Task-Based Synchronization in Complex Noisy Networks

H. Guclu (LANL)

G. Korniss (Rensselaer), Z. Toroczkai (Notre Dame)

Algorithms, inference and statistical mechanics
santa fe, may 1-4, 2007

Los Alamos
NATIONAL LABORATORY
Ideas That Change the World

Task-Completion System (General)

$G_{\mathcal{T}}$

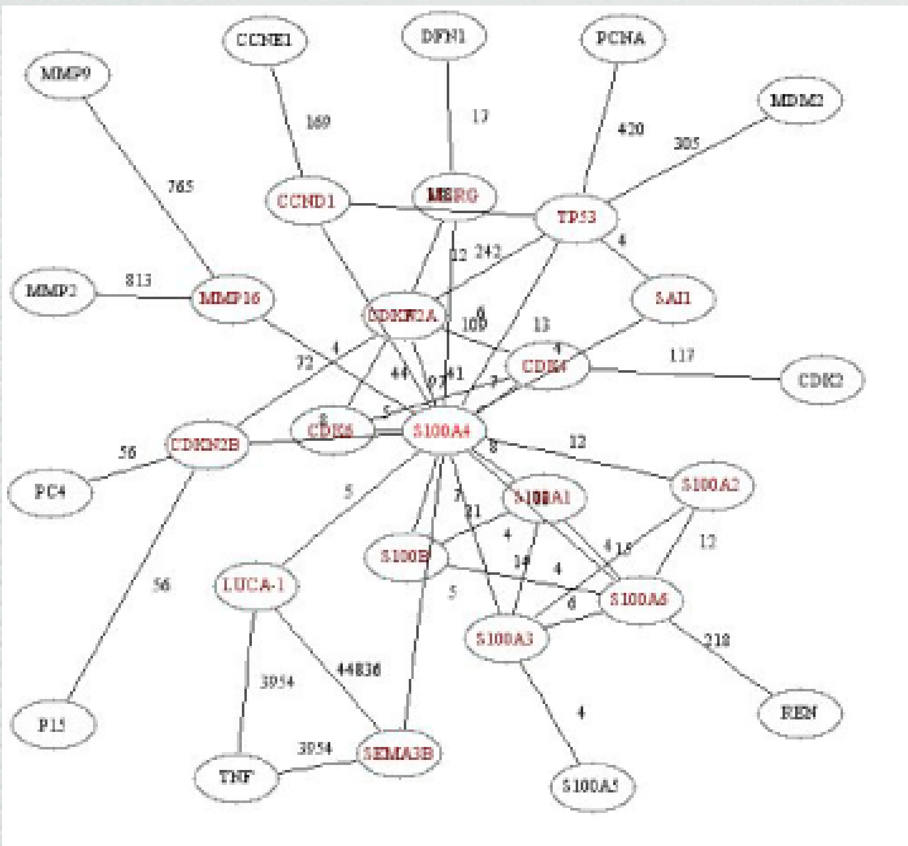

$G_{\mathcal{P}}$

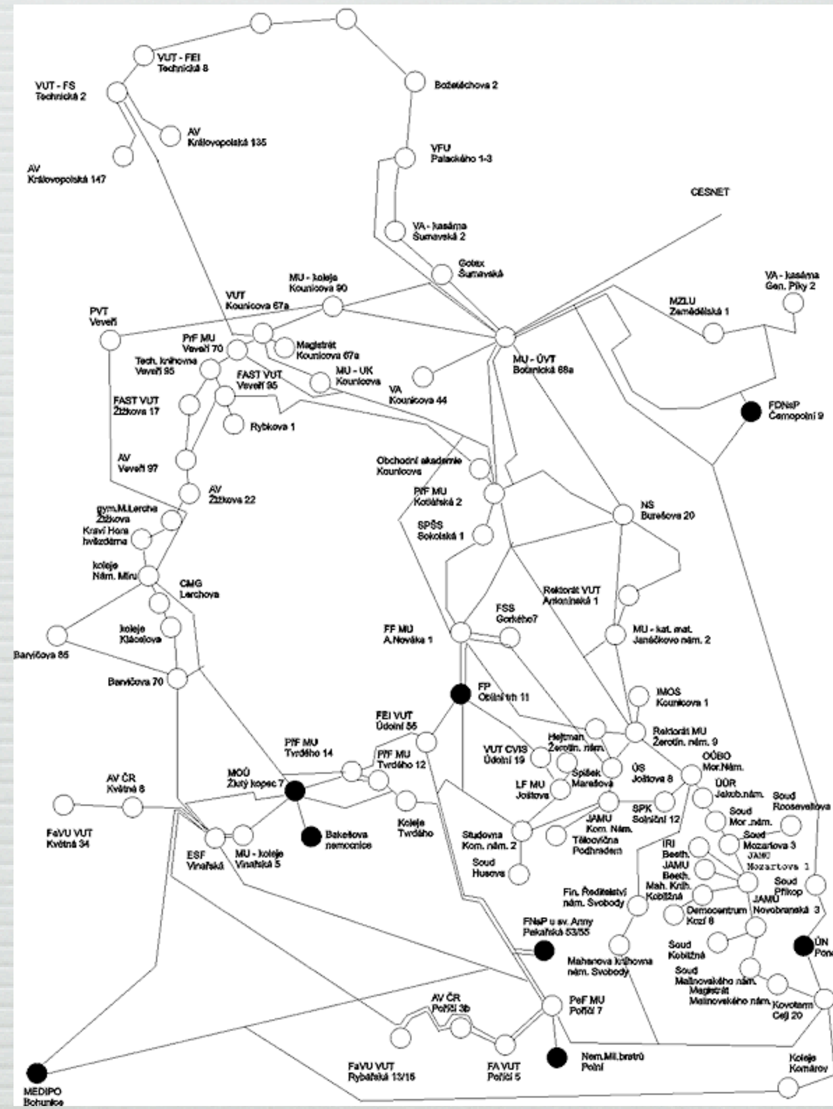

TAP, Spin glass

Task-Completion System (Particular)

$G_{\mathcal{T}}$

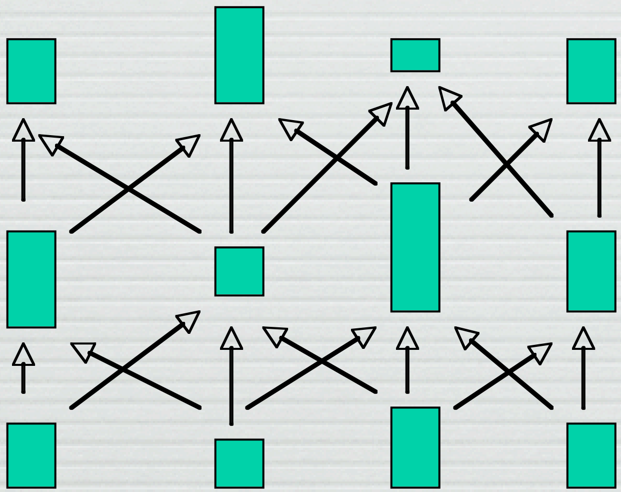

Magnetic systems,
factory physics, e-commerce.

