

# Apertures: Coordinated Partial Evaluation for Distributed Computation

Alexander Towell  
Department of Computer Science  
Southern Illinois University Edwardsville  
atowell@siue.edu

## Abstract

We present *apertures*, a coordination mechanism for distributed computation based on partial evaluation with explicit holes. Apertures (denoted `?variable`) represent unknown expressions that enable pausable and resumable evaluation across multiple parties. Unlike security-focused approaches, we make no claims about privacy preservation—apertures leak information through program structure and evaluation results. Instead, we position apertures as a lightweight coordination primitive for scenarios where parties can share computation structure but need to control when and where data is introduced. We demonstrate practical coordination patterns including multi-party progressive refinement and local computation workflows, where untrusted servers can optimize expressions without seeing private inputs or outputs. Our C++ implementation shows apertures add minimal overhead (2-5%) compared to direct evaluation while enabling novel distributed computation patterns. We are explicit about limitations: apertures are unsuitable for adversarial settings and should not be used when information leakage is unacceptable.

## 1 Introduction

Distributed computation often requires coordination between parties with different capabilities and trust levels. A common pattern involves untrusted servers with computational resources and clients with private data. Current solutions typically fall into two extremes: full encryption with massive overhead, or complete data sharing with no privacy.

We introduce *apertures*, a coordination mechanism that enables a middle ground through partial evaluation. An aperture, written `?variable`, is a hole in an expression representing an unknown value or subexpression. Programs with apertures can be partially evaluated, optimized, and transformed by parties that don't know the aperture values. When apertures are eventually filled, evaluation resumes from the partially evaluated state.

**Critical disclaimer:** Apertures are *not* a security mechanism. They leak information through program structure, evaluation patterns, and algebraic relationships. We make no privacy claims. Instead, apertures provide a coordination primitive for distributed systems where parties need to control the timing and location of data introduction while sharing computational structure.

Consider this local computation pattern:

```
1 ;; Server optimizes without data or results
2 (let ((plan (server-optimize ?query)))
3   (client-execute plan ?local-data))
4 ;; Results stay local
```

The server can optimize the query structure without seeing the actual query parameters or the local data. The client benefits from server-side optimization while keeping both inputs and outputs private. However, the server learns the query structure and optimization patterns, which may reveal information about the workload.

Our contributions:

- A simple calculus for partial evaluation with holes
- Coordination patterns for distributed computation

- Explicit characterization of when apertures should *not* be used
- Performance evaluation showing 2-5% overhead for coordination

## 2 The Aperture Language

### 2.1 Core Syntax

We define a minimal Lisp-like language with apertures:

$$e ::= v \mid x \mid ?h \mid (e e^*) \mid (\lambda (x^*) e) \quad (1)$$

$$\mid (\text{let } ((x e)^*) e) \mid (\text{if } e e e) \quad (2)$$

$$v ::= \text{nil} \mid n \mid s \mid \text{bool} \mid \text{list} \quad (3)$$

Where  $v$  ranges over values,  $x$  over variables,  $h$  over hole identifiers, and  $n, s$  over numbers and strings respectively. The key addition is  $?h$ , representing an aperture (hole).

**Terminology clarification:** While we write `?variable` for readability, apertures are not restricted to atomic values. An aperture can be filled with any expression, making them "unknown expressions" in the partial evaluation sense.

### 2.2 Evaluation Semantics

We define evaluation contexts and partial evaluation:

$$E ::= [\cdot] \mid (E e^*) \mid (v v^* E e^*) \quad (4)$$

$$\mid (\text{let } ((x v)^* (x E) b^*) e) \quad (5)$$

$$\mid (\text{if } E e e) \quad (6)$$

Standard evaluation rules apply, with apertures blocking evaluation:

$$E[?h] \rightarrow_p E[?h] \quad (\text{aperture blocks}) \quad (7)$$

$$E[(+ n_1 n_2)] \rightarrow n_1 + n_2 \quad (\text{arithmetic}) \quad (8)$$

$$E[(+ ?h n)] \rightarrow_p (+ ?h n) \quad (\text{partial}) \quad (9)$$

The key insight: evaluation proceeds until blocked by apertures, creating partially evaluated expressions that preserve computation structure.

### 2.3 Hole Filling

Apertures are filled via substitution:

$$\text{fill}(e, h, e') = e[e'/?h]$$

After filling, evaluation can resume:

$$\text{eval}(\text{fill}(e, h, v)) \rightarrow \text{eval}(e[v/?h])$$

Multiple apertures can be filled incrementally, enabling progressive refinement across multiple parties.

