



**MACIEJ BESTA, MARC FISCHER, TAL BEN-NUN, JOHANNES DE FINE LICHT, TORSTEN HOEFLER**

# Substream-Centric Maximum Matchings on FPGA




# Large graphs...

# Large graphs...



# Large graphs...


Why do we care?



# Large graphs...

Why do we care?

Useful model

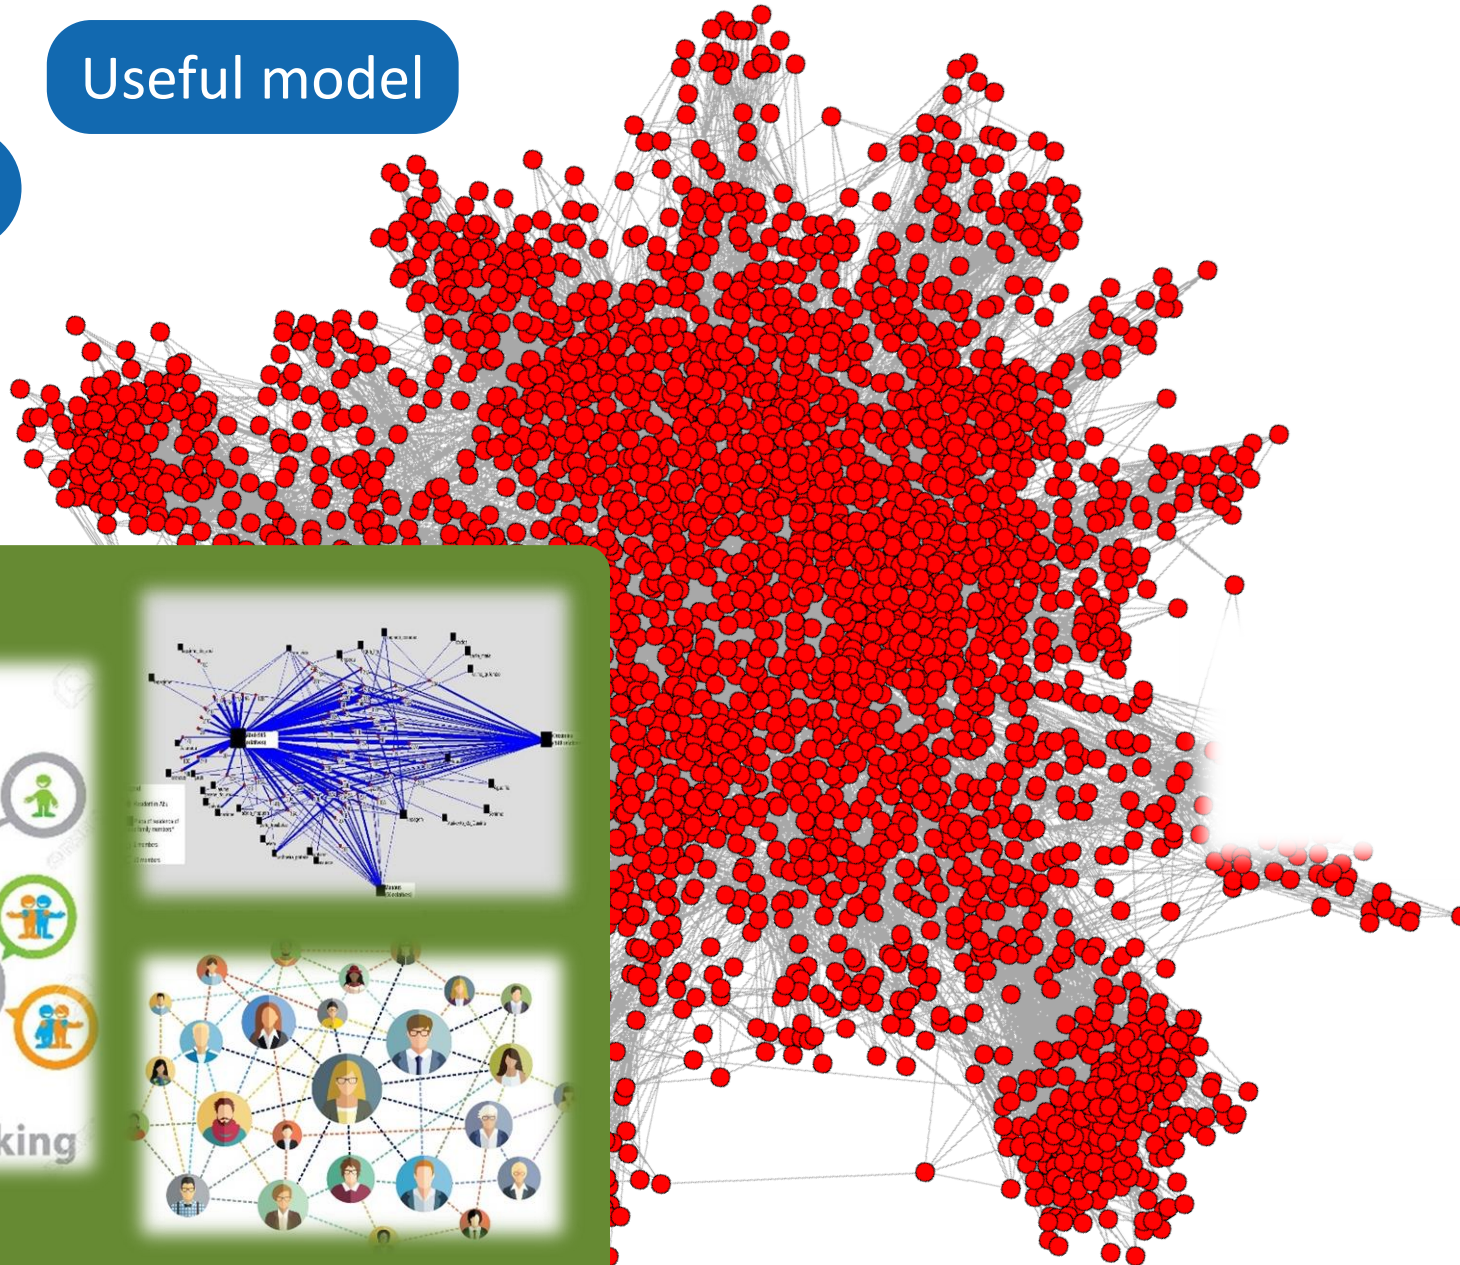
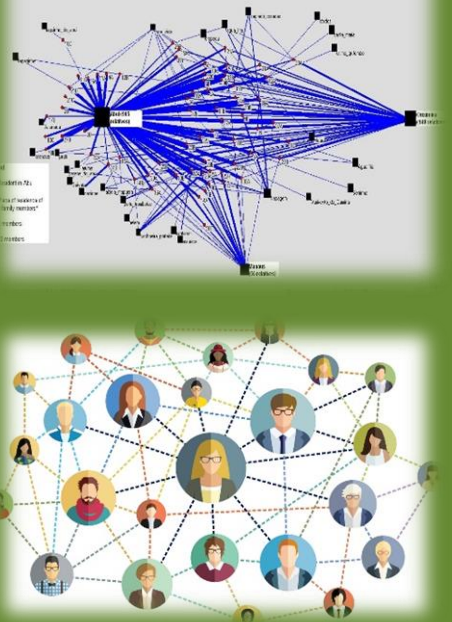


# Large graphs...

Useful model

Why do we care?

## Social networks

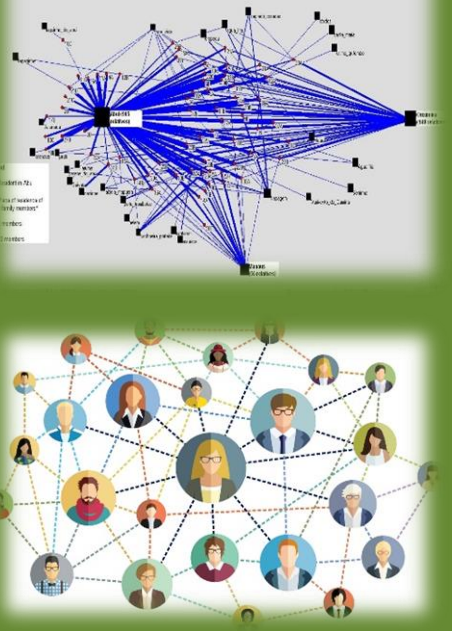


# Large graphs...

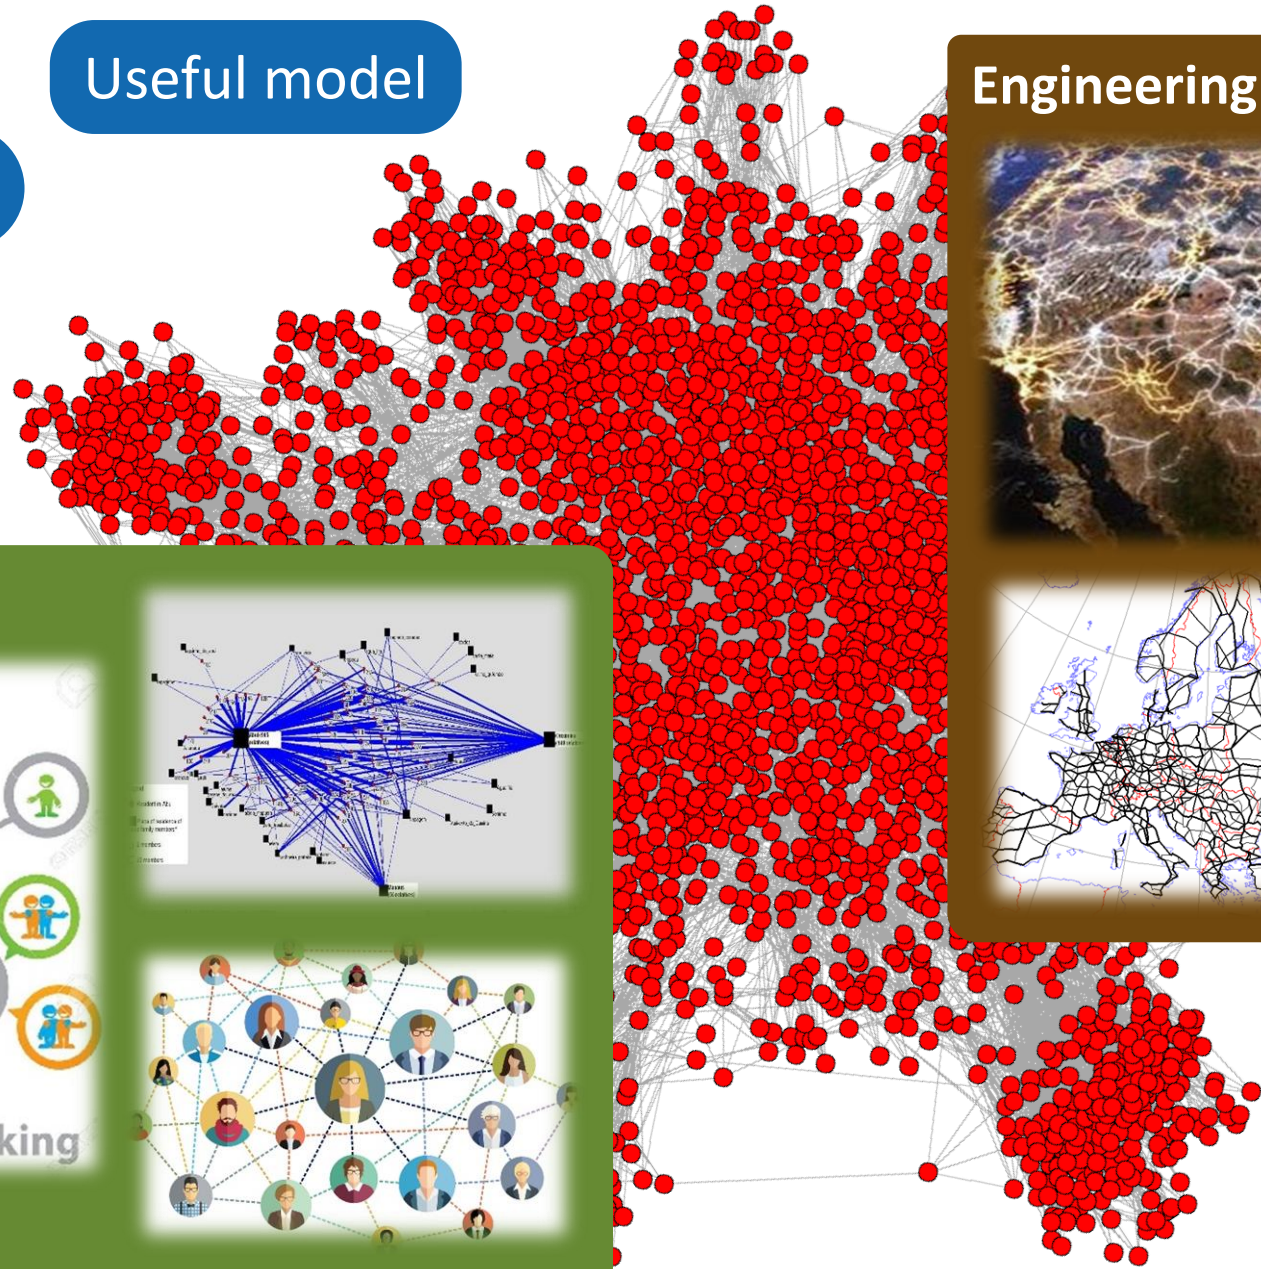

Useful model

Why do we care?

## Social networks



## Engineering networks



# Large graphs...




Useful model

Why do we care?

## Social networks



## Biological networks



## Engineering networks





# Large graphs...

# Useful model



# Why do we care?

# Social networks



Networking

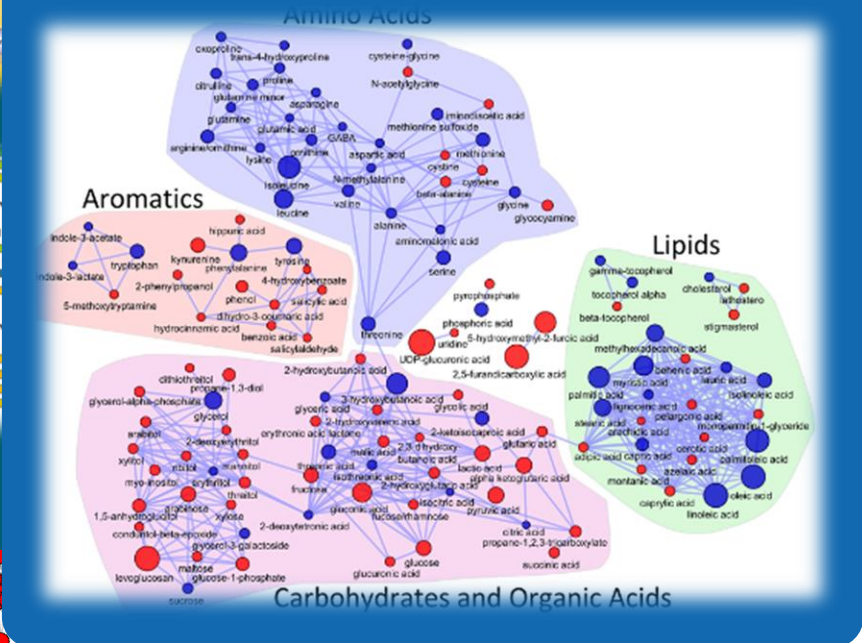
# Biological networks



# Engineering networks

# Physics, chemistry

$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



# Large graphs...

Why do we care?

Useful model

Engineering networks

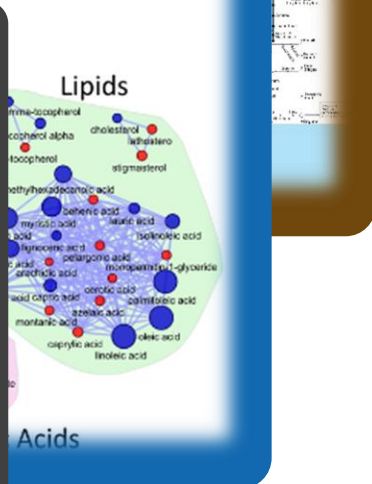

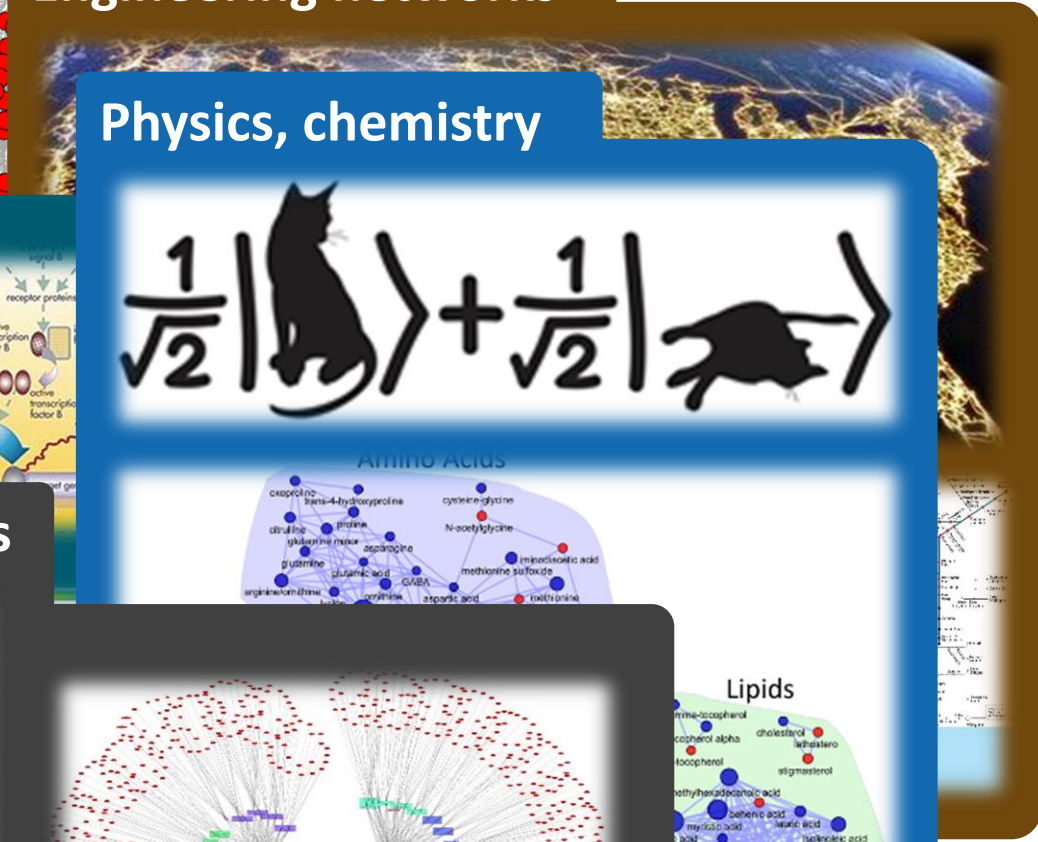


Physics, chemistry

Biological networks

Social networks

Communication networks

Networking



# Large graphs...

## Why do we care?

### Useful model

### Engineering networks






### Physics, chemistry

### Biological networks


### Machine learning

### Social


### Communication networks



$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



Acids



Large graphs...

Why do we care?

Useful model

Engineering networks




...even philosophy ☺

Physics, chemistry

Biological networks

Machine learning

Social



FOSDEM 2016 / [Schedule](#) / [Events](#) / [Developer rooms](#) / [Graph Processing](#) / [Modeling a Philosophical Inquiry: from MySQL to a graph database](#)


## Modeling a Philosophical Inquiry: from MySQL to a graph database

### The short story of a long refactoring process

- 📍 **Track:** Graph Processing devroom
- 📍 **Room:** AW1.126
- 📅 **Day:** Saturday
- 🕒 **Start:** 12:45
- 🕒 **End:** 13:35


[Bruno Latour](#) wrote a book about philosophy (an inquiry into modes of existence). He decided that the [paper book](#) was no place for the numerous footnotes, documentation or glossary, instead giving access to all this information surrounding the book through a [web application](#) which would present itself as a reading companion. He also offered to the community of readers to submit their contributions to his inquiry by writing new documents to be added to the platform. The first version

## Large graphs...



## Large graphs...


One particularly  
important problem



## Large graphs...

One particularly important problem


**Matching:** A set of edges such that no two edges share a common vertex



## Large graphs...

One particularly important problem

Matching: A set of edges such that no two edges share a common vertex






## Large graphs...

One particularly important problem


**Matching:** A set of edges such that no two edges share a common vertex



# Large graphs...

One particularly important problem


**Matching:** A set of edges such that no two edges share a common vertex



## Large graphs...

One particularly important problem

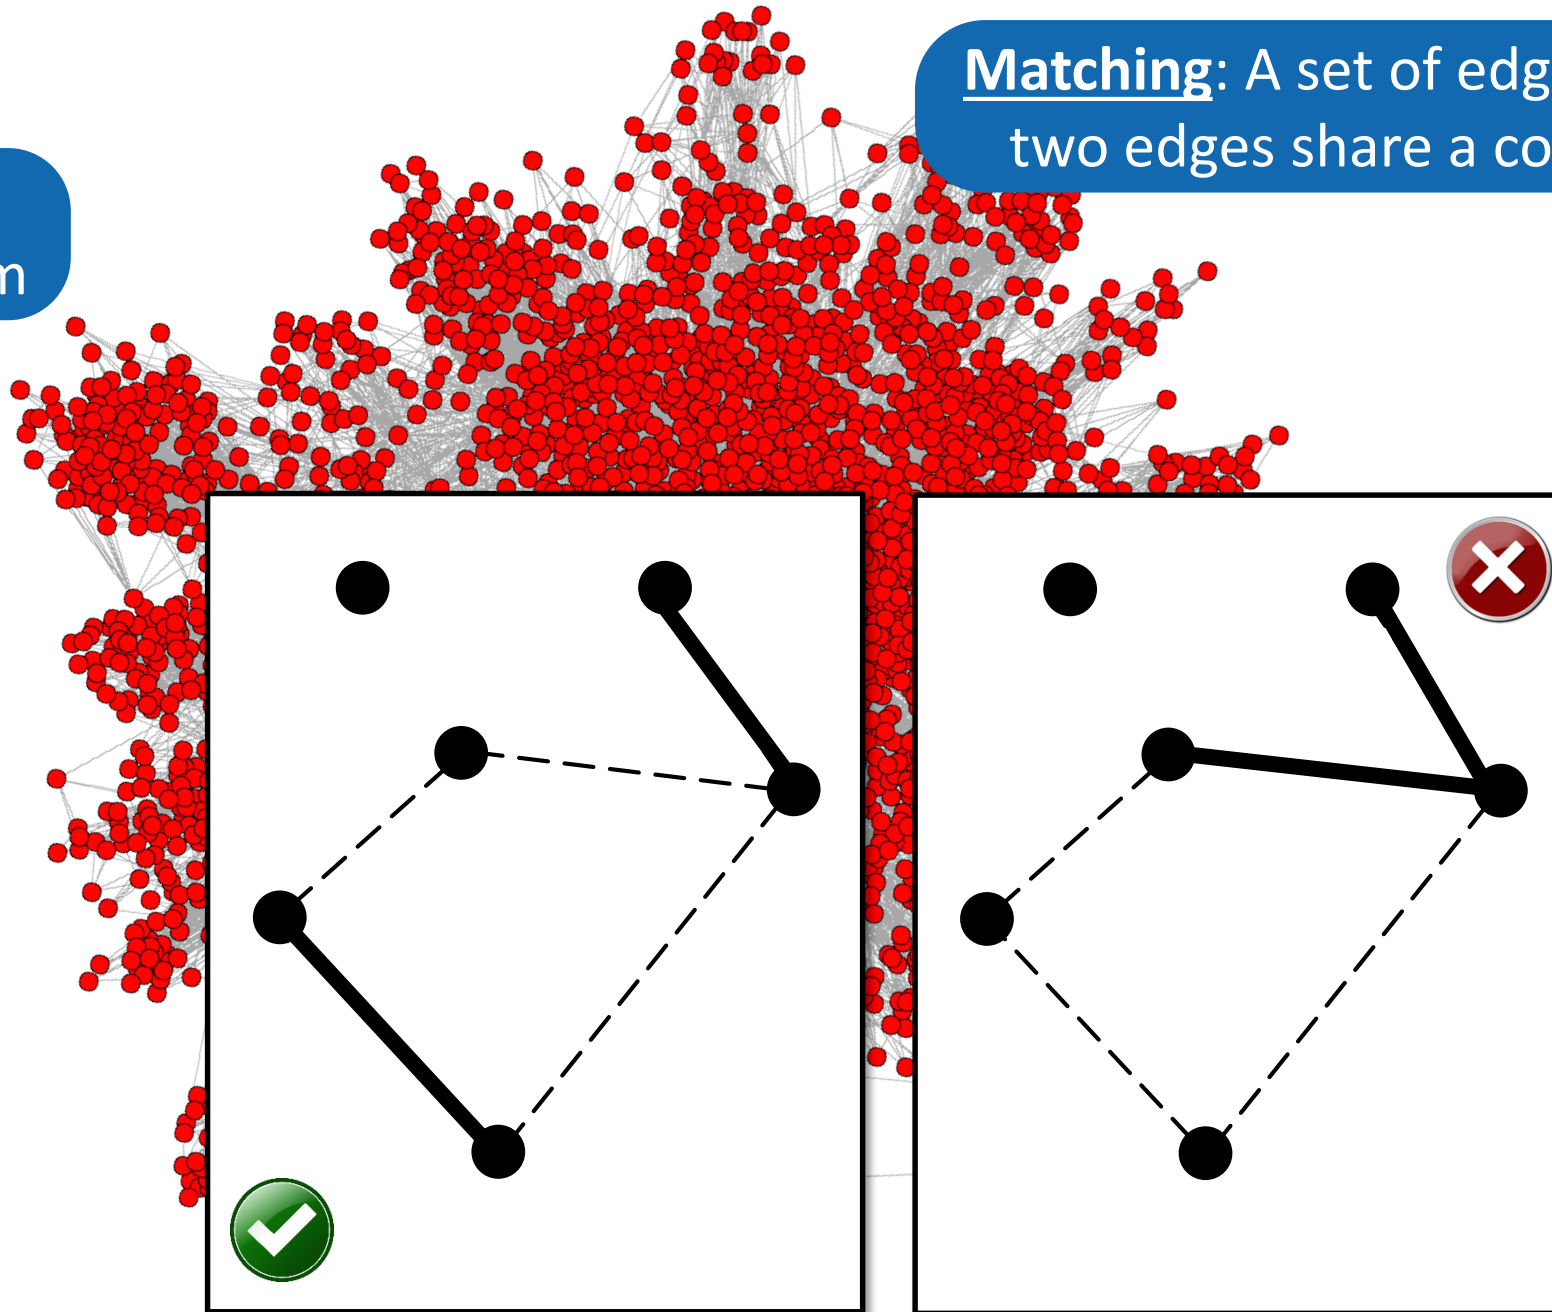
**Matching:** A set of edges such that no two edges share a common vertex



## Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

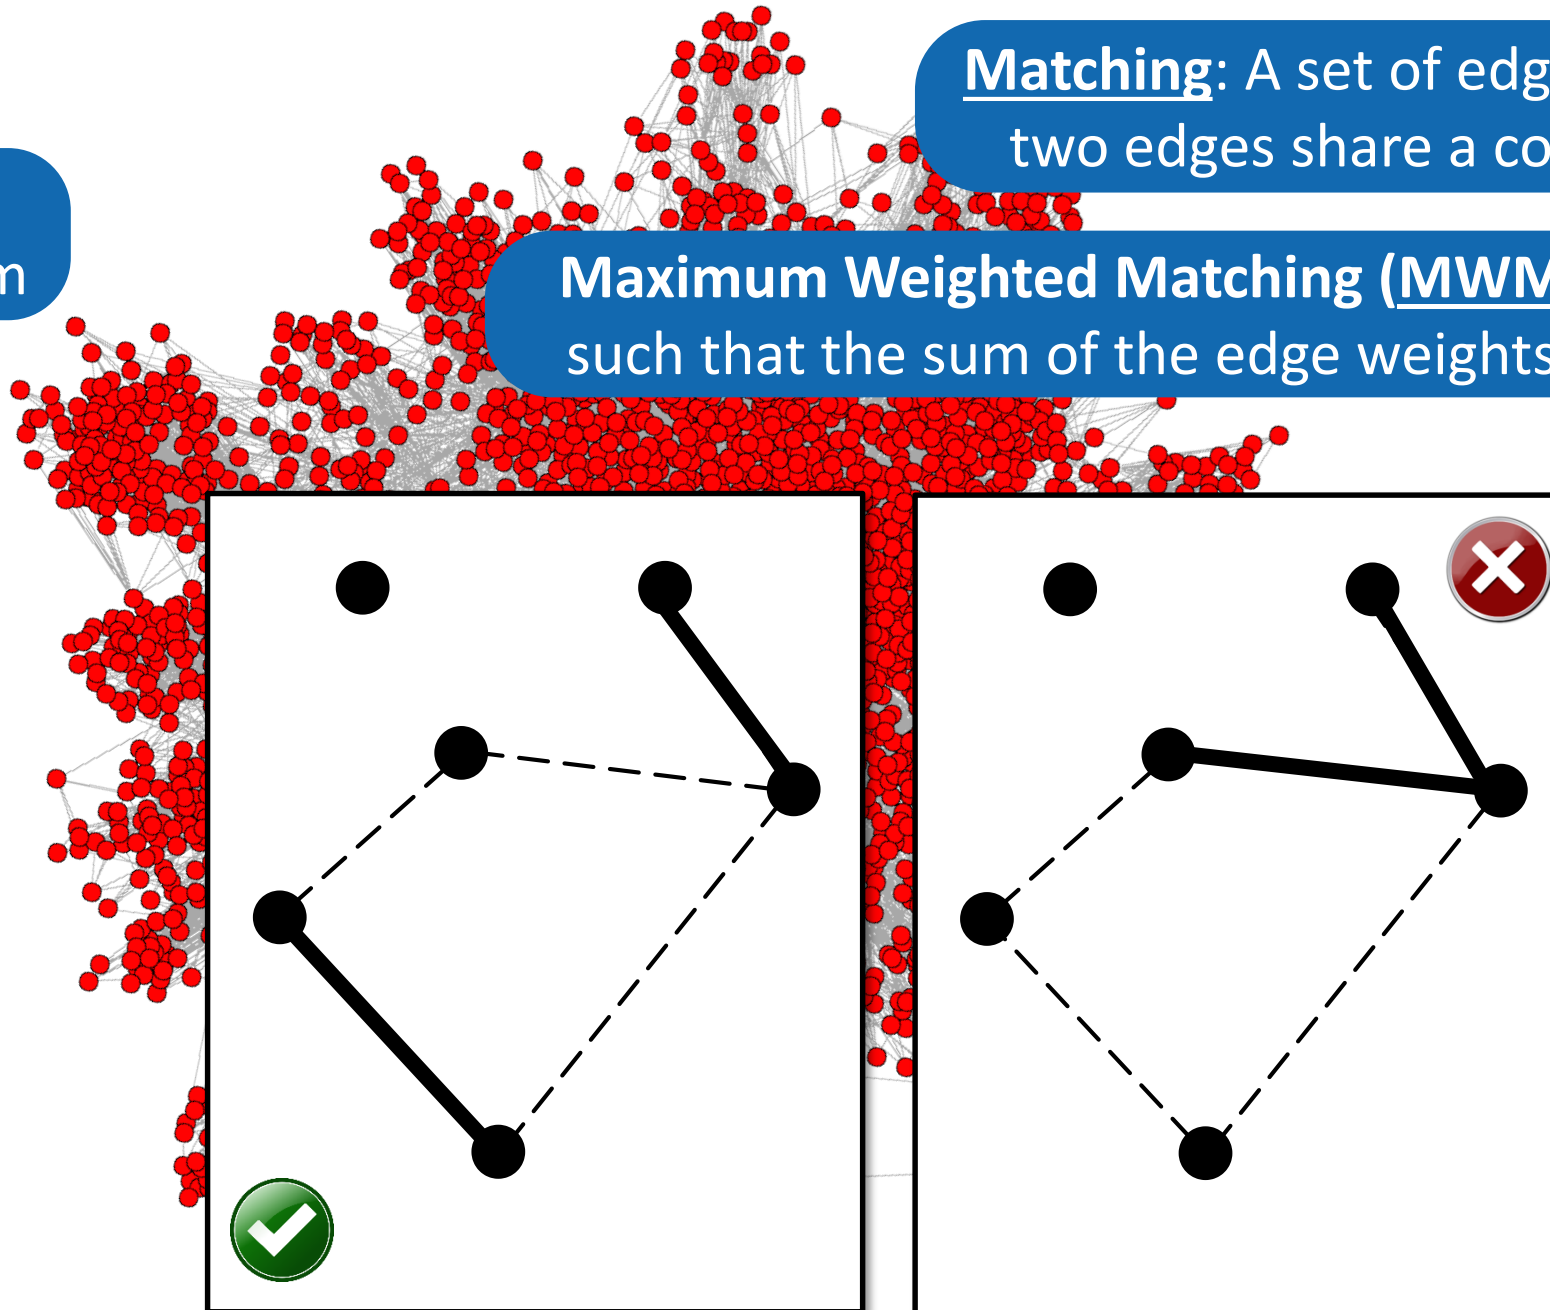


## Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized


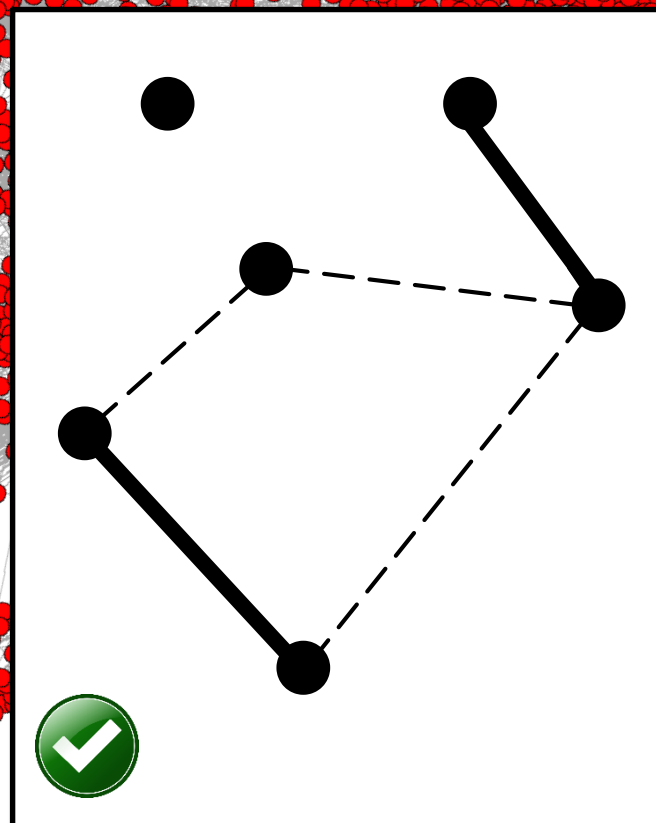



## Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized

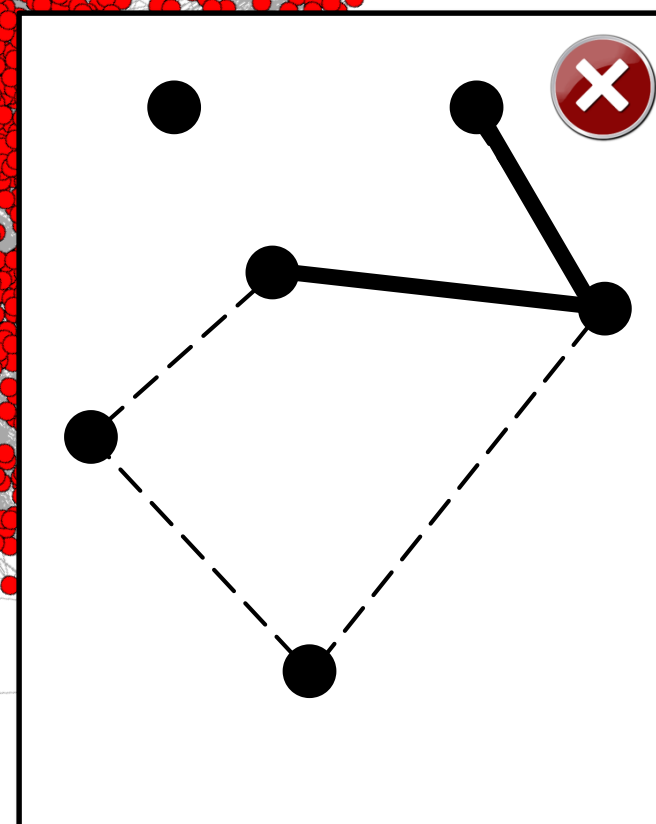




## Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized



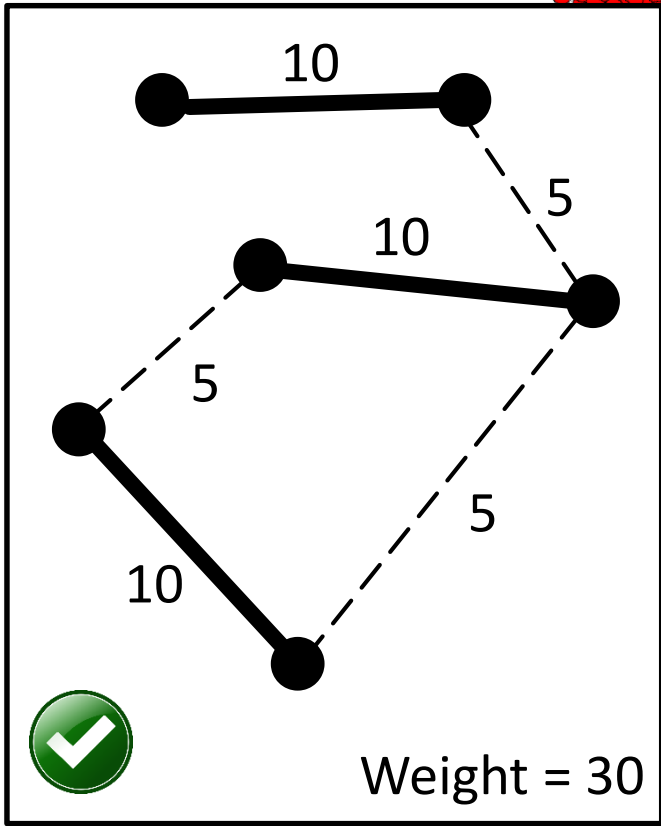


# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized








# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized






# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized


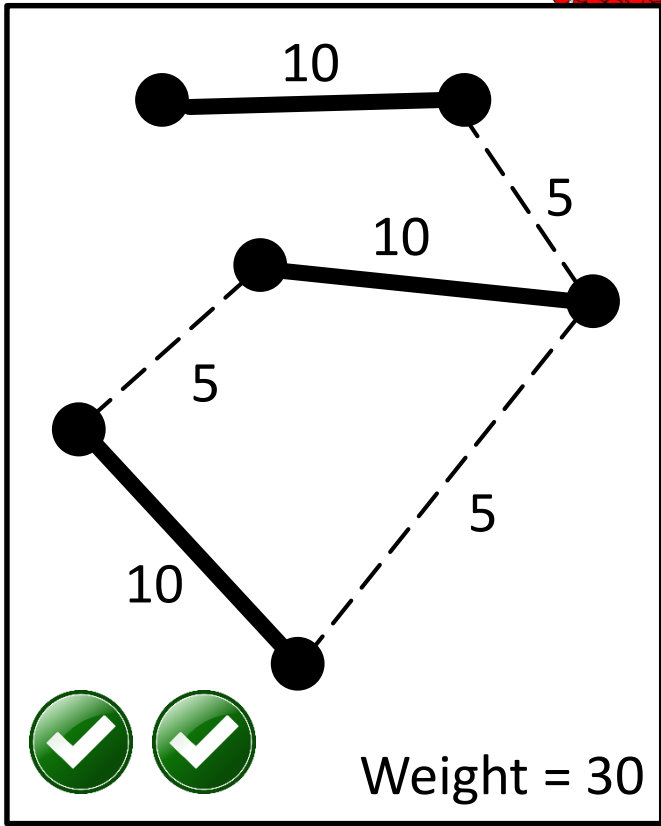


# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized





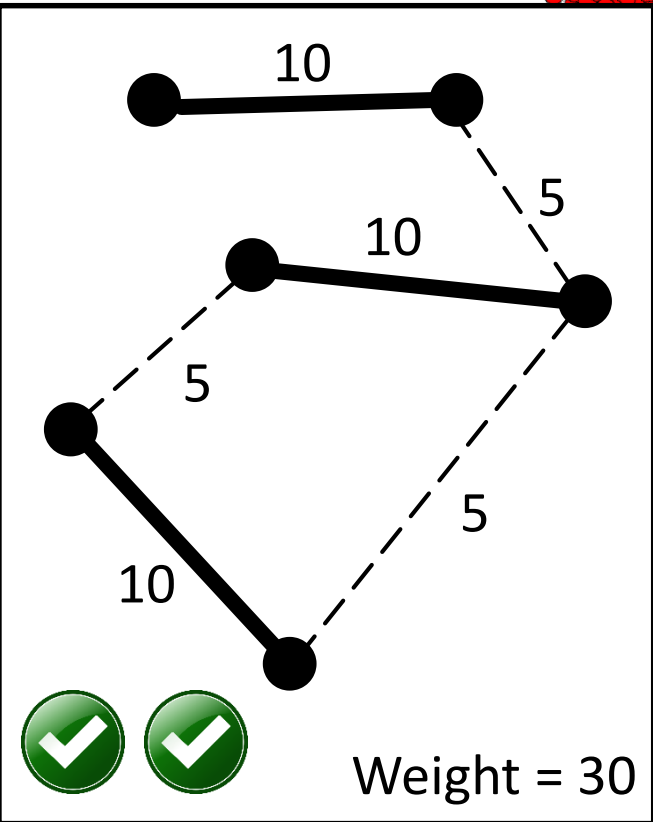
**2-approximation:** here, we have a 2-approximation ( $30/15 = 2$ )

# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized



**2-approximation:** here, we have a 2-approximation ( $30/15 = 2$ )




Approximate MWMs are crucial to many problems

# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized



2-approximation: here, we have a 2-approximation ( $30/15 = 2$ )

Approximate MWMs are crucial to many problems

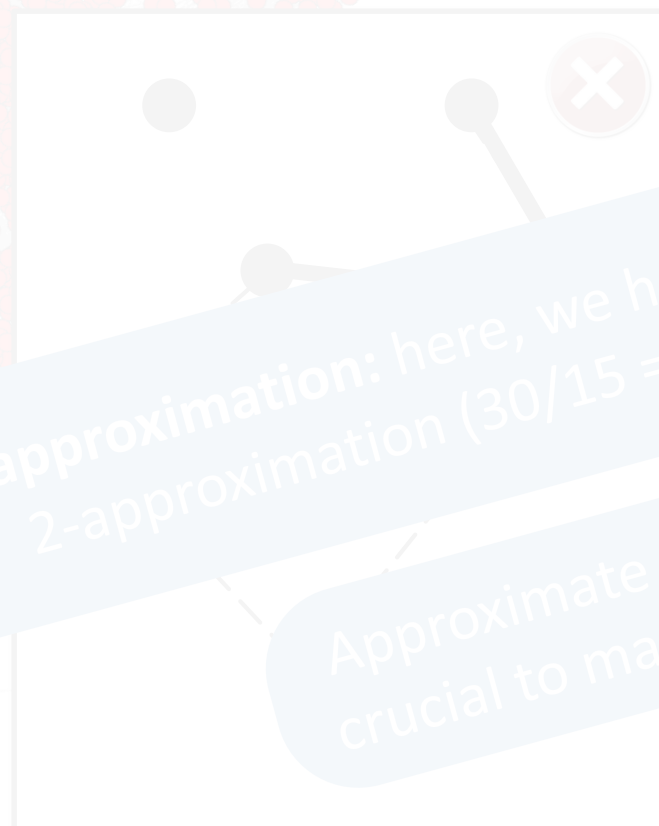

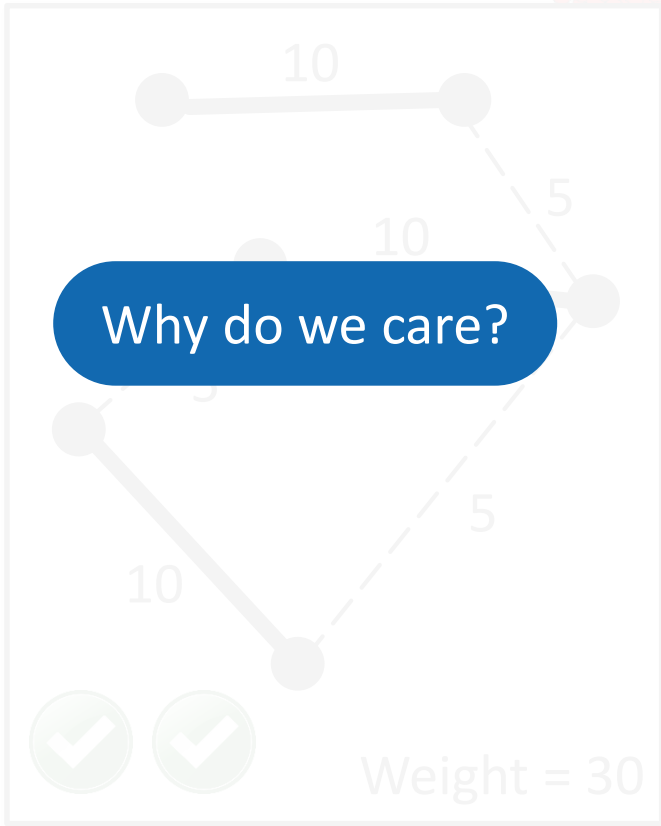


# Large graphs...

One particularly important problem

**Matching:** A set of edges such that no two edges share a common vertex

**Maximum Weighted Matching (MWM):** A matching such that the sum of the edge weights is maximized



2-approximation: here, we have a 2-approximation ( $30/15 = 2$ )

Approximate MWMs are crucial to many problems

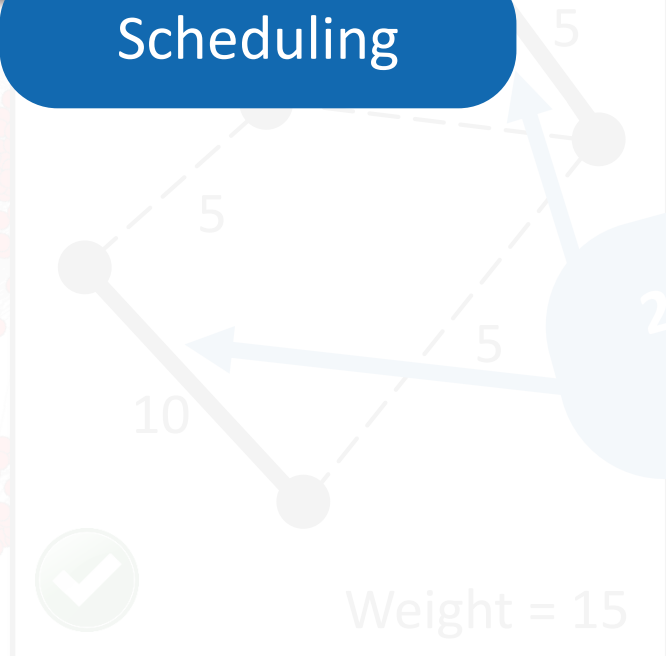
Large

On  
imp



Scheduling

Why do we care?




Matching: A set of edges such that no two edges share a common vertex

Maximum Weighted Matching (MWM): A matching such that the sum of the edge weights is maximized

2-approximation: here, we have a 2-approximation ( $30/15 = 2$ )

Approximate MWMs are crucial to many problems






Scheduling

Why do we care?

[Quantum] error correcting codes



Quantum Error-Correction

Matching: A set of edges such that no two edges share a common vertex

Maximum Weighted Matching (MWM): A matching such that the sum of the edge weights is maximized

2-approximation: here, we have a 2-approximation ( $30/15 = 2$ )

Approximate MWMs are crucial to many problems

Weight = 15



Large


On  
imp



Scheduling

Why do we care?

[Quantum] error correcting codes



Matching: A set of edges such that no two edges share a common vertex

Maximum Weighted Matching (MWM): A matching such that the sum of the edge weights is maximized



Transplant matching

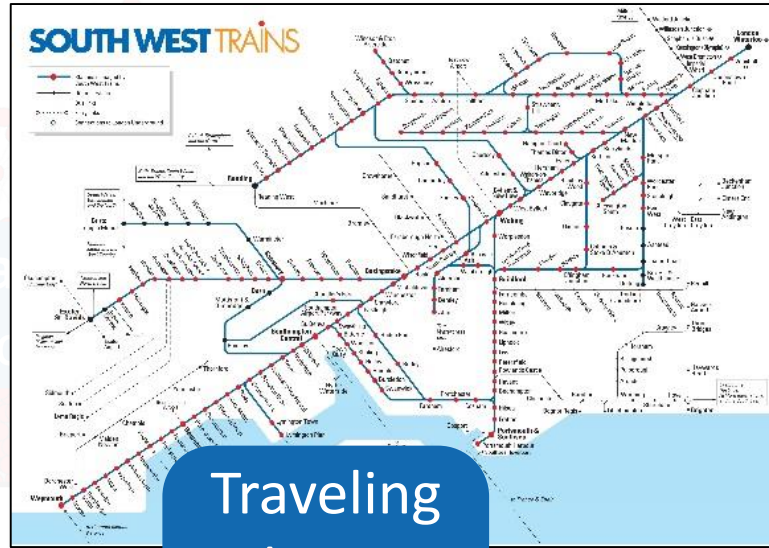


Weight = 15



Scheduling

Why do we care?



Traveling Salesman Problem


edges such that no  
a common vertex  
(MWM): A matching  
weights is maximized



Transplant matching




[Quantum] error correcting codes



Quantum Error-Correction

Weight = 15



Scheduling




Traveling Salesman Problem

edges such that no  
a common vertex  
MWM): A matching  
neighborhood

Many, many others...

Why do we care?



[Quantum] error correcting codes



Transplant matching





In all cases, approximations („reasonably” accurate) are useful



Traveling Salesman Problem


Scheduling

Why do we care?

“We live in a system of approximations” — Ralph Waldo Emerson

Many, many others...

[Quantum] error correcting codes



Quantum Error-Correction



Transplant matching



Weight = 15

# Research Questions

# Research Questions



Which programming paradigm to use for (approximate) MWM (and other graph problems)?

## Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Which programming paradigm to use for (approximate) MWM (and other graph problems)?

## Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Which programming paradigm to use for (approximate) MWM (and other graph problems)?



What is the HW FPGA design that ensures high performance?



## Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Which programming paradigm to use for (approximate) MWM (and other graph problems)?



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?




Which programming paradigm to use for (approximate) MWM (and other graph problems)?




What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?






What programming paradigm and why?



What programming paradigm and why?

Part 1: Seeking “the best paradigm”, we conducted a detailed analysis of graph processing on FPGAs



What programming paradigm and why?



What programming paradigm and why?

~15 FPGA graph processing frameworks

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k
Kapre [70] (GraphSoC)	ASAP'15	👍	SpMV	Vertex-Centric, BSP	C++ (HLS)	👍	BRAM	17k	126k
Dai [40] (FPGP)	FPGA'16	👍	BFS	None	unsp.	👍	DRAM	41.6M	1.4B
Oguntebi [93] (GraphOps)	FPGA'16	👍	BFS, SpMV, PR, Vertex Cover	None	MaxJ (HLS)	👎	BRAM	16M	128M
Zhou [134]	FCCM'16	👍	SSSP, WCC, MST	Edge-Centric	unsp.	👎	DRAM	4.7M	65.8M
Engelhardt [49] (GraVF)	FPL'16	👍	BFS, PR, SSSP, CC	Vertex-Centric	Migen (HLS)	👎	BRAM	128k	512k
Dai [41] (ForeGraph)	FPGA'17	👍	PR, BFS, WCC	None	unsp.	👍	DRAM	41.6M	1.4B
Zhou [136]	SBAC-PAD'17	👍	BFS, SSSP	Hybrid (Vertex- and Edge-Centric)	unsp.	👎	DRAM	10M	160M
Ma [85]	FPGA'17	👍	BFS, SSSP, CC, TC, BC	Transactional Memory [16, 59]	System-Verilog	👍	DRAM	24M	58M
Lee [79] (ExtraV)	FPGA'17	👍	BFS, PR, CC, AT* [60]	Graph Virtualization	C++ (HLS)	👎	DRAM	124M	1.8B
Zhou [135]	CF'18	👍	SpMV, PR	Edge-Centric, GAS	unsp.	👎	DRAM	41.6M	1.4B
Yang [125]	report (2018)	👍	BFS, PR, WCC	None	OpenCL	👎	DRAM	4.85M	69M
Yao [127]	report (2018)	👍	BFS, PR, WCC	None	unsp.	👎	BRAM	4.85M	69M

What programming paradigm and why?

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$	
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k	
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k	
Kapre [70] (GraphSoC)	ASAP'11		Babb [4] report (1996) 📄	SSSP	None	Verilog	👍	Hardwired	512	2051
Dai [40] (FPGP)	FPGA'11		Dandalis [43] report (1999) 📄	SSSP	None	unsp.	👍	Hardwired	2048	32k
Oguntebi [93] (GraphOps)	FPGA'11		Tommiska [116] report (2001) 📄	SSSP	None	VHDL	👎	BRAM	64	4096
Zhou [134] (GraVF)	FCCM'16		Mencer [87] FPL'02 📄	Reachability, SSSP	None	PAM-Bloks II	👎	Hardwired (3-state buffers)	88	7744
Engelhardt [49] (ForeGraph)	FPGA'11		Bondhugula [27] IPDPS'06 📄	APSP	Dynamic Program.	unsp.	👎	DRAM	unsp.	
Dai [41] (CyGraph)	FPGA'11		Sridharan [110] TENCN'09 📄	SSSP	None	VHDL	👎	BRAM	64	88
Zhou [136]	SBAC-F		Wang [121] ICFTP'10 📄	BFS	None	SystemC	👎	DRAM	65.5k	1M
Ma [85]	FPGA'11		Betkaoui [21] FTP'11 📄	GC	Vertex-Centric	Verilog	👍	DRAM	300k	3M
Lee [79] (ExtraV)	FPGA'11		Jagadeesh [65] report (2011) 📄	SSSP	None	VHDL	👎	Hardwired	128	466
Zhou [135]	CF'18		Betkaoui [22] FPL'12 📄	APSP	Vertex-Centric	Verilog	👍	≈ DRAM	38k	72M
Yang [125]	report (2016)		Betkaoui [23] ASAP'12 📄	BFS	Vertex-Centric	Verilog	👍	DRAM	16.8M	1.1B
Yao [127]	report (2016)		Attia [2] IPDPS'14 📄	BFS	Vertex-Centric	VHDL	👍	DRAM	8.4M	536M
			Ni [91] report (2014) 📄	BFS	None	Verilog	👎	DRAM, SRAM	16M	512M
			Zhou [132] IPDPS'15 📄	SSSP	None	unsp.	👎	DRAM	1M	unsp.
			Zhou [133] ReConFig'15 📄	PR	Edge-Centric	unsp.	👎	DRAM	2.4M	5M
			Umuroglu [117] FPL'15 📄	BFS	None	Chisel	👎	≈ DRAM	2.1M	65M
			Lei [80] report (2016) 📄	SSSP	None	unsp.	👎	DRAM	23.9M	58.2M
			Zhang [129] FPGA'17 📄	BFS	MapReduce	unsp.	👎	HMC	33.6M	536.9M
			Zhang [130] FPGA'18 📄	BFS	None	unsp.	👎	HMC		
			Kohram [76] FPGA'18 📄	BFS	None	unsp.	👎	HMC		
			Besta [13] FPGA'19 📄	MM	Substream-Centric	Verilog	👎	DRAM	4.8M	117M

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms



What programming paradigm and why?

Key techniques, paradigms, challenges, features, ...

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$	
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k	
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k	
Kapre [70] (GraphSoC)	ASAP'11		Babb [4] report (1996) 📄	SSSP	None	Verilog	👍	Hardwired	512	2051
Dai [40] (FPGP)	FPGA'11		Dandalis [43] report (1999) 📄	SSSP	None	unsp.	👍	Hardwired	2048	32k
Oguntebi [93] (GraphOps)	FPGA'11		Tommiska [116] report (2001) 📄	SSSP	None	VHDL	👎	BRAM	64	4096
Zhou [134] (GraVF)	FCCM'16		Mencer [87] FPL'02 📄	Reachability, SSSP	None	PAM-Bloks II	👎	Hardwired (3-state buffers)	88	7744
Engelhardt [49] (ForeGraph)	FPGA'11		Bondhugula [27] IPDPS'06 📄	APSP	Dynamic Program.	unsp.	👎	DRAM	unsp.	
Dai [41] (CyGraph)	FPGA'11		Sridharan [110] TENCN'09 📄	SSSP	None	VHDL	👎	BRAM	64	88
Zhou [136]	SBAC-F		Wang [121] ICFTP'10 📄	BFS	None	SystemC	👎	DRAM	65.5k	1M
Ma [85] (ExtraV)	FPGA'11		Betkaoui [21] FTP'11 📄	GC	Vertex-Centric	Verilog	👍	DRAM	300k	3M
Lee [79] (ExtraV)	FPGA'11		Jagadeesh [65] report (2011) 📄	SSSP	None	VHDL	👎	Hardwired	128	466
Zhou [135]	CF'18		Betkaoui [22] FPL'12 📄	APSP	Vertex-Centric	Verilog	👍	≈ DRAM	38k	72M
Yang [125]	report		Betkaoui [23] ASAP'12 📄	BFS	Vertex-Centric	Verilog	👍	DRAM	16.8M	1.1B
Yao [127]	report		Attia [2] IPDPS'14 📄	BFS	Vertex-Centric	VHDL	👍	DRAM	8.4M	536M
			Ni [91] report (2014) 📄	BFS	None	Verilog	👎	DRAM, SRAM	16M	512M
			Zhou [132] IPDPS'15 📄	SSSP	None	unsp.	👎	DRAM	1M	unsp.
			Zhou [133] ReConFig'15 📄	PR	Edge-Centric	unsp.	👎	DRAM	2.4M	5M
			Umuroglu [117] FPL'15 📄	BFS	None	Chisel	👎	≈ DRAM	2.1M	65M
			Lei [80] report (2016) 📄	SSSP	None	unsp.	👎	DRAM	23.9M	58.2M
			Zhang [129] FPGA'17 📄	BFS	MapReduce	unsp.	👎	HMC	33.6M	536.9M
			Zhang [130] FPGA'18 📄	BFS	None	unsp.	👎	HMC		
			Kohram [76] FPGA'18 📄	BFS	None	unsp.	👎	HMC		
			Besta [13] FPGA'19 📄	MM	Substream-Centric	Verilog	👎	DRAM	4.8M	117M

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$		
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k		
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [110]	Vertex-Centric	unsp.	👎	DRAM	110k	221k		
Kapre [70] (GraphSoC)	ASAP'11							Hardwired	512	2051	
Dai [40] (FPGP)	FPGA'11							Hardwired	2048	32k	
Oguntebi [93] (GraphOps)	FPGA'11							BRAM	64	4096	
Zhou [134]	FCCM'11							Hardwired (3-state buffers)	88	7744	
Engelhardt [49] (GraVF)	FPL'16		Bondhugula [27] IPDPS'06	👎	APSP	unsp.	👎	DRAM	unsp.		
Dai [41] (ForeGraph)	FPGA'11		Sridharan [110] TENCON'09	👎	SSSP	VHDL	👎	BRAM	64	88	
Zhou [136]	SBAC-F		Wang [121] ICFTP'10	👎	BFS	SystemC		DRAM	65.5k	1M	
Ma [85]	FPGA'11		Betkaoui [21] FTP'11	👎	GC	Vertex-Centric	👍	DRAM	300k	3M	
Lee [79] (ExtraV)	FPGA'11		Jagadeesh [65] report (2011)	👎	SSSP	None	👎	Hardwired	128	466	
Zhou [135]	CF'18		Betkaoui [22] FPL'12	👎	APSP	Vertex-Centric	👍	≈ DRAM	38k	72M	
Yang [125]	report (		Betkaoui [23] ASAP'12	👎	BFS	Vertex-Centric	👍	DRAM	16.8M	1.1B	
Yao [127]	report (		Attia [2] (CyGraph) IPDPS'14	👎	BFS	Vertex-Centric	👍	DRAM	8.4M	536M	
			Ni [91] report (2014)	👎	BFS	None	👎	DRAM, SRAM	16M	512M	
			Zhou [132] IPDPS'15	👎	SSSP	None	👎	DRAM	1M	unsp.	
			Zhou [133] ReConFig'15	👎	PR	Edge-Centric	👎	DRAM	2.4M	5M	
			Umuroglu [117] FPL'15	👎	BFS	None		Chisel	≈ DRAM	2.1M	65M
			Lei [80] report (2016)	👎	SSSP	None	👎	DRAM	23.9M	58.2M	
			Zhang [129] FPGA'17	👎	BFS	MapReduce	👎	HMC	33.6M	536.9M	
			Zhang [130] FPGA'18	👎	BFS	None	👎	HMC			
			Kohram [76] FPGA'18	👎	BFS	None	👎	HMC			
			Besta [13] FPGA'19	👎	MM	Substream-Centric	👎	DRAM	4.8M	117M	

Selected parts are in the FPGA paper, the rest is in...

? What programming paradigm and why?

Key techniques, paradigms, challenges, features, ...

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup>	Programming Model or Technique <sup>4</sup> (§ 2.4)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [110]	Vertex-Centric	unsp.	👎	DRAM	110k	221k

Selected parts are in the FPGA paper, the rest is in...

Kapre [70] (GraphSoC)	ASAP'11						Hardwired	512	2051
Dai [40] (FPGP)	FPGA'11						Hardwired	2048	32k
Oguntebi [93] (GraphOps)	FPGA'11						BRAM	64	4096
Zhou [134]	FCCM'11						Hardwired (3-state buffers)	88	7744
Engelhardt [49] (GraVF)	FPL'16						DRAM	unsp.	
Dai [41] (ForeGraph)							BRAM	64	88
							DRAM	65.5k	1M

Zhou [136]								300k	3M
Ma [85]								128	466
Lee [79] (ExtraV)								38k	72M
Zhou [135]								16.8M	1.1B
Yang [125]								8.4M	536M
Yao [127]								16M	512M
								1M	unsp.
								2.4M	5M
								2.1M	65M
								23.9M	58.2M
								33.6M	536.9M
								4.8M	117M

? What programming paradigm and why?

Key techniques, paradigms, challenges, features, ...

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field Programmable Gate Arrays (FPGAs) can be an energy-efficient solution to deliver specialized hardware for



Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [110]	Vertex-Centric	unsp.	👎	DRAM	110k	221k

Selected parts are in the FPGA paper, the rest is in...

<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
 (submitted to arXiv, will appear tonight)

Kapre [70] (GraphSoC)	ASAP'11						Hardwired	512	2051
Dai [40] (FPGP)	FPGA'11						Hardwired	2048	32k
Oguntebi [93] (GraphOps)	FPGA'11						BRAM	64	4096
Zhou [134]	FCCM'11						Hardwired		
Engelhardt [49] (GraVF)	FPGA'11						(3-state buffers)	88	7744
Dai [41] (ForeGraph)	FCCM'11						DRAM	unsp.	
Zhou [136]	FCCM'11						BRAM	64	88
Ma [85]	FCCM'11						BRAM	65.5k	1M
Lee [79] (ExtraV)	FCCM'11						DRAM	300k	3M
Zhou [135]	FCCM'11							128	466
Yang [125]	FCCM'11							38k	72M
Yao [127]	FCCM'11							16.8M	1.1B
								8.4M	536M
								16M	512M
								1M	unsp.
								2.4M	5M
								2.1M	65M
								23.9M	58.2M
								33.6M	536.9M
								4.8M	117M

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich


Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field Programmable Gate Arrays (FPGAs) can be an energy-efficient solution to deliver specialized hardware for

? What programming paradigm and why?

Key techniques, paradigms, challenges, features, ...

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms



What programming paradigm and why?

“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]


+ other issues

[1] S. Salihoglu and J. Widom,  
“Optimizing graph algorithms on  
Pregel-like systems”. VLDB. 2014.

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



What programming  
paradigm and why?



“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]

+ other issues


[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




What programming paradigm and why?

To be able to utilize pipelining well, we really want to use streaming (aka edge-centric)




What programming paradigm and why?






Use some form of streaming  
(aka **edge-centric**); we can  
use pipelining efficiently  
("streaming  $\approx$  pipelining")




Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming  $\approx$  pipelining”)

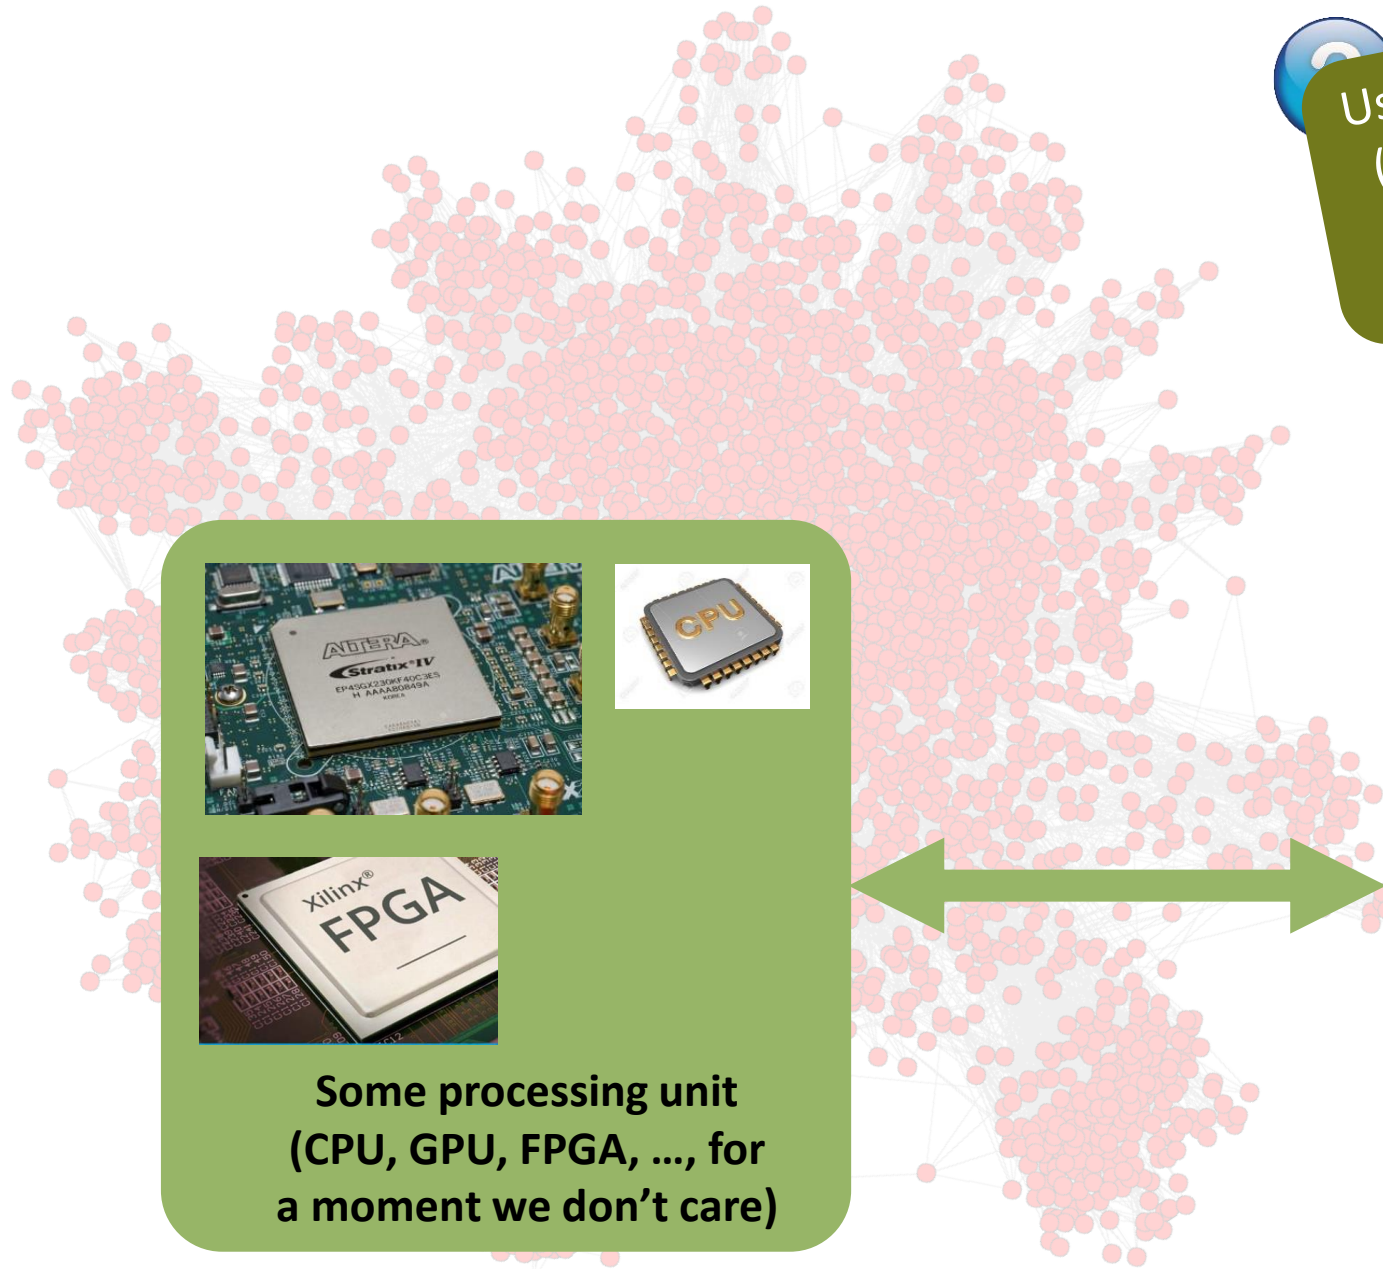




Use some form of streaming  
 (aka **edge-centric**); we can  
 use pipelining efficiently  
 ("streaming  $\approx$  pipelining")



Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

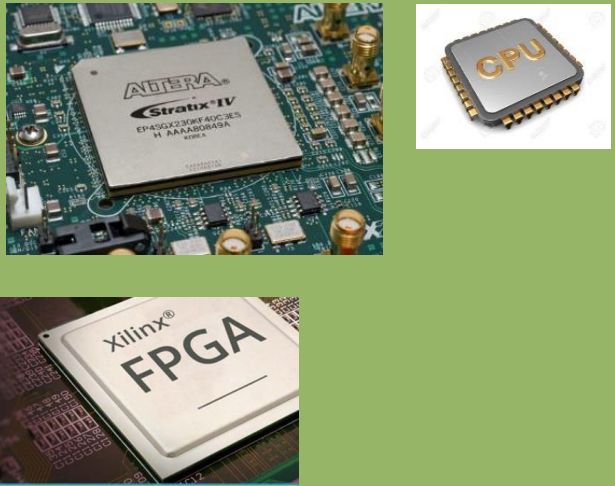
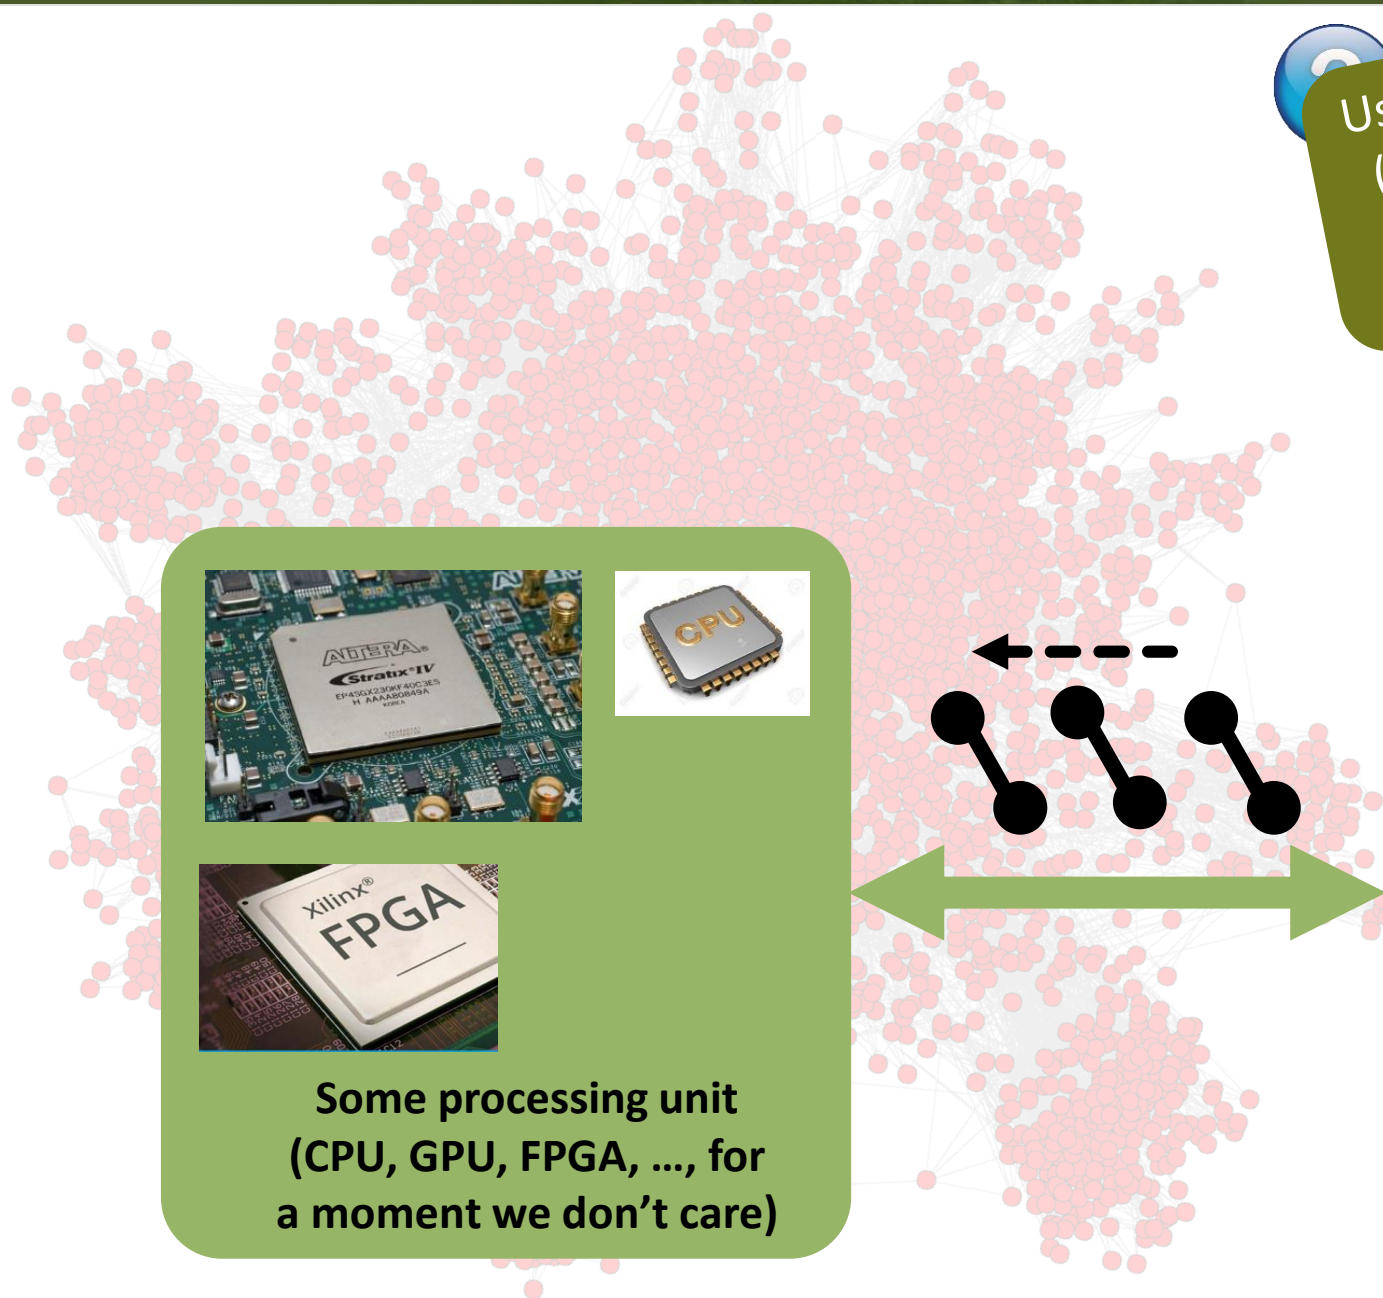


Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)


DRAM



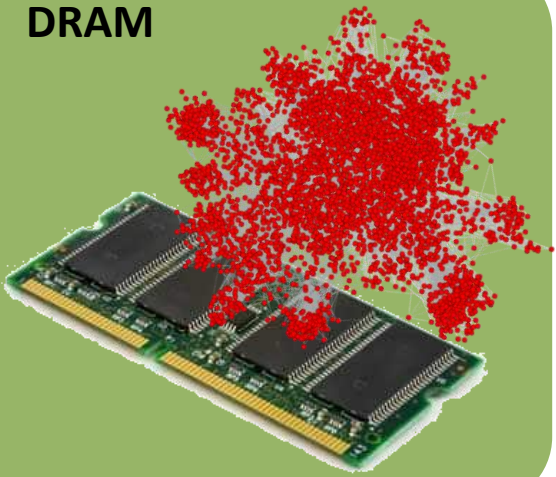
Use some form of streaming  
(aka **edge-centric**); we can  
use pipelining efficiently  
("streaming  $\approx$  pipelining")



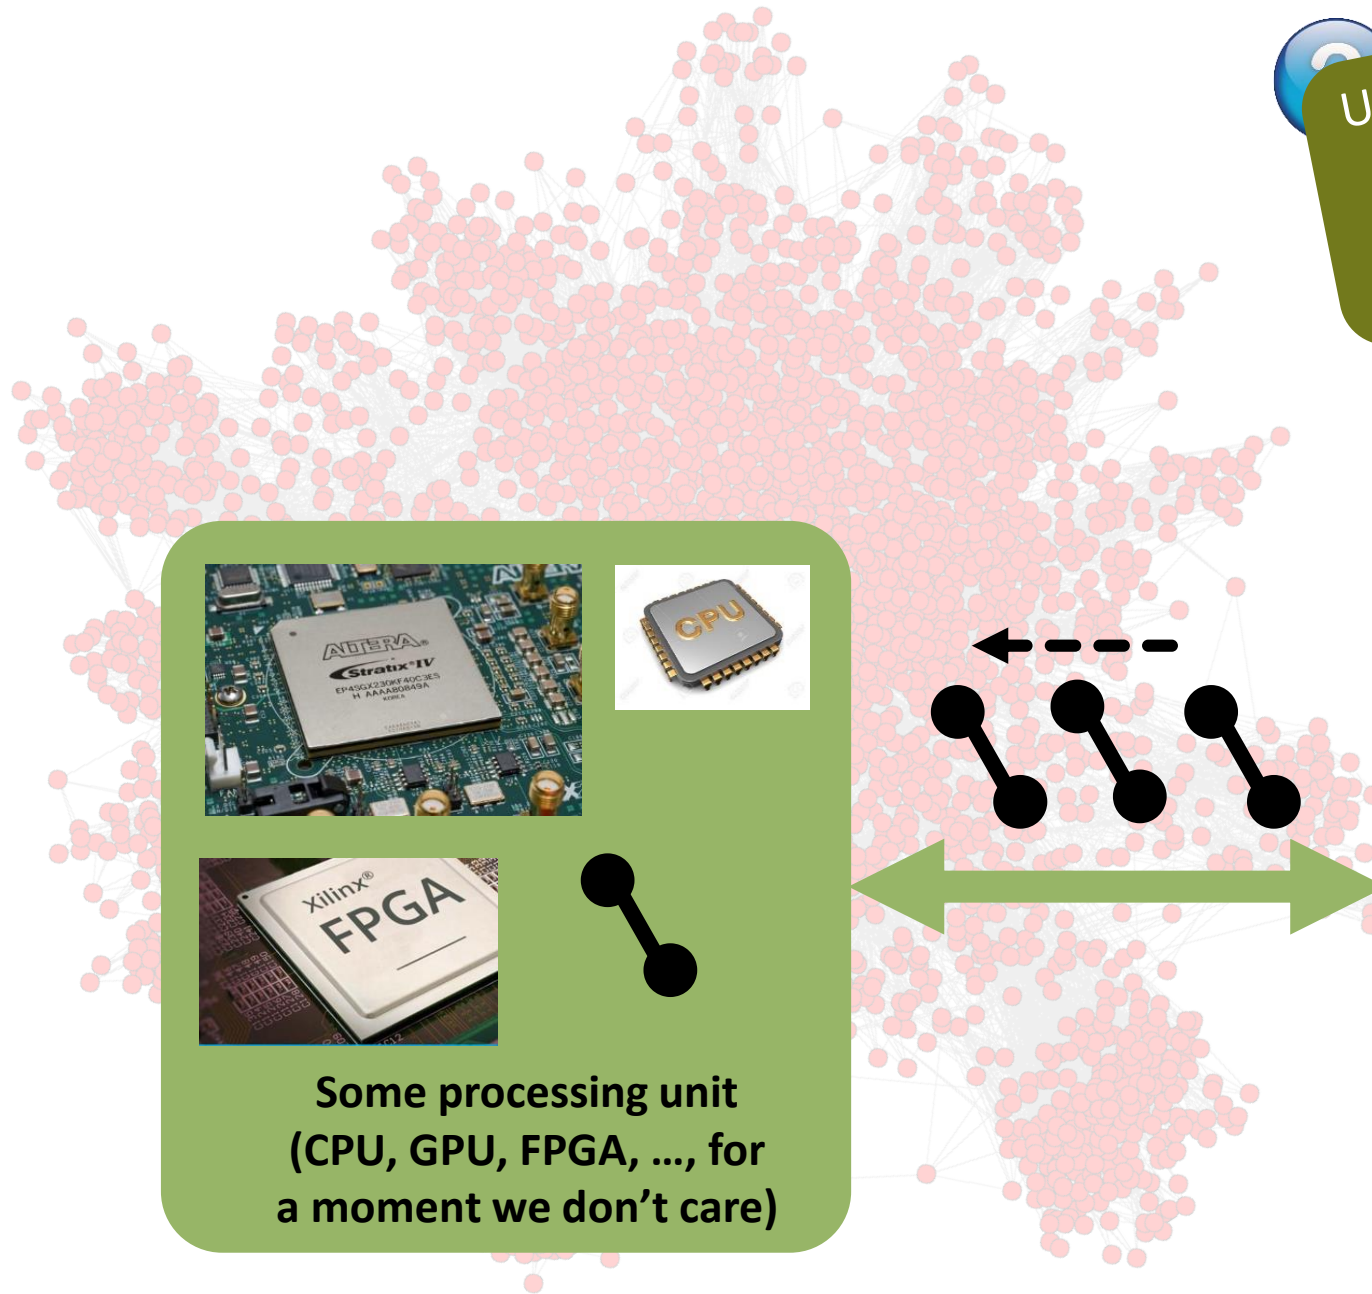
Some processing unit  
(CPU, GPU, FPGA, ..., for  
a moment we don't care)



DRAM





Use some form of streaming  
(aka **edge-centric**); we can  
use pipelining efficiently  
("streaming  $\approx$  pipelining")



Some processing unit  
(CPU, GPU, FPGA, ..., for  
a moment we don't care)


DRAM

Use some form of streaming  
(aka **edge-centric**); we can  
use pipelining efficiently  
("streaming  $\approx$  pipelining")

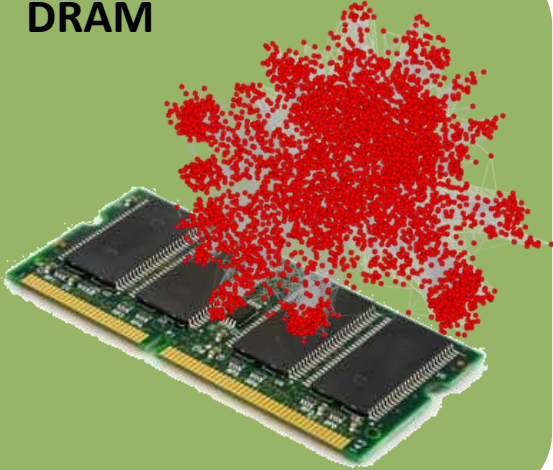


Some processing unit  
(CPU, GPU, FPGA, ..., for  
a moment we don't care)

This block contains images of an Altera Stratix IV FPGA, a CPU chip, and a Xilinx FPGA. A circular arrow icon is positioned between the FPGA images.

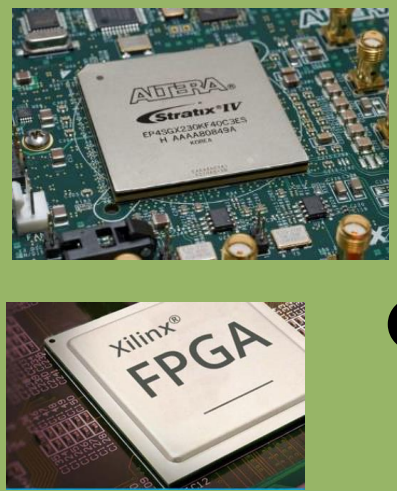
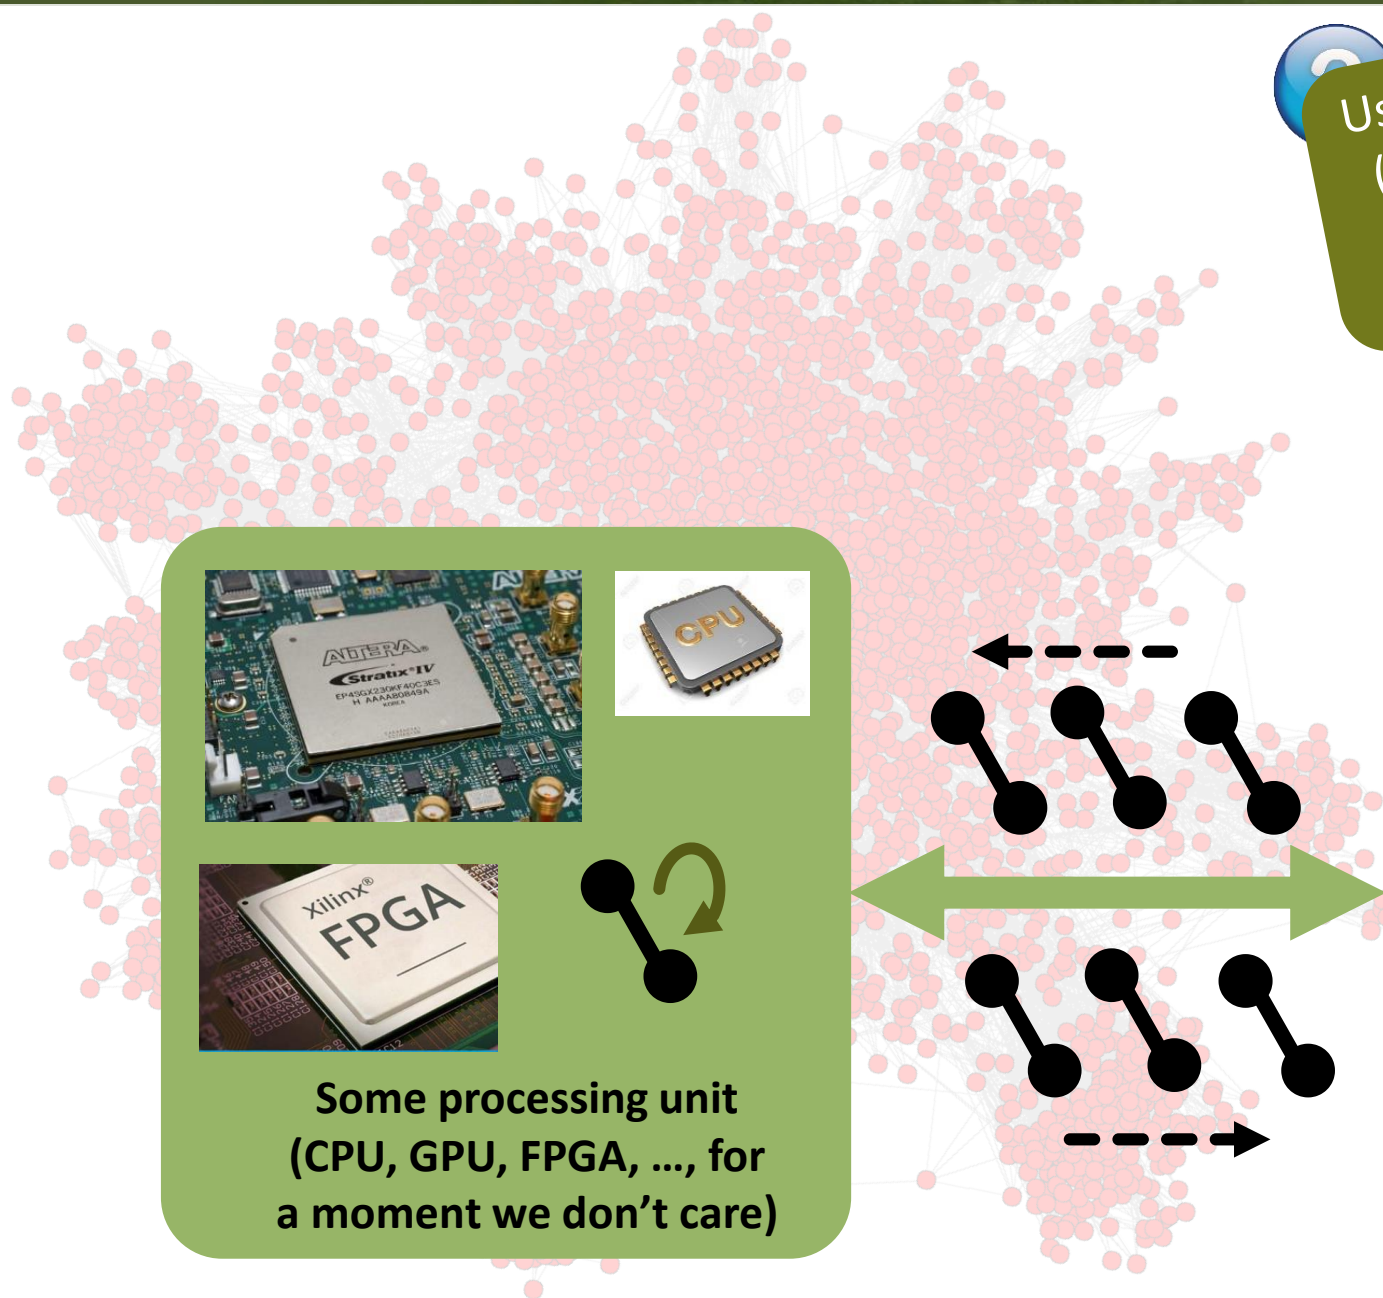


DRAM



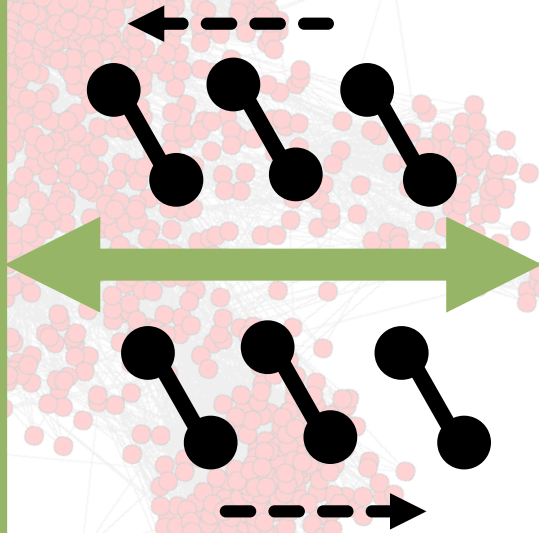
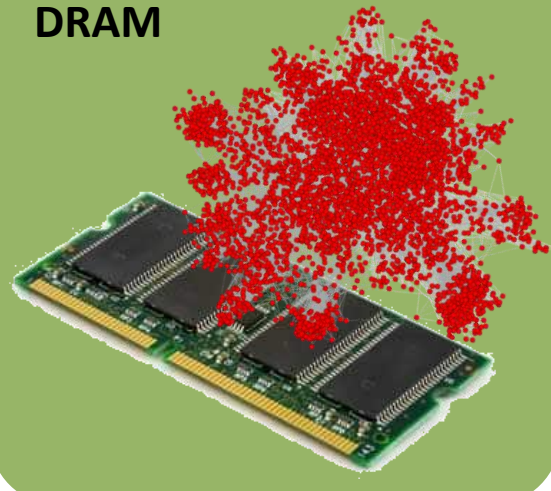
A DRAM module is shown with a cluster of red nodes and edges overlaid on top of it.

