

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Video Conferencing Web App Using WebRTC

Vinay Naik^a, Swarup Misal^a, Sarthak Mogane^a, Prasad Palkar^a, Bhagyashri Jadhav^b

^aStudents, Dr. D. Y. Patil Pratishthan's College of Engineering, Kolhapur, 416008. ^bAssistant Professor, Dr. D. Y. Patil Pratishthan's College of Engineering, Kolhapur, 416008. bhagyadyp27@gmail.com

ABSTRACT:

In recent years, video conferencing has become an essential tool for communication in work, education, and social contexts. This paper presents a WebRTC-based video conferencing application that enables real-time peer-to-peer video and audio communication directly within web browsers. The system is built using open web technologies – notably the WebRTC API for media streaming, a React.js front-end for user interface, and a Node.js/Express backend with Socket.io for signaling. Key features of the application include one-click meeting creation with unique IDs, direct peer-to-peer connectivity for high-quality audio/video, integrated text chat, screen sharing functionality, and the ability for users to toggle camera and microphone states during a call. We describe the design and architecture of the system with supporting diagrams (use case, sequence, state, and architecture diagrams), and discuss implementation details such as the use of STUN/TURN servers for NAT traversal and Firebase for user authentication. Results from our implementation demonstrate that a robust video conferencing experience can be achieved using WebRTC, with minimal latency and effective synchronization of media and data streams. We conclude with an analysis of the system's performance and discuss potential enhancements for scalability and additional features in future work.

Keywords: WebRTC; Video Conferencing; Peer-to-Peer Communication; Real-Time Communication; Signaling;

Introduction

The demand for reliable video conferencing solutions has surged, especially following the global shift to remote work and online collaboration. For example, a statistical analysis reported that 73% of internet users engaged in online meetings or lessons via video conferencing apps in 2020, up from 48% in 2019 irishexaminer.com

- Prominent platforms like Zoom experienced explosive growth a 1,900% increase in daily active users in early 2020, reaching ~200 million daily meeting participants usebubbles.com.
- This trend underscores the need for accessible, high-quality video communication tools. Traditional solutions often require dedicated
 installations or proprietary clients, whereas web-based approaches can offer instant connectivity through the browser. Web Real-Time
 Communication (WebRTC) has emerged as a standard solution for in-browser video conferencing
 medium.com
- WebRTC is a free, open-source project that enables web browsers to communicate directly with each other via real-time audio, video, and data streams.

productivityland.com

- By using WebRTC, developers can implement peer-to-peer (P2P) connectivity for voice and video calls without plugins, as all modern browsers natively support the required APIs
- medium.com
- This approach leverages the RTCPeerConnection API (for streaming media) and RTCDataChannel (for arbitrary data) to establish secure connections between clients. The chief advantage is that media is exchanged directly between peers, which can reduce latency and server load while ensuring end-to-end encryption by default. Given these benefits, we propose a web-based video conferencing application using WebRTC to allow users to initiate and join video calls with minimal friction. The purpose of our app is to demonstrate that open web technologies can be combined to achieve a rich conferencing experience comparable to popular commercial platforms. Our solution uses a front-end built with React.js for a dynamic user interface and a back-end powered by Node.js and Express.js to facilitate signaling and session management. A lightweight signaling server is necessary because, while WebRTC handles the media exchange, it does not define how peers initially find each other and negotiate the connection. A signaling mechanism (e.g., using WebSockets) is used to exchange Session Description Protocol (SDP) offer/answer messages and network candidate information between clients. developer.mozilla.org
- Once signaling is done, peers connect directly and stream audio/video content over a P2P channel. This paper details the design, architecture, and implementation of the proposed WebRTC video conferencing app. Section 2 reviews related work and existing solutions. Section 3 describes the system architecture and design with the help of diagrams. Section 4 covers implementation specifics, including the technologies

14523

and protocols employed. In Section 5, we discuss the results and features of the application, such as peer-to-peer media quality, text chat, screen sharing, and media control. Finally, Section 6 concludes the paper and suggests directions for future enhancements.

