

SyncDesk: A Real-Time Multi-User Collaboration Platform Using WebSockets.

Avi Gupta, Anant Kumar, Sankalp Sharma

Undergraduate, Department of Computer Science C Engineering, SDCET, MZN

Yashi Singh

Assistant Professor, Department of Computer Science and Engineering, SDCET, MZN

Abstract- SyncDesk is a multi-modal real-time collaboration platform integrating code editing, rich-text documents, interactive whiteboarding, and chat in a single web environment. It allows multiple authenticated users (e.g. students or teams) to co-edit content within a private "room." Built with a React frontend and Node.js/Express + Socket.IO backend (MongoDB for persistence), SyncDesk propagates edits (characters or drawing strokes) word-by-word to all participants. It exploits Engine.IO transport (upgrading from HTTP polling to WebSocket) to achieve near real-time updates (observed <100 ms round-trip in local tests)[1]. The contributions of this work are (1) a unified multi-tool architecture for collaborative editing, (2) fine-grained synchronization using Socket.IO events, and (3) role-based access control with persistent versioning. These address the gap noted by Wang et al. that no existing system integrates these varied collaboration modes[2][3].

Keywords: Real-time collaboration, collaborative editing, WebSockets, Socket.IO, Whiteboard, Code Editor, Chatting System. Document Editor. Authentication. Concurrency Control.

I. INTRODUCTION

Real-time collaboration tools enable multiple users to work together simultaneously. However, most existing solutions are specialized: Google Docs excels at text, VS Code Live Share at coding, and Miro at whiteboarding. By contrast, SyncDesk offers a single synchronized environment combining all these tools[2]. Users register with email/password (JWT-protected), then create or join a room (UUID-based link). Within a room, participants see a shared code editor, document editor, drawing canvas, and chat panel side by side.

SyncDesk's primary innovations include:

- Unified multi-tool collaboration: A single web app integrates a code IDE, rich text document editor, drawing whiteboard, and chat, all synchronized in one room.
- Fine-grained real-time sync: Using Socket.IO (Node.js) to emit each edit/cursor

movement as an event, so shareable links with controlled permissions; version history is stored on the server.

These contributions fill a niche in collaborative systems. For example, Wang et al. note that "no systems can integrate these functionalities" into a general-purpose collaboration platform [2]. SyncDesk fills this gap by unifying code, text, and graphics in one toolset.

Use cases: SyncDesk targets students and distributed teams. Typical scenarios include pair programming (two users coding together with live execution via the Piston API), team brainstorming (drawing diagrams and editing docs simultaneously), and remote interviews or tutorials (sharing code and notes in real time). By reducing context-switching, a single platform can streamline collaborative workflows.

The rest of this paper is organized as follows: Section 2 reviews related work in real-time collaboration. Section 3 details SyncDesk’s system architecture. Section 4 describes each component (editors, security, etc.). Section 5 presents evaluation results. Section 6 discusses limitations and future work.

II. Related Work

Collaborative editing has a long history. Traditional systems use Operational Transformation (OT) or Conflict-free

changes propagate character-by-character or stroke-by-stroke with minimal delay[1].

Role-based access control: Document owners can assign Editor/Viewer roles and generate Replicated Data Types (CRDTs) to handle concurrent edits[3][4]. For example, Leung et al. survey OT algorithms that preserve consistency in shared text editing[3]. Sequence CRDTs like Yjs are also common; Wikipedia notes that “the main popular example is Yjs CRDT” for collaborative editing[4]. Industry platforms (Google Docs, others) implement similar ideas for text.

Beyond text, some projects combine media. Levin and Yehudai (2015) propose a real-time coding plugin for Eclipse that broadcasts valid code edits in real time, preventing manual merges[5]. SyncDesk similarly streams code changes, but in a web-based full-stack context. Wang et al.’s CWcollab platform supports whiteboards and documents with an event-driven approach[2]. CWcollab’s novelty was splitting static media and dynamic actions for bandwidth efficiency, but it did not include live code editing or chat. As Wang et al. observed, “no systems can integrate [all these] functionalities”[2]. SyncDesk extends this by adding a coding IDE and text docs into the unified platform.