Use some form of streaming  
(aka **edge-centric**); we can  
use pipelining efficiently  
("streaming  $\approx$  pipelining")



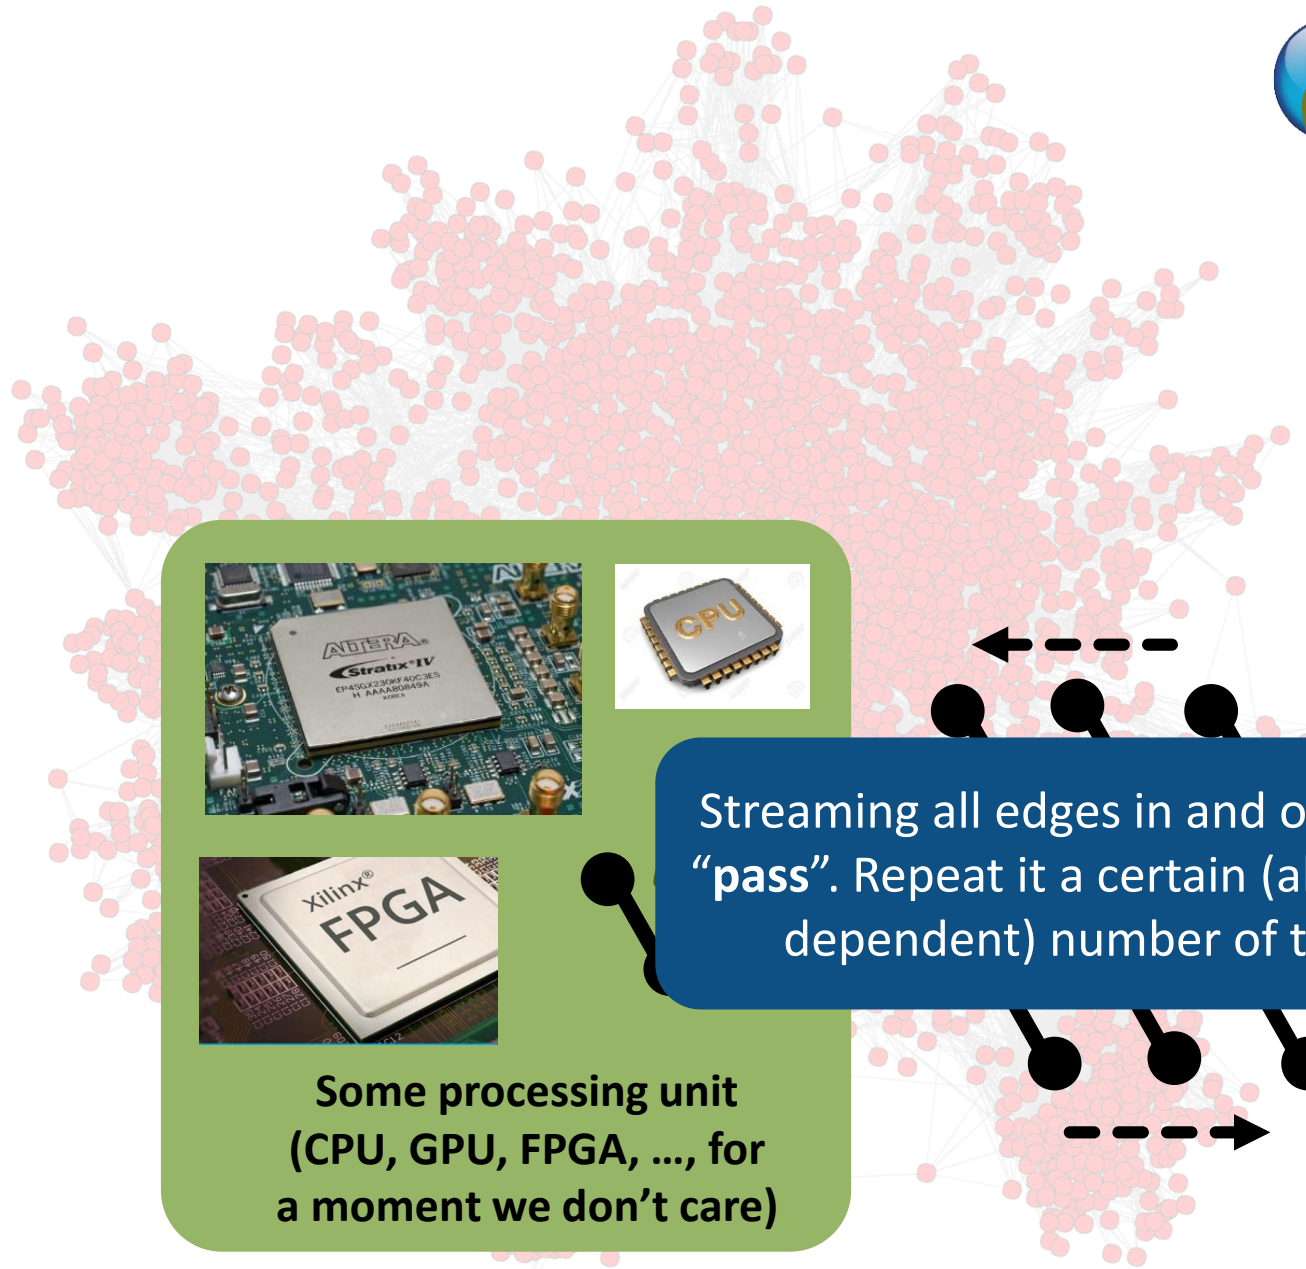
Some processing unit  
(CPU, GPU, FPGA, ..., for  
a moment we don't care)

DRAM





Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming  $\approx$  pipelining”)



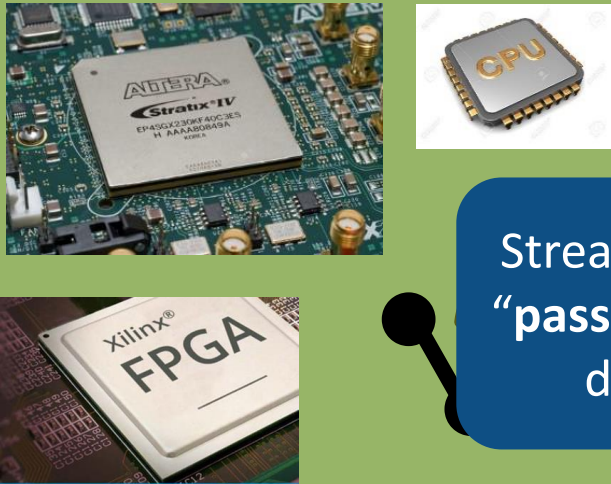
Streaming all edges in and out is one “**pass**”. Repeat it a certain (algorithm-dependent) number of times

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

DRAM

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming  $\approx$  pipelining”)

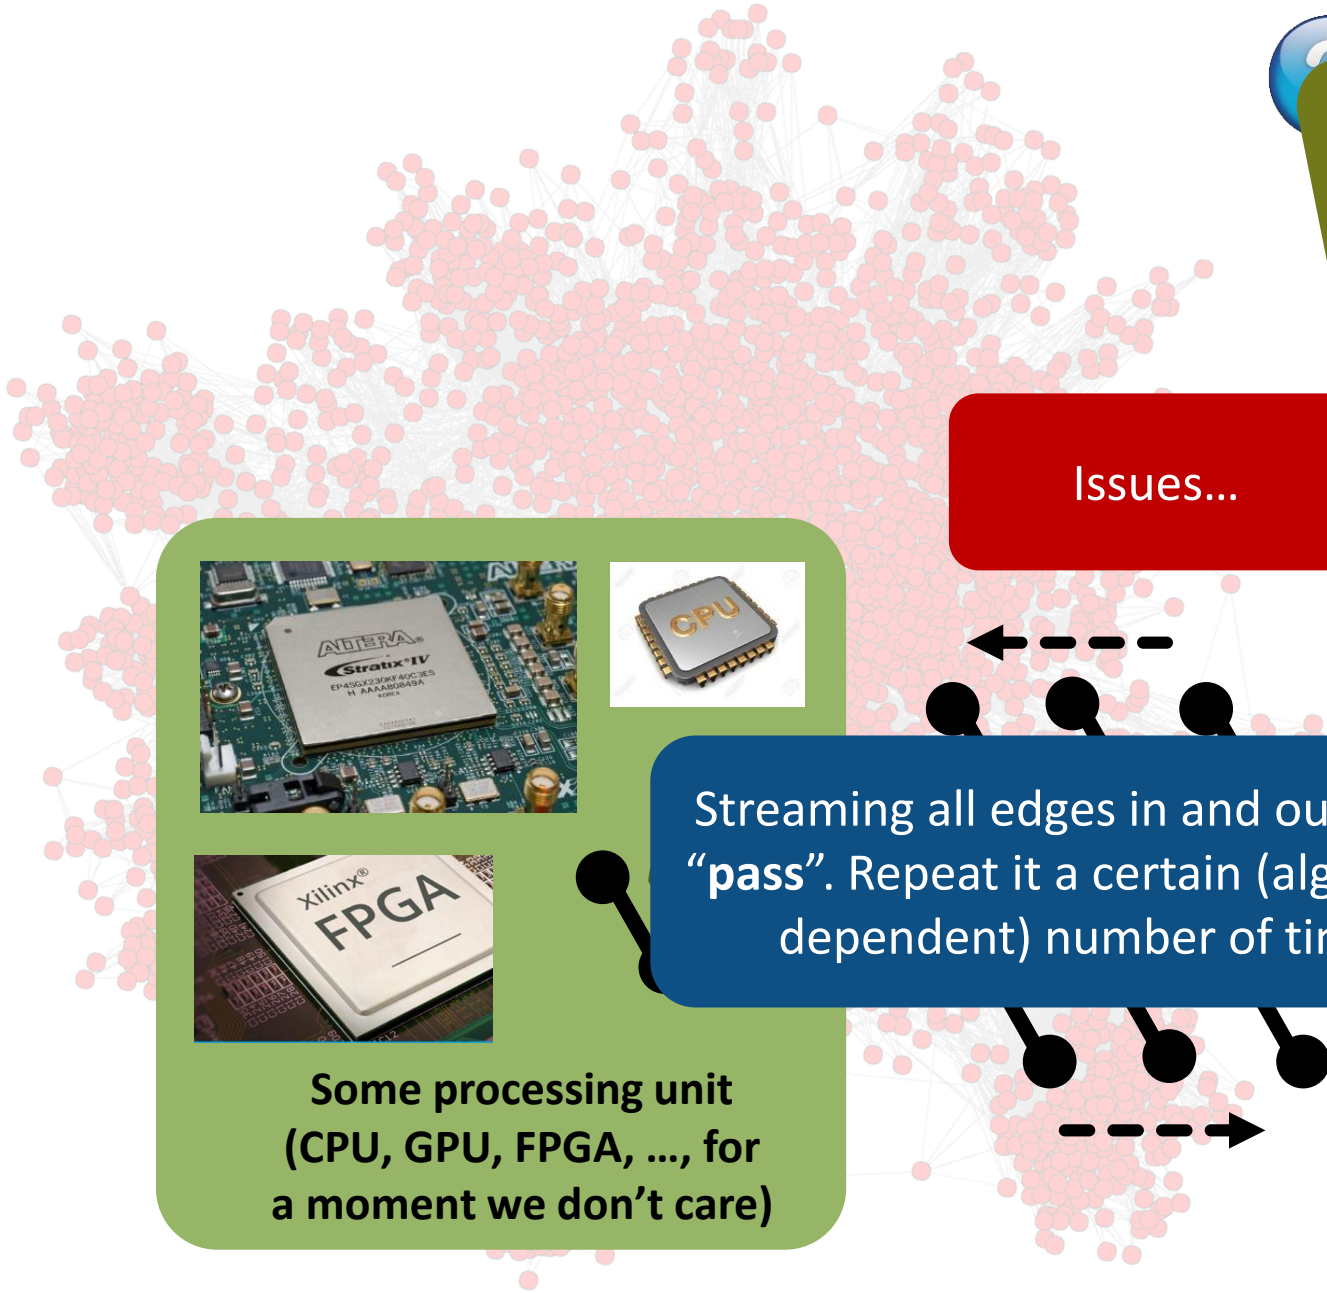
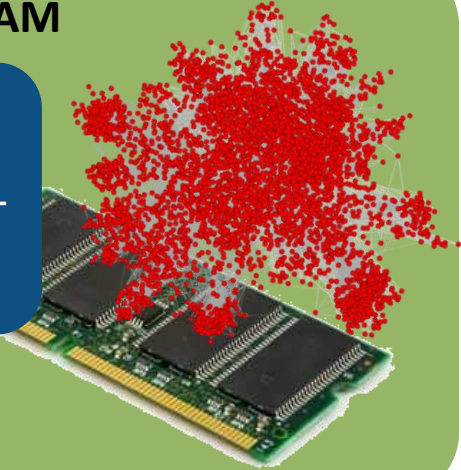
Issues...



Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

Streaming all edges in and out is one “pass”. Repeat it a certain (algorithm-dependent) number of times

DRAM



...How to minimize the number of “passes” over edges? This can get really bad in the “traditional” edge-centric approach (e.g., BFS needs  $D$  passes;  $D$  = diameter [1]).

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming  $\approx$  pipelining”)

...Processing edges is sequential – how to incorporate parallelism?

Issues...


Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

Streaming all edges in and out is one “pass”. Repeat it a certain (algorithm-dependent) number of times


DRAM

(aka edge-centric)  
use pipelining efficiently  
("streaming  $\approx$  pipelining")

...Processing edges  
is sequential – how  
to incorporate  
parallelism?



DRAM





Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming  $\approx$  pipelining”)

...Pro  
is seq  
to i  
pa



# Part 2: Substream-Centric: A new paradigm for processing graphs





# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – how to incorporate parallelism?




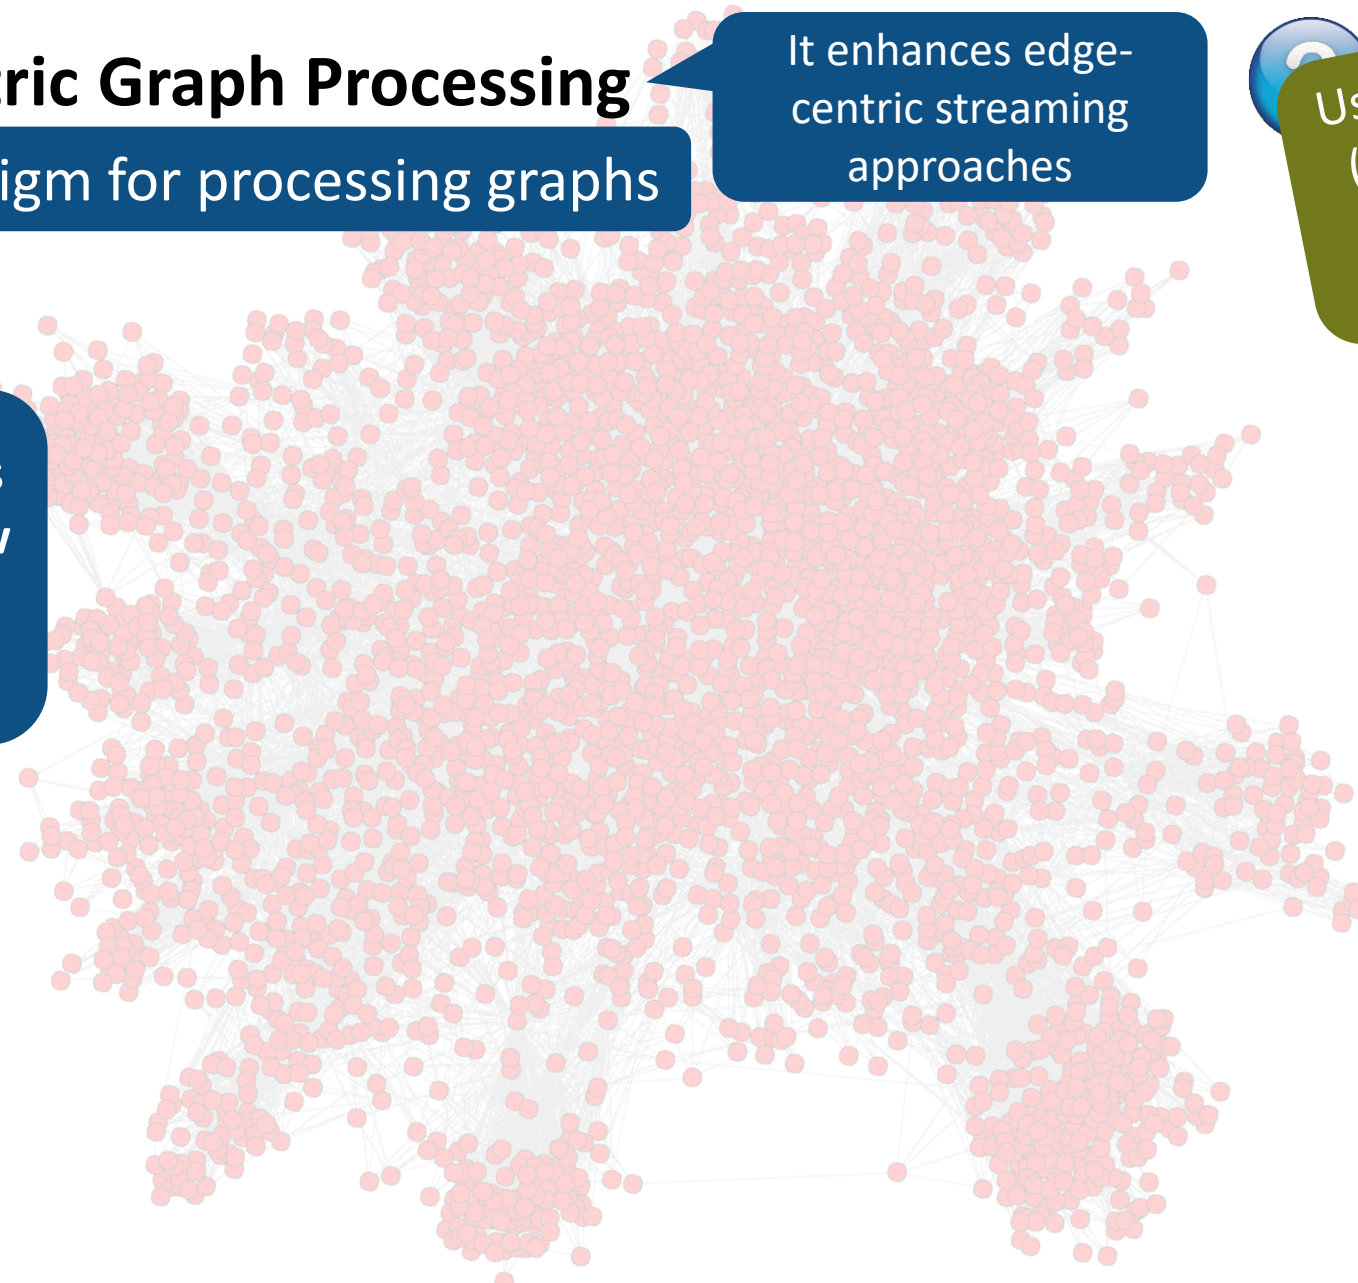
# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – how to incorporate parallelism?





# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – how to incorporate parallelism?





# Substream-Centric Graph Processing


## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches


Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Divide the input stream of edges according to some (algorithm-specific) pattern



Weighted edges



DRAM

# Substream-Centric Graph Processing



## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Divide the input stream of edges according to some (algorithm-specific) pattern



# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs



It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming ≈ pipelining")


...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Weighted edges



# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs


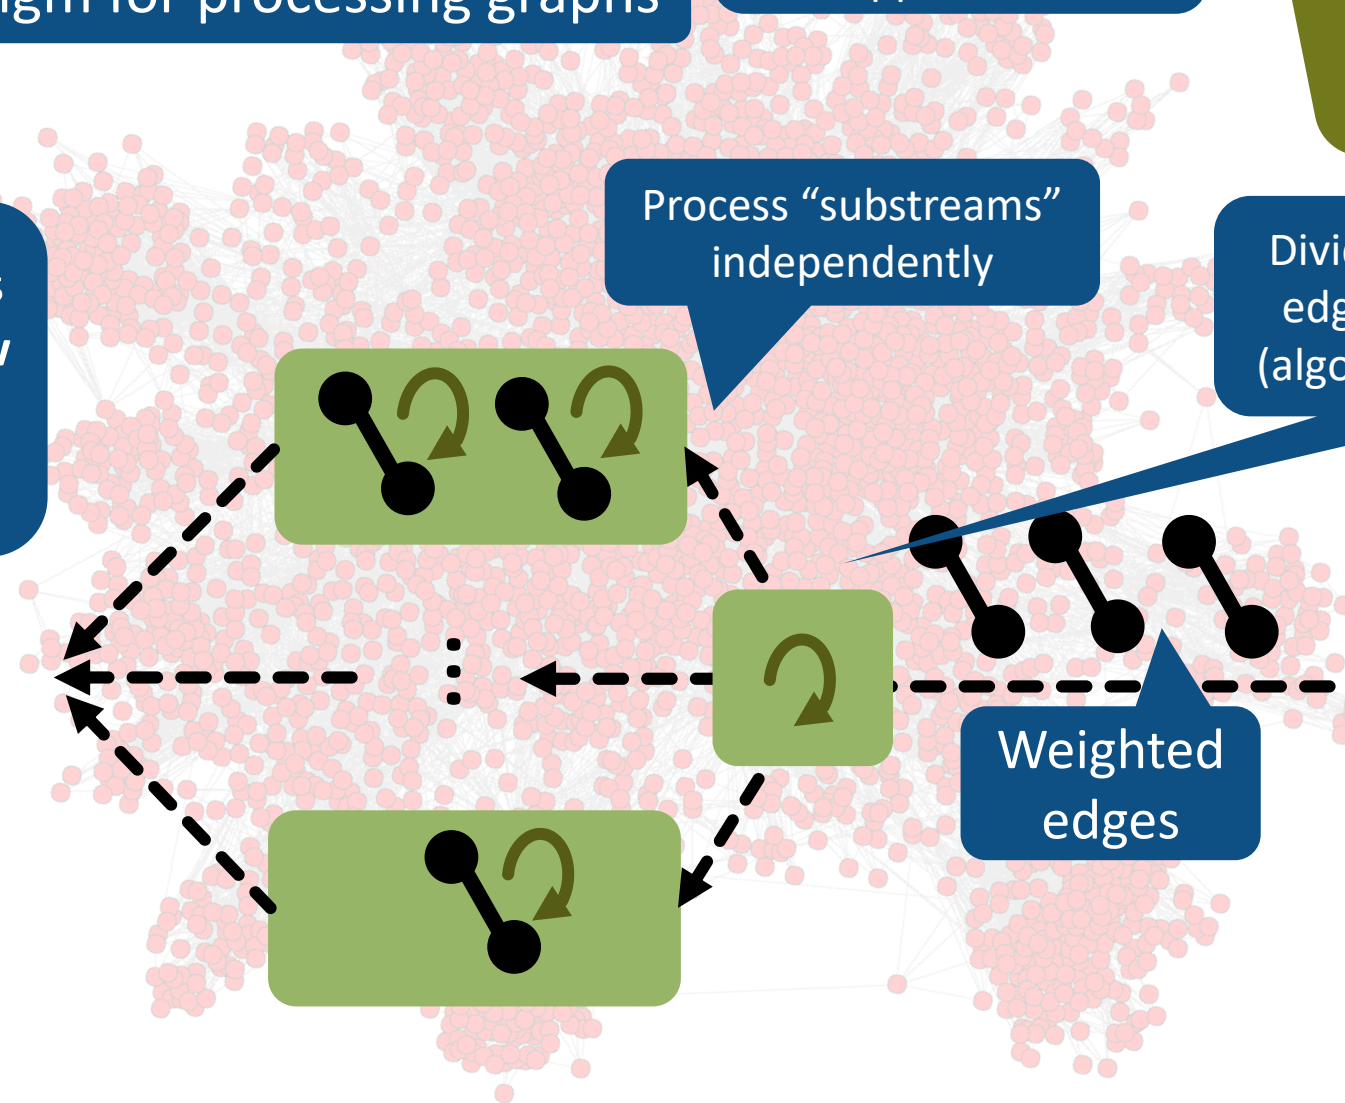
It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Weighted edges

# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming ≈ pipelining")


...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

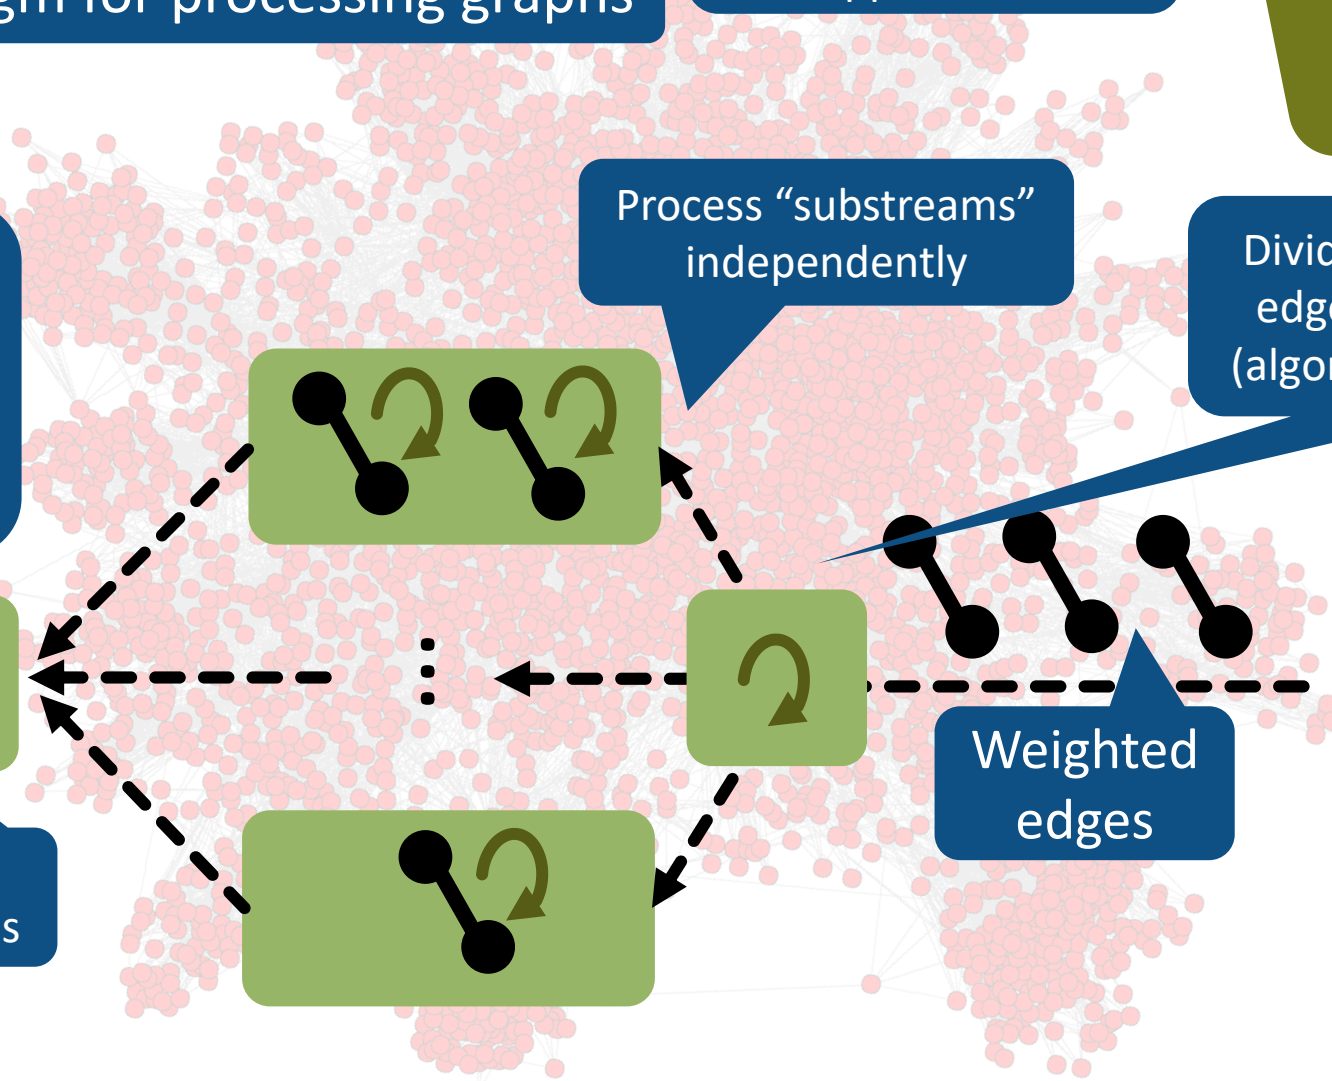

Divide the input stream of edges according to some (algorithm-specific) pattern



Merge substreams



Weighted edges



# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs


It enhances edge-centric streaming approaches

Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

...Processing edges is sequential – how to incorporate parallelism?


Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Merge substreams

Weighted edges



# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs


It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)


Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Merge substreams

Weighted edges



# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches

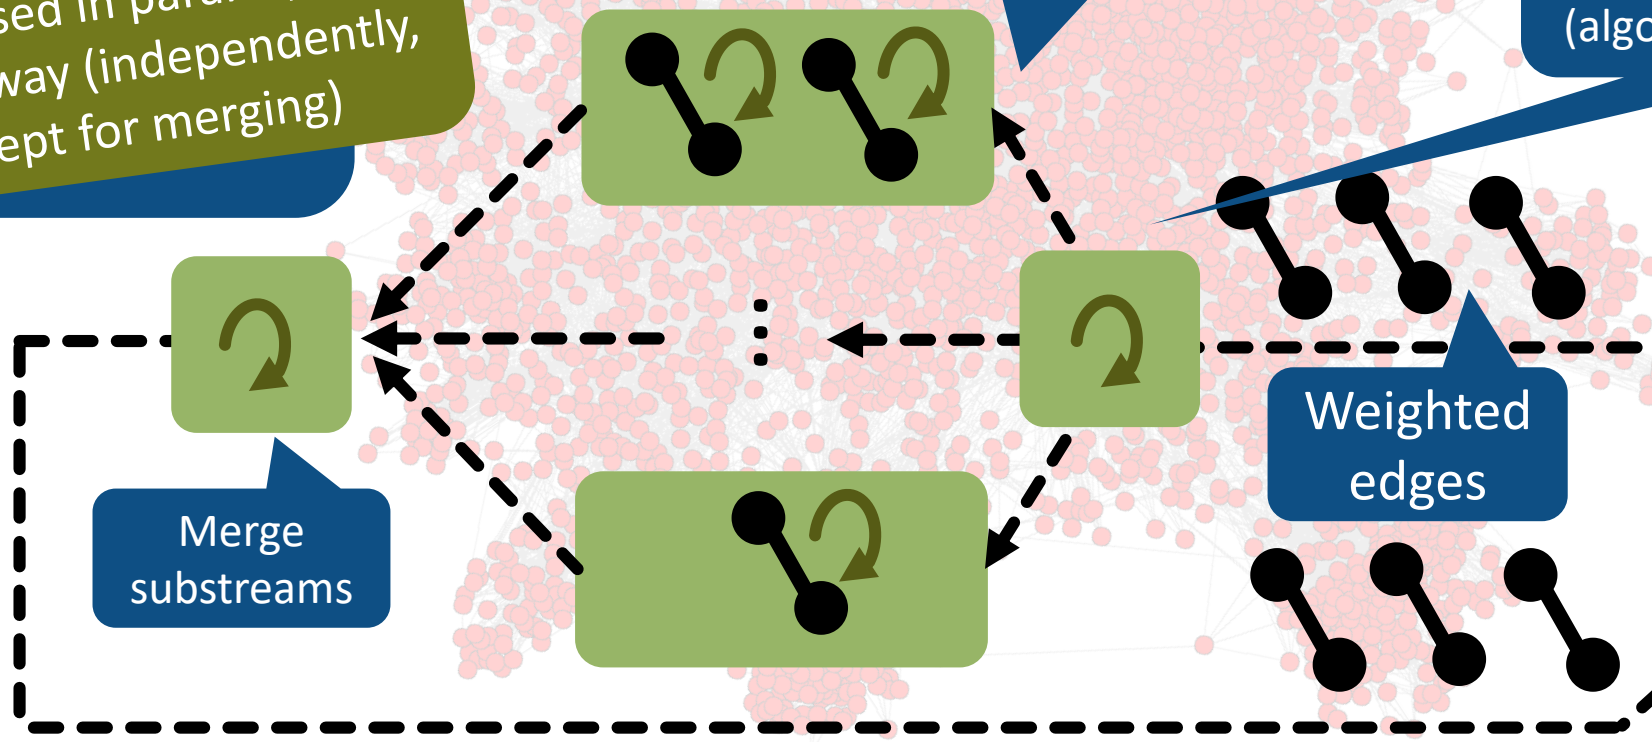
Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

Also, it enables (tunable) approximation and a (tunable) number of passes

Streams independently


Divide the input stream of edges according to some (algorithm-specific) pattern

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)



Merge substreams

Weighted edges





# Substream-Centric Graph Processing

## Part 2: A new paradigm for processing graphs

It enhances edge-centric streaming approaches


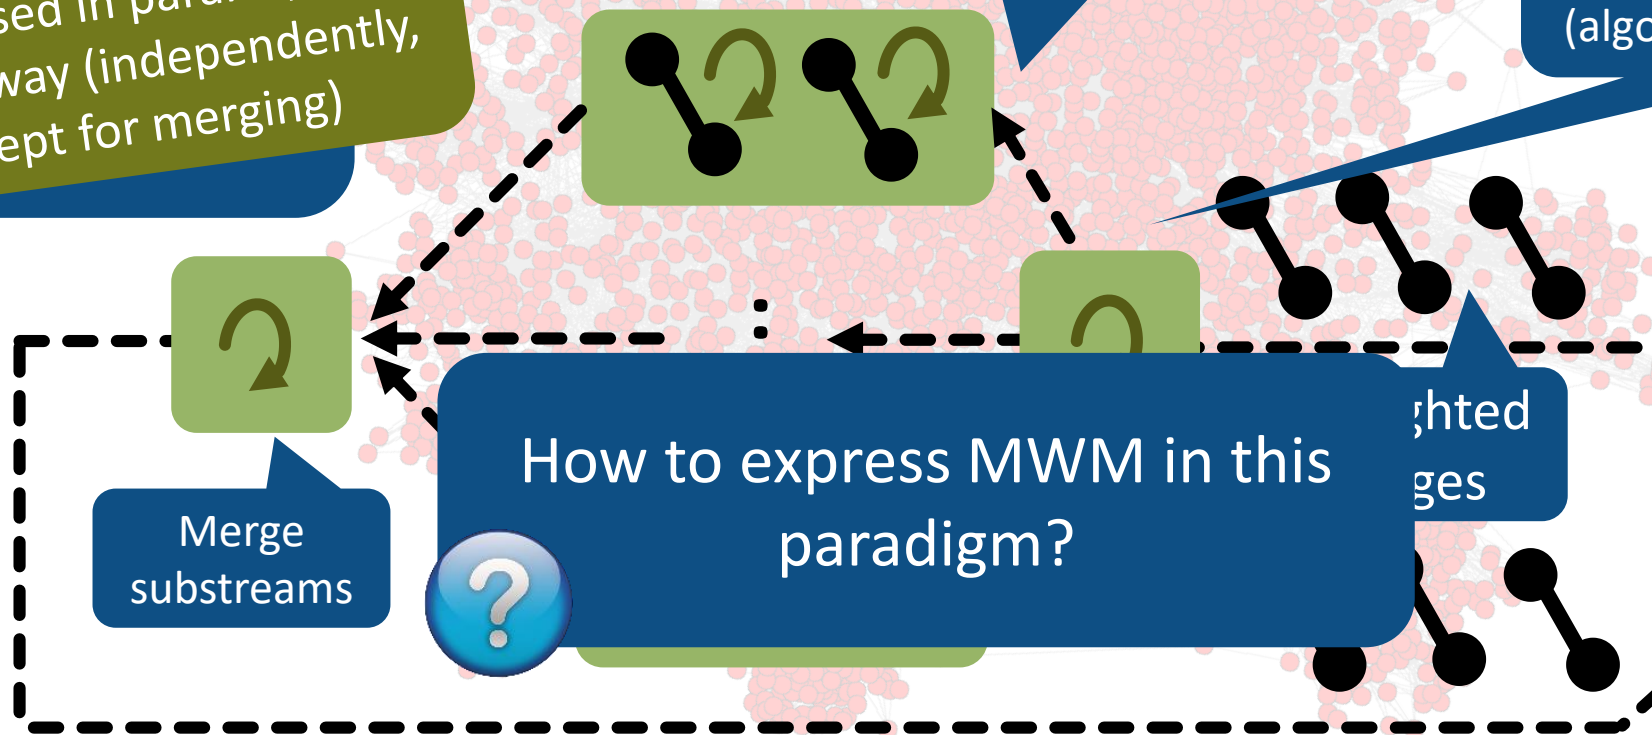
Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming  $\approx$  pipelining")

Also, it enables (tunable) approximation and a (tunable) number of passes

...streams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)



# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Which programming paradigm to use for (approximate) MWM (and many other problems)?



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Use substream-centric processing (exposes parallelism, enables easy pipelining, supports approximation)



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

## Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?




Use substream-centric processing (exposes parallelism, enables easy pipelining, supports approximation)



What is the HW FPGA design that ensures high performance?




What is the ultimate performance, power consumption, and the related tradeoffs?



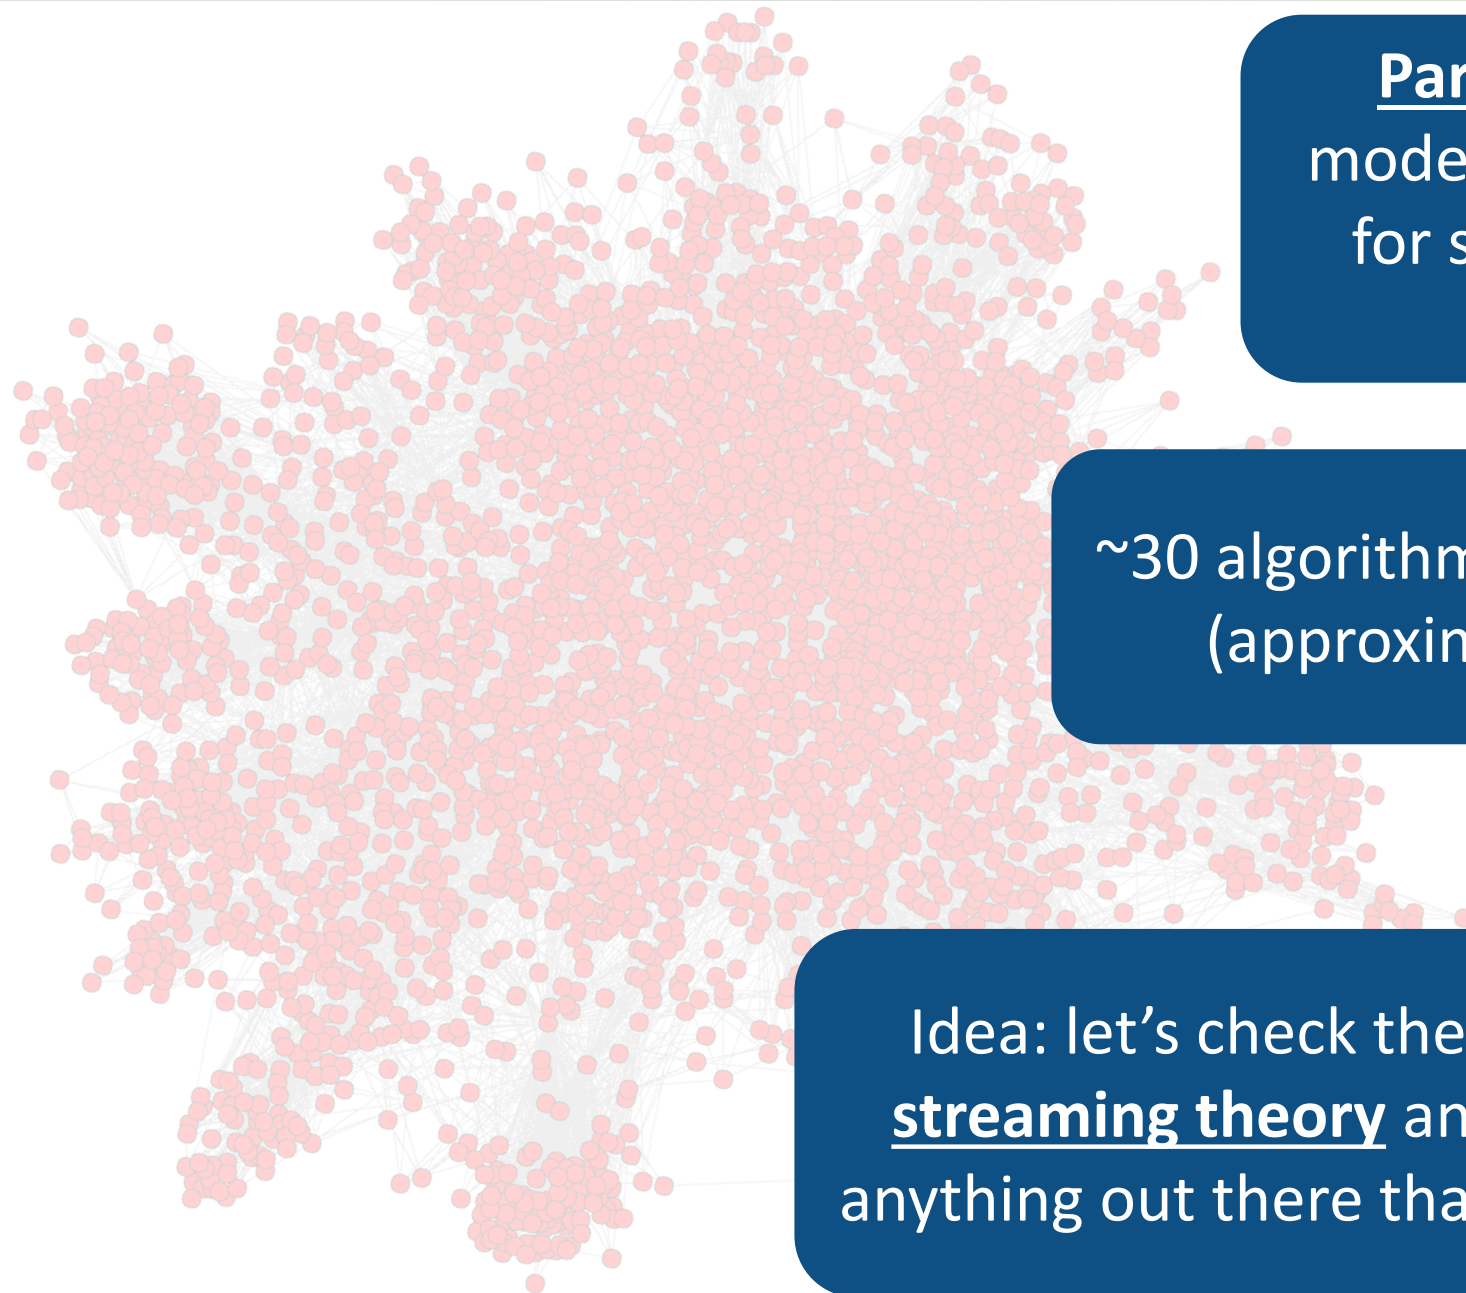


## Part 3: Analysis of models and algorithms for streaming graph processing



**Part 3:** Analysis of models and algorithms for streaming graph processing

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

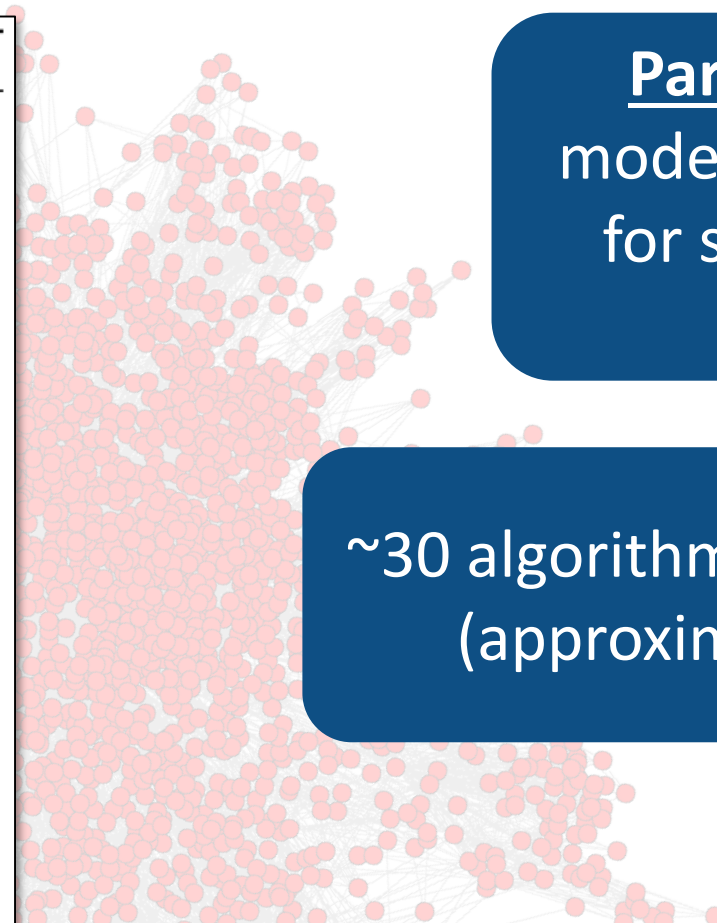


**Part 3:** Analysis of models and algorithms for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊






**Part 3: Analysis of models and algorithms for streaming graph processing**

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[26]	1/2	$O(n)$	1	👎	👍	👍
[41, Theorem 6]	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👎	👍	👍
[41, Theorem 2]	$1/2 + 0.003^*$	$O(n \text{ polylog}(n))$	1	👎	👍	👍
[36, Theorem 1.1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👎	👍	👍
[26, Theorem 1]	$2/3 - \epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👎	👎	👍
[6, Theorem 19]	$1 - \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👎	👎	👍
[41, Theorem 5]	$1/2 + 0.019$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👎	👎	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👎	👎	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👎	👎	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👎	👎	👍
[14]	1	$\tilde{O}(k^2)$	1	👎	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👎	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👎	👎	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	?
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	?
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	?
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	?
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	?
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	?
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	?
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	?
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	?
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👎	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍



**Part 3: Analysis of models and algorithms for streaming graph processing**

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[26]	1/2	$O(n)$	1	👍	👍	👍
[41, Theorem 6]	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[41, Theorem 2]	$1/2 + 0.003^*$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[36, Theorem 1.1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👍	👍	👍
[26, Theorem 1]	$2/3 - \epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👍	👍	👍
[6, Theorem 19]	$1 - \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👍	👍	👍
[41, Theorem 5]	$1/2 + 0.019$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👍	👍	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👍	👍	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👍	👍	👍
[14]	1	$\tilde{O}(k^2)$	1	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{2}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{3}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

No worries, no need to analyze it here, all the details are in the paper 😊

Most important goals:


Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👍	👍	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👍	👍	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👍	👍	👍
[14]	1	$\tilde{O}(k^2)$	1	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2/\epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{2}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{3}$	$O(n \log n)$	1	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

No worries, no need to analyze it here, all the details are in the paper 😊

**Part 3: Analysis of models and algorithms for streaming graph processing**

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[41, Theorem 1]	$1/\epsilon$	$O(n)$	1	👍	👍	👍
[41, Theorem 2]	$1/\epsilon$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[39]	$1/\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 1]	$1/\epsilon$	$O(\text{polylog}(n))$	1	👍	👍	👍
[28, Theorem 2]	$1/\epsilon$	$O(n \log n)$	$O(\log(1/\epsilon)/\epsilon)$	👍	👍	👍
[35, Theorem 2]	$1 - \frac{1}{(k-1)!}$	$O(n \text{ polylog}(n)/\epsilon^2)$	$O(\log \log(1/\epsilon)/\epsilon^2)$	👍	👍	👍
[14]	1	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(k^2)$	1	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2/\epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Most important goals:

Maximize accuracy

No worries, no need to analyze it here, all the details are in the paper 😊

**Part 3: Analysis of models and algorithms for streaming graph processing**

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[14]	1	$O(n)$	1	👍	👍	👍
[14]	1	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[14]	$1/\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$O(\text{polylog}(n))$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Most important goals:

Maximize accuracy


Minimize local space

No worries, no need to analyze it here, all the details are in the paper 😊

**Part 3:** Analysis of models and algorithms for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[41, Theorem 1]	$1/\epsilon$	$O(n)$	1	👍	👍	👍
[41, Theorem 2]	$1/\epsilon$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[39]	$1/\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 1]	$1/\epsilon$	$O(\text{polylog}(n))$	1	👍	👍	👍
[28, Theorem 2]	$1/\epsilon$	$O(n \log n)$	1	👍	👍	👍
[35, Theorem 2]	$1 - \frac{1}{(k-1)!}$	$O(n \text{ polylog}(n)/\epsilon^2)$	$O(\frac{1}{\epsilon^2})$	👍	👍	👍
[14]	1	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(k^2)$	1	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2/\epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \frac{\log n}{\epsilon \log n - \log \epsilon})$	$O(\frac{1}{\epsilon^2 \log(\epsilon^{-1})})$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \frac{\log n}{\epsilon \log n - \log \epsilon})$	$O(\frac{1}{\epsilon^2 \log(\epsilon^{-1})})$	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Most important goals:

Maximize accuracy

Minimize local space

Minimize #passes

No worries, no need to analyze it here, all the details are in the paper 😊

**Part 3: Analysis of models and algorithms for streaming graph processing**

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[14]	1	$O(n)$	1	👍	👍	👍
[14]	1	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[14]	$1/\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{2}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{3}$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Most important goals:

Maximize accuracy

Minimize local space

Minimize #passes

Expose parallelism (match substream-centric)

**Part 3:** Analysis of models and algorithms for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

No worries, no need to analyze it here, all the details are in the paper 😊

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[41, Theorem 1]	$1/\epsilon$	$O(n)$	1	👍	👍	👍
[41, Theorem 2]	$1/\epsilon$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[39]	$1/\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 1]	$1/\epsilon$	$O(\text{polylog}(n))$	1	👍	👍	👍
[28, Theorem 2]	$1/\epsilon$	$O(n \log n)$	1	👍	👍	👍
[35, Theorem 2]	$1 - \frac{1}{(k-1)!}$	$O(n \text{ polylog}(n)/\epsilon^2)$	$O(\frac{1}{\epsilon^2})$	👍	👍	👍
[14]	1	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(k^2)$	1	👍	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2/\epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n \log n)$	1	👍	👍	👍
Crouch and Stubbs [1]	$\epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Most important goals:

Maximize accuracy

Minimize local space

Minimize #passes

Expose parallelism (match substream-centric)

**Part 3:** Analysis of models and algorithms for streaming graph processing

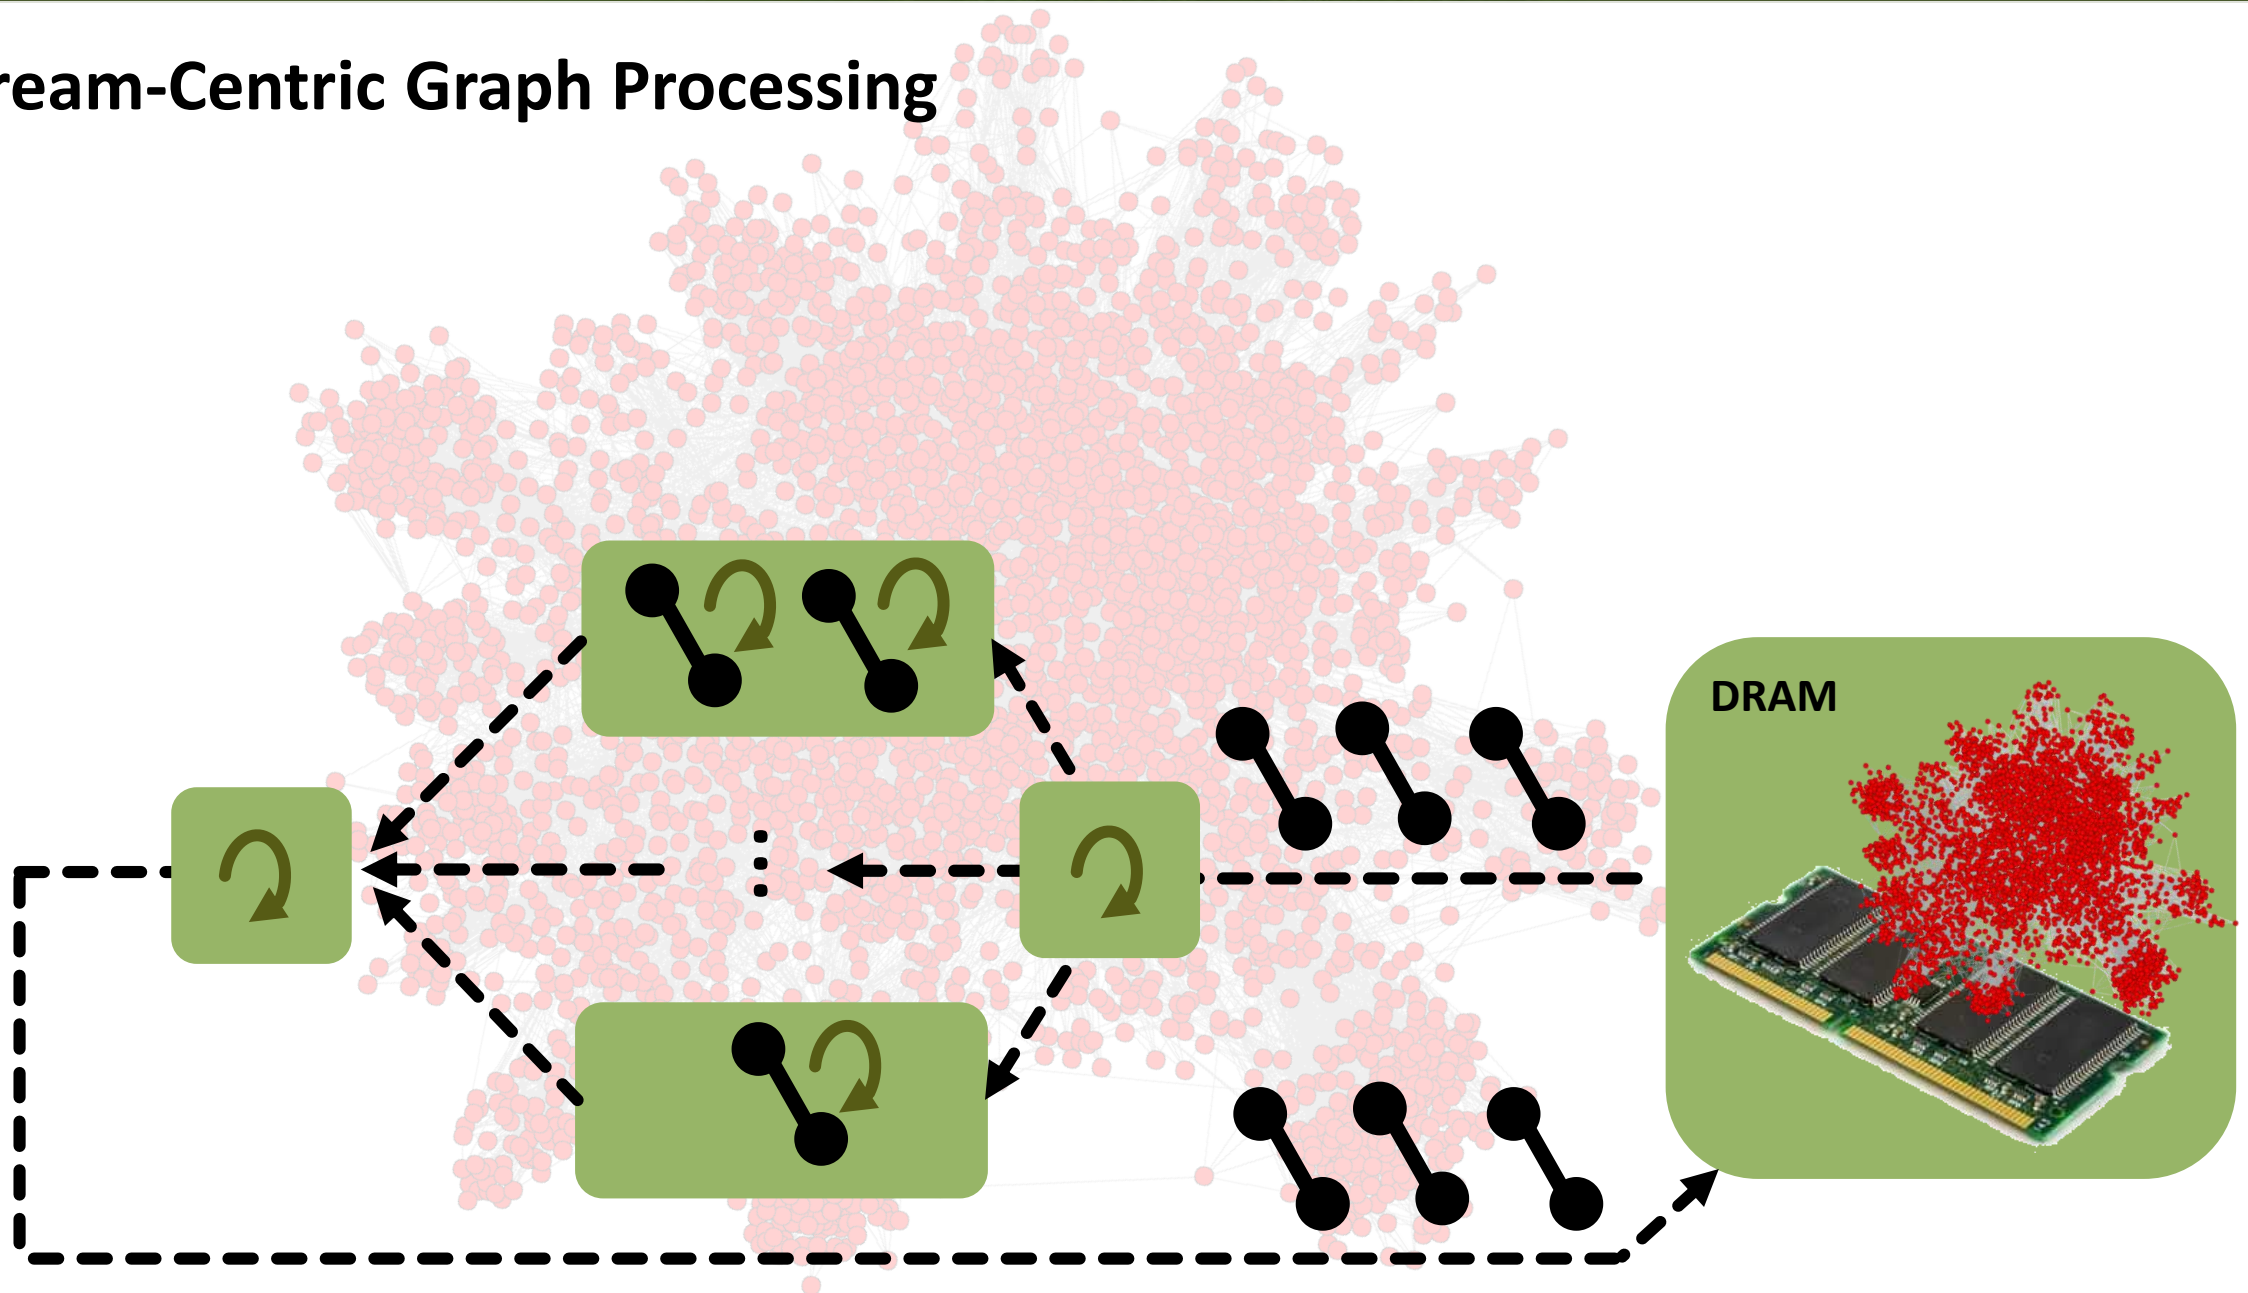
~30 algorithms for streaming (approximate) MWM

Idea: let's check the (rich) world of streaming theory and see if there is anything out there that we could use 😊

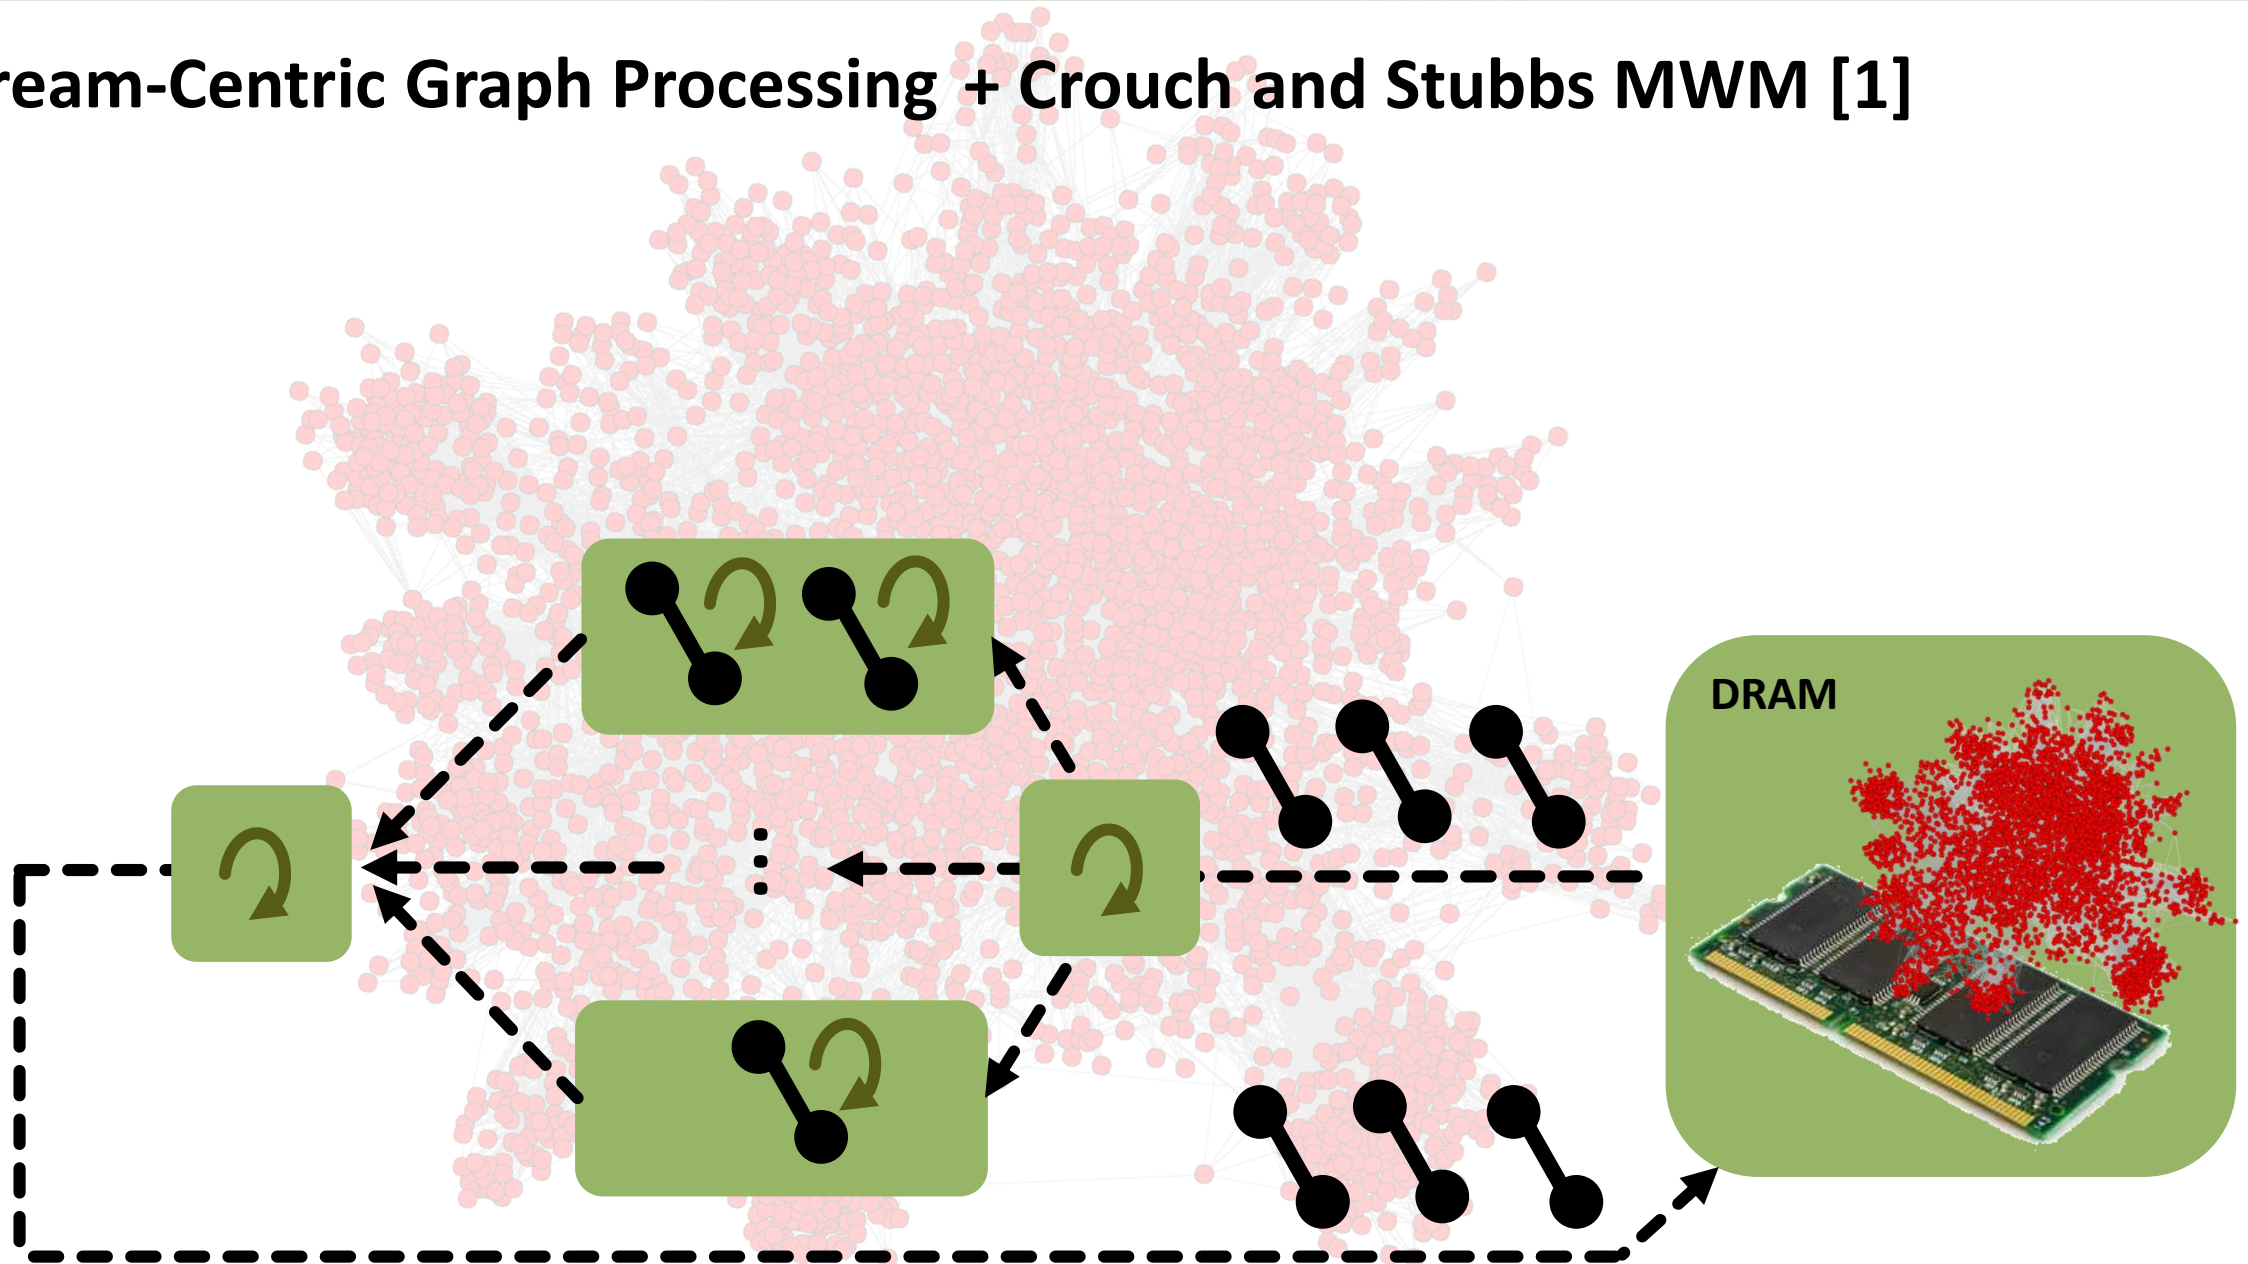
No worries, no need to analyze it here, all the details are in the paper 😊



# Substream-Centric Graph Processing

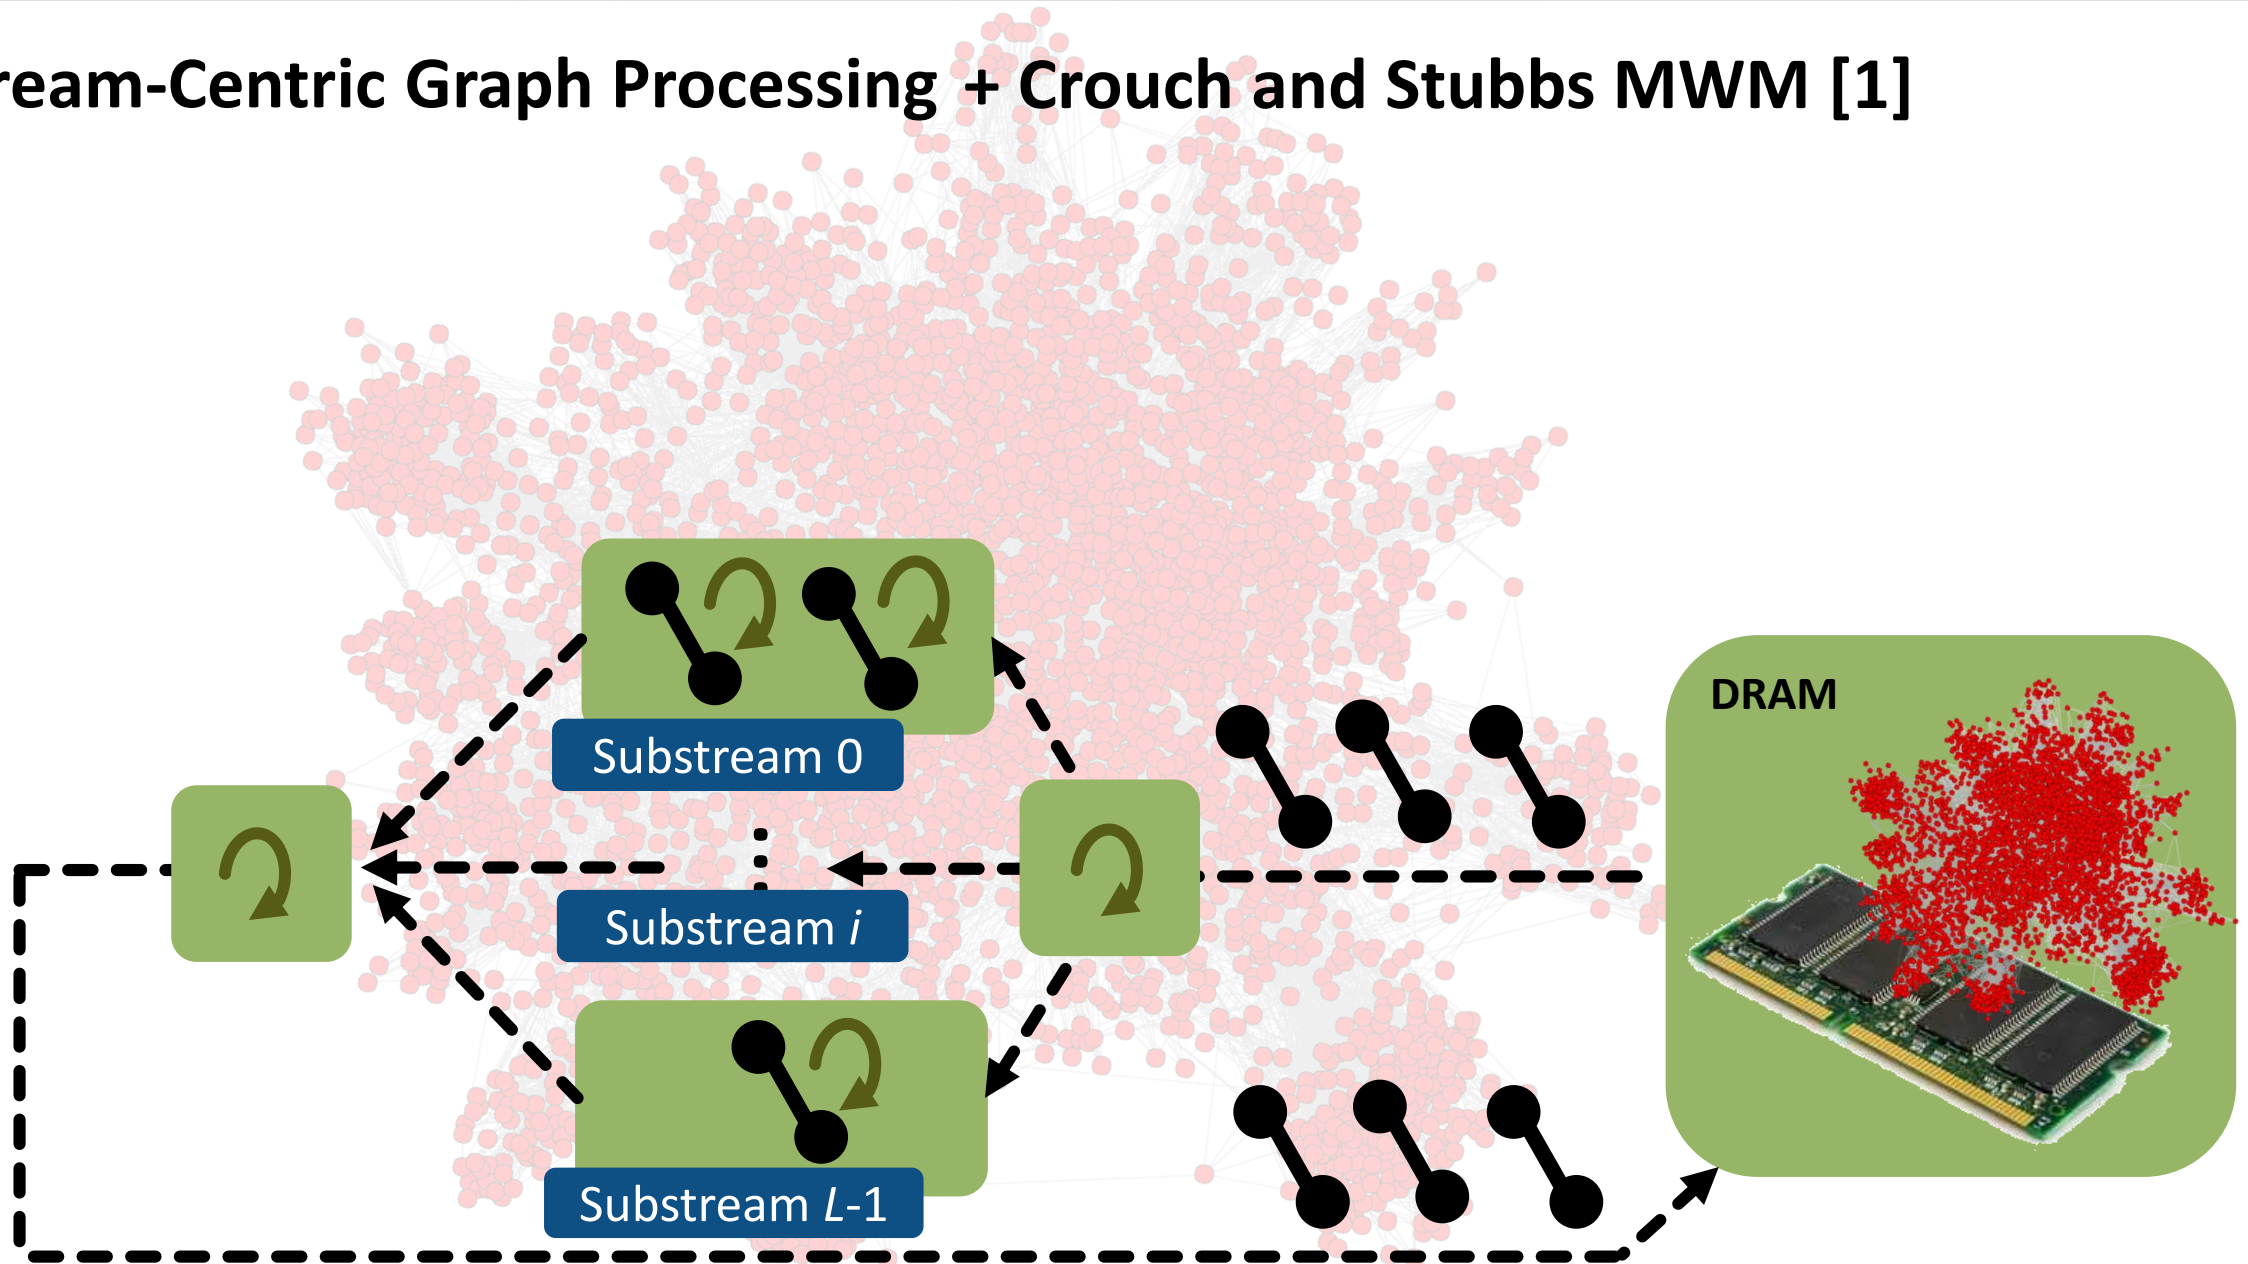


# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

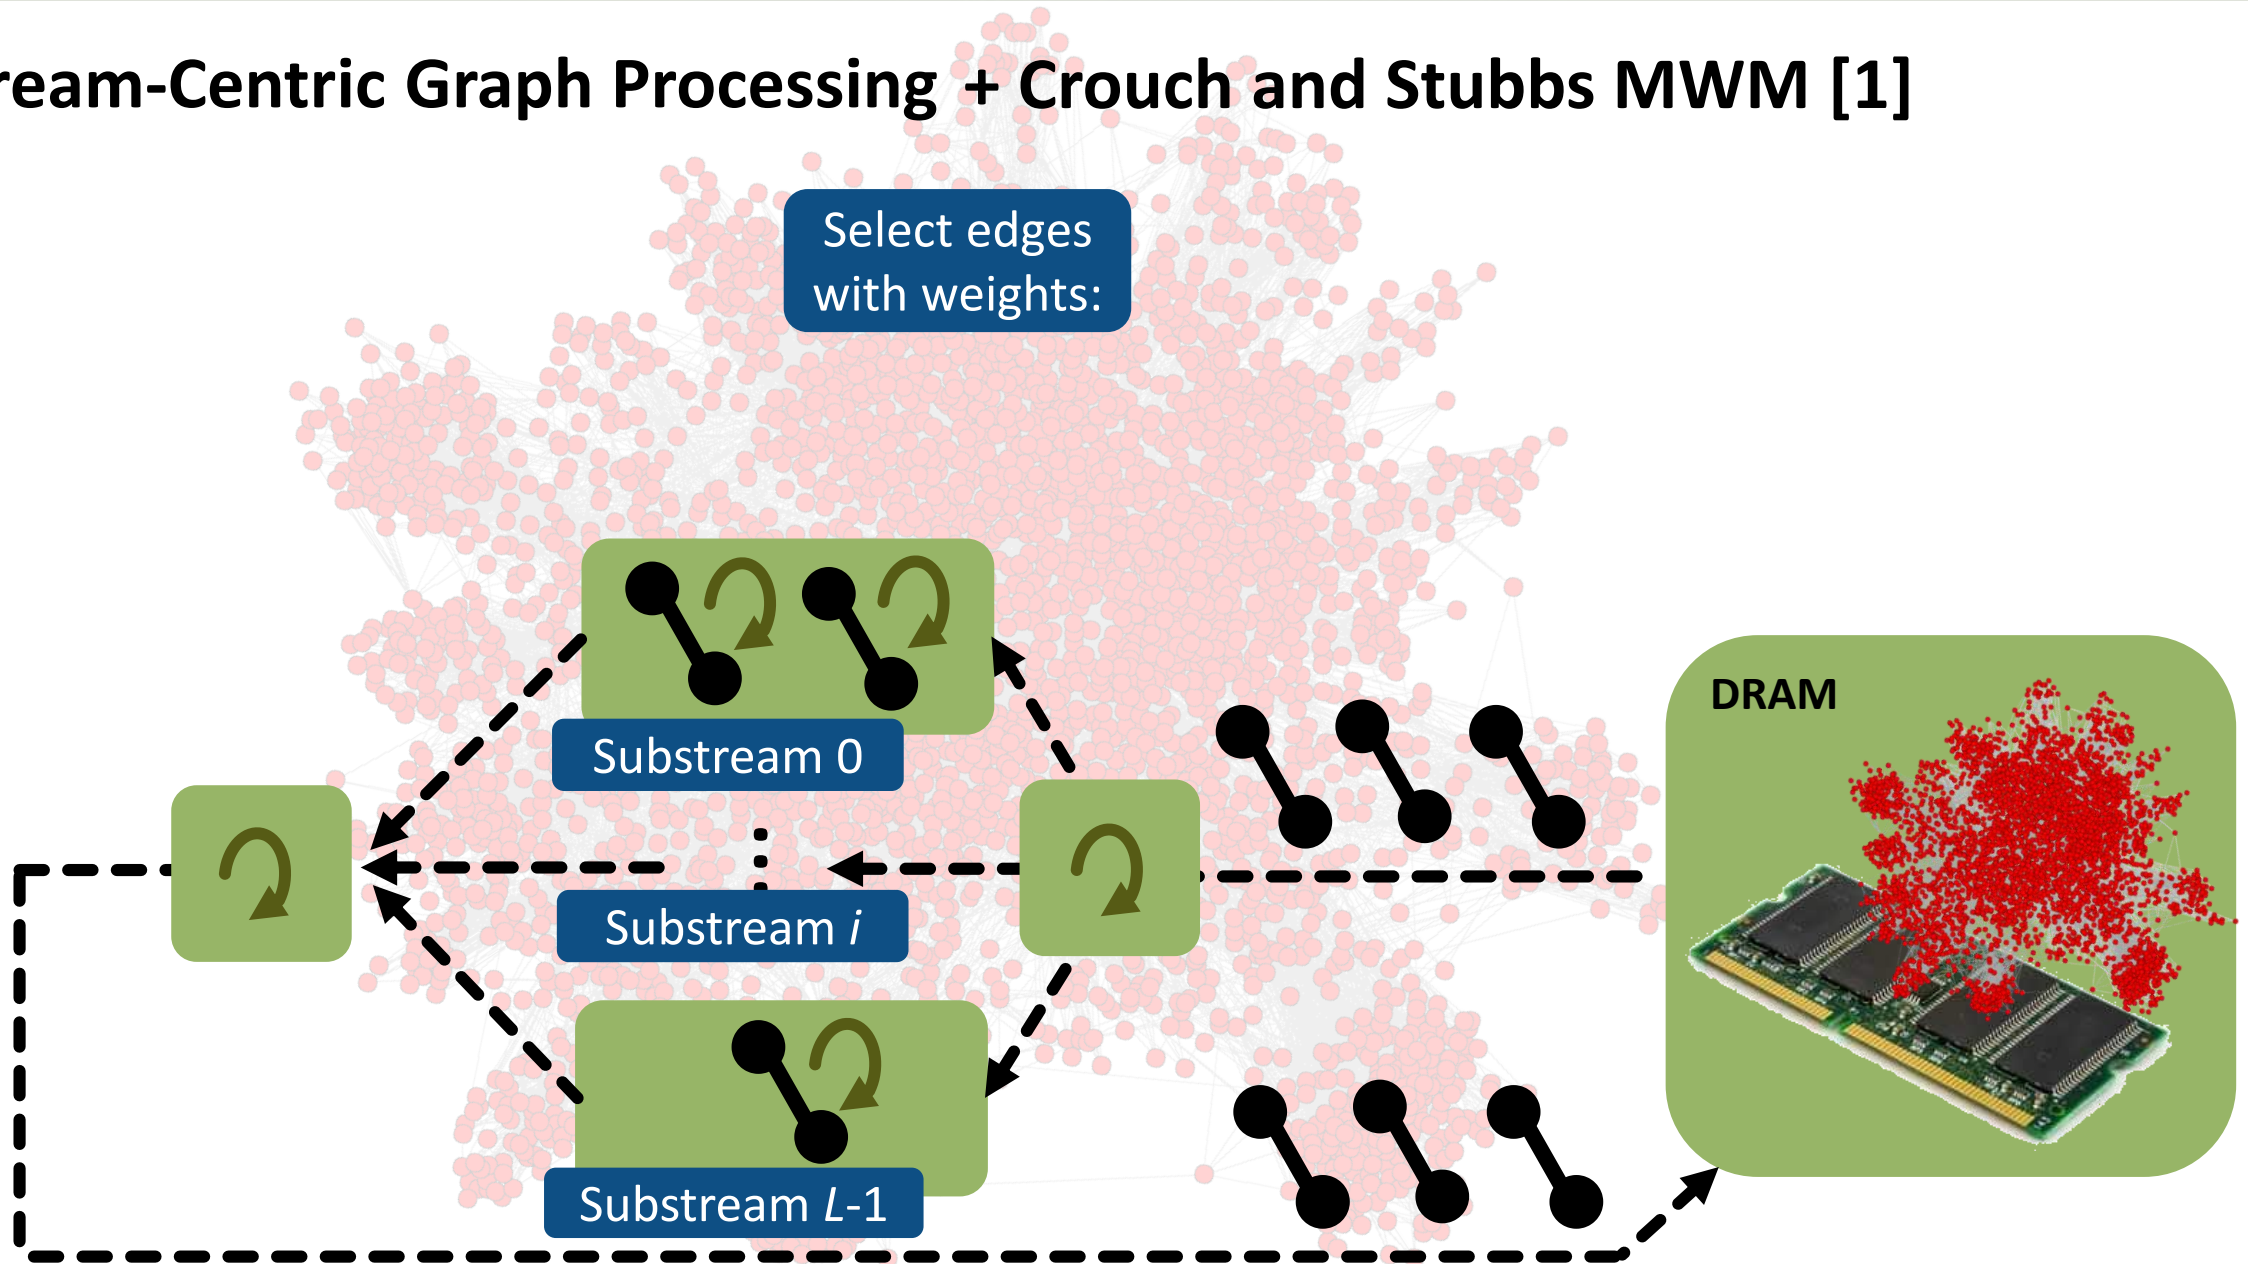


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

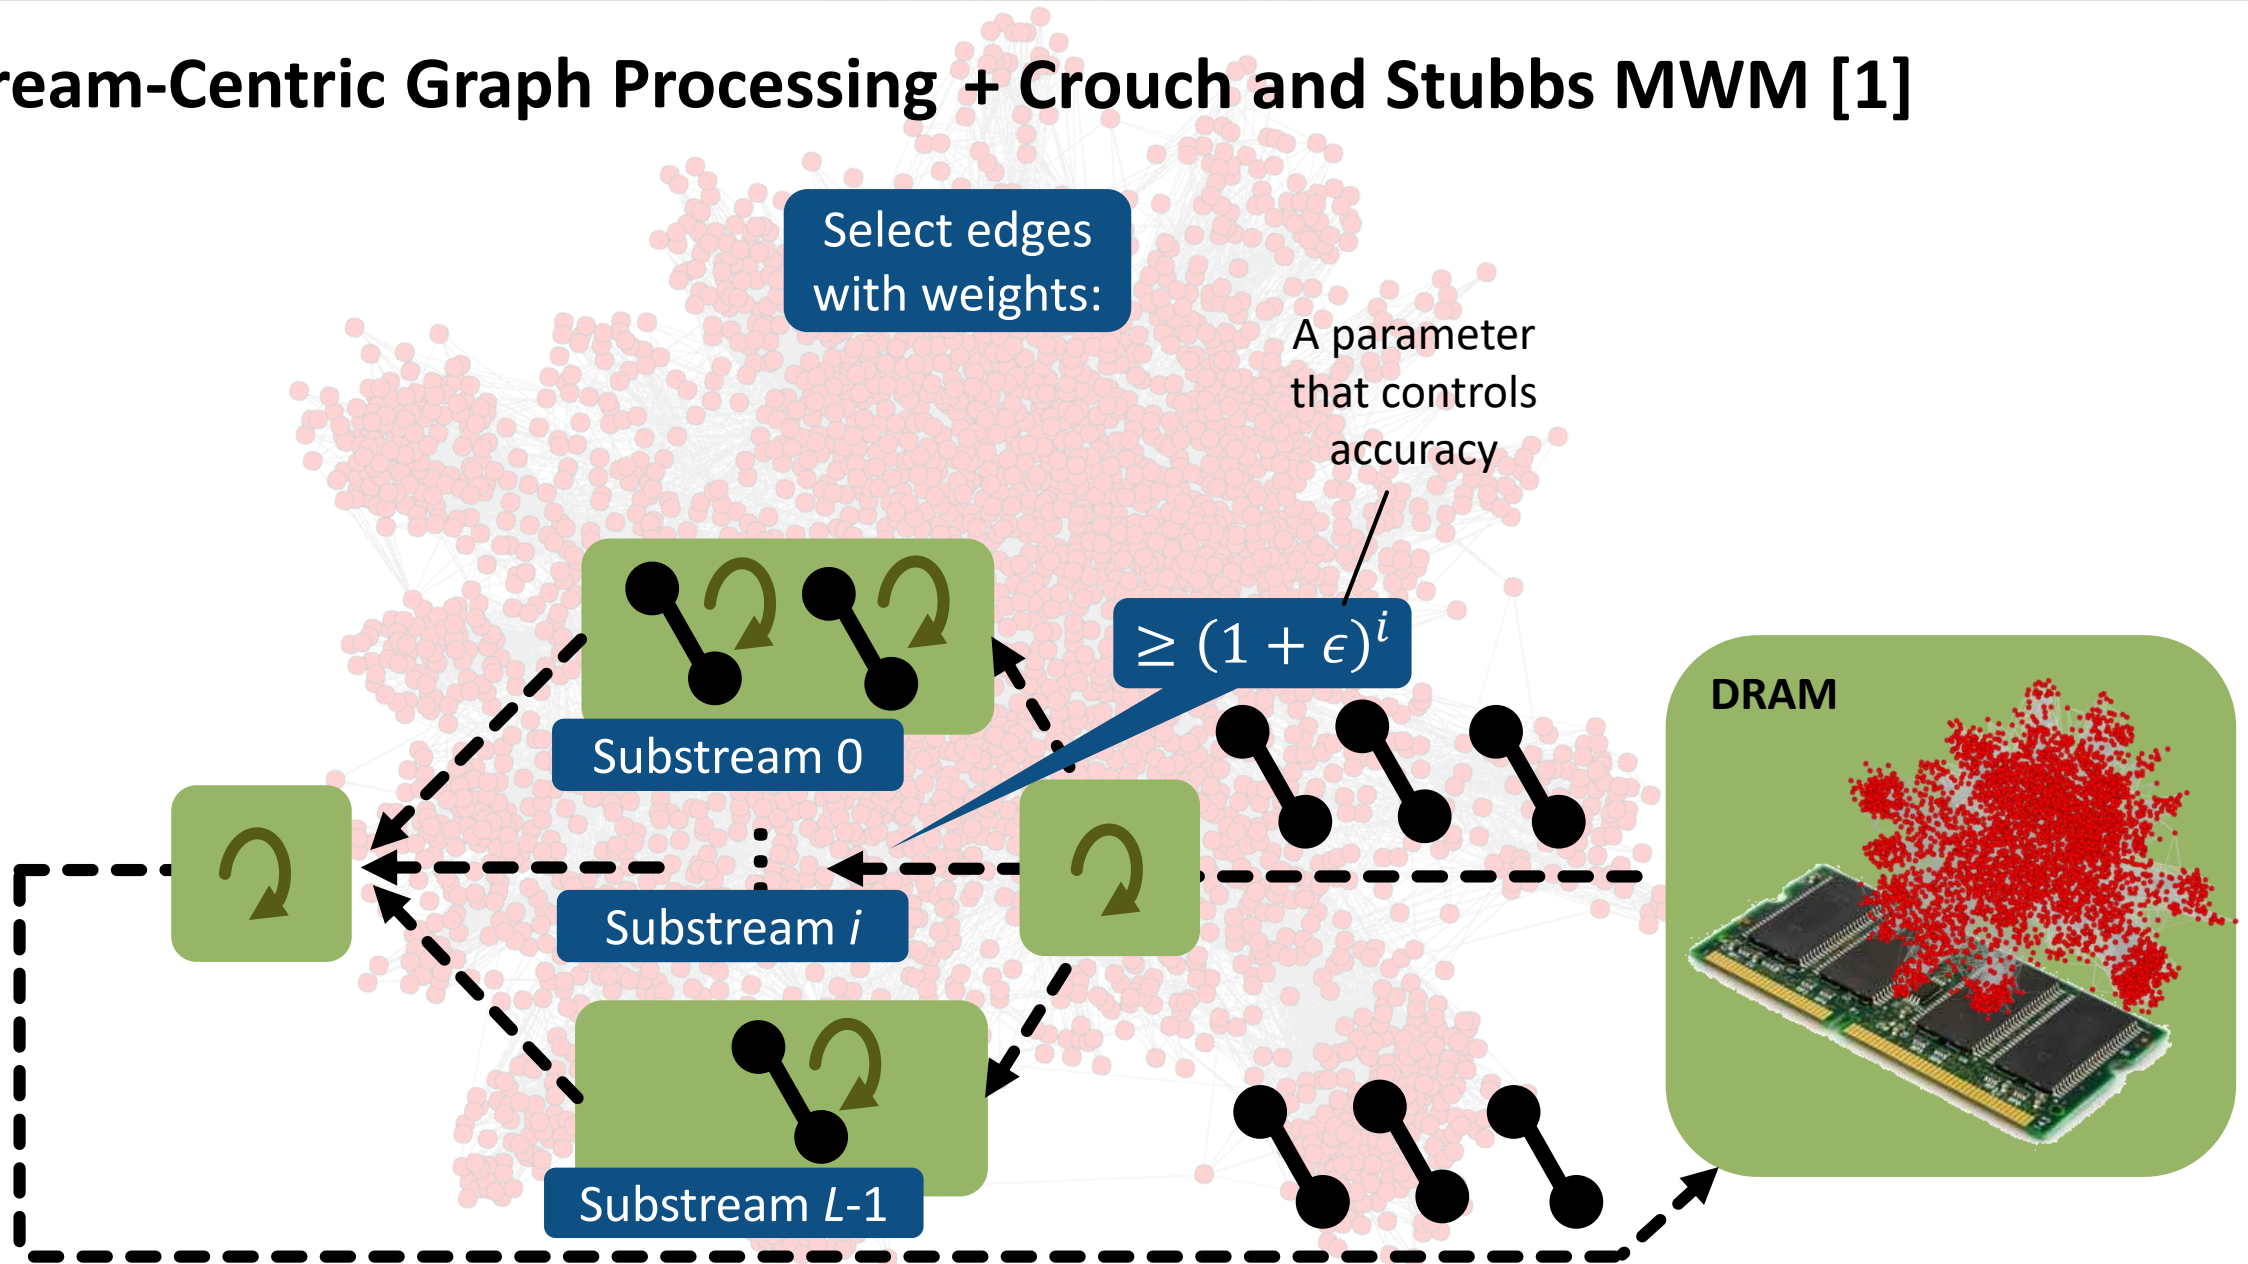


# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



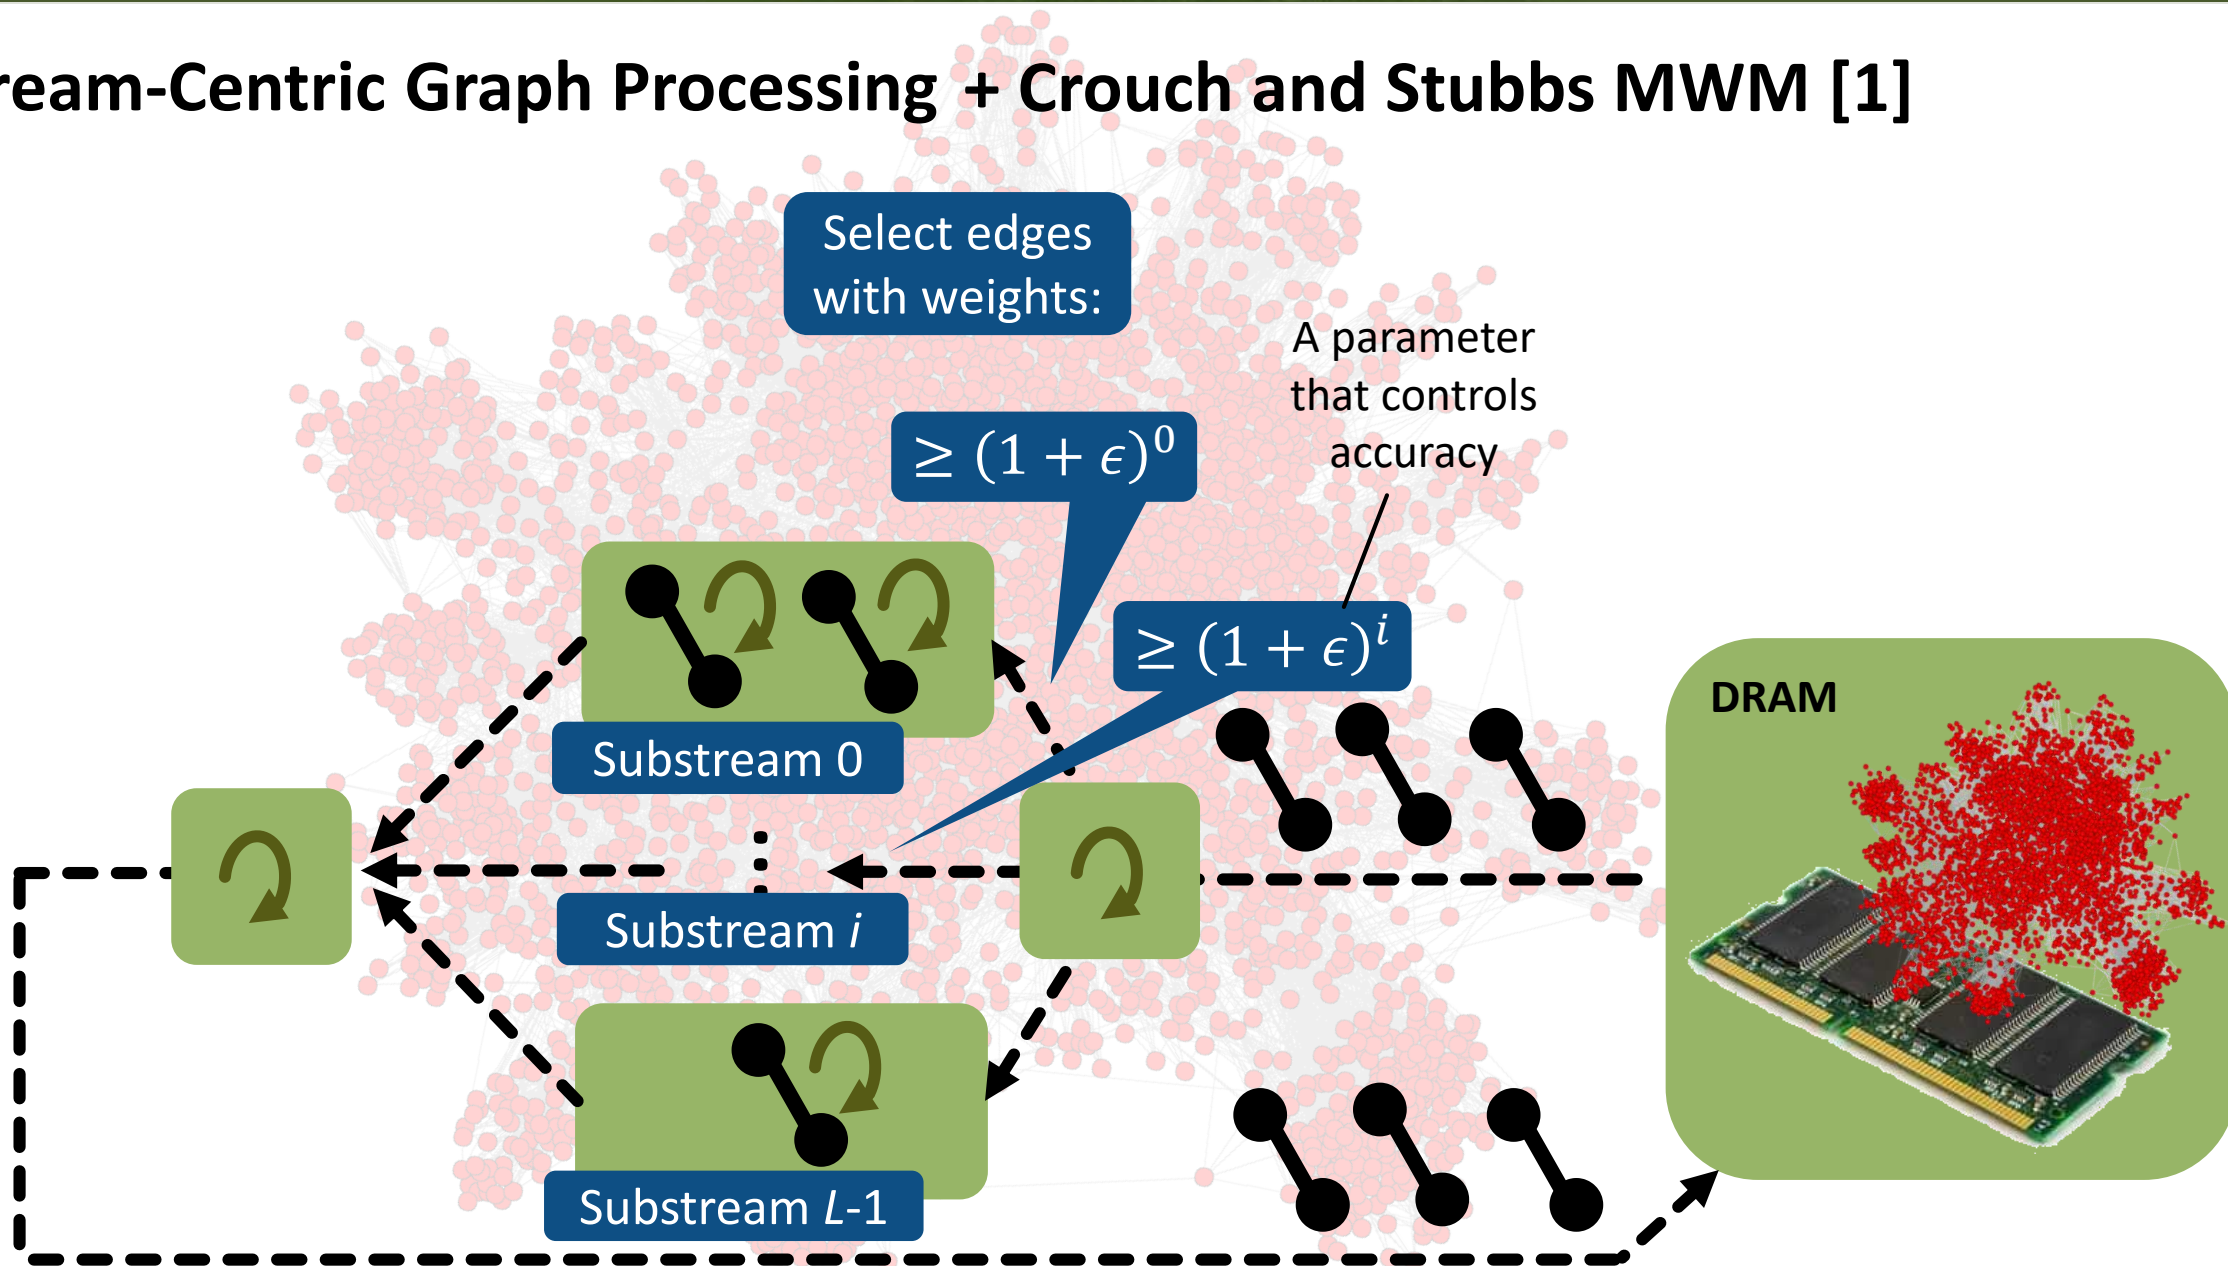
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



