Designing a Single Chip Chess Grandmaster While Knowing Nothing About Chess*

Feng-hsiung Hsu
IBM T. J. Watson Research Center
fhh@us.ibm.com

(*Well, I did know close to nothing about chess...)

Outline

- Project History
- Problem Descriptions
- Design Philosophy
- Chip Architecture
- System Performance
- World Chess Champion on a PC?

Project History

- Started as a one-person student project at Carnegie Mellon ('85)
- First Grandmaster Level Chess Machine (Deep Thought, '88)
- Became an IBM Project in '89
- New Chip Design in Full Swing from '92
- First Release Chip in '95 ('96 Match)
- Second Release Chip in '97 ('97 Rematch)

Problem Descriptions (Speed)

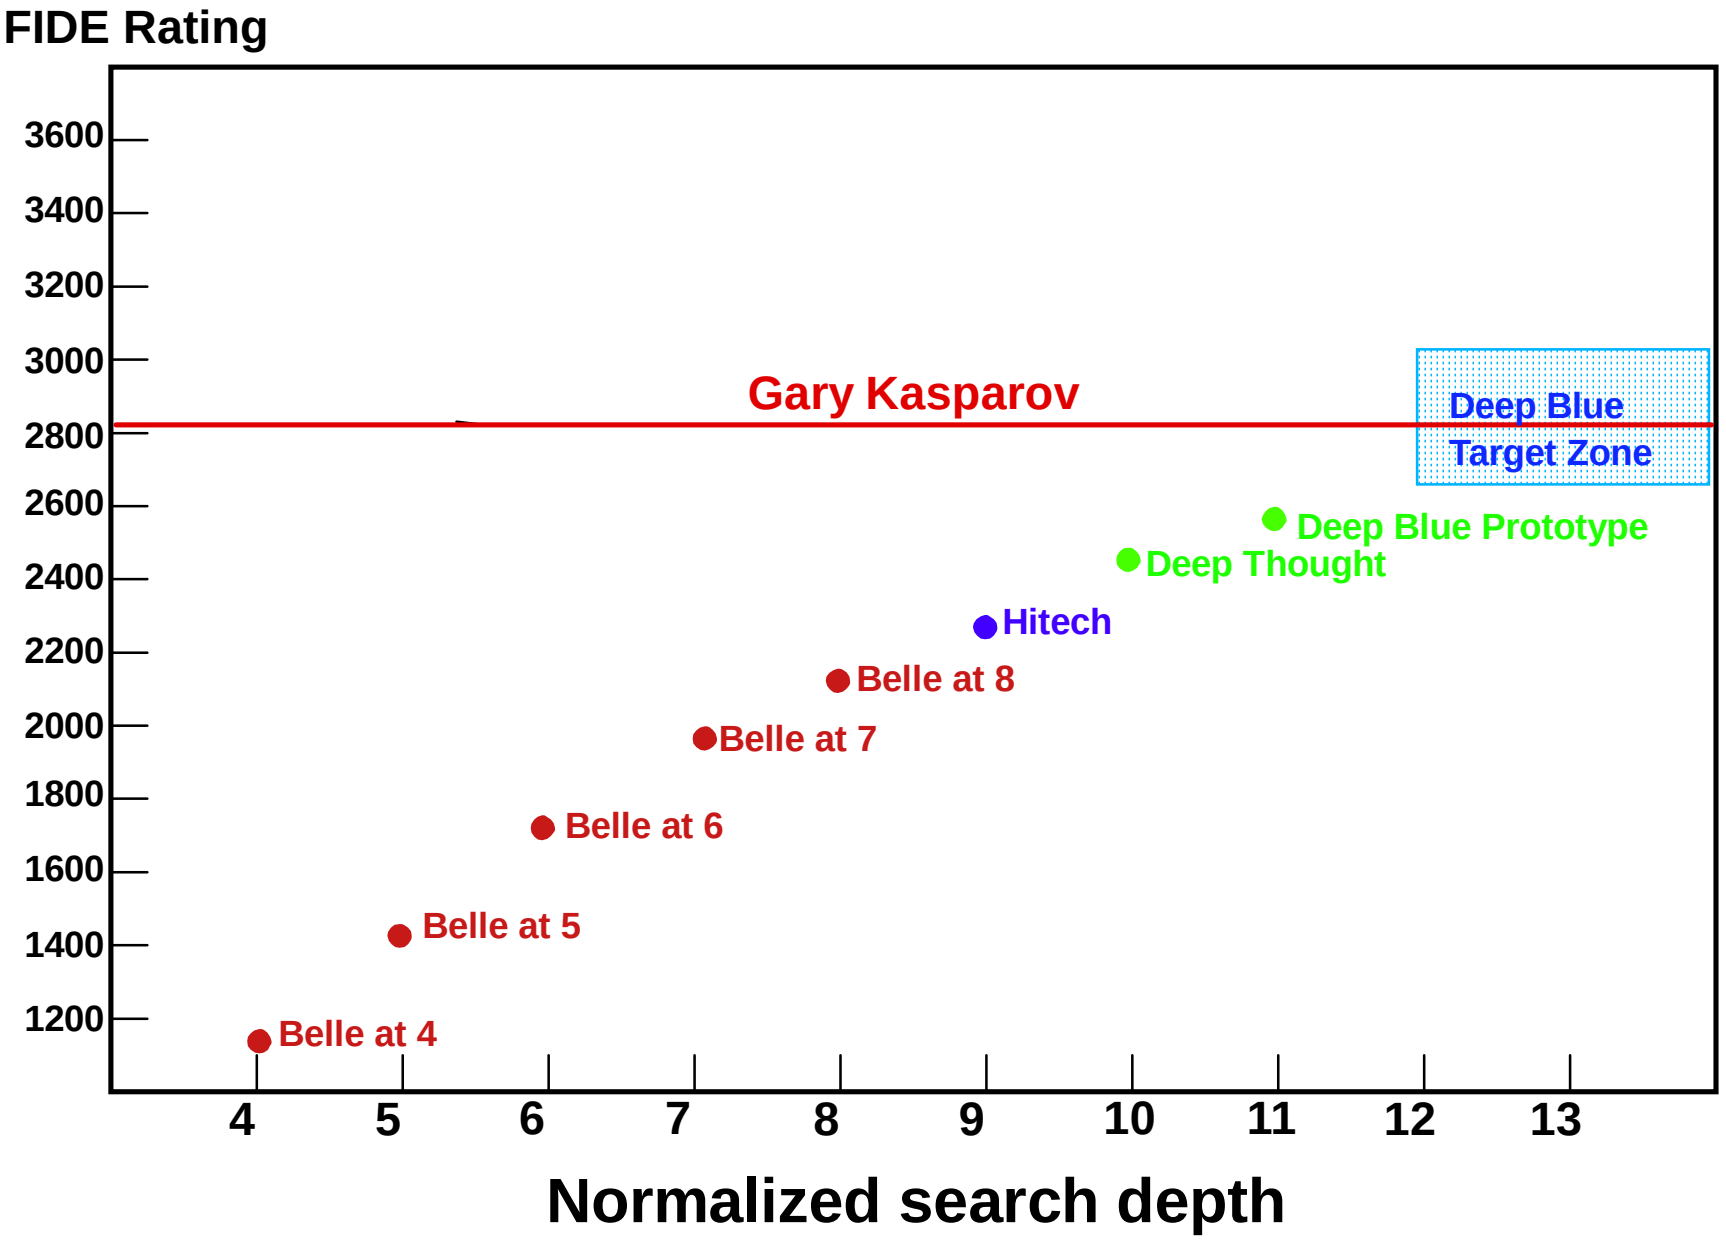

Problem Descriptions

- Search speed appears to be a necessary condition
- Speed can be obtained by increasing the level of integration (single chip chess machines) and by using massive parallelism
- But is speed alone enough?

