



# CloudShield<sup>®</sup>

*an SAIC company*

## **Managing Heterogeneous Processor Machine Dependencies in Computer Network Applications**

**Ralph Duncan, Peder Jungck,  
Kenneth Ross, Greg Triplett, Jim Frandeen**

**August 26, 2013 HeteroPar2013 (Aachen, Germany)**

# Heterogeneous Packet Processing



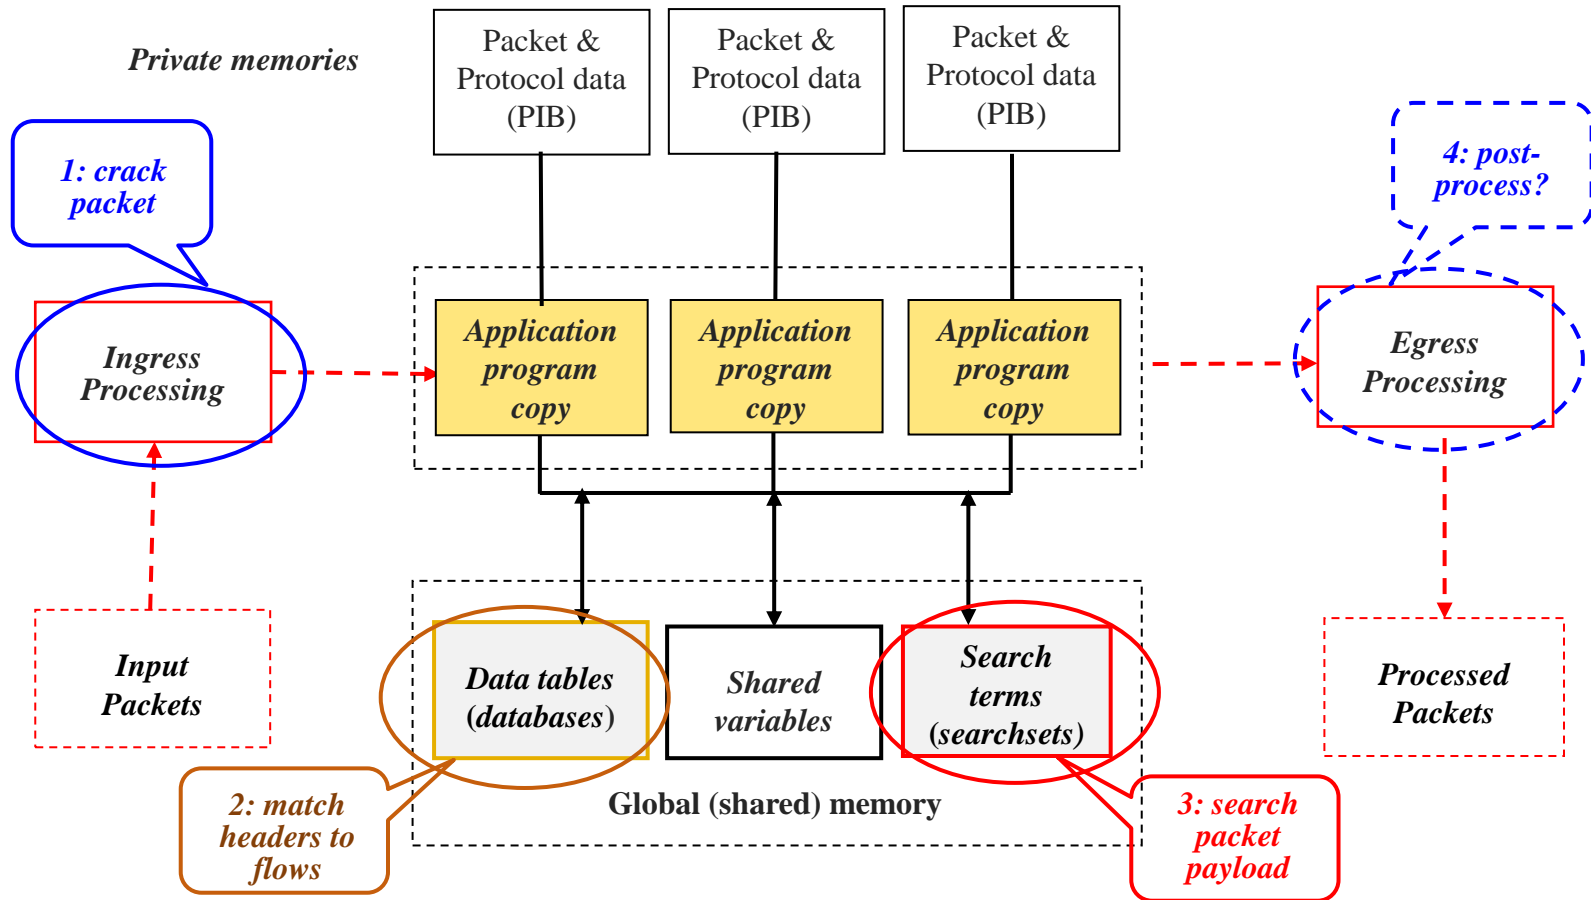
CloudShield®  
an SAIC company

- **Computer network packet processing** is the App Domain *par excellence* for **Heterogeneous parallelism**,
- Relentless **speed demands** 10 Gbps, 40, 100, 320?
- Predictable, specialized ops:
  - “*Cracking the packet*”
  - Matching **header fields**
  - Searching **packet payloads**
- Asynchronous, non-SIMD style processing
- Relevant chips: **NPU**s, **TCAM**s, **regex PE**s, **FPGA**s, ...

# Packet Processing Heterogeneous **Chip Types**

- Mapping **actions** to **Heterogeneous Processor Types**
  - ▶ “*crack the packet*” – FPGAs, ASICs
  - ▶ match *packet headers* to **Flows** – TCAMs  
(*Ternary Content Addressable Memory*)
  - ▶ Searching *packet payloads* – Regex Processors
  - ▶ **NPU**s (Network Processing Units) for many tasks

# Characterizing the Work



# Challenges for Using Heterogeneity in this Domain

- How to *orchestrate* Heterogeneous Processing
- How do users *program* Heterogeneous Processors?  