Related Work

Video conferencing has been extensively explored in both industry and academia, leading to a variety of solutions. Established platforms like Zoom, Google Meet, and Microsoft Teams dominate the space, especially after the pandemic-driven surge in remote communications. These platforms typically employ advanced infrastructures to handle large numbers of participants and to ensure call reliability. Google Meet, for instance, uses WebRTC under the hood with media servers to manage multi-party conferences, employing techniques like simulcast and scalable video coding for efficiency. developer.mozilla.org

- Zoom initially used proprietary protocols and native applications, but it also offers WebRTC-based web clients; the platform's rapid adoption demonstrated the scalability challenges and engineering complexities of real-time video at a global scale usebubbles.com
- Microsoft Teams similarly integrates WebRTC for browser clients and uses cloud-based routing to support group calls and recordings. Opensource projects provide alternative approaches to building video conferencing systems. Jitsi Meet is a notable example – it is an open-source WebRTC application that enables multi-user video conferences by using a Selective Forwarding Unit (Jitsi Videobridge) to relay media between participants.
 - jitsi.org
- This server-centric forwarding approach offloads the peer clients and allows dozens of users in a meeting, at the cost of requiring a powerful server. Other open-source frameworks and libraries, such as BigBlueButton (focused on online learning) and PeerJS (a simpler P2P WebRTC library), further illustrate the flexibility of WebRTC for various use cases. These solutions typically add features like recording, moderated controls, and integration with other services. Our proposed application situates itself in between these extremes it is lighter-weight than enterprise solutions but more tailored than a generic library. Rather than employing a multipoint control unit (MCU) or SFU server for media, our design uses the WebRTC peer-to-peer model for direct communication. This is suitable for calls with a modest number of participants (e.g., small team meetings or one-on-one discussions), where a fully peer-to-peer mesh can be utilized. By using a custom signaling server (built with Node.js and Socket.io), we maintain control over the connection setup process while relying on browsers to handle the media exchange. Similar architectures have been demonstrated in tutorials and prior projects, where a Node/Socket.io signaling back end coordinates WebRTC connections between clients in a web application
 - developer.mozilla.org.

System Architecture and Design

The video conferencing application is designed with a client-server architecture for signaling and a peer-to-peer architecture for media streaming. In essence, all participating clients (web browsers) communicate through a central signaling server only during the setup phase; once a call is established, audio and video streams flow directly between browsers via WebRTC. This section describes the overall architecture and key design diagrams that illustrate the system's functionality.

Overall Architecture:

The high-level system architecture is depicted in Figure 1, which we refer to as the proposed system design. It comprises three main parts: the end-user clients (browser applications), the server-side components, and external services. Each end-user interacts with the application through an intuitive User Interface (built in React) that manages local state (including any data stored in browser localStorage) and interacts with the WebRTC API for media capture and playback. When a user initiates or joins a meeting, the app engages the Media Devices API (getUserMedia) to access the user's camera and microphone streams, and the WebRTC API to prepare peer connection objects.

The server-side component includes a Signaling Server (built with Socket.io on Node.js/Express) that coordinates the exchange of connection metadata, an Authentication Server (using Firebase Authentication for user sign-up/login), and possibly a simple Room Management logic to handle meeting IDs. For NAT traversal and connectivity, standard STUN and TURN servers are utilized as external services alongside the signaling server. STUN servers enable clients to discover their public IP/port and any NAT restrictions. developer.mozilla.org, while TURN servers are a fallback to relay media when a direct peer-to-peer path cannot be established (e.g., due to symmetric NATs) developer.mozilla.org.

 Additionally, a cloud Firebase service is used for storing user credentials or meeting records, and optional Cloud Storage can be leveraged if features like recording are implemented.

The interactions among these components are illustrated by the arrows in Figure 1, showing how media streams and messages flow. Once two or more clients have exchanged the necessary signaling data (through the server and possibly via Firebase for authentication), they form direct peer connections among each other to carry the conference media streams.



Figure 1: Proposed system architecture for the WebRTC video conferencing app.

This diagram shows how an end user's browser (left, containing the user interface, local storage, and WebRTC API) interacts with external services (STUN/TURN servers for NAT traversal, Firebase for authentication/storage) and how multiple clients (right side, Client A/B/C) establish peer-to-peer connections for audio/video streaming. The signaling server and other server-side functions (center top) facilitate session setup, while media flows directly between clients once the conference is established.

Use Case Diagram: The functional scope of the application is captured in the use case diagram (Figure 2). It represents the primary actions a User can perform and how these actions relate to system components. The user can Sign Up and Login (handled via Firebase Authentication) to create an account or authenticate, ensuring that only authorized users can initiate or join meetings. Once authenticated, the user may Create Meeting (generating a new meeting room with a unique ID) or Join Meeting (by entering an existing meeting ID shared by a host).



