

EXEGY

Infinite data. One answer.

Exploiting New Computer Architectures for Modeling and Simulation

Roger D. Chamberlain

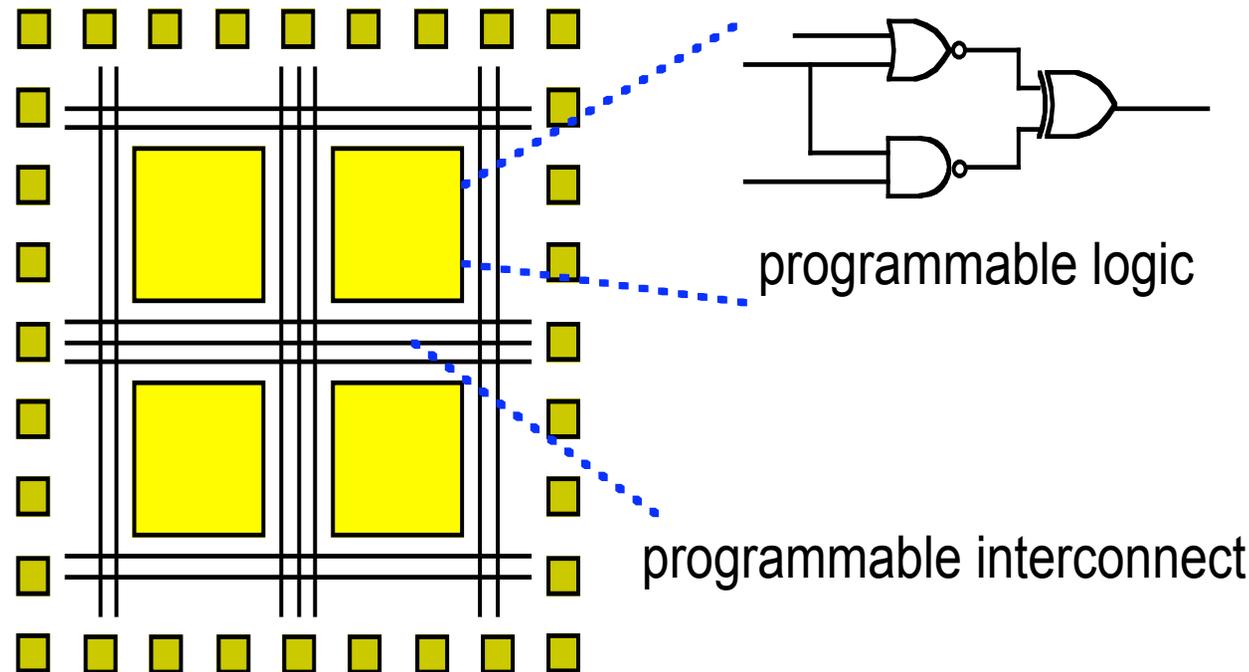
Dept. of Computer Science and Engineering

Washington University in St. Louis

and

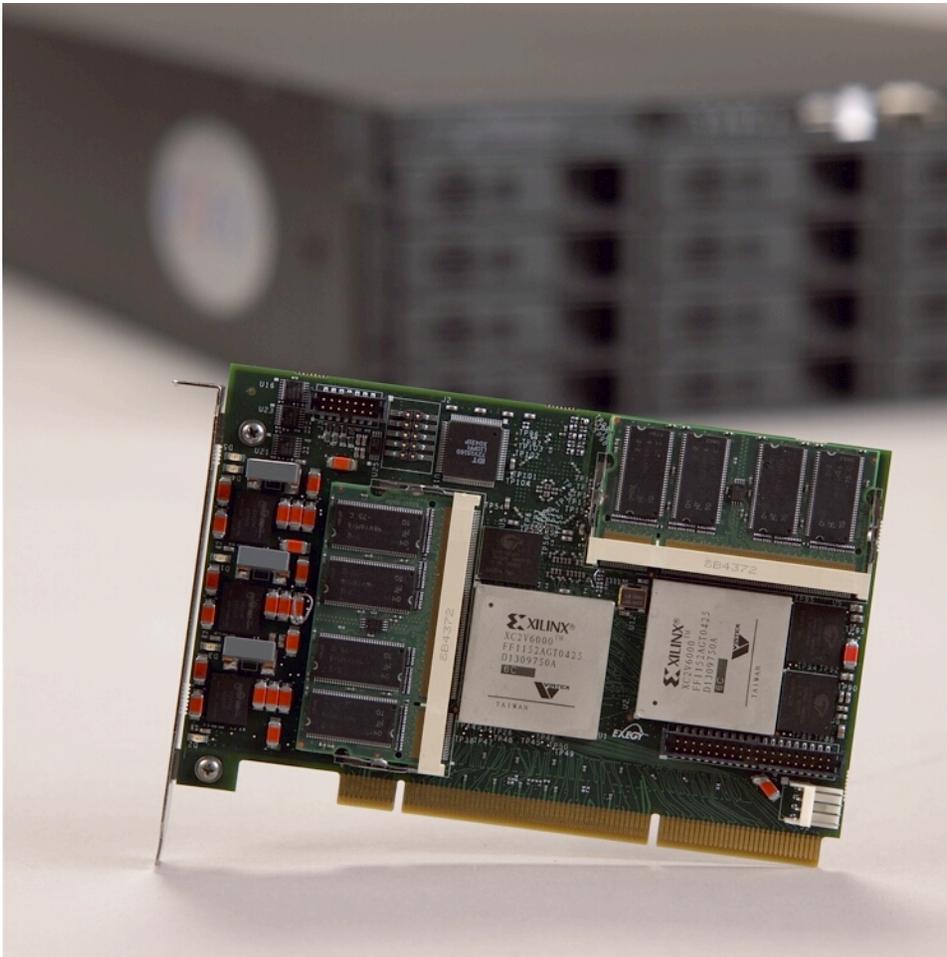
Exegy Inc.