(especially specialty chips, FGAs, etc.)
- Practical Considerations
  - ▶ *Hide* processor-specifics from **User Apps** (reusability).
  - ▶ *Encapsulate* processor-specifics in **Tool-chain**  
(maintainability)

# Response to Challenges: Our **Technical Approach**

- Hide **chips-specifics** from **Programmers** with *high-level language* **extended types** and **operators** – **packetC**.
- Hide **chips-specifics** from **Tool-chain** with *high-level* **bytecodes** to express operations.
- Encapsulate **chip-specifics** in **Microcoded Interpreters**:
  - ▶ Parallel copies of the **Interpreter** interpret the **bytecodes**
  - ▶ **Interpreter** knows **chip Command Languages**
  - ▶ **Interpreter** knows **chip Communication details**
- Orchestrate **processors** via the **Interpreters**

# Hide Chip-Specifics from Programmers with *packetC*

- Users program via **high-level packetC types, operators**
  - ▶ define *Data* as **extended, C-style types**
  - ▶ define *Actions* as C++-style methods on data objects
- E.g., make *masked databases* be **structs** (not TCAMs)

```
struct stype { short dest; short src;};  
database stype myDB[50];  
// each element of myDB has the structure: { stype data; stype mask;};
```

- E.g., express operations as C/C++-style *methods*

```
rownum = myDB.match( myStruct ); // do masked search
```