$G_{\mathcal{P}}$

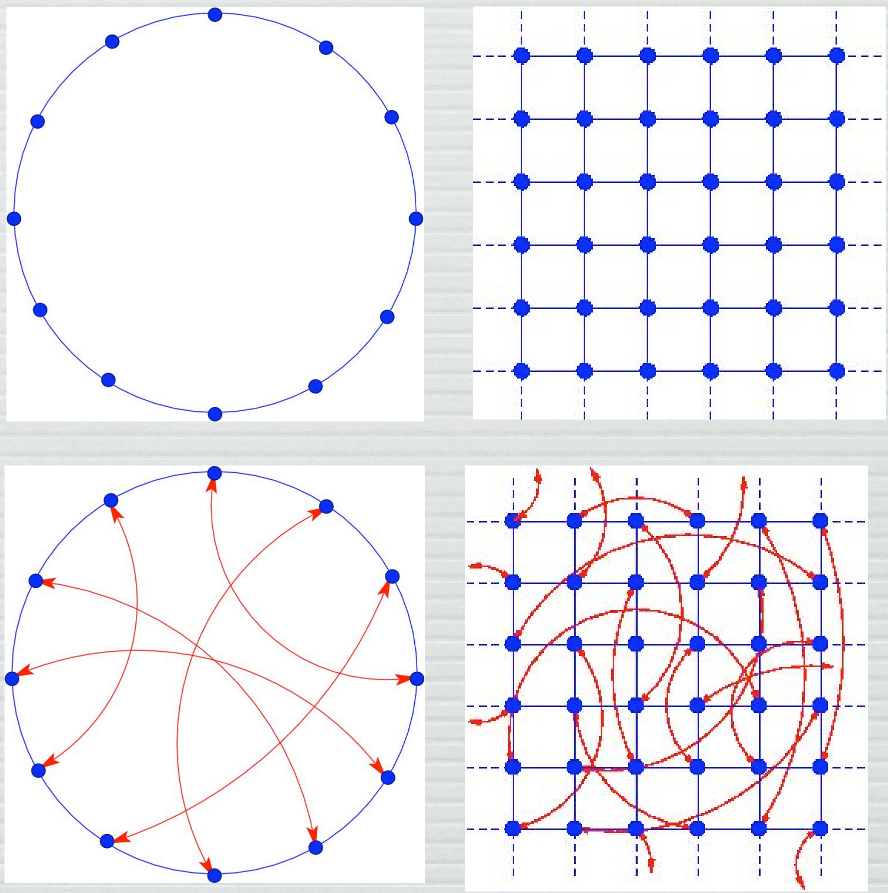

Conservative PDES

Lubachevsky '87, '88

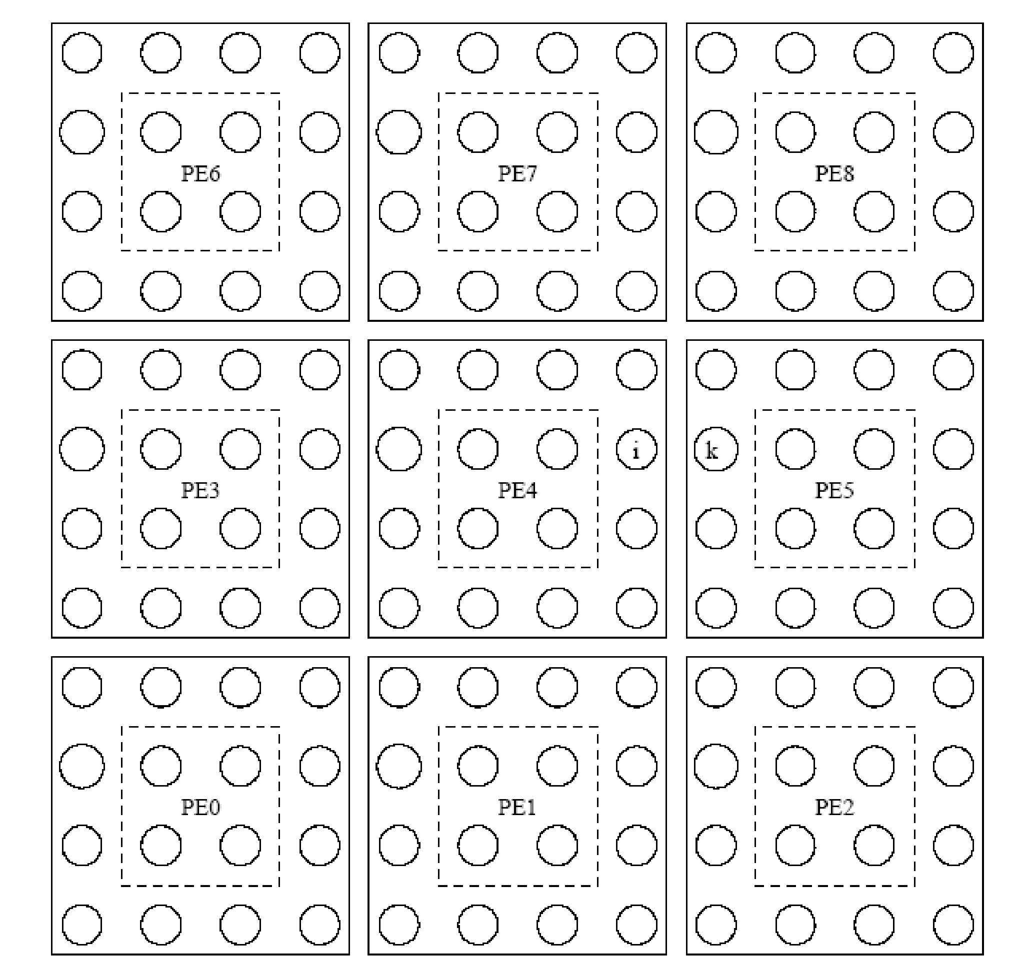

Korniss et al., '98

One spin per PE

A fundamental question of computing: Scalability

Mapping to surface growth problem

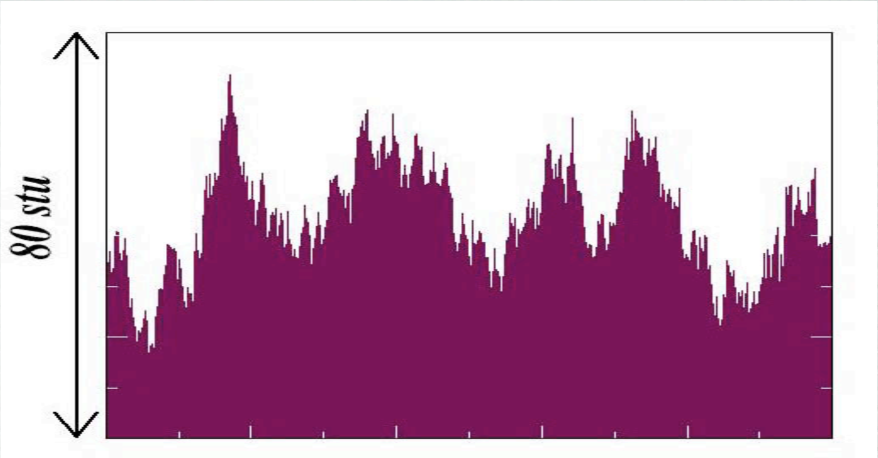

Computation Scalability

$$\langle u(L) \rangle = \left\langle \frac{\#\text{ of non-idling PEs}}{\text{total }\#\text{ of PEs}(L)} \right\rangle$$

$$\lim_{L\to\infty} \langle u(L) \rangle \to u(\infty) > 0$$

Measurement Scalability

$$\langle w^2(L) \rangle = \left\langle \frac{1}{L} \sum_{i=1}^{L} (\tau_i - \hat{\tau})^2 \right\rangle$$

$$\lim_{L\to\infty} \langle w^2(L) \rangle \to \text{const.}$$

BCS in 1D (Ring Topology)

