Bamboo Agent Update

Tadashi Maeno (BNL)

Contents

1. Bamboo V1 findings
2. Bamboo V2 plan
3. Scalable Intelligence with Bamboo and PanDA
4. Misc

Agentic AI 1/2
– MCP Tools + Knowledge Base + LLM –

- Start with a raw, high-level objective in NL
- RAG to retrieve relevant documentation, operational knowledge, and past experience
- MCP tools to obtain dynamic information from the live system
- Let LLM drive a ReAct loop
  - Retrieve context
  - Select and invoke tools
  - Observe results
  - Reason about the next step
- No predefined reasoning path: prompts and RAG influence LLM, which decides what to retrieve, what to do, and when to stop
- Hope that LLM is smart enough to discover the right reasoning path and achieve underlying scientific/operational goal

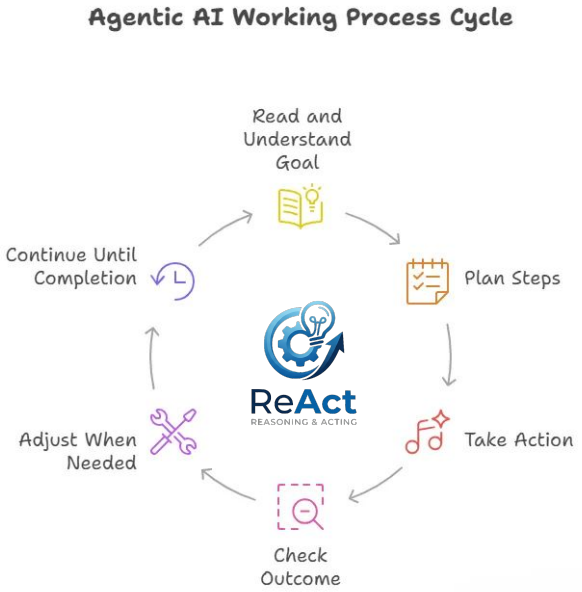

Agentic AI 2/2
– MCP Tools + Knowledge Base + LLM –

- Limitation
  - The performance heavily depends on the quality of the retrieved context
  - Optimizing RAG to consistently retrieve the right information is not straightforward
  - The same objective can lead to different reasoning paths
  - Difficult to guarantee that the agent follows the correct system behavior
  - Limited reproducibility and accountability
  - Important system knowledge may not be captured in documentation or the KB

Bamboo V1: Graph-guided Reasoning 1/2
– From free-form ReAct to structured reasoning –

- Graph becomes the backbone of the ReAct loop
  - Explicitly represents reasoning steps, dependencies, and decisions
  - Guides the agent toward relevant context and actions
- Graph accumulation through interaction with experts
  - Capture operational knowledge and expert reasoning
  - Preserve successful troubleshooting paths
  - Gradually build organizational knowledge
- Improved reproducibility
  - Same problem can follow an established reasoning path
  - Decisions and intermediate steps become inspectable
- Enhanced accountability
  - Easier to understand why a particular action was taken
  - Reasoning path can be reviewed after the fact

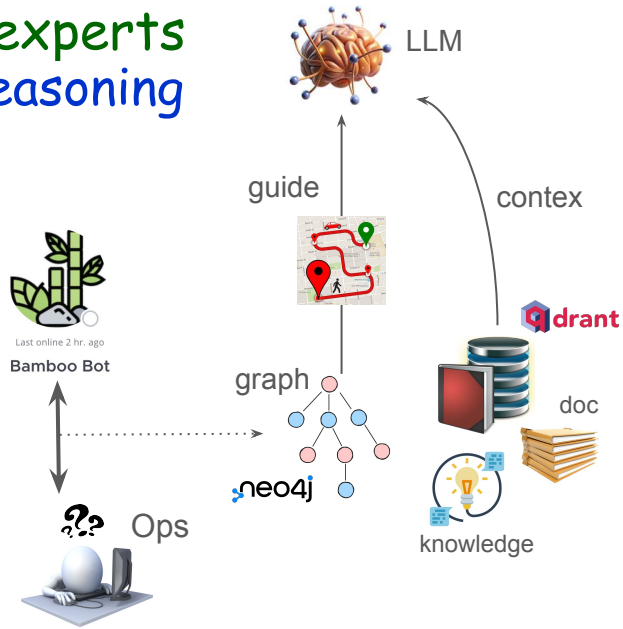

Bamboo V1: Graph-guided Reasoning 2/2
– From free-form ReAct to structured reasoning –