# Translate **packetC** source to Chip-Independent **Bytecodes**

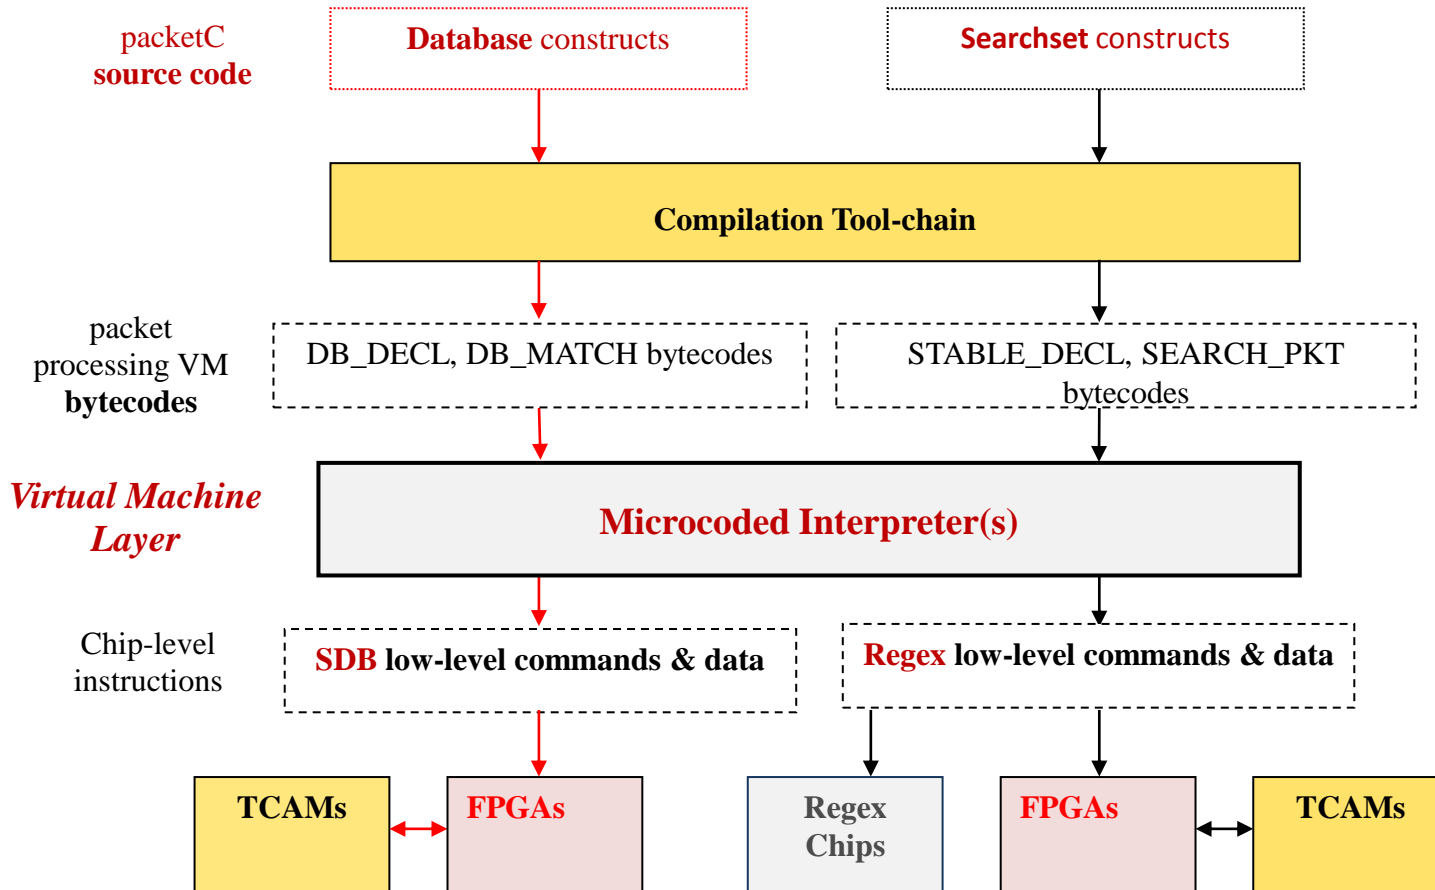
- Use complex, chip-independent **bytecodes** that...
  - ▶ Define a virtual ***packet processing machine***
  - ▶ Avoid **chip-specific** commands, formats, etc.

```
// Bytecode - database declaration and MATCH operation
DBASE_DECL myDB [500] 5 // where 5 is the database ID

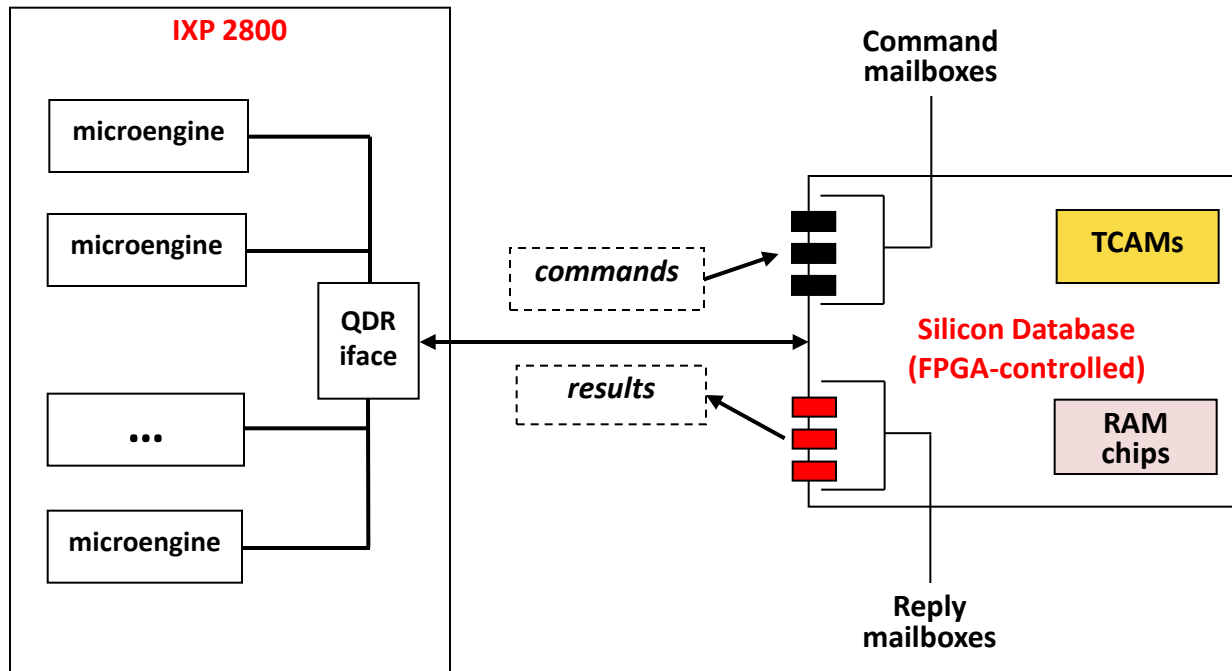
DB_MATCH Id, maskKind, inData, inMask, rownum
// Id: integer indicating with DB to use
// maskKind: scenario indicator
// inData: data to match
// inMask: to combine w/ each element mask
// rownum: 32bit integer to rec matching DB row #
```