This surface growth model is independent on the object of simulation, it corresponds to the massively parallel algorithm.

$$u(t, L) = \frac{\#\text{ of non-idling PEs}}{\text{total }\#\text{ of PEs }(L)} = \text{density of local minima}$$

$$\tau_i(t+1) = \tau_i(t) + \Theta[\tau_{i-1}(t) - \tau_i(t)]\Theta[\tau_{i+1}(t) - \tau_i(t)]\eta_i$$

$$\eta_i \text{ is independent of } t, i, \text{ and } \{\tau_i\}$$

BCS in 1D

Slope variables: $$\phi_i = \tau_i - \tau_{i-1}$$

$$\phi_i' - \phi_i = \Theta(-\phi_i)\Theta(\phi_{i+1})\eta_i - \Theta(-\phi_{i-1})\Theta(\phi_i)\eta_{i-1}$$

Biased diffusion

PBC: $$\sum_{i=1}^{L} \phi_i = 0$$

$$\langle \phi_i' \rangle - \langle \phi_i \rangle = -[\langle j_i \rangle - \langle j_{i-1} \rangle]$$

Continuity equation

$$\langle j_i \rangle = -\langle \Theta(-\phi_i)\Theta(\phi_{i+1}) \rangle$$

Average current

Utilization: $$\langle u \rangle = \frac{1}{L} \sum_{i=1}^{L} \langle \Theta(-\phi_i)\Theta(\phi_{i+1}) \rangle$$

BCS in 1D

Due to translational invariance:

$$\text{mean velocity of the surface} = |\,\langle j \rangle\,| = \langle u \rangle$$

Naïve coarse-graining: $$\Theta(\phi) = \lim_{\kappa \to 0} \Theta(\phi) = \lim_{\kappa \to 0} \frac{1}{2}[1 + \tanh \frac{\phi}{\kappa}]$$

To leading order in ϕ/κ

$$\langle \phi_i' \rangle - \langle \phi_i \rangle = \frac{1}{4\kappa} \langle \phi_{i+1} - 2\phi_i + \phi_{i-1} \rangle - \frac{1}{4\kappa^2} \langle \phi_i(\phi_{i+1} - \phi_{i-1}) \rangle$$

In the continuum limit:

$$\frac{\partial \hat{\phi}}{\partial t} = \frac{\partial^2 \hat{\phi}}{\partial x^2} - \lambda \frac{\partial}{\partial x}(\hat{\phi}^2)$$

Burger's equation for the coarse-grained field

BCS in 1D

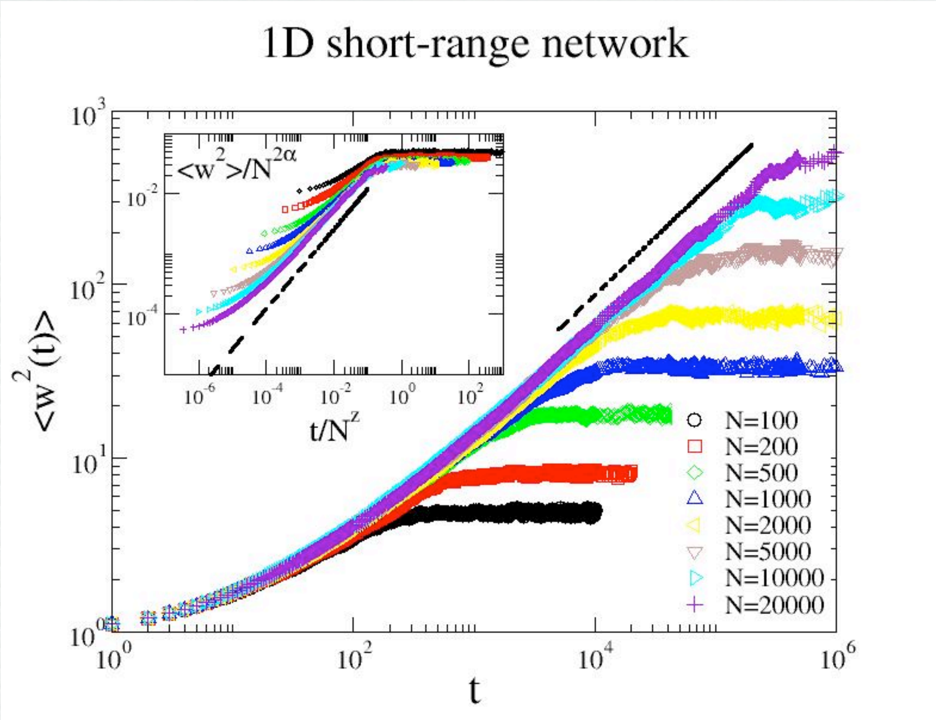

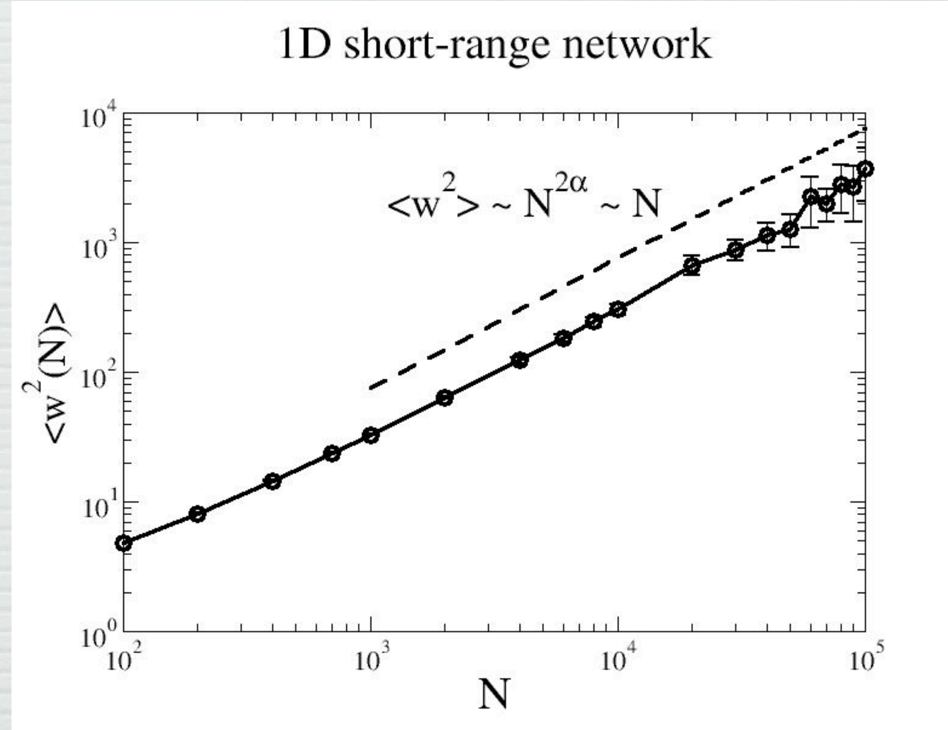

Width diverges ⇒ NOT measurement scalable

$$\alpha = \frac{1}{2},\ \beta = \frac{1}{3},\ z = \frac{\alpha}{\beta} = \frac{3}{2}$$

$$\langle w^2(N)\rangle \sim N^{2\alpha} = N \ \text{ in 1D}$$

$$\partial_i \hat{\tau}_i = \nabla^2 \hat{\tau}_i - \lambda(\nabla \hat{\tau}_i)^2 + \eta_i$$