During a conference, the participant can perform in-call actions: Toggle Camera and Toggle Microphone to enable/disable video or mute/unmute themselves, Share Screen to broadcast their screen to others, and use the Chat feature to send text messages. The user can also End Meeting to leave the call (or end it for everyone if they are the host).

Supporting systems are indicated below each set of use cases: authentication actions rely on Firebase Authentication, meeting state and user preferences may be stored in LocalStorage (for example, caching the user's name or camera preferences), and all real-time media streaming and toggling of devices are enabled by WebRTC APIs in the browser. This use case design ensures that the app covers all typical functionalities expected in a video conferencing tool.



Figure 2: Use case diagram illustrating the key user interactions in the WebRTC video conferencing system.

The single actor "User" encompasses both host and participant roles. Core use cases include user authentication (Sign Up, Login, Logout), meeting lifecycle management (Create Meeting, Join Meeting, End Meeting), and in-call features (Toggle Camera, Toggle Microphone, Share Screen, Chat). The diagram also highlights which system components or services facilitate each group of actions (e.g., Firebase for authentication, browser storage for user data, WebRTC for media streaming).

Sequence of Operations: To better understand the dynamic behaviour of the system, Figure 3 shows a sequence diagram for a typical call setup and teardown between two users (Browser A and Browser B) with the involvement of the Signaling Server. The process begins when User A opens the app and logs in (if not already logged in). User A then chooses to create a new meeting. Browser A sends a request to the signaling server to Create meeting, upon which the server generates a unique Meeting ID and returns it. User A shares this meeting ID with User B (e.g., via a link or invitation outside the system).

When User B opens the app and enters the provided meeting ID to join, Browser B contacts the signaling server to Join meeting. The server registers B as a participant of that meeting and notifies Browser A (the host) that a new participant has joined. At this point, the WebRTC peer-to-peer negotiation begins. Browser A, upon notification, creates an SDP offer (containing its media capabilities) and sends a connection offer via the signaling server to

Browser B. The server simply forwards the offer to Browser B. Browser B then generates an SDP answer (after setting up its end of the peer connection and possibly prompting the user to allow camera/microphone) and sends this answer back to Browser A through the server. Once A receives B's answer (forwarded by the server), both sides proceed to exchange ICE (Interactive Connectivity Establishment) candidates – network addresses and ports – to find the best path for direct connection

developer.mozilla.org

With the offer/answer and ICE exchange complete, the peers can establish the WebRTC connection directly. At this stage, audio/video streams flow between Browser A and Browser B over the peer-to-peer link (indicated by the bidirectional Audio/Video Stream arrows in the diagram). Throughout the call, either user can perform in-call actions: for example, if a user toggles their camera or mic, the app will update the media state (locally and via signaling to inform the other peer or simply via the media stream's track status).

Similarly, if a user sends a chat message, the message is delivered either through the data channel of the peer connection or relayed via the signaling server as a chat event (the design can use either approach; our implementation opts to use the existing Socket.io channel for simplicity).

The sequence diagram shows a Send chat message from Browser B, which the signaling server delivers to Browser A, enabling instant text communication alongside the video stream. Finally, when a user decides to end the call, they click "End call". Browser A (if ending) notifies the server that the call has ended, and the server in turn notifies Browser B that the meeting is terminated (Notify call ended).

Both peers then tear down their peer connection and return to an idle or post-call state. This sequence illustrates the complete lifecycle of a two-party conference in our system, from meeting creation and joining to real-time interaction and call termination.



Figure 3: Sequence diagram of the call setup and communication flow between two clients (Browser A and Browser B) using a signaling server.

It outlines the steps of opening the app, creating a meeting and sharing the meeting ID, joining a meeting, and the WebRTC offer/answer exchange through the signaling server to establish a peer-to-peer connection. Once connected, media (audio/video) streams directly between browsers, and features like toggling camera/mic or sending chat messages are handled either via the established peer connection or minimal signaling updates. The call teardown is also shown, where ending the call triggers notifications and closure of the connection.

State Diagram: Internally, the application's client-side logic can be described by a state machine reflecting the user's journey and the connection status (Figure 4). The state diagram starts at an Idle state, representing the initial condition when the app is opened and no action has been taken yet. Upon launching the app, the user enters an Authenticating state – for instance, triggering a Google Sign-In via Firebase or a similar login process (this is marked by an action "Google sign-in (Firebase)" in the diagram). Once authentication succeeds, the state moves to Authenticated, meaning the user is logged in and can now create or join meetings.

