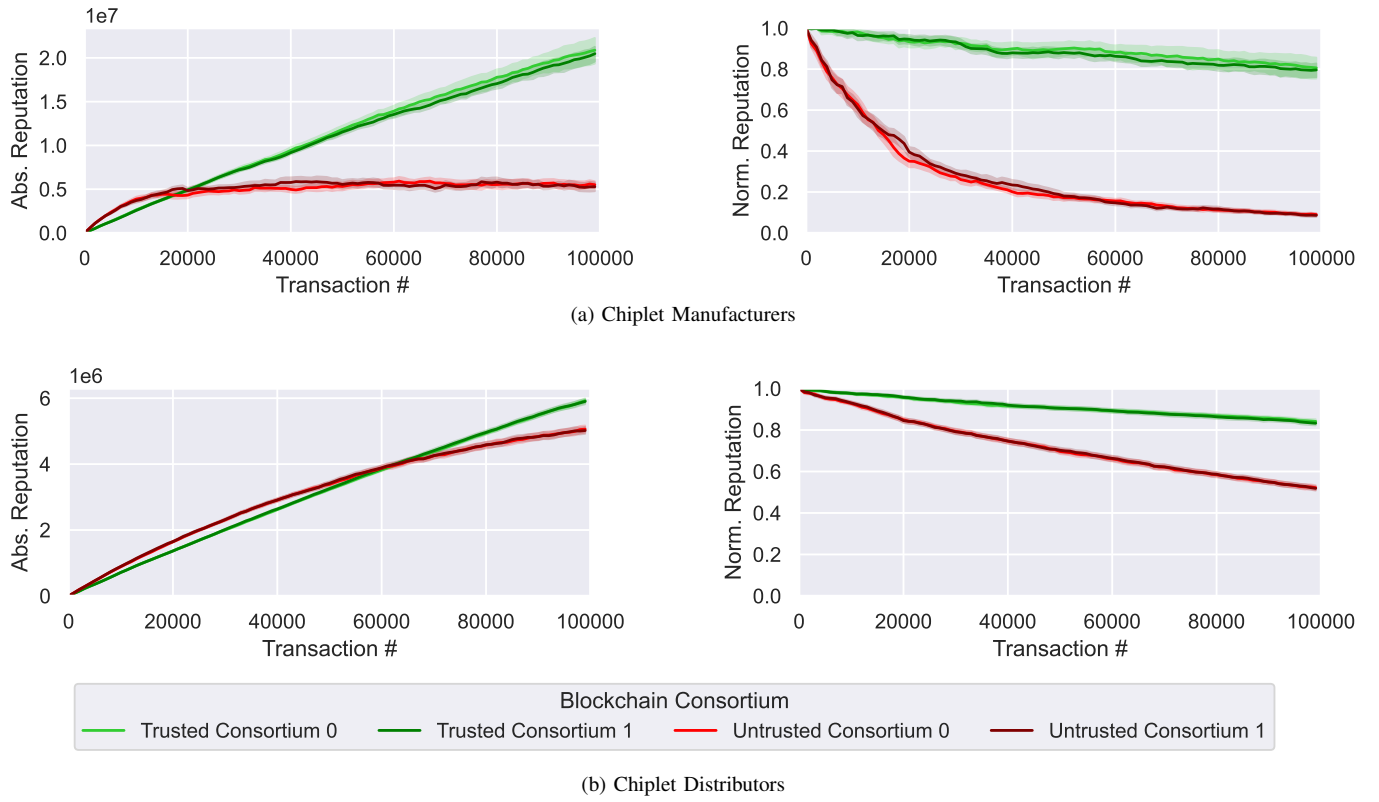


Figure 5: Evolution of absolute and normalized reputation of chiplet manufacturers and distributors, aggregated by consortiums, over an end-to-end simulation consisting of 10^6 transactions. Green and red lines indicate the mean reputation of trusted and untrusted consortiums, respectively. Shaded regions around the lines indicate the 95th percentile value. Similar trends are seen for IC manufacturers and distributors.