BCS in 1D

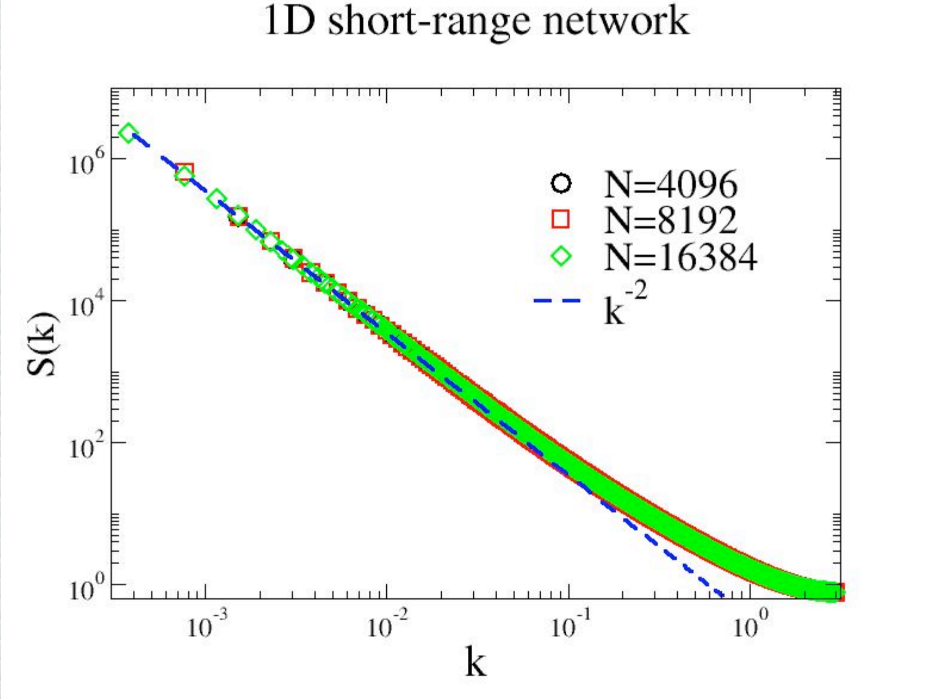

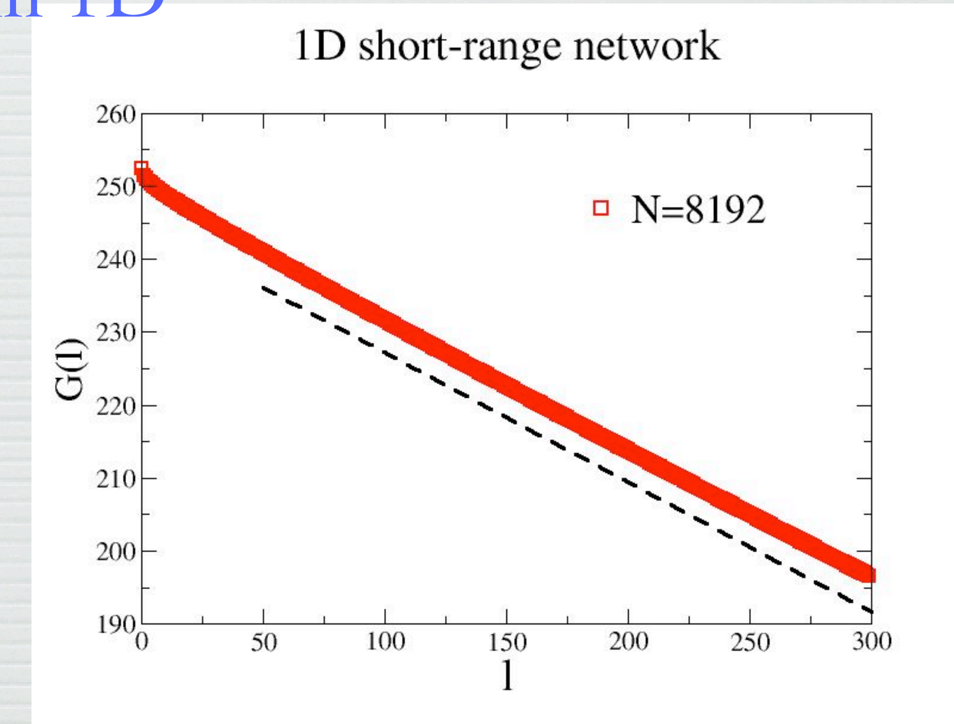

$$N\delta_{k,-k'}S(k,t) = \langle \tilde{\tau}_k(t)\tilde{\tau}_{k'}(t)\rangle$$

$$\tilde{\tau}_k = \sum_{j=1}^{N} e^{-ikj}\hat{\tau}_j$$

$$S(k) = \lim_{t\to\infty} S(k,t) = \frac{D}{2[1-\cos(k)]} \sim \frac{1}{k^2}$$

$$\langle w^2\rangle = \frac{1}{L}\sum_{k\neq 0} S(k) = G(0) \sim \frac{D}{12}L \sim L$$

$$G(l) = \frac{1}{L}\sum_{i=1}^{L} G_{i,i+l} = \frac{1}{L}\sum_{k\neq 0} e^{ikl}S(k)$$