Problem Descriptions

- Speed alone might not be enough
- Deep Thought was a tournament GM
- But not a match GM
- Human Grandmasters, in serious matches, learn from computers' mistakes, exploit the weaknesses, and drive a truck thru the gaping holes
- Need to have very few weaknesses, and only weaknesses that are difficult to exploit

Design Philosophy

- Chip speed is secondary
- Integration level is paramount
  - Encapsulate every chess evaluation terms from chess books
  - Create evaluation terms to deal with every known computer weaknesses
  - Add hooks to handle new weaknesses, if they appear, with external FPGA hardware
  - Put everything on one chip

Deep Blue Architecture

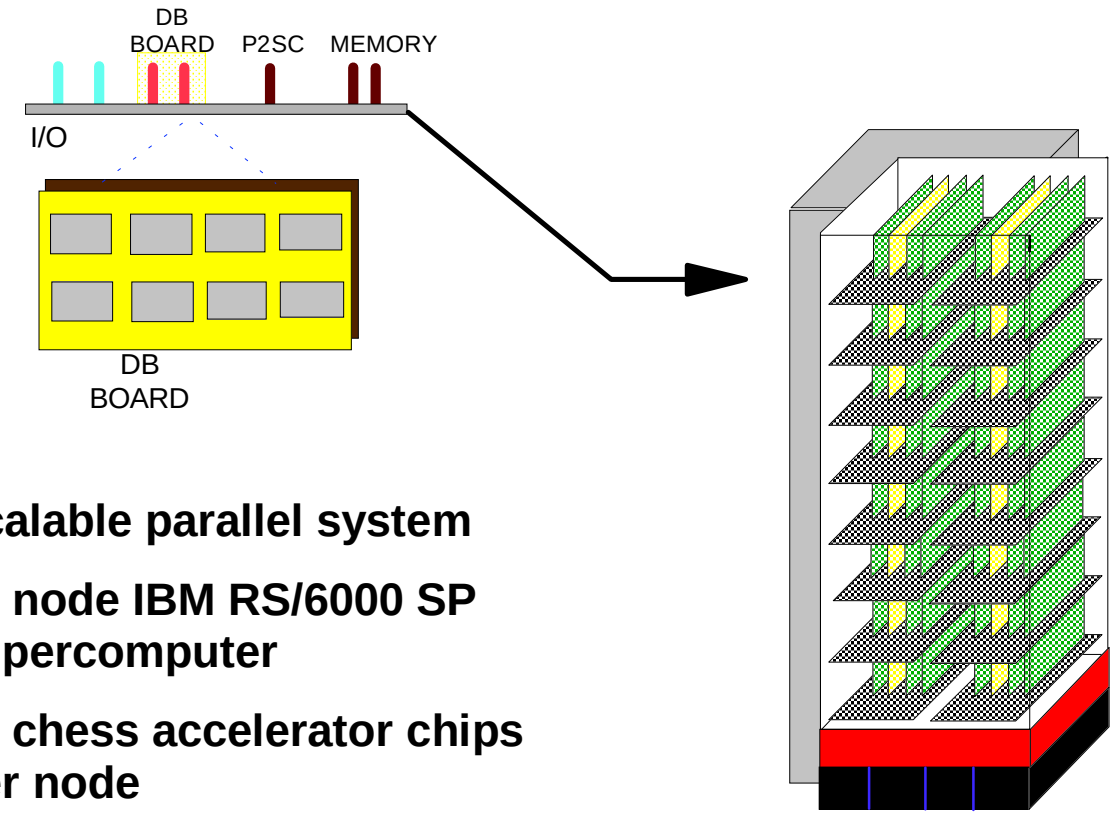

- Scalable parallel system
- 30 node IBM RS/6000 SP supercomputer
- 16 chess accelerator chips per node
- 200 million chess positions per second average

