

Mobile-MCP: Implementing the Model Context Protocol via Android’s Inter-Application Communication Mechanisms

Xiheng Li

Beijing University of Posts and Telecommunications

Chengcheng Wan

East China Normal University

Mengting He

The Pennsylvania State University

Linhai Song

SKLP, ICT, Chinese Academy of Sciences

Abstract

We argue that mobile AI assistants require a principled and secure interface for interacting with third-party applications. We present Mobile-MCP, an Android-native realization of the Model Context Protocol that enables applications to expose fine-grained, discoverable capabilities directly to large language models. By leveraging Android’s Intent mechanism, Mobile-MCP decouples AI assistants from tool implementations, allowing tools to evolve independently while preserving security boundaries. This work points toward a scalable ecosystem for building flexible and safe mobile AI assistants, moving beyond manually engineered integrations and unconstrained, GUI-based approaches.

1 Introduction

Mobile phones have become indispensable in modern daily life, supporting a wide range of activities such as entertainment [6, 12], communication [16, 23], office work [8, 11], and personal organization [3, 10]. With the advancement of large language models (LLMs), LLM-based AI assistants (*e.g.*, Apple Intelligence [2], Huawei XiaoYi [7]) have emerged as a key differentiator and selling point for smartphones. These assistants promise to help users accomplish complex tasks more efficiently through natural language interaction, thereby significantly enhancing the overall user experience.

Despite their rapid adoption, current architectures for mobile AI assistants remain fundamentally limited. Existing solutions generally fall into two categories. In the first category, AI assistants rely on coordinated APIs to interact with applications or tools [1, 2, 5, 17]. In many cases, both the assistants and their associated tools must be developed within the same company (*e.g.*, Qwen Chat [1]). This tightly coupled design significantly limits extensibility and hinders interoperability among AI assistants and tools from different developers. In the second category, AI assistants operate by observing the graphical user interface (GUI), often through screenshots, and relying on LLMs to infer appropriate actions, such as clicks or text input [4, 13–15, 21, 22]. These actions are then executed via OS-level accessibility and automation interfaces. Although more general, this approach effectively grants AI assistants the same privileges as end users, raising substantial security and privacy concerns. Moreover, it enables assistants to access all data exposed by the tools, potentially undermining the tools’ commercial interests.

We argue that neither approach offers a principled, secure, and scalable foundation for building mobile AI assistants. Instead, there should be a well-defined protocol between LLM-based assistants and mobile applications. Such a protocol would allow applications to expose selected functionalities in a controlled manner, while

enabling AI assistants to automatically understand and invoke these functionalities, without requiring custom development.

The Model Context Protocol (MCP) [9] is designed for this purpose, providing a standardized mechanism for tools to expose their functionalities and for LLMs to reason about and invoke them. However, MCP is primarily tailored to networked services and server-based tools. To date, there is no practical solution for natively using MCP on smartphones, nor for enabling LLMs—whether running on-device or backed by cloud computation—to safely and effectively interact with mobile applications.

In this paper, we propose Mobile-MCP, a practical design and implementation of MCP for Android smartphones that leverages Android’s existing inter-application invocation mechanism—the *Intent* framework. Specifically, Android applications can declare externally callable interfaces in their manifest files and provide textual descriptions that specify the functionality of each interface, as well as the semantics of its arguments and return values, analogous to how MCP servers expose services. LLM-based AI assistants can then query the Android package manager to discover all available MCP-style interfaces, interpret their descriptions, and supply appropriate parameters to invoke the corresponding APIs via the Intent mechanism, effectively mirroring how LLMs interact with standard MCP APIs.

Our approach eliminates the need for prior API coordination between AI assistants and tool applications. Instead, tools can be dynamically added to an AI assistant’s capability set and evolve independently, without requiring any modifications to the assistant’s code or internal prompts. Moreover, the approach allows each application to expose its capabilities in a fine-grained manner, enabling a clear distinction between interactions initiated by end users and those initiated by AI assistants. Together, these properties enable Mobile-MCP to support a new paradigm for building mobile AI assistants that can safely and flexibly interact with a broad ecosystem of third-party applications, rather than being limited to tightly integrated or manually engineered tools.

We have formally defined the data schema that governs how AI assistants communicate with tool applications, including how tools expose their APIs, how these APIs are discovered, how they are invoked, how requests are handled and responded to, and how results are interpreted. The formal specification is available at [20]. In this paper, we focus on the implementation aspects, and discuss how Mobile-MCP compatible AI assistants and tools should be implemented. In addition, we have developed a working prototype of Mobile-MCP to demonstrate its practical feasibility.