$$G_{i,i+l} = \langle(\tau_i - \bar{\tau})(\tau_{i+l} - \bar{\tau})\rangle$$

$$G(l) \sim \frac{D}{2}\left(\frac{L}{6} - l\right)$$

Toroczkai et al., PRE, '00

BCS in 1D

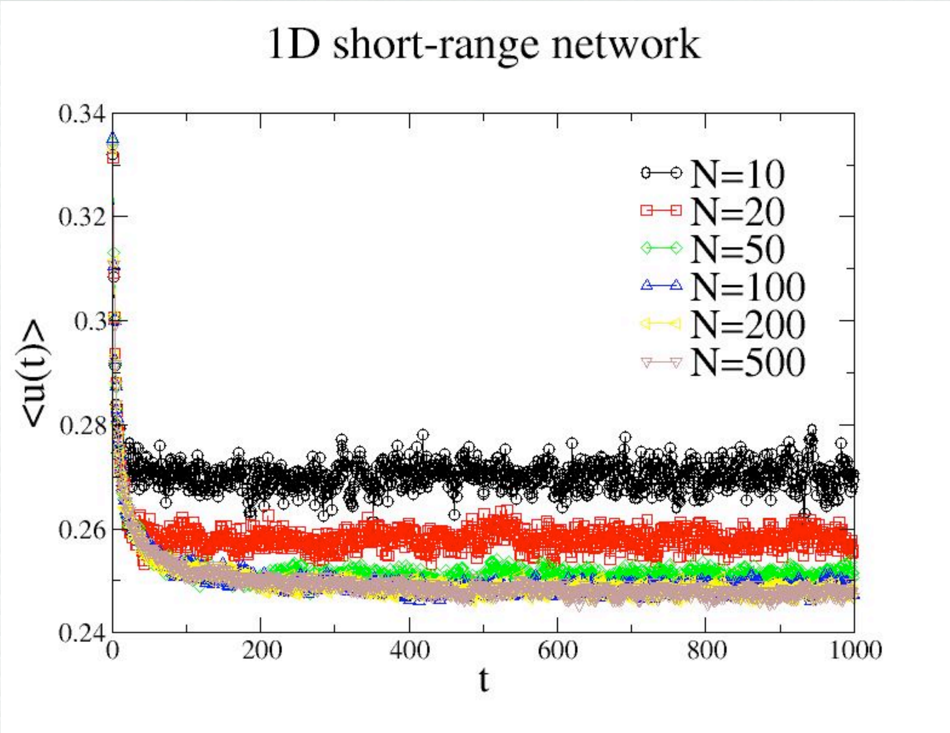

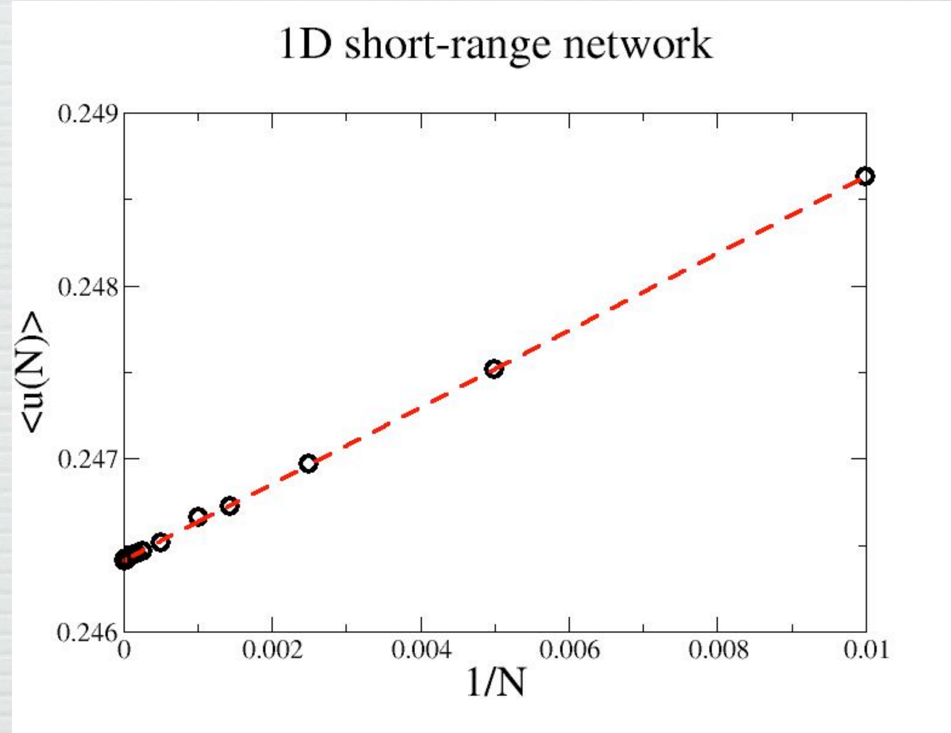

Finite utilization $\Rightarrow$ computationally scalable

$$\langle u(L)\rangle \simeq \langle u(\infty)\rangle + \frac{\text{const.}}{L}$$

$$\langle u(\infty)\rangle \simeq 0.2465$$

$$\text{const.} \simeq 0.2219$$

Regular Network

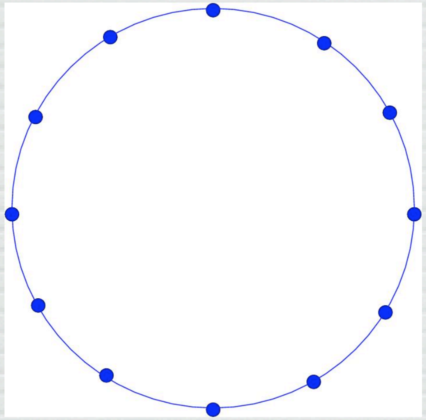

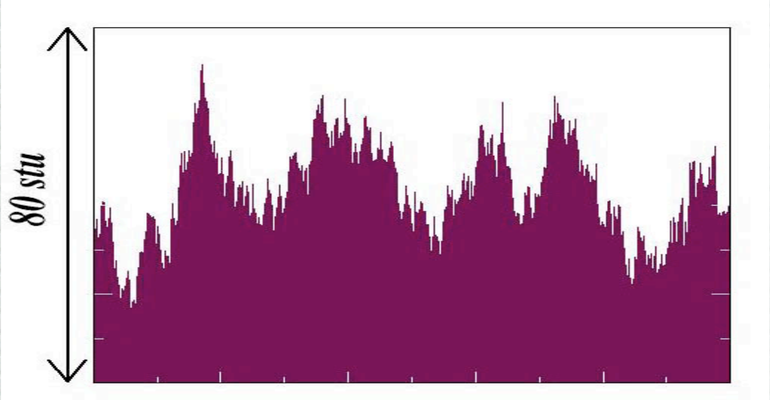

$$\lim_{L\to\infty} \langle w^2(N) \rangle \sim N$$

Small-World Network

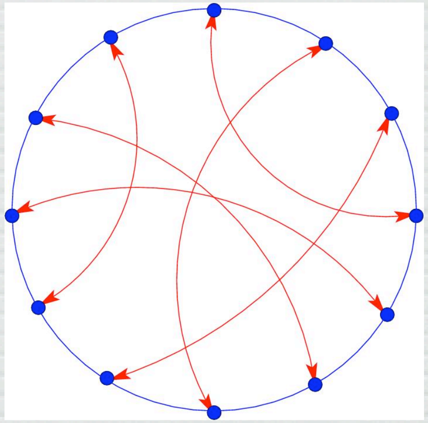

p

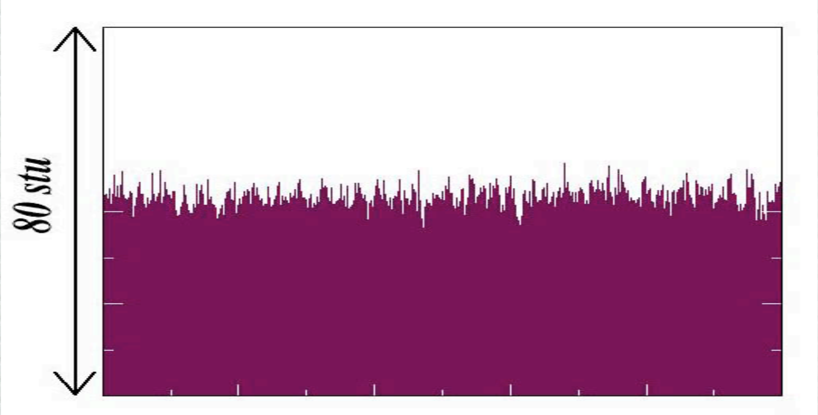

$$\lim_{L\to\infty} \langle w^2(N) \rangle \sim \text{const.}$$

Hastings, PRL (2003); Kozma et al., PRL (2003); Korniss et al., Science (2003)