Chess Machine Basics

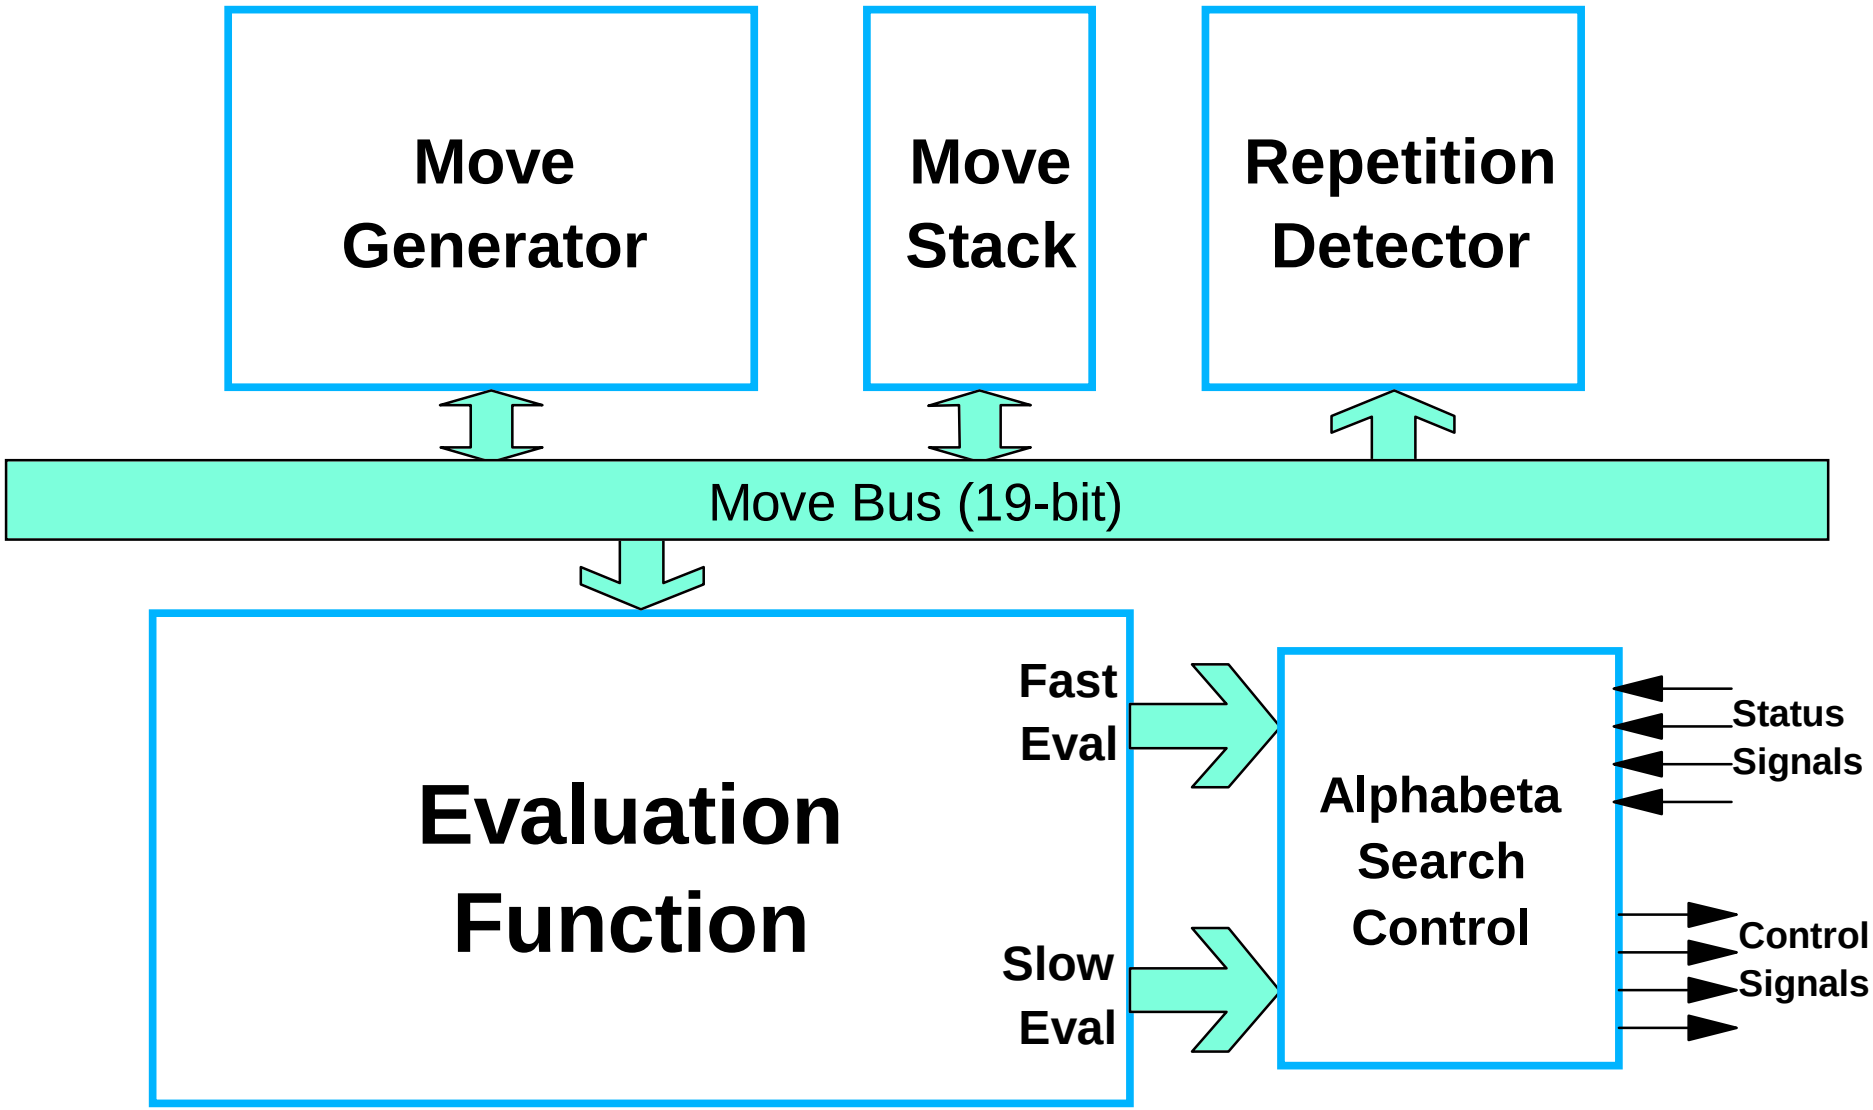

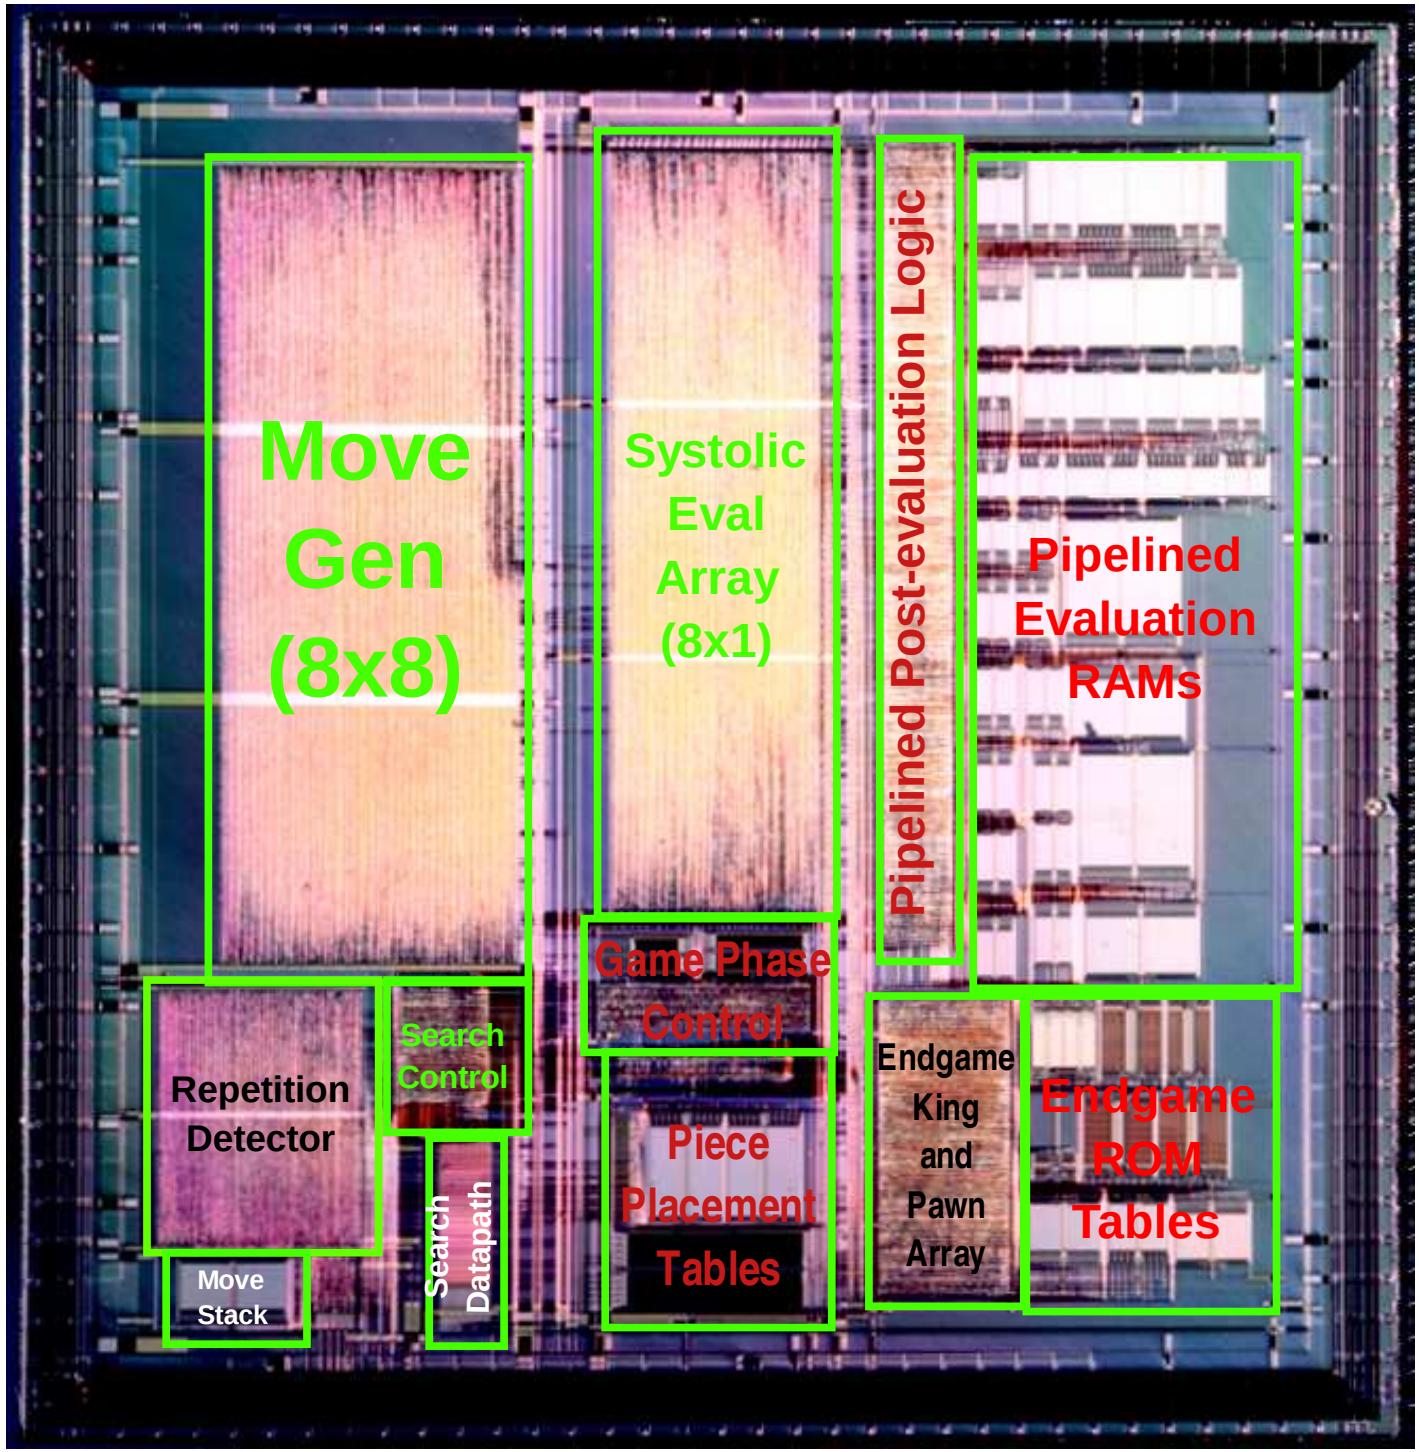
Move Gen (8x8)
Systolic Eval Array (8x1)
Pipelined Post-evaluation Logic
Pipelined Evaluation RAMs
Game Phase Control
Search Control
Repetition Detector
Move Stack
Search Datapath
Piece Placement Tables
Endgame King and Pawn Array
Endgame ROM Tables

