Control in Decentralized Industrial Organizations

Jay Bayne
jbayne@stratum4.org

Abstract

Blockchains offer improvements in the performance of distributed control systems regulating industrial processes, enabling coordination of production automation systems within and among decentralized industrial organizations (DIO). We introduce DIO requirements and a service-oriented architecture for distributed control in manufacturing ecosystems.

In industrial automation applications, blockchains based on *distributed ledger technology*[1] provide distributed, immutable and permissioned databases for the recording of monitoring and control transactions, preserving causal relationships and promoting coherence between and among cooperating production systems. As core architectural elements blockchains provide permanent cryptographically-secured records of sequences of measurements and subsequent control actions. Viewed as a multi-agent agreement protocol, blockchains help solve the challenging problem in realizing decentralized automation by guaranteeing consensus among federated peers, without need of a central authority, of shared knowledge concerning the unfolding states and behaviors of distributed physical and synthetic[2] *value production processes*.

Decentralized Control

In our service-oriented DIO model, intelligent, automated governance of quantifiable units of value production is the responsibility of designated *cyberspatial objects* (CSO, Figure 1) located in regions of *cyberspace* (Figure 2).

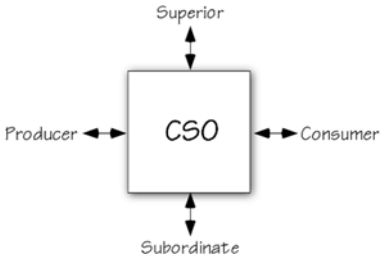

Figure 1 - Cyberspatial Object

Cyberspace is an abstract 4-dimensional space-time construct comprising time-dependent addresses for a CSO's *infospatial*, *geospatial*, and *sociospatial* service access points. Infospace is typically described in terms of Internet addresses (e.g., IPv4/IPv6). Geospace typically describes a CSO's physical location in earth-centered (e.g., GPS) coordinates of latitude, longitude, elevation. And sociospace identifies within a DIO's accountability hierarchy a CSO's authority over deployment of reusable assets (e.g., people, material, capital.)

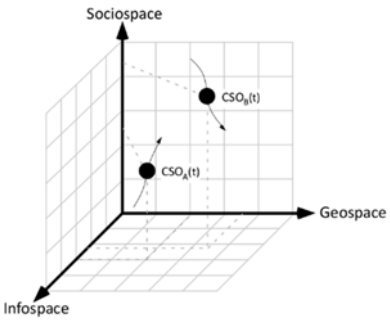

Figure 2 - Cyberspace

Within a DIO, CSO may be stationary or mobile in any of its cyberspatial dimensions. They have absolute and relative positions and if mobile have trajectories, velocities, and accelerations. A CSO's cyberspatial motion will contribute to its identity and permissions, requiring blockchains to record CSO positions as a function of time for cybersecurity purposes.

[1] NISTIR 8202 (DRAFT), Blockchain Technology Overview, https://csrc.nist.gov/publications/detail/nistir/8202/draft
[2] Objects are *synthetic* if realized by human effort (e.g., software) versus occurring naturally or synthesized organically.

As computational objects, CSO provide information processing services responsible for both *regulatory* and *supervisory* control of production processes under their authority, including interactions with collaborating CSO along their horizontal producer-consumer supply chains and along vertical superior-subordinate command chains.

To be aware, agile and reactive (i.e., *intelligent*), CSO subscribe to and concurrently process multiple input streams, recognizing and responding in *real-time*[3] to unfolding situations characterized by event *signatures*, patterns identified within and across the streams. As intelligent automation systems, CSO are considered *cyber-physical systems*[4] (CPS). Similarly, CSO embedded in and governing the behavior of intelligent web-connected devices are referred to variously as Internet *edge-systems,* "Internet of Things" (IoT) and "Industry 4.0" devices.