SW in 1D

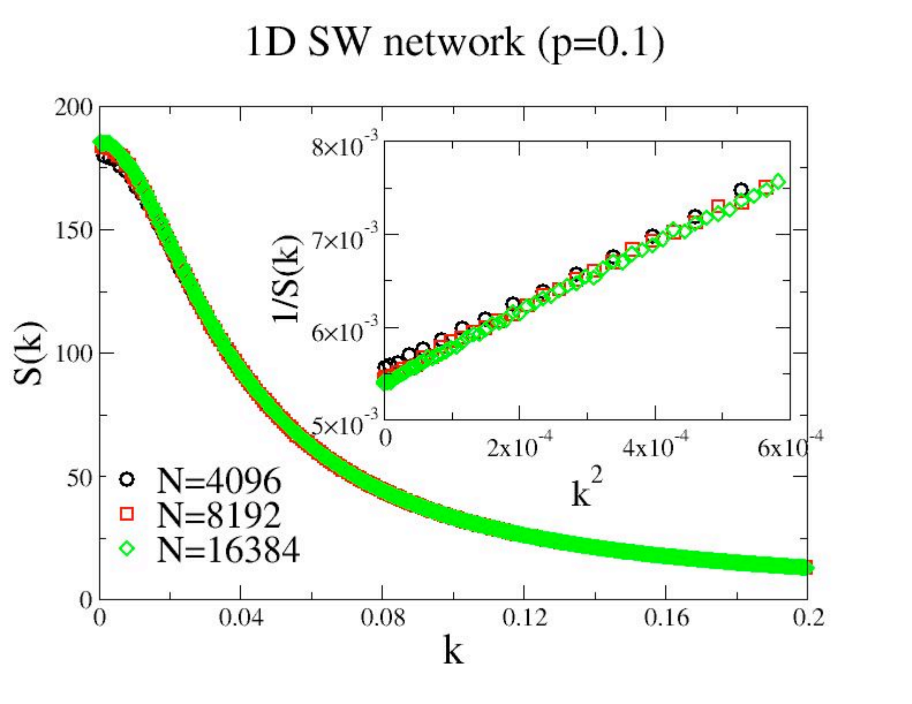

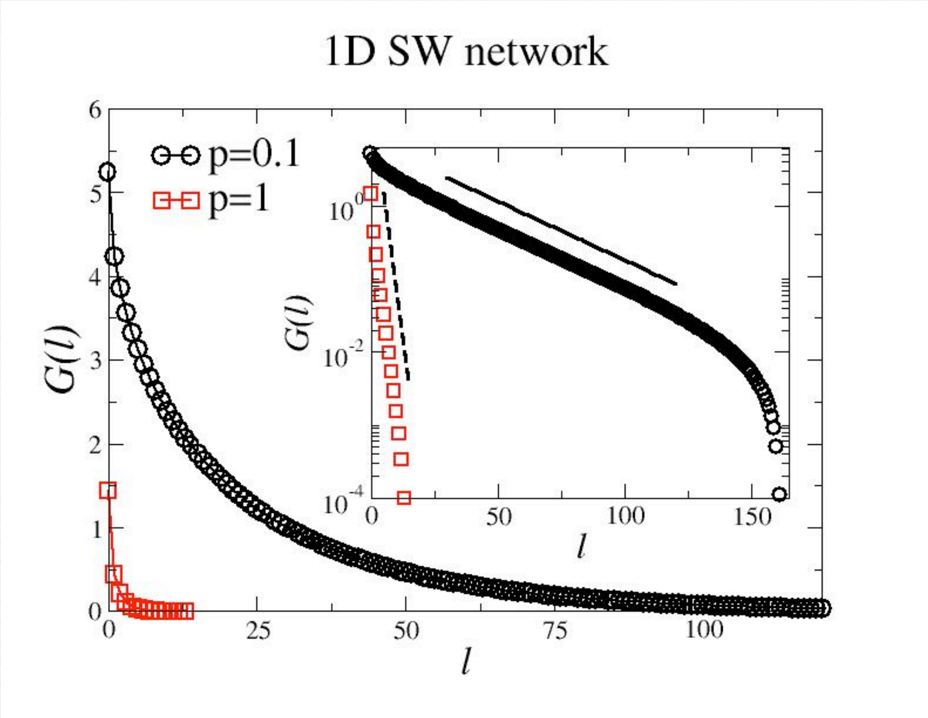

$$S(k) \sim \frac{1}{k^2 + \gamma} \quad \gamma = \gamma(p)$$

$$\langle w^2 \rangle = \frac{1}{L} \sum_{k \neq 0} S(k) \sim \frac{1}{\sqrt{\gamma}}$$