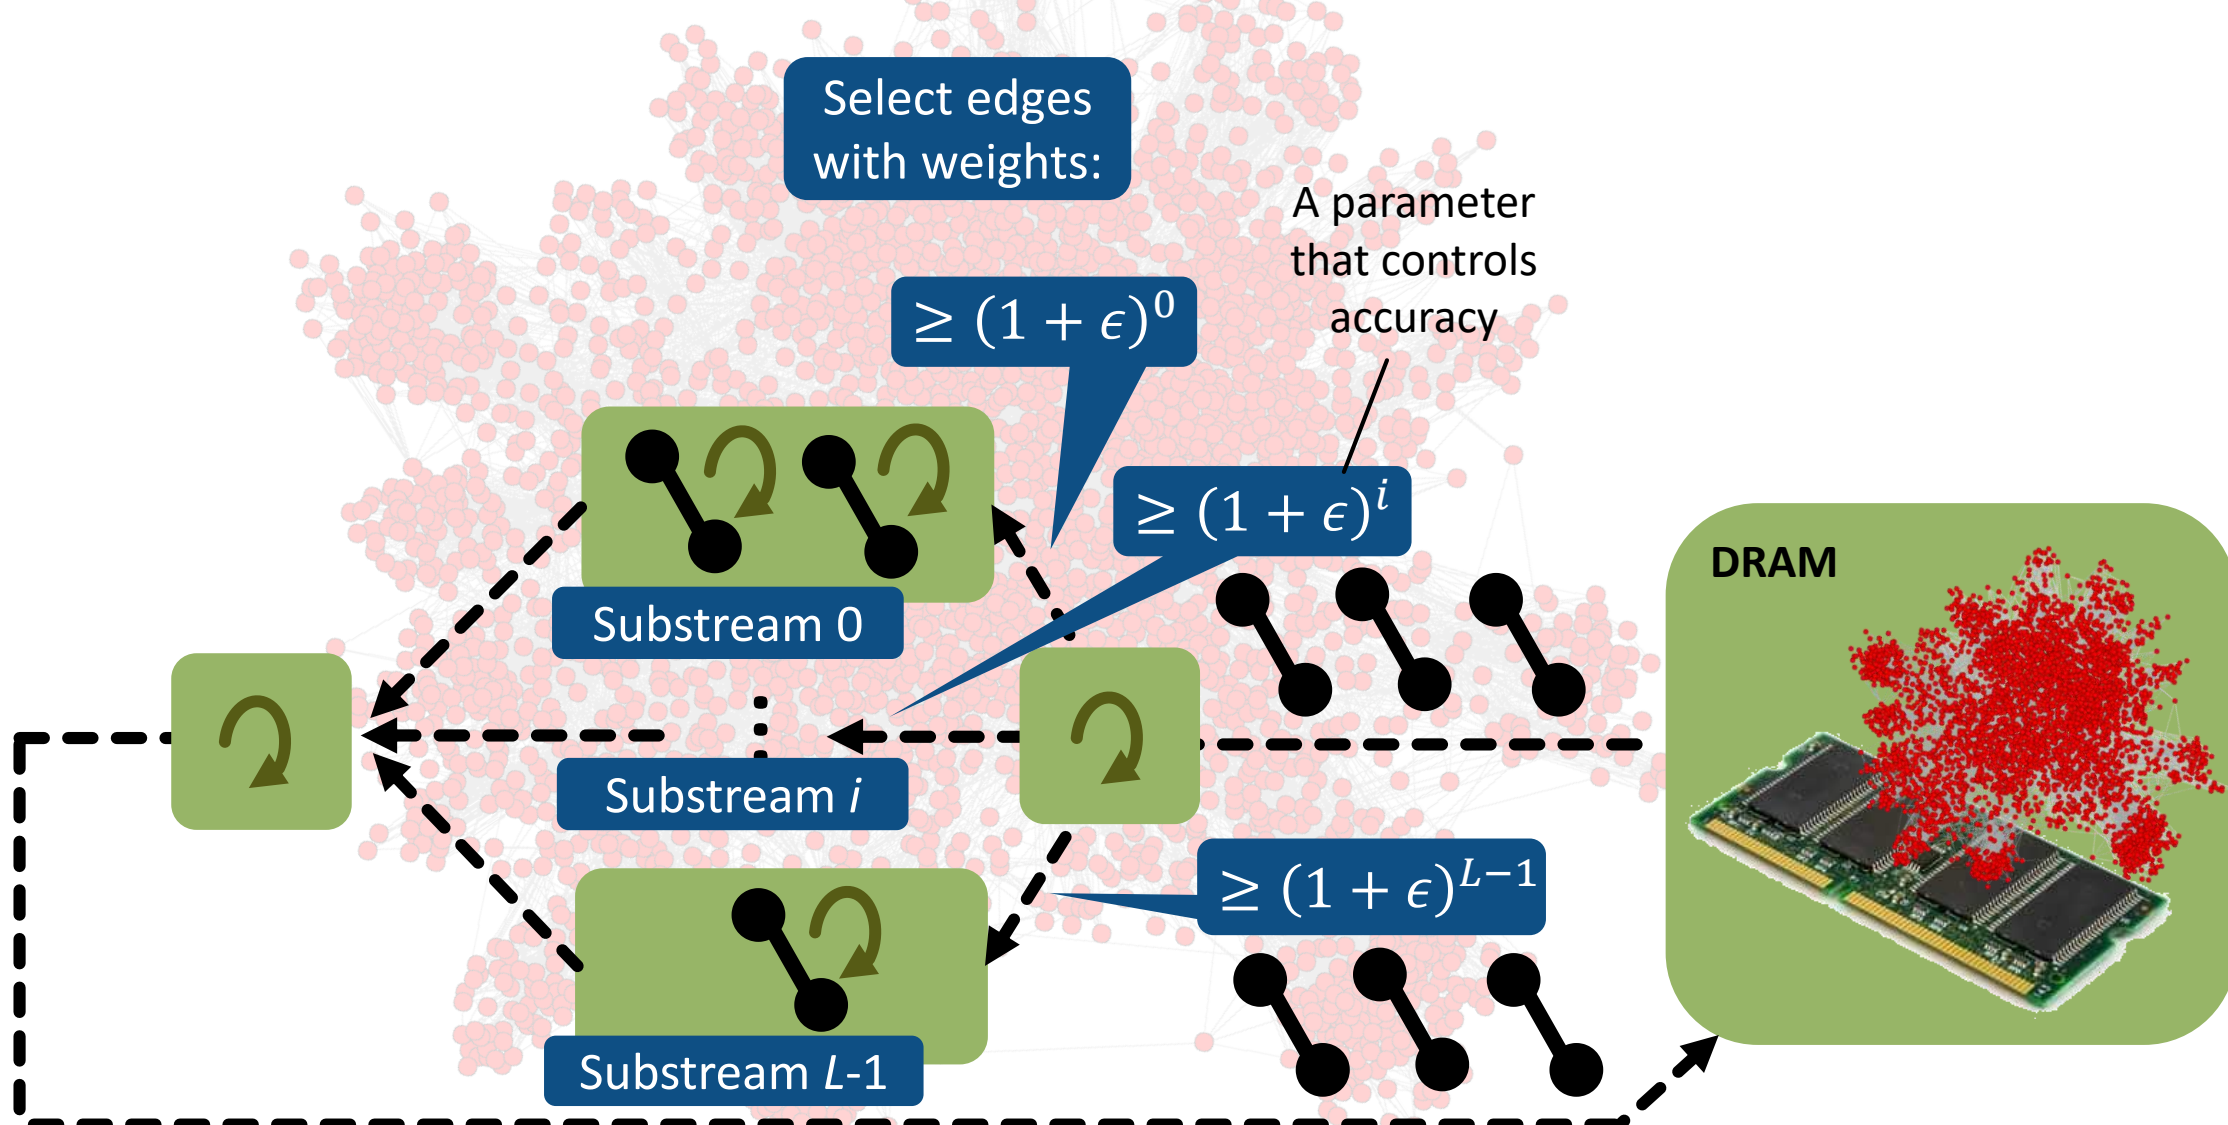
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



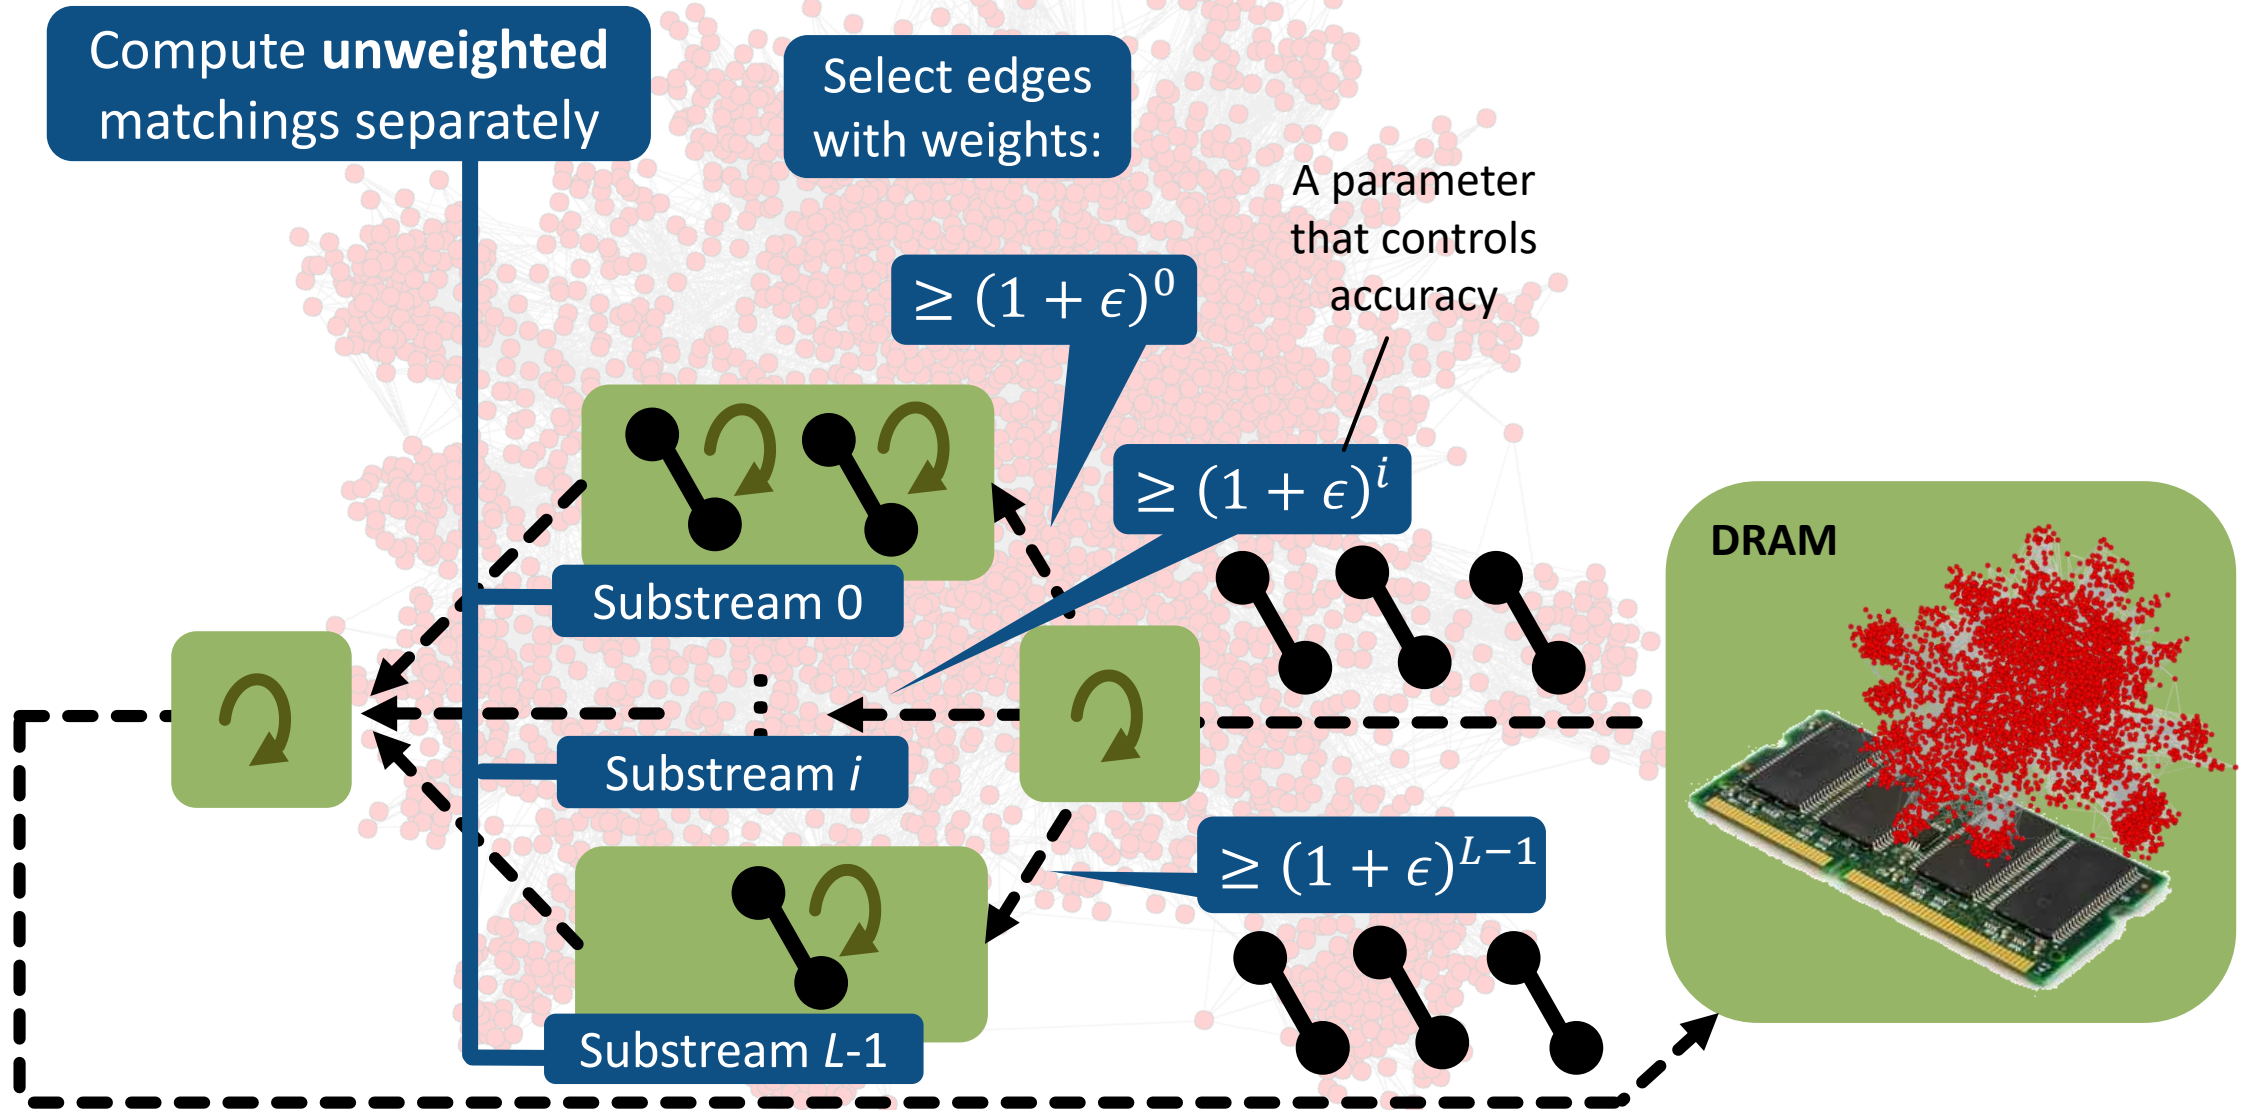
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

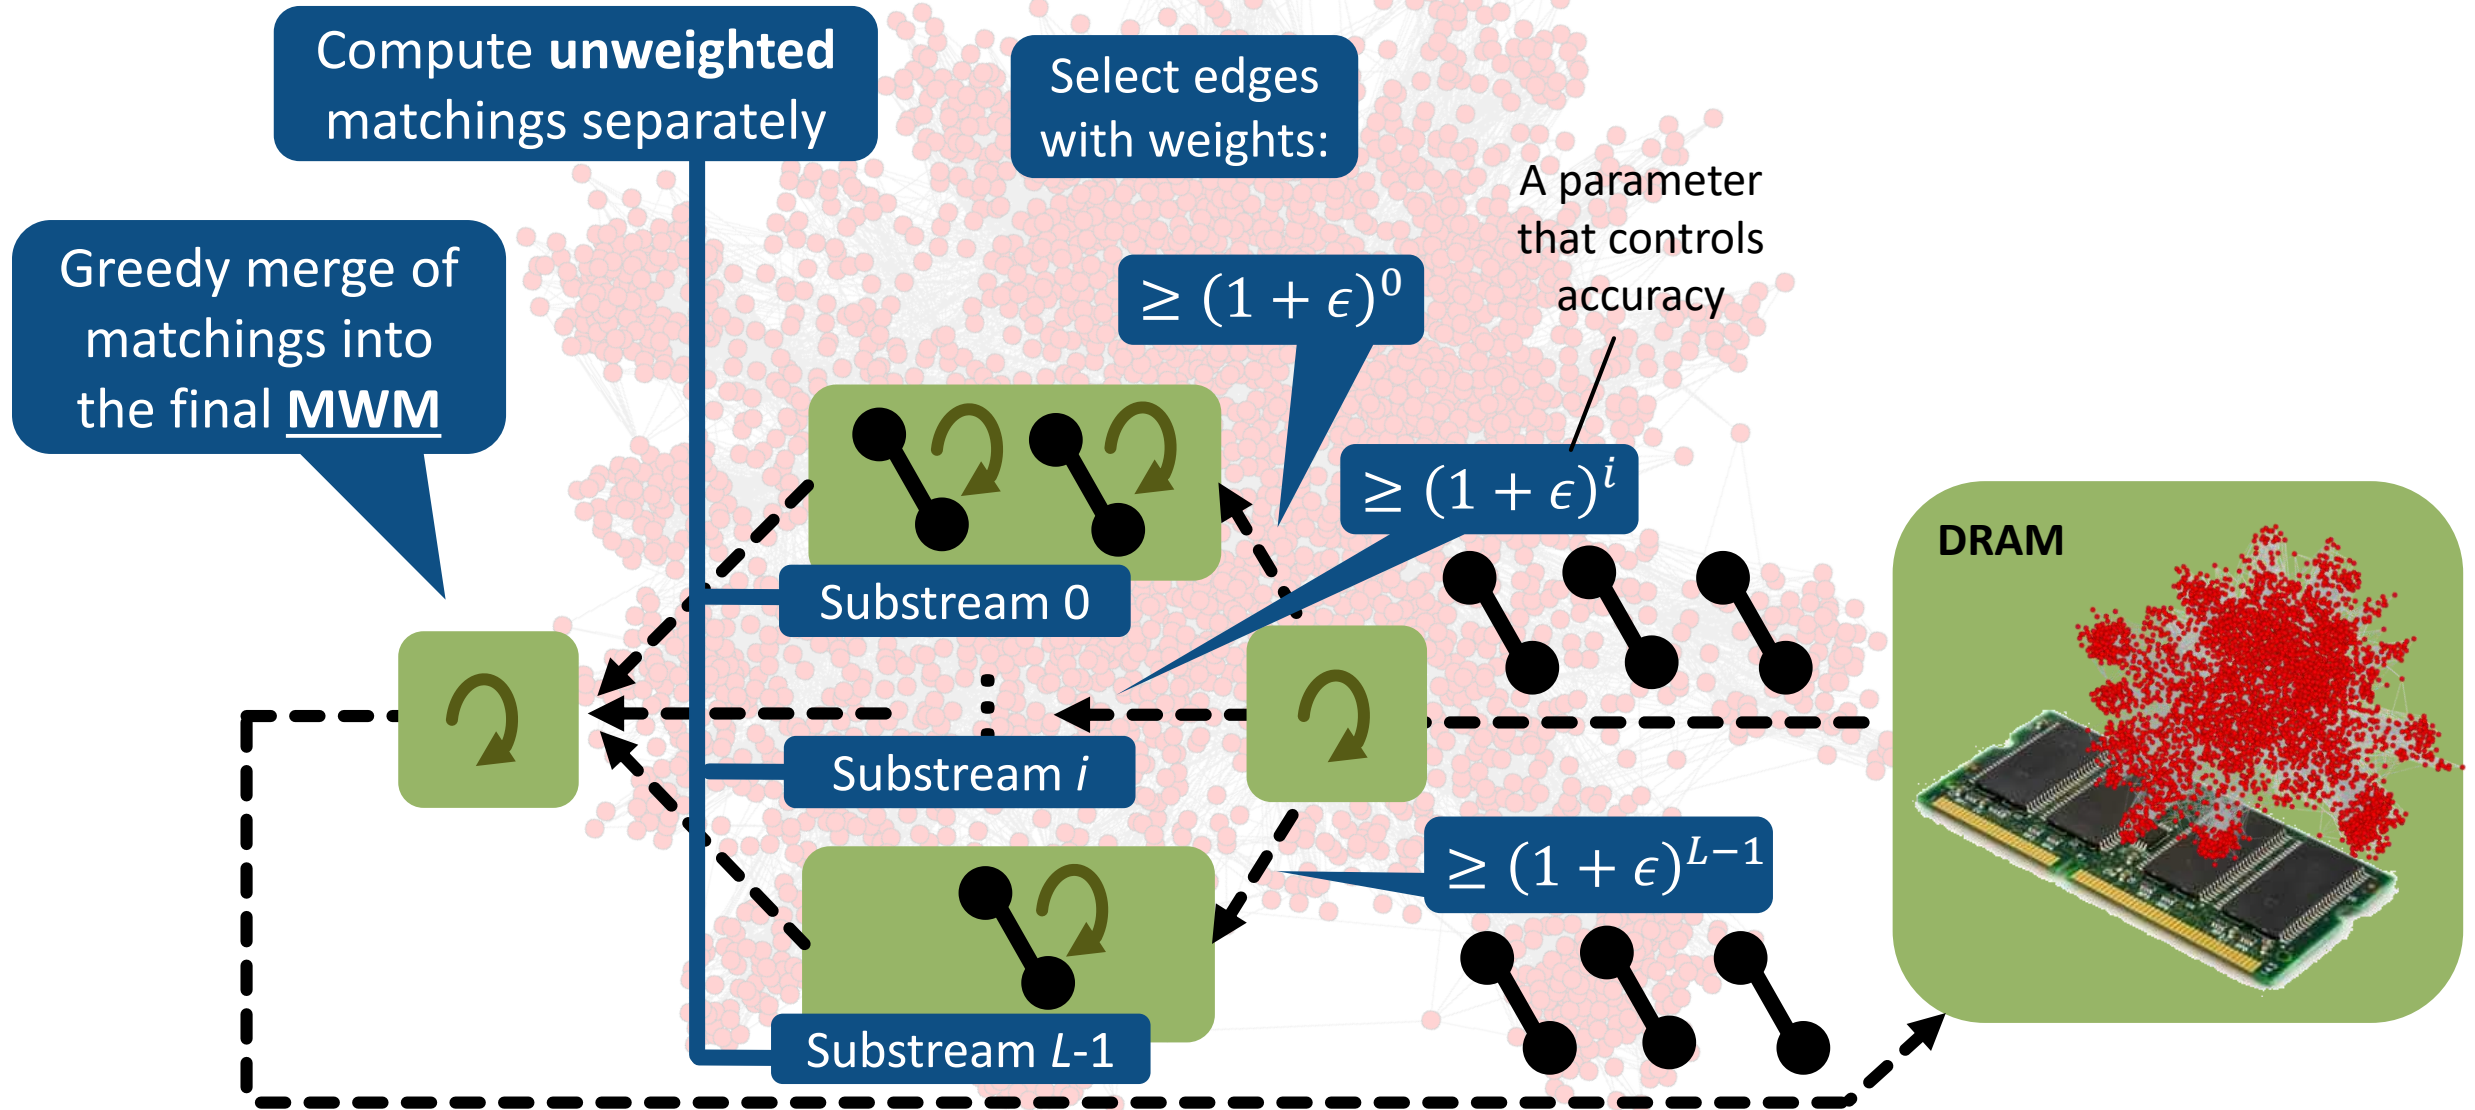
# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

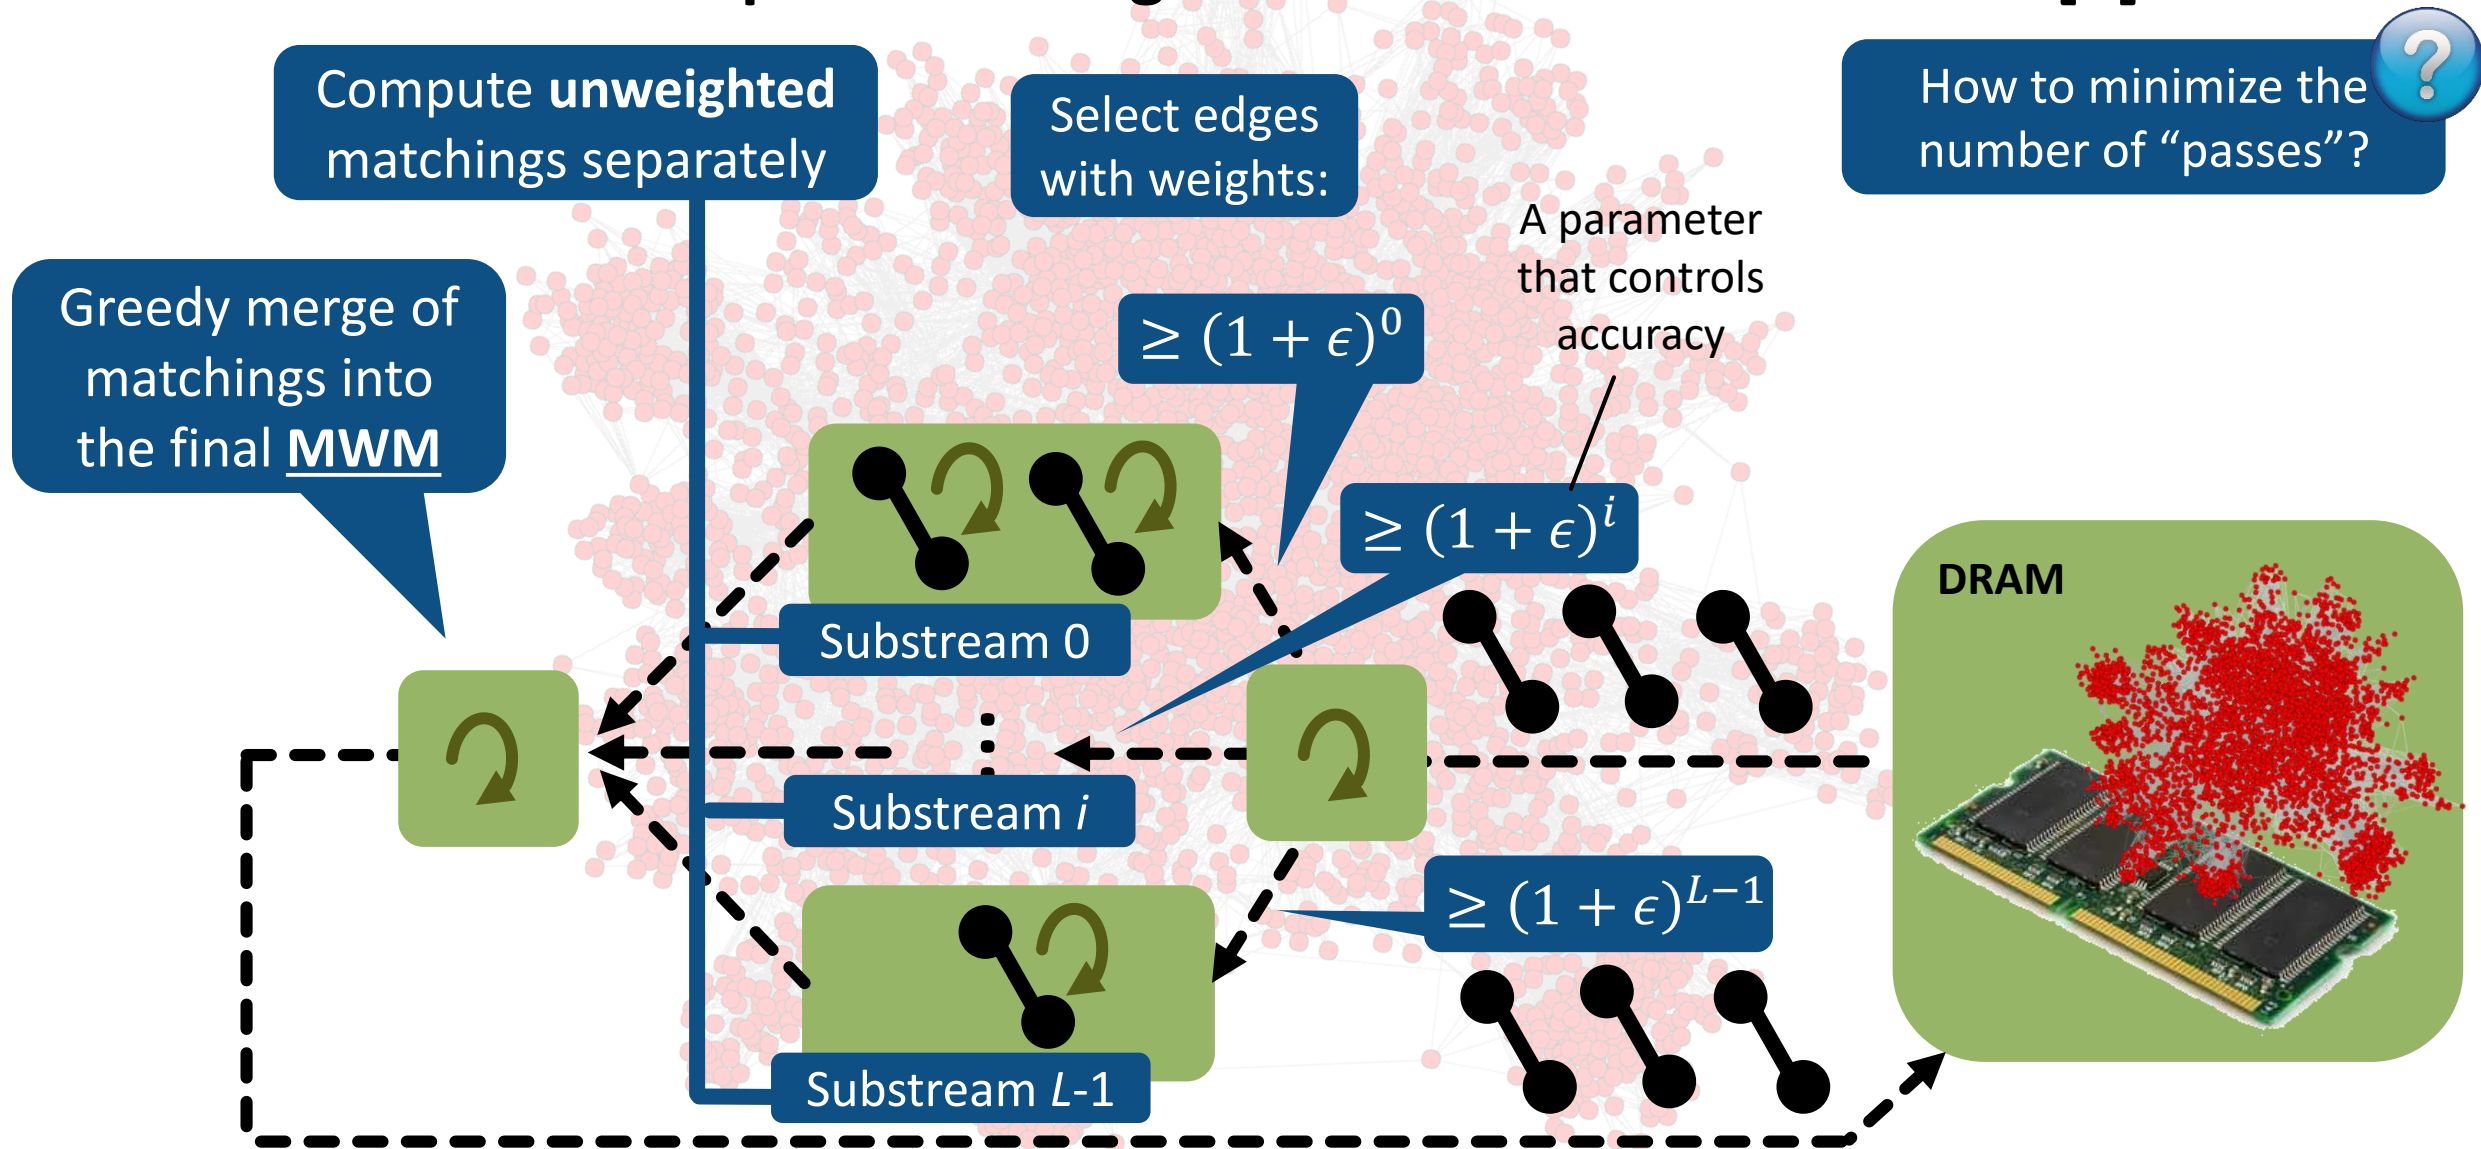


# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



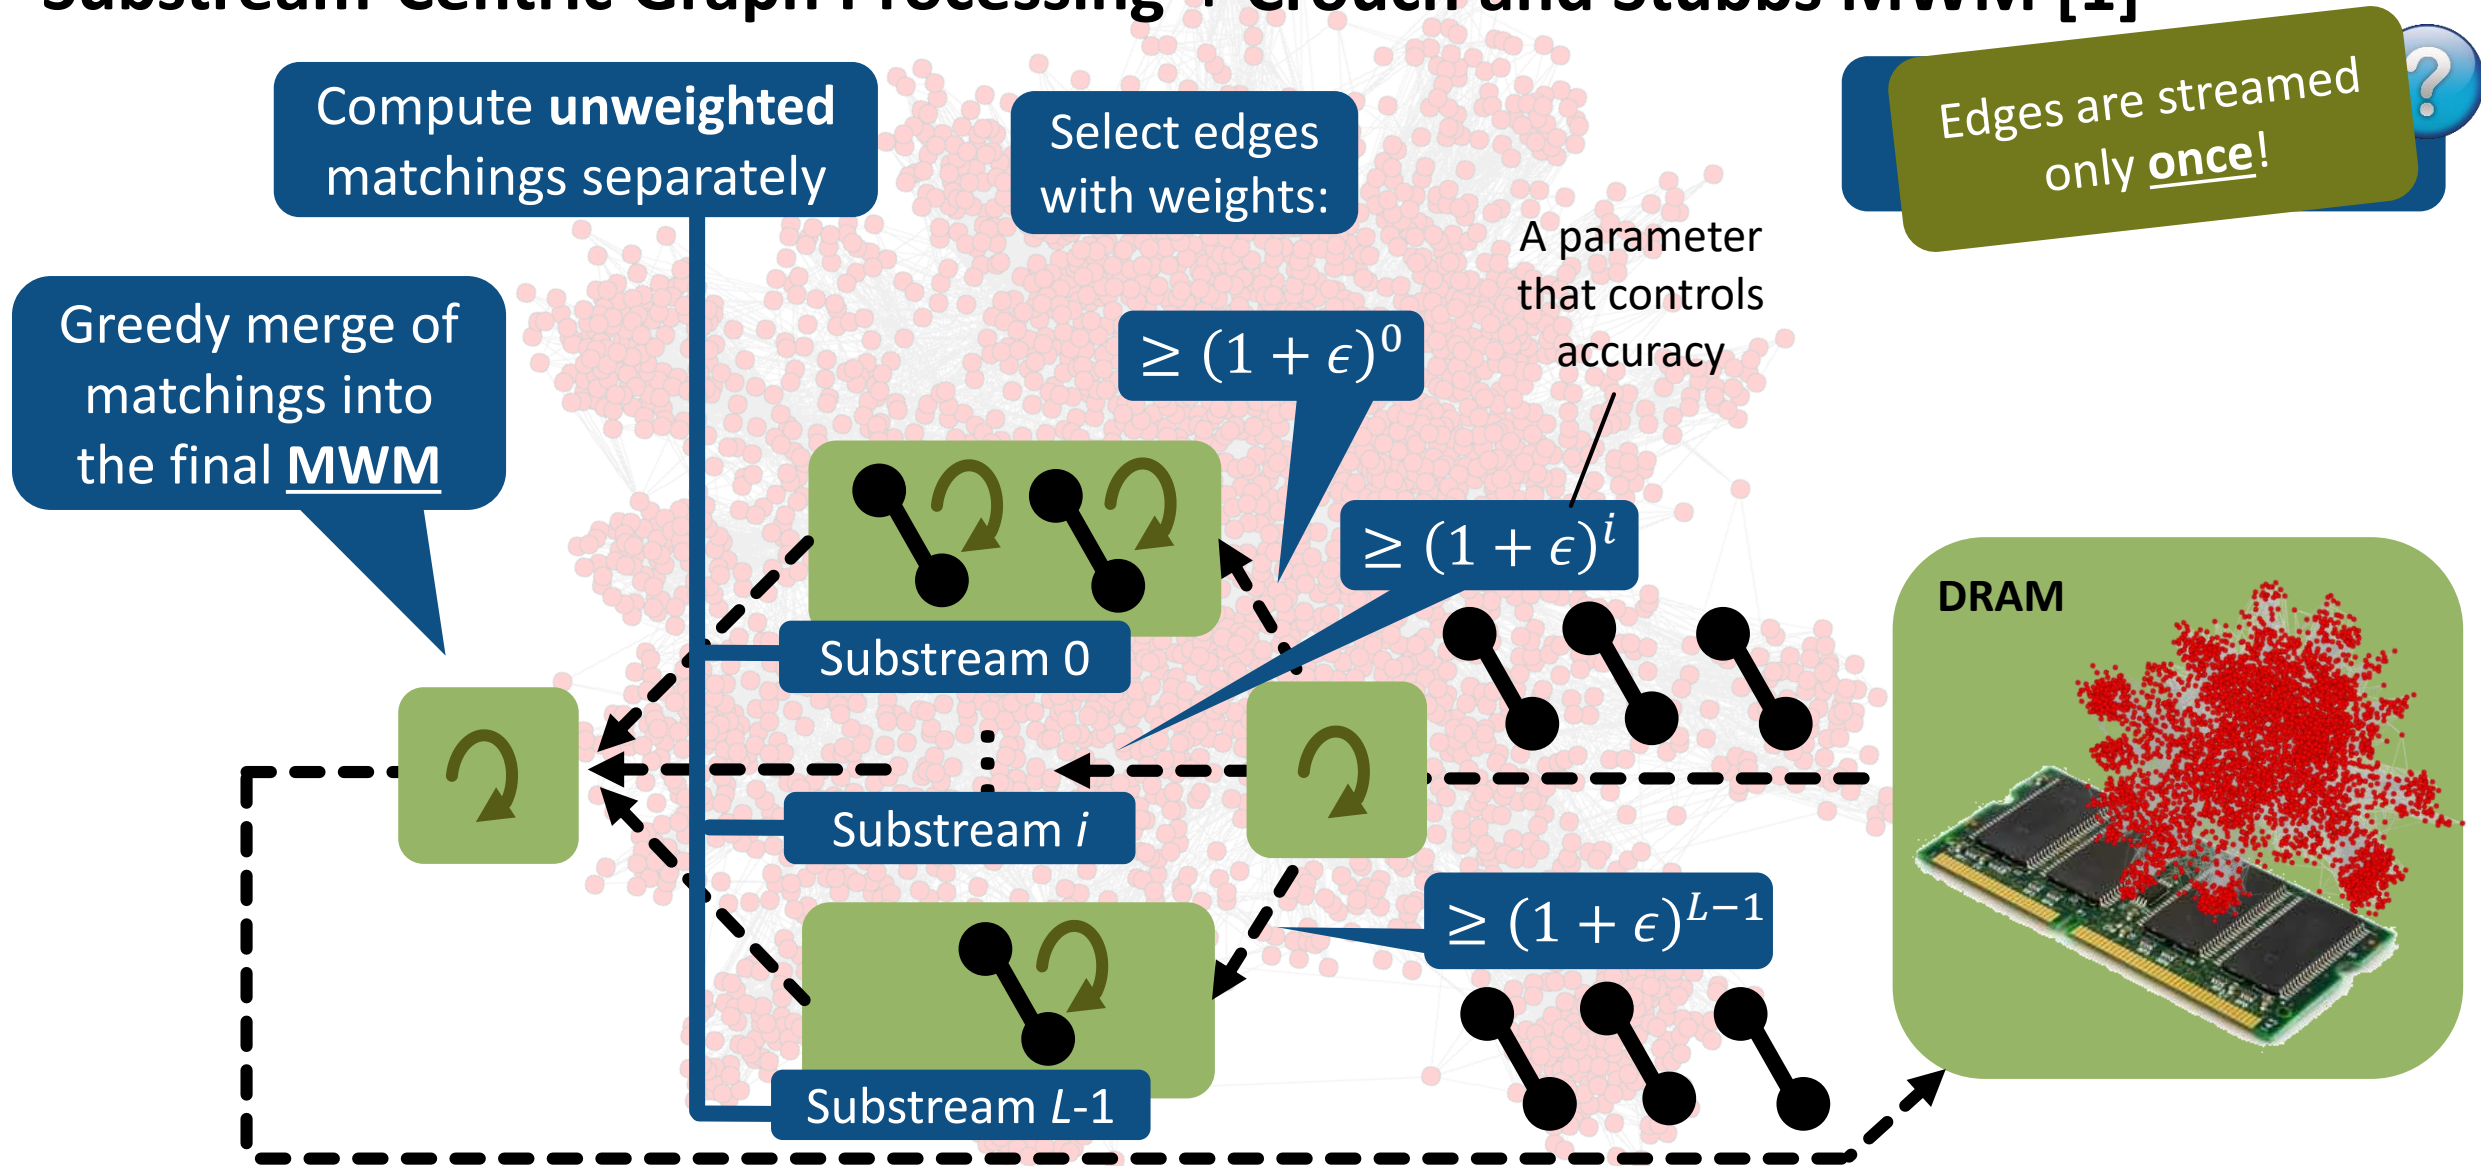
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



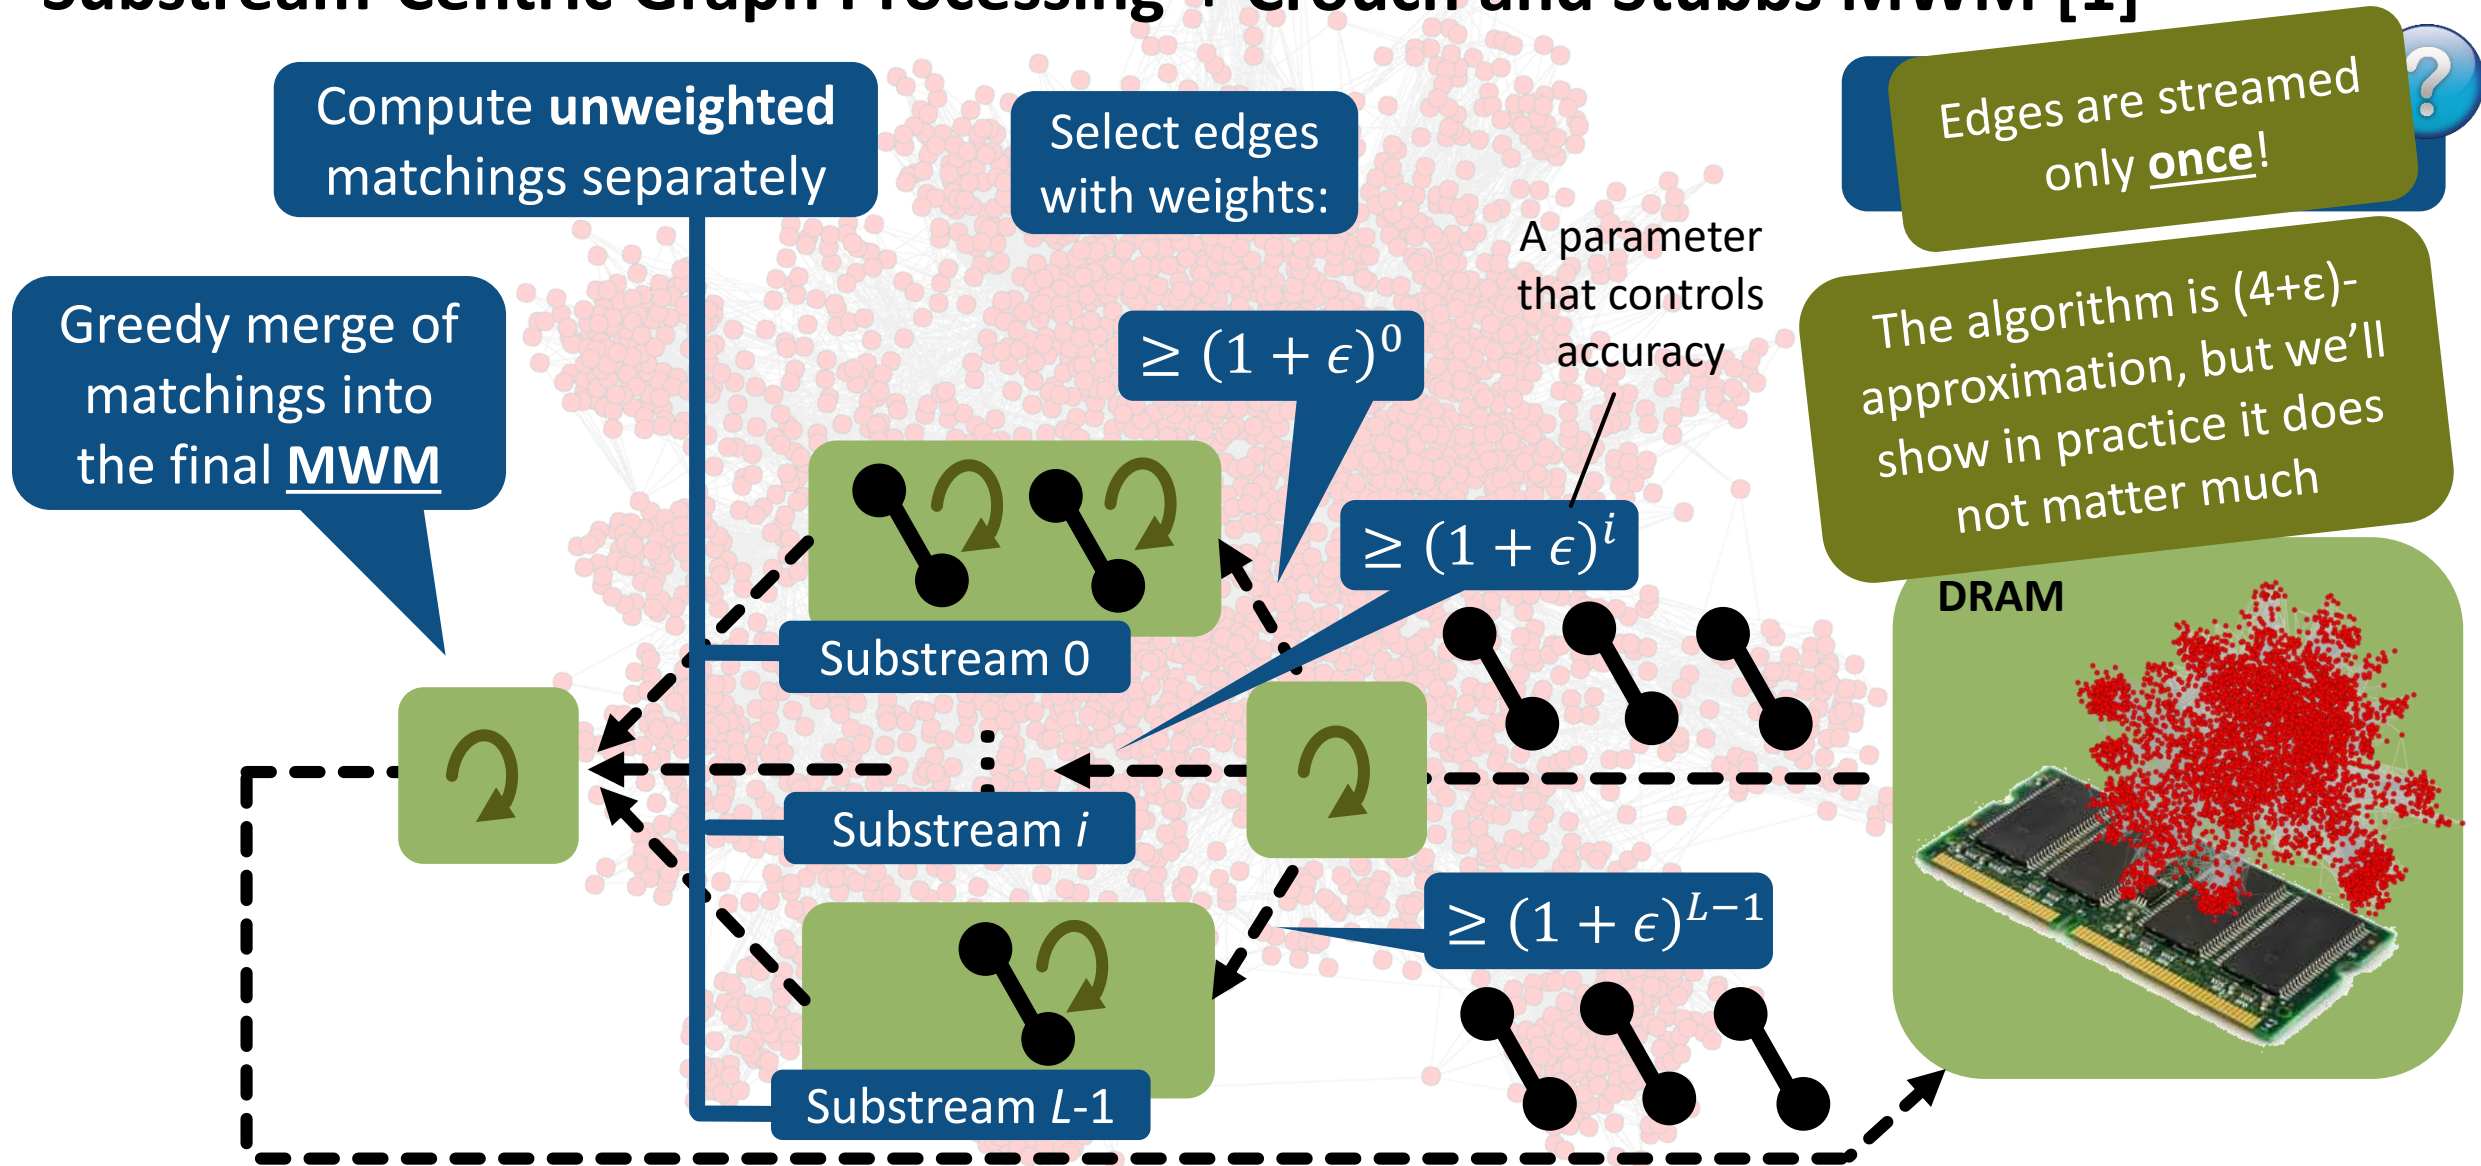
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



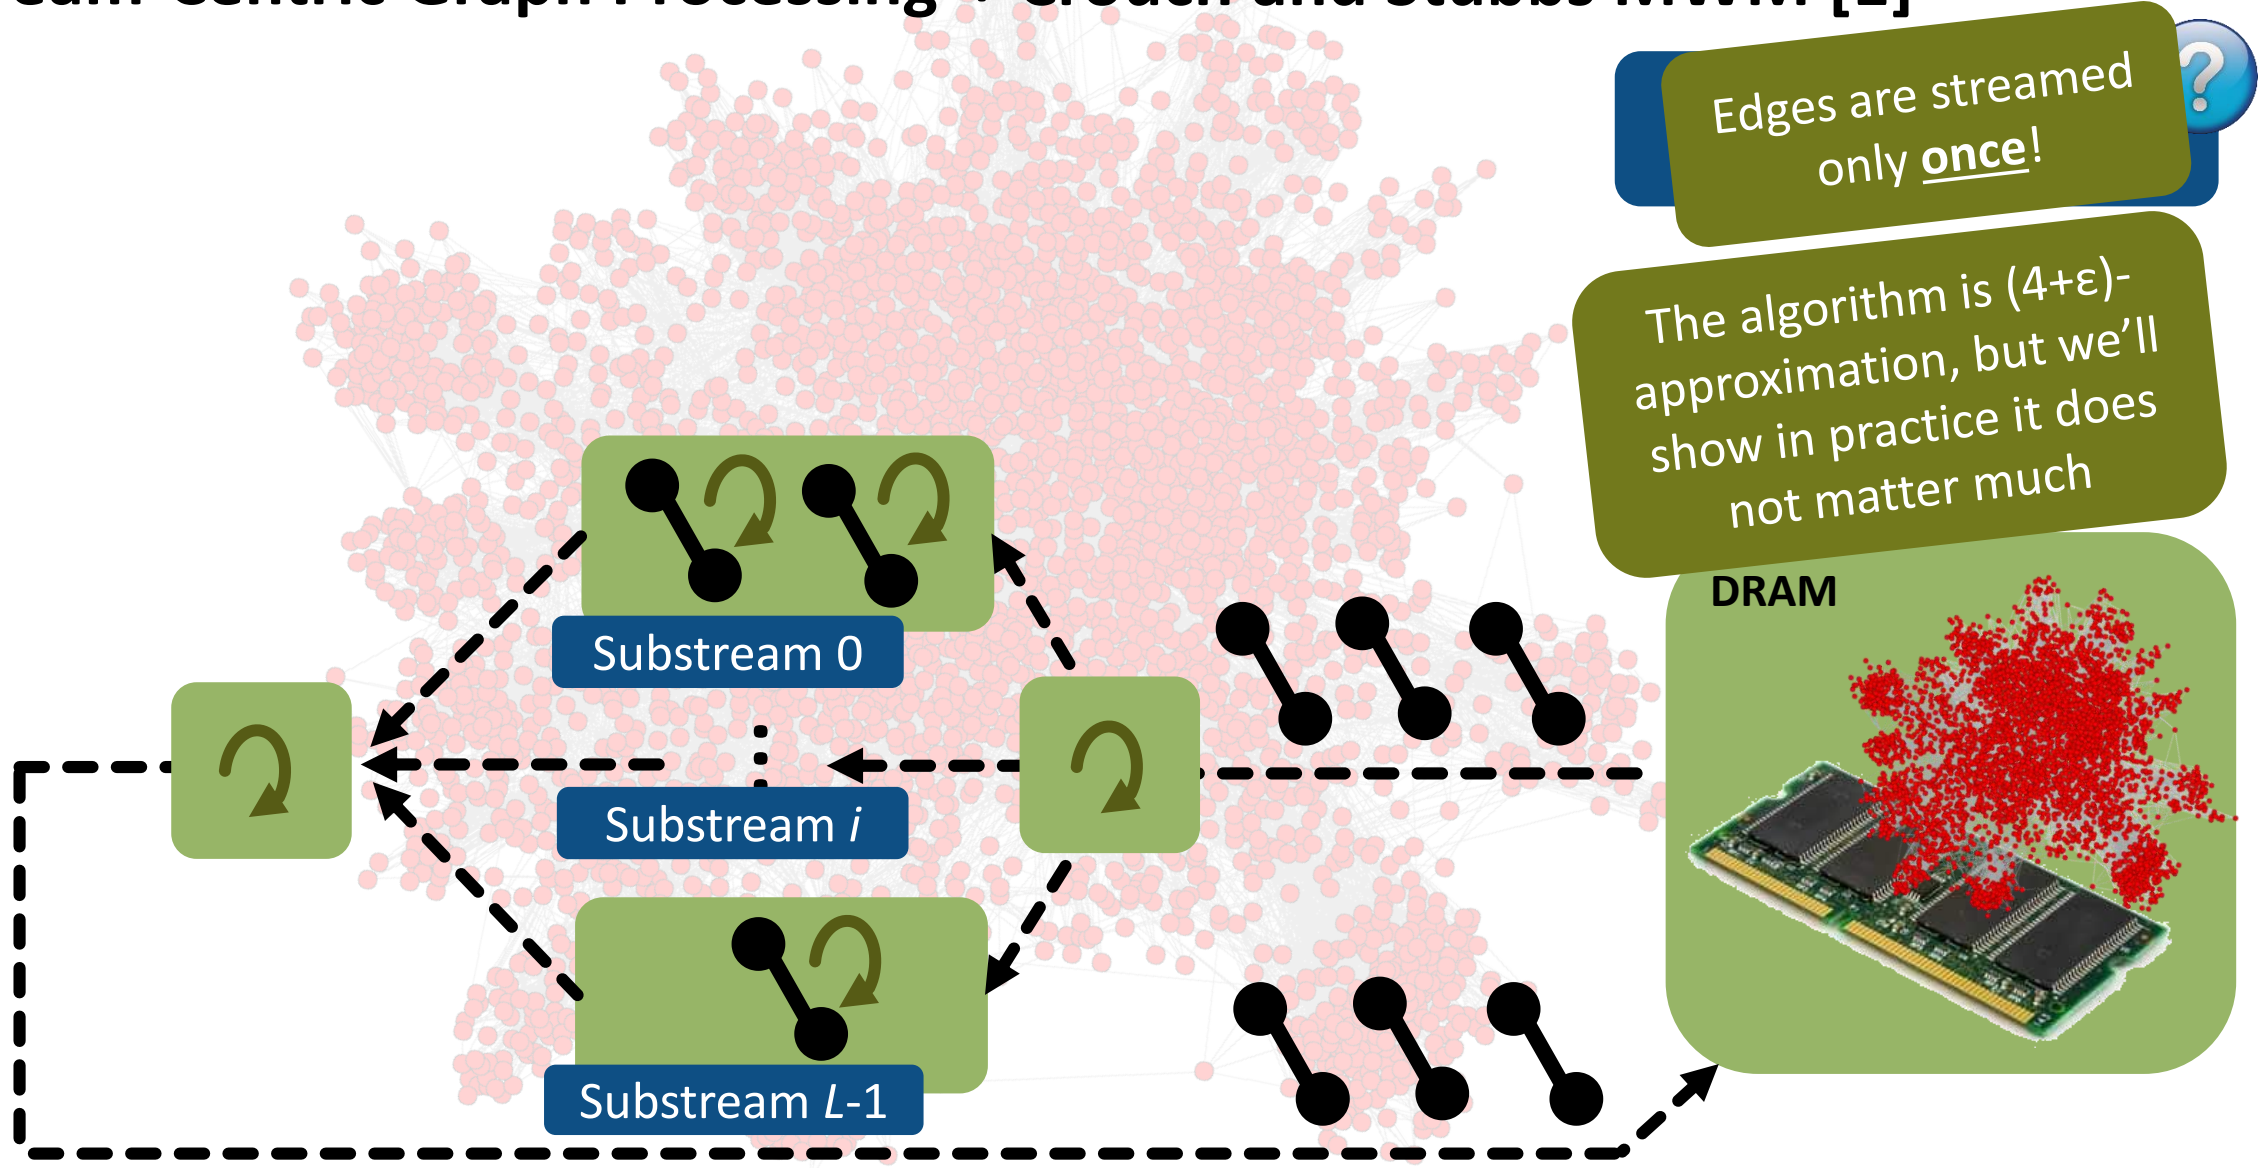
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]




[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

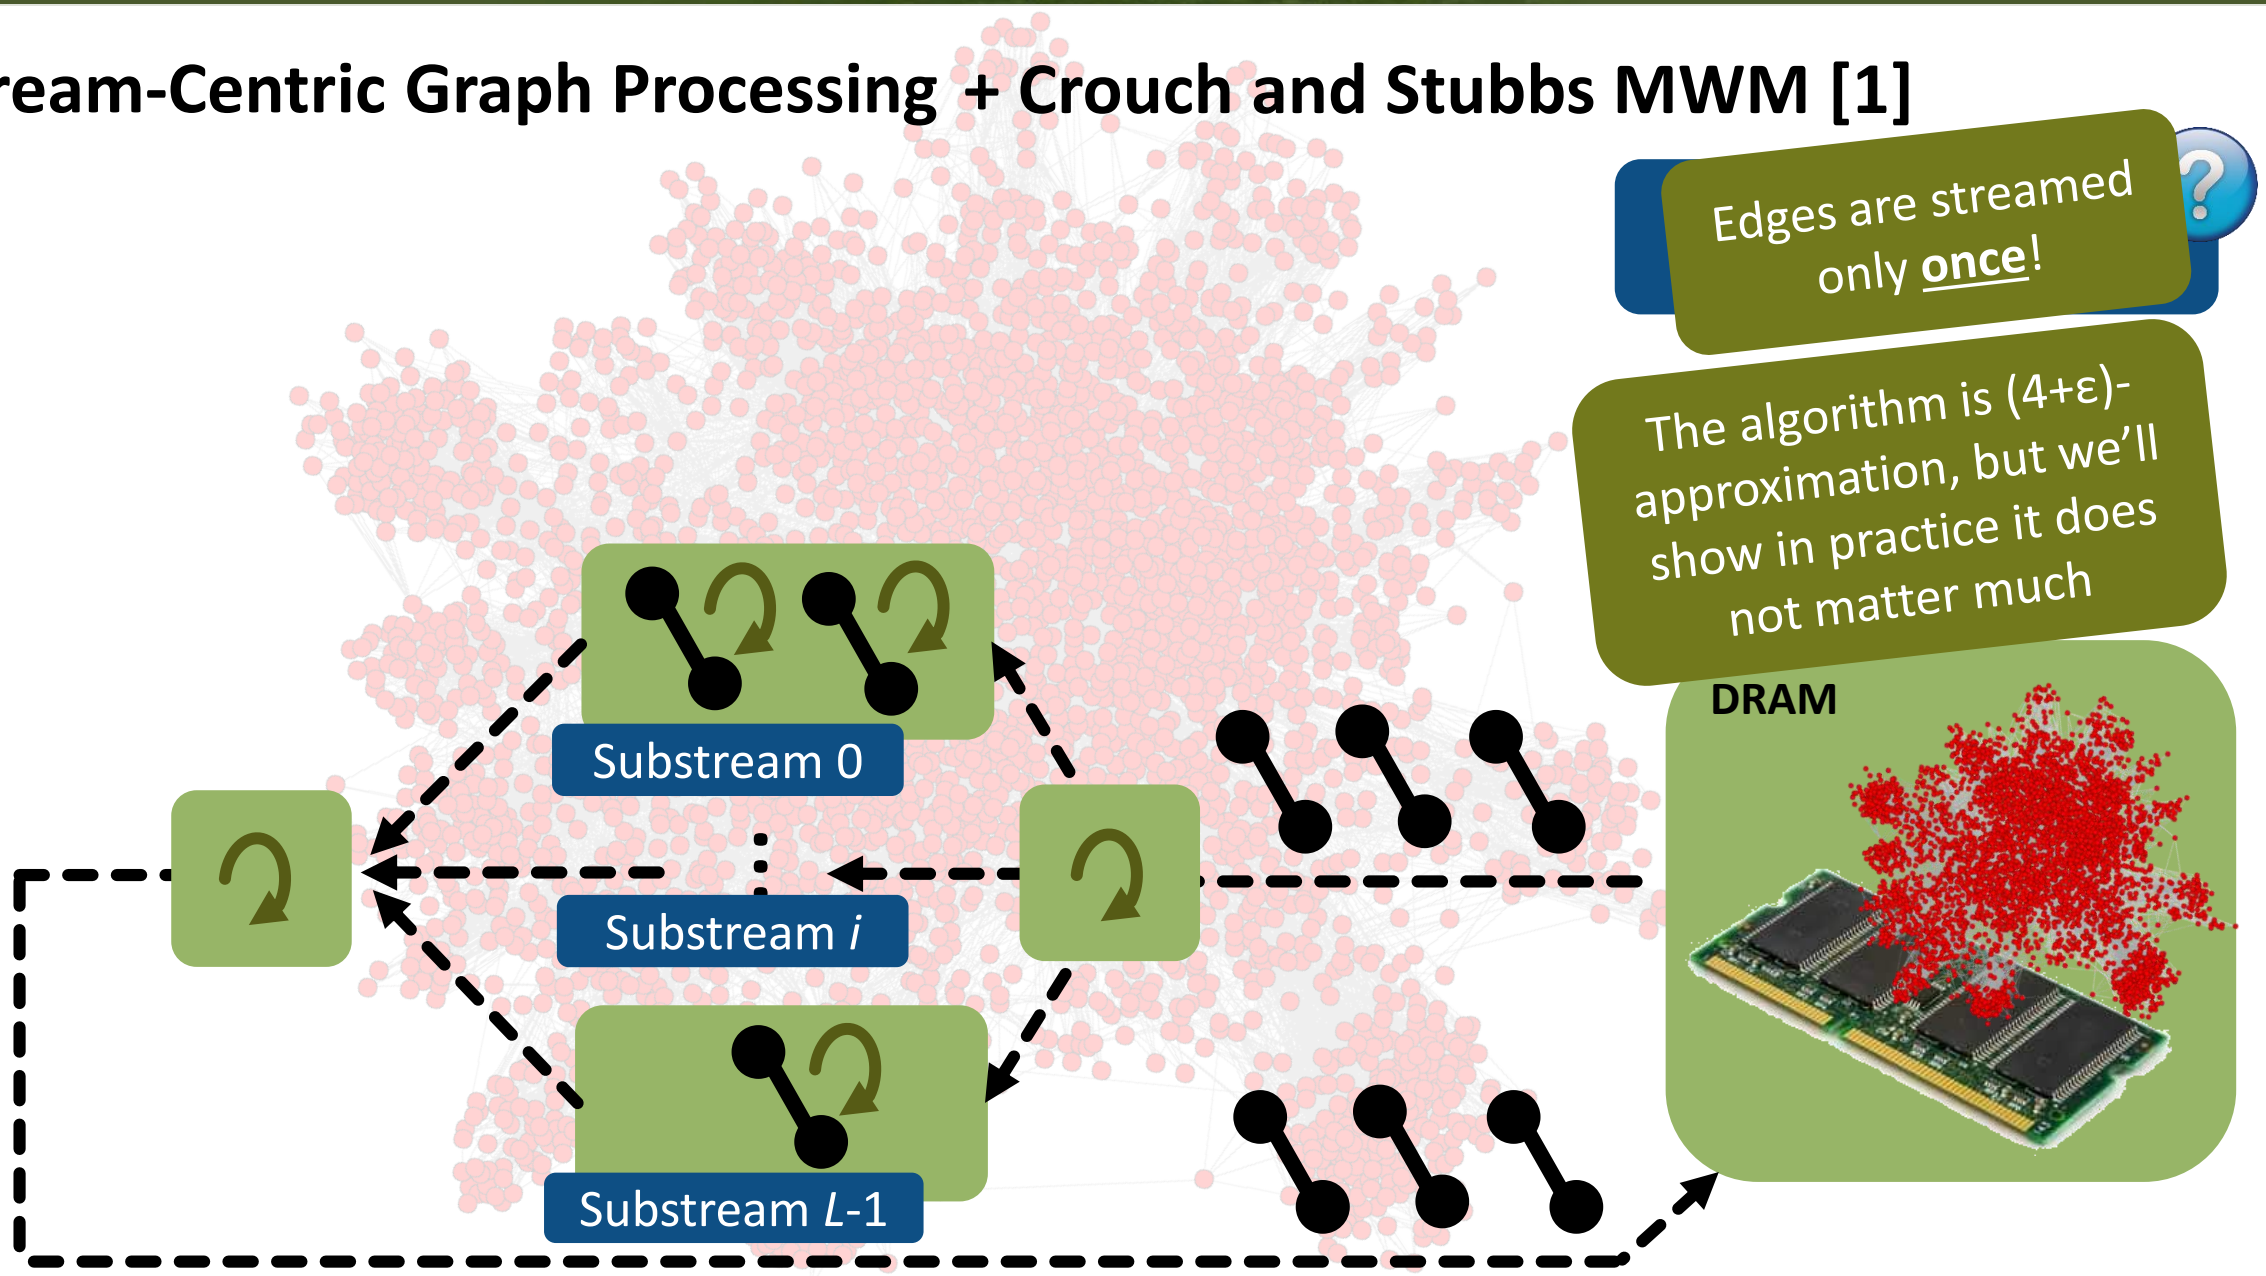
Edges are streamed only once!

Part 4: Mapping the algorithm to the „right“ hardware configuration




[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



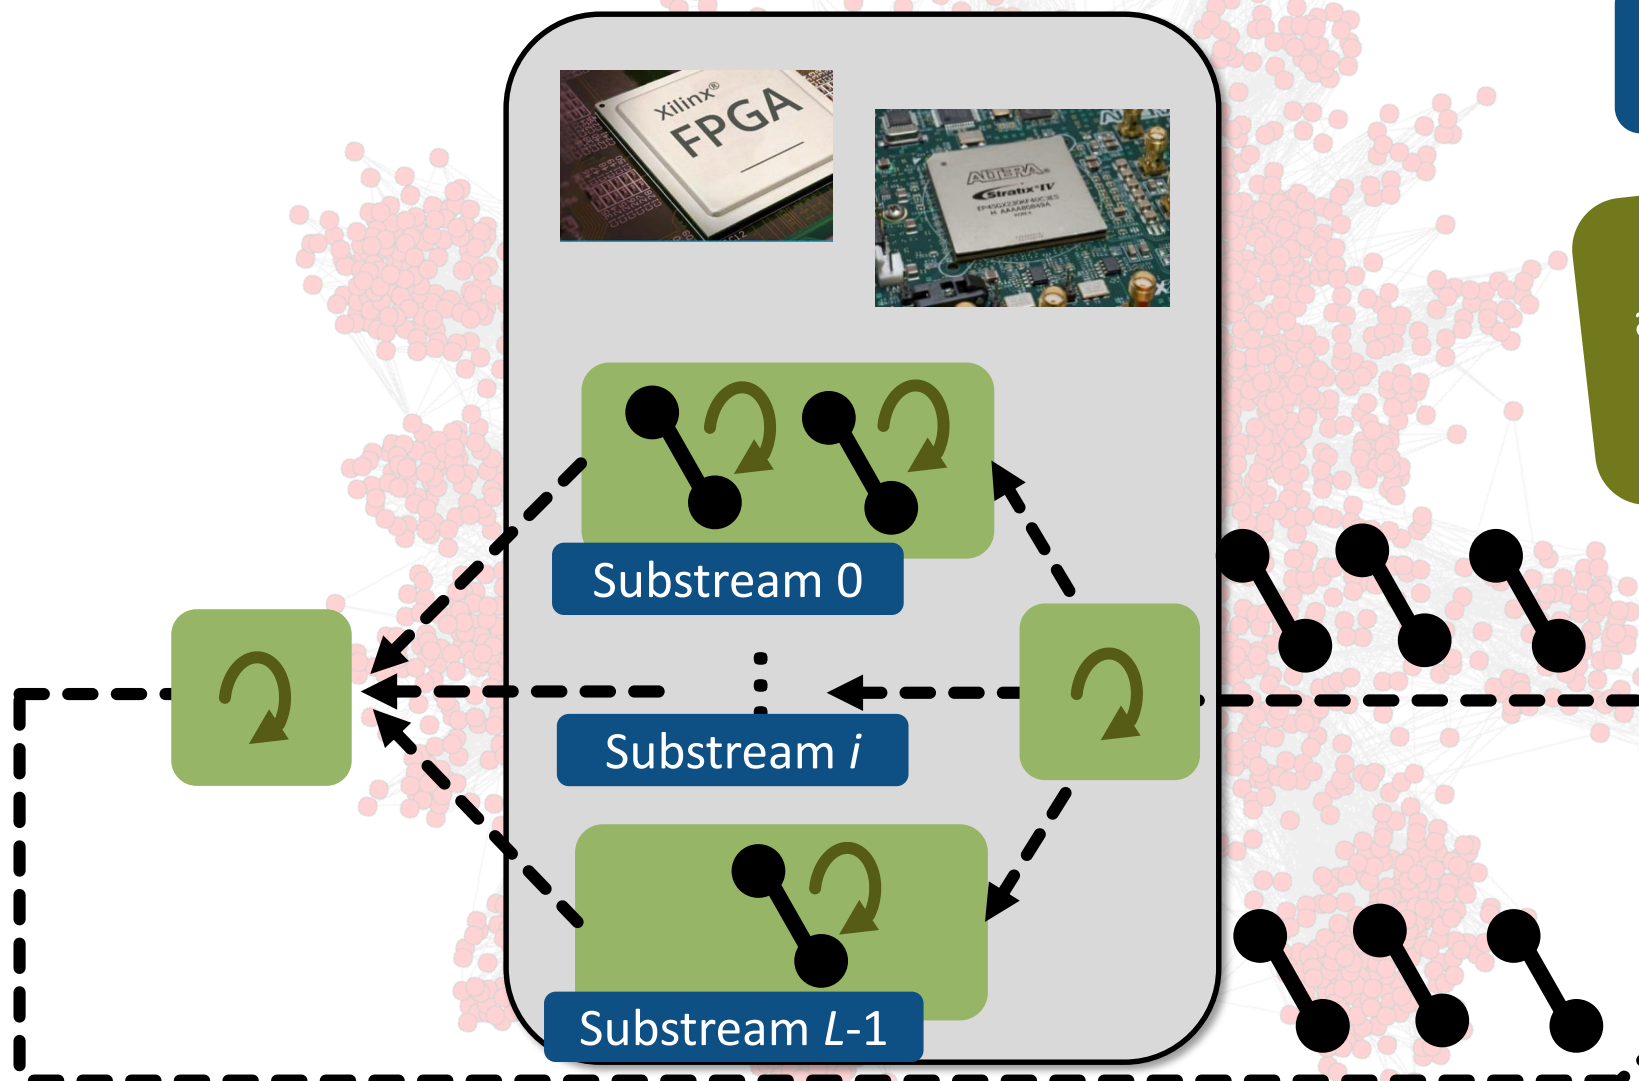
Edges are streamed only once!

The algorithm is  $(4+\epsilon)$ -approximation, but we'll show in practice it does not matter much




[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



Edges are streamed only once!