Chip Architecture

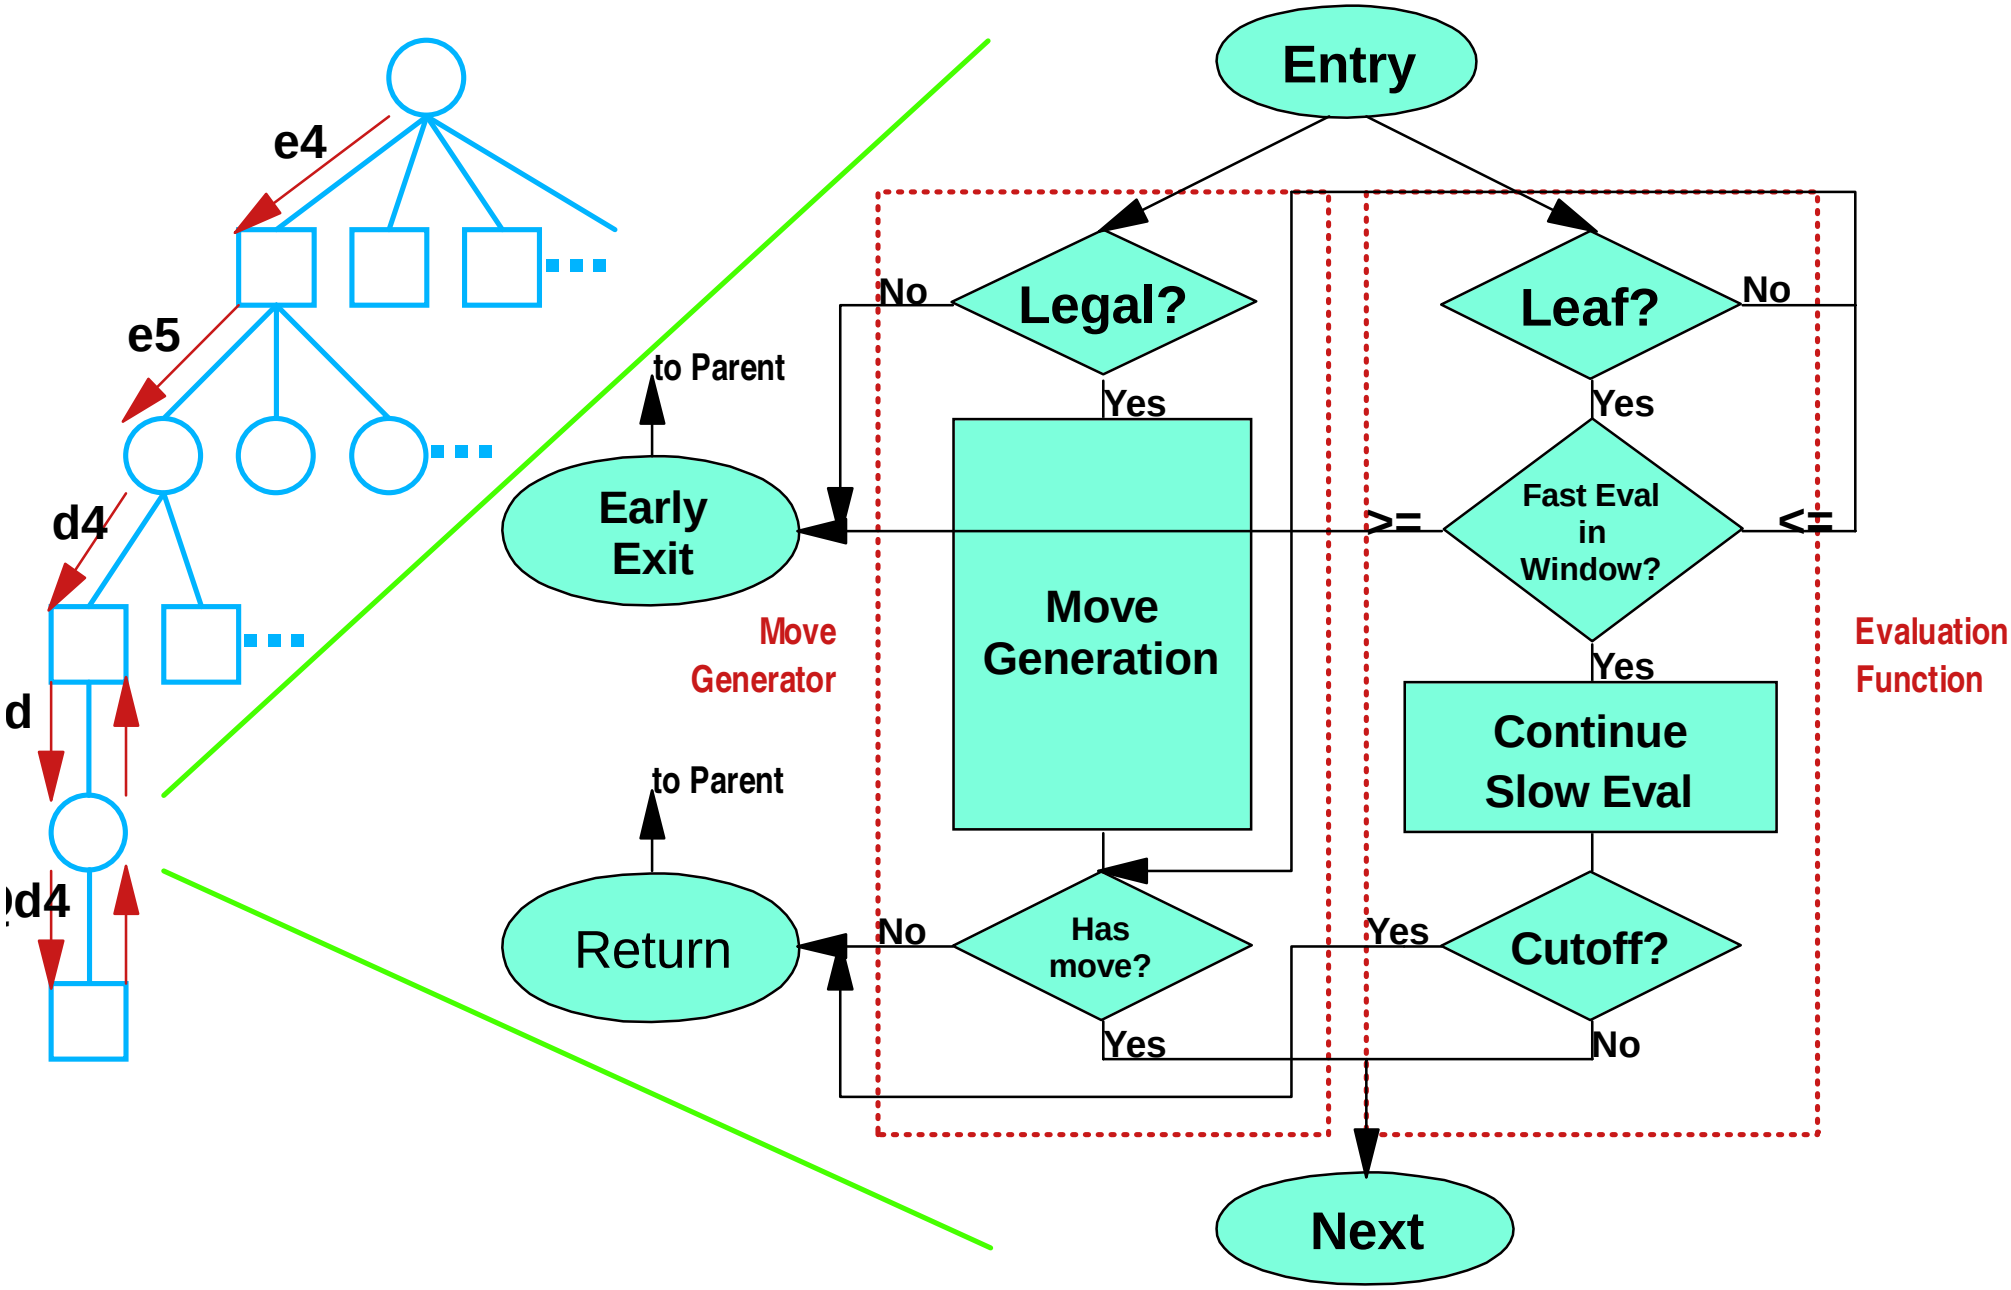

Chip Architecture (Move Gen)

- The move generator is an extension of the move generator used in Deep Thought
- 8 by 8 combinatorial array, one cell per square of the chess board
- Intercell wiring corresponding to the way chess pieces move
- Can generate checking and check evasion moves
- Also "generalized checking" moves

Chip Architecture (Eval)

- A fully parallel implementation of the evaluation function is too big
- Use a pipelined implementation
- Divide the evaluation into two parts
  - Fast evaluation in one cycle, including all the big terms
  - "Slow" evaluation computed in a 8-cycle sequence (+3 additional cycles), one cycle per column of the chess board
- Use small RAMs instead few large RAMs