$$\xi(p) \sim \frac{1}{\sqrt{\gamma(p)}} \sim \frac{1}{p^{0.84}}$$

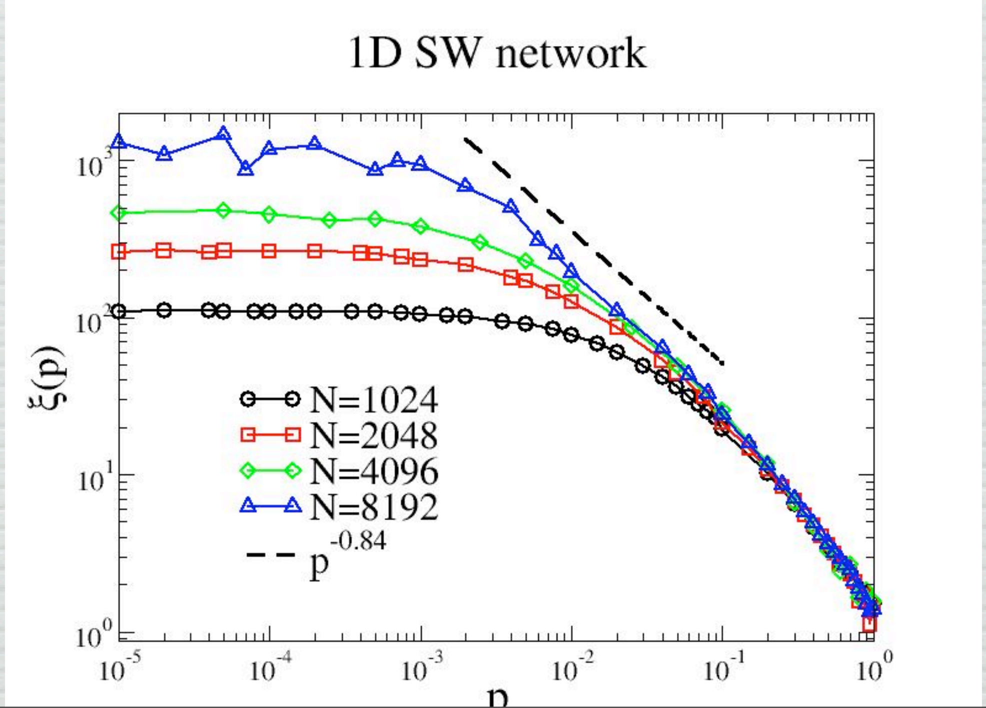

SW in 1D

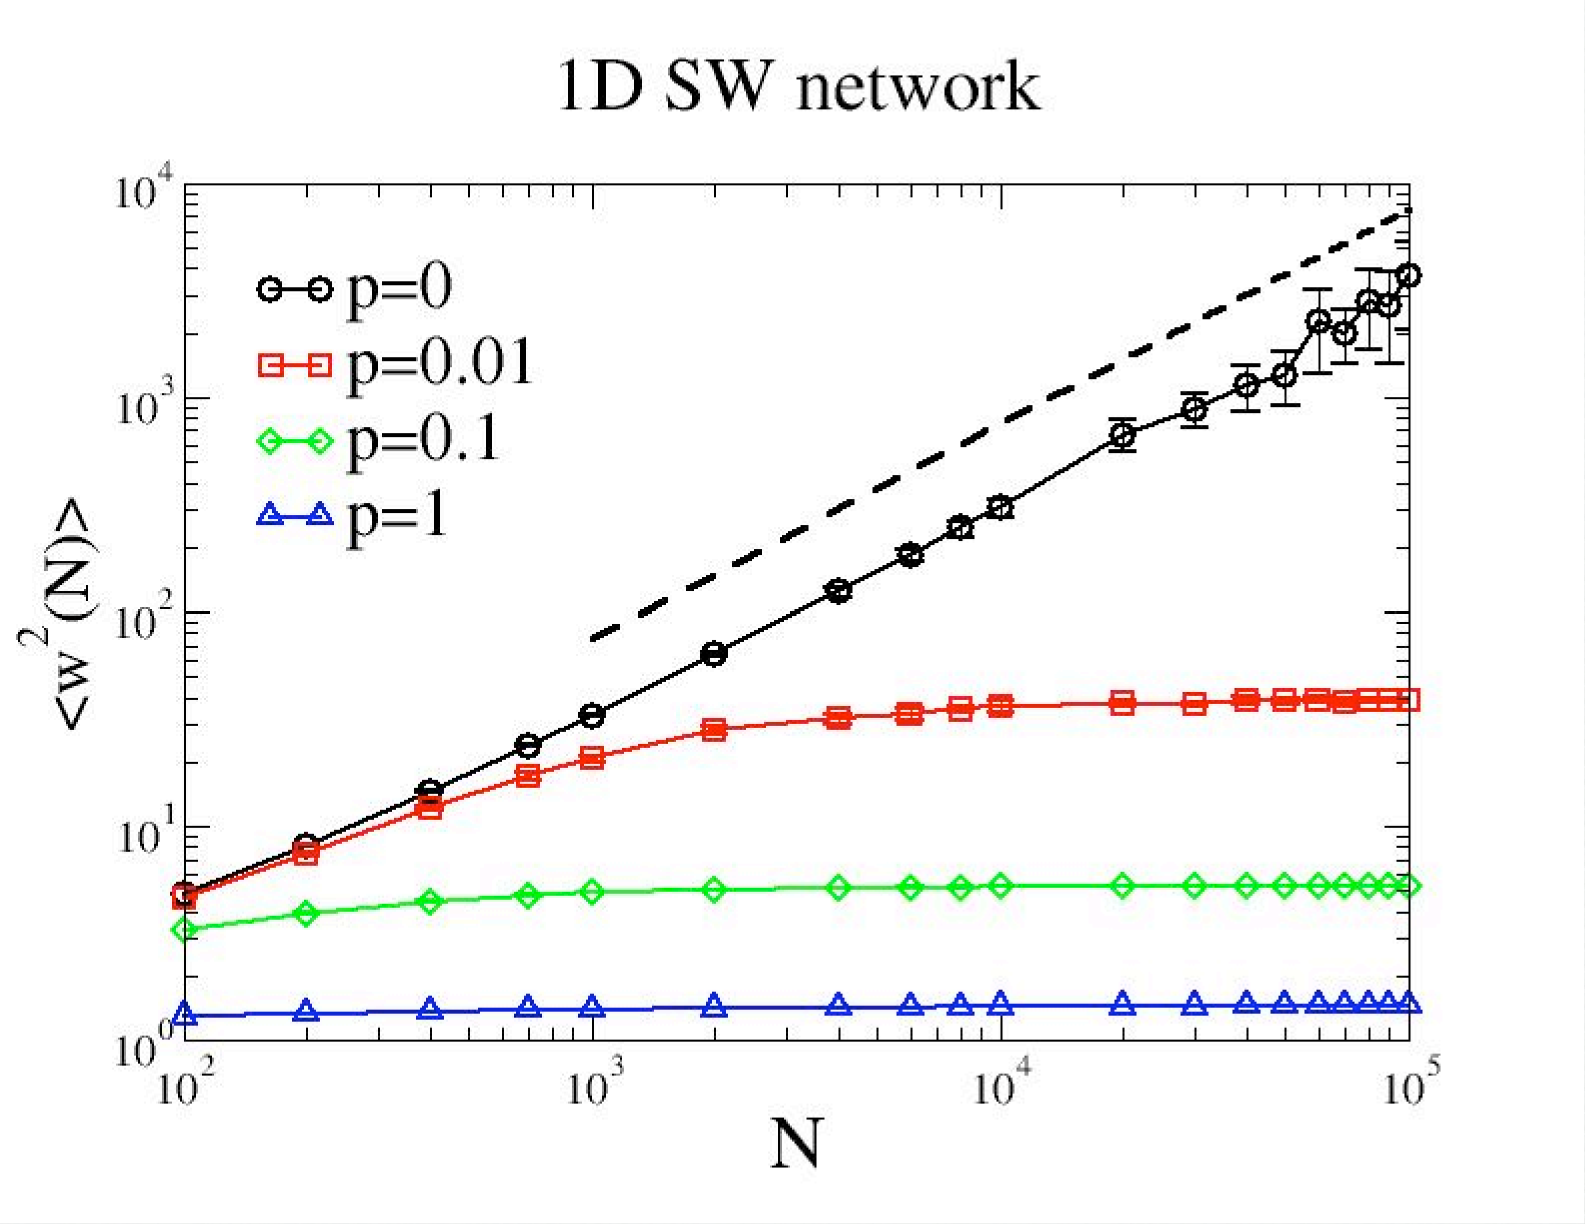

Finite width ⇒ measurement scalable

Korniss, et al., *Science, '03*

SW in 1D

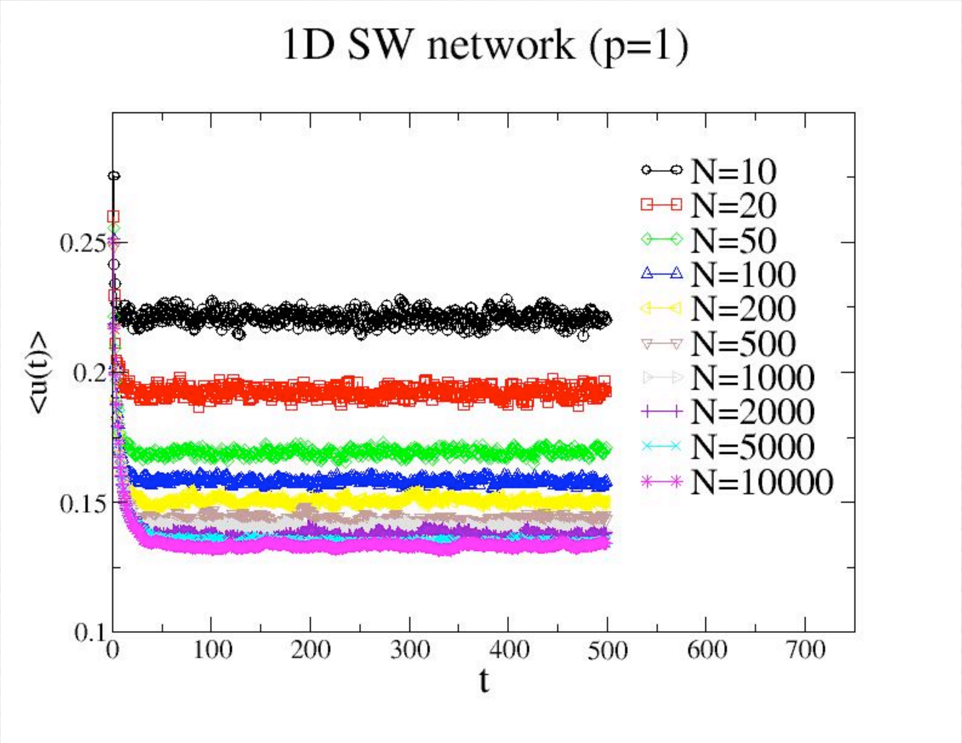

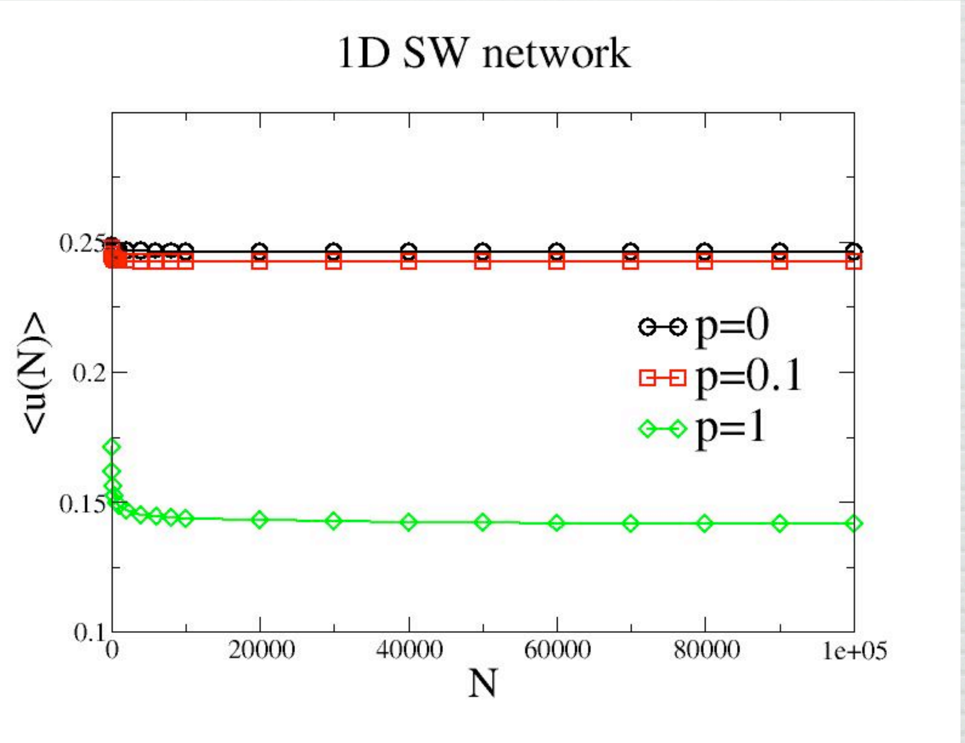

$$\langle u \rangle = (1-p)\left\langle \Theta(-\phi_{i-1})\Theta(\phi_i) \right\rangle + p\left\langle \Theta(-\phi_{i-1})\Theta(\phi_i)\Theta(-\phi_{r(i)}) \right\rangle$$

$$\langle u(\infty) \rangle_{p=0.1} \simeq 0.242 \qquad \langle u(\infty) \rangle_{p=1} \simeq 0.141$$

Finite utilization ⇒ computationally scalable

Edwards-Wilkinson Process on a network

G. Korniss et.al., cond-mat/0508056

Consider:

$$\partial_t h_i = -\sum_{j=1}^{N} A_{ij}(h_i - h_j) + \eta_i(t) \qquad (1)$$

where $h_i(t)$ is a scalar at a node (stochastic field variable such as virtual time)

$\eta_i(t)$ is delta-correlated white noise with zero mean and variance

$$\langle \eta_i(t)\eta_j(t')\rangle = 2\delta_{ij}\delta(t-t') \qquad (2)$$

$A_{ij} = A_{ji}$ is the effective coupling between nodes i and j, $A_{ii} = 0$

Defining the *Network Laplacian:*

$$\Gamma_{ij} = \delta_{ij}\sum_{l=1}^{N} A_{il} - A_{ij} \qquad (3)$$

(1) becomes:

$$\partial_t h_i = -\sum_{j=1}^{N} \Gamma_{ij} h_j + \eta_i(t) \qquad (4)$$

The steady-state 2-point equal time correlation function is given by:

$$G_{ij} \equiv \langle (h_i - \overline{h})(h_j - \overline{h}) \rangle = \hat{\Gamma}_{ij}^{-1} = \sum_{k=1}^{N-1} \frac{1}{\lambda_k} \psi_{ki} \psi_{kj} \qquad (5)$$

where $\overline{h} = \frac{1}{N} \sum_{i=1}^{N} h_i$ and $\langle \cdots \rangle$ denotes averaging over noise