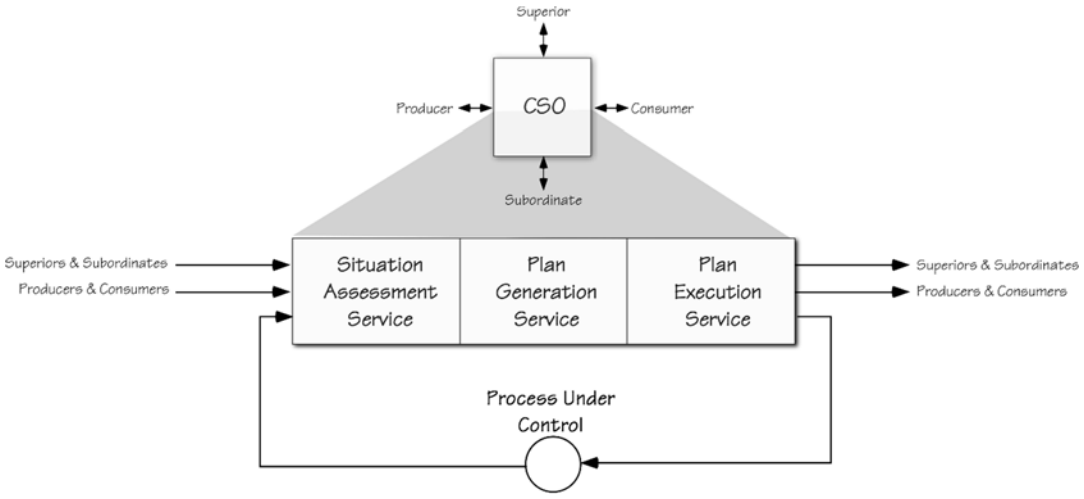

Figure 3 - CSO Stream Processing Pipeline

Figure 3 diagrams a CSO's data stream processing pipeline, emphasizing its three primary continuous and concurrently operating stages.

- The first, *Situation Assessment Service* (SAS), performs data acquisition, filtering, event detection (pattern recognition and machine learning[5]), data analytics, archiving and selection of feasible event response plans or *courses of action* (COA). SAS may also include modeling and simulation services employed to anticipate future situations and their respective COA.
- The second, *Plan Generation Service* (PGS), is responsible for converting COA into qualified *plans of action* (POA), where qualification means COA meet defined operating policies (i.e., "rules of engagement") and adhere to scheduling rules permitting coordinated sharing of reusable resources (e.g., task sequencing for optimal responds to high priority situations while avoiding resource deadlocks).
- And the third, *Plan Execution Service* (PES), converts POA into authorized *plans of record* (POR), sequences of executable tasks. PES services provide for accountability, task coordination and scheduling, task initiation and termination, and execution performance monitoring. PES effects on processes under control are fed back to SAS input to facilitate tracking, learning and adaptation.

[3] A system is considered *real-time* to the degree it meets its deadlines.
[4] See, for examples: https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503286
[5] https://www.gwern.net/docs/rl/2017-silver.pdf

DIO Control Structure

Effective DIO control strategies depend on intra and inter CSO coordination, a result of blockchain-enabled coherence and cooperation. With minimal "friction" collaborating CSO are permitted to link and synchronize their respective pipelines to form *federated systems*[6]. DIO federations require a shared governance framework. Figure 4 diagrams a CSO governance structure consistent with established architectures for object-oriented service systems[7], here enhanced by application of the Viable Systems Model[8,9], a *cybernetic*[10] framework employing blockchains along primary communications paths. This CSO model has the feature of being recursive, reusable, nested, scalable, and fractal in nature.

As shown, $CSO_{i,j,n}$ operates within DIO federation "i", at accountability level "n", and at supply chain position "j." *Supervisory* $CSO_{i,j,n}$ governs the behaviors of subordinate *regulatory* $CSO_{i,j,n-1}$, with each subordinate responsible for a specific embedded production process. The figure identifies $CSO_{i,j,n}$'s internal human or synthetic management *agents*, designated $E^p{}_n{}^q$, where index "p" denotes an agent's specific governance function (summarize in Table 1), and index "q," if present, identifies the agent's focus on one of the K_n subordinate value production processes.