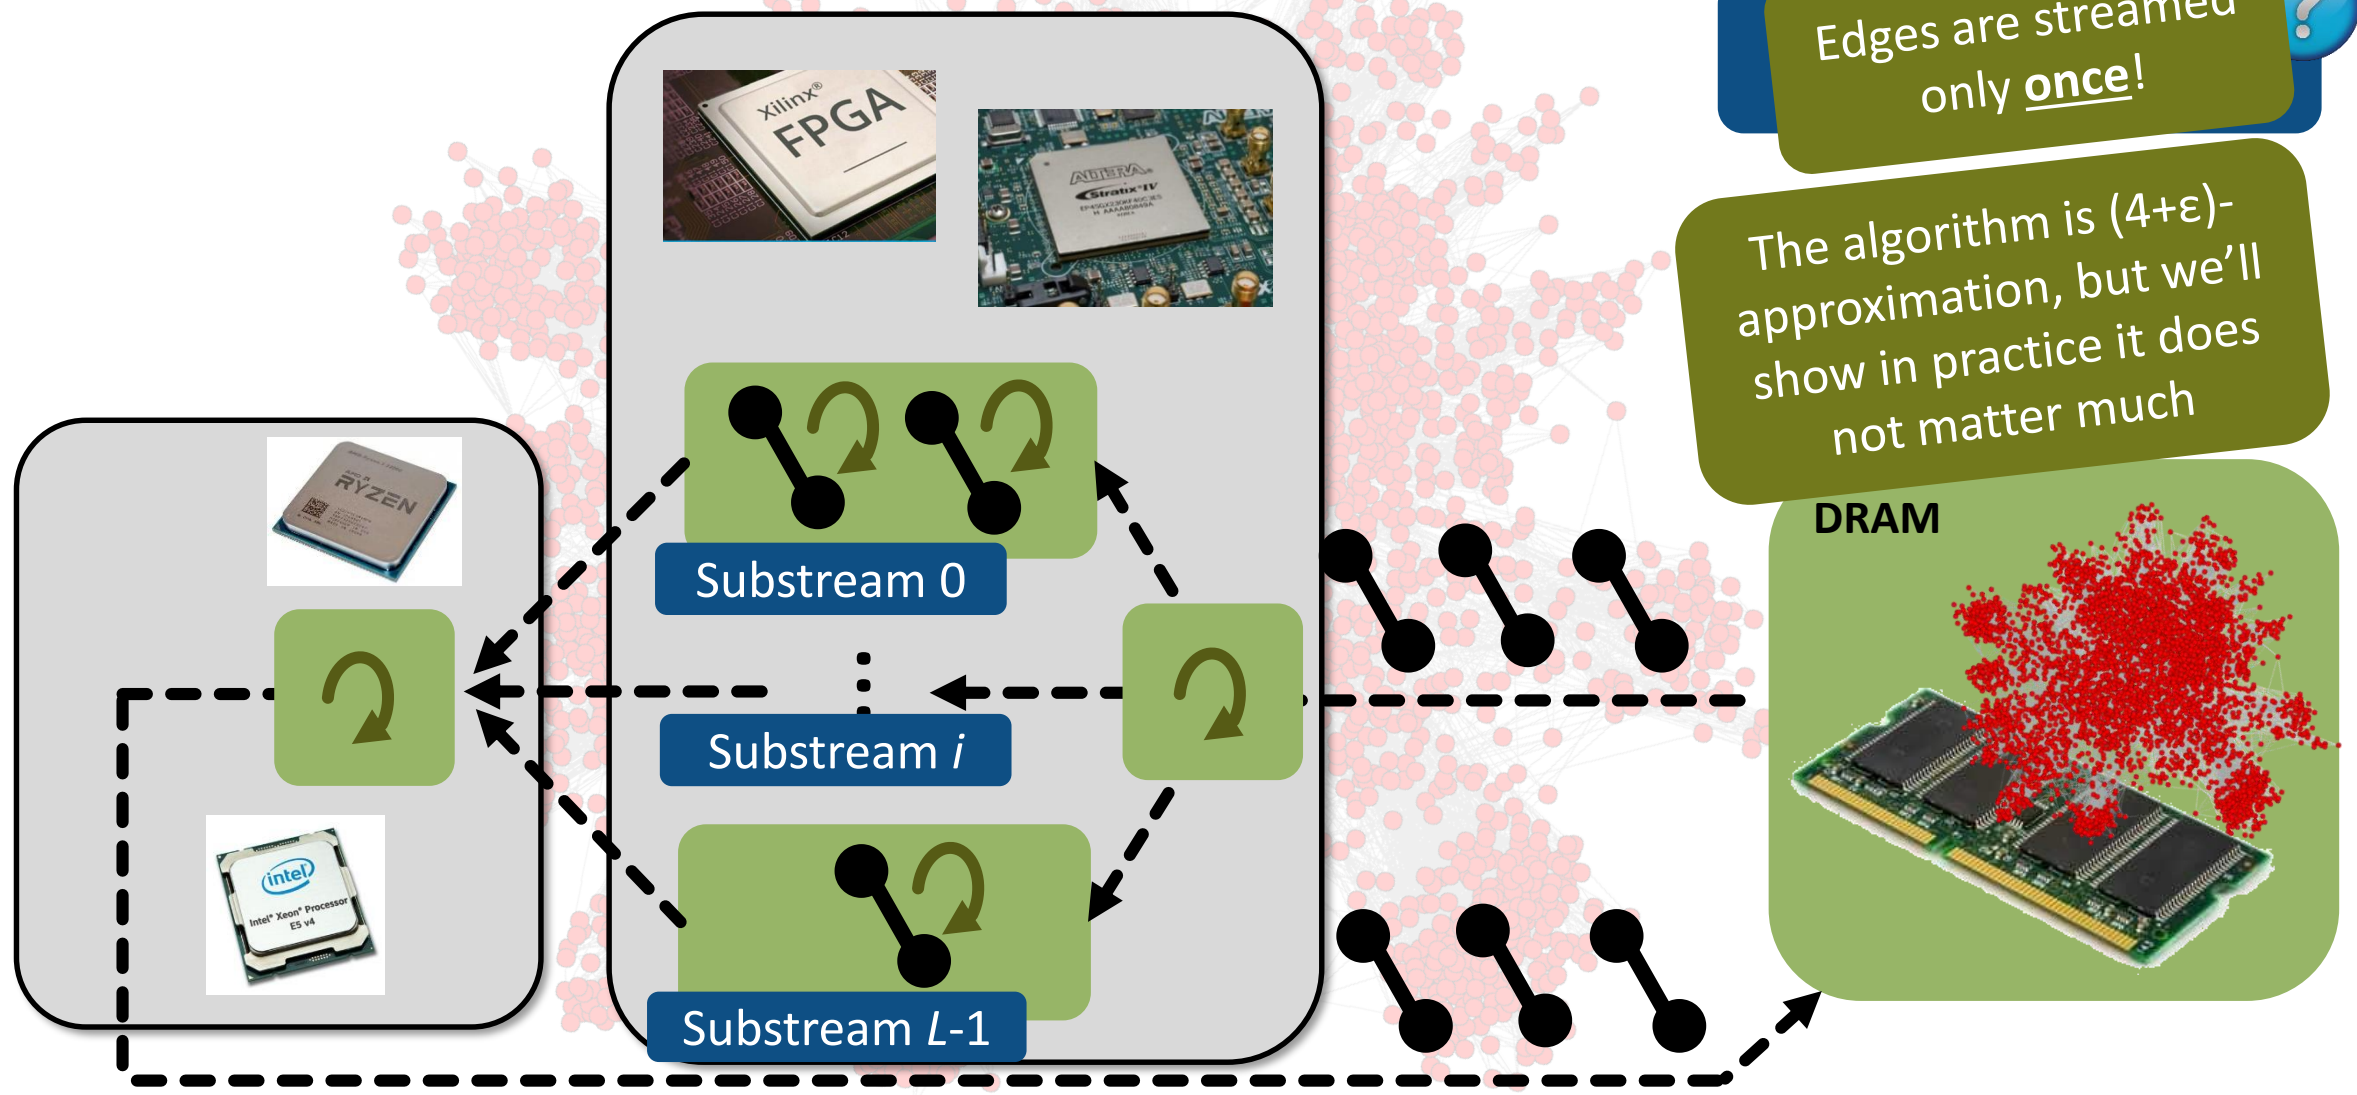
The algorithm is  $(4+\epsilon)$ -approximation, but we'll show in practice it does not matter much



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

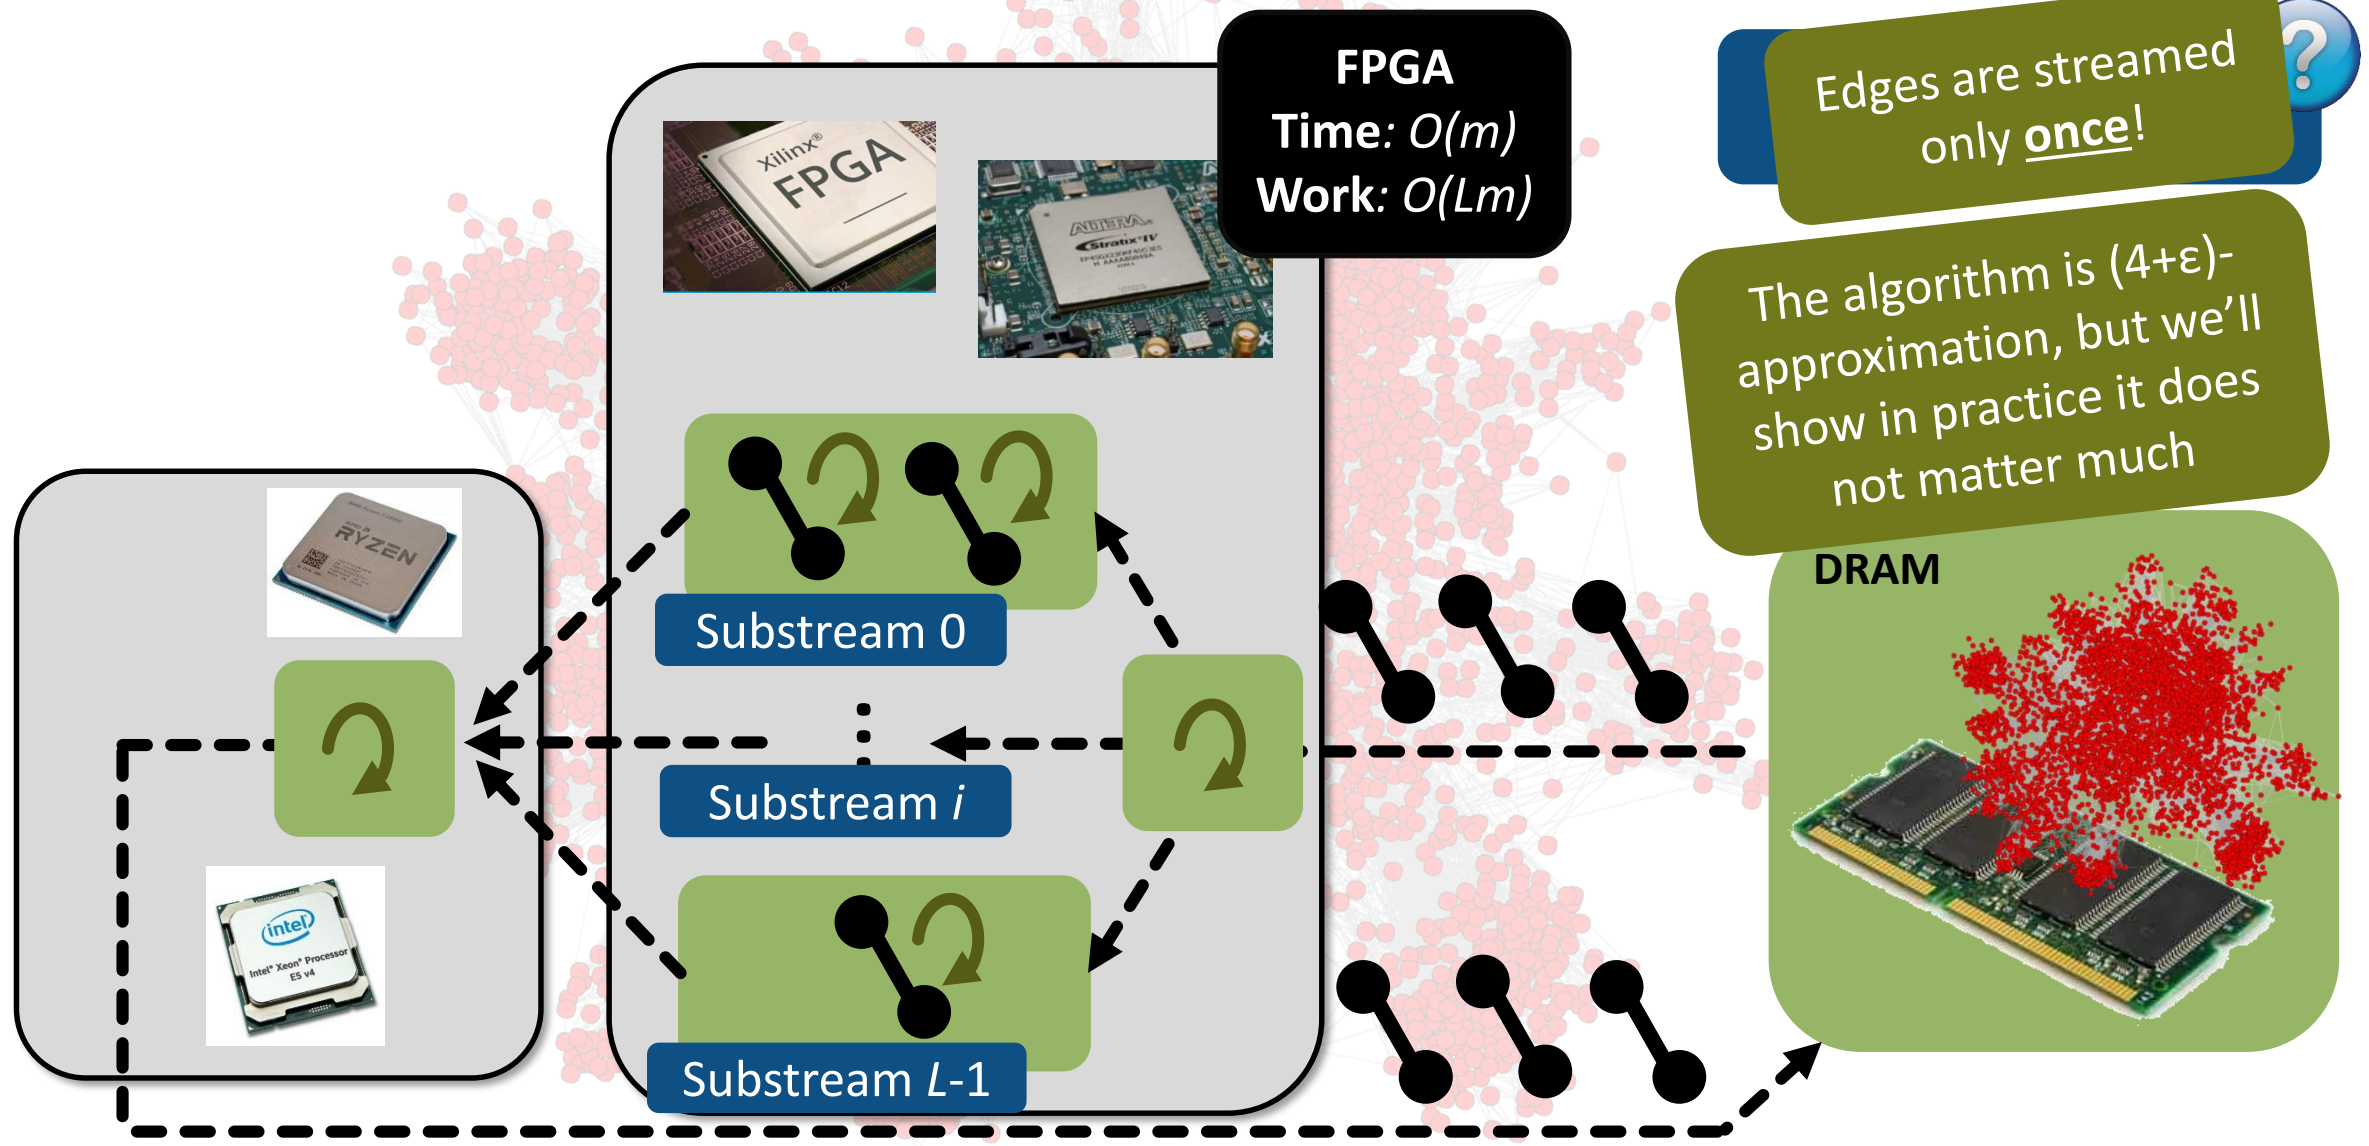


# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



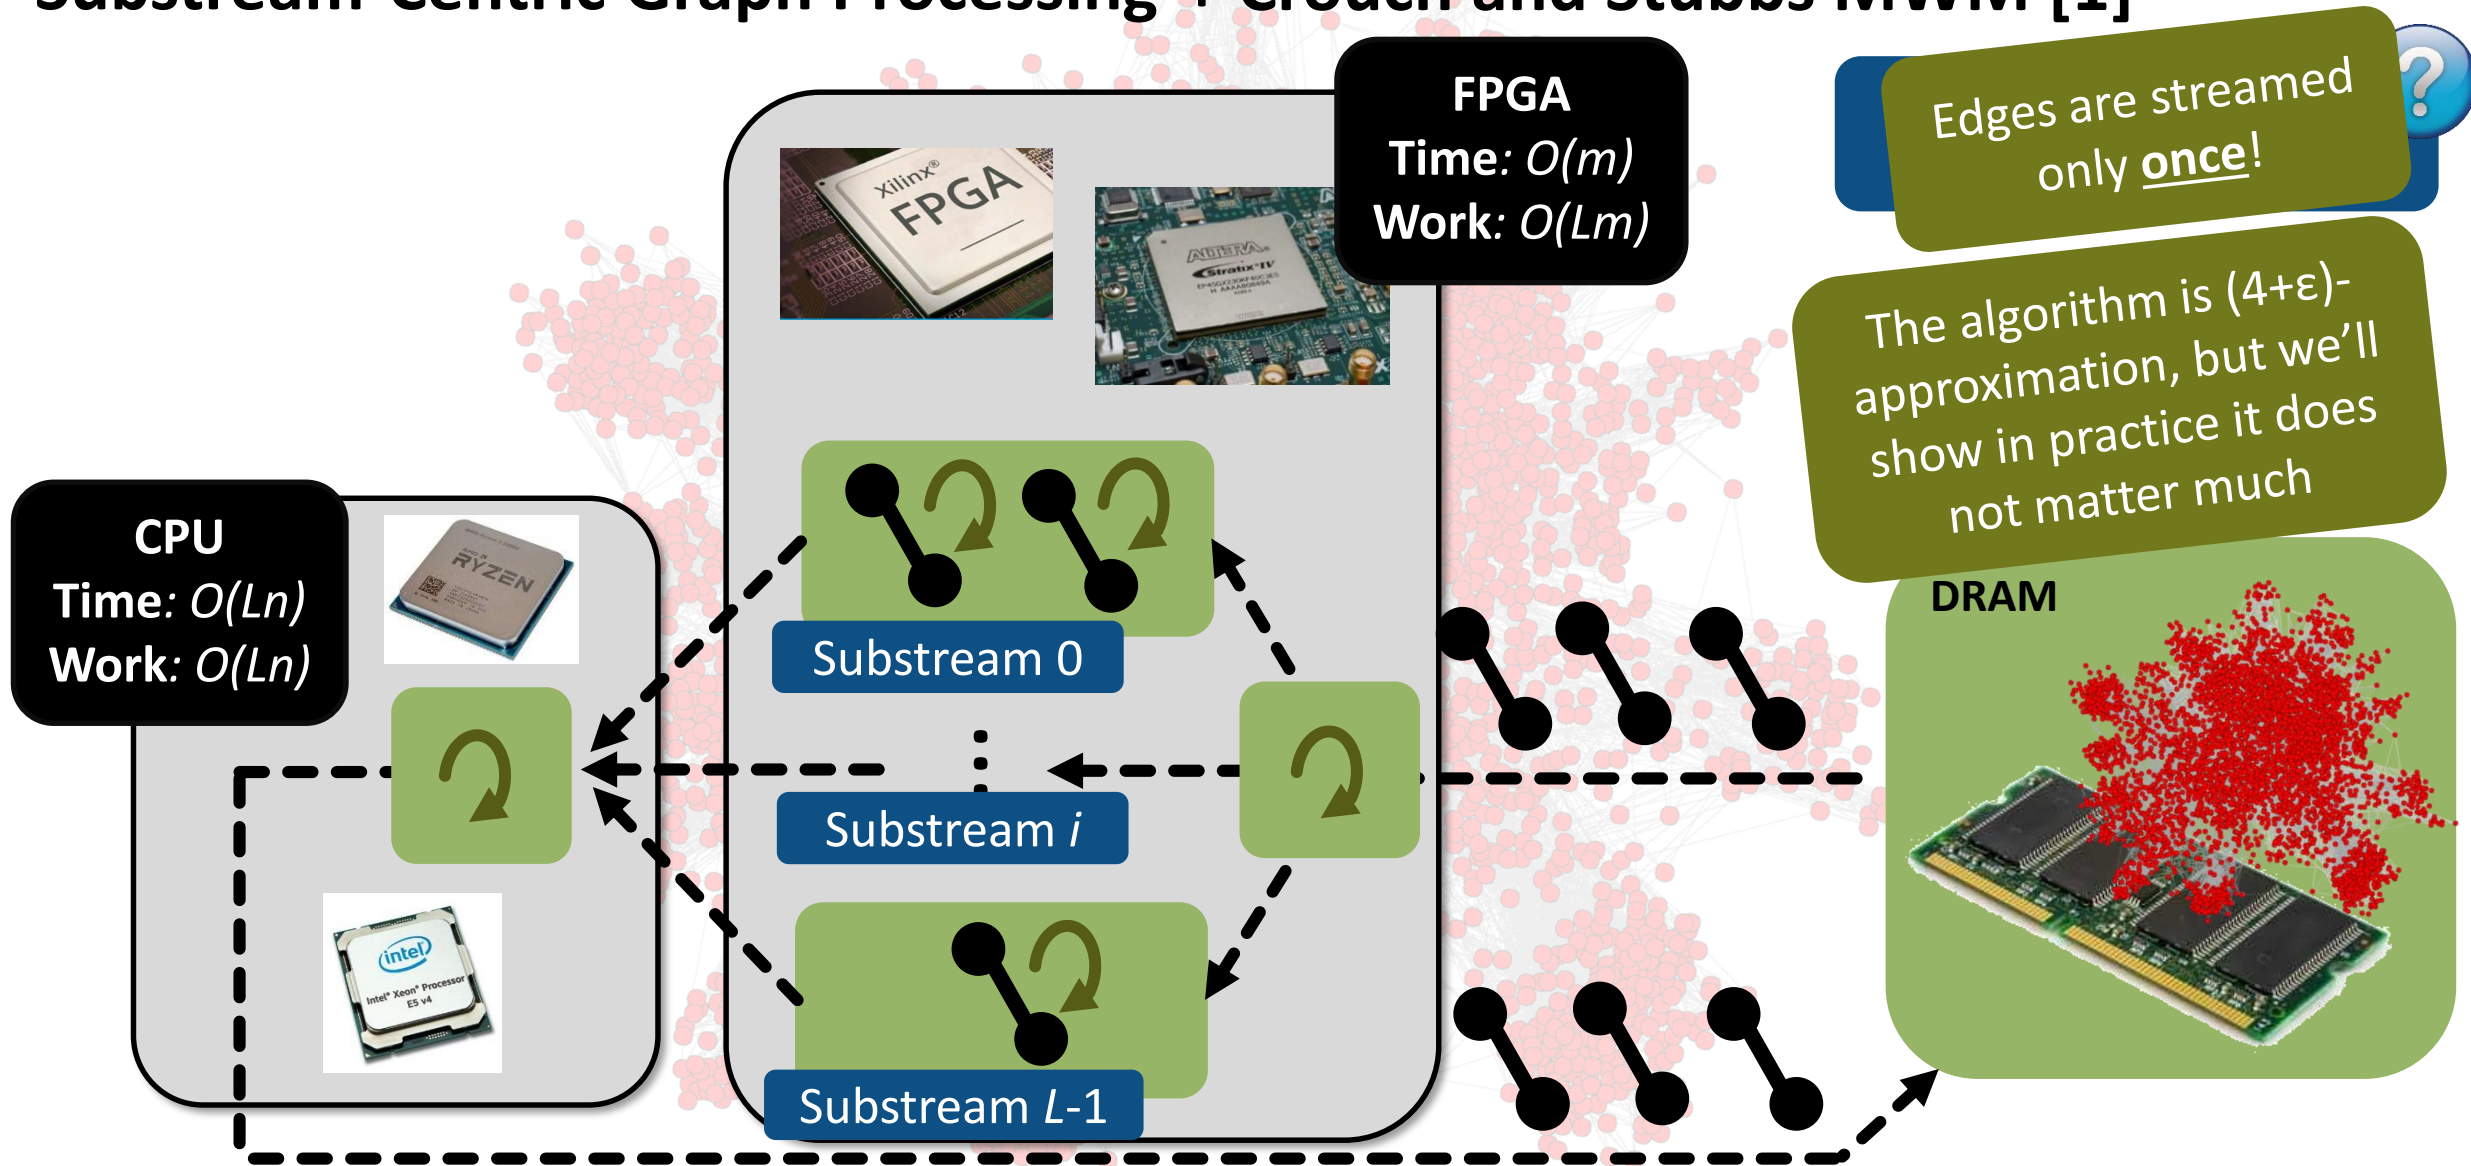
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

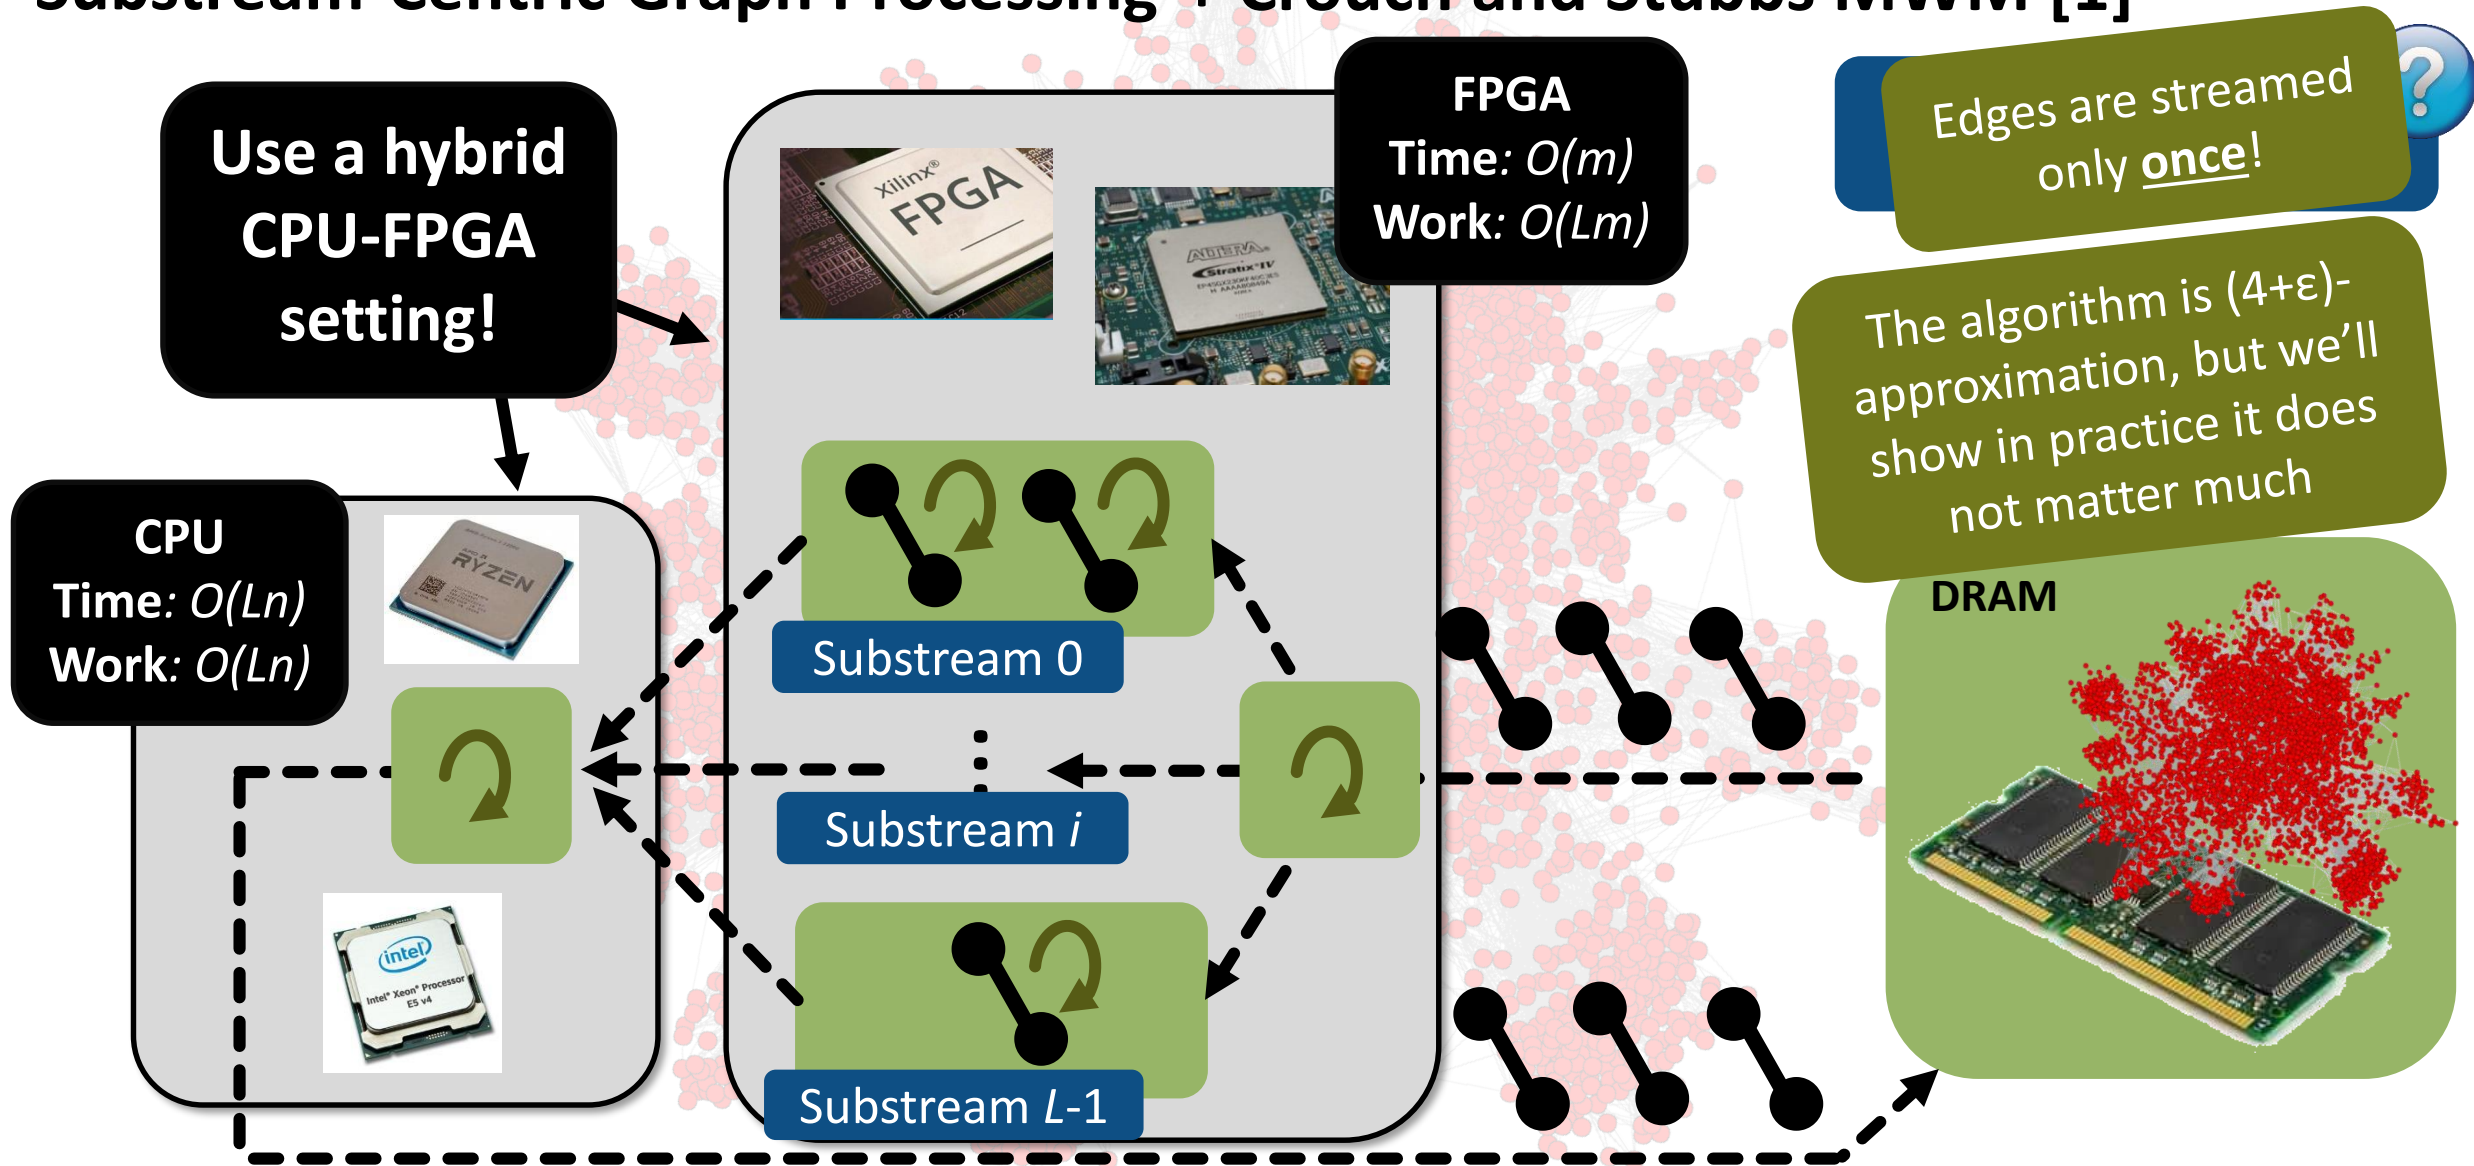


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

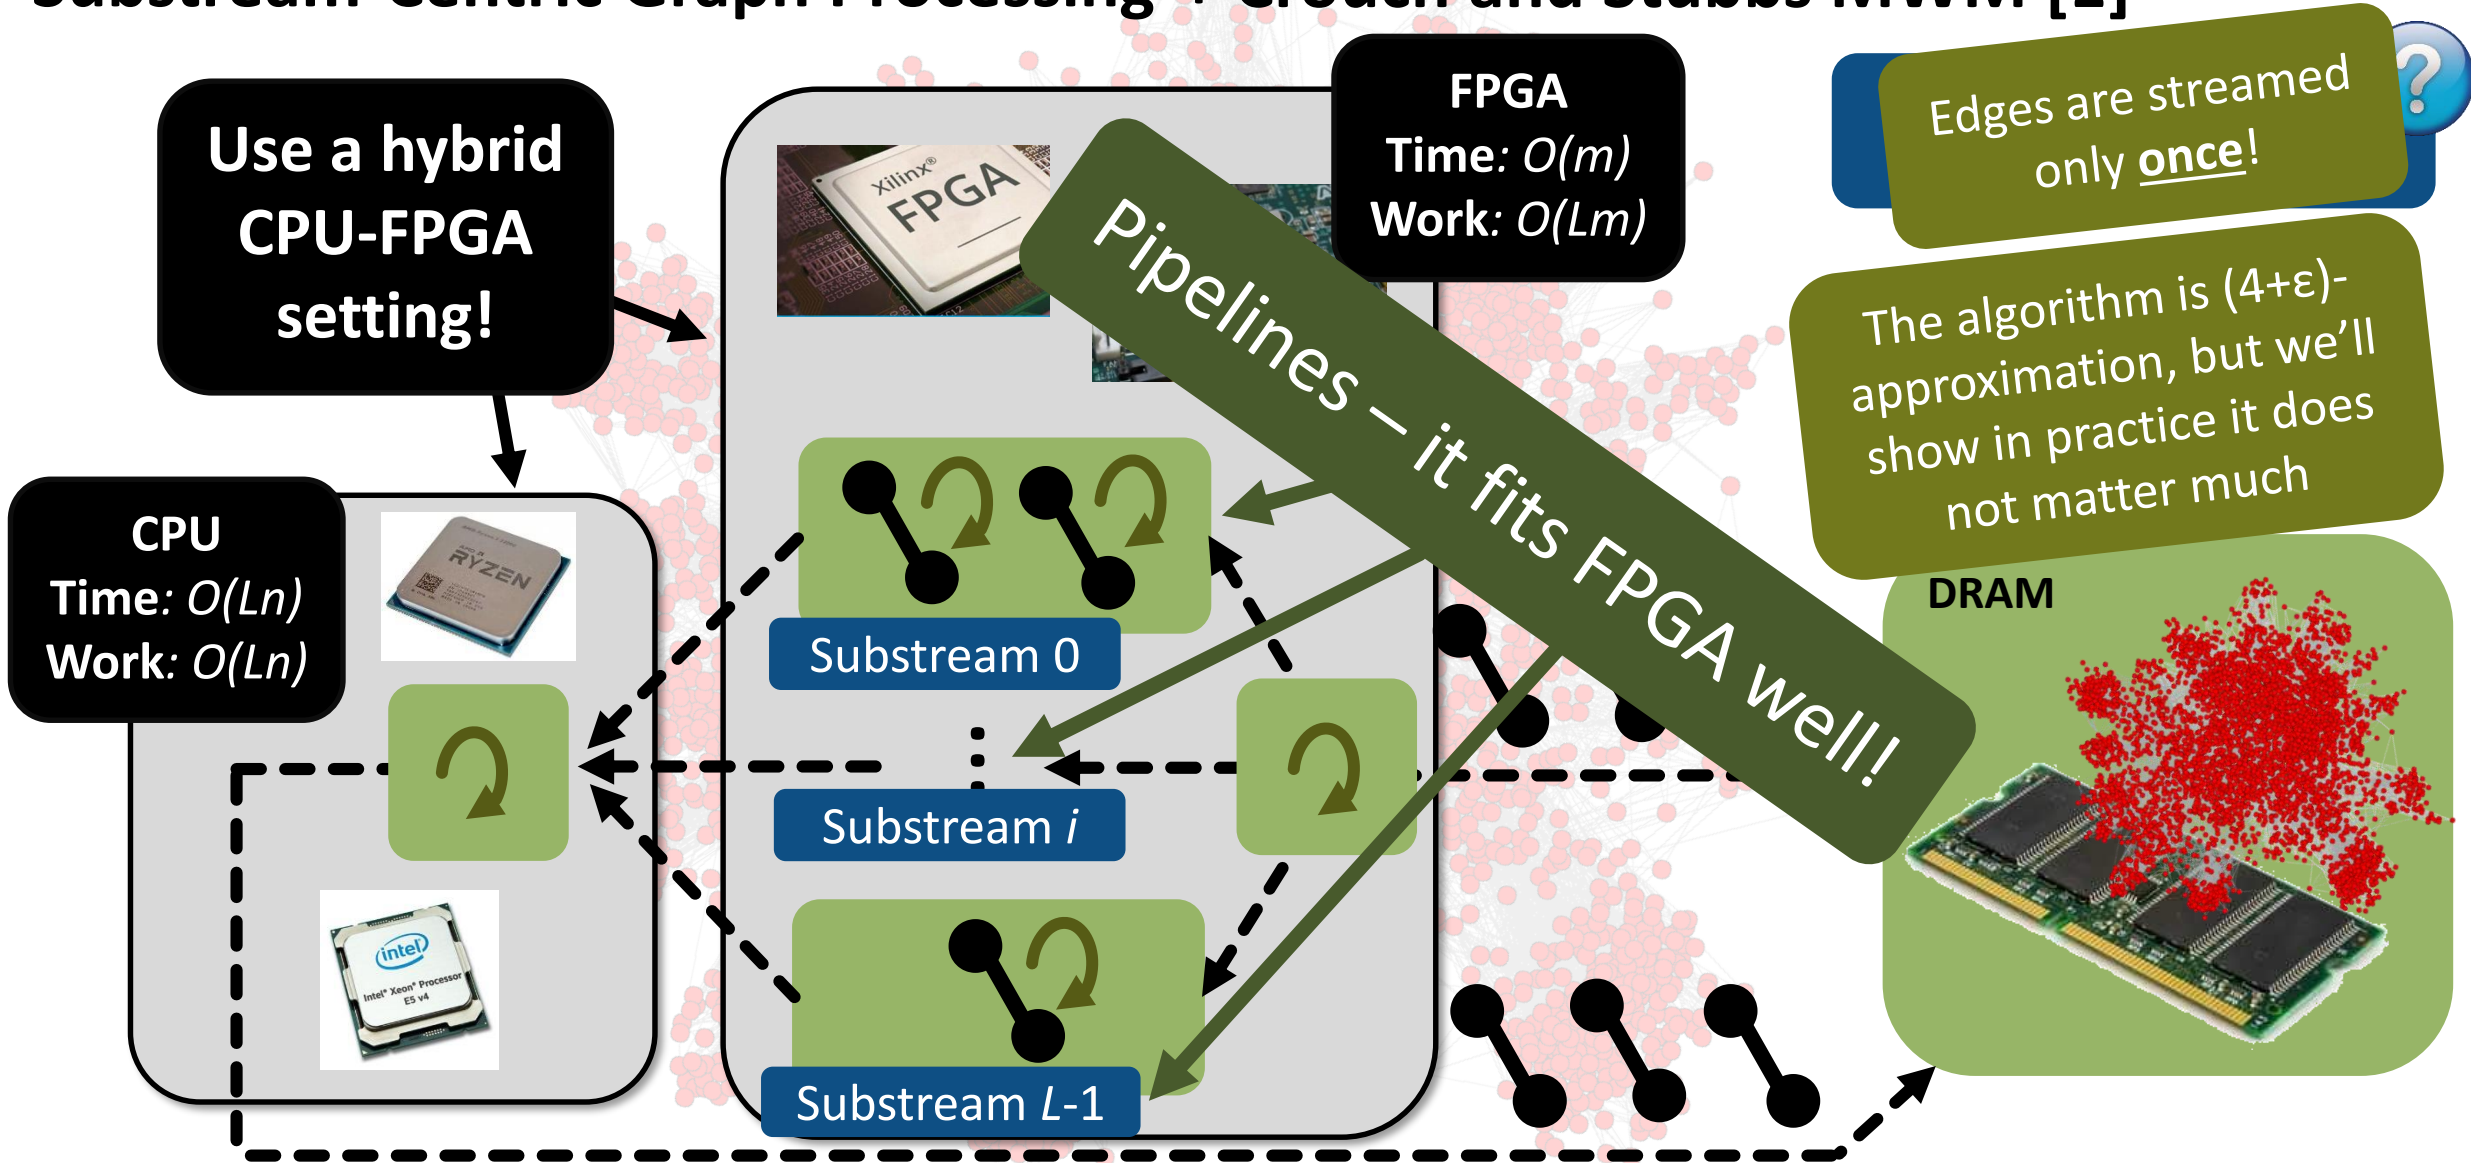


# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



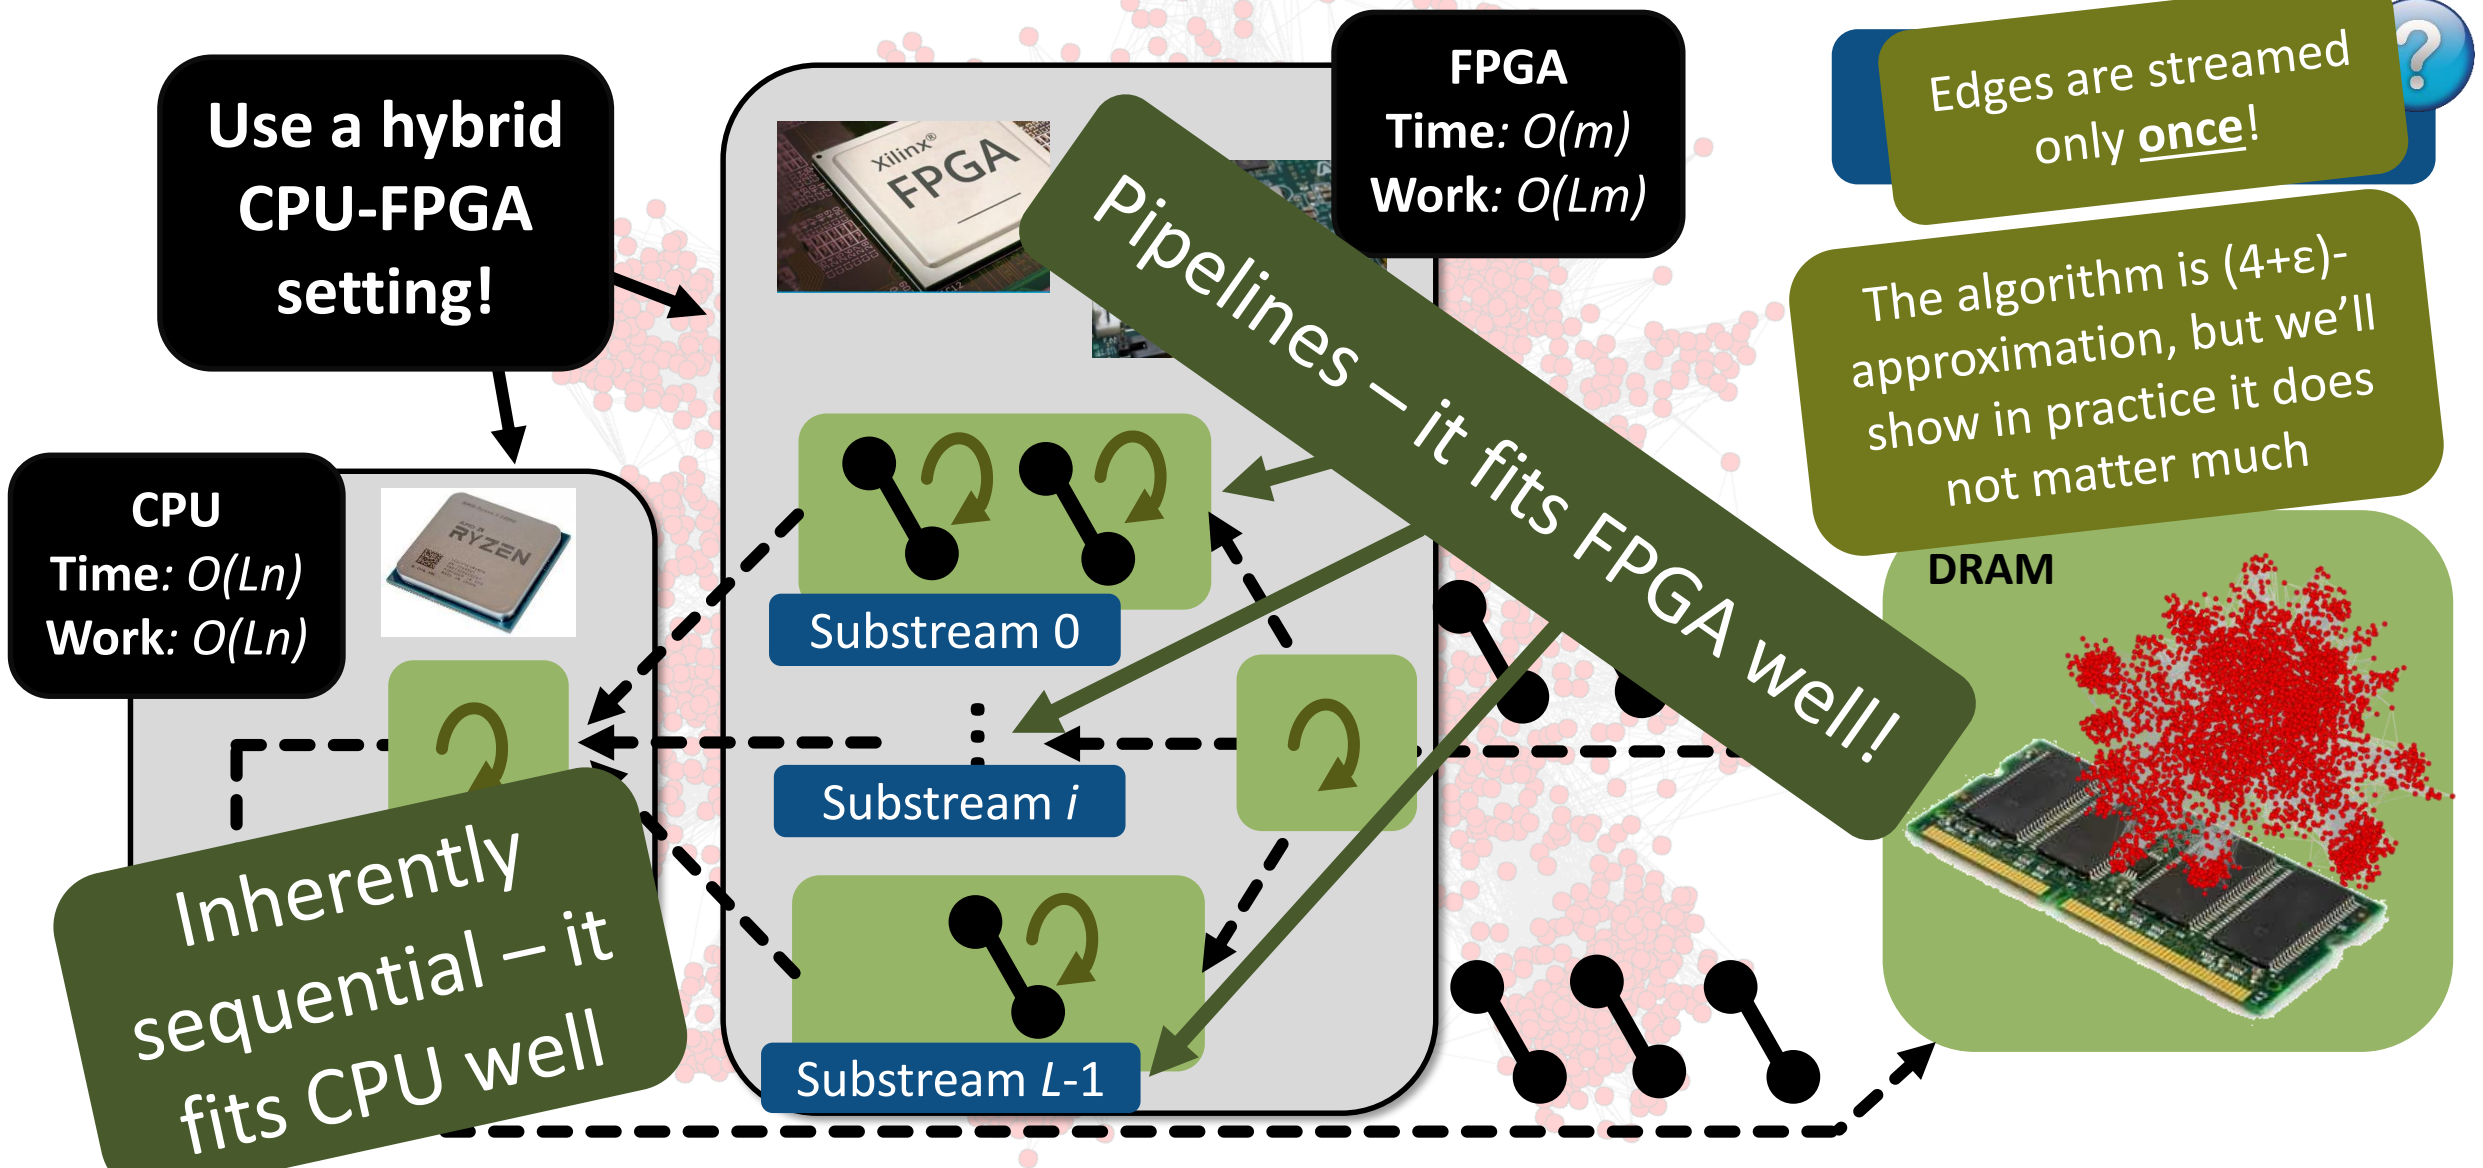
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]



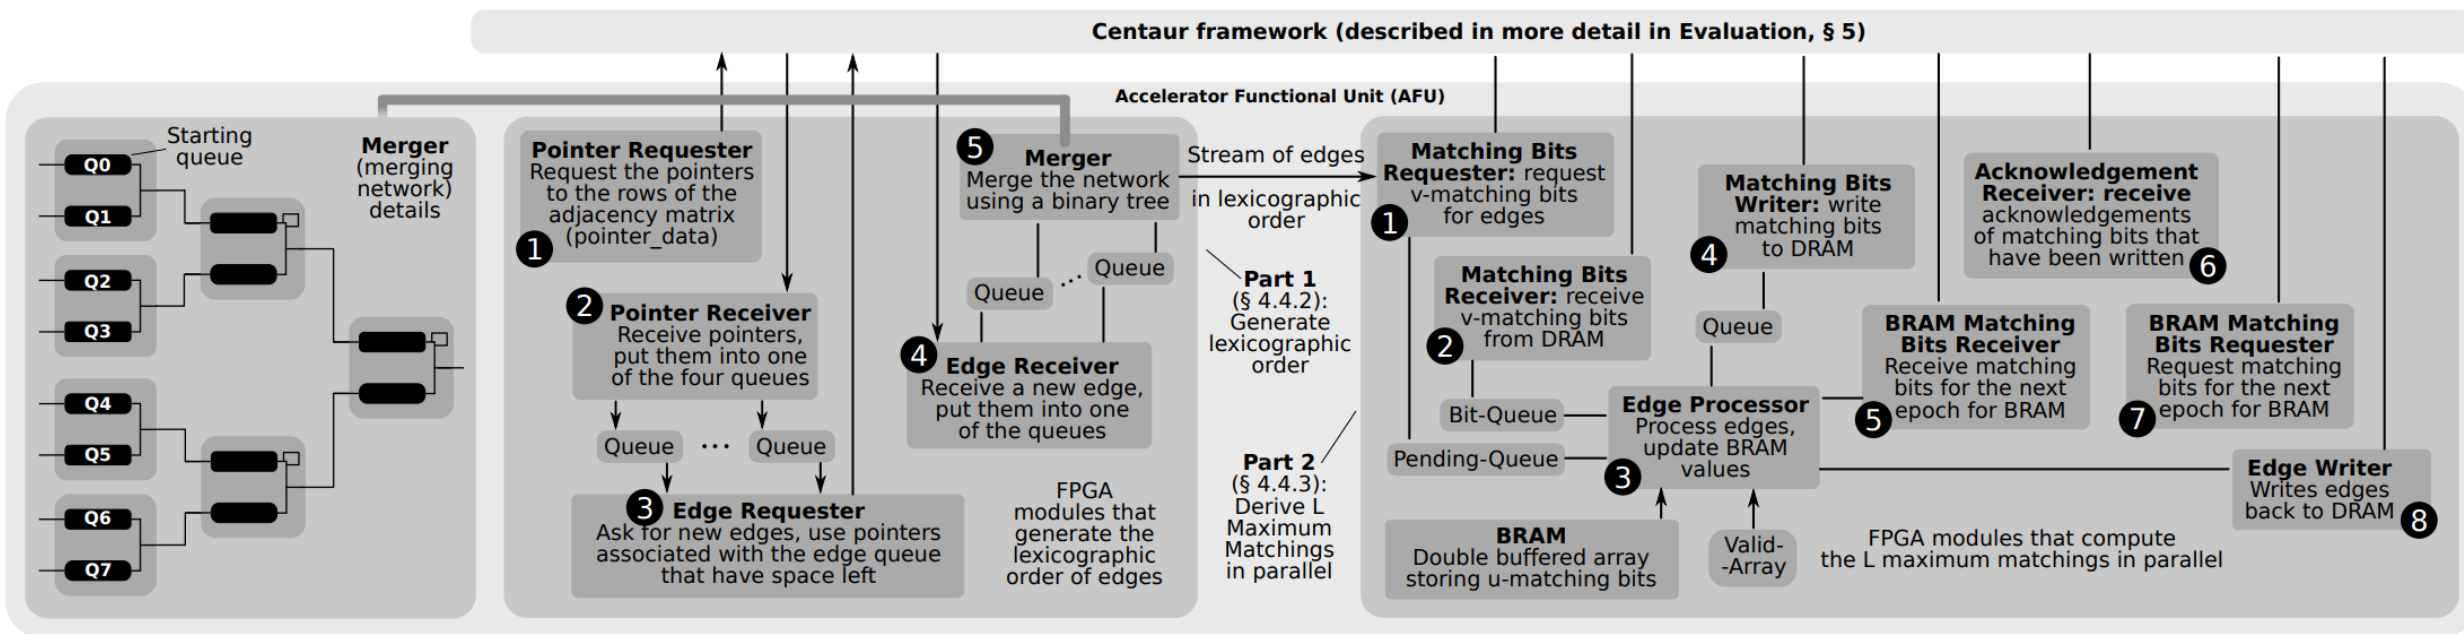
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

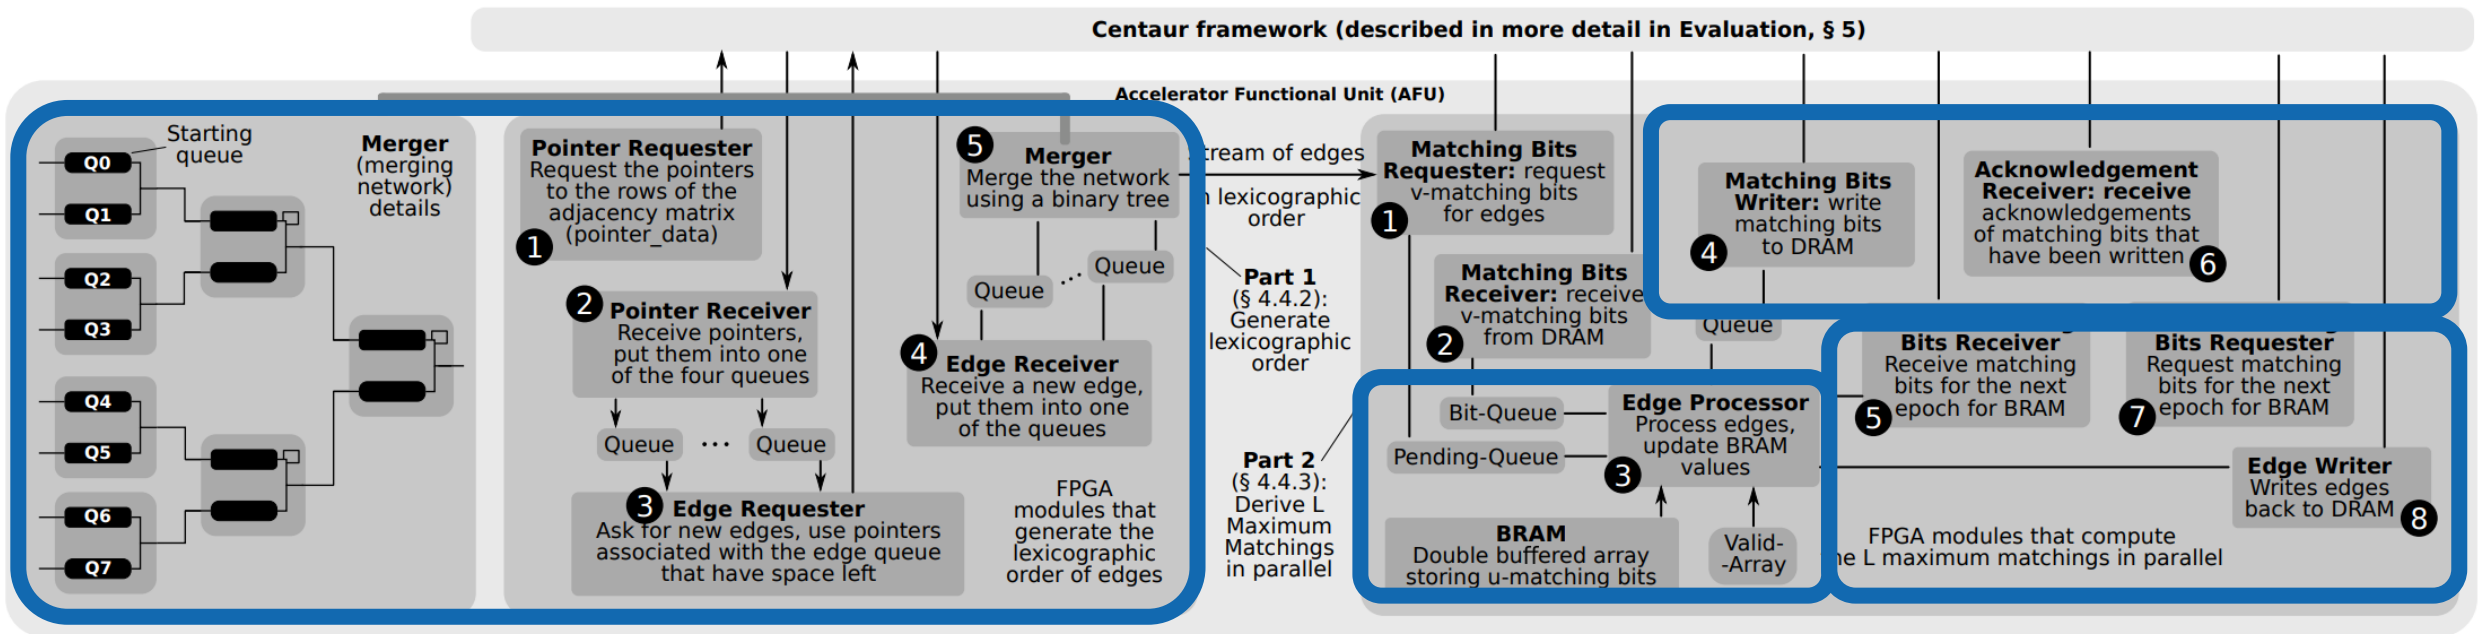


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

# Substream-Centric MWM: FPGA optimizations

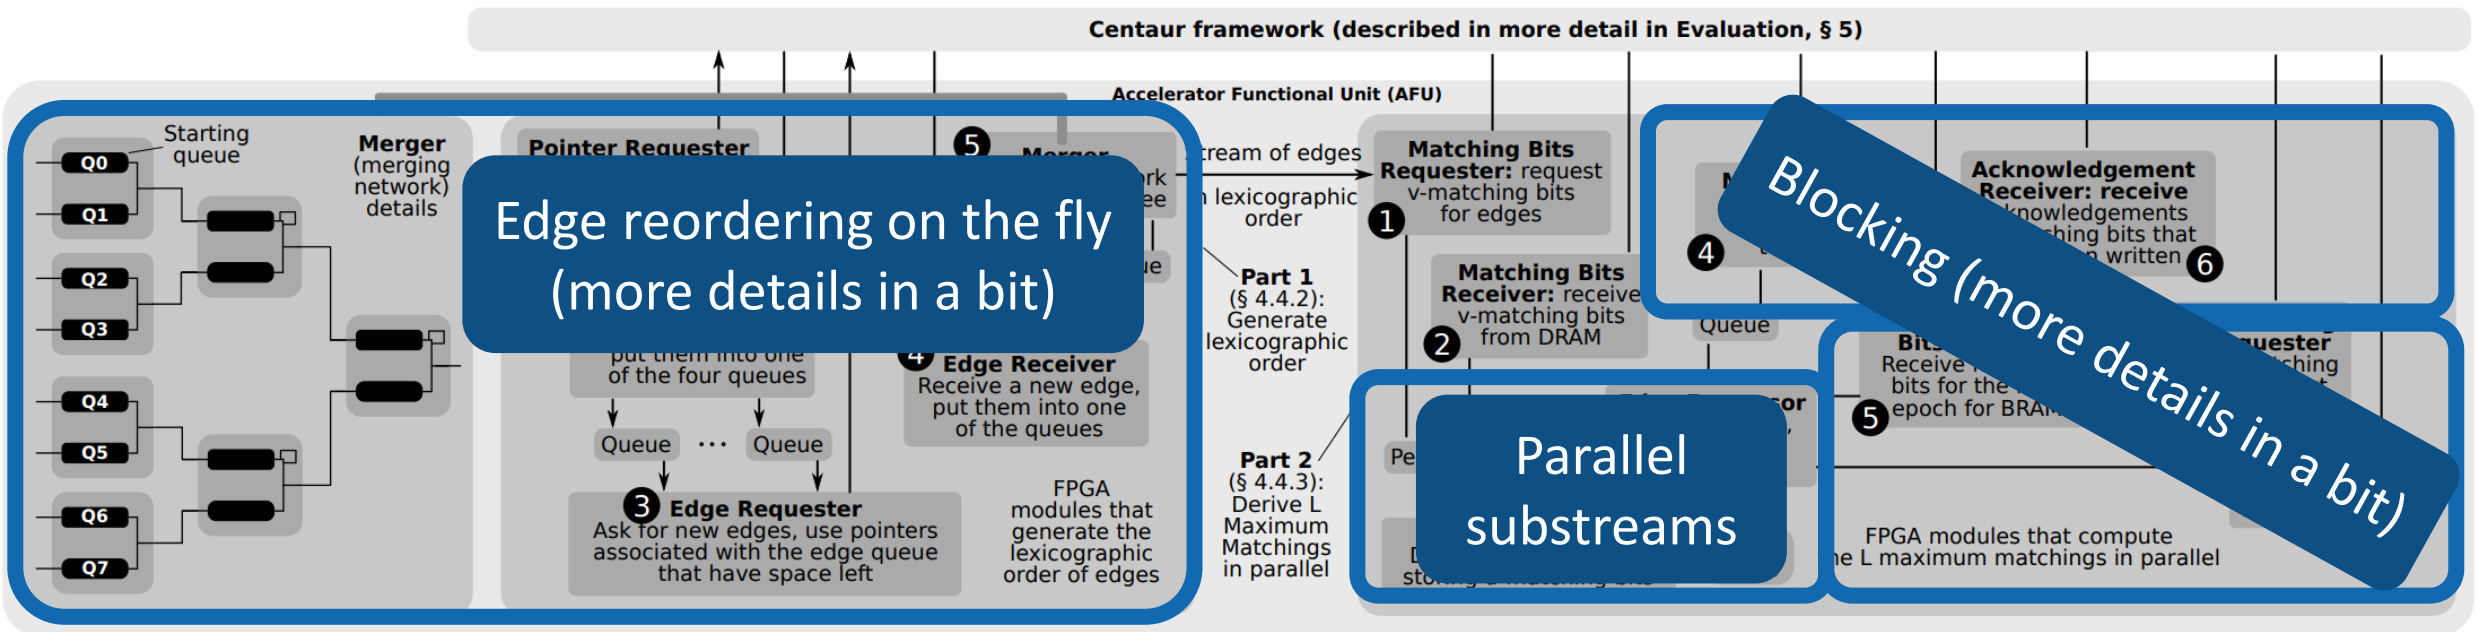


# Substream-Centric MWM: FPGA optimizations

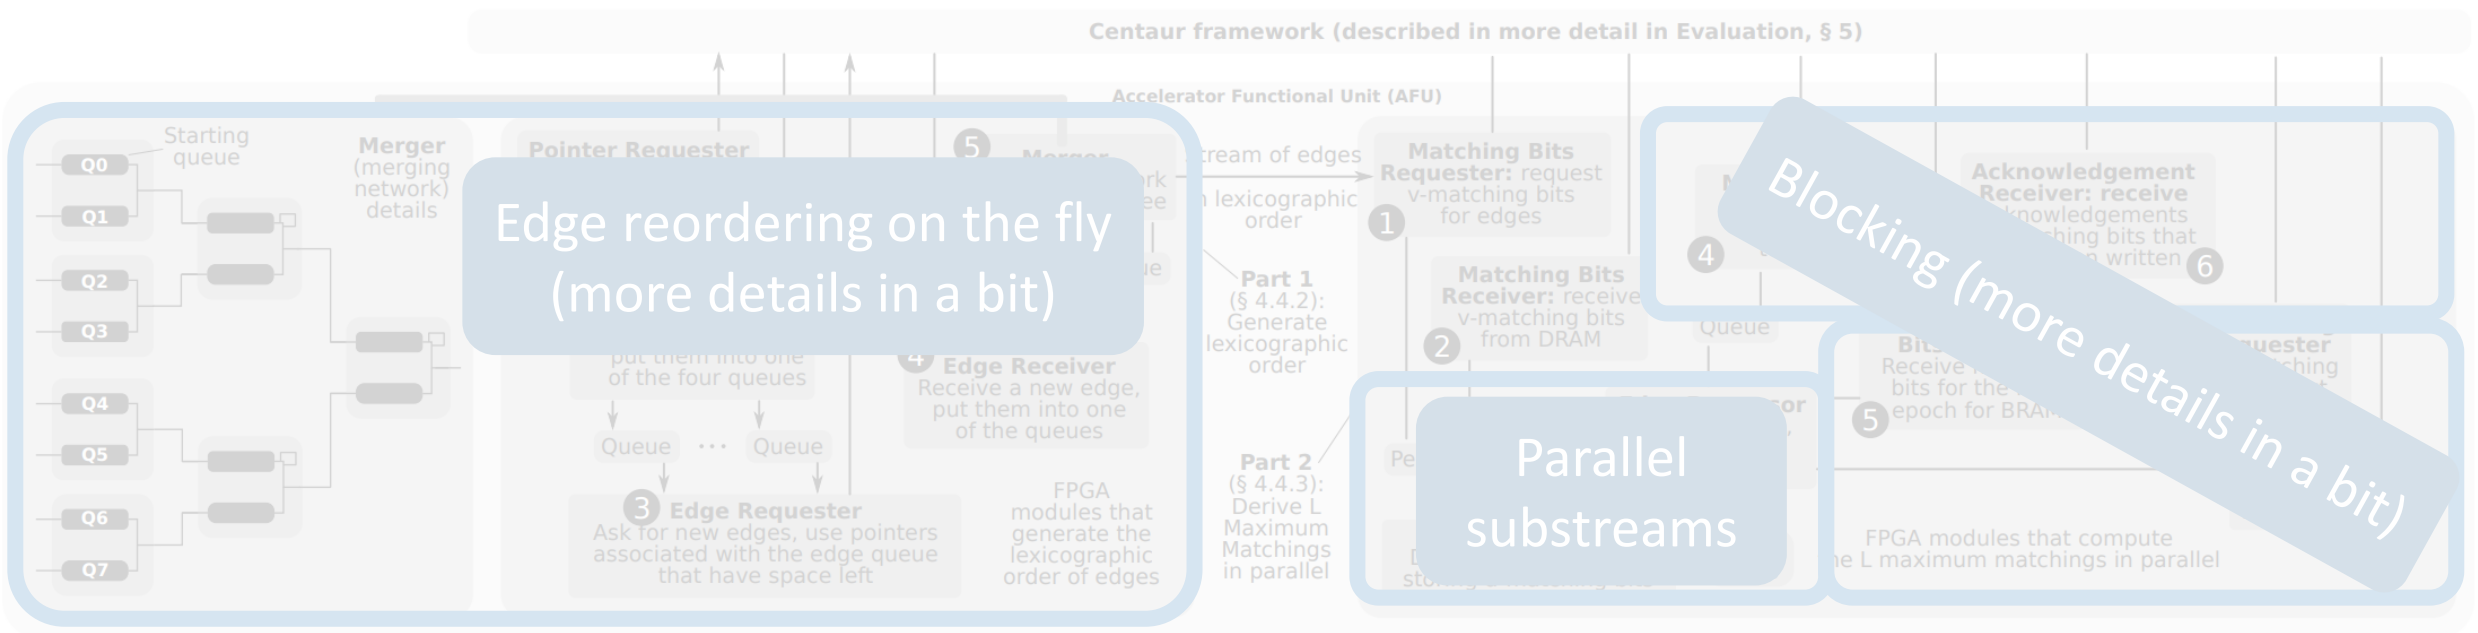




# Substream-Centric MWM: FPGA optimizations



# Substream-Centric MWM: FPGA optimizations




# Substream-Centric MWM: FPGA optimizations



# Substream-Centric MWM: FPGA optimizations

Blocking




All the details are in the paper. Let's focus on the key FPGA design ideas and optimizations

Prefetching

Pipelining

# Substream-Centric MWM: FPGA optimizations

Blocking



Vectorization

All the details are in the paper.  
Let's focus on the key FPGA  
design ideas and optimizations

details in a bit)


Prefetching

They are often used in graph processing schemes on FPGAs; we apply them as well.

Pipelining

# Substream-Centric MWM: FPGA optimizations

Blocking



All the details are in the paper. Let's focus on the key FPGA design ideas and optimizations

Vectorization

Prefetching

They are often used in graph processing schemes on FPGAs; we apply them as well.

Pipelining

# Substream-Centric MWM: FPGA optimizations

## Blocking

# Substream-Centric MWM: FPGA optimizations

## Blocking

$$\begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 1 \\ & 0 & 1 & 0 & 1 & 0 \\ & & 0 & 1 & 0 & 1 \\ & & & 0 & 1 & 0 \\ & & & & 0 & 1 \\ & & & & & 0 \end{bmatrix}$$

Adjacency Matrix



# Substream-Centric MWM: FPGA optimizations

## Blocking

	0	1	2	3	4	5
0	0	1	1	1	0	1
1		0	1	0	1	0
2			0	1	0	1
3				0	1	0
4					0	1
5						0

Adjacency Matrix

# Substream-Centric MWM: FPGA optimizations

## Blocking

Column IDs  
correspond  
to vertex IDs

	0	1	2	3	4	5
0	0	1	1	1	0	1
1		0	1	0	1	0
2			0	1	0	1
3				0	1	0
4					0	1
5						0

Row IDs  
correspond  
to vertex IDs

Adjacency Matrix

# Substream-Centric MWM: FPGA optimizations

## Blocking

An edge between vertices 0 and 1

Column IDs correspond to vertex IDs

	0	1	2	3	4	5
0	0	1	1	1	0	1
1		0	1	0	1	0
2			0	1	0	1
3				0	1	0
4					0	1
5						0

Row IDs correspond to vertex IDs

Adjacency Matrix

# Substream-Centric MWM: FPGA optimizations

## Blocking

An edge between vertices 0 and 1

Column IDs correspond to vertex IDs

	0	1	2	3	4	5
0	0	1	1	1	0	1
1		0	1	0	1	0
2			0	1	0	1
3				0	1	0
4					0	1
5						0

Row IDs correspond to vertex IDs

Adjacency Matrix

# Substream-Centric MWM: FPGA optimizations

## Blocking

An edge between vertices 0 and 1

Column IDs correspond to vertex IDs


	0	1	2	3	4	5
0	0	1	1	1	0	1
1		0	1	0	1	0
2			0	1	0	1
3				0	1	0
4					0	1
5						0

Row IDs correspond to vertex IDs

Adjacency Matrix

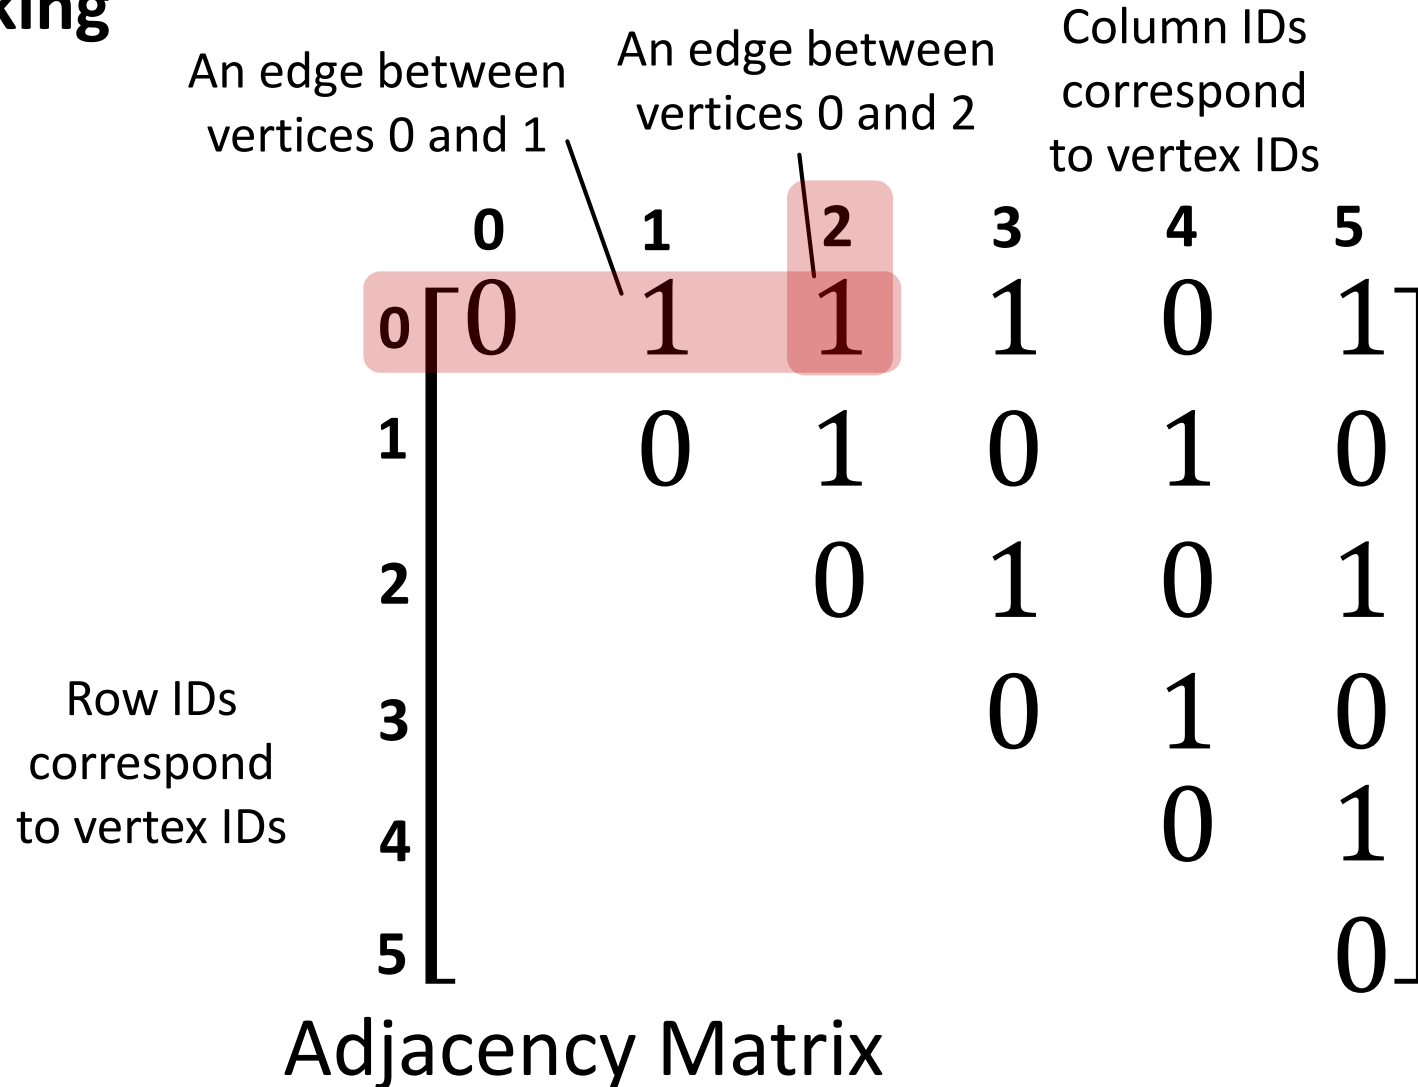
# Substream-Centric MWM: FPGA optimizations

## Blocking



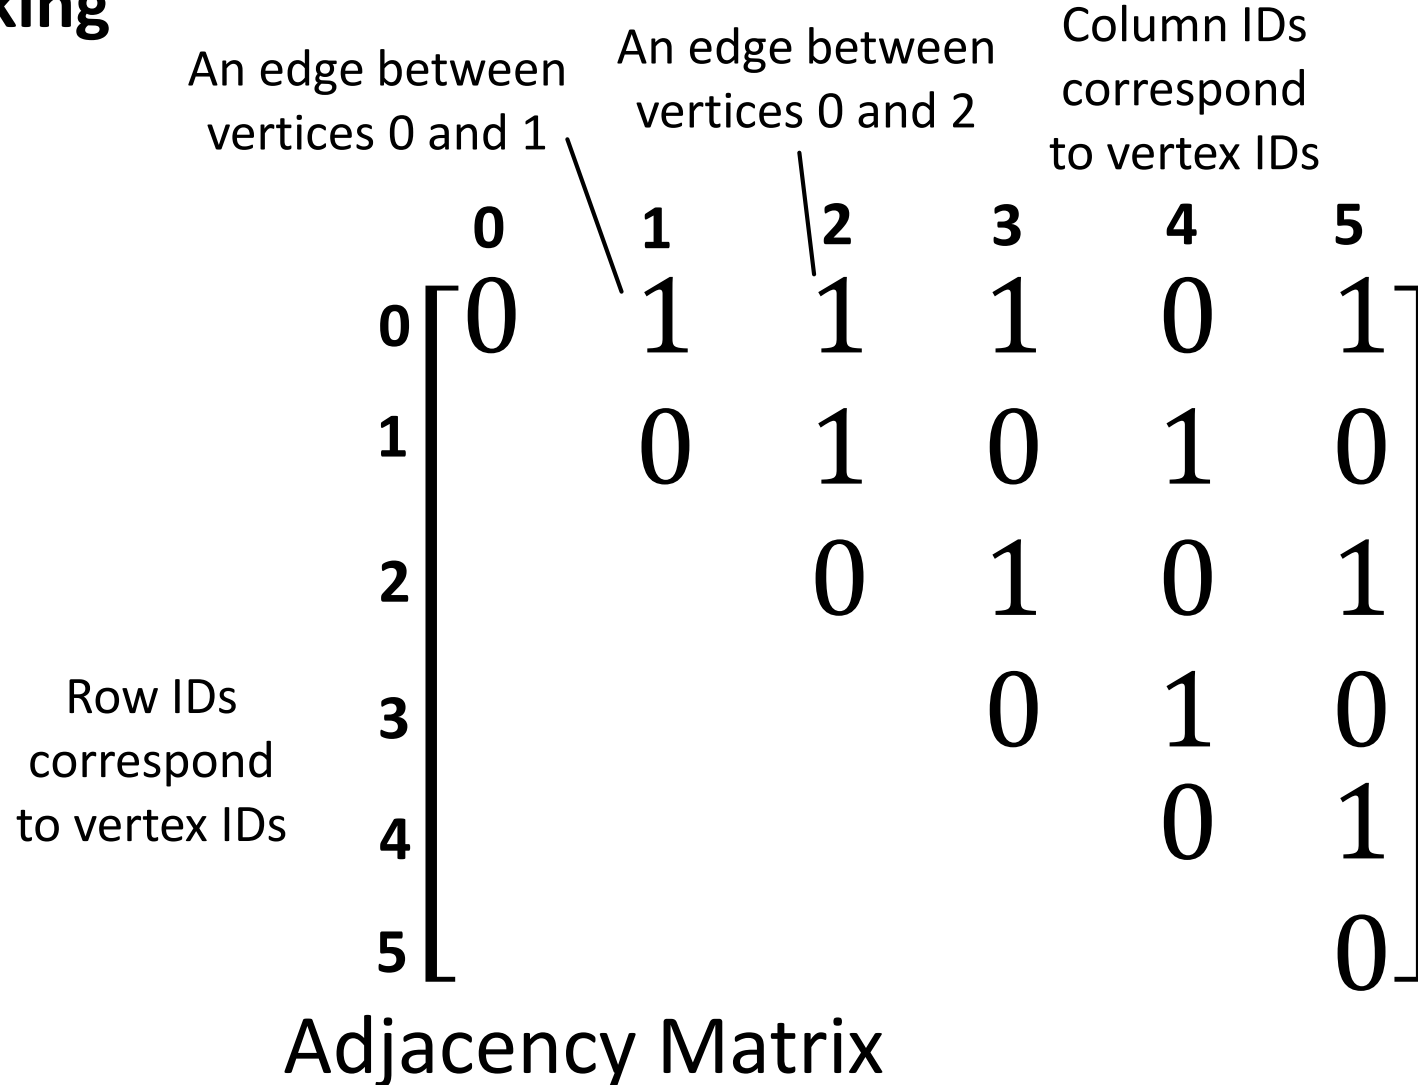
# Substream-Centric MWM: FPGA optimizations

## Blocking



# Substream-Centric MWM: FPGA optimizations

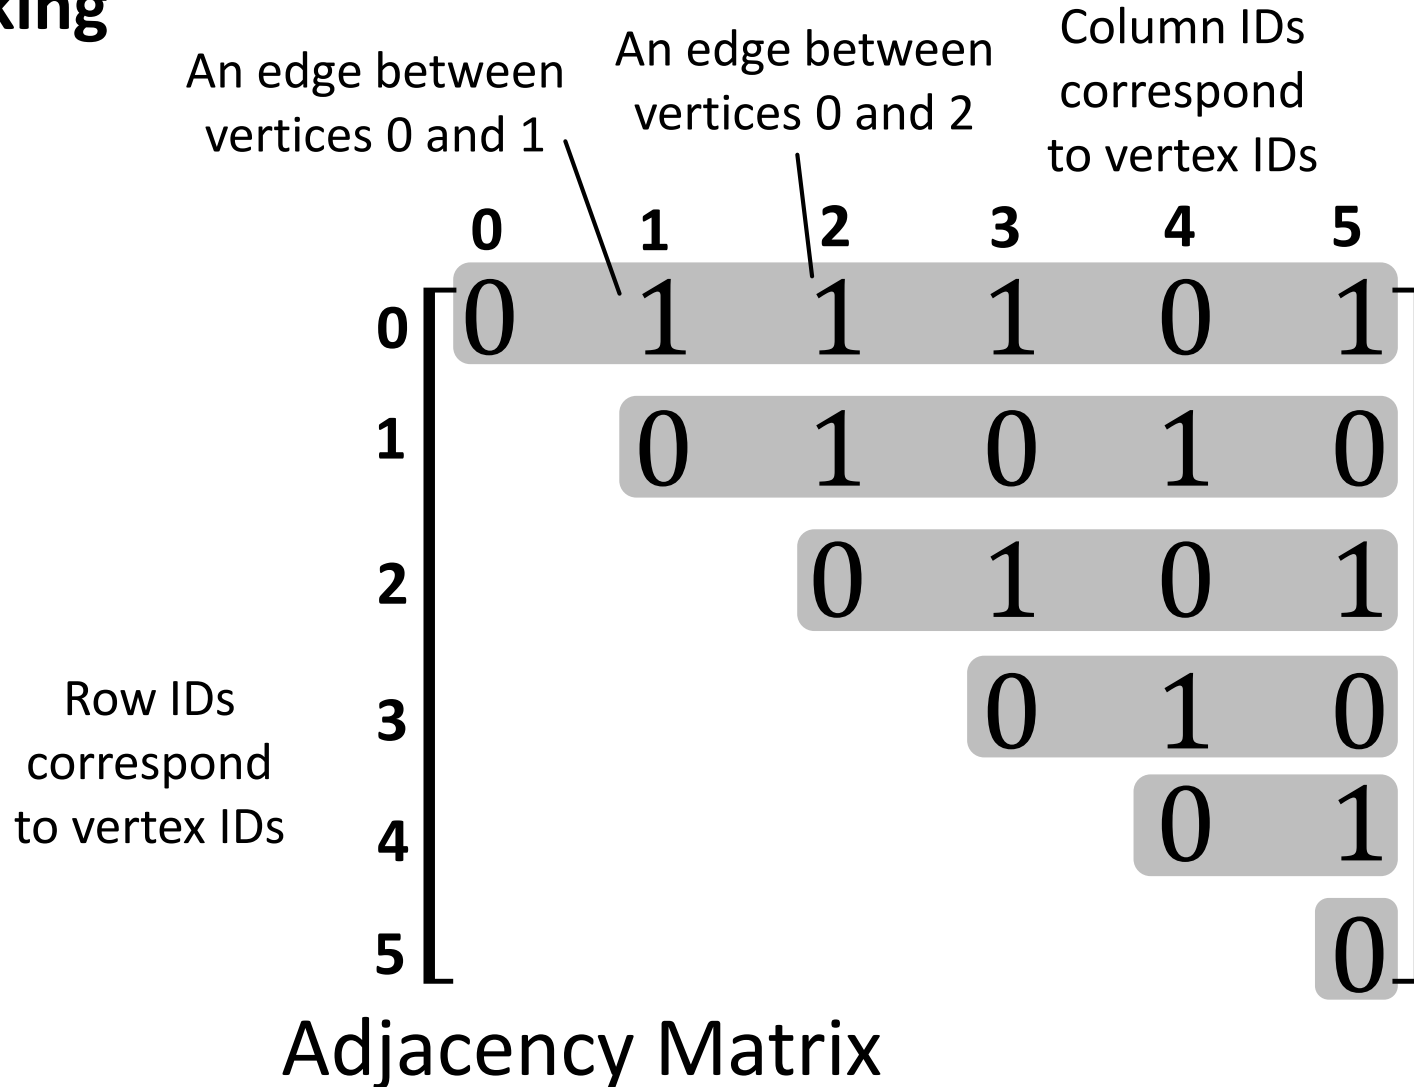
## Blocking





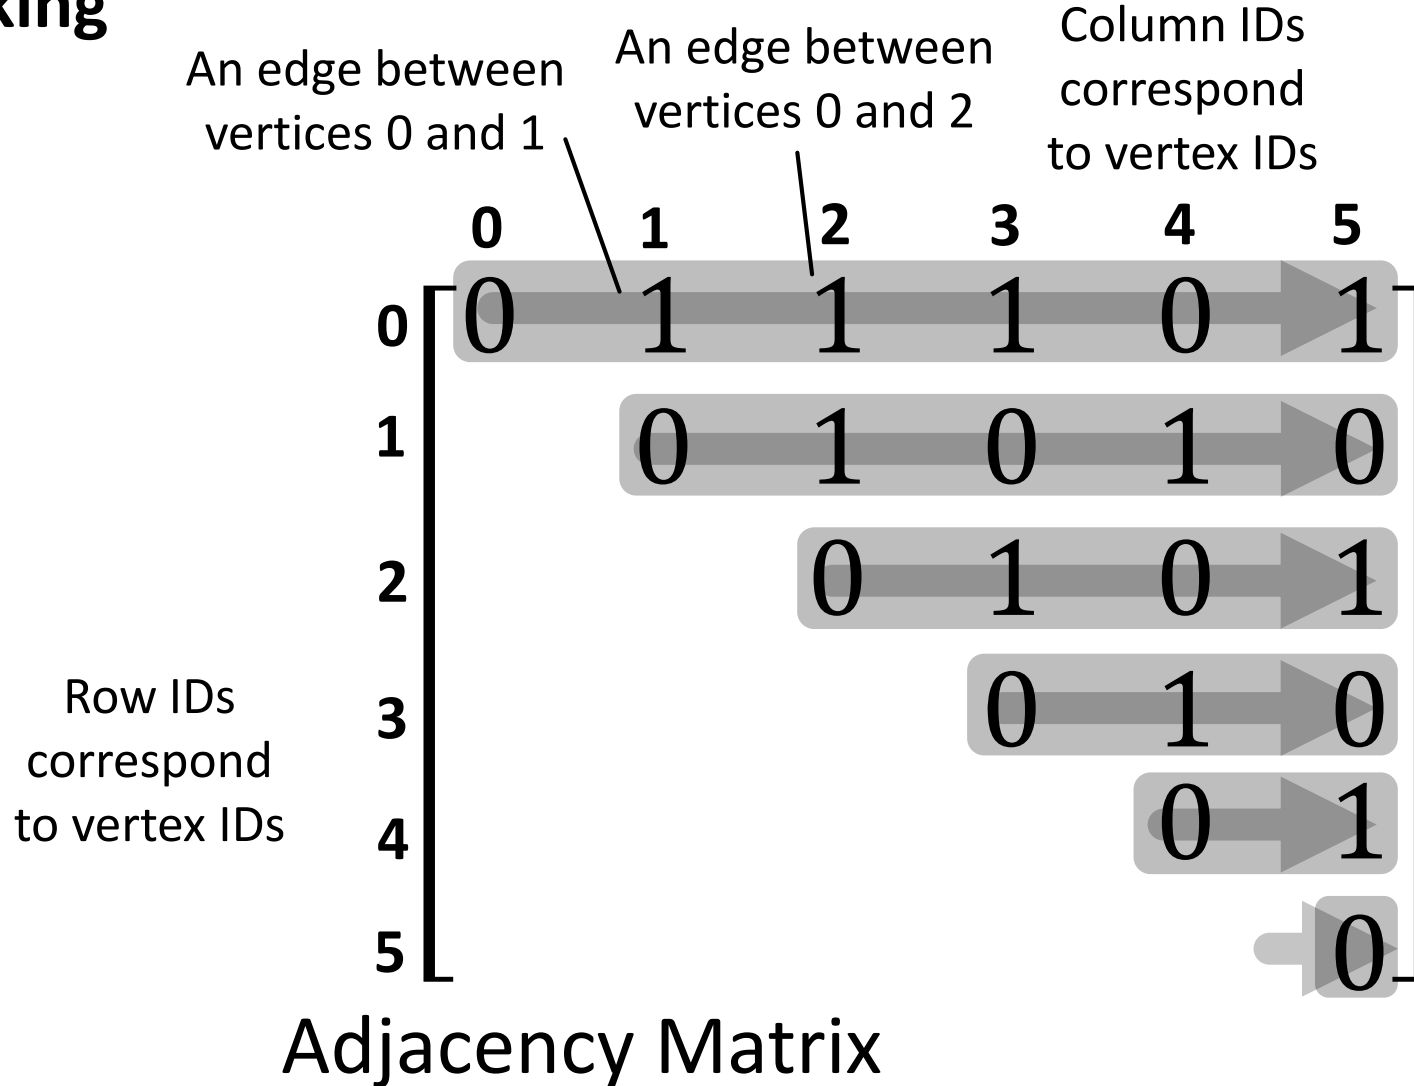
# Substream-Centric MWM: FPGA optimizations

## Blocking



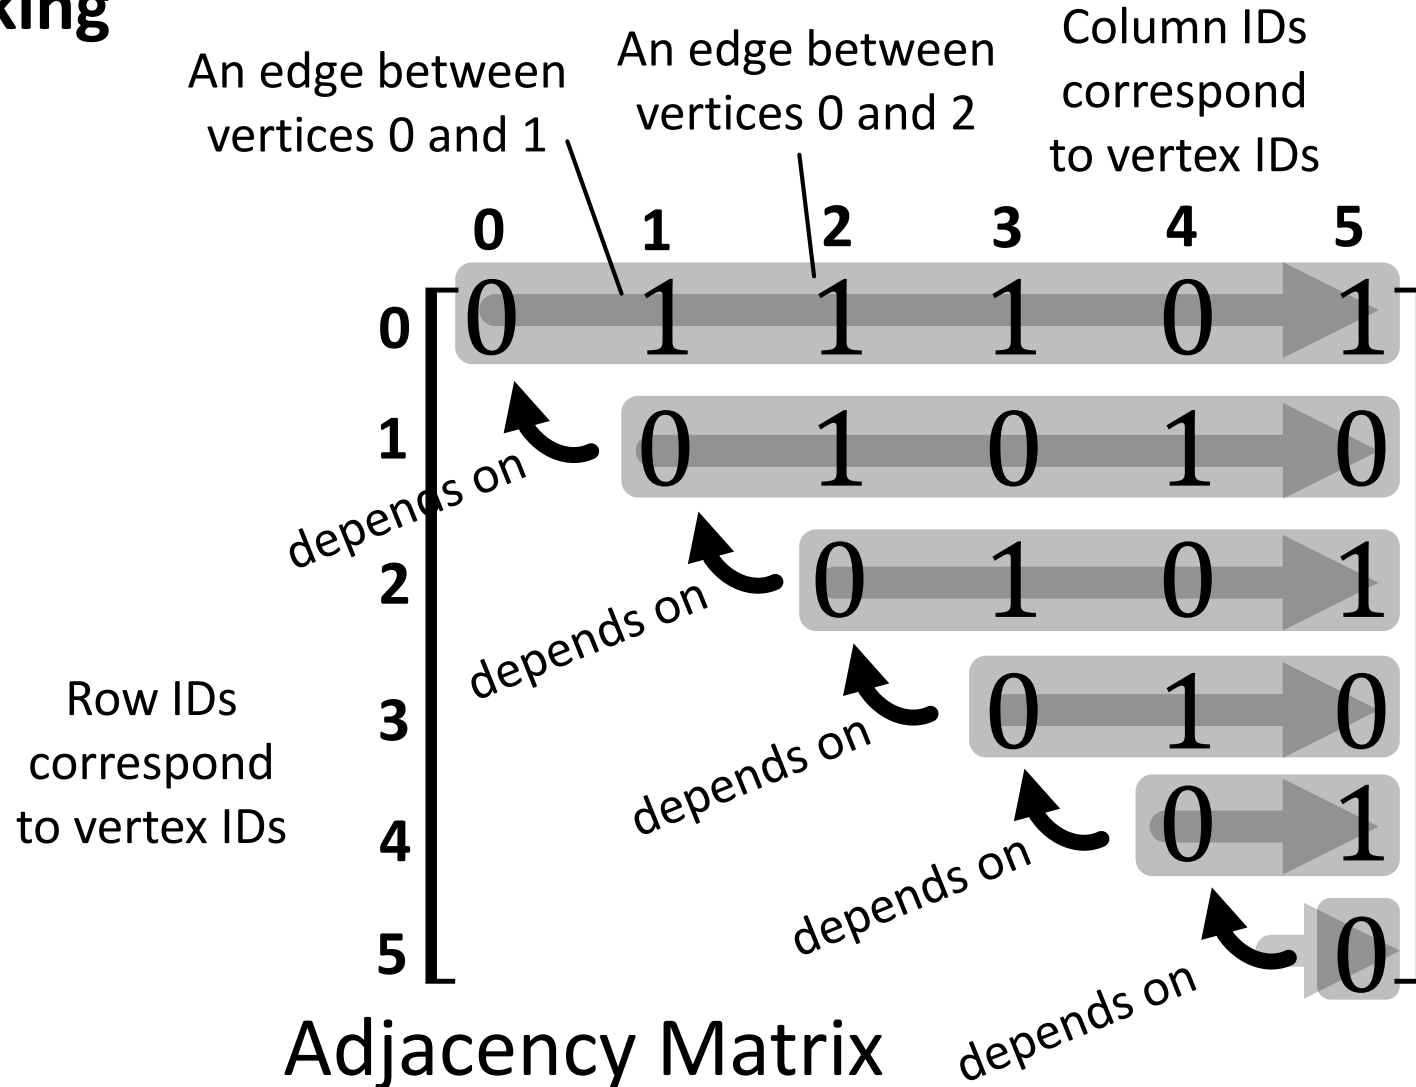
# Substream-Centric MWM: FPGA optimizations