## 2.4 Algebraic Simplification

During partial evaluation, we apply algebraic simplifications that preserve apertures:

$$(* 0 ?x) \rightarrow_p 0 \tag{10}$$

$$(+ 0 ?x) \rightarrow_p ?x \tag{11}$$

$$(* 1 ?x) \rightarrow_p ?x \tag{12}$$

$$(\text{if true } e_1 \ e_2) \rightarrow_p e_1 \tag{13}$$

These simplifications reduce expression size while maintaining evaluation semantics. However, they also leak information—for example, if  $(* ?x e) \rightarrow 0$  where  $e$  is a known non-zero constant, this reveals that  $?x = 0$ .

## 3 Coordination Patterns

Apertures enable several coordination patterns for distributed computation:

### 3.1 Progressive Refinement

Multiple parties can incrementally fill apertures:

```
1 ;; Initial expression with multiple holes
2 (let ((query (select ?table
3             (where (and (> ?col1 ?val1)
4                         (< ?col2 ?val2))))))
5
6 ;; Party A knows table and column names
7 (fill query {table: "users",
8             coll: "age",
9             col2: "salary"})
10
11 ;; Party B knows value constraints
12 (fill query {val1: 18, val2: 100000})
13
14 ;; Execute when all holes filled
15 (execute query))
```

Each party contributes their knowledge without seeing others' contributions until execution.

### 3.2 Local Computation Pattern

The most defensible use case: untrusted optimization with local execution:

```
1 ;; Client has private data
2 (define private-data (load-sensitive))
3
4 ;; Server optimizes generic computation
5 (define (optimize-computation expr)
6   ;; Algebraic simplification
7   ;; Common subexpression elimination
8   ;; Partial evaluation of known parts
9   (simplify expr))
10
11 ;; Client sends structure, not data
```

```

12 (define template
13   (sum (map (lambda (x)
14             (* ?weight (f x)))
15             ?data)))
16
17 ;; Server optimizes without seeing data
18 (define optimized
19   (server-optimize template))
20 ;; Returns: (let ((w ?weight))
21 ;;          (sum (map (lambda (x)
22 ;;                    (* w (f x)))
23 ;;                    ?data)))
24
25 ;; Client evaluates locally
26 (local-eval (fill optimized
27             {weight: 0.5,
28              data: private-data}))

```

Benefits:

- Server CPU used for optimization
- Data never leaves client
- Results stay local
- Reduced leakage channels

Limitations:

- Server learns computation structure
- Optimization patterns may leak workload information
- No protection against malicious servers

### 3.3 Speculative Compilation

Servers can compile multiple specializations:

```

1 ;; Template with configuration hole
2 (define (process-data config data)
3   (case ?config
4     [(fast) (quick-process data)]
5     [(accurate) (slow-process data)]
6     [(balanced) (hybrid-process data)]))
7
8 ;; Server compiles all branches
9 (define compiled
10  (compile-all-paths template))
11
12 ;; Client selects path at runtime
13 (execute compiled {config: 'accurate})

```

The server prepares optimized code paths without knowing which will be used.

## 4 Implementation Highlights

Our C++ implementation provides:

- S-expression parser with aperture syntax

- Partial evaluator with algebraic simplification
- Hole tracking and filling mechanism
- Fluent API for natural expression construction

Key design choices:

- Shared pointer memory management
- Visitor pattern for evaluation
- Variant-based value representation
- Stack-based evaluation with continuation frames

## 4.1 Performance Characteristics

We measured overhead compared to direct evaluation:

Operation	Direct ( $\mu s$ )	With Apertures ( $\mu s$ )	Overhead
Arithmetic (1000 ops)	127	131	3.1%
List processing	89	93	4.5%
Conditionals	43	44	2.3%
Lambda application	156	164	5.1%

Apertures add minimal overhead (2-5%) while enabling distributed coordination. The overhead comes from:

- Hole checking during evaluation
- Maintaining partial evaluation state
- Expression reconstruction after simplification

Partial evaluation can actually *improve* performance when expressions are reused with different aperture values, as the partially evaluated form serves as a template.

## 5 Case Study: Query Optimization

Consider a distributed database scenario where a client wants to run complex analytical queries on private data. Traditional approaches either:

1. Send data to server (privacy loss)
2. Send query to data (lose server optimization)
3. Use homomorphic encryption (1000× overhead)

With apertures, we enable a fourth option:

```

1 ;; Client query template
2 (define query-template
3   (select ?private-table
4     (where (and (> profit ?threshold)
5                (in region ?regions))))
6   (group-by department)

```

```

7      (having (> (count *) ?min-size))))
8
9      ;; Server optimizes structure
10     (define optimized-plan
11       (server-optimize query-template))
12     ;; Server can:
13     ;; - Reorder predicates
14     ;; - Plan aggregation strategy
15     ;; - Choose join algorithms
16     ;; - Prepare indexes
17
18     ;; Client executes locally
19     (define results
20       (local-execute
21         (fill optimized-plan
22           {private-table: my-data,
23            threshold: 1000000,
24            regions: ["NA", "EU"],
25            min-size: 10})))

```

The server provides optimization expertise without seeing:

- Actual table data
- Specific threshold values
- Selected regions
- Query results

However, the server learns:

- Query structure and complexity
- Types of operations performed
- Optimization applicability

This tradeoff is acceptable when:

- Query patterns are not sensitive
- Optimization benefits outweigh leakage
- Server is honest-but-curious, not malicious

## 6 Related Work

**Partial evaluation** [2]: Apertures extend partial evaluation with explicit holes for multi-party coordination. Unlike traditional partial evaluation which requires all static values upfront, apertures allow incremental filling.