If the user opts to create a new meeting or joins an existing one, the system transitions to a WaitingRoom state. In this state, the user might be alone in a newly created meeting (waiting for others to join) or is waiting for the host to start the call. The state diagram indicates that the user must take an action (e.g., click "Join Call") to initiate the WebRTC connection once all participants are ready.

Upon initiating the call, the client state goes into Connecting, during which the application is exchanging SDP offers/answers and ICE candidates with the peer(s) – this phase is labelled as Signaling in the diagram, reflecting the handshake process described earlier. After successful signaling, the peer connection is established and the state advances to ConnectingMedia, where the media streams are being set up.

Once the first audio/video streams are flowing and at least one remote track is received, the system enters the InCall state, indicating an active conference is underway. While in the InCall state, there are substates or concurrent states that represent optional modes: for example, muted (if the user has muted their microphone), ScreenSharing (if the user is currently sharing their screen), Chatting (if the chat panel is open or new messages are being exchanged), etc. These can be considered parallel states or flags since the user can be in the call and simultaneously have some of these features toggled on.

The diagram shows triggers like "User mutes mic" leading into the Muted sub-state, or "Start screen share" leading into ScreenSharing, and corresponding actions to exit those modes ("Stop screen share", "Close chat panel", etc.). Finally, when the call is ended – either the user leaves or all users disconnect – the state transitions to Disconnected, and from there the application may return to the Idle state or another appropriate state (for example, a post-call summary or simply ready to start/join another meeting). The state diagram ensures a clear understanding of how the application manages transitions from not being in a call to in-call and back, as well as how various in-call features are handled in terms of state.



Figure 4: State diagram for the client-side application, illustrating the progression of states from launching the app through authentication, meeting setup, and an active call.

The main states include Idle (no session), Authenticated (user logged in), WaitingRoom (in a lobby or newly created meeting), Connecting (handshaking via signaling), and InCall (active session). Transient signaling steps (exchange of SDP and ICE) are highlighted during the Connecting phase. The InCall state encompasses possible sub-states for features like being Muted, ScreenSharing, or Chatting. The diagram also shows the transition to a Disconnected

state when a call ends, after which the user can return to Idle or start a new session. The combination of these diagrams – use case, sequence, and state – guided the design and implementation of our system. They ensure that all functional requirements are accounted for and that the behaviour during the call setup and teardown is well-defined. Next, we delve into the implementation details, describing how we realized this design using specific technologies and libraries.

Implementation Details

The video conferencing application was implemented using a collection of modern web technologies and services. The choice of stack was influenced by the need for real-time performance, ease of development, and integration with WebRTC. Key technologies include WebRTC API in the browser, React.js for the front-end, Node.js with Express.js for the back-end server, Socket.io for real-time signaling, Firebase for authentication, and STUN/TURN servers for network traversal support. Figure 5 provides a structural overview of the technology stack and how these components interact in our implementation.



Figure 5: System implementation diagram illustrating the technology stack used in the application.

The front-end consists of HTML5/CSS3 for structure and styling, React.js as the UI framework, and JavaScript for client logic, which together interface with the browser's WebRTC API and utilize browser local storage. The back-end is built on Node.js (runtime) and Express.js (web framework), with Socket.io providing real-time communication over the WebSocket protocol. External services integrated include Firebase (for authentication) and STUN/TURN servers to facilitate peer-to-peer connectivity behind NATs.

Front-End (Client): The client side is a single-page application built with React.js. React was chosen for its component-based architecture and efficient state management, which are beneficial for an interactive UI that needs to reflect real-time events (like incoming calls, new chat messages, etc.). The UI comprises components for the lobby (meeting creation/join screens), the video call interface (video elements for local and remote streams, buttons for mute/unmute, etc.), and chat panel among others. We used standard HTML5 and CSS3 for layout and design, ensuring a responsive interface that can run in desktop browsers (and potentially mobile browsers with minor adjustments).

Client-side JavaScript logic manages interactions with the WebRTC API. For example, when the user clicks "Share Screen", the app invokes navigator.mediaDevices.getDisplayMedia() to capture the screen content as a stream developer.mozilla.org

This captured stream is then added to the peer connection so that it gets sent to the remote peer as an additional video track. Similarly, toggling the microphone or camera is implemented by controlling the media tracks – for instance, calling track.enabled = false on the audio track to mute (and toggling it true/false when user clicks the button). These operations are reflected in the UI via React state updates (e.g., showing a muted icon).