Enabling Technology: Reconfigurable Hardware



- Field Programmable Gate Arrays (FPGAs) provide custom logic function capability
- Operate at hardware speeds
- Can be altered (reconfigured) in the field to meet specific application needs

Specialized Processing on Custom Board



- FPGA accelerated custom board
- Permits massively parallel operations
- Offloads work from CPU
- Integrates with other system components enabling high-speed data ingress and egress
- Designed with common APIs to give user control of functionality
- Draws from a library of pre-defined modules used to perform certain operations
- New functional modules readily incorporated

Analogous to graphic accelerator cards

Exegy A2000 Appliance



Highly Optimized
Data Pipeline from
Input thru Output

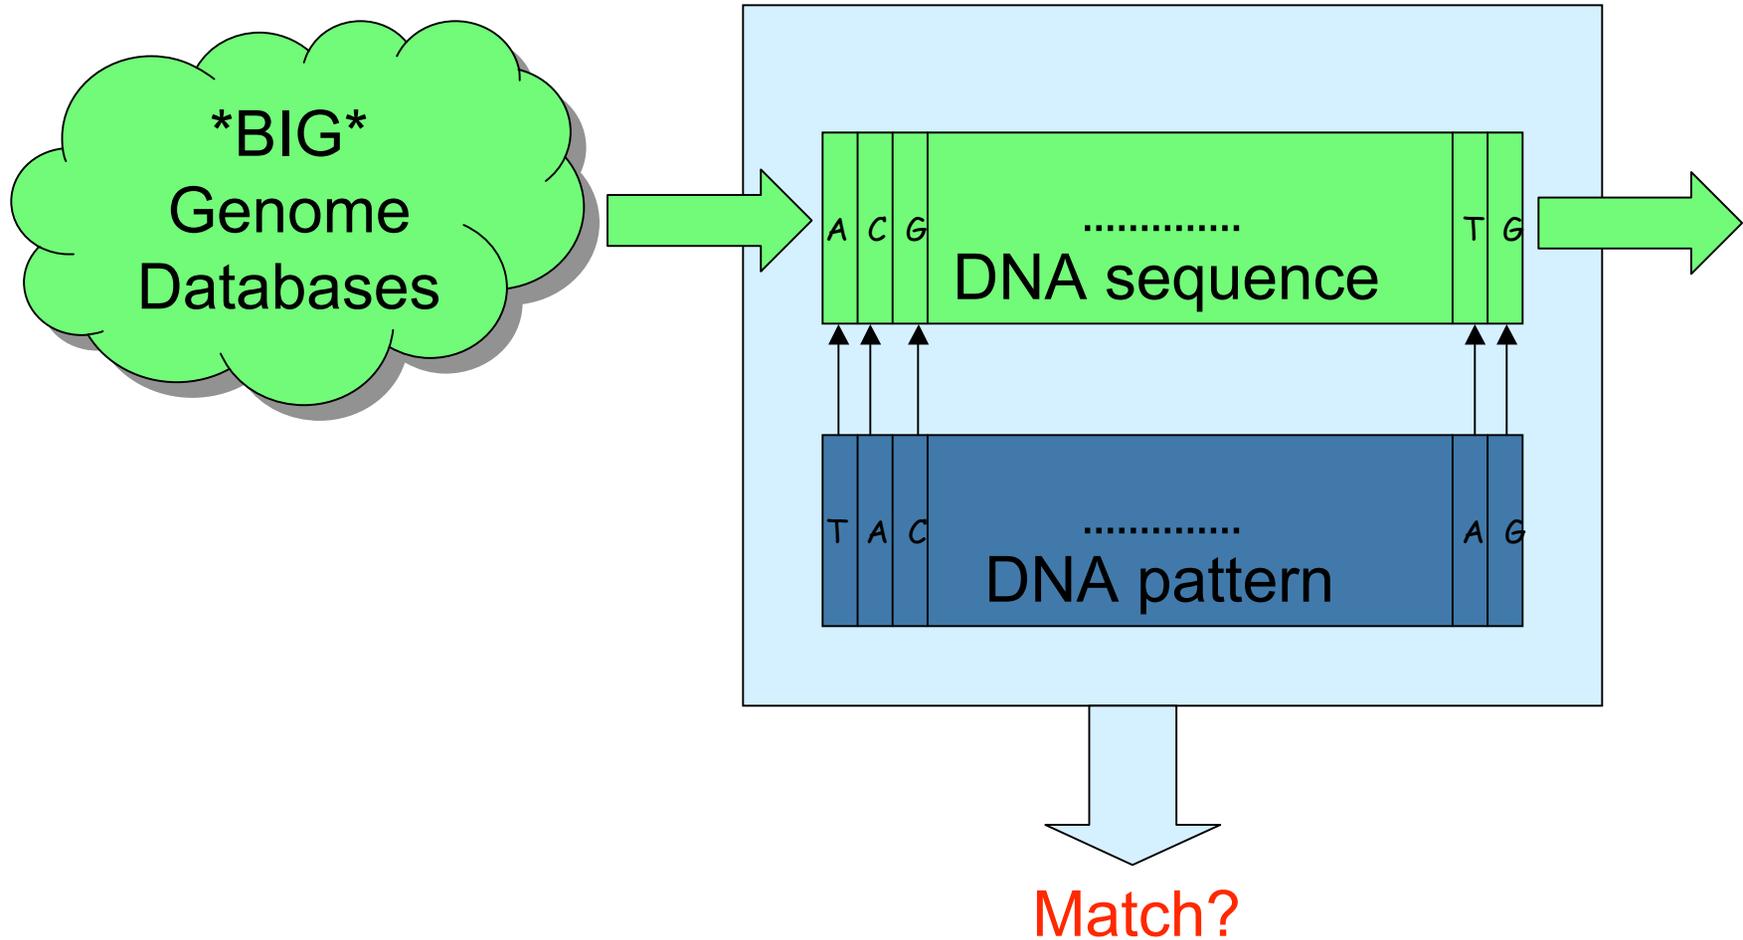
Specialized
Processing in Close
Proximity to Data

1-7 TB fast RAID;
RAM / FPGA
contiguous

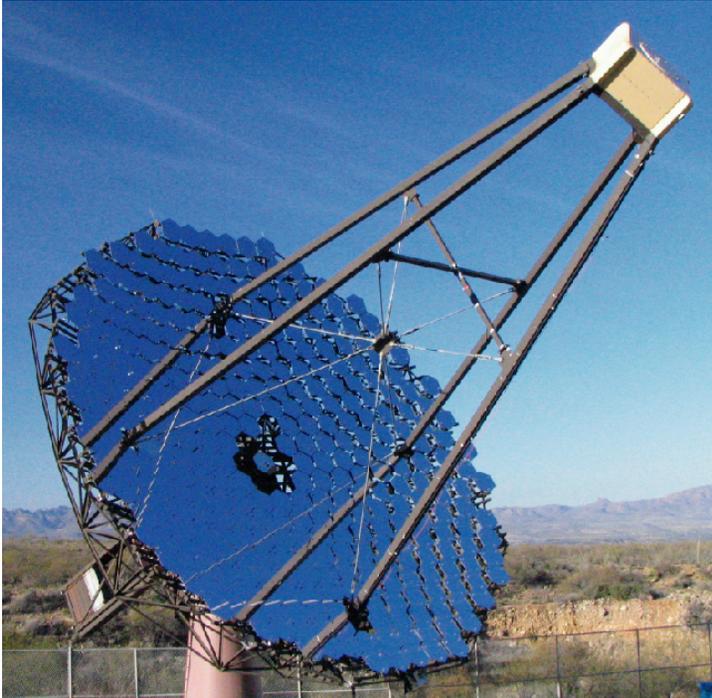
Applications Operational Today

- Biosequence search
- Science data mining
- Text search (exact and approximate)
- Compression/decompression
- Encryption/decryption
- Structured record search
- Signature hashing
- Financial data ticker plant