Fast and Slow Evaluation

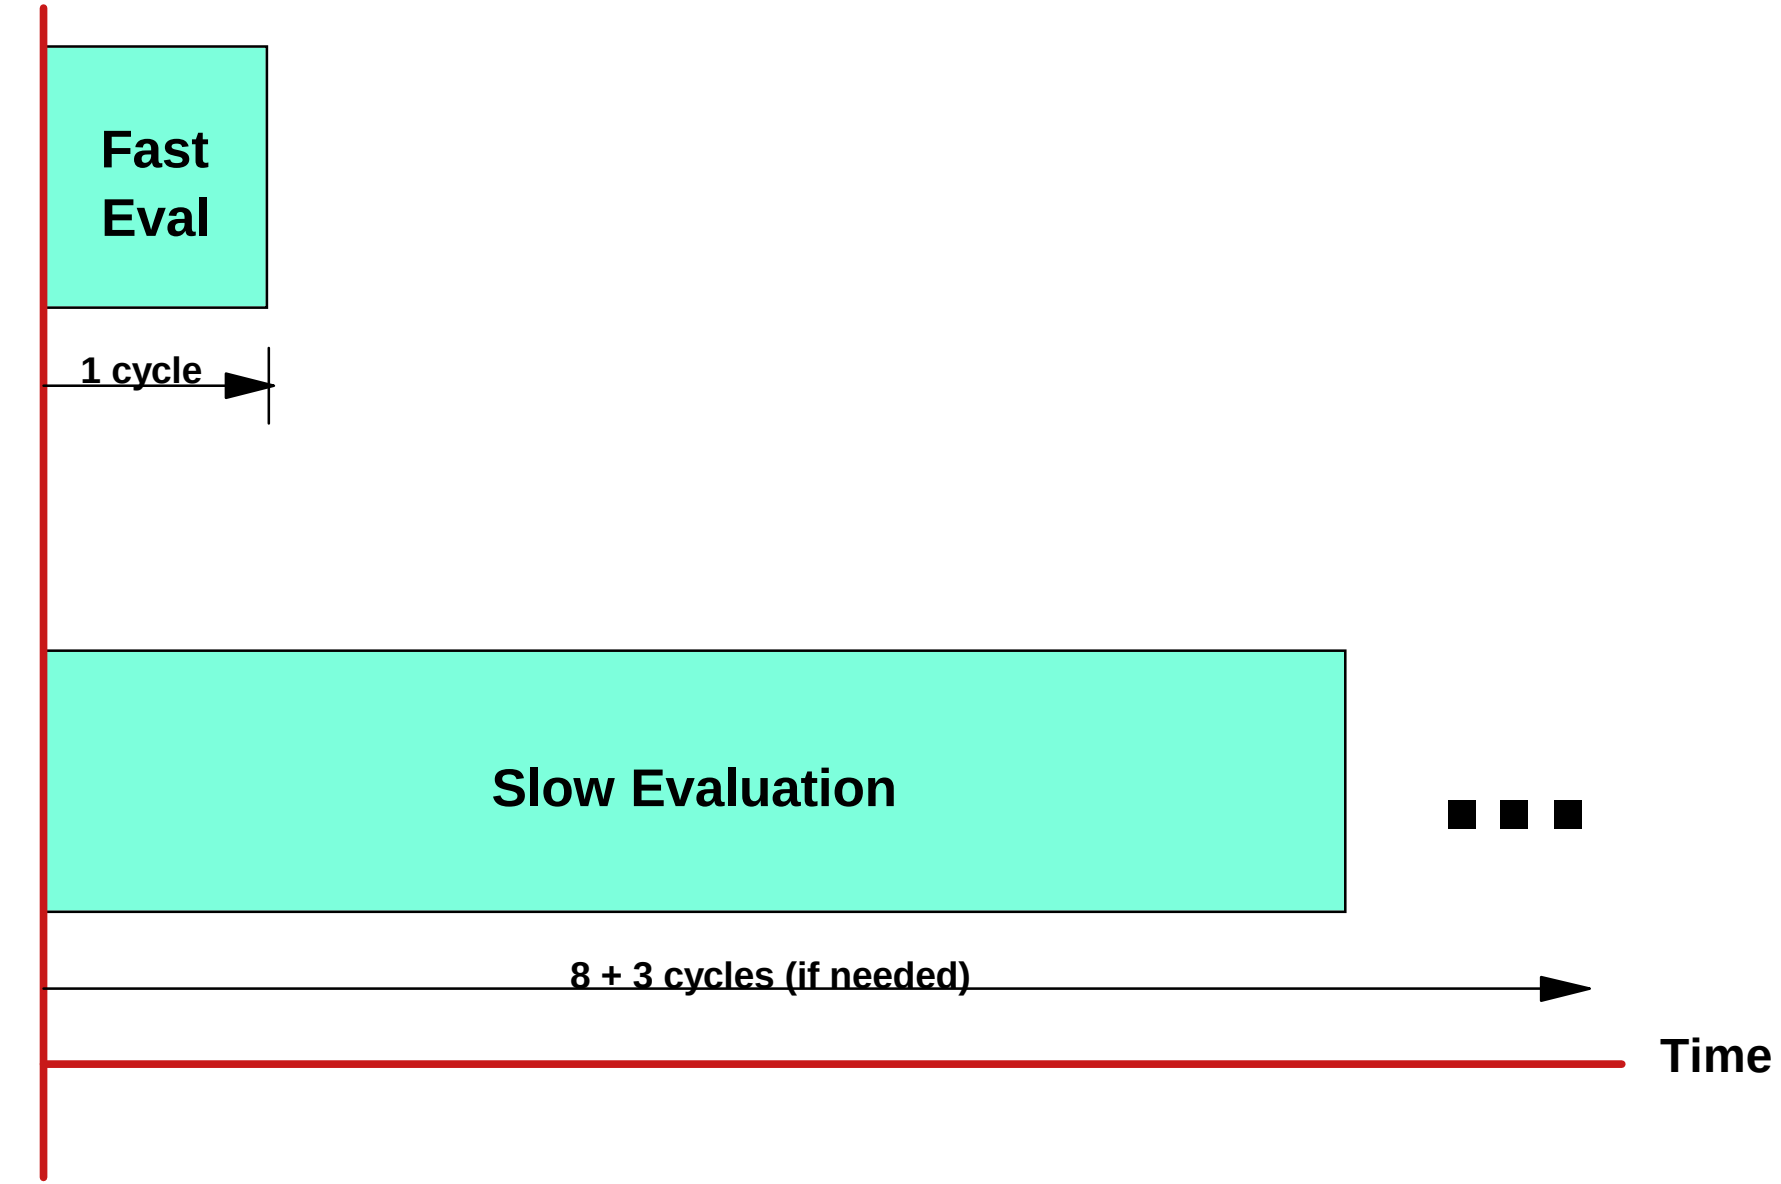

Fast Eval

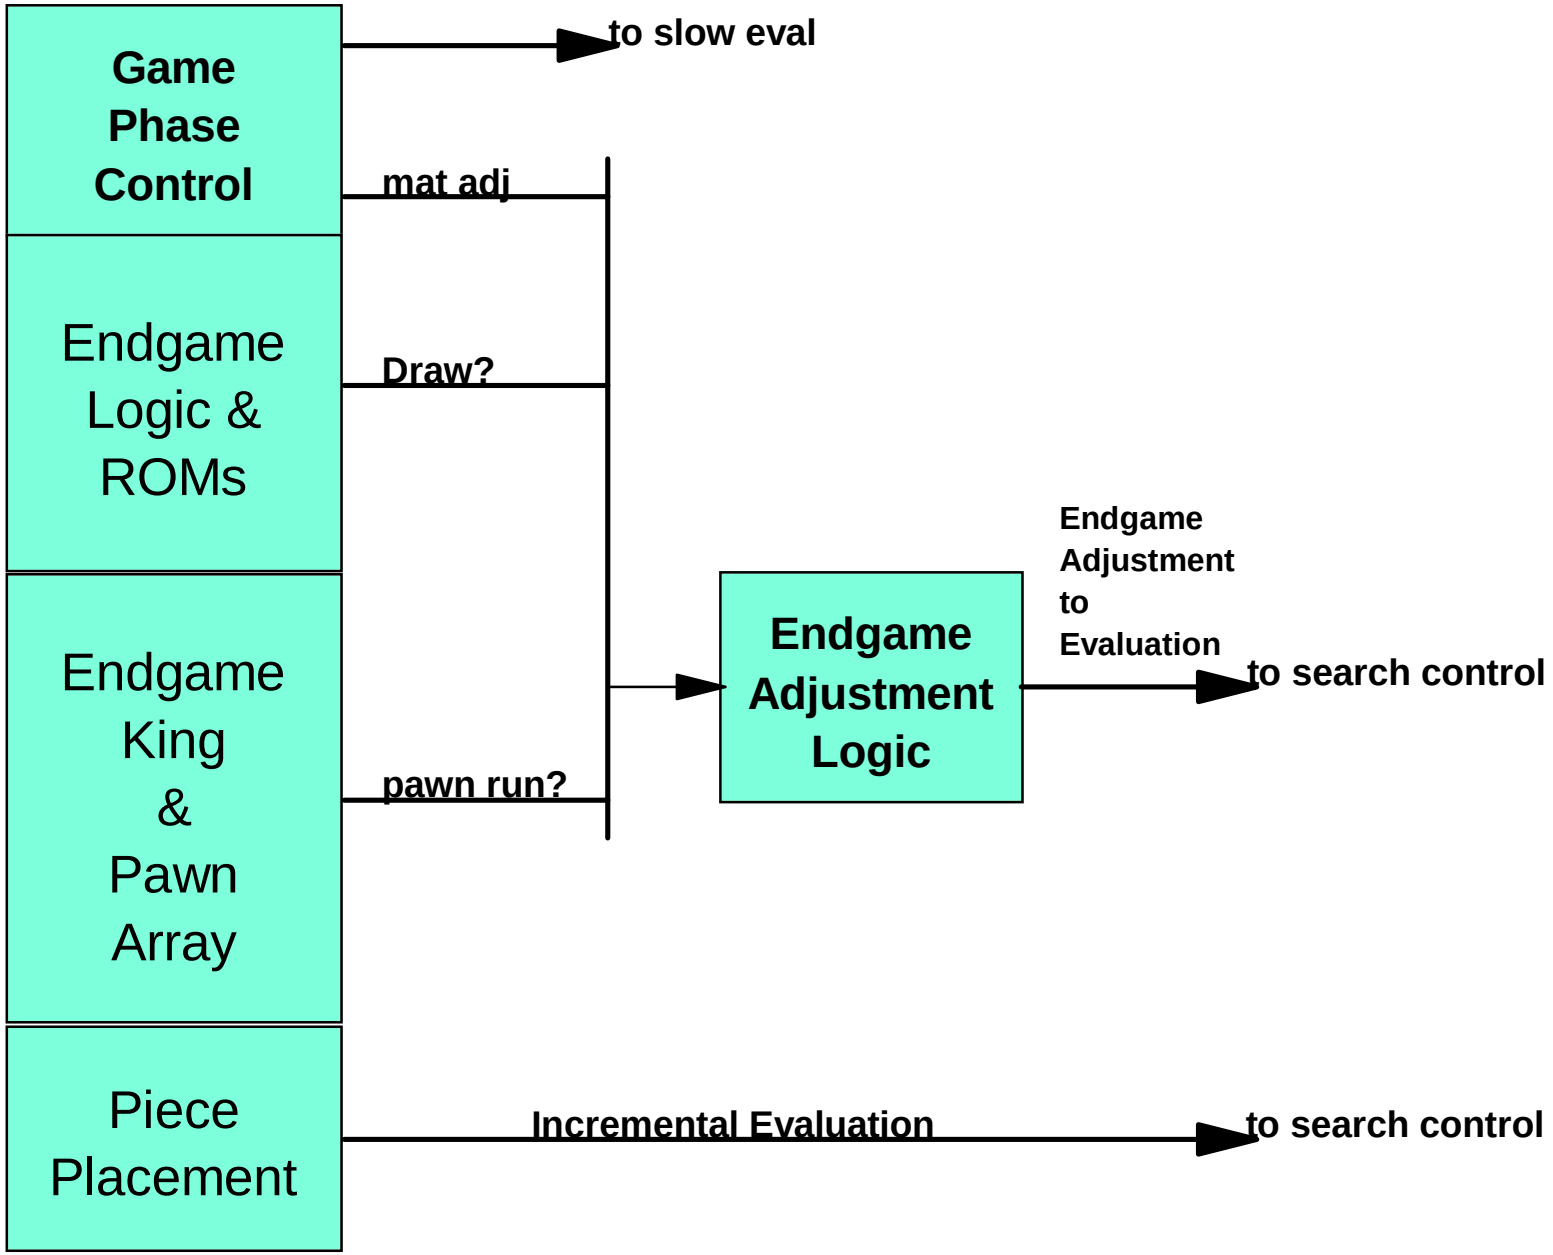
Game Phase Control
to slow eval
mat adj
Endgame Logic & ROMs
Draw?
Endgame King & Pawn Array
pawn run?
Endgame Adjustment Logic
Endgame Adjustment to Evaluation
to search control
Piece Placement
Incremental Evaluation
to search control