The front-end also handles establishing and managing the RTCPeerConnection object: when a user joins a call, the client creates a new RTCPeerConnection and sets up event handlers for ice candidate (to send ICE candidates to the server) and ontrack (to receive remote media streams and attach them to video elements).

Back-End (Server and Signaling): The signaling server is built with Node.js, using Express.js to handle HTTP requests and serve the React front-end (if deployed on the same server) as well as to define API endpoints if needed (for example, an endpoint to generate meeting IDs, though in our case meeting ID generation was handled on the fly in memory). Real-time bidirectional communication is enabled by Socket.io, a library that abstracts WebSocket connections and provides a simple event-driven API. Socket.io was integral in implementing the signaling protocol for our app: when a user creates or joins a meeting, the client emits Socket.io events (like "join-room" with the room ID). The Node.js server listens for these events and uses Socket.io's rooms feature to group sockets belonging to the same meeting. It then facilitates exchange of WebRTC signaling messages: for instance, when Client A wants to send an offer to Client B, it emits an "offer" event with the SDP data; the server, knowing which room B is in, forwards that offer to B's socket. The same happens for the answer and ICE candidates. It's worth noting that the signaling server does not interpret the contents of these messages (SDP or ICE), it simply relays them between clients

developer.mozilla.org

This keeps the server logic simple and lightweight.

Beyond signaling, the Node.js server also manages basic meeting state - e.g., keeping track of active rooms, and broadcasting participant join/leave events so that clients can update their UI (like showing when a new user has entered the call).

WebRTC and Networking: To establish peer-to-peer communication, our application relies on the WebRTC API provided by browsers. This includes using RTCPeerConnection to handle peer connections and getUserMedia/getDisplayMedia to access local media streams. A critical aspect of WebRTC implementation is handling NAT traversal. Users are often behind NAT routers, meaning direct P2P connections require a negotiation of network endpoints. We configured our peer connections with a list of ICE servers (STUN/TURN servers) – for example, using a public STUN server (like Google's STUN at stun.l.google.com) to get reflexive addresses

developer.mozilla.org,

and a TURN server (we set up a Coturn server for testing) to relay media if direct connection fails

developer.mozilla.org

In testing, most peers could connect via STUN (direct UDP) but in some restrictive networks the TURN relay was successfully used, confirming the importance of including a TURN fallback. The ICE framework automatically prioritizes direct routes and uses TURN only as a last resort medium.com

so, performance impact was minimal in typical scenarios.

Data Channel and Chat: WebRTC's RTCDataChannel API provides a way to send arbitrary data directly between peers, which can be utilized for chat messages or file transfer. In our implementation, we explored two approaches for the in-call chat feature: using a data channel vs. using Socket.io. Implementing via WebRTC data channel has the benefit of keeping all real-time interaction peer-to-peer; a data channel can be opened once the peer connection is established, allowing text to flow directly alongside audio/video. This approach ensures that chat messages are delivered with low latency and remain private between the participants (not even touching the server). WebRTC data channels are reliable by default, meaning they guarantee message delivery in order (similar to TCP)

developer.mozilla.org, which is appropriate for chat.

The alternative approach is to reuse the Socket.io signaling connection to transmit chat messages (since that is already a realtime path to all participants). We opted for the latter in our initial build for simplicity – the client sends a chat message event to the server, which broadcasts it to others in the room. This introduces a small dependency on the server during the call, but since the messages are low-bandwidth and our server can handle it, the trade-off was acceptable. In future iterations, enabling the RTCDataChannel for chat would be straightforward and could further decentralize the communication. Regardless of the method, from the user's perspective the chat is seamlessly integrated: messages appear in the chat sidebar in real time, even while video and audio are streaming.

Firebase Integration: We integrated Firebase primarily for its authentication service, which allowed us to implement user registration and login with minimal backend code. Users can sign up with email/password or use OAuth (e.g., Google sign-in) – our state diagram (Figure 4) highlights Google sign-in as an example. Firebase Auth handles secure credential storage and verification, returning an auth token that the client can use to identify itself. In the

context of the conferencing app, authentication isn't strictly required for WebRTC, but it adds a layer of access control (e.g., we could restrict creating meetings to logged-in users, and tag meetings with user IDs for ownership).