Biosequence Similarity Search



VERITAS Telescope



- Array of 12 m telescopes being constructed in Arizona.
- Looking for Cherenkov radiation from 50 GeV to 50 TeV gamma-ray interactions with upper atmosphere.
- Early indicator of supernovae, so timely data analysis is central to scientific mission.

Application Observations

- Previous two applications have much in common
- Large volumes of data
 - ♣ resulting from Human Genome Project, or
 - ♣ derived live from sensor array
- Searching for patterns in data
 - ♣ patterns are approximate (inexact)
 - ♣ pattern recognition algorithms are compute intensive
- Application to security sensing should be clear
 - ♣ all of above applies to problem of detecting threats in large volumes of sensor data

Modeling and Simulation

Stochastic discrete-event simulation loop:

```
while (not done)
  event ← next_event()
  time ← event.time
  process(event)
  for each new_event generated
    schedule(new_event)
  endfor
endwhile
```

Modeling and Simulation

Stochastic discrete-event simulation loop:

```
while (not done)
  event ← next_event()
  time ← event.time
  process(event)
  for each new_event generated
    schedule(new_event)
  endfor
endwhile
```

**Random numbers
needed here**

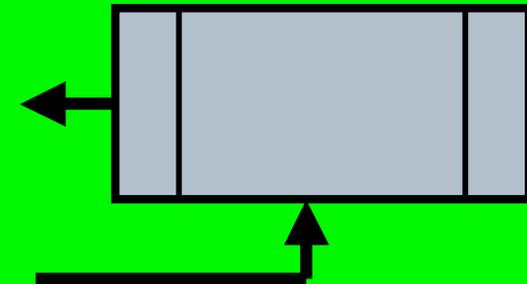
**FPGA can generate
266 Mvariates/sec**

Modeling and Simulation

Stochastic discrete-event simulation loop:

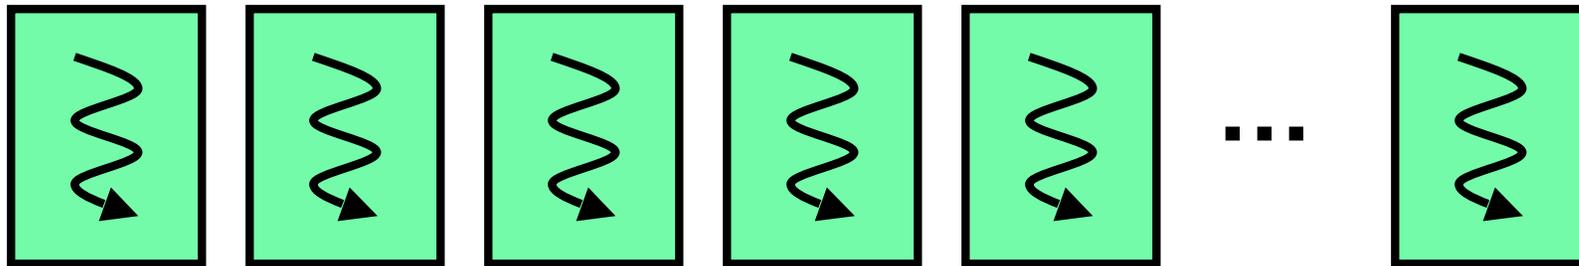
```
while (not done)
  event ← next_event()
  time ← event.time
  process(event)
  for each new_event generated
    schedule(new_event)
  endfor
endwhile
```

**Priority queue
needed here**



**FPGA can support
133 Mops/sec**

Monte Carlo Simulations



- Replicated, independent execution threads can improve statistical quality of results
- FPGA can support many concurrent threads
- Result is replicated results quality in single execution run time

Benefits of More Capable Simulation

- Higher fidelity models yield more relevant results
- Greater replications yield better statistical validity
- More capacity allows one to model more aspects of the actual system:
 - ♣ Queueing systems
 - ♣ Human behavior
 - ♣ Security infrastructure
- Faster than real time simulations can be used to support real time decision making