## Blocking



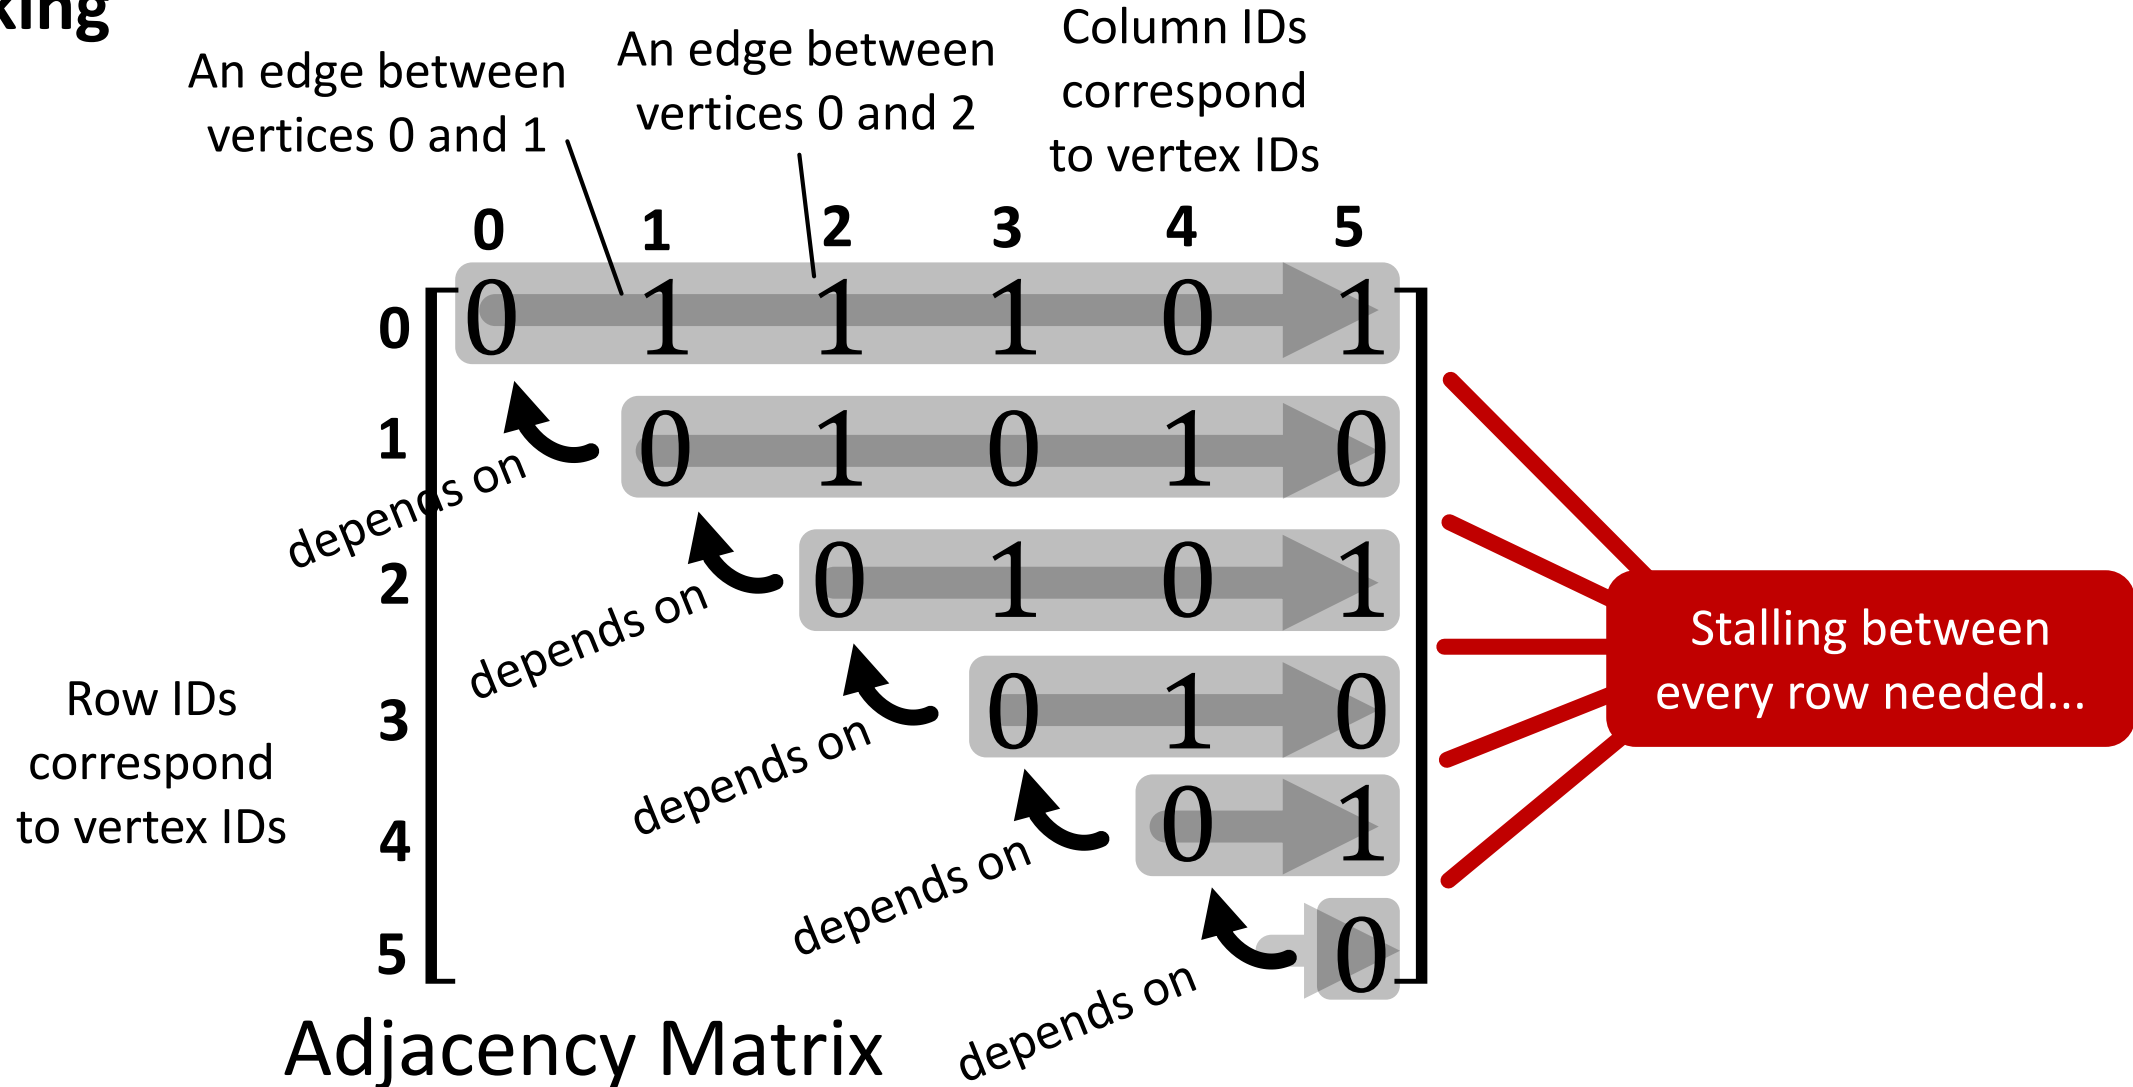
# Substream-Centric MWM: FPGA optimizations

## Blocking



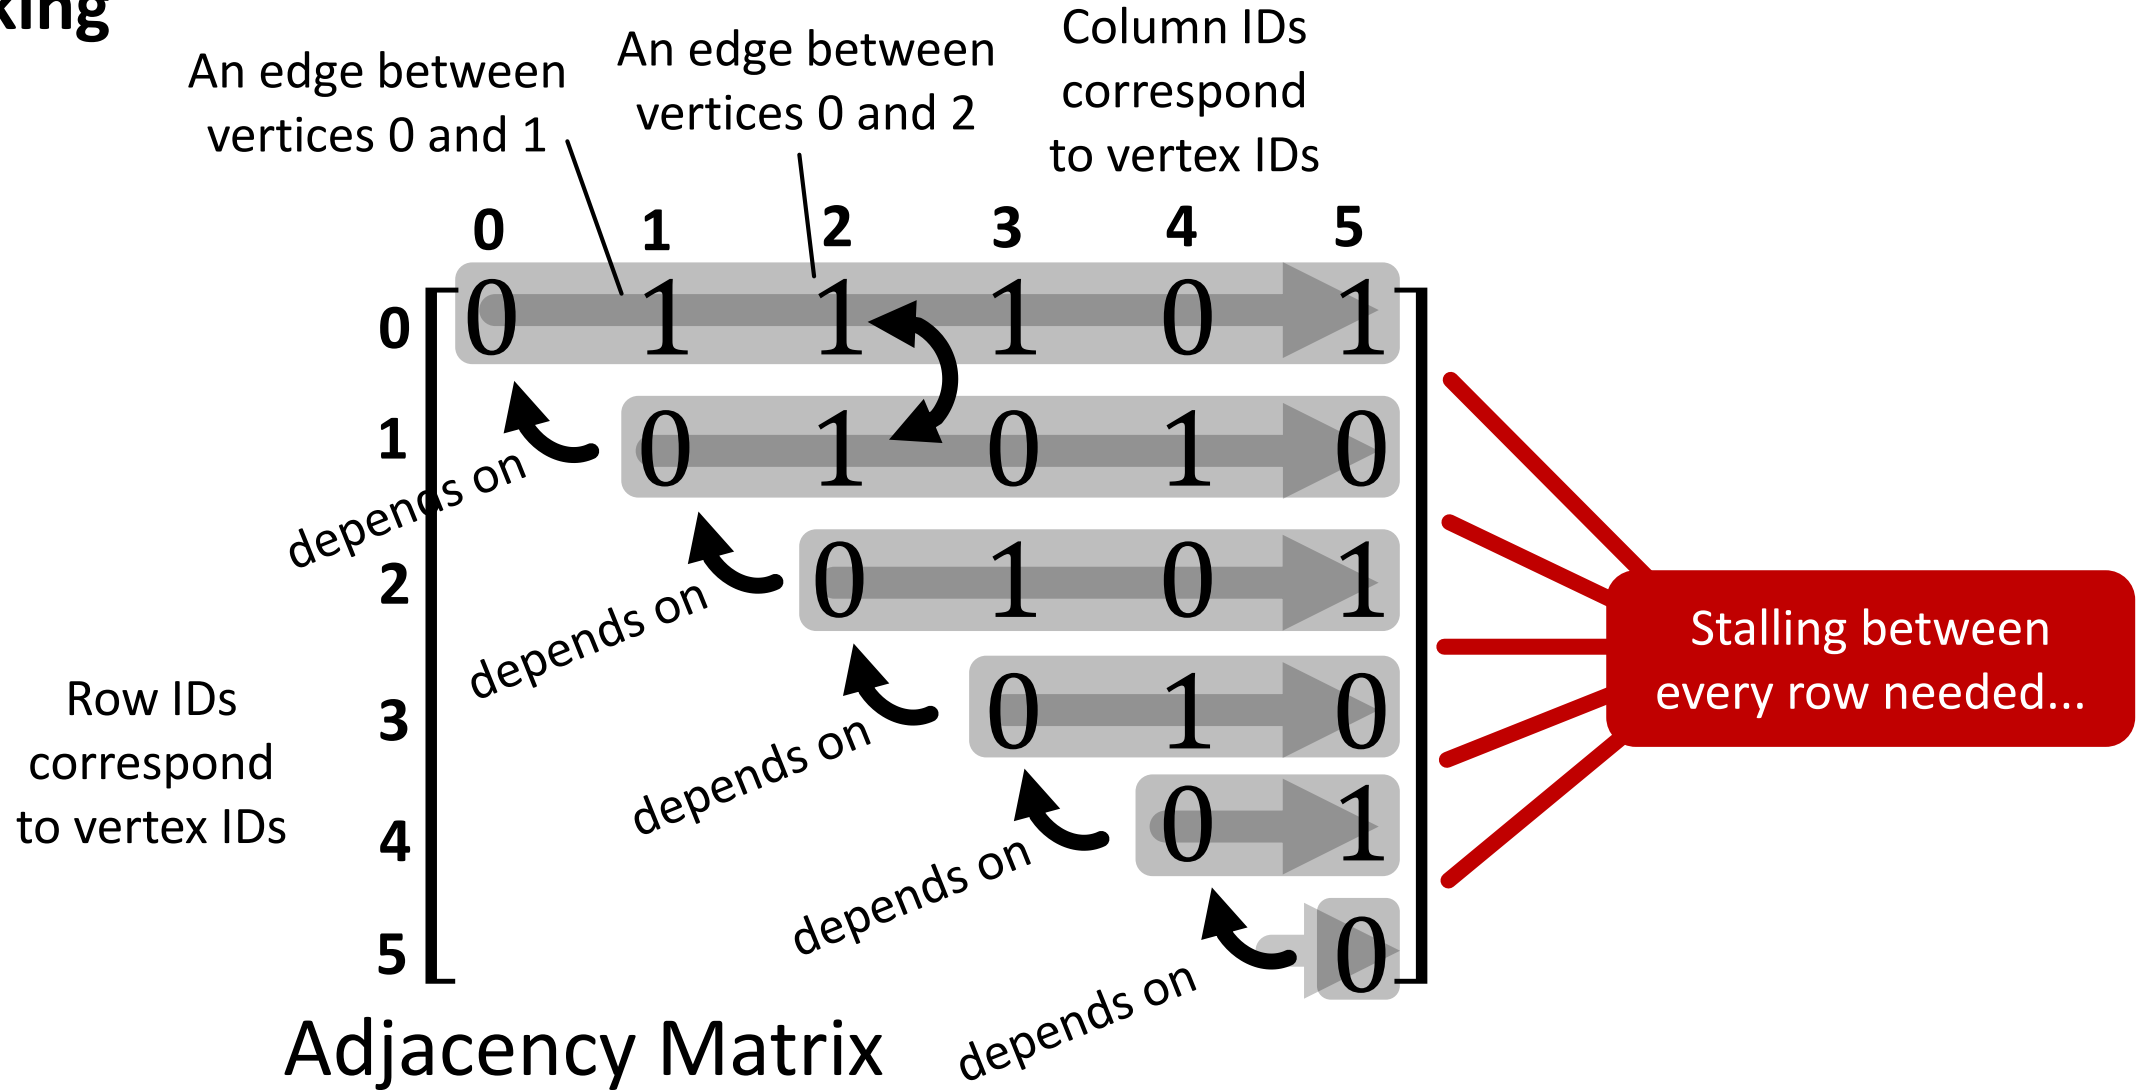
# Substream-Centric MWM: FPGA optimizations

## Blocking



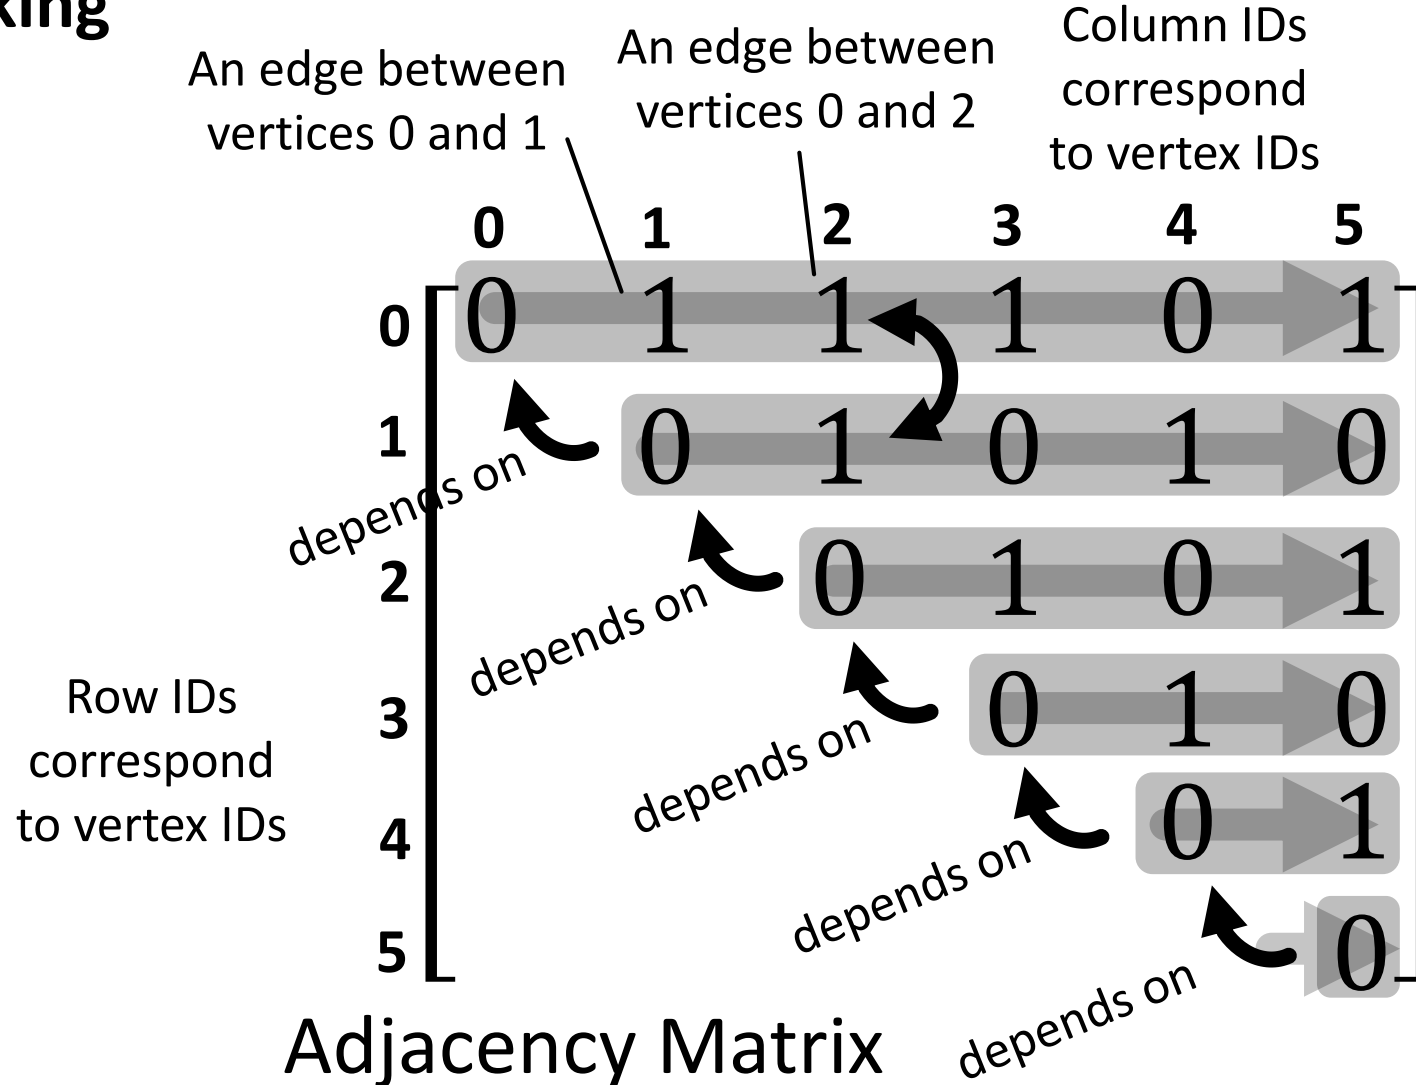
# Substream-Centric MWM: FPGA optimizations

## Blocking



# Substream-Centric MWM: FPGA optimizations

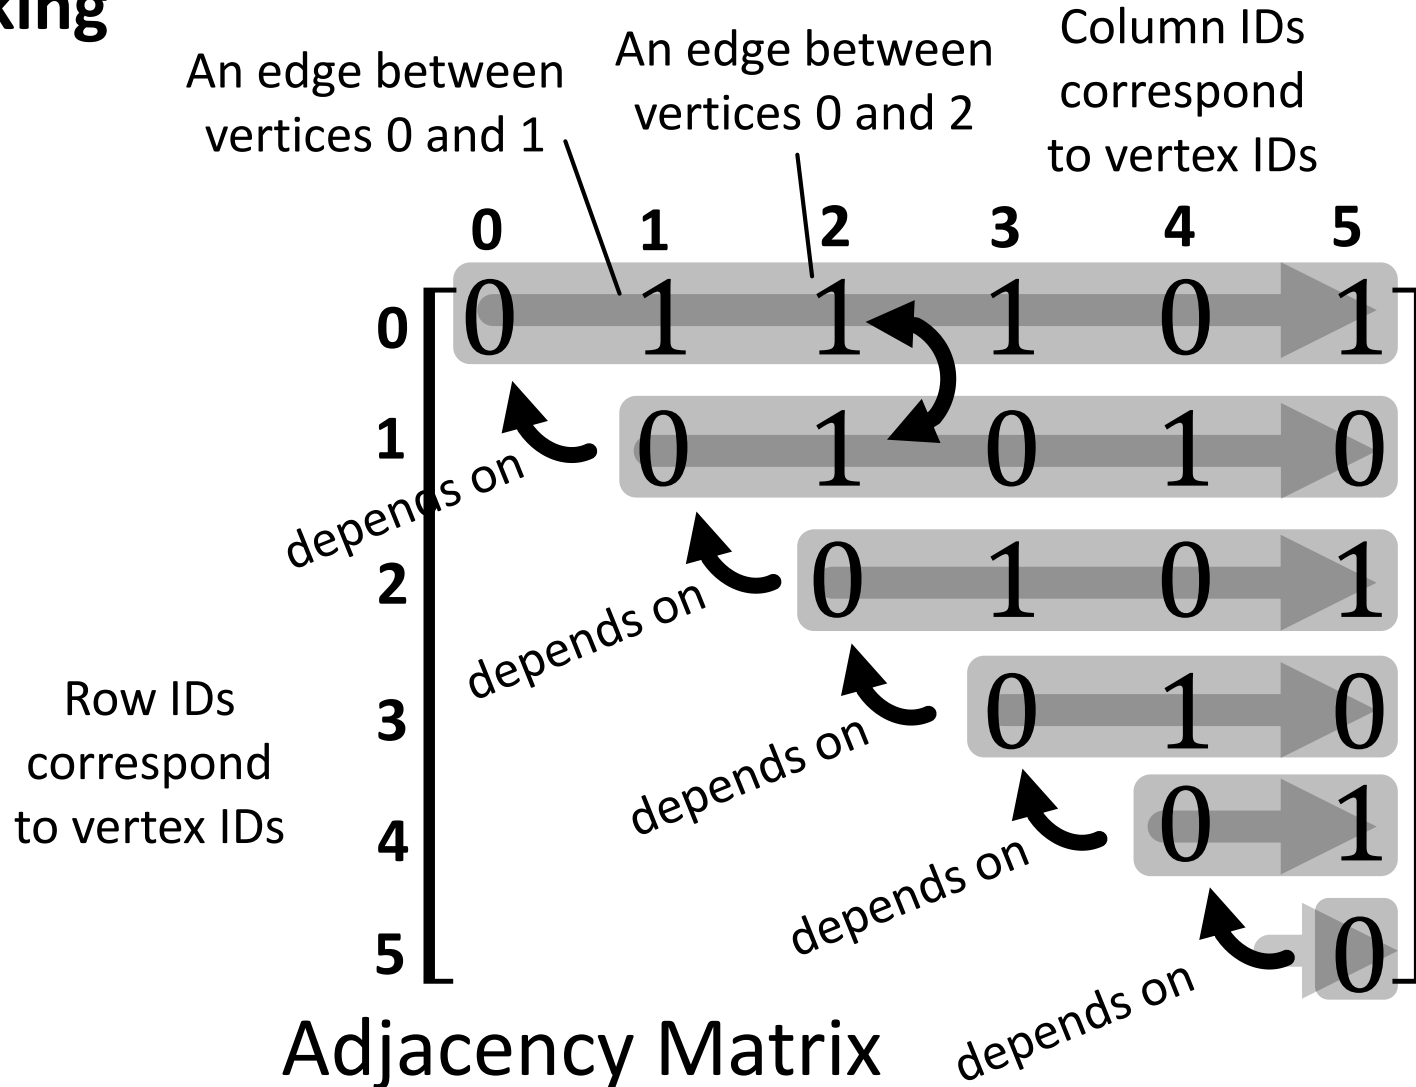
## Blocking



# Substream-Centric MWM: FPGA optimizations

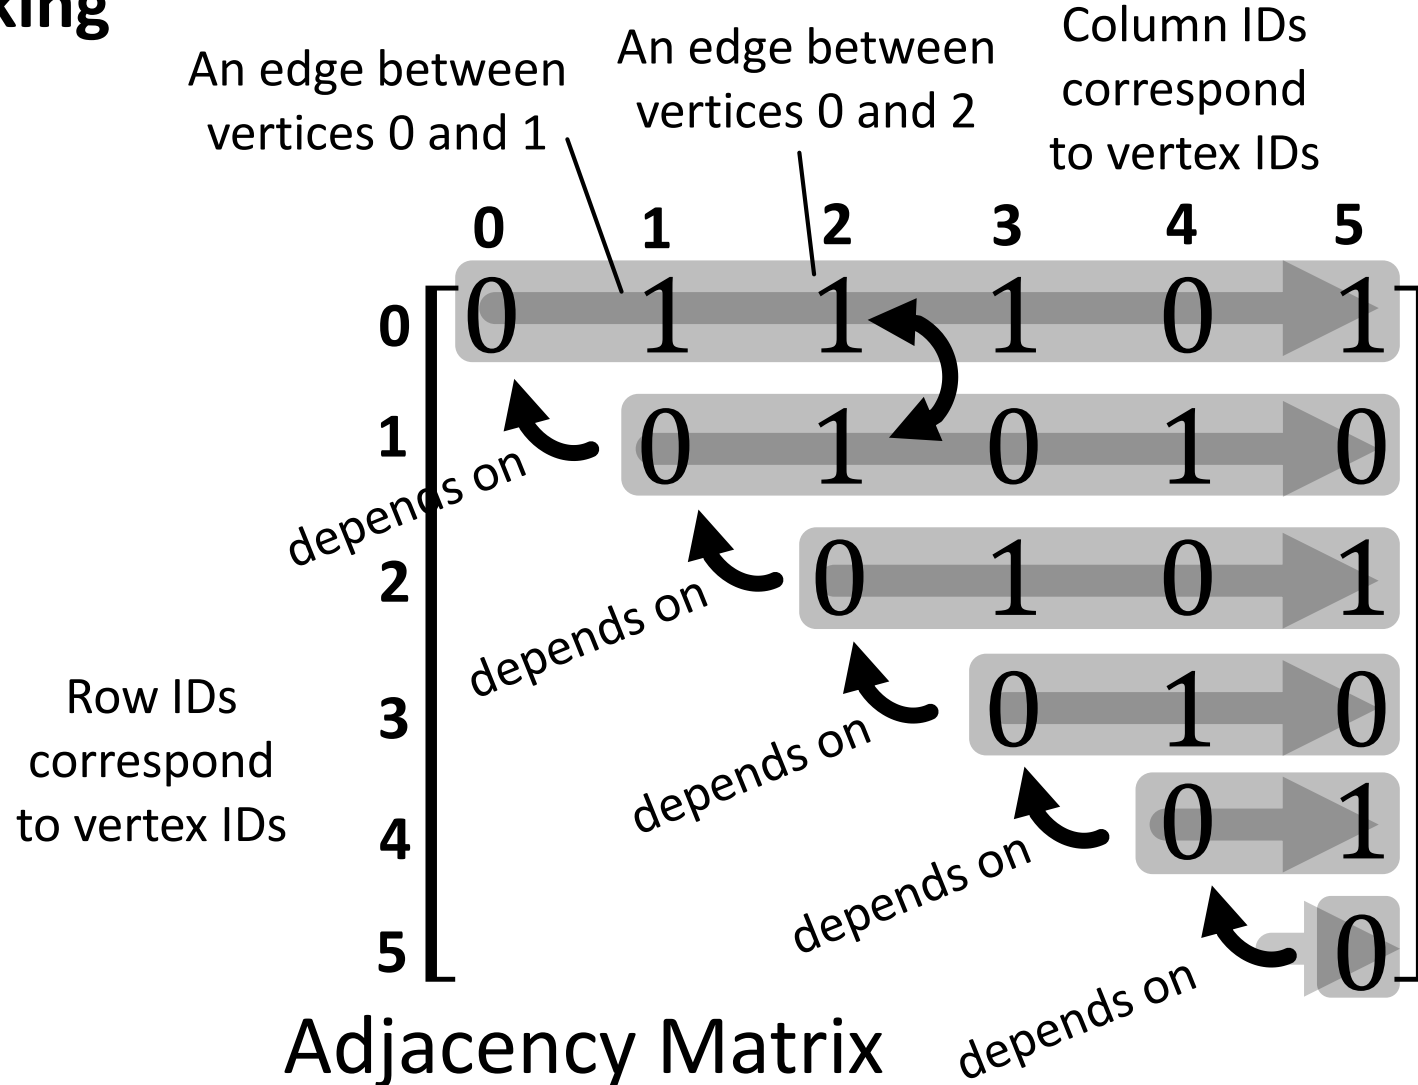
## Blocking

Introduce a (tunable) „blocking parameter” K



# Substream-Centric MWM: FPGA optimizations

## Blocking



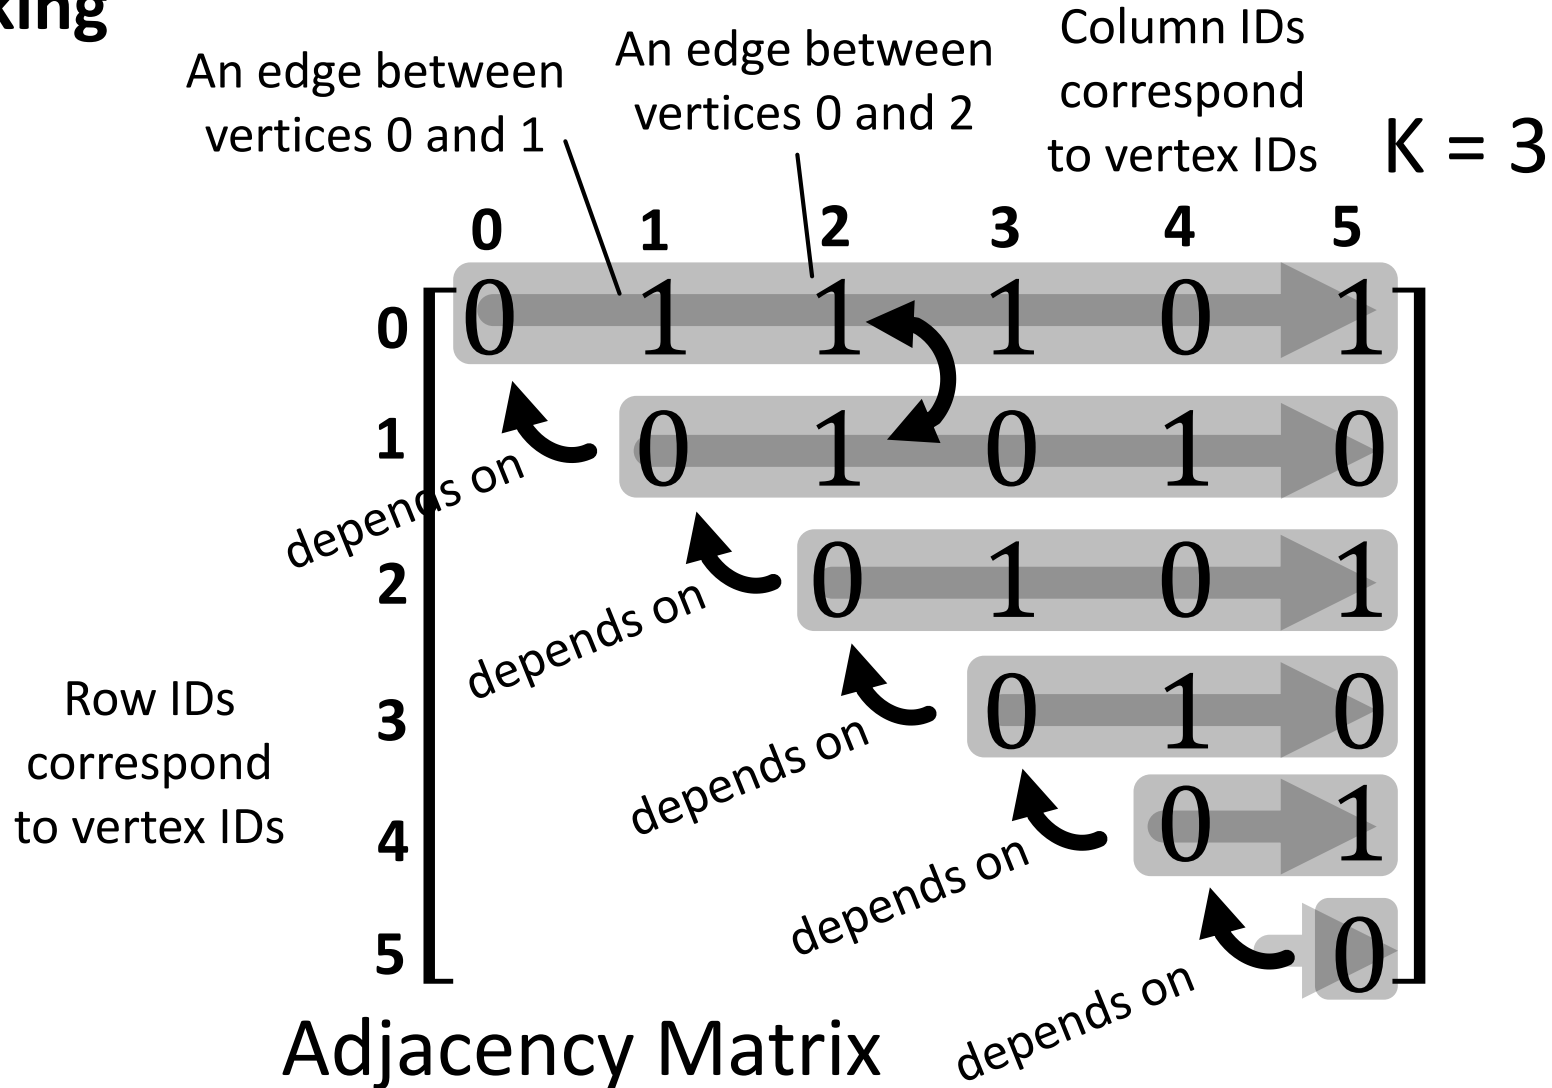
Introduce a (tunable) „blocking parameter” K

K determines how many stalls are allowed



# Substream-Centric MWM: FPGA optimizations

## Blocking




Introduce a (tunable) „blocking parameter”  $K$

$K$  determines how many stalls are allowed

# Substream-Centric MWM: FPGA optimizations

## Blocking

$K = 3$



Adjacency Matrix

Introduce a (tunable) „blocking parameter”  $K$

$K$  determines how many stalls are allowed

Portions of rows are ordered „lexicographically” (i.e., no strict ordering that enforces a stall is required)

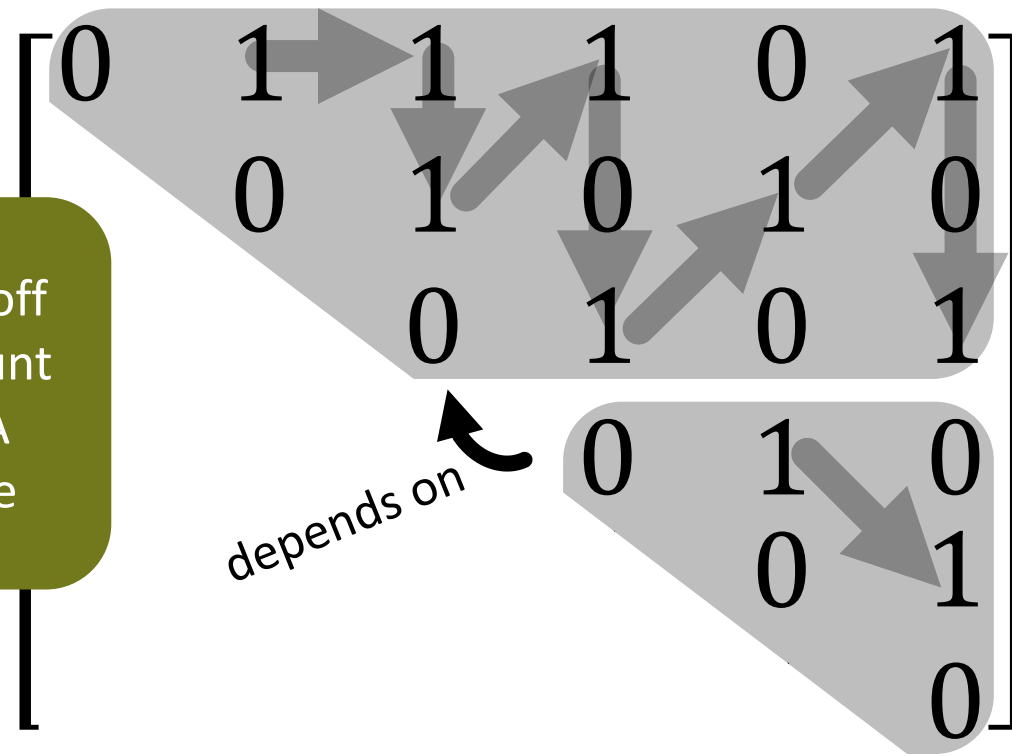
Algorithm still (provably) correct

# Substream-Centric MWM: FPGA optimizations

## Blocking

$K = 3$

K is tunable: it controls the tradeoff between the amount of the used FPGA resources and the performance



Adjacency Matrix

Introduce a (tunable) „blocking parameter” K

K determines how many stalls are allowed

Portions of rows are ordered „lexicographically” (i.e., no strict ordering that enforces a stall is required)

Algorithm still (provably) correct

# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Use substream-centric processing (exposes parallelism, enables easy pipelining, supports approximation)



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

## Research Questions



Combine the MWM algorithm by Crouch and Stubbs with a hybrid CPU-FPGA setting



Use substream-centric processing (exposes parallelism, enables easy pipelining, supports approximation)



The proper use of blocking, vectorization, pipelining, prefetching



What is the ultimate performance, power consumption, and the related tradeoffs?

## Research Questions



Combine the MWM algorithm by Crouch and Stubbs with a hybrid CPU-FPGA setting



Use substream-centric processing (exposes parallelism, enables easy pipelining, supports approximation)



The proper use of blocking, vectorization, pipelining, prefetching



What is the ultimate performance, power consumption, and the related tradeoffs?

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

# Part 5: Evaluation

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES



## PERFORMANCE ANALYSIS

### TYPES OF MACHINES

CPU: Intel Broadwell  
Xeon E5-2680 v4 @3.3 GHz  
14 Cores (28 Threads)



Altera Arria 10 @200MHz



# PERFORMANCE ANALYSIS

## TYPES OF GRAPHS

# PERFORMANCE ANALYSIS


## TYPES OF GRAPHS

Synthetic graphs


# PERFORMANCE ANALYSIS

## TYPES OF GRAPHS

### Synthetic graphs




Kronecker [1]




# PERFORMANCE ANALYSIS

## TYPES OF GRAPHS

### Synthetic graphs



Kronecker [1]



Real-world graphs (SNAP [2], KONECT [3], DIMACS [4])


[2] SNAP. <https://snap.stanford.edu>

[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.


# PERFORMANCE ANALYSIS

## TYPES OF GRAPHS


### Synthetic graphs



Kronecker [1]



### Real-world graphs (SNAP [2], KONECT [3], DIMACS [4])




Web graphs



Road networks



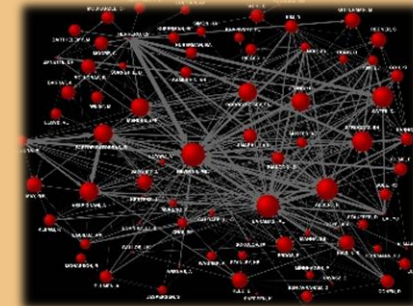
Social networks



Purchase networks



Communication graphs



Citation graphs

[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] SNAP. <https://snap.stanford.edu>

[3] KONECT. <https://konect.cc>

[4] DIMACS Challenge

# PERFORMANCE ANALYSIS

## ALGORITHMS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

[2] M. Ghaffari. Space-optimal semi-streaming for  $(2+\epsilon)$ -approximate matching. arXiv:1701.03730, 2017.

# PERFORMANCE ANALYSIS

## ALGORITHMS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
<b>Substream-Centric (SC-OPT)</b>	Hybrid

Our FPGA design,  
( $4+\epsilon$ )-approximation

[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

[2] M. Ghaffari. Space-optimal semi-streaming for  $(2+\epsilon)$ -approximate matching. arXiv:1701.03730, 2017.



# PERFORMANCE ANALYSIS

## ALGORITHMS

CPU implementations of the original Crouch scheme,  $(4+\epsilon)$ -approximation

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

Our FPGA design,  $(4+\epsilon)$ -approximation

[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

[2] M. Ghaffari. Space-optimal semi-streaming for  $(2+\epsilon)$ -approximate matching. arXiv:1701.03730, 2017.

# PERFORMANCE ANALYSIS

## ALGORITHMS

CPU implementations of the original Crouch scheme,  $(4+\epsilon)$ -approximation

State-of-the-art MWM algorithm, space-optimal, **time-optimal  $O(m)$** ,  $(2+\epsilon)$ -approximation

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

Our FPGA design,  $(4+\epsilon)$ -approximation

[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

[2] M. Ghaffari. Space-optimal semi-streaming for  $(2+\epsilon)$ -approximate matching. arXiv:1701.03730, 2017.

# PERFORMANCE ANALYSIS

## ALGORITHMS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

CPU implementations of the original Crouch scheme,  $(4+\epsilon)$ -approximation

State-of-the-art MWM algorithm, space-optimal, **time-optimal  $O(m)$** ,  $(2+\epsilon)$ -approximation

Our FPGA design,  $(4+\epsilon)$ -approximation

We test both CPU and hybrid (FPGA+CPU) platforms

[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

[2] M. Ghaffari. Space-optimal semi-streaming for  $(2+\epsilon)$ -approximate matching. arXiv:1701.03730, 2017.

### Parameters:

Blocking size (K) = 32,  
 #Substreams (L) = 64  
 #Threads = 4,  $\epsilon = 0.1$

# PERFORMANCE ANALYSIS

## VARIOUS GRAPHS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k; k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

### Parameters:



Blocking size ( $K$ ) = 32,  
 #Substreams ( $L$ ) = 64  
 #Threads = 4,  $\epsilon = 0.1$

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k; k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

# PERFORMANCE ANALYSIS

## VARIOUS GRAPHS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid



### Parameters:



Blocking size ( $K$ ) = 32,  
 #Substreams ( $L$ ) = 64  
 #Threads = 4,  $\epsilon = 0.1$

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k, k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

# PERFORMANCE ANALYSIS

## VARIOUS GRAPHS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid



# PERFORMANCE ANALYSIS

## VARIOUS GRAPHS



### Parameters:

Blocking size ( $K$ ) = 32,  
 #Substreams ( $L$ ) = 64  
 #Threads = 4,  $\epsilon = 0.1$

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k, k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

SC-OPT secures highest performance



# PERFORMANCE ANALYSIS



## VARIOUS GRAPHS

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

**Parameters:**  
 Blocking size (K) = 32,  
 #Substreams (L) = 64  
 #Threads = 4,  $\epsilon = 0.1$

**SC-OPT secures highest performance**

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k, k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770





# PERFORMANCE ANALYSIS

## APPROXIMATION (ACCURACY)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

### Parameters:

#Substreams (L) = 128,  
Blocking size (K) = 32,  
#threads = 4, #edges = 8M  
(Kronecker)


# PERFORMANCE ANALYSIS

## APPROXIMATION (ACCURACY)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

### Parameters:

#Substreams (L) = 128,  
 Blocking size (K) = 32,  
 #threads = 4, #edges = 8M  
 (Kronecker)




# PERFORMANCE ANALYSIS

## APPROXIMATION (ACCURACY)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

### Parameters:

#Substreams (L) = 128,  
 Blocking size (K) = 32,  
 #threads = 4, #edges = 8M  
 (Kronecker)



# PERFORMANCE ANALYSIS


## APPROXIMATION (ACCURACY)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

### Parameters:

#Substreams (L) = 128,  
 Blocking size (K) = 32,  
 #threads = 4, #edges = 8M  
 (Kronecker)

SC-OPT is comparable to the  $(2+\epsilon)$ -approximation by Ghaffari et al.



**Parameters:**

*L*: #Substreams (pipelines),

*K*: Blocking size,

*T*: #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

**PERFORMANCE ANALYSIS****ENERGY CONSUMPTION, RESOURCE UTILIZATION**

**Parameters:**

$L$ : #Substreams (pipelines),

$K$ : Blocking size,

$T$ : #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

# PERFORMANCE ANALYSIS

## ENERGY CONSUMPTION, RESOURCE UTILIZATION

Algorithm	Parameters	Energy Consumption [W]
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789
SC-OPT	$K = 32, L = 64$	14.657
CS-PAR	$T = 64$	120

**Parameters:**

$L$ : #Substreams (pipelines),

$K$ : Blocking size,

$T$ : #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

# PERFORMANCE ANALYSIS

## ENERGY CONSUMPTION, RESOURCE UTILIZATION

Algorithm	Parameters	Energy Consumption [W]
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789
SC-OPT	$K = 32, L = 64$	14.657
CS-PAR	$T = 64$	120

**Parameters:**

$L$ : #Substreams (pipelines),

$K$ : Blocking size,

$T$ : #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

# PERFORMANCE ANALYSIS

## ENERGY CONSUMPTION, RESOURCE UTILIZATION

Algorithm	Parameters	Energy Consumption [W]
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789
SC-OPT	$K = 32, L = 64$	14.657
CS-PAR	$T = 64$	120

SC-OPT (Hybrid) is ~8x more power-efficient than the CPU implementation



**Parameters:**

$L$ : #Substreams (pipelines),  
 $K$ : Blocking size,  
 $T$ : #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

# PERFORMANCE ANALYSIS

## ENERGY CONSUMPTION, RESOURCE UTILIZATION

Algorithm	Parameters	Energy Consumption [W]
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789
SC-OPT	$K = 32, L = 64$	14.657
CS-PAR	$T = 64$	120

SC-OPT (Hybrid) is ~8x more power-efficient than the CPU implementation

FPGA Algorithm	Parameters	Used BRAM	Used ALMs
SC-OPT	$K = 32, L = 512$	11.5 MBit (21%)	151,998 (32%)
SC-OPT	$K = 256, L = 128$	24.8 MBit (45%)	350,556 (82%)

**Parameters:**

$L$ : #Substreams (pipelines),

$K$ : Blocking size,

$T$ : #CPU threads

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

# PERFORMANCE ANALYSIS

## ENERGY CONSUMPTION, RESOURCE UTILIZATION

Algorithm	Parameters	Energy Consumption [W]
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789
SC-OPT	$K = 32, L = 64$	14.657
CS-PAR	$T = 64$	120

SC-OPT (Hybrid) is ~8x more power-efficient than the CPU implementation

FPGA Algorithm	Parameters	Used BRAM	Used ALMs
SC-OPT	$K = 32, L = 512$	11.5 MBit (21%)	151,998 (32%)
SC-OPT	$K = 256, L = 128$	24.8 MBit (45%)	350,556 (82%)

Blocking needs more resources (but is **tunable!**)


# PERFORMANCE ANALYSIS

## DESIGN SPACE EXPLORATION

# PERFORMANCE ANALYSIS

## DESIGN SPACE EXPLORATION

B – BRAM size allocated for matching data structures,  
L – number of substreams (pipelines)




x - the highest possible value of B for a given L

# PERFORMANCE ANALYSIS

## DESIGN SPACE EXPLORATION

B – BRAM size allocated for matching data structures,  
 L – number of substreams (pipelines)

Addition complexity grows linearly with  $L$



x - the highest possible value of B for a given L


# PERFORMANCE ANALYSIS

## DESIGN SPACE EXPLORATION

BRAM signal propagation limits the frequency


B – BRAM size allocated for matching data structures,  
L – number of substreams (pipelines)

Addition complexity grows linearly with  $L$




x - the highest possible value of B for a given L

# OTHER ALGORITHMS, PROBLEMS, ANALYSES




# OTHER ALGORITHMS, PROBLEMS, ANALYSES





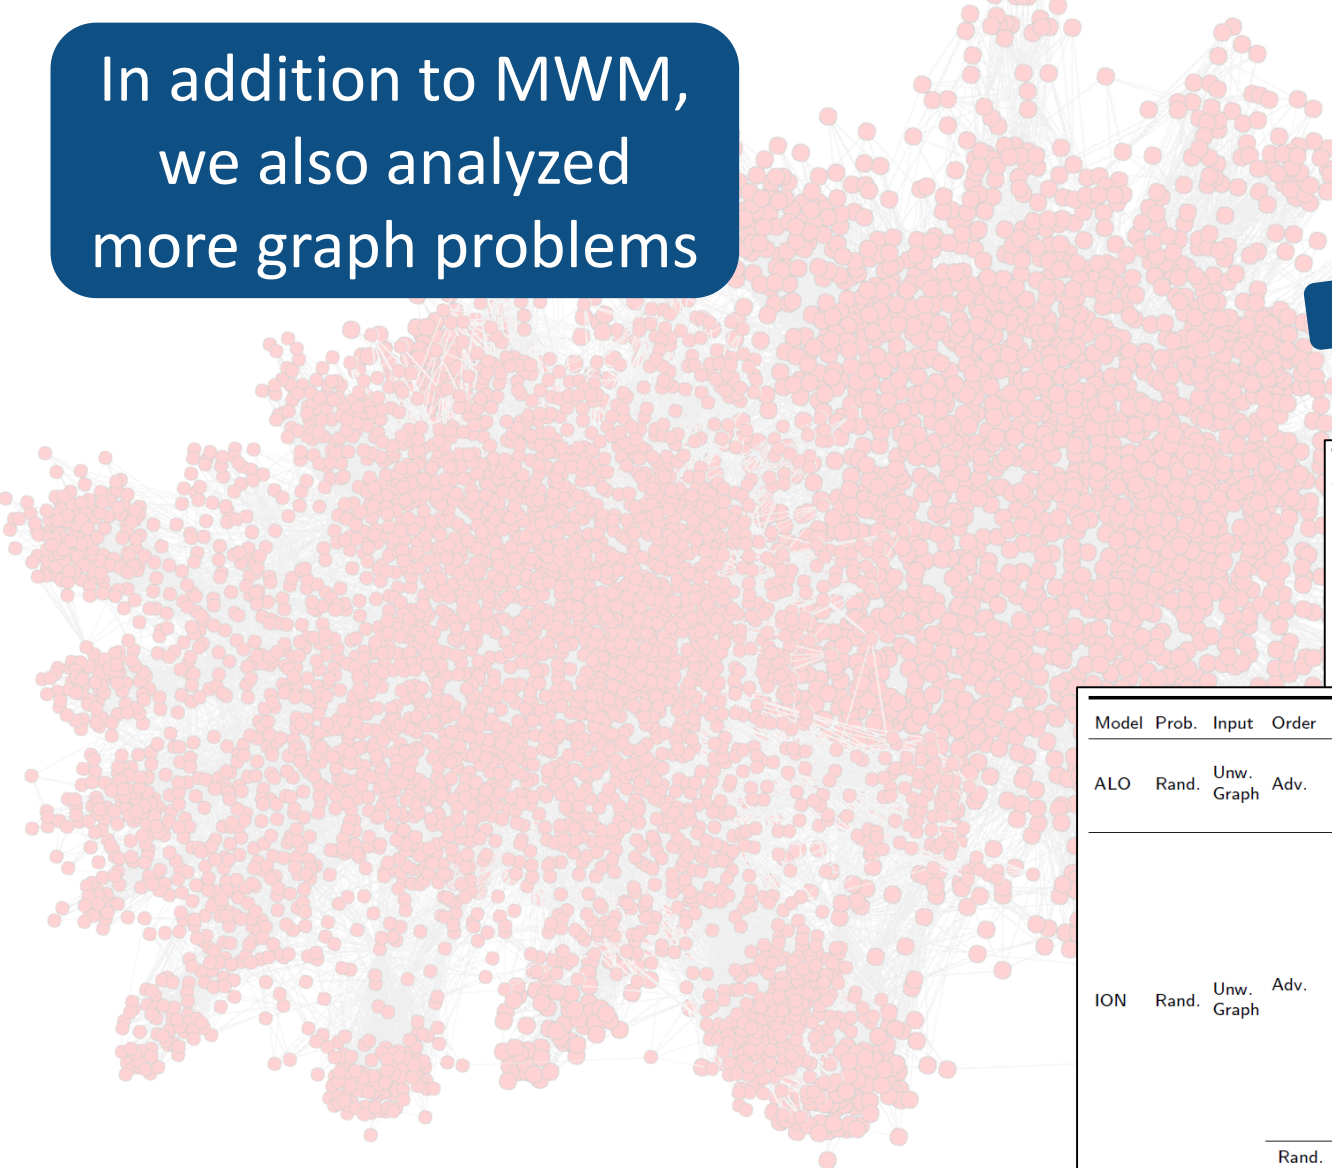
## OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM,  
we also analyzed  
more graph problems



# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems



- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Colorings
- Minimum spanning trees
- Random walks
- Spanners
- Connectivity
- Triangle counting
- Densest subgraphs
- K-edge connectivity
- K-vertex connectivity

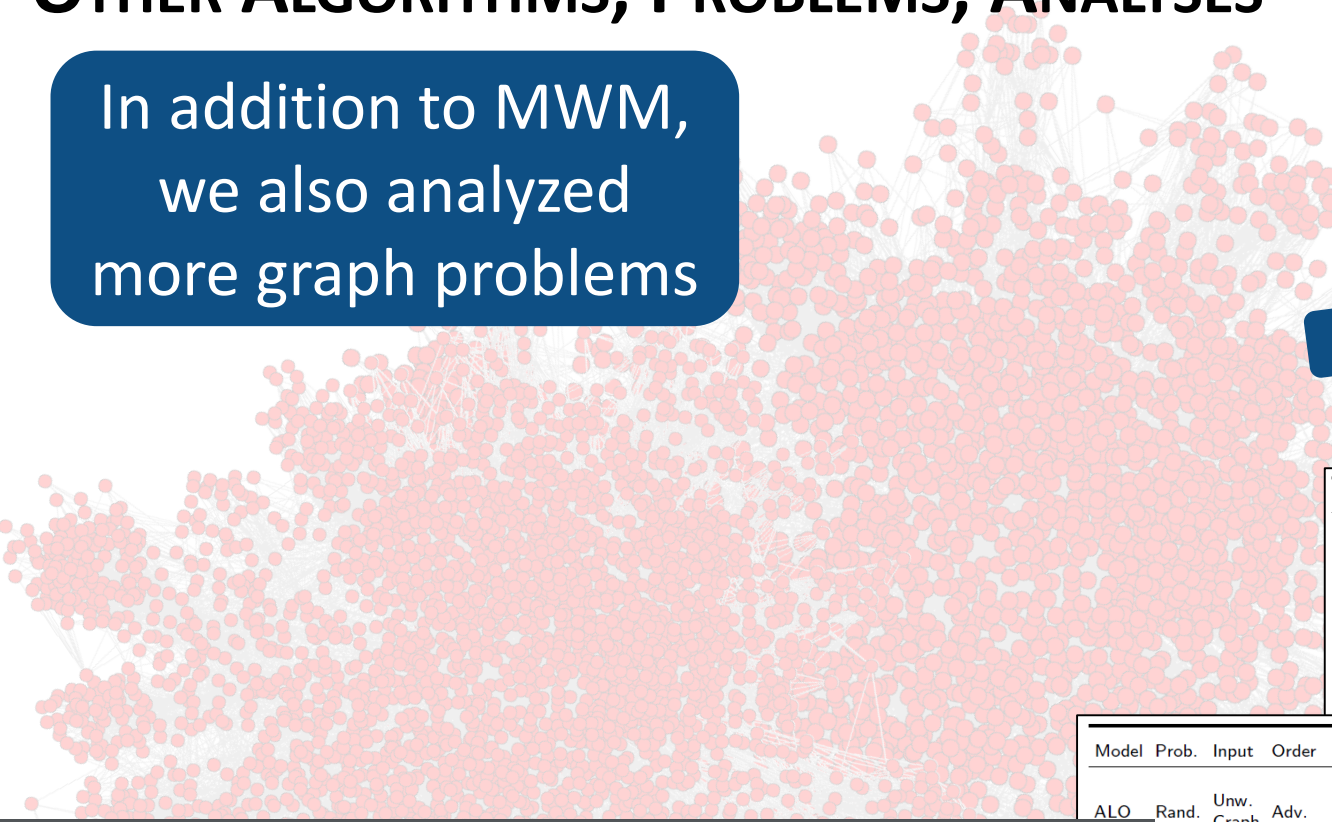
Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ION	Rand.	Unw. Graph	Adv.	[37]	$1/2$	$O(n)$	$O(1)$	1	-
DGS	Rand.	Unw. Graph	Adv.	[57]	$(\epsilon, \delta)$	$\tilde{O}(n \log(n))$	$O(1)$	1	Sampling
TN	Rand.	Unw. Graph	Adv.	[57]	$(\epsilon, \delta)$	$\tilde{O}(n \log(n))$	$O(1)$	1	Sampling

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ION	Rand.	Unw. Graph	Adv.	[22, Theorem 4]	$(\epsilon, \delta)$	$\tilde{O}(n(T_2 + T_3)/T_3 \log(2/\delta))$	-	1	Sampling
				[79, Theorem 3]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	1	Sampling
				[79, Theorem 5]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	2	Sampling
				[22]	$(\epsilon, \delta)$	$\tilde{O}((T_1 + 2T_2 + 3T_3)/T_3 \log(2/\delta))$	-	1	Sampling
				[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	2	Sampling
				[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	3	Sampling
				[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2} / T)$	-	4	Reserv. Sampling
				[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^2 / T)$	-	2	Reserv. Sampling
				[86, Theorem 3.3]	$(\epsilon, \delta)$	$O(6/\epsilon^2 m \Delta / T \log(2/\delta))$	-	1	-
				[55, Theorem 1]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) m \Delta^2 / T)$	1	Sampling	
[55, Theorem 2]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) ((m^3 + mC_4 + C_6) / T^2 + 1))$	1	Sampling					
[55, Theorem 3]	$(\epsilon, \delta)$	$O(n + 1/\epsilon^2 \log(1/\delta) (T_2/T + 1) \log(n))$	3	Sampling					
[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m / (\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	-	2	Sampling			
	Rand.		[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m / (\epsilon^{4.5} \sqrt{T}))$	-	1	Sampling	

[90, Theorem 3.1] -  $\tilde{O}(n\alpha + \sqrt{l/\alpha})$  -  $\tilde{O}(\sqrt{l/\alpha})$  Sampling

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems



Graph problems and algorithms:

- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Colorings
- Minimum spanning trees
- Random walks
- Spanners
- Connectivity
- Triangle counting
- Densest subgraphs
- K-edge connectivity
- K-vertex connectivity
- Sparsification

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ION	Rand.	Unw. Graph	Adv.	[37]	1/2	$O(n)$	$O(1)$	1	-
ION	Det.	Unw. Bi-Graph	Adv.	[37, Theorem 2]	6	$O(n \log(n))$	$O(1)^*$	1	Greedy
				[75, Theorem 3]	$2 + \epsilon$ (5.82)	$O(n \text{ polylog}(n))$	$O(1)^*$	$O(1)$ (1)	Greedy
				[102]	5.58	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-
				[34]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-
				[28]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	$\Omega(\log(n))^*$	1	-

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method	
ALO	Rand.	Unw. Graph	Adv.	[22, Theorem 4]	$(\epsilon, \delta)$	$\tilde{O}(n(T_2 + T_3)/T_3 \log(2/\delta))$	-	1	Sampling	
				[79, Theorem 3]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	1	Sampling	
				[79, Theorem 5]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	2	Sampling	
				[22]	$(\epsilon, \delta)$	$\tilde{O}((T_1 + 2T_2 + 3T_3)/T_3 \log(2/\delta))$	-	1	Sampling	
				[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	2	Sampling	
				[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	3	Sampling	
				[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2} / T)$	-	4	Reserv. Sampling	
				[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^2 / T)$	-	2	Reserv. Sampling	
				[86, Theorem 3.3]	$(\epsilon, \delta)$	$O(6/\epsilon^2 m \Delta / T \log(2/\delta))$	-	1	-	
				[55, Theorem 1]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) m \Delta^2 / T)$	1	1	Sampling	
				[55, Theorem 2]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta)((m^3 + mC_4 + C_6)/T^2 + 1))$	1	1	Sampling	
				[55, Theorem 3]	$(\epsilon, \delta)$	$O(n + 1/\epsilon^2 \log(1/\delta)(T_2/T + 1) \log(n))$	3	3	Sampling	
				[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m/(\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	-	2	Sampling
				[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m/(\epsilon^{4.5} \sqrt{T}))$	-	1	Sampling	

1

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges


Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

[90, Theorem 3.1] -  $\tilde{O}(n\alpha + \sqrt{l/\alpha})$  -  $\tilde{O}(\sqrt{l/\alpha})$  Sampling

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems



- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Colorings
- Minimum spanning trees
- Random walks
- Spanners
- Connectivity
- Triangle counting
- Densest subgraphs
- K-edge connectivity
- K-vertex connectivity

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ION	Rand.	Unw. Graph	Adv.	[37]	1/2	$O(n)$	$O(1)$	1	-
DGS	Rand.	Unw. Graph	Adv.	[57]					
TN	Rand.	Unw. Graph	Adv.	[57]					

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ION	Det.	Unw. Bi-Graph	Adv.	[37, Theorem 2]	6	$O(n \log(n))$	$O(1)^*$	1	Greedy
				[75, Theorem 3]	$2 + \epsilon$ (5.82)	$O(n \text{ polylog}(n))$	$O(1)^*$	$O(1)$ (1)	Greedy
				[102]	5.58	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-
				[34]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-
				[28]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	$\Omega(\log(n))^*$	1	-

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method
ALO	Rand.	Unw. Graph	Adv.	[22, Theorem 4]	$(\epsilon, \delta)$	$\tilde{O}(n(T_2 + T_3)/T_3 \log(2/\delta))$	-	1	Sampling
				[79, Theorem 3]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	1	Sampling
				[79, Theorem 5]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	2	Sampling
				[22]	$(\epsilon, \delta)$	$\tilde{O}((T_1 + 2T_2 + 3T_3)/T_3 \log(2/\delta))$	-	1	Sampling
				[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	2	Sampling
				[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	3	Sampling
				[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2} / T)$	-	4	Reserv. Sampling
				[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^2 / T)$	-	2	Reserv. Sampling
				[86, Theorem 3.3]	$(\epsilon, \delta)$	$O(6/\epsilon^2 m \Delta / T \log(2/\delta))$	-	1	-
				[55, Theorem 1]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) m \Delta^2 / T)$	1	Sampling	
[55, Theorem 2]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) ((m^3 + mC_4 + C_6) / T^2 + 1))$	1	Sampling					
[55, Theorem 3]	$(\epsilon, \delta)$	$O(n + 1/\epsilon^2 \log(1/\delta) (T_2/T + 1) \log(n))$	3	Sampling					
[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m / (\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	-	2	Sampling			
Rand.	[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m / (\epsilon^{4.5} \sqrt{T}))$	-	1	Sampling			

<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
 (submitted to arXiv, will appear tonight)

1

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems

To enable rigorous reasoning, we analyzed ~15 models for streaming graph processing (and selected the best for FPGAs)

<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
(submitted to arXiv, will appear tonight)

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph problems and algorithms highlighted in blue callouts:

- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Colorings
- Minimum spanning trees
- Random walks
- Spanners
- Connectivity
- Triangle counting
- Densest subgraphs
- K-edge connectivity
- K-vertex connectivity
- Sparsification

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method	
ION	Det.	Unw. Bi Graph		[37, Theorem 2]	6	$O(n \log(n))$	$O(1)^*$	1	Greedy	
				[75, Theorem 3]	$2 + \epsilon$ (5.82)	$O(n \text{ polylog}(n))$	$O(1)^*$	$O(1)$ (1)	Greedy	
				[102]	5.58	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-	
				[34]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)^*$	1	-	
				[28]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	$\Omega(\log(n))^*$	1	-	
ALO	Rand.	Unw. Graph	Adv.	[22, Theorem 4]	$(\epsilon, \delta)$	$\tilde{O}(n(T_2 + T_3)/T_3 \log(2/\delta))$	-	1	Sampling	
				[79, Theorem 3]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	1	Sampling	
				[79, Theorem 5]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	2	Sampling	
				[22]	$(\epsilon, \delta)$	$\tilde{O}((T_1 + 2T_2 + 3T_3)/T_3 \log(2/\delta))$	-	1	Sampling	
				[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	2	Sampling	
				[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	3	Sampling	
				[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2} / T)$	-	4	Reserv. Sampling	
				[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^2 / T)$	-	2	Reserv. Sampling	
				[86, Theorem 3.3]	$(\epsilon, \delta)$	$O(6/\epsilon^2 m \Delta / T \log(2/\delta))$	-	1	-	
				[55, Theorem 1]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) m \Delta^2 / T)$	1	1	Sampling	
				[55, Theorem 2]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) ((m^3 + mC_4 + C_6)/T^2 + 1))$	1	1	Sampling	
				[55, Theorem 3]	$(\epsilon, \delta)$	$O(n + 1/\epsilon^2 \log(1/\delta) (T_2/T + 1) \log(n))$	3	3	Sampling	
				[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m / (\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	-	2	Sampling
				[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m / (\epsilon^{4.5} \sqrt{T}))$	-	1	Sampling	
				[90, Theorem 3.1]	-	$\tilde{O}(n\alpha + \sqrt{l/\alpha})$	-	-	$\tilde{O}(\sqrt{l/\alpha})$ Sampling	

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems

To enable rigorous reasoning, we analyzed ~15 models for streaming graph processing (and selected the best for FPGAs)

<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
(submitted to arXiv, will appear tonight)

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich



**Graph Problems:** Bipartiteness, Motif counting, Triangle counting, Connected components, Minimum spanning trees, Random walks, Colorings, Spanners, Connectivity, Triangle counting, K-edge connectivity, Densest subgraphs, K-vertex connectivity, Sparsification.


**Streaming Models:** Cash-register, W-Stream, Semi-streaming, Simple streaming, Annotated streaming, MapReduce, Vertex-arrival, Online, MUD, Adjacency-list, Graph Sketching, Insert-only, Sliding window, Dynamic, Turnstile, StreamSort.

Model	Prob.	Input	Order	Reference	Approx.	Space	Time	Passes	Method	
ION	Rand.	Unw. Graph	Adv.	[22]	$(\epsilon, \delta)$	$O((T_1 + 2T_2 + 3T_3)/T_3 \log(2/\delta))$	-	1	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$	-	2	Sampling	
TN	Rand.	Unw. Graph	Adv.	[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m^{3/2} / T)$	-	3	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2} / T)$	-	4	Reserv. Sampling	
DGS	Rand.	Unw. Graph	Adv.	[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^2 / T)$	-	2	Reserv. Sampling	
DGS	Rand.	Unw. Graph	Adv.	[86, Theorem 3.3]	$(\epsilon, \delta)$	$O(6/\epsilon^2 m \Delta / T \log(2/\delta))$	-	1	-	
DGS	Rand.	Unw. Graph	Adv.	[55, Theorem 1]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) m \Delta^2 / T)$	1	1	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[55, Theorem 2]	$(\epsilon, \delta)$	$\tilde{O}(\log(1/\delta) ((m^3 + mC_4 + C_6)/T^2 + 1))$	1	1	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[55, Theorem 3]	$(\epsilon, \delta)$	$O(n + 1/\epsilon^2 \log(1/\delta) (T_2/T + 1) \log(n))$	3	3	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m / (\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	-	2	Sampling
DGS	Rand.	Unw. Graph	Adv.	[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m / (\epsilon^{4.5} \sqrt{T}))$	-	1	Sampling	
DGS	Rand.	Unw. Graph	Adv.	[90, Theorem 3.1]	-	$\tilde{O}(n\alpha + \sqrt{l/\alpha})$	-	-	$\tilde{O}(\sqrt{l/\alpha})$ Sampling	

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems

To enable rigorous reasoning, we analyzed ~15 models for streaming graph processing (and selected the best for FPGAs)



<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
(submitted to arXiv, will appear tonight)

**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**  
Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

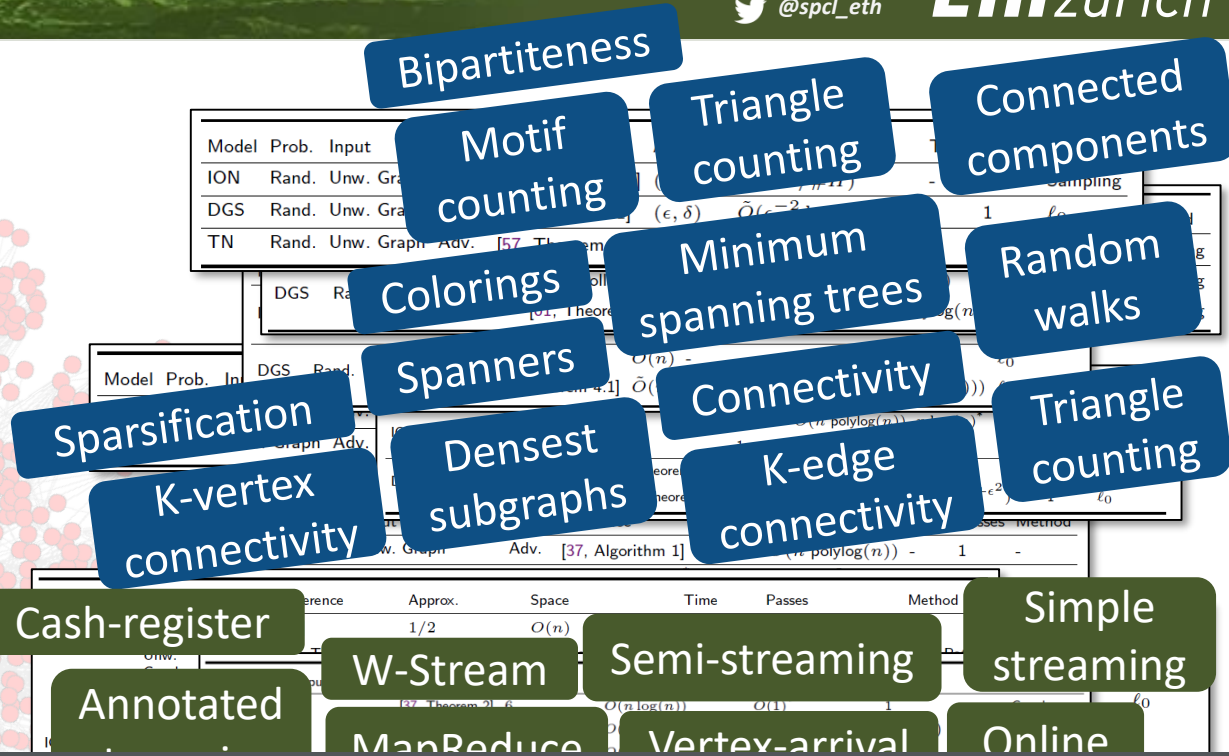
Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

In addition to MWM, we also analyzed more graph problems

To enable rigorous reasoning, we analyzed ~15 models for streaming graph processing (and selected the best for FPGAs)



<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
(submitted to arXiv, will appear tonight)

**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**  
Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich



# OTHER ALGORITHMS, PROBLEMS

In addition to MWM, we also analyzed more graph problems

Work in progress on the distributed setting 😊



To enable rigorous reasoning, we analyzed ~15 models for streaming graph processing (and selected the best for FPGAs)

- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Minimum spanning trees
- Random walks
- Colorings
- Spanners
- Connectivity
- Triangle counting
- Densest subgraphs
- K-edge connectivity
- K-vertex connectivity

Reference	Approx.	Space	Time	Passes	Method
Cash-register	1/2	$O(n)$			Simple streaming
Annotated		$O(n \log(n))$	$O(1)$		W-Stream
MapReduce					Semi-streaming
					Vertex-arrival
					Online

<http://spcl.inf.ethz.ch/Publications/.pdf/graphs-fpgas-survey.pdf>  
(submitted to arXiv, will appear tonight)

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges


Towards Understanding of Modern Graph Processing, Storage, and Analytics

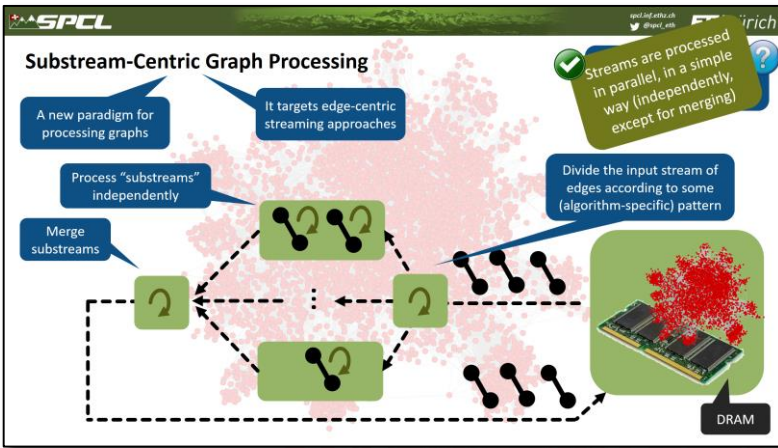
MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

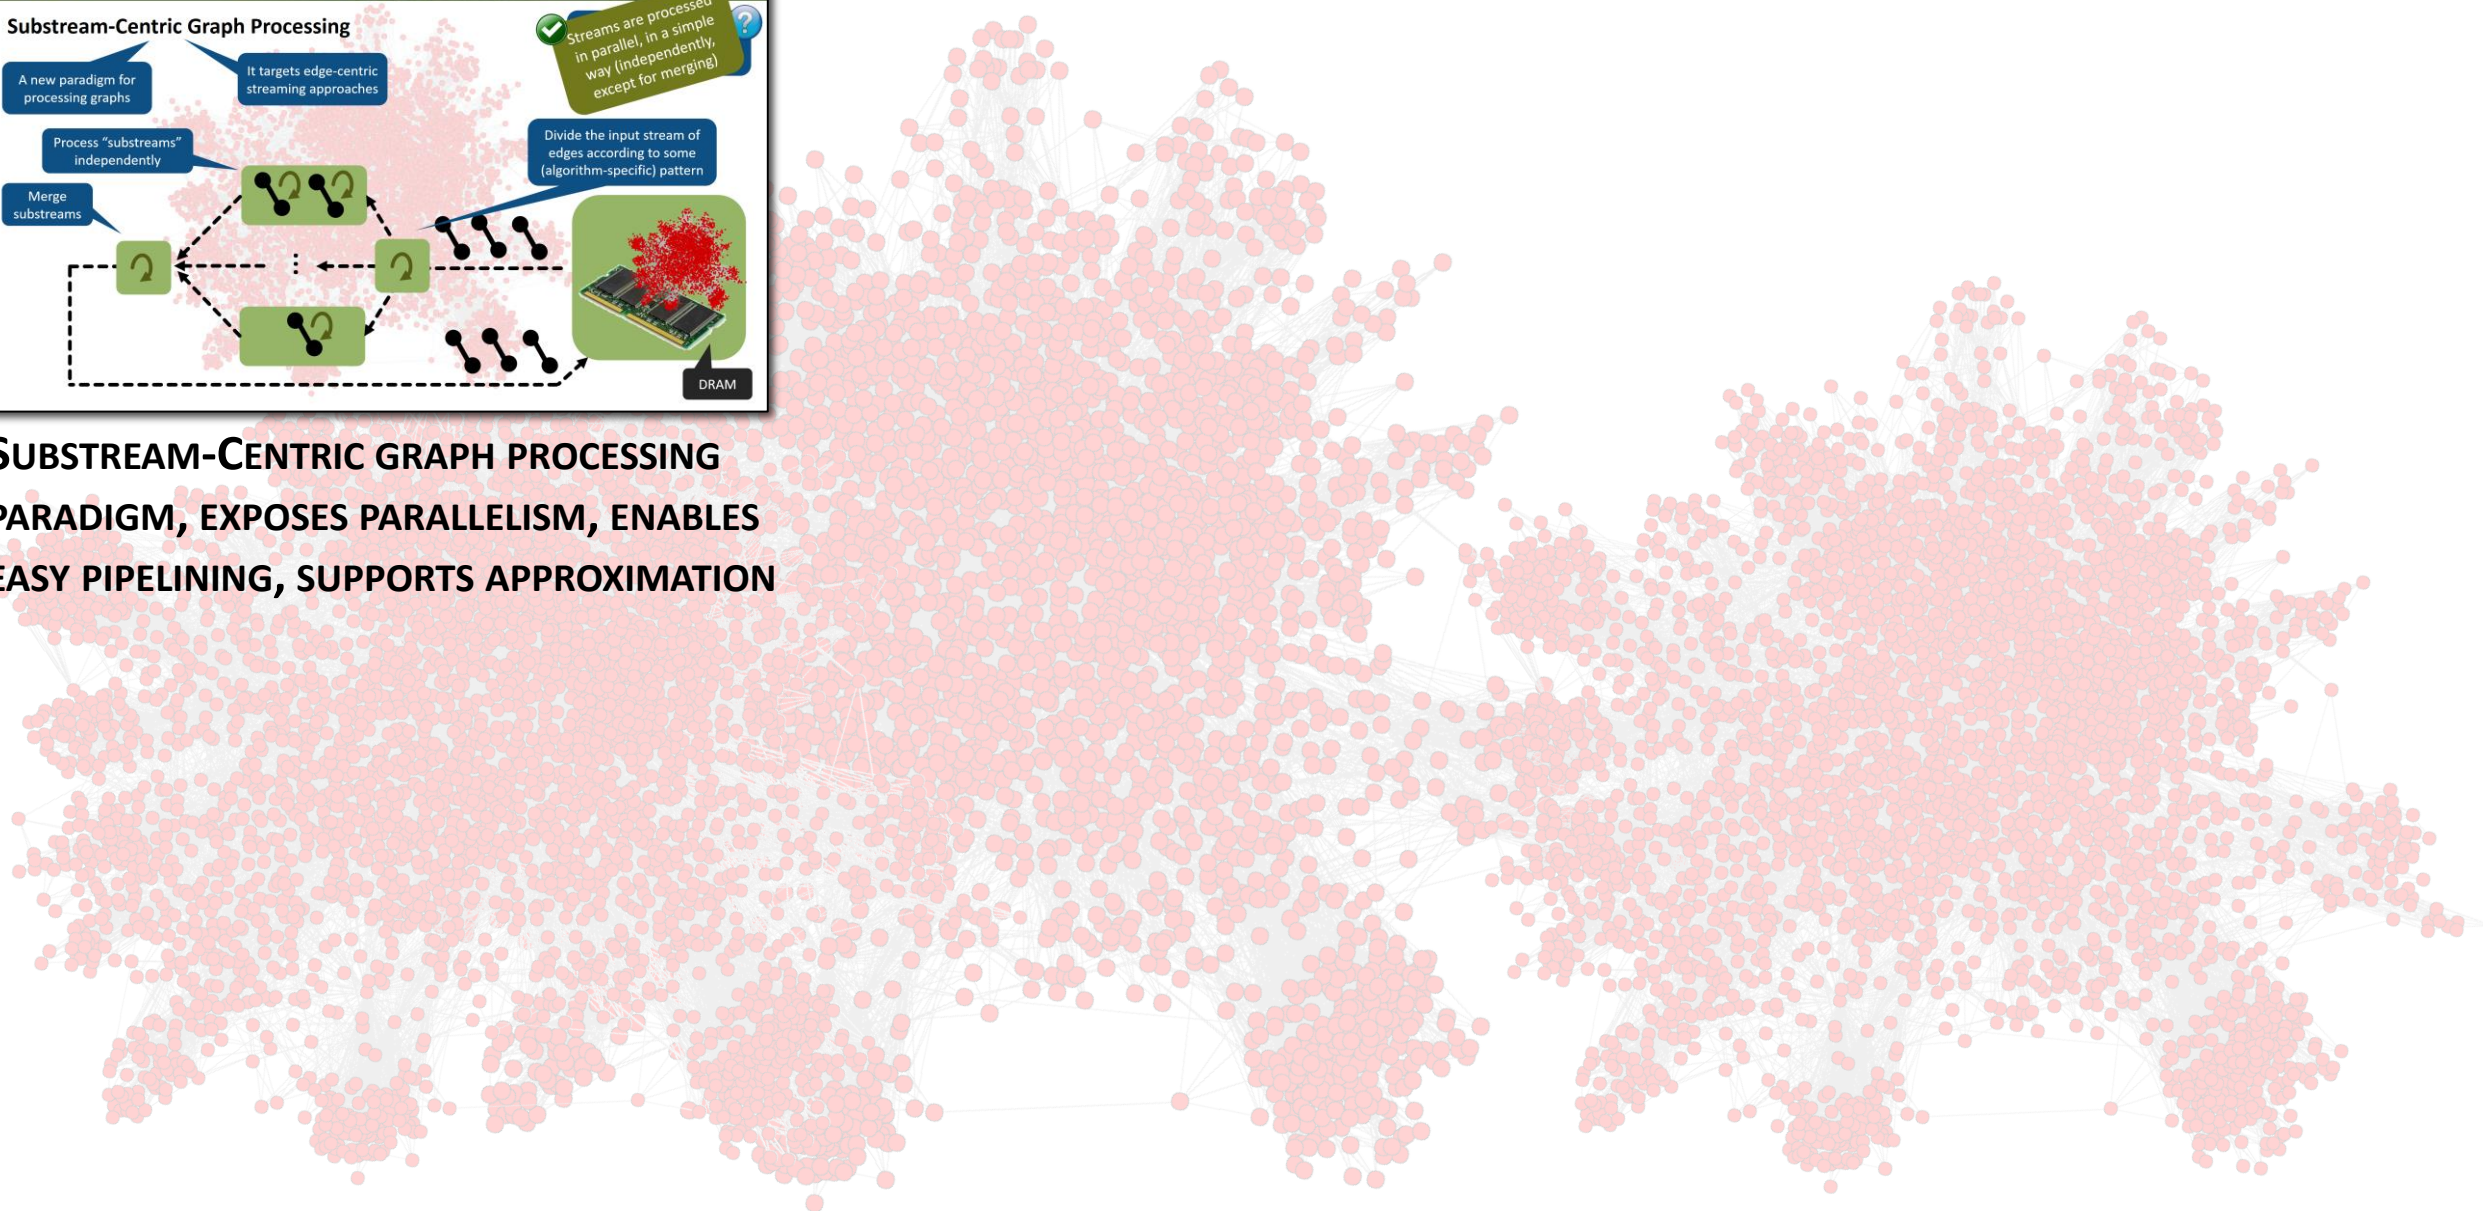
Towards Understanding of Modern Graph Processing and Storage


MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich






**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

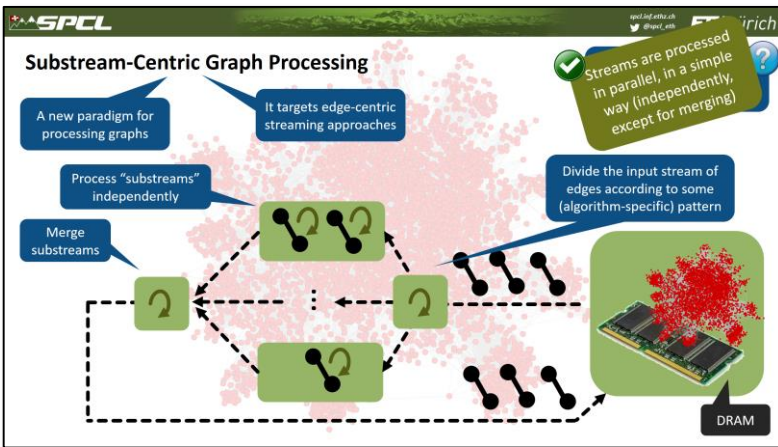




**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**





## SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION

## DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS

### Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing number of updates, with millions of edges added or removed per second. Graph streaming algorithms are crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to add


1


### Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field





**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**

**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

**Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing**

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
 MACIEJ BESTA, Department of Computer Science, ETH Zurich  
 TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

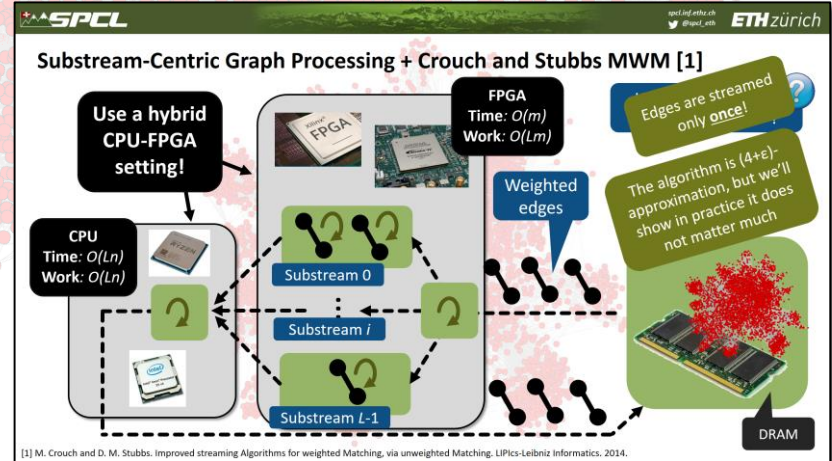
Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing updates, with millions of edges added or removed per second. Graph streaming algorithms are crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to ad


**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

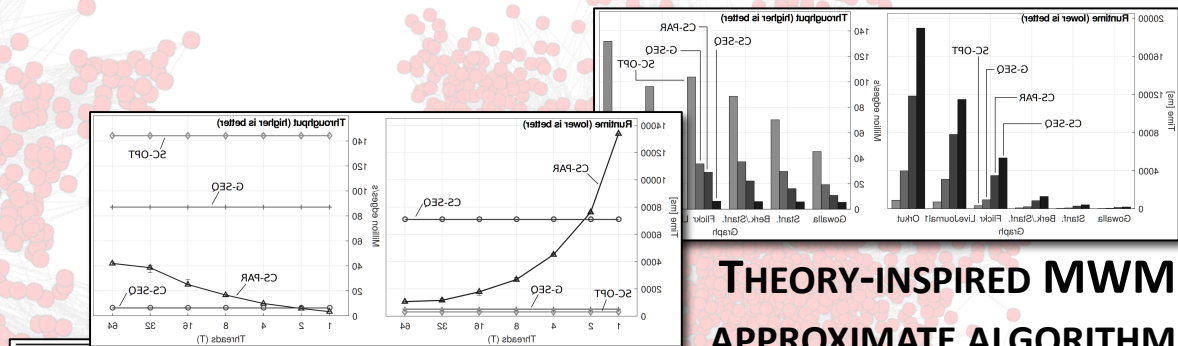

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field





**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**



**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

**Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing**

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing number of updates, with millions of edges added or removed per second. Graph streaming algorithms are crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to adapt

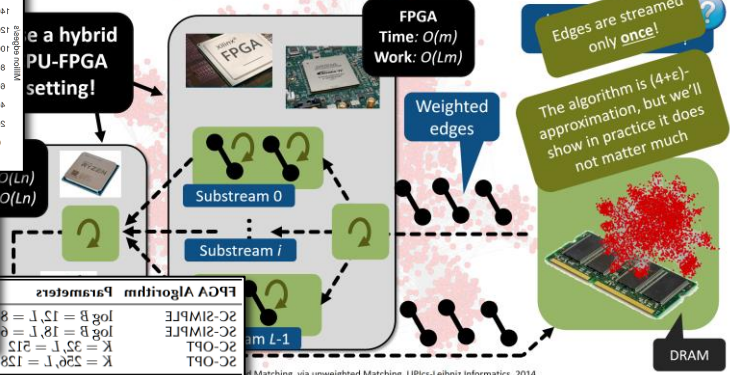
**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**

Towards Understanding of Modern Graph Processing, Storage, and Analytics

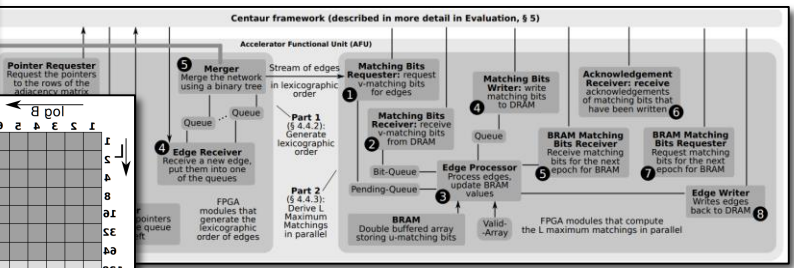
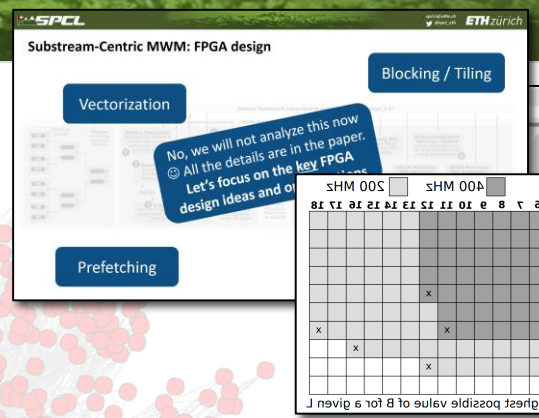

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field

**Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]**





\*Matching, via unweighted Matching. LIPIcs-Leibniz Informatics, 2014.



**GENERIC FPGA DESIGN,  
CODE AVAILABLE**

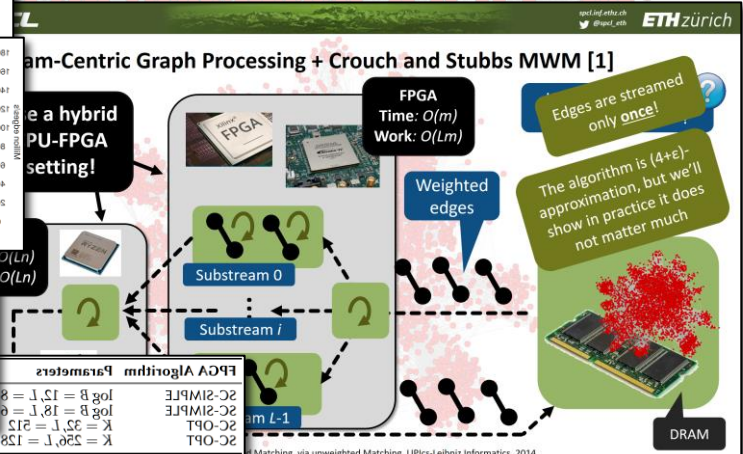
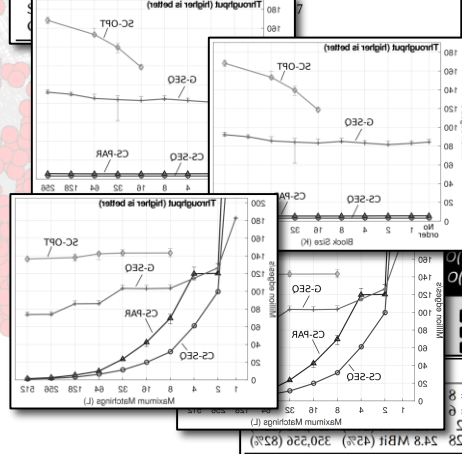
**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**



**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

Algorithm	Parameters	Energy Consumption [W]
SC-SIMPLE	$\log B = 18, L = 6$	14.714
SC-SIMPLE	$\log B = 12, L = 8$	14.598
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789



#### Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing updates, with millions of edges added or removed per second. Graph streaming algorithms are crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to adapt


#### Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field

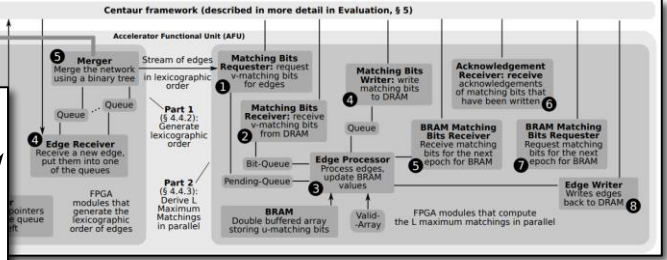





**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**

**GENERALIZABILITY TO OTHER GRAPH PROBLEMS AND SETTINGS**



**GENERIC FPGA DESIGN, CODE AVAILABLE**



**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

**Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing**

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
 MACIEJ BESTA, Department of Computer Science, ETH Zurich  
 TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

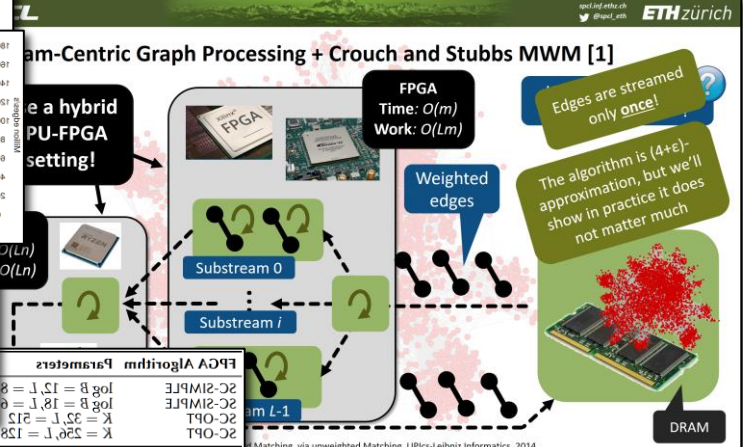
**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**

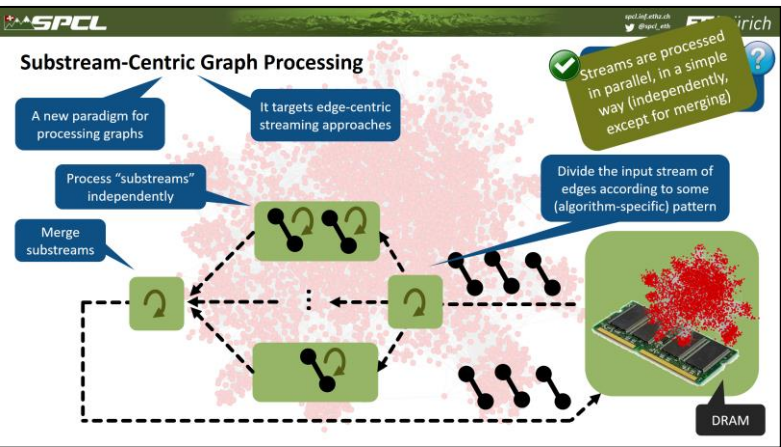
Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field

Algorithm	Parameters	Energy Consumption [W]
SC-SIMPLE	$\log B = 18, L = 6$	14.714
SC-SIMPLE	$\log B = 12, L = 8$	14.598
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789





**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**GENERALIZABILITY TO OTHER GRAPH PROBLEMS AND SETTINGS**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**

**OTHER ALGORITHMS, PROBLEMS, ANALYSES**

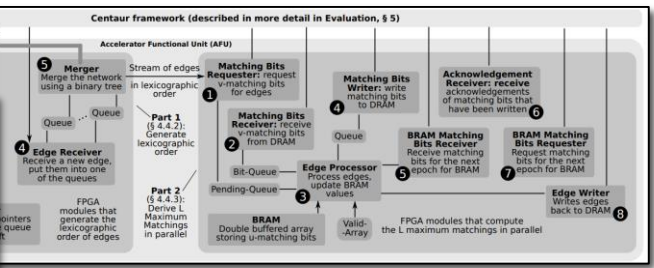
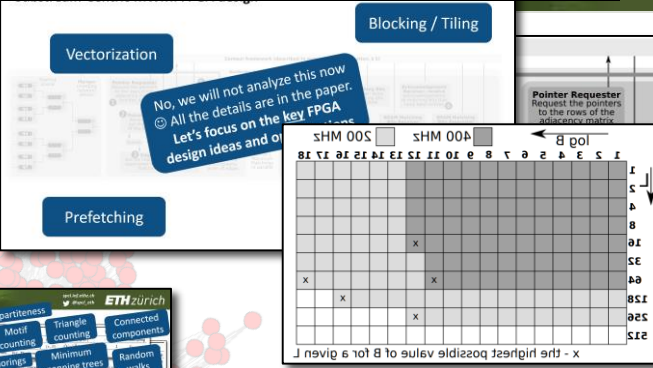
Well, not enough time to present :)

In addition to MWM, we also analyzed more graph problems


<https://arxiv.org/abs/...>

Graph Processing on FPGAs: Taxonomy, Survey, Challenges

ETH zürich



**GENERIC FPGA DESIGN, CODE AVAILABLE**



Algorithm	Parameters	Energy Consumption [W]
SC-SIMPLE	$\log B = 18, L = 6$	14.714
SC-SIMPLE	$\log B = 12, L = 8$	14.598
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789

**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

**Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing**

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

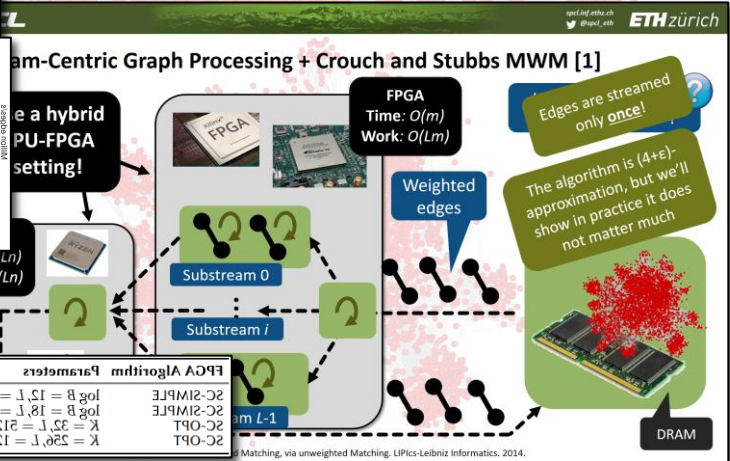
Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing updates, with millions of edges added or removed per second. Graph streaming algorithms are crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to adapt


**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field





**OTHER ALGORITHMS, PROBLEMS, ANALYSES**

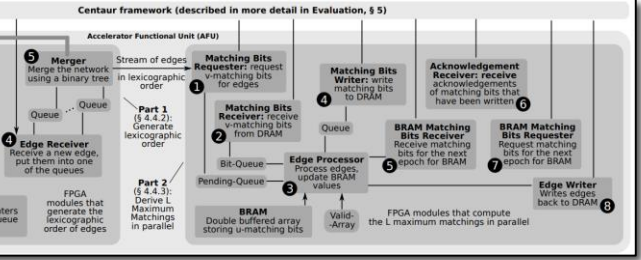
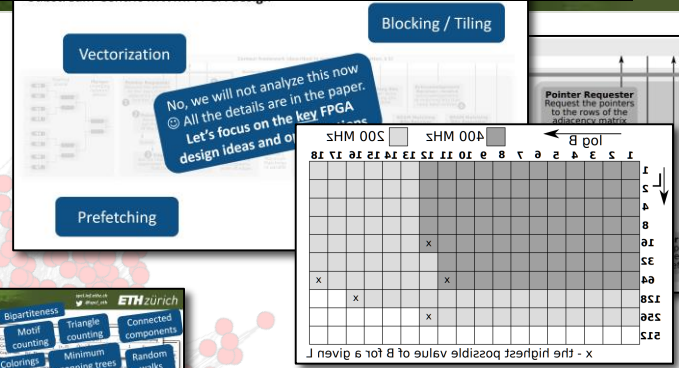
Well, not enough time to present ☹️

In addition to MWM, we also analyzed more graph problems

- Bipartiteness
- Motif counting
- Triangle counting
- Connected components
- Colorings
- Spanners
- Minimum spanning tree
- Random walks
- Sparification
- K-vertex connectivity
- Densest subgraphs
- K-edge connectivity
- Triangle counting

<https://arxiv.org/abs/...>

Graph Processing on FPGAs: Taxonomy, Survey, Challenges



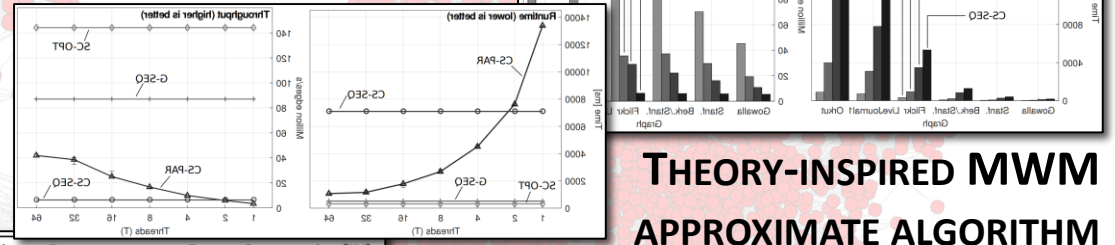
**GENERIC FPGA DESIGN, CODE AVAILABLE**

**SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM, EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION**

**DETAILED DOMAIN ANALYSIS, IDENTIFICATION OF SEMI-STREAMING MODEL AS FPGA BEST-FIT, 2 SURVEYS**

**GENERALIZABILITY TO OTHER GRAPH PROBLEMS AND SETTINGS**

**Thank you for your attention**



**THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING**

Algorithm	Parameters	Energy Consumption [W]
SC-SIMPLE	$\log B = 18, L = 6$	14.714
SC-SIMPLE	$\log B = 12, L = 8$	14.598
SC-OPT	$K = 32, L = 512$	14.789
SC-OPT	$K = 256, L = 128$	14.789

**Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]**

Use a hybrid CPU-FPGA setting!