We also used Firebase Cloud Firestore in a limited way to log active meeting IDs and some metadata, which could be useful for persistence (so that if the signaling server restarts, meeting info could be recovered). Additionally, if a recording feature were to be implemented, Firebase Cloud Storage could be a convenient destination to store recorded video files or snapshots. In our current implementation, recording is not yet active (though the design anticipates it as shown in the Proposed System diagram); if added, likely the approach would be to have one peer (or a server component) use the MediaStream Recording API to save the stream and upload it to cloud storage.

During development, we employed Browser Local Storage to retain certain user settings across sessions – for instance, remembering the user's name or whether their camera/mic were last left enabled, as well as caching the last used meeting ID for convenience. This improved the user experience when testing repeatedly. In summary, the implementation brings together a performant front-end with direct use of WebRTC capabilities, and a back-end that handles the crucial signaling logic. The use of Socket.io and Firebase greatly accelerated the development of reliable signaling and auth features, allowing us to focus on the core WebRTC integration. We next discuss how the implemented features perform in practice and examine the outcomes of our testing.

Results and Discussion

The completed WebRTC video conferencing application was tested in various scenarios to evaluate its functionality and performance. Overall, the results demonstrate that the app successfully enables real-time video communication with additional collaboration features, all within the browser environment. Key features and observations are discussed below:

Peer-to-Peer Connection Quality: Once the signaling handshake is completed, the media streams establish directly between peers, resulting in low latency and high video/audio quality. In one-on-one calls, the end-to-end delay for audio was barely noticeable (well under 200ms), and video latency was very low, on the order of a few hundred milliseconds or less (depending on network conditions). The direct P2P nature eliminates server relay for media, which is reflected in efficient bandwidth usage. In our tests across different networks, WebRTC's adaptation mechanisms adjusted stream quality based on bandwidth; for instance, when one user temporarily experienced limited network capacity, the video resolution and bitrate were automatically reduced to prevent stalling. This adaptive behaviour is built into WebRTC and helped maintain a stable conference experience. We also verified that the STUN/TURN setup was effective: in scenarios where two clients were behind different NAT types, the ICE candidates exchanged via the signaling server enabled them to find a working route. In one case of extremely restrictive network settings, the connection fell back to a TURN relay, which allowed the call to succeed albeit with a slight increase in latency (as expected when using a relay server).

Text Chat Functionality: The integrated chat worked reliably during calls, allowing participants to exchange messages alongside the video feed. Messages sent by one user appeared for the other user almost instantaneously (typically within a few tens of milliseconds delay, which is negligible). Using Socket.io for delivering chat messages proved sufficient — the message transmission is effectively real-time over the WebSocket connection. We also simulated scenarios of rapid message sending (to ensure our approach could handle bursts) and found no issues; Socket.io and the Node server handled the load well in these small-scale tests. The chat feature provides a useful communication channel for participants to share URLs, notes, or coordinate without interrupting the audio/video conversation. If we later switch to WebRTC data channels for chat, we expect similar performance, as data channels have been shown to handle low-latency text exchange in WebRTC applications developer.mozilla.org

• Screen Sharing: The ability to share one's screen was implemented using the WebRTC Screen Capture API (getDisplayMedia). When a user clicked "Share Screen", the browser prompted for permission to choose a screen or application window to share, as per security requirements. developer.mozilla.org



Figure 6: Room Chat

	Choose what to share with http://lo The site will be able to see the contents of yo	ocalhost:5173	
	Chrome Tab	Window Entire Screen	
	Entire screen		
Meeting Info ×		Share Cancel	
0	· · ·	🖵 × 🍕 🚱	

Figure 7: Screen Sharing

Upon user approval, the selected screen content was captured and a video stream of the screen was sent into the peer connection. In testing, the screen sharing feature allowed high-resolution display of the presenter's screen on the viewer's side, with a reasonable frame rate (typically 15-30 fps depending on content changes and network). This is adequate for things like presenting slides, demonstrating software, etc. The frame rate might drop for very fast-moving screen content (e.g., video playback) due to bandwidth constraints, but overall, it functioned as expected. Notably, WebRTC ensures that the screen stream is just another MediaStreamTrack – meaning all encryption and P2P benefits apply equally to it. The receiving end treated the screen share as another video track and we implemented logic to dynamically display it (e.g., switching the UI to show the shared screen prominently). We also confirmed that multiple peers could view the shared screen if more than two were in the meeting, as the track gets sent to each peer connection in a mesh scenario.