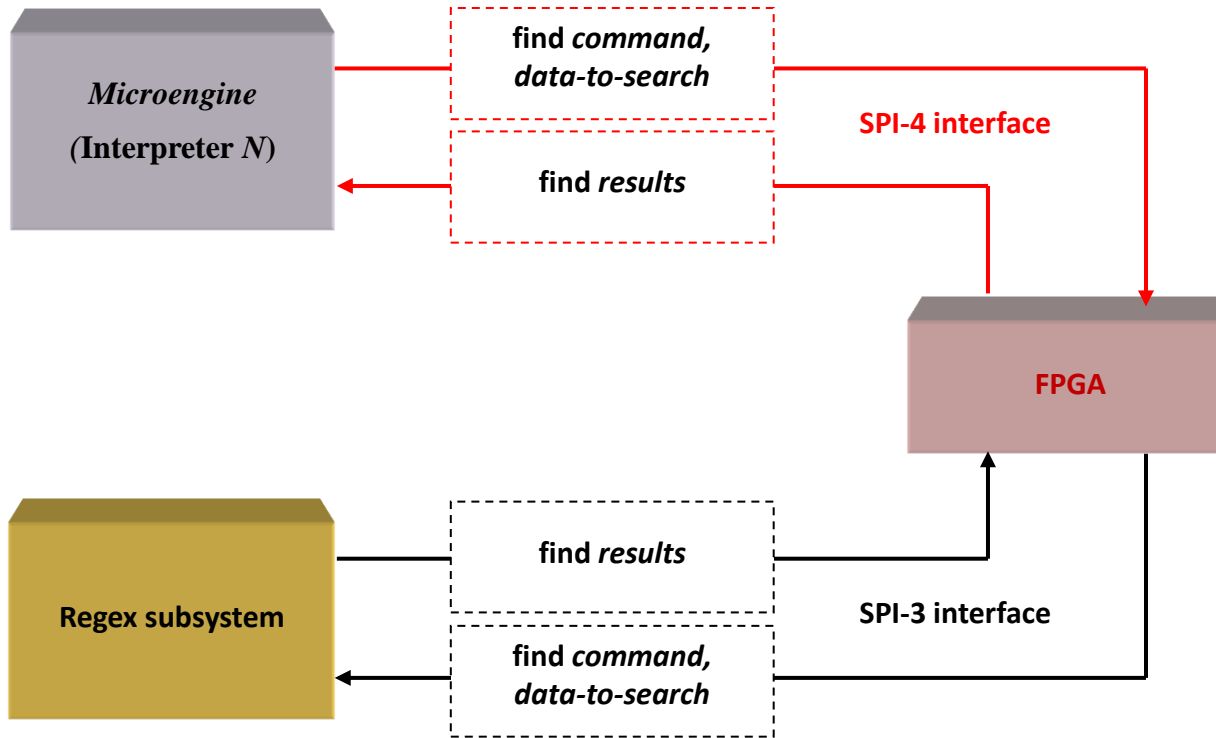
# Encapsulate Chip-Specifics in a *bytecode Interpreter*



# Orchestrate Using Specialized Processors: **FPGA/TCAMs** #1



# Orchestrate Using Specialized Processors: **FPGA/Regex** #2



# The Running System: a hive of **heterogeneous activity**

- A running CS-2000 or PN41 has:
  - ▶ **95 IXP contexts** running an **Interpreter** copy (**NPU**s)
  - ▶ Other **IXP contexts** doing house-keeping functions
  - ▶ An **FPGA** managing shared memory and **TCAM**s
  - ▶ **8 TCAM** chips (associative memories) running
  - ▶ Another **FPGA** managing access to a Regex System
  - ▶ An **IDT PAX Port regex system** running
  - ▶ Other **FPGAs** executing specialized tasks.
- This is a highly *asynchronous system*, driven by the arrival of **Packets**.