FPGA Time:  $O(m)$   
Work:  $O(Lm)$

Edges are streamed only once!

The algorithm is  $(4+\epsilon)$ -approximation, but we'll show in practice it does not matter much.

Weighted edges

Substream 0, Substream i

DRAM

**Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing**

Towards Understanding of Modern Graph Processing and Storage


MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**

Towards Understanding of Modern Graph Processing, Storage, and Analytics


MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field

A large, complex network graph is overlaid on the slide. It consists of numerous red circular nodes connected by thin, light gray lines. The nodes are densely packed in several areas, forming clusters that roughly outline the shape of Switzerland. The overall appearance is that of a large-scale social or research network visualization.

# BACKUP & EXTENDED SLIDES


# Large graphs...



## Large graphs...




Problems?



## Large graphs...



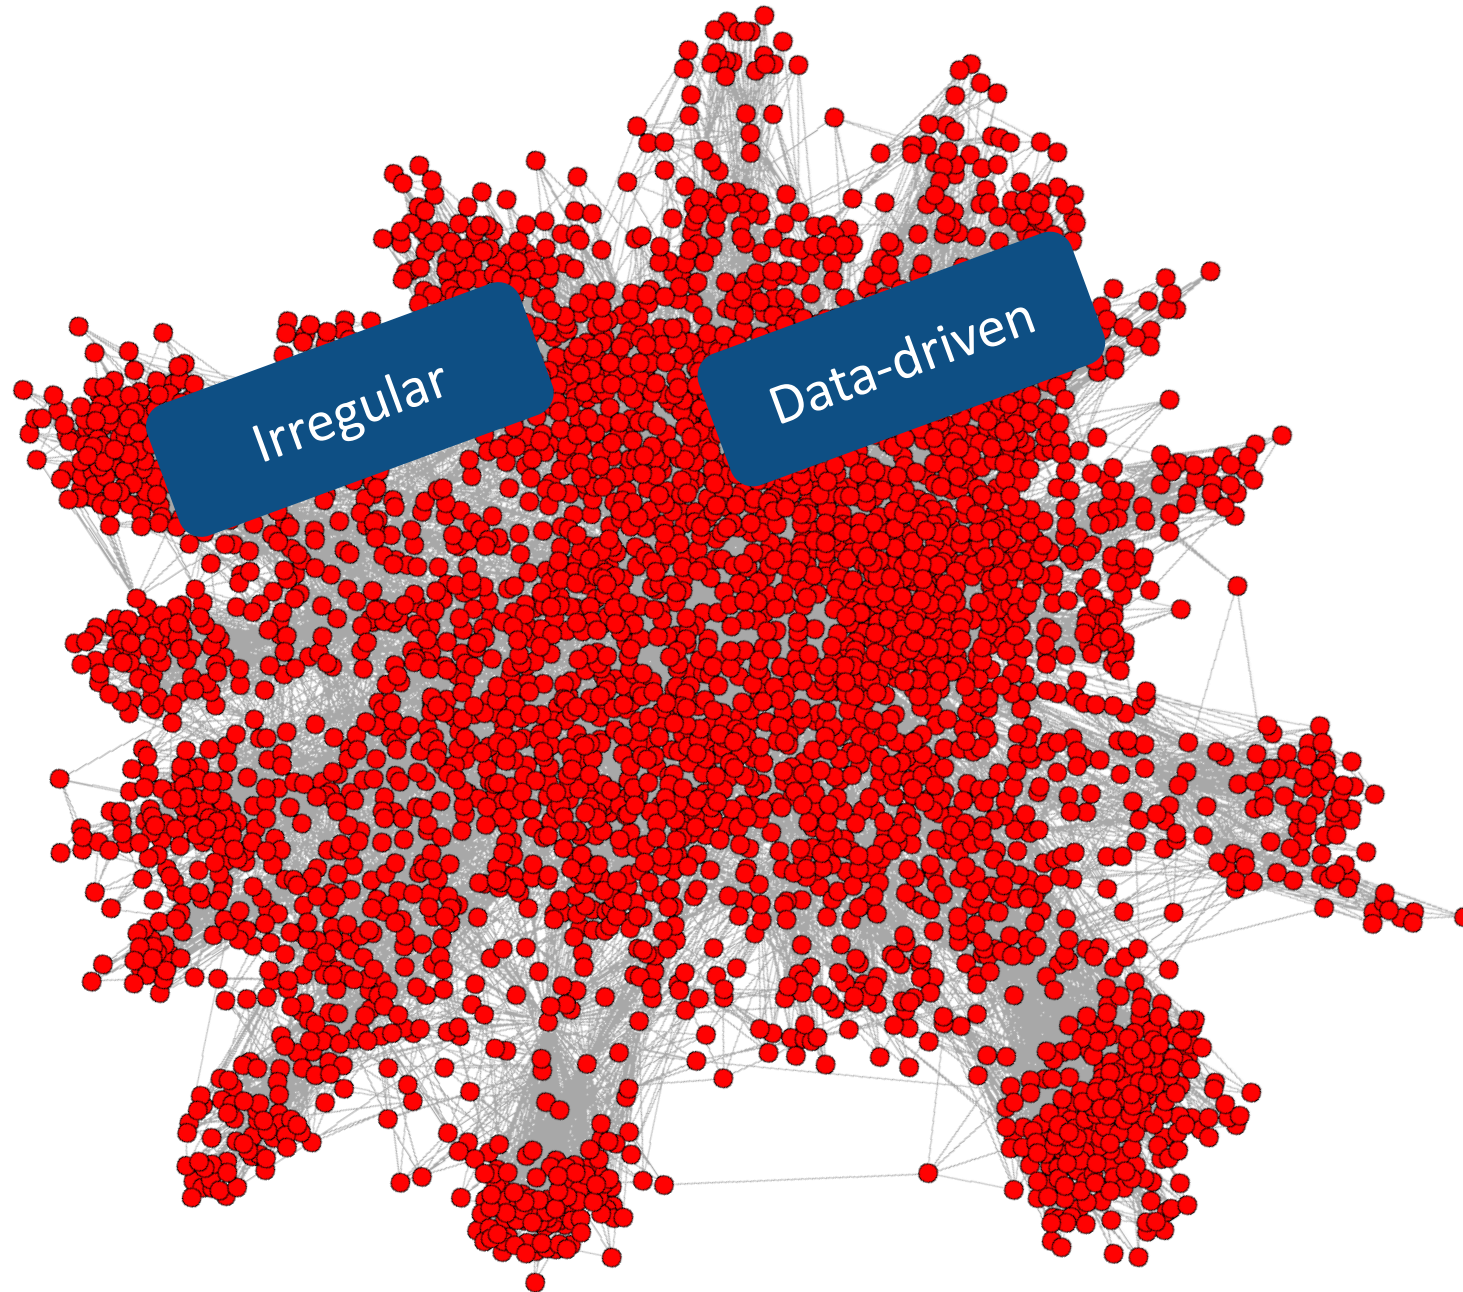
Problems?



## Large graphs...



Problems?






## Large graphs...

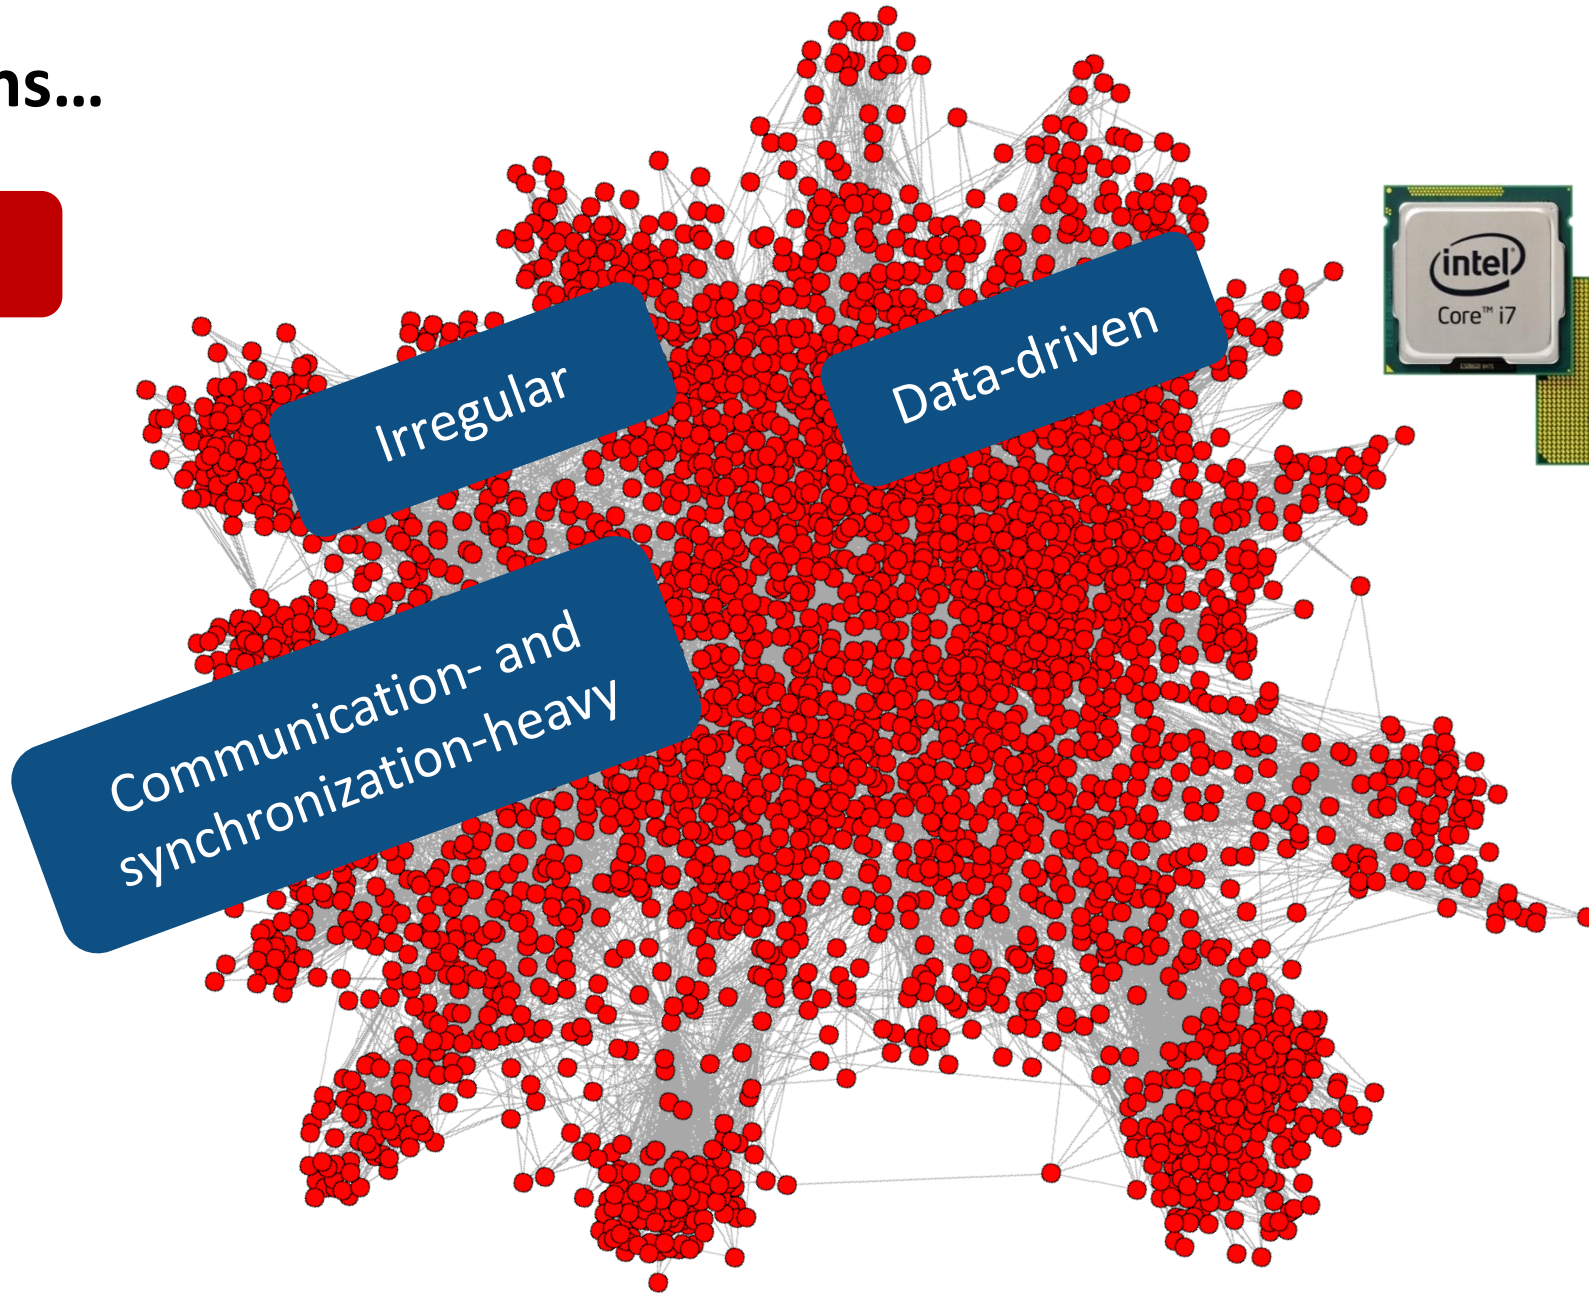


Problems?



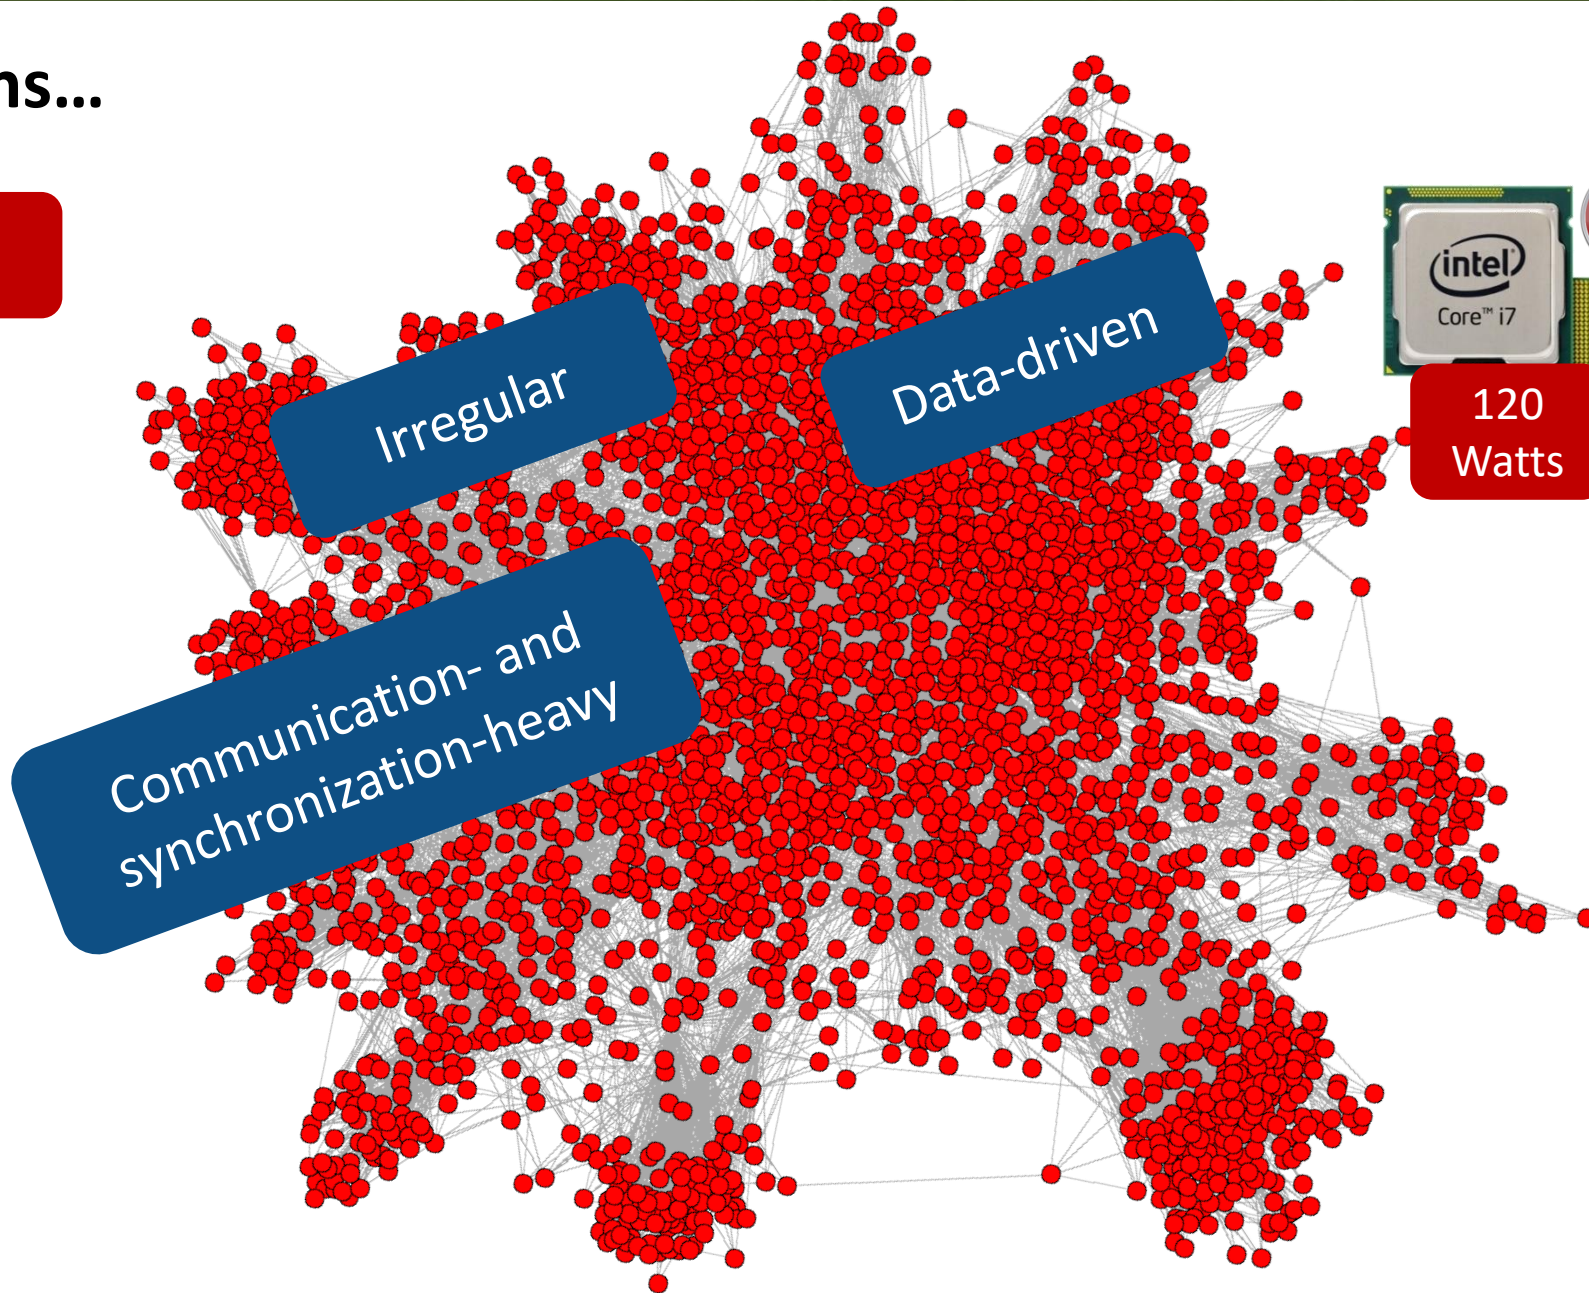
# Large graphs...

 Problems?




# Large graphs...

✗ Problems?



✗  
120  
Watts

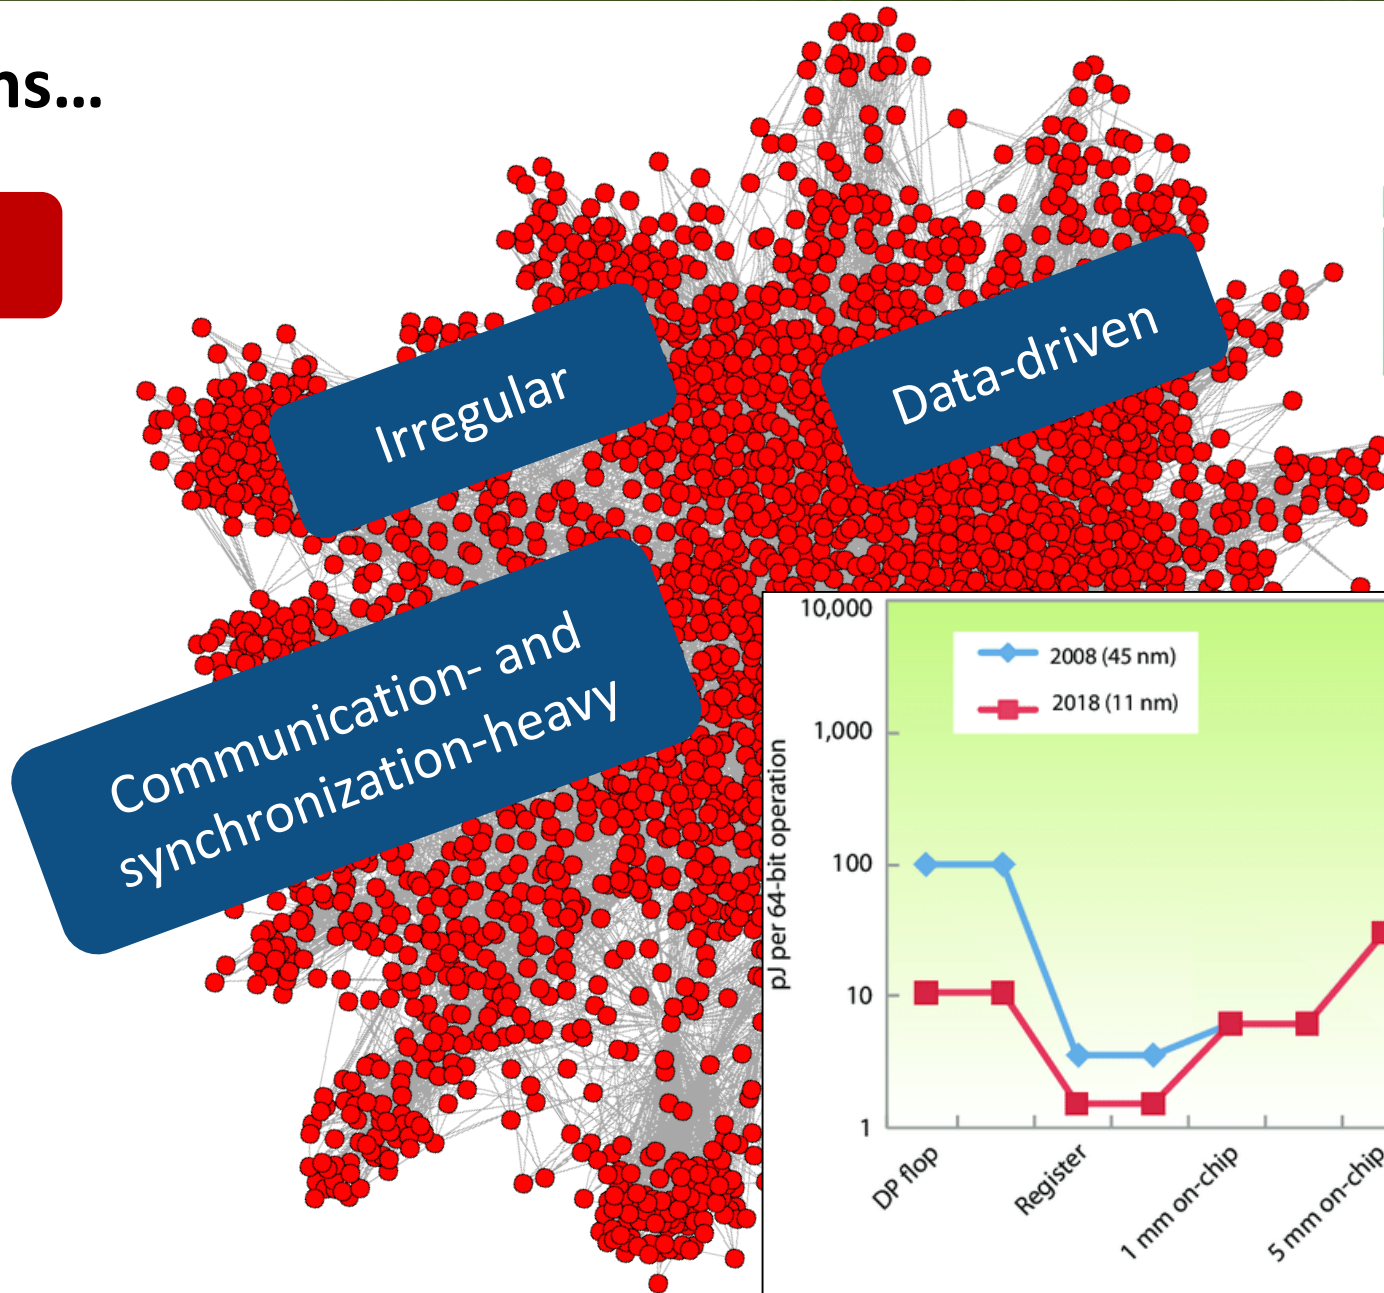


✗  
250  
Watts




# Large graphs...


**✗** Problems?



**✗**  
120  
Watts



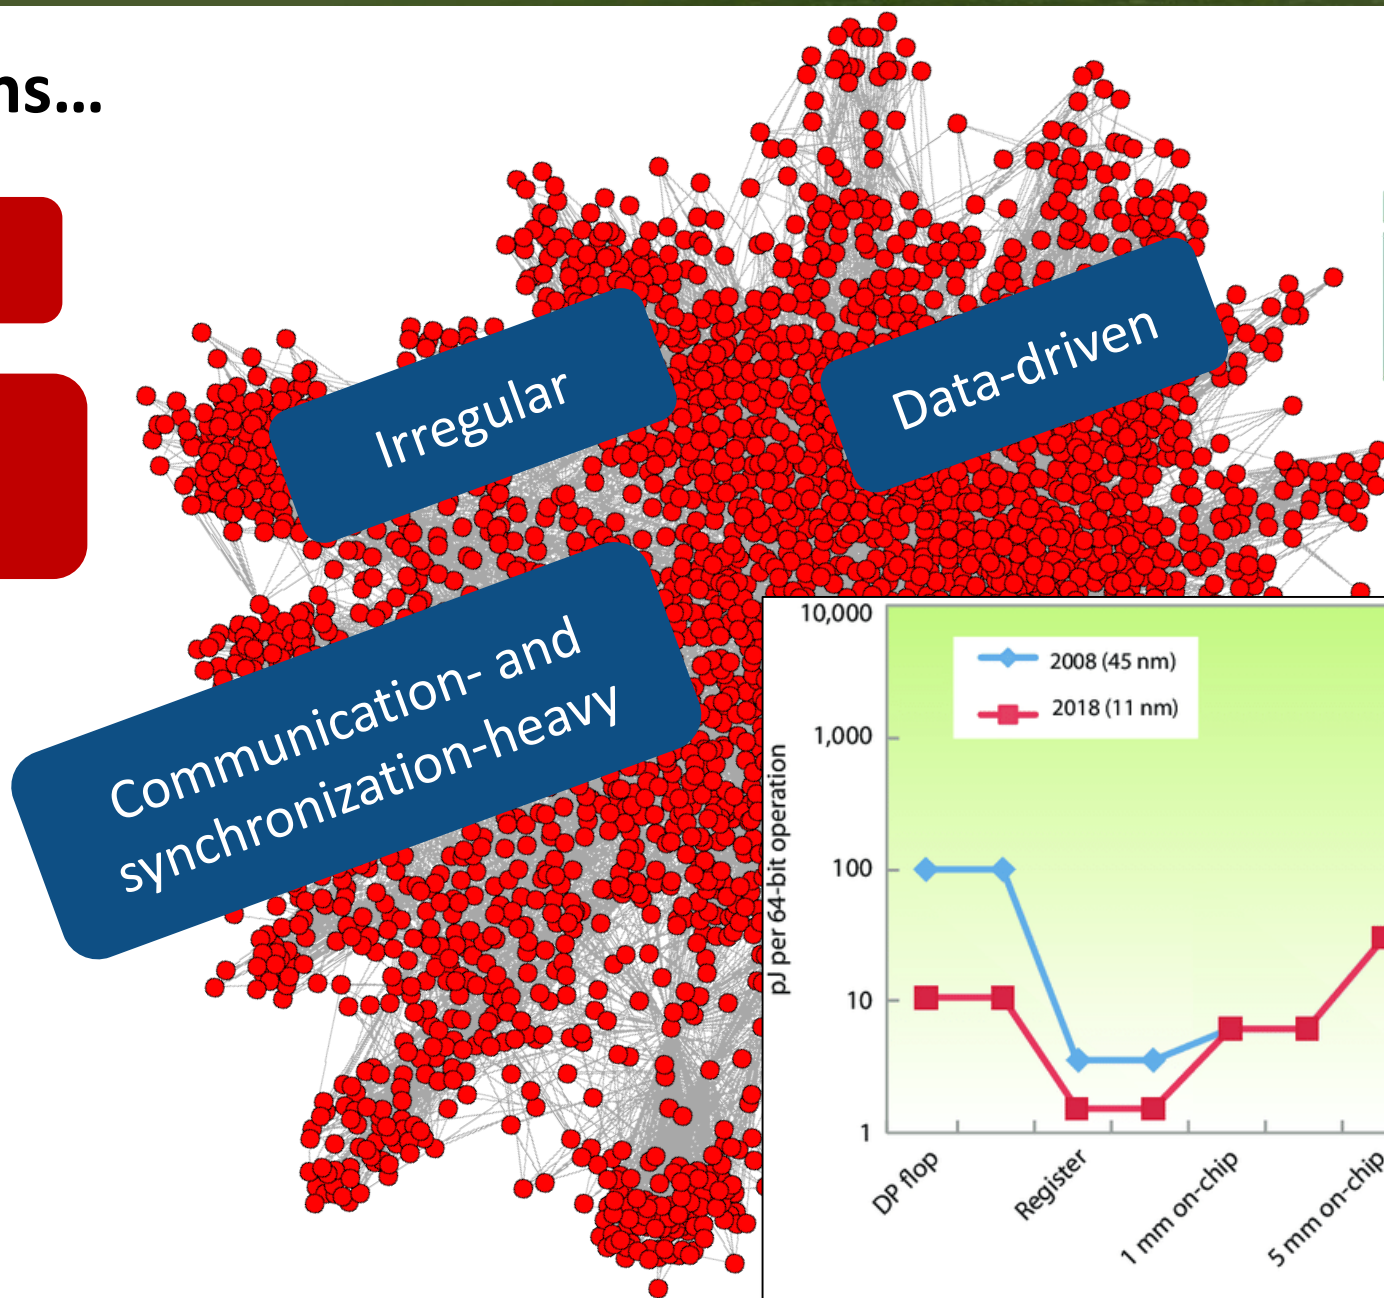
**✗**  
250  
Watts




# Large graphs...

**✗** Problems?


**✗** Low power efficiency!



**✗**  
120  
Watts



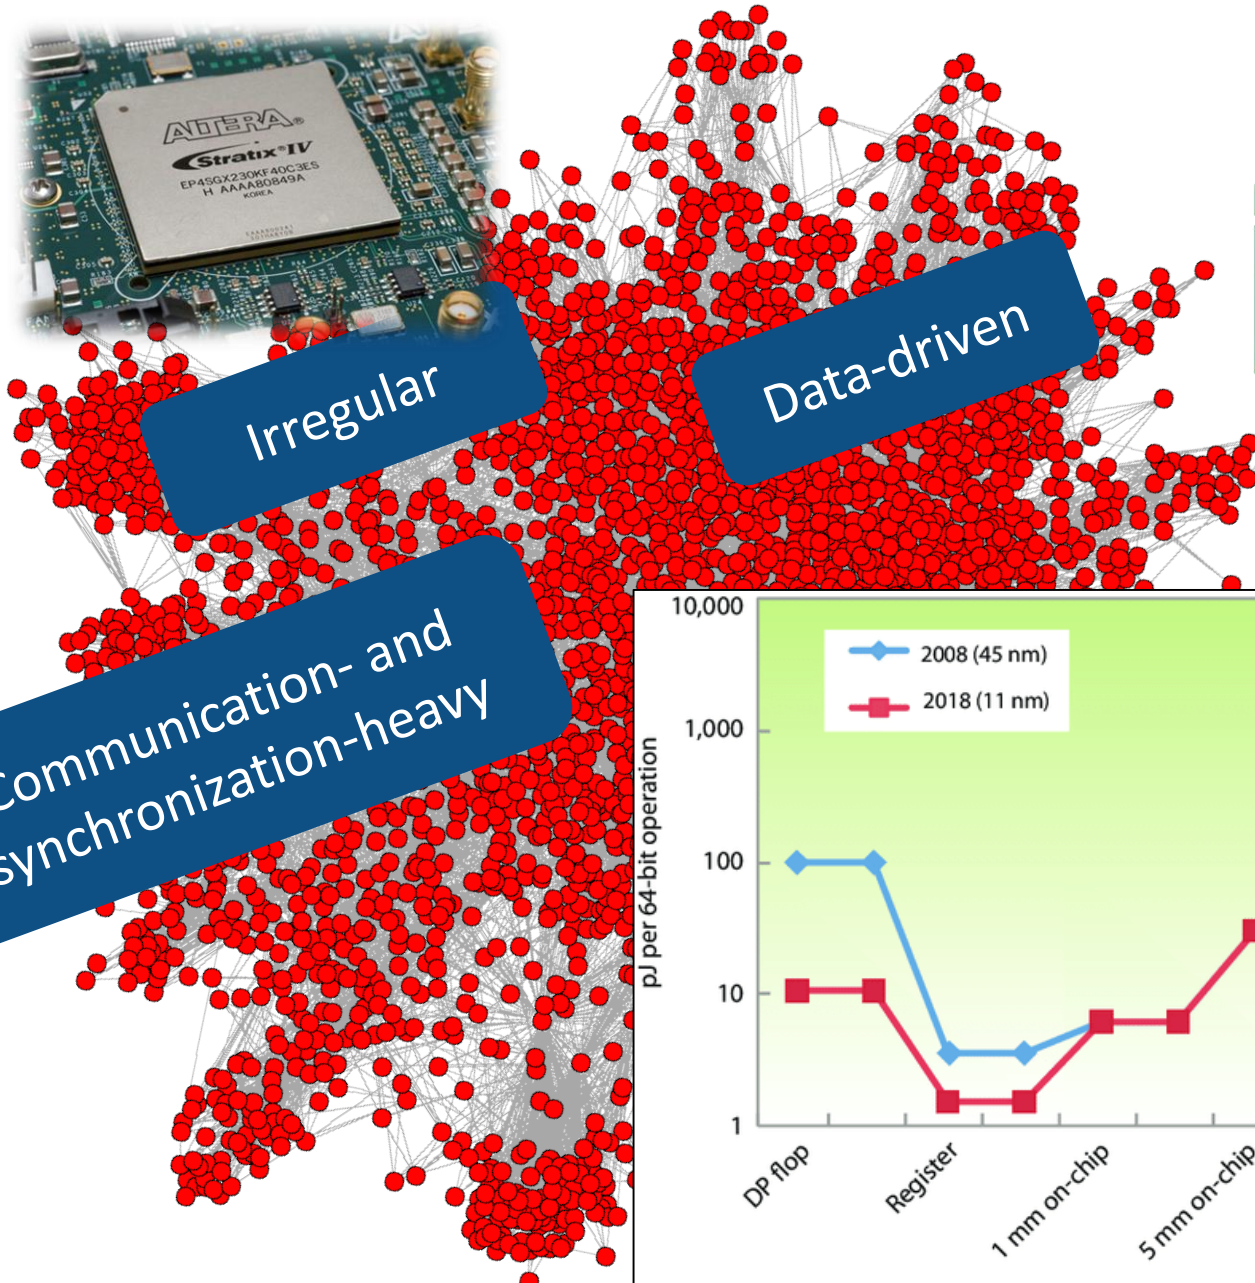
**✗**  
250  
Watts



# Large graphs...

**✗** Problems?

**✗** Low power efficiency!




**✗**

120 Watts

**✗**

250 Watts


**Communication- and synchronization-heavy**



# Large graphs...

**✗** Problems?


**✗** Low power efficiency!



**✗**  
120 Watts

**✗**  
250 Watts

**Comm. and synchron.**




Graph Problem	CPU (MTEPS/Watt) [1]	FPGA (MTEPS/Watt) [2]
SSSP	1.9	30.2
CC	0.5	48.1
MST	0.6	44.3

[1] A. Roy et al. X-stream: Edge-Centric Graph Processing using Streaming Partitions. ACM Symposium on Operating Syst. 2013.  
[2] S. Zhou et al. High-throughput and Energy-efficient Graph Processing on FPGA. FCCM. 2016.

# Large graphs...

**✗** Problems?

**✗** Low power efficiency!



**✗**

120 Watts

**✗**

250 Watts

Comm... and  
synchron...

Graph Problem	CPU (MTEPS/Watt) [1]	FPGA (MTEPS/Watt) [2]
SSSP	1.9	30.2
CC	0.5	48.1
MST	0.6	44.3

[1] A. Roy et al. X-stream: Edge-Centric Graph Processing using Streaming Partitions. ACM Symposium on Operating Syst. 2013.  
[2] S. Zhou et al. High-throughput and Energy-efficient Graph Processing on FPGA. FCCM. 2016.




# Large graphs...

✗ Problems?


✗ Low power efficiency!



Let's use FPGAs for Maximum Matchings...




Comm... and  
synchr...




Graph Problem	CPU (MTEPS/Watt) [1]	FPGA (MTEPS/Watt) [2]
SSSP	1.9	30.2
CC	0.5	48.1
MST	0.6	44.3

[1] A. Roy et al. X-stream: Edge-Centric Graph Processing using Streaming Partitions. ACM Symposium on Operating Syst. 2013.  
[2] S. Zhou et al. High-throughput and Energy-efficient Graph Processing on FPGA. FCCM. 2016.


# Large graphs...




# Large graphs...




# Large graphs...




# Large graphs...



# Large graphs...




# Large graphs...



# Large graphs...




# Large graphs...



# Large graphs...


Why do we care?



# Large graphs...

Why do we care?

Useful model





# Large graphs...

Useful model

Why do we care?

## Social networks




# Large graphs...

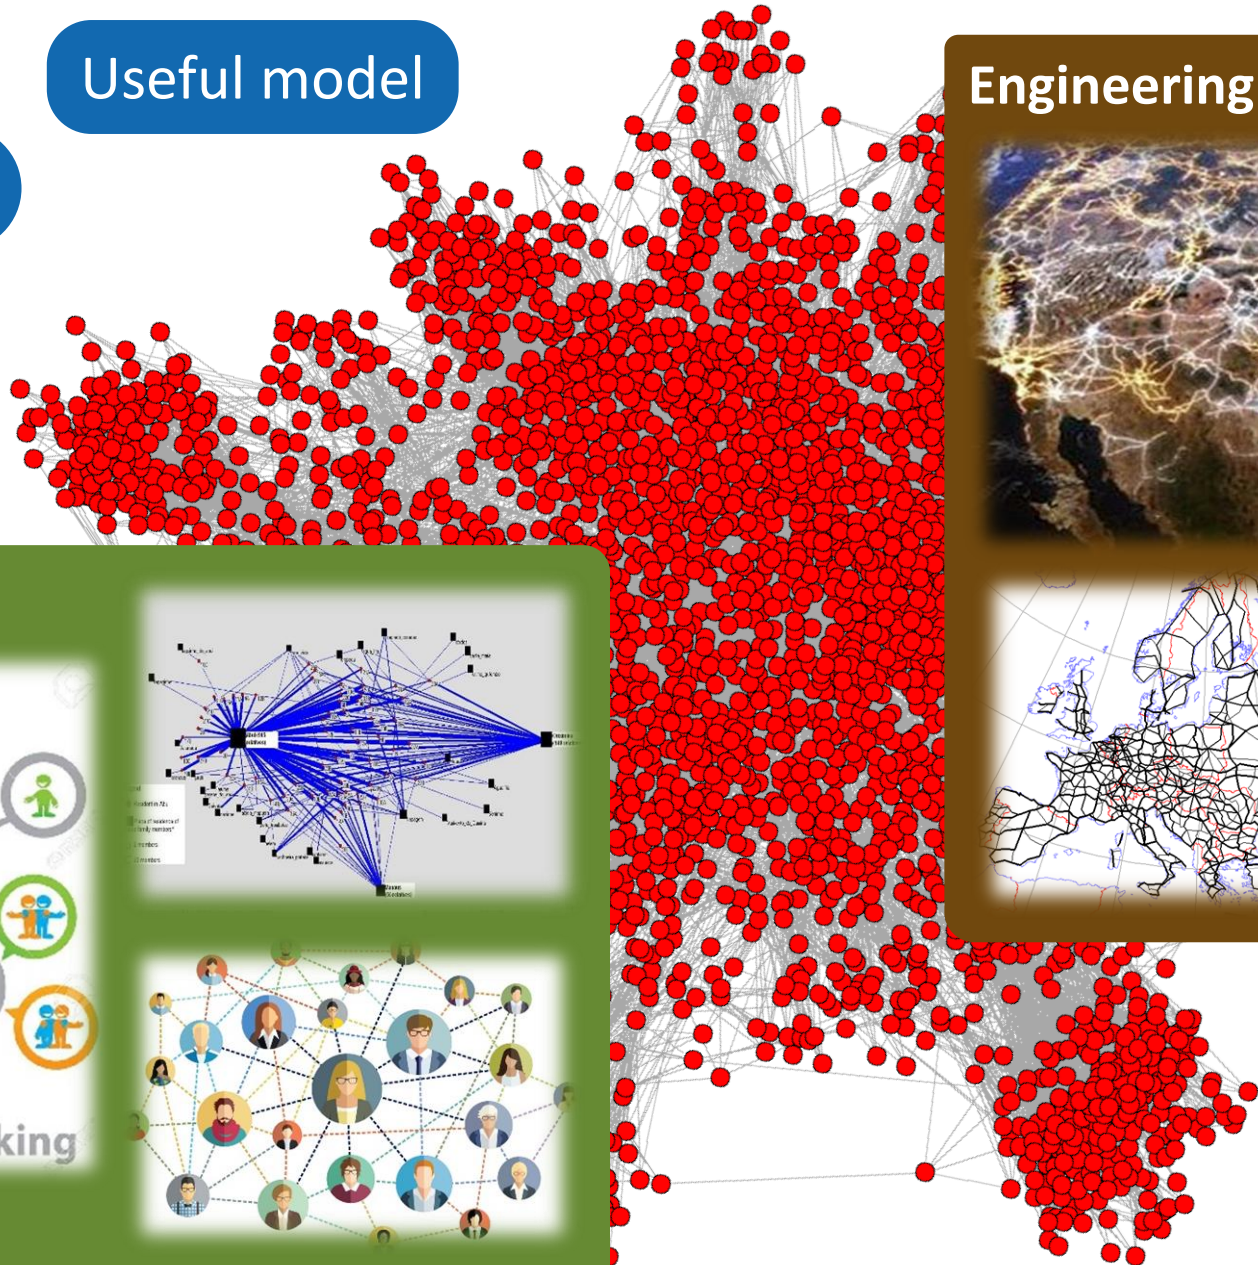
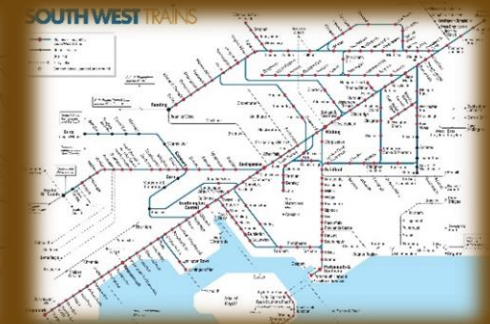
Useful model

Why do we care?

## Social networks



## Engineering networks



# Large graphs...




Useful model

Why do we care?

## Social networks



## Biological networks



## Engineering networks



# Large graphs...


# Useful model

# Why do we care?

# Social networks



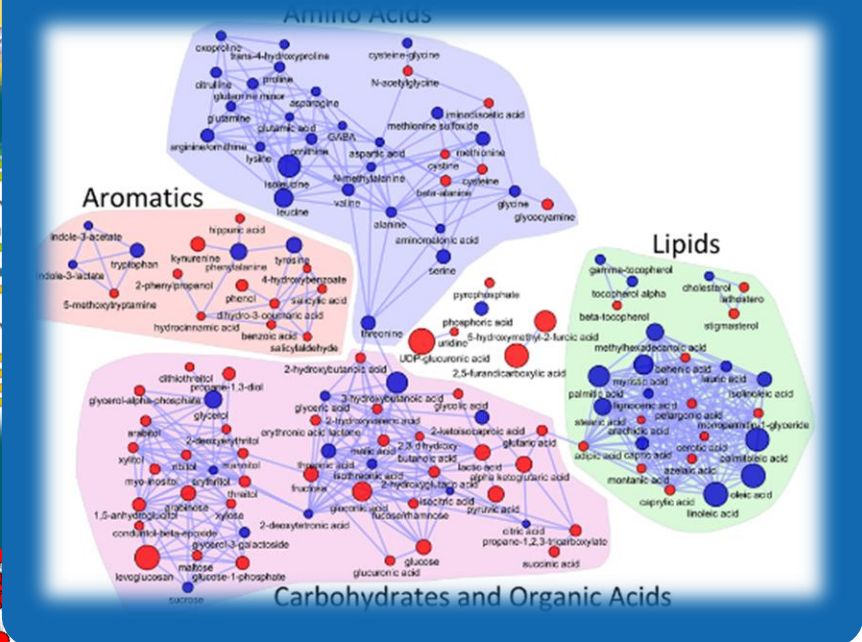
# Biological networks



# Engineering networks

# Physics, chemistry

$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



# Large graphs...

Why do we care?

Useful model

Engineering networks

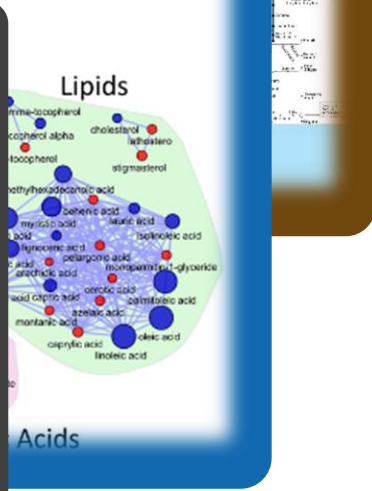
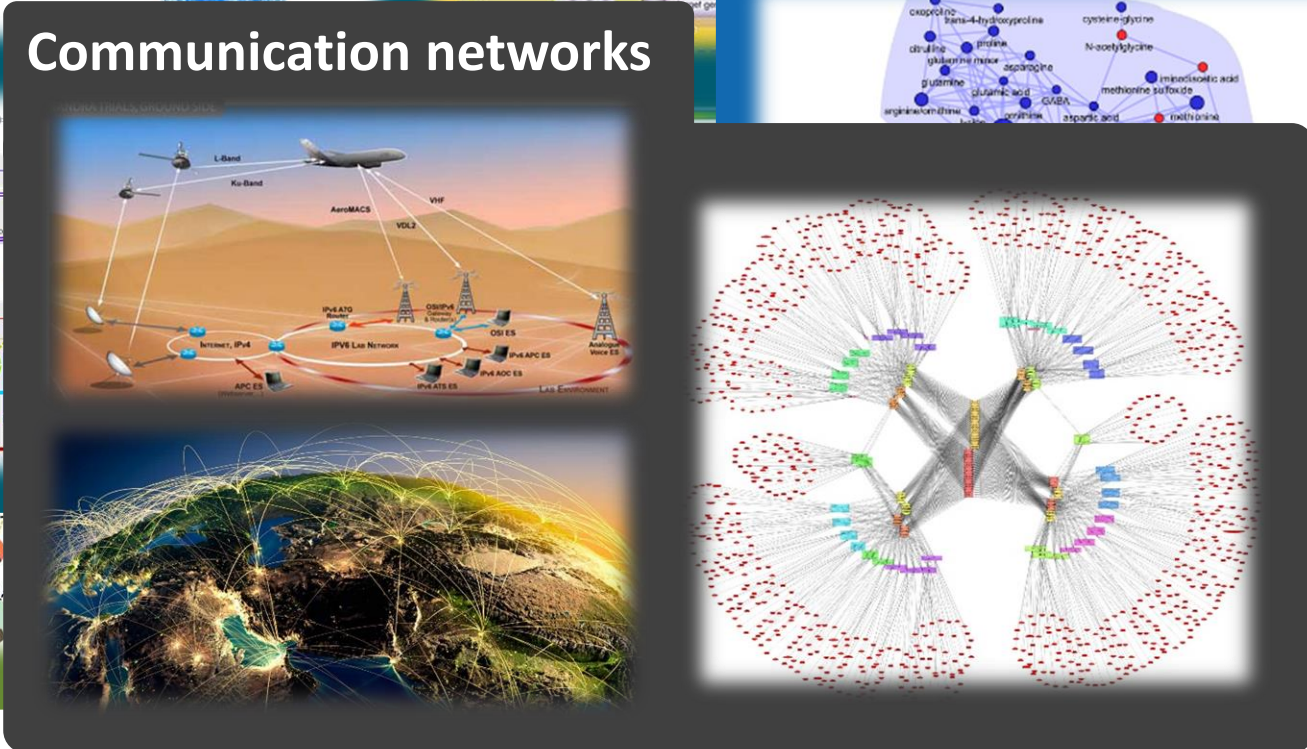


Physics, chemistry

Biological networks

Social networks

Communication networks

Networking





# Large graphs...

## Why do we care?

### Useful model

### Engineering networks






### Physics, chemistry

### Biological networks


### Machine learning

### Social


### Communication networks



$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



Acids



Large graphs...

Why do we care?

Useful model

Engineering networks




...even philosophy ☺

Physics, chemistry

Biological networks

Machine learning

Social



FOSDEM 2016 / Schedule / Events / Developer rooms / Graph Processing / Modeling a Philosophical Inquiry: from MySQL to a graph database

## Modeling a Philosophical Inquiry: from MySQL to a graph database

### The short story of a long refactoring process

- 📍 Track: Graph Processing devroom
- 📍 Room: AW1.126
- 📅 Day: Saturday
- 🕒 Start: 12:45
- 🕒 End: 13:35

Bruno Latour wrote a book about philosophy (an inquiry into modes of existence). He decided that the paper book was no place for the numerous footnotes, documentation or glossary, instead giving access to all this information surrounding the book through a web application which would present itself as a reading companion. He also offered to the community of readers to submit their contributions to his inquiry by writing new documents to be added to the platform. The first version

Large graphs...

Why do we care?

Useful model

Engineering networks



...even philosophy 😊

Physics, chemistry

Biological networks

Machine learning

Social



FOSDEM 2016 / [Schedule](#) / [Events](#) / [Developer rooms](#) / [Graph Processing](#) / Modeling a Philosophical Inquiry: from MySQL to a graph database

## Modeling a Philosophical Inquiry: from MySQL to a graph database

### The short story of a long refactoring process

-  **Track:** Graph Processing devroom
-  **Room:** AW1.126
-  **Day:** Saturday
-  **Start:** 12:45
-  **End:** 13:35



Bruno Latour wrote a book about philosophy (an inquiry into modes of existence). He decided that the paper book was no place for the numerous footnotes, documentation or glossary, instead giving access to all this information surrounding the book through a [web application](#) which would present itself as a reading companion. He also offered to the community of readers to submit their contributions to his inquiry by writing new documents to be added to the platform. The first version

# Large graphs...

Why do we care?

## Useful model

? Shortest network path?

## Engineering networks

? Gene alignment?

? Most reliable grid network?

## Machine learning

? Best phone connection?

? Disease spread channels?

? Least expensive computer network?

## Social

? Collaborator experience?

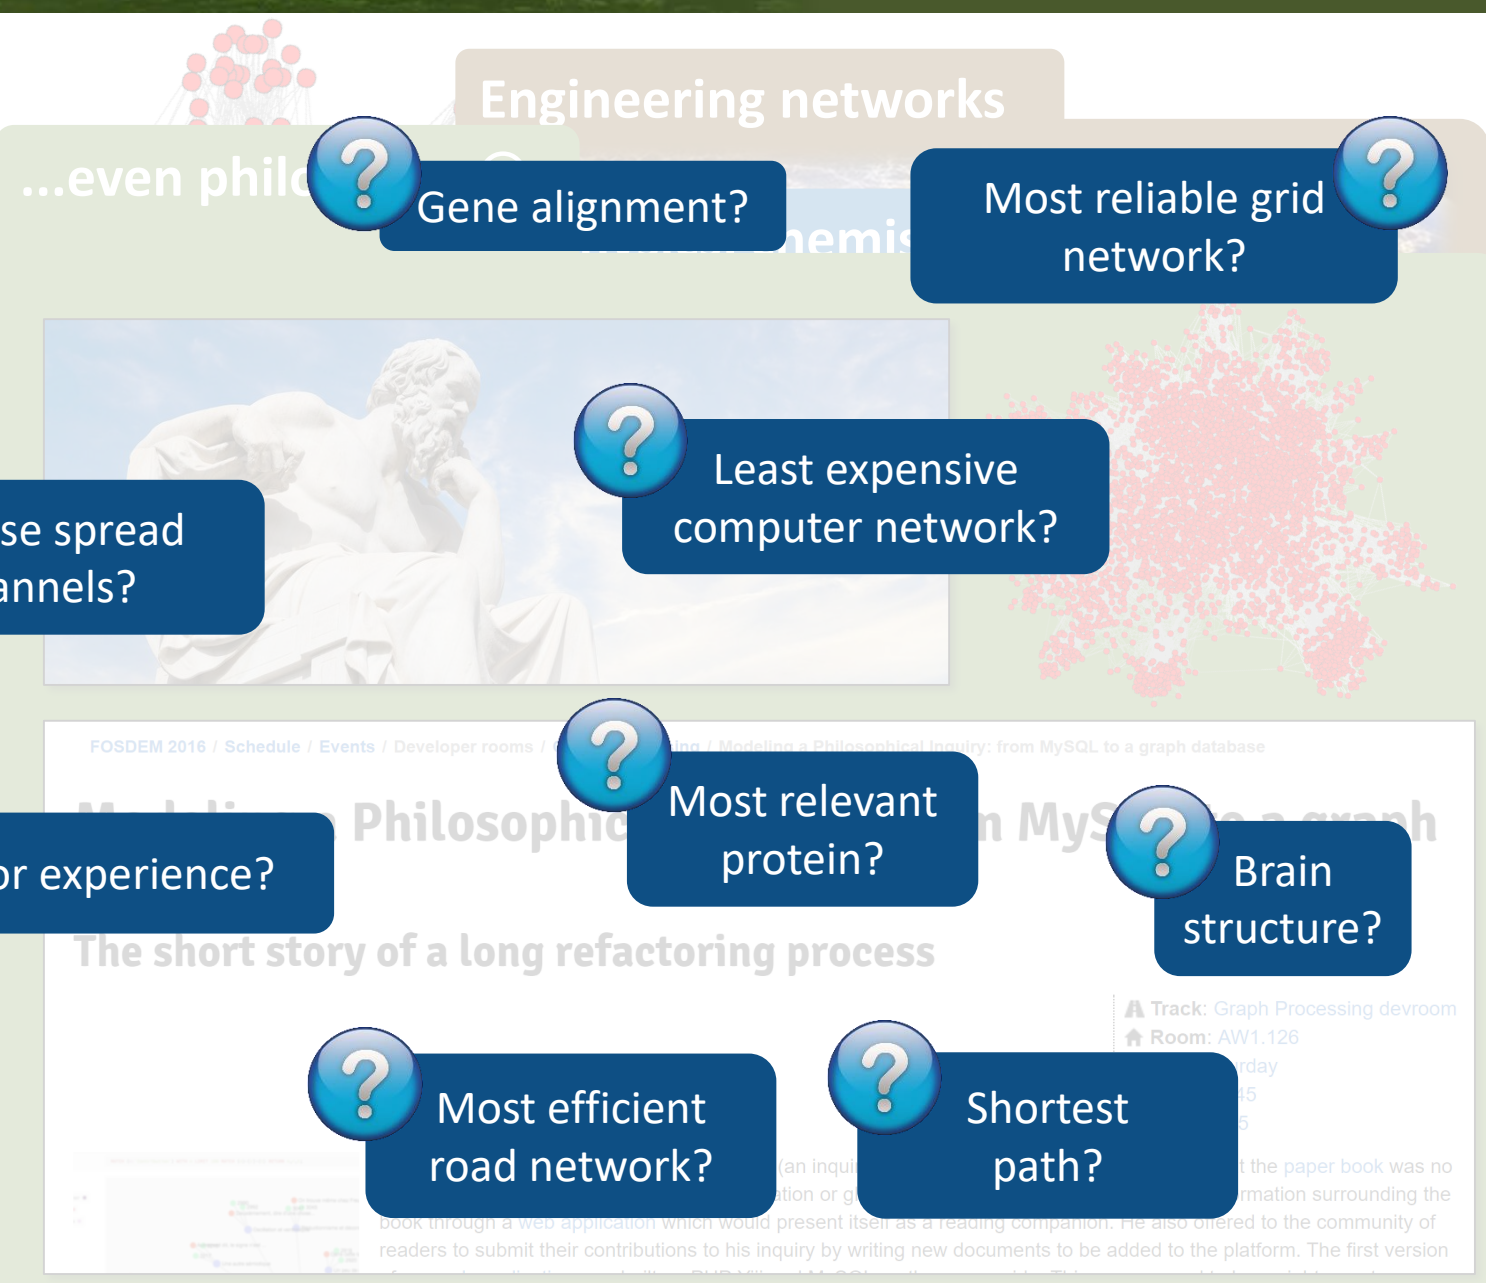
? Most relevant protein?

? Brain structure?


? Terrorism prevention?


? Most efficient road network?

? Shortest path?




...even philo  
Gene alignment?  
chemis  
Most reliable grid network?  
Least expensive computer network?  
FOSDEM 2016 / Schedule / Events / Developer rooms / ... / Modeling a Philosophical Inquiry: from MySQL to a graph database  
Philosophical Inquiry in MySQL  
The short story of a long refactoring process  
Track: Graph Processing devroom  
Room: AW1.126  
...the paper book was no  
...information surrounding the  
...to the community of  
...readers to submit their contributions to his inquiry by writing new documents to be added to the platform. The first version





What programming paradigm and why?



? What programming paradigm and why?

? What are the most promising techniques?



What programming paradigm and why?




What are the most promising techniques?

Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs



**Part 1:** To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs




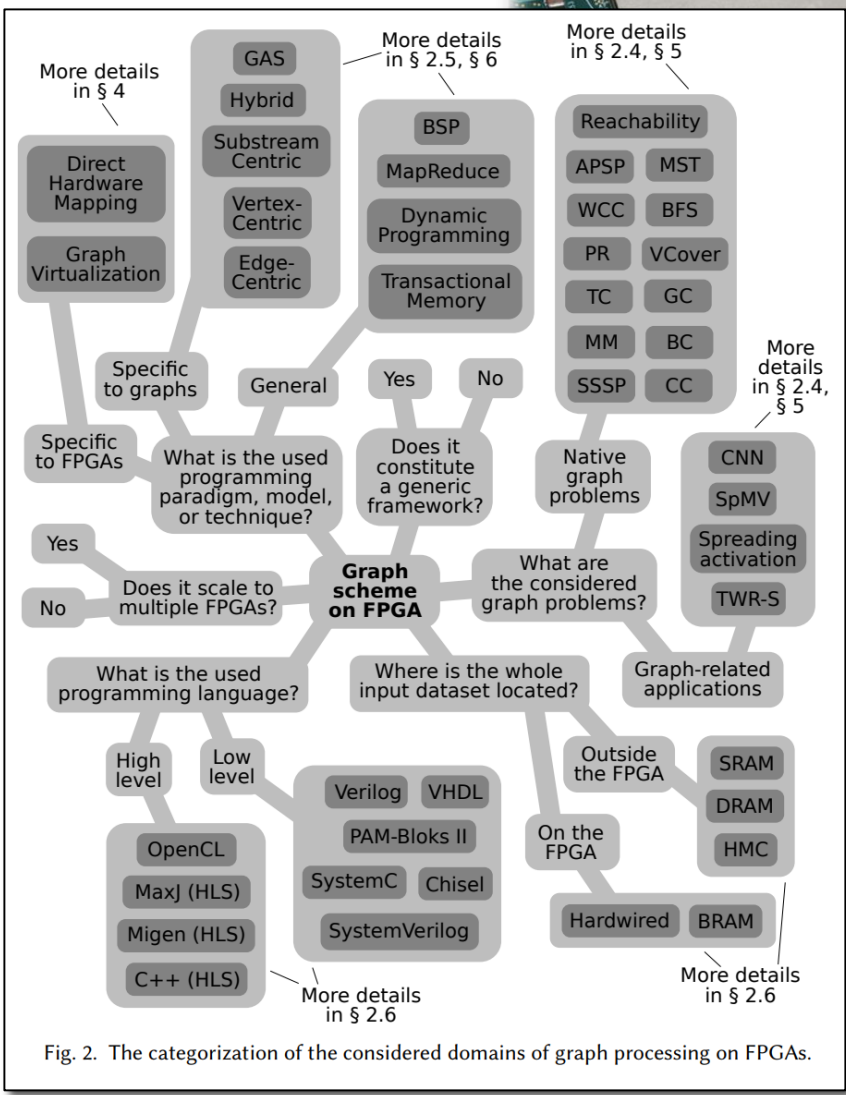
? What programming paradigm and why?

? What are the most promising techniques?

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

? What programming paradigm and why?

? What are the most promising techniques?



7 paradigms

Fig. 2. The categorization of the considered domains of graph processing on FPGAs.

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

? What programming paradigm and why?

? What are the most promising techniques?

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k
Kapre [70] (GraphSoC)	ASAP'15	👍	SpMV	Vertex-Centric, BSP	C++ (HLS)	👍	BRAM	17k	126k
Dai [40] (FPGP)	FPGA'16	👍	BFS	None	unsp.	👍	DRAM	41.6M	1.4B
Oguntebi [93] (GraphOps)	FPGA'16	👍	BFS, SpMV, PR, Vertex Cover	None	MaxJ (HLS)	👎	BRAM	16M	128M
Zhou [134]	FCCM'16	👍	SSSP, WCC, MST	Edge-Centric	unsp.	👎	DRAM	4.7M	65.8M
Engelhardt [49] (GraVF)	FPL'16	👍	BFS, PR, SSSP, CC	Vertex-Centric	Migen (HLS)	👎	BRAM	128k	512k
Dai [41] (ForeGraph)	FPGA'17	👍	PR, BFS, WCC	None	unsp.	👍	DRAM	41.6M	1.4B
Zhou [136]	SBAC-PAD'17	👍	BFS, SSSP	Hybrid (Vertex- and Edge-Centric)	unsp.	👎	DRAM	10M	160M
Ma [85]	FPGA'17	👍	BFS, SSSP, CC, TC, BC	Transactional Memory [16, 59]	System-Verilog	👍	DRAM	24M	58M
Lee [79] (ExtraV)	FPGA'17	👍	BFS, PR, CC, AT* [60]	Graph Virtualization	C++ (HLS)	👎	DRAM	124M	1.8B
Zhou [135]	CF'18	👍	SpMV, PR	Edge-Centric, GAS	unsp.	👎	DRAM	41.6M	1.4B
Yang [125]	report (2018)	👍	BFS, PR, WCC	None	OpenCL	👎		4.85M	69M
Yao [127]	report (2018)	👍	BFS, PR, WCC	None	unsp.	👎	BRAM	4.85M	69M

7 paradigms

~15 FPGA graph processing frameworks

Fig. 2. The categorization of the considered domains of graph processing on FPGAs.

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

? What programming paradigm and why?

? What are the most promising techniques?

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k
Kapre [4] (GraphStep)	Babb [4] report (1996)	👎	SSSP	None	Verilog	👍	Hardwired	512	2051
Dandalis [43]	report (1999)	👎	SSSP	None	unsp.	👍	Hardwired	2048	32k
Dai [4] (FPGA)	Tommiska [116] report (2001)	👎	SSSP	None	VHDL	👎	BRAM	64	4096
Ogunt [87] (GraphStep)	Mencer [87] FPL'02	👎	Reachability, SSSP	None	PAM-Bloks II	👎	Hardwired (3-state buffers)	88	7744
Bondhugula [27]	IPDPS'06	👎	APSP	Dynamic Program.	unsp.	👎	DRAM	unsp.	unsp.
Sridharan [110]	TENCON'09	👎	SSSP	None	VHDL	👎	BRAM	64	88
Engell [121]	ICFTP'10	👎	BFS	None	SystemC	👎	DRAM	65.5k	1M
Wang [21] (GraphStep)	FTP'11	👎	GC	Vertex-Centric	Verilog	👍	DRAM	300k	3M
Dai [4] (Fore)	Jagadeesh [65] report (2011)	👎	SSSP	None	VHDL	👎	Hardwired	128	466
Betkaoui [22]	FPL'12	👎	APSP	Vertex-Centric	Verilog	👍	≈ DRAM	38k	72M
Betkaoui [23]	ASAP'12	👎	BFS	Vertex-Centric	Verilog	👍	DRAM	16.8M	1.1B
Zhou [2] (CyGraph)	Attia [2] IPDPS'14	👎	BFS	Vertex-Centric	VHDL	👍	DRAM	8.4M	536M
Ma [8]	Ni [91] report (2014)	👎	BFS	None	Verilog	👎	DRAM, SRAM	16M	512M
Lee [7] (Extra)	Zhou [132] IPDPS'15	👎	SSSP	None	unsp.	👎	DRAM	1M	unsp.
Zhou [133]	ReConFig'15	👎	PR	Edge-Centric	unsp.	👎	DRAM	2.4M	5M
Umuroglu [117]	FPL'15	👎	BFS	None	Chisel	👎	≈ DRAM	2.1M	65M
Yang [80]	report (2016)	👎	SSSP	None	unsp.	👎	DRAM	23.9M	58.2M
Yao [129]	Zhang [129] FPGA'17	👎	BFS	MapReduce	unsp.	👎	HMC	33.6M	536.9M
Zhang [130]	FPGA'18	👎	BFS	None	unsp.	👎	HMC		
Kohram [76]	FPGA'18	👎	BFS	None	unsp.	👎	HMC		
Besta [13]	FPGA'19	👎	MM	Substream-Centric	Verilog	👎	DRAM	4.8M	117M

Fig. 2. The categorization of the considered domains of graph processing on FPGAs.

7 paradigms

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

? What programming paradigm and why?

? What are the most promising techniques?

Key techniques, challenges, features, ...

7 paradigms

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Considered Problems <sup>2</sup> (§ 2.4)	Programming Model or Technique <sup>4</sup> (§ 2.5)	Used Language	Multi FPGAs <sup>4</sup>	Input Location <sup>5</sup>	$n^\dagger$	$m^\dagger$
Kapre [71] (GraphStep)	FCCM'06	👍	spreading activation* [82]	BSP	unsp.	👍	BRAM	220k	550k
Weisz [92] (GraphGen)	FCCM'14	👍	TRW-S*, CNN* [112]	Vertex-Centric	unsp.	👎	DRAM	110k	221k
Kapre [71] (GraphStep)	Babb [4] report (1996)	👎	SSSP	None	Verilog	👍	Hardwired	512	2051
(GraphStep)	Dandalis [43] report (1999)	👎	SSSP	None	unsp.	👍	Hardwired	2048	32k
(FPGA)	Dai [4] Tommiska [116] report (2001)	👎	SSSP	None	VHDL	👎	BRAM	64	4096
(GraphStep)	Ogunniyi [87] FPL'02	👎	Reachability, SSSP	None	PAM-Bloks II	👎	Hardwired (3-state buffers)	88	7744
(GraphStep)	Bondhugula [27] IPDPS'06	👎	APSP	Dynamic Program.	unsp.	👎	DRAM	unsp.	unsp.
(GraphStep)	Sridharan [110] TENCN'09	👎	SSSP	None	VHDL	👎	BRAM	64	88
(GraphStep)	Engell Wang [121] ICFTP'10	👎	BFS	None	SystemC	👎	DRAM	65.5k	1M
(GraphStep)	Betkaoui [21] FTP'11	👎	GC	Vertex-Centric	Verilog	👍	DRAM	300k	3M
(Fore)	Dai [4] Jagadeesh [65] report (2011)	👎	SSSP	None	VHDL	👎	Hardwired	128	466
(Fore)	Betkaoui [22] FPL'12	👎	APSP	Vertex-Centric	Verilog	👍	≈ DRAM	38k	72M
(Fore)	Zhou Betkaoui [23] ASAP'12	👎	BFS	Vertex-Centric	Verilog	👍	DRAM	16.8M	1.1B
(CyGraph)	Ma [8] Attia [2] IPDPS'14	👎	BFS	Vertex-Centric	VHDL	👍	DRAM	8.4M	536M
(CyGraph)	Lee [7] Ni [91] report (2014)	👎	BFS	None	Verilog	👎	DRAM, SRAM	16M	512M
(Extra)	Zhou [132] IPDPS'15	👎	SSSP	None	unsp.	👎	DRAM	1M	unsp.
(Extra)	Zhou [133] ReConFig'15	👎	PR	Edge-Centric	unsp.	👎	DRAM	2.4M	5M
(Extra)	Yang Umuroglu [117] FPL'15	👎	BFS	None	Chisel	👎	≈ DRAM	2.1M	65M
(Extra)	Yao [1] Lei [80] report (2016)	👎	SSSP	None	unsp.	👎	DRAM	23.9M	58.2M
(Extra)	Zhang [129] FPGA'17	👎	BFS	MapReduce	unsp.	👎	HMC	33.6M	536.9M
(Extra)	Zhang [130] FPGA'18	👎	BFS	None	unsp.	👎	HMC		
(Extra)	Kohram [76] FPGA'18	👎	BFS	None	unsp.	👎	HMC		
(Extra)	Besta [13] FPGA'19	👎	MM	Substream-Centric	Verilog	👎	DRAM	4.8M	117M

Fig. 2. The categorization of the considered domains of graph processing on FPGAs.

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

Selected MWM-related parts are in the FPGA paper, the rest is in...

Reference (scheme name)	Venue	Generic Design <sup>1</sup>	Category	Algorithm	Language	Hardware	Area	Speed
Kapre [71] (GraphStep)	FCCM'06	👍	SSSP	None	Verilog	Hardwired	512	2051
Weisz [92] (GraphGen)	FCCM'14	👍	SSSP	None	unsp.	Hardwired	2048	32k
Kapre [4] (GraphStep)	report (1996)	👎	SSSP	None	VHDL	BRAM	64	4096
Dandalis [43]	report (1999)	👎	SSSP	None	unsp.	Hardwired	2048	32k
Dai [4] (FPGA)	Tommiska [116] report (2001)	👎	SSSP	None	VHDL	BRAM	64	4096
Ogunt [87] (GraphStep)	FPL'02	👎	Reachability, SSSP	None	PAM-Bloks II	Hardwired (3-state buffers)	88	7744
Bondhugula [27]	IPDPS'06	👎	APSP	Dynamic Program.	unsp.	DRAM	unsp.	unsp.
Sridharan [110]	TENCON'09	👎	SSSP	None	VHDL	BRAM	64	88
Engell [121]	ICFTP'10	👎	BFS	None	SystemC	DRAM	65.5k	1M
Betkaoui [21] (GraphStep)	FTP'11	👎	GC	Vertex-Centric	Verilog	DRAM	300k	3M
Dai [4] (Fore)	Jagadeesh [65] report (2011)	👎	SSSP	None	VHDL	Hardwired	128	466
Betkaoui [22]	FPL'12	👎	APSP	Vertex-Centric	Verilog	≈ DRAM	38k	72M
Zhou [23]	Betkaoui [23] ASAP'12	👎	BFS	Vertex-Centric	Verilog	DRAM	16.8M	1.1B
Ma [8] (CyGraph)	Attia [2] IPDPS'14	👎	BFS	Vertex-Centric	VHDL	DRAM	8.4M	536M
Lee [7] (Extra)	Ni [91] report (2014)	👎	BFS	None	Verilog	DRAM, SRAM	16M	512M
Zhou [132]	IPDPS'15	👎	SSSP	None	unsp.	DRAM	1M	unsp.
Zhou [133]	ReConFig'15	👎	PR	Edge-Centric	unsp.	DRAM	2.4M	5M
Umuroglu [117]	FPL'15	👎	BFS	None	Chisel	≈ DRAM	2.1M	65M
Lei [80]	report (2016)	👎	SSSP	None	unsp.	DRAM	23.9M	58.2M
Zhang [129]	FPGA'17	👎	BFS	MapReduce	unsp.	HMC	33.6M	536.9M
Zhang [130]	FPGA'18	👎	BFS	None	unsp.	HMC		
Kohram [76]	FPGA'18	👎	BFS	None	unsp.	HMC		
Besta [13]	FPGA'19	👎	MM	Substream-Centric	Verilog	DRAM	4.8M	117M

Fig. 2. The categorization of the considered domains of graph processing on FPGAs.

? What programming paradigm and why?

? What are the most promising techniques?

Key techniques, challenges, features, ...

7 paradigms

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

Selected MWM-related parts are in the FPGA paper, the rest is in...

? What programming paradigm and why?

? What are the most promising techniques?

Key techniques, challenges, features, ...

7 paradigms

~15 FPGA graph processing frameworks

~25 FPGA accelerators for specific algorithms

Reference (scheme name)	Venue	Generic Design <sup>1</sup>							
Kapre [71] (GraphStep)	FCCM'06	👍							
Weisz [92] (GraphGen)	FCCM'14	👍							
Kapre [4] (GraphStep)	Babb [4] report (1996)	👎	SSSP	None	Verilog	👍	Hardwired	512	2051
Dandalić [42]	report (1999)	👎	SSSP	None	unsp	👍	Hardwired	2048	22k
Dai [4] (FPGA)	Tomr								
Ogunt (Graph)	Menc								
Zhou (Graph)	Bond								
Engell (Graph)	Sridh								
Wang (Graph)	Wang								
Betka (Fore)	Betka								
Dai [4] (Fore)	Jagad								
Betka (Fore)	Betka								
Zhou (CyC)	Betka								
Ma [8] (CyC)	Attia								
Ni [9] (Extra)	Ni [9]								
Zhou (Extra)	Zhou								
Zhou (Extra)	Zhou								
Yang (Extra)	Umu								
Lei [8] (Extra)	Lei [8]								
Zhan (Extra)	Zhan								
Zhan (Extra)	Zhan								
Kohr (Extra)	Kohr								
Besta (Extra)	Besta								

## Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field Programmable Gate Arrays (FPGAs) can be an energy-efficient solution to deliver specialized hardware for graph processing. This is reflected by the recent interest in developing various graph algorithms and graph processing frameworks on FPGAs. To facilitate understanding of this emerging domain, we present the first survey and taxonomy on graph processing on FPGAs. Our survey describes and categorizes existing schemes

Fig. 2. The categorization of...

# Part 1: To understand the domain well, we conducted a detailed analysis of graph processing on FPGAs

Selected MWM-related parts are in the FPGA paper, the rest is in...

<https://arxiv.org/abs/...>

**1**

**Graph Processing on FPGAs: Taxonomy, Survey, Challenges**  
 Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
 JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas, such as machine learning, computational sciences, medical applications, social network analysis, and many others. Various graphs such as web or social networks may contain up to trillions of edges. The sheer size of such datasets, combined with the irregular nature of graph processing, poses unique challenges for the runtime and the consumed power. Field Programmable Gate Arrays (FPGAs) can be an energy-efficient solution to deliver specialized hardware for graph processing. This is reflected by the recent interest in developing various graph algorithms and graph processing frameworks on FPGAs. To facilitate understanding of this emerging domain, we present the first survey and taxonomy on graph processing on FPGAs. Our survey describes and categorizes existing schemes

? What programming paradigm and why?

? What are the most promising techniques?

Key techniques, challenges, features, ...


7 paradigms

~15 FPGA graph processing frameworks


~25 FPGA accelerators for specific algorithms

Fig. 2. The categorization of...






What programming paradigm and why?




What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?






What programming paradigm and why?



Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




? What programming paradigm and why?



Vertex-centric,  
Gather-Apply-  
Scatter, ... ?

? What programming paradigm and why?




Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




? What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




? What programming paradigm and why?




Vertex-centric,  
Gather-Apply-  
Scatter, ... ?





 What programming paradigm and why?




Vertex-centric,  
 Gather-Apply-  
 Scatter, ... ? 






 What programming paradigm and why?




Vertex-centric,  
 Gather-Apply-  
 Scatter, ... ? 


? What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




 What programming paradigm and why?




Vertex-centric,  
 Gather-Apply-  
 Scatter, ... ? 

Well...



What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




It was designed with the “batch” analytics in mind.

Well...



What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




It was designed with the “batch” analytics in mind.

Well...

? What programming paradigm and why?

Assumes the whole input graph is accessible... ❌

Vertex-centric, Gather-Apply-Scatter, ... ?



It was designed with the “batch” analytics in mind.


Well...

? What programming paradigm and why?

Assumes the whole input graph is accessible... ❌

...when in BRAM, size is severely limited ❌

Vertex-centric, Gather-Apply-Scatter, ... ?



It was designed with the “batch” analytics in mind.


...when in DRAM, accessing & pipelining become complex

? What programming paradigm and why?

Assumes the whole input graph is accessible...

...when in BRAM, size is severely limited

Vertex-centric, Gather-Apply-Scatter, ... ?





It was designed with the “batch” analytics in mind.

...when in DRAM, accessing & pipelining become complex


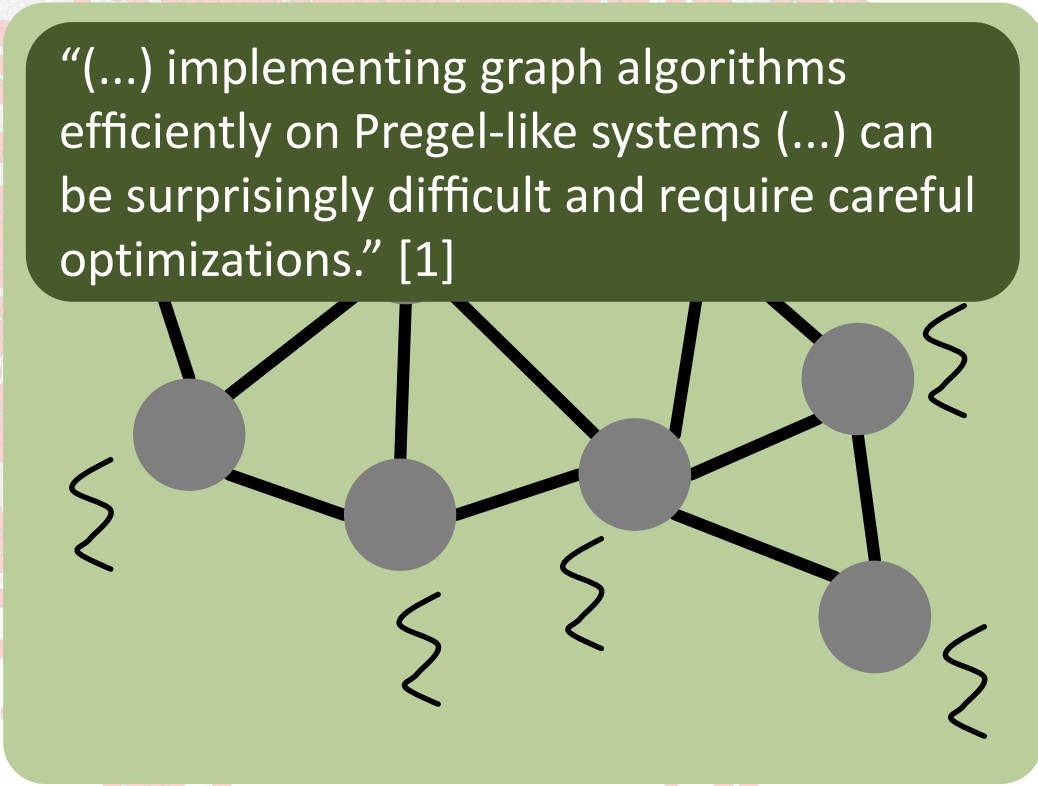
? What programming paradigm and why?

Assumes the whole input graph is accessible...

...when in BRAM, size is severely limited

Vertex-centric, Gather-Apply-Scatter, ... ?

“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]



[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.

It was designed with the “batch” analytics in mind.

...when in DRAM, accessing & pipelining become complex

? What programming paradigm and why?


Assumes the whole input graph is accessible...

...when in BRAM, size is severely limited


Vertex-centric, Gather-Apply-Scatter, ... ?

“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]


Vertex-Centric (aka Pregel-like) approach is complex for problems such as matchings, spanning trees, etc.



[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.



What programming paradigm and why?




What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?




? What programming paradigm and why?



Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



? What programming paradigm and why?





Vertex-centric,  
Gather-Apply-  
Scatter, ... ?





? What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?





 What programming paradigm and why?



Vertex-centric,  
 Gather-Apply-  
 Scatter, ... ? 





? What programming paradigm and why?




Vertex-centric,  
Gather-Apply-  
Scatter, ... ?





 What programming paradigm and why?




Vertex-centric,  
 Gather-Apply-  
 Scatter, ... ? 

? What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



? What programming paradigm and why?





Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



? What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?






- It was designed with the “batch” analytics in mind.
- It assumes the whole input is accesible. When in DRAM, accessing & pipelining becomes complex.



What programming paradigm and why?

Vertex-centric,  
Gather-Apply-  
Scatter, ... ?







- It was designed with the “batch” analytics in mind.
- It assumes the whole input is accesible. When in DRAM, accessing & pipelining becomes complex.



What programming paradigm and why?

“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]



Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.

- It was designed with the “batch” analytics in mind.
- It assumes the whole input is accesible. When in DRAM, accessing & pipelining becomes complex.




Vertex-centric,  
Gather-Apply-  
Scatter, ... ?



What programming  
paradigm and why?

“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]

Vertex-Centric (aka Pregel-like) approach is complex for problems such as matchings, spanning trees, etc.



[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.



- It was designed with the “batch” analytics in mind.
- It assumes the whole input is accesible. When in DRAM, accessing & pipelining becomes complex.



To be able to utilize pipelining well, we really want to use streaming (aka the edge-centric paradigm)




What programming paradigm and why?

Vertex-centric, Gather-Apply-Scatter, ... ?



“(...) implementing graph algorithms efficiently on Pregel-like systems (...) can be surprisingly difficult and require careful optimizations.” [1]

Vertex-Centric (aka Pregel-like) approach is complex for problems such as matchings, spanning trees, etc.



[1] S. Salihoglu and J. Widom, “Optimizing graph algorithms on Pregel-like systems”. VLDB. 2014.

Can be used but it was designed with the “batch” analytics in mind

...when in DRAM, accessing & parallelization become complex

? What programming paradigm and why?

Assumes the whole input graph is accessible...

...when in BRAM, size is severely limited

en	<a href="#">wikipedia edits (en)</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	50,757,442	572,591,272
TW	<a href="#">Twitter (WWW)</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	41,652,230	1,468,365,182
TF	<a href="#">Twitter (MPI)</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	52,579,682	1,963,263,821
FR	<a href="#">Friendster</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	68,349,466	2,586,147,869
UL	<a href="#">UK domain (2007)</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	105,153,952	3,301,876,564


KONECT graph datasets

Vertex-centric, Gather-Apply Scatter, ... ?




Webgraph datasets


Graph	Crawl date	Nodes	Arcs
<a href="#">uk-2014</a>	2014	787 801 471	47 614 527 250
<a href="#">eu-2015</a>	2015	1 070 557 254	91 792 261 600
<a href="#">gsh-2015</a>	2015	988 490 691	33 877 399 152
<a href="#">uk-2014-host</a>	2014	4 769 354	50 829 923
<a href="#">eu-2015-host</a>	2015	11 264 052	38 691 5963
<a href="#">gsh-2015-host</a>	2015	68 660 142	1 802 747 600
<a href="#">uk-2014-tpd</a>	2014	1 766 010	18 244 650
<a href="#">eu-2015-tpd</a>	2015	6 650 532	170 145 510
<a href="#">gsh-2015-tpd</a>	2015	30 809 122	602 119 716
<a href="#">clueweb12</a>	2012	978 408 098	42 574 107 469
<a href="#">uk-2002</a>	2002	18 520 486	298 113 762



What programming paradigm and why?



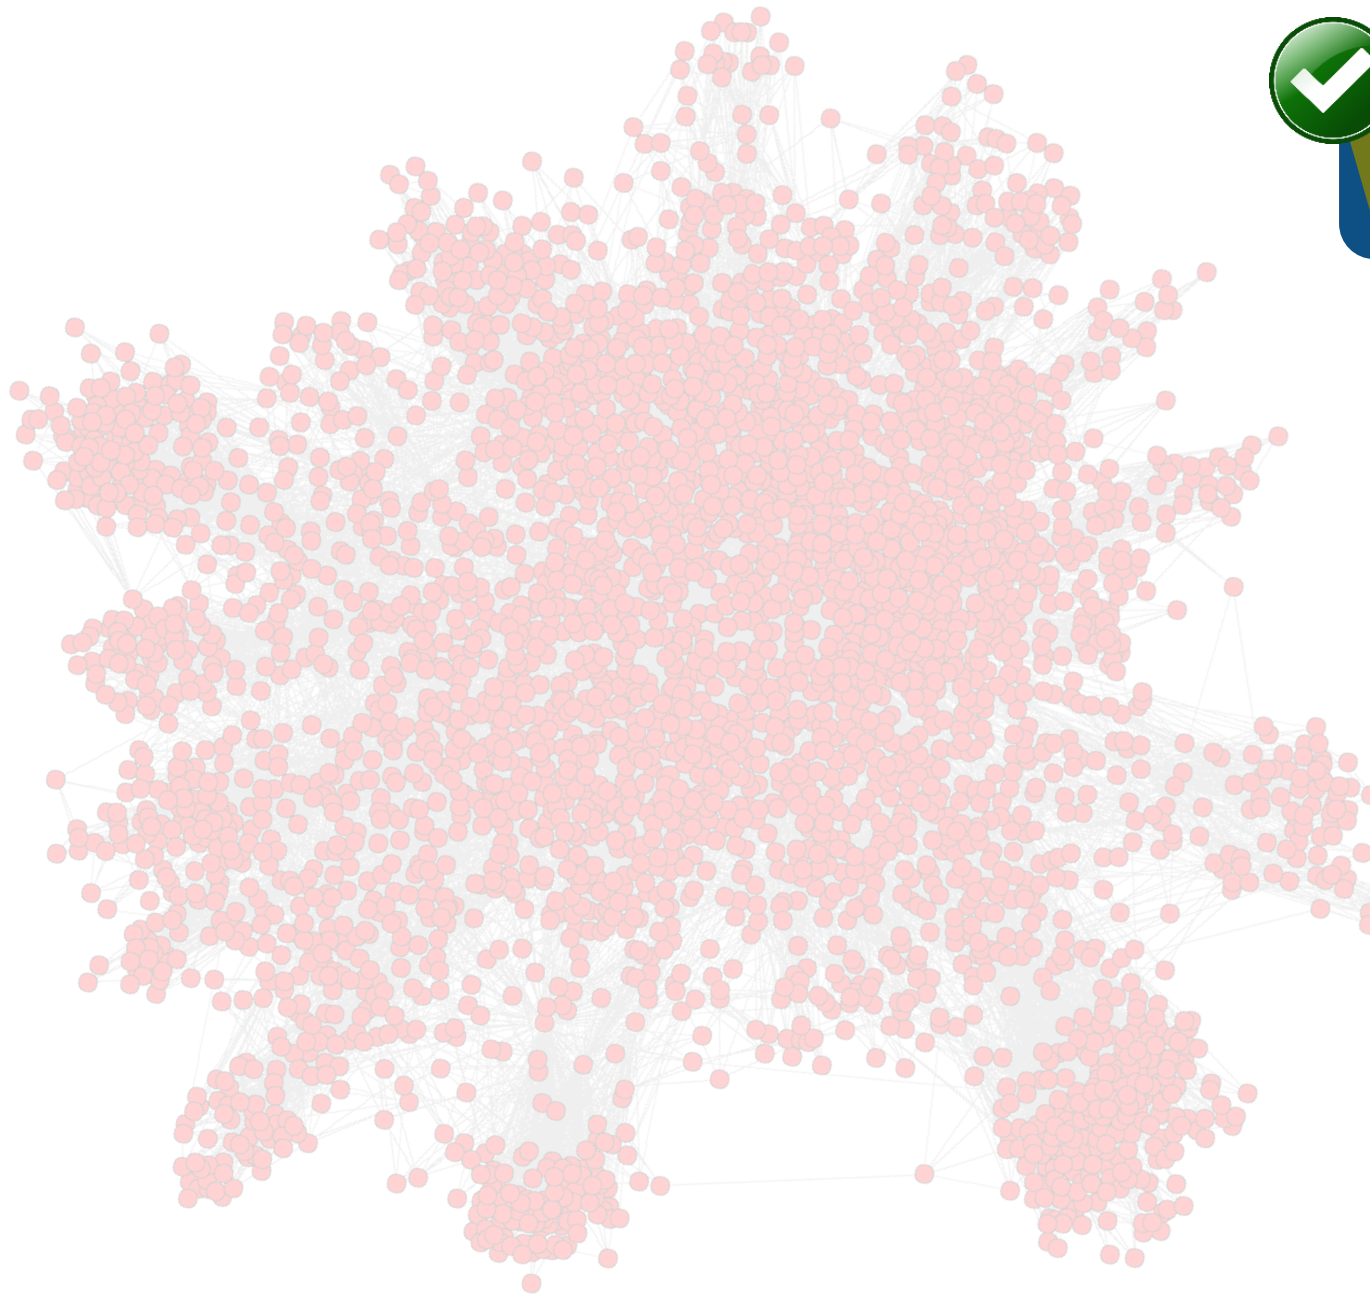
Use some form of  
streaming (aka  
edge-centric)




Use some form of streaming (aka edge-centric)



DRAM



Use some form of streaming (aka edge-centric)




DRAM



Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)





DRAM



Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)






DRAM



✓ Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)




DRAM

Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)

This green rounded rectangle represents the processing unit. It contains an image of an Altera Stratix IV FPGA chip, an image of an Intel Core i7 CPU chip, and a circular arrow icon indicating a processing or streaming operation.