Media Toggle (Mute/Unmute, Camera On/Off): Users can toggle their microphone and camera during the call. This was reflected both locally (e.g., showing a muted icon and not sending audio) and remotely. When a user muted their microphone, their audio track was effectively silenced for others. WebRTC handles this by still sending the track but marking it as disabled (or one can choose to stop the track); we chose to use the enabled/disabled toggle for quick resumption. Similarly, turning off the camera would freeze or hide the video on the other side. This feature is important for privacy and bandwidth management during calls. The responsiveness of toggling was instant – a user's own preview updated immediately, and the remote side saw the change within a second or so (depending on keyframe intervals for video when re-enabling). These controls map to typical user expectations (like "mute yourself when not speaking"), and our implementation confirmed they work without disrupting the underlying connection.



Figure 8: Media Toggle

Multi-Participant Calls: While most of our testing focused on one-on-one calls, we did trial the app with 3 participants in a mesh network. WebRTC natively supports group calls by each pair of peers establishing a connection (full mesh). In a 3-person meeting, this means each browser has two peer connections (one to each of the other two participants). We found that the application could handle this scenario – all three video streams were visible and audio flowed among all participants. However, as expected, the upstream bandwidth usage for each participant increased (since each had to send their video to two peers). In our test environment with moderate bandwidth, 3-way calls were acceptable, but quality could degrade if one user's upload capacity was limited. This underscores a common consideration: pure P2P meshes do not scale well to very large meetings (e.g., 10+ people) due to exponential growth in connections and bandwidth per client developer.mozilla.org

- Our system is geared towards small group meetings; for larger scale, a different architecture (like SFU) would be advisable. We include this
 observation to highlight the trade-off between a serverless mesh approach and a managed server approach like Jitsi's
 jitsi.org
- Reliability and Edge Cases: Throughout usage, the app remained stable. We tested edge cases such as: one user closing their browser tab
 unexpectedly (the other user got a disconnect event and the UI handled it gracefully), network interruption and recovery (if a user's connection
 dropped and then reconnected, in some cases the peer connection could be re-established using WebRTC's built-in ICE restart mechanisms;
 in other cases, a manual re-join was required improving automatic reconnection is an area for future improvement). The signaling server
 handled multiple sequential meetings and disconnections without issues, aided by the simplicity of our protocol. We also ensured that all
 resources (media tracks, peer connections) are properly closed on call end to avoid memory leaks or ghost streams.

Reliability and Edge Cases: Throughout usage, the app remained stable. We tested edge cases such as: one user closing their browser tab unexpectedly (the other user got a disconnect event and the UI handled it gracefully), network interruption and recovery (if a user's connection dropped and then reconnected, in some cases the peer connection could be re-established using WebRTC's built-in ICE restart mechanisms; in other cases, a manual re-join was required – improving automatic reconnection is an area for future improvement). The signaling server handled multiple sequential meetings and disconnections without issues, aided by the simplicity of our protocol. We also ensured that all resources (media tracks, peer connections) are properly closed on call end to avoid memory leaks or ghost streams.

In summary, the results confirm that a WebRTC-based web application can deliver a full-featured video conferencing experience. The direct peer-to-peer media path provided clear audio and video, the auxiliary features (chat, screen share) enhanced the usability, and the application interface (thanks to React) was able to update in real time to reflect the ongoing state of the conference. Users found the experience similar to mainstream video chat tools for small meetings. One important discussion point is that while our implementation excels in scenarios of up to 3-4 participants, scaling beyond that would require careful consideration – either imposing limits or integrating a server-side media relay for efficiency. This leads into our concluding remarks, where we outline what could be done to extend this project further.

Conclusion

In this paper, we presented a complete design and implementation of a Video Conferencing App using WebRTC, following the IJRPR format for doublecolumn academic papers. The application fulfils the need for a browser-based, plugin-free video conferencing solution by leveraging WebRTC's capabilities for real-time peer-to-peer communication. By integrating technologies such as React, Node.js/Express, Socket.io, and Firebase, we built a system that not only enables direct audio/video calls between users, but also provides supporting features like text chat, screen sharing, and media toggling.