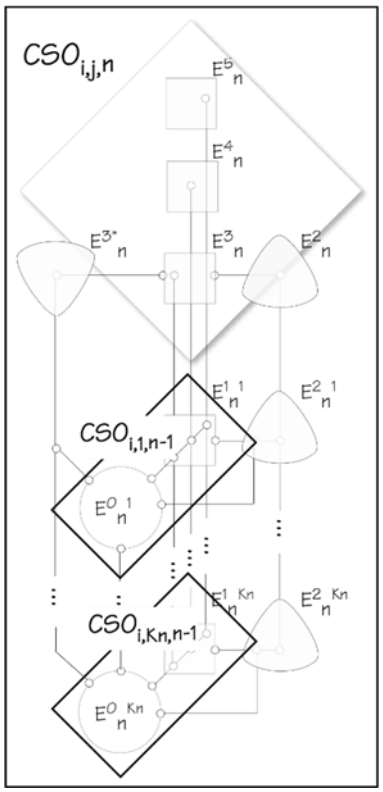

Figure 4 - CSO Governance Structure

Figure 5 is an example DIO accountability structure for an industrial enterprise (e.g., a manufacturer with multiple factories, a campus with multiple buildings, or a container shipping company operating multiple vessels.) $CSO_{1,1,4}$ might represent a business unit (P & L) with a single factory ($CSO_{1,1,3}$) having two automated production lines ($CSO_{1,1,2}$ and $CSO_{1,2,2}$), with each line containing two robot cells, $CSO_{1,1,1}$, $CSO_{1,2,1}$, and $CSO_{2,2,1}$, $CSO_{2,3,1}$, respectively.

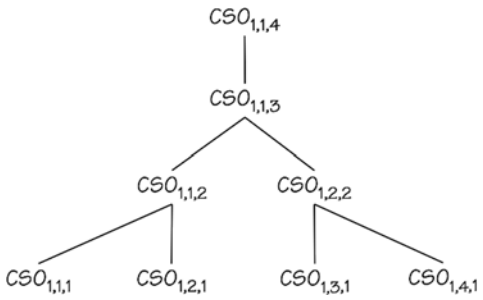

Figure 5 - DIO Accountability Graph

Blockchains provide for reliable permissioned P2P communication within and among DIO and their associated CSO, supporting secure transactions among their respective governance services. Included are chained transactions along Figure 4's central vertical "command axis" interconnecting CSOs into an operational chain of command. The supervisory E^5-E^4-E^3 triumvirate directly connects the E^3 Operator (e.g., a business unit manager) with its subordinate E^1 Process Directors (e.g., factory managers). The command axis provides for authorizing, issuing and acknowledging coordinated "tasking orders" governing CSO regulatory processes, subsequently returning operational status reports. The set of POA tasking orders is produced jointly by the E^5-E^4-E^3 triumvirate, with authorization by E^5, planning by E^4, and execution by E^3 and its E^1 subordinates.

---

[6] http://www.lewrockwell.com/vance/vance17.html, "Jeffersonian axioms of a free society"
[7] https://en.wikipedia.org/wiki/Service-oriented_architecture
[8] https://en.wikipedia.org/wiki/Viable_system_model
[9] https://github.com/rbcarleton/VSA
[10] http://www.ieeesmc.org/publications/transactions-on-smc-systems

Table 1 - CSO Internal Supervisory Agents

Label	Title	Agent Function
E^5	Executive	Highest level of authority, accountability
E^4	Navigator	Situation assessment, learning, adaptation and forward planning => (SAS, PGS)
E^3	Operator	Coordinated response plan execution => (PES)
E^{3*}	Auditor	Independent performance measurement and assessment
E^2	Coordinator	Shared resource management, scheduling and synchronization
E^1	Director	Production process supervisory control
E^0	Controller	Production process regulatory control