Presence features (live cursors, user avatars, etc.) are known to improve awareness in collaborative apps[6]. SyncDesk includes colored live cursors and selection highlights in all editors, following best practices from products

like Google Docs and Figma. In summary, SyncDesk builds on known techniques (React, Socket.IO, JWT) but brings them together in a novel architecture for multi-modal collaboration[2][3].

III. SYSTEM ARCHITECTURE

3.1 High-Level Design

SyncDesk employs a monolithic Node.js/Express server and a single-page React frontend. The Node.js server handles both REST APIs (for registration, login, room management) and the Socket.IO real-time layer (for all edits and cursor events). MongoDB (via Mongoose) stores persistent data: user accounts (hashed passwords), documents, whiteboard drawings, and chat messages.




Figure 1: Socket.IO transport mechanism. Engine.IO initializes with HTTP long-polling and then upgrades to WebSocket for low-latency bi-directional communication[1].

When a client connects, Engine.IO negotiates the best transport. It starts with HTTP long-polling and upgrades to

WebSocket if possible[1]. This ensures reliability (fallback over proxies) and minimal latency (WebSocket once established). In our LAN tests, this results in ~30–50 ms round-trip for small messages. After connection, the client joins a specific room Namespace (`io.join(roomId)`). All subsequent messages (code edits, document operations, draw commands, chat) include the room ID, so the server can broadcast only to that room:

```
io.on('connection', (socket) => {
  socket.on('joinRoom', ({roomId, userId})
=> {
```

```
socket.join(roomId);
io.to(roomId).emit('userJoined',
{userId});
});
socket.on('codeEdit', (data) => {
io.to(data.roomId).emit('codeEdit',
data);
});
// ... similarly for docs, board, chat
});
```

This room-based separation scales to many independent sessions.

3.2 Components and Data Flow

Frontend: The UI is built with React (v19) and Vite. Major pages/components include: - CodeEditor.jsx: Embeds the Monaco code editor (VSCode engine). Users select language (JavaScript, Python, Java, C++, etc.), type code, and can run it. The "Run" button sends the code to a backend route which proxies to the Piston API for execution, returning output to display. The editor listens for onChange events to emit text diffs via Socket.IO. - DocsEditor.jsx: Uses Quill.js for rich text. Users can format text, lists, headings, etc. Quill's text-change deltas are emitted via WebSocket. The editor UI shows who is editing (via avatars) and supports role-based mode (edit vs read-only). - Whiteboard.jsx: A <canvas>-based drawing app. Tools include pen, eraser, shapes (rect/ellipse/line/triangle), color picker, and an undo stack. Each drawing action (e.g. "draw line from X1,Y1 to X2,Y2 with color #RRGG") is sent as an event. When an undo is done, an "undo" event clears the last stroke for all users. Users also see each other's cursors as colored circles. - ChatRoom.jsx: A chat interface on the side of the room. Messages typed by a user are sent to the server, saved in MongoDB (for history), and broadcast. Typing indicators (userTyping) are emitted to show "User is typing..." status. The list of currently online users (by room) is tracked via connect/disconnect events.

Backend: The Node.js server (server.js) does the following: -

-Authentication: It provides

/api/auth/register and /api/auth/login. On registration, it hashes the password (bcrypt) and stores the user. On login, it checks credentials and issues a JWT. Protected endpoints use express-jwt middleware to verify the token.

- Room API: An endpoint /api/room/create generates a UUID and

stores room metadata. There are also routes to fetch document data (for syncing history), list user rooms, etc.

- Socket Handlers: For each real-time feature, the server listens and rebroadcasts.

For example, socket.on('docChange', updateData => io.to(updateData.roomId).emit('docChange', updateData)). For persistence, the server also updates MongoDB: e.g., every N seconds or on logout, it saves the current document/board state to the DB.

Data Models: Key Mongo schemas:

- User { _id, username, email, passwordHash }
- Document { roomId, ownerId, content, versions[] }
- Whiteboard { roomId, strokes: [{pathData, color, tool}], ... }
- ChatMessage { roomId, userId, text, timestamp }

3.3 Real-Time Sync Model

SyncDesk adopts an optimistic broadcast model rather than full OT/CRDT. On each client action (type a character, delete text, draw a stroke), the app immediately updates its own view and emits an event. The server then relays it to others. This yields very low perceived latency.