Figure 6: Comparison of benign, malicious, and benign-then-malicious behaviors. In this simulation, benign entities are assumed to have a defect probability of $d = 0.001$, while malicious entities have a defect probability of either $d = 0.0015$ or $d = 0.002$. A benign-then-malicious entity behaves benignly for the first 500K transactions (marked by the vertical red dashed line) before switching to malicious behavior.

of entities within a supply chain. Blockchain members can leverage this data for effective administrative decision-making, ensuring a more transparent and trustworthy supply chain ecosystem.

ACKNOWLEDGMENTS

This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-23-1-0312. Any opinions, findings, and conclusions, or recommendations expressed in this material are those of

the author(s) and do not necessarily reflect the views of the United States Air Force.

REFERENCES

- [1] N. Vashistha, M. L. Rahman, M. S. U. Haque, A. Uddin, M. S. U. I. Sami, A. M. Shuo, P. Calzada, F. Farahmandi, N. Asadizanjani, F. Rahman, and M. Tehranipoor, "Toshi-towards secure heterogeneous integration: Security risks, threat assessment, and assurance," *Cryptology ePrint Archive*, 2022.

- [2] Introducing TSMC 3DFabric: TSMC's Family of 3D Silicon Stacking, Advanced Packaging Technologies and Services, TSMC, 2020. <https://www.tsmc.com/english/news-events/blog-article-20200803>.
- [3] TSMC 3DFabric™. <https://3dfabric.tsmc.com>.
- [4] Samsung Electronics Develops Industry's First 12-Layer 3D-TSV Chip Packaging Technology, 2019.
- [5] UCIE™, Universal Chiplet Interconnect Express™, <https://www.uciexpress.org/>.
- [6] The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies, Bloomberg, 2018.
- [7] The Long Hack: How China Exploited a U.S. Tech Supplier, Bloomberg, 2021.
- [8] S. Bhunia and M. Tehranipoor, *Hardware security: a hands-on learning approach*. Morgan Kaufmann, 2018. ISBN: 0128124784.
- [9] M. Tehranipoor, U. Guin, and D. Forte, *Counterfeit Integrated Circuits: Detection and Avoidance*. Springer International Publishing, 2015.
- [10] Y. Zhong, A. Ebrahim, U. Guin, and V. Menon, "A modular blockchain framework for enabling supply chain provenance," in *IEEE Physical Assurance and Inspection of Electronics (PAINE)*, pp. 1–7, 2023.
- [11] P. Cui, J. Dixon, U. Guin, and D. DiMase, "A blockchain-based framework for supply chain provenance," *IEEE Access*, vol. 7, pp. 157113–157125, 2019.
- [12] X. Xu, F. Rahman, B. Shakya, A. Vassilev, D. Forte, and M. Tehranipoor, "Electronics supply chain integrity enabled by blockchain," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 24, no. 3, pp. 1–25, 2019.
- [13] M. Crosby, P. Pattanayak, S. Verma, V. Kalyanaraman, et al., "Blockchain technology: Beyond bitcoin," *Applied Innovation*, vol. 2, no. 6-10, p. 71, 2016.
- [14] M. Pilkington, "11 blockchain technology: principles and applications," *Research handbook on digital transformations*, vol. 225, 2016.
- [15] U. Guin, P. Cui, and A. Skjellum, "Ensuring Proof-of-Authenticity of IoT Edge Devices using Blockchain Technology," in *IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pp. 1042–1049, 2018.
- [16] M. N. Islam and S. Kundu, "Enabling ic traceability via blockchain pegged to embedded puf," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 24, no. 3, p. 36, 2019.
- [17] Y. Sun, R. Xue, R. Zhang, Q. Su, and S. Gao, "Rtchain: A reputation system with transaction and consensus incentives for e-commerce blockchain," *ACM Trans. Internet Technol.*, vol. 21, dec 2020.
- [18] S. Thottungal Valapu, T. Sarkar, J. Coleman, A. Avyukt, H. Embrechts, D. Torfs, M. Minelli, and B. Krishnamachari, "DARSAN: A Decentralized Review System Suitable for NFT Marketplaces," in *Blockchain – ICBC 2023* (Q. Wang, J. Feng, and L.-J. Zhang, eds.), (Cham), pp. 3–20, Springer Nature Switzerland, 2023.