The E^3-E^2 loop provides E^3 Operators with the ability to observe and synchronize the concurrent behavior of multiple subordinate E^1 production processes, coordinating task completion timing (i.e., multi process rendezvous), access to shared resources and the handling of exceptions (e.g., "faults"). The E^2 loop also allows individual E^1 Directors to observe and self-synchronize operations and share resources in relation to other E^1 peers. Each E^0 Controller regulates a quantifiable unit of value production that is visible to external producers and consumers through a supply chain "service access points" (SAP). E^1 Directors govern value production coordinated internally by E^3 through asset chain SAP. In addition to the reporting received by E^3 Operators from their subordinate E^1 Directors, E^3 receives independent consolidated assessments from its E^{3*} audit agents tasked with observing performance of E^0 supply chain endpoints. As noted in Figure 4, this cybernetic model is recursive. A subordinate E^1 Director at operational level "n," identified by its 45 degree rotation, represents the E^5-E^4-E^3 triumvirate governing activities at the next lower level, "n-1".

This CSO model and its use of distributed ledgers applies generally to all forms of enterprise (e.g., commercial v. noncommercial, public v. private, civilian v. military, product v. service). Its relevance to real-time control of continuous, discrete and batch processes operating within industrial enterprises at levels 0-3 is of special interest. In Figure 6, L_4 represents an enterprise, L_3 represents its profit centers (e.g., factories), L_2 value production areas, L_1 automated production cells and L_0 devices (e.g., sensors and actuators). Processes at each level have different timing constraints as well as requirements for retaining histories and planning for future actions.

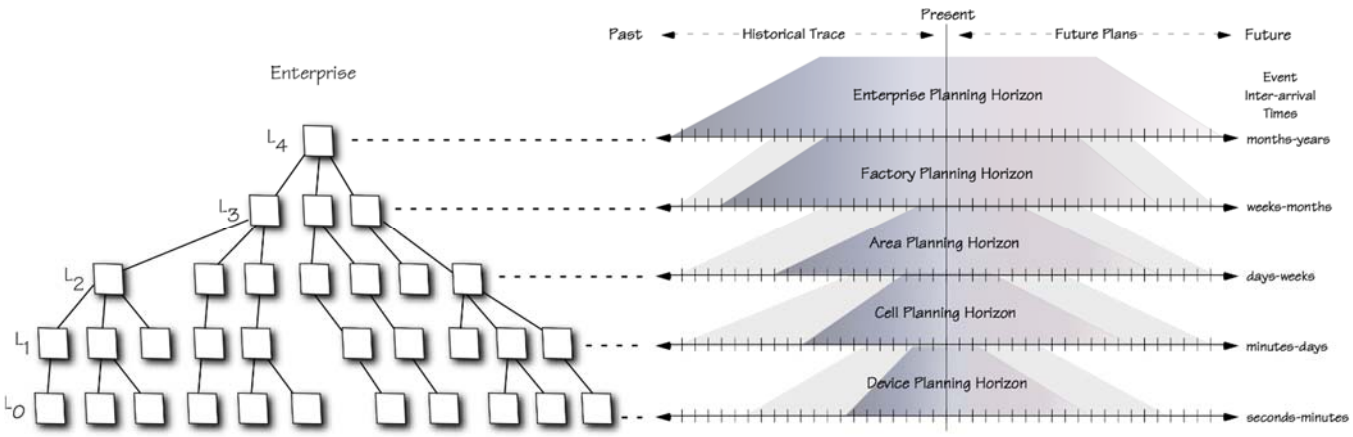

Figure 6 - Enterprise Governance Levels and Process Timing

Intelligent automation that is aided by machine learning will use operational histories to train models capable of predicting and adapting to future behaviors. Data sets used to train machine learning algorithms will grow in proportion to the rates at which CSO-governed processes unfold. For example, Table 2 shows daily data set

volumes for the hypothetical enterprise described in Figure 5. A Business Unit (BU) operates a factory with two production areas. Each area contains two automated cells, with each cell employing 100 sensor/actuator devices.