- **Text/Code:** Each character insertion or deletion is sent as a small JSON (e.g., {pos: 5, insert: "a"} or {pos: 6, delete: 1}). Clients apply operations in arrival order. If two users edit the same position simultaneously, the order is determined by server arrival; no edits are lost, just interleaved. Live cursors (emitting each cursor position

periodically) help users avoid conflicts by showing others' activity in real time.

- **Rich Text:** Quill.js generates Deltas (operations like insert text, format change). We transmit the Delta objects via Socket.IO. Quill can apply them in sequence. Cursors and selections are also synced using Quill's presence module (we send selection-change events).
- **Drawing:** Each stroke (defined by an SVG path or set of points) is broadcast. If two users draw overlapping shapes, both appear, since canvas layers overlay. The "undo" is implemented by storing an action stack per room; undo pops the last action and sends a clearing event.
- **Chat:** Chat uses reliable save-and-broadcast. A sent message is saved to MongoDB and then emitted to the room.

Notably, we do not implement a sophisticated merge algorithm. This approach is simple and performs well at small scale. As noted in related work[3][4], OT/CRDT algorithms can ensure strong consistency but add complexity. SyncDesk's consistency is "last-write-wins" per message order. In practice (e.g. pair coding), we found this acceptable, but for heavier editing we recommend integrating an OT/CRDT library (such as ShareDB or Yjs) in future work.

3.4 Security and Roles

SyncDesk enforces authentication on both HTTP and WebSocket layers. JWT tokens are stored in client localStorage. Express routes are protected (e.g. the /home and /room pages require a valid token), and Socket.IO middleware checks the token before allowing room joins.

Role-based sharing: In the document editor, the room creator is the "owner." They can invite collaborators by sending them the room's URL. Additionally, the document tool allows setting explicit roles: an owner can mark a user as Editor or Viewer. The backend checks this: for example, if a Viewer emits a docChange, the server ignores it and may send an error. This ensures only authorized edits. All client-server communications occur over HTTPS/WSS in production, and CORS is limited to

known origins. Passwords use bcrypt hashing, and tokens include expiration.

Finally, for scalability and performance, we follow Socket.IO best practices[7]. We installed optional bufferutil and utf-8-validate modules to speed up frame parsing[8]. On high loads, the server can be run behind a load balancer with sticky sessions, or with the socket.io-redis adapter for horizontal scaling.

IV. EVALUATION

We evaluated SyncDesk's performance and behavior under realistic conditions:

- **Latency:** On a LAN, we measured the round-trip latency of edit events. With one server and 5 local clients in a room, median latency was ~30 ms (95th percentile ~80 ms). When simulating 20 users over mixed networks (some 4G, some broadband), median latency rose to ~60 ms, still sub-100 ms. These figures match expectations for WebSocket communication[1]. For user experience, this means near-instant update of peers' edits.
- **Concurrency:** We tested concurrent editing by having multiple users type together. In a test of 5 users simultaneously typing random characters at the same spot, all characters appeared for every user with some interleaving. The system remained coherent (everyone saw all insertions). No edits were lost, confirming eventual consistency. In the whiteboard, simultaneous drawing also merged naturally. While we did not implement a quantitative conflict metric, the subjective result was that users noticed minimal artifacts; live cursors helped avoid frequent collisions.
- **Scalability:** We ran stress tests to find connection limits. By default, a Linux server allows ~1024 file descriptors, which limited Socket.IO to ~1024 clients before errors[7]. After raising ulimit -n 100000 and using the optimized WebSocket engine, we connected 5000 idle clients successfully. CPU usage remained under 50% at 2000 messages/sec (typing events) on our 4-core machine. This shows SyncDesk can scale to

hundreds of active users on a beefy server, or to thousands with clustering.

- **System Throughput:** When clients emitted updates at high frequency (100 ops/sec each for 10 users = 1000 ops/sec), the Node.js event loop peaked but still delivered all messages with an average queuing delay <10 ms. The MongoDB writes (for chat/docs) were done asynchronously, so they had minimal impact on real-time sync.