- [19] V. Dedeoglu, R. Jurdak, G. D. Putra, A. Dorri, and S. S. Kanhere, "A trust architecture for blockchain in iot," in *Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, MobiQuitous '19*, (New York, NY, USA), pp. 190–199, Association for Computing Machinery, 2020.
- [20] L. Clark, Y.-C. Tung, M. Clark, and L. Zapanta, "A blockchain-based reputation system for small satellite relay networks," in *2020 IEEE Aerospace Conference*, pp. 1–8, 2020.
- [21] E. Bellini, Y. Iraqi, and E. Damiani, "Blockchain-based distributed trust and reputation management systems: A survey," *IEEE Access*, vol. 8, pp. 21127–21151, 2020.
- [22] K. Abbas, M. Afaq, T. Ahmed Khan, and W.-C. Song, "A blockchain and machine learning-based drug supply chain management and recommendation system for smart pharmaceutical industry," *Electronics*, vol. 9, no. 5, p. 852, 2020.
- [23] S. Abdallah and N. Nizamuddin, "Blockchain-based solution for pharma supply chain industry," *Computers & Industrial Engineering*, vol. 177, p. 108997, 2023.
- [24] S. A. Bhat, N.-F. Huang, I. B. Sofi, and M. Sultan, "Agriculture-food supply chain management based on blockchain and iot: a narrative on enterprise blockchain interoperability," *Agriculture*, p. 40, 2021.
- [25] S. Guo, X. Sun, and H. K. Lam, "Applications of blockchain technology in sustainable fashion supply chains: Operational transparency and environmental efforts," *IEEE Transactions on Engineering Management*, vol. 70, no. 4, pp. 1312–1328, 2020.
- [26] B. Wang, W. Luo, A. Zhang, Z. Tian, and Z. Li, "Blockchain-enabled circular supply chain management: A system architecture for fast fashion," *Computers in Industry*, vol. 123, p. 103324, 2020.
- [27] I. A. Omar, R. Jayaraman, K. Salah, M. Debe, and M. Omar, "Enhancing vendor managed inventory supply chain operations using blockchain smart contracts," *IEEE access*, vol. 8, pp. 182704–182719, 2020.
- [28] T. Dasaklis and F. Casino, "Improving vendor-managed inventory strategy based on internet of things (iot) applications and blockchain technology," in *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pp. 50–55, IEEE, 2019.
- [29] P. E. Calzada, M. S. U. I. Sami, K. Z. Azar, F. Rahman, F. Farahmandi, and M. Tehranipoor, "Heterogeneous integration supply chain integrity through blockchain and chsm," *ACM Transactions on Design Automation of Electronic Systems*, vol. 29, no. 1, pp. 1–25, 2023.
- [30] Bill Eklow, "ECID vs Device ID," 2006.
- [31] B. Gassend, D. Clarke, M. Van Dijk, and S. Devadas, "Silicon physical random functions," in *Proceedings of ACM Conf. on Computer and Communications Sec. (CCS)*, ACM, 2002.
- [32] G. Suh and S. Devadas, "Physical Unclonable Functions for device authentication and secret key generation," in *Proc. of ACM/IEEE on Design Automation Conference*, 2007.
- [33] J. Guajardo, S. S. Kumar, G.-J. Schrijen, and P. Tuyls, "Fpga intrinsic pufs and their use for ip protection," in *International workshop on Cryptographic Hardware and Embedded Systems*, Springer, 2007.
- [34] W. Wang, A. Singh, U. Guin, and A. Chatterjee, "Exploiting power supply ramp rate for calibrating cell strength in SRAM PUFs," in *IEEE Latin-American Test Symposium*, 2018.
- [35] IEEE 1149.1-2013 - IEEE Standard for Test Access Port and Boundary-Scan Architecture, https://standards.ieee.org/standard/1149_1-2013.html.
- [36] E. Buchman, *Tendermint: Byzantine fault tolerance in the age of blockchains*. PhD thesis, University of Guelph, 2016.
- [37] E. Androuraki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, et al., "Hyperledger fabric: a distributed operating system for permissioned blockchains," in *Proceedings of the thirteenth EuroSys conference*, pp. 1–15, 2018.
- [38] D. Cason, E. Fynn, N. Milosevic, Z. Milosevic, E. Buchman, and F. Pedone, "The design, architecture and performance of the tendermint blockchain network," in *2021 40th International Symposium on Reliable Distributed Systems (SRDS)*, pp. 23–33, IEEE, 2021.