- Bottleneck: knowledge acquisition
  - Expert knowledge is difficult to extract at scale
  - Email / mailing-list discussions contain too much noise
    - Incomplete context
    - Unrelated discussion
    - Contradictory or obsolete information
    - Difficult to identify the actual reasoning chain
- Issues
  - Building a useful, high-quality graph requires requires continuous expert involvement and curation
  - Knowledge retrieval is still challenging and hard to optimize

Bamboo V2: Code-Grounded Graph 1/2
– From expert knowledge to system ground truth –

- The graph is generated from target system's source code
  - Code provides a definitive description of system behavior
  - Eliminates much of the ambiguity and noise of human-generated knowledge
- Graph as a system map
  - Junction points provide observable anchors
  - Traversal connects anchors to reconstruct complete execution traces
  - Traces expose actual system path
- Graph-guided multi-hop reasoning
  - LLM does not need to infer the system structure from scratch or documentation
  - Graph navigates reasoning to real paths in the system
  - Retrieved traces provide precise, execution-relevant context

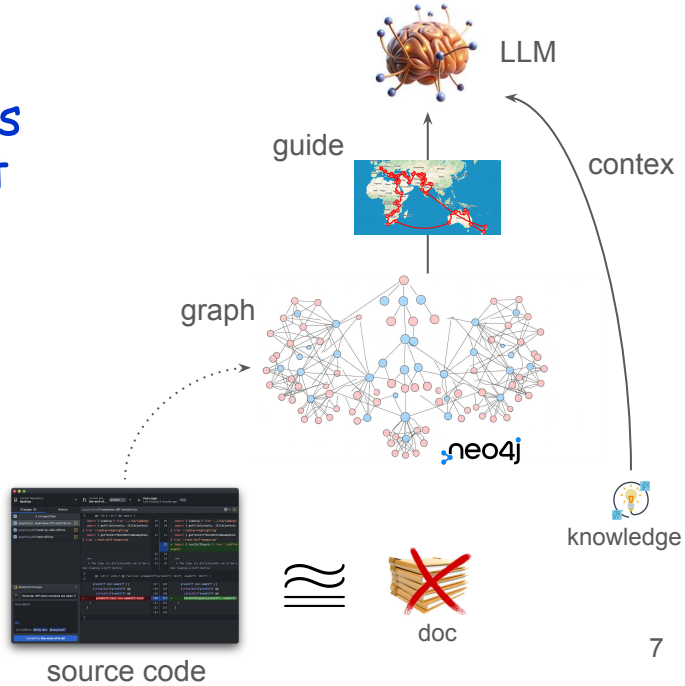

Bamboo V2: Code-Grounded Graph 2/2
– From expert knowledge to system ground truth –

- Goals
  - Much more accurate reasoning
  - Better reconstruction of what actually happened
  - More reliable diagnosis of root causes
  - Correct recovery actions become possible based on precise reasoning
  - Less dependence on manually curated documentation or expert explanations

V2 Status and Plans

- Massive PanDA codebase refactoring for type annotations
  - Established a code-level foundation for graph extraction

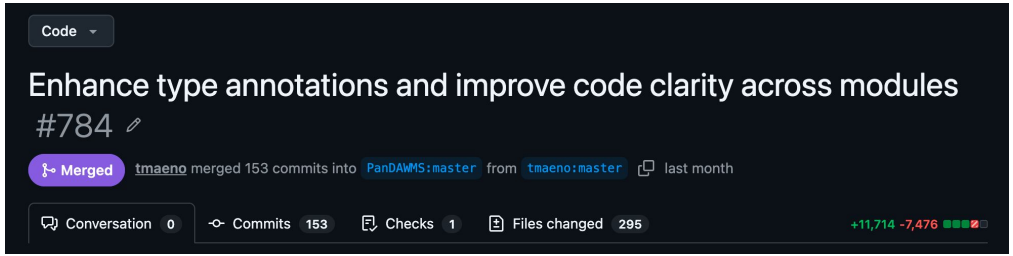

- AST-based graph extractor implemented
  - Automatically identifies important system entities and relationships
  - 140k lines of code → 500 subjects → 4k junction points
- Querier and analyzer under development
  - Self-driven traversal using graph-derived traces, logs, and live system info via API/MCP
  - Targeting end-to-end capability in ~1 month