Overall, this paper makes the following contributions:

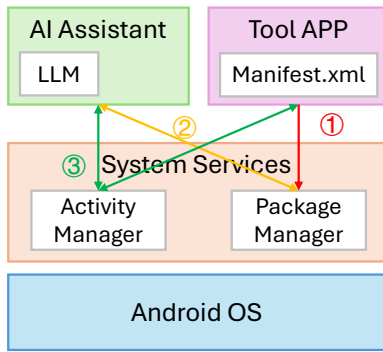


Figure 1: The workflow and architecture of Mobile-MCP.

- We propose Mobile-MCP, an Android-native realization of MCP that leverages the platform’s existing inter-application invocation mechanisms. Our design includes a formally defined protocol schema and a working prototype.
- We show that existing operating system mechanisms provide a strong foundation for building AI assistants, and we encourage future work that further explores this design space.

The source code [19] and a demonstration video [18] of Mobile-MCP have been publicly released for reference.

2 Design of Mobile-MCP

Mobile-MCP can be utilized via widely adopted programming languages in coding Android applications, such as Java and Kotlin. Its transport layer is based on Android’s Intent mechanism, and its data layer (message schema) is formally specified in [20].

As illustrated in Figure 1, applications participating in Mobile-MCP’s workflow assume one of the two roles: (1) AI assistants that leverage local or remote LLMs to discover Mobile-MCP compatible, external APIs available on the smartphone, select appropriate APIs, invoke them with suitable parameters, and interpret the results; and (2) tool applications that expose MCP-style APIs and handle incoming requests from AI assistants. The workflow consists of three stages: Mobile-MCP API registration, Mobile-MCP API discovery, and Mobile-MCP API invocation. Although multiple AI assistants and tool applications may coexist on a single device, we illustrate Mobile-MCP using a single assistant and a single tool application for clarity in Figure 1.

① *API Registration.* Each tool application exposes its MCP-style APIs by declaring a service component using a `<service>` element in its *AndroidManifest.xml*. This element includes three metadata entries: the application’s name, a textual description of its functionality, and the path to an XML file describing the application’s Mobile-MCP APIs. A dedicated intent filter marks the service as Mobile-MCP-enabled, allowing AI assistants to discover it. The XML file lists all available Mobile-MCP APIs. For each API, it includes a unique identifier, a natural-language description, and detailed descriptions of the API’s inputs and outputs. Each parameter is further annotated with its name, type, and semantic meaning. Upon installation, all declared Mobile-MCP APIs are automatically registered with the package manager via the *AndroidManifest.xml*.

② *API Discovery.* As illustrated in Figure 1, an AI assistant discovers available Mobile-MCP services by querying the package

manager using the `queryIntentServices(Intent, int)` method of class `PackageManager`. The `Intent` is constructed to match the intent filter used to mark Mobile-MCP services, ensuring that only Mobile-MCP-enabled services are returned. For each service, the package manager provides the application’s textual description, the list of exposed Mobile-MCP APIs, and the natural-language description of each API. Based on these descriptions and user intent, the AI assistant leverages an LLM to select an appropriate application, identify the relevant Mobile-MCP API, and determine suitable invocation parameters.

③ *API Invocation.* An AI assistant can invoke a Mobile-MCP API in two ways, depending on the interaction pattern. For repeated or stateful interactions, the assistant establishes a persistent binding to the tool application by invoking `bindService(Intent, ServiceConnection, int)` of class `Context`, where the `Intent` identifies the target Mobile-MCP service. Once bound, the assistant communicates with the tool via the `send(Message)` method of Android’s `Messenger` abstraction, with the `Message` to transmit the identifier of the Mobile-MCP API and parameter values encoded as JSON. For one-shot or infrequent interactions, the assistant instead invokes `startService(Intent)` of class `Context`, embedding the target tool, API identifier, and parameters directly in the `Intent`. As illustrated in Figure 1, in both cases the requests are processed by the activity manager, which enforces privilege checks and dispatches the requests to the tool application.

Demonstration. We implement a prototype system using a clock-in application as the tool. It exposes three Mobile-MCP APIs: clocking in for the current day, clocking in for a specified day, and querying clock-in records for a given day. These APIs respectively demonstrate Mobile-MCP calls without parameters, with parameters, and with both parameters and return values, covering the full range of Mobile-MCP invocation scenarios. The AI assistant leverages OpenAI’s LLM via its public APIs. The demonstration video in [18] shows that Mobile-MCP successfully supports all API invocation patterns. We further evaluate the performance of each step and find that over 99% of the execution time is spent interacting with the LLM, indicating that API invocations via Mobile-MCP do not constitute a performance bottleneck for AI assistants.