**Staged computation** [4]: Multi-stage programming separates computation phases. Apertures provide a simpler mechanism focused on coordination rather than performance.

**Homomorphic encryption** [1]: FHE provides cryptographic privacy but with 1000-10000× overhead. Apertures trade privacy for practical performance.

**Secure multi-party computation** [5]: MPC protocols provide strong security guarantees through cryptographic protocols. Apertures offer a lightweight alternative for non-adversarial settings.

**Information flow control** [3]: Type systems can track information flow statically. Apertures provide dynamic coordination without static guarantees.

## 7 Limitations: When NOT to Use Apertures

We explicitly enumerate scenarios where apertures are inappropriate:

### 7.1 Adversarial Settings

Apertures provide **no protection** against malicious parties. An adversary can:

- Infer aperture values from program structure
- Use algebraic relationships to constrain possibilities
- Exploit side channels in evaluation patterns
- Inject malicious code during partial evaluation

### 7.2 High-Sensitivity Data

When data sensitivity is high, apertures leak too much:

- Medical records: Disease patterns visible in query structure
- Financial data: Trading strategies revealed by operations
- Personal data: Behavioral patterns in access patterns

### 7.3 Regulatory Compliance

Apertures cannot provide guarantees required by:

- GDPR: No verifiable data protection
- HIPAA: Insufficient privacy safeguards
- Financial regulations: No audit trail of data access

### 7.4 Information Leakage Examples

Consider this expression:

```
1 (if (> ?age 65)
2   (discount 0.2)
3   (if (< ?age 18)
4     (discount 0.1)
5     (no-discount)))
```

Even without knowing `?age`, observers learn:

- System has age-based discrimination
- Exact discount tiers
- Business logic for pricing

Or this query:

```
1 (select users
2   (where (and (= status ?s)
3             (> last-login ?date)
4             (in department ?dept))))
```

Structure alone reveals:

- Organization has departments
- Tracks login times
- Has user statuses
- Performs temporal queries

## 7.5 When Apertures ARE Appropriate

Apertures work well when:

- Computation structure is not sensitive
- Parties are honest-but-curious, not adversarial
- Performance matters more than perfect privacy
- Coordination benefits outweigh leakage risks
- Local computation can limit result leakage

Examples of good use cases:

- Scientific computing with public algorithms
- Optimization of non-sensitive business logic
- Collaborative debugging and testing
- Educational and demonstration systems

## 8 Conclusion

Apertures provide a lightweight coordination mechanism for distributed computation through partial evaluation with holes. We make no security claims—apertures leak information through structure and evaluation patterns. Instead, they offer a practical tool for scenarios where parties need to coordinate computation while controlling when and where data is introduced.

The key insight is that many real-world scenarios don't require cryptographic privacy but do need coordination primitives. Apertures fill this gap with minimal overhead (2-5%) while enabling patterns like local computation with remote optimization.

Future work includes:

- Quantifying information leakage formally
- Combining apertures with oblivious algorithms
- Exploring apertures in specific domains (databases, ML)
- Building practical systems using aperture coordination

- **LLM-based automatic hole inference:** Integrating large language models to automatically fill apertures based on context. The hole name and surrounding computation could provide sufficient context for an LLM to infer appropriate values or implementations. For example, `?load-customer-data` could be automatically filled with an appropriate SQL query, or `?aggregation-fn` could be filled based on domain context (financial  $\rightarrow$  sum, health  $\rightarrow$  average). This would enable a new paradigm of *approximate program specification* where developers express intent through named holes and AI completes the implementation.

We hope apertures contribute to the discussion of practical coordination mechanisms that acknowledge rather than obscure their limitations. Not every distributed computation needs cryptographic security, but all need clear understanding of their tradeoffs.

## References

- [1] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of STOC*, pages 169–178. ACM, 2009.
- [2] Neil D Jones, Carsten K Gomard, and Peter Sestoft. Partial evaluation and automatic program generation. Prentice Hall, 1993.
- [3] Andrei Sabelfeld and Andrew C Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.
- [4] Walid Taha and Tim Sheard. Multi-stage programming: Axiomatization and type safety. *Proceedings of ICFP*, 1997.
- [5] Andrew C Yao. Protocols for secure computations. In *Proceedings of FOCS*, pages 160–164. IEEE, 1982.