Scalable Intelligence with Bamboo and PanDA

- Run a Bamboo agent as a PanDA job
- Package Bamboo as a self-contained container
  - Bamboo runtime, tools, and dependencies
  - LLM/model and database snapshot provided as job inputs
- Scale agents across distributed GPU resources
- Enable on-demand large-scale parallel investigation
  - Independent problem analysis sessions by separate agents
  - Result aggregation after Execution
- ~1k sessions per day

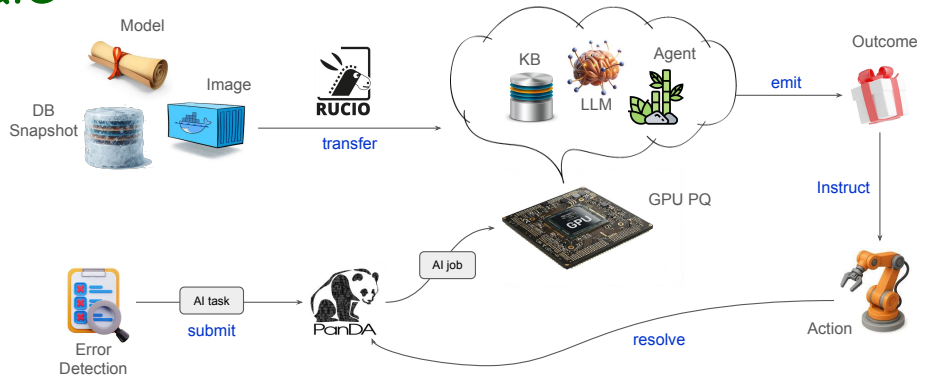

Job name: user.tmaeno.ecb32138-4245-409a-9cea-d625b7030512/.7239836631

PanDA ID	Owner	Task ID	Status	Transformation	Created Last modified	Time to start Duration [dd:hh:mm:ss]	Site	Cores	Priority	Attempt
7239836631	Tadashi Maeno	51796519	finished	runGen-00-00-02	2026-07-30 14:28:33 2026-07-30 14:54:38	0:0:14:18 0:0:11:44	CERN-GPU	8	1000	1
Datasets:	In: user.tmaeno:user.tmaeno.bamboo.sandbox.C14B5BF8-AF85-4F36-B5DA-0F5B38FCE1FC Out: user.tmaeno.ecb32138-4245-409a-9cea-d625b7030512.log.679203752.679203752									
Files summary:	input: 1, size: 23.48GB, #events: 0; output: 0; pseudo_input: 0; log: 1									