Table 2 - CSO Data Volume (Example)

Level	CSO Role	Transactions/ /CSO/Day	# CSO	MB/ Transaction	Data Volume (MB/Day)	
4	Bus Unit	1,000	1	1	1,000	0.01%
3	Factory	1,000	1	100	100,000	0.69%
2	Area	1,000	2	200	400,000	2.76%
1	Cell	100,000	4	10	4,000,000	27.58%
0	Device	1,000,000	100	0.1	10,000,000	68.96%
					14,501,000	100.00%

This example configuration generates up to 14 TB of data per day, all of which must be processed, but not all of which must be stored. This data is in some manner germane to tracking and learning about the performance of the BU enterprise and the state of products it manufactures. BU that employ discrete or batch manufacturing processes to produce regulated products (e.g., pharmaceuticals) have more stringent supply chain documentation, tracking and retention requirements, but involve slower production processes. BU that manufacture products with continuous processes (e.g., electricity, steel, petrochemicals) typically generate data at higher rates, but due to their "softer" regulatory and supply chain documentation requirements have less stringent requirements for process data retention.

Blockchains Applied to Industrial Control Systems

Contemporary discussions of applied DLT are primarily focused on non-manufacturing use cases, especially their potential use in retail, healthcare, insurance, and financial transaction. Today, nascent proprietary efforts[11] exist to apply blockchains to IoT use cases, primarily in supporting connections of simple edge devices to unregulated and non-safety-related cloud-based monitoring applications (e.g., "intelligent" homes and buildings).

This paper contributes to defining requirements for and architecture of DLT applications in manufacturing automation. In addition to this work, manufacturing use cases are being considered by some members of the Hyperledger Consortium[12], an umbrella project begun in December 2015 by members of the Linux Foundation[13]. Hyperledger projects support collaborative development of application-oriented blockchain-based distributed ledgers.

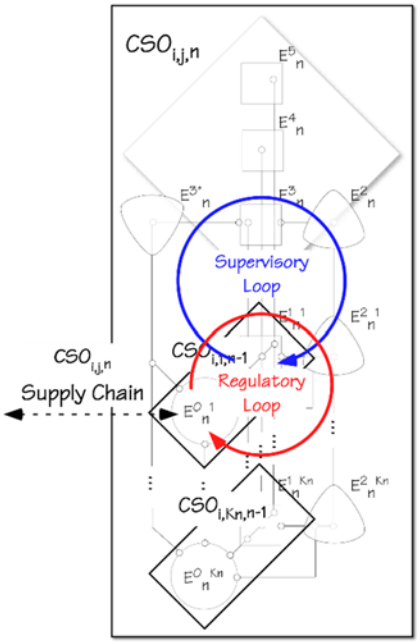

Figure 8 - CSO Control Loops

The objective of the Hyperledger project is to advance cross-industry collaboration by developing blockchains and distributed ledgers while maintaining focus on improving the performance and reliability of these systems (as compared to comparable cryptocurrency designs) so that they are capable of supporting global business transactions by major

[11] https://www.ibm.com/internet-of-things
[12] https://www.hyperledger.org/
[13] http://www.linuxfoundation.org/

technological, financial and supply chain companies. The project seeks to integrate independent open protocols and standards by means of a framework for use-specific modules, including blockchains with their own consensus and storage routines, as well as services for identity management, access control and smart contracts.

Blockchains are useful in supporting CSO control processing, for securing both supply chain regulatory and command chain supervisory control transactions (Figure 8). Regulatory control transactions are local to a given CSO and carried between E^1 Directors, E^0 Controls and E^2 Coordinators on permissioned control networks, and registered on dedicated ledger ($L^{1,0}$). Supervisory control transactions are local to a given CSO and carried between E^3 Operators, subordinate E^1 Directors and E^2 Coordinators on permissioned command networks, and registered on dedicated ledger ($L^{3,1}$). These blockchains provide immutable histories of the command and control operations and their effects on value production.