This green rounded rectangle represents the DRAM. It contains an image of a DRAM chip with a cluster of red particles on top, symbolizing data storage or a specific data structure.

DRAM



Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)



DRAM

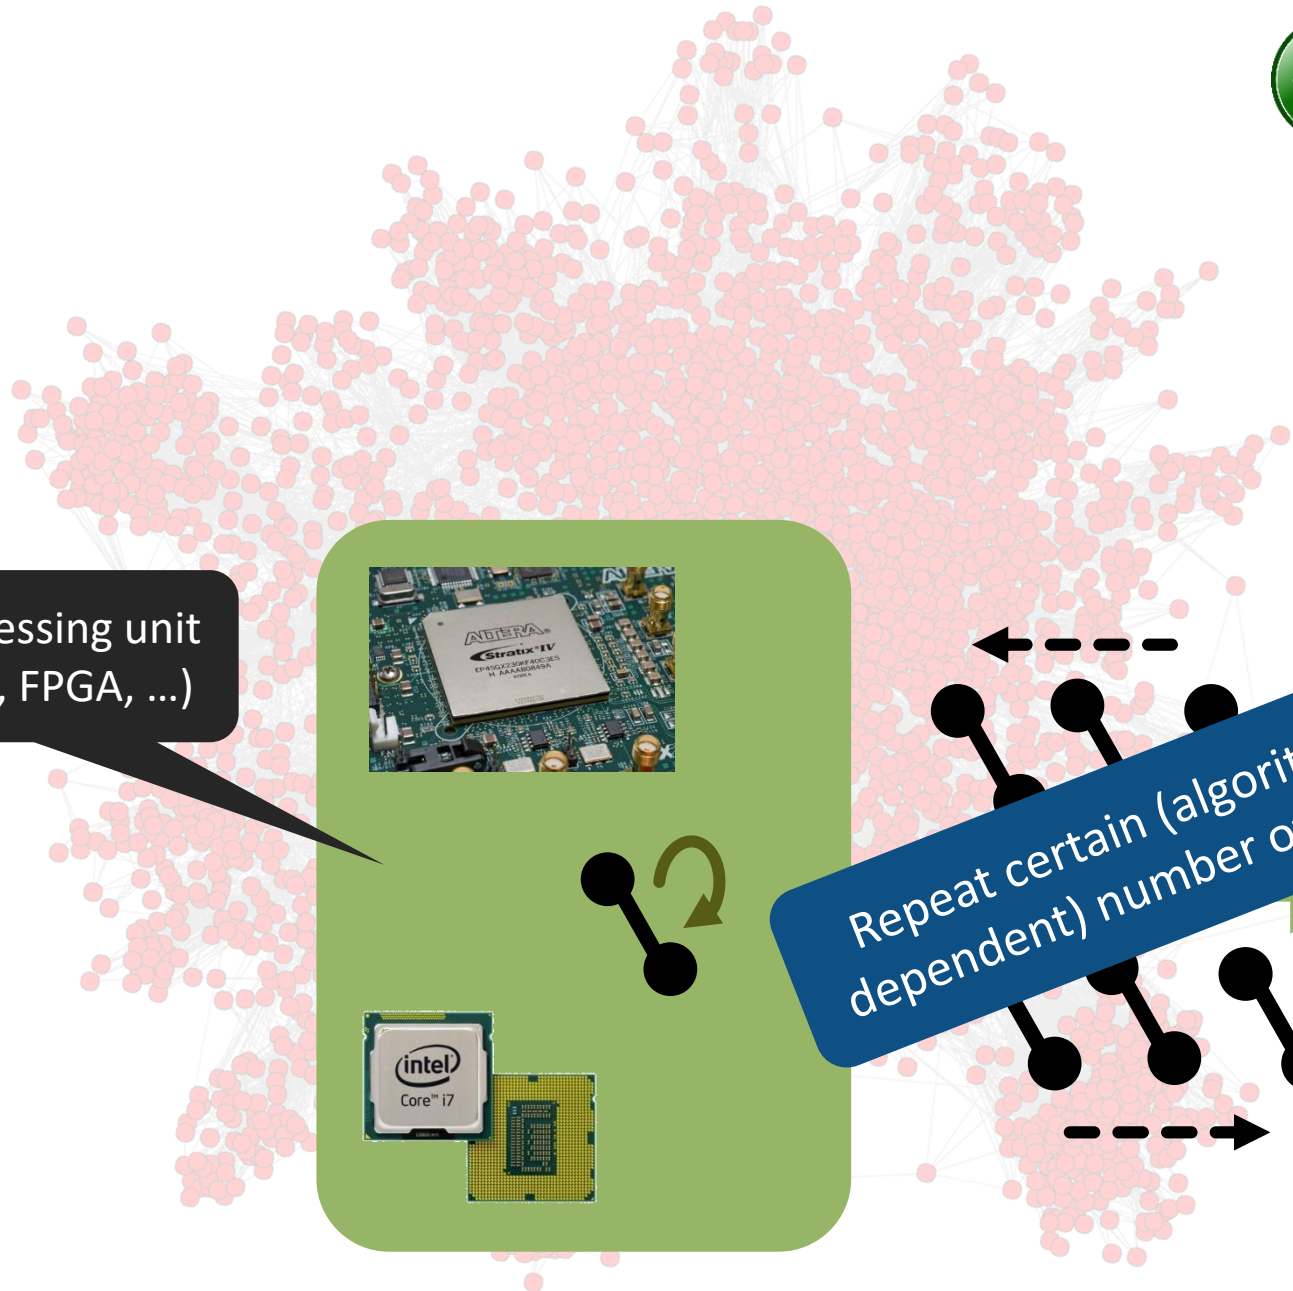


Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)

Repeat certain (algorithm-dependent) number of times


DRAM




How to implement efficiently on an FPGA?

Use some form of streaming (aka edge-centric)


Some processing unit (CPU, GPU, FPGA, ...)



Repeat certain (algorithm-dependent) number of times



DRAM




How to implement efficiently on an FPGA?


Processing edges is sequential – how to incorporate parallelism?

Use some form of streaming (aka edge-centric)

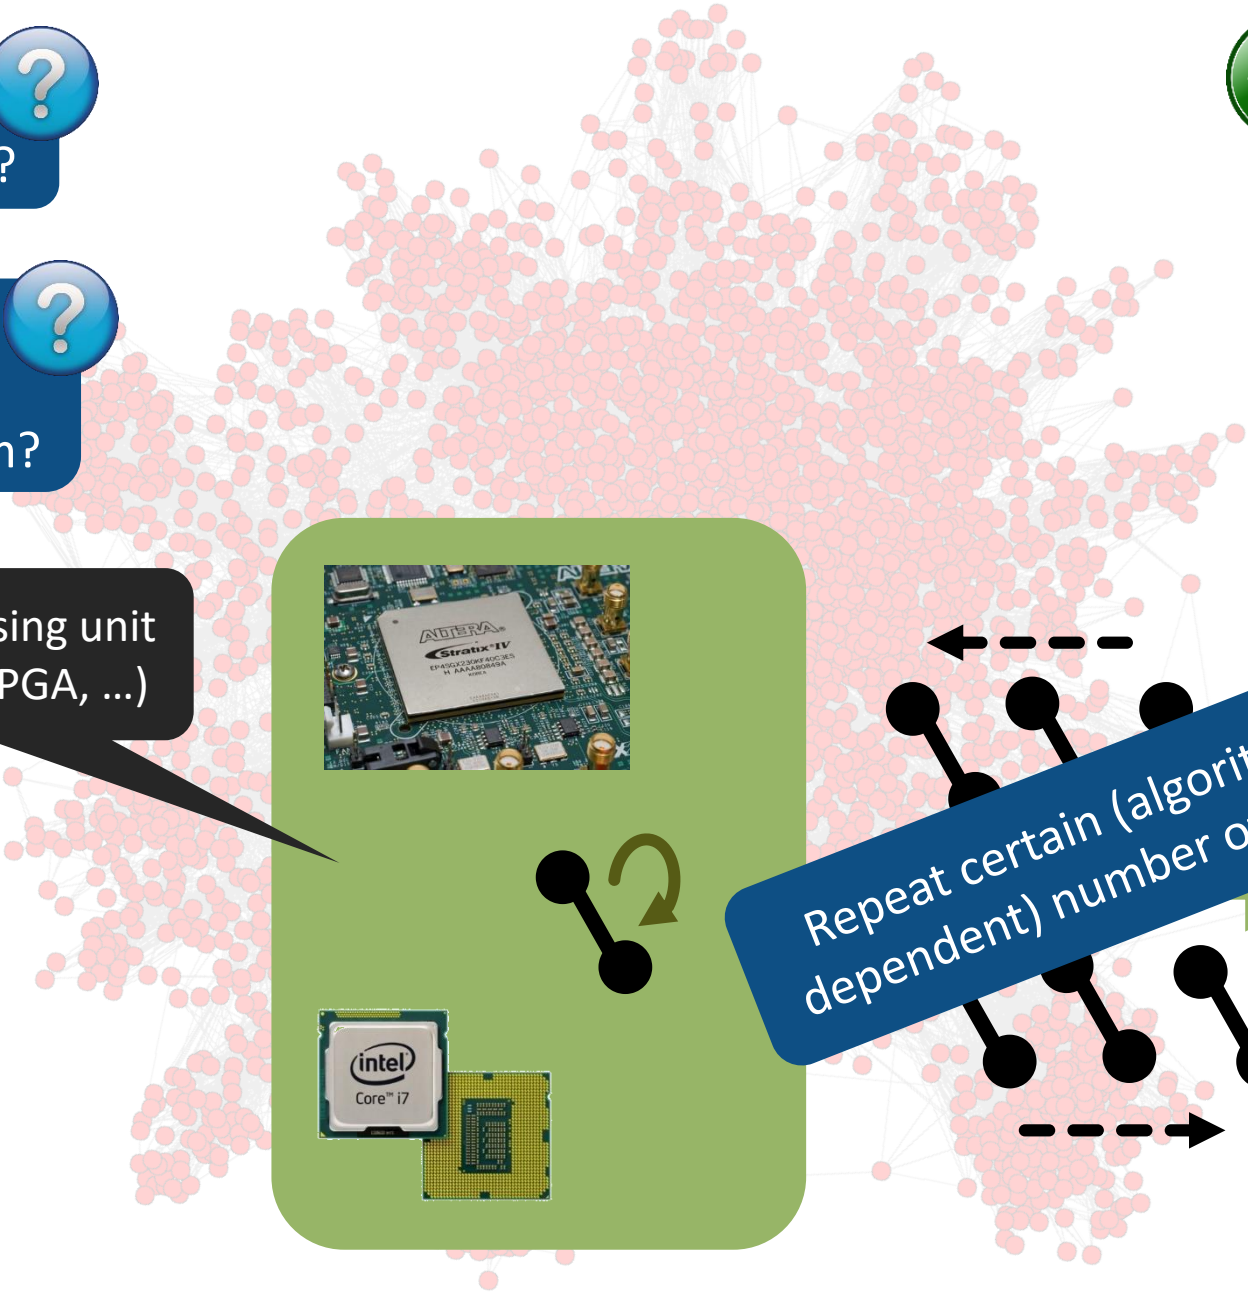
Some processing unit (CPU, GPU, FPGA, ...)



Repeat certain (algorithm-dependent) number of times



DRAM



How to implement efficiently on an FPGA?


Processing edges is sequential – how to incorporate parallelism?

How to minimize the number of “passes” over edges? (This can get really bad in the “traditional” edge-centric approach, e.g., BFS normally needs  $O(m+n)$  work, while in the edge-centric approach it takes  $O(D m)$  work ( $D$  passes [1]),


Use some form of streaming (aka edge-centric)

$m$ : #edges in a graph  
 $n$ : #vertices in a graph  
 $D$ : graph’s diameter (usually ~5-15)

Some processing unit (CPU, GPU, FPGA, ...)




Repeat certain (algorithm-dependent) number of times



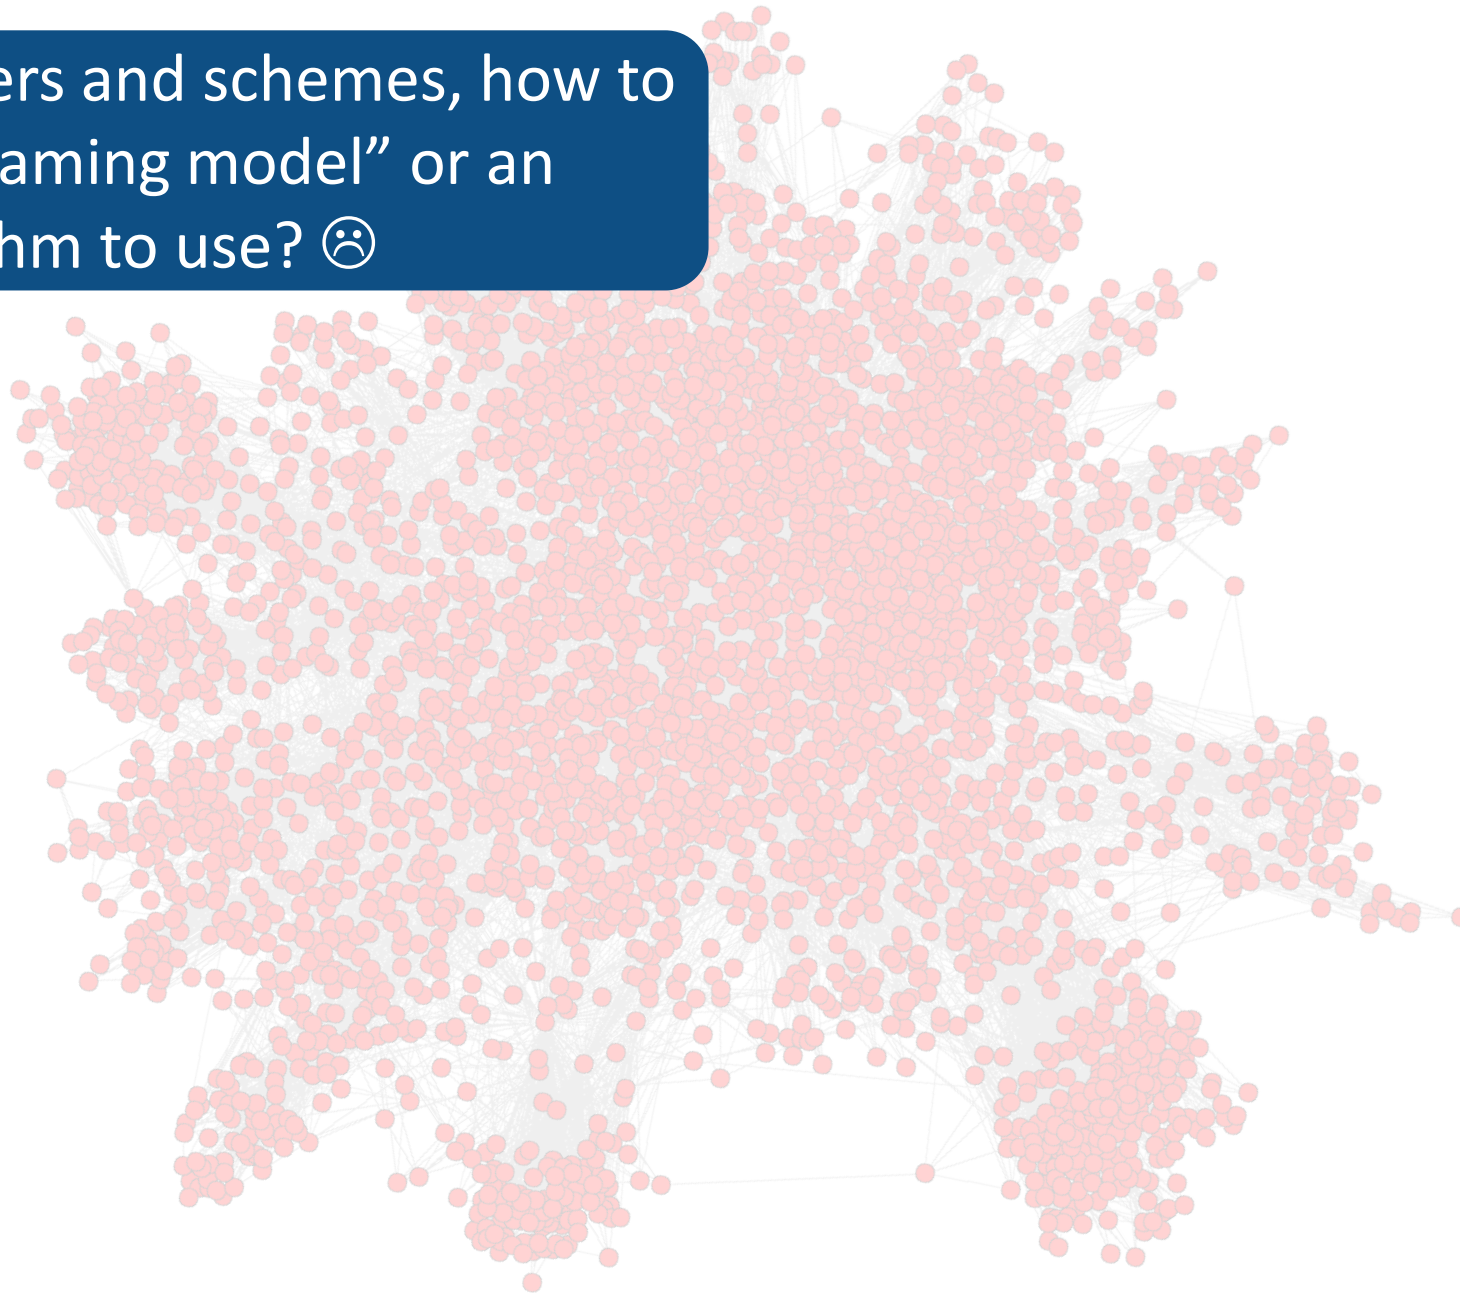
DRAM

[1] A. Roy et al. X-stream: Edge-Centric Graph Processing using Streaming Partitions. SOSP. 2013.





Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞



Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

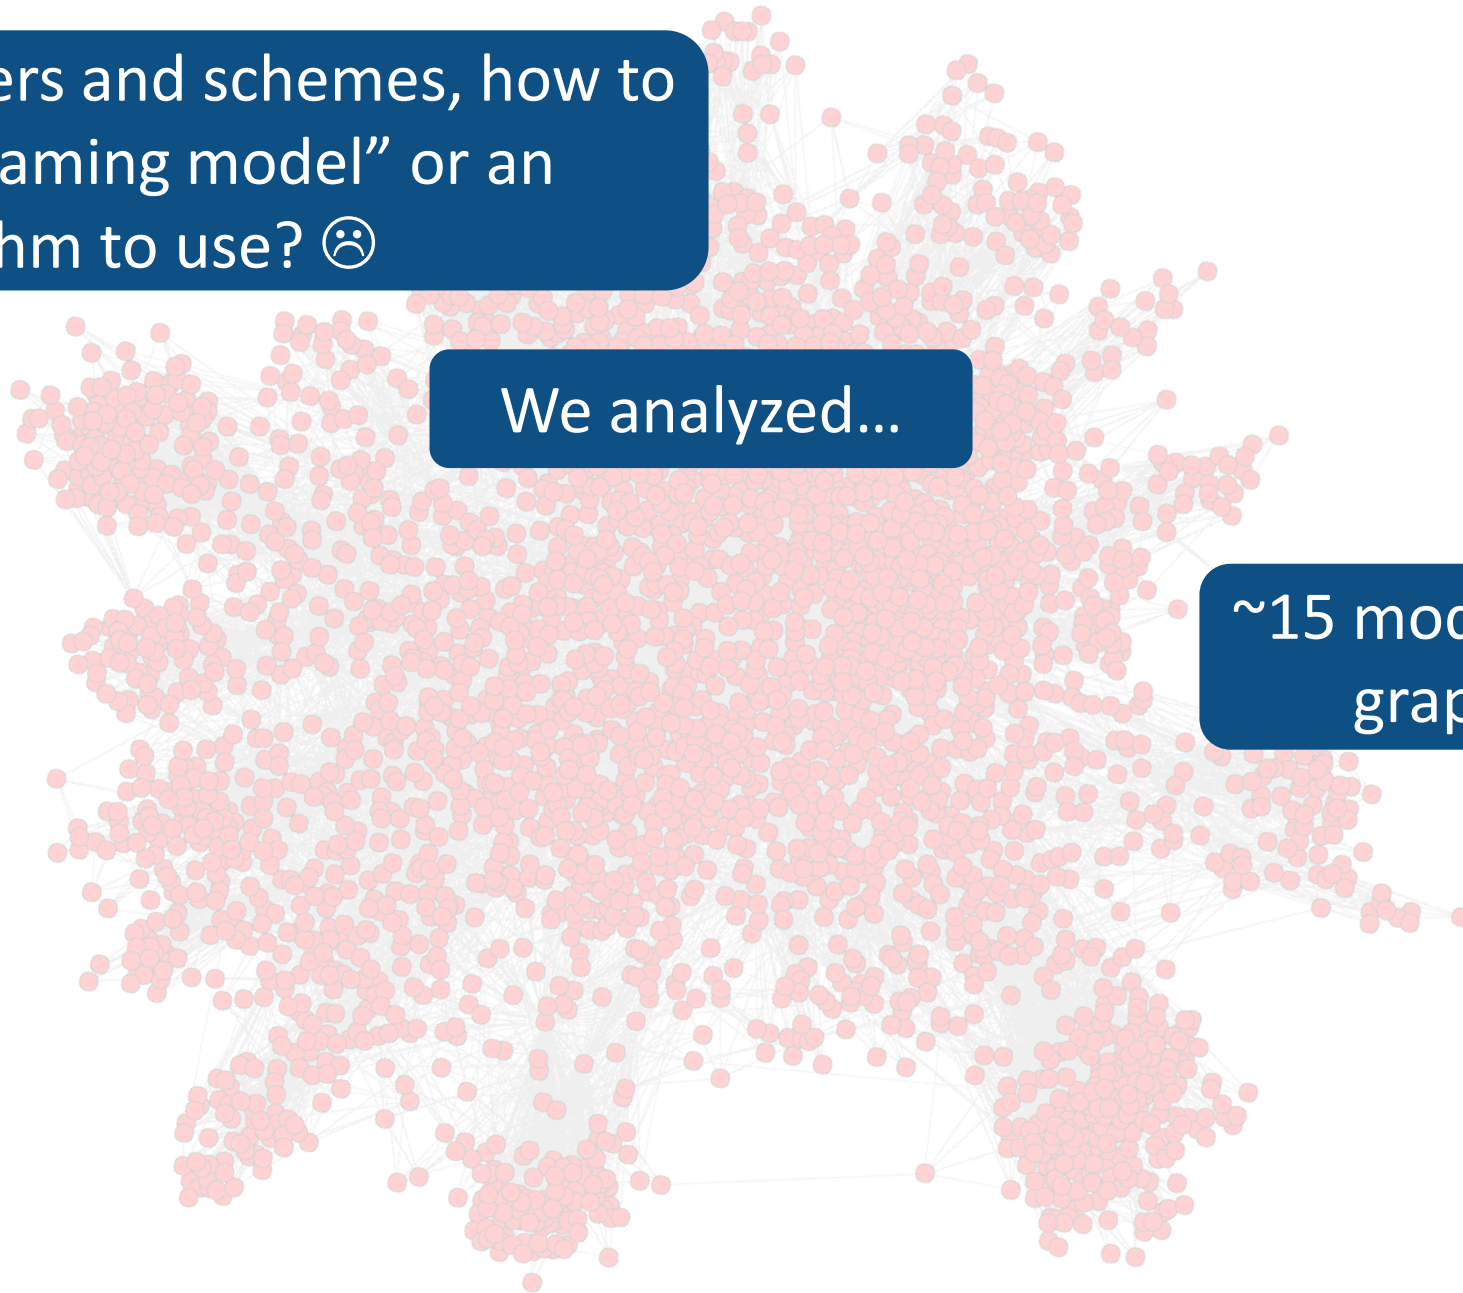
We analyzed...



Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We analyzed...

~15 models for streaming graph processing



Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We analyzed...

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We analyzed...

Which one to select?

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We analyzed...

Which one to select?

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Any interesting idea to use in the context of FPGAs and substream-centric processing?

Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We investigated the vast majority of cases, and... guess what happened 😊

We analyzed...

Which one to select?

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Any interesting idea to use in the context of FPGAs and substream-centric processing?

Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We investigated the vast majority of cases, and... guess what happened 😊

We analyzed...

Which one to select?

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Any interesting idea to use in the context of FPGAs and substream-centric processing?

1

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
 MACIEJ BESTA, Department of Computer Science, ETH Zurich  
 TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder accelerating graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing rate of graph updates, with millions of edges added or removed per second. Graph streaming algorithms are specifically crafted to eliminate these issues: The input graph is passed as a stream of updates, allowing to add and remove edges in a simple way. Recent years have seen the development of many such algorithms. However, they differ in the time needed to add or remove an edge, the required random access memory space, the number of passes



Hundreds of papers and schemes, how to select a “streaming model” or an algorithm to use? 😞

We investigated the vast majority of cases, and... guess what happened 😊

We analyzed...

Which one to select?

~15 models for streaming graph processing

~30 algorithms for streaming (approximate) MWM

Any interesting idea to use in the context of FPGAs and substream-centric processing?


It is almost ready (may take ~1 month more) – if you want to check it out earlier, let us know!

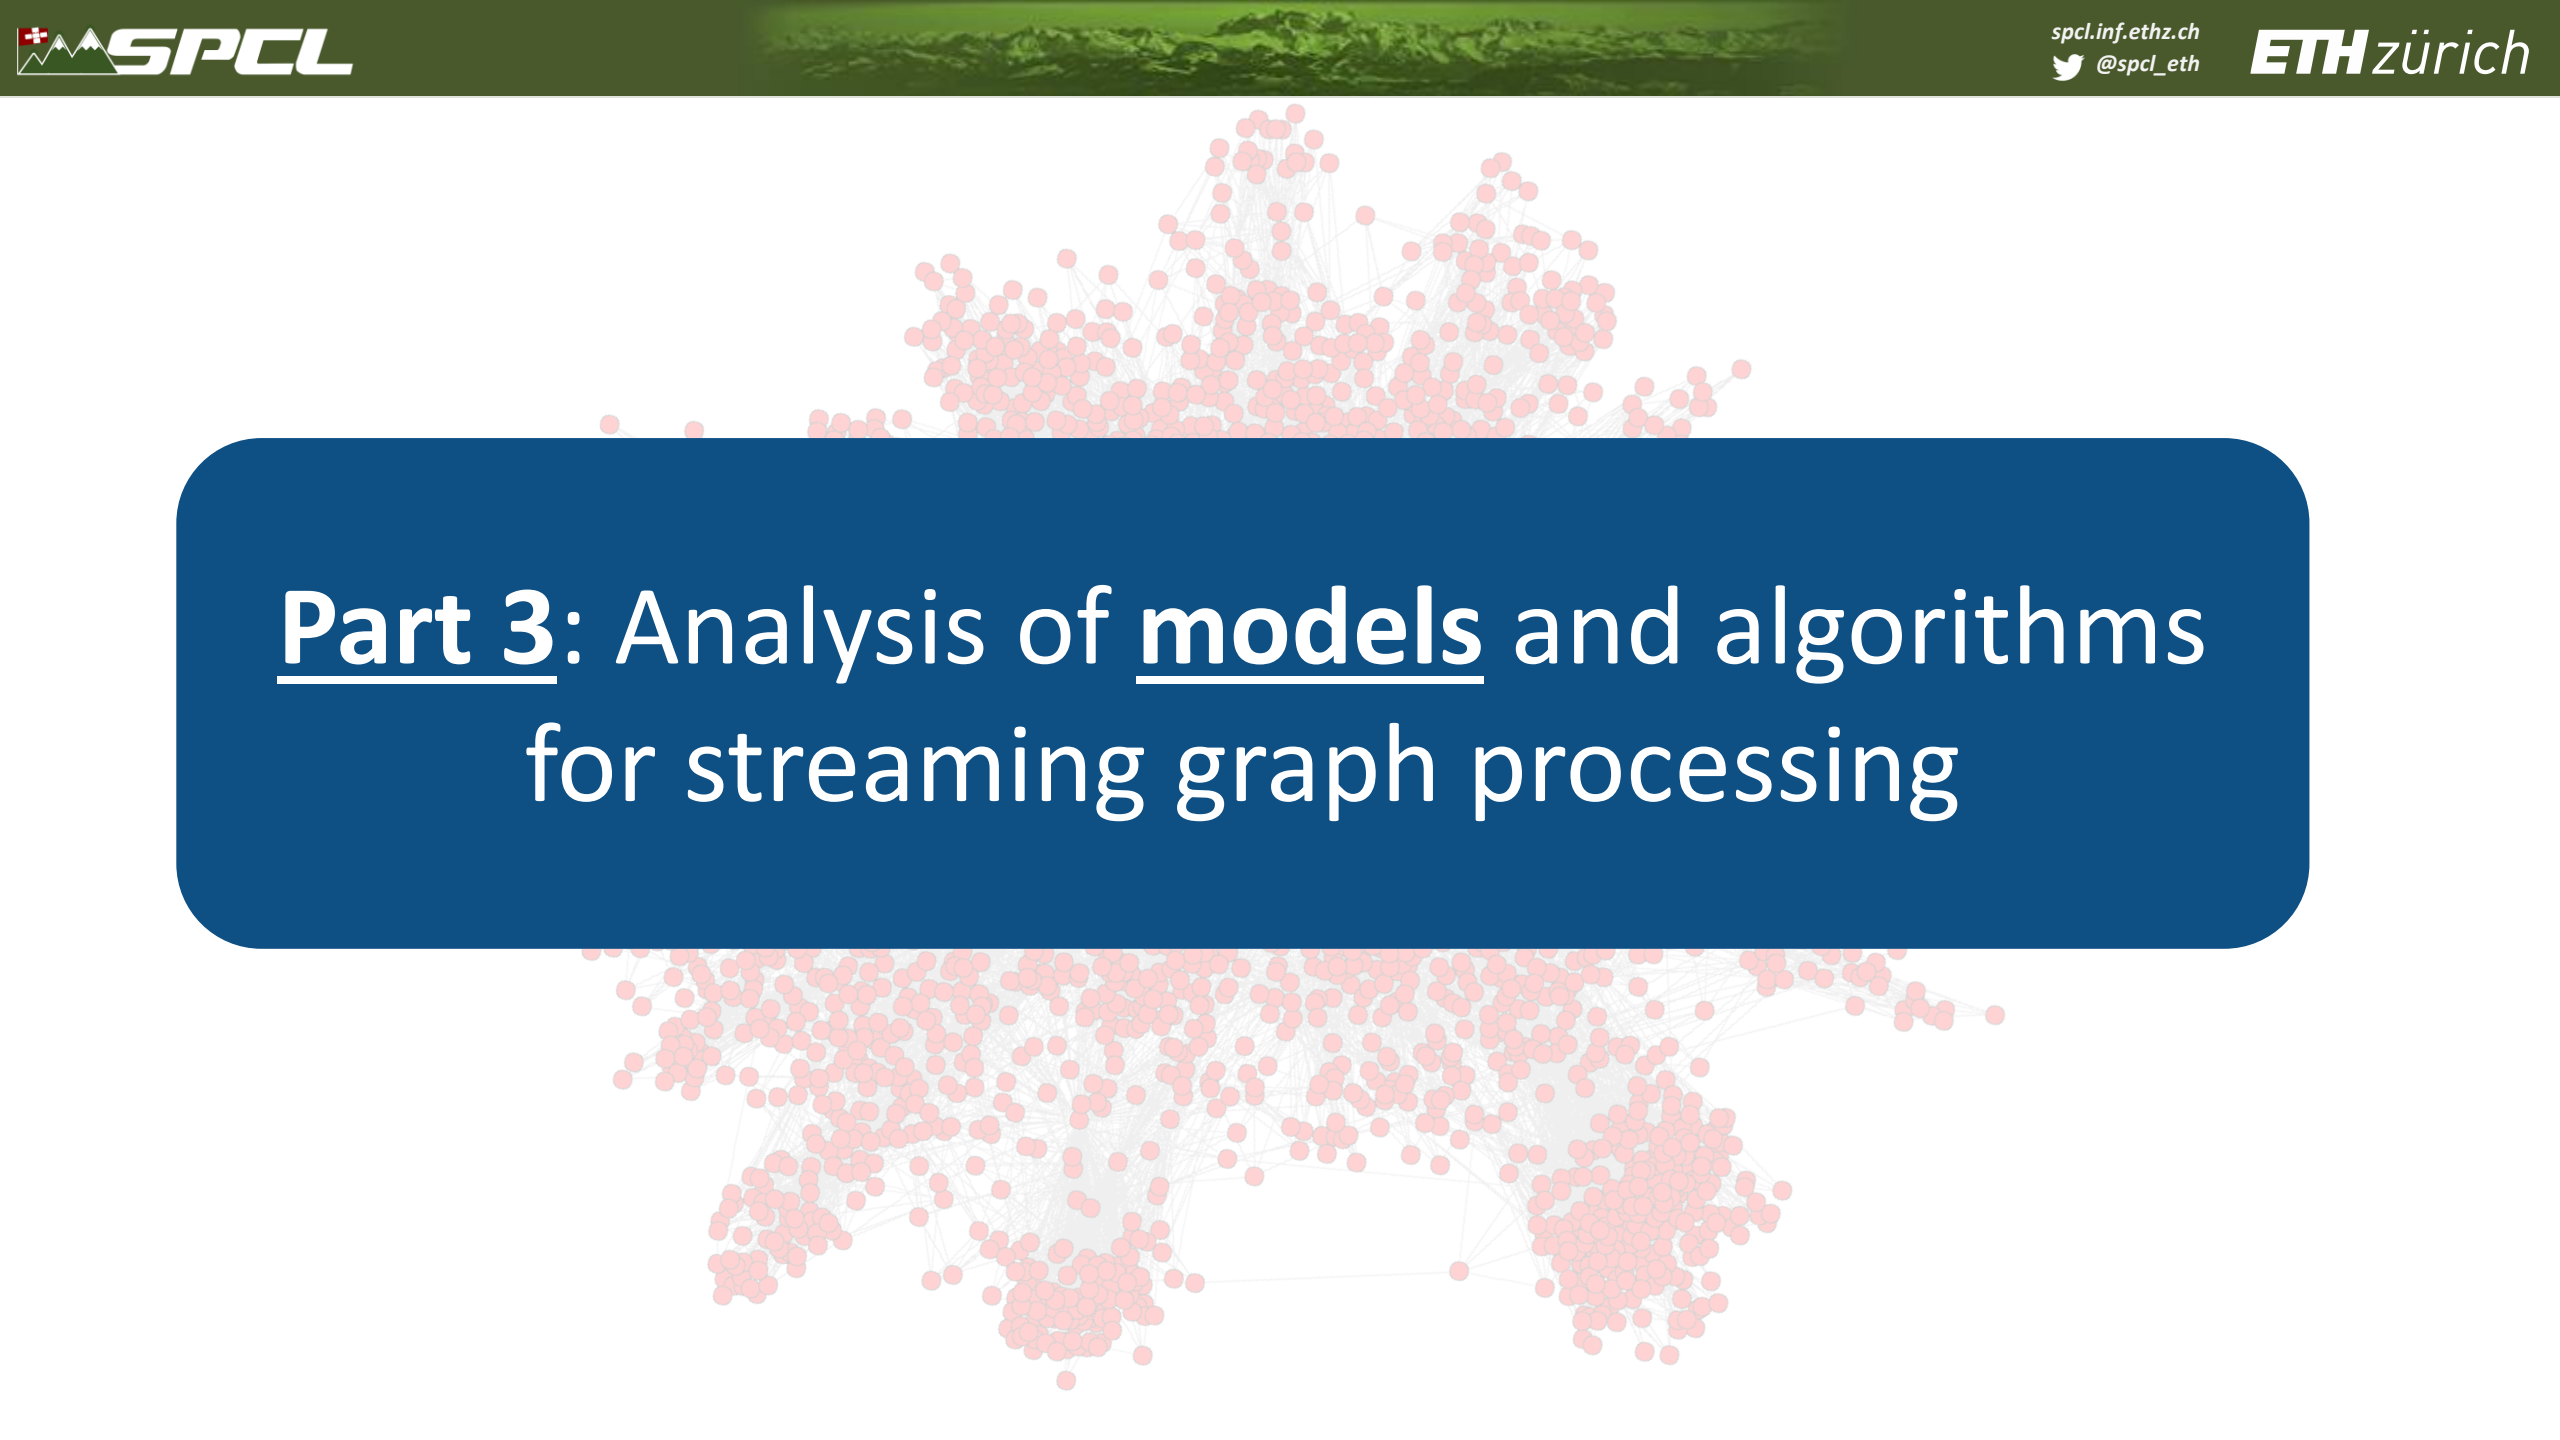
## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing

MARC FISCHER, Department of Computer Science  
MACIEJ BESTA, Department of Computer Science  
TAL BEN-NUN, Department of Computer Science  
TORSTEN HOEF, Department of Computer Science


Graph processing has long been a central part of machine learning, social network analysis, and many other domains. However, graph processing are (1) slow and (2) the growing rate of graph updates, with millions of updates per second. Graph streaming algorithms are specifically crafted to eliminate these bottlenecks. Graphs are passed as a stream of updates, allowing to add and remove edges in a simple way. Recently, we have seen the development of many such algorithms. However, they differ in the time needed to add or remove an edge, the required random access memory space, the number of passes






# Part 3: Analysis of models and algorithms for streaming graph processing

## Part 3: Analysis of models and algorithms for streaming graph processing



## Part 3: Analysis of models and algorithms for streaming graph processing

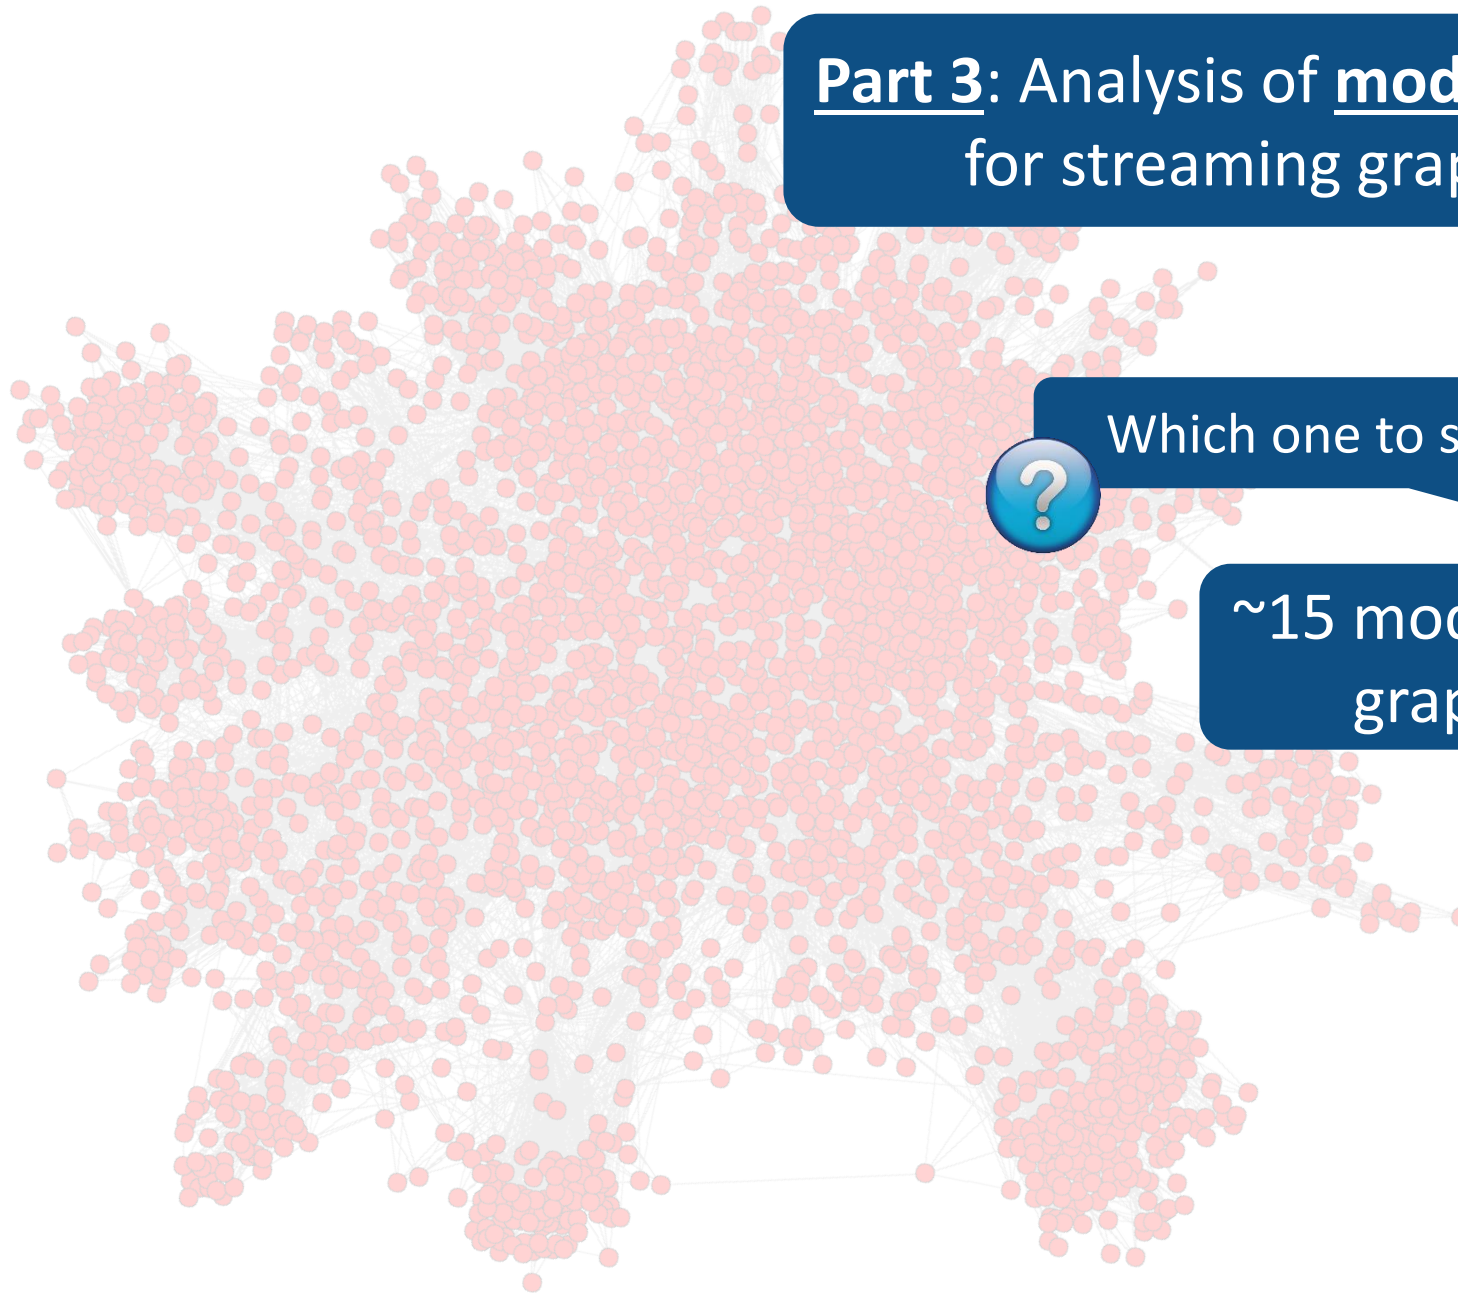


~15 models for streaming graph processing

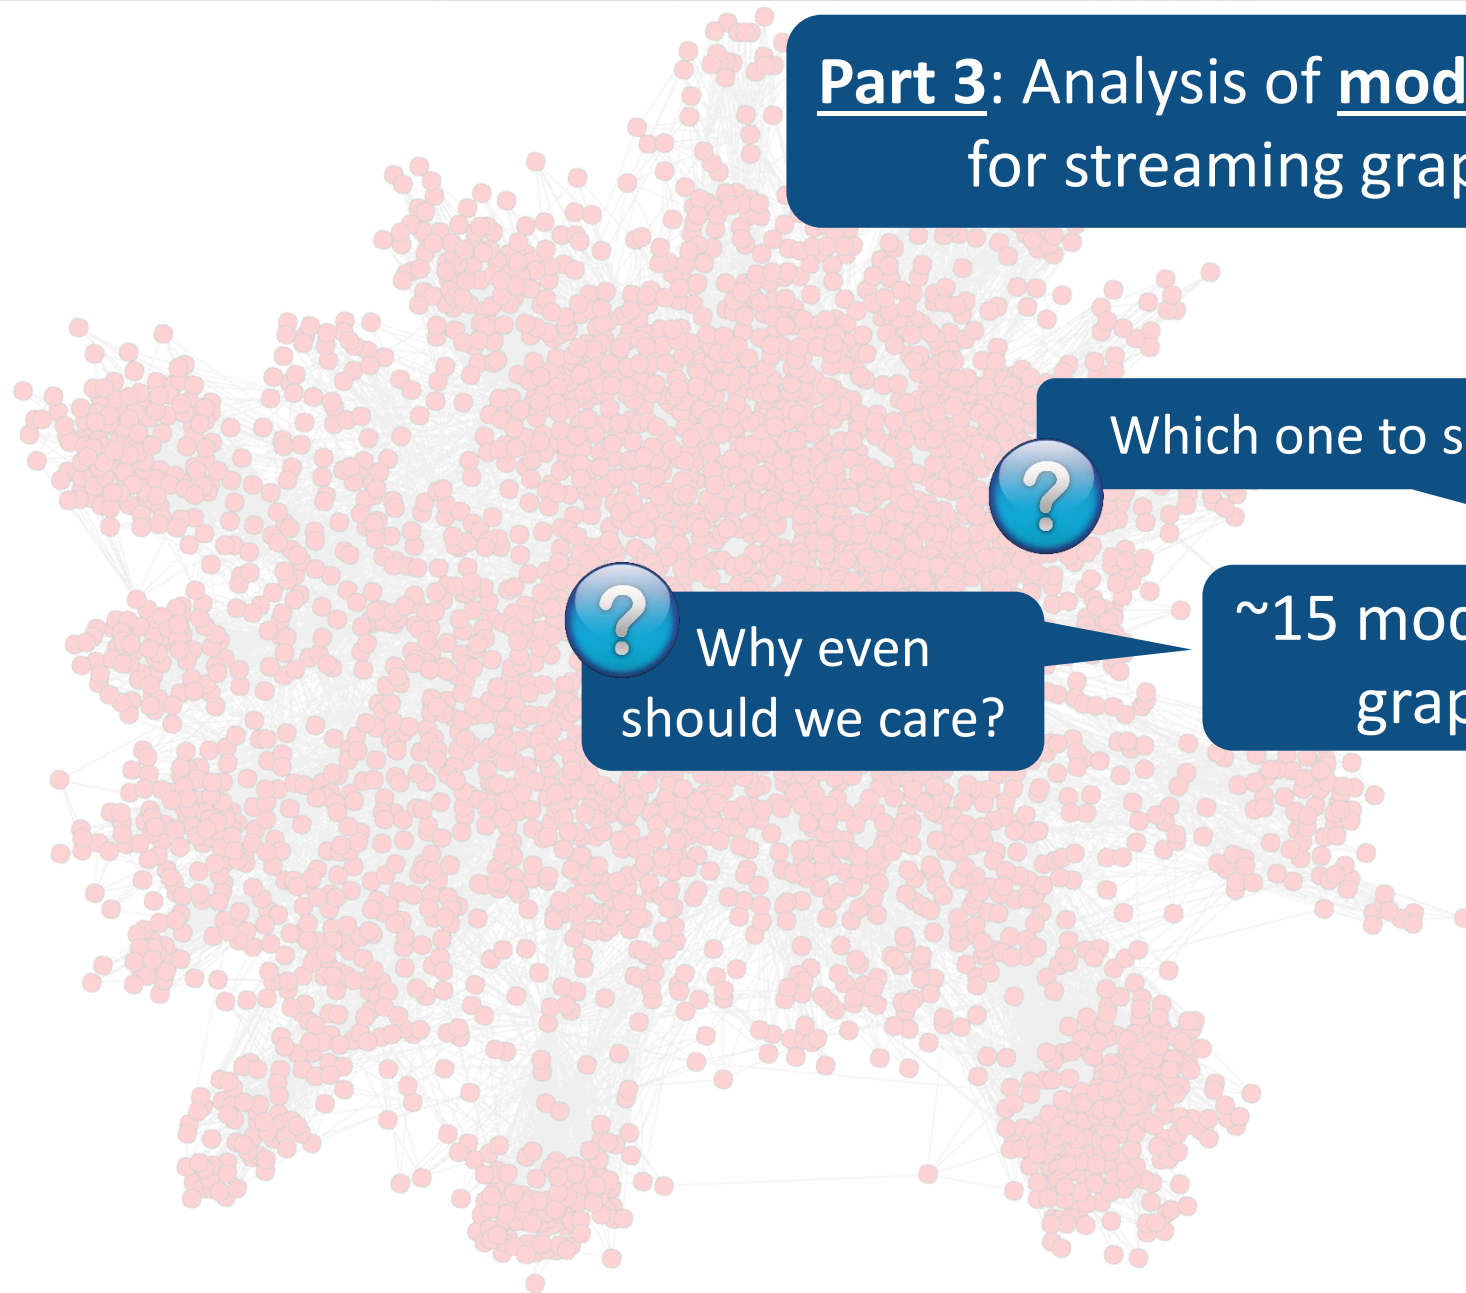
## Part 3: Analysis of models and algorithms for streaming graph processing

Which one to select?

~15 models for streaming graph processing



## Part 3: Analysis of models and algorithms for streaming graph processing



? Why even should we care?

? Which one to select?

~15 models for streaming graph processing

## Part 3: Analysis of models and algorithms for streaming graph processing

Using a model enables rigorous analysis of the algorithm behavior

? Why even should we care?

? Which one to select?

~15 models for streaming graph processing



## Part 3: Analysis of models and algorithms for streaming graph processing

Using a model enables rigorous analysis of the algorithm behavior

? Why even should we care?

? Which one to select?

~15 models for streaming graph processing

Selecting a right model enables more realistic predictions about the algorithm behavior on given hardware

## Part 3: Analysis of models and algorithms for streaming graph processing

Using a model enables rigorous analysis of the algorithm behavior


Limiting oneself to a particular model helps to select the best algorithm or technique (within that model)

? Why even should we care?

? Which one to select?

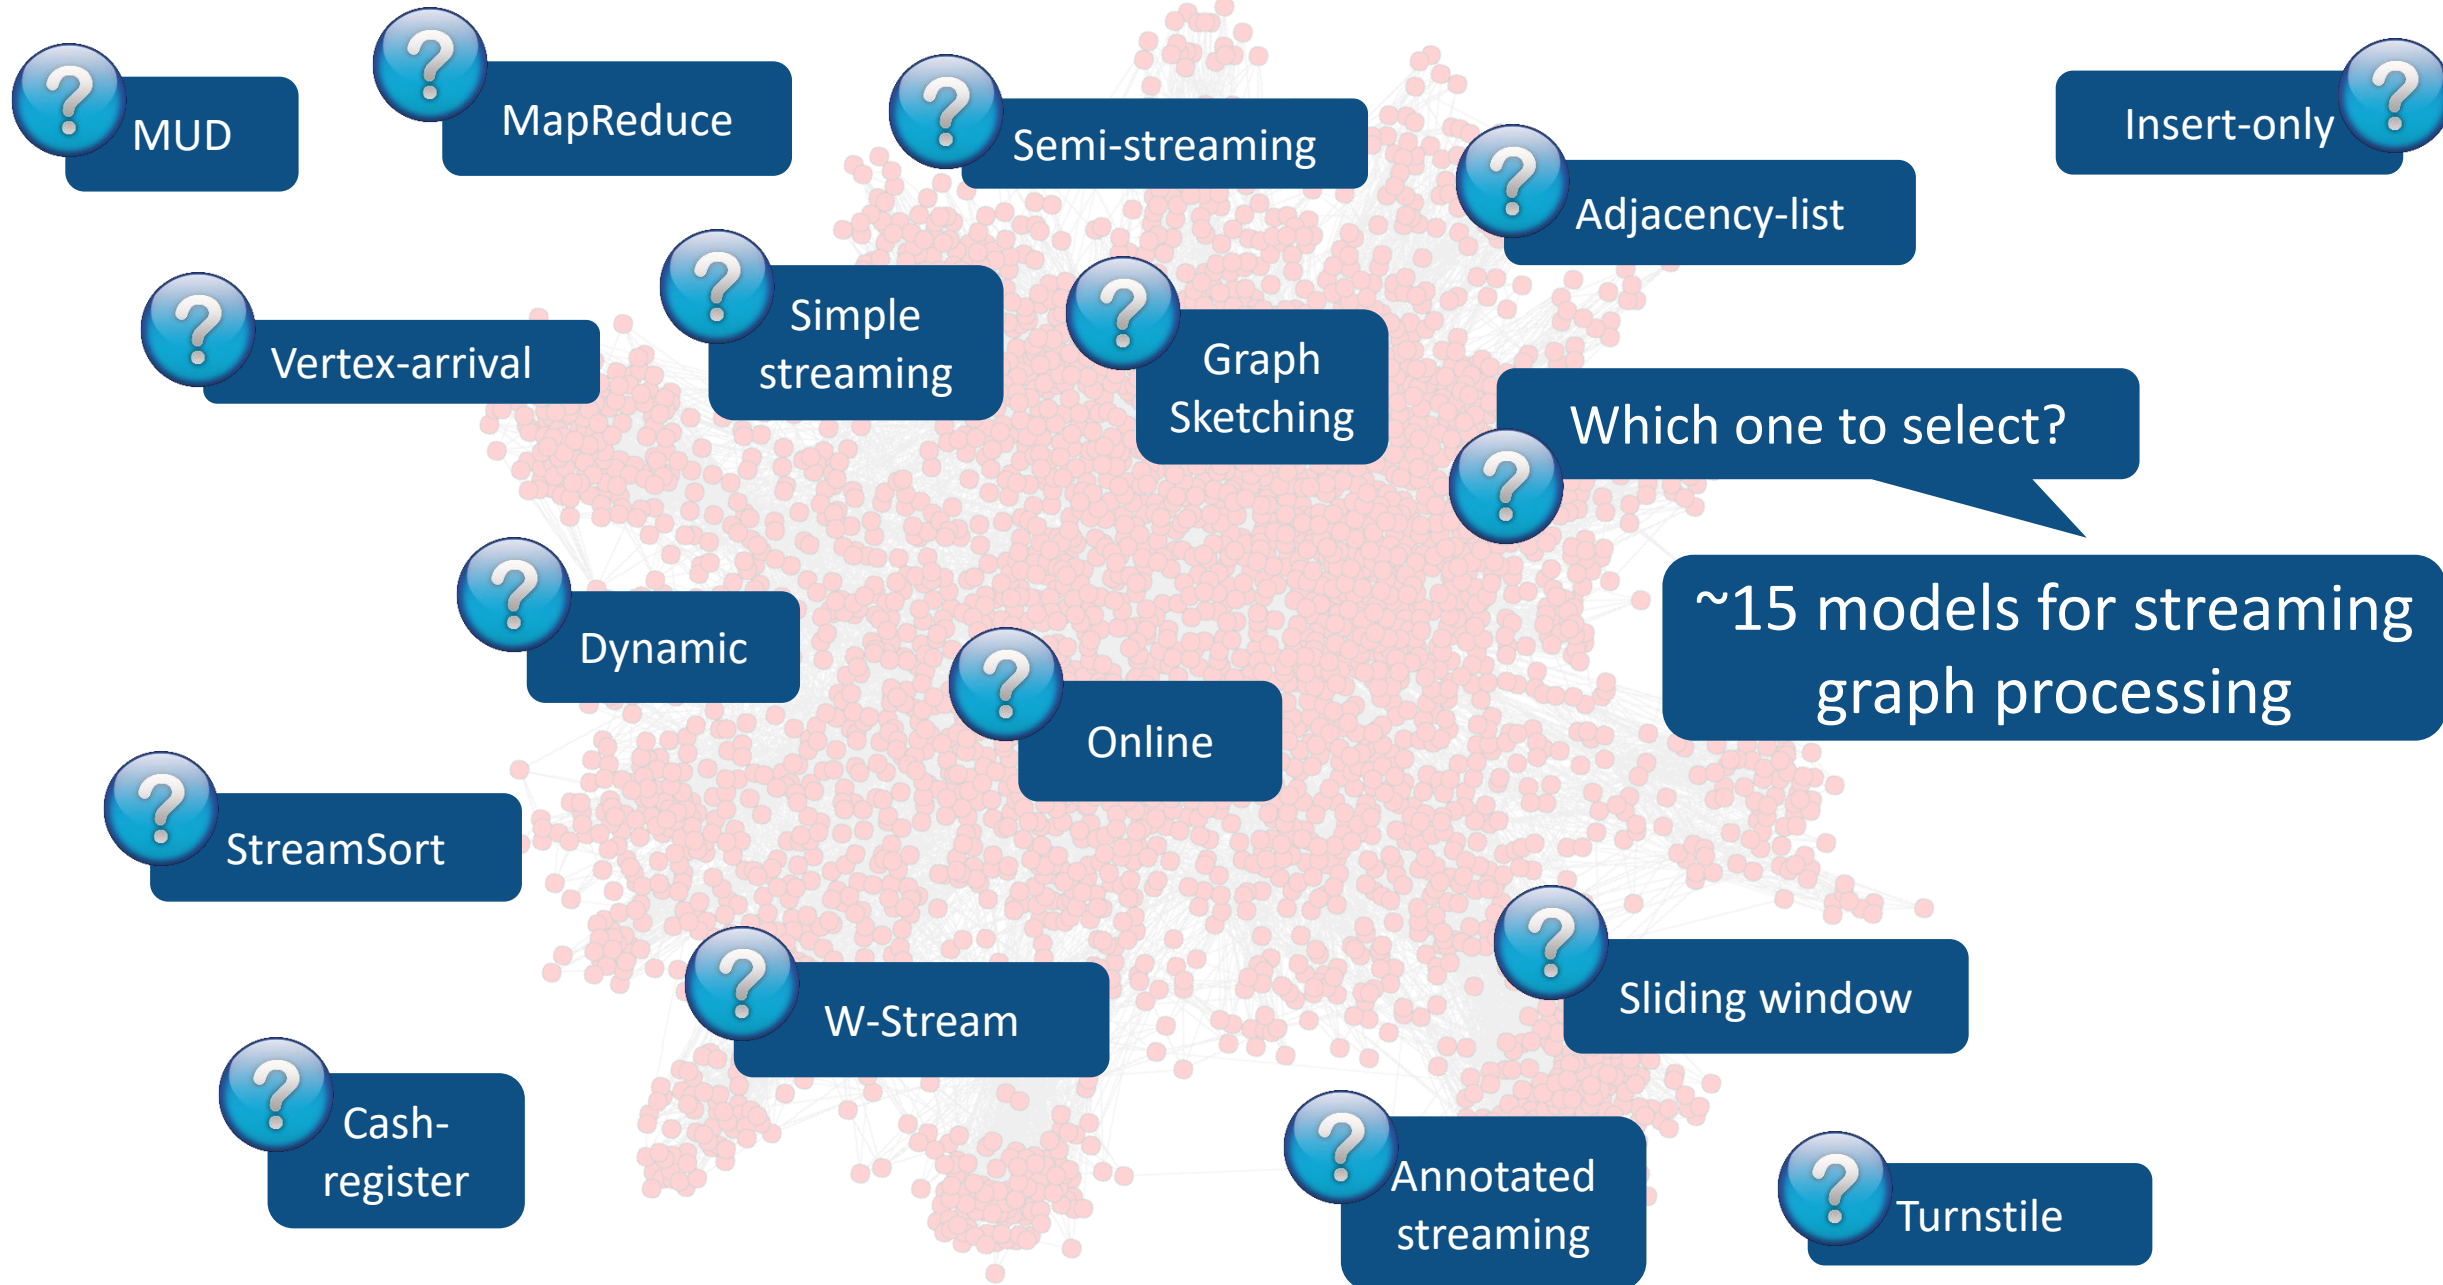
? ~15 models for streaming graph processing

Selecting a right model enables more realistic predictions about the algorithm behavior on given hardware



Which one to select?

~15 models for streaming  
graph processing



Any graph streaming algorithm belongs to the semi-streaming model if it uses at most  $O(n \text{ polylog}(n))$  space

No color change indicates the informal "instance" relationship: a model "below" is a particular instance of a streaming model

A change of color indicates the formal "reduction" relationship: a model "above" can be used to execute an algorithm that was developed in a model "below"

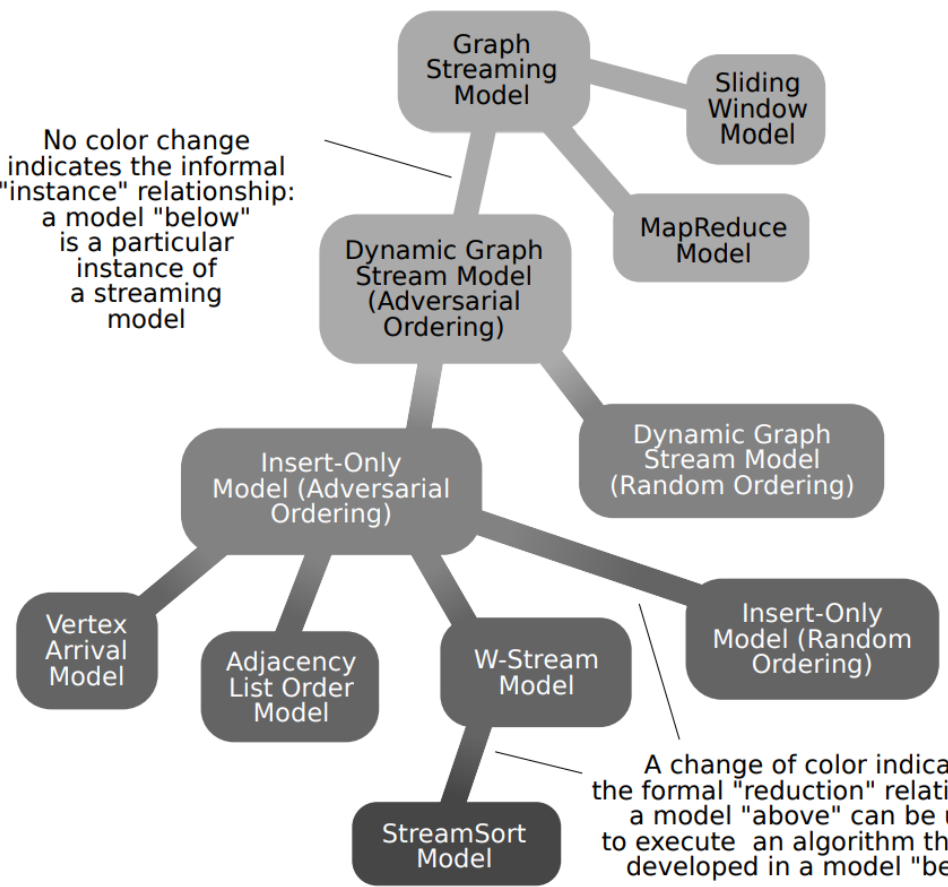


Fig. 1. The hierarchy of the graph streaming models.

register

semi-streaming

Insert-only

Graph Sketching

Adjacency-list

Which one to select?

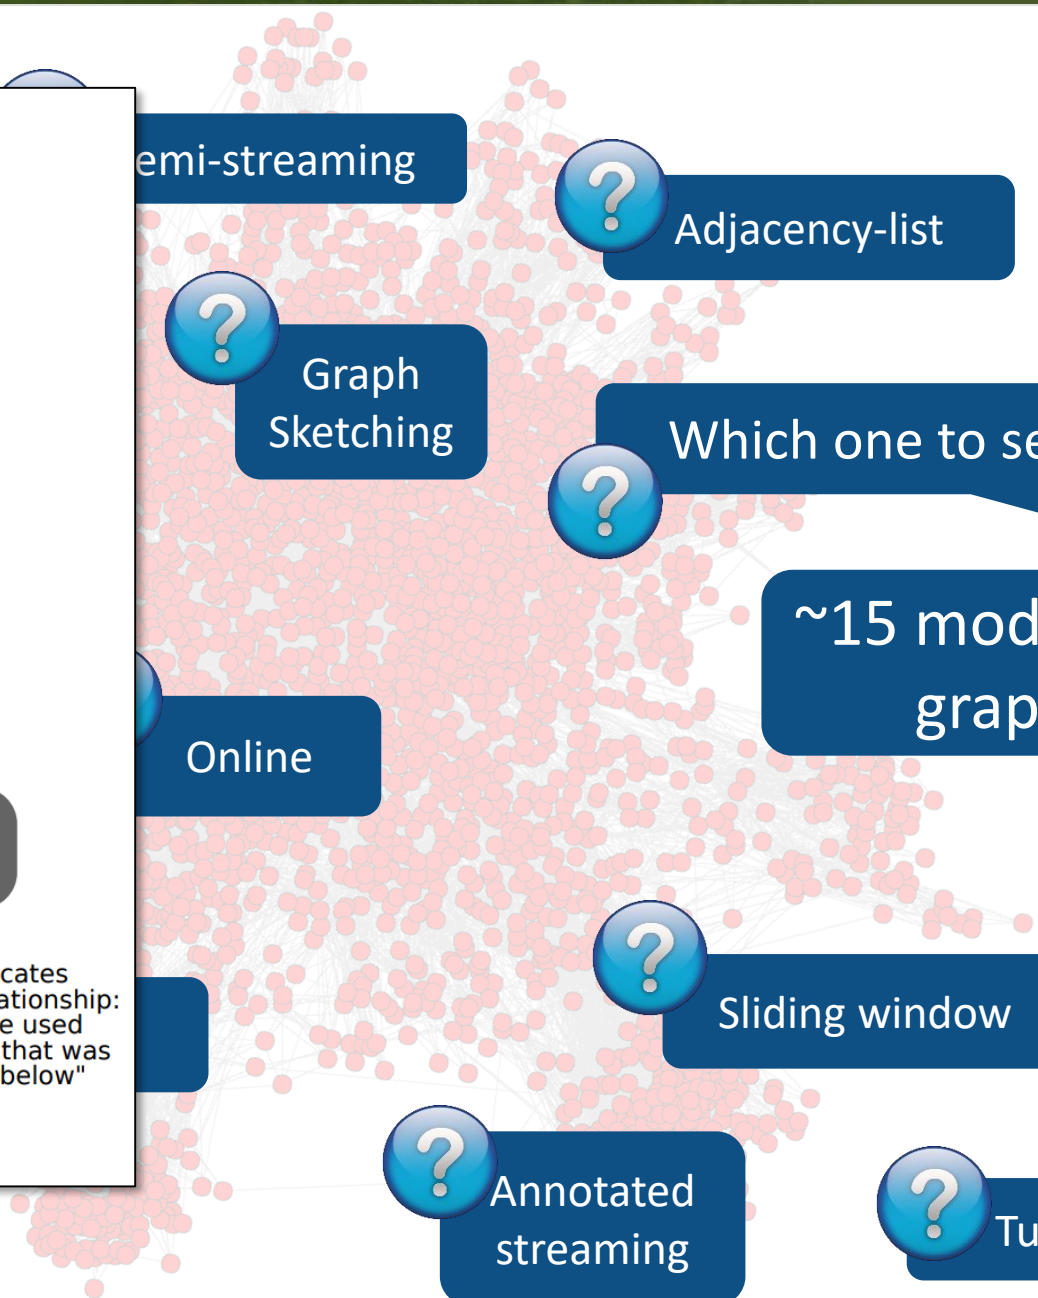
~15 models for streaming graph processing

Online

Sliding window

Annotated streaming

Turnstile



Any graph streaming algorithm belongs to the semi-streaming model if it uses at most  $O(n \text{ polylog}(n))$  space

To understand the models (and related caveats) well, we developed a formal taxonomy of the analyzed models with the aim of guiding future graph streaming designs (check the report for details 😊)

StreamSort Model

A change of color indicates the formal "reduction" relationship: a model "above" can be used to execute an algorithm that was developed in a model "below"

Fig. 1. The hierarchy of the graph streaming models.

register

semi-streaming

Insert-only

Adjacency-list

Graph Sketching

Which one to select?

~15 models for streaming graph processing

Online

Sliding window

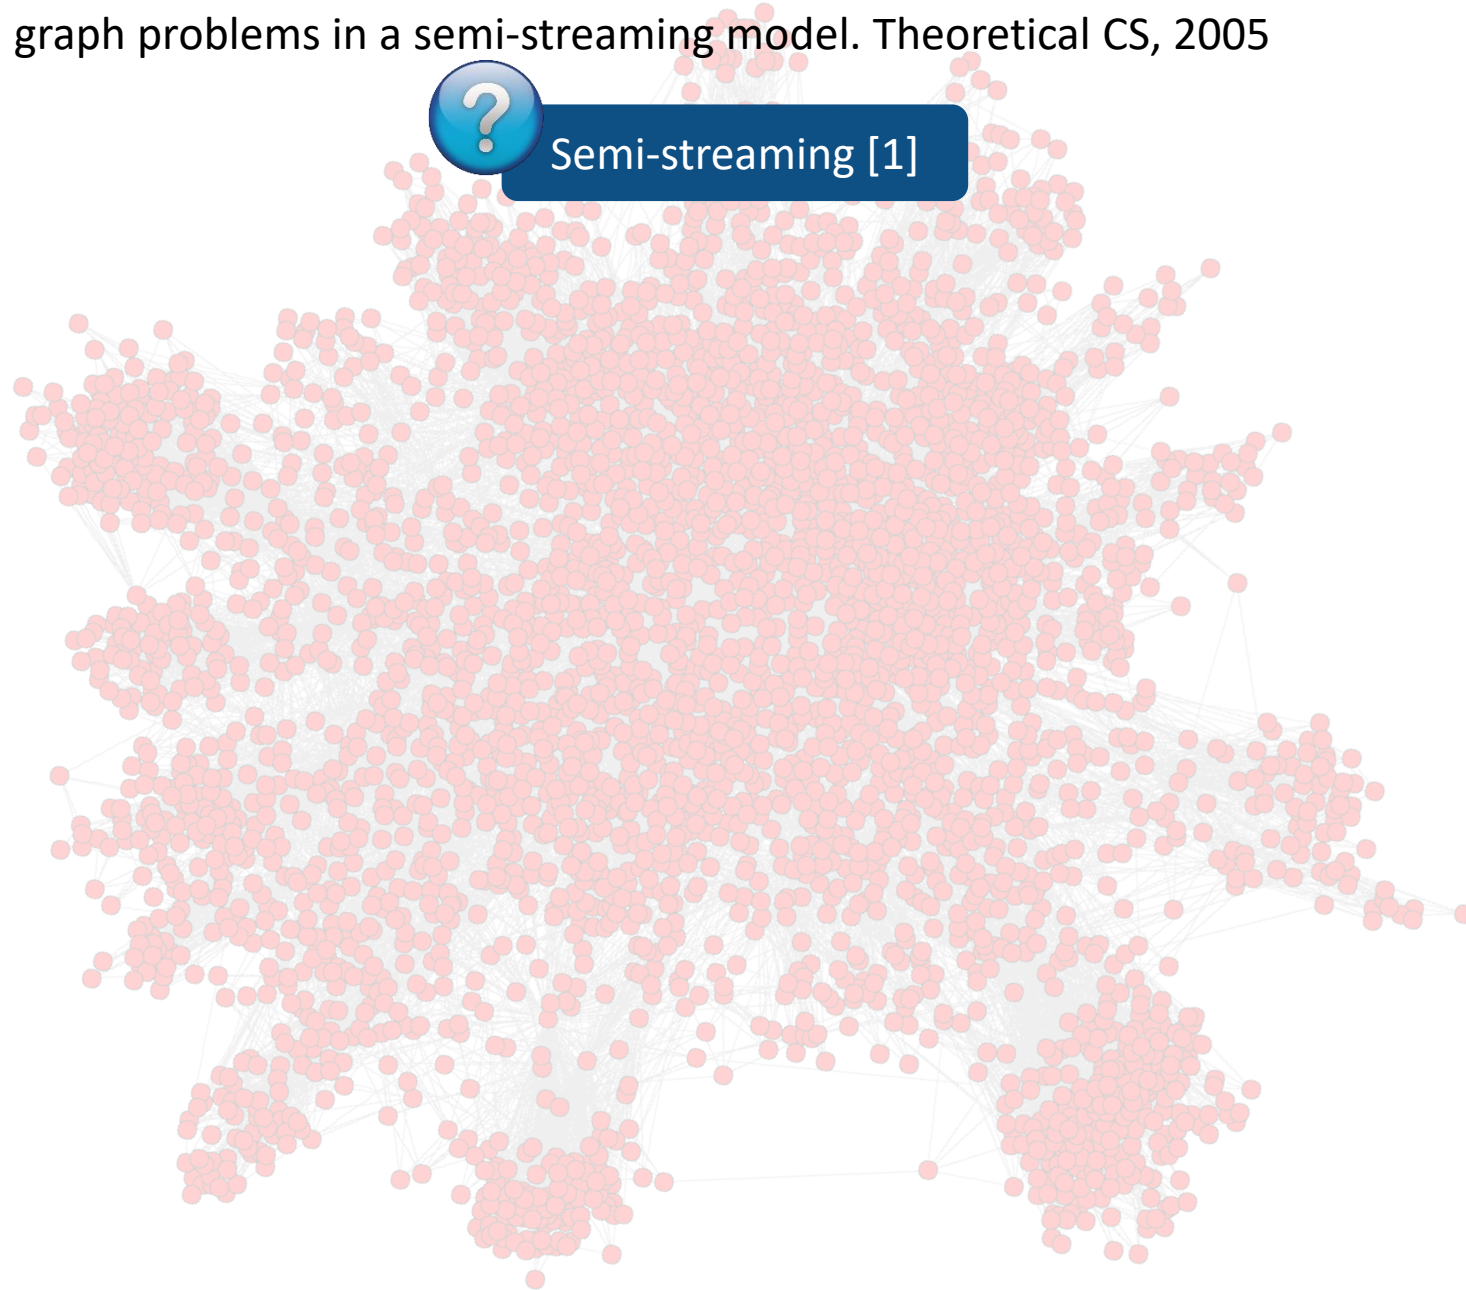
Annotated streaming

Turnstile

[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005



Semi-streaming [1]



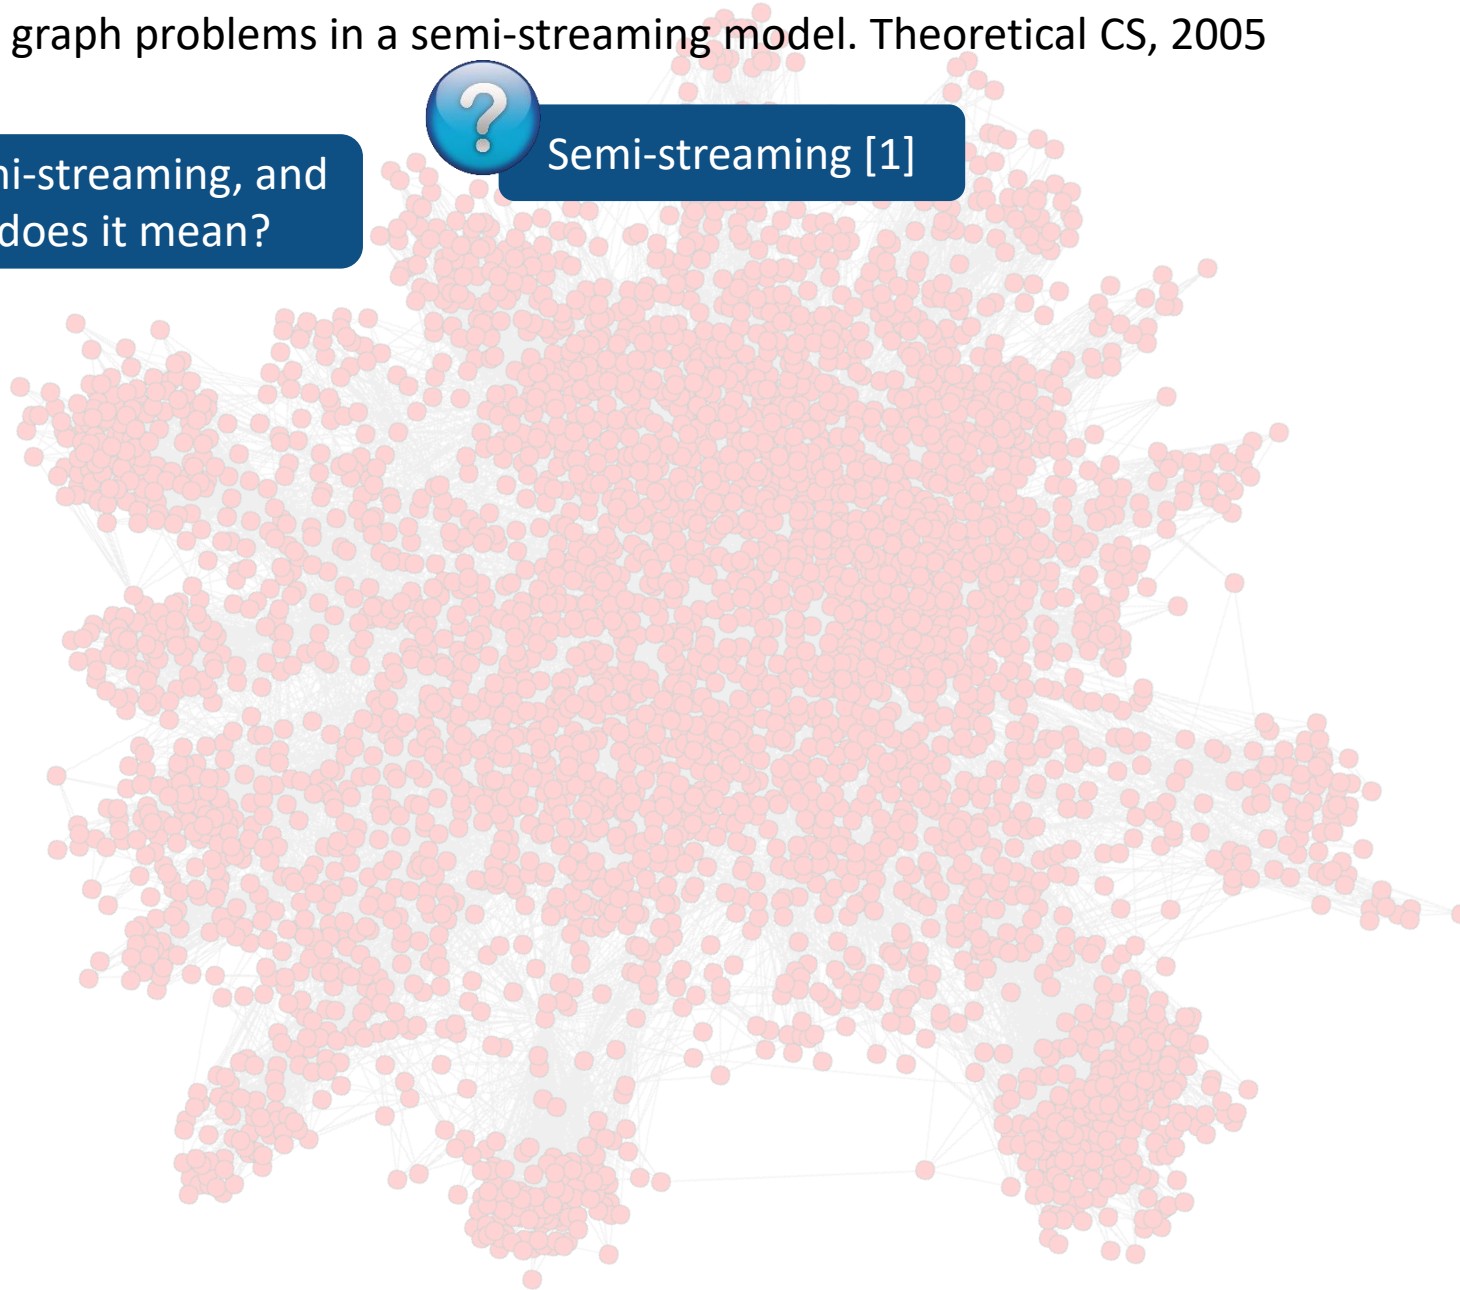
[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005



Why semi-streaming, and  
what does it mean?



Semi-streaming [1]





[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005




Why semi-streaming, and what does it mean?



Semi-streaming [1]



Covers a general streaming setting (= works for substream-centric)




[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005

? Why semi-streaming, and what does it mean?

? Semi-streaming [1]

✓ Covers a general streaming setting (= works for substream-centric)



? What programming paradigm and why?

[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005

Why semi-streaming, and what does it mean?

Semi-streaming [1]


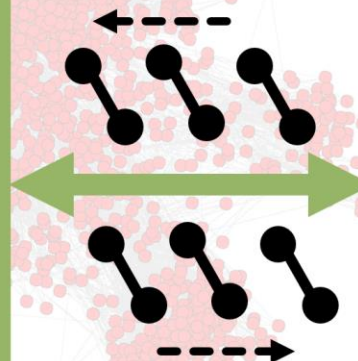

Covers a general streaming setting (= works for substream-centric)

Assumes  $O(n \log^c n)$  local space that can be used for processing an edge  $\rightarrow$  fits well FPGA BRAM constraints!

What programming paradigm and why?

Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)



DRAM

[1] J. Feigenbaum et al. On graph problems in a semi-streaming model. Theoretical CS, 2005

? Why semi-streaming, and what does it mean?

? Semi-streaming [1]

✓ Covers a general streaming setting (= works for substream-centric)




✓ Assumes  $O(n \log^c n)$  local space that can be used for processing an edge → fits well FPGA BRAM constraints!

? What programming paradigm and why?

✓ Offers (potentially powerful) MWM algorithms

Use some form of streaming (aka edge-centric)

Some processing unit (CPU, GPU, FPGA, ...)



DRAM

## Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



Which programming paradigm to use for (approximate) MWM?



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions

Use substream-centric processing (exposes parallelism)

... design to use for (approximate) MWM?



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?




A large, complex network graph visualization with numerous red circular nodes and thin grey edges, forming a dense, interconnected structure. The graph is centered in the background of the slide.

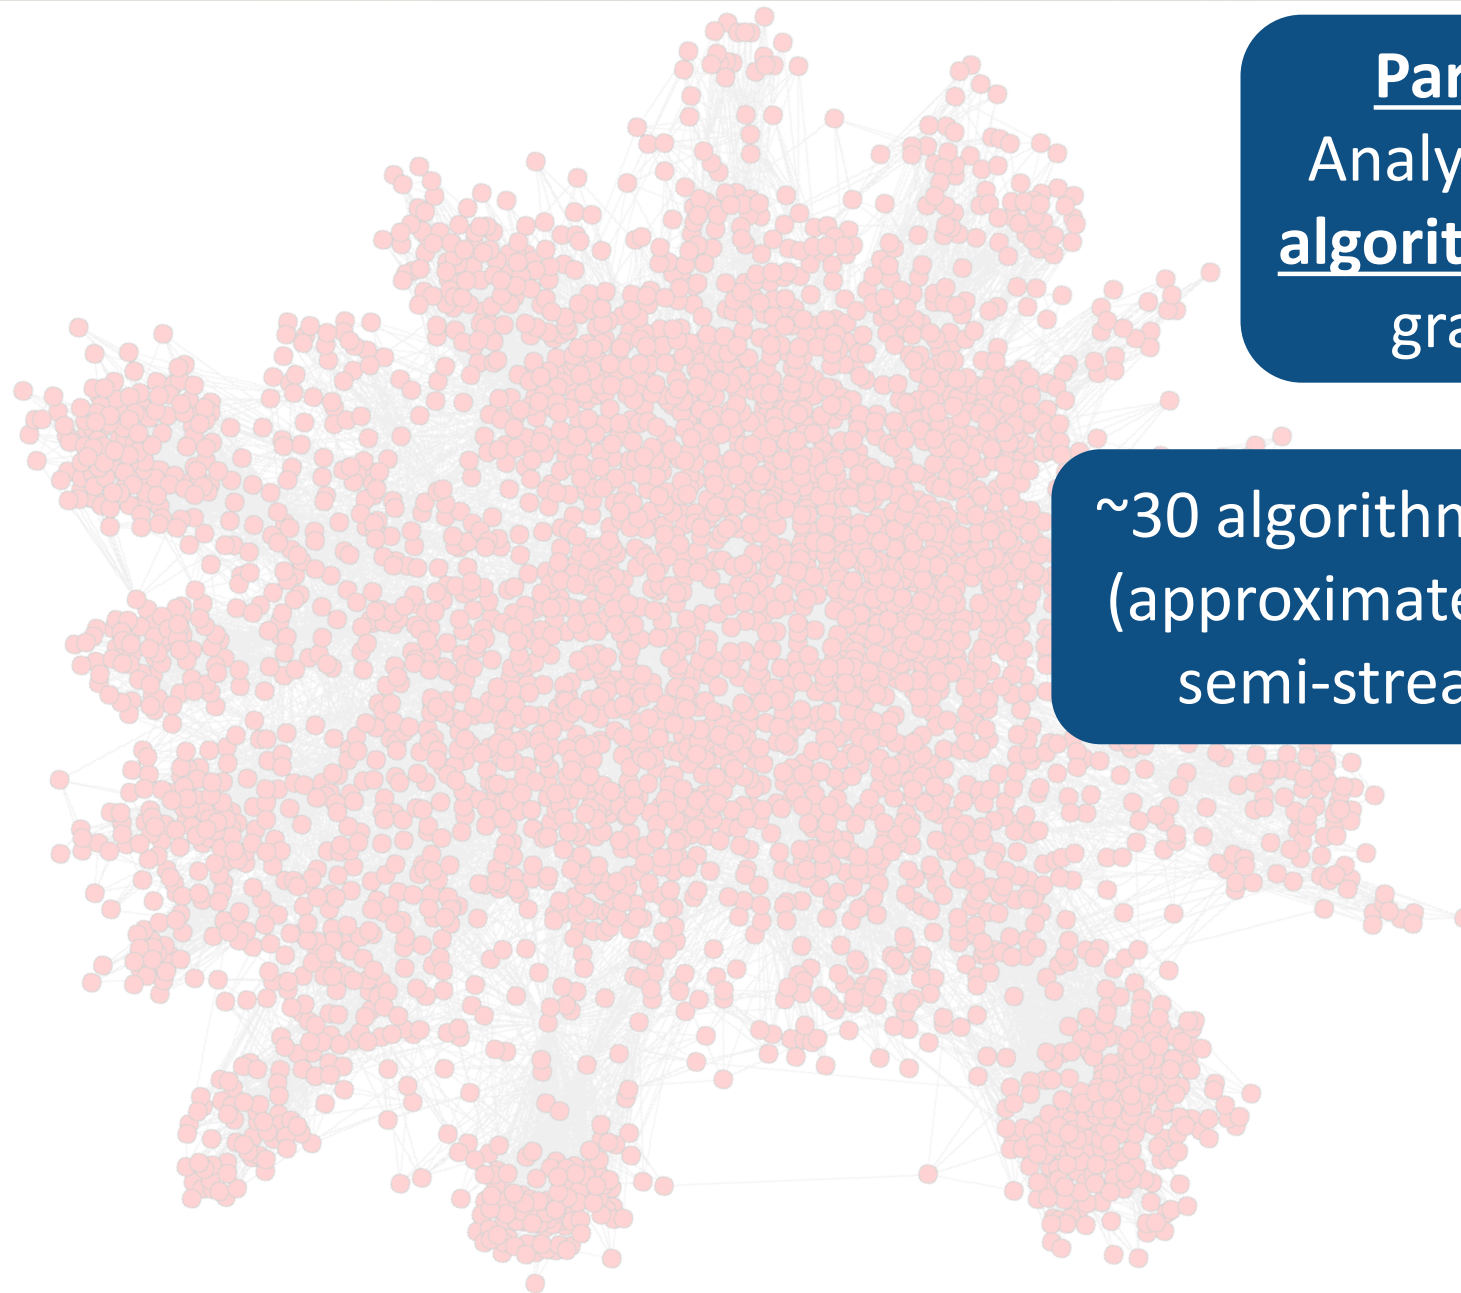
**Part 3 continued: Analysis of models and algorithms for streaming graph processing**

A large, complex network graph visualization with numerous red circular nodes and thin grey edges, forming a dense, interconnected structure. The graph is centered on the slide, with a dark blue rounded rectangle overlaid in the middle containing the title text.