3 Discussion and Conclusion

In this paper, we present Mobile-MCP, an Android-native implementation of MCP that leverages Android’s *Intent* framework for cross-application communication. We formally define the message schema used between AI assistants and tool applications. We also implement a prototype to demonstrate that our design is practical and feasible. In the future, we plan to extend this framework by integrating additional components that AI assistants commonly require, such as long-term memory support and on-device LLM inference. We also plan to explore how to implement the A2A protocol using the same cross-application invocation mechanism.

More broadly, this work demonstrates the value of reusing and extending existing operating system mechanisms when designing new types of applications, such as AI assistants. In particular, we aim to investigate whether Android’s existing privilege and permission system, especially features associated with the `Intent` framework, is sufficient to manage and restrict AI assistant behavior, or whether new abstractions and system mechanisms are necessary.

References

- [1] Alibaba Cloud. Qwen, 2026. <https://qwen.ai/home>.
- [2] Apple Inc. Apple intelligence, 2026. <https://www.apple.com/apple-intelligence/>.
- [3] Atlassian. Trello, 2024. <https://trello.com/>.
- [4] Droidrun. Droidrun, 2026. <https://droidrun.ai/>.
- [5] Google. Discover what google assistant is, 2026. <https://assistant.google.com/>.
- [6] GooglePlay. Youtube, 2026. https://play.google.com/store/apps/details?id=com.google.android.youtube&hl=en_US.
- [7] Huawei Corporation. HarmonyOS 5, 2026. https://en.wikipedia.org/wiki/HarmonyOS_5.
- [8] Microsoft. Outlook for ios and android, 2026. <https://www.microsoft.com/en-us/microsoft-365/outlook-mobile-for-android-and-ios>.
- [9] Model Context Protocol Project. Model context protocol, 2026. <https://modelcontextprotocol.io/docs/getting-started/intro>.
- [10] mylifeorganized.net. Mylifeorganized, 2026. <https://www.mylifeorganized.net/>.
- [11] Slack Technologies, LLC. Slack, 2026. <https://slack.com/downloads/android>.
- [12] TikTok. Tiktok, 2026. <https://www.tiktok.com/download>.
- [13] unimobile Contributors. Zhixing: A mobile agent development framework, 2026. <https://github.com/apbocldueo/unimobile>.
- [14] H. Wen, Y. Li, G. Liu, S. Zhao, T. Yu, T. J.-J. Li, S. Jiang, Y. Liu, Y. Zhang, and Y. Liu. Autodroid: Llm-powered task automation in android. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (MobiCom '24)*, Washington D.C., DC, USA, 2024.
- [15] H. Wen, S. Tian, B. Pavlov, W. Du, Y. Li, G. Chang, S. Zhao, J. Liu, Y. Liu, Y.-Q. Zhang, and Y. Li. Autodroid-v2: Boosting slm-based gui agents via code generation. In *Proceedings of the 23rd Annual International Conference on Mobile Systems, Applications and Services (MobiSys '25)*, Anaheim, California, US, 2025.
- [16] WhatsApp LLC. Whatsapp, 2026. <https://www.whatsapp.com/>.
- [17] Xiaomi. Xiaomi HyperOS, 2026. <https://www.mi.com/global/hyperos/>.
- [18] Xiheng Li, Mengting He, Chengcheng Wan and Linhai Song. The demonstration video of Mobile-MCP, 2026. <https://youtu.be/Bc2LG3sR1NY>.
- [19] Xiheng Li, Mengting He, Chengcheng Wan and Linhai Song. Mobile-MCP, 2026. <https://github.com/system-pclub/mobile-mcp>.
- [20] Xiheng Li, Mengting He, Chengcheng Wan and Linhai Song. Mobile-mcp specification, 2026. https://github.com/system-pclub/mobile-mcp/blob/main/spec/mobile-mcp_spec_v1.md.
- [21] J. Ye, X. Zhang, H. Xu, H. Liu, J. Wang, Z. Zhu, Z. Zheng, F. Gao, J. Cao, Z. Lu, J. Liao, Q. Zheng, F. Huang, J. Zhou, and M. Yan. Mobile-agent-v3: Fundamental agents for gui automation, 2025.
- [22] C. Zhang, Z. Yang, J. Liu, Y. Li, Y. Han, X. Chen, Z. Huang, B. Fu, and G. Yu. Appagent: Multimodal agents as smartphone users. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*, Yokohama, Japan, 2025.
- [23] Zoom Communications, Inc. Zoom meetings, 2026. <https://www.zoom.com/en/products/virtual-meetings/>.