Blocks the CSO blockchains are validated and linked as shown in Figure 9. At the start of operations, all chains initialize a root or "*genesis*" block containing in its transmitted data field "T_x" the initial conditions of production processes, in its Timestamp the production start time, in its "Nonce" field the seed for its cryptographic hash function, and in its "Hash" the encrypted production data (which is null in the genesis block.)

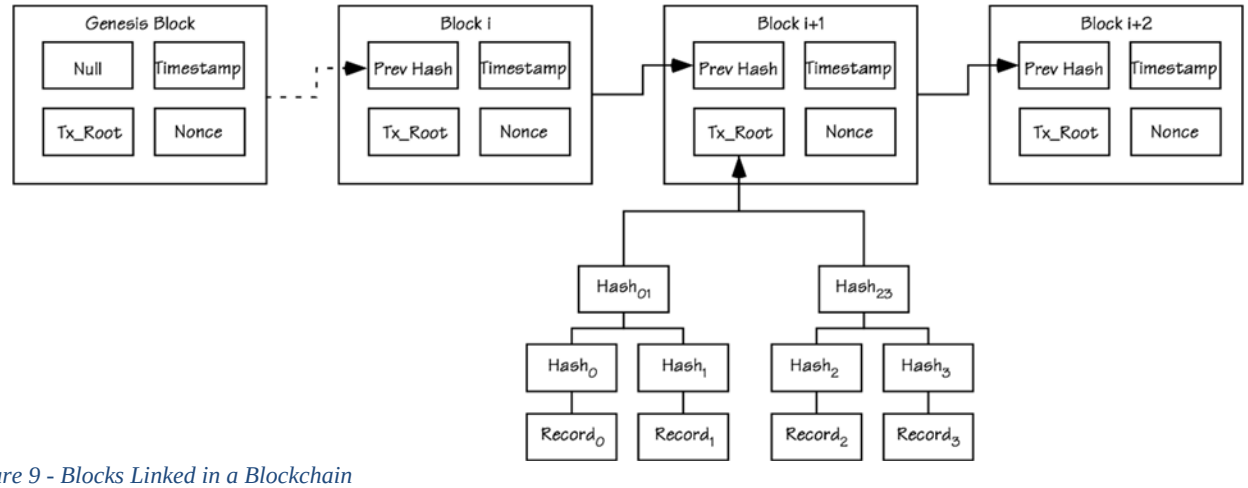

Figure 9 - Blocks Linked in a Blockchain

As production unfolds, each block points to the next block in the chain, establishing the causal ordering of events (situations) and associated encrypted data arising from CSO control actions. Blocks are only added to the chain if all (or at least a majority) of collaborating agents (E^i) agree to the validity of its content. In general, consensus is achieved quickly since the permissioned networks contain few nodes (i.e., K_n), even if subordinate CSO are cyberspacially distributed within the encapsulating CSO.

As noted in Figure 9, the blockchain data field T_x contains only a hash of the data produced by SAS, PGS and PES processing. The body of that data will typically reside in a companion database, accessible by pointers retrieved by decrypting previous hashes.

Conclusion

Details of the role of blockchains in DIO and their CSO automation services (SAS, PGS and PES) are the subject of other papers. The goal here was to introduce a decentralized ledger-based automation framework, including the CSO model, its operating structure and its principal agents. In the face of advancing DLT and AI techniques, industrial automation systems, their hardware and software resources, will continue to evolve to the point where

scalability (e.g., L_n, K_n) and security (e.g., hashed data, consensus[14], and smart contracts) will fundamentally alter traditional centralized industrial control systems, permitting decentralized control though collaboration among trusted peers.

[14] https://en.wikipedia.org/wiki/Byzantine_fault_tolerance