Across REDWOOD, MOSAIC, and LDRD

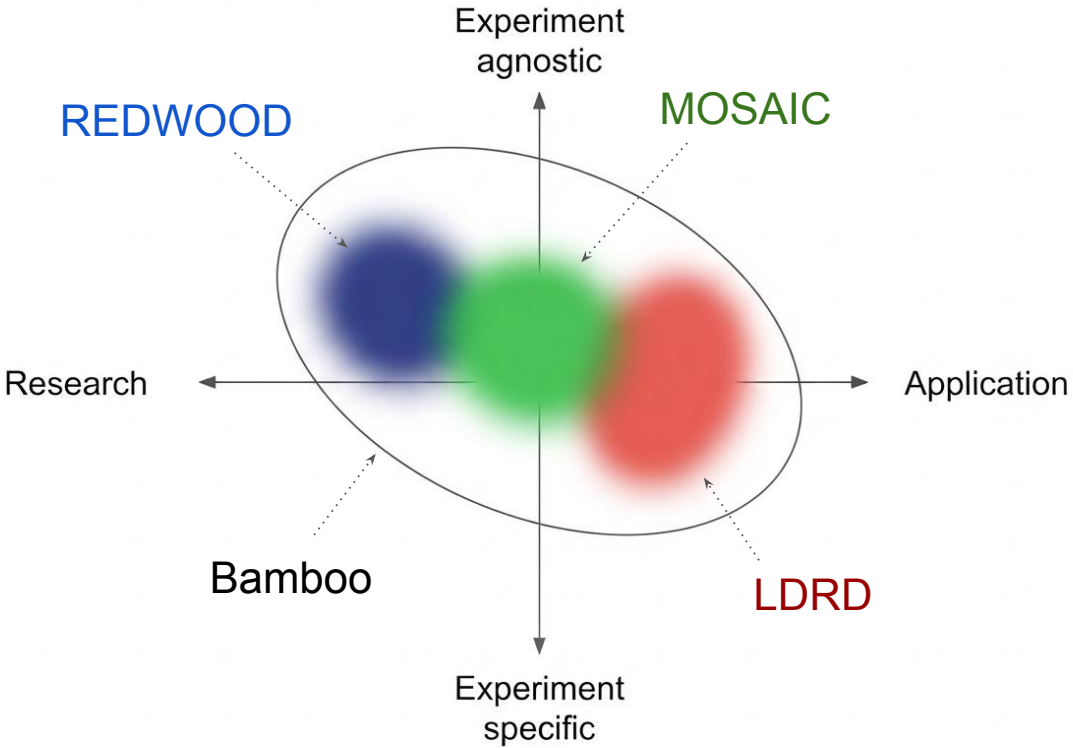

REDWOOD → Code-Grounded Graphs and Graph-Guided Reasoning

New Documentation with Starlight
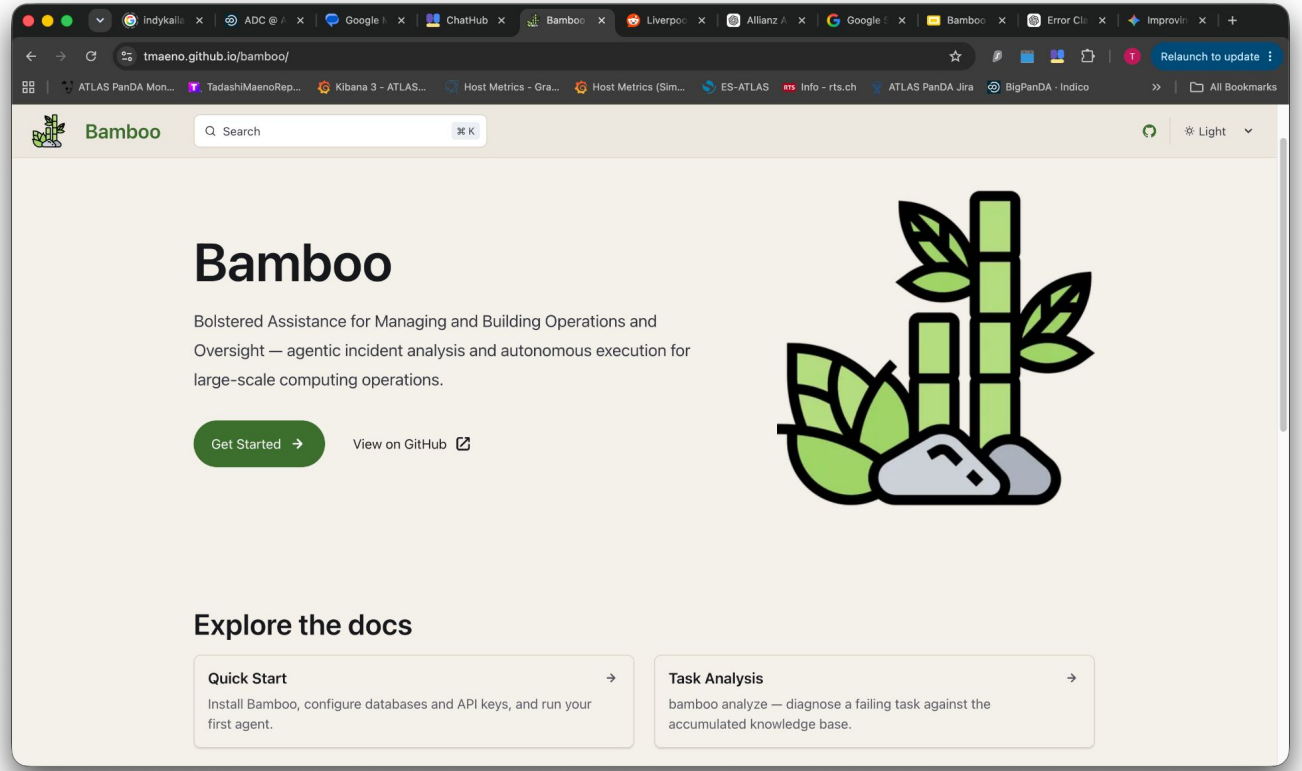
Link