Our architectural design cantered on simplicity and performance: using a lightweight signaling server to set up calls and then handing off the bulk of data transmission to peer-to-peer links. This approach was validated through our results – users can communicate with low latency and high fidelity using only their web browsers, which demonstrates the power of WebRTC as a platform for real-time communications.

The project also serves as a proof-of-concept for how modern web development frameworks can accelerate the creation of complex applications. The use of React.js made the front-end highly responsive to state changes (crucial for reflecting live call status), while Node.js and Socket.io simplified the real-time backend logic.

We found that many challenges of building a video conferencing app (such as NAT traversal, media handling, etc.) are effectively handled by existing standards and services: for instance, STUN/TURN servers (including open ones or affordable cloud services) can be used to address network obstacles, and the WebRTC API itself abstracts the heavy lifting of media negotiation and transport. As a result, developers can focus on user experience and specific features rather than reinventing lower-level protocols.

Future Work: While the current system works well for its intended scope, there are several avenues for enhancement and research. Firstly, scalability can be improved: as discussed, a mesh P2P network may not support large conferences efficiently. A future version could integrate a Selective Forwarding Unit (SFU) – for example, using an existing open-source SFU like Jitsi Videobridge or mediasoup – to handle distribution of streams in larger meetings. This would shift the architecture from pure P2P to a hybrid model, allowing dozens of participants by offloading some bandwidth requirements to a server.

Secondly, implementing recording and playback functionality would be valuable for users who want to save meetings. This could be achieved by having one of the peers (or a headless browser on the server) use the MediaStream Recording API to capture the conference and then upload it to a server or cloud storage.

Thirdly, security and moderation features could be added: for example, end-to-end encryption is already inherent to WebRTC, but additional encryption layers or authentication tokens for joining meetings could be explored for greater security in corporate use cases.

In terms of user interface, features like dynamic layout (auto-switching the video layout based on who's speaking) and virtual background or noise suppression (leveraging WebRTC extensions or client-side processing) could further enhance the user experience, aligning our app with modern trends in video conferencing tools.

Another interesting direction is optimizing for different network conditions. Techniques such as bandwidth estimation and adaptation (already partly covered by WebRTC's congestion control) could be tuned or exposed to the application for better control. Also, supporting mobile devices and ensuring compatibility across all major browsers (Chrome, Firefox, Safari, etc.) will increase the app's versatility – some minor adjustments may be needed for iOS Safari which has specific quirks with WebRTC.

In conclusion, the development of the WebRTC-based video conferencing app confirms that real-time peer-to-peer communication can be achieved effectively with web standards. The work bridges theoretical concepts (like SDP signaling and ICE servers) with a practical application that end-users can run with a single click of a link. As remote collaboration continues to be important, such browser-centric solutions offer a promising route for accessible and secure communication tools. We hope that this project can serve as a reference or stepping stone for further innovations in the realm of WebRTC applications and encourage the adoption of open technologies in building communication platforms.

Acknowledgements: The authors express gratitude to Dr. D. Y. Patil Pratishthan's College of Engineering, Kolhapur for providing the resources and environment to conduct this project. We thank our faculty mentors and colleagues for their guidance, feedback, and support throughout the development of the application. The valuable insights from the WebRTC developer community and open-source projects like Jitsi also helped shape our understanding and implementation.

References :

- Irish Examiner. (2021). Surge in video conference app use in 2020. Retrieved from IrishExaminer.com irishexaminer.com
- [2] Jitsi. (n.d.). Jitsi Meet Open-Source Video Conferencing. Retrieved from Jitsi.org jitsi.org
- [3] MDN Web Docs. (2025). Signaling and video calling WebRTC API. Mozilla Developer Network developer.mozilla.org
- [4] MDN Web Docs. (2025). Introduction to WebRTC protocols. Mozilla Developer Network developer.mozilla.org
- [5] MDN Web Docs. (2025). MediaDevices.getDisplayMedia() Screen Capture API. Mozilla Developer Network developer.mozilla.org
- [6] MDN Web Docs. (2024). WebRTC data channels. Mozilla Developer Network developer.mozilla.org
- [7] Moffatt, J. (2024). Zoom Statistics & Usage: Most Recent Stats, Trends, and Data. Usebubbles Blog usebubbles.com
- [8] Parker, S. (2025). The Technology Behind Real-Time Web Communication: WebRTC. ProductivityLand productivityland.com
- [9] Venkatachalam, G. (2023). How we do Google Meet. Medium medium.com