Piece Placement Table

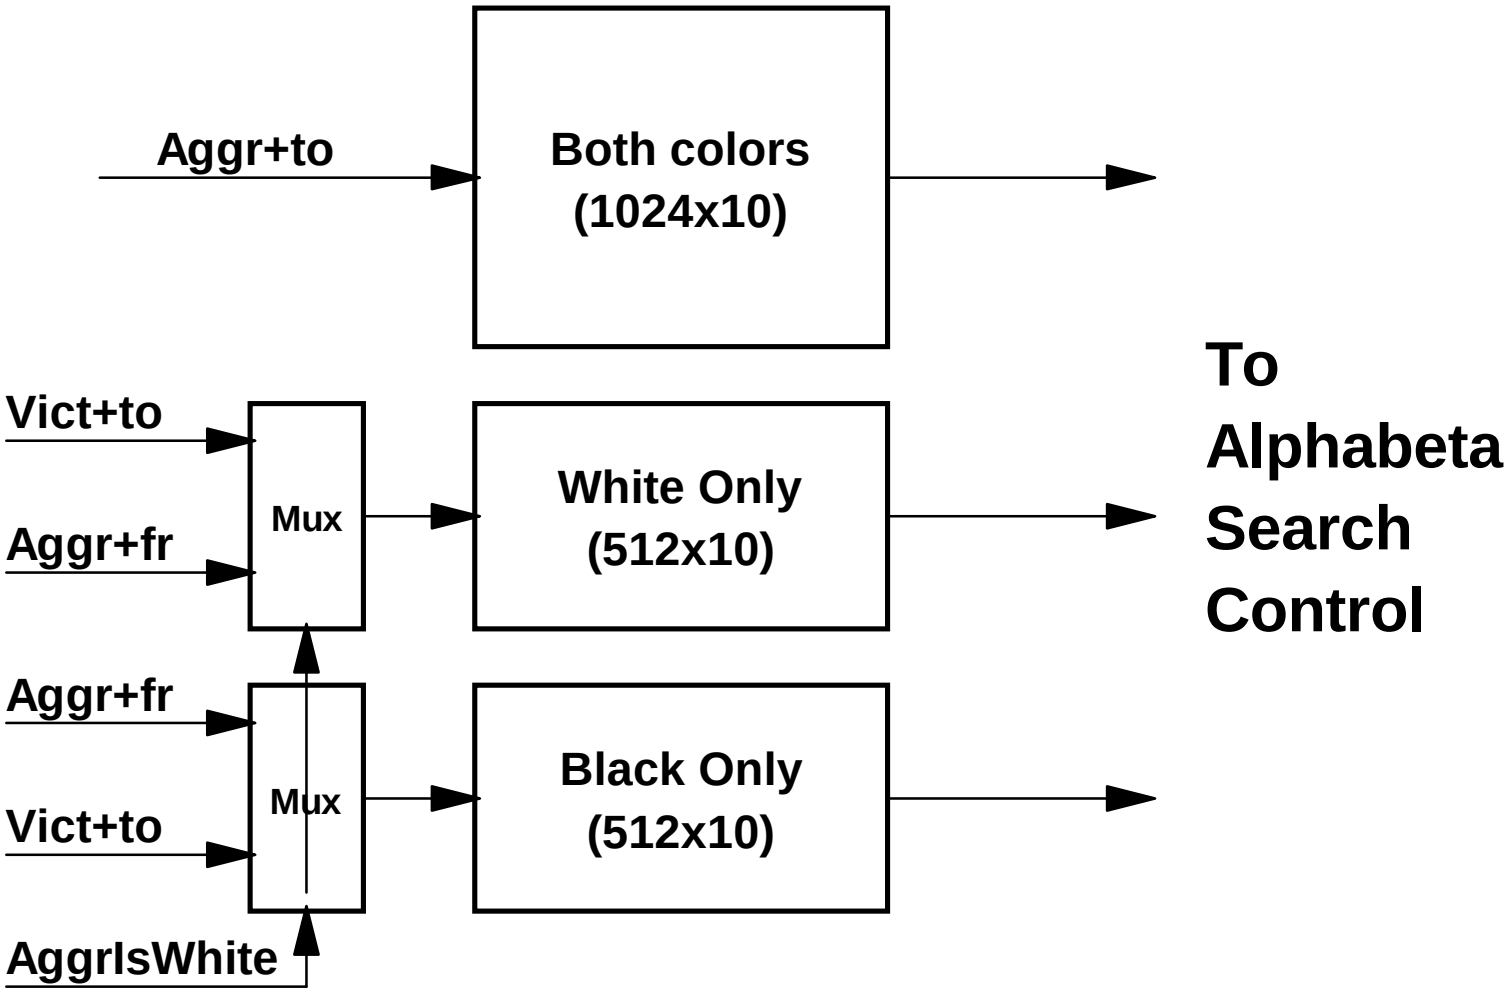

Endgame ROM Interface

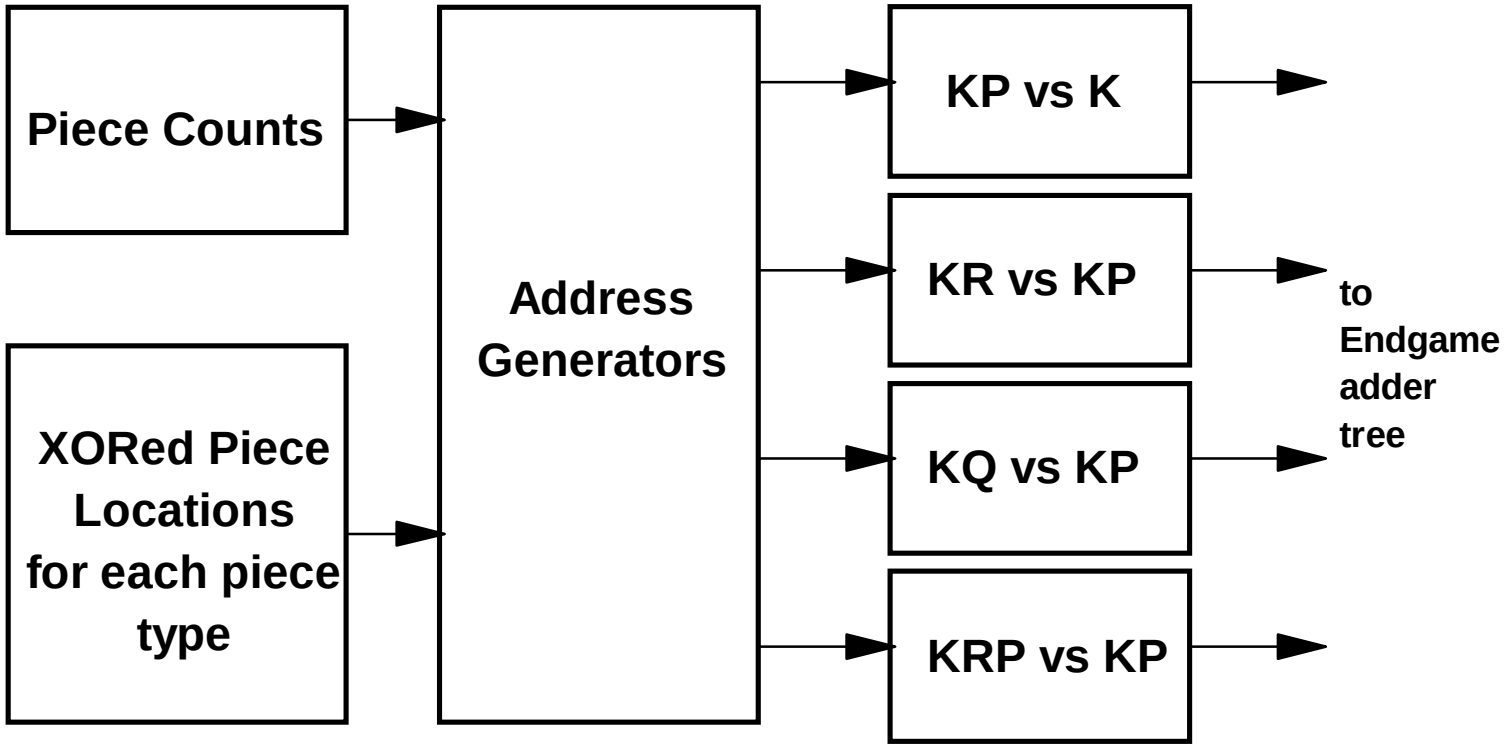

Main Slow Evaluation Flow

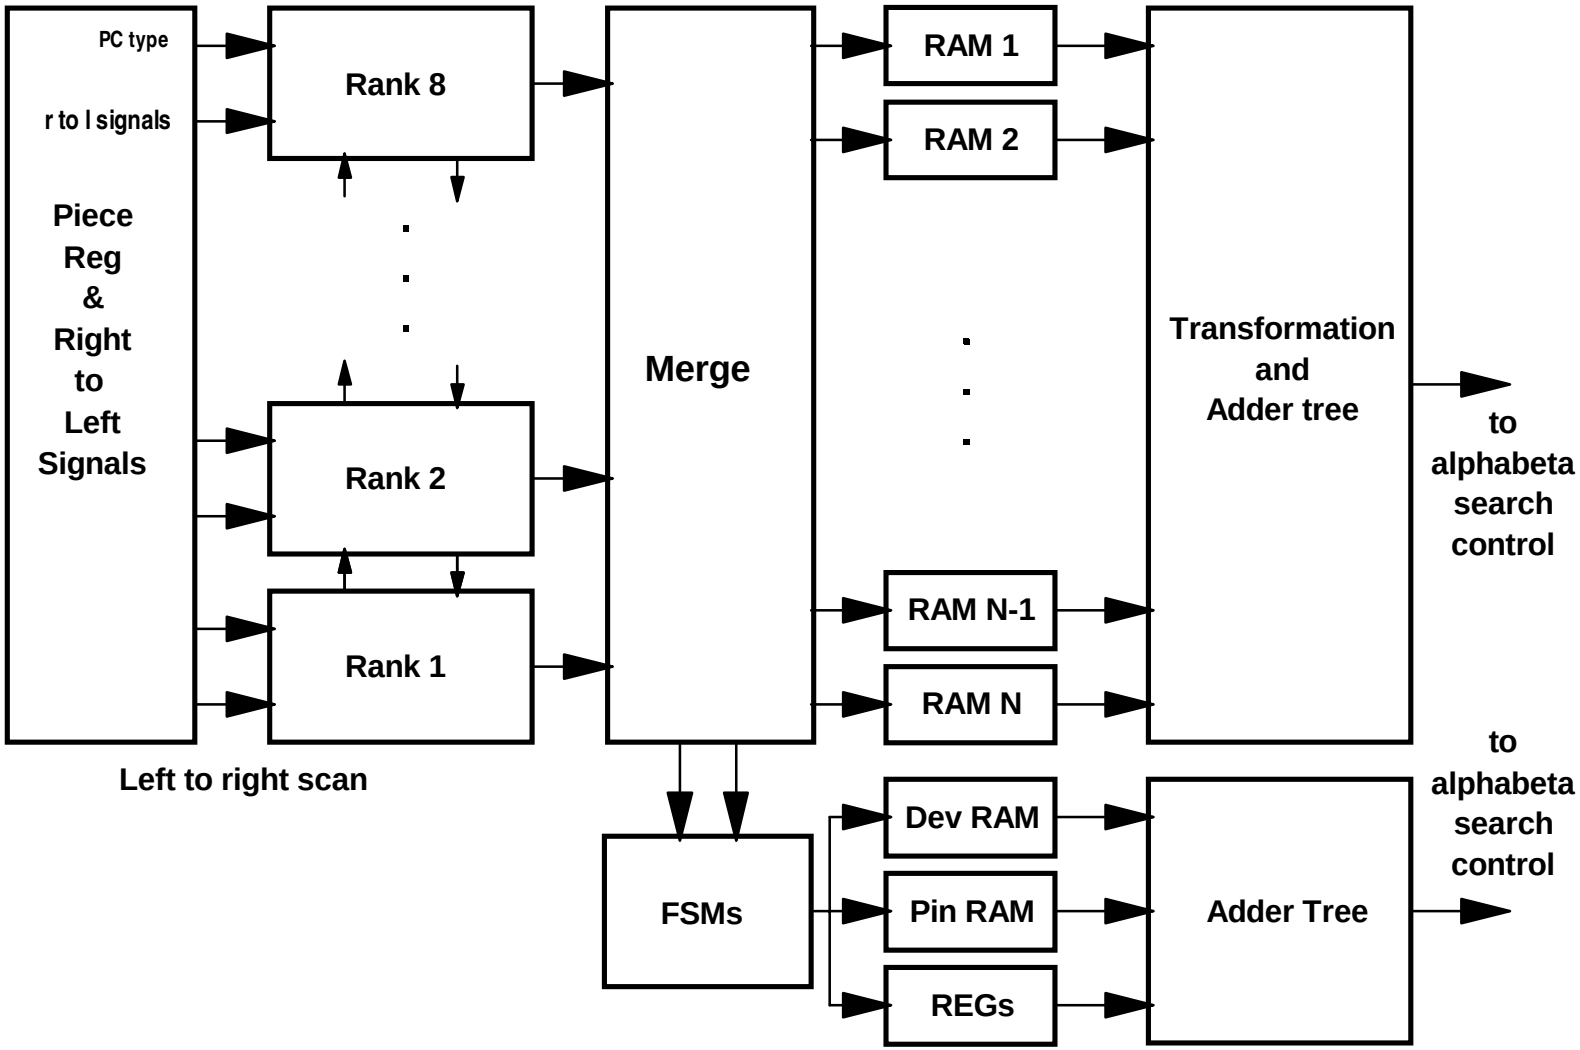

Chip Statistics

- 3-level metal 0.6-micron 5-V CMOS
- 1.5 million transistors, 1 W
- 1.4 cm x 1.4 cm
- 2.5 million chess positions/sec
  - 25 to 100 billion general purpose instructions per second
- A single chip system appears to play at strong Grandmaster and possibly Super Grandmaster level

System Performance

- A 480-chip system was assembled for the 1997 Rematch with Kasparov
- System sustained speed was about 200 million chess positions per second, or roughly 1/5 of the potential peak speed
- First match win ever by a computer over World Chess Champion (3.5 to 2.5)
- 78 million web hits (1997 figure)
- 5-7 billion impressions from Rematch alone
- IBM stock went up the next day

World Chess Champion on a PC?

- 0.6-micron CMOS not state of the art
- Design was not speed optimized
- 30 million chess positions/sec per chip appears possible with, say, a 0.35-micron process
- A small array of chips plugged into a PC could be sufficient to beat Kasparov, if a single chip is not sufficient
- Shogi (Japanese Chess) next?