$\hat{\Gamma}^{-1}$ is the inverse of Γ in the space orthogonal to the zero mode.

$\{ \lambda_k, \{\psi_{ki}\}_{i=1}^{i=N} \}, \quad k = \overline{0, N-1}$ are the k^{th} eigenvalues and normalized eigenvectors.

$k = 0$ represents the zero mode of the network where $\lambda_0 = 0$

Thus

$$\langle w^2 \rangle = \left\langle \frac{1}{N} \sum_{i=1}^{N} \left(h_i - \overline{h} \right)^2 \right\rangle = \frac{1}{N} \sum_{i=1}^{N} G_{ii} = \frac{1}{N} \sum_{k=1}^{N-1} \frac{1}{\lambda_k} \qquad (6)$$

For large systems and quenched network disorder, typically we have self-averaging:

$\langle w^2 \rangle \simeq [\langle w^2 \rangle]$ $\Longrightarrow$ calculate $[G_{ii}]$ get $N \to \infty$ limit.

Summary and Conclusions

- BCS exhibits KPZ-like roughening.
- SW-synchronized task-completion systems exhibit mean-field like characteristics.
- SW links generate an effective mass for the propagator of the virtual time horizon (in a field theory sense) corresponding to a finite correlation length and consequently the width becomes finite for an arbitrary small rate of synchronization through SW links while the utilization remains nonzero, yielding a fully scalable task-completion scheme.
- Systems exhibit (strict or anomalous) mean-field-like behavior when the original short-range interaction topology is modified to a SW network.
- When the interaction topology in a network is changed into SW, the extreme fluctuations diverge weakly (logarithmically) and in a power-law fashion with the system size when the noise is short-tailed and heavy-tailed, respectively.
- Our work is applicable to systems with "local" relaxation in a noisy environment.

Acknowledgment: NNSA (DOE) and NSF