Voice-Interactive Washing Machine Based on R2L100 SOM: Edge and Cloud AI in Harmony

By: Eoxys Systems India Pvt. Ltd.

Home appliances are no longer mechanical boxes with buttons and timers. In today's connected world, a washing machine can do much more than clean clothes. It can listen, respond, guide, and even resolve customer doubts all through natural voice interaction.

Imagine saying, ***"Start quick wash",*** and the machine immediately begins its cycle. Or asking, ***"How long does delicate wash take?",*** and hearing an instant spoken response. Even complex troubleshooting ***"What should I do if the drum is unbalanced?"***, could be answered conversational -ly, without opening a manual.

Fig 1: Voice-Interactive Washing Machine

This vision is becoming a reality with the **R2L100 SOM (**called as **Vision ML SOM)** which has **RZ/V2L MPU** from Renesas, powered by its **DRP-AI accelerator**. By embedding a complete voice-interactive AI system directly into a washing machine, we combine **edge intelligence for reliability** with **cloud intelligence for depth and reasoning**.

The Edge Advantage

At the heart of a voice-interactive washing machine is the ability to operate **offline, in real time, with low latency**. The **edge-only AI stack** achieves exactly that.

The system begins with **wake word detection and voice activity detection**, allowing the machine to remain in low-power mode until the user says **"Hello Washer."** Once activated, audio is processed through **automatic speech recognition (ASR)** using a compact model like ***Whisper-tiny***. This transforms the spoken words into text on the device itself.

Next, **natural language understanding (NLU)** models, such as ***TinyBERT*** **or** ***DistilBERT***, interpret user intent. If the command is operational **"Pause wash" or "Start spin cycle"** it is routed to the execution layer, directly controlling the washing machine hardware. For queries like ***"How long is quick wash?",*** the system runs a local **FAQ retrieval pipeline** powered by ***MiniLM*** *embeddings* and a small knowledge base of washing machine documentation.

Finally, the machine replies with speech, using a lightweight **text-to-speech (TTS)** engine like ***PicoTTS*** **or** ***eSpeak NG***. All of this runs comfortably within the R2L100 SOM's **2 GB RAM (with ~1.5 GB available for AI tasks)**, thanks to model quantization and DRP-AI acceleration.

The result: **an appliance that can understand, respond, and act without needing the cloud.**

Why Not Just Cloud?

Cloud AI, of course, offers impressive muscle. Large-scale models like ***LLaMA-2***, ***GPT-4***, or ***Claude*** can carry out advanced reasoning, multi-turn dialogue, and provide deeply personalized support. But these models require gigabytes of memory and strong connectivity constraints that make them impractical for an always-on household device.

On the R2L100 SOM, the models from the **LLaMA** family simply don't fit comfortably. Even heavily quantized, they demand more RAM and compute than the hardware can reliably provide. Lightweight models like ***TinyBERT*** **and** ***MiniLM*** strike the right balance compact enough for edge deployment, yet accurate enough for everyday voice commands and FAQs.

That said, the cloud still has its place. For troubleshooting beyond the scope of local knowledge say, diagnosing unusual error codes or providing long-term maintenance recommendations the system can offload queries to the cloud. Here, larger LLMs can process the request, access cloud-hosted vector databases, and send back answers.

The Hybrid Model – Best of Both Worlds

The most practical architecture is **hybrid**.

Everyday tasks like ***"Start wash," "Pause cycle,"*** **or** ***"What's the time left?"*** are handled locally, with blazing speed and privacy. Meanwhile, advanced queries ***"What's the optimal detergent for stained cotton?"*** can be sent to the cloud for deeper reasoning.

In this setup, the washing machine is never "dumb" without internet. It always remains functional, interactive, and responsive thanks to edge AI, while still gaining intelligence from the cloud when connected.