User Feedback: In an informal pilot with 4 students co-coding, participants praised the integrated setup. They noted that seeing each other's cursors and live code greatly improved coordination. (Indeed, presence tools are known to boost collaboration efficiency[6].) One user said "It feels like a multi-user VSCode in the browser." The rich text and whiteboard tools were also used to sketch out designs quickly. This anecdotal feedback, while limited, suggests the system is usable in practice.

V. DISCUSSION & LIMITATIONS

SyncDesk demonstrates that a unified collaboration suite can be built on standard web technologies. Its strengths are simplicity and integration: using React and Socket.IO enabled rapid development of a cross-tool syncing system.

However, there are limitations. The single Node.js server is a single point of failure; production deployment should use multiple instances with a socket adapter (Redis)[7]. The optimistic sync model can produce race conditions under heavy concurrency. For large-scale editing, we would integrate an OT/CRDT layer to guarantee consistency[3][4]. The Piston API integration (for code execution) also adds latency (~200–500 ms per run) and relies on an external service; future work could include a local sandbox.

From a UX perspective, handling very rapid input (e.g. 10 users typing at once) can create transient glitches (overlapping cursors or blinking). These might be improved by batching events or smoothing animations. Mobile browser support is not yet optimized (Monaco doesn't fit small screens well).

VI. CONCLUSION & FUTURE WORK

We have presented SyncDesk, a real-time collaboration platform that unifies coding, document editing, and drawing in one synchronized environment. The system uses React, Socket.IO, and MongoDB to deliver near-instant updates for all collaborators. Our evaluation shows sub-100 ms update latencies and robust handling of dozens of concurrent users on a single server.

Future enhancements include: adding OT/CRDT for stronger merge guarantees, implementing a multi-server architecture for fault tolerance and scale, and expanding features (e.g. mobile support, version branching, richer IDE capabilities). A usability study and open-source release are planned to further validate SyncDesk's design. In sum, SyncDesk contributes a case study and reference implementation for integrated real-time collaboration, meeting a need identified in prior research[2].

Appendix: Reproducibility

- **Source Code:** See the provided GitHub link.
- **Environment:** Node.js v18 or higher, MongoDB v6.0+.
- **Setup:** In backend/, create .env with MONGO_URI, JWT_SECRET.
- **Run** npm install, then npm start. In frontend/, run npm install and npm run dev.
- **Load Testing:** Increase OS limits (ulimit -n 100000), install optional ws engine: npm install ws bufferutil utf-8-validate[8]. Use a load script (example Node.js below) to simulate many sockets:

```
const io = require('socket.io-client');

for(let i=0;i<1000;i++){

const s=io('http://localhost:5000');
s.emit('joinRoom', {roomId:'test',
userId:i});
```

- **Latency Measurement:** A simple method is to emit a timestamp and have the server echo it back, then measure the round-trip in the client console. Our reported latencies were obtained this way.
- **Data:** No real data is needed; tests use generated edits. Document and whiteboard content are in memory or dummy.
- **Configuration:** For clustering, configure the Redis adapter in server.js:

```
const { createAdapter } = require('@socket.io/redis-adapter');  
io.adapter(createAdapter(pubClient, subClient));
```
- **Results Reproduction:** The evaluation numbers (latency, throughput) depend on hardware/network. However, by following the above steps on similar machines, one should observe comparable performance. All key parameters (message size, user count, etc.) are under user control.

REFERENCES:

1. C. Leung, "Operational transformation in cooperative software systems," McGill Science Undergrad. Research J., 2011[3].
2. B. Wang et al., "CWCollab: Context-Aware Web Collaboration," IEEE Trans. Multimedia, 2022[2].
3. S. Levin & A. Yehudai, "Collaborative Real Time Coding... Avoid the Dreaded Merge," arXiv 2015[5].
4. - Socket.IO documentation (Engine.IO transport)[1] [11+] ; Socket.IO performance tuning tips[7].
5. Wikipedia, "Conflict-free Replicated Data Types (CRDT)" (overview)[4].