# Timing the hive of activity: 2 Experiments

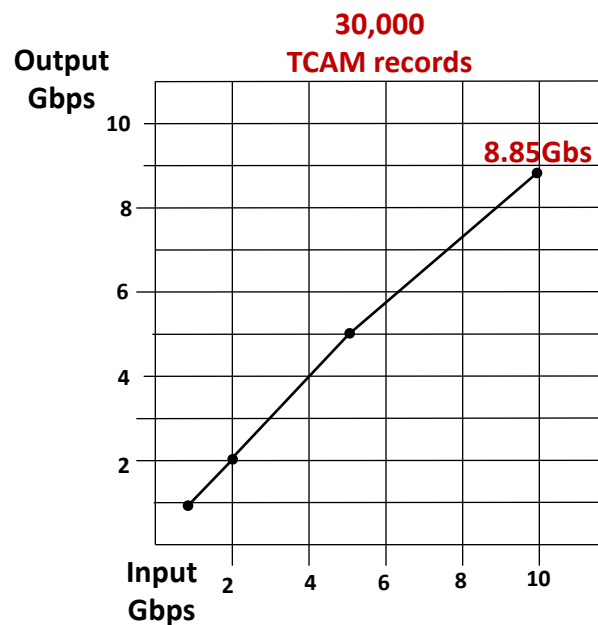
- Equipment
  - ▶ PN41 – (the heterogeneous CloudShield platform)
  - ▶ IXIA® XM12 *traffic generator* -> up to 10 Gbps traffic
- Experiment 1 – **matching**: NPUs, FPGA and TCAMs
  - ▶ **match** contents of a specific **payload location** vs.
  - ▶ **30,000 byte patterns** stored in **TCAMs**.
- Experiment 2 – **finding**: NPUs, FPGA and regex
  - ▶ **search** contents of the entire **packet** for
  - ▶ **10,000 strings (in regex form)** stored in **regex** memory.



# Performance

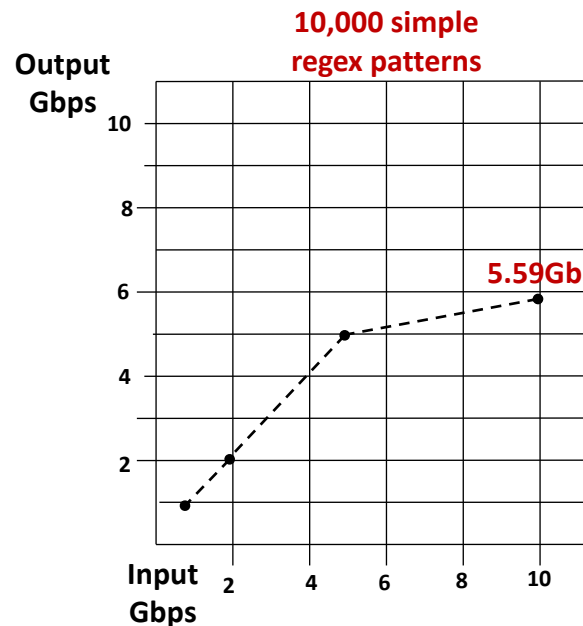
## Matching TCAM records

**Almost 9 Gbps:** TCAMs  
compare all records in  
parallel



## Searching for Regex Strings

**5.59 Gbps:** good for an  
app spending most of its  
time in a 5 Gbps regex



# CONCLUSIONS



CloudShield®  
an SAIC company

- **Speed** is at least state-of-practice (2-9 Gpbs)
  - ▶ *Microcoded interpretation* is low-level and fast
  - ▶ **SPI-3**, **SPI-4** and **QDR** are effective data movers.
- **packetC's** C/C++ style constructs
  - ▶ Are **intuitive** ways to express **classic network structures**
  - ▶ *Effectively disguise chip-specific* attributes.
- **High-level bytecodes** for packet-processing
  - ▶ **Hide** chip-specifics from the **Tool-chain**.
  - ▶ *Involves new Tool-chain design (working on refinements)*.