The Engineering Perspective

From a design standpoint, the R2L100 SOM proves ideal. Its **DRP-AI SDK** allows efficient acceleration of AI models, cutting latency and improving inference speeds. With careful selection of lightweight models ***Whisper-tiny for ASR (~200 MB), TinyBERT for NLU (~200 MB), MiniLM for embeddings (~100 MB), PicoTTS (~50 MB)*** the total AI footprint remains around **500–600 MB**, well within the 1.5 GB available.

This leaves ample headroom for OS overhead and washing machine control functions. And because all models are quantized and hardware-accelerated, response times stay comfortably real-time.

Recommendations for OEMs

The path forward is clear. Start with an **edge-only MVP** a washing machine that can recognize wake words, execute commands, and answer basic FAQs offline. Then, extend to **hybrid mode**, where cloud AI enriches the experience with complex reasoning, richer dialogue, and remote updates.

A key step is building the FAQ database: converting washing machine manuals and service documents into a structured text knowledge base, embedding them with MiniLM, and hosting them locally for retrieval. This enables the machine to not just "follow orders," but also "answer questions."

Note: Building more functionality on top of AI models for handling more complex workflows (retrieve + generate + Vector search + interact with onboard sensors + execute and schedule), the RAM requirements and fitting within 2GB need to be studied and validated in POC model.

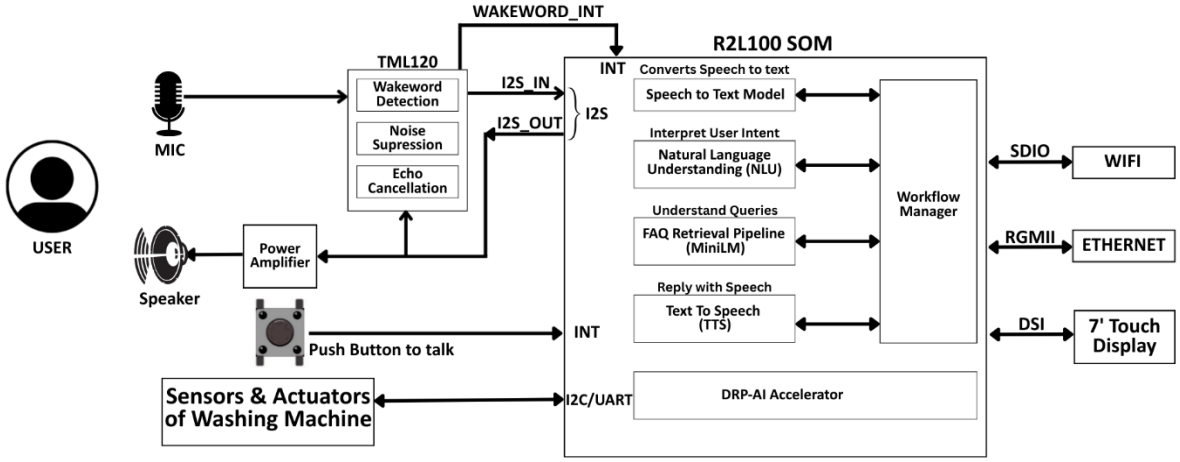

Fig 2: Voice-Interactive System Block Diagram

This block diagram shows the how Vision ML SOM, together with TML120, based product HW architecture along with the agentic AI models for a voice-interactive platform. By combining audio front-end processing with speech understanding, intent recognition, and system control, this design on Vision ML SOM offers customers a ready-to-adopt strategy in their product architecture. The product teams of OEMs can start from this reference architecture reducing integration effort, cutting development time, and accelerating the path to market.

Conclusion – A Conversational Future for Appliances

The washing machine is evolving from a simple appliance into a **voice-interactive partner**. By leveraging **RZ/V2L's DRP-AI acceleration** and a hybrid edge–cloud approach, it is possible to build machines that respond instantly, operate reliably offline, and still benefit from the depth of cloud AI when needed.

For consumers, this means appliances that are easier to use, more intuitive, and truly conversational. For manufacturers, it means a path to differentiation building products that stand out not just for performance, but for intelligence.

The future of home appliances is one where **AI listens, understands, and acts**. And with the R2L100 SOM at the core, that future is already within reach.