**Part 3 continued: Analysis of models and algorithms for streaming graph processing**

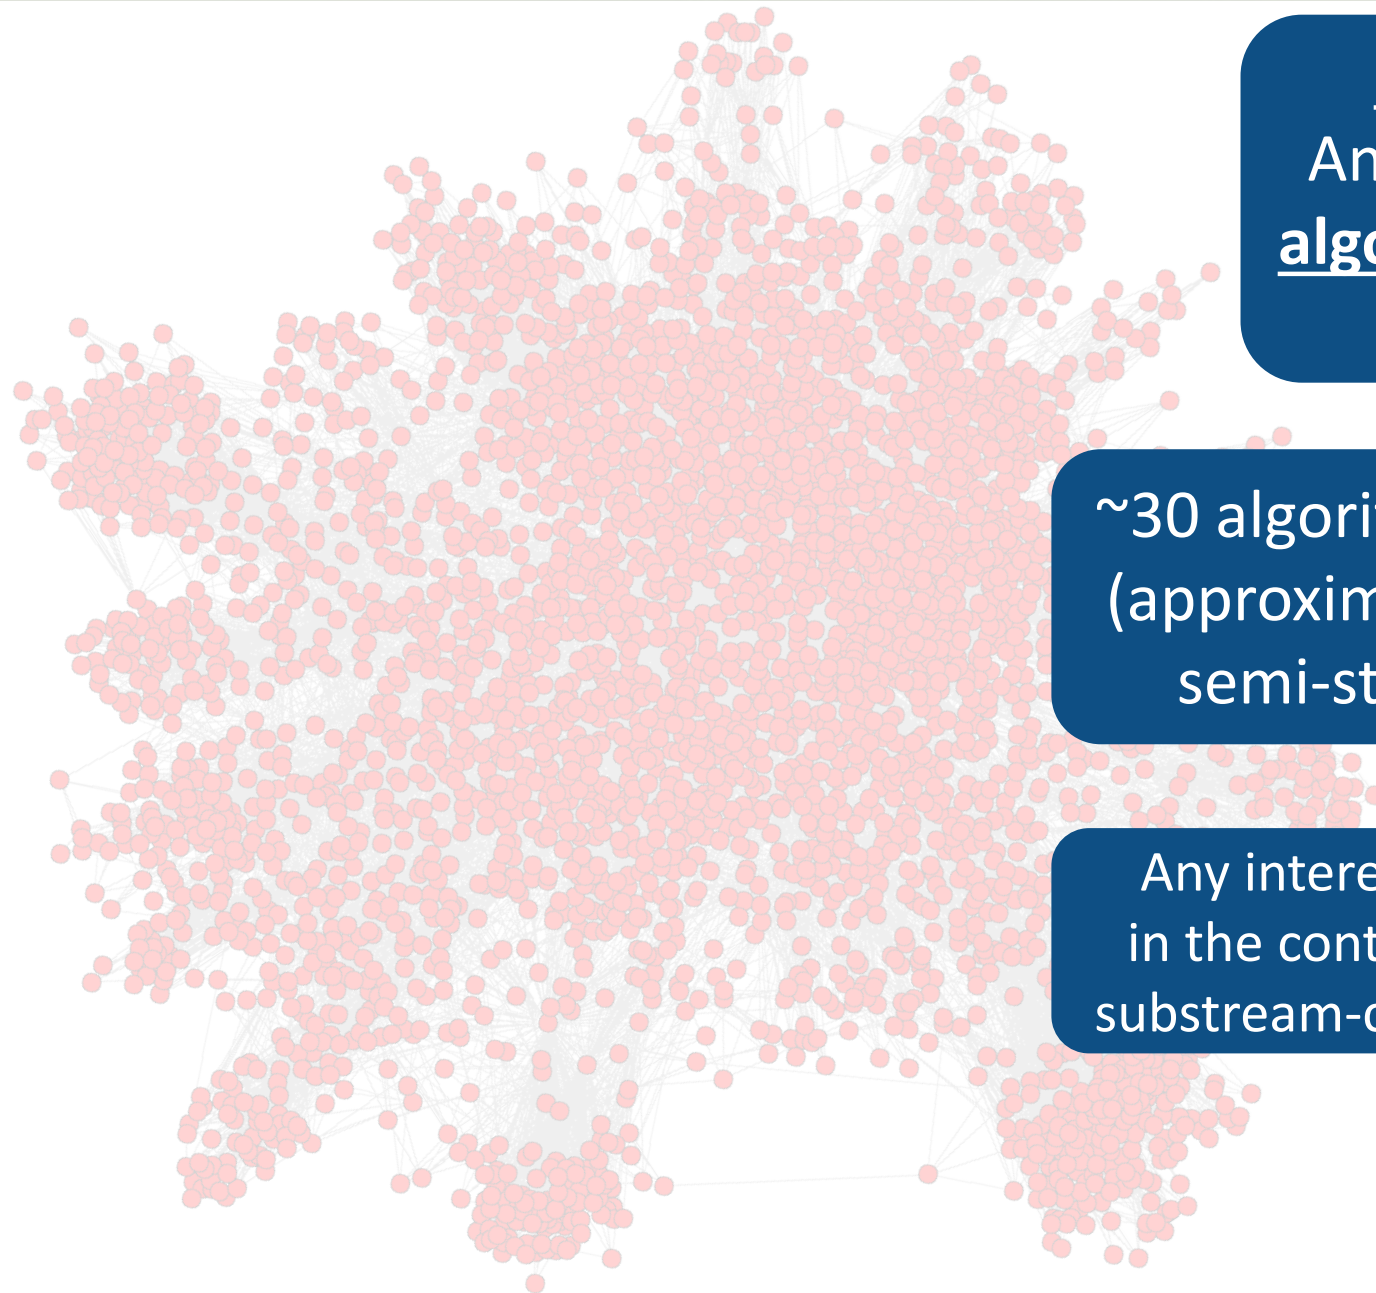


Part 3 continued :  
Analysis of models and  
algorithms for streaming  
graph processing



Part 3 continued :  
Analysis of models and  
algorithms for streaming  
graph processing

~30 algorithms for streaming  
(approximate) MWM (in the  
semi-streaming model)

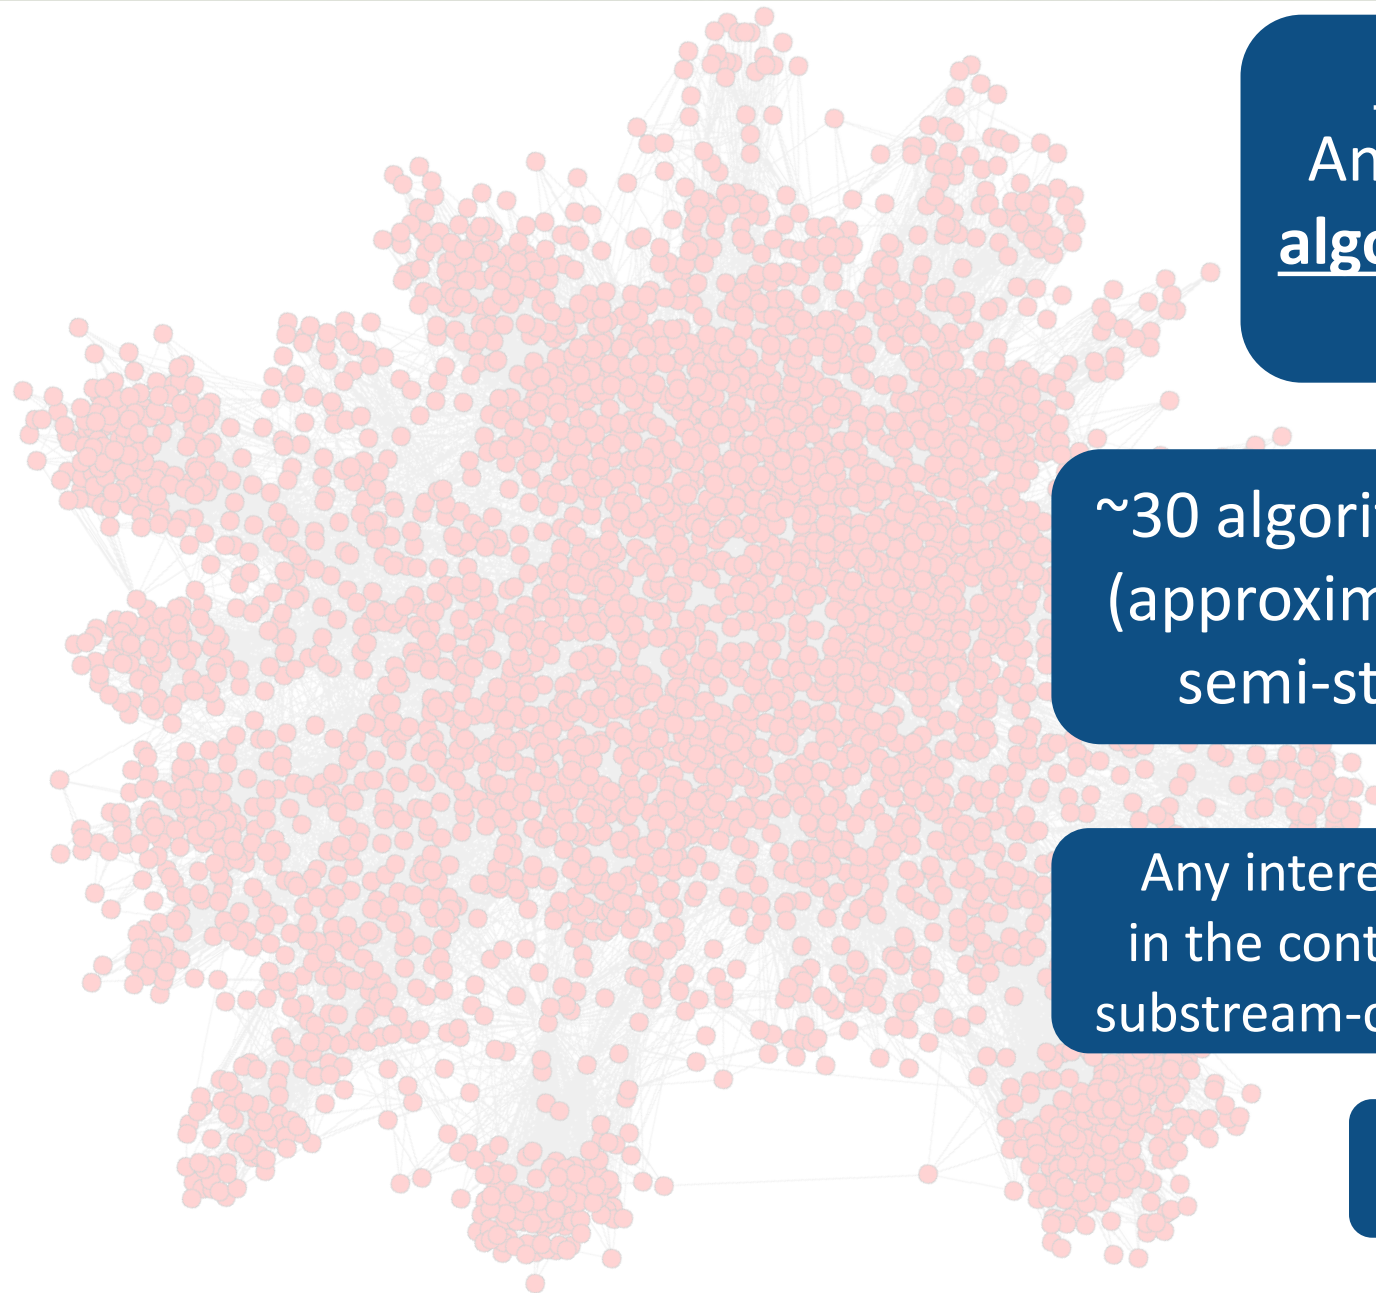


Part 3 continued :  
Analysis of models and  
algorithms for streaming  
graph processing

~30 algorithms for streaming  
(approximate) MWM (in the  
semi-streaming model)

Any interesting idea to use  
in the context of FPGAs and  
substream-centric processing?





Part 3 continued :  
Analysis of models and  
algorithms for streaming  
graph processing

~30 algorithms for streaming  
(approximate) MWM (in the  
semi-streaming model)

Any interesting idea to use  
in the context of FPGAs and  
substream-centric processing?



More specifically...

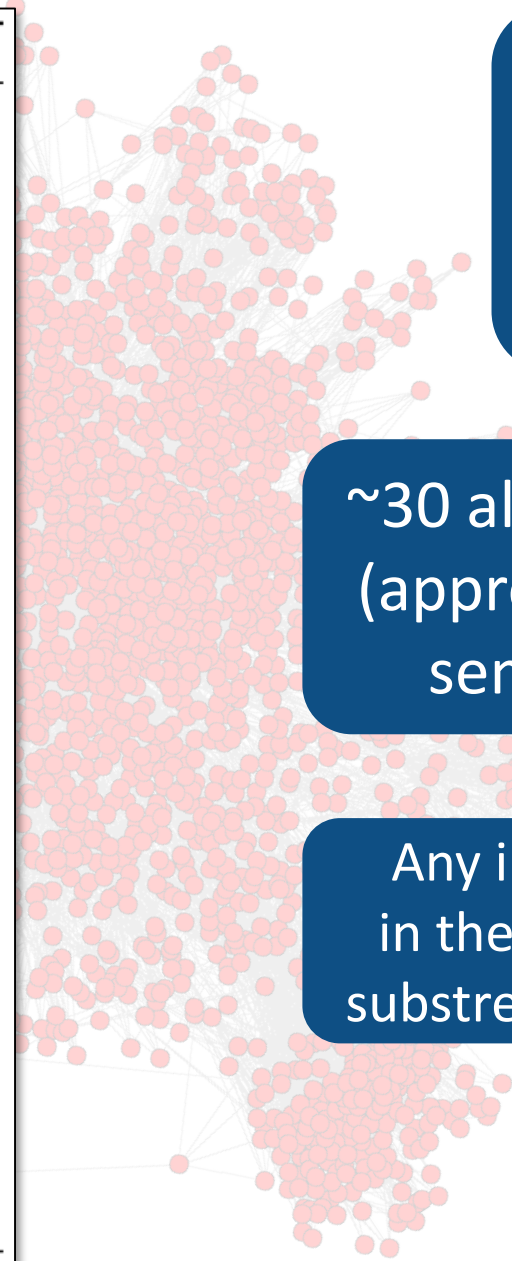
Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[26]	1/2	$O(n)$	1	👎	👍	👍
[41, Theorem 6]	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👎	👍	👍
[41, Theorem 2]	$1/2 + 0.003^*$	$O(n \text{ polylog}(n))$	1	👎	👍	👍
[36, Theorem 1.1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👎	👍	👍
[26, Theorem 1]	$2/3 - \epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👎	👎	👍
[6, Theorem 19]	$1 - \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👎	👎	👍
[41, Theorem 5]	$1/2 + 0.019$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👎	👎	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👎	👎	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👎	👎	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👎	👎	👍
[14]	1	$\tilde{O}(k^2)$	1	👎	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👎	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👎	👎	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	?
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	?
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	?
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	?
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	?
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	?
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	?
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	?
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	?
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👎	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

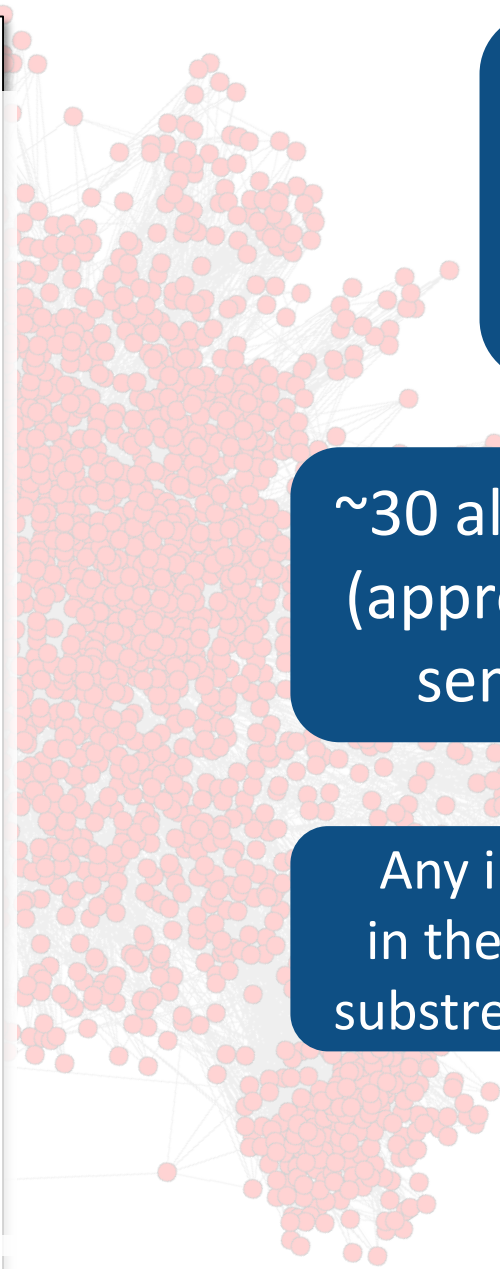
**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...





**Part 3 continued :**  
 Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...

Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[26]	1/2	$O(n)$	1	👎	👍	👍
[41, Theorem 6]	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👎	👍	👍
[41, Theorem 2]	$1/2 + 0.003^*$	$O(n \text{ polylog}(n))$	1	👎	👍	👍
[36, Theorem 1.1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👎	👍	👍
[26, Theorem 1]	$2/3 - \epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👎	👎	👍
[6, Theorem 19]	$1 - \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👎	👎	👍
[41, Theorem 5]	$1/2 + 0.019$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👎	👎	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👎	👎	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👎	👎	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👎	👎	👍
[14]	1	$\tilde{O}(k^2)$	1	👎	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👎	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👎	👎	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	🤔
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	🤔
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	🤔
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	🤔
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	🤔
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👎	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍



Our goals:

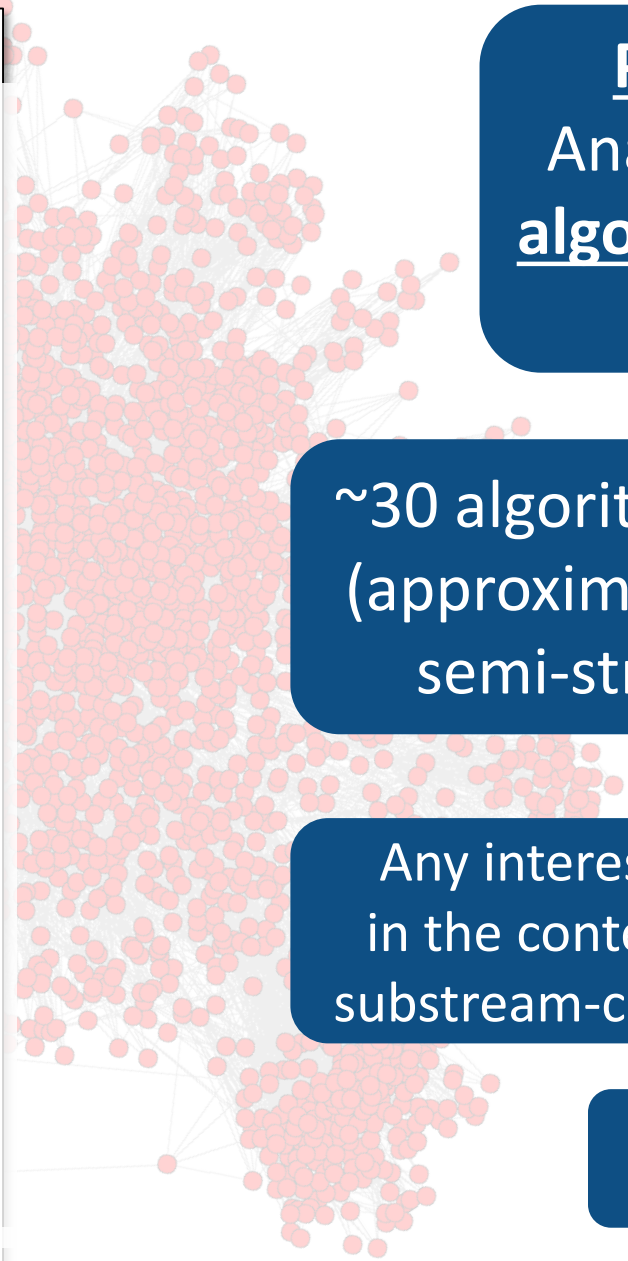
Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
	1/2	$O(n)$	1	👎	👍	👍
	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👎	👍	👍
	$1/2 + 0.003^*$	$O(n \text{ polylog}(n))$	1	👎	👍	👍
	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👎	👍	👍
	$2/3 - \epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👎	👎	👍
[6, Theorem 19]	$1 - \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👎	👎	👍
[41, Theorem 5]	$1/2 + 0.019$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[41, Theorem 1]	$1/2 + 0.005^*$	$O(n \log n)$	1	👎	👎	👍
[41, Theorem 4]	$1/2 + 0.0071^*$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[39]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👎	👎	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👎	👎	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👎	👎	👍
[14]	1	$\tilde{O}(k^2)$	1	👎	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👎	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👎	👎	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👎
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👎
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👎
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👎
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👎
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👎
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👎
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👎
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	👎
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👎	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

**Part 3 continued :**  
 Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Our goals:  
Maximize accuracy

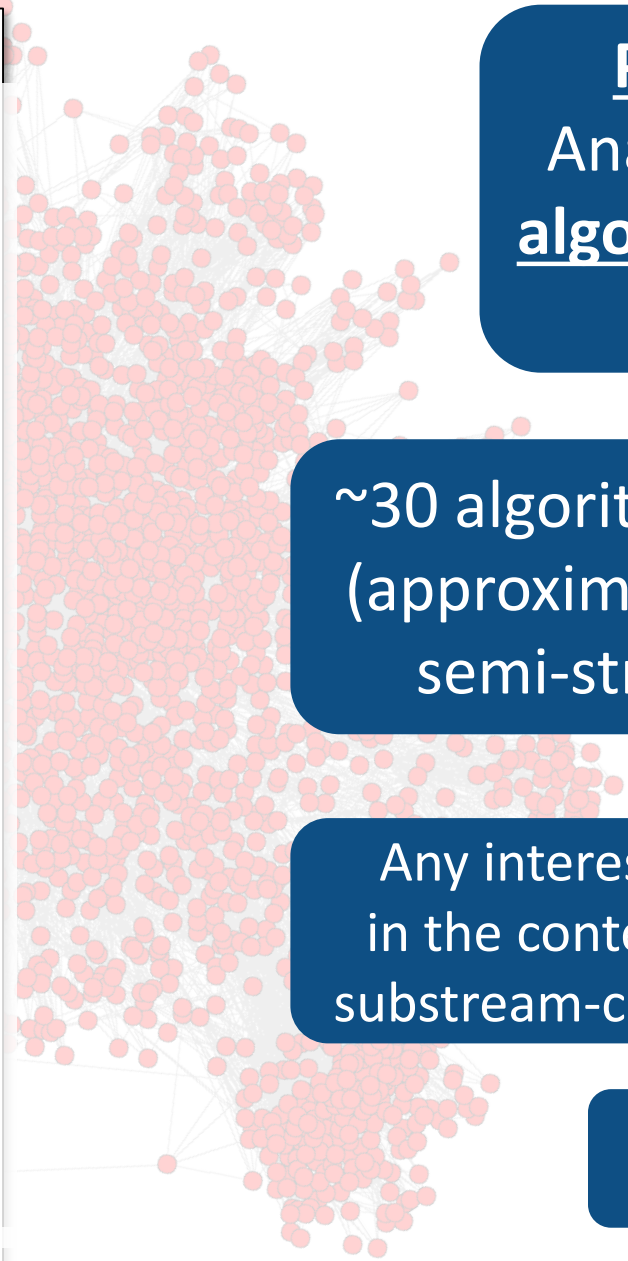
Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[12]	$1/2$	$O(n)$	1	👎	👍	👍
[12]	$1/2 + 0.0071$	$O(n \text{ polylog}(n))$	2	👎	👍	👍
[12]	$1/2 + 0.003$	$O(n \text{ polylog}(n))$	1	👎	👍	👍
[11]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	1	👎	👍	👍
[11]	$\epsilon$	$O(n \log n)$	$O(\log(1/\epsilon) / \epsilon)$	👎	👎	👍
[6, Theorem 19]		$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👎	👎	👍
[4]	$1/2$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[4]	$1/2 + \epsilon$	$O(n \log n)$	1	👎	👎	👍
[4]	$1/2 + \epsilon$	$O(n \text{ polylog}(n))$	2	👎	👎	👍
[35]	$1 - 1/e$	$O(n \text{ polylog}(n))$	1	👎	👎	👍
[28, Theorem 20]	$1 - 1/e$	$O(n)$	1	👎	👎	👍
[35, Theorem 2]	$1 - \frac{e^{-k} k^{k-1}}{(k-1)!}$	$O(n)$	$k$	👎	👎	👍
[14]	1	$\tilde{O}(k^2)$	1	👎	👍	👍
[14]	$1/\epsilon$	$\tilde{O}(n^2 / \epsilon^3)$	1	👎	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👎	👎	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	🤔
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	🤔
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	🤔
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	🤔
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	🤔
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	🤔
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👎	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Part 3 continued :  
Analysis of models and algorithms for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[12]	$1 + 0.0071$	$O(n)$	1	👍	👍	👍
[12]	$1 + 0.0071$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[12]	$1 + 0.003$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[12]	$1 + 0.003$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[12]	$1 + \epsilon$	$O(n \text{ polylog}(n))$	$O(\log(1/\epsilon) / \epsilon)$	👍	👍	👍
[6, Theorem 19]	$1 + \epsilon$	$O(n \text{ polylog}(n) / \epsilon^2)$	$O(\log \log(1/\epsilon) / \epsilon^2)$	👍	👍	👍
[4]	$1 + \epsilon$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[4]	$1 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[4]	$1 + \epsilon$	$O(n \text{ polylog}(n))$	2	👍	👍	👍
[35]	$1 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[28, Theorem 20]	$1 - 1/\epsilon$	$O(n)$	1	👍	👍	👍
[35, Theorem 2]	$1 + \epsilon$	$O(n)$	$k$	👍	👍	👍
[14]	$1 + \epsilon$	$O(n)$	1	👍	👍	👍
[14]	$1/\epsilon$	$O(n^2 / \epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n) / \epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[12]	$1.2$	$O(n)$	$O(\log(1/\epsilon)/\epsilon)$	👍	👍	👍
[12]	$1.2 + 0.0071$	$O(n \text{ polylog}(n))$	$O(\log \log(1/\epsilon)/\epsilon^2)$	👍	👍	👍
[12]	$1 + 0.003$	$O(n \text{ polylog}(n))$	$O(\log \log(1/\epsilon)/\epsilon^2)$	👍	👍	👍
[1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	$O(\log(1/\epsilon)/\epsilon)$	👍	👍	👍
[6, Theorem 19]	$1 - 1/\epsilon$	$O(n \text{ polylog}(n)/\epsilon^2)$	$O(\log \log(1/\epsilon)/\epsilon^2)$	👍	👍	👍
[4]	$2$	$O(n \text{ polylog}(n))$	$2$	👍	👍	👍
[4]	$5$	$O(n \log n)$	$1$	👍	👍	👍
[4]	$71$	$O(n \text{ polylog}(n))$	$2$	👍	👍	👍
[35]	$1 - 1/\epsilon$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[28, Theorem 20]	$1 - 1/\epsilon$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[35, Theorem 2]	$1 - 1/\epsilon$	$O(n \text{ polylog}(n))$	$k$	👍	👍	👍
[14]	$1/\epsilon$	$O(n^2/\epsilon^3)$	$1$	👍	👍	👍
[14]	$1/\epsilon$	$O(n^2/\epsilon^3)$	$1$	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	$1$	👍	👍	👍
[26, Theorem 2]	$6$	$O(n \log n)$	$1$	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	$5.82$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[63]	$5.58$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	$1$	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	$1$	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n)/\epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	$1$	👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

Minimize #passes

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O\left(n \log(n) / \epsilon^4\right)$	$O\left(\epsilon^{-4} \log n\right)$	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O\left(\log_{1+\epsilon/3} n\right)$	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O\left(n \log^2 n\right)$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}\left(n^{2-3\epsilon} + n^{1-\epsilon}\right)$	1	👍	👍	👍
[14]	$1/\epsilon$	$O\left(n^2/\epsilon^3\right)$	1	👍	👍	👍
[14]			1	👍	👍	👍
[14]			1	👍	👍	👍
[35, Theorem 2]			$k$	👍	👍	👍
[28, Theorem 20]	$1 - 1/\epsilon$			👍	👍	👍
[35]				👍	👍	👍
[4]				👍	👍	👍
[4]				👍	👍	👍
[4]				👍	👍	👍
[4]				👍	👍	👍
[6, Theorem 19]				👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

Minimize #passes

Accept weighted graphs

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[12]	$1 + 2\epsilon$	$O(n)$	$O(1/\epsilon)$	👍	👍	👍
[12]	$1.2 + 0.0071\epsilon$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[12]	$1 + 0.003\epsilon$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[1]	$O(\text{polylog}(n))$	$O(\text{polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[6, Theorem 19]	$1 - 1/e$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[4]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[4]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[4]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[35]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[28, Theorem 20]	$1 - 1/e$	$O(n \text{ polylog}(n))$	$O(1/\epsilon)$	👍	👍	👍
[35, Theorem 2]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	$k$	👍	👍	👍
[14]	$O(1/\epsilon)$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[14]	$1/\epsilon$	$O(n^2/\epsilon^3)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}(n^{2-3\epsilon} + n^{1-\epsilon})$	1	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O(n \log^2 n)$	1	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O(\log_{1+\epsilon/3} n)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O(n \log(n)/\epsilon^4)$	$O(\epsilon^{-4} \log n)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O(n (\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}))$	$O(\epsilon^{-2} \log(\epsilon^{-1}))$	👍	👍	👍
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

Minimize #passes

Accept weighted graphs

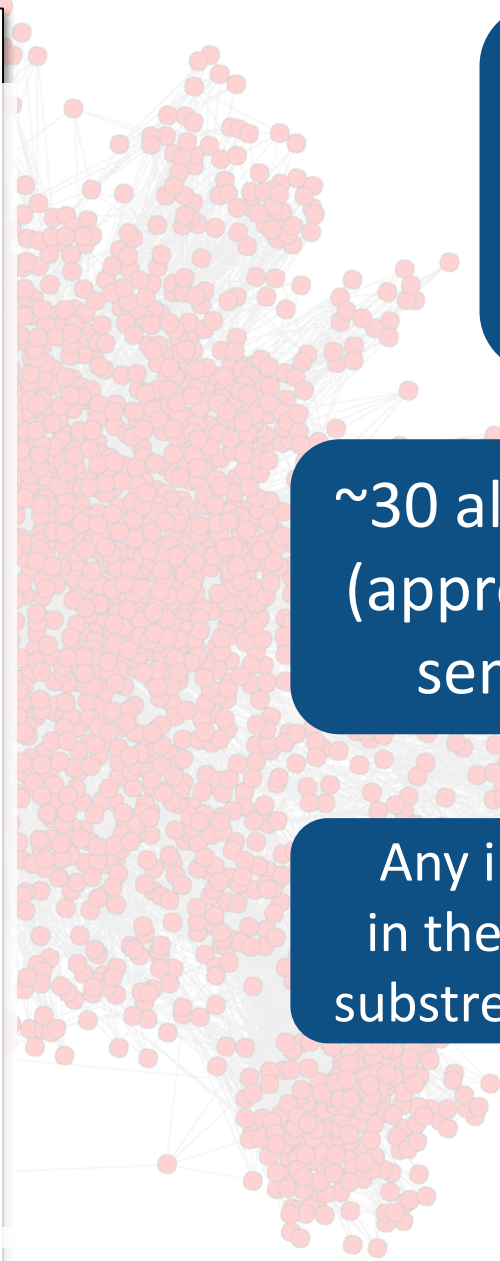
Accept general (not just bipartite) graphs

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O\left(n \log(n) / \epsilon^4\right)$	$O\left(\epsilon^{-4} \log n\right)$	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O\left(\log_{1+\epsilon/3} n\right)$	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O\left(n \log^2 n\right)$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}\left(n^{2-3\epsilon} + n^{1-\epsilon}\right)$	1	👍	👍	👍
[14]	$1/\epsilon$	$O\left(n^2/\epsilon^3\right)$	1	👍	👍	👍
[14]			1	👍	👍	👍
[35, Theorem 2]			$k$	👍	👍	👍
[28, Theorem 20]	$1 - 1/e$			👍	👍	👍
[3, ...]				👍	👍	👍
[4, ...]				👍	👍	👍
[4, ...]				👍	👍	👍
[4, ...]				👍	👍	👍
[4, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍
[1, ...]				👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

Minimize #passes

Accept weighted graphs

Accept general (not just bipartite) graphs

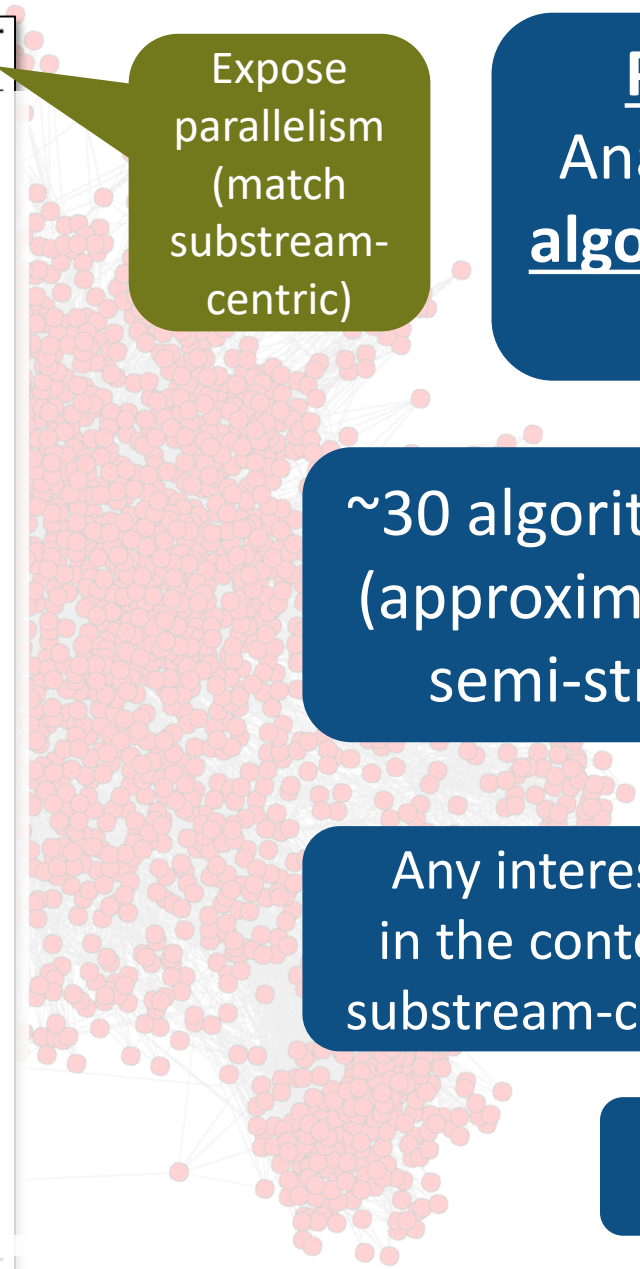
Expose parallelism (match substream-centric)

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...



Reference	Approx.	Space	#Passes	Wgh <sup>1</sup>	Gen <sup>2</sup>	Par <sup>3</sup>
[17]	$4 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[6, Theorem 22]	$\frac{1}{1-\epsilon}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 22]	$\frac{1}{\frac{2}{3}(1-\epsilon)}$	$O\left(n \left(\frac{\epsilon \log n - \log \epsilon}{\epsilon^2}\right)\right)$	$O\left(\epsilon^{-2} \log\left(\epsilon^{-1}\right)\right)$	👍	👍	👍
[6, Theorem 28]	$\frac{1}{1-\epsilon}$	$O\left(\frac{n \log(n)}{\epsilon^4}\right)$	$O\left(\epsilon^{-4} \log n\right)$	👍	👍	👍
[26, Section 3.2]	$2 + \epsilon$	$O(n \log n)$	$O\left(\log_{1+\epsilon/3} n\right)$	👍	👍	👍
[27]	$2 + \epsilon$	$O(n \log n)$	1	👍	👍	👍
[53]	$2 + \epsilon$	$O\left(n \log^2 n\right)$	1	👍	👍	👍
[29]	$3.5 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[25]	$4.911 + \epsilon$	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[63]	5.58	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	5.82	$O(n \text{ polylog}(n))$	1	👍	👍	👍
[44, Theorem 3]	$2 + \epsilon$	$O(n \text{ polylog}(n))$	$O(1)$	👍	👍	👍
[26, Theorem 2]	6	$O(n \log n)$	1	👍	👍	👍
[7, Theorem 1]	$n^\epsilon$	$\tilde{O}\left(n^{2-3\epsilon} + n^{1-\epsilon}\right)$	1	👍	👍	👍
[14]	$1/\epsilon$	$O\left(n^2/\epsilon^3\right)$	1	👍	👍	👍
[14]			1	👍	👍	👍
[35, Theorem 2]			$k$	👍	👍	👍
[28, Theorem 20]	$1 - 1/e$			👍	👍	👍

Our goals:

Maximize accuracy

Minimize local space

Minimize #passes

Accept weighted graphs

Accept general (not just bipartite) graphs

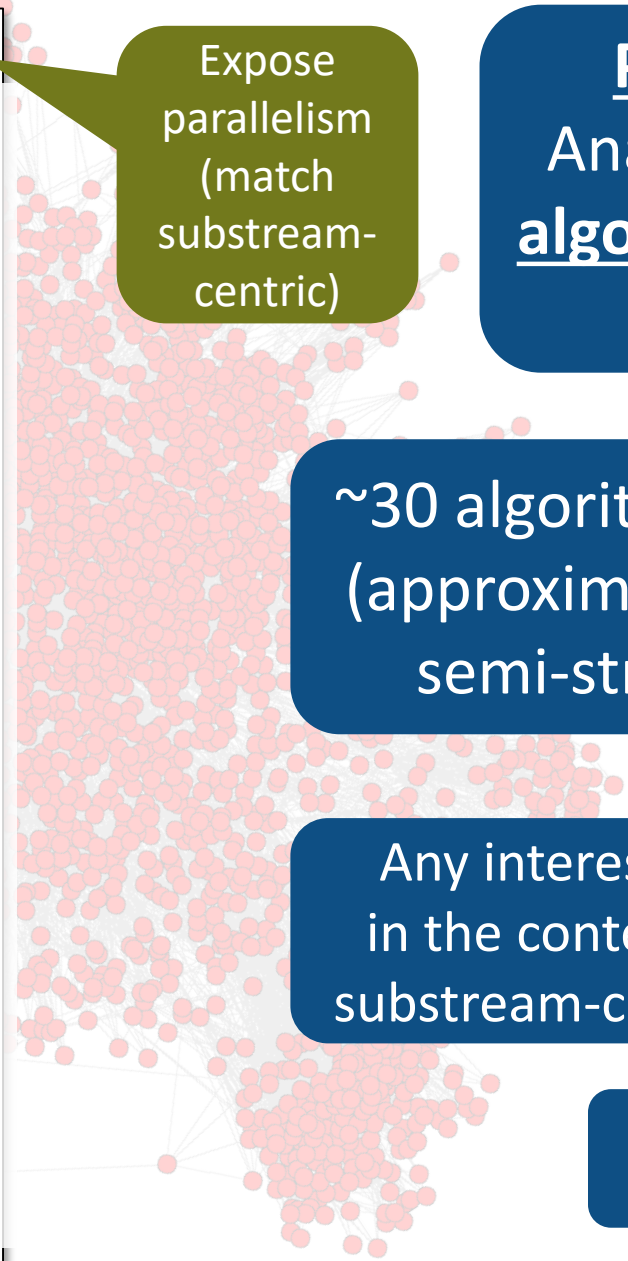
Expose parallelism (match substream-centric)

**Part 3 continued :**  
Analysis of models and **algorithms** for streaming graph processing

~30 algorithms for streaming (approximate) MWM (in the semi-streaming model)

Any interesting idea to use in the context of FPGAs and substream-centric processing?

More specifically...





# Research Questions



How to design a high-performance MWM algorithm (as dictated by the used paradigm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



Use the MWM algorithm by Crouch and Stubbs in the substream-centric paradigm in a hybrid CPU-FPGA setting (algorithm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



Use the MWM algorithm by Crouch and Stubbs in the substream-centric paradigm in a hybrid CPU-FPGA setting (algorithm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model

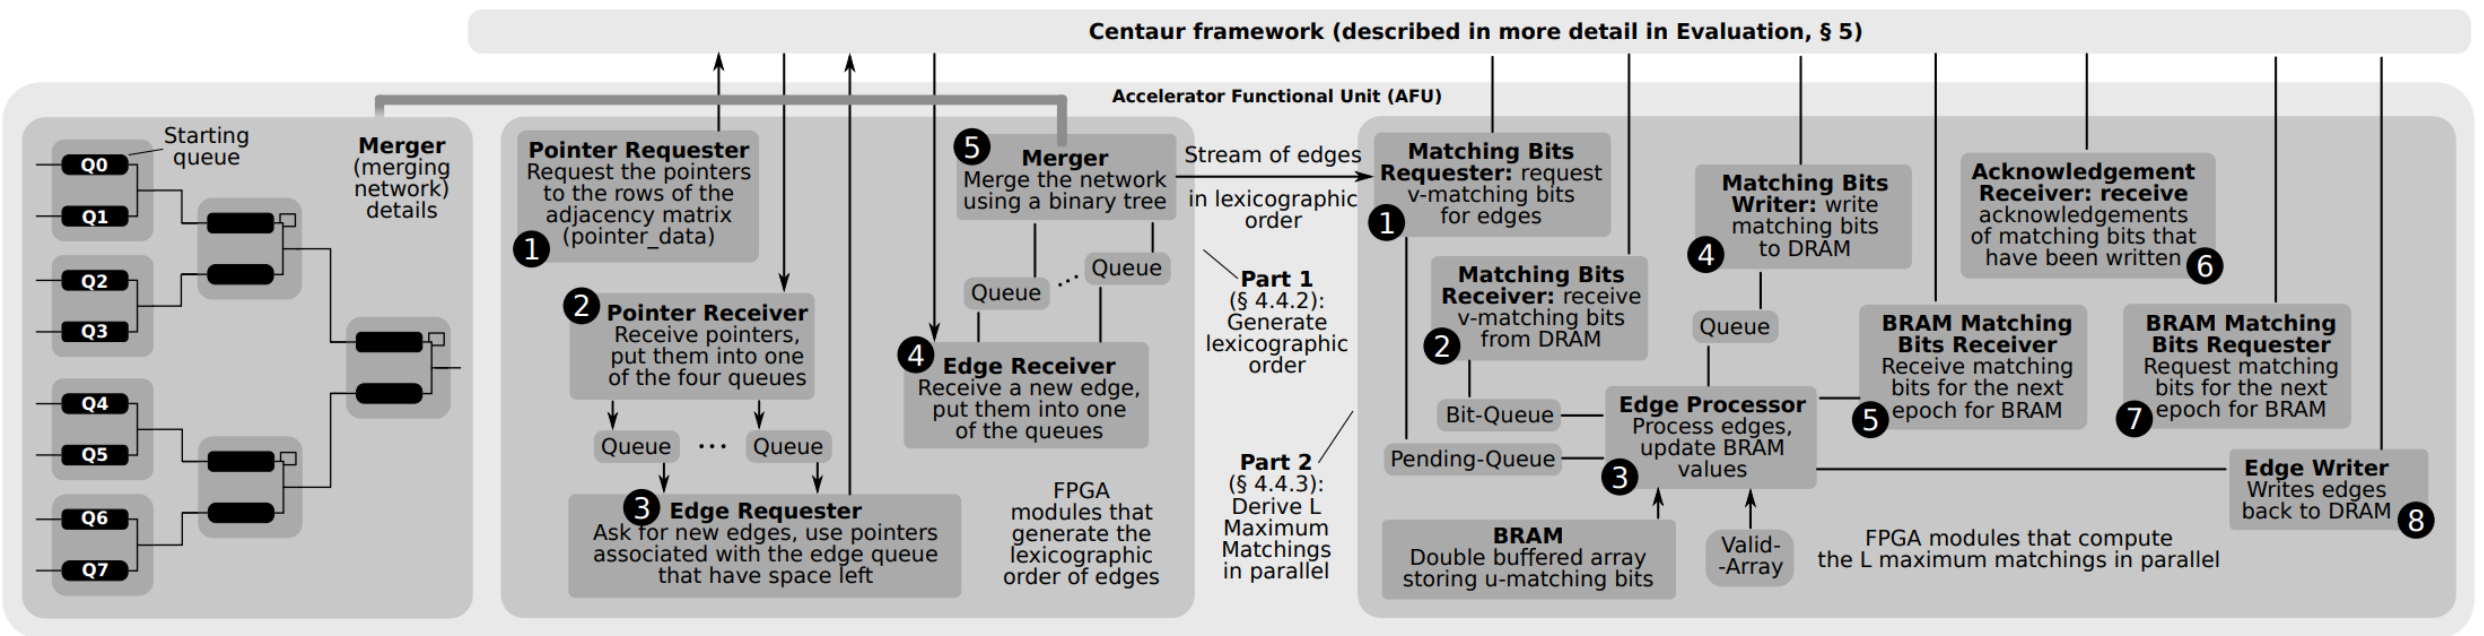


What is the HW FPGA design that ensures high performance?

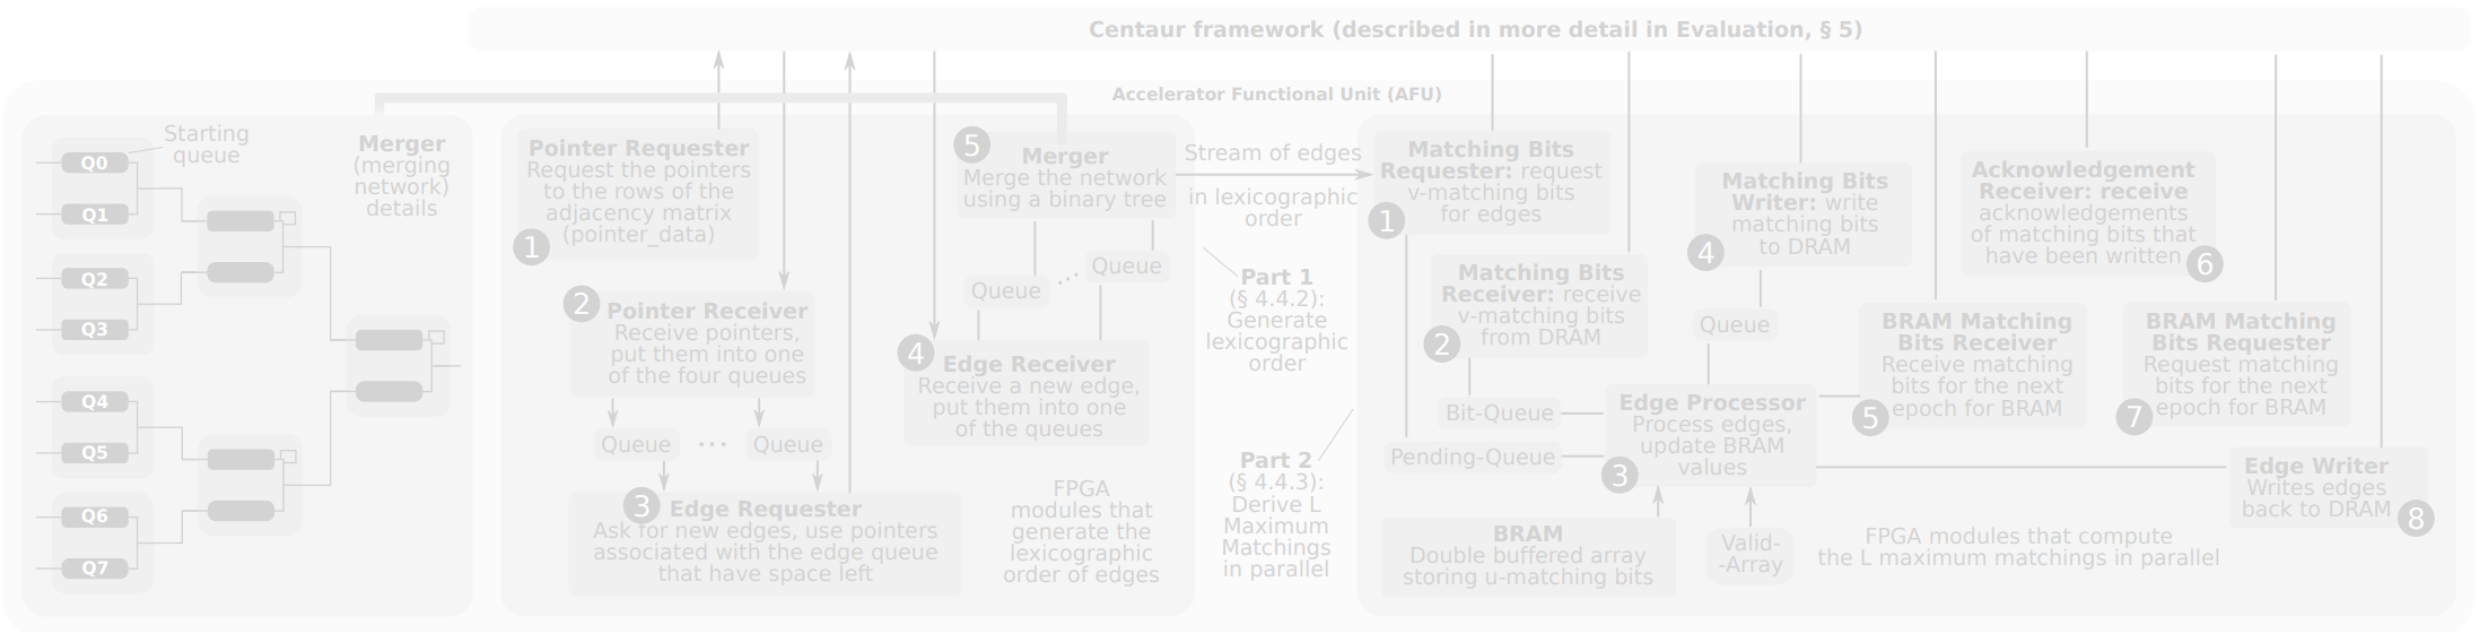


What is the ultimate performance, power consumption, and the related tradeoffs?

# Substream-Centric MWM: FPGA optimizations

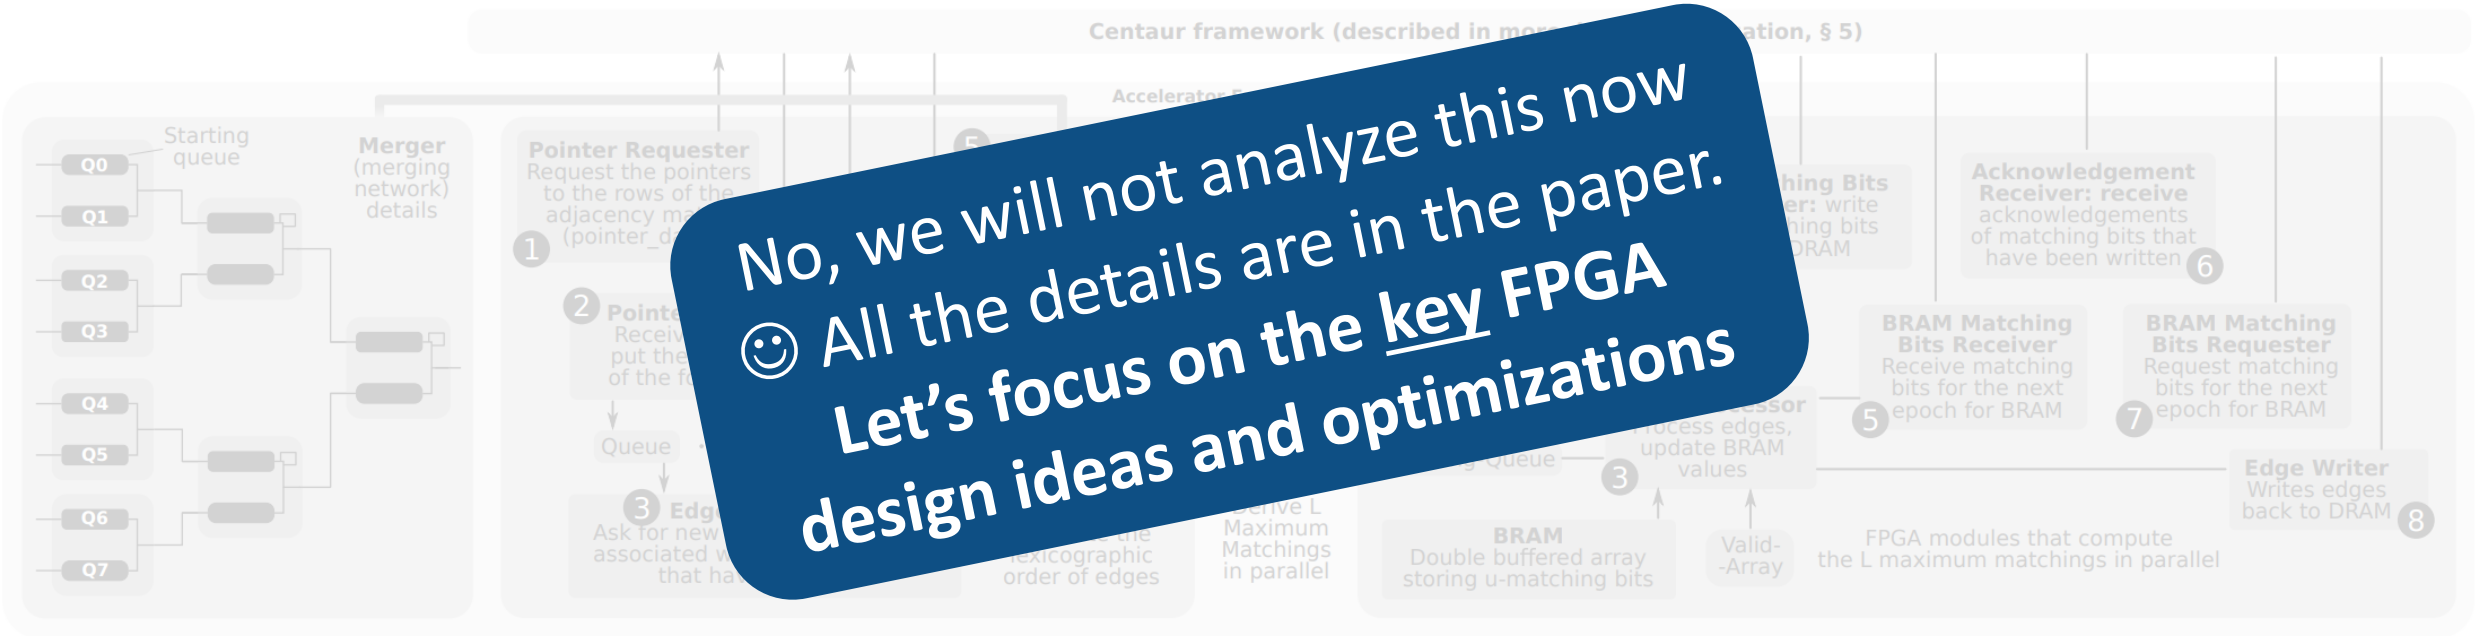


# Substream-Centric MWM: FPGA optimizations



# Substream-Centric MWM: FPGA optimizations

No, we will not analyze this now  
☺ All the details are in the paper.  
Let's focus on the key FPGA  
design ideas and optimizations




# Substream-Centric MWM: FPGA optimizations

Vectorization

Blocking / Tiling

No, we will not analyze this now  
😊 All the details are in the paper.  
Let's focus on the key FPGA  
design ideas and optimizations



Prefetching


Pipelining

# Substream-Centric MWM: FPGA optimizations

Blocking / Tiling

Vectorization

No, we will not analyze this now  
😊 All the details are in the paper.  
Let's focus on the key FPGA design ideas and optimizations



Prefetching

They are often used in graph processing schemes on FPGAs; we apply them as well.

Pipelining




# Substream-Centric MWM: FPGA optimizations

Blocking / Tiling

Vectorization

No, we will not analyze this now  
😊 All the details are in the paper.  
Let's focus on the key FPGA  
design ideas and optimizations

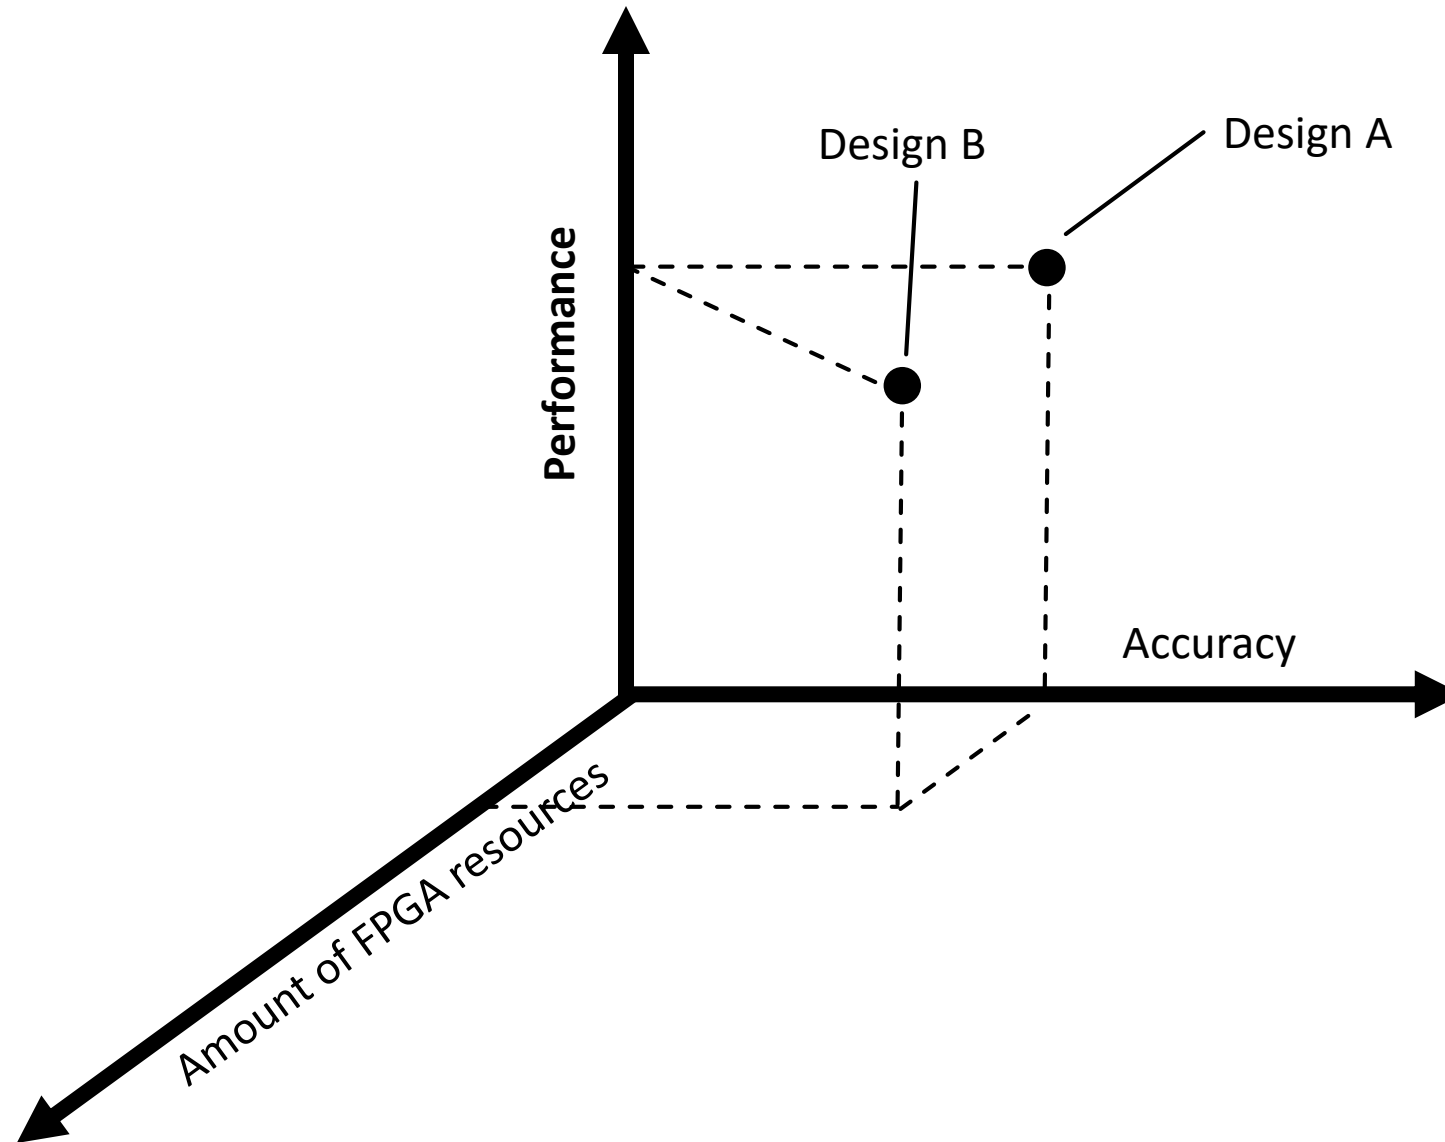


Prefetching

They are often used in graph processing schemes on FPGAs; we apply them as well.

Pipelining

# THE SPACE OF SUBSTREAM-CENTRIC



# Research Questions



Use the MWM algorithm by Crouch and Stubbs in the substream-centric paradigm in a hybrid CPU-FPGA setting (algorithm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model



What is the HW FPGA design that ensures high performance?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



Use the MWM algorithm by Crouch and Stubbs in the substream-centric paradigm in a hybrid CPU-FPGA setting (algorithm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model

The proper use of blocking, vectorization, and others (pipelining, prefetching) (performance)?



What is the ultimate performance, power consumption, and the related tradeoffs?

# Research Questions



Use the MWM algorithm by Crouch and Stubbs in the substream-centric paradigm in a hybrid CPU-FPGA setting (algorithm)?

Use substream-centric processing (exposes parallelism)

For enabling theoretical analysis and rigor in the context of FPGAs (small local space), use the semi-streaming model

The proper use of blocking, vectorization, and others (pipelining, prefetching) performance?




What is the ultimate performance, power consumption, and the related tradeoffs?

# PERFORMANCE ANALYSIS

## VARIOUS GRAPHS

Parameters:  
 Blocking size = 32, #Substreams = 64  
 #Threads = 4,  $\epsilon = 0.1$

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid



# PERFORMANCE ANALYSIS



## VARIOUS GRAPHS

### Parameters:

Blocking size = 32, #Substreams = 64  
 #Threads = 4,  $\epsilon = 0.1$

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k; k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid



# PERFORMANCE ANALYSIS



## VARIOUS GRAPHS

**Parameters:**  
 Blocking size = 32, #Substreams = 64  
 #Threads = 4,  $\epsilon = 0.1$

**SC-OPT secures highest performance**

Graph	Type	$m$	$n$
Kronecker	Synthetic power-law	$\approx 48n$	$2^k; k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid





# PERFORMANCE ANALYSIS

## VARIOUS THREAD (CPU) COUNTS

---

### Algorithm

---

Crouch et al. [1] Sequential (CS-SEQ)

Crouch et al. [1] Parallel (CS-PAR)

Ghaffari [2] Sequential (G-SEQ)

Substream-Centric, no blocking (SC-SIMPLE)

Substream-Centric, with blocking (SC-OPT)

---



Blocking size = 32, #Substreams = 64  
 #edges = 16M (Kronecker),  $\epsilon = 0.1$

# PERFORMANCE ANALYSIS

## VARIOUS THREAD (CPU) COUNTS

### Algorithm

- Crouch et al. [1] Sequential (CS-SEQ)
- Crouch et al. [1] Parallel (CS-PAR)
- Ghaffari [2] Sequential (G-SEQ)
- Substream-Centric, no blocking (SC-SIMPLE)
- Substream-Centric, with blocking (SC-OPT)





# PERFORMANCE ANALYSIS

## VARIOUS THREAD (CPU) COUNTS

Blocking size = 32, #Substreams = 64  
 #edges = 16M (Kronecker),  $\epsilon = 0.1$

SC-OPT secures highest performance for all considered numbers of threads

Algorithm
Crouch et al. [1] Sequential (CS-SEQ)
Crouch et al. [1] Parallel (CS-PAR)
Ghaffari [2] Sequential (G-SEQ)
Substream-Centric, no blocking (SC-SIMPLE)
Substream-Centric, with blocking (SC-OPT)



# PERFORMANCE ANALYSIS

## VARIOUS BLOCKING SIZE (K) AND #SUBSTREAMS (L)

---

### Algorithm

Crouch et al. [1] Sequential (CS-SEQ)

Crouch et al. [1] Parallel (CS-PAR)

Ghaffari [2] Sequential (G-SEQ)

Substream-Centric, no blocking (SC-SIMPLE)

Substream-Centric, with blocking (SC-OPT)

---

# PERFORMANCE ANALYSIS



## VARIOUS BLOCKING SIZE (K) AND #SUBSTREAMS (L)

### Algorithm

- Crouch et al. [1] Sequential (CS-SEQ)
- Crouch et al. [1] Parallel (CS-PAR)
- Ghaffari [2] Sequential (G-SEQ)
- Substream-Centric, no blocking (SC-SIMPLE)
- Substream-Centric, with blocking (SC-OPT)

Blocking size (K) = 32, #threads = 4,  
#edges = 16M (Kronecker),  $\epsilon = 0.1$

#Substreams (L) = 128, #threads = 4,  
#edges = 16M (Kronecker),  $\epsilon = 0.1$



# PERFORMANCE ANALYSIS

## VARIOUS BLOCKING SIZE (K) AND #SUBSTREAMS (L)



### Algorithm

- Crouch et al. [1] Sequential (CS-SEQ)
- Crouch et al. [1] Parallel (CS-PAR)
- Ghaffari [2] Sequential (G-SEQ)
- Substream-Centric, no blocking (SC-SIMPLE)
- Substream-Centric, with blocking (SC-OPT)

SC-OPT secures highest performance for all considered values of parameters

Blocking size (K) = 32, #threads = 4, #edges = 16M (Kronecker),  $\epsilon = 0.1$

#Substreams (L) = 128, #threads = 4, #edges = 16M (Kronecker),  $\epsilon = 0.1$



# PERFORMANCE ANALYSIS

## APPROXIMATION

---

### Algorithm

---

Crouch et al. [1] Sequential (CS-SEQ)  
Crouch et al. [1] Parallel (CS-PAR)  
Ghaffari [2] Sequential (G-SEQ)  
Substream-Centric, no blocking (SC-SIMPLE)  
Substream-Centric, with blocking (SC-OPT)

---

# PERFORMANCE ANALYSIS



## APPROXIMATION

### Algorithm

Crouch et al. [1] Sequential (CS-SEQ)  
 Crouch et al. [1] Parallel (CS-PAR)  
 Ghaffari [2] Sequential (G-SEQ)  
 Substream-Centric, no blocking (SC-SIMPLE)  
 Substream-Centric, with blocking (SC-OPT)

#Substreams (L) = 128, Blocking size (K) = 32,  
 #threads = 4, #edges = 8M (Kronecker)

Blocking size (K) = 32, #threads = 4,  
 #Substreams (L) = 128,  $\epsilon = 0.1$





# PERFORMANCE ANALYSIS


## APPROXIMATION

### Algorithm


- Crouch et al. [1] Sequential (CS-SEQ)
- Crouch et al. [1] Parallel (CS-PAR)
- Ghaffari [2] Sequential (G-SEQ)
- Substream-Centric, no blocking (SC-SIMPLE)
- Substream-Centric, with blocking (SC-OPT)

SC-OPT is comparable to the  $(2+\epsilon)$ -approximation by Ghaffari et al.


#Substreams (L) = 128, Blocking size (K) = 32, #threads = 4, #edges = 8M (Kronecker)




Blocking size (K) = 32, #threads = 4, #Substreams (L) = 128,  $\epsilon = 0.1$



# OTHER ALGORITHMS, PROBLEMS, ANALYSES




# OTHER ALGORITHMS, PROBLEMS, ANALYSES



# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊



## OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough  
time to present 😊

In addition to MWM, we  
also analyzed many more  
graph problems



# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems



Model	Prob.	Input	Order	Reference	Approx.	Time	Space	Method	Passes
ION	Rand.	Unw. Graph	Adv.	[17]	$n^\beta / \#$			Sampling	1
DGS	Rand.	Unw. Graph	Adv.	[6]	$n^{-2} \log$			Sampling	1
TN	Rand.	Unw. Graph	Adv.	[57]	$k \Delta^k /$			Sampling	1
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]	$1 + \epsilon$	$O(n^{1/\epsilon})$		Sampling	2
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]	$1 + \epsilon$	$O(1/\epsilon^2 n \text{ polylog}(n))$		Sampling	1
ION	Det.	Unw. Graph	Adv.	[53, Theorem 4.2]	$1^{***}$	$O(kn \log^3(n))$		Certificate	1
DGS	Rand.	Unw. Graph	Adv.	[53, Theorem 4.2]	$1^{***}$	$\tilde{O}(n^{1-\epsilon+\epsilon^2})$		Certificate	1
ION	Det.	Unw. Graph	Adv.	[37]	$1/2$			Method	1
ION	Det.	Unw. Graph	Adv.	[75, Theorem 3]	$2 + \epsilon$	$O(n \log(n))$		Greedy	1
ION	Unw. Bi. Graph	Adv.		[75, Theorem 3]	$2 + \epsilon$	$O(n \log(n))$		Greedy	1
ALO	Rand.	Unw. Graph	Adv.	[79, Theorem 5]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$		Sampling	1
ALO	Rand.	Unw. Graph	Adv.	[79, Theorem 5]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$		Sampling	1
ALO	Rand.	Unw. Graph	Adv.	[22]	$(\epsilon, \delta)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$		Sampling	1
ALO	Rand.	Unw. Graph	Adv.	[79, Theorem 10]	$(\epsilon, 1/50)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$		Sampling	2
ALO	Rand.	Unw. Graph	Adv.	[79, Theorem 14]	$(\epsilon, 1/100)$	$\tilde{O}(\epsilon^{-2} m / \sqrt{T})$		Sampling	2
ION	Rand.	Unw. Graph	Adv.	[17, Corollary 3.9]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	1
ION	Rand.	Unw. Graph	Adv.	[17, Theorem 3.1]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	1
ION	Rand.	Unw. Graph	Adv.	[86, Theorem 3.3]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	1
ION	Rand.	Unw. Graph	Adv.	[55, Theorem 1]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	1
ION	Rand.	Unw. Graph	Adv.	[55, Theorem 2]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	1
ION	Rand.	Unw. Graph	Adv.	[55, Theorem 3]	$(\epsilon, 1/3)$	$\tilde{O}(m^{3/2}/T)$		Sampling	3
ION	Rand.	Unw. Graph	Adv.	[26, Theorem 5]	$(\epsilon, 1/2)$	$O(m/(\epsilon^{2.5} \sqrt{T}))$	$\text{polylog}(n)^*$	Sampling	2
ION	Rand.	Unw. Graph	Adv.	[26, Corollary 6]	$(1/3 + \epsilon, -)$	$O(m/(\epsilon^{4.5} \sqrt{T}))$		Sampling	1
				[90, Theorem 3.1]		$\tilde{O}(n\alpha + \sqrt{l/\alpha})$		Sampling	1

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Spanners

Triangle counting

Connectivity

Densest subgraphs

K-edge connectivity

Colorings

Bipartiteness

K-vertex connectivity

Spectral sparsification

Cut sparsification

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

Model	Prob.	Input	Order	Reference	Approx.	Method
ION	Rand.	Unw. Graph	Adv.	[17]	$n^\beta / \#$	Sampling
DGS	Rand.	Unw. Graph	Adv.	[6]	$-2 \log$	Sampling
TN	Rand.	Unw. Graph	Adv.	[57]	$k \Delta^k /$	Sampling
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]	$1 + \epsilon O(n^{1/\epsilon})$	Sampling
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]	$O(1/\epsilon^2 n \text{ polylog}(n))$	Sampling
ION	Det.	Unw. Graph	Adv.		$\log^2$	Certificate
DGS	Rand.	Unw. Graph	Adv.		$\text{polylog}(n)$	Certificate
ION	Det.	Unw. Graph	Adv.		$O(kn \log^3(n))$	$\ell_0$
ION	Det.	Unw. Graph	Adv.	[53, Theorem 4.2]	$1^{***}$	$\ell_0$
ION	Det.	Unw. Graph	Adv.		1	

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

1

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
 MACIEJ BESTA, Department of Computer Science, ETH Zurich  
 TAL BEN-NUN, Department of Computer Science, ETH Zurich  
 TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

---

Graph processing has become an important part of various areas of computer science, including machine learning, social network analysis, computational sciences, and others. Two key challenges that hinder accelerating graph processing are (1) sizes of input datasets, reaching trillions of edges, and (2) the growing rate of graph

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

Model	Prob.	Input	Order	Reference	Approx.	Time	Method
ION	Rand.	Unw. Graph	Adv.	[17]		$n^\beta / \#$	Sampling
DGS	Rand.	Unw. Graph	Adv.	[6]		$-2 \log$	Sampling
TN	Rand.	Unw. Graph	Adv.	[57]		$k \Delta^k$	Sampling
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]		$1 + \epsilon O(n^{1/\epsilon})$	Sampling
ION	Det.	Unw. Graph	Adv.			$\log(n)$	Certificate
DGS	Rand.	Unw. Graph	Adv.			$\log(n)$	Certificate
ION	Det.	Unw. Graph	Adv.			$O(kn \log^3(n))$	$\ell_0$
ION	Det.	Unw. Graph	Adv.	[53, Theorem 4.2]		$1^{***}$	$\ell_0$

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science  
 MACIEJ BESTA, Department of Computer Science  
 TAL BEN-NUN, Department of Computer Science  
 TORSTEN HOEFLER, Department of Computer Science

Graph processing has become an important part of various applications, including machine learning, social network analysis, computational sciences, and data science. The challenges that hinder accelerating graph processing are (1) sizes of input datasets, reachability of edges, and (2) the growing rate of graph

Again, the most relevant parts are in the FPGA paper, the rest in this survey (ready in ~1-2 months)



# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

Preliminary substream-centric designs and results

Model	Prob.	Input	Order	Reference	Approx.	Time	Passes	Method
ION	Rand.	Unw. Graph	Adv.	[17]				Sampling
DGS	Rand.	Unw. Graph	Adv.	[6]				Sampling
TN	Rand.	Unw. Graph	Adv.	[57]				Sampling
DGS	Rand.	W. Graph	Adv.	[61, Theorem 1, Section 6]				Sampling
ION	Det.	Unw. Graph	Adv.					Certificate
DGS	Rand.	Unw. Graph	Adv.					Certificate
ION	Det.	Unw. Graph	Adv.					Method

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science  
 MACIEJ BESTA, Department of Computer Science  
 TAL BEN-NUN, Department of Computer Science  
 TORSTEN HOEFLER, Department of Computer Science

Graph processing has become an important part of various applications, including machine learning, social network analysis, computational sciences, and data science. The challenges that hinder accelerating graph processing are (1) sizes of input datasets, reaching billions of edges, and (2) the growing rate of graph

Again, the most relevant parts are in the FPGA paper, the rest in this survey (ready in ~1-2 months)

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

Preliminary substream-centric designs and results



Work in progress on the distributed setting 😊

Model	Prob.	Input	Order	Reference	Approx	Time	Passes	Method
DGS	Rand. W. Graph	Adv.	[17]					Sampling
		Adv.	[6]					Sampling
	Unw. Graph	Adv.	[57]					Sampling
			[61, Theorem 1, Section 6]					Sampling
								Certificate
								Certificate
								$\ell_0$
								$\ell_0$

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science  
 MACIEJ BESTA, Department of Computer Science  
 TAL BEN-NUN, Department of Computer Science  
 TORSTEN HOEFLER, Department of Computer Science

Graph processing has become an important part of various applications, including machine learning, social network analysis, computational sciences, and data science. The challenges that hinder accelerating graph processing are (1) sizes of input datasets, reachability of edges, and (2) the growing rate of graph

Again, the most relevant parts are in the FPGA paper, the rest in this survey (ready in ~1-2 months)

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

stream- results



Work in progress on the distributed setting 😊

Model	Prob.	Input	Order	Reference	Approx	Time	Passes	Method
DGS	Rand. W. Graph	Adv.	[17]					Sampling
		Adv.	[6,					Sampling
		Unw. Graph	Adv.	[57,				Sampling
				[61, Theorem 1, Section 6]				Sampling
								Certificate
								Certificate
				[53, Theorem 4.2]				

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

1

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various applications, including machine learning, social network analysis, computational sciences, and data science. The challenges that hinder accelerating graph processing are (1) sizes of input datasets, reachability of edges, and (2) the growing rate of graph

Again, the most relevant parts are in the FPGA paper, the rest in this survey (ready in ~1-2 months)

### Graph Processing on FPGAs: Taxonomy, Survey, Challenges

Towards Understanding of Modern Graph Processing, Storage, and Analytics

MACIEJ BESTA\*, DIMITRI STANOJEVIC\*, Department of Computer Science, ETH Zurich  
JOHANNES DE FINE LICHT, TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

# OTHER ALGORITHMS, PROBLEMS, ANALYSES

Well, not enough time to present 😊

In addition to MWM, we also analyzed many more graph problems

Work in progress on the distributed setting 😊



Model	Prob.	Input	Order	Reference	Approx	Time	Passes	Method
DGS	Rand. W. Graph	Adv.	[17]					Sampling
		Adv.	[6,					Sampling
		Unw. Graph	Adv.	[57,				Sampling
			[61, Theorem 1, Section 6]					Sampling
								Certificate
								Certificate
								1
								1
								1
								1

Triangle counting

Connected components

Motif counting

Minimum spanning trees

Random walks

Triangle counting

Connectivity

😊 <https://arxiv.org/abs/...>

1

## Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing

Towards Understanding of Modern Graph Processing and Storage

MARC FISCHER, Department of Computer Science, ETH Zurich  
MACIEJ BESTA, Department of Computer Science, ETH Zurich  
TAL BEN-NUN, Department of Computer Science, ETH Zurich  
TORSTEN HOEFLER, Department of Computer Science, ETH Zurich

Graph processing has become an important part of various applications, including machine learning, social network analysis, computational sciences, and data science. The challenges that hinder accelerating graph processing are (1) sizes of input datasets, reaching billions of edges, and (2) the growing rate of graph

Again, the most relevant parts are in the FPGA paper, the rest in this survey (ready in ~1-2 months)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...

## MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

1	0
0	0
1	0
0	0
0	0



# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

1	0
0	0
1	0
0	0
0	0

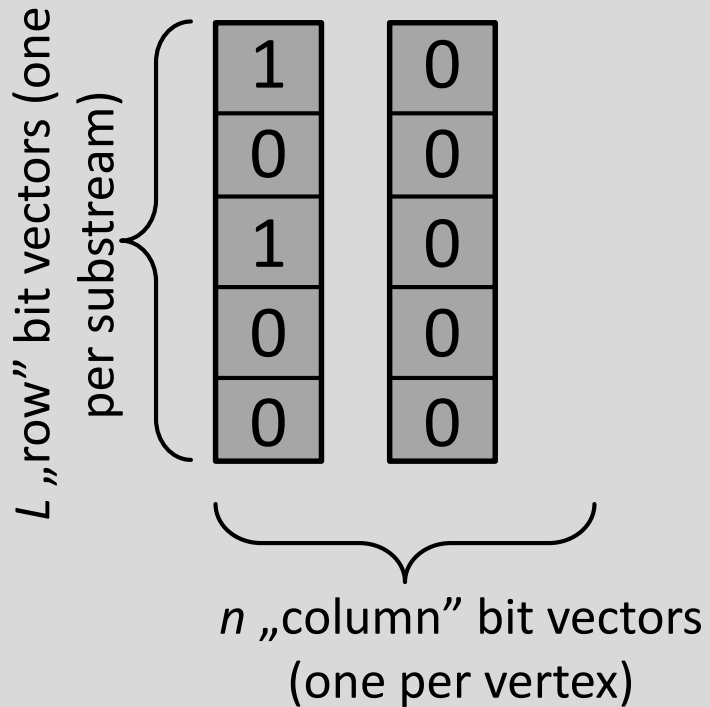


$n$  „column” bit vectors  
(one per vertex)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)



# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

( $L = 5$ )

$L$  „row“ bit vectors (one per substream)

1	0
0	0
1	0
0	0
0	0

$n$  „column“ bit vectors (one per vertex)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

( $L = 5$ )

$L$  „row” bit vectors (one per substream)

1	0
0	0
1	0
0	0
0	0

$u$ -th vector


$n$  „column” bit vectors (one per vertex)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

( $L = 5$ )

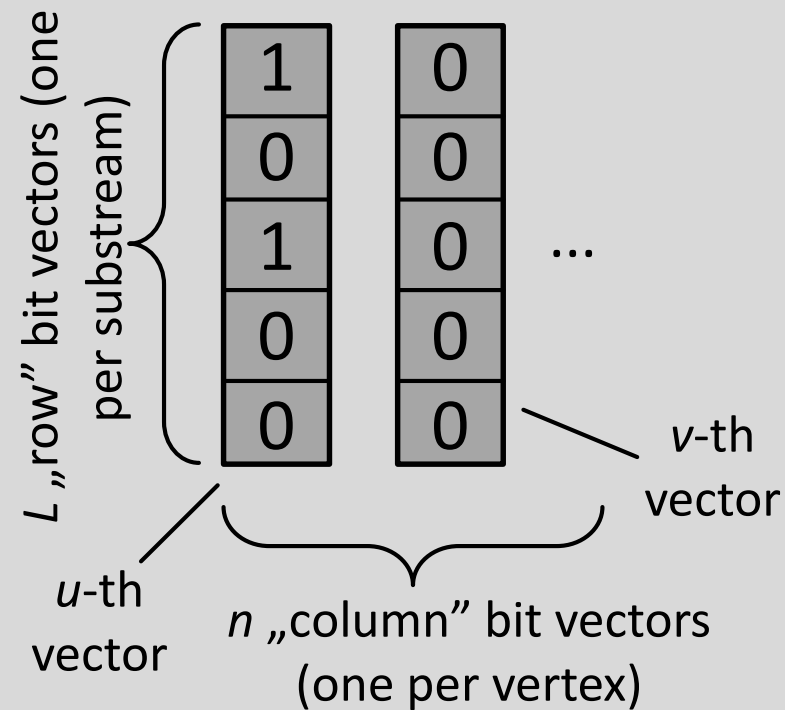


# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

Vertices + matchings (correctness)

( $L = 5$ )




# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)


Vertex  $u$  is a part of a matching in substreams 0, 2



# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)






# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)




## Edges + matchings (more performance)

# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)



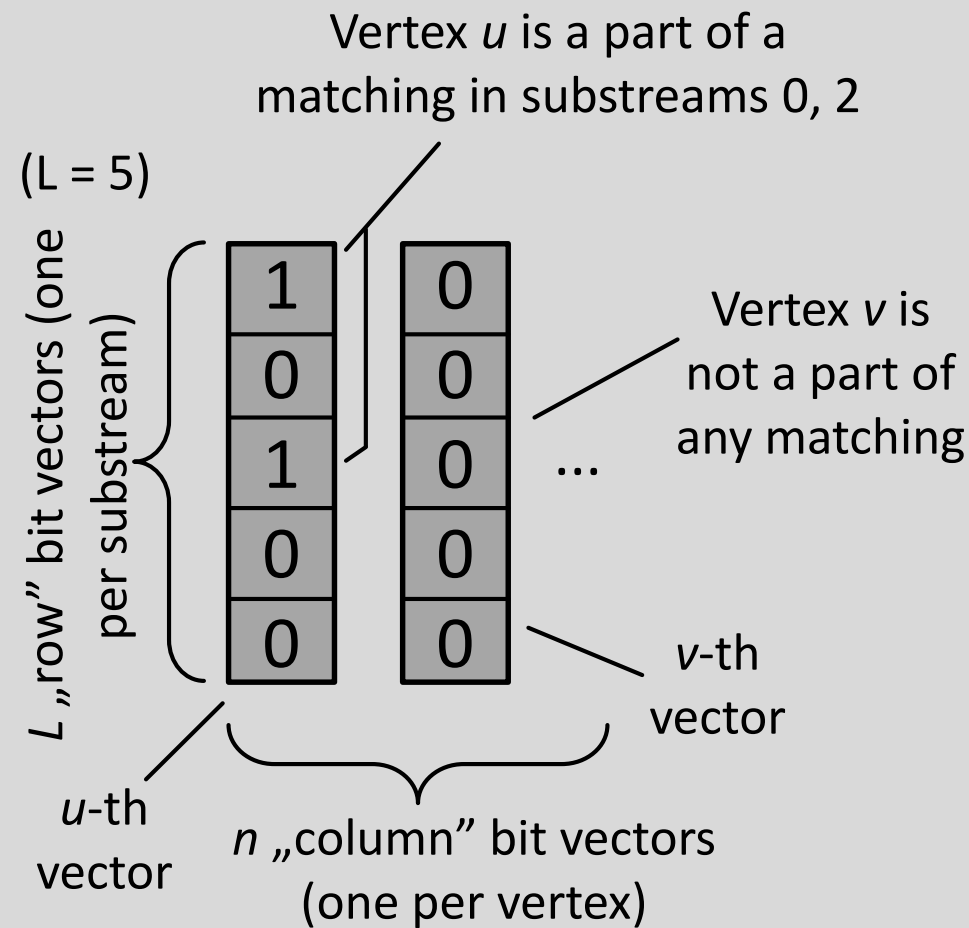
## Edges + matchings (more performance)



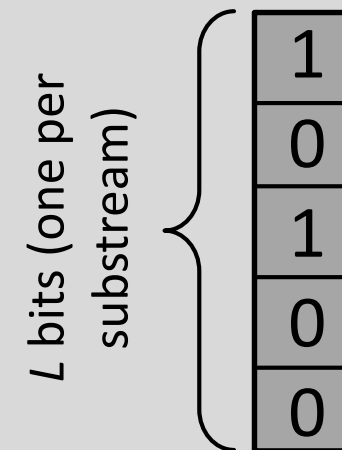
# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)




## Edges + matchings (more performance)




# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)



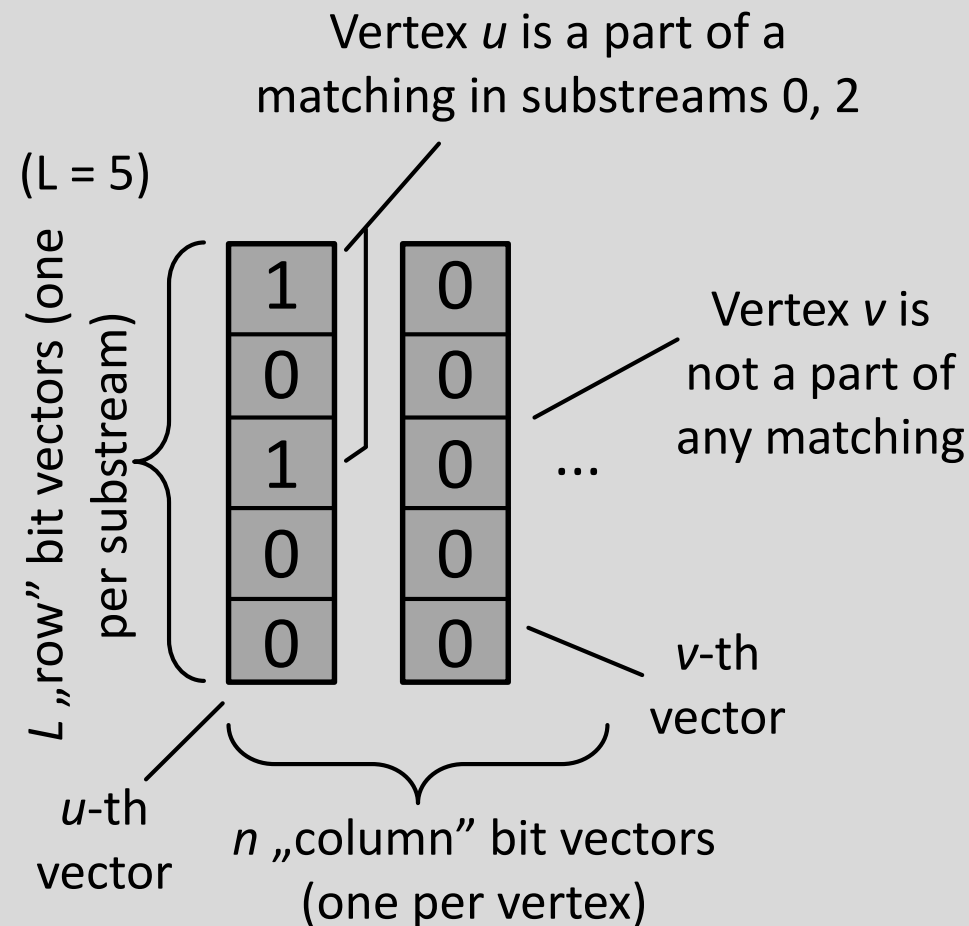
## Edges + matchings (more performance)



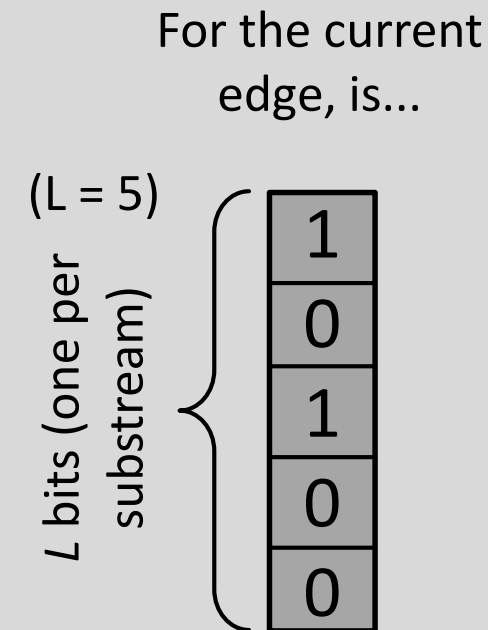
# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)



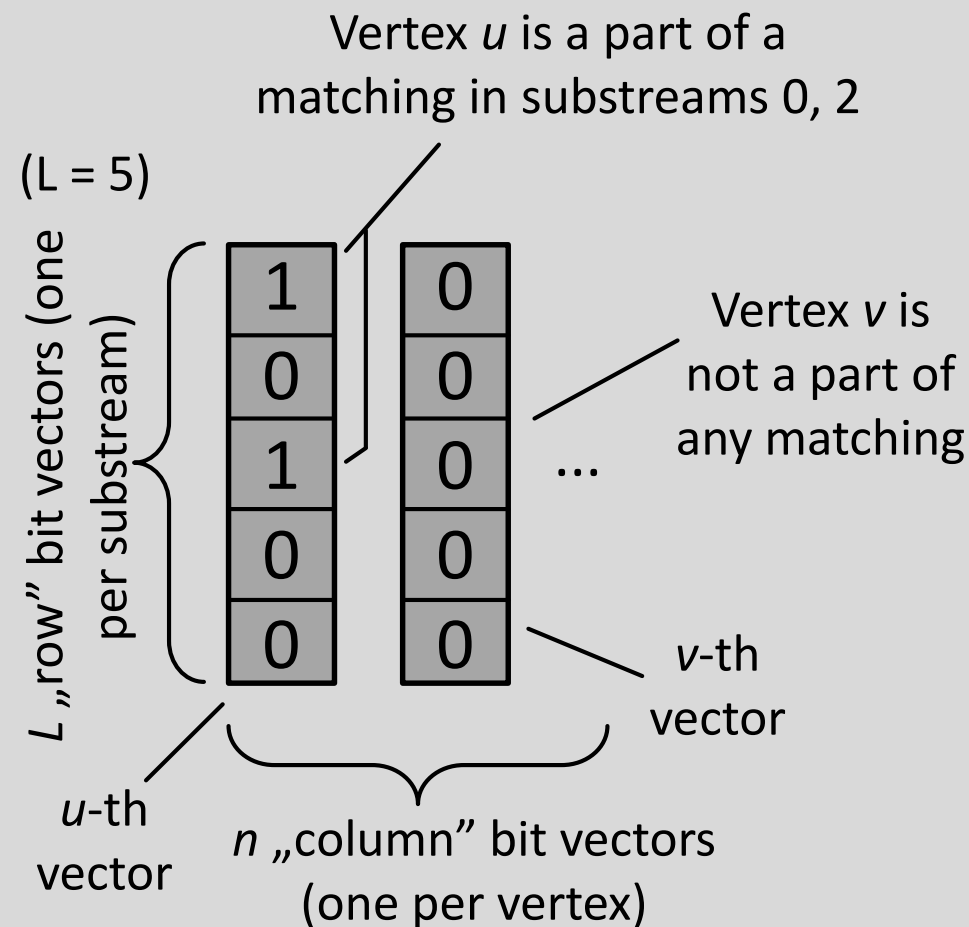
## Edges + matchings (more performance)



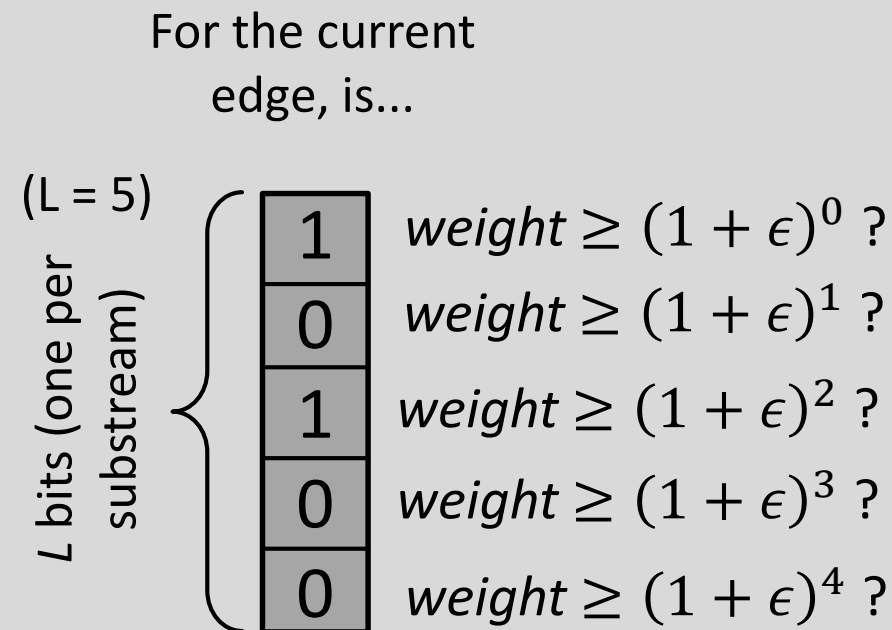
# MATCHING BITS: KEY DATA STRUCTURES FOR MAINTAINING INFORMATION ON MATCHINGS

An incoming edge...  $e = (u, v, weight)$

## Vertices + matchings (correctness)



## Edges + matchings (more performance)



# PERFORMANCE ANALYSIS

## VARIOUS #SUBSTREAMS (L)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

### Parameters:

Blocking size (K) = 32,

#threads = 4,

#edges = 16M

(Kronecker),  $\epsilon = 0.1$


#Substreams (pipelines)

# PERFORMANCE ANALYSIS

## VARIOUS #SUBSTREAMS (L)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

**Parameters:**  
 Blocking size (K) = 32,  
 #threads = 4,  
 #edges = 16M  
 (Kronecker),  $\epsilon = 0.1$






# PERFORMANCE ANALYSIS

## VARIOUS #SUBSTREAMS (L)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid

**Parameters:**  
 Blocking size (K) = 32,  
 #threads = 4,  
 #edges = 16M  
 (Kronecker),  $\epsilon = 0.1$



# PERFORMANCE ANALYSIS


## VARIOUS #SUBSTREAMS (L)

Algorithm	Platform
Crouch et al. [1] Sequential (CS-SEQ)	CPU
Crouch et al. [1] Parallel (CS-PAR)	CPU
Ghaffari [2] Sequential (G-SEQ)	CPU
Substream-Centric (SC-OPT)	Hybrid


### Parameters:

Blocking size (K) = 32,  
 #threads = 4,  
 #edges = 16M  
 (Kronecker),  $\epsilon = 0.1$

SC-OPT secures  
 highest performance  
 for all considered  
 values of parameters




## Large graphs...



## Large graphs...



# Large graphs...



# Large graphs...

